

13th International Conference on Typed Lambda Calculi and Applications

TLCA'15, July 1–3, 2015, Warsaw, Poland

Edited by

Thorsten Altenkirch



Editor

Thorsten Altenkirch
School of Computer Science
University of Nottingham
txa@cs.nott.ac.uk

ACM Classification 1998

D.1.1 Applicative (Functional) Programming, D.3.2 Language Classifications, D.3.3 Language Constructs and Features, E.1 Data Structures, F.3 Logics and Meanings of Programs, F.4 Mathematical Logic and Formal Languages, I.1 Symbolic and Algebraic Manipulation, I.2.3 Deduction and Theorem Proving

ISBN 978-3-939897-87-3

Published online and open access by

Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany. Online available at <http://www.dagstuhl.de/dagpub/978-3-939897-87-3>.

Publication date

July, 2015

Bibliographic information published by the Deutsche Nationalbibliothek

The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at <http://dnb.d-nb.de>.

License

This work is licensed under a Creative Commons Attribution 3.0 Unported license (CC-BY 3.0): <http://creativecommons.org/licenses/by/3.0/legalcode>.



In brief, this license authorizes each and everybody to share (to copy, distribute and transmit) the work under the following conditions, without impairing or restricting the authors' moral rights:

- Attribution: The work must be attributed to its authors.

The copyright is retained by the corresponding authors.

Digital Object Identifier: 10.4230/LIPIcs.TLCA.2015.i

ISBN 978-3-939897-87-3

ISSN 1868-8969

<http://www.dagstuhl.de/lipics>

LIPICs – Leibniz International Proceedings in Informatics

LIPICs is a series of high-quality conference proceedings across all fields in informatics. LIPICs volumes are published according to the principle of Open Access, i.e., they are available online and free of charge.

Editorial Board

- Susanne Albers (TU München)
- Chris Hankin (Imperial College London)
- Deepak Kapur (University of New Mexico)
- Michael Mitzenmacher (Harvard University)
- Madhavan Mukund (Chennai Mathematical Institute)
- Catuscia Palamidessi (INRIA)
- Wolfgang Thomas (*Chair*, RWTH Aachen)
- Pascal Weil (CNRS and University Bordeaux)
- Reinhard Wilhelm (Saarland University)

ISSN 1868-8969

<http://www.dagstuhl.de/lipics>

■ Contents

Preface	
<i>Thorsten Altenkirch</i>	vii
Herbrand Disjunctions, Cut Elimination and Context-Free Tree Grammars	
<i>Bahareh Afshari, Stefan Hetzl, and Graham E. Leigh</i>	1
Non-Wellfounded Trees in Homotopy Type Theory	
<i>Benedikt Ahrens, Paolo Capriotti, and Régis Spadotti</i>	17
Conservativity of Embeddings in the $\lambda\Pi$ Calculus Modulo Rewriting	
<i>Ali Assaf</i>	31
Models for Polymorphism over Physical Dimensions	
<i>Robert Atkey, Neil Ghani, Fredrik Nordvall Forsberg, Timothy Revell,</i> <i>and Sam Staton</i>	45
MALL Proof Equivalence is Logspace-Complete, via Binary Decision Diagrams	
<i>Marc Bagnol</i>	60
Mixin Composition Synthesis Based on Intersection Types	
<i>Jan Bessai, Andrej Dudenhefner, Boris Döder, Tzu-Chun Chen, Ugo de'Liguoro,</i> <i>and Jakob Rehof</i>	76
Non-Constructivity in Kan Simplicial Sets	
<i>Marc Bezem, Thierry Coquand, and Erik Parmann</i>	92
Logical Relations for Coherence of Effect Subtyping	
<i>Dariusz Biernacki and Piotr Polesiuk</i>	107
Observability for Pair Pattern Calculi	
<i>Antonio Bucciarelli, Delia Kesner, and Simona Ronchi Della Rocca</i>	123
Undecidability of Equality in the Free Locally Cartesian Closed Category	
<i>Simon Castellan, Pierre Clairambault, and Peter Dybjer</i>	138
The Inconsistency of a Brouwerian Continuity Principle with the Curry–Howard Interpretation	
<i>Martín Hötzel Escardó and Chuangjie Xu</i>	153
Curry-Howard for Sequent Calculus at Last!	
<i>José Espírito Santo</i>	165
Dependent Types for Nominal Terms with Atom Substitutions	
<i>Elliot Fairweather, Maribel Fernández, Nora Szasz, and Alvaro Tasistro</i>	180
Realizability Toposes from Specifications	
<i>Jonas Frey</i>	196
Standardization of a Call-By-Value Lambda-Calculus	
<i>Giulio Guerrieri, Luca Paolini, and Simona Ronchi Della Rocca</i>	211

13th International Conference on Typed Lambda Calculi and Applications (TLCA 2015).

Editor: Thorsten Altenkirch



Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Wild ω -Categories for the Homotopy Hypothesis in Type Theory <i>André Hirschowitz, Tom Hirschowitz, and Nicolas Tabareau</i>	226
Multivariate Amortised Resource Analysis for Term Rewrite Systems <i>Martin Hofmann and Georg Moser</i>	241
Termination of Dependently Typed Rewrite Rules <i>Jean-Pierre Jouannaud and Jianqi Li</i>	257
Well-Founded Recursion over Contextual Objects <i>Brigitte Pientka and Andreas Abel</i>	273
Polynomial Time in the Parametric Lambda Calculus <i>Brian F. Redmond</i>	288
Fibrations of Tree Automata <i>Colin Riba</i>	302
Multi-Focusing on Extensional Rewriting with Sums <i>Gabriel Scherer</i>	317
A Proof-theoretic Characterization of Independence in Type Theory <i>Yuting Wang and Kaustuv Chaudhuri</i>	332

■ Preface

This volume contains the papers of the 13th International Conference on Typed Lambda Calculi and Applications (TLCA 2015), which was held during 1–3 July 2015, in Warsaw, Poland. TLCA 2015 was part of the 8th International Conference on Rewriting, Deduction, and Programming (RDP 2015), together with the 26th International Conference on Rewriting Techniques and Applications (RTA 2015), the Workshop on Higher-Dimensional Rewriting and Applications (HDRA), the Workshop on Homotopy Type Theory / Univalent Foundations (HoTT/UF), the 29th International Workshop on Unification (UNIF 2015), the second International Workshop on Rewriting Techniques for Program Transformations and Evaluation (WPTE 2015), and the annual meeting of the IFIP Working Group 1.6 on Term Rewriting.

The TLCA series of conferences serves as a forum for presenting original research results that are broadly relevant to the theory and applications of lambda calculus. Previous TLCA conferences were held in Utrecht (1993), Edinburgh (1995), Nancy (1997), L’Aquila (1999), Kraków (2001), Valencia (2003), Nara (2005), Paris (2007), Brasília (2009), Novi Sad (2011), Eindhoven (2013) and Vienna (2014, merged with RTA).

A total of 23 papers were accepted out of 46 submissions for presentation at TLCA 2015 and for inclusion in the proceedings. I would like to thank everyone who submitted a paper and to express my regret that many interesting papers could not be included. Each submitted paper was reviewed by at least three members of the Programme Committee, who were assisted in their work by 71 external reviewers. I thank the members of the Programme Committee and the external reviewers for their review work, as well as Andrei Voronkov for providing the EasyChair system which proved invaluable throughout the review process and the preparation of this volume.

In addition to the contributed papers, the TLCA 2015 programme contained invited talks by Helene Kirchner (joint with RTA 2015), Herman Geuvers and Martin Hofmann.

Many people helped to make TLCA 2015 a success. I would like to thank RDP 2015 Chair Aleksy Schubert and the Organising Committee, the local organising team, TLCA Publicity Chair Luca Paolini, and the TLCA Steering Committee, especially Paweł Urzyczyn.

June 2015

Thorsten Altenkirch



■ TLCA 2015 Organisation

Program Committee

Thorsten Altenkirch (Chair)	University of Nottingham, UK
Steve Awodey	Carnegie Mellon University, USA
Stefano Berardi	University of Torino, Italy
James Chapman	Technical University of Tallinn, Estonia
Gilles Dowek	INRIA, France
Peter Dybjer	Chalmers University, Sweden
Silvia Ghilezan	University of Novi Sad, Serbia
Mauro Jaskelioff	University of Rosario, Argentina
Chantal Keller	Microsoft Research Cambridge, UK
Paul Levy	University of Birmingham, UK
Ralph Matthes	IRIT, CNRS & University of Toulouse, France
Keiko Nakata	FireEye Dresden, Germany
Damian Niwiński	University of Warsaw, Poland
Valeria de Paiva	Nuance Communications, USA
Matthieu Sozeau	INRIA, France
Wouter Swierstra	University of Utrecht, The Netherlands

RDP Conference Chair

Aleksy Schubert	University of Warsaw, Poland
-----------------	------------------------------

Organising Committee

Marcin Benke	University of Warsaw, Poland
Jacek Chrząszcz	University of Warsaw, Poland
Patryk Czarnik	University of Warsaw, Poland
Łukasz Czajka	University of Warsaw, Poland
Krystyna Jaworska	University of Warsaw, Poland
Paweł Urzyczyn	University of Warsaw, Poland
Daria Walukiewicz-Chrząszcz	University of Warsaw, Poland
Maciej Zielenkiewicz	University of Warsaw, Poland



TLCA Steering Committee

Sandra Alves	University of Porto, Portugal
Steve Awodey	Carnegie Mellon University, USA
Pierre-Louis Curien	CNRS and Université Paris Diderot, France
Ugo Dal Lago	University of Bologna, Italy
Gilles Dowek	INRIA, France
Masahito Hasegawa	Kyoto University, Japan
Hugo Herbelin	Université Paris Diderot, France
Martin Hofmann	Ludwig-Maximilians-University, Munich, Germany
Luke Ong	University of Oxford, UK
Michele Pagani	Université Paris Diderot, France
Jens Palsberg	University of California, Los Angeles,, USA
Jakob Rehof	University of Dortmund, Germany
Paweł Urzyczyn (Chair)	University of Warsaw, Poland
Philip Wadler	University of Edinburgh, UK

TLCA Honorary Advisors

Samson Abramsky	University of Oxford, UK
Henk Barendregt	Radboud University Nijmegen, The Netherlands
Roger Hindley	Swansea University, UK
Mariangiola Dezani-Ciancaglini	University of Turin, Italy
Simona Ronchi Della Rocca	University of Turin, Italy

TLCA Publicity Chair

Luca Paolini	University of Turin, Italy
--------------	----------------------------

External Reviewers

Andreas Abel
Ki Yung Ahn
Bruno Barras
Alexis Bernadet
Marc Bezem
Katalin Bimbó
James Brotherston
Ulrik Buchholtz
Jacek Chrzęszcz
Ranald Clouston
Thierry Coquand
Pierre-Louis Curien
Łukasz Czajka
Nils Anders Danielsson
Alejandro Diaz-Caro
Harley Eades Iii
Didier Galmiche
Ferruccio Guidi
Jennifer Hackett
Hugo Herbelin
Martin Hofmann
Chung-Kil Hur
Danko Ilik
Florent Jacquemard
Alan Jeffrey
Kentaro Kikuchi
Daisuke Kimura
Nicolai Kraus
Keiichirou Kusakari
Jarek Kusmieriek
James Laird
Pierre Lescanne
Silvia Likavec
Sam Lindley
Peter LeFanu Lumsdaine
Ian Mackie

Assia Mahboubi
Samuele Maschio
Conor McBride
Henryk Michalewski
Edward Morehouse
J. Garrett Morris
Guillaume Munch-Maccagnoni
Koji Nakazawa
Clive Newstead
Christine Paulin-Mohring
Iosif Petrakis
Elaine Pimentel
Andrew Pitts
Nicolas Pouillard
Jorge A. Parez
Colin Riba
Eike Ritter
Exequiel Rivas
Edmund Robinson
Reuben Rowe
Jorge Luis Sacchini
Alexis Saurin
Bas Spitters
Sam Staton
Martin Sulzmann
Nicolas Tabareau
Tarmo Uustalu
Benno van den Berg
Andrea Vezzosi
Ignas Vysniauskas
Benjamin Werner
Hongwei Xi
Noam Zeilberger
Yu Zhang
Florian Zuleger

Herbrand Disjunctions, Cut Elimination and Context-Free Tree Grammars

Bahareh Afshari, Stefan Hetzl, and Graham E. Leigh

Institute of Discrete Mathematics and Geometry
Vienna University of Technology
Wiedner Hauptstraße 8-10, 1040 Vienna, Austria
{bahareh.afshari, stefan.hetzl, graham.leigh}@tuwien.ac.at

Abstract

Recently a new connection between proof theory and formal language theory was introduced. It was shown that the operation of cut elimination for proofs in first-order predicate logic involving Π_1 -cuts corresponds to computing the language of a particular class of regular tree grammars. The present paper expands this connection to the level of Π_2 -cuts. Given a proof π of a Σ_1 formula with cuts only on Π_2 formulæ, we show there is associated to π a natural context-free tree grammar whose language is finite and yields a Herbrand disjunction for π .

1998 ACM Subject Classification F.4.1 Mathematical Logic, F.4.2 Grammars and Other Rewriting Systems, F.4.3 Formal Languages

Keywords and phrases Classical logic, Context-free grammars, Cut elimination, First-order logic, Herbrand's theorem, Proof theory

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.1

1 Introduction

The computational content of proofs is a central topic of proof theory. In intuitionistic first-order logic the existential witness property states that the provability of $\exists xF$ entails the existence of some ground term t such that $F(x/t)$ is provable. The analogue of this property in classical logic is Herbrand's theorem [12] (see also [5]). In its simplest version Herbrand's theorem states that if $\exists xF$ is valid and F quantifier-free there exist closed terms t_1, \dots, t_k such that $\bigvee_{i=1}^k F(x/t_i)$ is a tautology. Such formulæ, quantifier-free disjunctions of instances, are hence also called Herbrand disjunctions. Herbrand's theorem applies not only to existential but arbitrary first-order formulæ, providing a tautology by replacing each non-prenex quantifier with a suitable finite disjunction or conjunction of instances. Provided one is willing to speak about provability instead of validity Herbrand's theorem even extends to classical higher-order logic, see for example [25].

A Herbrand disjunction can be read off directly from a cut-free proof though proofs with cut may be non-elementarily smaller than the shortest Herbrand disjunction [30, 26, 27]. Therefore, in general, cut elimination (or an equivalent normalisation process) is necessary in order to obtain a Herbrand disjunction. However, if one is only interested in the witness terms of a Herbrand disjunction and not in the complete cut-free proof then it would be desirable to circumvent the cumbersome process of cut elimination.

For instance, in [13, 14] it was shown that a proof $\pi \vdash \exists xF$ in which all cut formulæ have at most one quantifier induces a (totally rigid acyclic) tree grammar \mathcal{G}_π , of size no greater than the size of the proof, with a finite language that, when interpreted as a collection of witness terms, forms a Herbrand disjunction for $\exists xF$ (see Figure 2).



© Bahareh Afshari, Stefan Hetzl, and Graham E. Leigh;
licensed under Creative Commons License CC-BY

13th International Conference on Typed Lambda Calculi and Applications (TLCA'15).

Editor: Thorsten Altenkirch; pp. 1–16



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

$$\begin{array}{ccc}
\frac{A(\alpha)}{\forall x A} & t_1, t_2, \dots, t_k & \pi \vdash \exists x F \xrightarrow{c.e.} \pi' \vdash_{cf} \exists x F \\
\vdots & \vdots & \downarrow \text{definition} \quad \downarrow \text{Herbrand extraction} \\
\Gamma, \forall x A & \Delta, \exists x \bar{A} & \mathcal{G}_\pi \text{ (of size } \leq |\pi|) \xrightarrow{\mathcal{L}} \mathcal{L}(\mathcal{G}_\pi) \supseteq \mathcal{H}(\pi') \\
\Gamma, \Delta & \text{cut} & \mathcal{L}(\mathcal{G}_\pi): \text{ language of } \mathcal{G}_\pi \quad \mathcal{H}(\pi'): \text{ Herbrand set of } \pi'
\end{array}$$

■ **Figure 1** Π_1 cut.

■ **Figure 2** Proof grammars.

Grammars for proofs with only cuts of the form $\exists x A$ or $\forall x A$ with A quantifier free are remarkably simple. Let π be such a proof with end sequent $\exists x F$ where F is quantifier free. Suppose u_1, u_2, \dots, u_m are the witnesses to the existential quantifier as they appear in π . The induced grammar \mathcal{G}_π consists of the production rules i) $\sigma \rightarrow F(u_1) \mid F(u_2) \mid \dots \mid F(u_m)$ where σ is the starting symbol of the grammar, and ii) $\alpha \rightarrow t_1 \mid t_2 \mid \dots \mid t_k$ for every Π_1 cut in π of the form given in Figure 1, where α is the eigenvariable of the cut and t_1, t_2, \dots, t_k are witness terms of existential quantifier in the right subproof.

1.1 Contributions

In this paper we show how the correspondence between proofs and grammars can be extended to the level of Π_2 cuts. This class of cuts is particularly important for computational applications: a Π_2 formula can be read as a specification of a program and its proof as providing an (non-deterministic) algorithm in line with the Curry–Howard correspondence.

We consider Π_2 -proofs in which all cut formulæ have at most one quantifier of each sort i.e. of the form $\forall x \exists y A$ or $\forall x A$ with A quantifier free. It turns out that these proofs correspond to (rigid) context-free tree grammars:

► **Theorem 1.** *Let π be a Π_2 -proof of a Σ_1 formula F . There is an associated rigid context-free tree grammar \mathcal{G}_π (whose number of production rules is bounded by the size of π) with a finite language yielding a Herbrand disjunction for F .*

In fact, if we consider proofs in which every universal introduction rule is immediately followed by a cut or an existential introduction (henceforth called *simple proofs*) we have

► **Theorem 2.** *Let $\pi_0, \pi_1, \dots, \pi_k$ be a sequence of simple Π_2 -proofs such that π_{i+1} is obtained from π_i by an application of a cut reduction rule (as in Figure 4) to a subproof of π_i . For every $i \leq k$, $\mathcal{L}(\mathcal{G}_{\pi_0}) \supseteq \mathcal{L}(\mathcal{G}_{\pi_i})$. In particular, if the proof π_k contains only quantifier-free cuts $\mathcal{L}(\mathcal{G}_{\pi_0}) \supseteq \mathcal{H}(\pi_k)$.*

Rigid tree languages were first introduced in [21, 22] with applications to verification in mind. Later in [14] rigid regular grammars were defined by restricting the admissible derivations in the grammar with an equality constraint. In this paper we extend the notion of rigidity to context-free grammars. It plays an important role in obtaining a concise grammar for our proofs. For a Π_2 -proof π , we prove the language of the induced rigid context-free grammar \mathcal{G}_π has a bound double exponential in the size of π (see Theorem 10). This bound is optimal as it matches the blow up from cut elimination.

Structure of the article. In Section 2 we fix the calculus and cut reduction steps, and define the class of proofs under consideration. Section 3 develops the theory of rigid context-free tree grammars. In Section 4 we present the proof grammars induced by Π_2 -proofs: an

Axioms:	A, \bar{A} for A an atomic formula	
Inference rules:	$\frac{\Gamma, A, B}{\Gamma, A \vee B} \vee$	$\frac{\Gamma, A \quad \Delta, B}{\Gamma, \Delta, A \wedge B} \wedge$
	$\frac{\Gamma, A[x/\alpha]}{\Gamma, \forall x A} \forall$	$\frac{\Gamma, A[x/s]}{\Gamma, \exists x A} \exists$
	$\frac{\Gamma}{\Gamma, \Delta} w$	$\frac{\Gamma, \Delta, \Delta^*}{\Gamma, \Delta} c$
		$\frac{\Gamma, A \quad \Delta, \bar{A}}{\Gamma, \Delta} \text{cut}$

■ **Figure 3** Axioms and rules of sequent calculus.

elementary grammar is given in Section 4.1 to motivate the definition laid out in Section 4.2. Section 5 establishes the preservation of the language of our proof grammars under transitive closure of cut reduction steps. In Section 6 we conclude by describing future work and potential applications of our results and techniques.

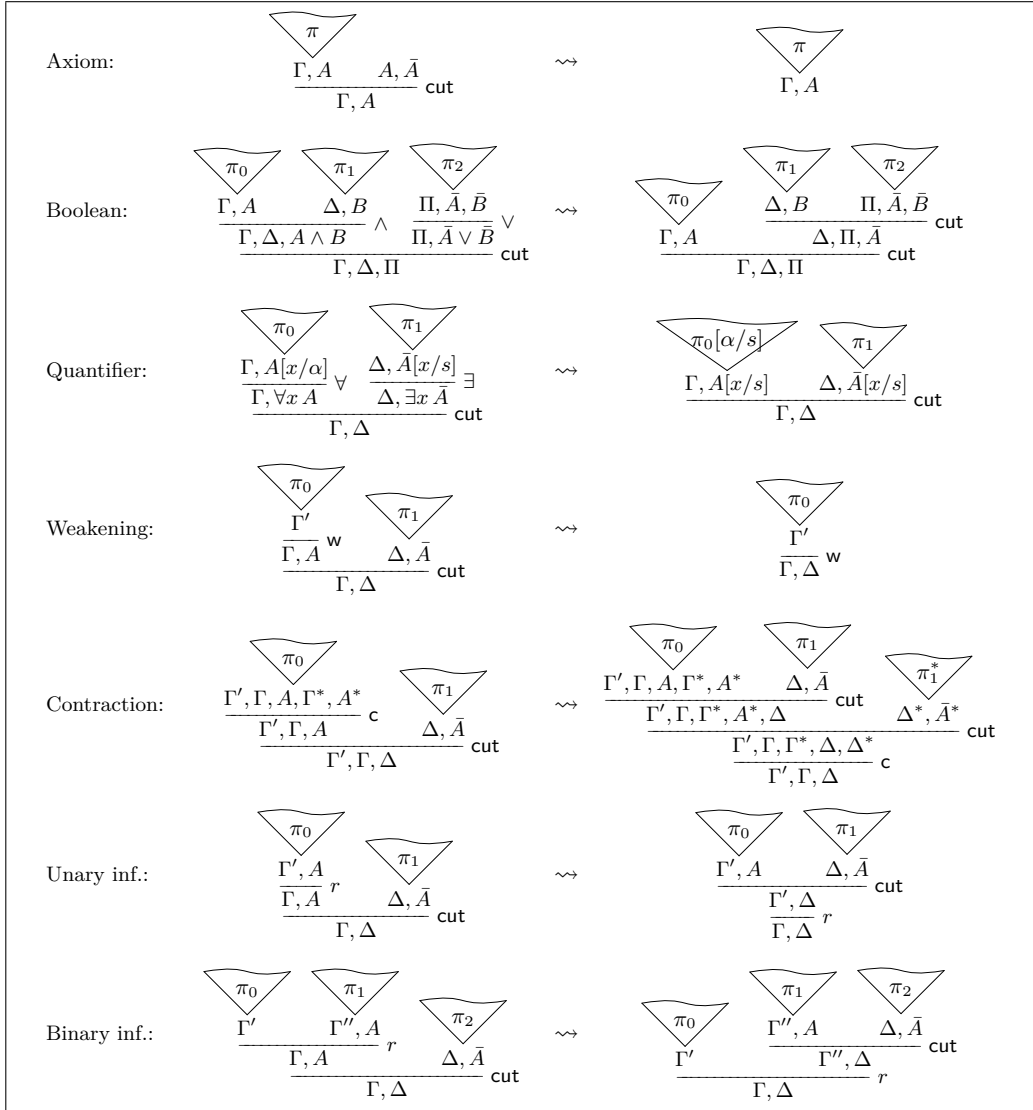
1.2 Related work

In [9] Gerhardy and Kohlenbach adapt Shoenfield’s variant of Gödel’s Dialectica interpretation to a system of pure predicate logic by explicitly adding decision-by-case constants to the target language. The resulting λ -term is first normalised and then used to directly read off a Herbrand disjunction. Heijltjes [10] and McKinley [24] study graphical formalisms of local reductions in classical first-order logic to derive a normal form corresponding to a cut-free proof from which a Herbrand disjunction is obtained. An approach similar to [10, 24] in the formalism of expansion trees [25] can be found in [19]. Historically, Hilbert’s ε -calculus [20] is the first formalism which allows a step-wise computation of a Herbrand disjunction in a way that abstracts from the propositional layer of predicate logic.

Like the aforementioned formalisms, the results presented in this paper allow the computation of a Herbrand disjunction in a way that bypasses literal cut elimination. The novelty of our approach lies in the fact that this is achieved via formal grammars. On the one hand this has the consequence that standard problems from formal language theory assume a proof-theoretic meaning and hence standard algorithms can be used to solve the corresponding proof-theoretic problems. For example, an algorithm for solving the membership problem for an adequate class of grammars also solves the following proof-theoretic problem (for the corresponding class of proofs): given a proof π and a term t , is t a witness obtainable by cut-elimination from π ? Usually, these ‘induced’ algorithms have a smaller asymptotic complexity (polynomial to at most exponential time; see e.g. [23, 22]) than the naive proof-theoretic algorithms which rely on explicitly computing the cut-free proof(s) (and need iterated-exponential time). On the other hand, the strong grip on the structure of a Herbrand disjunction afforded by a formal grammar opens the door to interesting theoretical and applied investigations (see Section 6).

2 Proof calculus

We work with a Tait-style (one-sided) sequent calculus for first-order logic with explicit weakening and contraction rules. A *proof* is a finite tree obtained from the axioms and rules laid out in Figure 3. We use capital Greek letters Γ, Δ , etc. for multisets of formulæ, \bar{A} to denote the dual of the formula A obtained by de Morgan laws, and $A[x/s]$ for the formula



■ **Figure 4** One-step cut reduction and permutation rules.

obtained from A by replacing x with the term s assuming this will not create any variable capture. Γ, Δ and Γ, A denote respectively the disjoint union of Γ, Δ and $\Gamma, \{A\}$.

In the (\forall) rule, α is called the *eigenvariable* and must not appear free in $\Gamma, \forall x A$; in (\exists) rule the term s is assumed to be free for x ; and in the contraction rule (c) , the set Δ^* denotes a renaming of Δ . Those formulæ which are explicitly mentioned in the premise of an inference rule are said to be *principal* in the rule applied, for example A and B are principal in (\wedge) rule, every formula from Δ^* is principal in (c) , and there are no principal formulæ in the weakening rule (w) .

We assume all proofs are *regular* namely, all quantifiers' eigenvariables are distinct and different from any free variables. We use the following naming convention for proofs throughout the paper: lower-case Greek letters α, β, γ , etc. represent eigenvariables; x, y for bound variable symbols; and π, π' , etc. for proofs. $\pi[\alpha/s]$ is the result of replacing throughout the proof π each occurrence of the variable symbol α by the term s . We write $\pi \vdash \Gamma$ to express

π is a proof of Γ and $\text{EV}(\pi)$ to denote the set of eigenvariables in π . For a position p in (the tree) π , $\pi|p$ denotes the subproof of π at position p with the convention that $\pi|\langle \rangle = \pi$, and $\pi|p0$ and $\pi|p1$ are respectively the left (or only) subproof and right subproof of $\pi|p$. We write $q < p$ if q is a proper prefix of p .

It is useful to isolate a particular class of Π_2 -proofs:

► **Definition 3** (Simple Π_2 -proof). A *simple Π_2 -proof* is a proof in which i) each sequent is a finite set of prenex Π_2 formulæ, ii) the conclusion is a set of closed Π_1 and Σ_1 formulæ, iii) cut formulæ feature at most one quantifier of each sort, and iv) every universally quantified formula appearing above a cut is principal in the inference directly after its introduction (either a cut or existential introduction).

Condition (ii) above is stipulated primarily for expository purposes. Note that every sequent can be transformed into a sequent containing only Σ_1 formulæ by Skolemization and prenexification. In contrast, it is impossible to Skolemize cut formulæ in a proof since the two (dual) occurrences of the cut formula would yield dual Skolemizations. Also note that condition (iv) does not restrict provability, as any proof satisfying (i)–(iii) can be modified using simple rule permutations to also satisfy (iv). This transformation will not increase the size (number of inference rules) of the proof.

Cut reduction and Herbrand sets

The standard cut reduction rules are given in Figure 4. For proofs π, π' we write $\pi \rightsquigarrow \pi'$ if π' is the result of applying one of the rules to π (and not to a subproof). Notice that in contraction reduction a subproof is duplicated and care is taken to rename the eigenvariables (expressed by annotating the proof/sequent/formula in question by an asterisk) to maintain regularity.

Let π be a cut-free proof of Γ and suppose $A \in \Gamma$ has the form $\exists x_1 \cdots \exists x_k B$ with B quantifier free. If $\bigvee_{(s_1, \dots, s_k) \in X} B[x_1/s_1, \dots, x_k/s_k]$ is the Herbrand disjunction for A read off from π , we call X the *Herbrand set of A in π* and define $\mathcal{H}(\pi, A) = X$. In addition, $\mathcal{H}(\pi) := \{\{A\} \times \mathcal{H}(\pi, A) \mid A \in \Gamma \cap \Sigma_1\}$.

Running example

We consider a formal proof of the pigeonhole principle for two boxes via the infinite pigeonhole principle. The question of the computational content of this proof is attributed to G. Stolzenberg in [6]. A variety of analytic methods have since been applied to this proof [11, 33, 3, 2] and its generalisations [29, 28]. Despite its relatively short and symmetric nature it allows us to adequately demonstrate grammars for proofs with Π_2 cuts.

Let $f: \mathbb{N} \rightarrow \{0, 1\}$ be a total Boolean function, let I_i express that there are infinitely many $m \in \mathbb{N}$ for which $f(m) = i$ and T express that there exists $m < n$ such that $f(m) = f(n)$. A consequence of the law of excluded middle is $I_0 \vee I_1$. Moreover I_i implies T for each i : assuming I_i , there exists $m \geq 0$ and $n \geq m + 1$ for which $f(m) = f(n) = i$. Combining these observations we conclude that T holds.

The following formalises the above argument. The formal language, Σ , consist of two unary function symbols \mathbf{f} , \mathbf{s} , one binary function symbol \mathbf{M} , a constant symbol $\mathbf{0}$ and a binary relation \leq . We make the following definitions and abbreviations:

- $T = \exists x \exists y (x < y \wedge \mathbf{f}x = \mathbf{f}y)$,
- $I_i = \forall x \exists y (x \leq y \wedge \mathbf{f}y = \underline{i})$ for $i \in \{0, 1\}$ where $\underline{0} = \mathbf{0}$ and $\underline{1} = \mathbf{s0}$,
- $\Gamma = \{\forall x \forall y (x \leq \mathbf{M}xy \wedge y \leq \mathbf{M}xy), \forall x (\mathbf{f}x = \mathbf{0} \vee \mathbf{f}x = \mathbf{s0})\}$,

$$\boxed{
\frac{
\frac{
\frac{
\frac{
\Gamma, I_0^{\alpha, M\alpha\hat{\alpha}}, I_1^{\hat{\alpha}, M\alpha\hat{\alpha}}
}{\Gamma, I_0^{\alpha}, I_1^{\hat{\alpha}}} \exists^*
}{\Gamma, I_0^{\alpha}, I_1} \forall
}{\Gamma, \Delta, T, I_0^{\alpha}} \forall
}{\Gamma, \Delta, T, I_0} \forall
}{
\frac{
\frac{
\Delta, T_{\hat{\beta}, \hat{\beta}'}, \bar{I}_1^{0, \hat{\beta}}, \bar{I}_1^{s\hat{\beta}, \hat{\beta}'}
}{\Delta, T, \bar{I}_1} \exists^*
}{\Gamma, \Delta, T, I_0^{\alpha}} \text{cut}
}{\Gamma, \Delta, T} \text{cut,c}
}
}
}
}
}$$

■ **Figure 5** Simple Π_2 -proof π_∞ of pigeonhole principle. The inference rules labelled \exists^* represent finitely many existential introduction rules (the order of which is unimportant).

- $\Delta = \{\overline{\forall x \forall y \forall z (x = y \wedge z = y \rightarrow x = z)}, \overline{\forall x \forall y (sx \leq y \rightarrow x < y)}\}$,
- I_i^s and $I_i^{s,t}$ denote, respectively, $\exists y (s \leq y \wedge fy = i)$ and $(s \leq t \wedge ft = i)$,
- $T_{s,t}$ denotes $(s < t \wedge fs = ft)$.

The intended interpretation of the symbols is: f represents the (arbitrary) function f , s the successor function on \mathbb{N} , \leq the standard ordering and M the binary max function.

A formal proof of the pigeonhole principle is given by the simple Π_2 -proof in Figure 5 which we name π_∞ . For brevity only eigenvariables and witnesses of the quantifiers and instances of the existential formula T are displayed in π_∞ . The proof uses about 50 application of the axioms and rules of the calculus but the only cuts in π_∞ are the two Π_2 cuts shown in the figure. Two normal forms of the proof (of size ~ 200) have been computed in a case study [33] from which one can read off the Herbrand set $\{(0, 1), (1, 2), (2, 3), (0, 2), (1, 3)\}$ (up to interpretation of the logical symbols) relative to the formula T . Once we have introduced our grammars we shall see how this Herbrand set can be directly computed from π_∞ .

3 Context-free tree grammars

In [14] and elsewhere a refinement of regular tree grammars was studied that mimics the construction of terms appearing in cut-elimination for first-order logic with Π_1 cuts. These grammars were called *rigid* tree grammars and are equivalent to the notion of rigid automata introduced and explored in [21, 22]. In this section we provide a generalisation to the class of context-free tree grammars corresponding to Π_2 -proofs.

Given a ranked alphabet Σ , we let $\text{Terms}(\Sigma)$ denote the set of terms in the simply-typed λ -calculus built from Σ . For a $T \in \text{Terms}(\Sigma)$ we write $\text{Pos}(T)$ for the set of positions in term (tree) T . For $p \in \text{Pos}(T)$, $T|p$ is the subterm of T at position p .

A *context-free tree grammar* (CFG) is a tuple $\mathcal{G} = \langle N, \Sigma, \sigma, \text{Pr} \rangle$ where N is a set of typed *non-terminals* of order at most 1, $\sigma \in N$ is a designated *start symbol* (of base-type ι), Σ is a ranked alphabet, called *terminals*, disjoint from N , and Pr consists of pairs (a, T) (called *production rules* and written $a \rightarrow T$) where $a \in N$ and $T \in \text{Terms}(\Sigma \cup N)$ that has the same type as a . Given a CFG \mathcal{G} we assume $\mathcal{G} = \langle N_{\mathcal{G}}, \Sigma_{\mathcal{G}}, \sigma_{\mathcal{G}}, \text{Pr}_{\mathcal{G}} \rangle$. If the set $N_{\mathcal{G}}$ of non-terminals contains only symbols of order 0, \mathcal{G} is a *regular tree grammar*.

Let d be a sequence $\langle \rho_i, p_i \rangle_{i < k}$ of pairs of production rules of a CFG \mathcal{G} and positions, and S and T terms. We call d a *derivation from S to T* , written $d: S \rightarrow T$, if there exist terms $(N_i)_{i \leq k}$ such that $N_0 = S$, $N_k = T$, and for each $0 \leq i < k$,

1. ρ_i is a production rule of \mathcal{G} and $p_i \in \text{Pos}(N_i)$,
2. For $\rho_i = (a \rightarrow S)$, we have $N_i|p_i = a$ and $N_{i+1} = N_i[p_i/S]$, namely the result of replacing the sub-term of N_i at position p_i by S (renaming bound variables if necessary).

The sequence of terms $(N_i)_{i \leq k}$ is uniquely determined by d and S , whence we may write $d(i)$ for N_i . The *length* of d , $\text{lh}(d)$, is k . We say T is *derivable from S* if there exists a derivation $d: S \rightarrow T$. \mathcal{G} is *acyclic* if for every non-terminal $a \in N_{\mathcal{G}}$ and every derivation $d: a \rightarrow T$ with $\text{lh}(d) > 0$, a does not appear in T .

The *language* of a CFG \mathcal{G} is defined as $\mathcal{L}(\mathcal{G}) = \{T \in \text{Terms}(\Sigma_{\mathcal{G}}) \mid \exists d: \sigma_{\mathcal{G}} \rightarrow T\}$. When comparing languages of CFGs it is convenient to work modulo β -convertibility. Thus for grammars \mathcal{G}, \mathcal{H} , we write $\mathcal{L}(\mathcal{G}) \subseteq \mathcal{L}(\mathcal{H})$ to express that for every $S \in \mathcal{L}(\mathcal{G})$ there exists a β -equivalent $T \in \mathcal{L}(\mathcal{H})$.

3.1 Rigidity

Let \mathcal{G} be a CFG, suppose \triangleleft is a transitive binary relation on $N_{\mathcal{G}}$, and R is a designated set of non-terminals. A derivation $d = \langle (a_i \rightarrow S_i), p_i \rangle_{i < \text{lh}(d)}: S \rightarrow T$ induces a natural equivalence relation on the positions in T corresponding to connectedness in parse trees: for $j_0, j_1 < \text{lh}(d)$, let $j_0 \sim_d j_1$ iff there exist $i_0 < j_0, j_1$ such that

1. $p_{i_0} \leq p_{j_0}, p_{j_1}$,
2. $a_{j_0} = a_{j_1} \in R$,
3. for every $k \in \{0, 1\}$ and $i_0 < i < j_k < \text{lh}(d)$, if $p_i \leq p_{j_k}$ then $a_{j_k} \not\triangleleft a_i$.

In other words, two occurrences of a non-terminal $a \in R$ in (the natural tree representation of) the derivation d are considered connected if there is no non-terminal of higher priority between them and their closest common ancestor.

Of particular interest to us is the class of derivations that respect their own \sim relation.

► **Definition 4.** Let \mathcal{G} be a CFG and suppose \triangleleft is a transitive ordering on $N_{\mathcal{G}}$ and $R \subseteq N_{\mathcal{G}}$. A derivation $d = \langle \rho_i, p_i \rangle_{i < k}: S \rightarrow T$ in \mathcal{G} is *rigid with respect to (\triangleleft, R)* if for every $i, j < k$, $i \sim_d j$ implies $T|p_i = T|p_j$.

A *rigid context-free tree grammar* is a structure $\mathcal{G} = \langle N_{\mathcal{G}}, R_{\mathcal{G}}, \triangleleft_{\mathcal{G}}, \Sigma_{\mathcal{G}}, \sigma_{\mathcal{G}}, \text{Pr}_{\mathcal{G}} \rangle$ such that $\langle N_{\mathcal{G}}, \Sigma_{\mathcal{G}}, \sigma_{\mathcal{G}}, \text{Pr}_{\mathcal{G}} \rangle$ is a CFG, $R_{\mathcal{G}} \subseteq N_{\mathcal{G}}$ and $\triangleleft_{\mathcal{G}}$ is a transitive order on $N_{\mathcal{G}}$. $R_{\mathcal{G}}$ is the set of *rigid* non-terminals of \mathcal{G} and $\triangleleft_{\mathcal{G}}$ is the *priority ordering* of \mathcal{G} . If $R_{\mathcal{G}} = N_{\mathcal{G}}$ then \mathcal{G} is called *totally rigid*.

Given a rigid CFG \mathcal{G} and a derivation d in its underlying CFG, we call d *rigid* if it is rigid with respect to $(\triangleleft_{\mathcal{G}}, R_{\mathcal{G}})$. The *language* of a rigid CFG \mathcal{G} , $\mathcal{L}(\mathcal{G})$, is the collection of terms derivable from rigid derivations starting from $\sigma_{\mathcal{G}}$:

$$\mathcal{L}(\mathcal{G}) = \{T \in \text{Terms}(\Sigma_{\mathcal{G}}) \mid \exists d: \sigma_{\mathcal{G}} \rightarrow T \text{ and } d \text{ is } (\triangleleft_{\mathcal{G}}, R_{\mathcal{G}})\text{-rigid}\}.$$

Note the language of a rigid CFG is a set of closed (well-typed) base-type λ -terms.

► **Example 5.** Let \mathcal{G} be the rigid CFG with start symbol σ , non-terminals σ, α and γ , rigid non-terminals R , ordering \triangleleft , terminal symbols of appropriate arity, and production rules $\sigma \rightarrow f(\alpha, \gamma, \gamma)$, $\gamma \rightarrow g(\gamma) \mid a$, and $\alpha \rightarrow \gamma$. If $\gamma \notin R$ we have, unsurprisingly, $\mathcal{L}(\mathcal{G}) = \{f(g^m(a), g^n(a), g^o(a)) \mid m, n, o \geq 0\}$. Otherwise,

1. if $\gamma \triangleleft \gamma \triangleleft \alpha$, $\mathcal{L}(\mathcal{G}) = \{f(g^m(a), g^n(a), g^n(a)) \mid m, n \geq 0\}$;
2. if $\gamma \triangleleft \gamma \not\triangleleft \alpha$, $\mathcal{L}(\mathcal{G}) = \{f(g^m(a), g^m(a), g^m(a)) \mid m \geq 0\}$;
3. if $\gamma \not\triangleleft \gamma \not\triangleleft \alpha$, $\mathcal{L}(\mathcal{G}) = \{f(a, a, a)\}$.

In contrast to the grammar presented in Example 5, acyclic tree grammars must have a finite language. If the grammar is also totally rigid its language has size essentially double exponential in the size of the grammar.

► **Lemma 6.** *If \mathcal{G} is an acyclic totally rigid CFG then $|\mathcal{L}(\mathcal{G})| \leq |\text{Pr}_{\mathcal{G}}|^{2^{|N_{\mathcal{G}}|-1}}$.*

Proof. Suppose \mathcal{G} satisfies the above requirements; let $m = |\text{Pr}_{\mathcal{G}}|$ and $n = |N_{\mathcal{G}}| - 1$. Since \mathcal{G} is acyclic we may further assume that $\triangleleft_{\mathcal{G}}$ is irreflexive. We argue, by induction on n , that the set $\{T \in \text{Terms}(\Sigma_{\mathcal{G}}) \mid \exists d: \sigma_{\mathcal{G}} \rightarrow T \text{ and } d \text{ is rigid}\}$ has size bounded by m^{2^n} .

The base case is $n = 0$. By the main assumption of the lemma every derivation $d: \sigma_{\mathcal{G}} \rightarrow T \in \text{Terms}(\Sigma_{\mathcal{G}})$ has length 1, of which there are no more than m . For the induction step, suppose $n = n_0 + 1$. Let $N = N_{\mathcal{G}} \setminus \{a\}$ where $a \neq \sigma_{\mathcal{G}}$ is chosen such that $a \not\triangleleft_{\mathcal{G}} b$ for all $b \in N_{\mathcal{G}} \setminus \{\sigma_{\mathcal{G}}\}$. Suppose $d: \sigma_{\mathcal{G}} \rightarrow T \in \text{Terms}(\Sigma_{\mathcal{G}})$ is a rigid derivation in \mathcal{G} . Since \mathcal{G} is acyclic d can be re-ordered to have the form $d_0 d_1$ where $d_0: \sigma_{\mathcal{G}} \rightarrow S$ and $d_1: S \rightarrow T = S[a/S']$ for appropriate terms S and S' , such that the non-terminal a is not rewritten in d_0 and not introduced by a production rule in d_1 . The induction hypothesis implies there is no more than $m^{2^{n_0}}$ possibilities for each of S and S' , whence there are $\leq m^{2^n}$ possibilities for T . \blacktriangleleft

4 Proof grammars

In this section we present rigid context-free grammars that arise from simple Π_2 -proofs. We first define *elementary proof grammars* which are a natural extension of the rigid regular grammars arising from Π_1 -proofs introduced in [14]. As we shall see, elementary proof grammars can have an infinite language and are not immediately suitable for producing Herbrand disjunctions. A proper form of *proof grammars* is then defined in Section 4.2 by adding further structure.

4.1 Elementary proof grammars

Let $\pi \vdash \Gamma$ be a simple Π_2 -proof and Γ a set of closed Σ_1 and Π_1 formulæ. The *elementary grammar for π* is a rigid CFG, denoted \mathcal{E}_{π} , of the form $\langle N_{\pi}, R_{\pi}, \triangleleft_{\pi}, \Sigma_{\pi}, \sigma, \text{Pr}_{\pi} \rangle$ where

- N_{π} consists of the eigenvariables appearing in π (of base-type) as well as a *starting symbol* σ and a symbol κ^p (of function type $\iota \rightarrow \iota$) for each position p in π at which the rule cut is applied;
- $R_{\pi} = \text{EV}(\pi)$;
- Σ_{π} is the term language of first-order logic expanded by
 - a symbol $\tau_{i,F}$ of base-type for each formula $\forall x_0 \cdots x_k F \in \Gamma \cap \Pi_1$ and each $i \leq k$,
 - a symbol F of function type with k arguments for each $\exists x_0 \cdots \exists x_k F \in \Gamma \cap \Sigma_1$;
- \triangleleft_{π} is the transitive ordering on non-terminals and Pr_{π} the set of production rules specified below.

Each cut occurring in π , as well as each quantified formula in the conclusion Γ , yields production rules in \mathcal{E}_{π} . In addition, the relative order of the quantifier introduction rules and each cut on a genuine Π_2 formula influence the rigidity ordering \triangleleft_{π} . We begin by specifying the rules induced by Γ .

For each formula $\exists x_0 \cdots \exists x_k F \in \Gamma$ with F quantifier-free, and each sequence of terms $\langle s_i \rangle_{i \leq k}$ that appear in π (together) as the witnesses of the sequence of existential quantifiers $\langle \exists x_i \rangle_{i \leq k}$, \mathcal{E}_{π} features a production rule $\sigma \rightarrow F s_0 \cdots s_k$. For each formula $\forall x_0 \cdots \forall x_k F \in \Gamma$ with F quantifier-free and each $i \leq k$, \mathcal{E}_{π} contains the production rule $\alpha_i \rightarrow \tau_{i,F}$ where α_i is the (unique) eigenvariable for the quantifier $\forall x_i$ appearing in π . See Figure 6.

The remaining non-terminals attain their production rules based on the structure of the cut in which they are used. Let p be a position in π , A a quantifier-free formula, α, δ eigenvariables and s, t terms. Suppose $\pi|p$ has the form of either cuts given in Figure 7. In the proof on the left the two distinguished appearances of the formula $\exists y A[x/\alpha]$ as well as

$\pi \vdash \Gamma$	$\frac{\frac{\frac{\Pi', F[x_0/s_0, \dots, x_k/s_k]}{\Pi', \exists x_k F[x_0/s_0, \dots, x_{k-1}/s_{k-1}]} \exists \quad \vdots}{\Gamma = \Pi, \exists x_0 \dots \exists x_k F}}{\Gamma = \Pi, \forall x_0 \dots \forall x_k F} \forall$	$\frac{\frac{\frac{\Pi', \forall x_{i+1} \dots \forall x_k F[x_0/\alpha_0, \dots, x_i/\alpha_i]}{\Pi', \forall x_i \dots \forall x_k F[x_0/\alpha_0, \dots, x_{i-1}/\alpha_{i-1}]} \forall \quad \vdots}{\Gamma = \Pi, \forall x_0 \dots \forall x_k F} \forall$
\mathcal{E}_π	$\sigma \rightarrow \mathbf{F}s_0 \dots s_k$	$\alpha_i \rightarrow \tau_{i,F}$
\mathcal{G}_π	$\sigma \rightarrow \mathbf{F}s_0^\sigma \dots s_k^\sigma$	$\alpha_i \rightarrow \tau_{i,F}$

■ **Figure 6** Start and terminal production rules in \mathcal{E}_π and \mathcal{G}_π .

Subproof πp	$\frac{\frac{\frac{\Pi'', A[x/\alpha, y/s]}{\Pi'', \exists y A[x/\alpha]} \exists \quad \vdots}{\Pi', \exists y A[x/\alpha]} \forall \quad \frac{\frac{\frac{\Delta', \bar{A}[x/t, y/\beta]}{\vdash \Delta', \forall y \bar{A}[x/t] (\dagger)} \exists \quad \vdots}{\Delta', \exists x \forall y \bar{A}} \forall \quad \frac{\frac{\frac{\Delta', \bar{A}[x/t]}{\Delta', \exists x \bar{A}} \exists \quad \vdots}{\Pi', A[x/\alpha]} \forall \quad \frac{\frac{\Pi', \forall x A}{\Pi', \forall x \exists y A} \forall \quad \frac{\Delta, \exists x \forall y \bar{A}}{\Delta, \exists x \forall y \bar{A}} \text{cut}}{\Pi, \Delta} \text{cut}}{\Pi, \Delta} \text{cut}$	$\frac{\frac{\frac{\Delta', \bar{A}[x/t]}{\Delta', \exists x \bar{A}} \exists \quad \vdots}{\Pi', A[x/\alpha]} \forall \quad \frac{\frac{\Pi', \forall x A}{\Pi', \forall x A} \forall \quad \frac{\Delta, \exists x \bar{A}}{\Delta, \exists x \bar{A}} \text{cut}}{\Pi, \Delta} \text{cut}}{\Pi, \Delta} \text{cut}$
Rules in \mathcal{E}_π	$\alpha \rightarrow t \quad \beta \rightarrow \kappa^p t \quad \kappa^p \rightarrow \lambda \alpha s$	$\alpha \rightarrow t$
Rules in \mathcal{G}_π	$\alpha \rightarrow \lambda x_1 \dots \lambda x_{k_\alpha} t^\alpha$ $\beta \rightarrow \lambda x_1 \dots \lambda x_{k_\beta} \cdot \kappa^p x_1 \dots x_{k_\alpha} t^\beta$ $\kappa^p \rightarrow \lambda x_1 \dots \lambda x_{k_{\alpha+1}} s^\alpha$	$\alpha \rightarrow \lambda x_1 \dots \lambda x_{k_\alpha} t^\alpha$

■ **Figure 7** Internal production rules in \mathcal{E}_π and \mathcal{G}_π .

the two distinguished occurrences of $\exists x \forall y \bar{A}$ are assumed to be on the same trace, as are, in the right proof, the two occurrences of $\exists x \bar{A}$. The grammar includes the production rules

$$\alpha \rightarrow t \quad \beta \rightarrow \kappa^p t \quad \kappa^p \rightarrow \lambda \alpha s$$

and we set $a \triangleleft_\pi \alpha \triangleleft_\pi \kappa^p \triangleleft_\pi b$ for each $a \in \text{EV}(\pi|q_\alpha)$ and $b \in \text{EV}(\pi|q_\beta)$ where q_α and q_β are, respectively, the positions in π at which the variables α and β were eliminated (that is p_0 and the position marked by (\dagger) in Figure 7). In the scenario pictured on the right in which the cut is performed on a Π_1 formula, only a single production rule for α is needed.

► **Example 7** (Elementary proof grammar for π_∞). Let $\mathcal{E}_{\pi_\infty} = \langle N_\infty, R_\infty, \triangleleft_\infty, \Sigma_\infty, \sigma, \text{Pr}_\infty \rangle$. The non-terminals of \mathcal{E}_{π_∞} comprise: starting symbol σ ; rigid non-terminals: $\alpha, \hat{\alpha}, \beta, \beta', \hat{\beta}$ and $\hat{\beta}'$; non-rigid non-terminals: κ^0 for the topmost cut on I_1 and $\kappa (= \kappa^\diamond)$ for the lower cut on I_0 . In Σ_∞ we have the term symbols and also terminal symbols for each formula in the conclusion (note there are no universals in the conclusion of π_∞). The induced priority ordering is $\hat{\alpha} \triangleleft_\infty \kappa^0 \triangleleft_\infty \{\hat{\beta}', \hat{\beta}\} \triangleleft_\infty \alpha \triangleleft_\infty \kappa^\diamond \triangleleft_\infty \{\beta', \beta\}$.

In this example we will focus only on the highlighted formula T and the corresponding symbol T of type $\iota \rightarrow \iota \rightarrow \iota$ in Σ_∞ . Pr_∞ has the following rules. The production rules on the left (right) hand side are read from the cut on I_0 (I_1).

$$\begin{array}{ll} \sigma \rightarrow \mathsf{T}\beta\beta' & \sigma \rightarrow \mathsf{T}\hat{\beta}\hat{\beta}' \\ \alpha \rightarrow 0 \mid \mathsf{s}\beta & \beta \rightarrow \kappa 0 \quad \hat{\alpha} \rightarrow 0 \mid \mathsf{s}\hat{\beta} \quad \hat{\beta} \rightarrow \kappa^0 0 \\ \kappa \rightarrow \lambda \alpha. \mathbf{M}\alpha\hat{\alpha} & \beta' \rightarrow \kappa(\mathsf{s}\beta) \quad \kappa^0 \rightarrow \lambda \hat{\alpha}. \mathbf{M}\alpha\hat{\alpha} \quad \hat{\beta}' \rightarrow \kappa^0(\mathsf{s}\hat{\beta}) \end{array}$$

In general, the production rules of the elementary proof grammar \mathcal{E}_π may be cyclic and therefore permit infinite derivations. In the case of \mathcal{E}_{π_∞} , for example, this is demonstrated by the sequence of non-terminals $\beta, \kappa, \hat{\alpha}, \hat{\beta}, \kappa^0, \alpha, \beta$.

To avoid enumerating unnecessary terms into the language of the grammar certain derivations should be disallowed. While it is possible to provide a characterisation of the derivations that yield Herbrand sets, the work is beyond the scope of this paper and will not be presented here. Instead, in the following section, we define acyclic proof grammars for which standard (rigid) derivations suffice. This is achieved by raising the types of non-terminals to make the necessary dependencies between non-terminals explicit.

4.2 Typed proof grammars

As before, let $\pi \vdash \Gamma$ be a simple Π_2 -proof and Γ a set of closed prenex Σ_1 and Π_1 formulæ. The (*proof*) *grammar* for π is a rigid CFG denoted by \mathcal{G}_π . $\mathcal{G}_\pi = \langle N_\pi, R_\pi, \triangleleft_\pi, \Sigma_\pi, \sigma, \text{Pr}_\pi \rangle$ has the same definition as the elementary grammar but with two essential differences:

1. Eigenvariables are no longer necessarily base-type symbols and their type depends on their relationship to other eigenvariables;
2. Production rules are modified accordingly by type-raising operations.

Types of non-terminals in \mathcal{G}_π

The type of a given non-terminal will be determined by the relation of its position in π to other eigenvariables in π . This relation will be defined as a well-founded ordering \prec_π on the elements of N_π and formalises the (potential) dependency of one non-terminal on another. As well as fixing the type of a non-terminal, the ordering can be seen as the basis of the priority ordering \triangleleft_π used in recognising rigid derivations.

For each position p in π , if $\pi|p$ has the form

$$\frac{\frac{\Pi', A[x/\alpha]}{\Pi', \forall x A} \forall \quad \Delta, \exists x \bar{A}}{\pi|p \vdash \Pi, \Delta} \text{cut}$$

we set $\alpha \prec_\pi \kappa^p$ and $\alpha \prec_\pi a$ for every $a \in \text{EV}(\pi|p0) \cup \{\kappa^q \mid p0 \leq q\}$. Notice that for $a, b \in \text{EV}(\pi)$, $a \prec_\pi b$ implies $b \triangleleft_\pi a$, and if we write \prec_a for the set $\{b \in N_\pi \mid b \prec_\pi a\}$, \prec_a is a set of eigenvariables linearly ordered by \prec_π . Moreover, for α appearing as above $\prec_{\kappa^p} = \prec_a \cup \{\alpha\}$.

The type of a non-terminal is chosen to be its order-type in \prec_π . Let $k_a = |\prec_a|$. The *type* of $a \in N_\pi$ is that of a function over the base-type taking k_a arguments. So in particular σ is of base-type, as is any eigenvariable relating to a universal quantifier in the conclusion.

► **Example 8** (Types of non-terminals in π_∞). In \mathcal{G}_{π_∞} we have the same set of non-terminals symbols as in \mathcal{E}_{π_∞} but they are now assigned the following types. The ordering \prec_∞ on N_∞ gives $\prec_\infty = \{(\alpha, \hat{\alpha}), (\hat{\alpha}, \kappa^0), (\alpha, \kappa^0), (\alpha, \hat{\beta}), (\alpha, \hat{\beta}'), (\alpha, \kappa^{\hat{\iota}})\}$ so α, β, β' and σ all have base-type, $\hat{\alpha}, \hat{\beta}, \hat{\beta}'$ and $\kappa^{\hat{\iota}}$ have type $\iota \rightarrow \iota$ and κ^0 has type $\iota \rightarrow \iota \rightarrow \iota$. Note, the ordering \triangleleft_∞ specifying rigidity is unchanged from \mathcal{E}_{π_∞} .

Production rules in \mathcal{G}_π

As mentioned earlier, the production rules of \mathcal{G}_π have the form as in \mathcal{E}_π . We now explain how the change in the type of non-terminals is to be taken into account. This is achieved by means of a type-lifting operation that either lifts an occurrence of a non-terminal to its appropriate type or replaces it with a variable symbol in case it should be abstracted.

For each $\gamma \in \text{EV}(\pi)$, let $\gamma_1 \prec_\pi \cdots \prec_\pi \gamma_{k_\gamma+1} = \gamma$ be the sequence of eigenvariables enumerating the elements of $\{\xi \mid \xi \preceq_\pi \gamma\}$. Also let $\{x_i\}$ be a fresh set of variable symbols.

Given $\alpha \in \text{EV}(\pi) \cup \{\sigma\}$, we define an operation $S \mapsto S^\alpha$ on terms in $\Sigma_\pi \cup N_\pi$ that lifts the occurrence of non-terminals to their required type: we set $(ST)^\alpha = S^\alpha T^\alpha$, $(\lambda y.S)^\alpha = \lambda y.S^\alpha$, $c^\alpha = c$ for $c \in \Sigma_\pi$, and for $\gamma \in \text{EV}(\pi) \cup \{\kappa^q \mid q \text{ a position in } \pi\}$,

$$\gamma^\alpha = \begin{cases} x_i, & \text{if } \gamma = \alpha_i \preceq_\pi \alpha, \\ \gamma\gamma_1^\alpha\gamma_2^\alpha \cdots \gamma_{k_\gamma}^\alpha, & \text{otherwise.} \end{cases}$$

The operation $S \mapsto S^\alpha$ replaces each non-terminal $\alpha_i \preceq_\pi \alpha$ by the variable x_i and to $\gamma \not\preceq_\pi \alpha$ makes explicit the dependence of γ on \prec_γ . For example, we have

$$\begin{aligned} \gamma^{\gamma_{k_\gamma}} &= \gamma x_1 x_2 \cdots x_{k_\gamma} \\ \gamma^\sigma &= \gamma\gamma_1(\gamma_2\gamma_1)(\gamma_3\gamma_1(\gamma_2\gamma_1)) \cdots (\gamma_{k_\gamma}\gamma_1(\gamma_2\gamma_1) \cdots (\gamma_{k_\gamma-1}\gamma_1 \cdots (\gamma_{k_\gamma-2}\gamma_1 \cdots))) \end{aligned}$$

Let p be a position in π , suppose $\pi|p$ has the form of either cuts in Figure 7. Then the grammar \mathcal{G}_π includes the production rules

$$\alpha \rightarrow \lambda x_1 \cdots \lambda x_{k_\alpha} t^\alpha \quad \beta \rightarrow \lambda x_1 \cdots \lambda x_{k_\beta} \cdot \kappa^p x_1 \cdots x_{k_\alpha} t^\beta \quad \kappa^p \rightarrow \lambda x_1 \cdots \lambda x_{k_\alpha+1} s^\alpha$$

Observe that $k_\alpha \leq k_\beta$, so $\alpha_i = \beta_i$ for each $0 < i \leq k_\alpha$ and the term in the central production rule is indeed closed. Also for α and κ^p above notice $k_{\kappa^p} = k_\alpha + 1$, hence the production rule for κ^p is well-typed. In the scenario pictured on the right side of Figure 7 in which the cut is performed on a Π_1 formula, naturally only the production rule for α is required.

Concerning the start symbol and formulæ in the conclusion Γ we add, in analogy to \mathcal{E}_π , production rules $\sigma \rightarrow \text{Fs}_0^\sigma \cdots s_k^\sigma$ and $\alpha_i \rightarrow \tau_{i,F}$ for appropriate formulæ F , terms $\langle s_i \rangle_{i \leq k}$ and eigenvariable α_i . This completes the definition of \mathcal{G}_π .

► **Example 9** (Proof grammar for π_∞). It is now possible to complete the definition of \mathcal{G}_{π_∞} . As in Example 7 we will focus on parts of the grammar that are relevant to the formula T and the computation of its language, which we denote $\mathcal{L}(\mathcal{G}_\infty, T)$. The non-terminals and their types were described in the previous example. Notice for instance $\xi^\sigma = \xi$ for $\xi \in \{\alpha, \beta, \beta'\}$, $\xi^\sigma = \xi\alpha$ and $\xi^\alpha = \xi x_1$ for $\xi \in \{\kappa, \hat{\alpha}, \hat{\beta}, \hat{\beta}'\}$, and $(M\alpha\hat{\alpha})^{\hat{\alpha}} = Mx_1x_2$. The production rules of the grammar (relating to the formula T) are therefore

$$\begin{aligned} \sigma &\rightarrow \text{T}\beta\beta' & \sigma &\rightarrow \text{T}(\hat{\beta}\alpha)(\hat{\beta}'\alpha) \\ \alpha &\rightarrow 0 \mid \mathfrak{s}\beta & \beta &\rightarrow \kappa 0 & \hat{\alpha} &\rightarrow \lambda x. 0 \mid \lambda x. \mathfrak{s}(\hat{\beta}x) & \hat{\beta} &\rightarrow \lambda x. \kappa^0 x 0 \\ \kappa &\rightarrow \lambda x. Mx(\hat{\alpha}x) & \beta' &\rightarrow \kappa(\mathfrak{s}\beta) & \kappa^0 &\rightarrow \lambda x \lambda y. Mxy & \hat{\beta}' &\rightarrow \lambda x. \kappa^0 x(\mathfrak{s}(\hat{\beta}x)) \end{aligned}$$

The computation of the language of $\mathcal{L}(\mathcal{G}_\infty, T)$ is not complicated. Nevertheless, for space considerations it is necessary to make a few simplifications. In particular, in accordance with the informal proof presented in Section 2, various terms will be evaluated according to the intended semantics, so $\mathfrak{s}S$ will be written as $S + 1$ and MST will be replaced by $\max\{S, T\}$. Also, $\text{T}ST$ will be presented as $\text{T}(S, T)$. In addition, implicit β -conversion of terms will be performed as this has no effect on rigidity of the considered derivations.

We begin by calculating the terms derivable from $\hat{\beta}$ and $\hat{\beta}'$ (with implicit β -conversion):

$$\hat{\beta}x \rightarrow \kappa^0 x 0 \rightarrow Mx 0 = x \quad \hat{\beta}'x \rightarrow^* Mx(\mathfrak{s}(\hat{\beta}x)) \rightarrow^* Mx(\mathfrak{s}(Mx 0)) = x + 1$$

Regarding $\hat{\alpha}$, modulo β -conversion it is easy to see we have $\hat{\alpha}x \rightarrow^* 0 \mid x + 1$. Thus we can also compute the derivations starting from β , β' and α :

$$\begin{aligned} \beta &\rightarrow^* M0(\hat{\alpha}0) \rightarrow^* 0 \mid 1 & \alpha &\rightarrow^* 0 \mid 1 \mid 2 \\ \beta' &\rightarrow^* M(\mathfrak{s}\beta)(\hat{\alpha}(\mathfrak{s}\beta)) \rightarrow \beta + 1 \mid \beta + 2 \end{aligned}$$

Thus $\mathsf{T}(\beta, \beta+1)$ and $\mathsf{T}(\beta, \beta+2)$ are the two terms derivable from $\mathsf{T}(\beta, \beta')$ without rewriting β . Combining these with the terms obtained from β above yields a total of eight derivations in the underlying non-rigid grammar. However, as $\kappa, \beta', \hat{\alpha}$ are the sole non-terminals used in deriving β and $\beta \not\sim \kappa, \beta', \hat{\alpha}$, only four of the derivations are rigid, leaving

$$\sigma \rightarrow \mathsf{T}(\beta, \beta') \rightarrow^* \mathsf{T}(0, 1) \mid \mathsf{T}(1, 2) \mid \mathsf{T}(0, 2) \mid \mathsf{T}(1, 3)$$

Concerning the terms derivable from σ via $\mathsf{T}(\hat{\beta}\alpha, \hat{\beta}'\alpha)$, rigid derivations yield

$$\sigma \rightarrow \mathsf{T}(\hat{\beta}\alpha, \hat{\beta}'\alpha) \rightarrow^* \mathsf{T}(0, 1) \mid \mathsf{T}(1, 2) \mid \mathsf{T}(2, 3)$$

Thus we conclude $\mathcal{L}(\mathcal{G}_\infty, T) = \{\mathsf{T}(0, 1), \mathsf{T}(1, 2), \mathsf{T}(2, 3), \mathsf{T}(0, 2), \mathsf{T}(1, 3)\}$. The reader may check that $\Gamma \cup \Delta \cup \{T_{m,n} \mid \mathsf{T}(m, n) \in \mathcal{L}(\mathcal{G}_\infty, T)\}$ is derivable.

► **Theorem 10** (Language bound). *Let π be a simple Π_2 -proof. The number of production rules in \mathcal{G}_π is bounded by the number of quantifier inferences in π and $|\mathcal{L}(\mathcal{G}_\pi)| < 2^{2^{|\pi|}}$.*

Proof. Let $\mathcal{G}' = \langle N_\pi, N_\pi, \triangleleft_\pi, \Sigma_\pi, \sigma, \text{Pr}_\pi \rangle$ be the modification of \mathcal{G}_π in which all non-terminals are marked as rigid. We observe $\mathcal{L}(\mathcal{G}') = \mathcal{L}(\mathcal{G}_\pi)$. Moreover, a study of paths in π reveals that \mathcal{G}' is acyclic, whereby $|\mathcal{L}(\mathcal{G}_\pi)| \leq |\text{Pr}_\pi|^{2^{|\pi|}}$ by Lemma 6. Let $k = \lfloor \frac{|\text{Pr}_\pi|}{2} \rfloor$. Then $N_\pi + k < |\pi|$ and so $|\mathcal{L}(\mathcal{G}_\pi)| \leq 2^{2^{N_\pi+k}} < 2^{2^{|\pi|}}$ as required. ◀

5 Language containment

► **Lemma 11** (Local reduction). *If $\pi \rightsquigarrow \pi'$ is a local one-step cut reduction between regular simple Π_2 -proofs then $\mathcal{L}(\mathcal{G}_{\pi'}) \subseteq \mathcal{L}(\mathcal{G}_\pi)$.*

Proof. We present the argument for two of the interesting cases, the case of Contraction Reduction and Quantifier Reduction; the remaining cases follow by a simple argument mirroring that of the former case.

Suppose, to begin, that the reduction $\pi \rightsquigarrow \pi'$ is an instance of Contraction Reduction and that the cut formula is principal in the contraction. Thus π and π' can be assumed to take the form given below, where $A = \forall xB$ is Π_2 and π_0^* is a renaming of π_0 so that π' is regular.

$$\frac{\frac{\frac{\Gamma, A}{\pi_0} \quad \frac{\frac{\Delta', \Delta, \bar{A}, \Delta^*, \bar{A}^*}{\pi_1} \text{ c}}{\Delta', \Delta, \bar{A}} \text{ c}}{\pi \vdash \Gamma, \Delta', \Delta} \text{ cut}}{\Gamma, A} \quad \frac{\frac{\frac{\Gamma, A}{\pi_0^*} \quad \frac{\frac{\Delta', \Delta, \bar{A}, \Delta^*, \bar{A}^*}{\pi_1} \text{ c}}{\Gamma, \Delta', \Delta, \Delta^*, \bar{A}^*} \text{ cut}}{\Gamma, \Gamma^*, \Delta', \Delta, \Delta^*} \text{ c}}{\pi' \vdash \Gamma, \Delta', \Delta} \text{ cut}}{\Gamma, \Delta', \Delta, \Delta^*, \bar{A}^*} \text{ cut}$$

Note that we may assume the contraction occurs in the ‘right’ sub-proof as π is simple and A is a universal formula.

We argue that every rigid derivation in $\mathcal{G}_{\pi'}$ starting from σ can be transformed into a rigid derivation in \mathcal{G}_π beginning at σ . Consider the function $f: N_{\pi'} \rightarrow N_\pi$ defined by

$$\begin{aligned} f(\sigma) &= \sigma & f(\kappa^0) &= f(\kappa^{01}) = \kappa^{\langle \rangle} \\ f(\gamma) &= f(\gamma^*) = \gamma \quad \text{for } \gamma \in \text{EV}(\pi_0) & f(\kappa^{00p}) &= f(\kappa^{010p}) = \kappa^{0p} \\ f(\delta) &= \delta \quad \text{for } \delta \in \text{EV}(\pi_1) & f(\kappa^{011p}) &= \kappa^{10p} \end{aligned}$$

We observe that for all $a, b \in N_{\pi'}$ we have $a \prec_{\pi'} b$ iff $f(a) \prec_{\pi} f(b)$, thus f is type-preserving and uniquely extends to a function mapping terms in the language of $\Sigma_{\pi'} \cup N_{\pi'}$ to (well-typed) terms in $\Sigma_{\pi} \cup N_{\pi}$. Moreover, if $a \rightarrow S$ is a production rule in $\mathcal{G}_{\pi'}$, $f(a) \rightarrow f(S)$ is a production rule in \mathcal{G}_{π} . So f transforms derivations in the former grammar to derivations in the latter grammar; all that remains is to check the operation preserves rigidity.

Rigidity is not immediate as f is not injective. Indeed, let $d = \langle (a_i \rightarrow S_i), p_i \rangle_{i < \text{lh}(d)} : \sigma \rightarrow S$ be a rigid derivation in $\mathcal{G}_{\pi'}$, let $d^f : \sigma \rightarrow f(S)$ be the derivation in \mathcal{G}_{π} induced by f and suppose $j_0, j_1 < \text{lh}(d)$ are such that $j_0 \not\prec_d j_1$ but $j_0 \sim_{d^f} j_1$, so $f(a_{j_0}) = f(a_{j_1}) \in R_{\mathcal{G}_{\pi}}$. Since f preserves the priority ordering, it follows that $a_{j_0} \neq a_{j_1}$ and so we may assume $a_{j_0} \in \text{EV}(\pi_0)$ and $a_{j_1} \in \text{EV}(\pi_0^*)$. But then d must utilise a production rule for a rigid non-terminal $\delta \in \text{EV}(\pi_1)$ at a position q such that i) $q < p_{j_0}$ iff $q \not< p_{j_1}$ and ii) either $a_{j_0} \triangleleft \delta$ or $a_{j_1} \triangleleft \delta$, contradicting $j_0 \sim_{d^f} j_1$.

Before proceeding with the second case, we remark that in all the other local reduction steps, the natural choice of the ‘renaming’ function f is injective and preservation of rigidity is immediate.

Suppose now $\pi \rightsquigarrow \pi'$ is an instance of Quantifier Reduction. We may assume π and π' have the form below.

$$\frac{\frac{\frac{\pi_0}{\Gamma, A[x/\alpha]} \quad \frac{\pi_1}{\Delta, \bar{A}[x/s]}}{\Gamma, \forall x A} \vee \quad \frac{\Delta, \bar{A}[x/s]}{\Delta, \exists x \bar{A}} \exists}{\pi \vdash \Gamma, \Delta} \text{cut} \qquad \frac{\frac{\pi_0[\alpha/s]}{\Gamma, A[x/s]} \quad \frac{\pi_1}{\Delta, \bar{A}[x/s]}}{\pi' \vdash \Gamma, \Delta} \text{cut}$$

As in the previous case we define a function mapping rigid derivations in $\mathcal{G}_{\pi'}$ to rigid derivations in \mathcal{G}_{π} . Although every rigid non-terminal of \mathcal{G}_{π} is a non-terminal in $\mathcal{G}_{\pi'}$, the non-terminals arising from $\text{EV}(\pi_0)$ have a higher type than their counterpart in π' as we have $\alpha \prec_{\pi} \xi$ for every $\xi \in \text{EV}(\pi_0)$. Notice, however, $k_{\alpha} = 0$ and $k_{\kappa(\cdot)} = 1$ in \mathcal{G}_{π} .

Let $d : \sigma \rightarrow S \in \mathcal{L}(\mathcal{G}_{\pi'})$ be a rigid derivation in $\mathcal{G}_{\pi'}$. As the formula $A[x/s]$ is Σ_1 it follows that d has one of following two forms.

In the first case, d is (up to simple renaming of non-terminals) a derivation in $\pi_0[\alpha/s]$ and $S \in \mathcal{L}(\mathcal{G}_{\pi_0[\alpha/s]})$. d induces a rigid derivation in \mathcal{G}_{π_0} which when augmented by the production rule $\alpha \rightarrow s$ (present in \mathcal{G}_{π}) provides a derivation of S in \mathcal{G}_{π} .

In the second case, (a permutation of) d has the form $d_0 d_1 d_2$ where $d_0 : \sigma \rightarrow S'$ is a derivation in \mathcal{G}_{π_1} ; $d_1 : S' \rightarrow S'[\beta/t^\beta]$ is a derivation using the single production rule $\beta \rightarrow t^\beta$ (note $k_{\beta} = 0$ in both $\mathcal{G}_{\pi'}$ and \mathcal{G}_{π}) where $\beta \in \text{EV}(\pi_1)$ is the unique eigenvariable for the universal quantifier in $\bar{A}[x/s]$ (if there is one) in $\mathcal{G}_{\pi'}$; and $d_2 : S'[\beta/t^\beta] \rightarrow S$ is a derivation in $\mathcal{G}_{\pi_0[\alpha/s]}$. We observe that the production rule $\beta \rightarrow t^\beta$ becomes, in \mathcal{G}_{π} , the derivation $\beta \rightarrow \kappa(\cdot) s \rightarrow (\lambda x_1 u^\alpha) s$ where $u[\alpha/s] = t$. Using these derivations in place of d_1 and making the appropriate modifications to the derivation d_2 yields a derivation $e : \sigma \rightarrow S''$ in \mathcal{G}_{π} such that S'' β -reduces to S . In both cases rigidity is also easily checked. \blacktriangleleft

Using the previous lemma it is not difficult to establish that the language of proof grammars respect also non-local cut reductions, provided simplicity is maintained.

► **Theorem 12 (Global Reduction).** *Suppose π and π' are regular simple Π_2 -proofs such that π differs from π' only in its subproof at position p . If $\pi'|p \rightsquigarrow \pi|p$ then $\mathcal{L}(\mathcal{G}_{\pi}) \subseteq \mathcal{L}(\mathcal{G}_{\pi'})$.*

We now re-state and prove our main results from the introduction. The first of these is a restatement of Theorem 2 and follows from Theorem 12; the second is a generalisation of Theorem 1.

► **Theorem 13.** Let $\langle \pi \rangle_{i \leq k}$ be a sequence of simple Π_2 -proofs such that for each $i < k$, π_{i+1} is obtained by application of one of the reduction rules of Figure 4 to a sub-proof of π_i . Then for every term $T \in \mathcal{L}(\mathcal{G}_{\pi_k})$ there exists a term $S \in \mathcal{L}(\mathcal{G}_{\pi_0})$ that β -reduces to T .

► **Theorem 14.** Let $\pi \vdash \Gamma, \Delta$ be a Π_2 -proof where $\Gamma \subseteq \Pi_1$ and $\Delta \subseteq \Sigma_1$ are sets of prenex formulæ and suppose $|\pi|$ denotes the number of inference rules occurring in π . There exists a totally rigid acyclic context-free tree grammar \mathcal{G} such that i) $|\text{Pr}_{\mathcal{G}}| \leq |\pi|$, ii) $|\mathcal{L}(\mathcal{G})| \leq 2^{2^{|\pi|}}$ and iii) there is a quantifier-free form of Γ, Γ' , such that the formula

$$\bigvee \Gamma' \vee \bigvee \{F[x_1/s_1, \dots, x_k/s_k] \mid (\exists x_1 \cdots \exists x_k F) \in \Delta \wedge \mathbf{F}s_1 \cdots s_k \in \mathcal{L}(\mathcal{G})\} \quad (1)$$

is a tautology.

Proof. By applying quantifiers inversion if necessary, π can be turned into a simple Π_2 -proof π' without an increase in size. Let $\mathcal{G} = \langle N_{\pi'}, N_{\pi'}, \triangleleft_{\pi'}, \Sigma_{\pi'}, \sigma, \text{Pr}_{\pi'} \rangle$ be the totally rigid grammar derived from $\mathcal{G}_{\pi'}$. Items (i) and (ii) follow from Theorem 10. Regarding (iii), let $\langle \pi_i \rangle_{i \leq k}$ be a reduction sequence of simple Π_2 -proofs (for example any obtained from the standard cut elimination algorithms of [31, 32]) starting from $\pi_0 = \pi'$ and leading to a cut-free proof π_k of Γ, Δ . By the previous theorem, $\mathcal{L}(\mathcal{G}_{\pi_k}) \subseteq \mathcal{L}(\mathcal{G}_{\pi})$.

Suppose Γ and Δ have the forms $\forall x_1 \cdots \forall x_{k_0} G_0, \dots, \forall x_1 \cdots \forall x_{k_m} G_m$ and $\exists x_1 \cdots \exists x_{l_0} F_0, \dots, \exists x_1 \cdots \exists x_{l_n} F_n$ respectively where $G_0, \dots, G_m, F_0, \dots, F_n$ are quantifier-free. The Herbrand disjunction read from π_k is the formula

$$X = \bigvee_{j \leq m} G_j(\alpha_j^1, \dots, \alpha_j^{k_j}) \vee \bigvee_{j \leq n} \bigvee_{(s_1, \dots, s_{l_j}) \in \mathcal{H}(\pi_k, F_j)} F_j(s_1, \dots, s_{l_j})$$

for appropriate choice of variables α_j^i . Since the formula in (1) is the result of replacing every α_j^i by τ_{i, G_j} in X we are done. ◀

6 Conclusion

This work provides an abstraction of proofs which focuses only on the aspects relevant to the extraction of Herbrand disjunctions. Compared to other approaches in the literature, including Herbrand nets [24], proof forests [10], expansion trees with cut [19] and functional interpretation [9], proof grammars offer a representation of Herbrand's theorem suitable for the following exploitation.

As carried out in [18], the result for Π_1 -proofs can be strengthened to the following: let π' be any cut-free proof obtained from a Π_1 -proof π via standard cut elimination, then $\mathcal{H}(\pi') \subseteq \mathcal{L}(\mathcal{G}_{\pi})$. Moreover, if π' is obtained from π by non-erasing reduction (which corresponds to the λI -calculus, see [4, Section 9]) then we even have $\mathcal{H}(\pi') = \mathcal{L}(\mathcal{G}_{\pi})$. Therefore all (possibly infinitely many) normal forms of the non-erasing reduction have the same Herbrand disjunction. This property of classical logic has been called *Herbrand-confluence* in [18] and provides a general way of defining the computational content of a classical proof in the sense that no witness is ruled out by a choice of reduction strategy. A deeper analysis of Π_2 -proofs carried out in [1] has also yielded a Herbrand-confluence result analogous to [18].

A notable application of proof grammars is in *cut introduction*, which is motivated by the aim to structure and compress automatically generated analytic proofs. As shown in [17, 16], the arrows of Figure 2 can be inverted in the sense that a grammar can be computed from a Herbrand disjunction and that from such a grammar one can compute cut formulæ which realise the compression of the grammar. This method has been implemented and empirically

evaluated with good results in [15]. An extension of these techniques to the case of proofs with Π_1 -induction has led to a new technique for inductive theorem proving [8]. A natural continuation of the present work is to find an analogous characterisation for proofs with Π_2 -induction as well as the development of techniques for the systematic introduction of Π_2 cuts.

An application of proof grammars to proof complexity consists in proving a lower bound on the length of proofs with cuts (which is notoriously difficult to control) by transferring a lower bound on the size of the corresponding grammar. This has been carried out for Π_1 cuts in [7] based on [14], and can potentially be extended to Π_2 cuts based on the results of this paper.

Acknowledgements. The authors' research was supported by the Wiener Wissenschafts-, Forschungs- und Technologiefonds (WWTF) project VRG12-04. The authors also wish to thank the anonymous referees for their helpful suggestions.

References

- 1 Bahareh Afshari, Stefan Hetzl, and Graham E. Leigh. On Herbrand confluence for first-order logic. Submitted, 2015.
- 2 Matthias Baaz, Stefan Hetzl, Alexander Leitsch, Clemens Richter, and Hendrik Spohr. Cut-Elimination: Experiments with CERES. In Franz Baader and Andrei Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning (LPAR) 2004*, volume 3452 of *Lecture Notes in Computer Science*, pages 481–495. Springer, 2005.
- 3 Franco Barbanera, Stefano Berardi, and Massimo Schivalocchi. “Classical” programming-with-proofs in λ_{PA}^{Sym} : An analysis of non-confluence. In Martín Abadi and Takayasu Ito, editors, *Theoretical Aspects of Computer Software*, volume 1281 of *Lecture Notes in Computer Science*, pages 365–390. Springer Berlin Heidelberg, 1997.
- 4 Hendrik Pieter Barendregt. *The Lambda Calculus*, volume 103 of *Studies in Logic and the Foundations of Mathematics*. Elsevier, 1984.
- 5 Samuel R. Buss. On Herbrand’s Theorem. In *Logic and Computational Complexity*, volume 960 of *Lecture Notes in Computer Science*, pages 195–209. Springer, 1995.
- 6 Thierry Coquand. A semantics of evidence for classical arithmetic. *Journal of Symbolic Logic*, 60(1):325–337, 1995.
- 7 Sebastian Eberhard and Stefan Hetzl. Compressibility of finite languages by grammars. preprint, available at <http://www.logic.at/people/hetzl/research/>, 2015.
- 8 Sebastian Eberhard and Stefan Hetzl. Inductive theorem proving based on tree grammars. to appear in the *Annals of Pure and Applied Logic*, preprint available at <http://www.logic.at/people/hetzl/research/>, 2015.
- 9 Philipp Gerhardy and Ulrich Kohlenbach. Extracting Herbrand Disjunctions by Functional Interpretation. *Archive for Mathematical Logic*, 44:633–644, 2005.
- 10 Willem Heijltjes. Classical proof forestry. *Annals of Pure and Applied Logic*, 161(11):1346–1366, 2010.
- 11 Hugo Herbelin. *Séquents qu’on calcule*. PhD thesis, Université Paris 7, 1995.
- 12 Jacques Herbrand. *Recherches sur la théorie de la démonstration*. PhD thesis, Université de Paris, 1930.
- 13 Stefan Hetzl. On the form of witness terms. *Archive for Mathematical Logic*, 49(5):529–554, 2010.
- 14 Stefan Hetzl. Applying Tree Languages in Proof Theory. In Adrian-Horia Dediu and Carlos Martín-Vide, editors, *Language and Automata Theory and Applications (LATA) 2012*, volume 7183 of *Lecture Notes in Computer Science*, pages 301–312. Springer, 2012.

- 15 Stefan Hetzl, Alexander Leitsch, Giselle Reis, Janos Tapolczai, and Daniel Weller. Introducing Quantified Cuts in Logic with Equality. In Stéphane Demri, Deepak Kapur, and Christoph Weidenbach, editors, *Automated Reasoning - 7th International Joint Conference, IJCAR*, volume 8562 of *Lecture Notes in Computer Science*, pages 240–254. Springer, 2014.
- 16 Stefan Hetzl, Alexander Leitsch, Giselle Reis, and Daniel Weller. Algorithmic introduction of quantified cuts. *Theoretical Computer Science*, 549:1–16, 2014.
- 17 Stefan Hetzl, Alexander Leitsch, and Daniel Weller. Towards Algorithmic Cut-Introduction. In *Logic for Programming, Artificial Intelligence and Reasoning (LPAR-18)*, volume 7180 of *Lecture Notes in Computer Science*, pages 228–242. Springer, 2012.
- 18 Stefan Hetzl and Lutz Straßburger. Herbrand-Confluence for Cut-Elimination in Classical First-Order Logic. In Patrick Cégielski and Arnaud Durand, editors, *Computer Science Logic (CSL) 2012*, volume 16 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 320–334. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2012.
- 19 Stefan Hetzl and Daniel Weller. Expansion trees with cut. preprint, available at <http://arxiv.org/abs/1308.0428>, 2013.
- 20 David Hilbert and Paul Bernays. *Grundlagen der Mathematik II*. Springer, 1939.
- 21 Florent Jacquemard, Francis Klay, and Camille Vacher. Rigid tree automata. In Adrian Horia Dediu, Armand-Mihai Ionescu, and Carlos Martín-Vide, editors, *Language and Automata Theory and Applications (LATA) 2009*, volume 5457 of *Lecture Notes in Computer Science*, pages 446–457. Springer, 2009.
- 22 Florent Jacquemard, Francis Klay, and Camille Vacher. Rigid tree automata and applications. *Information and Computation*, 209:486–512, 2011.
- 23 Markus Lohrey. Algorithmics on SLP-compressed strings: A survey. *Groups Complexity Cryptology*, 4(2):241–299, 2012.
- 24 Richard McKinley. Proof nets for Herbrand’s Theorem. *ACM Transactions on Computational Logic*, 14(1):5:1–5:31, 2013.
- 25 Dale Miller. A Compact Representation of Proofs. *Studia Logica*, 46(4):347–370, 1987.
- 26 V.P. Orevkov. Lower bounds for increasing complexity of derivations after cut elimination. *Zapiski Nauchnykh Seminarov Leningradskogo Otdeleniya Matematicheskogo Instituta*, 88:137–161, 1979.
- 27 Pavel Pudlák. The Lengths of Proofs. In Sam Buss, editor, *Handbook of Proof Theory*, pages 547–637. Elsevier, 1998.
- 28 Diana Ratiu and Trifon Trifonov. Exploring the Computational Content of the Infinite Pigeonhole Principle. *Journal of Logic and Computation*, 22(2):329–350, 2012.
- 29 Monika Seisenberger. *On the Constructive Content of Proofs*. PhD thesis, Ludwig-Maximilians-Universität München, 2003.
- 30 Richard Statman. Lower bounds on Herbrand’s theorem. *Proceedings of the American Mathematical Society*, 75:104–107, 1979.
- 31 G. Takeuti. *Proof Theory*. Dover Books on Mathematics Series. Dover Publications, Incorporated, 2013.
- 32 Anne S. Troelstra and Helmut Schwichtenberg. *Basic Proof Theory*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2000.
- 33 Christian Urban. *Classical Logic and Computation*. PhD thesis, University of Cambridge, 2000.

Non-Wellfounded Trees in Homotopy Type Theory*

Benedikt Ahrens¹, Paolo Capriotti², and Régis Spadotti¹

1 Institut de Recherche en Informatique de Toulouse
Université Paul Sabatier, Toulouse, France
{ahrens,spadotti}@irit.fr

2 Functional Programming Laboratory, School of Computer Science
University of Nottingham, UK
pvc@cs.nott.ac.uk

Abstract

We prove a conjecture about the constructibility of *coinductive types* – in the principled form of *indexed M-types* – in Homotopy Type Theory. The conjecture says that in the presence of *inductive types*, coinductive types are derivable. Indeed, in this work, we construct coinductive types in a subsystem of Homotopy Type Theory; this subsystem is given by Intensional Martin-Löf type theory with natural numbers and Voevodsky’s Univalence Axiom. Our results are mechanized in the computer proof assistant AGDA.

1998 ACM Subject Classification F.3.2 Semantics of Programming Languages

Keywords and phrases Homotopy Type Theory, coinductive types, computer theorem proving, Agda

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.17

1 Introduction


Coinductive data types are used in functional programming to represent infinite data structures. Examples include the ubiquitous data type of streams over a given base type, but also more sophisticated types; as an example we present an alternative definition of equivalence of types (Example 22).

From a categorical perspective, coinductive types are characterized by a *universal property*, which specifies the object with that property *uniquely* in a suitable sense. More precisely, a coinductive type is specified as the *terminal coalgebra* of a suitable endofunctor. In this category-theoretic viewpoint, coinductive types are dual to *inductive types*, which are defined as initial algebras.

Inductive, resp. coinductive, types are usually considered in the principled form of the family of *W-types*, resp. *M-types*, parametrized by a type A and a dependent type family B over A , that is, a family of types $(B(a))_{a:A}$. Intuitively, the elements of the coinductive type $M(A, B)$ are trees with nodes labeled by elements of A such that a node labeled by $a : A$ has $B(a)$ -many subtrees, given by a map $B(a) \rightarrow M(A, B)$; see Figure 1 for an example. The *inductive* type $W(A, B)$ contains only trees where any path within that tree eventually leads to a *leaf*, that is, to a node $a : A$ such that $B(a)$ is empty.

* The work of Benedikt Ahrens was partially supported by the CIMI (Centre International de Mathématiques et d’Informatique) Excellence program ANR-11-LABX-0040-CIMI within the program ANR-11-IDEX-0002-02 during a postdoctoral fellowship.





■ **Figure 1** Example of a tree (adapted from [14]).

In this work, we study coinductive types in Homotopy Type Theory (HoTT), an extension of intensional Martin-Löf type theory [11]; we give a brief overview in Section 2.

The universal properties defining inductive and coinductive types, respectively, can be expressed internally to intensional Martin-Löf type theory (and thus internally to HoTT). Awodey, Gambino, and Sojakova [6] use this facility when proving, within a subtheory \mathcal{H} of HoTT, a logical equivalence between

1. the existence of W -types (a.k.a. the existence of a universal object) and
2. the addition of a set of type-theoretic rules to their “base theory” \mathcal{H} .

We might call the W -types defined internally “internal W -types”, and those specified via type-theoretic rules “external” ones. In that sense, Awodey, Gambino, and Sojakova [6] prove a logical equivalence between the existence of internal and external W -types.

The universal property defining (internal) coinductive types in HoTT is dual to the one defining (internal) inductive types. One might hence assume that their existence is equivalent to a set of type-theoretic rules dual (in a suitable sense) to those given for external W -types as in Item 2 above. However, the rules for external W -types cannot be dualized in a naïve way, due to some asymmetry of HoTT related to dependent types as maps into a “type of types” (a *universe*), see the discussion in [10].

In this work, we show instead that coinductive types in the form of M -types can be derived from certain inductive types. (More precisely, only one specific W -type is needed: the type of natural numbers, which is readily specified as a W -type [6].)

The result presented in this work is not surprising; indeed, the constructibility of coinductive types from inductive types has been shown in extensional type theory (see Section 1.1) and was conjectured to work in HoTT during a discussion on the HoTT mailing list [10]. In this work, we give a formal proof of the constructibility of a class of coinductive types from inductive types, with a proof of correctness of the construction.

The theorem we prove here is actually more general than described above: instead of plain M -types as described above, we construct *indexed* M -types, which can be considered as a form of “(simply-)typed” trees, typed over a type of indices I . Plain M -types then correspond to the mono-typed indexed M -types, that is, to those for which $I = 1$. Since all the ideas are already contained in the case of plain M -types, we describe the construction of those extensively, and only briefly state the definitions and the main result for the indexed case. The formalisation in AGDA, however, is done for the more general, indexed, case. An example illustrates the need for these more general *indexed* M -types.

1.1 Related work

Inductive types in the form of W -types in HoTT have been studied by Awodey, Gambino, and Sojakova [6]. The content of that work is described above.

Van den Berg and De Marchi [14] study the existence of *plain* M -types in models of *extensional* type theory, that is, of type theory with a reflection rule identifying propositional and judgmental equality. They prove the derivability of M -types from W -types in such models, see Corollary 2.5 of the arXiv version of that article. A construction in extensional type theory of M -types from W -types is given by Abbott, Altenkirch, and Ghani [1].

Martin-Löf type theory without identity reflection, but with the principle of *Uniqueness of Identity Proofs* (Axiom K) can be identified with the 0-truncated fragment of HoTT (modulo the assumption of univalence and HITs). For such a type theory, a construction of (indexed) M -types from W -types is described by Altenkirch et al. [4], internalizing a standard result in 1-category theory [7]. The present work thus generalizes the construction described in [4] by extending it from the 0-truncated fragment to the whole of HoTT. More specifically, the main work in this generalization is to develop higher-categorical variants of the 1-categorical constructions used in [4] that are compatible with the higher-categorical structure (the coherence data) of types.

1.2 Synopsis

The paper is organized as follows: In Section 2 we present the type theory we are working in – a “subsystem” of HoTT as presented in [13]. In Section 3 we define signatures for *plain* M -types and, via a universal property, the M -type associated to a given signature. In Section 4 we construct the M -type of a given signature. In Section 5 we state the main result for the case of *indexed* M -types. Finally, in Section 6 we give an overview of the formalisation of our result in the proof assistant AGDA.

2 The type theory under consideration

The present work takes place within a type theory that is a subsystem of the type theory presented in the HoTT book [13]. The latter is often referred to as Homotopy Type Theory (HoTT); it is an extension of intensional Martin-Löf type theory (IMLTT) [11]. The extension is given by two data: firstly, the *Univalence Axiom*, introduced by Vladimir Voevodsky and proven consistent with IMLTT in the simplicial set model [9]. The second extension is given by *Higher Inductive Types* (HITs), the precise theory of which is still subject to active research. Preliminary results on HITs have been worked out by Sojakova [12] and Lumsdaine and Shulman – see [13, Chap. 6] for an introduction. In the present work, we use the Univalence Axiom, but do not make use of HITs.

The syntax of HoTT is extensively described in a book [13]; we only give a brief summary of the type constructors used in the present work, thus fixing notation. The fundamental objects are *types*, which have *elements* (“inhabitants”), written $a : A$. Types can be dependent on terms, which we write as $x : A \vdash B(x) : \mathcal{U}$. In the preceding judgment, we use a special type \mathcal{U} , the “universe” or “type of types”. In this work we assume any type being an element of \mathcal{U} in the sense of “typical ambiguity” [13, Chap. 1.3], without worrying about universe levels. The formalisation in AGDA ensures that everything works fine in that respect: as we will see later, the universe \mathcal{U} is closed under the construction of M -types.

We use the following type constructors: dependent products $\prod_{(x:A)} B(x)$, with non-dependent variant written $A \rightarrow B$, dependent sums $\sum_{(x:A)} B(x)$ with non-dependent variant

written $A \times B$, the identity type $x =_A y$ and the coproduct type $A + B$. In particular, we assume the empty type 0 and the singleton type 1 . Furthermore, we assume the existence of a type of natural numbers, given as an inductive type according to the rules given in [13, Chap. 1.9]. Finally, we assume the univalence axiom for the universe \mathcal{U} as presented in [13, Chap. 2.10].

Concerning terms, function application is denoted by parentheses as in $f(x)$ or, occasionally, simply by juxtaposition. We write dependent pairs as (a, b) for $b : B(a)$. Projections are indicated by a subscript, that is, for $x : \sum_{(a:A)} B(a)$ we have $x_0 : A$ and $x_1 : B(x_0)$. Indices are also used occasionally to specify earlier arguments of a function of several arguments; e.g., we write $B_i(a)$ instead of $B(i)(a)$.

We conclude this brief introduction by recalling two important internally definable properties of types: we call the type X **contractible**, if X is inhabited by a unique element, that is, if the following type is inhabited:

$$\text{isContr}(X) := \sum_{(x:X)} \prod_{(x':X)} x' = x .$$

We call the type Y a **proposition** if for all $y, y' : Y$, we have $y = y'$. Note that a type X is contractible iff X is a proposition and there is an element $x : X$ (see also [13, Lemma 3.11.3]).

3 Definition of M-types via universal property

Coinductive types represent *potentially infinite* data structures, such as streams or infinite lists. As such, they have to be contrasted to *inductive* datatypes, which represent structures that are *necessarily finite*, such as unary natural numbers or finite lists.

3.1 Signatures, a.k.a. containers

In order to analyze inductive and coinductive types systematically, one usually fixes a notion of “signature”: a signature specifies, in an abstract way, the rules according to which the instances of a data structure are built. In the following, we consider signatures to be given by “containers” [4]:

► **Definition 1.** A **container** (or **signature**) is a pair (A, B) of a type A and a dependent type $x : A \vdash B(x) : \mathcal{U}$ over A .

The container (A, B) then determines a type of “trees” built as follows: such a tree consists of a *root node*, labeled by an element $a : A$, and a family of trees – “subtrees” of the original tree – indexed by the type $B(a)$. A tree is *well-founded* if it does not have an infinite chain of subtrees.

To the container (A, B) one associates two types of trees built according to those rules: the type $W(A, B)$ of well-founded trees, and the type $M(A, B)$ of all trees, i.e., not necessarily well-founded.

The description of the inhabitants of $W(A, B)$ and $M(A, B)$ in terms of trees gives a suitable intuition; formally, those types are defined in terms of a *universal property*. Indeed, $M(A, B)$ will be defined as (the carrier of) a terminal object in a suitable sense.

3.2 Coalgebras for a signature

Any container (A, B) specifies an endomorphism on types as follows:

► **Definition 2.** Given a container (A, B) , define the **polynomial functor** $P : \mathcal{U} \rightarrow \mathcal{U}$ associated to (A, B) as

$$P(X) := P_{A,B}(X) := \sum_{a:A} (B(a) \rightarrow X) .$$

Given a map $f : X \rightarrow Y$, define $Pf : PX \rightarrow PY$ as the map

$$Pf(a, g) := (a, f \circ g) .$$

Note that Definition 2 does not really define a functor, and, more fundamentally, the universe \mathcal{U} is not a (pre-)category in the sense of [3]. Instead, the appropriate notion for P would be an ∞ -(endo)functor on the $(\infty, 1)$ -category \mathcal{U} [8]. However, we do not attempt to make any of these notions precise, and do not make use of any “functorial” properties of the defined maps. Our use of the word “functor” merely indicates an analogy to the 1-categorical case.

To any signature $S = (A, B)$ we associate a type of *coalgebras* \mathbf{Coalg}_S , and a family of types of morphisms between them:

► **Definition 3.** Given a signature $S = (A, B)$ as in Definition 2, an **S -coalgebra** is defined to be a pair (C, γ) consisting of a type $C : \mathcal{U}$ and a map $\gamma : C \rightarrow P_S C$. A map of coalgebras from (C, γ) to (D, δ) is defined to be a pair (f, p) of a map $f : C \rightarrow D$ and a path $p : \delta \circ f = P_S f \circ \gamma$. Put differently, we set

$$\mathbf{Coalg}_S := \sum_{C:\mathcal{U}} C \rightarrow P_S C$$

and

$$\mathbf{Coalg}_S((C, \gamma), (D, \delta)) := \sum_{f:C \rightarrow D} \delta \circ f = P_S f \circ \gamma .$$

There is an obvious composition of coalgebra morphisms, and the identity map $C \rightarrow C$ is the carrier of a coalgebra endomorphism on (C, γ) . We also write $(C, \gamma) \Rightarrow (D, \delta)$ for the type of coalgebra morphisms from (C, γ) to (D, δ) .

3.3 What is an M-type?

In this section we define (internal) M-types in HoTT via a universal property.

► **Definition 4.** Given a container (A, B) , the **(internal) M-type** $M_{A,B}$ associated to (A, B) is defined to be the pair $(M, \mathbf{out} : M \rightarrow P_{A,B} M)$ with the following universal property: for any coalgebra $(C, \gamma) : \mathbf{Coalg}_{(A,B)}$ of (A, B) , the type of coalgebra morphisms $(C, \gamma) \Rightarrow (M, \mathbf{out})$ from (C, γ) to (M, \mathbf{out}) is contractible.

The use of the definite article in Definition 4 is justified by the following lemma:

► **Lemma 5.** *The type*

$$\mathbf{Final}_S := \sum_{((X, \rho) : \mathbf{Coalg}_S)} \prod_{((C, \gamma) : \mathbf{Coalg}_S)} \mathbf{isContr}((C, \gamma) \Rightarrow (X, \rho))$$

is a proposition.

Proof. The proof that any two final coalgebras (L, out) and (L', out') have equivalent carriers is standard. The Univalence Axiom then implies that the carriers are (propositionally) equal, $L = L'$. It then remains to show that the coalgebra structure out , when transported along this identity, is equal to out' . We refer to the formalized proof for details. ◀

That is, any inhabitant of Final_S is necessarily unique up to propositional equality. We refer to this inhabitant as the “final coalgebra”. In Section 4 we construct the final coalgebra.

In the introduction, we use the adjectives “internal” and “external” to distinguish between types specified via universal properties and type-theoretic rules, respectively. Since we do not consider rules for M-types (that is, external M-types) in this work, we drop the adjective “internal” in what follows.

In the following example, we anticipate the result of the next section, namely the existence of a final coalgebra for any signature (A, B) :

► **Example 6.** The coinductive type $\text{Stream}(A_0)$ of streams over a base type A_0 is given by $\text{M}(A, B)$ with $A = A_0$ and $B(a) := 1$ for any $a : A_0$. The corresponding polynomial functor P satisfies $P(X) = A_0 \times X$.

Using finality of $\text{M}(A, B)$ we can define maps into streams and prove that they have the expected computational behaviour. For example, the zip function

$$\text{zip} : \text{Stream}(A) \times \text{Stream}(B) \rightarrow \text{Stream}(A \times B)$$

can be obtained from the universal property applied to the coalgebra

$$\begin{aligned} \theta : \text{Stream}(A) \times \text{Stream}(B) &\rightarrow (A \times B) \times (\text{Stream}(A) \times \text{Stream}(B)) \\ \theta(xs, ys) &:= ((\text{head}(xs), \text{head}(ys)), (\text{tail}(xs), \text{tail}(ys))) \end{aligned}$$

where $\text{head} : \text{Stream}(X) \rightarrow X$ and $\text{tail} : \text{Stream}(X) \rightarrow \text{Stream}(X)$ are the two components of the final coalgebra out . The computational behaviour of zip is expressed by the fact that zip is a coalgebra morphism

$$\text{zip}(xs, ys) = \text{cons}((\text{head}(xs), \text{head}(ys)), (\text{zip}(\text{tail}(xs), \text{tail}(ys)))),$$

where $\text{cons} = \text{out}^{-1}$.

4 Derivability of M-types


In Section 3 we defined the type Final_S of final coalgebras of a signature S , and showed that this type is a proposition (Lemma 5). In this section, we construct an element of Final_S , which, combined with Lemma 5, proves the following theorem per the remark at the end of Section 2:

► **Theorem 7.** *The type*

$$\text{Final}_S = \sum_{((X, \rho))} \prod_{((C, \gamma))} \text{isContr}((C, \gamma) \Rightarrow (X, \rho))$$

is contractible.

The construction of the final coalgebra is done in several steps, inspired by a construction of M-types from W-types in a type theory satisfying Axiom K by Altenkirch et al. [4]. Its carrier is defined as the limit of a *chain*:



■ **Figure 2** Universal property of L .

► **Definition 8.** A **chain** is a pair (X, π) of a family of types $X : \mathbb{N} \rightarrow \mathcal{U}$ and a family of functions $\pi_n : X_{n+1} \rightarrow X_n$. Here and below we write $X_n := X(n)$ for the n th component of the family X .

The (homotopy) limit of such a chain is given by the type of “compatible tuples”:

► **Definition 9.** The **limit** of the chain (X, π) is given by the type

$$L := \sum_{(x: \prod_{(n:\mathbb{N})} X_n)} \prod_{(n:\mathbb{N})} \pi_n x_{n+1} = x_n .$$

The limit is equipped with projections $p_n : L \rightarrow X_n$, and $\beta_n : \pi_n \circ p_{n+1} = p_n$. Sometimes, we simply write

$$L = \lim X,$$

when the maps π are clear.

Note that this limit (we drop the adjective “homotopy”) is an instance of the general construction of homotopy limits by Avigad, Kapulkin, and Lumsdaine [5].

► **Lemma 10.** *The type L satisfies the following universal property: for all types A , we have an equivalence of types between maps into L and “cones” over X :*

$$A \rightarrow L \simeq \sum_{(f: \prod_{(n:\mathbb{N})} A \rightarrow X_n)} \prod_{(n:\mathbb{N})} \pi_n \circ f_{n+1} = f_n \quad =: \text{Cone}(A) .$$

The equivalence, from left to right, maps a function $f : A \rightarrow L$ to its projections $p_n \circ f$, as shown in Figure 2.

The next lemma is about tuples in **cochains**, that is, tuples in chains with inverted arrows. Those tuples are determined by their first element:

► **Lemma 11.** *Let $X : \mathbb{N} \rightarrow \mathcal{U}$ be a family of types, and $l : \prod_{(n:\mathbb{N})} X_n \rightarrow X_{n+1}$ a family of functions. Let*

$$Z := \sum_{(x: \prod_{(n:\mathbb{N})} X_n)} \prod_{(n:\mathbb{N})} x_{n+1} = l_n(x_n) .$$

Then the projection $Z \rightarrow X_0$ is an equivalence.

Proof. Let G be the functor defined by $GY = 1 + Y$. Fix an element $z : Z$. Then z and l together define a G -algebra structure on X , regarded as a fibration over \mathbb{N} . Since \mathbb{N} is the homotopy initial algebra of G , the type Sz of algebra sections of X is contractible. But Z is equivalent to

$$\sum_{(z:X_0)} \sum_{(x:\prod_{(n:\mathbb{N})} X_n)} (x_0 = z) \times \left(\prod_{n:\mathbb{N}} x_{n+1} = l_n(x_n) \right),$$

which is exactly $\sum_{(z:X_0)} Sz \simeq X_0$. \blacktriangleleft

► **Lemma 12.** *Let (X, π) be a chain, and let (X', π') be the shifted chain, defined by $X'_n := X_{n+1}$ and $\pi'_n := \pi_{n+1}$. Then the two chains have equivalent limits.*

Proof. Let L and L' be the limits of (X, π) and (X', π') , respectively. We have

$$\begin{aligned} L' &\stackrel{(1)}{\simeq} \sum_{(y:\prod_{(n:\mathbb{N})} X_{n+1})} \prod_{(n:\mathbb{N})} \pi_{n+1} y_{n+1} = y_n \\ &\stackrel{(2)}{\simeq} \sum_{(x_0:X_0)} \sum_{(y:\prod_{(n:\mathbb{N})} X_{n+1})} (\pi_0 y_0 = x_0) \times \left(\prod_{n:\mathbb{N}} \pi_{n+1} y_{n+1} = y_n \right) \\ &\stackrel{(3)}{\simeq} \sum_{x:\prod_{(n:\mathbb{N})} X_n} (\pi_0 x_0 = x_0) \times \left(\prod_{n:\mathbb{N}} \pi_{n+1} x_{n+2} = x_{n+1} \right) \\ &\stackrel{(4)}{\simeq} \sum_{(x:\prod_{(n:\mathbb{N})} X_n)} \prod_{(n:\mathbb{N})} \pi_n x_{n+1} = x_n \\ &\stackrel{(5)}{\simeq} L, \end{aligned}$$

where (1) and (5) are by definition. Equivalence (2) is given by multiplying with the contractible type $\sum_{(x_0:X_0)} \pi_0 y_0 = x_0$ [13, Lem. 3.11.8] and subsequent swapping of components in a direct product. Equivalence (3) is given by joining the first two components, and similarly in (4) the last two components are joined. \blacktriangleleft

The next lemma says that polynomial functors (see Definition 2) commute with limits of chains. Let (A, B) be a container with associated polynomial functor $P = P_{A,B}$. Let (X, π) be a chain with limit (L, p) . Define the chain $(PX, P\pi)$ with $PX_n := P(X_n)$ and likewise for $P\pi$, and let L^P be its limit. The family of maps $Pp_n : PL \rightarrow PX_n$ determines a function $\alpha : PL \rightarrow L^P$.

► **Lemma 13.** *The function α is an isomorphism.*

Proof. By “equational” reasoning we have

$$\begin{aligned}
L^P &\stackrel{(6)}{\simeq} \sum_{(w: \prod_{(n:\mathbb{N})} \sum_{(a:A)} B(a) \rightarrow X_n)} \prod_{(n:\mathbb{N})} (P\pi_n)w_{n+1} = w_n \\
&\stackrel{(7)}{\simeq} \sum_{(a: \prod_{(n:\mathbb{N})} A)} \sum_{(u: \prod_{(n:\mathbb{N})} B(a_n) \rightarrow X_n)} \prod_{(n:\mathbb{N})} (a_{n+1}, \pi_n \circ u_{n+1}) = (a_n, u_n) \\
&\stackrel{(8)}{\simeq} \sum_{(a: \prod_{(n:\mathbb{N})} A)} \sum_{(p: \prod_{(n:\mathbb{N})} a_{n+1} = a_n)} \sum_{(u: \prod_{(n:\mathbb{N})} B(a_n) \rightarrow X_n)} \prod_{(n:\mathbb{N})} (p_n)_*(\pi_n \circ u_{n+1}) = u_n \\
&\stackrel{(9)}{\simeq} \sum_{(a:A)} \sum_{(u: \prod_{(n:\mathbb{N})} B(a) \rightarrow X_n)} \prod_{(n:\mathbb{N})} \pi_n \circ u_{n+1} = u_n \\
&\stackrel{(10)}{\simeq} \sum_{a:A} B(a) \rightarrow L \\
&\stackrel{(11)}{\simeq} PL
\end{aligned}$$

where (6) and (11) are by definition, (7) is by swapping Π and Σ , (8) by expanding equality of pairs, (9) by applying Lemma 11 and (10) by universal property of L . Verifying that the composition of these isomorphisms is α is straightforward. \blacktriangleleft

Proof of Theorem 7. We now construct a terminal coalgebra for a container (A, B) . Let $P = P_{A,B}$ the polynomial functor associated to the container. By recursion on \mathbb{N} we define the chain

$$1 \leftarrow \text{!} P1 \leftarrow \text{!} P^! P^2 1 \leftarrow \text{!} P^{2!} P^3 \leftarrow \text{!} P^{3!} \dots$$

which for brevity we call (W, π) , that is, $W_n := P^n 1$ and $\pi_n := P^n \text{!}$.

Let (L, p) be the limit of (W, π) . If L' is the limit of the shifted chain, we have a sequence of equivalences

$$PL \stackrel{(12)}{\simeq} L' \stackrel{(13)}{\simeq} L$$

where (12) is given by Lemma 13 and (13) by Lemma 12. We denote this equivalence by $\text{in} : PL \rightarrow L$, and its inverse by $\text{out} : L \rightarrow PL$.

It is worth noting that the construction of L “does not raise the universe level”, i.e., if A and B are contained in some universe \mathcal{U} , then L is contained in \mathcal{U} as well. In other words, we only need one universe to carry out our construction of the final coalgebra.

We will now show that (L, out) is a final (A, B) -coalgebra. For this, let (C, γ) be any coalgebra, i.e., $\gamma : C \rightarrow PC$. The type of coalgebra morphisms $(C, \gamma) \Rightarrow (L, \text{out})$ is given by

$$U := \sum_{f:C \rightarrow L} \text{out} \circ f = Pf \circ \gamma .$$

We need to show that U is contractible. We compute as follows (see below the math display for intermediate definitions):

$$\begin{aligned}
U &\stackrel{(14)}{\simeq} \sum_{f:C \rightarrow L} \text{out} \circ f = Pf \circ \gamma \\
&\stackrel{(15)}{\simeq} \sum_{f:C \rightarrow L} \text{out} \circ f = \text{step}(f) \\
&\stackrel{(16)}{\simeq} \sum_{f:C \rightarrow L} \text{in} \circ \text{out} \circ f = \text{in} \circ \text{step}(f) \\
&\stackrel{(17)}{\simeq} \sum_{f:C \rightarrow L} f = \Psi(f) \\
&\stackrel{(18)}{\simeq} \sum_{c:\text{Cone}} e(c) = \Psi(e(c)) \\
&\stackrel{(19)}{\simeq} \sum_{c:\text{Cone}} e(c) = e(\Phi(c)) \\
&\stackrel{(20)}{\simeq} \sum_{c:\text{Cone}} c = \Phi(c) \\
&\stackrel{(21)}{\simeq} \sum_{((u,q):\text{Cone})} \sum_{(p:u=\Phi_0(u))} p_*(q) = \Phi_1(u)(q) \\
&\stackrel{(22)}{\simeq} \sum_{(u:\text{Cone}_0)} \sum_{(p:u=\Phi_0 u)} \sum_{q:\text{Cone}_1 u} p_*(q) = \Phi_1(u)(q) \\
&\stackrel{(23)}{\simeq} \sum_{t:1} 1 \\
&\simeq 1
\end{aligned}$$

where we use the following definitions: The function $\text{step}_Y : (C \rightarrow Y) \rightarrow (C \rightarrow PY)$ is defined as $\text{step}_Y(f) := Pf \circ \gamma$ and $\Psi : (C \rightarrow L) \rightarrow (C \rightarrow L)$ is defined as $\Psi(f) := \text{in} \circ \text{step}_L(f)$. The map $\Phi : \text{Cone} \rightarrow \text{Cone}$ is the counterpart of Ψ on the side of cones. We define $\Phi(u, g) = (\Phi_0 u, \Phi_1 u(g)) : \text{Cone} \rightarrow \text{Cone}$ with

$$\begin{aligned}
(\Phi_0 u)_0 &:= x \mapsto tt : C \rightarrow 1 = W_0 \\
(\Phi_0 u)_{n+1} &:= \text{step}_{W_n}(u_n) : C \rightarrow W_{n+1} = PW_n
\end{aligned}$$

and analogously for Φ_1 on paths. By e we denote the equivalence of Lemma 10 from right to left, and $\text{Cone} = \sum_{(u:\text{Cone}_0)} \text{Cone}_1(u)$ is short for $\text{Cone}(C)$. The equivalence **(16)** follows from in being an equivalence, and **(17)** follows from in and out being inverse to each other. We pass from maps into L to cones in **(18)**, using the equivalence of Lemma 10, while **(19)** uses the commutativity of the following square:

$$\begin{array}{ccc}
\text{Cone} & \xrightarrow{e} & (C \rightarrow L) \\
\Phi \downarrow & & \downarrow \Psi \\
\text{Cone} & \xrightarrow{e} & (C \rightarrow L).
\end{array}$$

In **(21)**, identity in a sigma type is reduced to identity of the components, and in **(22)** the components are rearranged. Finally, step **(23)** consists of two applications of Lemma 11.

Altogether, this shows that for any coalgebra (C, γ) , the type of coalgebra morphisms $(C, \gamma) \Rightarrow (L, \text{out})$ is contractible. This concludes the construction of a final coalgebra for the (polynomial functor of the) signature (A, B) and thus, combined with Lemma 5, the proof of Theorem 7. \blacktriangleleft

From the construction of L we get the following corollary about the homotopy level of M -types:

► **Lemma 14.** *The homotopy level of the (carrier of the) M -type associated to the signature (A, B) is bounded by that of the type of nodes A , that is,*

$$\text{isofhlevel}_n(A) \rightarrow \text{isofhlevel}_n(M(A, B)) .$$

► **Example 15.** We continue the example of streams of Example 6, with $A = A_0$ the type of nodes. In that case, the chain considered in the proof of Theorem 7 is given by $W_n = P^n(1) = A^n$, and the map $\pi_n : A^{n+1} \rightarrow A^n$ chops off the $(n+1)$ th element of any $(n+1)$ -tuple. The limit L is hence given by $A^{\mathbb{N}} = (\mathbb{N} \rightarrow A)$. The type of streams over A has the same homotopy level as the type A of nodes.

We conclude this section with a proof of the *principle of coinduction*:

► **Definition 16** (Bisimulation). Let (C, γ) be a coalgebra for some signature S with associated polynomial functor P and let $\mathcal{R} : C \rightarrow C \rightarrow \mathcal{U}$ be a binary relation. Define

$$\overline{\mathcal{R}} := \sum_{(a:C)} \sum_{(b:C)} \mathcal{R}(a)(b)$$

along with two projections $\pi_1^{\overline{\mathcal{R}}}(a, b, p) := a$ and $\pi_2^{\overline{\mathcal{R}}}(a, b, p) := b$.

An S -bisimulation is a relation \mathcal{R} together with a map $\alpha_{\mathcal{R}} : \overline{\mathcal{R}} \rightarrow P(\overline{\mathcal{R}})$ such that both $\pi_1^{\overline{\mathcal{R}}}$ and $\pi_2^{\overline{\mathcal{R}}}$ are P -coalgebra morphisms:

$$\begin{array}{ccccc} C & \xleftarrow{\pi_1^{\overline{\mathcal{R}}}} & \overline{\mathcal{R}} & \xrightarrow{\pi_2^{\overline{\mathcal{R}}}} & C \\ \downarrow \gamma & & \downarrow \alpha_{\mathcal{R}} & & \downarrow \gamma \\ P(C) & \xleftarrow{P(\pi_1^{\overline{\mathcal{R}}})} & P(\overline{\mathcal{R}}) & \xrightarrow{P(\pi_2^{\overline{\mathcal{R}}})} & P(C) \end{array}$$

We say that a bisimulation is an *equivalence bisimulation* when the underlying relation is an equivalence relation.

► **Lemma 17.** *The identity relation $\cdot = \cdot$ over an S -coalgebra C is an equivalence bisimulation. We write Δ_C for $\cdot = \cdot$.*

► **Theorem 18** (Coinduction proof principle). *Let (L, out) be the final coalgebra for S . For any bisimulation $\overline{\mathcal{R}}$ over L , we have $\overline{\mathcal{R}} \subseteq \Delta_L$. That is, for any $m, m' : L$,*

$$\mathcal{R}(m)(m') \rightarrow m = m' .$$

Proof. Since (L, out) is the final coalgebra, for any coalgebra (C, γ) there exists a unique coalgebra morphism $\text{unfold}_C : C \rightarrow L$. It follows that $\pi_1^{\overline{\mathcal{R}}} = \text{unfold}_{\overline{\mathcal{R}}} = \pi_2^{\overline{\mathcal{R}}}$. Finally, given $r : \mathcal{R}(m)(m')$, we obtain $m = \pi_1^{\overline{\mathcal{R}}}(m, m', r) = \pi_2^{\overline{\mathcal{R}}}(m, m', r) = m'$. \blacktriangleleft

5 Indexed M-types

In this section, we briefly state the main definitions for *indexed* M-types. The difference to plain M-types is that the type of nodes of an indexed M-type is actually given by a family of types, indexed by a type of sorts I .

► **Definition 19.** An **indexed container** is given by a quadruple (I, A, B, r) such that $I : \mathcal{U}$ is a type, $i : I \vdash A(i) : \mathcal{U}$ is a family of types dependent on I , B is a family $i : I, a : A(i) \vdash B_i(a) : \mathcal{U}$ and r specifies the “sort” of the subtrees, i.e., $r : \prod_{(i:I)} \prod_{(a:A(i))} B_i(a) \rightarrow I$.

► **Definition 20.** The polynomial functor P associated to an indexed container (I, A, B, r) is an endofunction on the type $I \rightarrow \mathcal{U}$:

$$(PX)(i) := \sum_{(a:A(i))} \prod_{(b:B_i(a))} X(r_{i,a}(b)) .$$

The functorial action on morphisms is, analogously to Definition 2, given by postcomposition.

Coalgebras for indexed containers, and their morphisms, are defined completely analogously to Definition 3. Again, we prove that terminal coalgebras for indexed containers exist uniquely:

► **Theorem 21.** *Let (I, A, B, r) be an indexed container. Then the associated indexed M-type is uniquely specified and can be constructed in the type theory described in Section 2.*

An example of a coinductive type that needs indices to be expressed as an M-type is a coinductive formulation of *equivalence of types*, to our knowledge due to T. Altenkirch:

► **Example 22.** Let $I := \mathcal{U} \times \mathcal{U}$, and let $A : I \rightarrow \mathcal{U}$ be defined by $A(X, Y) := (X \rightarrow Y) \times (Y \rightarrow X)$. Define B as $B(X, Y)(f, g) := X \times Y$ and $r(X, Y)(f, g)(x, y) := (f(x) = y) \times (x = g(y))$. Then the associated M-type is a family $M : I \rightarrow \mathcal{U}$ and $M(A, B)$ is equivalent to $A \simeq B$.

6 Formalization

The proofs contained in this paper have been formalised in the proof assistant AGDA in a self-contained development. The proof has been type-checked by version 2.4.2.2 of AGDA. The source code as well as HTML documentation can be found on <https://github.com/HoTT/m-types/>. The source code is also archived with the arXiv version of this article [2].

The formalised proofs deal with the indexed case (Section 5) directly, but apart from that, they correspond closely to the informal proofs presented here. In particular, they make heavy use of the “equational reasoning” technique to prove equivalences between types.

In fact, it is often the case that proving an equivalence between types A and B “directly”, i.e. by defining functions $A \rightarrow B$ and $B \rightarrow A$, and then proving that they compose to identities in both directions, is unfeasibly hard, due to the complexity of the terms involved.

However, in most cases, we can construct an equivalence between A and B by composition of several simple equivalences. Those simple building blocks range from certain ad-hoc equivalences that make specific use of the features of the two types involved, to very general and widely applicable “rewriting rules”, like the fact that we can swap a Σ -type with a Π -type (sometimes called the *constructive axiom of choice* [13]).

By assembling elementary equivalences, then, we automatically get both a function $A \rightarrow B$ and a proof that it is an equivalence. However, sometimes care is needed to ensure that the resulting function has the desired computational properties.

An important consideration during the formalisation of the proof of Theorem 7 was keeping the size of the terms reasonably short. For example, in one early attempt, the innocent-looking term `in` was being normalised into a term spanning more than 12000 lines.

The explosion in term size was clearly causing performance issues during type-checking, which resulted in AGDA running out of memory while checking apparently trivial proofs.

We solved this problem by moving certain definitions (like that of `in` itself) into an *abstract* block, thereby preventing AGDA from expanding it at all. Of course, this means that we lost all the computational properties of certain functions, so we had to abstract out their computational behaviour in the form of propositional equalities, and manually use them in the proof of Theorem 7. This work-around is the source of most of the complications in the formal proof.

7 Conclusion and future work

We have shown how to construct a class of coinductive types from the basic type constructors and natural numbers in Homotopy Type Theory. Our construction follows a well-known pattern, that is known to work in type theory with identity reflection rule.

We work in a univalent universe, but we make minimal use of univalence itself: Lemma 5 is the only result that uses it directly, and elsewhere univalence only appears indirectly through the use of functional extensionality. We intend to make these dependencies more explicit in future developments of the formalisation.

Finally, the coinductive types we construct do not satisfy the expected computation rules *judgmentally*, but only *propositionally*. This fact would justify adding coinduction as a primitive rather than a derived notion – provided that judgmental computation rules are validated by the intended semantics. Those semantic questions are left for future work.

Acknowledgments. We are grateful to many people for helpful discussions about coinductive types, online as well as offline: Thorsten Altenkirch, Steve Awodey, Martín Escardó, Nicolai Kraus, Peter LeFanu Lumsdaine, Ralph Matthes, Paige North, Mike Shulman, and Vladimir Voevodsky. We thank Nicolai Kraus, Peter LeFanu Lumsdaine and Paige North for suggesting improvements and clarifications for this paper.

References

- 1 Michael Abbott, Thorsten Altenkirch, and Neil Ghani. Containers: Constructing strictly positive types. *Theor. Comput. Sci.*, 342(1):3–27, 2005. doi:10.1016/j.tcs.2005.06.002.
- 2 Benedikt Ahrens, Paolo Capriotti, and Régis Spadotti. Non-wellfounded trees in Homotopy Type Theory, 2015. arXiv:1504.02949.
- 3 Benedikt Ahrens, Krzysztof Kapulkin, and Michael Shulman. Univalent categories and the Rezk completion. *Mathematical Structures in Computer Science*, FirstView:1–30, 1 2015. doi:10.1017/S0960129514000486.
- 4 Thorsten Altenkirch, Neil Ghani, Peter Hancock, Conor McBride, and Peter Morris. Indexed Containers. available at <http://www.cs.nott.ac.uk/~txa/publ/jcont.pdf>.
- 5 Jeremy Avigad, Krzysztof Kapulkin, and Peter LeFanu Lumsdaine. Homotopy limits in type theory. *Mathematical Structures in Computer Science*, FirstView:1–31, 1 2015. doi:10.1017/S0960129514000498.
- 6 Steve Awodey, Nicola Gambino, and Kristina Sojakova. Inductive types in homotopy type theory. In *Proceedings of Symposium on Logic In Computer Science*, pages 95–104, 2012. doi:10.1109/LICS.2012.21.

- 7 Michael Barr. Terminal coalgebras in well-founded set theory. *Theor. Comput. Sci.*, 114(2):299–315, 1993. doi:10.1016/0304-3975(93)90076-6.
- 8 James Cranch. Concrete Categories in Homotopy Type Theory, 2013. arXiv:1311.1852.
- 9 Chris Kapulkin, Peter LeFanu Lumsdaine, and Vladimir Voevodsky. The Simplicial Model of Univalent Foundations, 2014. arXiv:1211.2851.
- 10 HoTT mailing list. Discussion on coinductive types on HoTT mailing list, <https://groups.google.com/d/msg/homotopytypetheory/tYRTcI20pyo/PIrI6t5me-oJ>.
- 11 Per Martin-Löf. *Intuitionistic Type Theory*. Bibliopolis, 1984.
- 12 Kristina Sojakova. Higher inductive types as homotopy-initial algebras. In *Proceedings of the Symposium on Principles of Programming Languages*, pages 31–42, 2015. doi:10.1145/2676726.2676983.
- 13 The Univalent Foundations Program. *Homotopy Type Theory: Univalent Foundations of Mathematics*. <http://homotopytypetheory.org/book>, Institute for Advanced Study, 2013.
- 14 Benno van den Berg and Federico De Marchi. Non-well-founded trees in categories. *Ann. Pure Appl. Logic*, 146(1):40–59, 2007. arXiv:math/0409158.

Conservativity of Embeddings in the $\lambda\Pi$ Calculus Modulo Rewriting

Ali Assaf^{1,2}

1 Inria, Paris, France

2 École polytechnique, Palaiseau, France

Abstract

The $\lambda\Pi$ calculus can be extended with rewrite rules to embed any functional pure type system. In this paper, we show that the embedding is conservative by proving a relative form of normalization, thus justifying the use of the $\lambda\Pi$ calculus modulo rewriting as a logical framework for logics based on pure type systems. This result was previously only proved under the condition that the target system is normalizing. Our approach does not depend on this condition and therefore also works when the source system is not normalizing.

1998 ACM Subject Classification F.4.1 Mathematical Logic

Keywords and phrases $\lambda\Pi$ calculus modulo rewriting, pure type systems, logical framework, normalization, conservativity

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.31

1 Introduction

The $\lambda\Pi$ *calculus modulo rewriting* is a logical framework that extends the $\lambda\Pi$ calculus [10] with rewrite rules. Through the Curry-de Bruijn-Howard correspondence, it can express properties and proofs of various logics. Cousineau and Dowek [6] introduced a general embedding of *functional pure type systems* (FPTS), a large class of typed λ -calculi, in the $\lambda\Pi$ calculus modulo rewriting: for any FPTS λS , they constructed the system $\lambda\Pi/S$ using appropriate rewrite rules, and defined two translation functions $|M|$ and $\|A\|$ that translate respectively the terms and the types of λS to $\lambda\Pi/S$. This embedding is complete, in the sense preserves typing: if $\Gamma \vdash_{\lambda S} M : A$ then $\|\Gamma\| \vdash_{\lambda\Pi/S} |M| : \|A\|$. From the logical point of view, it preserves provability. The converse property, called *conservativity*, was only shown partially: assuming $\lambda\Pi/S$ is strongly normalizing, if there is a term N such that $\|\Gamma\| \vdash_{\lambda\Pi/S} N : \|A\|$ then there is a term M such that $\Gamma \vdash_{\lambda S} M : A$.

1.1 Normalization and conservativity

Not much is known about normalization in $\lambda\Pi/S$. Cousineau and Dowek [6] showed that the embedding preserves reduction: if $M \longrightarrow M'$ then $|M| \longrightarrow^+ |M'|$. As a consequence, if $\lambda\Pi/S$ is strongly normalizing (i.e. every well-typed term normalizes) then so is λS , but the converse might not be true *a priori*. This was not enough to show the conservativity of the embedding, so the proof relied on the unproven assumption that $\lambda\Pi/S$ is normalizing. This result is insufficient if one wants to consider the $\lambda\Pi$ calculus modulo rewriting as a general logical framework for defining logics and expressing proofs in those logics, as proposed in [4, 5]. Indeed, if the embedding turns out to be inconsistent then checking proofs in the logical framework has very little benefit.



© Ali Assaf;

licensed under Creative Commons License CC-BY

13th International Conference on Typed Lambda Calculi and Applications (TLCA'15).

Editor: Thorsten Altenkirch; pp. 31–44

Leibniz International Proceedings in Informatics



LIPIC Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Consider the PTS λHOL that corresponds to higher order logic [1]:

$$\begin{aligned} \mathcal{S} &= \text{Prop, Type, Kind} \\ \mathcal{A} &= (\text{Prop} : \text{Type}), (\text{Type} : \text{Kind}) \\ \mathcal{R} &= (\text{Prop, Prop, Prop}), (\text{Type, Prop, Prop}), (\text{Type, Type, Type}) \end{aligned}$$

This PTS is strongly normalizing, and therefore consistent. A polymorphic variant of λHOL is specified by $U^- = HOL + (\text{Kind, Type, Type})$. It turns out that λU^- is inconsistent: there is a term ω such that $\vdash_{\lambda U^-} \omega : \Pi\alpha : \text{Prop}. \alpha$ and which is not normalizing [1]. We motivate the need for a proof of conservativity with the following example.

► **Example 1.** The polymorphic identity function $I = \lambda\alpha : \text{Type}. \lambda x : \alpha. x$ is *not* well-typed in λHOL , but it is well-typed in λU^- and so is its type:

$$\begin{aligned} \vdash_{\lambda U^-} I : \Pi\alpha : \text{Type}. \alpha \rightarrow \alpha \\ \vdash_{\lambda U^-} \Pi\alpha : \text{Type}. \alpha \rightarrow \alpha : \text{Type} \end{aligned}$$

However, the translation $|I| = \lambda\alpha : u_{\text{Type}}. \lambda x : \varepsilon_{\text{Type}} \alpha. x$ is well-typed in $\lambda\Pi/HOL$:

$$\begin{aligned} \vdash_{\lambda\Pi/HOL} |I| : \Pi\alpha : u_{\text{Type}}. \varepsilon_{\text{Type}} \alpha \rightarrow \varepsilon_{\text{Type}} \alpha \\ \vdash_{\lambda\Pi/HOL} \Pi\alpha : u_{\text{Type}}. \varepsilon_{\text{Type}} \alpha \rightarrow \varepsilon_{\text{Type}} \alpha : \text{Type} \end{aligned}$$

It seems that $\lambda\Pi/HOL$, just like λU^- , allows more functions than λHOL , even though the type of $|I|$ is not the translation of a λHOL type. Does that make $\lambda\Pi/HOL$ inconsistent?

1.2 Absolute normalization vs relative normalization

One way to answer the question is to prove strong normalization of $\lambda\Pi/S$ by constructing a model, for example in the algebra of *reducibility candidates* [9]. Dowek [7] recently constructed such a model for the embedding of higher-order logic (λHOL) and of the calculus of constructions (λC). However, this technique is still very limited. Indeed, proving such a result is, by definition, at least as hard as proving the consistency of the original system. It requires specific knowledge of λS and the construction of such a model can be very involved, such as for the calculus of constructions with an infinite universe hierarchy (λC^∞).

In this paper, we take a different approach and show that $\lambda\Pi/S$ is conservative in all cases, even when λS is *not* normalizing. Instead of showing that $\lambda\Pi/S$ is strongly normalizing, we show that it is weakly normalizing *relative to* λS , meaning that proofs in the target language can be reduced to proofs in the source language. That way we prove only what is needed to show conservativity, without having to prove the consistency of λS all over again. After identifying the main difficulties, we characterize a *PTS completion* [17, 16] S^* containing S , and define an inverse translation from $\lambda\Pi/S$ to λS^* . We then prove that λS^* is a conservative extension of λS using the *reducibility method* [18].

1.3 Outline

In Section 2, we recall the theory of pure type systems. In Section 3, we present the framework of the $\lambda\Pi$ calculus modulo rewriting. In Section 4, we introduce Cousineau and Dowek's embedding of functional pure type systems in the $\lambda\Pi$ calculus modulo rewriting. In Section 5, we prove the conservativity of the embedding using the techniques mentioned above. In Section 6, we summarize the results and discuss future work. Some long proofs have been omitted and can be found in the long version of this paper¹.

¹ Available online at arXiv:1504.05038.

$$\begin{array}{c}
\text{EMPTY} \\
\frac{}{\text{WF}_{\lambda S}(\cdot)} \\
\\
\text{DECLARATION} \\
\frac{\Gamma \vdash_{\lambda S} A : s \quad x \notin \Gamma}{\text{WF}_{\lambda S}(\Gamma, x : A)} \\
\\
\text{VARIABLE} \\
\frac{\text{WF}_{\lambda S}(\Gamma) \quad (x : A) \in \Gamma}{\Gamma \vdash_{\lambda S} x : A} \\
\\
\text{SORT} \\
\frac{\text{WF}_{\lambda S}(\Gamma) \quad (s_1 : s_2) \in \mathcal{A}}{\Gamma \vdash_{\lambda S} s_1 : s_2} \\
\\
\text{PRODUCT} \\
\frac{\Gamma \vdash_{\lambda S} A : s_1 \quad \Gamma, x : A \vdash_{\lambda S} B : s_2 \quad (s_1, s_2, s_3) \in \mathcal{R}}{\Gamma \vdash_{\lambda S} \Pi x : A. B : s_3} \\
\\
\text{ABSTRACTION} \\
\frac{\Gamma, x : A \vdash_{\lambda S} M : B \quad \Gamma \vdash_{\lambda S} \Pi x : A. B : s}{\Gamma \vdash_{\lambda S} \lambda x : A. M : \Pi x : A. B} \\
\\
\text{APPLICATION} \\
\frac{\Gamma \vdash_{\lambda S} M : \Pi x : A. B \quad \Gamma \vdash_{\lambda S} N : A}{\Gamma \vdash_{\lambda S} M N : B[x \setminus N]} \\
\\
\text{CONVERSION} \\
\frac{\Gamma \vdash_{\lambda S} M : A \quad \Gamma \vdash_{\lambda S} B : s \quad A \equiv_{\beta} B}{\Gamma \vdash_{\lambda S} M : B}
\end{array}$$

■ **Figure 1** Typing rules of λS .

2 Pure type systems

Pure type systems [1] are a general class of typed λ -calculi parametrized by a specification.

► **Definition 2.** A PTS *specification* is a triple $S = (\mathcal{S}, \mathcal{A}, \mathcal{R})$ where

- \mathcal{S} is a set of symbols called *sorts*
- $\mathcal{A} \subseteq \mathcal{S} \times \mathcal{S}$ is a set of *axioms* of the form $(s_1 : s_2)$
- $\mathcal{R} \subseteq \mathcal{S} \times \mathcal{S} \times \mathcal{S}$ is a set of *rules* of the form (s_1, s_2, s_3)

We write (s_1, s_2) as a short-hand for the rule (s_1, s_2, s_2) . The specification S is *functional* if the relations \mathcal{A} and \mathcal{R} are functional, that is $(s_1, s_2) \in \mathcal{A}$ and $(s_1, s'_2) \in \mathcal{A}$ imply $s_2 = s'_2$, and $(s_1, s_2, s_3) \in \mathcal{R}$ and $(s_1, s_2, s'_3) \in \mathcal{R}$ imply $s_3 = s'_3$. The specification is *full* if for all $s_1, s_2 \in \mathcal{S}$, there is a sort s_3 such that $(s_1, s_2, s_3) \in \mathcal{R}$.

► **Definition 3.** Given a PTS specification $S = (\mathcal{S}, \mathcal{A}, \mathcal{R})$ and a countably infinite set of variables \mathcal{V} , the abstract syntax of λS is defined by the following grammar:

$$\begin{array}{l}
\text{(terms)} \quad \mathcal{T} ::= \mathcal{S} \mid \mathcal{V} \mid \mathcal{T}\mathcal{T} \mid \lambda \mathcal{V} : \mathcal{T}. \mathcal{T} \mid \Pi \mathcal{V} : \mathcal{T}. \mathcal{T} \\
\text{(contexts)} \quad \mathcal{C} ::= \cdot \mid \mathcal{C}, \mathcal{V} : \mathcal{T}
\end{array}$$

We use lower case letters $x, y, \alpha, \beta, \dots$ to denote variables, uppercase letters such as M, N, A, B, \dots to denote terms, and uppercase Greek letters such as $\Gamma, \Delta, \Sigma, \dots$ to denote contexts. The set of free variables of a term M is denoted by $\text{FV}(M)$. We write $A \rightarrow B$ for $\Pi x : A. B$ when $x \notin \text{FV}(B)$.

The typing rules of λS are presented in Figure 1. We write $\Gamma \vdash M : A$ instead of $\Gamma \vdash_{\lambda S} M : A$ when the context is unambiguous. We say that M is a Γ -*term* when $\text{WF}(\Gamma)$ and $\Gamma \vdash M : A$ for some A . We say that A is a Γ -*type* when $\text{WF}(\Gamma)$ and either $\Gamma \vdash A : s$ or $A = s$ for some $s \in \mathcal{S}$. We write $\Gamma \vdash M : A : s$ as a shorthand for $\Gamma \vdash M : A \wedge \Gamma \vdash A : s$.

► **Example 4.** The following well-known systems can all be expressed as functional pure type systems using the same set of sorts $\mathcal{S} = \text{Type, Kind}$ and the same set of axioms $\mathcal{A} = (\text{Type} : \text{Kind})$:

- Simply-typed λ calculus ($\lambda\rightarrow$):
 $\mathcal{R} = (\text{Type}, \text{Type})$
- System F ($\lambda 2$):
 $\mathcal{R} = (\text{Type}, \text{Type}), (\text{Kind}, \text{Type})$
- $\lambda\Pi$ calculus (λP):
 $\mathcal{R} = (\text{Type}, \text{Type}), (\text{Type}, \text{Kind})$
- Calculus of constructions (λC):
 $\mathcal{R} = (\text{Type}, \text{Type}), (\text{Kind}, \text{Type}), (\text{Type}, \text{Kind}), (\text{Kind}, \text{Kind})$

► **Example 5.** Let $I = \lambda\alpha : \text{Type}. \lambda x : \alpha. x$ be the polymorphic identity function. The term I is not well-typed in the simply typed λ calculus but it is well-typed in the calculus of constructions λC :

$$\vdash_{\lambda C} I : \Pi\alpha : \text{Type}. \alpha \rightarrow \alpha$$

The following properties hold for all pure type systems [1].

► **Theorem 6 (Correctness of types).** *If $\Gamma \vdash_{\lambda S} M : A$ then $\text{WF}_{\lambda S}(\Gamma)$ and either $\Gamma \vdash_{\lambda S} A : s$ or $A = s$ for some $s \in \mathcal{S}$, i.e. A is a Γ -type.*

The reason why we don't always have $\Gamma \vdash_{\lambda S} A : s$ is that some sorts do not have an associated axiom, such as **Kind** in Example 4, which leads to the following definition.

► **Definition 7 (Top-sorts).** A sort $s \in \mathcal{S}$ is called a *top-sort* when there is no sort $s' \in \mathcal{S}$ such that $(s : s') \in \mathcal{A}$.

The following property is useful for proving properties about systems with top-sorts.

► **Theorem 8 (Top-sort types).** *If $\Gamma \vdash_{\lambda S} A : s$ and s is a top-sort then either $A = s'$ for some sort $s' \in \mathcal{S}$ or $A = \Pi x : B. C$ for some terms B, C .*

► **Theorem 9 (Confluence).** *If $M_1 \rightarrow_{\beta}^* M_2$ and $M_1 \rightarrow_{\beta}^* M_3$ then there is a term M_4 such that $M_2 \rightarrow_{\beta}^* M_4$ and $M_3 \rightarrow_{\beta}^* M_4$.*

► **Theorem 10 (Product compatibility).** *If $\Pi x : A. B \equiv_{\beta} \Pi x : A'. B'$ then $A \equiv_{\beta} A'$ and $B \equiv_{\beta} B'$.*

► **Theorem 11 (Subject reduction).** *If $\Gamma \vdash_{\lambda S} M : A$ and $M \rightarrow_{\beta}^* M'$ then $\Gamma \vdash_{\lambda S} M' : A$.*

Finally, we state the following property for functional pure type systems.

► **Theorem 12 (Uniqueness of types).** *Let \mathcal{S} be a functional specification. If $\Gamma \vdash_{\lambda S} M : A$ and $\Gamma \vdash_{\lambda S} M : B$ then $A \equiv_{\beta} B$.*

In the rest of the paper, all the pure type systems we will consider will be functional.

3 The $\lambda\Pi$ calculus modulo rewriting

The $\lambda\Pi$ calculus, also known as *LF* and as λP , is one of the simplest forms of λ calculus with dependent types, and corresponds through the Curry-de Bruijn-Howard correspondence to a minimal first-order logic of higher-order terms. As mentioned in Example 4, it can be defined as the functional pure type system λP with the following specification:

$$\begin{aligned} \mathcal{S} &= \text{Type}, \text{Kind} \\ \mathcal{A} &= \text{Type} : \text{Kind} \\ \mathcal{R} &= (\text{Type}, \text{Type}), (\text{Type}, \text{Kind}) \end{aligned}$$

$\frac{\text{EMPTY}}{\text{WF}_{\lambda\Pi/}(\cdot)}$	$\frac{\text{DECLARATION}}{\text{WF}_{\lambda\Pi/}(\Gamma, x : A)} \quad \Gamma \vdash_{\lambda\Pi/} A : s \quad x \notin \Sigma, \Gamma$	$\frac{\text{VARIABLE}}{\text{WF}_{\lambda\Pi/}(\Gamma)} \quad (x : A) \in \Sigma, \Gamma$ $\Gamma \vdash_{\lambda\Pi/} x : A$
$\frac{\text{SORT}}{\text{WF}_{\lambda\Pi/}(\Gamma)} \quad (s_1 : s_2) \in \mathcal{A}$ $\Gamma \vdash_{\lambda\Pi/} s_1 : s_2$	$\frac{\text{PRODUCT}}{\text{WF}_{\lambda\Pi/}(\Gamma, x : A)} \quad \Gamma \vdash_{\lambda\Pi/} A : s_1$	$\frac{\Gamma, x : A \vdash_{\lambda\Pi/} B : s_2 \quad (s_1, s_2, s_3) \in \mathcal{R}}{\Gamma \vdash_{\lambda\Pi/} \Pi x : A. B : s_3}$
$\frac{\text{ABSTRACTION}}{\text{WF}_{\lambda\Pi/}(\Gamma, x : A)} \quad \Gamma \vdash_{\lambda\Pi/} M : B \quad \Gamma \vdash_{\lambda\Pi/} \Pi x : A. B : s$ $\Gamma \vdash_{\lambda\Pi/} \lambda x : A. M : \Pi x : A. B$	$\frac{\text{APPLICATION}}{\text{WF}_{\lambda\Pi/}(\Gamma, x : A)} \quad \Gamma \vdash_{\lambda\Pi/} M : \Pi x : A. B \quad \Gamma \vdash_{\lambda\Pi/} N : A$ $\Gamma \vdash_{\lambda\Pi/} M N : B[x \setminus N]$	
$\frac{\text{CONVERSION}}{\text{WF}_{\lambda\Pi/}(\Gamma, x : A)} \quad \Gamma \vdash_{\lambda\Pi/} M : A \quad \Gamma \vdash_{\lambda\Pi/} B : s \quad A \equiv_{\beta R} B$ $\Gamma \vdash_{\lambda\Pi/} M : B$		

■ **Figure 2** Typing rules of $\lambda\Pi/(\Sigma, R)$.

The $\lambda\Pi$ *calculus modulo rewriting* extends the $\lambda\Pi$ calculus with rewrite rules. By equating terms modulo a set of rewrite rules R in addition to α and β equivalence, it can type more terms using the conversion rule, and therefore express theories that are more complex. The calculus can be seen as a variant of Martin-Löf's logical framework [13, 11] where equalities are expressed as rewrite rules.

We recall that a rewrite rule is a triple $[\Delta] M \rightsquigarrow N$ where Δ is a context and M, N are terms such that $\text{FV}(N) \subseteq \text{FV}(M)$. A set of rewrite rules R induces a reduction relation on terms, written \longrightarrow_R , defined as the smallest contextual closure such that if $[\Delta] M \rightsquigarrow N \in R$ then $\sigma(M) \longrightarrow_R \sigma(N)$ for any substitution σ of the variables in Δ . We define the relation $\longrightarrow_{\beta R}$ as $\longrightarrow_{\beta} \cup \longrightarrow_R$, the relation \equiv_R as the smallest congruence containing \longrightarrow_R , and the relation $\equiv_{\beta R}$ as the smallest congruence containing $\longrightarrow_{\beta R}$.

► **Definition 13.** A rewrite rule $[\Delta] M \rightsquigarrow N$ is *well-typed in a context* Σ when there is a term A such that $\Sigma, \Delta \vdash_{\lambda\Pi} M : A$ and $\Sigma, \Delta \vdash_{\lambda\Pi} N : A$.

► **Definition 14.** Let Σ be a well-formed $\lambda\Pi$ context and R a set of rewrite rules that are well-typed in Σ . The $\lambda\Pi$ *calculus modulo* (Σ, R) , written $\lambda\Pi/(\Sigma, R)$, is defined with the same syntax as the $\lambda\Pi$ calculus, but with the typing rules of Figure 2. We write $\lambda\Pi/$ instead of $\lambda\Pi/(\Sigma, R)$ when the context is unambiguous.

► **Example 15.** Let Σ be the context

$$\alpha : \text{Type}, c : \alpha, f : \alpha \rightarrow \text{Type}$$

and R be the following rewrite rule

$$[\cdot] f c \rightsquigarrow \Pi y : \alpha. f y \rightarrow f y$$

Then the term

$$\delta = \lambda x : f c . x c x$$

is well-typed in $\lambda\Pi/(\Sigma, R)$:

$$\vdash_{\lambda\Pi/(\Sigma, R)} \delta : f c \rightarrow f c$$

Note that the term δ would not be well-typed without the rewrite rule, even if we replace all the occurrences of $f c$ in δ by $\Pi y : \alpha. f y \rightarrow f y$.

The system $\lambda\Pi$ is a pure type system and therefore enjoys all the properties mentioned in Section 2. The behavior of $\lambda\Pi/(\Sigma, R)$ however depends on the choice of (Σ, R) . In particular, some properties analogous to those of pure type systems depend on the confluence of the relation $\rightarrow_{\beta R}$.

► **Theorem 16** (Correctness of types). *If $\Gamma \vdash_{\lambda\Pi/} M : A$ then $\text{WF}_{\lambda\Pi/}(\Gamma)$ and either $\Gamma \vdash_{\lambda\Pi/} A : s$ for some $s \in \{\text{Type}, \text{Kind}\}$ or $A = \text{Kind}$.*

► **Theorem 17** (Top-sort types). *If $\Gamma \vdash_{\lambda\Pi/} A : \text{Kind}$ then either $A = \text{Type}$ or $A = \Pi x : B. C$ for some terms B, C such that $\Gamma, x : B \vdash_{\lambda\Pi/} C : \text{Kind}$.*

Assuming $\rightarrow_{\beta R}$ is confluent, the following properties hold [3].

► **Theorem 18** (Product compatibility). *If $\Pi x : A. B \equiv_{\beta R} \Pi x : A'. B'$ then $A \equiv_{\beta R} A'$ and $B \equiv_{\beta R} B'$.*

► **Theorem 19** (Subject reduction). *If $\Gamma \vdash_{\lambda\Pi/} M : A$ and $M \rightarrow_{\beta R}^* M'$ then $\Gamma \vdash_{\lambda\Pi/} M' : A$.*

► **Theorem 20** (Uniqueness of types). *If $\Gamma \vdash_{\lambda\Pi/} M : A$ and $\Gamma \vdash_{\lambda\Pi/} M : B$ then $A \equiv_{\beta R} B$.*

4 Embedding FPTs's in the $\lambda\Pi$ calculus modulo

In this section, we present the embedding of functional pure type systems in the $\lambda\Pi$ calculus modulo rewriting as introduced by Cousineau and Dowek [6]. In this embedding, sorts are represented as *universes à la Tarski*, as introduced by Martin-Löf [12] and later developed by Luo [11] and Palmgren [14]. The embedding is done in two steps. First, given a pure type system λS , we construct $\lambda\Pi/S$ by giving an appropriate signature and rewrite system. Second, we define a translation from the terms and types of λS to the terms and types of $\lambda\Pi/S$. The proofs of the theorems in this section can be found in the original paper [6].

► **Definition 21** (The system $\lambda\Pi/S$). Consider a functional pure type system specified by $S = (\mathcal{S}, \mathcal{A}, \mathcal{R})$. Define Σ_S to be the well-formed context containing the declarations:

$$\begin{array}{ll} u_s : \text{Type} & \forall s \in \mathcal{S} \\ \varepsilon_s : u_s \rightarrow \text{Type} & \forall s \in \mathcal{S} \\ \dot{s}_1 : u_{s_2} & \forall s_1 : s_2 \in \mathcal{A} \\ \dot{\pi}_{s_1 s_2 s_3} : \Pi \alpha : u_{s_1}. (\varepsilon_{s_1} \alpha \rightarrow u_{s_2}) \rightarrow u_{s_3} & \forall (s_1, s_2, s_3) \in \mathcal{R} \end{array}$$

Let R_S be the well-typed rewrite system containing the rules

$$[\cdot] \varepsilon_{s_2} \dot{s}_1 \rightsquigarrow u_{s_1}$$

for all $s_1 : s_2 \in \mathcal{A}$, and

$$[\Delta_{s_1 s_2 s_3}] \varepsilon_{s_3} (\dot{\pi}_{s_1 s_2 s_3} A B) \rightsquigarrow \Pi x : (\varepsilon_{s_1} A). \varepsilon_{s_2} (B x)$$

for all $(s_1, s_2, s_3) \in \mathcal{R}$, where $\Delta_{s_1 s_2 s_3} = (A : u_{s_1}, B : (\varepsilon_{s_1} \alpha \rightarrow u_{s_2}))$. The system $\lambda\Pi/S$ is defined as the $\lambda\Pi$ calculus modulo (Σ_S, R_S) , that is, $\lambda\Pi/(\Sigma_S, R_S)$.

► **Theorem 22** (Confluence). *The relation $\longrightarrow_{\beta R}$ is confluent.*

The translation is composed of two functions, one from the terms of λS to the terms of $\lambda\Pi/S$, the other from the types of λS to the types of $\lambda\Pi/S$.

► **Definition 23.** The translation $|M|_{\Gamma}$ of Γ -terms and the translation $\|A\|_{\Gamma}$ of Γ -types are mutually defined as follows.

$$\begin{aligned}
|s|_{\Gamma} &= \dot{s} \\
|x|_{\Gamma} &= x \\
|MN|_{\Gamma} &= |M|_{\Gamma} |N|_{\Gamma} \\
|\lambda x:A. M|_{\Gamma} &= \lambda x:\|A\|_{\Gamma}. |M|_{\Gamma, x:A} \\
|\Pi x:A. B|_{\Gamma} &= \dot{\pi}_{s_1 s_2 s_3} |A|_{\Gamma} (\lambda x:\|A\|_{\Gamma}. |B|_{\Gamma, x:A}) \\
&\quad \text{where } \Gamma \vdash A : s_1 \\
&\quad \text{and } \Gamma, x : A \vdash B : s_2 \\
&\quad \text{and } (s_1, s_2, s_3) \in \mathcal{R} \\
\|s\|_{\Gamma} &= u_s \\
\|\Pi x:A. B\|_{\Gamma} &= \Pi x:\|A\|_{\Gamma}. \|B\|_{\Gamma, x:A} \\
\|A\|_{\Gamma} &= \varepsilon_s |A|_{\Gamma} \text{ where } \Gamma \vdash A : s
\end{aligned}$$

Note that this definition is redundant but it is well-defined up to $\equiv_{\beta R}$. In particular, because some Γ -types are also Γ -terms, there are two ways to translate them, but they are equivalent:

$$\begin{aligned}
\varepsilon_{s_2} \dot{s}_1 &\equiv_{\beta R} u_{s_1} \\
\varepsilon_{s_3} |\Pi x:A. B|_{\Gamma} &\equiv_{\beta R} \Pi x:\|A\|_{\Gamma}. \|B\|_{\Gamma, x:A}
\end{aligned}$$

This definition is naturally extended to well-formed contexts as follows.

$$\begin{aligned}
\|\cdot\| &= \cdot \\
\|\Gamma, x : A\| &= \|\Gamma\|, x : \|A\|_{\Gamma}
\end{aligned}$$

► **Example 24.** The polymorphic identity function of the Calculus of constructions λC is translated as

$$|I| = \lambda\alpha : u_{\text{Type}}. \lambda x : \varepsilon_{\text{Type}} \alpha. x$$

and its type $A = \Pi\alpha : \text{Type}. \alpha \rightarrow \alpha$ is translated as:

$$|A| = \dot{\pi}_{\text{Kind}, \text{Type}, \text{Type}} \dot{\text{Type}} (\lambda\alpha : u_{\text{Type}}. |A_{\alpha}|)$$

where $A_{\alpha} = \alpha \rightarrow \alpha$ and

$$|A_{\alpha}| = \dot{\pi}_{\text{Type}, \text{Type}, \text{Type}} \alpha (\lambda x : \varepsilon_{\text{Type}} \alpha. \varepsilon_{\text{Type}} \alpha)$$

The identity function applied to itself is translated as:

$$|IAI| = |I| |A| |I|$$

The embedding is complete, in the sense that all the typing relations of λS are preserved by the translation.

► **Theorem 25** (Completeness). *For any context Γ and terms M and A , if $\Gamma \vdash_{\lambda S} M : A$ then $\|\Gamma\| \vdash_{\lambda\Pi/S} |M|_{\Gamma} : \|A\|_{\Gamma}$.*

5 Conservativity

In this section, we prove the converse of the completeness property. One could attempt to prove that if $\|\Gamma\| \vdash_{\lambda\Pi/S} |M|_{\Gamma} : \|A\|_{\Gamma}$ then $\Gamma \vdash_{\lambda S} M : A$. However, that would be too weak because the translation $|M|_{\Gamma}$ is only defined for well-typed terms. A second attempt would be to define inverse translations $\varphi(M)$ and $\psi(A)$ and prove that if $\Gamma \vdash_{\lambda\Pi/S} M : A$ then $\psi(\Gamma) \vdash_{\lambda S} \varphi(M) : \psi(A)$, but that would not work either because not all terms and types of $\lambda\Pi/S$ correspond to valid terms and types of λS , as was shown in Example 1. Therefore the property that we want to prove is: if there is a term N such that $\|\Gamma\| \vdash_{\lambda\Pi/S} N : \|A\|_{\Gamma}$ then there is a term M such that $\Gamma \vdash_{\lambda S} M : A$.

The main difficulty is that some of these *external* terms can be involved in witnessing valid λS types, as illustrated by the following example.

► **Example 26.** Consider the context $nat : \text{Type}$. Even though the polymorphic identity function I and its type are not well-typed in λHOL , they can be used in $\lambda\Pi/HOL$ to construct a witness for $nat \rightarrow nat$.

$$nat : u_{\text{Type}} \vdash_{\lambda\Pi/HOL} (|I| \text{ nat}) : (\varepsilon_{\text{Type}} \text{ nat} \rightarrow \varepsilon_{\text{Type}} \text{ nat})$$

We can normalize the term $|I| \text{ nat}$ to $\lambda x : \varepsilon_{\text{Type}} \text{ nat}. x$ which is a term that corresponds to a valid λHOL term: it is the translation of the term $\lambda x : nat.x$. However, as discussed previously, we cannot restrict ourselves to normal terms because we do not know if $\lambda\Pi/S$ is normalizing.

To prove conservativity, we will therefore need to address the following issues:

1. The system $\lambda\Pi/S$ can type more terms than λS .
2. These terms can be used to construct proofs for the translation of λS types.
3. The $\lambda\Pi/S$ terms that inhabit the translation of λS types can be reduced to the translation of λS terms.

We will proceed as follows. First, we will eliminate β -redexes at the level of Kind by reducing $\lambda\Pi/S$ to a subset $\lambda\Pi^-/S$. Then, we will extend λS to a *minimal completion* λS^* that can type more terms than λS , and show that $\lambda\Pi^-/S$ corresponds to λS^* using inverse translations $\varphi(M)$ and $\psi(A)$. Finally, we will show that λS^* terms inhabiting λS types can be reduced to λS terms. The procedure is summarized in the following diagram.

$$\begin{array}{ccc}
 \lambda\Pi/S & \xrightarrow[\beta^*]{\text{(Lemma 28)}} & \lambda\Pi^-/S \\
 \uparrow \text{(Theorem 25)} \quad |M| \quad \|A\| & & \downarrow \varphi(M) \quad \psi(A) \quad \text{(Lemma 39)} \\
 \lambda S & \xleftarrow[\beta^*]{\text{(Lemma 47)}} & \lambda S^*
 \end{array}$$

5.1 Eliminating β -redexes at the level of Kind

In $\lambda\Pi/S$, we can have β -redexes at the level of Kind such as $(\lambda x : A. u_s) M$. These redexes are artificial and are never generated by the forward translation of any PTS. We show here that they can always be safely eliminated.

► **Definition 27.** A Γ -term M of type C is at the level of Kind (resp. Type) if $\Gamma \vdash C : \text{Kind}$ (resp. $\Gamma \vdash C : \text{Type}$). We define $\lambda\Pi^-/S$ terms as the subset of well-typed $\lambda\Pi/S$ terms that do not contain any Kind -level β -redexes.

► **Lemma 28.** *For any $\lambda\Pi/S$ context Γ and Γ -term M , there is a $\lambda\Pi^-/S$ term M^- such that $M \longrightarrow_{\beta}^* M^-$.*

Proof. Reducing a Kind-level β -redex $(\lambda x : A. B)N$ does not create other Kind-level β -redexes because N is at the level of **Type**. Indeed, in the $\lambda\Pi$ calculus modulo rewriting the only Kind rule is (Type, Kind, Kind). Therefore $N : A : \mathbf{Type}$. If N reduces to a λ -abstraction then the only redexes it can create are at the level of **Type**. Therefore, the number of Kind-level β -redexes strictly decreases, so any Kind-level β -reduction strategy will terminate. ◀

► **Example 29.** The term

$$I_1 = \lambda\alpha : u_{\mathbf{Type}}. \lambda x : \varepsilon_{\mathbf{Type}} ((\lambda\beta : u_{\mathbf{Type}}. \beta) \alpha). x$$

is in $\lambda\Pi^-/HOL$. The term

$$I_2 = \lambda\alpha : u_{\mathbf{Type}}. \lambda x : ((\lambda\beta : u_{\mathbf{Type}}. \varepsilon_{\mathbf{Type}} \beta) \alpha). x$$

is not in $\lambda\Pi^-/HOL$ but

$$I_2 \longrightarrow_{\beta} \lambda\alpha : u_{\mathbf{Type}}. \lambda x : \varepsilon_{\mathbf{Type}} \alpha. x$$

which is in $\lambda\Pi^-/HOL$.

5.2 Minimal completion

To simplify our reducibility proof in the next section, we will translate $\lambda\Pi/S$ back to a pure type system, but since it cannot be λS we will define a slightly larger PTS called λS^* that contains λS and that will be easier to manipulate than $\lambda\Pi/S$.

The reason we need a larger PTS is that we have types that do not have a type, such as top-sorts because there is no associated axiom. Similarly, we can sometimes prove $\Gamma, x : A \vdash_{\lambda S} M : B$ but cannot abstract over x because there is no associated product rule. Completions of pure type systems were originally introduced by Severi [17, 16] to address these issues by injecting λS into a larger pure type system.

► **Definition 30** (Completion [16]). A specification $S' = (\mathcal{S}', \mathcal{A}', \mathcal{R}')$ is a *completion* of S if

1. $\mathcal{S} \subseteq \mathcal{S}', \mathcal{A} \subseteq \mathcal{A}', \mathcal{R} \subseteq \mathcal{R}'$, and
2. for all sorts $s_1 \in \mathcal{S}$, there is a sort $s_2 \in \mathcal{S}'$ such that $(s_1 : s_2) \in \mathcal{A}'$, and
3. for all sorts $s_1, s_2 \in \mathcal{S}'$, there is a sort $s_3 \in \mathcal{S}'$ such that $(s_1, s_2, s_3) \in \mathcal{R}'$.

Notice that all the top-sorts of λS are typable in $\lambda S'$ and that $\lambda S'$ is full, meaning that all products are typable. These two properties reflect exactly the discrepancy between λS and $\lambda\Pi^-/S$. Not all completions are conservative though, so we define the following completion.

► **Definition 31** (Minimal completion). We define the *minimal completion* of S , written S^* , to be the following specification:

$$\begin{aligned} \mathcal{S}^* &= \mathcal{S} \cup \{\tau\} \\ \mathcal{A}^* &= \mathcal{A} \cup \{(s_1 : \tau) \mid s_1 \in \mathcal{S}, \nexists s_2, (s_1 : s_2) \in \mathcal{A}\} \\ \mathcal{R}^* &= \mathcal{R} \cup \{(s_1, s_2, \tau) \mid s_1, s_2 \in \mathcal{S}^*, \nexists s_3, (s_1, s_2, s_3) \in \mathcal{R}\} \end{aligned}$$

where $\tau \notin \mathcal{S}$.

We add a new top-sort τ and axioms $s : \tau$ for all previous top-sorts s , and complete the rules to obtain a PTS full. The new system is a completion by Definition 30 and it is minimal in the sense that we generically added the smallest number of sorts, axioms, and rules so that the result is guaranteed to be conservative. Any well-typed term of λS is also well-typed in λS^* , but just like $\lambda\Pi^-/S$, this system allows more functions than λS .

► **Example 32.** The polymorphic identity function is well-typed in λHOL^* .

$$\vdash_{\lambda HOL^*} I : \Pi\alpha : \mathbf{Type}. \alpha \rightarrow \alpha$$

$$\vdash_{\lambda HOL^*} \Pi\alpha : \mathbf{Type}. \alpha \rightarrow \alpha : \tau$$

Next, we define inverse translations that translate the terms and types of $\lambda\Pi^-/S$ to the terms and types of λS^* .

► **Definition 33 (Inverse translations).** The inverse translation of terms $\varphi(M)$ and the inverse translation of types $\psi(A)$ are mutually defined as follows.

$$\begin{aligned} \varphi(\dot{s}) &= s \\ \varphi(\dot{\pi}_{s_1 s_2 s_3}) &= \lambda\alpha : s_1. \lambda\beta : (\alpha \rightarrow s_2). \Pi x : \alpha. \beta x \\ \varphi(x) &= x \\ \varphi(M N) &= \varphi(M) \varphi(N) \\ \varphi(\lambda x : A. M) &= \lambda x : \psi(A). \varphi(M) \\ \\ \psi(u_s) &= s \\ \psi(\varepsilon_s M) &= \varphi(M) \\ \psi(\Pi x : A. B) &= \Pi x : \psi(A). \psi(B) \end{aligned}$$

Note that this is only a partial definition, but it is total for $\lambda\Pi^-/S$ terms. In particular, it is an inverse of the forward translation in the following sense.

► **Lemma 34.** For any Γ -term M and Γ -type A ,

1. $\varphi(|M|_\Gamma) \equiv_\beta M$,
2. $\psi(|A|_\Gamma) \equiv_\beta A$.

Proof. By induction on M or A . We show the product case where $M = \Pi x : A. B$. By induction hypothesis, $\varphi(|A|) \equiv_\beta A$ and $\varphi(|B|) \equiv_\beta B$. Therefore

$$\begin{aligned} \varphi(|M|) &= (\lambda\alpha. \lambda\beta. \Pi x : \alpha. \beta x) \varphi(|A|) (\lambda x. \varphi(|B|)) \\ &\xrightarrow{\tau_\beta^*} \Pi x : \varphi(|A|). \varphi(|B|) \\ &\equiv_\beta \Pi x : A. B \end{aligned}$$

◀

Next we show that the inverse translations preserve typing.

► **Lemma 35.**

1. $\varphi(M[x \setminus N]) = \varphi(M)[x \setminus \varphi(N)]$
2. $\psi(A[x \setminus N]) = \psi(A)[x \setminus \varphi(N)]$

Proof. By induction on M or A . We show the product case $A = \Pi y : B. C$. Without loss of generality, $y \neq x$ and $y \notin N$ and $y \notin \varphi(N)$. Then $\Pi y : B. C[x \setminus N] = \Pi y : B[x \setminus N]. C[x \setminus N]$.

By induction hypothesis, $\psi(B[x \setminus N]) = \psi(B)[x \setminus \varphi(N)]$ and $\psi(C[x \setminus N]) = \psi(C)[x \setminus \varphi(N)]$. Therefore

$$\begin{aligned} \psi(A[x \setminus N]) &= \Pi y : \psi(B)[x \setminus \varphi(N)]. \psi(C)[x \setminus \varphi(N)] \\ &= \Pi x : \psi(B). \psi(C)[x \setminus \varphi(N)] \\ &= \psi(\Pi x : B. C)[x \setminus \varphi(N)] \end{aligned}$$

◀

► **Lemma 36.**

1. If $M \rightarrow_{\beta R} N$ then $\varphi(M) \rightarrow_{\beta}^* \varphi(N)$
2. If $A \rightarrow_{\beta R} B$ then $\psi(A) \rightarrow_{\beta}^* \psi(B)$

Proof. By induction on M or A . We show the base cases.

- Case $M = (\lambda x : A_1. M_1) N_1$, $N = M_1[x \setminus N_1]$. Then $\varphi(M) = (\lambda x : \psi(A_1). \varphi(M_1)) \varphi(N_1)$. Therefore $\varphi(M) \rightarrow_{\beta} \varphi(M_1)[x \setminus \varphi(N_1)]$ which is equal to $\varphi(M_1[x \setminus N_1])$ by Lemma 35.
- Case $A = \varepsilon_s \dot{s}$, $B = u_s$. Then $\psi(A) = s = \psi(B)$.
- Case $A = \varepsilon_{s_1} (\tilde{\pi}_{s_1 s_2 s_3} A_1 B_1)$, $B = \Pi x : \varepsilon_{s_1} A_1. \varepsilon_{s_2} (B_1 x)$. Then

$$\begin{aligned} \psi(A) &= (\lambda \alpha. \lambda \beta. \Pi x : \alpha. \beta x) \varphi(A_1) \varphi(B_1) \\ &\rightarrow_{\beta}^* \Pi x : \varphi(A_1). \varphi(B_1) x \\ &= \psi(\Pi x : A_1. B_1 x) \end{aligned}$$

◀

► **Lemma 37.**

1. If $M \equiv_{\beta R} N$ then $\varphi(M) \equiv_{\beta} \varphi(N)$
2. If $A \equiv_{\beta R} B$ then $\psi(A) \equiv_{\beta} \psi(B)$

Proof. Follows from Lemma 36. ◀

Because the forward translation of contexts does not introduce any type variable, we define the following restriction on contexts.

► **Definition 38** (Object context). We say that Γ is an *object context* if $\Gamma \vdash_{\lambda\Pi/S} A : \text{Type}$ for all $x : A \in \Gamma$. If $\Gamma = (x_1 : A_1, \dots, x_n : A_n)$ is an object context, we define $\psi(\Gamma)$ as $(x_1 : \psi(A_1), \dots, x_n : \psi(A_n))$.

► **Lemma 39.** For any $\lambda\Pi^-/S$ object context Γ and terms M, A :

1. If $\text{WF}_{\lambda\Pi/S}(\Gamma)$ then $\text{WF}_{\lambda S^*}(\psi(\Gamma))$.
2. If $\Gamma \vdash_{\lambda\Pi/S} M : A : \text{Type}$ then $\psi(\Gamma) \vdash_{\lambda S^*} \varphi(M) : \psi(A)$.
3. If $\Gamma \vdash_{\lambda\Pi/S} A : \text{Type}$ then $\psi(\Gamma) \vdash_{\lambda S^*} \psi(A) : s$ for some sort $s \in \mathcal{S}^*$.

Proof. By induction on the derivation. The details of the proof can be found in the long version of this paper. ◀

5.3 Reduction to λS

In order to show that λS^* is a conservative extension of λS , we prove that β -reduction at the level of τ terminates. A straightforward proof by induction would fail because contracting a τ -level β -redex can create other such redexes. To solve this, we adapt Tait's *reducibility method* [18]. The idea is to strengthen the induction hypothesis of the proof by defining a predicate by induction on the type of the term.

► **Definition 40.** The predicate $\Gamma \models_S M : A$ is defined as $\text{WF}_{\lambda S}(\Gamma)$ and $\Gamma \vdash_{\lambda S^*} M : A : s$ for some sort s and:

- if $s \neq \tau$ or $A = s'$ for some $s' \in \mathcal{S}$ then $\Gamma \models_S M : A$ iff $M \longrightarrow_{\beta}^* M'$ and $A \longrightarrow_{\beta}^* A'$ for some M', A' such that $\Gamma \vdash_{\lambda S} M' : A'$,
- if $s = \tau$ and $A = \Pi x : B.C$ for some B, C then $\Gamma \models_S M : A$ iff for all N such that $\Gamma \models_S N : B$, $\Gamma \models_S MN : C[x \setminus N]$.

Note that recursive definition covers all cases thanks to Theorem 8. To show that it is well-founded, we define the following measure of A .

► **Definition 41.** If $\text{WF}_{\lambda S}(\Gamma)$ and $\Gamma \vdash_{\lambda S^*} A : s$ then $\mathcal{H}_{\tau}(A)$ is defined as:

$$\begin{aligned} \mathcal{H}_{\tau}(A) &= 0 && \text{if } s \neq \tau \\ \mathcal{H}_{\tau}(s') &= 0 && \text{if } s = \tau \\ \mathcal{H}_{\tau}(\Pi x : B.C) &= 1 + \max(\mathcal{H}_{\tau}(B) + \mathcal{H}_{\tau}(C)) && \text{if } s = \tau \end{aligned}$$

► **Lemma 42.** If $\Gamma, x : B \vdash_{\lambda S^*} C : \tau$ and $\Gamma \vdash_{\lambda S^*} N : B$ then $\mathcal{H}_{\tau}(C[x \setminus N]) = \mathcal{H}_{\tau}(C)$.

Proof. By induction on C . ◀

► **Corollary 43.** Definition 40 is well-founded.

Proof. The measure $\mathcal{H}_{\tau}(A)$ strictly decreases in the definition. ◀

The predicate we defined is compatible with β -equivalence.

► **Lemma 44.** If $\Gamma \models_S M : A$ and $\Gamma \vdash_{\lambda S^*} M' : A$ and $M \equiv_{\beta} M'$ then $\Gamma \models_S M' : A$.

Proof. By induction on the height of A .

- If $s \neq \tau$ or $A = s'$ for some $s' \in \mathcal{S}$ then $M \longrightarrow_{\beta}^* M''$ and $A \longrightarrow_{\beta}^* A'$ for some M'', A' such that $\Gamma \vdash_{\lambda S} M'' : A'$. By confluence and subject reduction, $M' \longrightarrow_{\beta}^* M'''$ such that $\Gamma \vdash_{\lambda S} M''' : A'$.
- If $s = \tau$ and $A = \Pi x : B.C$ for some B, C then for all N such that $\Gamma \models_S N : B$, $\Gamma \models_S MN : C[x \setminus N]$. By induction hypothesis, $\Gamma \models_S M'N : C[x \setminus N]$. Therefore $\Gamma \models_S M' : \Pi x : B.C$. ◀

► **Lemma 45.** If $\Gamma \models_S M : A$ and $\Gamma \vdash_{\lambda S^*} A' : s$ and $A \equiv_{\beta} A'$ then $\Gamma \models_S M : A'$.

Proof. By induction on the height of A .

- If $s \neq \tau$ or $A = s'$ for some $s' \in \mathcal{S}$ then $M \longrightarrow_{\beta}^* M'$ and $A \longrightarrow_{\beta}^* A''$ for some M', A'' such that $\Gamma \vdash_{\lambda S} M' : A''$. By conversion, $\Gamma \vdash_{\lambda S^*} M : A'$, so by subject reduction $\Gamma \vdash_{\lambda S^*} M' : A'$. By confluence, subject reduction, and conversion, $A' \longrightarrow_{\beta}^* A'''$ such that $\Gamma \vdash_{\lambda S} M' : A'''$.
- If $s = \tau$ and $A = \Pi x : B.C$ for some B, C then for all N such that $\Gamma \models_S N : B$, $\Gamma \models_S MN : C[x \setminus N]$. By product compatibility, $A' = \Pi x : B'.C'$ such that $B \equiv_{\beta} B'$ and $C \equiv_{\beta} C'$. By induction hypothesis, $\Gamma \models_S MN : C'[x \setminus N]$. Therefore $\Gamma \models_S M : \Pi x : B'.C'$. ◀

We extend the definition of the inductive predicate to contexts and substitutions before proving the main general lemma.

► **Definition 46.** If $\text{WF}_{\lambda S^*}(\Gamma)$, $\text{WF}_{\lambda S}(\Gamma')$, and σ is a substitution for the variables of Γ , then $\Gamma' \models_S \sigma : \Gamma$ when $\Gamma' \models_S \sigma(x) : \sigma(A)$ for all $(x : A) \in \Gamma$.

► **Lemma 47.** *If $\Gamma \vdash_{\lambda S^*} M : A : s$ then for any context Γ' and substitution σ such that $\text{WF}_{\lambda S}(\Gamma')$ and $\Gamma' \models_S \sigma : \Gamma$, $\Gamma' \models_S \sigma(M) : \sigma(A)$.*

Proof. By induction on the derivation of $\Gamma \vdash_{\lambda S^*} M : A$. The details of the proof can be found in the long version of this paper. ◀

► **Corollary 48.** *Suppose $\text{WF}_{\lambda S}(\Gamma)$ and either $\Gamma \vdash_{\lambda S} A : s$ or $A = s$ for some $s \in S$. If $\Gamma \vdash_{\lambda S^*} M : A$ then $M \rightarrow_{\beta}^* M'$ such that $\Gamma \vdash_{\lambda S} M' : A$.*

Proof. Taking σ as the identity substitution, there are terms M' and A' such that $M \rightarrow_{\beta}^* M'$ and $A \rightarrow_{\beta}^* A'$ and $\Gamma \vdash_{\lambda S} M' : A'$. If $A = s \in S$ then $A' = s$ and we are done. Otherwise by conversion we get $\Gamma \vdash_{\lambda S} M' : A$. ◀

We now have all the tools to prove the main theorem.

► **Theorem 49 (Conservativity).** *For any Γ -type A of λS , if there is a term N such that $\|\Gamma\| \vdash_{\lambda \Pi/S} N : \|A\|_{\Gamma}$ then there is a term M such that $\Gamma \vdash_{\lambda S} M : A$.*

Proof. By Lemma 28, there is a $\lambda \Pi^-/S$ term N^- such that $N \rightarrow_{\beta}^* N^-$. By subject reduction, $\|\Gamma\| \vdash_{\lambda \Pi/S} N^- : \|A\|_{\Gamma}$. By Lemmas 39 and 34, $\Gamma \vdash_{\lambda S^*} \varphi(N^-) : A$. By Corollary 48, there is a term M such that $\varphi(N^-) \rightarrow_{\beta}^* M$ and $\Gamma \vdash_{\lambda S} M : A$. ◀

6 Conclusion

We have shown that $\lambda \Pi/S$ is conservative even when λS is not normalizing. Even though $\lambda \Pi/S$ can construct more functions than λS , it preserves the semantics of λS . This effect is similar to various conservative extensions of pure type systems such as *pure type systems with definitions* [17], *pure type systems without the Π -condition* [16], or *predicative (ML) polymorphism* [15]. Inconsistency in pure type systems usually does not come from the ability to type more functions, but from the possible impredicativity caused by assigning a sort to the type of these functions. It is clear that no such effect arises in $\lambda \Pi/S$ because there is no constant $\tilde{\pi}_{s_1 s_2 s_3}$ associated to the type of illegal abstractions.

One could ask whether the techniques we used are adequate. While the construction of λS^* is not absolutely necessary, we feel that it simplifies the proof and that it helps us better understand the behavior of $\lambda \Pi/S$ by reflecting it back into a pure type system. The relative normalization steps of Section 5.3 correspond to the normalization of a simply typed λ calculus. Therefore, it is not surprising that we had to use Tait's reducibility method. However, our proof can be simplified in some cases. A PTS is *complete* when it is a completion of itself. In that case, the construction of S^* is unnecessary. The translations $\varphi(M)$ and $\psi(A)$ translate directly into λS , and Section 5.3 can be omitted. This is the case for example for the calculus of constructions with *infinite type hierarchy* (λC^∞) [17], which is the basis for proof assistants such as Coq and Matita.

The results of this paper can be extended in several directions. They could be adapted to show the conservativity of other embeddings, such as that of the *calculus of inductive constructions* (CIC) [4]. They also indirectly imply that $\lambda \Pi/S$ is weakly normalizing when λS is weakly normalizing because the image of a λS term is normalizing [6]. The strong normalization of $\lambda \Pi/S$ when λS is strongly normalizing is still an open problem. The Barendregt-Geuvers-Klop conjecture states that any weakly normalizing PTS is also strongly normalizing [8]. There is evidence that this conjecture is true [2], in which case we hope that its proof could be adapted to prove the strong normalization of $\lambda \Pi/S$. Weak normalization could also be used as an intermediary step for constructing models by induction on types in order to prove strong normalization.

Acknowledgments. We thank Gilles Dowek and Guillaume Burel for their support and feedback, as well as Frédéric Blanqui, Raphaël Cauderlier, and the various anonymous referees for their comments and suggestions on previous versions of this paper.

References

- 1 H. P. Barendregt. Lambda calculi with types. In *Handbook of Logic in Computer Science*, volume 2. Oxford University Press, 1992.
- 2 Gilles Barthe, John Hatcliff, and Morten Heine Sørensen. Weak normalization implies strong normalization in a class of non-dependent pure type systems. *Theoretical Computer Science*, 269(1-2):317–361, 2001.
- 3 Frédéric Blanqui. Definitions by rewriting in the calculus of constructions. *Mathematical Structures in Computer Science*, 15(01):37–92, 2005.
- 4 M. Boespflug and G. Burel. CoqInE: translating the calculus of inductive constructions into the $\lambda\Pi$ -calculus modulo. In *Proof Exchange for Theorem Proving - Second International Workshop, PxTP 2012*, pages 44–50, 2012.
- 5 M. Boespflug, Q. Carbonneaux, and O. Hermant. The $\lambda\Pi$ -calculus modulo as a universal proof language. In *Proof Exchange for Theorem Proving - Second International Workshop, PxTP 2012*, pages 28–43, 2012.
- 6 Denis Cousineau and Gilles Dowek. Embedding pure type systems in the lambda-Pi-calculus modulo. In Simona Ronchi Della Rocca, editor, *Typed Lambda Calculi and Applications*, number 4583 in Lecture Notes in Computer Science, pages 102–117. Springer Berlin Heidelberg, 2007.
- 7 Gilles Dowek. Models and termination of proof-reduction in the $\lambda\Pi$ -calculus modulo theory. arXiv:1501.06522, hal-01101834, 2014.
- 8 Herman Geuvers. *Logics and type systems*. PhD thesis, University of Nijmegen, 1993.
- 9 Jean-Yves Girard. *Interprétation fonctionnelle et élimination des coupures de l'arithmétique d'ordre supérieur*. Thèse de doctorat, Université Paris VII, 1972.
- 10 Robert Harper, Furio Honsell, and Gordon Plotkin. A framework for defining logics. *J. ACM*, 40(1):143–184, 1993.
- 11 Zhaohui Luo. *Computation and Reasoning: A Type Theory for Computer Science*. Oxford University Press, Inc., New York, NY, USA, 1994.
- 12 Per Martin-Löf and Giovanni Sambin. *Intuitionistic type theory*, volume 17. Bibliopolis Naples, 1984.
- 13 Bengt Nordström, Kent Petersson, and Jan M. Smith. *Programming in Martin-Löf's type theory*, volume 200. Oxford University Press Oxford, 1990.
- 14 Erik Palmgren. On universes in type theory. In *Twenty-five years of constructive type theory*, pages 191–204. Oxford University Press, 1998.
- 15 Cody Roux and Floris van Doorn. The structural theory of pure type systems. In Gilles Dowek, editor, *Rewriting and Typed Lambda Calculi*, number 8560 in Lecture Notes in Computer Science, pages 364–378. Springer International Publishing, 2014.
- 16 Paula Severi. Pure type systems without the Pi-condition. *Proceedings of 7th Nordic Workshop on Programming Theory*, 1995.
- 17 Paula Severi and Erik Poll. Pure type systems with definitions. In Anil Nerode and Yu V. Matiyasevich, editors, *Logical Foundations of Computer Science*, number 813 in Lecture Notes in Computer Science, pages 316–328. Springer Berlin Heidelberg, 1994.
- 18 W. W. Tait. Intensional interpretations of functionals of finite type I. *The Journal of Symbolic Logic*, 32(2):198–212, 1967.

Models for Polymorphism over Physical Dimensions

Robert Atkey¹, Neil Ghani¹, Fredrik Nordvall Forsberg¹,
Timothy Revell¹, and Sam Staton²

¹ University of Strathclyde, UK

² University of Oxford, UK

Abstract

We provide a categorical framework for models of a type theory that has special types for physical quantities. The types are indexed by the physical dimensions that they involve. Fibrations are used to organize this index structure in the models of the type theory. We develop some informative models of this type theory: firstly, a model based on group actions, which captures invariance under scaling, and secondly, a way of constructing new models using relational parametricity.

1998 ACM Subject Classification F.3.2 Semantics of Programming Languages

Keywords and phrases Category Theory, Units of Measure, Dimension Types, Type Theory

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.45

1 Introduction

This paper is about semantic models of programs that manipulate physical quantities. Physical quantities are organized into dimensions, such as *Length* or *Time*. A fundamental principle is that it is not meaningful to add or compare quantities of different dimensions, but they can be multiplied. To measure a physical quantity we use units, such as metres for length and seconds for time. We understand these units as chosen constant quantities of given dimensions.

Here is a simple polymorphic program that is defined for all dimensions; it takes a quantity x of a given dimension X , and returns its double, which has the same dimension.

$$f := (\Lambda X. \lambda x : \text{Quantity}(X). x + x) : \forall X. \text{Quantity}(X) \rightarrow \text{Quantity}(X) \quad (1)$$

To illustrate, we can use the polymorphic function f to double a length of 5 metres.

$$f_{\text{Length}}(5\text{m}) = 10\text{m} : \text{Quantity}(\text{Length}) \quad (2)$$

There are a few key points that are worth emphasising about examples (1) and (2) above:

- There are two kinds of variable, X and x . The first variable X stands for a dimension whereas x stands for an inhabitant of a type. To emphasise this distinction, we use different abstraction symbols (λ and Λ) for the two kinds of variable.
- The type $\text{Quantity}(X)$ depends on a dimension X , and it is inhabited by quantities of that dimension. For example, the standard unit of measurement for length, the metre, is a quantity of that dimension, i.e. a constant $\text{m} : \text{Quantity}(\text{Length})$.

Several authors have developed programming languages with type systems that support physical quantities [8, 12, 17, 10, 6]. The type systems are often motivated as static analyses that help to prevent disasters by accommodating dimensions. For example, the Mars Climate Orbiter was lost as a result of a unit mismatch in the software [13].



© Robert Atkey, Neil Ghani, Fredrik Nordvall Forsberg, Timothy Revell, and Sam Staton;
licensed under Creative Commons License CC-BY

13th International Conference on Typed Lambda Calculi and Applications (TLCA'15).

Editor: Thorsten Altenkirch; pp. 45–59



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Our starting point in this paper is the work of Kennedy [10] who developed techniques for reasoning about these kinds of programs. However, we take a different approach by developing a general categorical notion of model for a programming language of this form, and by developing ways of building models. The main contributions of this paper are:

1. We provide a general notion of a model for a programming language with physical dimension types by introducing the concept of a λD -model (Section 3). The basic idea is that for each context of dimension variables, there is a model of the simply typed λ -calculus extended with types of quantities of the dimensions definable in the context ($\text{Quantity}(D)$ etc.). Moreover these models of the simply typed λ -calculus are related by substituting for dimension variables, and this also defines a universal property for polymorphic quantification over dimension variables.
2. An important example of a λD -model is built from group actions (Example 9). A difficulty with set-theoretic models of dimension polymorphism is as follows: how does one understand $\text{Quantity}(X)$ as a set, if the dimension X is not specified and we have no fixed units of measure for X ? We resolve this by interpreting $\text{Quantity}(X)$ as the set of magnitudes, i.e. positive real numbers, thought of as quantities of some unspecified unit of measure, but then by equipping $\text{Quantity}(X)$ with an action of the scaling group, to explain how to change the units of measure. We can then ask that any function $\text{Quantity}(X) \rightarrow \text{Quantity}(X)$ is invariant under changing that unspecified unit of measure, more precisely, invariant under scaling.
3. We show how the λD -model built from group actions supports a diverse range of parametricity-like theorems, without the need to define a separate relational semantics (Section 4). This results in simple proofs of theorems that would otherwise require more heavy machinery.
4. We explore the relationship between the parametricity-like theorems of the λD -model built from group actions, and a natural notion of a relational model (Section 5). Formally, we show that when interpreting the syntax these two notions coincide.

2 Types with Physical Dimensions

We begin by recalling a simple type theory, which we call λD , indexed by dimensions based on Kennedy's work [10]. Within this type theory we can express programs such as (1) and (2). Since there are two kinds of variable, we have two kinds of context.

Dimensions and Dimension Contexts. A dimension context Δ is a finite list of distinct dimension variables. A dimension-expression-in-context $\Delta \vdash D \text{ dim}$ is a monomial D in the variables Δ . More precisely, if $\Delta = X_1, \dots, X_n$ and $k_i \in \mathbb{Z}$ then $\Delta \vdash X_1^{k_1} \dots X_n^{k_n} \text{ dim}$. We can make the set $\{D \mid \Delta \vdash D \text{ dim}\}$ an Abelian group under addition of exponents, and indeed this is the free Abelian group on Δ . This universal property gives a notion of substitution on dimension expressions. For example, $X, Z \vdash (X^2Y^3)[(XZ^2)/Y] = X^5Z^6 \text{ dim}$.

Types. Well-formed types are given by judgements of the form $\Delta \vdash T \text{ type}$ where Δ is a dimension context. The judgements are generated by the following rules.

$$\frac{\Delta \vdash D \text{ dim}}{\Delta \vdash \text{Quantity}(D) \text{ type}} \quad \frac{\Delta, X \vdash T \text{ type}}{\Delta \vdash \forall X. T \text{ type}} \quad \frac{\Delta \vdash T \text{ type} \quad \Delta \vdash U \text{ type}}{\Delta \vdash T \rightarrow U \text{ type}}$$

$$\frac{}{\Delta \vdash 1 \text{ type}} \quad \frac{\Delta \vdash T \text{ type} \quad \Delta \vdash U \text{ type}}{\Delta \vdash T \times U \text{ type}} \quad \frac{}{\Delta \vdash 0 \text{ type}} \quad \frac{\Delta \vdash T \text{ type} \quad \Delta \vdash U \text{ type}}{\Delta \vdash T + U \text{ type}}$$

Notice that we do not have System-F-style polymorphism, but instead dimension polymorphism: types can be parameterised by dimensions, but they cannot be parameterised by types, since we do not have type variables. From the Curry-Howard perspective this is a first-order-logic where the domain of discourse is the theory of Abelian groups and where there is a single atomic predicate, `Quantity`.

Terms and Typing Contexts. Well-formed typing contexts are given by judgements $\Delta \vdash \Gamma \text{ ctx}$ where Δ is a dimension context, Γ is of the form $x_1 : T_1, \dots, x_n : T_n$ and there is a well-formed typing judgement $\Delta \vdash T_i \text{ type}$ for every i . Well-formed terms are given by judgements $\Delta; \Gamma \vdash t : T$ where there is a well-formed typing context $\Delta \vdash \Gamma \text{ ctx}$ and a well-formed type $\Delta \vdash T \text{ type}$. The rules for the type formers `1`, `_ × _`, `0`, `_ + _` and `_ → _` are the usual ones from simply typed λ -calculus.

$$\frac{\Delta \vdash \Gamma, \Gamma' \text{ ctx} \quad \Delta \vdash T \text{ type}}{\Delta; \Gamma, x : T, \Gamma' \vdash x : T} \quad \frac{\Delta \vdash \Gamma \text{ ctx} \quad (\text{op} : T) \in \text{Ops}}{\Delta; \Gamma \vdash \text{op} : T}$$

$$\frac{\Delta; \Gamma, x : T \vdash t : U}{\Delta; \Gamma \vdash \lambda x. t : T \rightarrow U} \quad \frac{\Delta; \Gamma \vdash t : T \rightarrow U \quad \Delta; \Gamma \vdash u : T}{\Delta; \Gamma \vdash t u : U}$$

$$\frac{\Delta \vdash \Gamma \text{ ctx}}{\Delta; \Gamma \vdash () : 1} \quad \frac{\Delta; \Gamma \vdash t_1 : T_1 \quad \Delta; \Gamma \vdash t_2 : T_2}{\Delta; \Gamma \vdash (t_1, t_2) : T_1 \times T_2} \quad \frac{\Delta; \Gamma \vdash t : T_1 \times T_2}{\Delta; \Gamma \vdash \text{pr}_i(t) : T_i}$$

$$\frac{\Delta; \Gamma \vdash t : 0 \quad \Delta \vdash T \text{ type}}{\Delta; \Gamma \vdash \text{case } t \text{ of } \{ \} : T} \quad \frac{\Delta; \Gamma \vdash t : T_i}{\Delta; \Gamma \vdash \text{inj}_i t : T_1 + T_2}$$

$$\frac{\Delta; \Gamma \vdash t : T_1 + T_2 \quad (\Delta; \Gamma, x_i : T_i \vdash u_i : U)_{i \in \{1,2\}}}{\Delta; \Gamma \vdash \text{case } t \text{ of } \{ \text{inj}_1 x_1 \mapsto u_1; \text{inj}_2 x_2 \mapsto u_2 \} : U}$$

In addition, we have the introduction and elimination rules for quantification over a dimension variable.

$$\frac{\Delta, X; \Gamma \vdash t : T}{\Delta; \Gamma \vdash \Lambda X. t : \forall X. T} \quad \frac{\Delta \vdash D \text{ dim} \quad \Delta; \Gamma \vdash t : \forall X. T}{\Delta; \Gamma \vdash t_D : T[D/X]}$$

We use `Bool` as an abbreviation for `1 + 1`. Our calculus is parameterised by a collection `Ops` of primitive operation typings (`op : Top`) where for each primitive operation `op : Top`, its type `Top` is closed (i.e., $\vdash T_{\text{op}} \text{ type}$). An example set of primitive operations includes dimension-polymorphic arithmetic and test operations on quantities:

$$\begin{aligned} \text{Ops} = & (+ : \forall X. \text{Quantity}(X) \times \text{Quantity}(X) \rightarrow \text{Quantity}(X), \\ & \times : \forall X_1. \forall X_2. \text{Quantity}(X_1) \times \text{Quantity}(X_2) \rightarrow \text{Quantity}(X_1 \cdot X_2), \quad 1 : \text{Quantity}(1), \\ & \text{inv} : \forall X. \text{Quantity}(X) \rightarrow \text{Quantity}(X^{-1}), \\ & < : \forall X. \text{Quantity}(X) \times \text{Quantity}(X) \rightarrow \text{Bool}). \end{aligned}$$

One could also define a type of signed/zero quantities $\text{Real}(X) := \text{Quantity}(X) + 1 + \text{Quantity}(X)$, and then extend the language with further arithmetic term constants such as signed addition $+ : \forall X. \text{Real}(X) \times \text{Real}(X) \rightarrow \text{Real}(X)$.

To write terms that make use of common sets of dimensions and units, we judge terms in a context $(\Delta_{\text{dim}}, \Gamma_{\text{units}})$. For example, $\Delta_{\text{dim}} = (\text{Length}, \text{Time})$ and $\Gamma_{\text{units}} = (\text{m} : \text{Quantity}(\text{Length}), \text{ft} : \text{Quantity}(\text{Length}), \text{s} : \text{Quantity}(\text{Time}))$.

3 Categorical Semantics of Dimension Types

Next up we give a general categorical semantics for the λD type theory. Central to this is the notion of a $\lambda\forall$ -fibration.

► **Definition 1.** A $\lambda\forall$ -fibration is a bicartesian closed fibration with simple products.

It is well-known that $\lambda\forall$ -fibrations give a categorical model of the fragment of first-order logic without existential quantifiers. Nevertheless, we briefly introduce the basic notions now, since they are central to our development. We refer to Jacobs [9] for the full details.

A fibration $p : \mathcal{E} \rightarrow \mathcal{B}$ is a functor between categories satisfying certain conditions. These conditions (along with the structure in Definition 2) allow us to model the λD type theory. The basic idea is that dimension contexts will be interpreted as objects in a category \mathcal{B} . For each $B \in \mathcal{B}$, we consider the fibre \mathcal{E}_B , i.e. the subcategory of \mathcal{E} with objects $E \in \mathcal{E}$ for which $p(E) = B$. The objects of \mathcal{E}_B will be used to interpret types in dimension context B , and the morphisms in \mathcal{E}_B will be used to interpret terms. We can substitute dimension expressions for dimension variables, and this substitution will be interpreted using morphisms in \mathcal{B} . Since p is a fibration, one can form a reindexing functor $f^* : \mathcal{E}_{B'} \rightarrow \mathcal{E}_B$ for each morphism $f : B \rightarrow B'$ in \mathcal{B} , which describes substitution for dimension variables in types and terms.

A fibration is said to be bicartesian closed if \mathcal{E}_B is a Cartesian closed category with coproducts for all B , and each reindexing functor $f^* : \mathcal{E}_{B'} \rightarrow \mathcal{E}_B$ preserves products, exponentials and coproducts. This bicartesian closed structure is needed to interpret the product, function and coproduct types.

Concatenation of dimension contexts will be interpreted using products in the category \mathcal{B} . The reindexing functors $\pi^* : \mathcal{E}_B \rightarrow \mathcal{E}_{B \times B'}$ for the product projections $\pi : B \times B' \rightarrow B$ correspond to context-weakening. A fibration $p : \mathcal{E} \rightarrow \mathcal{B}$ is said to have simple products if \mathcal{B} has products and the reindexing functors for the product projections have right adjoints $\forall : \mathcal{E}_{B \times B'} \rightarrow \mathcal{E}_B$ that are compatible with reindexing (‘Beck-Chevalley’). A fibration is said to have products if this condition holds for all morphisms in the base, not just projections. These right adjoints are needed to interpret universal quantification of dimension variables in types.

► **Definition 2.** A λD -model (p, G, Q) is a $\lambda\forall$ -fibration $p : \mathcal{E} \rightarrow \mathcal{B}$, an Abelian group object G in \mathcal{B} , and an object Q in the fibre \mathcal{E}_G .

Recall that an Abelian group object in a category \mathcal{B} with products is given by an object G together with maps $e : 1 \rightarrow G$, $m : G \times G \rightarrow G$ and $i : G \rightarrow G$ satisfying the laws of Abelian groups. This group structure is needed to interpret dimension expressions: for each vector of n integers we have a morphism $G^n \rightarrow G$.

An equivalent way to define Abelian group objects if \mathcal{B} has chosen products is as follows. Recall that the Lawvere theory for Abelian groups is the category \mathbf{L}_{Ab} whose objects are natural numbers, and where a morphism $m \rightarrow n$ is an $m \times n$ matrix of integers. Composition of morphisms is given by matrix multiplication, and categorical products are given by arithmetic addition of natural numbers. An Abelian group object in \mathcal{B} is an object G of \mathcal{B} together with a strictly-product-preserving functor $F : \mathbf{L}_{\text{Ab}} \rightarrow \mathcal{B}$ such that $F(1) = G$.

We remark that the object Q in a λD -model is analogous to the generic object in a model of System F.

In order to ascertain the value of Definition 2, we now do three things: i) we show that a λD -model in fact does provide categorical models of dimension types, ii) we give examples of λD -models, and iii) we prove theorems that show the viability of reasoning at this level of abstraction.

3.1 Modelling Dimension Types

To show that λD -models provide a categorical semantics for dimension types, we must show how to interpret the syntax given in Section 2 in any given λD -model. We will use the λV -fibration to separate the indexing information (the dimensions) from the indexed information (the types and terms). This means that the base category of the fibration will be used to interpret dimension contexts, and types and terms will be interpreted as objects and morphisms in the fibres above the dimension contexts in which they are defined. Bicartesian closure of the fibres will allow us to inductively interpret types built from 1 , \times , 0 , $+$ and \rightarrow , and we will take the standard approach in categorical logic to interpret quantification of dimensions by using right adjoints. Finally, since dimension expressions for a dimension context are defined as elements of the free Abelian group on that dimension context, we will use the Abelian group object structure to interpret such expressions. Formally, we interpret the syntax as follows.

- Dimension contexts $\Delta = X_1, \dots, X_n$ are interpreted as the product of the Abelian group object $\llbracket \Delta \rrbracket = G^n$ in \mathcal{B} .
- Dimension expressions $\Delta \vdash D \text{ dim}$ are interpreted as morphisms $G^n \rightarrow G$ in the base \mathcal{B} , by using the structure of the Abelian group object G . For example, $\llbracket X, Y \vdash X \cdot Y^{-1} \rrbracket = G \times G \xrightarrow{\text{id}_G \times i} G \times G \xrightarrow{m} G$.
- Well-formed types $\Delta \vdash T \text{ type}$ are interpreted as objects $\llbracket T \rrbracket$ in the fibre above $\llbracket \Delta \rrbracket$, defined by induction on the structure of T . We interpret 1 , \times , $+$ and \rightarrow using the bicartesian closed structure of the fibres, and quantification of a dimension variable $\llbracket \Delta \vdash \forall X. T \rrbracket$ is defined by right adjoint to reindexing along the projection $\pi : \llbracket \Delta \vdash \Gamma, X \rrbracket \rightarrow \llbracket \Delta \vdash \Gamma \rrbracket$. Quantities $\text{Quantity}(D)$ are interpreted by reindexing the object Q along the interpretation of D , i.e. $\llbracket \Delta \vdash \text{Quantity}(D) \rrbracket = \llbracket \Delta \vdash D \text{ dim} \rrbracket^*(Q)$.
- Well-formed typing contexts $\Delta \vdash \Gamma \text{ ctx}$ are interpreted as products in the fibre above $\llbracket \Delta \rrbracket$, i.e. $\llbracket \Delta \vdash x_1 : T_1, \dots, x_n : T_n \rrbracket = \llbracket \Delta \vdash T_1 \rrbracket \times \dots \times \llbracket \Delta \vdash T_n \rrbracket$.
- Well-formed terms $\Delta; \Gamma \vdash t : T$ are interpreted as morphisms $\llbracket t \rrbracket : \llbracket \Gamma \rrbracket \rightarrow \llbracket T \rrbracket$ in the fibre above $\llbracket \Delta \rrbracket$. We assume that there is an interpretation $\llbracket \text{op} \rrbracket : 1 \rightarrow \llbracket T \rrbracket$ for each primitive operation $(\text{op} : T) \in \text{Ops}$.

In this paper we have only considered universal quantification of units but existential quantification can be given just as easily. Existential quantification is interpreted as the left adjoint to reindexing along a projection. Properties of existential quantification can be proven by dualising the relevant proofs of properties about universal quantification.

Care is needed to ensure that the denotational semantics properly respects substitution. This can be done either by requiring that the fibrations are *split*, or by adding explicit coercions to the language [5]. All the examples of fibrations in this paper are split.

3.2 First Examples of λD -Models

We now give some examples of λD -models. We begin by noting that in Kennedy's paper [10], a simpler approach is taken to the semantics of dimensions: the dimensions are simply thrown away in a *dimension-erasure semantics*. From the categorical perspective, this means that the calculus is stripped of its fibred structure leaving only a simply typed λ -calculus which Kennedy models, as is to be expected, within a Cartesian closed category. In particular, he chooses the Cartesian closed category of complete partial orders which he needs for recursion. Nevertheless, Kennedy's model can be viewed as a λD -model.

► **Example 3** (Dimension-Erasure Models). Let \mathcal{C} be a bicartesian closed category. Then the functor $\mathcal{C} \rightarrow \mathbf{1}$ is a $\lambda\forall$ -fibration. The unique object of $\mathbf{1}$ is a trivial Abelian group object. By taking \mathcal{C} to be the category of complete partial orders and continuous functions, and by choosing the flat pointed cpo \mathbb{Q}_\perp to interpret **Quantity** we obtain a model corresponding to Kennedy’s dimension-erasure model. This model supports a plethora of primitive operations, including all the standard arithmetical ones. However, the model also contains many functions which are not dimensionally invariant, i.e. they do not scale appropriately under change of units – Kennedy uses relational parametricity [15] to remove these unwanted elements; we will come back to his relational model in Section 5.

► **Example 4** (Syntactical Models). We can construct a $\lambda\forall$ -fibration $\mathcal{C}\ell(\lambda D)$ from the syntax in a standard way. The base category \mathcal{B} is the Lawvere theory of Abelian groups \mathbf{L}_{Ab} . The fibre $\mathcal{C}\ell(\lambda D)_n$ over n is the category whose objects are types with n dimension variables, and whose morphisms are terms in context, modulo a standard notion of conversion. The object $\mathbf{1}$ in \mathbf{L}_{Ab} is an Abelian group object, and $(\mathcal{C}\ell(\lambda D) \rightarrow \mathbf{L}_{\text{Ab}}, \mathbf{1}, (X \vdash \text{Quantity}(X) \text{ type}))$ is a λD -model.

► **Example 5** (The Dimension-Indexed Families Model). Let $\text{Fam}(\text{Set})$ be the category whose objects are pairs $(I, \{X_i\}_{i \in I})$ of a set I and an I -indexed family of sets $\{X_i\}_{i \in I}$. A morphism $(I, \{X_i\}_{i \in I}) \rightarrow (J, \{Y_j\}_{j \in J})$ is a pair $(f, \{\phi_i\}_{i \in I})$ where f is a function $f : I \rightarrow J$ and ϕ_i is a function $\phi_i : X_i \rightarrow Y_{f(i)}$ for all $i \in I$. It is well known that the forgetful functor $(I, \{X_i\}_{i \in I}) \mapsto I : \text{Fam}(\text{Set}) \rightarrow \text{Set}$, taking a family to its index set, is a $\lambda\forall$ -fibration (see e.g. Jacobs [9, Lemma 1.9.5]). For any given set B of fundamental dimensions (e.g. *Length*, *Time*, *Mass* etc.), let G be the free Abelian group on B . Suppose that we also have a set Q_d of quantities for each dimension $d \in G$ (for instance, we can choose $Q_d = \mathbb{R}^+ \times \{\bar{d}\}$ where \bar{d} is a unit of measure for the dimension d , e.g. $\overline{\text{Length}} = \text{m}$, $\overline{\text{Time}} = \text{s}$, $\overline{d \cdot d'} = \bar{d} \cdot \bar{d}'$ etc). We then have a λD -model with **Quantity** interpreted as $(G, \{Q_d\}_{d \in G})$.

In this model, a dimension expression $X_1, \dots, X_n \vdash D \text{ dim}$ is interpreted as a function $G^n \rightarrow G$ using the free Abelian group structure on G : for each valuation of the dimension variables as physical dimensions, we have an interpretation of the expression as a physical dimension. A type with a free dimension variable $X \vdash T \text{ type}$ is interpreted as a family of sets, indexed by the dimensions in G . Similarly a term with a free dimension variable is interpreted as a family of functions, one for each dimension in G . This model does support many primitive operations, but it does not support dimension invariant polymorphism. For instance, the model supports adding a term $\text{eq} : \forall X_1. \forall X_2. \text{Bool}$ which tests whether two dimensions are the same, which is clearly not invariant under change of representation.

Related examples include the relations fibration $\text{Rel} \rightarrow \text{Set}$ and the subobject fibration $\text{Sub}(\text{Set}) \rightarrow \text{Set}$. This example can also be generalised to the fibration $\text{Fam}(\mathcal{C}) \rightarrow \text{Set}$, which is a $\lambda\forall$ -fibration if \mathcal{C} is bicartesian closed.

A Source of Fibrations with Simple Products

We next look at a particular class of λD models, where the fibres in the $\lambda\forall$ -fibration are functors. We prove a general theorem for such fibrations, and instantiate it to construct several examples. We first introduce some notation. Let \mathcal{S} be a category (typically $\mathcal{S} = \text{Set}$), and consider the category $\text{Cat} // \mathcal{S}$. The objects are pairs $(\mathcal{C}, P : \mathcal{C} \rightarrow \mathcal{S})$, where \mathcal{C} is a small category and $P : \mathcal{C} \rightarrow \mathcal{S}$ a functor. Morphisms $(F, \phi) : (\mathcal{C}, P) \rightarrow (\mathcal{D}, Q)$ are pairs of a functor $F : \mathcal{C} \rightarrow \mathcal{D}$ and a natural transformation $\phi : P \rightarrow Q \circ F$. The obvious projection functor $(\mathcal{C}, P) \mapsto \mathcal{C} : \text{Cat} // \mathcal{S} \rightarrow \text{Cat}$ is a fibration. The fibre over a small category \mathcal{C} is the category

$\mathcal{S}^{\mathcal{C}}$ of functors $[\mathcal{C} \rightarrow \mathcal{S}]$ and natural transformations between them. Reindexing is given by precomposition of functors.

► **Theorem 6.** *If \mathcal{S} has all small limits then the fibration $\text{Cat}/\mathcal{S} \rightarrow \text{Cat}$ has simple products.*

This result appears to be fairly well-known folklore (see e.g. Lawvere [11, end of §3], Melliès and Zeilberger [14]), but since it is important in what follows we sketch a proof.

Proof Sketch. For any functor $F : \mathcal{C} \rightarrow \mathcal{D}$, the reindexing functor $F^* : \mathcal{S}^{\mathcal{D}} \rightarrow \mathcal{S}^{\mathcal{C}}$ has a right adjoint $F_* : \mathcal{S}^{\mathcal{C}} \rightarrow \mathcal{S}^{\mathcal{D}}$, known as the ‘right Kan extension along F ’, which always exists when \mathcal{S} has limits. For simple products, we are only interested in a right adjoint to weakening, i.e. in the functor $\forall_{\mathcal{C}} : \mathcal{S}^{\mathcal{C} \times \mathcal{D}} \rightarrow \mathcal{S}^{\mathcal{C}}$ which is the right Kan extension along the projection functor $\pi_{\mathcal{C}} : \mathcal{C} \times \mathcal{D} \rightarrow \mathcal{C}$. Expanding the definitions, we see that $\forall_{\mathcal{C}}(P) : \mathcal{C} \rightarrow \mathcal{S}$ is a point-wise limit:

$$(\forall_{\mathcal{C}} P)(c) = \lim_{d \in \mathcal{D}} P(c, d) . \quad (3)$$

The Beck-Chevalley condition requires that the canonical map $F^* \forall_{\mathcal{C}'} \rightarrow \forall_{\mathcal{C}}(F \times \text{id}_{\mathcal{D}})^*$ is a natural isomorphism for all functors $F : \mathcal{C} \rightarrow \mathcal{C}'$. Indeed, for any $P : \mathcal{C} \times \mathcal{D} \rightarrow \mathcal{S}$, $c \in \mathcal{C}$:

$$\begin{aligned} (F^*(\forall_{\mathcal{C}'} P))(c) &= (\forall_{\mathcal{C}} P)(F(c)) \cong \lim_{d \in \mathcal{D}} (F(c), d) = \lim_{d \in \mathcal{D}} (((F \times \text{id}_{\mathcal{D}})^*(P))(c, d)) \\ &\cong (\forall_{\mathcal{C}}((F \times \text{id}_{\mathcal{D}})^*(P)))(c) . \quad \blacktriangleleft \end{aligned}$$

A Source of Models by Change of Base

In general, a useful way of building fibrations is by changing the base. If $p : \mathcal{E} \rightarrow \mathcal{B}$ is a fibration, and $F : \mathcal{A} \rightarrow \mathcal{B}$ is a functor, then the pullback of p along F in Cat , denoted $F^*p : F^*\mathcal{E} \rightarrow \mathcal{A}$, is again a fibration. The same is true of λD -models.

► **Theorem 7.** *Let $p : \mathcal{E} \rightarrow \mathcal{B}$ be a fibration, and let $F : \mathcal{A} \rightarrow \mathcal{B}$ be a functor.*

- (i) *If p has simple products and F preserves products, then $F^*p : F^*\mathcal{E} \rightarrow \mathcal{A}$ has simple products.*
- (ii) *If p is bicartesian closed then $F^*p : F^*\mathcal{E} \rightarrow \mathcal{A}$ is bicartesian closed.*
- (iii) *If G is an Abelian group object in \mathcal{A} and $(p, F(G), Q)$ is a λD -model, then also $(F^*p, G, (G, Q))$ is a λD -model.*

Proof. For item (i): for any $A \in \mathcal{A}$, reindexing along a projection $\pi_A : A \times A' \rightarrow A$ in \mathcal{A} is by construction reindexing along $F(\pi_A)$ in \mathcal{B} , which (as F preserves finite products) is the same as reindexing along a projection $\pi_{FA} : FA \times FA' \rightarrow FA$, which has a right adjoint and satisfies the Beck-Chevalley condition, since p has simple products.

For item (ii): F^*p is a bicartesian closed fibration since each fibre $(F^*\mathcal{E})_A$ is by construction of the form \mathcal{E}_{FA} and hence bicartesian closed, and reindexing by f in \mathcal{A} is by construction defined to be reindexing by Ff in \mathcal{B} , which preserves the structure.

Item (iii) is an immediate corollary. ◀

For a simple illustration of the change of base result, notice that the dimension-erasure fibration $\mathcal{C} \rightarrow 1$ arises from pulling back the families fibration $\text{Fam}(\mathcal{C}) \rightarrow \text{Set}$ along the unique product-preserving functor $1 \rightarrow \text{Set}$.

► **Example 8** (Models over the Lawvere theory \mathbf{L}_{Ab}). Let $(p : \mathcal{E} \rightarrow \mathcal{B}, G, Q)$ be a λD -model. Recall that the Abelian group object G in \mathcal{B} gives rise to a unique product-preserving functor $F : \mathbf{L}_{\text{Ab}} \rightarrow \mathcal{B}$ such that $F(1) = G$. By Theorem 7, we have a λD -model $(F^*p, 1, (1, Q))$.

► **Example 9** (A Model Built from Group Actions). Let G be a group. Recall that a G -set consists of a set A together with a group action, i.e. a function $\cdot_A : G \times A \rightarrow A$ such that $e \cdot_A a = a$ and $(gh) \cdot_A a = g \cdot_A (h \cdot_A a)$. The category $\mathbf{Grp//Set}$ has as objects pairs (G, A) where G is a group and A is a G -set. A morphism $(G, A) \rightarrow (H, B)$ in $\mathbf{Grp//Set}$ is given by a group homomorphism $\phi : G \rightarrow H$ and a function $f : A \rightarrow B$ such that for any $g \in G$ and $a \in A$ we have $f(g \cdot_A a) = (\phi g) \cdot_B (f a)$. Let \mathbf{Grp} be the category of groups and homomorphisms. We call the forgetful functor $p : \mathbf{Grp//Set} \rightarrow \mathbf{Grp}$ the $\mathbf{Grp//Set}$ fibration.

► **Proposition 10.** *Let G be an Abelian group, and let Q be a G -set. Then $(p : \mathbf{Grp//Set} \rightarrow \mathbf{Grp}, G, Q)$ is a λD -model.*

Proof. For any group H , the fibre above H is the category of H -sets and equivariant functions. This is isomorphic to the functor category \mathbf{Set}^H , where we consider the group H as a category with one object \star and a morphism for each element of H . Indeed, there is a product-preserving, full and faithful functor $\mathbf{Grp} \rightarrow \mathbf{Cat}$, taking a group to the corresponding one-object category. The fibration $\mathbf{Grp//Set} \rightarrow \mathbf{Grp}$ is thus the pullback of the fibration $\mathbf{Cat//Set} \rightarrow \mathbf{Cat}$ along this embedding $\mathbf{Grp} \rightarrow \mathbf{Cat}$. Thus, by Theorem 6 and Theorem 7(i), $\mathbf{Grp//Set} \rightarrow \mathbf{Grp}$ has simple products.

Each fibre is bicartesian closed. Products and coproducts are inherited from \mathbf{Set} . For the function space, let A and B be G -sets; then the set of functions $(A \rightarrow B)$ is also a G -set, with action given by $(g \cdot_{(A \rightarrow B)} f)(x) := g \cdot_B (f(g^{-1} \cdot_A x))$. It follows that reindexing preserves the bicartesian closed structure. (This is not the case more generally for $\mathbf{Cat//Set} \rightarrow \mathbf{Cat}$, so Theorem 7(ii) does not apply.) Finally, an Abelian group object in \mathbf{Grp} is the same thing as an Abelian group. Hence (p, G, Q) is a λD -model. ◀

In the $\mathbf{Grp//Set}$ fibration, the Abelian group G can be thought of as a group of scaling factors, and the G -set Q is a set of quantities together with a scaling action. For instance, let $Q = G = (\mathbb{R}^+, \times, 1)$, the positive reals. We model a type with a free dimension variable $X \vdash T$ type as a G -set. A term with a free dimension variable is interpreted as a function that is invariant under G . We explore this model in more detail in Section 4.

The $\mathbf{Grp//Set}$ model supports several primitive operations, which we discuss after Theorem 13.

More generally, instead of having sets and group actions, we also have λD -models built from actions of groupoids.

► **Example 11** (A Model Built from Groupoid Actions). Recall that a *groupoid* is a small category \mathcal{C} where every morphism is an isomorphism, and that a functor $\mathcal{C} \rightarrow \mathbf{Set}$ is called a groupoid action (or presheaf). The category $\mathbf{Gpd//Set}$ has as objects pairs (\mathcal{A}, ϕ) where \mathcal{A} is a groupoid and $\phi : \mathcal{A} \rightarrow \mathbf{Set}$ is a functor. A morphism $(\mathcal{A}, \phi) \rightarrow (\mathcal{B}, \psi)$ in $\mathbf{Gpd//Set}$ is given by a functor $F : \mathcal{A} \rightarrow \mathcal{B}$ and a natural transformation $\eta : \phi \rightarrow \psi \circ F$ between functors $\mathcal{A} \rightarrow \mathbf{Set}$. Let \mathbf{Gpd} be the category of groupoids and functors. Then the forgetful functor $\mathbf{Gpd//Set} \rightarrow \mathbf{Gpd}$, which we call the $\mathbf{Gpd//Set}$ fibration, is a $\lambda \forall$ -fibration, and the proof of this is very similar to the proof in Example 9.

On theme with this subsection, the $\mathbf{Gpd//Set}$ fibration is related to the other fibrations by change of base:

- The families fibration $\mathbf{Fam}(\mathbf{Set}) \rightarrow \mathbf{Set}$ from Example 5 arises from pulling back the groupoid action fibration $\mathbf{Gpd//Set} \rightarrow \mathbf{Gpd}$ along the discrete-groupoid-functor $\mathbf{Set} \rightarrow \mathbf{Gpd}$.
- The group action fibration $\mathbf{Grp//Set} \rightarrow \mathbf{Grp}$ from Example 9 arises from pulling back the groupoid action fibration $\mathbf{Gpd//Set} \rightarrow \mathbf{Gpd}$ along the functor $\mathbf{Grp} \rightarrow \mathbf{Gpd}$ that regards each group as a groupoid with one object.

We now discuss λD -models in $\mathbf{Gpd//Set}$. Let $f : G \rightarrow H$ be a homomorphism of Abelian groups. This induces a groupoid whose objects are the elements of H , and where the hom-sets are $\text{mor}(h, h') = \{g \in G \mid f(g) \cdot_H h = h'\}$. The group operation in G provides composition of morphisms. This groupoid can be given the structure of an Abelian group object in \mathbf{Gpd} , and, moreover, every Abelian group in \mathbf{Gpd} arises in this way [3].

We have already seen that the $\mathbf{Gpd//Set}$ fibration subsumes the families and group actions fibrations. It also subsumes them as λD -models. To recover group actions (Example 9), let G be an Abelian group of scale factors. The Abelian group object induced by the unique homomorphism $G \rightarrow 1$ is a one-object groupoid, and hence we build the λD -models of group actions. To recover the families example (Example 5), fix a set of dimension constants and let H be the free Abelian group on that set. The unique homomorphism $1 \rightarrow H$ induces the discrete groupoid whose objects are H , and hence we build the λD -models of families of sets.

4 Group Actions and Dimension Types

In this section we will look in greater detail at the λD -model given by the $\mathbf{Grp//Set}$ fibration. It turns out that many interesting theorems can be proven in this model, and so to aid us in this task we first concretely spell out the reindexing and simple product structure.

Given a group G , we write \mathcal{G} (with a different font) for the corresponding one-element category, which has morphisms given by elements of G and composition given by group multiplication. Suppose that $\phi : \mathcal{G} \rightarrow \mathbf{Set}$ is a G -set (considered as a functor). We write $|\phi| := \phi(\star)$ for the underlying carrier set of ϕ . Reindexing along $\pi : \mathcal{G} \times \mathcal{H} \rightarrow \mathcal{G}$ yields the $G \times H$ -set given by $\phi \circ \pi$. In other words, $\pi^*\phi$ is a $G \times H$ -set with the same underlying carrier $|\pi^*\phi| = |\phi|$ as the G -set ϕ , and action given by $(g, h) \cdot_{\pi^*\phi} x = g \cdot_{\phi} x$.

Now suppose that $\psi : \mathcal{G} \times \mathcal{H} \rightarrow \mathbf{Set}$ is a $G \times H$ -set. According to Theorem 6 (equation (3)), the underlying set of $\forall_{\pi}\psi$ is given by $|\forall_{\pi}\psi| = \lim_{y \in \mathcal{H}} \psi(\star, y)$. By the universal property of limits,

$$\lim_{y \in \mathcal{H}} \psi(\star, y) \cong \mathbf{Set}(1, \lim_{y \in \mathcal{H}} \psi(\star, y)) \cong [\mathcal{H}, \mathbf{Set}](K1, \psi(\star, _)) ,$$

hence $|\forall_{\pi}\psi| = \{y \in |\psi| \mid \forall h \in H . (e_A, h) \cdot_{\psi} y = y\}$, and the action is given by $g \cdot_{\forall_{\pi}\psi} x = (g, e_H) \cdot_{\psi} x$. Notice that to give the group action of $\forall_{\pi}\psi$, we had to make a particular choice of an element in H , namely the identity element e_H . However, any element of H would have given the same result, since for all $y \in |\forall_{\pi}\psi|$,

$$(g, h) \cdot_{\psi} y = ((g, e_H)(e_G, h)) \cdot_{\psi} y = (g, e_H) \cdot_{\psi} ((e_G, h) \cdot_{\psi} y) = (g, e_H) \cdot_{\psi} y .$$

Many of the properties of dimension types that Kennedy proves using parametricity can be shown to hold in the $\mathbf{Grp//Set}$ -fibration, without having to define a separate relational semantics and this is the content of Theorems 13–18. Before we formally state and prove these we introduce a substitution lemma, which holds in any model.

► **Lemma 12** (Substitution Lemma). *Suppose that $\Delta, X \vdash T$ type and that $\Delta \vdash D$ dim denotes a dimension expression. Then $\llbracket T[D/X] \rrbracket \cong (\text{id}_{\llbracket \Delta \rrbracket}, \llbracket D \rrbracket)^* \llbracket T \rrbracket$.*

Proof. By induction on the structure of T . ◀

Explicitly, Lemma 12 says that the semantics of substituting a dimension expression for a dimension variable is given by reindexing along the identity paired with the dimension expression. Since reindexing is given by precomposition we have that

$$(\text{id}_{\llbracket \Delta \rrbracket}, \llbracket D \rrbracket)^* \llbracket T \rrbracket \cong \llbracket T \rrbracket (\text{id}_{\llbracket \Delta \rrbracket}, \llbracket D \rrbracket) ,$$

i.e., substitution of the n^{th} dimension variable is given by precomposition at the n^{th} component.

For the rest of this section, we will use semantic brackets $\llbracket _ \rrbracket$ to refer only to the Grp//Set interpretation.

► **Theorem 13.** *Suppose that $X_1, \dots, X_n, X \vdash S, T$ type. Then*

$$\llbracket \forall X. S \rightarrow T \rrbracket \cong [\mathcal{G}, \text{Set}](\underbrace{\llbracket S \rrbracket(\star, \dots, \star, _)}_{n \text{ times}}, \underbrace{\llbracket T \rrbracket(\star, \dots, \star, _)}_{n \text{ times}})$$

Proof. By the Kan extension formula and Yoneda. ◀

This theorem says that in the Grp//Set model a universally quantified variable over an arrow type can be considered as a natural transformation between the domain and codomain of the arrow type, with the first n components fixed. In other words, it is interpreted as the set of functions that are equivariant in the last argument.

In particular, if $X_1 \dots X_n \vdash S, T$ type then the type $(\forall \vec{X}. S \rightarrow T)$ is interpreted as the set of all homomorphisms $[\mathcal{G}^n, \text{Set}](\llbracket S \rrbracket, \llbracket T \rrbracket)$. We use this fact to conclude that the group actions model supports several primitive operations. For any $q \in Q$, we can accommodate a term constant $q : \text{Quantity}(1)$, which is interpreted by $\llbracket q \rrbracket = q$. When $Q = G$, we can also accommodate a term constant for multiplication

$$\times : \forall X. \forall Y. \text{Quantity}(X) \times \text{Quantity}(Y) \rightarrow \text{Quantity}(X \cdot Y)$$

which is interpreted as the group operation. When $Q = G = (\mathbb{R}^+, \times, 1)$, the positive reals, we also have addition, $+$: $\forall X. \text{Quantity}(X) \times \text{Quantity}(X) \rightarrow \text{Quantity}(X)$, which is equivariant since $q(r + s) = qr + qs$.

► **Theorem 14.** *Suppose that $\Delta, X \vdash T$ type. Then $\llbracket \forall X. \text{Quantity}(X) \rightarrow T \rrbracket \cong \llbracket T[1/X] \rrbracket$.*

Proof. By Theorem 13, Lemma 12 and Yoneda. ◀

We now prove some theorems about the Grp//Set fibration that are parametricity results in Kennedy's original paper. The proofs here involve applications of Lemma 12, Theorem 13 and Theorem 14. First, we take a look at the interplay between scaling factors and polymorphic functions.

► **Theorem 15 (Scaling Factors).** *Suppose $\Delta_{\text{dim}}; \Gamma_{\text{units}} \vdash t : \forall X. \text{Quantity}(X) \rightarrow \text{Quantity}(X^n)$, where $n \in \mathbb{N}$. Then for all $g \in G$ and $x \in \llbracket \text{Quantity}(X) \rrbracket$, we have $\llbracket t \rrbracket(g \cdot x) = g^n \cdot \llbracket t \rrbracket x$.*

Proof. We know from Theorem 13 that $\llbracket t \rrbracket \in [\mathcal{G}, \text{Set}](\llbracket \text{Quantity}(X) \rrbracket, \llbracket \text{Quantity}(X^n) \rrbracket)$. In other words, $\llbracket t \rrbracket(g \cdot x) = g^n \cdot \llbracket t \rrbracket x$ for all $x \in \llbracket \text{Quantity}(X) \rrbracket$, as required. ◀

This theorem tells us that polymorphic functions are *invariant under scaling*. Intuitively we see that scaling factors must be changed in an appropriately polymorphic way. If we apply Theorem 14 to the type $\forall X. \text{Quantity}(X) \rightarrow \text{Quantity}(X^n)$, we see that

$$\llbracket \forall X. \text{Quantity}(X) \rightarrow \text{Quantity}(X^n) \rrbracket \cong \llbracket \llbracket \text{Quantity}(1^n) \rrbracket \rrbracket \cong \llbracket \llbracket \text{Quantity}(1) \rrbracket \rrbracket \cong Q,$$

Putting $Q = G$, we conclude that all the terms of type $\forall X. \text{Quantity}(X) \rightarrow \text{Quantity}(X^n)$ are of the form $\lambda X. \lambda q : \text{Quantity}(X). r \times q^n$ for $r \in G$.

► **Theorem 16.** *There is no ground term $\vdash t : \forall X. \text{Quantity}(X^2) \rightarrow \text{Quantity}(X)$, i.e., we cannot write a polymorphic square root function.*

Proof. To see this we exhibit a model where the existence of such a term is impossible. Consider the λD -model $(p : \text{Grp} // \text{Set} \rightarrow \text{Grp}, \mathbb{Z}_2, \mathbb{Z}_2)$ where the Abelian group $\mathbb{Z}_2 = (\{-1, 1\}, \cdot, 1)$ is used to interpret both dimensions and units. Theorem 13 says that the interpretation of the type $\forall X. \text{Quantity}(X^2) \rightarrow \text{Quantity}(X)$ is given by

$$|\llbracket \forall X. \text{Quantity}(X^2) \rightarrow \text{Quantity}(X) \rrbracket| \cong [\mathbb{Z}_2, \text{Set}](|\llbracket \text{Quantity}(X^2) \rrbracket|, |\llbracket \text{Quantity}(X) \rrbracket|)$$

i.e. any element f of $|\llbracket \forall X. \text{Quantity}(X^2) \rightarrow \text{Quantity}(X) \rrbracket|$, satisfies for all $g, x \in \mathbb{Z}_2$

$$f(g^2 \cdot x) = g \cdot (fx) \quad (*)$$

If f exists, then either $f(-1) = -1$ or $f(-1) = 1$, but both lead to contradictions. To this end suppose that $f(-1) = -1$, then by $(*)$ we have $f((-1)^2 \cdot -1) = (-1) \cdot f(-1)$, which is a contradiction since the left-hand side is equal to -1 and the right-hand side is equal to 1 . A similar argument shows that $f(-1) = 1$ is also not possible, and hence there exists no such f . \blacktriangleleft

This result can be extended to also include terms t using primitive operations Ops , as long as these operations can be interpreted in the model in question. For example, the result holds in the presence of multiplication

$$\times : \forall X. \forall Y. \text{Quantity}(X) \times \text{Quantity}(Y) \rightarrow \text{Quantity}(X \cdot Y).$$

However, this model does not support a polymorphic zero constant $0 : \forall X. \text{Quantity}(X)$, as such a primitive would of course give rise to a trivial counterexample to the theorem.

Next, we can prove a theorem that relates a dimensionally invariant function to a dimensionless one. This is a simplified version of the *Buckingham Pi Theorem* of dimensional analysis [4] (for a more modern introduction, see Sonin [16]).

► **Theorem 17.** *We have a bijection*

$$|\llbracket \forall X. \text{Quantity}(X) \times \text{Quantity}(X) \rightarrow \text{Quantity}(1) \rrbracket| \cong |\llbracket \text{Quantity}(1) \rightarrow \text{Quantity}(1) \rrbracket|$$

Proof. This is a consequence of Theorem 14, after currying. \blacktriangleleft

We finish this section with another uninhabitedness result, this time about a higher order type.

► **Theorem 18.** *There is no term*

$$\vdash t : \forall X_1. \forall X_2. (\text{Quantity}(X_1) \rightarrow \text{Quantity}(X_2)) \rightarrow \text{Quantity}(X_1 \cdot X_2) .$$

Proof. Choose G and Q to be \mathbb{Z}_2 . Interpreting the type of t , we have

$$\begin{aligned} & |\llbracket \forall X_1. \forall X_2. (\text{Quantity}(X_1) \rightarrow \text{Quantity}(X_2)) \rightarrow \text{Quantity}(X_1 \cdot X_2) \rrbracket| \\ &= \{t \in (\mathbb{Z}_2 \rightarrow \mathbb{Z}_2) \rightarrow \mathbb{Z}_2 \mid \\ & \quad \forall g_1, g_2 \in \mathbb{Z}_2, f : \mathbb{Z}_2 \rightarrow \mathbb{Z}_2. (g_1 g_2) \cdot (t(f)) = t(\lambda q \in \mathbb{Z}_2. g_2 \cdot (f(g_1^{-1} \cdot q)))\} \end{aligned}$$

Hence for any $t \in |\llbracket \forall X_1. \forall X_2. (\text{Quantity}(X_1) \rightarrow \text{Quantity}(X_2)) \rightarrow \text{Quantity}(X_1 \cdot X_2) \rrbracket|$, instantiating $f = \text{id}_Q$ we get that $(g_1 g_2) \cdot (t(\text{id}_Q)) = t(\lambda q \in \mathbb{Z}_2. g_2 \cdot (g_1^{-1} \cdot q))$ for all g_1 and g_2 , but this is not possible. If $g_1 = 1$ and $g_2 = -1$, then the equation reduces to $-1 \cdot t(\text{id}_Q) = t(\text{id}_Q)$, which is a contradiction since $t(\text{id}_Q) \in \mathbb{Z}_2 = \{-1, 1\}$. \blacktriangleleft

Again, the result can be extended to terms that use a set of primitive operations Ops , as long as all primitive operations in Ops can be interpreted in the model used in the proof.

5 Relational Models

In Section 4 we pointed out that many of the results that we proved in the $\mathbf{Grp//Set}$ λD -model are results that Kennedy [10] proves using parametricity. It is curious how the parametricity-style proofs in the $\mathbf{Grp//Set}$ λD -model are simple and slick and do not require a separate relational semantics. One cannot help but wonder, is the $\mathbf{Grp//Set}$ λD -model really as good as having full-blown parametricity at one's finger tips?

To answer this question we look at a general method of attaching a (fibrational) logic to a λD -model to give a notion of a *relational* λD -model. This allows us to reconstruct Kennedy's relational parametricity in our setting (Example 21), as well as to talk about a relational version of the $\mathbf{Grp//Set}$ λD -model (Example 22).

To begin this section, we first recall a theorem about the composition of fibred structure.

► **Theorem 19.** *Suppose that $p : \mathcal{A} \rightarrow \mathcal{B}$ and $q : \mathcal{B} \rightarrow \mathcal{C}$ are fibrations and let $u : \mathcal{A} \rightarrow \mathcal{C}$ denote the composite $q \circ p$ (hence u is also a fibration). Suppose further that q has simple products. For any projection map $\pi_{q(B)} : q(B) \times Y \rightarrow q(B)$ in \mathcal{C} , denote the Cartesian morphism in \mathcal{B} above it by $\pi_{q(B)}^{\S} : \pi^* B \rightarrow B$. Then u has simple products that are preserved by p if and only if for any projection map $\pi : q(B) \times Y \rightarrow q(B)$ in \mathcal{C} , the functor $(\pi_{q(B)}^{\S})^* : \mathcal{A}_B \rightarrow \mathcal{A}_{\pi^* B}$ has right adjoints for all $B \in \mathcal{B}$, satisfying the Beck-Chevalley condition.*

$$\begin{array}{ccc} \mathcal{A} & \xrightarrow{p} & \mathcal{B} \\ & \searrow u & \downarrow q \\ & & \mathcal{C} \end{array}$$

Proof. This theorem is proven by using the factorisation and lifting properties of the 2-category \mathbf{Fib} as outlined by Hermida [7]. Though the proof is not too difficult, it does require 2-categorical technology, which we do not introduce here. Hence, we leave the proof as an exercise for the 2-category-savvy reader. ◀

We now put this theorem to use. Given a λD -model $q : \mathcal{A} \rightarrow \mathcal{L}$ and a logic $p : \mathcal{E} \rightarrow \mathcal{B}$, there is a natural way to glue them together to provide a relational semantics.

► **Theorem 20.** *Let $(q : \mathcal{A} \rightarrow \mathcal{L}, G, Q_0)$ be a λD -model, $F : \mathcal{A} \rightarrow \mathcal{B}$ a product preserving functor and $p : \mathcal{E} \rightarrow \mathcal{B}$ a bicartesian closed fibration with products. Consider the pullback of p along F , and let Q_R denote an object in the fibre $\mathcal{E}_{F(Q_0)}$. Then $(q \circ F^*p : F^*\mathcal{E} \rightarrow \mathcal{L}, G, (Q_0, Q_R))$ is a λD -model.*

$$\begin{array}{ccc} F^*(\mathcal{E}) & \xrightarrow{\quad} & \mathcal{E} \\ F^*p \downarrow & \lrcorner & \downarrow p \\ \mathcal{A} & \xrightarrow{F} & \mathcal{B} \\ q \downarrow & & \\ \mathcal{L} & & \end{array}$$

Proof. Clearly G is an Abelian group object in \mathcal{L} , and (Q_0, Q_R) is in the fibre $(F^*\mathcal{E})_G$. To check that $q \circ F^*p$ is a bicartesian closed fibration is a simple exercise. Finally, since p has all products, so does F^*p . Hence, $q \circ F^*p$ has simple products by Theorem 19. ◀

Next, we look at an example that uses Theorem 20 to generate Kennedy's original relationally parametric model of dimension types [10] from essentially the dimension-erasure model back in Example 3.

► **Example 21.** Let G be an Abelian group. Then using the notation from Theorem 20, let \mathcal{L} be the Lawvere theory of Abelian groups \mathbf{L}_{Ab} , \mathcal{A} be the category $\mathbf{L}_{\text{Ab}} \times \mathbf{Set}$, $q : \mathbf{L}_{\text{Ab}} \times \mathbf{Set} \rightarrow \mathbf{L}_{\text{Ab}}$ be the fibration given by the first projection, and $p : \mathbf{Sub}(\mathbf{Set}) \rightarrow \mathbf{Set}$ be the subset fibration. Define $F : \mathbf{L}_{\text{Ab}} \times \mathbf{Set} \rightarrow \mathbf{Set}$ to be the product preserving functor defined on objects

$(n, X) \in \mathbf{L}_{\text{Ab}} \times \mathbf{Set}$ by $F(n, X) = G^n \times X \times X$, and on morphisms $(f, g) : (n, X) \rightarrow (m, Y)$ by $F(f, g) = (G^f, g, g)$. Finally, let $Q_0 = G$, and $Q_R = \{(g, g_1, g_2) \mid gg_1 = g_2\} \subseteq G \times G \times G$.

In this model, each type $\Delta \vdash T$ is interpreted as a triple $(|\Delta|, \llbracket T \rrbracket_o, \llbracket T \rrbracket_r) \in \mathbf{L}_{\text{Ab}} \times \mathbf{Set} \times \mathbf{Sub}(\mathbf{Set})$, where $\llbracket T \rrbracket_r \subseteq G^n \times \llbracket T \rrbracket_o \times \llbracket T \rrbracket_o$. Spelling this out explicitly, we have the following interpretations, which are equivalent to Kennedy's original relationally parametric model for dimension types:

$$\begin{aligned} \llbracket \Delta \vdash \text{Quantity}(D) \rrbracket &= (|\Delta|, G, \{(g, g_1, g_2) \mid (\llbracket D \rrbracket g)g_1 = g_2\}) \\ \llbracket \Delta \vdash T \times U \rrbracket &= (|\Delta|, \llbracket T \rrbracket_o \times \llbracket U \rrbracket_o, \\ &\quad \{(g, (t_1, u_1), (t_2, u_2)) \mid (g, t_1, t_2) \in \llbracket T \rrbracket_r, (g, u_1, u_2) \in \llbracket U \rrbracket_r\}) \\ \llbracket \Delta \vdash T + U \rrbracket &= (|\Delta|, \llbracket T \rrbracket_o + \llbracket U \rrbracket_o, \\ &\quad \{(g, \text{inj}_1 t, \text{inj}_1 t') \mid (g, t, t') \in \llbracket T \rrbracket_r\} \\ &\quad \cup \{(g, \text{inj}_2 u, \text{inj}_2 u') \mid (g, u, u') \in \llbracket U \rrbracket_r\}) \\ \llbracket \Delta \vdash T \rightarrow U \rrbracket &= (|\Delta|, \llbracket T \rrbracket_o \rightarrow \llbracket U \rrbracket_o, \\ &\quad \{(g, f_1, f_2) \mid \forall t_1, t_2. (g, t_1, t_2) \in \llbracket T \rrbracket_r \implies (g, f_1 t_1, f_2 t_2) \in \llbracket U \rrbracket_r\}) \\ \llbracket \Delta \vdash \forall X. T \rrbracket &= (|\Delta|, \llbracket T \rrbracket_o, \{(g, t_1, t_2) \mid \forall g' \in G. ((g, g'), t_1, t_2) \in \llbracket T \rrbracket_r\}) \end{aligned}$$

Note that in the interpretation of $\forall X. T$, the ‘‘carrier’’ (i.e., the second component) is exactly the carrier of the interpretation of T .

We can also apply Theorem 20 to obtain a natural relational model for the $\mathbf{Grp}/\mathbf{Set}$ λD -model (Example 9).

► **Example 22.** As before, let G be an Abelian group and \mathcal{L} be the Lawvere theory of Abelian groups \mathbf{L}_{Ab} . Let $q : \mathcal{A} \rightarrow \mathbf{L}_{\text{Ab}}$ be the pullback of the fibration $\mathbf{Grp}/\mathbf{Set} \rightarrow \mathbf{Grp}$ along the unique product-preserving functor $M : \mathbf{L}_{\text{Ab}} \rightarrow \mathbf{Grp}$ with $M(1) = G$, as in Example 8, so that the objects of \mathcal{A} are triples (n, X, ϕ) with (X, ϕ) a G^n -set. Let $p : \mathbf{Sub}(\mathbf{Set}) \rightarrow \mathbf{Set}$ be the subset fibration. Define $F : \mathcal{A} \rightarrow \mathbf{Set}$ to be the product preserving functor defined on objects by $F(n, X, \phi) = G^n \times X \times X$ and on morphisms $(f, \alpha) : (n, X, \phi) \rightarrow (m, Y, \psi)$ by $F(f, \alpha) = (\alpha, f, f)$. Finally, we let $Q_0 = (G, \phi)$, where ϕ denotes group multiplication, and $Q_R = \{(g, g_1, g_2) \mid gg_1 = g_2\} \subseteq G \times G \times G$.

Then each type $\Delta \vdash T$ is again interpreted as a triple $(|\Delta|, \llbracket T \rrbracket_o, \llbracket T \rrbracket_r) \in \mathbf{L}_{\text{Ab}} \times \mathbf{Sub}(\mathbf{Set})$, with $\llbracket T \rrbracket_r \subseteq G^n \times \llbracket T \rrbracket_o \times \llbracket T \rrbracket_o$. The only difference between the interpretation of types in this example and Example 21 is the second component of the interpretation of dimension quantification:

$$\begin{aligned} \llbracket \Delta \vdash \forall X. T \rrbracket_r &= (|\Delta|, \{t \in \llbracket T \rrbracket_r \mid \forall g \in G. ((e_{G^{|\Delta|}}, g), t, t) \in \llbracket T \rrbracket_r\}, \\ &\quad \{(g, t_1, t_2) \mid \forall g' \in G. ((g, g'), t_1, t_2) \in \llbracket T \rrbracket_r\}) \end{aligned}$$

This interpretation, in contrast to the interpretation in Example 21, has ‘‘cut-down’’ the carrier of the interpretation of \forall -types to only include the ‘‘parametric’’ elements. As a consequence, this interpretation satisfies an analogue of the *Identity Extension* lemma from relationally parametric models of System F [15].

► **Proposition 23.** *For all type interpretations $(|\Delta|, \llbracket T \rrbracket_o, \llbracket T \rrbracket_r)$, we have:*

$$\forall x_1, x_2 \in \llbracket T \rrbracket_o. (e, x_1, x_2) \in \llbracket T \rrbracket_r \Leftrightarrow x_1 = x_2$$

Compare this to the identity extension property for System F models, which states that if we instantiate the relational interpretation of a type with the equality relation for all of its free variables, then the resulting relation is the equality relation. In the current setting, equality relations for the free variables are replaced by the unit element of the groups $G^{|\Delta|}$. Indeed, this model is equivalent to the restriction to one-dimensional scalings of the reflexive graph model for System F ω with geometric symmetries presented by Atkey [1].

We end this discussion of relational models by showing the relationships between the models in Examples 21 and 22 and the Grp//Set model we considered in detail in Section 4. By construction, the carriers of the interpretations of each type in the model in Example 22 and the Grp//Set model are identical. Moreover, the relational interpretation in Example 22 and the group action in the Grp//Set model are related as follows.

► **Theorem 24.** *For all types $\Delta \vdash T$ type, if the interpretation of T in the model of Example 22 is $(|\Delta|, A, P \subseteq G^{|\Delta|} \times A \times A)$ and the Grp//Set model interpretation is (G^n, A, ψ) , then $(g, a_1, a_2) \in P \Leftrightarrow g \cdot_\psi a_1 = a_2$.*

Proof. By induction on the derivation of $\Delta \vdash T$ type. ◀

Using Theorem 24, we can see that we could have used the relationally parametric model to derive the results in Section 4. There is literally no difference between the two models for the purposes of interpreting the types of our calculus.

We can also relate the relationally parametric model from Example 22 to the dimension-erasure semantics in Example 3. By constructing a logical relation between the two models, we can show:

► **Theorem 25.** *For any closed term $\vdash t : \text{Bool}$, the interpretation of t in the dimension-erasure model of Example 3 is equal to the interpretation of t in the relationally parametric model of Example 22.*

By the compositionality of both interpretations, this theorem means that if we can show that two open terms s and t are equal in the model of Example 22 (and equivalently, the Grp//Set model), then they will be contextually equivalent for the dimension erasure model.

It remains to discuss the relationship between Kennedy’s original relational model (Example 21), and the relational model in Example 22 that satisfies the identity extension property. As noted above, the difference between these interpretations lies in the semantics of the \forall -type. Kennedy’s model does not restrict the carrier of the interpretation to just the “parametric” elements, i.e., the elements that preserve all relations. Therefore, the interpretations of types that contain nested \forall s are not directly comparable. We might expect that we could observe a difference between the two models when proving statements about terms whose types contain negatively nested \forall -types. However, Kennedy’s original work does not present any results involving terms with such types, and we have not found any natural examples. This is in contrast with the situation with relationally parametric models of System F, where the proof that final coalgebras can be represented crucially relies on the restriction of the interpretation of quantified types to the parametric elements [2].

Therefore, our Grp//Set model and the equivalent relational model in Example 22 practically coincides with Kennedy’s original model, but offer the advantage of not requiring a separate relational semantics to prove important theorems. This in many cases makes proofs of these theorems clearer. Additionally, the Grp//Set model offers an interpretation that directly links the semantics to symmetry.

6 Concluding Remarks

To conclude, we have studied a typed λ -calculus with polymorphism over physical dimensions, which we called λD (Section 2) and we have developed a model theory for the calculus. Under the Curry-Howard correspondence, the λD -calculus is a fragment of first-order logic where the domain of discourse is an unspecified Abelian group, and so our notion of model (Definition 2) is based on the standard fibrational techniques in categorical logic.

One particular model turned out to be particularly straightforward and yet informative – the model based on group actions (Example 9). Of course, automorphisms and group actions play a key role in the classical model theory of first order logic, but in this paper we have shown that these techniques are also useful on the other side of the Curry-Howard correspondence. Many arguments about the λD -calculus, including type isomorphisms and definability arguments, can be made in this model (Section 4).

Parametricity is most often studied using relational techniques, and in this paper we have developed a method for building relational λD -models (Theorem 20). Using this method we were able to reconstruct two particular relational models: a relational model due to Kennedy (Example 21, [10]) and a restriction of a relational model due to Atkey (Example 22, [1]). Although the group-actions model is different in style, we showed (formally) that it is actually closely related to the two relational models (Theorems 24 and 25).

Acknowledgements. We are grateful to participants in a discussion on the categories mailing list about the `Grp/Set` fibration. Research supported by the EPSRC Grant EP/K023837/1 (NG, FNF) and ERC Grant QCLS and a Royal Society Fellowship (SS).

References

- 1 Robert Atkey. From parametricity to conservation laws, via Noether’s theorem. In *Proc. POPL 2014*, pages 491–502, 2014.
- 2 Lars Birkedal and Rasmus E Møgelberg. Categorical models for Abadi and Plotkin’s logic for parametricity. *Mathematical Structures in Computer Science*, 15(04):709–772, 2005.
- 3 Ronald Brown and Christopher B Spencer. G-groupoids, crossed modules and the fundamental groupoid of a topological group. *Proc. Indag. Math.*, 79(4):296–302, 1976.
- 4 Edgar Buckingham. On physically similar systems; illustrations of the use of dimensional equations. *Physical Review*, 4(4):345–376, 1914.
- 5 P-L Curien, Richard Garner, and Martin Hofmann. Revisiting the categorical interpretation of dependent type theory. *Theoret. Comput. Sci.*, 546:99–119, 2014.
- 6 Martin Erwig and Margaret Burnett. Adding apples and oranges. In *Practical Aspects of Declarative Languages*, pages 173–191. Springer, 2002.
- 7 Claudio Hermida. Some properties of `Fib` as a fibred 2-category. *Journal of Pure and Applied Algebra*, 134(1):83–109, 1999.
- 8 Ronald T. House. A proposal for an extended form of type checking of expressions. *The Computer Journal*, 26(4):366–374, 1983.
- 9 Bart Jacobs. *Categorical logic and type theory*, volume 141. Elsevier, 1999.
- 10 Andrew J. Kennedy. Relational parametricity and units of measure. In *POPL’97*, 1997.
- 11 F. William Lawvere. Adjointness in foundations. *Dialectica*, 23(3-4):281–296, 1969.
- 12 R Männer. Strong typing and physical units. *ACM Sigplan Notices*, 21(3):11–20, 1986.
- 13 Mars Climate Orbiter Mishap Investigation Board. Phase I report. NASA, 1999.
- 14 Paul-André Melliès and Noam Zeilberger. Functors are type refinement systems. In *Proc. POPL 2015*, pages 3–16, 2015.
- 15 John Reynolds. Types, abstraction and parametric polymorphism. In *Information Processing*, 1983.
- 16 Ain A Sonin. The physical basis of dimensional analysis. Department of Mechanical Engineering, MIT, 2001.
- 17 Mitchell Wand and Patrick O’Keefe. Automatic dimensional inference. In *Computational Logic – Essays in Honor of Alan Robinson*, pages 479–483, 1991.

MALL Proof Equivalence is Logspace-Complete, via Binary Decision Diagrams

Marc Bagnol

Department of Mathematics and Statistics, University of Ottawa, Canada

Abstract

Proof equivalence in a logic is the problem of deciding whether two proofs are equivalent *modulo* a set of permutation of rules that reflects the commutative conversions of its cut-elimination procedure. As such, it is related to the question of proofnets: finding canonical representatives of equivalence classes of proofs that have good computational properties. It can also be seen as the word problem for the notion of free category corresponding to the logic.

It has been recently shown that proof equivalence in MLL (the multiplicative with units fragment of linear logic) is Pspace-complete, which rules out any low-complexity notion of proofnet for this particular logic.

Since it is another fragment of linear logic for which attempts to define a fully satisfactory low-complexity notion of proofnet have not been successful so far, we study proof equivalence in MALL (multiplicative-additive without units fragment of linear logic) and discover a situation that is totally different from the MLL case. Indeed, we show that proof equivalence in MALL corresponds (under AC_0 reductions) to equivalence of binary decision diagrams, a data structure widely used to represent and analyze Boolean functions efficiently.

We show these two equivalent problems to be Logspace-complete. If this technically leaves open the possibility for a complete solution to the question of proofnets for MALL, the established relation with binary decision diagrams actually suggests a negative solution to this problem.

1998 ACM Subject Classification F.1.3 Complexity Measures and Classes, F.4.1 Mathematical Logic

Keywords and phrases linear logic, proof equivalence, additive connectives, proofnets, binary decision diagrams, logarithmic space, AC_0 reductions

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.60

1 Introduction

Proofnets: from commutative conversions to canonicity

From the perspective of the Curry-Howard (or propositions-as-types) correspondence [5], a proof of $A \Rightarrow B$ in a logic enjoying a cut-elimination procedure can be seen as a program that inputs (through the cut rule) a proof of A and outputs a cut-free proof of B .

Coming from this dynamic point of view, linear logic [6] makes apparent the distinction between data that can or cannot be copied/erased via its exponential modalities and retains the symmetry of classical logic: the linear negation $(\cdot)^*$ is an involutive operation. The study of cut-elimination is easier in this setting thanks to the linearity constraint. However, in its sequent calculus presentation, the cut-elimination procedure of linear logic still suffers from the common flaw of these type of calculi: commutative conversions.



© Marc Bagnol;

licensed under Creative Commons License CC-BY

13th International Conference on Typed Lambda Calculi and Applications (TLCA'15).

Editor: Thorsten Altenkirch; pp. 60–75



Leibniz International Proceedings in Informatics

LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

$$\frac{\frac{\langle \pi \rangle}{A^*, B^*, C, D, \Gamma} \wp}{A^* \wp B^*, C, D, \Gamma} \wp \quad \frac{\langle \mu \rangle \quad \langle \nu \rangle}{A \quad B} \otimes}{\frac{A^* \wp B^*, C \wp D, \Gamma}{A \otimes B} \text{cut}} \wp \quad \rightarrow \quad \frac{\frac{\langle \pi \rangle}{A^*, B^*, C, D, \Gamma} \wp \quad \frac{\langle \mu \rangle \quad \langle \nu \rangle}{A \quad B} \otimes}{\frac{C, D, \Gamma}{C \wp D, \Gamma} \wp} \text{cut}$$

In the above reduction, one of the two formulas related by the *cut* rule is introduced deeper in the proof, making it impossible to perform an actual elimination step right away: one needs first to *permute* the rules in order to be able to go on.

This type of step is called a *commutative conversion* and their presence complexify a lot the study of the *cut*-elimination procedure, as one needs to work *modulo* an equivalence relation on proofs that is not orientable into a rewriting procedure in an obvious way: there are for instance situations of the form

$$\frac{\frac{\langle \pi_1 \rangle}{A^*, B^*, \Gamma} \text{cut} \quad \frac{\langle \pi_2 \rangle}{A} \quad \langle \pi_3 \rangle}{\vdash B^*, \Gamma} \text{cut} \quad \leftrightarrow \quad \frac{\frac{\langle \pi_1 \rangle}{A^*, B^*, \Gamma} \quad \frac{\langle \pi_3 \rangle}{B} \text{cut} \quad \langle \pi_2 \rangle}{\vdash A^*, \Gamma} \text{cut}$$


where it is not possible to favor one side of the equivalence without further non-local knowledge of the proof. The point here is that, as a language for describing proofs, sequent calculus is somewhat *too explicit*. For instance, the fact that the two proofs

$$\frac{\frac{\langle \pi \rangle}{A, B, C, D, \Gamma} \wp}{A \wp B, C, D, \Gamma} \wp \quad \text{and} \quad \frac{\frac{\langle \pi \rangle}{A, B, C, D, \Gamma} \wp}{A, B, C \wp D, \Gamma} \wp}{A \wp B, C \wp D, \Gamma} \wp$$

are different objects from the point of view of sequent calculus generates the first commutative conversion we saw above.

A possible solution to this issue is to look for more intrinsic description of proofs, to find a language that is more *synthetic*; if possible to the point where we have no commutative conversions to perform anymore.

Introduced at the same time as linear logic, the theory of *proofnets* [6, 7] partially addresses this issue. The basic idea is to describe proofs as graphs rather than trees, where application of logical rules become local graph construction, thus erasing some inessential sequential informations. Indeed, the two proofs above would translate into the same proofnet:



(where \mathcal{R}_π is the proofnet translation of the rest of the proof) and the corresponding commutative conversion disappears.

For the multiplicative without units fragment of linear logic (MLL⁻), proofnets yield an entirely satisfactory solution to the problem, and constitute a low-complexity canonical representation of proofs based on local operations on graphs.

By canonical, we mean here that two proofs are equivalent *modulo* the permutations of rules induced by the commutative conversions if and only if they have the same proofnet

translation. From a categorical perspective, this means that proofnets constitute a syntactical presentation of the free semi- $*$ -autonomous category and a solution to the associated word problem [10].

Contrastingly, the linear logic community has struggled to extend the notion of proofnets to wider fragment: even the question of MLL (that is, MLL $^\perp$ plus the multiplicative units) could not find a satisfactory answer. A recent result [9] helps to understand this situation: proof equivalence of MLL is actually a Pspace-complete problem. Hence, there is no hope for a satisfactory notion of low-complexity proofnet for this fragment¹.

In this article, we consider the same question, but in the case of MALL $^\perp$: the multiplicative-additive without units fragment of linear logic. Indeed, this fragment has so far also resisted the attempts to build a notion of proofnet that at the same time characterizes proof equivalence and has basic operations of tractable complexity: we have either canonical nets of exponential size [13] or tractable nets that are not canonical [7]. Therefore, it would have not been too surprising to have a similar result of completeness for some untractable complexity class. An obvious candidate in that respect would be coNP: as we will see, one of these two approaches to proofnets for MALL $^\perp$ is related to Boolean formulas, which equivalence problem is coNP-complete.

It turns out in the end that this is not the case: our investigation concludes that the equivalence problem in MALL $^\perp$ is Logspace-complete under AC₀ reductions. But maybe more importantly, we uncover in the course of the proof an unexpected connexion of this theoretical problem with a very practical issue: indeed we show that MALL $^\perp$ proofs are closely related to binary decision diagrams.

Binary decision diagrams

The problem of the representation of Boolean functions is of central importance in circuit design and has a large range of practical applications. Over the years, binary decision diagrams (BDD) [2] became the most widely used data structure to treat this question.

Roughly speaking, a BDD is a binary tree with nodes labeled by Boolean variables and leaves labeled by values 0 and 1. Such a tree represents a Boolean function in the sense that once an assignment of the variable is chosen, then following the left or right path at each node according to the value 0 or 1 chosen for its variable eventually leads to a leaf, which is the output of the function.

This representation has many advantages which justify its popularity [15]: most basic operations (negation and other logical connectives) on BDD can be implemented efficiently, in many practical cases the size of the BDD representing a Boolean function remains compact (thanks to the possibility to have shared subtrees) and once a variable ordering is chosen they enjoy a notion of normal form.

In this article, we consider both BDD and ordered BDD with no sharing of subtrees and write them as special kinds of Boolean functions by introducing an `IfThenElse` constructor. However, when manipulating them from a complexity point of view we will keep the binary tree presentation in mind.

¹ Of course, this applies only to the standard formulation of units: the equivalence problem for any notion of multiplicative units enjoying less permutations of rules could potentially still be tractable via proofnets: see for instance the work of S. Guerrini and A. Masini [8] and D. Hughes [11].

AC₀ reductions

To show that a problem is complete for some complexity class C , one needs to specify the notion of reduction functions considered, and of course this needs to be a class of functions supposed to be smaller than C itself (indeed *any* problem in C is complete under C reductions).

A standard notion of reduction for the class Logspace is (uniform) AC_0 reduction [3], formally defined in terms of uniform circuits of fixed depth and unbounded fan-in. We will not be getting in the details about this complexity class and, as we will consider only graph transformations, we will rely on the following intuitive principle: if a graph transformation locally replaces each vertex by a bounded number of vertices and the replacement depends *only* on the vertex considered and eventually its direct neighbors, then the transformation is in AC_0 . Typical examples of such a transformation are certain simple cases of so-called “gadget” reductions used in complexity theory to prove hardness results.

Outline of the paper

Section 2 covers some background material on MALL^- and notions of proofnet for this fragment: monomial proofnets and the associated vocabulary for Boolean formulas and BDD in Section 2.1 and the notion of slicing in Section 2.2. Then, we introduce in Section 2.3 an intermediary notion of proof representation that will help us to relate proofs in MALL^- and BDD. In Section 3, we prove that proof equivalence in MALL^- and equivalence of BDD relate to each other through AC_0 reductions and that they are both Logspace -complete.

2 Proof equivalence in MALL^-

► **Notation 1.** *The formulas of MALL^- are built inductively from atoms which we write $\alpha, \beta, \gamma, \dots$ their duals $\alpha^*, \beta^*, \gamma^*, \dots$ and the binary connectives $\wp, \otimes, \&_x, \oplus$ (we consider that the $\&$ connectives carry a label x to simplify some reasonings, but we will omit it when it is not relevant). We write formulas as uppercase letters A, B, C, \dots unless we want to specify they are atoms. Sequents are sequences of formulas, written as greek uppercase letters $\Gamma, \Delta, \Lambda, \dots$ such that all occurrences of the connective $\&$ in a sequent carry a different label. The concatenation of two sequents Γ and Δ is simply written Γ, Δ .*

Let us recall the rules of MALL^{-2} . We do not include the cut rule in our study, since in a static situation (we are not looking at the cut-elimination procedure of MALL^-) it can always be encoded w.l.o.g. using the \otimes rule.

$$\alpha, \alpha^* \quad \frac{\Gamma, A, B}{\Gamma, A \wp B} \wp \quad \frac{\Gamma, A \quad \Delta, B}{\Gamma, \Delta, A \otimes B} \otimes \quad \frac{\Gamma, A}{\Gamma, A \oplus B} \oplus_1 \quad \frac{\Gamma, B}{\Gamma, A \oplus B} \oplus_r \quad \frac{\Gamma, A \quad \Gamma, B}{\Gamma, A \&_x B} \&_x$$

(by convention, we leave the axiom rule implicit to lighten notations. Also, we will use the notation $\frac{\langle \pi \rangle}{\Gamma}$ for “the proof π of conclusion Γ ”.)

► **Remark 2.** Any time we will look at a MALL^- proof from a complexity perspective, we will consider they are represented as trees with nodes corresponding to rules, labeled by the connective introduced and the sequent that is the conclusion of the rule.

² We consider a η -expanded version of MALL^- , which simplifies proofs and definitions, but the extension of our results to a version with non-atomic axioms would be straightforward. Also, we work *modulo* the exchange rule.

Two MALL \perp proofs π and ν are said to be *equivalent* (notation $\pi \sim \nu$) if one can pass from one to the other via permutations of rules [14]. We have an associated decision problem.

► **Definition 3** (MALL \perp -equ). MALL \perp -equ is the decision problem:

“Given two MALL \perp proofs π and ν with the same conclusion, do we have $\pi \sim \nu$?”

We will not go through all the details about this syntactic way to define proof equivalence in MALL \perp . The reason for this is that we already have an available equivalent characterization in terms of *slicing* [14] which we review in Section 2.2. Instead, let us focus only on the most significant case.

$$\frac{\frac{\langle \pi \rangle}{\Gamma, A, C} \quad \frac{\langle \mu \rangle}{\Gamma, B, C} \quad \frac{\langle \nu \rangle}{\Delta, D}}{\Gamma, A \& B, C} \& \quad \otimes \quad \sim \quad \frac{\frac{\langle \pi \rangle}{\Gamma, A, C} \quad \frac{\langle \nu \rangle}{\Delta, D} \quad \frac{\langle \mu \rangle}{\Gamma, B, C} \quad \frac{\langle \nu \rangle}{\Delta, D}}{\Gamma, \Delta, A, C \otimes D} \otimes \quad \frac{\otimes}{\Gamma, \Delta, B, C \otimes D} \& \quad \otimes \quad \frac{\otimes}{\Gamma, \Delta, A \& B, C \otimes D} \&$$

In the above equivalence, the \otimes rule gets lifted above the $\&$ rule. But doing so, notice that we created two copies of ν instead of one, therefore the size of the proof tree has grown. Iterating on this observation, it is not hard to build pairs of proofs that are equivalent, but one of which is exponentially bigger than the other. This is indeed where the difficulty of proof equivalence in MALL \perp lies. As a matter of fact, this permutation of rules *alone* would be enough to build the encoding of the equivalence problem of binary decision diagrams presented in Section 3.1.

A way to attack proof equivalence in a logic, as we exposed in Section 1, is to try to setup a notion of proofnet for this logic. In the following, we will review the main two approaches to this idea in the case of MALL \perp : *monomial proofnets* by J.-Y. Girard [7, 16] and *slicing proofnets* by D. Hughes and R. van Glabbeek [13, 14]. We will then design an intermediate notion of BDD *slicing* that will be more suited to our needs.

2.1 Monomial proofnets

The first attempt in the direction of a notion of proofnet for MALL \perp is due to J.-Y. Girard [7], followed by a version with a full cut-elimination procedure by O. Laurent and R. Maielli [16].

While proofnets for multiplicative linear logic without units were introduced along linear logic itself [6], extending the notion to the multiplicative-additive without units fragment proved to be a true challenge, mainly because of the *superposition* at work in the $\&$ rule.

Girard’s idea was to represent the superposed “versions” of the proofnet by attributing a Boolean formula (called a *weight*) to each link, with one Boolean variable for each $\&$ connective in the conclusion Γ . To retrieve the version of the proofnet corresponding to some selection of the left/right branches of each $\&$, one then just needs to evaluate the Boolean formulas with the corresponding valuation of their variables.

This is the occasion to introduce the vocabulary to speak about Boolean formulas.

► **Definition 4** (Boolean formula). Given a finite set of variables $V = \{x_1, \dots, x_n\}$, a *Boolean formula* over V is inductively defined from the elements of V ; the constants 0 (“false”) and 1 (“true”); the unary symbol $\bar{\cdot}$ (“negation”); the binary symbols $+$ and \cdot (“sum/or/disjunction” and “product/and/conjunction” respectively).

We consider a syntactic notion of *equality* of Boolean formulas: for instance $0.x \neq 0$. The real important notion, that we therefore state separately, is *equivalence*: the fact that if we replace the variables with actual values, we get the same output.

► **Definition 5** (equivalence). A *valuation* v of V is a choice of 0 or 1 for any element of V . A valuation induces an *evaluation* function $v(\cdot)$ from Boolean formulas over V to $\{0, 1\}$ in the obvious way. Two Boolean formulas ϕ and ψ over V are *equivalent* (notation $\phi \sim \psi$) when for any valuation v of V , we have $v(\phi) = v(\psi)$.

► **Definition 6** (monomial). We write $\bar{V} = \{\bar{x}_1, \dots, \bar{x}_n\}$. A *monomial* over V is a Boolean formula of the form $y_1 \dots y_k$ with $\{y_1, \dots, y_k\} \subseteq V \cup \bar{V}$.

Two monomials m and m' are in *conflict* if $m.m' \sim 0$.

► **Remark 7**. Two monomials are in conflict if and only if there is a variable x such that x appears in one of them and \bar{x} appears in the other.

While monomials are a specific type of Boolean formula, the *binary decision diagrams* we are about to introduce are not defined directly as Boolean formulas. Of course, they relate to each other in an obvious way, but having a specific syntax for binary decision diagrams will prove more convenient to solve the problems we will be facing.

► **Definition 8** (BDD). A *binary decision diagram* (BDD) is defined inductively as:

- The constants 0 and 1 are BDD
- If ϕ, ψ are BDD and X is either a variable, 0 or 1, **If X Then ϕ Else ψ** is a BDD
- If ϕ is a BDD and X is either a variable, 0 or 1, **DontCare X Then ϕ** is a BDD

Moreover, suppose we have an ordered set of variables $V = \{x_1, \dots, x_n\}$ with the convention that variables are listed in the reverse order: x_n is first, then x_{n-1} , etc. We define a subclass of BDD which we call *ordered binary decision diagrams over V* (${}^\circ\text{BDD}/V$) by restricting to the following inductive cases (we let $V' = \{x_1, \dots, x_{n-1}\}$)

- The constants 0 and 1 are ${}^\circ\text{BDD}/\emptyset$
- If ϕ and ψ are ${}^\circ\text{BDD}/V'$, **If x_n Then ϕ Else ψ** is a ${}^\circ\text{BDD}/V$
- If ϕ is a ${}^\circ\text{BDD}/V'$, **DontCare x_n Then ϕ** is a ${}^\circ\text{BDD}/V$

The notions of valuation and equivalence are extended to BDD and ${}^\circ\text{BDD}$ the obvious way so that **DontCare X Then ϕ** $\sim \phi$ and **If X Then ϕ Else ψ** $\sim X.\phi + \bar{X}.\psi$.

► **Remark 9**. Any time we will look at BDD and ${}^\circ\text{BDD}$ from a complexity perspective, we will consider they are represented as labeled trees. The cases of **If 0, DontCare 1, ...** will be useful to obtain AC_0 reductions in Section 2.3 and Section 3.2, since erasing a whole subpart of a graph of which we do not know the address in advance is not something that is doable in this complexity class. The absence of sharing implied by the tree representation is also crucial to get low-complexity reductions.

► **Example 10**. The Boolean formula $x.y.\bar{z}$ is a monomial, while $x.y + z$ and $x.1$ are not.

The BDD **If x_2 Then (If x_1 Then 0 Else 1) Else 1** (which is *not* a ${}^\circ\text{BDD}/\{x_1, x_2\}$ by the way) is equivalent to the Boolean formula $x_2.\bar{x}_1 + \bar{x}_2$: both evaluate to 1 only for the valuations $\{x_1 \mapsto 1, x_2 \mapsto 0\}$, $\{x_1 \mapsto 0, x_2 \mapsto 0\}$ and $\{x_1 \mapsto 0, x_2 \mapsto 1\}$. There exist an equivalent ${}^\circ\text{BDD}/\{x_1, x_2\}$: **If x_2 Then (If x_1 Then 0 Else 1) Else (DontCare x_1 Then 1)**.

► **Definition 11**. **BDDequ** is the following decision problem:

“given two BDD ϕ and ψ , do we have $\phi \sim \psi$?”

${}^\circ\text{BDDequ}$ is the following decision problem:

“given two ${}^\circ\text{BDD}/V$ ϕ and ψ , do we have $\phi \sim \psi$?”

- If $\pi = \frac{\langle \mu \rangle \quad \langle \nu \rangle}{\frac{\Gamma, A \quad \Gamma, B}{\Gamma, \Delta, A \& B} \&}$ then $\mathcal{S}_\pi = \mathcal{S}_\mu \cup \mathcal{S}_\nu$

► **Remark 13.** In the \otimes rule, it is clearly seen that the number of slices are multiplied. This is just what is needed in order to have a combinatorial explosion: for any n , a proof π_n of

$$\underbrace{\alpha^* \otimes \cdots \otimes \alpha^*}_{n \text{ times}}, \underbrace{\alpha \& \alpha, \dots, \alpha \& \alpha}_{n \text{ times}}$$

obtained by combining with the \otimes rule n copies of the proof

$$\frac{\alpha^*, \alpha \quad \alpha^*, \alpha}{\alpha^*, \alpha \& \alpha} \&$$

will be of linear size in n , but with a slicing \mathcal{S}_n containing 2^n linkings.

Slicings (associated to a proof) correspond exactly to the notion of proofnets elaborated by D. Hughes and R. van Glabbeek [13]. While their study was mainly focused on the problems of finding a correctness criterion and designing a cut-elimination procedure for these, it also covers the proof equivalence problem. The proof that their notion of proofnet characterizes MALL⁻ proof equivalence can be found in an independent note [14].

► **Theorem 14** (slicing equivalence [14, Theorem 1]). *Let π and ν be two MALL⁻ proofs. We have that $\pi \sim \nu$ if and only if $\mathcal{S}_\pi = \mathcal{S}_\nu$.*

Let us also end this section with a graphical representation of an example of proofnet from the article of Hughes and van Glabbeek, encoding the proof on the left with three linkings $\lambda_1, \lambda_2, \lambda_3$:

$$\frac{\frac{\frac{P^*, P}{P^* \oplus Q^*, P^{\oplus 1}} \quad \frac{Q^*, Q}{P^* \oplus Q^*, Q^{\oplus 1}}}{(P^* \oplus Q^*) \oplus R^*, P^{\oplus 1}} \& \quad \frac{R^*, R}{(P^* \oplus Q^*) \oplus R^*, P^{\oplus 1}} \&}{(P^* \oplus Q^*) \oplus R^*, (P \& Q) \& R} \&$$

λ_1 :

λ_2 :

λ_3 :

$(P^\perp \oplus Q^\perp) \oplus R^\perp, P \& (Q \& R)$

$(P^\perp \oplus Q^\perp) \oplus R^\perp, P \& (Q \& R)$

$(P^\perp \oplus Q^\perp) \oplus R^\perp, P \& (Q \& R)$

2.3 BDD slicings

We finally introduce an intermediate notion of representation of proofs which will be a central tool in the next section. In a sense, it is a synthesis of monomial proofnets and slicings: acknowledging the fact that slicing makes the size of the representation explode, we rely on BDD to keep things more compact.

Of course, the canonicity property is lost. But this is exactly the point! Indeed, deciding whether two “BDD slicings” are equivalent is the reformulation of proof equivalence in MALL⁻ we rely on in the reductions between MALL⁻equ (Definition 3) and BDDequ (Definition 11).

► **Definition 15** (BDD slicing). Given a MALL⁻ sequent Γ , a BDD *slicing* of Γ is a function \mathcal{B} that associates a BDD to every element $[\gamma, \gamma^*]$ of the set of (unordered) pairs of occurrences of dual atoms in Γ .

We say that two BDD slicings M, N of the same Γ are *equivalent* (notation $M \sim N$) if for any pair $[\gamma, \gamma^*]$, we have $M[\gamma, \gamma^*] \sim N[\gamma, \gamma^*]$ in the sense of Definition 5.

To any MALL⁻ proof π , we associate a BDD slicing \mathcal{B}_π by induction:

- If $\pi = \alpha, \alpha^*$ then $\mathcal{B}_\pi[\alpha, \alpha^*] = 1$.

- If $\pi = \frac{\langle \mu \rangle}{\Gamma, A, B} \wp$ then $\mathcal{B}_\pi[\gamma, \gamma^*] = \mathcal{B}_\mu[\gamma, \gamma^*]$ where we see atoms of $A \wp B$ as the corresponding atoms of A and B
- If $\pi = \frac{\langle \mu \rangle \quad \langle \nu \rangle}{\Gamma, A \quad \Delta, B} \otimes$ then $\mathcal{B}_\pi[\gamma, \gamma^*] = \begin{cases} \mathcal{B}_\mu[\gamma, \gamma^*] & \text{if } \gamma, \gamma^* \text{ are atoms of } \Gamma, A \\ \mathcal{B}_\nu[\gamma, \gamma^*] & \text{if } \gamma, \gamma^* \text{ are atoms of } \Delta, B \\ 0 & \text{otherwise}^3 \end{cases}$
- If $\pi = \frac{\langle \mu \rangle}{\Gamma, A} \oplus_1$ then $\mathcal{B}_\pi[\gamma, \gamma^*] = \begin{cases} \mathcal{B}_\mu[\gamma, \gamma^*] & \text{if } \gamma, \gamma^* \text{ are atoms of } \Gamma, A \\ 0 & \text{otherwise} \end{cases}$ and likewise for \oplus_r .
- If $\pi = \frac{\langle \mu \rangle \quad \langle \nu \rangle}{\Gamma, A \quad \Gamma, B} \wp_x$ then

$$\mathcal{B}_\pi[\gamma, \gamma^*] = \begin{cases} \text{If } x \text{ Then } \mathcal{B}_\mu[\gamma, \gamma^*] \text{ Else } 0 & \text{if } \gamma \text{ or } \gamma^* \text{ is an atom of } A \\ \text{If } x \text{ Then } 0 \text{ Else } \mathcal{B}_\nu[\gamma, \gamma^*] & \text{if } \gamma \text{ or } \gamma^* \text{ is an atom of } B \\ \text{If } x \text{ Then } \mathcal{B}_\mu[\gamma, \gamma^*] \text{ Else } \mathcal{B}_\nu[\gamma, \gamma^*] & \text{otherwise} \end{cases}$$

► **Remark 16.** The BDD we obtain this way are actually of a specific type: they are usually called *read-once* BDD: from the root to any leaf, one never crosses two **IfThenElse** nodes asking for the value of the same variable.

► **Example 17.** The weight of the pairs $[\alpha, \alpha^*]$ and $[\delta, \delta^*]$ in the BDD slicing of the proof

$$\pi = \frac{\frac{\alpha, \alpha^*}{\alpha \oplus \beta, \alpha^*} \oplus_1 \quad \frac{\beta, \beta^*}{\alpha \oplus \beta, \beta^*} \oplus_r}{\frac{\alpha \oplus \beta, \alpha^* \wp_x \beta^*}{\alpha \oplus \beta, (\alpha^* \wp_x \beta^*)} \wp_x \quad \delta, \delta^*}{\alpha \oplus \beta, (\alpha^* \wp_x \beta^*) \otimes \delta, \delta^*} \otimes$$

are $\mathcal{B}_\pi[\alpha, \alpha^*] = \text{If } x \text{ Then } 1 \text{ Else } 0$ and $\mathcal{B}_\pi[\delta, \delta^*] = 1$.

It is not hard to see that proof equivalence matches the equivalence of BDD slicings by relating them to slicings from the previous section.

► **Theorem 18** (BDD slicing equivalence). *Let π and ν be two MALL proofs. We have that $\pi \sim \nu$ if and only if $\mathcal{B}_\pi \sim \mathcal{B}_\nu$.*

Proof. We show in fact that $\mathcal{S}_\pi = \mathcal{S}_\nu$ if and only if $\mathcal{B}_\pi \sim \mathcal{B}_\nu$, with Theorem 14 in mind.

To a BDD slicing \mathcal{B} , we can associate a linking $v(\mathcal{B})$ for each valuation v of the variables occurring in \mathcal{B} by setting $v(\mathcal{B}) = \{ [\alpha, \alpha^*] \mid v(\mathcal{B}[\alpha, \alpha^*]) = 1 \}$ and then a slicing $f(\mathcal{B}) = \{ v(\mathcal{B}) \mid v \text{ valuation} \}$. By definition, it is clear that if \mathcal{B} and \mathcal{B}' involve the same variables and $\mathcal{B} \sim \mathcal{B}'$ then $f(\mathcal{B}) = f(\mathcal{B}')$.

Conversely, suppose $\mathcal{B}_\pi \not\sim \mathcal{B}_\nu$, so that there is a v such that $v(\mathcal{B}_\pi) \neq v(\mathcal{B}_\nu)$. To conclude that $f(\mathcal{B}_\pi) \neq f(\mathcal{B}_\nu)$, we must show that there is no other v' such that $v'(\mathcal{B}_\nu) = v(\mathcal{B}_\pi)$.

To do this, we can extend the notion of valuation to proofs: if v is a valuation of the labels x of the \wp_x in π , $v(\pi)$ is defined by keeping only the left or right branch of \wp_x according to the value of x . Now we can consider the set $v^{\text{ax}}(\pi)$ of axiom rules in $v(\pi)$ and we can show by

³ Remember we consider *occurrences* of atoms, and as the \otimes rule splits the context into two independent parts that and no axiom rule can cross this splitting.

induction that $v^{\text{ax}}(\pi)$ must contain at least one pair with one atom which is a subformula of the side of each $\&_x$ that has been kept. Therefore for any π and ν with the same conclusion, if $v \neq v'$ we have $v^{\text{ax}}(\nu) \neq v'^{\text{ax}}(\pi)$ no matter what. Then we can remark that $v^{\text{ax}}(\pi)$ is just another name for $v(\mathcal{B}_\pi)$ so that in the end, there cannot be $v \neq v'$ such that $v'(\mathcal{B}_\nu) = v(\mathcal{B}_\pi)$.

Finally, an easy induction shows that $f(\mathcal{B}_\pi) = \mathcal{S}_\pi$ and therefore we are done. \blacktriangleleft

Also, a BDD equivalent to the BDD associated to a pair can be computed in AC_0 .

► **Lemma 19** (computing BDD slicings). *For any MALL^- proof π and any pair $[\gamma, \gamma^*]$, we can compute in AC_0 a BDD ϕ such that $\phi \sim \mathcal{B}_\pi[\gamma, \gamma^*]$.*

Proof. As we see proofs as labeled trees (Remark 2), we will only locally replace the rules of the proof the following way to obtain the corresponding BDD ϕ :

- Replace axiom rules γ, γ^* by 1 and other axiom rules by 0
- Replace all \wp and \oplus rules by **DontCare** 1 **Then** \cdot nodes
- In the \otimes case, test which side the atoms γ, γ^* are attributed to (*this can be done locally by looking at the conclusions of the premise of the rule*) and replace it by a **If** 0 **Then** \cdot **Else** \cdot or a **If** 1 **Then** \cdot **Else** \cdot node accordingly
- Replace $\&_x$ rules by a **If** x **Then** \cdot **Else** \cdot nodes

We can see by induction that the resulting BDD is equivalent to $\mathcal{B}_\pi[\gamma, \gamma^*]$. All these operations can be performed by looking only at the rule under treatment (and its immediate neighbors in the case of \otimes) and always replaces one rule by exactly one node. Therefore it is in AC_0 . \blacktriangleleft

► **Corollary 20** (reduction). *$\text{MALL}^- \text{equ}$ reduces to BDDequ in AC_0 .*

In the next section we focus on the equivalence of BDD and $^\circ\text{BDD}$, proving first that the case of $^\circ\text{BDD}$ can be reduced to proof equivalence in MALL^- . Then, we will show the problem of equivalence of BDD to be in Logspace , and that of $^\circ\text{BDD}$ to be Logspace-hard , thus characterizing the intrinsic complexity of proof equivalence in MALL^- as Logspace-complete .

Note that this contrasts with the classical result that equivalence of general Boolean formulas is coNP-complete . It turns out indeed that the classes of BDD we consider enjoy a number of properties that allow to solve equivalence more easily.

3 Equivalence of BDD

3.1 Equivalence of $^\circ\text{BDD}$ reduces to proof equivalence in MALL^-

We now show that the converse of Lemma 19 holds for $^\circ\text{BDD}$.

To do this, we rely on a formula B which will serve as the placeholder of a $^\circ\text{BDD}/V$ ϕ we want to encode; and a context Γ which contains one $\&_x$ connective for each variable x in V , organized in a way that allows for an inductive specification of $^\circ\text{BDD}$.

Given an $^\circ\text{BDD}$ ϕ , we wish to obtain a proof π_ϕ of B, Γ such that dual pairs with one element in B will receive either the value ϕ or a value equivalent to $\bar{\phi}$ in the BDD slicing of π_ϕ ; and on the other hand, the other dual pairs of Γ will receive equivalent values whatever the $^\circ\text{BDD}$ we encode is. This will lead to the fact that two such encoding proofs are equivalent if and only if the $^\circ\text{BDD}$ they encode are equivalent.

► **Notation 21.** *We fix atomic formulas $\alpha_1, \dots, \alpha_n \dots$ and β and write $B = \beta \oplus \beta$. In what follows, we will use β^l and β^r to refer respectively to the left and right copies of β in B ; and likewise α_i^l and α_i^r for copies of α_i in $\alpha_i \& \alpha_i$.*

We write respectively π_0 and π_1 the proofs

$$\frac{\beta, \beta^*}{B, \beta^*}^{\oplus_1} \quad \frac{\beta, \beta^*}{B, \beta^*}^{\oplus_r}$$

and for any n , we write $\pi_{\&}^n$ the proof

$$\frac{\alpha_n^*, \alpha_n \quad \alpha_n^*, \alpha_n}{\alpha_n^*, \alpha_n \& \alpha_n}^{\&}$$

► **Definition 22** (encoding a \circ BDD). Let ϕ be an \circ BDD over the variables $V = \{x_1, \dots, x_n\}$. We define the sequent

$$\Gamma^n = (\dots(\beta^* \otimes \alpha_1^*) \otimes \alpha_2^*) \otimes \dots) \otimes \alpha_n^*, \alpha_1 \&_{x_1} \alpha_1, \alpha_2 \&_{x_2} \alpha_2, \dots, \alpha_{n-1} \&_{x_{n-1}} \alpha_{n-1}$$

with $\Gamma^0 = \beta^*$ and set $\Delta^n = \Gamma^n, \alpha_n \&_{x_n} \alpha_n$ with also $\Delta^0 = \beta^*$.

We define the proof π_ϕ of conclusion B, Δ^n by induction on n

- The base cases are 0 and 1, encoded respectively as π_0 and π_1 .
- Otherwise, if $\phi = \text{If } x_n \text{ Then } \psi \text{ Else } \zeta$, with both ψ and ζ being \circ BDD/ $\{x_1, \dots, x_{n-1}\}$, we have π_ψ and π_ζ defined by induction, and then

$$\pi_\phi = \frac{\frac{\langle \pi_\psi \rangle}{B, \Delta^{n-1} \quad \alpha_n^*, \alpha_n} \otimes \frac{\langle \pi_\zeta \rangle}{B, \Delta^{n-1} \quad \alpha_n^*, \alpha_n}}{B, \Gamma^n, \alpha_n} \otimes \frac{B, \Delta^n}{B, \Gamma^n, \alpha_n \&_{x_n}}^{\&}$$

- The last case is $\phi = \text{DontCare } x_n \text{ Then } \psi$, with ψ being a \circ BDD/ $\{x_1, \dots, x_{n-1}\}$ so that we have π_ψ defined by induction, and then

$$\pi_\phi = \frac{\langle \pi_\psi \rangle \quad \langle \pi_{\&}^n \rangle}{B, \Delta^{n-1} \quad \alpha_n^*, \alpha_n \&_{x_n} \alpha_n} \otimes \frac{B, \Delta^n}{B, \Delta^n}$$

We still need to state in what sense π_ϕ is an encoding of ϕ : we turn the statement about the value of atoms of B we made in the beginning of this section into a precise property.

► **Lemma 23** (associated BDD slicing). *Writing \mathcal{B} the BDD slicing of π_ϕ , we have*

$$\mathcal{B}[\beta^1, \beta^*] = \phi \quad \mathcal{B}[\beta^r, \beta^*] \sim \bar{\phi} \quad \mathcal{B}[\alpha_i^*, \alpha_i^1] \sim x_i \quad \mathcal{B}[\alpha_i^*, \alpha_i^r] \sim \bar{x}_i$$

Proof. By a routine inspection of induction cases. Let us only review $\phi = \text{If } x_n \text{ Then } \psi \text{ Else } \zeta$. Let us write \mathcal{B}_ϕ , \mathcal{B}_ψ and \mathcal{B}_ζ the BDD slicings respectively associated to the proofs π_ϕ , π_ψ and π_ζ .

By induction we have that $\mathcal{B}_\psi[\beta^1, \beta^*] = \psi$ and $\mathcal{B}_\zeta[\beta^1, \beta^*] = \zeta$, therefore by definition of the BDD slicing associated to a $\&$ rule, we have that $\mathcal{B}_\phi[\beta^1, \beta^*] = \text{If } x_n \text{ Then } \psi \text{ Else } \zeta = \phi$. The case of β^r is similar, but for its use of Lemma 27 from the next section.

As for the other pairs of occurrences of dual atoms in the conclusion, let us have a look at the case of $[\alpha_i^*, \alpha_i^1]$: if by induction $\mathcal{B}_\psi[\alpha_i^*, \alpha_i^1] \sim x_i$ and $\mathcal{B}_\zeta[\alpha_i^*, \alpha_i^1] \sim x_i$, then by definition $\mathcal{B}_\phi[\alpha_i^*, \alpha_i^1] \sim \text{If } x_n \text{ Then } x_i \text{ Else } x_i \sim x_i$. The case of α_i^r is similar. ◀

► **Corollary 24** (equivalence). *If ϕ and ψ are two \circ BDD/ $\{x_1, \dots, x_n\}$, then $\pi_\phi \sim \pi_\psi$ if and only if $\phi \sim \psi$.*

► **Lemma 25** (computing the encoding). *Given a \circ BDD/ $\{x_1, \dots, x_n\}$ ϕ , π_ϕ can be computed in AC_0 .*

Proof. As in the proof of Lemma 19, we show that the inductive definition can in fact be seen as a local graph transformation introducing nodes of bounded size:

- Replace each 0 node by B, β^* and 1 node by B, β^* .

$$\begin{array}{ccc} \langle \pi_0 \rangle & & \langle \pi_1 \rangle \\ \vdots & & \vdots \end{array}$$
- Replace each `DontCare` x_i `Then` \cdot by $\frac{\langle \pi_{\&}^{x_i} \rangle}{B, \Delta^i} \otimes$

$$\begin{array}{c} \vdots \\ \alpha_i^*, \alpha_i \&_{x_i} \alpha_i \\ \vdots \end{array}$$
- Replace each `If` x_i `Then` \cdot `Else` \cdot by $\frac{\frac{\langle \alpha_i^*, \alpha_i \rangle}{B, \Gamma^i, \alpha_i} \otimes \frac{\langle \alpha_i^*, \alpha_i \rangle}{B, \Gamma^i, \alpha_i}}{B, \Delta^i} \&_{x_i}$

$$\begin{array}{c} \vdots \\ \alpha_i^*, \alpha_i \\ \vdots \end{array}$$

For any of these replacements, we see that the choice of the case to apply and the label of the resulting block (of bounded size) of rules replacing a node depends only on the label of the node we are replacing, therefore the transformation is in AC_0 . ◀

► **Corollary 26** (reduction). $\circ\text{BDDequ}$ reduces to MALL^{equ} in AC_0 .

To sum up, we have so far the following chain of AC_0 reductions:

$$\circ\text{BDDequ} \rightarrow \text{MALL}^{\text{equ}} \rightarrow \text{BDDequ}$$

3.2 Logspace-completeness

We prove in this section that all these equivalence problems are **Logspace**-complete. We begin by listing a few useful properties of BDD that will allow to design a **Logspace** decision procedure for their equivalence. Then, we prove the **Logspace**-hardness by reducing to $\circ\text{BDDequ}$ a **Logspace**-complete problem on line graph orderings.

The starting point is the good behavior of BDD with respect to negation. In the following lemma, we consider the negation of a BDD which is not strictly speaking a BDD itself: we think of it as the equivalent Boolean formula, the point being precisely to show that this Boolean formula can be easily expressed as a BDD.

► **Lemma 27** (negation). *If ϕ and ψ are BDD and X is either 0, 1 or a variable, we have*

$$\overline{\text{If } X \text{ Then } \phi \text{ Else } \psi} \sim \text{If } X \text{ Then } \overline{\phi} \text{ Else } \overline{\psi}$$

$$\overline{\text{DontCare } X \text{ Then } \phi} \sim \text{DontCare } X \text{ Then } \overline{\phi}$$

Proof. First we can transform our expression by

$$\overline{\text{If } X \text{ Then } \phi \text{ Else } \psi} \sim \overline{X \cdot \phi + \overline{X} \cdot \psi} \sim (\overline{X + \overline{\phi}}) \cdot (X + \overline{\psi}) \sim X \cdot \overline{\psi} + \overline{X} \cdot \overline{\phi} + \psi \cdot \phi$$

then we can apply the so-called “consensus rule” of Boolean formulas

$$X \cdot \overline{\psi} + \overline{X} \cdot \overline{\phi} + \psi \cdot \phi \sim X \cdot \overline{\phi} + \overline{X} \cdot \overline{\psi} \sim \text{If } X \text{ Then } \overline{\phi} \text{ Else } \overline{\psi}$$

The case of `DontCare` is obvious. ◀

► **Corollary 28.** *If ϕ is a BDD, then there is a BDD $\widehat{\phi}$ such that $\overline{\phi} \sim \widehat{\phi}$. Moreover $\widehat{\phi}$ can be computed in logarithmic space.*

Proof. An induction on the previous lemma shows that we can obtain the negation of a BDD simply by flipping the 0 nodes to 1 nodes and conversely. Hence the transformation is even in AC_0 . ◀

Then, we show that a BDD can be seen as a sum of monomials through a **Logspace** transformation.

► **Lemma 29** (BDD as sums of monomials). *If ϕ is a BDD, then there is a formula ϕ_m which is a sum of monomials and is such that $\phi \sim \phi_m$.*

Moreover ϕ_m can be computed in logarithmic space.

Proof. For each 1 node in ϕ , go down to the root of ϕ and output one by one the variables of any **If x Then \cdot Else \cdot** encountered: this produces the monomial associated to this 1 node. Then ϕ_m is the sum of all the monomials obtained this way and is clearly equivalent to ϕ . The procedure is in **Logspace** because we only need to remember which 1-leave we are treating and where we are in the tree (when going down) at any point. ◀

Putting all this together, we finally obtain a space-efficient decision procedure. Note however that it is totally sub-optimal in terms of time: to keep with the logarithmic space bound, we have to recompute a lot of things rather than store them.

► **Corollary 30** (BDDequ is in Logspace). *There is a logarithmic space algorithm that, given two BDD ϕ and ψ , decides whether they are equivalent.*

Proof. The BDD ϕ and ψ are equivalent if and only if $\phi \Leftrightarrow \psi = (\bar{\phi} + \psi) \cdot (\bar{\psi} + \phi) \sim 1$ that is to say (by passing to the negation) $(\bar{\psi} \cdot \phi) + (\bar{\phi} \cdot \psi) \sim 0$, which holds if and only if both $\bar{\psi} \cdot \phi \sim 0$ and $\bar{\phi} \cdot \psi \sim 0$.

But then, considering the first one (the other being similar) we can rewrite it in logarithmic space using the two above lemmas as $(\widehat{\psi})_m \cdot \phi_m \sim 0$. This holds if and only if for all pairs (m, m') of one monomial in $(\widehat{\psi})_m$ and one monomial in ϕ_m , m and m' are in conflict; which can be checked in logarithmic space using Remark 7. ◀

Let us now introduce an extremely simple, yet **Logspace**-complete problem [4], which will ease the **Logspace**-hardness part of our proof.

► **Definition 31** (order between vertices). *Order between vertices (ORD) is the following decision problem:*

“Given a directed graph $G = (V, E)$ that is a line⁴ and two vertices $f, s \in V$ do we have $f < s$ in the total order induced by G ?”

► **Lemma 32.** *ORD reduces to ${}^\circ$ BDDequ in AC_0 .*


Proof. Again we are going to build a local graph transformation that is in AC_0 .

First, we assume w.l.o.g. that the begin b and the exit e vertices of G are different from f and s . We write f^+ and s^+ the vertices immediately after f and s in G .

Then, we perform a first transformation by replacing the graph with three copies of itself (this can be done by locally scanning the graph and create labeled copies of the vertices


⁴ We use the standard definition of graph as a pair (V, E) of sets of vertices and edges (oriented couples of vertices $x \rightarrow y$). A graph is a *line* if it is connected and all the vertices have in-degree and out-degree 1, except the *begin* vertex which has in-degree 0 and out-degree 1 and the *exit* vertex which has in-degree 1 and out-degree 0. A line induces a total order on vertices through its transitive closure.

and edges). We write x_i to refer to the copy of the vertex x in the graph i . The second transformation is a rewiring of the graph as follows: erase the edges going out of the f_i and s_i and replace them as pictured in the two first subgraphs:



Let us call G_r the rewired graph and G_n the non-rewired graph. To each of them we add two binary nodes x and y connected to the begin vertices b_i as pictured in the third graph above. Then we can produce two corresponding \circ BDD ϕ_r and ϕ_n by replacing the exit vertices e_1, e_2, e_3 by 1, 0, 0 respectively, x and y by a `If x Then (DontCare y Then ·) Else (If y Then · Else ·)` block of nodes; and any other v_i vertex by a `DontCare v Then ·`. It is then easy to see that if $f < s$ in the order induced by G if and only if ϕ_r and ϕ_n are equivalent.

Let us illustrate graphically what happens in the case where $f < s$: we draw the resulting \circ BDD as a labeled graph with the convention that a node labeled with z with out-degree 1 is a `DontCare z Then ·` and a node labeled with z with out-degree 2 is a `If z Then · Else ·` node with the upper edge corresponding to the `Then` branch and the lower edge corresponding to the `Else` branch.



► **Remark 33.** The above construction relies on the fact that there are non-commuting permutations on the set of three elements: in a sense we are just attributing two non-commuting σ and τ to f and s and make sure that the order in which they intervene affects the equivalence class of the resulting \circ BDD. An approach quite similar in spirit with the idea of *permutation branching program* [1].

We can now extend our chain of reductions with the two new elements from this section

$$\text{ORD (Logspace-hard)} \rightarrow \circ\text{BDDequ} \rightarrow \text{MALL}^-\text{equ} \rightarrow \text{BDDequ} (\in \text{Logspace})$$

so in the end we get our main result:

► **Theorem 34 (Logspace-completeness).** *The decision problems $\circ\text{BDDequ}$, MALL^-equ and BDDequ are Logspace-complete under AC_0 reductions.*

4 Conclusion

In this work, we characterized precisely the complexity of proof equivalence in MALL^- as Logspace -complete, contrasting greatly with the situation for the MLL fragment which has a Pspace -complete equivalence problem. We did so by establishing a correspondence between MALL^- proofs and specific classes of BDD .

This path we took for proving our result is interesting in itself since the established correspondence allows a transfer of ideas in both directions. In particular, any progress in the theoretical problem of finding a correct notion of proofnet for MALL^- would yield potential

applications to BDD, a notion of widespread practical use. An idea to explore might be the notion of *conflict net* defined by D. Hughes in an unpublished note [12]. Roughly speaking, the principle is to consider a presentation of proofnets with the information of the links that cannot be present at the same time, rather than giving an explicit formula to compute their presence, as it is the case with monomial proofnets or the BDD slicings we introduced.

On the other hand, since many optimization problems regarding BDD are known to be NP-complete, a finer view at the encoding of Section 3.1 in addition to basic constraints about what we expect from a notion of proofnet should lead to an impossibility result, even though the equivalence problem for MALL is only Logspace-complete.

Acknowledgements to people from `cstheory.stackexchange.com` for pointing the author to the notion of BDD; to Dominic Hughes for the live feedback during the redaction of the article; to Clément Aubert, for his help in understanding AC_0 reductions.

References

- 1 David A. Barrington. Bounded-width polynomial-size branching programs recognize exactly those languages in NC1. *Journal of Computer and System Sciences*, 38(1):150 – 164, 1989.
- 2 Randal E. Bryant. Graph-based algorithms for Boolean function manipulation. *IEEE Trans. Comput.*, 35(8):677–691, August 1986.
- 3 Ashok K. Chandra, Larry J. Stockmeyer, and Uzi Vishkin. Constant depth reducibility. *SIAM J. Comput.*, 13(2):423–439, 1984.
- 4 Kousha Etessami. Counting quantifiers, successor relations, and logarithmic space. *Journal of Computer and System Sciences*, 54(3):400 – 411, 1997.
- 5 Jean Gallier. On the correspondence between proofs and lambda-terms. In Philippe de Groote, editor, *The Curry-Howard isomorphism*, Cahiers du Centre de Logique, pages 55–138. Academia, 1995.
- 6 Jean-Yves Girard. Linear logic. *Theoret. Comput. Sci.*, 50(1):1–101, 1987.
- 7 Jean-Yves Girard. Proof-nets: The parallel syntax for proof-theory. *Logic and Algebra*, 180:97–124, May 1996.
- 8 Stefano Guerrini and Andrea Masini. Parsing MELL Proof Nets. *Theoretical Computer Science*, 254(1-2):317–335, 2001.
- 9 Willem Heijltjes and Robin Houston. No Proof Nets for MLL with Units: Proof Equivalence in MLL is PSPACE-complete. In *Proceedings of the Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*, CSL-LICS '14, pages 50:1–50:10, New York, NY, USA, 2014. ACM.
- 10 Willem Heijltjes and Lutz Straßburger. Proof nets and semi-star-autonomous categories. *Mathematical Structures in Computer Science*, FirstView:1–40, 11 2014.
- 11 Dominic J. D. Hughes. Simple multiplicative proof nets with units. Technical report, 2005.
- 12 Dominic J. D. Hughes. Abstract p-time proof nets for MALL: Conflict nets. *arXiv: math.LO/0801.2421v1*, 2008.
- 13 Dominic J. D. Hughes and Rob J. van Glabbeek. Proof nets for unit-free multiplicative-additive linear logic. *ACM Trans. Comput. Log.*, 6(4):784–842, 2005.
- 14 Dominic J. D. Hughes and Rob J. van Glabbeek. MALL proof nets identify proofs modulo rule commutation. <http://boole.stanford.edu/~dominic/MALL-equiv.pdf>, to appear, 2015.
- 15 Donald E. Knuth. *The Art of Computer Programming, Volume 4, Fascicle 1: Bitwise Tricks & Techniques; Binary Decision Diagrams*. Addison-Wesley Professional, 12th edition, 2009.

- 16 Olivier Laurent and Roberto Maieli. Cut elimination for monomial MALL proof nets. In *Proceedings of the Twenty-Third Annual IEEE Symposium on Logic in Computer Science, LICS 2008, 24-27 June 2008, Pittsburgh, PA, USA*, pages 486–497, 2008.

Mixin Composition Synthesis Based on Intersection Types*

Jan Bessai¹, Andrej Dudenhefner¹, Boris Döder¹, Tzu-Chun Chen², Ugo de'Liguoro³, and Jakob Rehof¹

1 Technical University of Dortmund, Dortmund, Germany

{jan.bessai, boris.duedder, andrej.dudenhefner, jakob.rehof}@cs.tu-dortmund.de

2 Technical University of Darmstadt, Darmstadt, Germany

tcchen@rbg.informatik.tu-darmstadt.de

3 University of Torino, Torino, Italy

ugo.deliguoro@unito.it

Abstract

We present a method for synthesizing compositions of mixins using type inhabitation in intersection types. First, recursively defined classes and mixins, which are functions over classes, are expressed as terms in a lambda calculus with records. Intersection types with records and record-merge are used to assign meaningful types to these terms without resorting to recursive types. Second, typed terms are translated to a repository of typed combinators. We show a relation between record types with record-merge and intersection types with constructors. This relation is used to prove soundness and partial completeness of the translation with respect to mixin composition synthesis. Furthermore, we demonstrate how a translated repository and goal type can be used as input to an existing framework for composition synthesis in bounded combinatory logic via type inhabitation. The computed result corresponds to a mixin composition typed by the goal type.

1998 ACM Subject Classification F.4.1 Mathematical Logic – λ Calculus and Related Systems

Keywords and phrases Record Calculus, Combinatory Logic, Type Inhabitation, Mixin, Intersection Type

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.76

1 Introduction

Starting with Cardelli's pioneering work [13], various typed λ -calculi extended with records have been thoroughly studied to model sophisticated features of object-oriented programming languages, like recursive objects and classes, object extension, method overriding and inheritance (see e.g. [1, 11, 23]).

Here, we focus on the synthesis of mixin compositions. In the object-oriented paradigm, mixins [9, 10] have been introduced as an alternative construct for code reuse that improves over the limitations of multiple inheritance, e.g. connecting incompatible base classes and semantic ambiguities caused by the diamond problem. Together with abstract classes and traits, mixins (functions over classes) can be considered as an advanced construct to obtain flexible implementations of module libraries and to enhance code reusability; many popular

* This work was partially supported by EU COST Action IC1201: BETTY and MIUR PRIN CINA Prot. 2010LHT4KM, San Paolo Project SALT.



programming languages miss native support for mixins, but they are an object of intensive study and research (e.g. [8, 24]). In this setting we aim at synthesizing classes from a library of mixins that can be used in programming languages like Java, which do not natively support mixins. Our particular modeling approach is inspired by modern language features (e.g. ECMAScript “`bind`”) to preserve contexts in order to prevent programming errors [21].

We formalize synthesis of classes from a library of mixins as an instance of the relativized type inhabitation problem in bounded combinatory logic with intersection types [20]. Relativized type inhabitation is the decision problem: given a combinatory type context Δ and a type τ does there exist an applicative term e such that e has type τ under the type assumptions in Δ ? We denote type inhabitation by $\Delta \vdash_{\text{BCL}} ? : \tau$ and implicitly include the problem of constructing a term inhabiting τ .

Relativized type inhabitation, which is undecidable in general [20], is decidable in k -bounded combinatory logic $\text{BCL}_k(\rightarrow, \cap)$, that is, combinatory logic typed with arrow and intersection types of depth at most k , and hence an algorithm for semi-deciding type inhabitation for $\text{BCL}(\rightarrow, \cap) = \bigcup_k \text{BCL}_k(\rightarrow, \cap)$ can be obtained by iterative deepening over k and solving the corresponding decision problem in $\text{BCL}_k(\rightarrow, \cap)$ [20]. In the present paper, we enable combinatory synthesis of classes via intersection typed mixin combinators. Intersection types [4] play an important rôle in combinatory synthesis, because they allow for semantic specification of components and synthesis goals [20, 5].

Now, looking at $\{C_1 : \sigma_1, \dots, C_p : \sigma_p, M_1 : \tau_1, \dots, M_q : \tau_q\} \subseteq \Delta$ as the abstract specification of a library including classes C_i and mixins M_j with interfaces σ_i and τ_j respectively, and given a type τ specifying an unknown class, we may identify the class synthesis problem with the type inhabitation problem $\Delta \vdash_{\text{BCL}} ? : \tau$. To make this feasible, we have to bridge the gap between the expressivity of highly sophisticated type systems used for typing classes and mixins, for instance F -bounded polymorphism used in [12, 15], and the system of intersection types from [4]. In doing so, we move from the system originally presented in [16], consisting of a type assignment system of intersection and record types to a λ -calculus which we enrich here with record merge operation (called “**with**” in [15]), to allow for expressive mixin combinators. The type system is modified by reconstructing record types $\langle l_i : \sigma_i \mid i \in I \rangle$ as intersection of unary record types $\langle l_i : \sigma_i \rangle$, and considering a subtype relation extending the one in [4]. This is however not enough for typing record merge, for which we consider a type-merge operator $+$. The problem of typing extensible records and merge, faced for the first time in [27, 26], is notoriously hard; to circumvent difficulties the theory of record subtyping in [15] (where a similar type-merge operator is considered) allows just for “exact” record typing, which involves subtyping in depth, but not in width. Such a restriction, that has limited effects w.r.t. a rich and expressive type system like F -bounded polymorphism, would be too severe in our setting. Therefore, we undertake a study of the type algebra of record types with intersection and type-merge, leading to a type assignment system where exact record typing is required only for the right-hand side operand of the term merge operator, which is enough to ensure soundness of typing.

The next challenge is to show that we can type in a meaningful way in our system classes and mixins, where the former are essentially recursive records and the latter are made of a combination of fixed point combinators and record merge. Such combinators, which usually require recursive types, can be typed in our system by means of an iterative method exploiting the ability of intersection types to represent approximations of the potentially infinite unfolding of recursive definitions.

The final problem we face is the encoding of intersection types with record types and type-merge into the language of $\text{BCL}(\rightarrow, \cap)$. For this purpose we consider a conservative extension

of bounded combinatory logic, called $\text{BCL}(\mathbb{T}_{\mathbb{C}})$, where we allow unary type constructors that are monotonic and distribute over intersection. We show that the (semi) algorithm solving inhabitation for $\text{BCL}(\rightarrow, \cap)$ can be adapted to $\text{BCL}(\mathbb{T}_{\mathbb{C}})$, by proving that the key properties necessary to solve the inhabitation problem in $\text{BCL}(\rightarrow, \cap)$ are preserved in $\text{BCL}(\mathbb{T}_{\mathbb{C}})$ and showing how the type-merge operator can be simulated in $\text{BCL}(\mathbb{T}_{\mathbb{C}})$. In fact, type-merge is not monotonic in its second argument, due to the lack of negative information caused by the combination of $+$ and \cap . Our work culminates in two theorems that ensure soundness and completeness of the so obtained method w.r.t. synthesis of classes by mixins composition.

Related works. This work evolves from the contributions [5, 17, 6] to the workshop ITRS'14. The papers that have inspired our work, mainly by Cook and others, have been cited above. The theme of using intersection types and bounded-polymorphism for typing object-oriented languages and inheritance has been treated in [14, 25]. Type inhabitation has been recently used for synthesis of object oriented code [22, 18], but to our best knowledge the present paper provides, for the first time, a theory of type-safe mixin composition synthesis based on the component-oriented approach of combinatory logic synthesis.

2 Intersection Types for Mixins and Classes

2.1 Intersection and record types

We consider a type-free λ -calculus of extensible records, equipped with a merge operator. The term syntax is defined by the following grammar:

$$\begin{aligned} \Lambda_R \ni M, N, M_i & ::= x \mid (\lambda x.M) \mid (MN) \mid (M.l) \mid R \mid (M \oplus R) & \text{terms} \\ R & ::= \langle l_i = M_i \mid i \in I \rangle & \text{records} \end{aligned}$$

where $x \in \mathbf{Var}$ and $l \in \mathbf{Label}$ range over denumerably many *term variables* and *labels* respectively, and the sets of indexes I are finite. Free and bound variables are defined as usual for ordinary λ -calculus, and we name Λ_R^0 the set of all closed terms in Λ_R ; terms are identified up to renaming of bound variables and $M\{N/x\}$ denotes capture avoiding substitution of N for x in M . We adopt notational conventions from [3]; in particular application associates to the left and external parentheses are omitted when unnecessary; also the dot notation for record selection takes precedence over λ , so that $\lambda x.M.l$ reads as $\lambda x.(M.l)$. If not stated otherwise \oplus also associates to the left, and we avoid external parentheses when unnecessary.

Terms $R \equiv \langle l_i = M_i \mid i \in I \rangle$ (writing \equiv for syntactic identity) represent *records*, with fields l_i and M_i as the respective values; we set $\text{lbl}(\langle l_i = M_i \mid i \in I \rangle) = \{l_i \mid i \in I\}$. The term $M.l$ is *field selection* and $M \oplus R$ is *record merge*. In particular if R_1 and R_2 are records then $R_1 \oplus R_2$ is the record with as fields the union of the fields of R_1 and R_2 and as values those of the original records but in case of ambiguity, where the values in R_2 prevail. The syntactic constraint that R is a record in $M \oplus R$ is justified after Definition 8.

The actual meaning of these operations is formalized by the following reduction relation:

► **Definition 1** (Λ_R reduction). Reduction $\longrightarrow \subseteq \Lambda_R^2$ is the least compatible relation such that:

$$\begin{aligned} (\beta) \quad & (\lambda x.M)N \longrightarrow M\{N/x\} \\ (sel) \quad & \langle l_i = M_i \mid i \in I \rangle.l_j \longrightarrow M_j \quad \text{if } j \in I \\ (\oplus) \quad & \langle l_i = M_i \mid i \in I \rangle \oplus \langle l_j = N_j \mid j \in J \rangle \longrightarrow \langle l_i = M_i, l_j = N_j \mid i \in I \setminus J, j \in J \rangle \end{aligned}$$

We claim that \longrightarrow^* is Church-Rosser, and that $(M \oplus R_1) \oplus R_2$ is equivalent to $M \oplus (R_1 \oplus R_2)$ under any reasonable observational semantics (e.g. by extending to Λ_R applicative bisimulation from the lazy λ -calculus).

Record merge subsumes field update: $M.l := N \equiv M \oplus (l = N)$, but merge is not uniformly definable in terms of update as long as labels are not expressions in the calculus.

In the spirit of Curry's assignment of polymorphic types and of intersection types in particular, types are introduced as a syntactical tool to capture semantic properties of terms, rather than as constraints to term formation.

► **Definition 2** (Intersection types for Λ_R).

$$\begin{aligned} \mathbb{T} \ni \sigma, \sigma_i &::= a \mid \omega \mid \sigma_1 \rightarrow \sigma_2 \mid \sigma_1 \cap \sigma_2 \mid \rho & \text{types} \\ \mathbb{T}_{\langle \rangle} \ni \rho, \rho_i &::= \langle \rangle \mid \langle l : \sigma \rangle \mid \rho_1 + \rho_2 \mid \rho_1 \cap \rho_2 & \text{record types} \end{aligned}$$

where a ranges over *type constants*, $l \in \mathbf{Label}$.

We use σ, τ , possibly with sub and superscripts, for types in \mathbb{T} and ρ, ρ_i , possibly with superscripts, for record types in $\mathbb{T}_{\langle \rangle}$ only. Note that \rightarrow associates to the right, and \cap binds stronger than \rightarrow . As with intersection type systems for the λ -calculus, the intended meaning of types are sets, provided a set theoretic interpretation of constants a .

Following [4], type semantics is given axiomatically by means of the subtyping relation \leq , that can be interpreted as subset inclusion. It is the least pre-order over \mathbb{T} such that:

► **Definition 3** (Type inclusion: arrow and intersection types).

$$\begin{aligned} \sigma &\leq \omega, & \omega &\leq \omega \rightarrow \omega, \\ \sigma \cap \tau &\leq \sigma, & \sigma \cap \tau &\leq \tau, & \sigma \leq \tau_1 \ \& \ \sigma \leq \tau_2 &\Rightarrow \sigma \leq \tau_1 \cap \tau_2, \\ (\sigma \rightarrow \tau_1) \cap (\sigma \rightarrow \tau_2) &\leq \sigma \rightarrow \tau_1 \cap \tau_2 & \sigma_2 &\leq \sigma_1 \ \& \ \tau_1 \leq \tau_2 &\Rightarrow \sigma_1 \rightarrow \tau_1 \leq \sigma_2 \rightarrow \tau_2 \end{aligned}$$

We write $\sigma = \tau$ for $\sigma \leq \tau$ and $\tau \leq \sigma$.

► **Definition 4** (Type inclusion: record types).

$$\begin{aligned} \langle l : \sigma \rangle \cap \langle l : \tau \rangle &\leq \langle l : \sigma \cap \tau \rangle, & \sigma &\leq \tau &\Rightarrow \langle l : \sigma \rangle &\leq \langle l : \tau \rangle, \\ \langle l : \sigma \rangle &\leq \langle \rangle, & \langle l : \sigma \rangle + \langle \rangle &= \langle l : \sigma \rangle = \langle \rangle + \langle l : \sigma \rangle, \\ \langle l : \sigma \rangle + \langle l : \tau \rangle &= \langle l : \tau \rangle, & \langle l : \sigma \rangle + \langle l' : \tau \rangle &= \langle l : \sigma \rangle \cap \langle l' : \tau \rangle & \quad (l \neq l'), \\ \langle l : \sigma \rangle + (\langle l : \tau \rangle \cap \rho) &= \langle l : \tau \rangle \cap \rho, & \langle l : \sigma \rangle + (\langle l' : \tau \rangle \cap \rho) &= \langle l' : \tau \rangle \cap (\langle l : \sigma \rangle + \rho) & \quad (l \neq l'). \end{aligned}$$

While Definition 3 is standard after [4], comments on Definition 4 are in order. Type $\langle \rangle$ is the type of all records. Type $\langle l : \sigma \rangle$ is a unary record type, whose meaning is the set of records having at least a field labeled by l , with value of type σ ; therefore $\langle l : \sigma \rangle \cap \langle l : \tau \rangle$ is the type of records having label l with values both of type σ and τ , that is of type $\sigma \cap \tau$. In fact the equation $\langle l : \sigma \rangle \cap \langle l : \tau \rangle = \langle l : \sigma \cap \tau \rangle$ is derivable. On the other hand $\langle l : \sigma \rangle \cap \langle l' : \tau \rangle$, with $l \neq l'$, is the type of records having fields labeled by l and l' , with values of type σ and τ respectively. It follows that intersection of record types can be used to express properties of records with arbitrary (though finitely) many fields, which justifies the abbreviation $\langle l_i : \sigma_i \mid i \in I \neq \emptyset \rangle = \bigcap_{i \in I} \langle l_i : \sigma_i \rangle$ and $\langle l_i : \sigma_i \mid i \in \emptyset \rangle = \langle \rangle$, where we assume that the l_i are pairwise distinct. Finally, as it will be apparent from Definition 8 below, $\rho_1 + \rho_2$ is the type of all records obtained by merging a record of type ρ_1 with a record of type ρ_2 , which is intended to type \oplus that is at the same time a record extension and field updating operation. Since this is the distinctive feature of the system introduced here, we comment on this by means of a few lemmas, illustrating its properties.

► **Lemma 5.**

1. $(\forall j \in J \subseteq I. \sigma_j \leq \tau_j) \Rightarrow \langle l_i : \sigma_i \mid i \in I \rangle \leq \langle l_j : \tau_j \mid j \in J \rangle$,
2. $\langle l_i : \sigma_i \mid i \in I \rangle + \langle l_j : \tau_j \mid j \in J \rangle = \langle l_i : \sigma_i, m_j : \tau_j \mid i \in I \setminus J, j \in J \rangle$,
3. $\forall \rho \in \mathbb{T}_{\langle \rangle}. \exists \langle l_i : \sigma_i \mid i \in I \rangle. \rho = \langle l_i : \sigma_i \mid i \in I \rangle$.

Part (1) of Lemma 5 states that subtyping among intersection of unary record types subsumes subtyping in width and depth of ordinary record types from the literature. Part (2) shows that the $+$ type constructor reflects at the level of types the operational behavior of the merge operator \oplus . Part (3) says that any record type is equivalent to an intersection of unary record types; this implies that types of the form $\rho_1 + \rho_2$ are eliminable in principle. However they play a key role in typing mixins, motivating the issue of control of negative information in the synthesis process: see sections 2.2 and 4. More properties of subtyping record types w.r.t. $+$ and \cap are listed in the next lemma. Let us preliminary define the map $lbl : \mathbb{T}_{\langle \rangle} \rightarrow \wp(\mathbf{Label})$ (where $\wp(\mathbf{Label})$ is the powerset of \mathbf{Label}) by:

$$lbl(\langle l : \sigma \rangle) = \{l\}, \quad lbl(\rho_1 \cap \rho_2) = lbl(\rho_1 + \rho_2) = lbl(\rho_1) \cup lbl(\rho_2).$$

Then we immediately have:

► **Lemma 6.**

1. $\rho_1 = \rho_2 \Rightarrow lbl(\rho_1) = lbl(\rho_2)$,
2. $lbl(\rho_1) \cap lbl(\rho_2) = \emptyset \Rightarrow \rho_1 + \rho_2 = \rho_1 \cap \rho_2$.

In (1) above $\rho_1 = \rho_2$ is $\rho_1 \leq \rho_2 \leq \rho_1$. About (2) note that condition $lbl(\rho_1) \cap lbl(\rho_2) = \emptyset$ is essential, since $\rho_1 + \rho_2 \neq \rho_2 + \rho_1$ in general, as it immediately follows by Lemma 5.2.

► **Lemma 7.**

1. $(\rho_1 + \rho_2) + \rho_3 = \rho_1 + (\rho_2 + \rho_3)$,
2. $(\rho_1 \cap \rho_2) + \rho_3 = (\rho_1 + \rho_3) \cap (\rho_2 + \rho_3)$,
3. $\rho_1 \leq \rho_2 \Rightarrow \rho_1 + \rho_3 \leq \rho_2 + \rho_3$,
4. $\rho_1 + \rho_2 = \rho_1 \cap \rho_2 \Leftrightarrow \rho_1 + \rho_2 \leq \rho_1$.

► **Remark.** In general $\rho_1 + (\rho_2 \cap \rho_3) \neq (\rho_1 + \rho_3) \cap (\rho_2 + \rho_3)$: take $\rho_1 \equiv \langle l_1 : \sigma_1, l_2 : \sigma_2 \rangle$, $\rho_2 \equiv \langle l_1 : \sigma'_1 \rangle$ and $\rho_3 \equiv \langle l_2 : \sigma'_2 \rangle$, with $\sigma_1 \neq \sigma'_1$ and $\sigma_2 \neq \sigma'_2$. Then we have: $\rho_1 + (\rho_2 \cap \rho_3) = \rho_1 + \langle l_1 : \sigma'_1, l_2 : \sigma'_2 \rangle = \langle l_1 : \sigma'_1, l_2 : \sigma'_2 \rangle$, while $(\rho_1 + \rho_3) \cap (\rho_2 + \rho_3) = \langle l_1 : \sigma_1, l_2 : \sigma'_2 \rangle \cap \langle l_1 : \sigma'_1, l_2 : \sigma_2 \rangle = \langle l_1 : \sigma_1 \cap \sigma'_1, l_2 : \sigma_2 \cap \sigma'_2 \rangle$. The last example suggests that $(\rho_1 + \rho_3) \cap (\rho_2 + \rho_3) \leq \rho_1 + (\rho_2 \cap \rho_3)$. On the other hand $\rho_2 \leq \rho_3 \not\Rightarrow \rho_1 + \rho_2 \leq \rho_1 + \rho_3$. Indeed:

$$\begin{aligned} \langle l_0 : \sigma_1, l_1 : \sigma_2 \rangle + \langle l_1 : \sigma_3, l_2 : \sigma_4 \rangle &= \langle l_0 : \sigma_1, l_1 : \sigma'_1, l_2 : \sigma_2 \rangle \\ &\not\leq \langle l_0 : \sigma_0, l_1 : \sigma_1, l_2 : \sigma_2 \rangle \quad \text{if } \sigma'_1 \not\leq \sigma_1 \\ &= \langle l_0 : \sigma_1, l_1 : \sigma_2 \rangle + \langle l_2 : \sigma_4 \rangle \end{aligned}$$

even if $\langle l_1 : \sigma_1, l_2 : \sigma_2 \rangle \leq \langle l_2 : \sigma_2 \rangle$. From this and (3) of Lemma 7, we conclude that $+$ is monotonic in its first argument, but not in its second one.

We come now to the type assignment system. A *basis* (also called a context in the literature) is a finite set $\Gamma = \{x_1 : \sigma_n, \dots, x_n : \sigma_n\}$, where the variables x_i are pairwise distinct; we set $dom(\Gamma) = \{x \mid \exists \sigma. x : \sigma \in \Gamma\}$ and we write $\Gamma, x : \sigma$ for $\Gamma \cup \{x : \sigma\}$ where $x \notin dom(\Gamma)$. Then we consider the following extension of the system in [4], also called **BCD** in the literature.

► **Definition 8** (Type Assignment). The rules of the assignment system are:

$$\begin{array}{c}
\frac{x : \sigma \in \Gamma}{\Gamma \vdash x : \sigma} (Ax) \qquad \frac{\Gamma, x : \sigma \vdash M : \tau}{\Gamma \vdash \lambda x.M : \sigma \rightarrow \tau} (\rightarrow I) \\
\frac{\Gamma \vdash M : \sigma \rightarrow \tau \quad \Gamma \vdash N : \sigma}{\Gamma \vdash MN : \tau} (\rightarrow E) \qquad \frac{\Gamma \vdash M : \sigma \quad \Gamma \vdash M : \tau}{\Gamma \vdash M : \sigma \cap \tau} (\cap) \\
\frac{}{\Gamma \vdash M : \omega} (\omega) \qquad \frac{\Gamma \vdash M : \sigma \quad \sigma \leq \tau}{\Gamma \vdash M : \tau} (\leq) \\
\frac{}{\Gamma \vdash \langle l_i = M_i \mid i \in I \rangle : \langle \rangle} (\langle \rangle) \qquad \frac{\Gamma \vdash M_k : \sigma \quad k \in I}{\Gamma \vdash \langle l_i = M_i \mid i \in I \rangle : \langle l_k : \sigma \rangle} (rec) \\
\frac{\Gamma \vdash M : \langle l : \sigma \rangle}{\Gamma \vdash M.l : \sigma} (sel) \qquad \frac{\Gamma \vdash M : \rho_1 \quad \Gamma \vdash R : \rho_2 \quad (*)}{\Gamma \vdash M \oplus R : \rho_1 + \rho_2} (+)
\end{array}$$

where $(*)$ in rule $(+)$ is the side condition: $lbl(R) = lbl(\rho_2)$.

Using Lemma 5.1, the following rule is easily shown to be admissible:

$$\frac{\Gamma \vdash M_j : \sigma_j \quad \forall j \in J \subseteq I}{\Gamma \vdash \langle l_i = M_i \mid i \in I \rangle : \langle l_j : \sigma_j \mid j \in J \rangle} (rec')$$

Contrary to this, the side condition $(*)$ of rule $(+)$ is equivalent to “exact” record typing in [10], disallowing record subtyping in width. Such a condition is necessary for soundness of typing. Indeed suppose that $\Gamma \vdash M_0 : \sigma$ and $\Gamma \vdash M'_0 : \sigma'_0$ but $\Gamma \not\vdash M'_0 : \sigma_0$; then without $(*)$ we could derive:

$$\frac{\Gamma \vdash \langle l_0 = M_0 \rangle : \langle l_0 : \sigma_0 \rangle \quad \Gamma \vdash \langle l_0 = M'_0, l_1 : \sigma_1 \rangle : \langle l_1 : \sigma_1 \rangle}{\Gamma \vdash \langle l_0 = M_0 \rangle \oplus \langle l_0 = M'_0, l_1 : \sigma_1 \rangle : \langle l_0 : \sigma_0, l_1 : \sigma_1 \rangle}$$

from which we obtain that $\Gamma \vdash (\langle l_0 = M_0 \rangle \oplus \langle l_0 = M'_0, l_1 : \sigma_1 \rangle).l_0 : \sigma_0$ breaking subject reduction, since $(\langle l_0 = M_0 \rangle \oplus \langle l_0 = M'_0, l_1 : \sigma_1 \rangle).l_0 \rightarrow^* M'_0$. The essential point is that proving that $\Gamma \vdash N : \langle l : \sigma \rangle$ doesn't imply that $l' \notin lbl(R')$ for any $l' \neq l$, which follows only by the uncomputable (not even r.e.) statement that $\Gamma \not\vdash N : \langle l' : \omega \rangle$, a negative information.

This explains the restriction to record terms as the second argument of \oplus : in fact allowing $M \oplus N$ to be well formed for an arbitrary N we might have $N \equiv x$ in $\lambda x.(M \oplus x)$. But extending lbl to all terms in Λ_R is not possible without severely limiting the expressiveness of the assignment system. In fact to say that $lbl(N) = lbl(R)$ if $N \rightarrow^* R$ would make the lbl function non computable; on the other hand putting $lbl(x) = \emptyset$, which is the only reasonable and conservative choice as we do not know of possible substitutions for x in $\lambda x.(M \oplus x)$, implies that the latter term has type $\omega \rightarrow \omega = \omega$ at best.

As a final remark, let us observe that we do not adopt exact typing of records in general, but only for typing the right-hand side of \oplus -terms, a feature that will be essential when typing mixins.

► **Lemma 9.** *Let $\sigma \neq \omega$:*

1. $\Gamma \vdash x : \sigma \iff \exists \tau. x : \tau \in \Gamma \ \& \ \tau \leq \sigma$,
2. $\Gamma \vdash \lambda x.M : \sigma \iff \exists I, \sigma_i, \tau_i. \Gamma, x : \sigma_i \vdash M : \tau_i \ \& \ \bigcap_{i \in I} \sigma_i \rightarrow \tau_i \leq \sigma$,
3. $\Gamma \vdash MN : \sigma \iff \exists \tau. \Gamma \vdash M : \tau \rightarrow \sigma \ \& \ \Gamma \vdash N : \tau$,
4. $\Gamma \vdash \langle l_i = M_i \mid i \in I \rangle : \sigma \iff \forall i \in I \ \exists \sigma_i. \Gamma \vdash M_i : \sigma_i \ \& \ \langle l_i : \sigma_i \mid i \in I \rangle \leq \sigma$,
5. $\Gamma \vdash M.l : \sigma \iff \Gamma \vdash M : \langle l : \sigma \rangle$,
6. $\Gamma \vdash M \oplus R : \sigma \iff \exists \rho_1, \rho_2. \Gamma \vdash M : \rho_1 \ \& \ \Gamma \vdash R : \rho_2 \ \& \ lbl(R) = lbl(\rho_2) \ \& \ \rho_1 + \rho_2 \leq \sigma$.

► **Theorem 10** (Subject reduction). $\Gamma \vdash M : \sigma \ \& \ M \longrightarrow N \Rightarrow \Gamma \vdash N : \sigma$.

Proof Sketch. The proof is by cases of reduction rules, using Lemma 9. The only relevant case is when $M \equiv R_1 \oplus R_2$, with $R_1 \equiv \langle l_i = M_i \mid i \in I \rangle$ and $R_2 \equiv \langle l_j = N_j \mid j \in J \rangle$, and $N \equiv \langle l_i = M_i, l_j = N_j \mid i \in I \setminus J, j \in J \rangle$. By Lemma 9.6 we may suppose w.l.o.g. that $\sigma = \rho_1 + \rho_2$, and that $\rho_1 = \langle l_i : \sigma_i \mid i \in I' \rangle$ for some $I' \subseteq I$, and $\rho_2 = \langle l_j : \tau_j \mid j \in J' \rangle$ for some $J' \subseteq J$; but also we know that $J' = J$, because of condition (*).

Now by Lemma 5.2, $\rho_1 + \rho_2 = \langle l_i : \sigma_i, l_j : \tau_j \mid i \in I' \setminus J, j \in J \rangle$; on the other hand by Lemma 9.4, we know that $\Gamma \vdash M_i : \sigma_i$ for all $i \in I'$, $\Gamma \vdash N_j : \tau_j$ for all $j \in J$, and therefore we conclude that $\Gamma \vdash N : \langle l_i : \sigma_i, l_j : \tau_j \mid i \in I' \setminus J, j \in J \rangle$ by multiple applications of rules (*rec*) and (\cap). ◀

2.2 Class and Mixin combinators

The following definition of classes and mixins is inspired by [15] and [10] respectively, though with some departures to be discussed below. To make the description more concrete, in the examples we add constants to Λ_R .

Recall that a *combinator* is a term in Λ_R^0 , namely a closed term. Let \mathbf{Y} be Curry's fixed point combinator: $\lambda f.(\lambda x.f(xx))(\lambda x.f(xx))$ (the actual definition of \mathbf{Y} is immaterial, since all fixed point combinators have the same types in **BCD**).

► **Definition 11.** Let `myClass`, `state` and `argClass` be (term) variables and \mathbf{Y} be a fixed point combinator; then we define the following sets of combinators:

$$\begin{aligned} \text{Class: } C & ::= \mathbf{Y}(\lambda \text{myClass } \lambda \text{state}. \langle l_i = N_i \mid i \in I \rangle) \\ \text{Mixin: } M & ::= \lambda \text{argClass}. \mathbf{Y}(\lambda \text{myClass } \lambda \text{state}. (\text{argClass } \text{state}) \oplus \langle l_i = N_i \mid i \in I \rangle) \end{aligned}$$

We define \mathcal{C} and \mathcal{M} as the sets of classes and mixins respectively.

To illustrate this definition let us use the abbreviation `let $x = N$ in $M \equiv M\{N/x\}$` . Then a class combinator $C \in \mathcal{C}$ can be written in a more perspicuous way as follows:

$$C \equiv \mathbf{Y}(\lambda \text{myClass } \lambda \text{state}. \text{let } \text{self} = (\text{myClass } \text{state}) \text{ in } \langle l_i = N_i \mid i \in I \rangle). \quad (1)$$

A *class* is the fixed point of a function, the *class definition*, mapping a recursive definition of the class itself and a state S , that is the value or a record of values in general, for the instance variables of the class, into a record $\langle l_i = N_i \mid i \in I \rangle$ of *methods*. A class C is instantiated to an *object* $O \equiv C S$ by applying the class C to a state S . Hence we have:

$$O \equiv C S \longrightarrow^* \text{let } \text{self} = (C S) \text{ in } \langle l_i = N_i \mid i \in I \rangle,$$

where the variable `self` is used in the method bodies N_i to call other methods from the same object. Note that the recursive parameter `myClass` might occur in the N_i in subterms other than `(myClass state)`, and in particular $N_i\{C/\text{myClass}\}$ might contain a subterm $C S'$, where S' is a state possibly different than S ; even C itself might be returned as the value of a method. Classes are the same as in [15] §4, but for the explicit identification of `self` with `(myClass state)`.

We come now to typing of classes. Let $R = \langle l_i = N_i \mid i \in I \rangle$, and suppose that $C \equiv \mathbf{Y}(\lambda \text{myClass } \lambda \text{state}. R) \in \mathcal{C}$. To type C we must find a type σ (a type of its state) and a sequence of types $\rho_1, \dots, \rho_n \in \mathbb{T}_{\langle \rangle}$ such that for all $i < n$:

$$\text{myClass} : \sigma \rightarrow \rho_i, \text{state} : \sigma \vdash R : \rho_{i+1}.$$

Note that this is always possible for any n : in the worst case, we can take $\rho_i = \langle l_i : \omega \mid i \in I \rangle$ for all $0 < i \leq n$. In general one has more expressive types, depending on the typings of the N_i in R (see example 12 below). It follows that:

$$\vdash \lambda \text{myClass } \lambda \text{state}. R : (\omega \rightarrow \rho_1) \cap \bigcap_{1 \leq i < n} (\sigma \rightarrow \rho_i) \rightarrow (\sigma \rightarrow \rho_{i+1}),$$

and therefore, by using the fact that $\vdash \mathbf{Y} : (\omega \rightarrow \tau_1) \cap \dots \cap (\tau_{n-1} \rightarrow \tau_n) \rightarrow \tau_n$ for arbitrary types τ_1, \dots, τ_n , we conclude that the typing of classes has the following shape (where $\rho = \rho_n$):

$$\vdash C \equiv \mathbf{Y}(\lambda \text{myClass } \lambda \text{state}. \langle l_i = N_i \mid i \in I \rangle) : \sigma \rightarrow \rho \quad (2)$$

In conclusion the type of a class C is the arrow from the type of the state σ to a type ρ of its instances.

► **Example 12.** The class `Point` has an integer state and contains the method `get` to retrieve the state, `set` to update the current state and `shift` to add a value to the current state.

```
Point =  $\mathbf{Y}(\lambda \text{myClass}. \lambda \text{state}. \text{let } \text{self} = \text{myClass } \text{state } \text{in}$ 
   $\langle \text{get} = \text{state}, \text{set} = \lambda \text{state}'. \text{state}', \text{shift} = \lambda d. \text{self}. \text{set}(\text{self}. \text{get} + d) \rangle)$ 
```

Note that in a setting without references, we rely on purely functional state updates. Therefore, every function returns a new state that can be used to construct the new object. An example for this is `set`, which just returns the new state.

To type `Point` we have $\vdash \text{Point} : \text{Int} \rightarrow \rho_{\text{Point}}$ where $\rho_{\text{Point}} = \rho_2$ and

$$\begin{aligned} \rho_1 &= \langle \text{get} : \text{Int}, \text{set} : \text{Int} \rightarrow \text{Int}, \text{shift} : \omega \rangle \\ \rho_2 &= \langle \text{get} : \text{Int}, \text{set} : \text{Int} \rightarrow \text{Int}, \text{shift} : \text{Int} \rightarrow \text{Int} \rangle \end{aligned}$$

using $\mathbf{Y} : (\omega \rightarrow \text{Int} \rightarrow \rho_1) \cap ((\text{Int} \rightarrow \rho_1) \rightarrow \text{Int} \rightarrow \rho_2) \rightarrow \text{Int} \rightarrow \rho_2$

A *mixin* $M \in \mathcal{M}$ is a combinator such that, if $C \in \mathcal{C}$ then MC reduces to a new class $C' \in \mathcal{C}$, inheriting from C . Writing M in a more explicit way we obtain:

```
 $M \equiv \lambda \text{argClass}. \mathbf{Y}(\lambda \text{myClass } \lambda \text{state}. \text{let } \text{super} = (\text{argClass } \text{state}) \text{ in}$ 
   $\text{let } \text{self} = (\text{myClass } \text{state}) \text{ in}$ 
   $\text{super} \oplus \langle l_i = N_i \mid i \in I \rangle)$ 
```

In words, a mixin merges an instance CS of the input class C with a new state S together with a *difference* record $R \equiv \langle l_i = N_i \mid i \in I \rangle$, that would be written $\Delta(CS)$ in terms of [10]. Note that our mixins are not the same as class modifiers (also called wrappers e.g. in [9]) because the latter do not take the instantiation of a class as the value of `super`, but the class definition, namely the function defining the class *before* taking its fixed point.

Let $M \equiv \lambda \text{argClass}. \mathbf{Y}(\lambda \text{myClass } \lambda \text{state}. (\text{argClass } \text{state}) \oplus R) \in \mathcal{M}$; to type M we have to find types $\sigma^1, \sigma^2, \rho^1$ and a sequence $\rho_1^2, \dots, \rho_n^2 \in \mathbb{T}_{\langle \rangle}$ of record types such that for all $1 \leq i < n$ it is true that $\text{lbl}(R) = \text{lbl}(\rho_i^2)$ and such that, setting $\Gamma_0 = \{\text{argClass} : \sigma^1 \rightarrow \rho^1, \text{myClass} : \omega, \text{state} : \sigma^1 \cap \sigma^2\}$ and $\Gamma_i = \{\text{argClass} : \sigma^1 \rightarrow \rho^1, \text{myClass} : (\sigma^1 \cap \sigma^2) \rightarrow \rho^1 + \rho_i^2, \text{state} : \sigma^1 \cap \sigma^2\}$ for all $1 \leq i < n$, we may deduce for all $0 \leq i < n$:

$$\frac{\frac{\Gamma_i \vdash \text{state} : \sigma^1 \cap \sigma^2}{\Gamma_i \vdash \text{state} : \sigma^1} (\leq)}{\Gamma_i \vdash \text{argClass } \text{state} : \rho^1} (\rightarrow E) \quad \frac{\Gamma_i \vdash R : \rho_{i+1}^2 \quad \text{lbl}(R) = \text{lbl}(\rho_{i+1}^2)}{\Gamma_i \vdash (\text{argClass } \text{state}) \oplus R : \rho^1 + \rho_{i+1}^2} (+)$$

Hence for all $0 \leq i < n$ we can derive the typing judgment:

$$\begin{aligned} \text{argClass} : \sigma^1 \rightarrow \rho^1 \vdash \lambda \text{myClass } \lambda \text{state. } (\text{argClass } \text{state}) \oplus R : \\ ((\sigma^1 \cap \sigma^2) \rightarrow (\rho^1 + \rho_i^2)) \rightarrow (\sigma^1 \cap \sigma^2) \rightarrow (\rho^1 + \rho_{i+1}^2) \end{aligned}$$

and therefore, by reasoning as for classes, we get (setting $\rho^2 = \rho_n^2$):

$$\begin{aligned} \vdash M \equiv \lambda \text{argClass. } \mathbf{Y}(\lambda \text{myClass } \lambda \text{state. } (\text{argClass } \text{state}) \oplus R) : \\ (\sigma^1 \rightarrow \rho^1) \rightarrow (\sigma^1 \cap \sigma^2) \rightarrow (\rho^1 + \rho^2) \quad (3) \end{aligned}$$

Spelling out this type, we can say that σ^1 is a type of the state of the argument-class of M ; $\sigma^1 \cap \sigma^2$ is the type of the state of the resulting class, that refines σ^1 . ρ^1 expresses the requirements of M about the methods of the argument-class, i.e. what is assumed to hold for the usages of **super** and **argClass** in R to be properly typed; $\rho^1 + \rho^2$ is a type of the record of methods of the refined class, resulting from the merge of the methods of the argument-class with those of the difference R ; since in general there will be overridden methods, whose types might be incompatible, the $+$ type constructor cannot be replaced by intersection.

► **Example 13.** The mixin `Movable`, provided an argument class that contains a `set` and a `shift` method, creates a new class with (potentially overwritten) methods `set` and `move` along with delegated methods. The method `set` is fixed to set the underlying state to 1 and the method `move` shifts the state by 1.

```
Movable = λargClass.Y(λmyClass.λstate.let super = argClass state in
                    let self = myClass state in
                    super ⊕ ⟨set = super.set(1), move = self.shift(1)⟩)
```

Note that to update `self` and `super` one can use `myClass` and `argClass`. Generalizing for all $\rho \in \mathbb{T}_{\langle \rangle}$ following the argumentation above we can choose:

$$\begin{aligned} \rho^1 &= \rho \cap \langle \text{set} : \text{Int} \rightarrow \text{Int}, \text{shift} : \text{Int} \rightarrow \text{Int} \rangle & \rho^2 &= \langle \text{set} : \text{Int}, \text{move} : \text{Int} \rangle \\ \sigma^1 &= \text{Int} & \sigma^2 &= \omega \\ \rho_1^2 &= \langle \text{set} : \text{Int}, \text{move} : \omega \rangle & \rho_2^2 &= \langle \text{set} : \text{Int}, \text{move} : \text{Int} \rangle \end{aligned}$$

Using these choices we can type \mathbf{Y} by

$$\vdash \mathbf{Y} : (\omega \rightarrow (\text{Int} \rightarrow \rho^1 + \rho_1^2)) \cap ((\text{Int} \rightarrow \rho^1 + \rho_1^2) \rightarrow (\text{Int} \rightarrow \rho^1 + \rho_2^2)) \rightarrow (\text{Int} \rightarrow \rho^1 + \rho_2^2)$$

and obtain the typings of `Movable` for all ρ :

$$\frac{\dots}{\frac{\vdash \text{Movable} : (\text{Int} \rightarrow \rho \cap \langle \text{set} : \text{Int} \rightarrow \text{Int}, \text{shift} : \text{Int} \rightarrow \text{Int} \rangle) \rightarrow (\text{Int} \rightarrow \rho^1 + \rho_2^2)}{\vdash \text{Movable} : (\text{Int} \rightarrow \rho \cap \langle \text{set} : \text{Int} \rightarrow \text{Int}, \text{shift} : \text{Int} \rightarrow \text{Int} \rangle) \rightarrow (\text{Int} \rightarrow \rho + \langle \text{set} : \text{Int}, \text{move} : \text{Int} \rangle)} (\leq)}$$

3 Encoding of Record Types in Bounded Combinatory Logic

Our main goal is to combine type information given by intersection types for Λ_R and the capabilities of the logical programming language given by BCL inhabitation to synthesize meaningful mixin compositions as terms of a combinatory logic. Such *combinatory terms* are formed by application of combinators from a *repository* (combinatory logic context) Δ .

► **Definition 14** (Combinatory Term). $E, E' ::= C \mid (E E'), C \in \text{dom}(\Delta)$

We create repositories of typed combinators that can be considered logic programs for the existing BCL synthesis framework (CL)S [7] to reason about semantics of such compositions. The underlying type system of (CL)S is an extension of the intersection type system **BCD** [4] by covariant constructors. The extended type system \mathbb{T}_C , while suited for synthesis, is flexible enough to encode record types and features of $+$. While (CL)S implements covariant constructors of arbitrary arity, we only use unary constructors, which are sufficient for our encoding.

► **Definition 15** (Intersection Types with Constructors \mathbb{T}_C). The set \mathbb{T}_C is given by:

$$\mathbb{T}_C \ni \sigma, \tau, \tau_1, \tau_2 ::= a \mid \alpha \mid \omega \mid \tau_1 \rightarrow \tau_2 \mid \tau_1 \cap \tau_2 \mid c(\tau)$$

where a ranges over constants, α over type variables and c over unary constructors \mathbb{C} .

\mathbb{T}_C adds the following two subtyping axioms to the **BCD** system

$$\tau_1 \leq \tau_2 \Rightarrow c(\tau_1) \leq c(\tau_2) \quad c(\tau_1) \cap c(\tau_2) \leq c(\tau_1 \cap \tau_2)$$

The additional axioms ensure *constructor distributivity*, i.e., $c(\tau_1) \cap c(\tau_2) = c(\tau_1 \cap \tau_2)$.

► **Definition 16** (Type Assignment in \mathbb{T}_C).

$$\frac{C : \tau \in \Delta \quad S \text{ Substitution}}{\Delta \vdash_{\text{BCL}} C : S(\tau)} \text{ (Var)} \quad \frac{\Delta \vdash_{\text{BCL}} E : \sigma \rightarrow \tau \quad \Delta \vdash_{\text{BCL}} E' : \sigma}{\Delta \vdash_{\text{BCL}} EE' : \tau} (\rightarrow E)$$

$$\frac{\Delta \vdash_{\text{BCL}} E : \sigma \quad \Delta \vdash_{\text{BCL}} E : \tau}{\Delta \vdash_{\text{BCL}} E : \sigma \cap \tau} (\cap) \quad \frac{\Delta \vdash_{\text{BCL}} E : \sigma \quad \sigma \leq \tau}{\Delta \vdash_{\text{BCL}} E : \tau} (\leq)$$

We extend the necessary property of *beta-soundness* and a notion of *paths* and *organized types* from [20] to constructors in the following way.

► **Lemma 17** (Extended Beta-Soundness). *If $\bigcap_{i \in I} (\sigma_i \rightarrow \tau_i) \cap \bigcap_{j \in J} c_j(\tau_j) \cap \bigcap_{k \in K} \alpha_k \cap \bigcap_{k' \in K'} a_{k'} \leq c(\tau)$, then $\{j \in J \mid c_j = c\} \neq \emptyset$ and $\bigcap \{\tau_j \mid j \in J, c_j = c\} \leq \tau$.*

► **Definition 18** (Path). A *path* π is a type of the form: $\pi ::= a \mid \alpha \mid \tau \rightarrow \pi \mid c(\omega) \mid c(\pi)$, where α is a variable, τ is a type, c is a constructor and a is a constant.

► **Definition 19** (Organized Type). A type τ is called *organized*, if it is an intersection of paths $\tau \equiv \bigcap_{i \in I} \tau_i$, where τ_i for $i \in I$ are paths.

Similarly, we obtain the following property of subtyping w.r.t. organized types.

► **Lemma 20**. *Given two organized types $\tau \equiv \bigcap_{i \in I} \tau_i$ and $\sigma \equiv \bigcap_{j \in J} \sigma_j$, we have $\tau \leq \sigma$ iff for all $j \in J$ there exists an $i \in I$ with $\tau_i \leq \sigma_j$.*

Note that for all intersection types there exists an equivalent organized intersection type coinciding with the notion of strict intersection types [2].

For a set of typed combinators Δ and a type $\tau \in \mathbb{T}_C$ we say τ is *inhabitable* in Δ , if there exists a combinatory term E such that $\Delta \vdash_{\text{BCL}} E : \tau$. In the following we fix a finite set of labels $\mathcal{L} \subseteq \text{Label}$ that are used in the particular domain of interest for mixin composition synthesis.

Records as Unary Covariant Distributing Constructors

We define constructors $\langle\langle \cdot \rangle\rangle$ and $l(\cdot)$ for $l \in \mathcal{L}$ to represent record types using the following partial translation function

$$\llbracket \cdot \rrbracket : \mathbb{T} \rightarrow \mathbb{T}_{\mathbb{C}}, \llbracket \tau \rrbracket = \begin{cases} \tau & \text{if } \tau \equiv \omega \text{ or } \tau \equiv a \\ \llbracket \tau_1 \rrbracket \rightarrow \llbracket \tau_2 \rrbracket & \text{if } \tau \equiv \tau_1 \rightarrow \tau_2 \\ \llbracket \tau_1 \rrbracket \cap \llbracket \tau_2 \rrbracket & \text{if } \tau \equiv \tau_1 \cap \tau_2 \\ \langle\langle l(\tau) \rangle\rangle & \text{if } \tau \equiv \langle l : \tau \rangle \\ \langle\langle \omega \rangle\rangle & \text{if } \tau \equiv \langle \rangle \\ \text{undefined} & \text{else} \end{cases}$$

Since atomic records are covariant and distribute over \cap , the presented translation preserves subtyping. We have $\llbracket \langle l_i : \tau_i \mid i \in I \rangle \rrbracket = \llbracket \cap_{i \in I} \langle l_i : \tau_i \rangle \rrbracket = \cap_{i \in I} \langle\langle l_i(\llbracket \tau_i \rrbracket) \rangle\rangle$ if $I \neq \emptyset$.

Note that the translation function $\llbracket \cdot \rrbracket$ is not defined for types containing $+$ in \mathbb{T} . Additionally, $+$ has non-monotonic properties and therefore cannot be immediately represented by a covariant type constructor. Simply applying Lemma 5(3) is impossible, if the left-hand side of $+$ is all-quantified. There are two possibilities to deal with this situation. The first option is extending the type-system used for inhabitation. Here, the main difficulty is that existing versions of the inhabitation algorithm crucially rely on the separation of intersections into paths [20]. As demonstrated in the remark accompanying Lemma 7, it becomes unclear how to perform such a separation in the presence of the non-monotonic $+$ operation. The second option, pursued in the rest of this section, is to use the expressiveness of the logical programming language given by $\text{BCL}(\mathbb{T}_{\mathbb{C}})$ inhabitation. Specifically, encoding \mathbb{T} types containing $+$ as $\mathbb{T}_{\mathbb{C}}$ types accompanied by following repositories $\Delta_{\mathcal{L}}$ and $\Delta_{\emptyset(\mathcal{L})}$ suited for $\text{BCL}(\mathbb{T}_{\mathbb{C}})$ inhabitation. We introduce $|\mathcal{L}|$ distinct variables $\alpha_{l'}$ indexed by $l' \in \mathcal{L}$ and $2^{|\mathcal{L}|}$ distinct constructors $w_L(\cdot)$ indexed by $L \subseteq \mathcal{L}$.

$$\Delta_{\mathcal{L}} = \{W_{\{l\}} : w_{\{l\}}(\bigcap_{l' \in \mathcal{L} \setminus \{l\}} \langle\langle l'(\alpha_{l'}) \rangle\rangle) \mid l \in \mathcal{L}\}$$

$$\Delta_{\emptyset(\mathcal{L})} = \{W_L : w_{\{l_1\}}(\alpha) \rightarrow w_{\{l_2\}}(\alpha) \rightarrow \dots \rightarrow w_{\{l_k\}}(\alpha) \rightarrow w_L(\alpha) \mid k \geq 2, \{l_1, \dots, l_k\} = L \subseteq \mathcal{L}\}.$$

These repositories are purely logical in a sense that they do not represent terms in Λ_R but encode necessary side conditions in the logic program. In particular, we formalize the conditions for the absence of a label in a record type by the following Lemma 21. This encoding of negative information is crucial to encode non-monotonic properties of $+$.

► **Lemma 21.** *Let $l \in \mathcal{L}$ be a label and let $\tau \in \mathbb{T}$ be a type such that $\llbracket \tau \rrbracket$ is defined. $w_{\{l\}}(\llbracket \tau \rrbracket)$ is inhabitable in $\Delta_{\mathcal{L}} \cup \Delta_{\emptyset(\mathcal{L})}$ iff $\tau \in \mathbb{T}_{\langle \rangle} \cup \{\omega\}$ with $l \notin \text{lbl}(\tau)$.*

Proof Sketch. Let $l \in \mathcal{L}$ be a label.

(\Rightarrow) Let wlog. $\llbracket \tau \rrbracket \equiv \cap_{i \in I} \tau_i$ be an organized intersection type such that $w_{\{l\}}(\llbracket \tau \rrbracket)$ is inhabitable

in $\Delta_{\mathcal{L}} \cup \Delta_{\emptyset(\mathcal{L})}$. The only combinator with a matching target to inhabit $w_{\{l\}}(\llbracket \tau \rrbracket)$ is $W_{\{l\}}$. Due to distributivity of type constructors there exists a substitution S such that $S(w_{\{l\}}(\bigcap_{l' \in \mathcal{L} \setminus \{l\}} \langle\langle l'(\alpha_{l'}) \rangle\rangle)) \leq w_{\{l\}}(\llbracket \tau \rrbracket)$. By Lemma 20 followed by Lemma 17 we obtain

for each $i \in I$ that there exists an $l' \in \mathcal{L} \setminus \{l\}$ such that $\tau_i = \langle\langle l'(\llbracket \sigma_i \rrbracket) \rangle\rangle$ for some type $\sigma_i \in \mathbb{T}$. By definition of $\llbracket \cdot \rrbracket$ and distributivity we obtain $\tau \in \mathbb{T}_{\langle \rangle} \cup \{\omega\}$ with $l \notin \text{lbl}(\tau)$.

(\Leftarrow) Let $\tau = \bigcap_{l' \in L} \langle l' : \tau_{l'} \rangle$ (resp. $\langle \rangle$) for some $L \subseteq \mathcal{L} \setminus \{l\}$ and types $\tau_{l'} \in \mathbb{T}$ for $l' \in L$. We have

$$W_{\{l\}} : S(w_{\{l\}}(\bigcap_{l' \in \mathcal{L} \setminus \{l\}} \langle\langle l'(\alpha_{l'}) \rangle\rangle)) \leq w_{\{l\}}(\llbracket \tau \rrbracket) \text{ for a substitution } S(\alpha_{l'}) = \begin{cases} \llbracket \tau_{l'} \rrbracket & \text{if } l' \in L \\ \omega & \text{else} \end{cases} \blacktriangleleft$$

► **Lemma 22.** *Let $L \subseteq \mathcal{L}$ be a non-empty set of labels and let $\tau \in \mathbb{T}$ be a type such that $\llbracket \tau \rrbracket$ is defined. $w_L(\llbracket \tau \rrbracket)$ is inhabitable in $\Delta_{\mathcal{L}} \cup \Delta_{\wp(\mathcal{L})}$ iff $\tau \in \mathbb{T}_{\langle \rangle} \cup \{\omega\}$ with $L \cap \text{lbl}(\tau) = \emptyset$.*

Proof Sketch. Using W_L and Lemma 21 for each argument of W_L . ◀

Our intermediate goal is to use inhabitation in $\Delta_{\mathcal{L}} \cup \Delta_{\wp(\mathcal{L})}$ to translate types of the shape $\rho + \bigcap_{l \in L} \langle l : \tau_l \rangle$ into $\rho \cap \bigcap_{l \in L} \langle l : \tau_l \rangle$, which in general is incorrect. The following Lemma 23 describes sufficient circumstances where this translation holds.

► **Lemma 23** ($\Delta_{\mathcal{L}} \cup \Delta_{\wp(\mathcal{L})}$ Translation Soundness). *Let $L \subseteq \mathcal{L}$ be a non-empty set of labels, let $\rho \in \mathbb{T}_{\langle \rangle}$ be a type such that $\llbracket \rho \rrbracket$ is defined and let $\tau_l \in \mathbb{T}$ for $l \in L$ be types. If $w_L(\llbracket \rho \rrbracket)$ is inhabitable in $\Delta_{\mathcal{L}} \cup \Delta_{\wp(\mathcal{L})}$ then $\rho + \bigcap_{l \in L} \langle l : \tau_l \rangle = \rho \cap \bigcap_{l \in L} \langle l : \tau_l \rangle$.*

Proof Sketch. Since $w_L(\llbracket \rho \rrbracket)$ is inhabitable in $\Delta_{\mathcal{L}} \cup \Delta_{\wp(\mathcal{L})}$ we have $\text{lbl}(\rho) \cap L = \emptyset$. The result follows from Lemma 6 (2). ◀

Next, we show by the following Lemma 24 that inhabitation in $\Delta_{\mathcal{L}}$ is not too restrictive.

► **Lemma 24** ($\Delta_{\mathcal{L}} \cup \Delta_{\wp(\mathcal{L})}$ Translation Completeness). *Let $L \subseteq \mathcal{L}$ be a non-empty set of labels, let $\tau_l \in \mathbb{T}$ for $l \in L$ be types, let $\rho \in \mathbb{T}_{\langle \rangle}$ be a type such that $\llbracket \rho \rrbracket$ is defined. There exists a type $\rho' \in \mathbb{T}_{\langle \rangle}$ such that $w_L(\llbracket \rho' \rrbracket)$ is inhabitable in $\Delta_{\mathcal{L}} \cup \Delta_{\wp(\mathcal{L})}$ and $\rho' \cap \bigcap_{l \in L} \langle l : \tau_l \rangle \leq \rho + \bigcap_{l \in L} \langle l : \tau_l \rangle$.*

Proof Sketch. Given $\rho = \bigcap_{l' \in L'} \langle l' : \tau_{l'} \rangle$ choose $\rho' = \bigcap_{l' \in L' \setminus L} \langle l' : \tau_{l'} \rangle$ (resp. $\rho' = \langle \rangle$ if $L' \setminus L = \emptyset$). By Lemma 22 $w_L(\llbracket \rho' \rrbracket)$ is inhabitable in $\Delta_{\mathcal{L}} \cup \Delta_{\wp(\mathcal{L})}$ and $\rho' \cap \bigcap_{l \in L} \langle l : \tau_l \rangle \leq \rho + \bigcap_{l \in L} \langle l : \tau_l \rangle$. ◀

Note that in Lemma 24 the type ρ' can be chosen greater than ρ only due to the non-monotonic properties of $+$.

4 Mixin Composition Synthesis by Type Inhabitation

In this section we denote type assignment in Λ_R by $\vdash_{\langle \rangle}$ and fix the following ingredients:

- A finite set of classes \mathcal{C} .
- For each $C \in \mathcal{C}$ types $\sigma_C \in \mathbb{T}$, $\rho_C \in \mathbb{T}_{\langle \rangle}$ such that $\llbracket \sigma_C \rightarrow \rho_C \rrbracket$ is defined and $\emptyset \vdash_{\langle \rangle} C : \sigma_C \rightarrow \rho_C$.
- A finite set of mixins \mathcal{M} .
- For each $M \in \mathcal{M}$ types $\sigma_M \in \mathbb{T}$ and $\rho_M^1, \rho_M^2 \in \mathbb{T}_{\langle \rangle}$ such that $\llbracket \sigma_M \rrbracket, \llbracket \rho_M^1 \rrbracket, \llbracket \rho_M^2 \rrbracket$ are defined and for all types $\rho \in \mathbb{T}_{\langle \rangle}$ we have $\emptyset \vdash_{\langle \rangle} M : (\sigma_M \rightarrow \rho \cap \rho_M^1) \rightarrow (\sigma_M \rightarrow \rho + \rho_M^2)$.
- For each $M \in \mathcal{M}$ the non-empty set of labels $L_M = \text{lbl}(\rho_M^2) \subseteq \mathcal{L}$.

We translate given classes and mixins to the following a repository $\Delta_{\mathcal{L}}^{\mathcal{C}, \mathcal{M}}$ of combinators

$$\begin{aligned} \Delta_{\mathcal{L}}^{\mathcal{C}, \mathcal{M}} = & \{C : \llbracket \sigma_C \rightarrow \rho_C \rrbracket \mid C \in \mathcal{C}\} \\ & \cup \{M : w_{L_M}(\alpha) \rightarrow (\llbracket \sigma_M \rrbracket \rightarrow \alpha \cap \llbracket \rho_M^1 \rrbracket) \rightarrow (\llbracket \sigma_M \rrbracket \rightarrow \alpha \cap \llbracket \rho_M^2 \rrbracket) \mid M \in \mathcal{M}\} \\ & \cup \Delta_{\mathcal{L}} \cup \{W_{L_M} \in \Delta_{\wp(\mathcal{L})} \mid M \in \mathcal{M}, |L_M| > 1\} \end{aligned}$$

To simplify notation, we introduce the infix metaoperator \gg such that $x \gg f = f x$. It is right associative and has the lowest precedence. Accordingly, $x \gg f \gg g = g (f x)$.

Although types in $\Delta_{\mathcal{L}}^{\mathcal{C}, \mathcal{M}}$ do not contain record-merge, we show by following Theorem 25 that types of mixin compositions in $\text{BCL}(\rightarrow, \cap)$ are sound.

► **Theorem 25** (Soundness). *Let $M_1, \dots, M_n \in \mathcal{M}$ be mixins, let $L_1, \dots, L_n \subseteq \mathcal{L}$ be sets of labels, let $C \in \mathcal{C}$ be a class and let $\sigma \in \mathbb{T}, \rho \in \mathbb{T}_{\langle \rangle}$ be types such that $\llbracket \sigma \rightarrow \rho \rrbracket$ is defined.*

If $\Delta_{\mathcal{L}}^{\mathcal{C}, \mathcal{M}} \vdash_{BCL} C \gg (M_1 W_{L_1}) \gg (M_2 W_{L_2}) \gg \dots \gg (M_n W_{L_n}) : \llbracket \sigma \rightarrow \rho \rrbracket$, then $\emptyset \vdash_{\langle \rangle} C \gg M_1 \gg M_2 \gg \dots \gg M_n : \sigma \rightarrow \rho$.

Proof Sketch. Induction on n proving $L_i = L_{M_i}$ followed by Lemma 23 and $(\rightarrow E)$. ◀

Complementary, we show by the following Theorem 26 that typing of mixin compositions in $BCL(\rightarrow, \cap)$ is complete with respect to previously described typing in \mathbb{T} .

► **Theorem 26** (Partial Completeness). *Let $\Gamma \subseteq \{x_C : \sigma_C \rightarrow \rho_C \mid C \in \mathcal{C}\} \cup \{x_M^\rho : (\sigma_M \rightarrow \rho \cap \rho_M^1) \rightarrow (\sigma_M \rightarrow \rho + \rho_M^2) \mid M \in \mathcal{M}, \rho \in \mathbb{T}_{\langle \rangle}, \llbracket \rho \rrbracket \text{ is defined}\}$ be a finite context and let $\sigma \in \mathbb{T}, \rho \in \mathbb{T}_{\langle \rangle}$ be types such that $\llbracket \sigma \rightarrow \rho \rrbracket$ is defined.*

If $\Gamma \vdash_{\langle \rangle} x_C \gg x_{M_1}^{\rho_1} \gg x_{M_2}^{\rho_2} \gg \dots \gg x_{M_n}^{\rho_n} : \sigma \rightarrow \rho$,

then $\Delta_{\mathcal{L}}^{\mathcal{C}, \mathcal{M}} \vdash_{BCL} C \gg (M_1 W_{L_{M_1}}) \gg (M_2 W_{L_{M_2}}) \gg \dots \gg (M_n W_{L_{M_n}}) : \llbracket \sigma \rightarrow \rho \rrbracket$.

Proof Sketch. Induction on n choosing for each x_M^ρ where $\rho = \bigcap_{l \in L} \langle l : \tau_l \rangle$ (resp. $\rho = \langle \rangle$) the substitution $S_i(\alpha) = \bigcap_{l \in L \setminus L_M} \langle l(\tau_l) \rangle$ to type $M \in \Delta_{\mathcal{L}}^{\mathcal{C}, \mathcal{M}}$ and using $(\rightarrow E)$ and Lemma 24. ◀

Coming back to our running example, we obtain

$$\begin{aligned} \Delta_{\{get, set, shift, move\}}^{\{Point\}, \{Movable\}} = \{ & \text{Point} : \quad \text{Int} \rightarrow \langle \langle get(\text{Int}) \cap set(\text{Int} \rightarrow \text{Int}) \cap shift(\text{Int} \rightarrow \text{Int}) \rangle \rangle, \\ & \text{Movable} : \quad w_{\{set, move\}}(\alpha) \\ & \quad \rightarrow (\text{Int} \rightarrow \alpha \cap \langle \langle set(\text{Int} \rightarrow \text{Int}) \cap shift(\text{Int} \rightarrow \text{Int}) \rangle \rangle) \\ & \quad \rightarrow (\text{Int} \rightarrow \alpha \cap \langle \langle set(\text{Int}) \cap move(\text{Int}) \rangle \rangle), \\ & W_{\{get\}} : \quad w_{\{get\}}(\langle \langle set(\alpha_1) \cap shift(\alpha_2) \cap move(\alpha_3) \rangle \rangle), \\ & W_{\{set\}} : \quad w_{\{set\}}(\langle \langle get(\alpha_1) \cap shift(\alpha_2) \cap move(\alpha_3) \rangle \rangle), \\ & W_{\{shift\}} : \quad w_{\{shift\}}(\langle \langle get(\alpha_1) \cap set(\alpha_2) \cap move(\alpha_3) \rangle \rangle), \\ & W_{\{move\}} : \quad w_{\{move\}}(\langle \langle get(\alpha_1) \cap set(\alpha_2) \cap shift(\alpha_3) \rangle \rangle), \\ & W_{\{set, move\}} : \quad w_{\{set\}}(\alpha) \rightarrow w_{\{move\}}(\alpha) \rightarrow w_{\{set, move\}}(\alpha) \} \end{aligned}$$

We may ask inhabitation questions such as

$$\Delta_{\{get, set, shift, move\}}^{\{Point\}, \{Movable\}} \vdash_{BCL} ? : \llbracket \text{Int} \rightarrow \langle shift : \text{Int} \rightarrow \text{Int}, move : \text{Int} \rangle \rrbracket$$

and obtain the combinatory term “Movable $W_{\{set, move\}}$ Point” as a synthesized result. From Theorem 25 we know

$$\emptyset \vdash_{\langle \rangle} \text{Movable Point} : \text{Int} \rightarrow \langle shift : \text{Int} \rightarrow \text{Int}, move : \text{Int} \rangle$$

On the other hand, if we want to inhabit $\llbracket \text{Int} \rightarrow \langle set : \text{Int} \rightarrow \text{Int}, move : \text{Int} \rangle \rrbracket$ we obtain no results. From Lemma 26 we know that, restricted to the previously described typing in \mathbb{T} , there is no mixin composition applied to a class with the resulting type $\text{Int} \rightarrow \langle set : \text{Int} \rightarrow \text{Int}, move : \text{Int} \rangle$.

The presented encoding has several benefits with respect to scalability. First, the size of the presented repositories is polynomial in $|\mathcal{L}| * |\mathcal{C}| * |\mathcal{M}|$. Second, expanding the label set \mathcal{L} requires only to update combinators in $\Delta_{\mathcal{L}}$ leaving existing types of classes and mixins untouched. Third, adding a class/mixin to an existing repository is as simple as adding one typed combinator for the class/mixin and at most one logical combinator. Again, it is important that the existing combinators in the repository remain untouched. As an example, we add the following mixin MovableBy to $\Delta_{\{get, set, shift, move\}}^{\{Point\}, \{Movable\}}$.

$$\begin{aligned} \text{MovableBy} &= \lambda \text{argClass. } \mathbf{Y}(\lambda \text{myClass. } \lambda \text{state.} \\ &\quad \text{let super} = \text{argClass state in} \\ &\quad \text{let self} = \text{myClass state in super} \oplus \langle \text{move} = \text{super.shift} \rangle \end{aligned}$$

In Λ_R for all types $\rho \in \mathbb{T}_{\langle \rangle}$ we have

$$\emptyset \vdash_{\langle \rangle} \text{MovableBy} : (\text{Int} \rightarrow \rho \cap \langle \text{shift} : \text{Int} \rightarrow \text{Int} \rangle) \rightarrow (\text{Int} \rightarrow \rho + \langle \text{move} : \text{Int} \rightarrow \text{Int} \rangle)$$

We obtain the following extended repository

$$\begin{aligned} \Delta_{\{\text{Point}\}, \{\text{Movable}, \text{MovableBy}\}} &= \Delta_{\{\text{Point}\}, \{\text{Movable}\}} \cup \{\text{MovableBy} : w_{\{\text{move}\}}(\alpha) \\ &\rightarrow (\text{Int} \rightarrow \alpha \cap \langle \langle \text{shift}(\text{Int} \rightarrow \text{Int}) \rangle \rangle) \rightarrow (\text{Int} \rightarrow \alpha \cap \langle \langle \text{move}(\text{Int} \rightarrow \text{Int}) \rangle \rangle)\} \end{aligned}$$

Asking the inhabitation question

$$\Delta_{\{\text{Point}\}, \{\text{Movable}, \text{MovableBy}\}} \vdash_{\text{BCL}^?} : \llbracket \text{Int} \rightarrow \langle \text{set} : \text{Int}, \text{move} : \text{Int} \rightarrow \text{Int} \rangle \rrbracket$$

synthesizes “Point \gg (Movable $W_{\{\text{set}, \text{move}\}}$) \gg (MovableBy $W_{\{\text{move}\}}$)”. Note that even in such a simplistic scenario the order in which mixins are applied can be crucial mainly because \oplus is not commutative. Moreover, the early binding of self and the associated preservation of overwritten methods may make multiple applications of a single mixin meaningful.

5 Conclusion and Future Work

We presented a theory for automatic compositional construction of object oriented classes by combinatory synthesis. This theory is based on the λ -calculus with records and \oplus typed by intersection types with records and $+$. It is capable of modeling classes as states to records (i.e. objects), and mixins as functions from classes to classes. Mixins can be assigned meaningful types using $+$ expressing their compositional character. However, non-monotonic properties of $+$ are incompatible with the existing well-studied theory of $\text{BCL}(\rightarrow, \cap)$ synthesis. Therefore, we designed a translation to repositories of combinators typed in $\text{BCL}(\mathbb{T}_C)$. We have proven this translation to be sound (Theorem 25) and partially complete (Theorem 26). A notable feature is the encoding of negative information (the absence of labels). It exploits the logic programming capabilities of inhabitation, by adding sets of combinators serving as witnesses for the non-presence of labels. In section 4 we also showed that this encoding scales wrt. extension of repositories.

Future work includes further studies on the possibilities to encode predicates exploiting patterns similar to the negative information encoding. The partial completeness result indicates a more expressive power of type constructors compared to records. Another direction of future work is to extend types of mixins and classes by semantic as well as modal types [19], a development initiated in [5]. In particular, the expressiveness of semantic types can be used to assign meaning to multiple applications of a single mixin and allow to reason about object oriented code on a higher abstraction level as well as higher semantic accuracy.

Acknowledgments. The authors would like to thank the anonymous reviewers for their valuable comments.

References

- 1 Martín Abadi and Luca Cardelli. *A Theory of Objects*. Springer, 1996.
- 2 Steffen Van Bakel. Strict intersection types for the lambda calculus. *ACM Comput. Surv.*, 43(3):20:1–20:49, April 2011.
- 3 H. Barendregt. *The Lambda Calculus – Its Syntax and Semantics*, volume 103 of *Studies in Logic and the Foundations of Mathematics*. North-Holland, 1984.
- 4 H. Barendregt, M. Coppo, and M. Dezani-Ciancaglini. A Filter Lambda Model and the Completeness of Type Assignment. *Journal of Symbolic Logic*, 48(4):931–940, 1983.
- 5 Jan Bessai, Boris Döder, Andrej Dudenhefer, and Moritz Martens. Delegation-based mixin composition synthesis. <http://www-seal.cs.tu-dortmund.de/seal/downloads/papers/paper-ITRS2014.pdf>, 2014.
- 6 Jan Bessai, Boris Döder, Andrej Dudenhefer, Tzu-Chun Chen, and Ugo de'Liguoro. Typing classes and mixins with intersection types. *arXiv preprint arXiv:1503.04911*, 2015.
- 7 Jan Bessai, Andrej Dudenhefer, Boris Döder, Moritz Martens, and Jakob Rehof. Combinatory Logic Synthesizer. In Tiziana Margaria and Bernhard Steffen, editors, *ISoLA'14*, volume 8802, pages 26–40, 2014.
- 8 Viviana Bono, Amit Patel, and Vitaly Shmatikov. A core calculus of classes and mixins. In *ECOOP*, volume 1628 of *Lecture Notes in Computer Science*, pages 43–66, 1999.
- 9 Gilad Bracha. *The Programming Language JIGSAW: Mixins, Modularity and Multiple Inheritance*. PhD thesis, Univeristy of Utha, 1992.
- 10 Gilad Bracha and William R. Cook. Mixin-based inheritance. In *OOPSLA/ECOOP*, pages 303–311, 1990.
- 11 Kim B. Bruce. *Foundations of Object-Oriented Languages – Types and Semantics*. MIT Press, 2002.
- 12 Peter S. Canning, William R. Cook, Walter L. Hill, Walter G. Olthoff, and John C. Mitchell. F-bounded polymorphism for object-oriented programming. In *FPCA*, pages 273–280, 1989.
- 13 Luca Cardelli. A semantics of multiple inheritance. In *Semantics of Data Types*, volume 173, pages 51–67, 1984.
- 14 Adriana B. Compagnoni and Benjamin C. Pierce. Higher-order intersection types and multiple inheritance. *Mathematical Structures in Computer Science*, 6(5):469–501, 1996.
- 15 William R. Cook, Walter L. Hill, and Peter S. Canning. Inheritance is not subtyping. In *POPL'90*, pages 125–135. ACM Press, 1990.
- 16 Ugo de'Liguoro. Characterizing convergent terms in object calculi via intersection types. In *TLCA*, pages 315–328, 2001.
- 17 Ugo de'Liguoro and Tzu chun Chen. Semantic Types for Classes and Mixins. <http://www.di.unito.it/~deligu/papers/UdLTC14.pdf>, 2014.
- 18 Boris Döder. *Automatic Synthesis of Component & Connector-Software Architectures with Bounded Combinatory Logic*. Dissertation, TU Dortmund, 2014.
- 19 Boris Döder, Moritz Martens, and Jakob Rehof. Staged composition synthesis. In *ESOP*, volume 8410 of *Lecture Notes in Computer Science*, pages 67–86, 2014.
- 20 Boris Döder, Moritz Martens, Jakob Rehof, and Paweł Urzyczyn. Bounded Combinatory Logic. In *Proceedings of CSL'12*, volume 16, pages 243–258. Schloss Dagstuhl, 2012.
- 21 Standard Ecma. ECMA-262 ECMAScript Language Specification, 2011.
- 22 Tihomir Gvero, Viktor Kuncak, Ivan Kuraj, and Ruzica Piskac. Complete Completion using Types and Weights. *SIGPLAN Notices*, 48(6):27–38, 2013.
- 23 Oleg Kiselyov and Ralf Lämmel. Haskell's overlooked object system. *CoRR*, abs/cs/0509027, 2005.
- 24 Martin Odersky and Matthias Zenger. Scalable component abstractions. In *OOPSLA*, pages 41–57. ACM, 2005.

- 25 Benjamin C. Pierce. Intersection types and bounded polymorphism. *Mathematical Structures in Computer Science*, 7(2):129–193, 1997.
- 26 Didier Rémy. Typing record concatenation for free. In *POPL'92*, pages 166–176, 1992.
- 27 Mitchell Wand. Type inference for record concatenation and multiple inheritance. *Inf. Comput.*, 93(1):1–15, 1991.

Non-Constructivity in Kan Simplicial Sets

Marc Bezem¹, Thierry Coquand², and Erik Parmann³

- 1 Department of Informatics, University of Bergen, Norway, bezem@ii.uib.no
- 2 Department of Computer Science and Engineering, Chalmers/University of Gothenburg, Sweden, thierry.coquand@cse.gu.se
- 3 Department of Informatics, University of Bergen, Norway, Erik.Parmann@ii.uib.no

Abstract

We give an analysis of the non-constructivity of the following basic result: if X and Y are simplicial sets and Y has the Kan extension property, then Y^X also has the Kan extension property. By means of Kripke countermodels we show that even simple consequences of this basic result, such as edge reversal and edge composition, are not constructively provable. We also show that our unprovability argument will have to be refined if one strengthens the usual formulation of the Kan extension property to one with explicit horn-filler operations.

1998 ACM Subject Classification F.4.1 Mathematical Logic

Keywords and phrases Constructive logic, simplicial sets, semantics of simple types

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.92

1 Introduction

Brouwer's Programme is the constructive reformulation of (as much as possible of) classical mathematics. In [2] it has been shown that the following theorem, though classically true (cf. [10, Corollary 7.11]), cannot be proved constructively.

► **Theorem 1** (classical). *The fibers of 0 and 1 of a Kan fibration $p : E \rightarrow \Delta^1$ are homotopy equivalent.*

In this paper we show that the following basic theorems cannot be proved constructively.

► **Theorem 2** (classical). *If X and Y are Kan simplicial sets, then any edge in Y^X can be reversed.*

► **Theorem 3** (classical). *If X and Y are Kan simplicial sets, then compatible edges in Y^X can be composed.*

The above two theorems follow immediately and constructively from the following.

► **Theorem 4** (classical). *If X and Y are Kan simplicial sets, then also Y^X is so.*

Hence we obtain that also Theorem 4, though classically true even without requiring that X is Kan (cf. [10, Theorem 6.9]), cannot be proved constructively.

The importance of these results is twofold. First, it is of evident importance for Brouwer's Programme to understand which results of classical mathematics already are constructive and which results are not. Second, Theorem 4 plays a crucial role in the construction of models of type theory with the Univalence Axiom, see [7]. The use of classical logic in proving this crucial property implies in particular that the model construction cannot be used to give a computational interpretation of univalence. Actually, Theorem 4 is a necessary step in the



© Marc Bezem, Thierry Coquand, and Erik Parmann;
licensed under Creative Commons License CC-BY

13th International Conference on Typed Lambda Calculi and Applications (TLCA'15).

Editor: Thorsten Altenkirch; pp. 92–106



Leibniz International Proceedings in Informatics

LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

semantics of the simply typed λ -calculus based on Kan simplicial sets. In what follows we expand on these points; for more motivation we refer to [2].

We would like to use the occasion to say a few not-too-technical words on the role of the Kan extension property of simplicial sets in relation to univalence. Let MLTT be Martin-Löf type theory with universe U and inductive equality $=_U$ on U . Assume we have two distinct copies of the natural numbers, inductively defined by constructors $0 : N$ ($0' : N'$) and $S : N \rightarrow N$ ($S' : N' \rightarrow N'$). MLTT proves $N =_U N$ and $N' =_U N'$, but not $N =_U N'$. The Univalence Axiom (UA) implies that (homotopy) equivalent types are equal, and in particular $N =_U N'$. By the Leibniz property of inductive equality, this implies that N and N' have the same properties and that all structure on N can be transported to N' and vice versa. This holds even uniformly, for example, $\Pi P : U \rightarrow U. (PN \rightarrow PN')$ is inhabited under UA. On the other hand, without UA, $\Pi P : U \rightarrow U. (PN \rightarrow PN')$ is not inhabited in MLTT. (One reason is that MLTT has models in which $N \neq N'$, so one can take $P \equiv \lambda X : U. (N =_U X)$ and get PN but not PN').

The above observation concerns not only the rather artificial type N' but also any other type that is equivalent to N , such as the type of lists over a unit type with one object. In fact the observation concerns all equivalent types. A less artificial example is perhaps the equivalence of the unit type to $\Sigma x : A. (a =_A x)$ for given $a : A : U$. The upshot is that validating UA requires an interpretation of $=_U$ that carries much more information than in MLTT without UA, since the elimination rule for $=_U$ (roughly, the Leibniz property, or substitutivity of equals for equals) has to be much stronger. In our simple example, the interpretation of $=_U$ must be leveraged to give an inhabitant of $\Pi P : U \rightarrow U. (PN \rightarrow PN')$.

Simplicial sets can be used to build a presheaf-style [6] model of MLTT. In this model the interpretation of $=_U$ does not validate UA. It turns out that if one builds a model of MLTT based on Kan simplicial sets, then it is possible to validate UA. The crucial notion here is that of a Kan fibration. A Kan *fibration* $p : E \rightarrow B$ is a map of simplicial sets with a specific lifting property. This lifting property *lifts* a path from b_0 to b_1 in B to a transport function from the fiber $p^{-1}(b_0)$ to the fiber $p^{-1}(b_1)$. In the model based on Kan simplicial sets, an inhabitant of $N =_U N'$ is interpreted as a path from N to N' in U . (Here and below we omit the correct but tedious phrase *the interpretation of N, N', U, \dots*). Any $P : U \rightarrow U$ is interpreted as a Kan fibration with fibers PT for any $T : U$. (NB the fibration, being a projection on the base type, has a direction opposite to the arrow in $P : U \rightarrow U$). Then the transport function obtained from the lifting property is the desired function $PN \rightarrow PN'$. In short, one can say that the transport functions interpret substitutivity of equals by equals.

Finally, to come back to the topic of this paper: if all types are to have Kan structure, one has to prove this inductively following the rules of type formation. One of the induction steps is Theorem 4. The unprovability of Theorem 4 shows that, from the constructive point of view, there is a problem with using the exponent Y^X in the category of Kan simplicial sets to interpret function types $X \rightarrow Y$.

The type theoretic (synthetic) formulation of homotopy equivalence and the Univalence Axiom, as well as the model of MLTT plus UA using Kan simplicial sets are all due to Voevodsky [14, 7]. This model confirms the homotopical interpretation proposed by Awodey and Warren [1].

Theorem 4 (without requiring that X is Kan) has an interesting history. The first appearance seems to be [12, Appendix A, p. 1A-8, Theorem 3]. Moore credits A. Heller for the definition of the function space Y^X on page 1A-4. Moore's proof is combinatorial, using the excluded middle in distinguishing the cases *a non/degenerate* on page 1A-9, l. 17ff. (Typo: on page 1A-7, l. 12 and 15, the map F is missing on the rhs; evidently $F_{(\mu, \nu)}$ was

intended to depend on F .) The proof in [10, Theorem 6.9] is much the same as the one by Moore (with the F 's in place). Several variations of this argument can be found in the literature.

An essentially more abstract proof using anodyne extensions is given by Gabriel and Zisman in [4, Chapter Four, 3.1.2] (take $B = \Delta^0$). Here the classical reasoning shows up when in 2.1.2 amalgamated sums over sets of non-degenerate simplices are taken.

The results of Moore and Heller imply that Kan simplicial sets form a cartesian closed category, which can be seen as a germ of the fact that they model dependent type theory.

The rest of the paper is structured as follows. In Section 2 we give an introduction to simplicial sets, and in Section 3 we provide several examples of simplicial sets which will be in use in the rest of the article. In Section 4 we take a closer look at Theorem 2, and provide a Kripke model showing that a constructive consequence, Lemma 14 cannot be proven constructively. Section 5 deals with edge composition, much in the same way as Section 4 deals with edge reversal. A summary and evaluation of the results obtained so far is given in Section 6. In Section 7 we strengthen the Kan condition and prove constructively a weak version of Lemma 14. This shows that our unprovability argument will have to be refined for the stronger Kan condition. We sum up our findings and discuss further research in Section 8.

2 Preliminaries

► **Definition 5** (Simplicial set). A *simplicial set* A is a collection of sets $A[i]$ for $i \in \mathbb{N}$ such that for every $0 < n$ and $j \leq n$ we have a function (*face map*) $d_j^n : A[n] \rightarrow A[n-1]$, and for every $0 \leq n$ and $j \leq n$ we have a function (*degeneracy map*) $s_j^n : A[n] \rightarrow A[n+1]$, satisfying the following *simplicial identities* for all suitable superscripts, which we happily omit:

$$d_i d_j = d_{j-1} d_i \quad \text{if } i < j \quad (1)$$

$$d_i s_j = s_{j-1} d_i \quad \text{if } i < j \quad (2)$$

$$d_i s_j = id \quad \text{for } i = j, j+1 \quad (3)$$

$$d_i s_j = s_j d_{i-1} \quad \text{if } i > j+1 \quad (4)$$

$$s_i s_j = s_j s_{i-1} \quad \text{if } i > j \quad (5)$$

An element of $A[i]$ is called an *i -simplex*, or just simplex when we don't wish to stipulate the dimension. A *degenerate* element is any element $a \in A[i+1]$ in the image of a degeneracy map.

Note that a simplicial identity like, e.g., $d_i^n d_j^{n+1} = d_{j-1}^n d_i^{n+1}$ actually means

$$\forall x \in A[n+1]. d_i^n(d_j^{n+1}(x)) = d_{j-1}^n(d_i^{n+1}(x)).$$

With a countably infinite signature, the above definition can be expressed completely in many-sorted first-order logic. That means that we can see first-order models which satisfy the above requirement as simplicial sets, and instead of simplicial sets we could talk about first-order models satisfying the above requirements.

Simplicial sets form a category. For two simplicial sets A and B , $Hom_S(A, B)$ is the set of all natural transformations from A to B . A natural transformation is a collection of maps $g[n] : A[n] \rightarrow B[n]$ commuting with the face and degeneracy maps of A and B : $g[n]s_i = s_i g[n-1]$ for all $0 \leq i < n$ and $g[n+1]d_i = d_i g[n]$ for all $0 \leq i \leq n+1$. We freely omit the dimension $[n]$ when it can be inferred from the other arguments. For more information on simplicial sets we refer to, for example, [10, 5, 3].

► **Definition 6** (Kan simplicial set). A simplicial set Y satisfies the Kan condition if for any collection of simplices $y_0, \dots, y_{k-1}, y_{k+1}, \dots, y_n$ in $Y[n-1]$ such that $d_i y_j = d_{j-1} y_i$ for any $i < j$ with $i \neq k$ and $j \neq k$, there is an n -simplex y in Y such that $d_i y = y_i$ for all $i \neq k$. The Kan condition is also called the Kan extension property, and a simplicial set is called a *Kan simplicial set* if it satisfies the Kan condition.

► **Definition 7** (Kan graph). A *reflexive multigraph* consists of C_1, C_0, d_0, d_1, s where C_0 is a set of points, C_1 a set of edges, $d_i : C_1 \rightarrow C_0$, d_1 the *source* and d_0 the *target* function, and $s : C_0 \rightarrow C_1$ the function mapping each $c \in C_0$ to a selfloop of c . We write $e : a \rightarrow b$ if e is in C_1 such that $d_1(e) = a$ and $d_0(e) = b$ (note the direction!). In particular we have $d_i(s(c)) = c$ for all $c \in C_0$. A *Kan graph* is a reflexive multigraph having the property that for all a, b, c in C_0 , if $e : a \rightarrow b$ and $f : a \rightarrow c$, then there exists an edge $g : b \rightarrow c$ in C_1 .

Kan graphs can be viewed as truncated Kan simplicial sets, modelling a truncated proof-relevant equality relation. Note that we don't require the Kan graph to have explicit functions giving the required edges like in [2], we merely require that the edges exists. We discuss this distinction further in Section 7. The special requirement of the edges for the Kan graph is in the literature often called *Euclidean*. Euclidean combined with reflexivity gives both transitivity and symmetry.

3 Examples of simplicial sets

We give some examples of simplicial sets that are used in the sequel.

3.1 Standard simplicial k -simplex Δ^k

Δ^k is the simplicial set with $\Delta^k[j]$ consisting of all non-decreasing sequences of numbers $0, \dots, k$ of length $j+1$. Equivalently, $\Delta^k[j]$ is the set of order-preserving functions $[j] \rightarrow [k]$, where $[i]$ denotes $0, \dots, i$ with the natural ordering. Examples are $\Delta^1[0] = \{0, 1\}$, $\Delta^1[1] = \{00, 01, 11\}$, $\Delta^2[1] = \{00, 01, 02, 11, 12, 22\}$ and

$$\Delta^2[2] = \{000, 001, 002, 011, 012, 022, 111, 112, 122, 222\}.$$

The degeneracy map $s_k^j : \Delta^i[j] \rightarrow \Delta^i[j+1]$ duplicates the k -th element in its input. So, $s_k^j(x_0 \dots x_k \dots x_{j+1}) = x_0 \dots x_k x_k \dots x_{j+1}$. The face map $d_k^j : \Delta^i[j] \rightarrow \Delta^i[j-1]$ deletes the k -th element. So, $d_k^j(x_0 \dots x_j) = x_0 \dots x_{k-1} x_{k+1} \dots x_j$.

3.2 The k -horns Λ_j^k

Λ_j^k is the j 'th horn of the standard k -simplex Δ^k , and defined by $\Lambda_j^k[n] = \{f \in \Delta^k[n] \mid [k] - \{j\} \not\subseteq \text{Im}(f)\}$. Alternatively, it is $\Delta^k[n]$ except every element must avoid some element not equal to j . For example, $\Lambda_0^2[1] = \{00, 01, 02, 11, \cancel{12}, 22\} = \Delta^2[1] - \{12\}$ (excluding 12, since 12 does not avoid any element not equal to 0). We also have:

$$\Lambda_0^2[2] = \{000, 001, 002, 011, \cancel{012}, 022, 111, \cancel{112}, \cancel{122}, 222\}.$$

The Kan extension condition for a simplicial set Y can also be formulated as: every map $F : \Lambda_j^k \rightarrow Y$ can be extended to a map $F' : \Delta^k \rightarrow Y$. This is equivalent to Definition 6.

3.3 Cartesian products

For two simplicial sets A and B , $A \times B$ is the simplicial set given by $(A \times B)[i] = A[i] \times B[i]$, and the structural maps d and s use d^A and d^B component-wise (and likewise for s^A and s^B). So if $a \in A[i]$ and $b \in B[i]$ then $(a, b) \in (A \times B)[i]$, and $d_i((a, b)) = (d_i^A(a), d_i^B(b))$. In particular, the degenerate simplices of $A \times B$ are pairs $(s_j^A(a), s_j^B(b)) \in (A \times B)[i + 1]$. (Caveat: this is stronger than both components being degenerate.)

3.4 Function spaces

We give the standard definition [12, p. 1A-4]: Y^X is the simplicial set given by $Y^X[i] = \text{Hom}_S(\Delta^i \times X, Y)$, where Hom_S denotes morphisms (natural transformations) of simplicial sets, and structural maps as follows. The face maps $d_k[i] : Y^X[i] \rightarrow Y^X[i - 1]$ need to map elements of $\text{Hom}_S(\Delta^i \times X, Y)$ to $\text{Hom}_S(\Delta^{i-1} \times X, Y)$ and the degeneracy maps vice versa. For their definition it is convenient to view a k -simplex in Δ^i as an order-preserving function $a : [k] \rightarrow [i]$. Let d_k^* be the strictly increasing function on natural numbers such that $d_k^*(n) = n$ if $n < k$ and $d_k^*(n) = n + 1$ otherwise (d_k^* ‘jumps’ over k). Given $F \in \text{Hom}_S(\Delta^i \times X, Y)$, define $(d_k F)[i](a, x) = F[i](d_k^* a, x)$. For the degeneracy maps, let s_k^* be the weakly increasing function on natural numbers such that $s_k^*(n) = n$ if $n \leq k$ and $s_k^*(n) = n - 1$ otherwise (s_k^* ‘duplicates’ k). Then define $(s_k F)[i](a, x) = F[i](s_k^* a, x)$.

3.5 The simplicial set defined by a reflexive multigraph

The following definition from [2] gives the general construction of a simplicial set from a reflexive multigraph. It is important to note that, even if the reflexive multigraph is transitive, its simplicial set is not the same as the nerve [5, Example 1.4] of the category defined by the multigraph. The difference is subtle: if we have edges $f : x \rightarrow y$, $g : y \rightarrow z$, $h, h' : x \rightarrow z$, where the composition $gf = h$, then the nerve does not contain the 2-simplex with f, g, h' , in contrast to below.

► **Definition 8.** Given a reflexive multigraph C we define the simplicial set $S(C)$ as follows. $S(C)[0] = C_0$, $S(C)[1] = C_1$ and $S(C)[n]$, for $n \geq 2$, consisting of all tuples of the form $(u_0, \dots, u_n; \dots, e_{ij}, \dots)$ such that

$$e_{ij} : u_i \rightarrow u_j \text{ in } C_1 \text{ for all } 0 \leq i < j \leq n.$$

The maps d_k in $S(C)$ are defined by removing from $(u_0, \dots, u_n; \dots, e_{ij}, \dots)$ the point u_k and all edges e_{ik} and e_{kj} . The maps s_k in $S(C)$ are defined by duplicating the point u_k in $(u_0, \dots, u_n; \dots, e_{ij}, \dots)$, adding an edge $e_{k(k+1)} = s(u_k)$, and duplicating edges and incrementing indices of edges as appropriate. This completes the construction of the simplicial set $S(C)$.

We now see why Kan graphs are named as they are: the S construction above turns them into Kan simplicial sets.


► **Lemma 9.** $S(Y)$ is a Kan simplicial set whenever Y is a Kan graph.

Proof. Consider Λ_k^n for some $n \geq 1$ and $0 \leq k \leq n$ and let $f : \Lambda_k^n \rightarrow S(Y)$. We have to define a lifting $h : \Delta^n \rightarrow S(Y)$. Δ^n consists of elements in every dimension, but we only need to specify h for every element in $\Delta^n[n]$. This since both the higher and lower dimensional objects are the (possibly repeated) s_i or d_j images of objects in $\Delta^n[n]$, and h must commute with both s_i and d_j , which determines h .

If $n = 1$, note that that Λ_k^1 only consists of one point, and degenerations of that point in the higher dimensions. E.g., if $k = 0$ then $\Lambda_0^1[0] = \{0\}$, $\Lambda_0^1[1] = \{00\}$ etc. In that case we extend f to $h : \Delta^1 \rightarrow S(Y)$ by mapping $h(1) = h(0) = f(0)$, which determines h in higher dimensions.

If $n = 2$ we use the fact that Y is a Kan graph, so for any two edges $f : a \rightarrow b$ and $g : a \rightarrow c$ there is an edge from b to c . The 2-horn gives two edges in the graph with at least one common point, and the fact that the graph is both reflexive, symmetric and transitive (because of the Kan property) enables us to find a third edge with compatible endpoints. The procedure depends on the value of k . We will here give the procedure for $k = 2$; $k = 0, 1$ are just simple adaptations.

Given $f : \Lambda_2^2 \rightarrow S(Y)$ we have edges $f(02) : f(0) \rightarrow f(2)$ and edges $f(12) : f(1) \rightarrow f(2)$, and we need to find a value for $h(01) : f(0) \rightarrow f(1)$ such that $d_1 h(01) = d_1 f(02)$ and $d_0 h(01) = d_0 f(12)$. In other words, we need to find that the dotted edge in the diagram actually exists (self-loops are not displayed).



Recall that $S(Y)[1] = Y[1]$, so both $f(01)$ and $f(12)$ are actual edges in Y . By applying the Kan property on $s(f(0))$ and $f(02)$ we get an edge $e_1 : f(2) \rightarrow f(0)$. Similarly we get an edge $e_2 : f(2) \rightarrow f(1)$. Now, by using the Kan property on e_1 and e_2 we get an edge $e_3 : f(0) \rightarrow f(1)$, and we put $h(01) = e_3$.

Finally, if $n \geq 3$ we observe that the horn Λ_k^n contains all points and edges of Δ^n , and we define the lifting by

$$h(q) = (f_{[0]}(q(0)), \dots, f_{[0]}(q(m)); \dots, f_{[1]}(e_{ij}), \dots).$$

Here $q : [m] \rightarrow [n]$ is order-preserving and e_{ij} is the edge from $q(i)$ to $q(j)$ in $\Delta^n[1] = \Lambda_k^n[1]$. ◀

4 Edge reversal

In this section we give the classical proof of Theorem 2 and show that there is no constructive proof.

4.1 Edge reversal, definition and classical proof

► **Definition 10** (Edge reversal). A simplicial set Y is said to have *edge reversal* when for every edge $e \in Y[1]$ there exists an edge $f \in Y[1]$ with $d_1(f) = d_0(e)$ and $d_0(f) = d_1(e)$.

► **Lemma 11.** *Kan simplicial sets have edge reversal.*

Proof. Given an arbitrary Kan simplicial set Y and an edge $e \in Y[1]$ we can make a map $G : \Lambda_0^2 \rightarrow Y$ by letting $G(0) = G(2) = d_1(e)$, $G(1) = d_0(e)$, $G(01) = e$ and $G(02) = s(d_1(e))$. Since Y is Kan we can extend G to $G : \Delta^2 \rightarrow Y$, giving us a value for $G(12) \in Y[1]$, which must be an edge between $G(1)$ and $G(2) = G(0)$, giving the reverse edge. ◀

We introduce some convenient ad-hoc terminology for later use.

► **Definition 12** (Y^X -good). Let X and Y be reflexive multigraphs and $F_{01} : X[1] \rightarrow Y[1]$. Define $F_0 = d_1 F_{01} s : X[0] \rightarrow Y[0]$ and $F_1 = d_0 F_{01} s : X[0] \rightarrow Y[0]$. We say that F_{01} is Y^X -good when the following two requirements hold for $i = 0, 1$:

- For all $e, e' \in X[1]$, if $d_i(e) = d_i(e')$ then $d_i F_{01}(e) = d_i F_{01}(e')$;
- For all $e \in X[1]$, $F_i d_0(e) = F_i d_1(e)$.

The first requirement expresses that F_{01}, F_0, F_1 respect endpoints, that is, if $e : a \rightarrow b$ in $X[1]$, then $F_{01}(e) : F_0(a) \rightarrow F_1(b)$ in $Y[1]$. The second requirement ensures that F_0 and F_1 are constant on each weakly connected component of X . (Notice that $F_{01} s(y)$ for $y \in Y[0]$ does not need to map to a degenerate edge, so F_0 and F_1 are not necessarily identical.)

► **Lemma 13.** *If X and Y are reflexive multigraphs and $F_{01} : X[1] \rightarrow Y[1]$ is Y^X -good, then we can extend F_{01} to a 1-simplex in $S(Y)^{S(X)}$.*

Proof. To be a map in $S(Y)^{S(X)}$ we need to extend $F_{01} : X[1] \rightarrow Y[1]$ to a family of maps $F'_{01}[n] : (\Delta^1 \times S(X))[n] \rightarrow S(Y)[n]$ which commute with d_i and s_j . Recall the definitions $F_0 = d_1 F_{01} s$ and $F_1 = d_0 F_{01} s$. We define $F'_{01}[n]$ depending on n . If $n = 0$ then the input will have the form (i, x) where $0 \leq i \leq 1$ and $x \in X[0]$, and we put $F'_{01}[0](i, x) = F_i(x)$. If $n = 1$ the input will have the form (ij, e) where $0 \leq i < j \leq 1$ and $e \in X[1]$. If $i = j$ we put $F'_{01}(ij, e) = s F_i(d_0(e))$. Note that since F_{01} is Y^X -good, we know that $F_i(d_0(e)) = F_i(d_1(e))$, justifying our choice of the degenerate edge as the output. If $i < j$ we let $F'_{01}(01, e) = F_{01}(e)$. If $n > 1$ any input to $F'_{01}[n]$ will have the form $(0^a 1^b, (x_0, \dots, x_n; \dots, e_{ij}, \dots))$ such that $a + b = n + 1$. We let $F'_{01}[n]$ map this element to the tuple

$$(F_0(x_0), \dots, F_0(x_{a-1}), F_1(x_a), \dots, F_1(x_{a+b-1}); \dots, e'_{ij}, \dots),$$

where $e'_{ij} = s(F_0(x_a))$ if $i < j < a$, $e'_{ij} = F_{01}(e_{ij})$ if $i < a \leq j$, and $e'_{ij} = s(F_1(x_a))$ if $a \leq i < j$. That is, the $F'_{01}[n]$ images are sequences of a number of F_0 images followed by b number of F_1 images, with all edges being degenerate, except the bridges between the two nodes. Since each of the derived F_i functions are constant on each connected component, and the input consists exactly of sequences of nodes in the same connected component, all of the elements $F_0(x_0), \dots, F_0(x_{a-1})$ are the same element in $Y[0]$, and likewise for $F_1(x_a), \dots, F_1(x_{a+b-1})$. This justifies our choice of e'_{ij} as the degenerate edges.

It should be clear that this map does indeed commute with d_i and s_j , completing the proof. ◀

► **Lemma 14** (classical). *For all Kan graphs Y and X , if $F_{01} : X[1] \rightarrow Y[1]$ is Y^X -good, then there is an $F_{10} : X[1] \rightarrow Y[1]$ such that $d_0 F_{01} = d_1 F_{10}$ and $d_1 F_{01} = d_0 F_{10}$.*

Proof. Let X and Y be Kan graphs. The $S(Y)$ and $S(X)$ are Kan simplicial sets by Lemma 9. By applying the classical Theorem 2 we get that $S(Y)^{S(X)}$ has edge reversal. Since F_{01} is Y^X -good we extend F_{01} to an edge $F'_{01} \in S(Y)^{S(X)}[1]$ as defined in the proof of Lemma 13. By edge reversal in $S(Y)^{S(X)}$ we get an $F'_{10} \in S(Y)^{S(X)}[1]$ satisfying $d_1(F'_{10}) = d_0(F'_{01})$ and $d_0(F'_{10}) = d_1(F'_{01})$. We put $F_{10}(x) = F'_{10}(01, x)$. By expanding the definition of d_k from Section 3.4, we get the following properties: $F'_{10}(00, e) = F'_{01}(11, e)$ and $F'_{10}(11, e) = F'_{01}(00, e)$, giving $F'_{10}(0, d_i(e)) = F'_{01}(1, d_i(e))$ and $F'_{10}(1, d_i(e)) = F'_{01}(0, d_i(e))$. We calculate $d_0 F_{01}(e) = d_0 F'_{01}(01, e) = F'_{01}(1, d_0(e)) = F_1 d_0(e)$. Since F_{01} is Y^X -good (2nd requirement) we have $F_1 d_0(e) = F_1 d_1(e)$. We continue the calculation: $F_1 d_1(e) = F'_{01}(1, d_1(e)) = F'_{10}(0, d_1(e))$ where the last step is justified above. We continue: $F'_{10}(0, d_1(e)) = d_1 F'_{10}(01, e) = d_1 F_{10}(e)$. In total we have proved $d_0 F_{01}(e) = d_1 F_{10}(e)$ for all $e \in X[1]$. Hence $d_0 F_{01} = d_1 F_{10}$. The other equation is proved symmetrically. ◀

■ **Table 1** Kripke (counter)model for edge reversal.

Day 1	
X_0	$\{x, x'\}$
X_1	$\{s(x), s(x')\}$
Y_0	$\{y_0, y_1, y'_0, y'_1\}$
Y_1	$\{s(y_0), s(y_1), s(y'_0), s(y'_1), y_0y_1 : y_0 \rightarrow y_1, y'_0y'_1 : y'_0 \rightarrow y'_1, a : y_1 \rightarrow y_0, b : y'_1 \rightarrow y'_0\}$

Day 2		Input	Output	
X_0	$\{x=x'\}$	F_0	x	y_0
X_1	$\{s(x)=s(x')\}$	F_0	x'	y'_0
Y_0	$\{y_0 = y'_0, y_1 = y'_1\}$	F_1	x	y_1
Y_1	$\{s(y_0) = s(y'_0), s(y_1) = s(y'_1), y_0y_1 = y'_0y'_1, a, b\}$	F_1	x'	y'_1
		F_{01}	$s(x)$	y_0y_1
		F_{01}	$s(x')$	$y'_0y'_1$

Kripke [8] showed that constructive logic is sound for Kripke models, so the existence of a Kripke countermodel of a statement gives the non-existence of a constructive proof of that statement. We will now, by the means of a Kripke model, see that Lemma 14 does not hold constructively.

4.2 Edge reversal, the Kripke countermodel

We describe a Kripke model containing a Y^X -good F_{01} such that there cannot be a function $F_{10} : X[1] \rightarrow Y[1]$ with $d_0F_{01} = d_1F_{10}$ and $d_1F_{01} = d_0F_{10}$, even though X and Y are Kan graphs.

For clarity, the functions $F_0 = d_1F_{01}s$ and $F_1 = d_0F_{01}s$ as defined in Definition 12 are also made explicit in this model. Face maps are part of the model, but not made explicit.

The model consists of two days, with an X and a Y part each. On day 1 both X and Y consist of two separate components, which get merged on day 2. We give the model both in Table 1 and, graphically, in Figure 1 and 2.


It is easy to see that both X and Y are Kan graphs by simply observing that each of their two components are strongly connected. It is also clear that we cannot define a consistent F_{10} . In day 1 we would have to set $F_{10}(s(x)) = a$ and $F_{10}(s(x')) = b$ to satisfy the requirement that $d_0F_{01} = d_1F_{10}$ and $d_1F_{01} = d_0F_{10}$. The problem occurs in day 2, where we have that $s(x) = s(x')$, but $a \neq b$, making it impossible for F_{10} to respect equality. Note that all other functions, F_0 , F_1 , F_{01} , s , d_0 , and d_1 remain consistent after collapsing, that is, they still map equal elements to equal elements.

5 Edge composition


In this section we give the classical proof of Theorem 3 and show that there is no constructive proof.

► **Definition 15** (Edge composition). A simplicial set Y is said to have *edge composition* when for every edge $e_1, e_2 \in Y[1]$, if $d_0(e_1) = d_1(e_2)$ then there exists an edge $f \in Y[1]$ with $d_1(f) = d_1(e_1)$ and $d_0(f) = d_0(e_2)$.

► **Lemma 16.** *Kan simplicial sets have edge composition.*



■ **Figure 1** Kripke (counter)model for edge reversal, day 1.




■ **Figure 2** Kripke (counter)model for edge reversal, day 2.

Proof. Given an arbitrary Kan simplicial set Y and edges $e_1, e_2 \in Y[1]$ with $d_0(e_1) = d_1(e_2)$, we can make a map $G : \Lambda_1^2 \rightarrow Y$ by putting $G(0) = d_1(e_1)$, $G(1) = d_0(e_1)$, $G(2) = d_0(e_2)$, $G(01) = e_1$ and $G(12) = e_2$. Since Y is Kan we can extend G to $G : \Delta^2 \rightarrow Y$, giving us a simplex $G(02) : G(0) \rightarrow G(2)$ in $Y[1]$, the composition of e_1 and e_2 . ◀


By a proof essentially identical to the proof of Lemma 14 we get the following lemma.

► **Lemma 17** (classical). *For all Kan graphs Y and X , if $F_{01} : X[1] \rightarrow Y[1]$ and $F_{12} : X[1] \rightarrow Y[1]$ are Y^X -good maps satisfying $d_0 F_{01} = d_1 F_{12}$, then there is an $F_{02} : X[1] \rightarrow Y[1]$ such that $d_0 F_{01} = d_0 F_{02}$ and $d_1 F_{12} = d_1 F_{02}$.*

In Figure 3 and 4 we see that Lemma 17 is not constructively provable. We have two Y^X -good functions F_{01} and F_{12} , satisfying the requirement, and both X and Y are Kan graphs. If $S(Y)^{S(X)}$ had edge composition we would get a function F_{02} that $d_1 F_{01} = d_1 F_{02}$ and $d_0 F_{12} = d_0 F_{02}$. However, such a function is not definable in the Kripke model. The reason is analogous to the case of edge-reversal: from day 1 to day 2 we have equated objects in the domain of F_{02} while keeping the images distinct. Specifically, on day 1 we are forced to set $F_{02}(s(x)) = a$ and $F_{02}(s(x')) = b$, but on day 2 we have $s(x) = s(x')$, but $a \neq b$.



■ **Figure 3** Kripke (counter)model for edge composition, day 1.




■ **Figure 4** Kripke (counter)model for edge composition, day 2.

6 Evaluation of the results

The results up to now are summarized in Figure 5.

Having concrete, finite Kripke countermodels against Lemma 14 and 17 allows for a further simplification: everything remains valid under the condition that X has at most two points. Likewise, explicit bounds read off from the Kripke models can be imposed on the number of points of Y and on the number of edges in X and in Y . The simplified results are denoted by postfixing the number of the result by a ‘b’ for bounded, so Lemma 14b is the bounded version of Lemma 14.



■ **Figure 5** Summary of results, all implications constructive.

With explicit bounds on the size of the domain, functions are completely determined by a finite number of function values. For example, if we have $\forall z \in X. (z = x \vee z = x')$ for $x, x' \in X$, then the binary predicate $\text{fun}(y, y') \equiv (x = x' \rightarrow y = y')$ on Y completely describes all functions $X \rightarrow Y$, in evidence $x \mapsto y, x' \mapsto y'$. With this in mind it is not difficult to express Lemma 14b as a first-order classical tautology Φ that is not true in all Kripke models.

Now fix a constructive framework that is sufficiently expressive for the results in Figure 5. For example, IZF (Zermelo-Fraenkel set theory in IPL, intuitionistic predicate logic) will do. Let $\llbracket \Phi \rrbracket$ be the Tarski interpretation of Φ expressed in IZF. The following fundamental property of IZF could be called the semantic conservativity of IZF over IPL:

If $\llbracket \Phi \rrbracket$ is provable in IZF, then Φ is true in all Kripke models.

Lubarsky [9] and McCarty [11] independently provided constructive proofs of the above conservativity property of IZF. We gratefully acknowledge their prompt answers to our question.¹

Empowered by the proofs of Lubarsky and McCarty we can now conclude that Lemma 14b cannot be proved in IZF. The same is true for Lemma 17b, and for all other results in Figure 5, as well as for their bounded versions.

7 Kan graphs with explicit filler functions

Let us first give an intuitive explanation of our countermodels. They actually exploit the undecidability of equality: on day 1 we don't know what will be equal on day 2. (This is different from the decidability of degeneracy, but the two are related: for example, an edge e is degenerate iff $e = s_0(d_1(e))$.) In Figure 1 and 2, the point is that $y_0 \neq y'_0$ on day 1, so one cannot put $F_{10}(s(x)) = F_{10}(s(x')) = a$ since this conflicts with $d_0 F_{10} = d_1 F_{01}$. One is thus forced to a choice that turns out to be wrong on day 2.

One attempt to deal with this lack of information is to give Kan simplicial sets more structure. One could for example change Definition 6 of a Kan simplicial set into one where we not only know that the required n -simplex exists, but actually have *functions* producing them. In the formulation using horns as in Section 3.2 this would amount to a dependent function $\text{fill}(k, j, F)$ such that $\text{fill}(k, j, F) : \Delta^k \rightarrow Y$ extends $F : \Lambda_j^k \rightarrow Y$, for any k, j, F . This form of Kan simplicial set has been introduced by Nikolaus in [13] under the name of *algebraic Kan complex*. The definition with explicit fill-functions has certain advantages, both classically and constructively, as we will see below. However, one should be careful in defining Y^X : morphisms in the category of algebraic Kan complexes are required to map chosen fillers in X to chosen fillers in Y . As a consequence, there are less maps from X to Y

¹ Strengthening the semantic conservativity to syntactic conservativity, that is, concluding that Φ is provable in intuitionistic predicate logic, by using the completeness of the Kripke semantics implicates some classical logic. Although not needed for this paper, we think there is some general interest in a constructive proof that $\text{IPL} \vdash \Psi$ whenever $\text{IZF} \vdash \llbracket \Psi \rrbracket$, for any first-order sentence Ψ .

as algebraic Kan complexes than as just simplicial sets. What we propose could be called a *functional* Kan simplicial set, with explicit fill-functions but with maps as for ordinary simplicial sets. As a consequence the exponential Y^X of simplicial sets can be used.

To be able to prove an analogue of Lemma 9 we have to strengthen the notion of Kan graph to also include such *filler functions*, cf. [2].

► **Definition 18** (Kan fill-graph). A *Kan fill-graph* is a reflexive multigraph with a partial function $\text{fill} : Y[1] \times Y[1] \rightarrow Y[1]$ such that for all $e_1, e_2 \in Y[1]$, if $e_1 : a \rightarrow b$ and $e_2 : a \rightarrow c$, then $\text{fill}(e_1, e_2) : b \rightarrow c$.

As noted earlier, the Kan property together with reflexivity implies symmetry and transitivity. We can now define the corresponding functions.

► **Definition 19** (Edge reversal). For all $e \in Y[1]$ where Y is a Kan fill-graph let

$$e^{-1} = \text{fill}(e, sd_1(e)).$$

If $e : a \rightarrow b$, then $sd_1(e) : a \rightarrow a$, and $\text{fill}(e, sd_1(e)) : b \rightarrow a$.

Note that we in general don't have $(e^{-1})^{-1} = e$, but we do have that $d_i((e^{-1})^{-1}) = d_i(e)$.

► **Definition 20** (Edge composition). Using the inverse for edges in Y we define the composition of two edges $e_1 : a \rightarrow b$ and $e_2 : b \rightarrow c$ as

$$\text{trans}(e_1, e_2) = \text{fill}(e_1^{-1}, e_2).$$

Again we are in no way guaranteed that $\text{trans}(e_1, s(b)) = e_1$ or $\text{trans}(s(x), s(x)) = s(x)$.

We immediately see that the addition of explicit functions adds power, as we can now prove constructively and trivially an analogue of Lemma 14.

► **Lemma 21.** For all Kan fill-graphs Y, X and for every $F : X[1] \rightarrow Y[1]$, the function $F^{-1} : X[1] \rightarrow Y[1]$ defined by $F^{-1}(e) = F(e)^{-1}$ satisfies $d_0 F = d_1 F^{-1}$ and $d_1 F = d_0 F^{-1}$.

Note how using explicit functions rules out the Kripke counter-example we gave of Lemma 14. If $s(x) = s(x')$ on day 2, then we immediately get $a = F_{01}^{-1}(s(x)) = F_{01}^{-1}(s(x')) = b$ since equality has to be preserved.


We can even use the above fact to show that:

► **Lemma 22.** For any reflexive multigraph X and Kan fill-graph Y , $S(Y)^{S(X)}$ has edge reversal.

Proof. Assume an edge $F \in S(Y)^{S(X)}[1]$, we proceed to define F^{-1} such that $d_0(F) = d_1(F^{-1})$ and $d_1(F) = d_0(F^{-1})$. As $F \in S(Y)^{S(X)}[1]$ we have $F[n] : \Delta^1[n] \times X[n] \rightarrow Y[n]$. We start with $n = 0$, defining $F^{-1}[0] : (\Delta^1 \times X)[0] \rightarrow Y[0]$ by letting $F^{-1}[0](0, x) = F[0](1, x)$ and $F^{-1}[0](1, x) = F[0](0, x)$. Likewise for $n = 1$ we define $F^{-1}(00, e) = F(11, e)$ and $F^{-1}(11, e) = F(00, e)$, these are directly enforced by $d_0(F) = d_1(F^{-1})$ and $d_1(F) = d_0(F^{-1})$. For the case of $F^{-1}(01, e)$ we need to find an edge $F^{-1}(01, e) : F^{-1}(0, d_1 e) \rightarrow F^{-1}(1, d_0 e)$, which from the way we defined $F^{-1}[0]$ is the same as an edge

$$F^{-1}(01, e) : F(1, d_1 e) \rightarrow F(0, d_0 e).$$

The diagram in Figure 6 shows $e \in S(X)[1]$ with its endpoints on the left, and the nodes and edges we have directly reachable in $S(Y)$ using only F on the right.



■ **Figure 6** Reversing F .

Reading off the figure we can define $F^{-1}(01, e)$ as follows:

$$F^{-1}(01, e) = \text{trans}(F(11, e), \text{fill}(F(01, e), F(00, e)))$$

Note that F^{-1} is well-defined since the functions involved in the definition are. Moreover, F^{-1} commutes with s_0, d_0, d_1 by construction.

Having defined F^{-1} for dimension 0 and 1, F^{-1} is also determined in higher dimensions, because of the truncation in $S(X), S(Y)$. In the case of $n > 1$ any input to $F^{-1}[n]$ will have the form

$$F^{-1}(0^a 1^b, (x_0, \dots, x_n; \dots e_{ij}, \dots))$$

where $a + b = n + 1$. We let $F^{-1}[n]$ map this element to the tuple

$$(F^{-1}(0, x_0), \dots, F^{-1}(0, x_{a-1}), F^{-1}(1, x_a), \dots, F^{-1}(1, x_{a+b-1}); \dots e'_{ij}, \dots),$$

where $e'_{ij} = F^{-1}(00, e_{ij})$ if $i < j < a$, $e'_{ij} = F^{-1}(01, e_{ij})$ if $i < a \leq j$, and $e'_{ij} = F^{-1}(11, e_{ij})$ if $a \leq i < j$. This commutes with face and degeneracy maps. ◀

Using the same techniques we can constructively prove the following variant of Lemma 17.

► **Lemma 23.** *For any Kan graph X and Kan fill-graph Y , if $F_{01} : X[1] \rightarrow Y[1]$ and $F_{12} : X[1] \rightarrow Y[1]$ satisfy $d_0 F_{01} = d_1 F_{12}$, then there is a $F_{02} : X[1] \rightarrow Y[1]$ such that $d_1 F_{01} = d_1 F_{02}$ and $d_0 F_{12} = d_0 F_{02}$.*

► **Lemma 24.** *For any reflexive multigraph X and Kan fill-graph Y , $S(Y)^{S(X)}$ has edge composition.*

Proof. Assume edges $F_{01} \in S(Y)^{S(X)}$, $F_{12} \in S(Y)^{S(X)}$ such that $d_0(F_{01}) = d_1(F_{12})$, and we proceed to define $F_{02} \in S(Y)^{S(X)}$ such that $d_1(F_{02}) = d_1(F_{01})$ and $d_0(F_{02}) = d_0(F_{12})$.


As was the case in the proof of Lemma 22, we are forced on $F_{02}(0, x) = F_{01}(0, x)$, $F_{02}(1, x) = F_{12}(1, x)$, $F_{02}(00, e) = F_{01}(00, e)$, and $F_{02}(11, e) = F_{12}(11, e)$.

For the case of $F_{02}(01, e)$ we need to find an edge $F_{02}(01, e) : F_{02}(0, d_1 e) \rightarrow F_{02}(1, d_0 e)$, which from the way we defined $F_{02}[0]$ is the same as an edge

$$F_{02}(01, e) : F_{01}(0, d_1 e) \rightarrow F_{12}(1, d_0 e).$$

We note that $d_0(F_{01}) = d_1(F_{12})$ enforces $F_{01}(11, e) = F_{12}(00, e)$, which again enforces $F_{01}(1, d_1(e)) = F_{12}(0, d_1(e))$. This gives the diagram in Figure 7, enabling us to read off:

$$F_{02}(01, e) = \text{fill}(\text{trans}(F_{01}(11, e), F_{01}(01, e)^{-1}), F_{12}(01, e)). \quad \blacktriangleleft$$



■ **Figure 7** Filling the horn Λ_1^2 .

8 Conclusions and Future Research

We have given a thorough analysis of the non-constructivity of the basic result that the Kan extension property is preserved under the usual operation of exponentiation of simplicial sets. An important step in this analysis, also employed in [2], is the truncation of simplicial sets to dimension 1. This allows us to study the basic result in the simplified situation of Kan graphs. Once one has shown the constructive unprovability of the basic result in the situation of Kan graphs, one obtains *a fortiori* its unprovability for Kan simplicial sets.

The much simpler notion of Kan graph (as compared to Kan simplicial set) invites to further thought experiments. One of those is the study of simple, constructive consequences of the Kan extension property, such as edge reversal and edge composition. It turns out that already these consequences cannot be proven constructively.

Another experiment is to strengthen the Kan extension property from existence of an n -simplex as in Definition 6 to having a function, called a *filler*, yielding these n -simplices. This makes quite a difference. None of the Kripke models we have introduced is able to deal with such fillers, since equating objects in X and Y implies that filler-values such as a and b in Figure 1 also have to be equal. The question arises whether this is necessary so, or just coincidental in the particular Kripke model. This question is answered in Section 7, where we prove constructively that, if X is a graph and Y a Kan-fill graph, then $S(Y)^{S(X)}$ has edge reversal and edge composition. This result may be of independent interest. It suggests that showing the (expected) constructive unprovability of Theorem 4 for *algebraic Kan complexes* as in [13] will require more complicated structures than graphs. The above expectation is based on an analysis of filling a 2-horn in Y^X , which requires defining $F(001, t)$. As F has to commute with s_0 , one must know whether the 2-simplex t is an s_0 -image or not. This can in general only be decided by an appeal to classical logic. We have to leave this to future research.

References

- 1 Steve Awodey and Michael Warren. Homotopy theoretic models of identity types. *Mathematical Proceedings of the Cambridge Philosophical Society*, 146:45–55, 2009.

- 2 Marc Bezem and Thierry Coquand. A Kripke model for simplicial sets. *TCS*, 2015. doi:10.1016/j.tcs.2015.01.035.
- 3 Greg Friedman. An elementary illustrated introduction to simplicial sets. Preprint, <http://arxiv.org/abs/0809.4221>, 2008.
- 4 Peter Gabriel and Michel Zisman. *Calculus of fractions and homotopy theory*. Springer, 1967.
- 5 Paul Goerss and J.F. Jardine. *Simplicial Homotopy Theory*. Modern Birkhäuser Classics. Birkhauser Verlag GmbH, 2009. Reprint of Vol. 174 of Progress in Mathematics, 1999.
- 6 Martin Hofmann. Syntax and semantics of dependent types. In A.M. Pitts and P. Dybjer, editors, *Semantics and logics of computation*, volume 14 of *Publ. Newton Inst.*, pages 79–130. Cambridge University Press, Cambridge, 1997.
- 7 Chris Kapulkin, Peter LeFanu Lumsdaine, and Vladimir Voevodsky. The simplicial model of univalent foundations. Preprint, <http://arxiv.org/abs/1211.2851>, 2012.
- 8 Saul Kripke. Semantical analysis of intuitionistic logic I. In M. Dummett and J.N. Crossley, editors, *Formal Systems and Recursive Functions*, pages 92–130. North-Holland, Amsterdam, 1965.
- 9 Bob Lubarsky, March 2015. Personal communication.
- 10 Jon Peter May. *Simplicial Objects in Algebraic Topology*. Chicago Lectures in Mathematics. University of Chicago Press, 2nd edition, 1993.
- 11 Charles McCarty. Two questions about IZF and intuitionistic validity. Pdf file, personal communication, March 2015.
- 12 John C. Moore. Algebraic homotopy theory. Lectures at Princeton, <http://faculty.tcu.edu/gfriedman/notes/aht1.pdf>, 1956.
- 13 Thomas Nikolaus. Algebraic models for higher categories. Preprint, <http://arxiv.org/abs/1003.1342>, 2010.
- 14 Vladimir Voevodsky. Notes on type systems. http://www.math.ias.edu/~vladimir/Site3/Univalent_Foundations_files/expressions_current.pdf, 2009.

Logical Relations for Coherence of Effect Subtyping

Dariusz Biernacki and Piotr Polesiuk

Institute of Computer Science, University of Wrocław
Joliot-Curie 15, 50-383 Wrocław, Poland
{dabi,ppolesiuk}@cs.uni.wroc.pl

Abstract

A coercion semantics of a programming language with subtyping is typically defined on typing derivations rather than on typing judgments. To avoid semantic ambiguity, such a semantics is expected to be coherent, i.e., independent of the typing derivation for a given typing judgment. In this article we present heterogeneous, biorthogonal, step-indexed logical relations for establishing the coherence of coercion semantics of programming languages with subtyping. To illustrate the effectiveness of the proof method, we develop a proof of coherence of a type-directed, selective CPS translation from a typed call-by-value lambda calculus with delimited continuations and control-effect subtyping. The article is accompanied by a Coq formalization that relies on a novel shallow embedding of a logic for reasoning about step-indexing.

1998 ACM Subject Classification D.3.3 Language Constructs and Features, F.3.3 Studies of Program Constructs

Keywords and phrases type system, coherence of subtyping, logical relation, control effect, continuation-passing style

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.107

1 Introduction

Programming languages that allow for subtyping, i.e., a mechanism facilitating coercions of expressions of one type to another, are usually given either a subset semantics, where one type is considered a subset of another type, or – a coercion semantics, where expressions are explicitly converted from one type to another. In the presence of subtyping, typing derivations depend on the occurrences of the subtyping judgments and, therefore, typing judgments do not have unique typing derivations. Consequently, a coercion semantics that interprets subtyping judgment by introducing explicit type coercions is defined on typing derivations rather than on typing judgments. But then a natural question arises as to whether such a semantics is coherent, i.e., whether it does not depend on the typing derivation.

The problem of coherence has been considered in a variety of typed lambda calculi. Reynolds proved the coherence of the denotational semantics for intersection types [22]. Breazu-Tannen et al. proved the coherence of a coercion translation from the lambda calculus with polymorphic, recursive and sum types to system F [8], by showing that any two derivations of the same judgment are normalizable to a unique normal derivation where the correctness of the normalization steps is justified by an equational theory in the target calculus. Curien and Ghelli introduced a translation from system F_{\leq} to a calculus with explicit coercions and showed that any two derivations of the same judgment are translated to terms that are normalizable to a unique normal form [9]. Finally, Schwinghammer followed Breazu-Tannen et al.'s approach to prove the coherence of coercion translation from Moggi's computational lambda calculus with subtyping [24].



© Dariusz Biernacki and Piotr Polesiuk;
licensed under Creative Commons License CC-BY

13th International Conference on Typed Lambda Calculi and Applications (TLCA'15).

Editor: Thorsten Altenkirch; pp. 107–122



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The normalization-based proofs consist in finding a normal form for a representation of the derivation and they hinge on showing that such normal forms are unique for a given typing judgment. When the source calculus under consideration is presented in the spirit of the lambda calculus à la Church, i.e., the lambda abstractions are type annotated, as is the case in all the aforementioned articles that follow the normalization-based approach, the term and the typing context indeed determine the shape of the normal derivation (modulo a top level coercion that depends on the type of the term) [18]. However, in calculi à la Curry this is no longer the case and the method cannot be directly applied. Still, if the calculus is at least weakly normalizing, one can hope to recover the uniqueness property for normal typing derivations for source terms in normal form, assuming that term normalization preserves the coercion semantics. For instance, in the simply typed λ -calculus the typing context uniquely determines the type of the term in the function position in applications building a β -normal form, and, hence, derivations in normal form for such terms are unique. This line of reasoning cannot be used when the calculus includes recursion.

In this article, we consider the coherence problem in simply-typed lambda calculi with general recursion, control effects and with no type annotations. The coercion semantics we study translate typing derivations in the source calculus to a corresponding target calculus with explicit type coercions (that in most cases can be further replaced with equivalent lambda-term representations) and our criterion for coherence of the translation is contextual equivalence [19] in the target calculus.

The main result of this work is a construction of logical relations for establishing so construed coherence of coercion semantics, applicable in a variety of calculi. In particular, we address the problem of coherence of a type-directed CPS translation from the call-by-value λ -calculus with delimited-control operators and control-effect subtyping introduced by Materzok and the first author [16], extended with recursion. While the translation for the calculus with explicit type annotations has been shown to be coherent in terms of an equational theory in a target calculus [15], no CPS coercion translation for the original version, let alone extended with recursion, has been proven coherent.

The reasons why coherence in this calculus is important are twofold. First of all, it is very expressive and therefore interesting from the theoretical point of view. In particular, the calculus has been shown to generalize the canonical type-and-effect system for Danvy and Filinski's shift and reset control operators [10, 11], and, furthermore, that it is strictly more expressive than the CPS hierarchy of Danvy and Filinski [17]. These results heavily rely on the effect subtyping relation that, e.g., allows to coerce pure expressions to effectful ones. From a more practical point of view, the selective CPS translation, that leaves pure (i.e., control-effect free) expressions in direct style and introduces explicit coercions to interpret effect subtyping in the source calculus, is a good candidate for embedding the control operators in an existing programming language, such as Scala [23].

In order to deal with the complexity of the source calculus and of the translation itself, we introduce binary logical relations on terms of the target calculus that are: heterogeneous, biorthogonal [14, 20, 13], and step-indexed [3, 2, 1]. Heterogeneity allows us to relate terms of different types, and in particular those in continuation-passing style with those in direct style. This is a crucial property, since the same term can have a pure type, resulting in a direct-style term through the translation and another, impure type, resulting in a term in continuation-passing style. Relating such terms requires quantification over types and to assure well-foundedness of the construction, we need to use step-indexing, which also supports reasoning about recursion, even if not in a critical way. We follow Dreyer et al. [12] in using logical step-indexed logical relations in our presentation of step-indexing. Biorthogonality, by

imposing a particular order of evaluation on expressions, simplifies the construction of the logical relations. It also facilitates reasoning about continuations represented as evaluation contexts.

Apart from the calculus with effect subtyping, we have used the ideas presented in this article to show coherence of subtyping in several other calculi, including the simply typed lambda calculus with subtyping [18] extended with recursion, the calculus of intersection types [22], and the lambda calculus with subtyping and the control operator `call/cc`.

The article is accompanied by a Coq development that presently consists of a library that provides a new shallow embedding of the logic for reasoning about step-indexed logical relations, and complete formalization of the proofs presented in the rest of the article. The code is available at <http://www.ii.uni.wroc.pl/~ppolesiuk/lrcoherence>.

The rest of this article is structured as follows. In Section 2, we briefly present Dreyer et al.'s logic for reasoning about step indexing [12] on which we base our presentation. We also describe the main ideas behind our Coq formalization of the logic. In Section 3, we introduce the construction of the logical relations in a simple yet sufficiently interesting scenario – the simply typed lambda calculus à la Curry with natural numbers, type `Top`, general recursion and standard subtyping. The goal of this section is to introduce the basic ingredients of the proof method before embarking on a considerably more challenging journey in the subsequent section. In Section 4, we present the main result of the article – the logical relations for establishing the coherence of the CPS translation from the calculus of delimited control with effect subtyping. In Section 5, we summarize the article.

2 Reasoning about step-indexed logical relations

Step-indexed logical relations [3, 2, 1] are a powerful tool for reasoning about programming languages. Instead of describing a general behavior of program execution, they focus on the first n computation steps, where the step index n is an additional parameter of the relation. This additional parameter makes it possible to define logical relations inductively not only on the structure of types, but also on the number of computation steps that are allowed for a program to make and, therefore, they provide an elegant way to reason about features that introduce non-termination to the programming language, including recursive types [2] and references [1].

However, reasoning directly about step-indexed logical relations is tedious because proofs become obscured by step-index arithmetic. Dreyer et al. [12] proposed logical step-indexed logical relations (LSLR) to avoid this problem. The LSLR logic is an intuitionistic logic for reasoning about one particular Kripke model: where possible worlds are natural numbers (step-indices) and where future worlds have smaller indices than the present one. All formulas are interpreted as monotone (non-increasing) sequences of truth values, whereas the connectives are interpreted as usual. In particular, in the case of implication we quantify over all future worlds to ensure monotonicity, so the formula $\varphi \Rightarrow \psi$ is valid at index n (written $n \models \varphi \Rightarrow \psi$) iff $k \models \varphi$ implies $k \models \psi$ for every $k \leq n$. In contrast to Dreyer et al. we do not assume that all formulas are valid in world 0, because it is not necessary.

The LSLR logic is also equipped with a modal operator \triangleright (later), to provide access to strictly future worlds. The formula $\triangleright\varphi$ means φ holds in any future world, or formally $\triangleright\varphi$ is always valid at world 0, and $n + 1 \models \triangleright\varphi$ iff φ is valid at n (and other future worlds by monotonicity). The later operator comes with two inference rules:

$$\frac{\Gamma, \Sigma \vdash \varphi}{\Gamma, \triangleright\Sigma \vdash \triangleright\varphi} \triangleright\text{-intro} \qquad \frac{\Gamma, \triangleright\varphi \vdash \varphi}{\Gamma \vdash \varphi} \text{Löb}$$

$$\begin{array}{ll} \tau ::= \text{Nat} \mid \text{Top} \mid \tau \rightarrow \tau & \text{(types)} \\ e ::= x \mid \lambda x.e \mid e e \mid \text{fix } x(x).e \mid n & \text{(expressions)} \end{array}$$

$$\begin{array}{c} \frac{}{\tau \leq \tau} \text{S-REFL} \quad \frac{\tau_2 \leq \tau_3 \quad \tau_1 \leq \tau_2}{\tau_1 \leq \tau_3} \text{S-TRANS} \quad \frac{}{\tau \leq \text{Top}} \text{S-TOP} \\ \frac{\tau'_2 \leq \tau'_1 \quad \tau_1 \leq \tau_2}{(\tau'_1 \rightarrow \tau_1) \leq (\tau'_2 \rightarrow \tau_2)} \text{S-ARR} \quad \frac{(x : \tau) \in \Gamma}{\Gamma \vdash x : \tau} \text{T-VAR} \quad \frac{\Gamma, x : \tau_1 \vdash e : \tau_2}{\Gamma \vdash \lambda x.e : \tau_1 \rightarrow \tau_2} \text{T-ABS} \\ \frac{\Gamma \vdash e_1 : \tau_2 \rightarrow \tau_1 \quad \Gamma \vdash e_2 : \tau_2}{\Gamma \vdash e_1 e_2 : \tau_1} \text{T-APP} \quad \frac{\Gamma, f : \tau_1 \rightarrow \tau_2, x : \tau_1 \vdash e : \tau_2}{\Gamma \vdash \text{fix } f(x).e : \tau_1 \rightarrow \tau_2} \text{T-FIX} \\ \frac{}{\Gamma \vdash n : \text{Nat}} \text{T-CONST} \quad \frac{\Gamma \vdash e : \tau \quad \tau \leq \tau'}{\Gamma \vdash e : \tau'} \text{T-SUB} \end{array}$$

■ **Figure 1** The source language – λ -calculus with subtyping.

The first rule allows one to shift reasoning to a future world, making the assumptions about the future world available. The Löb rule expresses an induction principle for indices. Note that the premise of the rule also captures the base case, because the assumption $\triangleright\varphi$ is trivial in the world 0. The later operator comes with no general elimination rule.

Predicates in LSLR logic as well as step-indexed logical relations can be defined inductively on indices. More generally, we can define a recursive predicate $\mu r.\varphi(r)$, provided all occurrences of r in φ are guarded by the later operator, to guarantee well-foundedness of the definition. For the sake of readability, in this paper we define recursive predicates and relations by giving a set of clauses instead of using the μ operator.

Since the logic is developed for reasoning about one particular model, we can freely add new inference rules for the logic if we prove they are valid in the model. We can also add new relations or predicates to the logic if we provide their monotone interpretation. In particular, constant functions are monotone, so we can safely use predicates defined outside of the logic, such as typing or reduction relations.

Our Coq formalization accompanying this article is built on our IxFree library that contains a shallow embedding of the LSLR logic similar to Appel et al.’s formalization of the “very modal model” [4]. Logical connectives including the later operator are functions on a special type `IProp` of “indexed propositions” defined as a type of monotone functions from `nat` to `Prop`. The library provides tactics representing the most important inference rules. One of the main differences between our library and Appel et al.’s formalization is a way of keeping track of the assumptions. Instead of interpreting a sequent $\varphi_1, \dots, \varphi_n \vdash \psi$ directly, we treat it as $k \models \psi$ with the standard Coq assumptions $k \models \varphi_1, \dots, k \models \varphi_n$. This approach is very convenient since it allows for reusing a number of existing Coq tactics.

3 Introducing the logical relations

In this section we prove the coherence of subtyping in the simply-typed call-by-value lambda calculus extended with recursion, where the coercion semantics is given by a standard translation to the simply-typed lambda calculus with explicit coercions [9]. Our goal here is to introduce the proof method in a simple scenario, so that in Section 4 we can focus on issues specific to control effects. The logical relations we present in this section are biorthogonal and step-indexed, which is not strictly necessary but it makes the development more elegant. Furthermore, biorthogonality and step-indexing become crucial in handling more complicated calculi such as the one of Section 4 and, therefore, are essential for the method to scale.

$\tau ::= \text{Nat} \mid \text{Unit} \mid \tau \rightarrow \tau$	(types)
$c ::= \text{id} \mid c \circ c \mid \text{top} \mid c \rightarrow c$	(coercions)
$e ::= x \mid \lambda x.e \mid ee \mid ce \mid \text{fix } x(x).e \mid n \mid \langle \rangle$	(expressions)
$v ::= x \mid \lambda x.e \mid \text{fix } x(x).e \mid (c \rightarrow c)v \mid n \mid \langle \rangle$	(values)
$E ::= \square \mid Ee \mid vE \mid cE$	(evaluation contexts)

$\frac{}{\text{id} :: \tau \triangleright \tau}$ S-REFL	$\frac{c_1 :: \tau_2 \triangleright \tau_3 \quad c_2 :: \tau_1 \triangleright \tau_2}{c_1 \circ c_2 :: \tau_1 \triangleright \tau_3}$ S-TRANS	$\frac{}{\text{top} :: \tau \triangleright \text{Unit}}$ S-TOP
$\frac{c_1 :: \tau'_2 \triangleright \tau'_1 \quad c_2 :: \tau_1 \triangleright \tau_2}{c_1 \rightarrow c_2 :: (\tau'_1 \rightarrow \tau_1) \triangleright (\tau'_2 \rightarrow \tau_2)}$ S-ARR	$\frac{}{\Gamma \vdash n : \text{Nat}}$ T-CONST	$\frac{(x : \tau) \in \Gamma}{\Gamma \vdash x : \tau}$ T-VAR
$\frac{}{\Gamma \vdash \langle \rangle : \text{Unit}}$ T-UNIT	$\frac{\Gamma, x : \tau_1 \vdash e : \tau_2}{\Gamma \vdash \lambda x.e : \tau_1 \rightarrow \tau_2}$ T-ABS	$\frac{c :: \tau \triangleright \tau' \quad \Gamma \vdash e : \tau}{\Gamma \vdash ce : \tau'}$ T-CAPP
$\frac{\Gamma \vdash e_1 : \tau_2 \rightarrow \tau_1 \quad \Gamma \vdash e_2 : \tau_2}{\Gamma \vdash e_1 e_2 : \tau_1}$ T-APP	$\frac{\Gamma, f : \tau_1 \rightarrow \tau_2, x : \tau_1 \vdash e : \tau_2}{\Gamma \vdash \text{fix } f(x).e : \tau_1 \rightarrow \tau_2}$ T-FIX	

$E[(\lambda x.e) v] \rightarrow_\beta E[e\{v/x\}]$	$E[\text{id } v] \rightarrow_\iota E[v]$
$E[(\text{fix } f(x).e) v] \rightarrow_\beta E[e\{\text{fix } f(x).e/f, v/x\}]$	$E[(c_1 \circ c_2) v] \rightarrow_\iota E[c_1 (c_2 v)]$
	$E[\text{top } v] \rightarrow_\iota E[\langle \rangle]$
	$E[(c_1 \rightarrow c_2) v_1 v_2] \rightarrow_\iota E[c_2 (v_1 (c_1 v_2))]$

■ **Figure 2** The target language – λ -calculus with explicit coercions.

3.1 The simply-typed lambda calculus with subtyping

The syntax and typing rules for the source language are given in Figure 1. The language is the simply-typed lambda calculus with recursive functions ($\text{fix } f(x).e$) and natural numbers (n). For brevity we do not consider any primitive operations on natural numbers or other basic types, but they could seamlessly be added to the language. We include the type **Top**, to make the subtyping relation interesting. The typing and subtyping rules are standard [18].

3.2 Coercion semantics

The semantics of the source language is given by a translation of the typing derivations to a target language that extends the source language with explicit type coercions (and replaces **Top** with **Unit**). The coercions express conversion of a term from one type to another, according to the subtyping relation. Figure 2 contains syntax, typing rules and reduction rules of the target language. The type coercions c and their typing rules correspond exactly to the subtyping rules of the source language. The grammar of terms contains explicit coercion application of the form ce and it is worth noting that terms of the form $(c \rightarrow c)v$ are considered values, since they represent a coercion expecting another value as argument (witness the last reduction rule).

The operational semantics of the target language distinguishes between β -rules that perform actual computations and ι -rules that rearrange coercions. Both of them are used during program evaluation. We say that program e terminates (written $e \downarrow$) when it can be reduced to a value using both sorts of reduction rules, according to the evaluation strategy determined by the evaluation contexts.

General contexts are closed terms with one hole (possibly under some binders), and are denoted by the metavariable C . We write $\vdash C : (\Gamma; \tau_1) \rightsquigarrow \tau_2$ if for any e with $\Gamma \vdash e : \tau_1$ we have $\Gamma \vdash C[e] : \tau_2$. Contextual approximation, written $\Gamma \vdash e_1 \lesssim_{ctx} e_2 : \tau$, means that

$$\llbracket \tau_1 \rightarrow \tau_2 \rrbracket = \llbracket \tau_1 \rrbracket \rightarrow \llbracket \tau_2 \rrbracket \quad \llbracket \text{Nat} \rrbracket = \text{Nat} \quad \llbracket \text{Top} \rrbracket = \text{Unit}$$

$$\begin{aligned} \mathcal{S}[\tau \leq \tau]_{\text{S-REFL}} &= \text{id} \\ \mathcal{S}[\text{Top} \leq \tau]_{\text{S-TOP}} &= \text{top} \\ \mathcal{S}[\tau_1 \leq \tau_3]_{\text{S-TRANS}(D_1, D_2)} &= \mathcal{S}[\tau_2 \leq \tau_3]_{D_1} \circ \mathcal{S}[\tau_1 \leq \tau_2]_{D_2} \\ \mathcal{S}[\tau'_1 \rightarrow \tau_1 \leq \tau'_2 \rightarrow \tau_2]_{\text{S-ARR}(D_1, D_2)} &= \mathcal{S}[\tau'_2 \leq \tau'_1]_{D_1} \rightarrow \mathcal{S}[\tau_1 \leq \tau_2]_{D_2} \end{aligned}$$

$$\begin{aligned} \mathcal{T}[x]_{\text{T-VAR}} &= x & \mathcal{T}[\text{fix } f(x).e]_{\text{T-FIX}(D)} &= \text{fix } f(x). \mathcal{T}[e]_D \\ \mathcal{T}[\lambda x.e]_{\text{T-ABS}(D)} &= \lambda x. \mathcal{T}[e]_D & \mathcal{T}[e]_{\text{T-SUB}(D_1, D_2)} &= \mathcal{S}[\tau \leq \tau']_{D_2} \mathcal{T}[e]_{D_1} \\ \mathcal{T}[e_1 e_2]_{\text{T-APP}(D_1, D_2)} &= \mathcal{T}[e_1]_{D_1} \mathcal{T}[e_2]_{D_2} & \mathcal{T}[n]_{\text{T-CONST}} &= n \end{aligned}$$

■ **Figure 3** Coercion semantics for the λ -calculus with subtyping.

for any context C and type τ' , such that $\vdash C : (\Gamma; \tau) \rightsquigarrow \tau'$ if $C[e_1]$ terminates, then so does $C[e_2]$. If $\Gamma \vdash e_1 \lesssim_{ctx} e_2 : \tau$ and $\Gamma \vdash e_2 \lesssim_{ctx} e_1 : \tau$, then we say that e_1 and e_2 are contextually equivalent.

The coercion semantics of the source language is given in Figure 3. The function $\mathcal{S}[\cdot]$ translates subtyping proofs into coercions, and function $\mathcal{T}[\cdot]$ translates typing derivations into terms of the target language.

► **Lemma 1.** *Coercion semantics preserves types.*

1. If $D :: \tau_1 \leq \tau_2$ then $\mathcal{S}[\tau_1 \leq \tau_2]_D :: \llbracket \tau_1 \rrbracket \triangleright \llbracket \tau_2 \rrbracket$.
2. If $D :: \Gamma \vdash e : \tau$ then $\llbracket \Gamma \rrbracket \vdash \mathcal{T}[e]_D : \llbracket \tau \rrbracket$.

3.3 Logical relations

In order to reason about contextual equivalence in the target language, we define logical relations (Figure 4). Relations are expressed in the LSLR logic described in Section 2, so they are implicitly step-indexed.

We call these relations heterogeneous because they are parameterized by two types, one for each of the arguments. This property is important for our coherence proof, since it makes it possible to relate the results of the translation of two typing derivations which assign different types to the same term. When both types τ_1 and τ_2 are Nat or both are arrow types, the value relation $\mathcal{V}[\tau_1; \tau_2]$ is standard. Two values are related for type Nat if they are the same constant, and two functions are related when they map related arguments to related results. The most interesting are the cases when type parameters of the relation are different. When one of these types is Unit , then any values are in the relation, because we do not expect them to carry any information – Unit is the result of translating the Top type. Functions are never related with natural numbers.

The relation $\mathcal{E}[\tau_1; \tau_2]$ for closed terms is defined by biorthogonality. Two terms are related if they behave the same in related contexts, and contexts are related (relation $\mathcal{K}[\tau_1; \tau_2]$) if they yield the same observations when plugged with related values. Yielding the same observations (relation \lesssim) is defined for each step-index separately: $e_1 \lesssim_k e_2$ is valid at k if termination of e_1 using at most k β -steps (and any number of ι -steps) implies termination of e_2 in any number of steps. This interpretation is monotone, so the relation \lesssim can be added to the LSLR logic. The relation $\mathcal{E}[\tau_1; \tau_2]$ is extended to open terms as usual: two open

$$\begin{aligned}
(v_1, v_2) \in \mathcal{V}[\mathbf{Nat}; \mathbf{Nat}] &\iff \exists n, v_1 = v_2 = n \\
(v_1, v_2) \in \mathcal{V}[\tau'_1 \rightarrow \tau_1; \tau'_2 \rightarrow \tau_2] &\iff \forall (a_1, a_2) \in \mathcal{V}[\tau'_1; \tau'_2]. (v_1 a_1, v_2 a_2) \in \mathcal{E}[\tau_1; \tau_2] \\
(v_1, v_2) \in \mathcal{V}[\tau_1; \tau_2] &\iff \top \quad \text{if } \tau_1 = \mathbf{Unit} \text{ or } \tau_2 = \mathbf{Unit} \\
(v_1, v_2) \in \mathcal{V}[\tau_1; \tau_2] &\iff \perp \quad \text{otherwise} \\
(e_1, e_2) \in \mathcal{E}[\tau_1; \tau_2] &\iff \forall (E_1, E_2) \in \mathcal{K}[\tau_1; \tau_2]. E_1[e_1] \lesssim E_2[e_2] \\
(E_1, E_2) \in \mathcal{K}[\tau_1; \tau_2] &\iff \forall (v_1, v_2) \in \mathcal{V}[\tau_1; \tau_2]. E_1[v_1] \lesssim E_2[v_2] \\
k \models e_1 \lesssim e_2 &\iff e_1 \downarrow^k \implies e_2 \downarrow \\
(\gamma_1, \gamma_2) \in \mathcal{G}[\Gamma_1; \Gamma_2] &\iff \forall x, (\gamma_1(x), \gamma_2(x)) \in \mathcal{V}[\Gamma_1(x); \Gamma_2(x)] \\
\Gamma_1; \Gamma_2 \vdash e_1 \lesssim_{log} e_2 : \tau_1; \tau_2 &\iff \forall (\gamma_1, \gamma_2) \in \mathcal{G}[\Gamma_1; \Gamma_2]. (e_1 \gamma_1, e_2 \gamma_2) \in \mathcal{E}[\tau_1; \tau_2]
\end{aligned}$$

■ **Figure 4** Logical relations for the λ -calculus with explicit coercions.

terms are related (written $\Gamma_1; \Gamma_2 \vdash e_1 \lesssim_{log} e_2 : \tau_1; \tau_2$) when every pair of related closing substitutions makes them related.

Notice that we do not assume that related terms have valid types. Our relations may include some “garbage”, e.g., $(1, \lambda x.x) \in \mathcal{V}[\mathbf{Unit}; \mathbf{Nat}]$, but it is non-problematic. One can mechanically prune these relations to well-typed terms, but this change complicates formalization and we did not find it useful.

In this presentation we consider languages with only one base type. Adding more base types and some subtyping between them will not change the general shape of the proof, but defining logical relations for such a case is a little trickier. We would stipulate that two values v_1 and v_2 are related for base types b_1 and b_2 iff for every common supertype b of b_1 and b_2 , coercing v_1 and v_2 to b yields the same constant.

The relation \lesssim is preserved by reductions in the following sense, where the third assertion expresses an elimination rule of the later modality that is crucial in the subsequent proofs.

► **Lemma 2.** *The following assertions hold:*

1. If $e_1 \rightarrow_i e'_1$ and $e'_1 \lesssim e_2$ then $e_1 \lesssim e_2$.
2. If $e_2 \rightarrow_i e'_2$ and $e_1 \lesssim e'_2$ then $e_1 \lesssim e_2$.
3. If $e_1 \rightarrow_\beta e'_1$ and $\triangleright e'_1 \lesssim e_2$ then $e_1 \lesssim e_2$.
4. If $e_2 \rightarrow_\beta e'_2$ and $e_1 \lesssim e'_2$ then $e_1 \lesssim e_2$.

The proof of soundness of the logical relations follows closely the standard technique for biorthogonal logical relations [20, 13]. First, we need to show compatibility lemmas, which state that the relation is preserved by every language construct. Most of them are standard and we omit them due to lack of space. The only compatibility lemma specific to our relations is the following lemma for coercion application.

► **Lemma 3 (Coercion compatibility).** *The logical relation is preserved by coercion application.*

1. If $c :: \tau_1 \triangleright \tau_2$ and $\Gamma_1; \Gamma_2 \vdash e_1 \lesssim_{log} e_2 : \tau_1; \tau_0$ then $\Gamma_1; \Gamma_2 \vdash c e_1 \lesssim_{log} e_2 : \tau_2; \tau_0$.
2. If $c :: \tau_1 \triangleright \tau_2$ and $\Gamma_1; \Gamma_2 \vdash e_1 \lesssim_{log} e_2 : \tau_0; \tau_1$ then $\Gamma_1; \Gamma_2 \vdash e_1 \lesssim_{log} c e_2 : \tau_0; \tau_2$.

Proof. We prove both cases by induction on the typing derivation of the coercion c . ◀

► **Theorem 4 (Fundamental property).** *If $\Gamma \vdash e : \tau$ then $\Gamma; \Gamma \vdash e \lesssim_{log} e : \tau; \tau$.*

Proof. By induction on the derivation $\Gamma \vdash e : \tau$. In each case we apply the corresponding compatibility lemma. ◀

► **Lemma 5** (Precongruence). *If $\vdash C : (\Gamma; \tau) \rightsquigarrow \tau_0$ and $\Gamma; \Gamma \vdash e_1 \lesssim_{log} e_2 : \tau; \tau$ then $(C[e_1], C[e_2]) \in \mathcal{E}[\tau_0; \tau_0]$.*

Proof. By induction on the derivation of context typing, using the appropriate compatibility lemma in each case. For contexts containing subterms we also need the fundamental property. For the empty context we use the empty substitution, since the empty substitutions are in relation $\mathcal{G}[\emptyset; \emptyset]$. ◀

► **Lemma 6** (Adequacy). *If $(e_1, e_2) \in \mathcal{E}[\tau; \tau]$ then $e_1 \lesssim e_2$.*

Proof. Let us show $\square[e_1] \lesssim \square[e_2]$. Using the assertion $(e_1, e_2) \in \mathcal{E}[\tau; \tau]$, it suffices to show $(\square, \square) \in \mathcal{K}[\tau; \tau]$, which is trivial, since values always terminate. ◀

► **Theorem 7** (Soundness). *If $k \models \Gamma; \Gamma \vdash e_1 \lesssim_{log} e_2 : \tau; \tau$ holds for every k , then $\Gamma \vdash e_1 \lesssim_{ctx} e_2 : \tau$.*

Proof. Suppose $\vdash C : (\Gamma; \tau) \rightsquigarrow \tau_0$ and $C[e_1] \downarrow$, we need to show $C[e_2] \downarrow$. By Lemma 5 and Lemma 6 we know $k \models C[e_1] \lesssim C[e_2]$ for every k . Taking k to be the number of steps in which $C[e_1]$ terminates, we have that $C[e_2]$ also terminates, by the definition of \lesssim . ◀

3.4 Coherence of the coercion semantics

Having established soundness of the logical relations, we are in a position to prove the main coherence lemma, phrased in terms of the logical relations, and the coherence theorem.

► **Lemma 8.** *If $D_i :: \Gamma_i \vdash e : \tau_i$ for $i = 1, 2$ are two typing judgments for the same term e of the source language, then $[\Gamma_1]; [\Gamma_2] \vdash \mathcal{T}[e]_{D_1} \lesssim_{log} \mathcal{T}[e]_{D_2} : [\tau_1]; [\tau_2]$.*

Proof. The proof follows by induction on the structure of both derivations D_1 and D_2 . At least one of these derivations is decreased in every case. When one of derivations starts with subsumption rule (T-SUB), we apply Lemma 3. The coercion that we get after translation is well-typed by Lemma 1. In other cases we just apply the appropriate compatibility lemma. ◀

► **Theorem 9** (Coherence). *If D_1 and D_2 are derivations of the same typing judgment $\Gamma \vdash e : \tau$, then $[\Gamma] \vdash \mathcal{T}[e]_{D_1} \lesssim_{ctx} \mathcal{T}[e]_{D_2} : [\tau]$.*

Proof. Immediately from Lemma 8 and Theorem 7. ◀

Coercion semantics described here translates the source language into the language with explicit coercions. We chose coercions to be a separate syntactic category, because we found it very convenient, especially for proving Lemma 3. However, one can define a coercion semantics which translates subtyping proofs directly to λ -expressions. Our result can be easily extended for such a translation. Let $|e|$ be a term e with all the coercions replaced by the corresponding expressions. To prove that for any contextually equivalent terms e_1 and e_2 in the language with coercions, terms $|e_1|$ and $|e_2|$ are contextually equivalent in the language without coercions, we need three simple facts that can be easily verified:

1. every well-typed term in the language without coercions is well typed in the language with coercions,
2. term e terminates iff $|e|$ terminates,
3. if context C does not contain coercions then $C[|e|] = |C[e]|$.

In the next section we show that the results presented in this section can be adapted to a considerably more complex calculus – a calculus of delimited control with control-effect subtyping.

$\tau ::= \text{Nat} \mid \tau \rightarrow T$	(pure types)
$T ::= \tau \mid \tau[T]T$	(types)
$e ::= x \mid \lambda x.e \mid ee \mid \text{fix } x(x).e \mid \mathcal{S}_0 x.e \mid \langle e \rangle \mid n$	(expressions)

$\frac{}{T \leq T}$	$\frac{T_2 \leq T_3 \quad T_1 \leq T_2}{T_1 \leq T_3}$	$\frac{\tau_1 \leq \tau_2 \quad T_2 \leq T_1 \quad U_1 \leq U_2}{\tau_1[T_1]U_1 \leq \tau_2[T_2]U_2}$
$\frac{T_1 \leq T_2}{\tau \leq \tau[T_1]T_2}$	$\frac{\tau_2 \leq \tau_1 \quad T_1 \leq T_2}{(\tau_1 \rightarrow T_1) \leq (\tau_2 \rightarrow T_2)}$	$\frac{\Gamma \vdash e : T \quad T \leq U}{\Gamma \vdash e : U}$
$\frac{(x : \tau) \in \Gamma}{\Gamma \vdash x : \tau}$	$\frac{\Gamma, x : \tau \vdash e : T}{\Gamma \vdash \lambda x.e : \tau \rightarrow T}$	$\frac{\Gamma \vdash e_1 : \tau \rightarrow T \quad \Gamma \vdash e_2 : \tau}{\Gamma \vdash e_1 e_2 : T}$
$\frac{\Gamma \vdash e_1 : (\tau_2 \rightarrow \tau_1[U_4]U_3)[U_2]U_1 \quad \Gamma \vdash e_2 : \tau_2[U_3]U_2}{\Gamma \vdash e_1 e_2 : \tau_1[U_4]U_1}$		$\frac{}{\Gamma \vdash n : \text{Nat}}$
$\frac{\Gamma, f : \tau \rightarrow T, x : \tau \vdash e : T}{\Gamma \vdash \text{fix } f(x).e : \tau \rightarrow T}$	$\frac{\Gamma, x : \tau \rightarrow T \vdash e : U}{\Gamma \vdash \mathcal{S}_0 x.e : \tau[T]U}$	$\frac{\Gamma \vdash e : \tau[\tau]T}{\Gamma \vdash \langle e \rangle : T}$

■ **Figure 5** The source language – λ -calculus with delimited control and effect subtyping.

4 Coherence of a CPS translation of control-effect subtyping

The calculus of delimited control studied in this section is the λ -calculus extended with recursion and the control operators shift_0 (\mathcal{S}_0) and reset_0 ($\langle \cdot \rangle$) [11] that can explore and reorganize an arbitrary portion of the stack of delimited continuations. It was built around a central idea that types for such a calculus should contain information about the stack and that the amount of that information should be governed by a subtyping mechanism [16]. We will define the semantics of the calculus by a CPS translation to a target calculus endowed with a reduction semantics, but if we were to directly give a reduction rule for shift_0 , it would be:

$$F[\langle E[\mathcal{S}_0 x.e] \rangle] \rightarrow F\{e\{\lambda y.\langle E[y] \rangle/x\}$$

where E is a pure evaluation context representing the current delimited continuation (delimited by $\langle \cdot \rangle$ and captured by \mathcal{S}_0) and F is a metacontext, i.e., a list of such pure evaluation contexts separated by control delimiters, representing the current metacontinuation [6]. In terms of abstract-machine semantics E together with F represent the entire control stack. We can see from the reduction rule that nothing prevents the expression e from capturing another delimited continuation from F and, therefore, the types of the calculus express requirements on the shape of the control stack, so that expressions can safely perform control operations exploring the stack. However, under some conditions, an expression that imposes certain requirements on the stack can be used with a stack of which more is known or assumed. Such coercions are possible thanks to the subtyping relation that lies at the heart of the calculus.

4.1 The lambda calculus with delimited control and effect subtyping

The syntax and typing rules of the calculus of delimited control are shown in Figure 5. Types are either pure (τ) or effect annotated ($\tau[T_1]T_2$). An expression of type $\tau[T_1]T_2$ can be used with the top-most delimited continuation that when given a value of type τ behaves as specified by its answer type T_1 , and with the rest of the stack whose type is T_2 .

The calculus comprises the simply typed lambda calculus with the standard subtyping rules, extended with typing and subtyping rules for effectful computations. The most interesting from the point of view of effect subtyping are the rules S-CONS and S-LIFT. The

$\tau ::= \text{Nat} \mid \tau \rightarrow T$	(pure types)
$T ::= \tau \mid (\tau \rightarrow T) \Rightarrow T$	(types)
$c ::= \text{id} \mid c \circ c \mid c \rightarrow c \mid \uparrow c \mid c[c]c$	(coercions)
$e ::= x \mid \lambda x.e \mid ee \mid ce \mid \text{fix } x(x).e \mid n$	(expressions)
$v ::= x \mid \lambda x.e \mid \text{fix } x(x).e \mid (c \rightarrow c) v \mid \uparrow c v \mid (c[c]c) v \mid n$	(values)
$E ::= \square \mid E e \mid v E \mid c E$	(evaluation contexts)

$\frac{}{\text{id} :: T \triangleright T}$	S-REFL	$\frac{c_1 :: T_2 \triangleright T_3 \quad c_2 :: T_1 \triangleright T_2}{c_1 \circ c_2 :: T_1 \triangleright T_3}$	S-TRANS	
$\frac{c_1 :: \tau_2 \triangleright \tau_1 \quad c_2 :: T_1 \triangleright T_2}{c_1 \rightarrow c_2 :: (\tau_1 \rightarrow T_1) \triangleright (\tau_2 \rightarrow T_2)}$				S-ARR
$\frac{c :: T_1 \triangleright T_2}{\uparrow c :: \tau \triangleright ((\tau \rightarrow T_1) \Rightarrow T_2)}$	S-LIFT	$\frac{c :: \tau_1 \triangleright \tau_2 \quad c_1 :: T_2 \triangleright T_1 \quad c_2 :: U_1 \triangleright U_2}{c[c_1]c_2 :: ((\tau_1 \rightarrow T_1) \Rightarrow U_1) \triangleright ((\tau_2 \rightarrow T_2) \Rightarrow U_2)}$	S-CONS	
$\frac{(x : \tau) \in \Gamma}{\Gamma \vdash x : \tau}$	T-VAR	$\frac{\Gamma, x : \tau \vdash e : T}{\Gamma \vdash \lambda x.e : \tau \rightarrow T}$	T-ABS	
$\frac{\Gamma \vdash e_1 : \tau \rightarrow T \quad \Gamma \vdash e_2 : \tau}{\Gamma \vdash e_1 e_2 : T}$		T-APP		
$\frac{\Gamma, x : \tau \rightarrow T \vdash e : U}{\Gamma \vdash \lambda x.e : (\tau \rightarrow T) \Rightarrow U}$		T-KABS	$\frac{\Gamma \vdash e : (\tau \rightarrow T) \Rightarrow U \quad \Gamma \vdash v : \tau \rightarrow T}{\Gamma \vdash e v : U}$	T-KAPP
$\frac{c :: T \triangleright U \quad \Gamma \vdash e : T}{\Gamma \vdash c e : U}$	T-CAPP	$\frac{\Gamma, f : \tau \rightarrow T, x : \tau \vdash e : T}{\Gamma \vdash \text{fix } f(x).e : \tau \rightarrow T}$	T-FIX	
$\frac{}{\Gamma \vdash n : \text{Nat}}$		T-CONST		

$E[(\lambda x.e) v] \rightarrow_\beta E\{v/x\}$	$E[\text{id } v] \rightarrow_\iota E[v]$
$E[(\text{fix } f(x).e) v] \rightarrow_\beta E\{e\{\text{fix } f(x).e/f, v/x\}\}$	$E[(c_1 \circ c_2) v] \rightarrow_\iota E[c_1 (c_2 v)]$
$E[\uparrow c v_1 v_2] \rightarrow_\beta E[c (v_2 v_1)]$	$E[(c_1 \rightarrow c_2) v_1 v_2] \rightarrow_\iota E[c_2 (v_1 (c_1 v_2))]$
	$E[(c[c_1]c_2) v_1 v_2] \rightarrow_\iota E[c_2 (v_1 ((c \rightarrow c_1) v_2))]$

■ **Figure 6** The target language – λ -calculus with explicit coercions of control effects.

rule S-CONS is a direct consequence of the interpretation of effect-annotated types given above that actually follows from the CPS interpretation of delimited continuations – a type $\tau[T_1]T_2$ is interpreted in CPS as $(\tau \rightarrow T_1) \Rightarrow T_2$, where \Rightarrow means an effectful function space (see Section 4.2). The rule S-LIFT is more interesting and it says that a pure computation can be considered impure, provided the answer type of the top-most delimited continuation can be coerced into the type of the rest of the stack. We have, for an example, $\text{Nat} \leq \text{Nat}[\text{Nat}]\text{Nat}$. It is worth noting that the calculus provides two rules for function application. They are necessary for expressivity reasons, but at the same time they are an additional source of difficulty when it comes to establishing the coherence of the subtyping. For more detailed presentation of the calculus we refer the reader to [16]. (The presentation of the type system in [16] differs from the one shown in Figure 5 in some inessential details, but the two type systems are equally expressive.)

4.2 Coercion semantics: a type-directed selective CPS translation

The syntax and typing rules of the target language are presented in Figure 6. There are two kinds of arrow type: the usual one $\tau \rightarrow T$ for regular functions and the effectful one $(\tau \rightarrow T) \Rightarrow U$ for expressions in CPS. We make this distinction to express the fact that the CPS translation (see Figure 7) yields expressions with strong restrictions on the occurrence of terms in CPS: they are never passed as arguments (typing environment consists of only pure types) and they can be applied only to values (witness the rule T-KAPP) representing continuations.

$$\begin{aligned} \llbracket \text{Nat} \rrbracket_p &= \text{Nat} & \llbracket \tau \rrbracket &= \llbracket \tau \rrbracket_p \\ \llbracket \tau \rightarrow T \rrbracket_p &= \llbracket \tau \rrbracket_p \rightarrow \llbracket T \rrbracket & \llbracket \tau[T]U \rrbracket &= (\llbracket \tau \rrbracket_p \rightarrow \llbracket T \rrbracket) \Rightarrow \llbracket U \rrbracket \end{aligned}$$

$$\begin{aligned} \mathcal{S}[\llbracket T \leq T \rrbracket]_{\text{S-REFL}} &= \text{id} \\ \mathcal{S}[\llbracket T_1 \leq T_3 \rrbracket]_{\text{S-TRANS}(D_1, D_2)} &= \mathcal{S}[\llbracket T_2 \leq T_3 \rrbracket]_{D_1} \circ \mathcal{S}[\llbracket T_1 \leq T_2 \rrbracket]_{D_2} \\ \mathcal{S}[\llbracket \tau_1 \rightarrow T_1 \leq \tau_2 \rightarrow T_2 \rrbracket]_{\text{S-ARR}(D_1, D_2)} &= \mathcal{S}[\llbracket \tau_2 \leq \tau_1 \rrbracket]_{D_1} \rightarrow \mathcal{S}[\llbracket T_1 \leq T_2 \rrbracket]_{D_2} \\ \mathcal{S}[\llbracket \tau \leq \tau[T]U \rrbracket]_{\text{S-LIFT}(D)} &= \uparrow \mathcal{S}[\llbracket T \leq U \rrbracket]_D \\ \mathcal{S}[\llbracket \tau_1[T_1]U_1 \leq \tau_2[T_2]U_2 \rrbracket]_{\text{S-CONS}(D, D_1, D_2)} &= \mathcal{S}[\llbracket \tau_1 \leq \tau_2 \rrbracket]_D \mathcal{S}[\llbracket T_2 \leq T_1 \rrbracket]_{D_1} \mathcal{S}[\llbracket U_1 \leq U_2 \rrbracket]_{D_2} \end{aligned}$$

$$\begin{aligned} \mathcal{T}[\llbracket x \rrbracket]_{\text{T-VAR}} &= x & \mathcal{T}[\llbracket \mathcal{S}_0 x.e \rrbracket]_{\text{T-SFT}(D)} &= \lambda x. \mathcal{T}[\llbracket e \rrbracket]_D \\ \mathcal{T}[\llbracket \lambda x.e \rrbracket]_{\text{T-ABS}(D)} &= \lambda x. \mathcal{T}[\llbracket e \rrbracket]_D & \mathcal{T}[\llbracket \langle e \rangle \rrbracket]_{\text{T-RST}(D)} &= \mathcal{T}[\llbracket e \rrbracket]_D (\lambda x.x) \\ \mathcal{T}[\llbracket \text{fix } f(x).e \rrbracket]_{\text{T-FIX}(D)} &= \text{fix } f(x). \mathcal{T}[\llbracket e \rrbracket]_D & \mathcal{T}[\llbracket n \rrbracket]_{\text{T-CONST}} &= n \\ \mathcal{T}[\llbracket e \rrbracket]_{\text{T-SUB}(D_1, D_2)} &= \mathcal{S}[\llbracket T \leq U \rrbracket]_{D_2} \mathcal{T}[\llbracket e \rrbracket]_{D_1} \\ \mathcal{T}[\llbracket e_1 e_2 \rrbracket]_{\text{T-PAPP}(D_1, D_2)} &= \mathcal{T}[\llbracket e_1 \rrbracket]_{D_1} \mathcal{T}[\llbracket e_2 \rrbracket]_{D_2} \\ \mathcal{T}[\llbracket e_1 e_2 \rrbracket]_{\text{T-APP}(D_1, D_2)} &= \lambda k. \mathcal{T}[\llbracket e_1 \rrbracket]_{D_1} (\lambda f. \mathcal{T}[\llbracket e_2 \rrbracket]_{D_2} (\lambda x. f x k)) \end{aligned}$$

■ **Figure 7** Type-directed selective CPS translation.

Again, the operational semantics distinguishes between β -rules and ι -rules. We classified the last β -rule as “actual computation” because it does not only rearrange coercions. It translates back a lifted value v_1 and applies to it a given continuation v_2 . This rule and the last ι -rule reduce a coerced value applied to a continuation, so terms of the form $(\uparrow c v)$ and $(c[c]c v)$ are considered values. Notice that these values have effectful types. We extend the notion of ι -reduction to evaluation contexts: $E_1 \rightarrow_{\iota} E_2$ holds iff $E_1[v] \rightarrow_{\iota} E_2[v]$ for every value v .

As in Section 3, the metavariable C denotes general closed contexts. We also define typing of general contexts $\vdash C : (\Gamma; T) \rightsquigarrow T_0$ as before. The definition of contextual approximation $\Gamma \vdash e_1 \lesssim_{ctx} e_2 : T$ is slightly weaker, because we allow only contexts with pure answer type. Indeed, an expression that requires a continuation to trigger computation can hardly be considered a complete program.

The coercion semantics of the source language is given by the type-directed selective CPS translation presented in Figure 7. The translation is selective because it leaves terms of pure type in direct style – witness, e.g, the equations for variable or pure application. Effectful applications are translated according to Plotkin’s call-by-value CPS translation [21], whereas the translation of shift_0 and reset_0 is surprisingly straightforward – shift_0 is turned into a lambda-abstraction expecting a delimited continuation, and reset_0 is interpreted by providing its subexpression with the reset delimited continuation, represented by the identity function.

► **Example 10.** Consider the program $(\text{fix } f(x).f x) 1$ in the source language. We derive the type $\text{Nat}[T]T$ for it in two ways: let D_1 be the derivation

$$\frac{\frac{\frac{\dots}{f : \text{Nat} \rightarrow \text{Nat}[T]T, x : \text{Nat} \vdash f x : \text{Nat}[T]T} \text{T-FIX}}{\vdash \text{fix } f(x).f x : \text{Nat} \rightarrow \text{Nat}[T]T} \text{T-CONST}}{\vdash (\text{fix } f(x).f x) 1 : \text{Nat}[T]T} \text{T-PAPP}}{\vdash 1 : \text{Nat}}$$

and D_2 be the derivation

$$\begin{array}{lcl}
(v_1, v_2) \in \mathcal{V}[\mathbf{Nat}; \mathbf{Nat}] & \iff & \exists n, v_1 = v_2 = n \\
(v_1, v_2) \in \mathcal{V}[\tau_1 \rightarrow T_1; \tau_2 \rightarrow T_2] & \iff & \forall (a_1, a_2) \in \mathcal{V}[\tau_1; \tau_2]. (v_1 a_1, v_2 a_2) \in \mathcal{E}[\mathbf{T}_1; \mathbf{T}_2] \\
(v_1, v_2) \in \mathcal{V}[\tau_1; \tau_2] & \iff & \perp \quad \text{otherwise} \\
(e_1, e_2) \in \mathcal{E}[\mathbf{T}_1; \mathbf{T}_2] & \iff & \forall (E_1, E_2) \in \mathcal{K}[\mathbf{T}_1; \mathbf{T}_2]. E_1[e_1] \lesssim E_2[e_2] \\
(E_1, E_2) \in \mathcal{K}[\tau_1; \tau_2] & \iff & \forall (v_1, v_2) \in \mathcal{V}[\tau_1; \tau_2]. E_1[v_1] \lesssim E_2[v_2] \\
(E_1, E_2) \in \mathcal{K}[\tau_1; (\tau_2 \rightarrow T_2) \Rightarrow U_2] & \iff & \exists T, (E, k) \in \mathcal{K}\mathcal{V}[\tau_1 \rightsquigarrow T; \tau_2 \rightarrow T_2], \\
& & (E'_1, E'_2) \in \mathcal{K}[T; U_2]. \\
& & E_1 \rightarrow_{\iota}^* E'_1[E[]] \wedge E_2 \rightarrow_{\iota}^* E'_2[[] k] \\
(E_1, E_2) \in \mathcal{K}[(\tau_1 \rightarrow T_1) \Rightarrow U_1; \tau_2] & \iff & \exists T, (k, E) \in \triangleright \mathcal{K}\mathcal{V}[\tau_1 \rightarrow T_1; \tau_2 \rightsquigarrow T], \\
& & (E'_1, E'_2) \in \triangleright \mathcal{K}[U_1; T]. \\
& & E_1 \rightarrow_{\iota}^* E'_1[[] k] \wedge E_2 \rightarrow_{\iota}^* E'_2[E[]] \\
(E_1, E_2) \in \mathcal{K}[(\tau_1 \rightarrow T_1) \Rightarrow U_1; \\
(\tau_2 \rightarrow T_2) \Rightarrow U_2] & \iff & \exists (k_1, k_2) \in \mathcal{V}[\tau_1 \rightarrow T_1; \tau_2 \rightarrow T_2], \\
& & (E'_1, E'_2) \in \mathcal{K}[U_1; U_2]. \\
& & E_1 \rightarrow_{\iota}^* E'_1[[] k_1] \wedge E_2 \rightarrow_{\iota}^* E'_2[[] k_2] \\
(E, v) \in \mathcal{K}\mathcal{V}[\tau_1 \rightsquigarrow T_1; \tau_2 \rightarrow T_2] & \iff & \forall (a_1, a_2) \in \mathcal{V}[\tau_1; \tau_2]. (E[a_1], v a_2) \in \mathcal{E}[\mathbf{T}_1; \mathbf{T}_2] \\
(v, E) \in \mathcal{V}\mathcal{K}[\tau_1 \rightarrow T_1; \tau_2 \rightsquigarrow T_2] & \iff & \forall (a_1, a_2) \in \mathcal{V}[\tau_1; \tau_2]. (v a_1, E[a_2]) \in \mathcal{E}[\mathbf{T}_1; \mathbf{T}_2] \\
(\gamma_1, \gamma_2) \in \mathcal{G}[\Gamma_1; \Gamma_2] & \iff & \forall x. (\gamma_1(x), \gamma_2(x)) \in \mathcal{V}[\Gamma_1(x); \Gamma_2(x)] \\
\Gamma_1; \Gamma_2 \vdash e_1 \lesssim_{\log} e_2 : T_1; T_2 & \iff & \forall (\gamma_1, \gamma_2) \in \mathcal{G}[\Gamma_1; \Gamma_2]. (e_1 \gamma_1, e_2 \gamma_2) \in \mathcal{E}[\mathbf{T}_1; \mathbf{T}_2]
\end{array}$$

■ **Figure 8** Logical relations for the λ -calculus with explicit coercions of control effects.

$$\frac{\frac{\vdash \text{fix } f(x).f x : \mathbf{Nat} \rightarrow \mathbf{Nat}[T]T \quad \dots}{\vdash \text{fix } f(x).f x : (\mathbf{Nat} \rightarrow \mathbf{Nat}[T]T)[T]T} \text{T-SUB} \quad \frac{\vdash 1 : \mathbf{Nat} \quad \dots}{\vdash 1 : \mathbf{Nat}[T]T} \text{T-CONST}}{\vdash (\text{fix } f(x).f x) 1 : \mathbf{Nat}[T]T} \text{T-APP}$$

Then

$$\begin{aligned}
\mathcal{T}[(\text{fix } f(x).f x) 1]_{D_1} &= (\text{fix } f(x).f x) 1 \\
\mathcal{T}[(\text{fix } f(x).f x) 1]_{D_2} &= \lambda k. (\uparrow \text{id } (\text{fix } f(x).f x)) (\lambda g. (\uparrow \text{id } 1) (\lambda y. g y k))
\end{aligned}$$

By the results of the next sections, these terms are contextually equivalent.

► **Lemma 11.** *Coercion semantics preserves types.*

1. If $D :: T_1 \leq T_2$ then $\mathcal{S}[\mathbf{T}_1 \leq \mathbf{T}_2]_D :: \llbracket \mathbf{T}_1 \rrbracket \triangleright \llbracket \mathbf{T}_2 \rrbracket$.
2. If $D :: \Gamma \vdash e : T$ then $\llbracket \Gamma \rrbracket \vdash \mathcal{T}[e]_D : \llbracket T \rrbracket$.

4.3 Logical relations

The logical relations are defined in Figure 8. The relation $\mathcal{V}[\tau_1; \tau_2]$ for pure values and the relation $\mathcal{E}[\mathbf{T}_1; \mathbf{T}_2]$ for expressions are similar to the relations defined in Section 3.3. All information about control effects is captured in the relation $\mathcal{K}[\mathbf{T}_1; \mathbf{T}_2]$ for contexts. If T_1 and T_2 are computation arrow types, then two contexts are related iff they can be decomposed as applications to related continuations in related contexts. In general, this application to a continuation does not have to be the top-most element of the context and some rearrangement of coercions is needed, so such a decomposition is defined by ι -reduction of contexts.

The most interesting are the cases that relate pure and impure contexts. As previously, the impure context should be decomposed to a continuation k and the rest of the context. Then the pure context should be decomposed in such a way that the continuation k is

related with some portion E of the pure context. The answer type of E cannot be retrieved from the type of the initial pure context, so we quantify over all possible types. In order to make the construction well-founded, the relations are defined by nested induction on step indices and on the structure of the second type. Notice that step indices play a role only in one case – when we quantify over the second type and the later operator guards the non-structural use of the relations $\mathcal{VK}[\tau_1 \rightarrow T_1; \tau_2 \rightsquigarrow T]$ and $\mathcal{K}[[U_1; T]]$. The auxiliary relations $\mathcal{KV}[\tau_1 \rightsquigarrow T_1; \tau_2 \rightarrow T_2]$ and $\mathcal{VK}[\tau_1 \rightarrow T_1; \tau_2 \rightsquigarrow T_2]$ relate a portion of an evaluation context with a value of an arrow type and they are defined analogously to the value relation for functions.

The relations of this section possess properties analogous to the ones of Section 3.3, in particular the relation \lesssim is preserved by reduction (Lemma 2) and the compatibility lemmas (including Lemma 3) hold. However, the proof of the compatibility lemmas requires the following results that establish the preservation of relations with respect to ι -reductions of evaluation contexts.

► **Lemma 12.** *The following assertions hold:*

1. *If $E \rightarrow_i^* E'$ and $E'[e_1] \lesssim e_2$ then $E[e_1] \lesssim e_2$.*
2. *If $E \rightarrow_i^* E'$ and $e_1 \lesssim E'[e_2]$ then $e_1 \lesssim E[e_2]$.*
3. *If $E_1 \rightarrow_i^* E'_1$ and $(E'_1, E_2) \in \mathcal{K}[[T_1; T_2]]$ then $(E_1, E_2) \in \mathcal{K}[[T_1; T_2]]$.*
4. *If $E_2 \rightarrow_i^* E'_2$ and $(E_1, E'_2) \in \mathcal{K}[[T_1; T_2]]$ then $(E_1, E_2) \in \mathcal{K}[[T_1; T_2]]$.*

The rest of the soundness proof follows the same lines as in Section 3.3. Interestingly, the adequacy lemma can be proved only for pure types, which is in harmony with the notion of contextual equivalence in the target calculus.

► **Theorem 13** (Fundamental property). *If $\Gamma \vdash e : T$ then $\Gamma; \Gamma \vdash e \lesssim_{log} e : T; T$.*

► **Lemma 14** (Precongruence). *If $\vdash C : (\Gamma; T) \rightsquigarrow \tau$ and $\Gamma; \Gamma \vdash e_1 \lesssim_{log} e_2 : T; T$, then $(C[e_1], C[e_2]) \in \mathcal{E}[[\tau; \tau]]$.*

► **Lemma 15** (Adequacy). *If $(e_1, e_2) \in \mathcal{E}[[\tau; \tau]]$ then $e_1 \lesssim e_2$.*

► **Theorem 16** (Soundness). *If $k \models \Gamma; \Gamma \vdash e_1 \lesssim_{log} e_2 : T; T$ holds for every k , then $\Gamma \vdash e_1 \lesssim_{ctx} e_2 : T$.*

4.4 Coherence of the CPS translation

Although standard compatibility lemmas and coercion compatibility suffice to prove soundness of logical relations, we need another kind of compatibility to prove coherence, since there is another source of ambiguity. Two typing derivations in the source language can be different not only because of the subsumption rule, but also because of two rules for application.

► **Lemma 17** (Mixed application compatibility). *The following assertions hold:*

1. *If $\Gamma_1; \Gamma_2 \vdash f_1 \lesssim_{log} f_2 : ((\tau'_1 \rightarrow (\tau_1 \rightarrow U_4) \Rightarrow U_3) \rightarrow U_2) \Rightarrow U_1; \tau'_2 \rightarrow T_2$
and $\Gamma_1; \Gamma_2 \vdash e_1 \lesssim_{log} e_2 : (\tau'_1 \rightarrow U_3) \Rightarrow U_2; \tau'_2$
then $\Gamma_1; \Gamma_2 \vdash \lambda k. f_1 (\lambda f. e_1 (\lambda x. f x k)) \lesssim_{log} f_2 e_2 : (\tau'_1 \rightarrow U_4) \Rightarrow U_1; T$.*
2. *If $\Gamma_1; \Gamma_2 \vdash f_1 \lesssim_{log} f_2 : \tau'_1 \rightarrow T_1; ((\tau'_2 \rightarrow (\tau_2 \rightarrow U_4) \Rightarrow U_3) \rightarrow U_2) \Rightarrow U_1$
and $\Gamma_1; \Gamma_2 \vdash e_1 \lesssim_{log} e_2 : \tau'_1; (\tau'_2 \rightarrow U_3) \Rightarrow U_2$
then $\Gamma_1; \Gamma_2 \vdash f_1 e_1 \lesssim_{log} \lambda k. f_2 (\lambda f. e_2 (\lambda x. f x k)) : T; (\tau'_2 \rightarrow U_4) \Rightarrow U_1$.*

Proof. Both cases are similar, so we show only the first one. We have to show that both terms closed by substitutions have the same observations in related contexts $(E_1, E_2) \in$

$\mathcal{K}[(\tau'_1 \rightarrow U_4) \Rightarrow U_1; T]$. Since context E_1 is in relation for effectful type, by the definition of logical relations and Lemma 12, it can be decomposed as a continuation k and the rest of the context. Now we have the missing continuation k that can trigger computation in the first term, so the rest of the proof consists in simple context manipulations, applying definitions and performing reductions. \blacktriangleleft

► **Lemma 18.** *If $D_i :: \Gamma_i \vdash e : T_i$ for $i = 1, 2$ are two typing judgments for the same term e of the source language, then $[\Gamma_1]; [\Gamma_2] \vdash \mathcal{T}[e]_{D_1} \lesssim_{log} \mathcal{T}[e]_{D_2} : [T_1]; [T_2]$.*

► **Theorem 19 (Coherence).** *If D_1 and D_2 are derivations of the same typing judgment $\Gamma \vdash e : T$, then $[\Gamma] \vdash \mathcal{T}[e]_{D_1} \lesssim_{ctx} \mathcal{T}[e]_{D_2} : [T]$.*

In contrast to the calculus considered in Section 3.4, such coherence theorem does not imply coherence of translation directly to the simply typed λ -calculus (where coercions are expressed as λ -terms). As a counterexample, the derivations from Example 10 give us terms that can be distinguished by the context $C = (\lambda x.1) []$. This is because types in the target language carry more information than simple types, and in particular, an expression of a type $(\tau \rightarrow T) \Rightarrow U$ is not a usual function, but a computation waiting for a continuation, as explained in Section 4.2.

But still we can prove some interesting properties of a direct translation to the simply typed λ -calculus in two cases: when control effects do not leak to the context or when we relate only whole programs. Let $|e|$ be a term e with all coercions replaced by corresponding expressions.

► **Corollary 20.** *If $D_1, D_2 :: \Gamma \vdash e : \tau$ and τ does not contain any type of the form $\tau'[T]U$, then $| \mathcal{T}[e]_{D_1} |$ and $| \mathcal{T}[e]_{D_2} |$ are contextually equivalent.*

► **Corollary 21.** *If $D_1, D_2 :: \Gamma \vdash e : \tau$ then $| \mathcal{T}[e]_{D_1} |$ terminates iff $| \mathcal{T}[e]_{D_2} |$ terminates. Moreover, if $\tau = \text{Nat}$ and one of the expressions terminates to a constant, then the other term evaluates to the same constant.*

5 Conclusion

We have shown that the technique of logical relations can be used for establishing the coherence of subtyping, when it is phrased in terms of contextual equivalence in the target of the coercion translation. In particular, we have demonstrated that a combination of heterogeneity, biorthogonality and step-indexing provides a sufficiently powerful tool for establishing coherence of effect subtyping in a calculus of delimited control with the coercion semantics given by a type-directed selective CPS translation. However, we have successfully applied the presented approach also to other calculi with subtyping, e.g., as demonstrated in this article for the simply-typed λ -calculus with recursion. The Coq development accompanying this paper is based on a new embedding of Dreyer et al.'s logic for reasoning about step-indexing that, we believe, considerably improves the presentation and formalization of the logical relations.

Regarding logical relations for type-and-effect systems, there has been a work on proving correctness of a partial evaluator for shift and reset by Asai [5], and on termination of evaluation of the λ -calculi with delimited-control operators by Biernacka et al. [6, 7] and by Materzok and the first author [16]. Unsurprisingly, all these results, like ours, are built on the notion of biorthogonality, even if not mentioned explicitly. The distinctive feature of our construction is a combination of heterogeneity and step-indexing that supports reasoning about the observational equivalence of terms of different types whose structure is very distant from each other, e.g., about direct-style and continuation-passing-style terms.

Acknowledgments. We thank Andrés A. Aristizábal, Małgorzata Biernacka, and the anonymous reviewers for helpful comments on the presentation of this work.

References

- 1 Amal Ahmed, Derek Dreyer, and Andreas Rossberg. State-dependent representation independence. In *POPL'09*, Savannah, GA, USA, 2009, pp. 340–353.
- 2 Amal J. Ahmed. Step-indexed syntactic logical relations for recursive and quantified types. In *ESOP'06*, Vienna, Austria, March 2006, *LNCS* 3924, pp. 69–83.
- 3 Andrew W. Appel and David McAllester. An indexed model of recursive types for foundational proof-carrying code. *ACM TOPLAS*, 23(5):657–683, 2001.
- 4 Andrew W. Appel, Paul-André Melliès, Christopher D. Richards, and Jérôme Vouillon. A very modal model of a modern, major, general type system. In *POPL'07*, Nice, France, 2007, pp. 109–122.
- 5 Kenichi Asai. Logical relations for call-by-value delimited continuations. In *Trends in Functional Programming 2005*, Tallinn, Estonia, 2005, pp. 413–428.
- 6 Małgorzata Biernacka and Dariusz Biernacki. Context-based proofs of termination for typed delimited-control operators. In *PPDP'09*, Coimbra, Portugal, 2009, pp. 289–300.
- 7 Małgorzata Biernacka, Dariusz Biernacki, and Sergueï Lenglet. Typing control operators in the CPS hierarchy. In *PPDP'11*, Odense, Denmark, 2011, pp. 149–160.
- 8 Val Breazu-Tannen, Thierry Coquand, Carl A. Gunter, and Andre Scedrov. Inheritance as implicit coercion. *Information and Computation*, 93(1):172–221, 1991.
- 9 Pierre-Louis Curien and Giorgio Ghelli. Coherence of subsumption, minimum typing and type-checking in F_{\leq} . *Mathematical Structures in Computer Science*, 2(1):55–91, 1992.
- 10 Olivier Danvy and Andrzej Filinski. A functional abstraction of typed contexts. DIKU Rapport 89/12, DIKU, University of Copenhagen, Copenhagen, Denmark, 1989.
- 11 Olivier Danvy and Andrzej Filinski. Abstracting control. In *Lisp and Functional Programming 1990*, Nice, France, 1990, pp. 151–160.
- 12 Derek Dreyer, Amal Ahmed, and Lars Birkedal. Logical step-indexed logical relations. *Logical Methods in Computer Science*, 7(2:16):1–37, 2011.
- 13 Derek Dreyer, Georg Neis, and Lars Birkedal. The impact of higher-order state and control effects on local relational reasoning. *Journal of Functional Programming*, 22(4-5):477–528, 2012.
- 14 Jean-Louis Krivine. Classical logic, storage operators and second-order lambda-calculus. *Annals of Pure and Applied Logic*, 68(1):53–78, 1994.
- 15 Marek Materzok. Axiomatizing subtyped delimited continuations. In *CSL'13*, Torino, Italy, 2013, *LIPICs* 23, pp. 521–539.
- 16 Marek Materzok and Dariusz Biernacki. Subtyping delimited continuations. In *ICFP'11*, Tokyo, Japan, 2011, pp. 81–93.
- 17 Marek Materzok and Dariusz Biernacki. A dynamic interpretation of the CPS hierarchy. In *APLAS'12*, *LNCS* 7705, Kyoto, Japan, 2012, pp. 296–311.
- 18 John C. Mitchell. *Foundations for Programming Languages*. MIT Press, 1996.
- 19 James H. Morris. *Lambda Calculus Models of Programming Languages*. PhD thesis, Massachusetts Institute of Technology, 1968.
- 20 Andrew Pitts and Ian Stark. Operational reasoning for functions with local state. In *Higher Order Operational Techniques in Semantics*, pp. 227–273. 1998.
- 21 Gordon D. Plotkin. Call-by-name, call-by-value and the λ -calculus. *Theoretical Computer Science*, 1:125–159, 1975.
- 22 John C. Reynolds. The coherence of languages with intersection types. In *Theoretical Aspects of Computer Software 1991*, Sendai, Japan, 1991, *LNCS* 526, pp. 675–700.

- 23 Tiark Rompf, Ingo Maier, and Martin Odersky. Implementing first-class polymorphic delimited continuations by a type-directed selective CPS-transform. In *ICFP'09*, Edinburgh, UK, 2009, pp. 317–328.
- 24 Jan Schwinghammer. Coherence of subsumption for monadic types. *Journal of Functional Programming*, 19(2):157–172, 2009.

Observability for Pair Pattern Calculi

Antonio Bucciarelli¹, Delia Kesner¹, and
Simona Ronchi Della Rocca²

1 Univ Paris Diderot, Sorbonne Paris Cité, PPS, UMR 7126, CNRS, Paris,
France

2 Dipartimento di Informatica, Università di Torino, Italy

Abstract

Inspired by the notion of solvability in the λ -calculus, we define a notion of observability for a calculus with pattern matching. We give an intersection type system for such a calculus which is based on non-idempotent types. The typing system is shown to characterize the set of terms having canonical form, which properly contains the set of observable terms, so that typability alone is not sufficient to characterize observability. However, the inhabitation problem associated with our typing system turns out to be decidable, a result which – together with typability – allows to obtain a full characterization of observability.

1998 ACM Subject Classification F.4.1 Lambda calculus and related systems, F.3.2 Operational Semantics, F.4.1 Proof theory

Keywords and phrases Solvability, pattern calculi, intersection types, inhabitation

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.123

1 Introduction

In these last years there has been a growing interest in *pattern λ -calculi* [16, 11, 6, 12, 10, 15] which are used to model the pattern-matching primitives of functional programming languages (*e.g.* OCAML, ML, Haskell) and proof assistants (*e.g.* Coq, Isabelle). These calculi are extensions of λ -calculus, where abstractions are written as $\lambda p.t$, where p is a *pattern* specifying the expected structure of the argument. In this paper we restrict our attention to *pair* patterns, which are expressive enough to illustrate the challenging notion of solvability/observability in the framework of pattern λ -calculi.

In order to implement different *evaluation strategies*, the use of *explicit pattern-matching* becomes appropriate, giving rise to different languages with explicit pattern-matching [6, 7, 1]. In all of them, an application $(\lambda p.t)u$ reduces to $t[p/u]$, where $[p/u]$ is an explicit matching, defined by means of suitable reduction rules, which are used to decide if the argument u matches the pattern p . If the matching is possible, the evaluation proceeds by computing a substitution which is applied to the body t . Otherwise, two cases arise: either a successful matching is not possible at all, and then the term $t[p/u]$ reduces to a *failure*, denoted by the constant `fail`, or it could become possible after the application of some pertinent substitution to the argument u , in which case the reduction is simply *blocked*. An example of failure is caused by the term $(\lambda\langle z_1, z_2 \rangle.z_1)(\lambda y.y)$, while a blocked reduction is caused by the term $(\lambda\langle z_1, z_2 \rangle.z_1)y$.

Inspired by the notion of solvability in the λ -calculus, we define a notion of *observability* for a pair pattern calculus with explicit matching. A term t is said to be observable if there is a *head-context* C such that $C[t]$ reduces to a pair, which is the only data structure of the language. This notion is conservative with respect to the notion of solvability in the λ -calculus, *i.e.* t is solvable in the λ -calculus if and only if t is observable in our calculus.



© Antonio Bucciarelli, Delia Kesner, and Simona Ronchi Della Rocca;
licensed under Creative Commons License CC-BY

13th International Conference on Typed Lambda Calculi and Applications (TLCA'15).

Editor: Thorsten Altenkirch; pp. 123–137



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Solvability in the λ -calculus is of course undecidable, but it has been characterized at least in three different ways: syntactically by the notion of head-normal form [2], operationally by the notion of head-reduction [2], and logically by an intersection type assignment system [3, 13]. The problem becomes harder when changing from the call-by-name to the call-by-value setting. Indeed, in the call-by-value λ -calculus, there are normal forms that are unsolvable, like the term $(\lambda z.\Delta)(xI)\Delta$, where $\Delta = \lambda x.xx$. The problem for the pair pattern calculus is similar to that for the call-by-value, but even harder. As in the call-by-value setting, an argument needs to be partially evaluated before being consumed. Indeed, in order to evaluate an application $(\lambda p.t)u$, it is necessary to verify if u matches the pattern p , and thus the subterm u can be forced to be partially evaluated. However, while only discrimination between values and non-values are needed in the call-by-value setting, the possible shapes of patterns are infinite here.

The difficulty of the problem depends on two facts. First, there is no simple *syntactical* characterization of observability: indeed, we supply a notion of *canonical form* such that reducing to some canonical form is a necessary condition for being observable. But this is not sufficient: canonical forms may contain blocking explicit matchings, so that we need to know whether or not there exists a substitution being able to unblock simultaneously all these blocked forms.

This theoretical complexity is reflected in the logical characterization we supply for observability: a term t turns out to be observable if and only if it is typable, say with a type of the shape $A_1 \rightarrow A_2 \rightarrow \dots \rightarrow A_n \rightarrow \alpha$ (where α is a product type), and all the types A_i ($1 \leq i \leq n$) are *inhabited*. The inhabitation problem for idempotent intersection types is known to be undecidable [17], but it has recently been proved that it is decidable in the non-idempotent case [5]. More precisely, there is a sound and complete algorithm solving the inhabitation problem of non-idempotent intersection types for the λ -calculus. In this paper, we supply a type assignment system, based on non-idempotent intersection, which assigns types to terms of our pair pattern calculus. We then extend the inhabitation algorithm given in [5] to this framework, that is substantially more complicated, due to the explicit pattern matching and the use of structural information of patterns in the typing rules. However, the paper does not only show decidability of inhabitation for the pair pattern calculus, but it *uses* the decidability result to derive a full characterization of observability, which is the main result of the paper. We thus combine typability with inhabitation in order to obtain an interesting characterization of the set of *meaningful* terms of the pair pattern calculus.

The paper is organized as follows. Sec. 2 introduces the pattern calculus. Sec. 3 presents the type system and proves a characterization of terms having canonical forms by means of typability. Sec. 4 discusses the relationship between observability and inhabitation and Sec. 5 presents a sound and complete algorithm for the inhabitation problem associated to our typing system. Sec. 6 shows a complete characterization of observability, and Sec. 7 concludes by discussing some future work.

2 The Pair Pattern Calculus

We now introduce the Λ_p -calculus, a generalization of the λ -calculus where abstraction is extended to *patterns* and terms to pairs. Pattern matching is specified by means of an *explicit* operation. Reduction is performed only if the argument matches the abstracted pattern.

Terms and contexts of the Λ_p -calculus are defined as follows:

$$\begin{array}{lll}
\text{(Patterns)} & p, q & ::= x \mid \langle p, p \rangle \\
\text{(Terms)} & t, u, v & ::= x \mid \lambda p. t \mid \langle t, t \rangle \mid tt \mid t[p/t] \mid \text{fail} \\
\text{(Contexts)} & C & ::= \square \mid \lambda p. C \mid \langle C, t \rangle \mid \langle t, C \rangle \mid Ct \mid tC
\end{array}$$

where x, y, z range over a countable set of variables, and every pattern p is *linear*, *i.e.* every variable appears at most once in p . We denote by Id the **identity function** $\lambda x. x$. As usual we use the abbreviation $\lambda p_1 \dots p_n. t_1 \dots t_m$ for $\lambda p_1 (\dots (\lambda p_n ((t_1 t_2) \dots t_m)) \dots)$, $n \geq 0$, $m \geq 1$. Remark that every λ -term is a Λ_p -term.

The operator $[p/t]$ is called an **explicit matching**. The constant **fail** denotes the failure of the matching operation. **Free** and **bound variables** of terms are defined as expected, in particular $\text{fv}(\lambda p. t) := \text{fv}(t) \setminus \text{fv}(p)$ and $\text{fv}(t[p/u]) := (\text{fv}(t) \setminus \text{fv}(p)) \cup \text{fv}(u)$. We write $p \# q$ iff $\text{fv}(p)$ and $\text{fv}(q)$ are disjoint. As usual, terms are considered modulo α -conversion. Given a context C and a term t , $C[t]$ denotes the term obtained by replacing the unique occurrence of \square in C by t , allowing the capture of free variables of t . A **head-context** is a context of the shape $(\lambda p_1 \dots p_n. \square) t_1 \dots t_m$ ($n, m \geq 0$).

The **reduction relation** of the Λ_p -calculus, denoted by \rightarrow , is the contextual closure of the following reduction rules:

$$\begin{array}{llll}
(r_1) & (\lambda p. t)u & \mapsto & t[p/u] & (r_6) & t[\langle p_1, p_2 \rangle / \lambda y. u] & \mapsto & \text{fail} \\
(r_2) & t\{x/u\} & \mapsto & t\{x/u\} & (r_7) & t[\langle p_1, p_2 \rangle / \text{fail}] & \mapsto & \text{fail} \\
(r_3) & t[\langle p_1, p_2 \rangle / \langle u_1, u_2 \rangle] & \mapsto & t[p_1/u_1][p_2/u_2] & (r_8) & \text{fail } t & \mapsto & \text{fail} \\
(r_4) & t[p/v]u & \mapsto & (tu)[p/v] & (r_9) & \text{fail}[p/t] & \mapsto & \text{fail} \\
(r_5) & t[\langle p_1, p_2 \rangle / u[q/v]] & \mapsto & t[\langle p_1, p_2 \rangle / u][q/v] & (r_{10}) & \lambda p. \text{fail} & \mapsto & \text{fail} \\
& & & & (r_{11}) & \langle t, u \rangle v & \mapsto & \text{fail}
\end{array}$$

where $t\{x/u\}$ denotes the substitution of all the free occurrences of x in t by u . By α -conversion, and without loss of generality, no reduction rule captures free variables. Thus for example in rule r_4 the bound and free variables of the term $t[p/v]u$ are supposed to be disjoint, so that the variables of p (which are bound in the whole term) cannot be free in u . The reflexive and transitive closure of \rightarrow is written \rightarrow^* .

The rule (r_1) triggers the pattern operation while rule (r_2) performs substitution, rules (r_3) , (r_6) and (r_7) implement (successful or unsuccessful) pattern matching. Rules (r_8) , (r_9) and (r_{10}) deal with propagation of failure. Rules (r_4) and (r_5) may seem unnecessary, and the calculus would be also confluent without them, but they are particularly useful for the design of the inhabitation algorithm (see Sec. 5). Indeed, rule (r_4) pushes *head* explicit matchings out, and rule (r_5) eliminates *nested* explicit matchings, *i.e.* matchings of the form $t[\langle p_1, p_2 \rangle / u[q/v]]$. Notice that confluence would be lost if we allow (r_5) on the more general form: $t[p/u[q/v]] \mapsto t[p/u][q/v]$. Indeed, the following critical pair could not be closed: $y[\langle z_1, z_2 \rangle / z] \xrightarrow{r_5, r_2^*} y[x/u[\langle z_1, z_2 \rangle / z]] \xrightarrow{r_2} y$.

► **Lemma 1.**

1. *The reduction relation \rightarrow is confluent.*
2. *Every infinite \rightarrow -reduction sequence contains an infinite number of \rightarrow_{r_2} -reduction steps.*

The proof of the first item relies on the decreasing diagram technique [18]; that of the second one is by induction on a suitable syntactic measure.

Canonical forms are terms defined by the following grammar:

$$\mathcal{J} ::= \lambda p. \mathcal{J} \mid \langle t, t \rangle \mid \mathcal{K} \mid \mathcal{J}[\langle p, q \rangle / \mathcal{K}] \qquad \mathcal{K} ::= x \mid \mathcal{K}t$$

A term \mathfrak{t} is in **canonical form** (or it is **canonical**), written *cf*, if it is generated by \mathcal{J} , and it **has a canonical form** if it reduces to a term in *cf*. Note that the *cf* of a term is not unique, *e.g.* both $\langle \text{Id}, \text{Id Id} \rangle$ and $\langle \text{Id}, \text{Id} \rangle$ are *cfs* of $(\lambda xy. \langle x, y \rangle) \text{Id} (\text{Id Id})$. It is worth noticing that *cfs* and normal forms do not coincide. For example, the terms $\lambda \langle x, y \rangle. \langle x(\Delta\Delta) \rangle [\langle z_1, z_2 \rangle / y \text{Id}]$ and $\langle \text{Id}, \text{Id Id} \rangle$ are in *cf*, but not in normal form, while **fail** is in normal form but not in *cf*. Every head normal-form in the λ -calculus is a *cf* in the $\Lambda_{\mathfrak{p}}$ -calculus.

On the pathway towards the definition of an adequate notion of *solvability* for the $\Lambda_{\mathfrak{p}}$ -calculus, we first recall the notion of solvability for the λ -calculus. A term \mathfrak{t} is solvable iff there is a head-context \mathbf{C} such that $\mathbf{C}[\mathfrak{t}]$ reduces to **Id**. It is clear that pairs have to be taken into account in order to extend the notion of solvability to the pair pattern calculus. When should a pair be considered as meaningful? At least two choices are possible: the *lazy* semantics considers a pair as meaningful in itself, the *strict* one requires both of its components to be meaningful. The first choice is adopted in this paper, since being a pair is already an observable property, particularly sufficient to unblock an explicit matching, independently from the observability of its components.

Thus, a term \mathfrak{t} is said to be **observable** iff there is a head-context \mathbf{C} such that $\mathbf{C}[\mathfrak{t}]$ reduces to a pair, *i.e.* $\mathbf{C}[\mathfrak{t}] \rightarrow^* \langle \mathfrak{t}_1, \mathfrak{t}_2 \rangle$, for some terms $\mathfrak{t}_1, \mathfrak{t}_2 \in \Lambda_{\mathfrak{p}}$. Thus for example, the term $\langle \Delta\Delta, \Delta\Delta \rangle$, consisting of a pair of unsolvable terms $\Delta\Delta$, is observable. This notion of observability turns out to be conservative with respect to that of solvability for the λ -calculus (see Theorem 23).

3 The Type System \mathcal{P}

In this section we present a type system for the $\Lambda_{\mathfrak{p}}$ -calculus, and we show that it characterizes terms having canonical form.

The set \mathcal{T} of types is generated by the following grammar:

$$\begin{array}{ll} \alpha & ::= o \mid \times_1(\tau) \mid \times_2(\tau) \quad (\text{product types}) \\ \sigma, \tau, \pi, \rho & ::= \alpha \mid \mathbf{A} \rightarrow \sigma \quad (\text{strict types}) \\ \mathbf{B} & ::= [\sigma_i]_{i \in I} \quad (I \neq \emptyset) \quad (\text{non-empty multiset types}) \\ \mathbf{A} & ::= [] \mid \mathbf{B} \quad (\text{multiset types}) \end{array}$$

where I is a finite set of indices. The arrow constructor is right associative. We consider a unique type constant o , which can be assigned to any pair.

We write $\text{supp}(\mathbf{A})$ to denote the *support set* of the multiset \mathbf{A} , \sqcup for multiset union and \in to denote multiset membership. The **product** operation \mathbb{X} on multisets is defined as follows:

$$\begin{array}{ll} [] & \mathbb{X} \quad [] & := & [o] \\ [\sigma_i]_{i \in I} & \mathbb{X} \quad [\rho_j]_{j \in J} & := & [\times_1(\sigma_i)]_{i \in I} \sqcup [\times_2(\rho_j)]_{j \in J} \quad \text{if } I \neq \emptyset \text{ or } J \neq \emptyset \end{array}$$

Remark that $\sqcup_{i \in I} \mathbf{A}_i \mathbb{X} \sqcup_{i \in I} \mathbf{A}'_i \sqsubseteq \sqcup_{i \in I} (\mathbf{A}_i \mathbb{X} \mathbf{A}'_i)$, the multiset inclusion being strict for example in the following case: $([] \sqcup []) \mathbb{X} ([] \sqcup []) = [o] \sqsubset [o, o] = ([] \mathbb{X} []) \sqcup ([] \mathbb{X} [])$.

The **structure** of a pattern describes its shape, it is defined as follows:

$$\begin{array}{ll} \mathcal{S}(x) & := \quad [] \\ \mathcal{S}(\langle p_1, p_2 \rangle) & := \quad \mathcal{S}(p_1) \mathbb{X} \mathcal{S}(p_2) \end{array}$$

E.g. $\mathcal{S}(\langle x, y \rangle) = [o]$, $\mathcal{S}(\langle x, \langle y, z \rangle \rangle) = [\times_2(o)]$ and $\mathcal{S}(\langle \langle x, w \rangle, \langle y, z \rangle \rangle) = [\times_1(o), \times_2(o)]$. Notice that $\mathcal{S}(\mathfrak{p})$ is nothing but a description of \mathfrak{p} seen as a binary tree whose leaves are distinct variables, and whose nodes are labeled by the pair constructor. Indeed, each element of $\mathcal{S}(\mathfrak{p})$ specifies a maximal branch of such a tree, *i.e.* a branch whose last node is a pair constructor, and whose children are both leaves (*i.e.* variables). $\mathcal{S}(\mathfrak{p})$ should be understood as the multiset

$$\begin{array}{c}
\frac{}{x : B \Vdash x : B} \text{ (varpat)} \quad \frac{}{\Vdash x : []} \text{ (weakpat)} \quad \frac{\Gamma \Vdash p : A_1 \quad \Delta \Vdash q : A_2 \quad p \# q}{\Gamma + \Delta \Vdash \langle p, q \rangle : A_1 \times A_2} \text{ (pairpat)} \\
\\
\frac{}{x : [\pi] \vdash x : \pi} \text{ (var)} \quad \frac{\Gamma \vdash t : \pi \quad \Gamma|_p \Vdash p : [\sigma_i]_{i \in I} \quad (\sigma_j \in \mathcal{S}(p))_{j \in J} \quad I \cap J = \emptyset}{\Gamma \setminus \Gamma|_p \vdash \lambda p. t : [\sigma_k]_{k \in I \cup J} \rightarrow \pi} (\rightarrow i) \\
\\
\frac{\Gamma \vdash t : [\sigma_i]_{i \in I} \rightarrow \pi \quad (\Delta_i \vdash u : \sigma_i)_{i \in I}}{\Gamma +_{i \in I} \Delta_i \vdash tu : \pi} (\rightarrow e) \\
\\
\frac{}{\vdash \langle t, u \rangle : o} \text{ (emptypair)} \quad \frac{\Gamma \vdash t : \sigma}{\Gamma \vdash \langle t, u \rangle : \times_1(\sigma)} \text{ (pair1)} \quad \frac{\Gamma \vdash u : \tau}{\Gamma \vdash \langle t, u \rangle : \times_2(\tau)} \text{ (pair2)} \\
\\
\frac{\Gamma \vdash t : \sigma \quad \Gamma|_p \Vdash p : [\sigma_i]_{i \in I} \quad (\sigma_j \in \mathcal{S}(p))_{j \in J} \quad (\Delta_k \vdash u : \sigma_k)_{k \in I \cup J} \quad I \cap J = \emptyset}{(\Gamma \setminus \Gamma|_p) +_{k \in I \cup J} \Delta_k \vdash t[p/u] : \sigma} \text{ (sub)}
\end{array}$$

■ **Figure 1** The type assignment system \mathcal{P} .

of *non depletable* resources associated with p ; the persistent character of these resources is highlighted in the forthcoming typing system.

Typing environments, written Γ, Δ , are functions from variables to multiset types, assigning the empty multiset to almost all the variables. The **domain** of Γ , written $\text{dom}(\Gamma)$, is the set of variables whose image is different from $[]$. We write $\Gamma \# \Delta$ iff $\text{dom}(\Gamma) \cap \text{dom}(\Delta) = \emptyset$.

► **Notation 2.** Given the environments $\{\Gamma_i\}_{i \in I}$, we write $+_{i \in I} \Gamma_i$ for the environment which maps x to $\sqcup_{i \in I} \Gamma_i(x)$. If $I = \emptyset$, the resulting environment is the one having an empty domain. Note that $\Gamma + \Delta$ and $\Gamma +_{i \in I} \Delta_i$ are just particular cases of the previous general definition. When $\Gamma \# \Delta$ we write $\Gamma; \Delta$ instead of $\Gamma + \Delta$. We write $\Gamma \setminus x$ for the environment assigning $[]$ to x , and acting as Γ otherwise; $x_1 : A_1; \dots; x_n : A_n$ is the environment assigning A_i to x_i , for $1 \leq i \leq n$, and $[]$ to any other variable; $\Gamma|_p$ denotes the environment such that $\Gamma|_p(x) = \Gamma(x)$, if $x \in \text{fv}(p)$, $[]$ otherwise.

The **type assignment system** \mathcal{P} (see Fig. 1) is a set of typing rules assigning strict types of \mathcal{T} to terms of Λ_p . We write $\Pi \triangleright \Gamma \vdash t : \sigma$ (resp. $\Pi \triangleright \Gamma \Vdash p : A$) to denote a **typing derivation** ending in the sequent $\Gamma \vdash t : \sigma$ (resp. $\Gamma \Vdash p : A$), in which case t (resp. p) is called the **subject** of Π ; by abuse of notation, $\Gamma \vdash t : \sigma$ (resp. $\Gamma \Vdash p : A$) also denotes the existence of some typing derivation ending in this sequent, in which case t (resp. p) is said to be **typable**. The **measure** of a typing derivation Π , written $\text{meas}(\Pi)$, is the number of typing rules in Π .

Rules (var) and $(\rightarrow e)$ are those used for λ -calculus in [5, 8]. Linearity of patterns is guaranteed by the clause $p \# q$ in rule (pairpat). Rule (weakpat) is essential to type *erasing* functions such as for example $\lambda x. \text{Id}$. The rule (emptypair) types for example $\langle \Delta \Delta, \Delta \Delta \rangle$, and thus $(\lambda \langle x, y \rangle. \text{Id}) \langle \Delta \Delta, \Delta \Delta \rangle$. Rules (pair1) and (pair2) type pairs having just one typed component, whereas standard typed calculi with pairs (e.g. [6]) requires both components to be typed. This is necessary to type terms like $(\lambda \langle x, y \rangle. x) \langle \text{Id}, \Delta \Delta \rangle$. Moreover, the standard policy can be easily recovered from ours by typing a pair whose components are both typed using (pair1) and (pair2) successively.

The rules $(\rightarrow \mathfrak{i})$ and (sub) are the most subtle ones¹. Here is where the structural types come into play: they can be used *ad libitum* (whence the notation $(\sigma_j \in \mathcal{S}(\mathfrak{p}))_{j \in J}$), thanks to *non depletable* nature of the information provided by the structure of patterns (whereas the type information of variables should be understood as *depletable*). Concerning more specifically the rule (sub) : in order to type $\mathfrak{t}[\mathfrak{p}/\mathfrak{u}]$, on one hand we need to type \mathfrak{t} and on the other one we need to check that \mathfrak{p} and \mathfrak{u} can be assigned the same types. Since the system is relevant, we need to collect the environments used in all the premises typing \mathfrak{p} and \mathfrak{u} . Remark however that there is a lack of symmetry between patterns and terms: while the only information we can use about terms is the one concerning their types, a pattern \mathfrak{p} has not only a type (description of its depletable resources), but also an intrinsic shape that is completely described by the structural (non depletable) types in the set $\mathcal{S}(\mathfrak{p})$.

Actually the structural information on patterns is necessary, in particular, to guarantee subject reduction for rule (r_5) . Indeed, given $\mathfrak{t} = \lambda \mathfrak{w}.(\mathfrak{z}\mathfrak{z}')[\langle \mathfrak{z}, \mathfrak{z}' \rangle / (\mathfrak{y}\mathfrak{x})][\langle \mathfrak{x}, \mathfrak{x}' \rangle / \mathfrak{w}] \rightarrow_{r_5} \lambda \mathfrak{w}.(\mathfrak{z}\mathfrak{z}')[\langle \mathfrak{z}, \mathfrak{z}' \rangle / (\mathfrak{y}\mathfrak{x})][\langle \mathfrak{x}, \mathfrak{x}' \rangle / \mathfrak{w}] = \mathfrak{t}'$, and $\Gamma = \mathfrak{y} : [[] \rightarrow \times_1(\tau), [\pi] \rightarrow \times_2(\sigma)]$, we have that $\Gamma \vdash \mathfrak{t} : [o, \times_1(\pi)] \rightarrow \sigma$, but $\Gamma \vdash \mathfrak{t}' : [o, \times_1(\pi)] \rightarrow \sigma$ holds only by using the fact that $o \in \mathcal{S}(\langle \mathfrak{x}, \mathfrak{x}' \rangle)$. This counterexample shows that a clear tension appears between the rewriting rule (r_5) and the use of the structural set $\mathcal{S}(\mathfrak{p})$ in the typing rules $(\rightarrow \mathfrak{i})$ and (sub) . Eliminating (r_5) from the reduction system would certainly simplify the typing system, but would significantly complicate the inhabitation algorithm that will be presented in Sec. 5.

► **Example 3.** The following (partially described) derivation is valid:

$$\frac{\begin{array}{l} (a) \mathfrak{x} : [\alpha] \vdash \mathfrak{x} : \alpha \quad (b) \mathfrak{x} : [\alpha] \Vdash \langle \mathfrak{x}, \mathfrak{y} \rangle : [\times_1(\alpha)] \\ (c) (o \in \mathcal{S}(\langle \mathfrak{x}, \mathfrak{y} \rangle)) \quad (d) \mathfrak{z} : [o] \vdash \mathfrak{z} : o \quad (e) \mathfrak{z} : [\times_1(\alpha)] \vdash \mathfrak{z} : \times_1(\alpha) \end{array}}{\mathfrak{z} : [o, \times_1(\alpha)] \vdash \mathfrak{x}[\langle \mathfrak{x}, \mathfrak{y} \rangle / \mathfrak{z}] : \alpha} \text{ (sub)}$$

Using only the hypothesis (a), (b) and (e) we get another valid typing derivation ending in $\mathfrak{z} : [\times_1(\alpha)] \vdash \mathfrak{x}[\langle \mathfrak{x}, \mathfrak{y} \rangle / \mathfrak{z}] : \alpha$ which does not use structural information about the pattern $\langle \mathfrak{x}, \mathfrak{y} \rangle$.

The system is relevant, in the sense that only the used premises are registered in the typing environments. This property, formally stated in the following lemma, will be an important technical tool used to develop the inhabitation algorithm.

► **Lemma 4 (Relevance).**

- If $\Gamma \Vdash \mathfrak{p} : \mathfrak{A}$, then $\text{dom}(\Gamma) \subseteq \text{fv}(\mathfrak{p})$.
- If $\Gamma \vdash \mathfrak{t} : \sigma$, then $\text{dom}(\Gamma) \subseteq \text{fv}(\mathfrak{t})$.

Proof. By induction on the typing derivations. ◀

Some useful properties will be needed in the sequel. In particular, the next technical lemma says that, given different types \mathfrak{A}_i for a given pattern \mathfrak{p} , it is always possible to split $\sqcup_{i \in I} \mathfrak{A}_i$ into a bunch of resource types \mathfrak{A} and another one of structural types \mathfrak{A}' .

► **Lemma 5.** Let $I \neq \emptyset$. If $(\Gamma_i \Vdash \mathfrak{p} : \mathfrak{A}_i)_{i \in I}$, then there exist $\mathfrak{A}, \mathfrak{A}'$ such that

1. $\mathfrak{A} \sqcup \mathfrak{A}' = \sqcup_{i \in I} \mathfrak{A}_i$,
2. $+_{i \in I} \Gamma_i \Vdash \mathfrak{p} : \mathfrak{A}$

¹ Notice that “ Γ ”, “ $\Gamma|_{\mathfrak{p}}$ ” and “ $\Gamma \setminus \Gamma|_{\mathfrak{p}}$ ” in rules $(\rightarrow \mathfrak{i})$ and (sub) could be replaced by “ $\Gamma_1; \Gamma_2$ ”, “ Γ_2 ” and “ Γ_1 ”, respectively, only if $\text{dom}(\Gamma_1) \cap \text{fv}(\mathfrak{p}) = \emptyset$. Otherwise, for instance, $\lambda \mathfrak{x}. \mathfrak{x}$ would be typable with type $[] \rightarrow \sigma$.

3. $A = []$ implies $A' = []$,
4. $\text{supp}(A') \subseteq \mathcal{S}(p)$,
5. $\text{meas}(+_{i \in I} \Gamma_i \Vdash p : A) \leq \sum_{i \in I} \text{meas}(\Gamma_i \Vdash p : A_i)$.

Proof. By induction on p . ◀

The following lemma can be shown by induction on typing derivations; it is used in the forthcoming subject reduction property.

► **Lemma 6** (Substitution Lemma). *If $\Pi \triangleright \Gamma; \mathbf{x} : [\rho_i]_{i \in I} \vdash \mathbf{t} : \tau$, and $(\Theta_i \triangleright \Delta_i \vdash \mathbf{u} : \rho_i)_{i \in I}$ then $\Pi' \triangleright \Gamma +_{i \in I} \Delta_i \vdash \mathbf{t}\{\mathbf{x}/\mathbf{u}\} : \tau$ where $\text{meas}(\Pi') < \text{meas}(\Pi) + \sum_{i \in I} \text{meas}(\Theta_i)$.*

Notice that, in the process of assigning a type to a term \mathbf{t} , some subterms of \mathbf{t} may be left untyped. Typically, this happens when \mathbf{t} contains occurrences of non typable terms, like in $\lambda \mathbf{x}. \mathbf{x}(\Delta \Delta)$. We are then going to define the notion of **typed occurrence** of a typing derivation, which plays an essential role in the rest of this paper: indeed, thanks to the use of non-idempotent intersection types, a combinatorial argument based on a measure on typing derivations (*cf.* Lem. 9.1), allows to prove the termination of reduction of redexes occurring in typed occurrences of their respective typing derivations.

Let us then define an **occurrence** of a subterm \mathbf{u} in a term \mathbf{t} as a context \mathbf{C} such that $\mathbf{C}[\mathbf{u}] = \mathbf{t}$. Then, given a typing derivation $\Pi \triangleright \Gamma \vdash \mathbf{t} : \sigma$, an occurrence of a subterm of \mathbf{t} is a typed occurrence of Π if and only if it is the subject of a subderivation of Π . More precisely:

► **Definition 7.** Given a type derivation Π , the set of **typed occurrences** of Π , written $\text{toc}(\Pi)$, by induction on the last rule of Π .

- If Π ends with (**var**), then $\text{toc}(\Pi) := \{\square\}$.
- If Π ends with (**pair1**) with subject $\langle \mathbf{u}, \mathbf{v} \rangle$ and premise Π' , then $\text{toc}(\Pi) := \{\square\} \cup \{\langle \mathbf{C}, \mathbf{v} \rangle \mid \mathbf{C} \in \text{toc}(\Pi')\}$.
- If Π ends with (**pair2**) with subject $\langle \mathbf{u}, \mathbf{v} \rangle$ and premise Π' then $\text{toc}(\Pi) := \{\square\} \cup \{\langle \mathbf{u}, \mathbf{C} \rangle \mid \mathbf{C} \in \text{toc}(\Pi')\}$.
- If Π ends with (\rightarrow **i**) with subject $\lambda \mathbf{p}. \mathbf{u}$ and premise Π' then $\text{toc}(\Pi) := \{\square\} \cup \{\lambda \mathbf{p}. \mathbf{C} \mid \mathbf{C} \in \text{toc}(\Pi')\}$.
- If Π ends with (\rightarrow **e**) with subject $\mathbf{t}\mathbf{u}$ and premises Π_1 and Π_k ($k \in K$) with subjects \mathbf{t} and \mathbf{u} respectively, then $\text{toc}(\Pi) := \{\square\} \cup \{\mathbf{t}\mathbf{C} \mid \mathbf{C} \in \text{toc}(\Pi_k), k \in K\} \cup \{\mathbf{C}\mathbf{u} \mid \mathbf{C} \in \text{toc}(\Pi_1)\}$.
- If Π ends with (**sub**) with subject $\mathbf{t}[\mathbf{p}/\mathbf{u}]$ and premises Π_1 and Π_k ($k \in K$) with subjects \mathbf{t} and \mathbf{u} respectively, then $\text{toc}(\Pi) := \{\square\} \cup \{\mathbf{C}[\mathbf{p}/\mathbf{u}] \mid \mathbf{C} \in \text{toc}(\Pi_1)\} \cup \{\mathbf{t}[\mathbf{p}/\mathbf{C}] \mid \mathbf{C} \in \text{toc}(\Pi_k), k \in K\}$.

► **Example 8.** Given the following derivations Π and Π' , the occurrences \square and $\square \mathbf{y}$ belong to both $\text{toc}(\Pi)$ and $\text{toc}(\Pi')$ while $\mathbf{x}\square$ belongs to $\text{toc}(\Pi)$ but not to $\text{toc}(\Pi')$.

$$\Pi \triangleright \frac{\mathbf{x} : [[\tau] \rightarrow \tau] \vdash \mathbf{x} : [\tau] \rightarrow \tau \quad \mathbf{y} : [\tau] \vdash \mathbf{y} : \tau}{\mathbf{x} : [[\tau] \rightarrow \tau], \mathbf{y} : [\tau] \vdash \mathbf{x}\mathbf{y} : \tau} \quad \Pi' \triangleright \frac{\mathbf{x} : [[\] \rightarrow \tau] \vdash \mathbf{x} : [\] \rightarrow \tau}{\mathbf{x} : [[\] \rightarrow \tau] \vdash \mathbf{x}\mathbf{y} : \tau}$$

Given $\Pi \triangleright \Gamma \vdash \mathbf{t} : \tau$, \mathbf{t} is said to be in **Π -normal form**, also written **Π -nf**, if for every typed occurrence $\mathbf{C} \in \text{toc}(\Pi)$ such that $\mathbf{t} = \mathbf{C}[\mathbf{u}]$, the subterm \mathbf{u} is not a redex.

The system \mathcal{P} enjoys both subject reduction and subject expansion. In particular, thanks to the use of multisets, subject reduction decreases the measure of the derivation, in case a substitution is performed by rule (r_2) and the redex is typed. This property allows for a simple proof of the “only if” part of the characterization theorem.

► **Lemma 9.**

1. **(Weighted Subject Reduction)** If $\Pi \triangleright \Gamma \vdash \mathfrak{t} : \tau$ and $\mathfrak{t} \rightarrow \mathfrak{v}$, then $\Pi' \triangleright \Gamma \vdash \mathfrak{v} : \tau$ and $\text{meas}(\Pi') \leq \text{meas}(\Pi)$. Moreover, if the reduced redex is (r_2) and it occurs in a typed occurrence of Π , then $\text{meas}(\Pi') < \text{meas}(\Pi)$.
2. **(Subject Expansion)** If $\Gamma \vdash \mathfrak{v} : \sigma$ and $\mathfrak{t} \rightarrow \mathfrak{v}$, then $\Gamma \vdash \mathfrak{t} : \sigma$.

Proof. 1. By induction on $\mathfrak{t} \rightarrow \mathfrak{v}$ using Lemmas 5, 6 and 4.

2. By induction on $\mathfrak{t} \rightarrow \mathfrak{v}$. ◀

We are now ready to provide the logical characterization of terms having canonical form.

► **Theorem 10 (Characterization).** *A term \mathfrak{t} is typable iff \mathfrak{t} has a canonical form.*

Proof. ■ (if) We reason by induction on the grammar defining the canonical forms. We first prove that for all type σ and for all \mathcal{K} -canonical form \mathfrak{t} , \mathfrak{t} can be typed by σ . In fact every \mathcal{K} -canonical form is of the shape $\mathfrak{x}\mathfrak{t}_1 \dots \mathfrak{t}_n$, for $n \geq 0$. It is easy to check that $\mathfrak{x} : [\] \rightarrow \dots \rightarrow [\] \rightarrow \sigma \vdash \mathfrak{x}\mathfrak{t}_1 \dots \mathfrak{t}_n : \sigma$. Let \mathfrak{t} be a \mathcal{J} -canonical form. If $\mathfrak{t} = \langle \mathfrak{u}, \mathfrak{v} \rangle$ then by

rule (**emptypair**) $\vdash \langle \mathfrak{u}, \mathfrak{v} \rangle : \sigma$. If $\mathfrak{t} = \lambda \mathfrak{p}. \mathfrak{u}$, then by induction \mathfrak{u} can be typed and the result follows from rule (\rightarrow I). Let $\mathfrak{t} = \mathfrak{t}'[\langle \mathfrak{p}, \mathfrak{q} \rangle / \mathfrak{v}]$, where \mathfrak{t}' (resp. \mathfrak{v}) is a \mathcal{J} (resp. \mathcal{K}) canonical form. By the *i.h.* there are Γ, σ such that $\Gamma \vdash \mathfrak{t}' : \sigma$. Moreover, it is easy to see that $\Gamma|_{\langle \mathfrak{p}, \mathfrak{q} \rangle} \Vdash \langle \mathfrak{p}, \mathfrak{q} \rangle : [\sigma_i]_{i \in I}$, for some $[\sigma_i]_{i \in I}$. Since \mathfrak{v} is a \mathcal{K} -canonical form, then $\Delta_i \vdash \mathfrak{v} : \sigma_i$ for all $i \in I$, as shown above. Thus $\Gamma +_{i \in I} \Delta_i \vdash \mathfrak{t}'[\langle \mathfrak{p}, \mathfrak{q} \rangle / \mathfrak{v}] : \sigma$ by rule (**sub**) with $J = \emptyset$.

- (only if) Let \mathfrak{t} be a typable term, *i.e.* $\Pi \triangleright \Gamma \vdash \mathfrak{t} : \sigma$. Consider a reduction strategy \mathcal{ST} that always chooses a *typed* redex occurrence. By Lem. 9.1 and Lem. 1.2 the strategy \mathcal{ST} always terminates. Let \mathfrak{t}' be a normal-form of \mathfrak{t} for the strategy \mathcal{ST} , *i.e.* \mathfrak{t} reduces to \mathfrak{t}' using \mathcal{ST} , and \mathcal{ST} applied to \mathfrak{t}' is undefined. We know that $\Pi' \triangleright \Gamma \vdash \mathfrak{t}' : \sigma$ by Lem. 9.1. Then, by definition of \mathcal{ST} , \mathfrak{t}' has no typed redex occurrence. A simple induction on \mathfrak{t}' allows to conclude that it is a canonical form. ◀

4 From canonicity to observability

We proved in the previous section that system \mathcal{P} gives a complete characterization of terms having canonical forms. The next theorem proves that system \mathcal{P} is complete with respect to observability.

► **Theorem 11.** *Observability implies typability.*

Proof. If \mathfrak{t} is observable, then there is a head context \mathfrak{C} such that $\mathfrak{C}[\mathfrak{t}]$ reduces to $\langle \mathfrak{u}, \mathfrak{v} \rangle$, for some \mathfrak{u} and \mathfrak{v} . Since all pairs are typable, the term $\mathfrak{C}[\mathfrak{t}]$ is typable by Lem. 9.2. Remember that $\mathfrak{C}[\mathfrak{t}] = (\lambda \mathfrak{p}_1 \dots \mathfrak{p}_n. \mathfrak{t})\mathfrak{t}_1 \dots \mathfrak{t}_m$ so that \mathfrak{t} is typable too, by easy inspection of the typing system. ◀

Unfortunately, soundness does not hold, *i.e.* the set of observable terms is strictly included in the set of terms having canonical form, as shown below.

► **Example 12.** The term $\mathfrak{t}_1 = \lambda \mathfrak{x}. \text{Id}[\langle \mathfrak{y}, \mathfrak{z} \rangle / \mathfrak{x}][\langle \mathfrak{y}', \mathfrak{z}' \rangle / \mathfrak{x}\text{Id}]$ is canonical, hence typable (by Thm. 10), but not observable. In fact, it is easy to see that there is no term \mathfrak{u} such that both \mathfrak{u} and $\mathfrak{u}\text{Id}$ reduce to pairs. A less trivial example is the term $\mathfrak{t}_2 = \lambda \mathfrak{x}. \text{Id}[\langle \mathfrak{y}, \mathfrak{z} \rangle / \mathfrak{x}\langle \text{Id}, \text{Id} \rangle][\langle \mathfrak{y}', \mathfrak{z}' \rangle / \mathfrak{x}\text{IdId}]$, which is canonical, hence typable, but not observable, as proved in the next lemma.

► **Lemma 13.** *There is no closed term u s.t. both $u\langle \text{Id}, \text{Id} \rangle$ and $u\text{IdId}$ reduce to pairs.*

Proof. By contradiction. Indeed, assume that there exist a closed term u such that both $u\langle \text{Id}, \text{Id} \rangle$ and $u\text{IdId}$ reduce to pairs. Since pairs are always typable, then $u\langle \text{Id}, \text{Id} \rangle$ and $u\text{IdId}$ are typable by Lem. 9.2. In any of the typing derivations of such terms, u occurs in a typed position, so that u turns out to be also typable.

Now, since u is typable and closed, then it reduces to a (typable and closed) canonical form $v \in \mathcal{J}$ by Thm. 10. But v cannot be in \mathcal{K} , which only contains open terms. Moreover, v cannot be a pair, otherwise $u\langle \text{Id}, \text{Id} \rangle \rightarrow^* v\langle \text{Id}, \text{Id} \rangle \rightarrow^* \langle v_1, v_2 \rangle \langle \text{Id}, \text{Id} \rangle \rightarrow^* \text{fail}$ which contradicts (by Lem. 1) the fact that $u\langle \text{Id}, \text{Id} \rangle$ reduces to a pair. We then have two possible forms for v .

If $v = s[\langle p_1, p_2 \rangle / k]$, where $s \in \mathcal{J}$ and $k \in \mathcal{K}$. Then k is an open term which implies v is an open term. Contradiction.

If $v = \lambda p.s$, where $s \in \mathcal{J}$, then p is necessarily a variable, say z , since otherwise $v\text{Id}$ reduces to **fail**, and hence $u\text{IdId} \rightarrow^* v\text{IdId} \rightarrow^* \text{fail}$, which contradicts (by Lem. 1) the fact that $u\text{IdId}$ reduces to a pair. We analyze the possible forms of s .

- If s is a pair, then $u\text{IdId} \rightarrow^* (\lambda z.s)\text{IdId} \rightarrow^* \text{fail}$, which contradicts (by Lem. 1) the fact that $u\text{IdId}$ reduces to a pair.
- If s is an abstraction, then $u\langle \text{Id}, \text{Id} \rangle \rightarrow^* (\lambda z.s)\langle \text{Id}, \text{Id} \rangle$ which reduces to an abstraction, contradicting (by Lem. 1) the fact that $u\langle \text{Id}, \text{Id} \rangle$ reduces to a pair.
- If s is in \mathcal{K} , then $s = x\tau_1 \dots \tau_n$ with $n \geq 0$. Remark that $z \neq x$ is not possible since $v = \lambda z.s$ is closed. Then $z = x$. If $s = z$, then $u\text{IdId}$ reduces to **Id** which contradicts (by Lem. 1) the fact that $u\text{IdId}$ reduces to a pair. Otherwise, $s = z\tau_1 \dots \tau_n$ with $n \geq 1$, and thus $u\langle \text{Id}, \text{Id} \rangle$ reduces to $\langle \text{Id}, \text{Id} \rangle \tau_1 \dots \tau_n \rightarrow^* \text{fail}$, which contradicts again (by Lem. 1) the fact that $u\langle \text{Id}, \text{Id} \rangle$ reduces to a pair.
- If s is $s'[\langle p_1, p_2 \rangle / k]$, with $k \in \mathcal{K}$, then $k = z\tau_1 \dots \tau_n$ with $n \geq 0$, since any other head variable for k would contradict v closed. Now, in the first case we have $u\text{IdId}$ reduces to **fail** which contradicts (by Lem. 1) the fact that $u\text{IdId}$ reduces to a pair. Otherwise, $k = z\tau_1 \dots \tau_n$ with $n \geq 1$ implies $u\langle \text{Id}, \text{Id} \rangle$ reduces to **fail** which contradicts (by Lem. 1) the fact that $u\langle \text{Id}, \text{Id} \rangle$ reduces to a pair. ◀

The first non-observable term τ_1 in Ex. 12 could be ruled out by introducing a notion of *compatibility* between types and requiring multiset types to be composed only by compatible strict types. Unfortunately, we claim that a compatibility relation defined *syntactically*, let us call it **comp**, cannot lead to a sound and complete characterization of observability. By “defined syntactically” we mean that the value of $\text{comp}(\sigma \rightarrow \sigma', \rho \rightarrow \rho')$ should only depend on the values of $\text{comp}(\sigma, \rho)$ and $\text{comp}(\sigma', \rho')$. Another basic requirement of **comp** would be that every product type is incompatible with any functional type. The second non-observable term τ_2 in Ex. 12 is appropriate to illustrate our claim, by keeping in mind that any pair of types assignable to x in any typing derivation for τ_2 need to be incompatible.

Indeed, the shortest typing for τ_2 above is obtained by assigning to x the two types $[] \rightarrow o$ and $[] \rightarrow [] \rightarrow o$, and in order to state the incompatibility between them it would be necessary to define that $\text{comp}(\sigma, \rho)$ and $\neg \text{comp}(\sigma', \rho')$ imply $\neg \text{comp}([\sigma] \rightarrow \sigma', [\rho] \rightarrow \rho')$. Another typing for τ_2 is obtained by assigning to x the two types $[o] \rightarrow o$ and $[\tau] \rightarrow [\tau] \rightarrow o$ respectively, where $\tau = [o] \rightarrow o$, so that $\neg \text{comp}(\sigma, \rho)$ and $\neg \text{comp}(\sigma', \rho')$ should imply $\neg \text{comp}([\sigma] \rightarrow \sigma', [\rho] \rightarrow \rho')$. We conclude that $\neg \text{comp}(\sigma', \rho')$ alone should imply $\neg \text{comp}([\sigma] \rightarrow \sigma', [\rho] \rightarrow \rho')$. However, arrow types $[\sigma] \rightarrow \sigma'$ and $[\rho] \rightarrow \rho'$ having incompatible right-hand sides may very well be compatible. For instance, letting $\sigma = \sigma' = o$ and $\rho = \rho' = [o] \rightarrow o$, one gets two types for **Id**

which need of course to be compatible. Hence, a *syntactic* characterization of such a notion of compatibility seems out of reach.

Fortunately, there exists a sound and complete *semantical* notion of compatibility between types, obtained *a posteriori* as follows: given two strict types π_1 and π_2 , build the corresponding sets of inhabitants $\mathsf{T}(\emptyset, \pi_1)$ and $\mathsf{T}(\emptyset, \pi_2)$, using the inhabitation algorithm presented in Sec. 5. Then π_1 and π_2 are *semantically* compatible if and only if $\mathsf{T}(\emptyset, \pi_1) \cap \mathsf{T}(\emptyset, \pi_2)$ is non-empty.

While the inhabitation problem for (idempotent) intersection types is undecidable [17], it becomes decidable for non-idempotent intersection types [5], which is just a subsystem of our typing system \mathcal{P} introduced in Sec. 3. We will prove in the following that inhabitation is also decidable for the non-trivial extension \mathcal{P} . We will then use this result for characterizing observability in the pattern calculus without referring to a complete syntactic characterization, which is not possible in this framework, as illustrated by Example 12.

5 Inhabitation for System \mathcal{P}

We now show a sound and complete algorithm to solve the inhabitation problem for System \mathcal{P} . Given a strict type σ , the inhabitation problem consists in finding a closed term \mathbf{t} such that $\vdash \mathbf{t} : \sigma$ is derivable. We extend the problem to multiset types by defining \mathbf{A} to be inhabited if and only if there is a closed term \mathbf{t} such that $\vdash \mathbf{t} : \sigma_i$ for every $\sigma_i \in \mathbf{A}$. These notions will naturally be generalized later to non-closed terms.

We already noticed that the system \mathcal{P} allows to type terms containing untyped subterms through the rule $(\rightarrow \mathbf{e})$ with $I = \emptyset$ and the rule (\mathbf{sub}) with $I = J = \emptyset$. In order to identify inhabitants in such cases we introduce a term constant Ω to denote a generic untyped subterm. Our inhabitation algorithm produces **approximate normal forms** $(\mathbf{a}, \mathbf{b}, \mathbf{c})$, also written *anf*, defined as follows:

$$\begin{aligned} \mathbf{a}, \mathbf{b}, \mathbf{c} &::= \Omega \mid \mathcal{N} & \mathcal{N} &::= \lambda \mathbf{p}. \mathcal{N} \mid \langle \mathbf{a}, \mathbf{b} \rangle \mid \mathcal{L} \mid \mathcal{N}[\langle \mathbf{p}, \mathbf{q} \rangle / \mathcal{L}] \\ & & \mathcal{L} &::= \mathbf{x} \mid \mathcal{L} \mathbf{a} \end{aligned}$$

Note that *anfs* do not contain redexes, differently from canonical forms. In particular, thanks to the reduction rule (r_4) (resp. (r_5)), they do not contain *head* (resp. *nested*) explicit matchings. This makes the inhabitation algorithm much more intuitive and simpler.

► **Example 14.** the term $\lambda \langle \mathbf{x}, \mathbf{y} \rangle. (\mathbf{x}(\mathbf{IdId}))[\langle \mathbf{z}_1, \mathbf{z}_2 \rangle / \mathbf{yId}]$ is canonical but not an *anf*, while $\lambda \langle \mathbf{x}, \mathbf{y} \rangle. (\mathbf{x}\Omega)[\langle \mathbf{z}_1, \mathbf{z}_2 \rangle / \mathbf{yId}]$ is an *anf*.

Anfs are ordered by the smallest contextual order \leq such that $\Omega \leq \mathbf{a}$, for any \mathbf{a} . We also write $\mathbf{a} \leq \mathbf{t}$ when the term \mathbf{t} is obtained from \mathbf{a} by replacing each occurrence of Ω by a term of $\Lambda_{\mathbf{p}}$: For example $\mathbf{x}\Omega \leq \mathbf{x}(\mathbf{Id}\Delta)(\Delta\Delta)$ is obtained by replacing the first (resp. second) occurrence of Ω by $\mathbf{Id}\Delta$ (resp. $\Delta\Delta$).

Let $\mathcal{A}(\mathbf{t}) = \{\mathbf{a} \mid \exists \mathbf{u} \mathbf{t} \rightarrow^* \mathbf{u} \text{ and } \mathbf{a} \leq \mathbf{u}\}$ be the set of **approximants** of the term \mathbf{t} , and let \bigvee denote the least upper bound with respect to \leq . We write $\uparrow_{i \in I} \mathbf{a}_i$ to denote the fact that $\bigvee \{\mathbf{a}_i\}_{i \in I}$ does exist. It is easy to check that, for every \mathbf{t} and $\mathbf{a}_1, \dots, \mathbf{a}_n \in \mathcal{A}(\mathbf{t})$, $\uparrow_{i \in \{1, \dots, n\}} \mathbf{a}_i$. An *anf* \mathbf{a} is a **head subterm** of \mathbf{b} if either $\mathbf{b} = \mathbf{a}$ or $\mathbf{b} = \mathbf{c}\mathbf{c}'$ and \mathbf{a} is a head subterm of \mathbf{c} . System \mathcal{P} can also be trivially extended to give types to *anfs*, simply assuming that no type can be assigned to the constant Ω . It is easy to check that, if $\Gamma \vdash \mathbf{a} : \sigma$ and $\mathbf{a} \leq \mathbf{b}$ (resp. $\mathbf{a} \leq \mathbf{t}$) then $\Gamma \vdash \mathbf{b} : \sigma$ (resp. $\Gamma \vdash \mathbf{t} : \sigma$).

Given $\Pi \triangleright \Gamma \vdash \mathbf{t} : \tau$, where \mathbf{t} is in Π -nf (cf. Sec. 3), $\mathcal{A}(\Pi)$ is the minimal approximant \mathbf{b} of \mathbf{t} such that $\Pi \triangleright \Gamma \vdash \mathbf{b} : \tau$. Formally, given $\Pi \triangleright \Gamma \vdash \mathbf{t} : \sigma$, where \mathbf{t} is in Π -nf, the **minimal approximant** of Π , written $\mathcal{A}(\Pi)$, is defined by induction on $\mathbf{meas}(\Pi)$ as follows:

- $\mathcal{A}(\Gamma \vdash x : \rho) = x$; $\mathcal{A}(\Gamma \vdash \langle t, u \rangle : o) = \langle \Omega, \Omega \rangle$.
- If $\Pi \triangleright \Gamma \vdash \lambda p.t : A \rightarrow \rho$ follows from $\Pi' \triangleright \Gamma' \vdash t : \rho$, then $\mathcal{A}(\Pi) = \lambda p.\mathcal{A}(\Pi')$, t being in Π' -nf.
- If $\Pi \triangleright \Gamma \vdash \langle t, u \rangle : \times_1(\tau)$ follows from $\Pi' \triangleright \Gamma \vdash t : \tau$, then $\mathcal{A}(\Pi) = \langle \mathcal{A}(\Pi'), \Omega \rangle$, t being in Π' -nf. Similarly for a pair of type $\times_2(\tau)$.
- If $\Pi \triangleright \Gamma = \Gamma' +_{i \in I} \Delta_i \vdash tu : \rho$ follows from $\Pi' \triangleright \Gamma' \vdash t : [\sigma_i]_{i \in I} \rightarrow \rho$ and $(\Pi'_i \triangleright \Delta_i \vdash u : \sigma_i)_{i \in I}$, then $\mathcal{A}(\Pi) = \mathcal{A}(\Pi')(\bigvee_{i \in I} \mathcal{A}(\Pi'_i))$
- If $\Pi \triangleright \Gamma = \Gamma' +_{i \in I} \Delta_i \vdash t[p/u] : \tau$ follows from $\Pi' \triangleright \Gamma'' \vdash t : \tau$ and $(\Psi_i \triangleright \Delta_i \vdash u : \rho_i)_{i \in I}$, then $\mathcal{A}(\Pi) = \mathcal{A}(\Pi')[p/\bigvee_{i \in I} \mathcal{A}(\Psi_i)]$

Remark that, in the application case of the definition above, the *anf* corresponding to $I = \emptyset$ is $\mathcal{A}(\Pi')\Omega$. Moreover, in the last case, p cannot be a variable, t being in Π -nf. A simple inspection of the typing rules for \Vdash shows that in this case $I \neq \emptyset$.

► **Example 15.** Consider the following derivation Π :

$$\frac{\frac{y : [[] \rightarrow o] \vdash y : [] \rightarrow o}{y : [[] \rightarrow o] \vdash y(\Delta\Delta) : o} \quad \Vdash \langle z_1, z_2 \rangle : o \quad \frac{x : [[] \rightarrow o] \vdash x : [] \rightarrow o}{x : [[] \rightarrow o] \vdash x\text{Id} : o}}{x : [[] \rightarrow o]; y : [[] \rightarrow o] \vdash y(\Delta\Delta)[\langle z_1, z_2 \rangle/x\text{Id}] : o}}{\vdash \lambda xy.y(\Delta\Delta)[\langle z_1, z_2 \rangle/x\text{Id}] : [[] \rightarrow o] \rightarrow [[] \rightarrow o] \rightarrow o}$$

The minimal approximant of Π is $\lambda xy.y\Omega[\langle z_1, z_2 \rangle/x\Omega]$.

A simple induction on $\text{meas}(\Pi)$ allows to show the following:

► **Lemma 16.** *If $\Pi \triangleright \Gamma \vdash t : \sigma$ and t is in Π -nf, then $\Pi \triangleright \Gamma \vdash \mathcal{A}(\Pi) : \sigma$.*

5.1 The inhabitation algorithm

The inhabitation algorithm is presented in Fig. 2. As usual, in order to solve the problem for closed terms, it is necessary to extend the algorithm to open ones, so, given an environment Γ and a strict type σ , the algorithm builds the set $\mathbf{T}(\Gamma, \sigma)$ containing *all* the *anfs* \mathbf{a} such that there exists a derivation $\Pi \triangleright \Gamma \vdash \mathbf{a} : \sigma$, with $\mathbf{a} = \mathcal{A}(\Pi)$, then stops². Thus, our algorithm is not an extension of the classical inhabitation algorithm for simple types [4, 9]. In particular, when restricted to simple types, it constructs all the *anfs* inhabiting a given type, while the original algorithm reconstructs just the *long η -normal forms*. The algorithm uses four auxiliary predicates, namely

- $\mathbf{P}_{\mathcal{V}}(\mathbf{A})$, where \mathcal{V} is a finite set of variables, contains the pairs (Γ, \mathbf{p}) such that (i) $\Gamma \Vdash \mathbf{p} : \mathbf{A}$, and (ii) \mathbf{p} does not contain any variable in \mathcal{V} .
- $\mathbf{TI}(\Gamma, [\sigma_i]_{i \in I})$, contains all the *anfs* $\mathbf{a} = \bigvee_{i \in I} \mathbf{a}_i$ such that $\Gamma = +_{i \in I} \Gamma_i$, $\mathbf{a}_i \in \mathbf{T}(\Gamma_i, \sigma_i)$ for all $i \in I$, and $\uparrow_{i \in I} \mathbf{a}_i$.
- $\mathbf{H}_{\mathbf{b}}^{\Delta}(\Gamma, \sigma) \triangleright \tau$ contains all the *anfs* \mathbf{a} such that \mathbf{b} is a head subterm of \mathbf{a} , and such that if $\mathbf{b} \in \mathbf{T}(\Delta, \sigma)$ then $\mathbf{a} \in \mathbf{T}(\Gamma + \Delta, \tau)$.
- $\mathbf{HI}_{\mathbf{b}}^{\Delta}(\Gamma, [\sigma_i]_{i \in I}) \triangleright [\rho_i]_{i \in I}$ contains all the *anf* $\mathbf{a} = \bigvee_{i \in I} \mathbf{a}_i$ such that $\Delta = +_{i \in I} \Delta_i$, $\Gamma = +_{i \in I} \Gamma_i$, $\mathbf{a}_i \in \mathbf{H}_{\mathbf{b}}^{\Delta}(\Gamma, \sigma_i) \triangleright \rho_i$ and $\uparrow_{i \in I} \mathbf{a}_i$.

² It is worth noticing that, given Γ and σ , the set of *anfs* \mathbf{a} such that there exists a derivation $\Pi \triangleright \Gamma \vdash \mathbf{a} : \sigma$ is possibly infinite. However, the subset of those verifying $\mathbf{a} = \mathcal{A}(\Pi)$ is finite; they are the minimal ones, those generated by the inhabitation algorithm (this is proved in Lem. 19).

Note that the algorithm has different kinds of non-deterministic behaviours, *i.e.* different choices of rules can produce different results. Indeed, given an input (Γ, σ) , the algorithm may apply a rule like **(Abs)** in order to decrease the type σ , or a rule like **(Head)** in order to decrease the environment Γ . Moreover, every rule (R) which is based on some decomposition of the environment and/or the type, like **(Subs)**, admits different applications. In what follows we illustrate the non-deterministic behaviour of the algorithm. For that, we represent a **run of the algorithm** as a tree whose nodes are labeled with the name of the rule applied.

► **Example 17.** We consider different inputs of the form (\emptyset, σ) , for different strict types σ . For every such input, we give an output and the corresponding run.

1. $\sigma = [[\alpha] \rightarrow \alpha] \rightarrow [\alpha] \rightarrow \alpha$.
 - a. output: $\lambda xy.xy$, run: $\text{Abs}(\text{Abs}(\text{Head}(\text{Prefix}(\text{TUn}(\text{Head}(\text{Final}))), \text{Final})), \text{Varp}), \text{Varp})$.
 - b. output: $\lambda x.x$, run: $\text{Abs}(\text{Head}(\text{Final}), \text{Varp})$.
2. $\sigma = [[[] \rightarrow \alpha] \rightarrow \alpha$. output: $\lambda x.x\Omega$, run: $\text{Abs}(\text{Head}(\text{Prefix}(\text{TUn}, \text{Final})), \text{Varp})$.
3. $\sigma = [[o] \rightarrow o, o] \rightarrow o$.
 - a. output: $\lambda x.xx$, run: $\text{Abs}(\text{Head}(\text{Prefix}(\text{TUn}(\text{Head}(\text{Final}))), \text{Final})), \text{Varp})$.
 - b. Explicit substitutions may be used to consume some, or all, the resources in $[[o] \rightarrow o, o]$
 output: $\lambda x.x[\langle y, z \rangle / x](\Omega, \Omega)$, run:
 $\text{Abs}(\text{Subs}(\text{HUn}(\text{Prefix}(\text{TUn}(\text{Pair}), \text{Final}))), \text{Pairp}(\text{Weakp}, \text{Weakp}), \text{Head}(\text{Final})), \text{Varp})$.
 - c. There are four additional runs, producing the following outputs:
 - $\lambda x.x(\Omega, \Omega)[\langle y, z \rangle / x]$,
 - $\lambda x.\langle \Omega, \Omega \rangle[\langle y, z \rangle / xx]$,
 - $\lambda x.\langle \Omega, \Omega \rangle[\langle y, z \rangle / x][\langle w, s \rangle / x(\Omega, \Omega)]$,
 - $\lambda x.\langle \Omega, \Omega \rangle[\langle y, z \rangle / x(\Omega, \Omega)][\langle w, s \rangle / x]$.

Along the recursive calls of the inhabitation algorithm, the parameters (type and/or environment) decrease strictly, for a suitable notion of measure, so that every run is finite:

► **Lemma 18.** *The inhabitation algorithm terminates.*

5.2 Soundness and completeness

We now prove soundness and completeness of our inhabitation algorithm.

► **Lemma 19.** $\mathbf{a} \in \mathsf{T}(\Gamma, \sigma) \Leftrightarrow \exists \Pi \triangleright \Gamma \vdash \mathbf{a} : \sigma$ such that $\mathbf{a} = \mathcal{A}(\Pi)$.

Proof. The “only if” part is proved by induction on the rules in Fig. 2, and the “if” part is proved by induction on the definition of $\mathcal{A}(\Pi)$. In both parts, additional statements concerning the predicates of the inhabitation algorithm other than **T** are required, in order to strengthen the inductive hypothesis. ◀

► **Theorem 20 (Soundness and Completeness).**

1. If $\mathbf{a} \in \mathsf{T}(\Gamma, \sigma)$ then, for all \mathbf{t} such that $\mathbf{a} \leq \mathbf{t}$, $\Gamma \vdash \mathbf{t} : \sigma$.
2. If $\Pi \triangleright \Gamma \vdash \mathbf{t} : \sigma$ then there exists $\Pi' \triangleright \Gamma \vdash \mathbf{t}' : \sigma$ such that \mathbf{t}' is in Π' -nf, and $\mathcal{A}(\Pi') \in \mathsf{T}(\Gamma, \sigma)$.

Proof. Soundness follows from Lem. 19 (\Rightarrow) and the fact that $\Gamma \vdash \mathbf{a} : \sigma$ and $\mathbf{a} \leq \mathbf{t}$ imply $\Gamma \vdash \mathbf{t} : \sigma$. For completeness we first apply Lem. 9.1 that guarantees the existence of $\Pi' \triangleright \Gamma \vdash \mathbf{t}' : \sigma$ such that \mathbf{t}' is in Π' -nf, and then Lem. 16 and Lem. 19 (\Leftarrow). ◀

$$\begin{array}{c}
\frac{x \notin \mathcal{V}}{(\emptyset; x) \in_0 P_{\mathcal{V}}([\])} \text{ (Weakp)} \quad \frac{x \notin \mathcal{V}}{(x : B; x) \in_0 P_{\mathcal{V}}(B)} \text{ (Varp)} \\
\\
\frac{(\Gamma; p) \in_i P_{\mathcal{V}}(A_1) \quad (\Delta; q) \in_j P_{\mathcal{V}}(A_2) \quad p \# q}{(\Gamma; \Delta; \langle p, q \rangle) \in_1 P_{\mathcal{V}}(A_1 \boxtimes A_2)} \text{ (Pairp)} \\
\\
\frac{a \in T(\Gamma; \Delta, \tau) \quad A = A_1 \sqcup A_2 \quad (\Delta; p) \in_i P_{\text{dom}(\Gamma)}(A_1) \quad \text{supp}(A_2) \subseteq \mathcal{S}(p)}{\lambda p. a \in T(\Gamma, A \rightarrow \tau)} \text{ (Abs)} \\
\\
\frac{(\mathbf{a}_i \in T(\Gamma_i, \sigma_i))_{i \in I} \quad \uparrow_{i \in I} \mathbf{a}_i}{\bigvee_{i \in I} \mathbf{a}_i \in \text{TI}(+_{i \in I} \Gamma_i, [\sigma_i]_{i \in I})} \text{ (TUn)} \quad \frac{(\mathbf{a}_i \in H_b^{\Delta_i}(\Gamma_i, \sigma_i) \triangleright \rho_i)_{i \in I} \quad \uparrow_{i \in I} \mathbf{a}_i}{\bigvee_{i \in I} \mathbf{a}_i \in H_b^{+_{i \in I} \Delta_i}(+_{i \in I} \Gamma_i, [\sigma_i]_{i \in I}) \triangleright [\rho_i]_{i \in I}} \text{ (HUn)} \\
\\
\frac{}{\langle \Omega, \Omega \rangle \in T(\emptyset, o)} \text{ (Pair)} \quad \frac{a \in T(\Gamma, \tau)}{\langle \mathbf{a}, \Omega \rangle \in T(\Gamma, \times_1(\tau))} \text{ (Prod1)} \quad \frac{a \in T(\Gamma, \tau)}{\langle \Omega, \mathbf{a} \rangle \in T(\Gamma, \times_2(\tau))} \text{ (Prod2)} \\
\\
\frac{a \in H_x^{\mathbf{x}:[\sigma]}(\Gamma, \sigma) \triangleright \tau}{a \in T(\Gamma + (x : [\sigma]), \tau)} \text{ (Head)} \quad \frac{\sigma = \tau}{a \in H_a^{\Delta}(\emptyset, \sigma) \triangleright \tau} \text{ (Final)} \\
\\
\frac{\Gamma = \Gamma_0 + \Gamma_1 \quad b \in \text{TI}(\Gamma_0, A) \quad a \in H_{cb}^{\Delta + \Gamma_0}(\Gamma_1, \sigma) \triangleright \tau}{a \in H_c^{\Delta}(\Gamma, A \rightarrow \sigma) \triangleright \tau} \text{ (Prefix)} \\
\\
\frac{\Gamma = \Gamma_0 + \Gamma_1 \quad c \in H_x^{\mathbf{x}:B}(\Gamma_0, B) \triangleright F(B) \quad F(B) = A_1 \sqcup A_2 (*) \quad (\Delta, p) \in_1 P_{\text{dom}(\Gamma_0 + \Gamma_1 + (x:B))}(A_1) \quad \text{supp}(A_2) \subseteq \mathcal{S}(p) \quad b \in T(\Gamma_1; \Delta, \tau)}{b[p/c] \in T(\Gamma + (x : B), \tau)} \text{ (Subs)} \\
\\
(*) \text{ where the operator } F() \text{ is defined as follows:} \\
F(\alpha) := \alpha \quad F(A \rightarrow \tau) := F(\tau) \quad F([\]) := [\] \quad F([\sigma_i]_{i \in I}) := [F(\sigma_i)]_{i \in I}
\end{array}$$

■ **Figure 2** The inhabitation algorithm.

6 Characterizing Observability

We are now able to state the main result of this paper, *i.e.* the characterization of observability for the pattern calculus. The following lemma assures that types reflect correctly the structure of the data types.

- **Lemma 21.** *Let \mathfrak{t} be a closed and typable term, then*
- *If \mathfrak{t} has functional type, then \mathfrak{t} reduces to an abstraction.*
 - *If \mathfrak{t} has product type, then \mathfrak{t} reduces to a pair.*

Proof. Let \mathfrak{t} be a closed and typable term. By Thm. 10 we know that \mathfrak{t} reduces to a (closed) canonical form in \mathcal{J} . The proof is by induction on the maximal length of such reduction sequences.

- If \mathfrak{t} is already a canonical form, we analyze all the cases.
- If \mathfrak{t} is a variable, then this gives a contradiction with \mathfrak{t} closed.
 - If \mathfrak{t} is a function, then the property trivially holds.

- If τ is a pair, then the property trivially holds.
- If τ is an application, then τ has the form $x\tau_1 \dots \tau_n$. Therefore at least x belongs to the set of free variables of τ , which leads to a contradiction with τ closed.
- If τ is a closure, *i.e.* $\tau = u[\langle p_1, p_2 \rangle / v]$, where $v \in \mathcal{K}$ has the form $x\tau_1 \dots \tau_n$, then at least x belongs to the set of free variables of τ , which leads to a contradiction with τ closed.

Otherwise, $\tau \rightarrow \tau' \rightarrow^* u$, where u is in \mathcal{J} . The term τ' is also closed and typable (Lem. 9.1), then the *i.h.* gives the desired result for τ' , so the property holds also for τ . ◀

► **Theorem 22** (Characterizing Observability). *A term τ is observable iff $\Pi \triangleright x_1 : A_1; \dots; x_n : A_n \vdash \tau : B_1 \rightarrow \dots \rightarrow B_m \rightarrow \alpha$, where $n \geq 0, m \geq 0$, α is a product type and all $A_1, \dots, A_n, B_1, \dots, B_m$ are inhabited.*

Proof. The left-to-right implication: if τ is observable, then there exists a head-context C such that $C[\tau] \rightarrow^* \langle u, v \rangle$. Since $\vdash \langle u, v \rangle : o$, we get $\Pi' \triangleright \vdash C[\tau] : o$ by Lem. 9.2. By definition $C[\tau] = (\lambda p_1 \dots \lambda p_n. \tau) u_1 \dots u_m$, so Π has a subderivation $\Pi' \triangleright \vdash \lambda p_1 \dots \lambda p_n. \tau : B_1 \rightarrow \dots \rightarrow B_m \rightarrow o$ (by rule $(\rightarrow e)$), where B_i is inhabited by u_i ($1 \leq i \leq m$). Since $n \leq m$, Π' has a subderivation $\Pi'' \triangleright \Gamma \vdash \tau : B_{n+1} \rightarrow \dots \rightarrow B_m \rightarrow o$ (by rule $(\rightarrow i)$), where $\Gamma|_{p_i} \Vdash p_i : B_i$ ($1 \leq i \leq n$). The result follows since $x_1 : A_1, \dots, x_l : A_l \Vdash p : B$ and B is inhabited implies that all the A_i are inhabited. The right-to-left implication: if $A_1, \dots, A_n, B_1, \dots, B_m$ are all inhabited, then there exist $u_1, \dots, u_n, v_1, \dots, v_m$ such that $\vdash u_i : \sigma_i^j$ for every type σ_i^j of A_i ($1 \leq i \leq n$) and $\vdash v_i : \rho_i^j$ for every type ρ_i^j of B_i ($1 \leq i \leq m$). Let $C = (\lambda x_1 \dots x_n. \square) u_1 \dots u_n v_1 \dots v_n$ be a head-context. We have $\vdash C[\tau] : \alpha$, which in turn implies that $C[\tau]$ reduces to a pair, by Lem. 21. Then the term τ is observable by definition. ◀

The notion of observability is conservative with respect to that of solvability in λ -calculus.

► **Theorem 23** (Conservativity). *A λ -term τ is solvable in the λ -calculus if and only if τ is observable in Λ_p .*

- Proof.** ■ (if) Take an unsolvable λ -term τ so that τ does not have head normal-form. Then τ (seen as a term of our calculus) has no canonical form, and thus τ is not typable by Thm. 10. It turns out that τ is not observable in Λ_p by Thm. 22.
- (only if) Take a solvable λ -term τ so that there exist a head-context C such that $C[\tau]$ reduces to Id , then it is easy to construct a head context C' such that $C'[\tau]$ reduces to a pair (just take $C' = C \langle \tau_1, \tau_2 \rangle$ for some terms τ_1, τ_2). ◀

7 Conclusion and Further Work

We propose a notion of observability for pair pattern calculi which is conservative with respect to the notion of solvability for λ -calculus.

We provide a logical characterization of observable terms by means of typability *and* inhabitation.

Further work will be developed in different directions. As we already discussed in Sec. 2, different definitions of observability would be possible. We explored the one based on a lazy semantics, but it would be also interesting to obtain a full characterization based on a strict semantics. Another point to be developed is the definition of a suitable notion of head reduction, which, despite its relative simplicity, turn out to be quite cumbersome. On the semantical side, it is well known that non-idempotent intersection types can be used to supply a logical description of the relational semantics of λ -calculus [8, 14]. We would like to start from our type assignment system for building a denotational model of the pattern

calculus. Last but not least, a challenging question is related to the characterization of observability in a more general framework of pattern λ -calculi allowing the patterns to be dynamic [10].

References

- 1 Thibaut Balabonski. On the implementation of dynamic patterns. In Eduardo Bonelli, editor, *HOR*, volume 49 of *EPTCS*, pages 16–30, 2010.
- 2 Henk Barendregt. *The Lambda Calculus: Its Syntax and Semantics*, volume 103 of *Studies in logic and the foundation of mathematics*. North-Holland, Amsterdam, revised edition, 1984.
- 3 Henk Barendregt, Mario Coppo, and Mariangiola Dezani-Ciancaglini. A filter lambda model and the completeness of type assignment. *The Journal of Symbolic Logic*, 48(4):931–940, 1983.
- 4 C. Ben-Yelles. *Type-assignment in the lambda-calculus; syntax and semantics*. PhD thesis, University of Wales Swansea, 1979.
- 5 Antonio Bucciarelli, Delia Kesner, and Simona Ronchi Della Rocca. The inhabitation problem for non-idempotent intersection types. In Josep Díaz, Ivan Lanese, and Davide Sangiorgi, editors, *TCS, LNCS*. Springer, 2014. To appear.
- 6 Serenella Cerrito and Delia Kesner. Pattern matching as cut elimination. *Theoretical Computer Science*, 323(1-3):71–127, 2004.
- 7 Horatiu Cirstea, Germain Faure, and Claude Kirchner. A rho-calculus of explicit constraint application. *Higher-Order and Symbolic Computation*, 20(1-2):37–72, 2007.
- 8 Daniel de Carvalho. Execution time of lambda-terms via denotational semantics and intersection types. *CoRR*, abs/0905.4251, 2009.
- 9 J. Roger Hindley. *Basic Simple Type Theory*. Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, Amsterdam, 2008.
- 10 Barry Jay and Delia Kesner. First-class patterns. *Journal of Functional Programming*, 19(2):191–225, 2009.
- 11 Wolfram Kahl. Basic pattern matching calculi: A fresh view on matching failure. In Yuki Yoshi Kameyama and Peter Stuckey, editors, *FLOPS*, volume 2998 of *LNCS*, pages 276–290. Springer, 2004.
- 12 Jan-Willem Klop, Vincent van Oostrom, and Roel de Vrijer. Lambda calculus with patterns. *Theoretical Computer Science*, 398(1-3):16–31, 2008.
- 13 Jean Louis Krivine. *Lambda-Calculus, Types and Models*. Masson, Paris, and Ellis Horwood, Hemel Hempstead, 1993.
- 14 Luca Paolini, Mauro Piccolo, and Simona Ronchi Della Rocca. Logical relational lambda-models. To appear in *Mathematical Structures in Computer Science*.
- 15 Barbara Petit. A polymorphic type system for the lambda-calculus with constructors. In Pierre-Louis Curien, editor, *Typed Lambda Calculi and Applications, 9th International Conference, TLCA 2009, Brasilia, Brazil, July 1-3, 2009. Proceedings*, volume 5608 of *Lecture Notes in Computer Science*, pages 234–248. Springer, 2009.
- 16 Simon Peyton-Jones. *The Implementation of Functional Programming Languages*. Prentice-Hall, Inc., 1987.
- 17 Pawel Urzyczyn. The emptiness problem for intersection types. *Journal of Symbolic Logic*, 64(3):1195–1215, 1999.
- 18 Vincent van Oostrom. Confluence by decreasing diagrams. *Theoretical Computer Science*, 126(2):259–280, 1994.

Undecidability of Equality in the Free Locally Cartesian Closed Category

Simon Castellan¹, Pierre Clairambault¹, and Peter Dybjer²

- 1 ENS de Lyon, CNRS, Inria, UCBL, Université de Lyon
LIP, 46 allée d'Italie, 69364 Lyon, France
2 Chalmers University of Technology, Sweden

Abstract

We show that a version of Martin-Löf type theory with extensional identity, a unit type N_1, Σ, Π , and a base type is a free category with families (supporting these type formers) both in a 1- and a 2-categorical sense. It follows that the underlying category of contexts is a free locally cartesian closed category in a 2-categorical sense because of a previously proved biequivalence. We then show that equality in this category is undecidable by reducing it to the undecidability of convertibility in combinatory logic.

1998 ACM Subject Classification F.4.1 Mathematical Logic, F.3.2 Semantics of Programming Languages

Keywords and phrases Extensional type theory, locally cartesian closed categories, undecidability

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.138

1 Introduction

In previous work [5, 6] we showed the biequivalence of locally cartesian closed categories (lcccs) and the I, Σ, Π -fragment of extensional Martin-Löf type theory. More precisely, we showed the biequivalence of the following two 2-categories.

- The first has as *objects* lcccs, as *arrows* functors which preserve the lccc-structure (up to isomorphism), and as *2-cells* natural transformations.
- The second has as *objects* categories with families (cwfs) [8] which support extensional identity types (I-types), Σ -types, Π -types, and are *democratic*, as *arrows* pseudo cwf-morphisms (preserving structure up to isomorphism), and as *2-cells* pseudo cwf-transformations. A cwf is democratic iff there is an equivalence between its category of contexts and its category of closed types.

This result is a corrected version of a result by Seely [13] concerning the equivalence of the category of lcccs and the category of Martin-Löf type theories. Seely's paper did not address the coherence problem caused by the interpretation of substitution as pullbacks [7]. As Hofmann showed [9], this coherence problem can be solved by extending a construction of Bénabou [2]. Our biequivalence is based on this construction.

Cwfs are models of the most basic rules of dependent type theory; those dealing with substitution, assumption, and context formation, the rules which come before any rules for specific type formers. The distinguishing feature of cwfs, compared to other categorical notions of model of dependent types, is that they are formulated in a way which makes the connection with the ordinary syntactic formulation of dependent type theory transparent. They can be defined purely equationally [8] as a generalised algebraic theory (gat) [3], where each sort symbol corresponds to a judgment form, and each operator symbol corresponds to



© Simon Castellan, Pierre Clairambault, and Peter Dybjer;
licensed under Creative Commons License CC-BY

13th International Conference on Typed Lambda Calculi and Applications (TLCA'15).
Editor: Thorsten Altenkirch; pp. 138–152



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

an inference rule in a variable free formulation of Martin-Löf's explicit substitution calculus for dependent type theory [11, 15].

Cwfs are not only models of dependent type theory, but also suggest an answer to the question what dependent type theory is as a mathematical object. Perhaps surprisingly, this is a non-trivial question, and Voevodsky has remarked that “a type system is not a mathematical notion”. There are numerous variations of Martin-Löf type theory in the literature, even of the formulation of the most basic rules for dependent types. There are systems with explicit and implicit substitutions, variations in assumption, context formation, and substitution rules. There are formulations with de Bruijn indices and with ordinary named variables, etc. In fact, there are so many rules that most papers do not try to provide a complete list; and if you do try to list all of them how can you be sure that you haven't forgotten any? Nevertheless, there is a tacit assumption that most variations are equivalent and that a complete list of rules could be given if needed. However, from a mathematical point of view this is neither clear nor elegant.

To remedy this situation we suggest to define Martin-Löf type theory (and other dependent type theories) abstractly as the initial cwf (with extra structure). The category of cwfs and morphisms which preserve cwf-structure on the nose was defined by Dybjer [8]. We suggest that the correctness of a definition or an implementation of dependent type theory means that it gives rise to an initial object in this category of cwfs (with extra structure). Here we shall construct the initial object in this category explicitly in the simplest possible way following closely the definition of the generalised algebraic theory of cwfs. Note however that the notion of a generalised algebraic theory is itself based on dependent type theory, that is, on cwf-structure. So just defining the initial cwf as the generalised algebraic theory of cwfs would be circular. Instead we construct the initial cwf explicitly by giving grammar and inference rules which follow closely the operators of the gat of cwfs. However, we must also make equality reasoning explicit. To decrease the number of rules, we present a “per-style” system rather than an ordinary one. We will mutually define four partial equivalence relations (pers): for the judgments of context equality $\Gamma = \Gamma'$, substitution equality $\Delta \vdash \gamma = \gamma' : \Gamma$, type equality $\Gamma \vdash A = A'$, and term equality $\Gamma \vdash a = a' : A$. The ordinary judgments will be defined as the reflexive instances, for example, $\Gamma \vdash a : A$ will be defined as $\Gamma \vdash a = a : A$.

Our only optimisation is the elimination of some redundant arguments of operators. For example, the composition operator in the gat of cwfs has five arguments: three objects and two arrows. However, the three object arguments can be recovered from the arrows, and can hence be omitted. This method is also used in *D-systems*, the essentially algebraic formulation of cwfs by Voevodsky.

The goal of the present paper is to prove the undecidability of equality in the free lccc. To this end we extend our formal system for cwfs with rules for extensional I-types, N_1 , Σ , Π , and a base type. Now we want to show that this yields a free lccc on one object, by appealing to our biequivalence theorem. (Since the empty context corresponds to the unit type N_1 and context extension to Σ , it follows that our free cwf is democratic.) However, it does not suffice to show that we get a free cwf in the 1-category of cwfs and strict cwf-morphism, but we must show that it is also free (“bifree”) in the 2-category of cwfs and pseudo cwf-morphisms. Although informally straightforward, this proof is technically more involved because of the complexity of the notion of pseudo cwf-morphism.

Once we have constructed the free lccc (as a cwf-formulation of Martin-Löf type theory with extensional I-types, N_1 , Σ , Π , and one base type) we will be able to prove undecidability. It is a well-known folklore result that extensional Martin-Löf type theory with one universe

has undecidable equality, and we only need to show that a similar construction can be made without a universe, provided we have a base type. We do this by encoding untyped combinatory logic as a context, and use the undecidability of equality in this theory.

Related work. Palmgren and Vickers [12] show how to construct free models of essentially algebraic theories in general. We could use this result to build a free cwf, but this only shows freeness in the 1-categorical sense. We also think that the explicit construction of the free (and bifree) cwf is interesting in its own right.

Plan. In Section 2 we prove a few undecidability theorems, including the undecidability of equality in Martin-Löf type theory with extensional I-types, Π , and one base type. In Section 3 we construct a free cwf on one base type. We show that it is free both in a 1-categorical sense (where arrows preserve cwf-structure on the nose) and in a 2-categorical sense (where arrows preserve cwf-structure up to isomorphism). In Section 4 we construct a free cwf with extensional identity types, N_1, Σ, Π , and one base type. We then use the biequivalence result to conclude that this yields a free lccc in a 2-categorical sense.

2 Undecidability in Martin-Löf type theory

Like any other single-sorted first order equational theory, combinatory logic can be encoded as a context in Martin-Löf type theory with I-types, Π -types, and a base type o . The context Γ_{CL} for combinatory logic is the following:

$$\begin{array}{ll} k & : o, & ax_k & : \Pi xy : o. I(o, k \cdot x \cdot y, x), \\ s & : o, & ax_s & : \Pi xyz : o. I(o, s \cdot x \cdot y \cdot z, x \cdot z \cdot (y \cdot z)) \\ \cdot & : o \rightarrow o \rightarrow o, \end{array}$$

Here we have used the left-associative binary infix symbol “ \cdot ” for application. Note that k, s, \cdot, ax_k, ax_s are all variables.

► **Theorem 1.** *Type-inhabitation in Martin-Löf type theory with (intensional or extensional) identity-types, Π -types and a base type is undecidable.*

This follows from the undecidability of convertibility in combinatory logic, because the type

$$\Gamma_{\text{CL}} \vdash I(o, M, M')$$

is inhabited iff the closed combinatory terms M and M' are convertible. Clearly, if the combinatory terms are convertible, it can be formalized in this fragment of type theory. For the other direction we build a model of the context Γ_{CL} where o is interpreted as the set of combinatory terms modulo convertibility.

► **Theorem 2.** *Judgmental equality in Martin-Löf type theory with extensional identity-types, Π -types and a base type is undecidable.*

With extensional identity types [10] the above identity type is inhabited iff the corresponding equality judgment is valid:

$$\Gamma_{\text{CL}} \vdash M = M' : o$$

This theorem also holds if we add N_1 and Σ -types to the theory. The remainder of the paper will show that the category of contexts for the resulting fragment of Martin-Löf type theory is bifree in the 2-category of lcccs (Theorem 20). Our main result follows:

► **Theorem 3.** *Equality of arrows in the bifree lccc on one object is undecidable.*

We would like to remark that the following folklore theorem can be proved in the same way.

► **Theorem 4.** *Judgmental equality in Martin-Löf type theory with extensional identity-types, Π -types and a universe U is undecidable.*

If we have a universe we can instead work in the context

$$\begin{array}{ll} X & : \quad U \\ k & : \quad X, \\ s & : \quad X, \end{array} \quad \begin{array}{ll} \cdot & : \quad X \rightarrow X \rightarrow X, \\ ax_k & : \quad \Pi xy : X. I(X, k \cdot x \cdot y, x), \\ ax_s & : \quad \Pi xyz : X. I(X, s \cdot x \cdot y \cdot z, x \cdot z \cdot (y \cdot z)) \end{array}$$

and prove undecidability for this theory (without a base type) in the same way as above.

Note that we don't need any closure properties at all for U – only the ability to quantify over small types. Hence we prove a slightly stronger theorem than the folklore theorem which assumes that U is closed under function types, and then uses the context

$$\begin{array}{ll} X & : \quad U, \\ x & : \quad I(U, X, X \rightarrow X) \end{array}$$

so that X is a model of the untyped lambda calculus.

3 A free category with families

In this section we define a free cwf syntactically, as a *term model* consisting of derivable contexts, substitutions, types and terms modulo derivable equality. To this end we give syntax and inference rules for a cwf-calculus, that is, a variable free explicit substitution calculus for dependent type theory.

We first prove that this calculus yields a free cwf in the category where morphisms preserve cwf-structure on the nose. The free cwf on one object is a rather degenerate structure, since there are no non-trivial dependent types. However, we have nevertheless chosen to present this part of the construction separately. Cwfs model the common core of dependent type theory, including all generalised algebraic theories, pure type systems [1], and fragments of Martin-Löf type theory. The construction of a free pure cwf is thus the common basis for constructing free and initial cwfs with appropriate extra structure for modelling specific dependent type theories.

In Section 3.4 we prove that our free cwf is also bifree. We then extend this result to cwfs supporting N_1 , Σ , and Π -types. By our biequivalence result [5, 6] it also yields a bifree lccc.

3.1 The 2-category of categories with families

The 2-category of cwfs and pseudo-morphisms which preserve cwf-structure up to isomorphism was defined in [5, 6]. Here we only give an outline.

► **Definition 5** (Category with families). A cwf \mathcal{C} is a pair (\mathcal{C}, T) of a category \mathcal{C} and a functor $T : \mathcal{C}^{\text{op}} \rightarrow \mathbf{Fam}$ where \mathbf{Fam} is the category of families of sets. We write $\text{Ctx}_{\mathcal{C}} = |\mathcal{C}|$ and $\text{Sub}_{\mathcal{C}}(\Delta, \Gamma) = \text{Hom}_{\mathcal{C}}(\Delta, \Gamma)$. For $\Gamma \in \text{Ctx}_{\mathcal{C}}$ we write $T\Gamma = (\text{Tm}_{\mathcal{C}}(\Gamma, A))_{A \in \text{Ty}_{\mathcal{C}}\Gamma}$. The functorial action of T on a type A is written $A[_]$: if $\gamma : \text{Sub}_{\mathcal{C}}(\Gamma, \Delta)$ and $A \in \text{Ty}_{\mathcal{C}}(\Delta)$, $A[\gamma] \in \text{Ty}_{\mathcal{C}}(\Gamma)$. Similarly if $a \in \text{Tm}_{\mathcal{C}}(\Delta, A)$, we write $a[\gamma] \in \text{Tm}_{\mathcal{C}}(\Gamma, A[\gamma])$ for the functorial action of T on a .

We assume that \mathcal{C} has a terminal object $1_{\mathcal{C}}$. Moreover we assume that for each $\Gamma \in \text{Ctx}_{\mathcal{C}}$ and $A \in \text{Ty}_{\mathcal{C}}(\Gamma)$ there exists $\Gamma.A \in \text{Ctx}_{\mathcal{C}}$ with a map $\mathbf{p}_A : \text{Sub}_{\mathcal{C}}(\Gamma.A, \Gamma)$ and a term $\mathbf{q}_A \in \text{Tm}_{\mathcal{C}}(\Gamma.A, A[\mathbf{p}_A])$, such that for every pair $\gamma : \text{Sub}_{\mathcal{C}}(\Delta, \Gamma)$ and $a \in \text{Tm}_{\mathcal{C}}(\Delta, A[\gamma])$ there exists a unique map $\langle \gamma, a \rangle : \text{Sub}_{\mathcal{C}}(\Delta, \Gamma.A)$ such that $\mathbf{p}_A \circ \langle \gamma, a \rangle = \gamma$ and $\mathbf{q}_A[\langle \gamma, a \rangle] = a$.

Note that with the notation $\text{Ty}_{\mathcal{C}}$ and $\text{Tm}_{\mathcal{C}}$ there is no need to explicitly mention the functor T when working with categories with families, and we will often omit it. Given a substitution $\gamma : \Gamma \rightarrow \Delta$, and $A \in \text{Ty}_{\mathcal{C}}(\Delta)$, we write $\gamma \uparrow A$ or γ^+ (when A can be inferred from the context) for the lifting of γ to A : $\langle \gamma \circ \mathbf{p}, \mathbf{q} \rangle : \Gamma.A[\gamma] \rightarrow \Delta.A$.

The indexed category. In [5, 6] it is shown that any cwf \mathcal{C} induces a functor $\mathbf{T} : \mathcal{C}^{\text{op}} \rightarrow \mathbf{Cat}$ assigning to each context Γ the category whose objects are types in $\text{Ty}_{\mathcal{C}}(\Gamma)$ and morphisms from A to B are substitutions $\gamma : \Gamma.A \rightarrow \Gamma.B$ such that $\mathbf{p} \circ \gamma = \mathbf{p}$. (They are in bijection with terms of type $\Gamma.A \vdash B[\mathbf{p}]$.) Any morphism γ in $\mathbf{T}\Gamma$ from a type A to B induces a function on terms of that type written $\{\gamma\} : \text{Tm}_{\mathcal{C}}(\Gamma, A) \rightarrow \text{Tm}_{\mathcal{C}}(\Gamma, B)$ defined by $\{\gamma\}(a) = \mathbf{q}[\gamma \circ \langle \text{id}, a \rangle]$. We will write $\theta : A \cong_{\Gamma} B$ for an isomorphism in $\mathbf{T}\Gamma$.

The functorial action is given by $\mathbf{T}(\gamma)(\varphi) = \langle \mathbf{p}, \mathbf{q}[\varphi \circ \gamma \uparrow A] \rangle : \Gamma.A[\gamma] \rightarrow \Gamma.B[\gamma]$, from which we deduce the action on terms $\{\mathbf{T}(\gamma)(\varphi)\}(a) = \mathbf{q}[\varphi \circ \langle \gamma, a \rangle]$.

► **Definition 6** (Pseudo cwf-morphisms). A pseudo-cwf morphism from a cwf (\mathcal{C}, T) to a cwf (\mathcal{C}', T') is a pair (F, σ) where $F : \mathcal{C} \rightarrow \mathcal{C}'$ is a functor and for each $\Gamma \in \mathcal{C}$, σ_{Γ} is a **Fam**-morphism from $T\Gamma$ to $T'F\Gamma$ preserving the structure up to isomorphism. For example, there are isomorphisms

- $\rho_{\Gamma, A} : F(\Gamma.A) \cong F\Gamma.FA$
- $\theta_{A, \gamma} : F\Gamma.FA[F\gamma] \cong F\Gamma.F(A[\gamma])$ for $\Gamma \vdash \gamma : \Delta$.

satisfying some coherence diagrams, see [6] for the complete definition.

As σ_{Γ} is a **Fam**-morphism from $(\text{Tm}_{\mathcal{C}}(\Gamma, A))_{A \in \text{Ty}_{\mathcal{C}}(\Gamma)}$ to $(\text{Tm}_{\mathcal{D}}(F\Gamma, B))_{B \in \text{Ty}_{\mathcal{D}}(F\Gamma)}$, we write FA for the image of A by $\text{Ty}_{\mathcal{C}}(\Gamma) \rightarrow \text{Ty}_{\mathcal{D}}(F\Gamma)$ induced by σ_{Γ} and Fa for the image of $\Gamma \vdash a : A$ through $\text{Tm}_{\mathcal{C}}(\Gamma, A) \rightarrow \text{Tm}_{\mathcal{D}}(F\Gamma, FA)$ induced by σ_{Γ} .

A pseudo cwf-morphism is strict whenever $\theta_{A, \gamma}$ and $\rho_{\Gamma, A}$ are both identities and $F1 = 1$. Cwfs and strict cwf-morphisms form a category \mathbf{CwF}_s .

► **Definition 7** (Pseudo cwf-transformation). A pseudo cwf-transformation between functors (F, σ) and (G, τ) is a pair (φ, ψ) where $\varphi : F \Rightarrow G$ is a natural transformation, and for each $\Gamma \in \mathcal{C}$ and $A \in \text{Ty}_{\mathcal{C}}(\Gamma)$ $\psi_{\Gamma, A}$ is a type isomorphism $FA \cong_{F\Gamma} GA[\varphi_{\Gamma}]$ satisfying:

$$\varphi_{\Gamma, A} = F(\Gamma.A) \xrightarrow{\rho_F} F\Gamma.FA \xrightarrow{\psi_{\Gamma, A}} F\Gamma.GA[\varphi_{\Gamma}] \xrightarrow{\varphi_{\Gamma}^+} G\Gamma.GA \xrightarrow{\rho_G^{-1}} G(\Gamma.A)$$

We will write \mathbf{CwF} for the resulting 2-category.

3.2 Syntax and inference rules for the free category with families

3.2.1 Raw terms

In this section we define the syntax and inference rules for a minimal dependent type theory with one base type o . This theory is closely related to the generalised algebraic theory of cwfs [8], but here we define it as a usual logical system with a grammar and a collection of inference rules. The grammar has four syntactic categories: contexts Ctx , substitutions Sub , types Ty and terms Tm :

$$\begin{array}{ll} \Gamma ::= 1 \mid \Gamma.A & A ::= o \mid A[\gamma] \\ \gamma ::= \gamma \circ \gamma \mid \text{id}_{\Gamma} \mid \langle \rangle_{\Gamma} \mid \mathbf{p}_A \mid \langle \gamma, a \rangle_A & a ::= a[\gamma] \mid \mathbf{q}_A \end{array}$$

These terms have as few annotations as possible, only what is needed to recover the domain and codomain of a substitution, the context of a type, and the type of a term:

$$\begin{array}{ll}
\text{dom}(\gamma \circ \gamma') = \text{dom}(\gamma') & \text{cod}(\gamma \circ \gamma') = \text{cod}(\gamma) \\
\text{dom}(\text{id}_\Gamma) = \Gamma & \text{cod}(\text{id}_\Gamma) = \Gamma \\
\text{dom}(\langle \rangle_\Gamma) = \Gamma & \text{cod}(\langle \rangle_\Gamma) = 1 \\
\text{dom}(\mathbf{p}_A) = \text{ctx-of}(A).A & \text{cod}(\mathbf{p}_A) = \text{ctx-of}(A) \\
\text{dom}(\langle \gamma, a \rangle_A) = \text{dom}(\gamma) & \text{cod}(\langle \gamma, a \rangle_A) = \text{cod}(\gamma).A \\
\\
\text{ctx-of}(o) = 1 & \text{type-of}(a[\gamma]) = (\text{type-of}(a))[\gamma] \\
\text{ctx-of}(A[\gamma]) = \text{dom}(\gamma) & \text{type-of}(\mathbf{q}_A) = A[\mathbf{p}_A]
\end{array}$$

These functions will be used in the freeness proof.

3.2.2 Inference rules

We simultaneously inductively define four families of partial equivalence relations (pers) for the four forms of equality judgment:

$$\Gamma = \Gamma' \vdash \quad \Gamma \vdash A = A' \quad \Delta \vdash \gamma = \gamma' : \Gamma \quad \Gamma \vdash a = a' : A$$

In the inference rules which generate these pers we will use the following abbreviations for the basic judgment forms: $\Gamma \vdash$ abbreviates $\Gamma = \Gamma \vdash$, $\Gamma \vdash A$ abbreviates $\Gamma \vdash A = A$, $\Delta \vdash \gamma : \Gamma$ abbreviates $\Delta \vdash \gamma = \gamma : \Gamma$, and $\Gamma \vdash a : A$ abbreviates $\Gamma \vdash a = a : A$.

Per-rules for the four forms of judgments:

$$\begin{array}{c}
\frac{\Gamma = \Gamma' \vdash \quad \Gamma' = \Gamma'' \vdash}{\Gamma = \Gamma'' \vdash} \quad \frac{\Gamma = \Gamma' \vdash}{\Gamma' = \Gamma \vdash} \quad \frac{\Delta \vdash \gamma = \gamma' : \Gamma \quad \Delta \vdash \gamma' = \gamma'' : \Gamma}{\Delta \vdash \gamma = \gamma'' : \Gamma} \quad \frac{\Delta \vdash \gamma = \gamma' : \Gamma}{\Delta \vdash \gamma' = \gamma : \Gamma} \\
\\
\frac{\Gamma \vdash A = A' \quad \Gamma \vdash A' = A''}{\Gamma \vdash A = A''} \quad \frac{\Gamma \vdash A = A'}{\Gamma \vdash A' = A} \quad \frac{\Gamma \vdash a = a' : A \quad \Gamma \vdash a' = a'' : A}{\Gamma \vdash a = a'' : A} \\
\\
\frac{\Gamma \vdash a = a' : A}{\Gamma \vdash a' = a : A}
\end{array}$$

Preservation rules for judgments:

$$\begin{array}{c}
\frac{\Gamma = \Gamma' \vdash \quad \Delta = \Delta' \vdash \quad \Gamma \vdash \gamma = \gamma' : \Delta}{\Gamma' \vdash \gamma = \gamma' : \Delta'} \quad \frac{\Gamma = \Gamma' \vdash \quad \Gamma \vdash A = A'}{\Gamma' \vdash A = A'} \\
\\
\frac{\Gamma = \Gamma' \vdash \quad \Gamma \vdash A = A' \quad \Gamma \vdash a = a' : A}{\Gamma' \vdash a = a' : A'}
\end{array}$$

Congruence rules for operators:

$$\begin{array}{c}
\frac{\Gamma \vdash \delta = \delta' : \Delta \quad \Delta \vdash \gamma = \gamma' : \Theta}{\Gamma \vdash \gamma \circ \delta = \gamma' \circ \delta' : \Theta} \quad \frac{\Gamma = \Gamma' \vdash}{\Gamma \vdash \text{id}_\Gamma = \text{id}_{\Gamma'} : \Gamma} \quad \frac{\Gamma \vdash A = A' \quad \Delta \vdash \gamma = \gamma' : \Gamma}{\Delta \vdash A[\gamma] = A'[\gamma']} \\
\frac{\Gamma \vdash a = a' : A \quad \Delta \vdash \gamma = \gamma' : \Gamma}{\Delta \vdash a[\gamma] = a'[\gamma'] : A[\gamma']} \quad \frac{}{1 = 1 \vdash} \quad \frac{\Gamma = \Gamma' \vdash}{\Gamma \vdash \langle \rangle_\Gamma = \langle \rangle_{\Gamma'} : 1} \quad \frac{\Gamma = \Gamma' \vdash \quad \Gamma \vdash A = A'}{\Gamma.A = \Gamma'.A' \vdash} \\
\frac{\Gamma \vdash A = A'}{\Gamma.A \vdash \mathbf{p}_A = \mathbf{p}_{A'} : \Gamma} \quad \frac{\Gamma \vdash A = A'}{\Gamma.A \vdash \mathbf{q}_A = \mathbf{q}_{A'} : A[\mathbf{p}_A]} \\
\frac{\Gamma \vdash A = A' \quad \Delta \vdash \gamma = \gamma' : \Gamma \quad \Delta \vdash a = a' : A[\gamma]}{\Delta \vdash \langle \gamma, a \rangle_A = \langle \gamma', a' \rangle_{A'} : \Gamma.A}
\end{array}$$

Conversion rules:

$$\begin{array}{c}
\frac{\Delta \vdash \theta : \Theta \quad \Gamma \vdash \delta : \Delta \quad \Xi \vdash \gamma : \Gamma}{(\theta \circ \delta) \circ \gamma = \theta \circ (\delta \circ \gamma)} \quad \frac{\Gamma \vdash \gamma : \Delta}{\Gamma \vdash \gamma = \text{id}_\Delta \circ \gamma : \Delta} \quad \frac{\Gamma \vdash \gamma : \Delta}{\Gamma \vdash \gamma = \gamma \circ \text{id}_\Gamma : \Delta} \\
\frac{\Gamma \vdash A \quad \Delta \vdash \gamma : \Gamma \quad \Theta \vdash \delta : \Delta}{\Theta \vdash A[\gamma \circ \delta] = (A[\gamma])[\delta]} \quad \frac{\Gamma \vdash A}{\Gamma \vdash A[\text{id}_\Gamma] = A} \quad \frac{\Gamma \vdash a : A \quad \Delta \vdash \gamma : \Gamma \quad \Theta \vdash \delta : \Delta}{\Theta \vdash a[\gamma \circ \delta] = (a[\gamma])[\delta] : (A[\gamma])[\delta]} \\
\frac{\Gamma \vdash a : A}{\Gamma \vdash a[\text{id}_\Gamma] = a : A} \quad \frac{\Gamma \vdash \gamma : 1}{\Gamma \vdash \gamma = \langle \rangle_\Gamma : 1} \quad \frac{\Gamma \vdash A \quad \Delta \vdash \gamma : \Gamma \quad \Delta \vdash a : A[\gamma]}{\Delta \vdash \mathbf{p}_A \circ \langle \gamma, a \rangle_A = \gamma : \Gamma} \\
\frac{\Gamma \vdash A \quad \Delta \vdash \gamma : \Gamma \quad \Delta \vdash a : A[\gamma]}{\Delta \vdash \mathbf{q}_A[\langle \gamma, a \rangle_A] = a : A[\gamma]} \quad \frac{\Delta \vdash \gamma : \Gamma.A}{\Delta \vdash \gamma = \langle \mathbf{p}_A \circ \gamma, \mathbf{q}_A[\gamma] \rangle_A : \Gamma.A}
\end{array}$$

Rule for the base type:

$$\frac{}{1 \vdash o = o}$$

3.2.3 The syntactic cwf \mathcal{T}

We can now define a term model as the syntactic cwf obtained by the well-formed contexts, etc, modulo judgmental equality. We use brackets for equivalence classes in this definition. (Note that brackets are also used for substitution in types and terms, but this should not cause confusion since we will soon drop the equivalence class brackets.)

► **Definition 8.** The term model \mathcal{T} is given by:

- $\text{Ctx}_{\mathcal{T}} = \{\Gamma \mid \Gamma \vdash \} / \equiv^c$, where $\Gamma \equiv^c \Gamma'$ if $\Gamma = \Gamma' \vdash$ is derivable.
- $\text{Sub}_{\mathcal{T}}([\Gamma], [\Delta]) = \{\gamma \mid \Gamma \vdash \gamma : \Delta\} / \equiv_{\Delta}^{\Gamma}$ where $\gamma \equiv_{\Delta}^{\Gamma} \gamma'$ iff $\Gamma \vdash \gamma = \gamma' : \Delta$ is derivable. Note that this makes sense since it only depends on the equivalence class of Γ (morphisms and morphism equality are preserved by object equality).
- $\text{Ty}_{\mathcal{T}}([\Gamma]) = \{A \mid \Gamma \vdash A\} / \equiv^{\Gamma}$ where $A \equiv^{\Gamma} B$ if $\Gamma \vdash A = B$.
- $\text{Tm}_{\mathcal{T}}([\Gamma], [A]) = \{a \mid \Gamma \vdash a : A\} / \equiv_A^{\Gamma}$ where $a \equiv_A^{\Gamma} a'$ if $\Gamma \vdash a = a' : A$.

The cwf-operations on \mathcal{T} can now be defined in a straightforward way. For example, if $\Delta \vdash \theta : \Theta$, $\Gamma \vdash \delta : \Delta$, we define $[\theta] \circ_{\mathcal{T}} [\delta] = [\theta \circ \delta]$, which is well-defined since composition preserves equality.

3.3 Freeness of \mathcal{T}

We shall now show that \mathcal{T} is the free cwf on one base type, in the sense that given a cwf \mathcal{C} and a type $A \in \text{Ty}_{\mathcal{C}}(1)$, there exists a unique strict cwf morphism $\mathcal{T} \rightarrow \mathcal{C}$ which maps $[o]$ to A . Such a morphism can be defined by first defining a partial function for each sort of raw terms (where Ctx denotes the set of raw contexts, Sub the set of raw substitutions, and so on defined by the grammar of Section 3.2.1), cf Streicher [14].

$$\begin{aligned} \llbracket - \rrbracket & : \text{Ctx} \rightarrow \text{Ctx}_{\mathcal{C}} \\ \llbracket - \rrbracket_{\Gamma, \Delta} & : \text{Sub} \rightarrow \text{Sub}_{\mathcal{C}}(\llbracket \Gamma \rrbracket, \llbracket \Delta \rrbracket) \\ \llbracket - \rrbracket_{\Gamma} & : \text{Ty} \rightarrow \text{Ty}_{\mathcal{C}}(\llbracket \Gamma \rrbracket) \\ \llbracket - \rrbracket_{\Gamma, A} & : \text{Tm} \rightarrow \text{Tm}_{\mathcal{C}}(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket_{\Gamma}) \end{aligned}$$

These functions are defined by mutual induction on the structure of raw terms:

$$\begin{aligned} \llbracket 1 \rrbracket & = 1_{\mathcal{C}} & \llbracket \Gamma.A \rrbracket & = \llbracket \Gamma \rrbracket ._{\mathcal{C}} \llbracket A \rrbracket \\ \llbracket \gamma' \circ \gamma \rrbracket_{\Gamma, \Theta} & = \llbracket \gamma' \rrbracket_{\Delta, \Theta} \circ_{\mathcal{C}} \llbracket \gamma \rrbracket_{\Gamma, \Delta} & \llbracket \text{id}_{\Gamma} \rrbracket_{\Gamma, \Gamma} & = \text{id}_{\mathcal{C}} \llbracket \Gamma \rrbracket \\ \llbracket \langle \gamma, a \rangle_A \rrbracket_{\Gamma, \Delta.A} & = \langle \llbracket \gamma \rrbracket_{\Gamma, \Delta}, \llbracket a \rrbracket_{\Gamma, A[\gamma]} \rangle & \llbracket \langle \rangle_{\Gamma} \rrbracket_{\Gamma, 1} & = \langle \rangle_{\Gamma} \\ \llbracket A[\gamma] \rrbracket_{\Gamma} & = \llbracket A \rrbracket_{\Delta} \llbracket \llbracket \gamma \rrbracket_{\Gamma, \Delta} \rrbracket & \llbracket a[\gamma] \rrbracket_{\Gamma, A[\gamma]} & = \llbracket a \rrbracket_{\Delta, A} \llbracket \llbracket \gamma \rrbracket_{\Gamma, \Delta} \rrbracket \\ \llbracket p_A \rrbracket_{\Gamma, A, \Gamma} & = p_A & \llbracket q_A \rrbracket_{\Gamma, A, A[p]} & = q_A \\ \llbracket o \rrbracket_1 & = A \end{aligned}$$

Note that $\Delta = \text{dom}(\gamma') = \text{cod}(\gamma)$ in the equation for \circ , etc.

We then prove by induction on the inference rules that:

- **Lemma 9.** ■ *If $\Gamma = \Gamma' \vdash$, then $\llbracket \Gamma \rrbracket = \llbracket \Gamma' \rrbracket$ and both are defined.*
- *If $\Gamma \vdash \gamma = \gamma' : \Delta$, then $\llbracket \gamma \rrbracket_{\Gamma, \Delta} = \llbracket \gamma' \rrbracket_{\Gamma, \Delta}$ and both are defined.*
- *If $\Gamma \vdash A = A'$, then $\llbracket A \rrbracket_{\Gamma} = \llbracket A' \rrbracket_{\Gamma}$ and both are defined.*
- *If $\Gamma \vdash a = a' : A$, then $\llbracket a \rrbracket_{\Gamma, A} = \llbracket a' \rrbracket_{\Gamma, A}$ and both are defined.*

Hence the partial interpretation lifts to the quotient of syntax by judgmental equality:

$$\begin{aligned} \overline{\llbracket - \rrbracket} & : \text{Ctx}_{\mathcal{T}} \rightarrow \text{Ctx}_{\mathcal{C}} \\ \overline{\llbracket - \rrbracket}_{\llbracket \Gamma \rrbracket, \llbracket \Delta \rrbracket} & : \text{Sub}_{\mathcal{T}}(\llbracket \Gamma \rrbracket, \llbracket \Delta \rrbracket) \rightarrow \text{Sub}_{\mathcal{C}}(\llbracket \llbracket \Gamma \rrbracket \rrbracket, \llbracket \llbracket \Delta \rrbracket \rrbracket) \\ \overline{\llbracket - \rrbracket}_{\llbracket \Gamma \rrbracket} & : \text{Ty}_{\mathcal{T}}(\llbracket \Gamma \rrbracket) \rightarrow \text{Ty}_{\mathcal{C}}(\llbracket \llbracket \Gamma \rrbracket \rrbracket) \\ \overline{\llbracket - \rrbracket}_{\llbracket \Gamma \rrbracket, \llbracket A \rrbracket} & : \text{Tm}_{\mathcal{T}}(\llbracket \Gamma \rrbracket, \llbracket A \rrbracket) \rightarrow \text{Tm}_{\mathcal{C}}(\llbracket \llbracket \Gamma \rrbracket \rrbracket, \llbracket \llbracket A \rrbracket \rrbracket_{\llbracket \Gamma \rrbracket}) \end{aligned}$$

This defines a strict cwf morphism $\mathcal{T} \rightarrow \mathcal{C}$ which maps $[o]$ to A . It is easy to check that it is the unique such strict cwf-morphism. Hence we have proved:

- **Theorem 10.** *\mathcal{T} is the free cwf on one object, that is, for every other cwf \mathcal{C} and $A \in \text{Ty}_{\mathcal{C}}(1)$ there is a unique strict cwf morphism $\mathcal{T} \rightarrow \mathcal{C}$ which maps $[o]$ to A .*

From now on we will uniformly drop the equivalence class brackets and for example write Γ for $\llbracket \Gamma \rrbracket$. There should be no risk of confusion, but we remark that proofs by induction on syntax and inference rules are on representatives rather than equivalence classes.

3.4 Bifreeness of \mathcal{T}

We recall that an object I is bi-initial in a 2-category iff for any object A there exists an arrow $I \rightarrow A$ and for any two arrows $f, g : I \rightarrow A$ there exists a unique 2-cell $\theta : f \Rightarrow g$. It follows that θ is invertible. It also follows that bi-initial objects are equivalent.

► **Definition 11.** A cwf \mathcal{C} is bifree on one base type iff it is bi-initial in the 2-category \mathbf{CwF}^o :

- *Objects:* Pairs $(\mathcal{C}, o_{\mathcal{C}})$ of a CwF and a chosen $o_{\mathcal{C}} \in \text{Ty}_{\mathcal{C}}(1)$,
- *1-cells from $(\mathcal{C}, o_{\mathcal{C}})$ to $(\mathcal{D}, o_{\mathcal{D}})$:* pseudo cwf-morphisms $F : \mathcal{C} \rightarrow \mathcal{D}$ such that there exists $\varphi_F : F(o_{\mathcal{C}}) \cong o_{\mathcal{D}}$ in $\text{Ty}_{\mathcal{D}}(1)$,
- *2-cells from F to G with type $(\mathcal{C}, o_{\mathcal{C}}) \rightarrow (\mathcal{D}, o_{\mathcal{D}})$:* pseudo cwf-transformations (φ, ψ) from F to G satisfying $\psi_{o_{\mathcal{C}}} = \alpha_G^{-1} \circ \alpha_F : F(o_{\mathcal{C}}) \cong G(o_{\mathcal{C}})$.

► **Theorem 12.** \mathcal{T} is a bifree cwf on one base type.

We have showed that for every cwf \mathcal{C} , $A \in \text{Ty}_{\mathcal{C}}(1)$, the interpretation $\overline{[-]}$ is a strict cwf-morphism mapping o to A . Hence it is a morphism in \mathbf{CwF}^o . It remains to show that for any other morphism $F : \mathcal{T} \rightarrow \mathcal{C}$ in \mathbf{CwF}^o , there is a unique 2-cell (pseudo cwf-transformation) $(\varphi, \psi) : \overline{[-]} \rightarrow F$, which happens to be an isomorphism. This asymmetric version of bi-initiality is equivalent to that given above because the 2-cell we build is an isomorphism.

Existence of (φ, ψ)

We construct (φ, ψ) by induction on the inference rules and simultaneously prove their naturality properties:

- If $\Gamma = \Gamma' \vdash$, then there exist $\varphi_{\Gamma} = \varphi_{\Gamma'} : \overline{[\Gamma]} \cong F\Gamma$.
- If $\Gamma \vdash A = A'$, then there exist $\psi_A = \psi_{A'} : \overline{[\Gamma.A]} \cong_{\overline{[\Gamma]}} \overline{[\Gamma]}.FA[\varphi_{\Gamma}]$.
- If $\Gamma \vdash \gamma = \gamma' : \Delta$, then $F\gamma \circ \varphi_{\Gamma} = \varphi_{\Delta} \circ \overline{[\gamma]}$
- If $\Gamma \vdash a = a' : A$, then $\{\psi_A\}(\overline{[a]}) = Fa[\varphi_{\Gamma}]$

Since it also follows that $\varphi_{\Gamma.A} = \rho^{-1} \circ \varphi_{\Gamma}^+ \circ \psi_A$ we conclude that (φ, ψ) is a pseudo cwf-transformation. For space reasons we only present the proofs of the first two items and refer the reader to the long version of the paper [4] for the other two.

Empty context. F preserves terminal objects, thus we let $\phi_1 : \overline{[1]} = 1_{\mathcal{C}} \cong F1$.

Context extension. By induction, we have $\psi_A : \overline{[A]} \cong_{\Gamma} FA[\varphi_{\Gamma}]$. We define $\varphi_{\Gamma.A}$ as the following composition of isomorphisms:

$$\varphi_{\Gamma.A} = \overline{[\Gamma.A]} \xrightarrow{\psi_A} \overline{[\Gamma]}.FA[\varphi_{\Gamma}] \xrightarrow{\langle \varphi_{\Gamma}, \mathbf{q} \rangle} F\Gamma.FA \xrightarrow{\rho_{\Gamma.A}^{-1}} F(\Gamma.A)$$

Type substitution. Let $\Gamma \vdash \gamma : \Delta$ and $\Delta \vdash A$. By induction we get $\varphi_{\Delta} \circ \overline{[\gamma]} = F\gamma \circ \varphi_{\Gamma}$ and $\psi_A : \overline{[A]} \cong_{\Delta} FA[\varphi_{\Delta}]$. Since \mathbf{T} is a functor, $\mathbf{T}\gamma$ is a functor from $\mathbf{T}\Delta$ to $\mathbf{T}\Gamma$ thus,

$$\mathbf{T}(\overline{[\gamma]})(\psi_A) : \overline{[A[\gamma]]} \cong_{\Gamma} FA[\varphi_{\Delta} \circ \gamma] = FA[F\gamma][\varphi_{\Gamma}]$$

by induction hypothesis on γ . So we define

$$\psi_{A[\gamma]} = \mathbf{T}(\varphi_{\Gamma})(\theta_{A,\gamma}) \circ \mathbf{T}(\overline{[\gamma]})(\psi_A) : \overline{[A[\gamma]]} \cong_{\overline{[\Gamma]}} (F(A[\gamma]))[\varphi_{\Gamma}]$$

Using the previous case we can get a simpler equation for $\varphi_{\Gamma.A[\gamma]}$:

$$\varphi_{\Gamma.A[\gamma]} = \langle \varphi_{\Gamma} \circ \mathbf{p}, \mathbf{q}[\rho \circ \varphi_{\Delta.A} \circ \gamma \uparrow A] \rangle : \overline{[\Gamma.A[\gamma]]} \rightarrow F(\Gamma.A[\gamma])$$

Base type. By definition, F is equipped with $\alpha_F : \overline{[o]} \cong F(o)$. We define $\psi_o = \alpha_F^{-1} : \overline{[o]} \cong F(o)$ in $\text{Ty}_{\mathcal{C}}(1)$.

Uniqueness of (φ, ψ)

Let $(\varphi', \psi') : \overline{\llbracket \cdot \rrbracket} \rightarrow F$ be another pseudo cwf-transformation in \mathbf{CwF}^o . We prove the following by induction:

- If $\Gamma \vdash$, then $\varphi_\Gamma = \varphi'_\Gamma$
- If $\Gamma \vdash A$, then $\psi_A = \psi'_A$

Empty context. There is a unique morphism between the terminal objects $\overline{\llbracket 1 \rrbracket}$ and $F1$, so $\varphi_1 = \varphi'_1$.

Context extension. Assume by induction $\varphi_\Gamma = \varphi'_\Gamma$ and $\psi_A = \psi'_A$. By the coherence law of pseudo cwf-transformations, we have $\varphi'_{\Gamma.A} = \rho^{-1} \circ \varphi'_\Gamma \circ \psi'_A$ from which the equality $\varphi'_{\Gamma.A} = \varphi_{\Gamma.A}$ follows.

Type substitution. Assume we have $\Delta \vdash A$ and $\Gamma \vdash \gamma : \Delta$, and consider $\psi_{A[\gamma]}$ and $\psi'_{A[\gamma]}$. By definition of pseudo cwf-transformations, one has $\mathbf{T}(\varphi'_\Gamma)(\theta_{A,\gamma}^{-1}) \circ \psi'_{A[\gamma]} = \mathbf{T}(F\gamma)(\psi'_A)$. Since we know $\varphi_\Gamma = \varphi'_\Gamma$ we know φ'_Γ is an isomorphism and thus $\psi'_{A[\gamma]}$ depends only on φ'_Γ and ψ'_A from which it follows that $\psi'_{A[\gamma]} = \psi_{A[\gamma]}$.

Base type. The definition of 2-cells in \mathbf{CwF}^o entails $\psi'_o = \alpha_F^{-1} : \overline{\llbracket o \rrbracket} \rightarrow F(\llbracket o \rrbracket)$.

4 A free lccc**4.1 From cwfs to lcccs**

We now extend our cwf-calculus with extensional I-types, N_1, Σ , and Π and prove that it yields a free cwf supporting these type formers. To show that this yields a free lccc we apply the biequivalence [6] between lcccs and *democratic* cwfs supporting these type formers.

► **Definition 13** (Democratic cwfs). A cwf \mathcal{C} is democratic when for each context Γ there is a type $\bar{\Gamma} \in \text{Ty}(1)$ such that $\Gamma \cong 1.\bar{\Gamma}$. A pseudo cwf morphism $F : \mathcal{C} \rightarrow \mathcal{D}$ between democratic cwf preserves democracy when there is an isomorphism $F(\bar{\Gamma}) \cong \overline{F\bar{\Gamma}}[\langle \rangle_\Gamma]$ satisfying a coherence diagram stated in [6] (Definition 8).

The free cwf with N_1, Σ , and Π -types is democratic since the empty context can be represented by the unit type N_1 and context extension by a Σ -type.

4.2 Cwfs with support for type constructors

A cwf supports a certain type former when it has extra structure and equations corresponding to the to formation, introduction, elimination, and equality rules for the type former in question. We only spell out what it means for a cwf to support and preserve extensional identity types and refer the reader to [6] for the definitions wrt Σ - and Π -types. The definition of what it means to support and preserve N_1 is analogous.

► **Definition 14** (Cwf with identity types). A cwf \mathcal{C} is said to support extensional identity types when for each $a, a' \in \text{Tm}_{\mathcal{C}}(\Gamma, A)$ there is a type $I(A, a, a') \in \text{Ty}_{\mathcal{C}}(\Gamma)$ satisfying the following condition:

- $I(A, a, a')[\gamma] = I(A[\gamma], a[\gamma], a'[\gamma])$ for any $\gamma : \Delta \rightarrow \Gamma$
- For $a \in \text{Tm}_{\mathcal{C}}(\Gamma, A)$, there exists $r(a) \in \text{Tm}_{\mathcal{C}}(\Gamma, I(A, a, a))$. Moreover, if $c \in \text{Tm}_{\mathcal{C}}(\Gamma, I(A, a, a'))$ then $a = a'$ and $c = r(a)$.

A pseudo cwf morphism F preserves identity types when

$$I(FA, Fa, Fa') \cong_{\Gamma} F(I(A, a, a')).$$

We write $\mathbf{Cwf}_d^{I,\Sigma,\Pi}$ for the 2-category of democratic cwf's supporting I, Σ , and Π with morphisms preserving them, and $\mathbf{Cwf}_{s,d}^{\Sigma,\Pi,I}$ for the strict version. Note that by democracy, any democratic cwf has a unit type representing the empty context.

Σ and Π are functorial and preserve isomorphisms:

► **Lemma 15** (Functoriality of Σ). *Let $f_A : A \cong_{\Gamma} A'$ and $f_B : B \cong_{\Gamma.A} B'[f_A]$ with $\Gamma.A \vdash B$ and $\Gamma.A' \vdash B'$. Then $\Sigma(A, B) \cong \Sigma(A', B')$ functorially: if $g_A : A' \cong A''$ and $g_B : B' \cong B''[g_A]$, then $\Sigma(g_A \circ f_A, g_B \circ f_B) = \Sigma(g_A, g_B) \circ \Sigma(f_A, f_B) : \Sigma(A, B) \rightarrow \Sigma(A'', B'')$.*

► **Lemma 16** (Functoriality of Π). *Let $f_A : A \cong_{\Gamma} A'$ and $f_B : B \cong_{\Gamma.A} B'[f_A]$. Then there is a type isomorphism $\Pi(f_A, f_B) : \Pi(A, B) \cong_{\Gamma} \Pi(A', B')$ functorially.*

4.3 The syntactic cwf with extensional I, N_1, Σ , and Π

We extend the grammar and the set of inference rules with rules for I, N_1, Σ , and Π -types:

$$\begin{aligned} A & ::= \dots \mid I(a, a) \mid N_1 \mid \Sigma(A, A) \mid \Pi(A, A) \\ a & ::= \dots \mid r(a) \mid 0_1 \mid \text{fst}(A, a) \mid \text{snd}(A, A, a) \mid \text{pair}(A, A, a, a) \mid \text{ap}(A, A, a, a) \mid \lambda(A, a) \end{aligned}$$

For each type we define its context:

$$\begin{aligned} \text{ctx-of}(I(a, a')) &= \text{ctx-of}(\text{type-of}(a)) & \text{ctx-of}(\Sigma(A, B)) &= \text{ctx-of}(A) \\ \text{ctx-of}(\Pi(A, B)) &= \text{ctx-of}(A) \end{aligned}$$

For each term we define its type:

$$\begin{aligned} \text{type-of}(\text{fst}(A, c)) &= A & \text{type-of}(r(a)) &= I(a, a) \\ \text{type-of}(\text{snd}(A, B, c)) &= B \langle \text{id}_{\text{ctx-of}(A)}, \text{fst}(A, c) \rangle & \text{type-of}(\lambda(A, c)) &= \Pi(A, \text{type-of}(c)) \\ \text{type-of}(\text{pair}(A, B, a, b)) &= \Sigma(A, B) & \text{type-of}(\text{ap}(A, B, c, a)) &= B \langle \text{id}_{\text{ctx-of}(A)}, a \rangle \\ \text{type-of}(\text{pair}(A, B, a, b)) &= \Sigma(A, B) \end{aligned}$$

4.3.1 Inference rules

Rules for I -types:

$$\begin{aligned} \frac{\Gamma \vdash a = a' : A \quad \Gamma \vdash b = b' : A}{\Gamma \vdash I(a, b) = I(a', b')} & \quad \frac{\Gamma \vdash a = a' : A}{\Gamma \vdash r(a) = r(a') : I(a, a')} & \quad \frac{\Gamma \vdash c : I(a, a')}{\Gamma \vdash a = a' : \text{type-of}(a)} \\ \frac{\Gamma \vdash c : I(a, a')}{\Gamma \vdash c = r(a) : I(a, a')} & \quad \frac{\Gamma \vdash a : A \quad \Gamma \vdash a' : A \quad \Delta \vdash \gamma : \Gamma}{\Gamma \vdash I(a, a')[\gamma] = I(a[\gamma], a'[\gamma])} \end{aligned}$$

Rules for N_1 :

$$\frac{}{\vdash N_1} \quad \frac{}{\vdash 0_1 : N_1} \quad \frac{\vdash a : N_1}{\vdash a = 0_1 : N_1}$$

Rules for Σ -types:

$$\frac{\Gamma \vdash A = A' \quad \Gamma.A \vdash B = B'}{\Gamma \vdash \Sigma(A, B) = \Sigma(A', B')}$$

$$\frac{\Gamma \vdash A = A' \quad \Gamma \vdash c = c' : \Sigma(A, B)}{\Gamma \vdash \text{fst}(A, c) = \text{fst}(A', c') : A}$$

$$\frac{\Gamma \vdash A = A' \quad \Gamma.A \vdash B = B' \quad \Gamma \vdash c = c' : \Sigma(A, B)}{\Gamma \vdash \text{snd}(A, B, c) = \text{snd}(A', B', c') : B \langle \text{id}_\Gamma, \text{fst}(A, c) \rangle}$$

$$\frac{\Gamma \vdash A = A' \quad \Gamma.A \vdash B = B' \quad \Gamma \vdash a = a' : A' \quad \Gamma \vdash b = b' : B \langle \text{id}_\Gamma, \text{fst}(A, c) \rangle}{\Gamma \vdash \text{pair}(A, B, a, b) = \text{pair}(A', B', a', b') : \Sigma(A, B)}$$

$$\frac{\Gamma \vdash A \quad \Gamma.A \vdash B \quad \Gamma \vdash a : A' \quad \Gamma \vdash b : B \langle \text{id}_\Gamma, \text{fst}(A, c) \rangle}{\Gamma : \text{fst}(A, \text{pair}(A, B, a, b)) = a : A}$$

$$\frac{\Gamma \vdash A \quad \Gamma.A \vdash B \quad \Gamma \vdash a : A' \quad \Gamma \vdash b : B \langle \text{id}_\Gamma, \text{fst}(A, c) \rangle}{\Gamma : \text{snd}(A, B, \text{pair}(A, B, a, b)) = b : B \langle \text{id}_\Gamma, \text{fst}(A, c) \rangle}$$

$$\frac{\Gamma \vdash A \quad \Gamma.A \vdash B \quad \Gamma \vdash a : A \quad \Gamma \vdash b : B[\langle \text{id}_\Gamma, \text{fst}(A, c) \rangle]}{\Gamma \vdash c = \text{pair}(A, B, \text{fst}(A, c), \text{snd}(A, B, c)) : \Sigma(A, B)}$$

$$\frac{\Gamma \vdash A \quad \Gamma.A \vdash B \quad \Delta \vdash \gamma : \Gamma}{\Gamma \vdash \Sigma(A, B)[\gamma] = \Sigma(A[\gamma], B[\gamma^+])}$$

$$\frac{\Gamma \vdash A \quad \Gamma \vdash c : \Sigma(A, B) \quad \Delta \vdash \gamma : \Gamma}{\Gamma \vdash \text{fst}(A, c)[\gamma] = \text{fst}(A[\gamma], c[\gamma]) : A}$$

$$\frac{\Gamma \vdash A \quad \Gamma.A \vdash B \quad \Gamma \vdash c : \Sigma(A, B) \quad \Delta \vdash \gamma : \Gamma}{\Gamma \vdash \text{snd}(A, B, c)[\gamma] = \text{snd}(A[\gamma], B[\gamma^+], c[\gamma]) : B[\langle \gamma, \text{fst}(A, c)[\gamma] \rangle]}$$

$$\frac{\Gamma \vdash A \quad \Gamma.A \vdash B \quad \Gamma \vdash a : A \quad \Gamma \vdash b : B[\langle \text{id}_\Gamma, \text{fst}(A, c) \rangle] \quad \Delta \vdash \gamma : \Gamma}{\Gamma \vdash \text{pair}(A, B, a, b)[\gamma] = \text{pair}(A[\gamma], B[\gamma^+], a[\gamma], b[\gamma]) : \Sigma(A, B)[\gamma]}$$

Rules for Π -types:

$$\frac{\Gamma \vdash A = A' \quad \Gamma.A \vdash B = B'}{\Gamma \vdash \Pi(A, B) = \Pi(A', B')}$$

$$\frac{\Gamma \vdash A = A' \quad \Gamma.A \vdash b = b' : B}{\Gamma \vdash \lambda(A, b) = \lambda(A', b') : \Pi(A, B)}$$

$$\frac{\Gamma \vdash A = A' \quad \Gamma.A \vdash B = B' \quad \Gamma \vdash c = c' : \Pi(A, B) \quad \Gamma \vdash a = a' : A}{\Gamma \vdash \text{app}(A, B, c, a) = \text{app}(A', B', c', a') : B[\langle \text{id}, a \rangle]}$$

$$\frac{\Gamma \vdash c : \Pi(A, B) \quad \Gamma \vdash a : A}{\Gamma \vdash \text{app}(A, B, \lambda(A, b), a) = b[\langle \text{id}, a \rangle] : B[\langle \text{id}, a \rangle]}$$

$$\frac{\Gamma \vdash c : \Pi(A, B)}{\Gamma \vdash \lambda(\text{app}(c[p], q)) = c : \Pi(A, B)}$$

$$\frac{\Gamma \vdash A \quad \Gamma.A \vdash B \quad \Delta \vdash \gamma : \Gamma}{\Delta \vdash \Pi(A, B)[\gamma] = \Pi(A[\gamma], B[\gamma^+])}$$

$$\frac{\Gamma \vdash c : \Pi(A, B) \quad \Delta \vdash \gamma : \Gamma}{\Delta \vdash \lambda(b)[\gamma] = \lambda(b[\gamma^+]) : \Pi(A, B)[\gamma]}$$

$$\frac{\Gamma \vdash c : \Pi(A, B) \quad \Gamma \vdash a : A \quad \Delta \vdash \gamma : \Gamma}{\Delta \vdash \text{app}(c, a)[\gamma] = \text{app}(c[\gamma], a[\gamma]) : B[\langle \gamma, a[\gamma] \rangle]}$$

4.3.2 The syntactic cwf supporting \mathbf{I} , \mathbf{N}_1 , $\mathbf{\Sigma}$, and $\mathbf{\Pi}$

It is straightforward to extend the definition of the term model \mathcal{T} with \mathbf{I} , \mathbf{N}_1 , $\mathbf{\Sigma}$, and $\mathbf{\Pi}$ -types, to form a cwf $\mathcal{T}^{\mathbf{I}, \mathbf{N}_1, \mathbf{\Sigma}, \mathbf{\Pi}}$ supporting these type constructors. As we already explained it is democratic.

We want to show that $\mathcal{T}^{I, N_1, \Sigma, \Pi}$ is free, not only in the 2-category of cwfs supporting I, Σ, Π but in the subcategory of the democratic ones. (Democracy entails that N_1 is also supported.) It is straightforward to extend the interpretation functor and prove its uniqueness. It is also easy to check that it preserves democracy.

► **Theorem 17.** $\mathcal{T}^{I, N_1, \Sigma, \Pi}$ is the free democratic cwf supporting I, Σ, Π on one object.

4.4 Bifreeness of $\mathcal{T}^{I, N_1, \Sigma, \Pi}$

We now prove the key result:

► **Theorem 18.** $\mathcal{T}^{I, N_1, \Sigma, \Pi}$ is the bifree democratic cwf supporting I, Σ, Π on one object.

This means that $\mathcal{T}^{I, N_1, \Sigma, \Pi}$ is bi-initial in the 2-category $\mathbf{CwF}_d^{I, \Sigma, \Pi, o}$ where objects are democratic cwfs which support I, Σ, Π , and a base type o , and where morphisms preserve these type formers up to isomorphism.

4.4.1 Existence

We resume our inductive proof from Section 3.4 with the cases for I, N_1, Σ , and Π .

Unit type. Since F preserves democracy and terminal object it follows that $1.F(N_1) = 1.F(\bar{1}) \cong \overline{F\bar{1}} \cong \bar{1}$.

Identity type. Assume $\Gamma \vdash a, a' : A$. By induction, we have $\psi_A : \overline{[A]} \cong_{\overline{[\Gamma]}} FA[\varphi_\Gamma]$. We know I-types preserve isomorphisms in the indexed category (Lemma 10, page 35 of [6]) yielding (over $\overline{[\Gamma]}$):

$$\begin{aligned} \psi_{I(a, a')} : \overline{[I(a, a')]} &\cong I_C(FA[\varphi_\Gamma], \{\psi_A\}(\overline{[a]}), \{\psi_A\}(\overline{[a']})) \\ &= I_C(FA[\varphi_\Gamma], F(a)[\varphi_\Gamma], F(a')[\varphi_\Gamma]). \end{aligned}$$

Σ -types. Assume we have $\Gamma \vdash A$ and $\Gamma.A \vdash B$. By induction we have the isomorphisms $\psi_A : \overline{[A]} \cong_\Gamma FA[\varphi_\Gamma]$ and $\psi_B : \overline{[B]} \cong_{\Gamma.A} FB[\varphi_{\Gamma.B}]$. We let

$$\begin{aligned} \psi_{\Sigma(A, B)} &= \Gamma.\Sigma(A, B) \xrightarrow{\Sigma(\psi_A, \psi_B)} \Gamma.\Sigma(FA[\varphi_\Gamma], FB[\rho^{-1} \circ \varphi_\Gamma^{-1+}]) \\ &\xrightarrow{\mathbf{T}(\varphi_\Gamma)(s_{A, B}^{-1})} \Gamma.F(\Sigma(A, B))[\varphi_\Gamma] \end{aligned}$$

$\psi_{\Sigma(A, B)}$ can be related to $\varphi_{\Gamma.A.B}$:

$$\begin{aligned} \psi_{\Sigma(A, B)} &= \mathbf{T}(\varphi_\Gamma)(s_{A, B}^{-1}) \circ \Sigma(\psi_A, \psi_B) \\ &= \varphi_\Gamma^{-1+} \circ s_{A, B}^{-1} \circ \varphi_\Gamma^+ \circ \chi^{-1} \circ \psi_A^+ \circ \psi_B \circ \chi_{A, B} \\ &= \varphi_\Gamma^{-1+} \circ s_{A, B}^{-1} \circ \chi_{A, B} \circ \varphi_\Gamma^{++} \circ \psi_A^+ \circ \psi_B \circ \chi_{A, B} \\ &= \varphi_\Gamma^{-1+} \circ \rho \circ F(\chi_{A, B}) \circ \rho^{-1} \circ \rho^{-1+} \circ \varphi_\Gamma^{++} \circ \psi_A^+ \circ \psi_B \circ \chi_{A, B} \\ &= \varphi_\Gamma^{-1+} \circ \rho \circ F(\chi_{A, B}^{-1}) \circ \rho^{-1} \circ \varphi_{\Gamma.A}^+ \circ \psi_B \circ \chi_{A, B} \\ &= \varphi_\Gamma^{-1+} \circ \rho \circ F(\chi_{A, B}^{-1}) \circ \varphi_{\Gamma.A.B} \circ \chi_{A, B} \end{aligned}$$

From that calculation, we deduce $\varphi_{\Gamma.\Sigma(A, B)} = F(\chi_{A, B}^{-1}) \circ \varphi_{\Gamma.A.B} \circ \chi_{A, B}$.

Π -types. Define $\psi_{\Pi(A, B)}$ as follows

$$\begin{aligned} \overline{[\Gamma.\Pi(A, B)]} &\xrightarrow{\Pi(\psi_A, \psi_B)} \overline{[\Gamma]}. \Pi(FA[\varphi_\Gamma], FB[\rho \circ \varphi_{\Gamma \uparrow FA}]) \\ &= \overline{[\Gamma]}. \Pi(FA, FB[\rho])[\varphi_\Gamma] \\ &\xrightarrow{T(\varphi_\Gamma)(\xi_{A, B}^{-1})} \overline{[\Gamma]}. F(\Pi(A, B))[\varphi_\Gamma] \end{aligned}$$

4.4.2 Uniqueness

We resume the uniqueness proof from 3.4.

The unit type. It follows from the preservation of democracy of F .

Identity types. We need to show $\psi'_{\mathbb{I}(a,a')} = \psi_{\mathbb{I}(a,a')} : \Gamma.\mathbb{I}(a,a') \rightarrow \Gamma.F(\mathbb{I}(a,a'))[\varphi_\Gamma]$. Let $A = \text{type-of}(a)$. By post-composing with the coherence isomorphism $F(\mathbb{I}(a,a')) \cong_{F\Gamma} \mathbb{I}(FA, Fa, Fa')$, we get a morphism between identity types. In an extensional type theory, identity types are either empty or singletons, thus there is at most one morphism between two identity types (which is an isomorphism). This implies that $\psi_{\mathbb{I}(a,a')} = \psi'_{\mathbb{I}(a,a')}$.

Σ -types. By induction, we assume that $\varphi_{\Gamma.A.B} = \varphi'_{\Gamma.A.B}$. By naturality of φ' , we have $\varphi'_{\Sigma(A,B)} = F(\chi_{A,B}^{-1}) \circ \varphi'_{\Gamma.A.B} \circ \chi_{A,B} = \varphi_{\Gamma.\Sigma(A,B)}$. Hence $\psi_{\Sigma(A,B)} = \psi'_{\Sigma(A,B)}$.

Π -types. As in the previous section, by induction we assume $\varphi_{\Gamma.A.B} = \varphi'_{\Gamma.A.B}$. Let ev be the obvious map $\Gamma.A.\Pi(A,B)[\mathbb{p}] \rightarrow \Gamma.A.B$. Proposition 11 of [6] entails:

► **Lemma 19.** *Assume $\Gamma.A \vdash B$. The only automorphism ω of $\Pi(A,B)$ (in \mathbf{TT}) such that $T\mathbb{p}(\omega) : \Gamma.A.\Pi(A,B)[\mathbb{p}] \cong \Gamma.A.\Pi(A,B)[\mathbb{p}]$ satisfies $\text{ev} \circ T\mathbb{p}(\omega) = \text{ev}$ is the identity.*

It remains to show that $\psi_{\Pi(A,B)}^{-1} \circ \psi'_{\Pi(A,B)}$ satisfies the condition. But we have:

$$\begin{aligned} & F(\text{ev}) \circ \rho^{-1} \circ \theta_{\Pi(A,B),\mathbb{p}} \circ \varphi_{\Gamma.A} \circ T\mathbb{p}(\psi'_{\Pi(A,B)}) \\ &= F(\text{ev}) \circ \rho^{-1} \circ \varphi_{\Gamma.A} \circ T(\varphi_{\Gamma.A})(\theta) \circ T\mathbb{p}(\psi'_{\Pi(A,B)}) \\ &= F(\text{ev}) \circ \rho^{-1} \circ \varphi_{\Gamma.A} \circ \psi_{\Pi(A,B)[\mathbb{p}]} \\ &= F(\text{ev}) \circ \varphi_{\Gamma.A.\Pi(A,B)[\mathbb{p}]} \\ &= \varphi'_{\Gamma.A.B} \circ \text{ev} \end{aligned}$$

By only using naturality conditions on φ' and ψ' . Write $\tau : \Gamma.A.F(\Pi(A,B))[\varphi_\Gamma][\mathbb{p}] \rightarrow F(\Gamma.A.B)$ for the map $F(\text{ev}) \circ \rho^{-1} \circ \theta_{\Pi(A,B),\mathbb{p}} \circ \varphi_{\Gamma.A}$. Since φ and ψ are natural, we can do the same reasoning, and have $\tau \circ T\mathbb{p}(\psi_{\Pi(A,B)}^{-1}) = \varphi_{\Gamma.A.B} \circ \text{ev}$. Thus, we get:

$$\varphi_{\Gamma.A.B}^{-1} \circ \tau = \text{ev} \circ T\mathbb{p}(\psi_{\Pi(A,B)}^{-1})$$

Using our induction hypothesis on B ($\varphi_{\Gamma.A.B} = \varphi'_{\Gamma.A.B}$) we have

$$\text{ev} \circ T\mathbb{p}(\psi_{\Pi(A,B)}^{-1}) \circ T\mathbb{p}(\psi'_{\Pi(A,B)}) = \varphi_{\Gamma.A.B}^{-1} \circ \tau \circ T\mathbb{p}(\psi'_{\Pi(A,B)}) = \text{ev}$$

as desired. Hence $\psi_{\Pi(A,B)} = \psi'_{\Pi(A,B)}$.

4.5 The free lccc

Let \mathbf{LCC} be the 2-category of lcccs. The biequivalence of [6] yields pseudofunctors:

$$U : \mathbf{Cwf}_d^{\Sigma,\Pi,I} \rightarrow \mathbf{LCC} \quad H : \mathbf{LCC} \rightarrow \mathbf{Cwf}_d^{\Sigma,\Pi,I}$$

such that $UH = \mathbb{I}$ and $HU \cong \mathbb{I}$. In particular there are adjunctions $H \dashv U$ and $U \dashv H$.

► **Theorem 20.** *$UT^{\mathbb{I},N_1,\Sigma,\Pi}$ is the bifree lccc on one object, that is, it is bi-initial in \mathbf{LCC}^o .*

Proof. Let \mathbb{C} be an lccc with a chosen object $o_{\mathbb{C}} \in \mathbb{C}$. By democracy $o_{\mathbb{C}}$ can be viewed as a type over the empty context in the cwf $H\mathbb{C}$, thus $(H\mathbb{C}, o)$ is in $\mathbf{Cwf}_d^{\Sigma,\Pi,I,o}$. Thus we have a pseudo cwf functor $\llbracket \cdot \rrbracket : \mathcal{T}^{\Sigma,\Pi,N_1,I} \rightarrow H\mathbb{C}$ satisfying $\llbracket o \rrbracket \cong o_{\mathbb{C}}$. Hence $F : UT^{\Sigma,\Pi,N_1,I} \rightarrow \mathbb{C}$ in \mathbf{LCC} because of the adjunction.

Assume we have another $G : UT^{\Sigma,\Pi,N_1,I} \rightarrow \mathbb{C}$. Because of the adjunction we get $G^* : \mathcal{T}^{\Sigma,\Pi,N_1,I} \rightarrow H\mathbb{C}$. Thus by bifreeness of $\mathcal{T}^{\Sigma,\Pi,N_1,I}$ we have $\varphi : \llbracket \cdot \rrbracket \cong G^*$, thus $F \cong G$. Moreover, any other morphism $F \rightarrow G$ yields a morphism $\llbracket \cdot \rrbracket \rightarrow G^*$ and is equal to φ . ◀

References

- 1 Henk P. Barendregt. Lambda calculi with types. In Samson Abramsky, Dov Gabbay, and Tom Maibaum, editors, *Handbook of Logic in Computer Science*, volume 2, pages 118–310. Oxford University Press, 1992.
- 2 Jean Benabou. Fibered categories and the foundations of naive category theory. *J. Symb. Log.*, 50(1):10–37, 1985.
- 3 John Cartmell. Generalized algebraic theories and contextual categories. *Annals of Pure and Applied Logic*, 32:209–243, 1986.
- 4 Simon Castellan, Pierre Clairambault, and Peter Dybjer. Undecidability of equality in the free locally cartesian closed category. Available at <http://arxiv.org/abs/1504.03995>, 2015.
- 5 Pierre Clairambault and Peter Dybjer. The biequivalence of locally cartesian closed categories and martin-löf type theories. In *Typed Lambda Calculi and Applications – 10th International Conference, TLCA 2011, Novi Sad, Serbia, June 1-3, 2011. Proceedings*, pages 91–106, 2011.
- 6 Pierre Clairambault and Peter Dybjer. The biequivalence of locally cartesian closed categories and Martin-Löf type theories. *Mathematical Structures in Computer Science*, 24(6), 2014.
- 7 Pierre-Louis Curien. Substitution up to isomorphism. *Fundamenta Informaticae*, 19(1,2):51–86, 1993.
- 8 Peter Dybjer. Internal type theory. In *TYPES’95, Types for Proofs and Programs*, number 1158 in Lecture Notes in Computer Science, pages 120–134. Springer, 1996.
- 9 Martin Hofmann. Interpretation of type theory in locally cartesian closed categories. In *Proceedings of CSL*. Springer LNCS, 1994.
- 10 Per Martin-Löf. Constructive mathematics and computer programming. In *Logic, Methodology and Philosophy of Science, VI, 1979*, pages 153–175. North-Holland, 1982.
- 11 Per Martin-Löf. Substitution calculus. Notes from a lecture given in Göteborg, November 1992.
- 12 Erik Palmgren and Steve J. Vickers. Partial horn logic and cartesian categories. *Annals of Pure and Applied Logic*, 145(3):314 – 353, 2007.
- 13 Robert Seely. Locally cartesian closed categories and type theory. *Math. Proc. Cambridge Philos. Soc.*, 95(1):33–48, 1984.
- 14 Thomas Streicher. *Semantics of Type Theory*. Number 12 in Progress in Theoretical Computer Science. Basel: Birkhauser Verlag, 1991.
- 15 Alvaro Tasistro. Formulation of Martin-Löf’s theory of types with explicit substitutions. Technical report, Department of Computer Sciences, Chalmers University of Technology and University of Göteborg, 1993. Licentiate Thesis.

The Inconsistency of a Brouwerian Continuity Principle with the Curry–Howard Interpretation

Martín Hötzel Escardó and Chuangjie Xu

School of Computer Science, University of Birmingham, UK

Abstract

If all functions $(\mathbb{N} \rightarrow \mathbb{N}) \rightarrow \mathbb{N}$ are continuous then $0 = 1$. We establish this in intensional (and hence also in extensional) intuitionistic dependent-type theories, with existence in the formulation of continuity expressed as a Σ type via the Curry–Howard interpretation. But with an intuitionistic notion of anonymous existence, defined as the propositional truncation of Σ , it is consistent that all such functions are continuous. A model is Johnstone’s topological topos. On the other hand, any of these two intuitionistic conceptions of existence give the same, consistent, notion of uniform continuity for functions $(\mathbb{N} \rightarrow 2) \rightarrow \mathbb{N}$, again valid in the topological topos. It is open whether the consistency of (uniform) continuity extends to homotopy type theory. The theorems of type theory informally proved here are also formally proved in Agda, but the development presented here is self-contained and doesn’t show Agda code.

1998 ACM Subject Classification F.4.1 Mathematical Logic

Keywords and phrases Dependent type, intensional Martin-Löf type theory, Curry–Howard interpretation, constructive mathematics, Brouwerian continuity axioms, anonymous existence, propositional truncation, function extensionality, homotopy type theory, topos theory

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.153

1 Introduction

We show that a continuity principle that holds in Brouwerian intuitionistic mathematics becomes false when we move to its Curry–Howard interpretation. We formulate and prove this in an intensional version of intuitionistic type theory (Section 2). Another Brouwerian (uniform) continuity principle, however, is logically equivalent to its Curry–Howard interpretation (Section 4).

In order to be able to formulate and prove this logical equivalence, we need a type theory in which both a formula and its Curry–Howard interpretation can be expressed (Section 3). For example, toposes admit both \forall, \exists (via the subobject classifier) and Π, Σ (via their local cartesian closedness) and hence qualify. We adopt the HoTT-book [17] approach of working with propositional truncation $\| - \|$ to express $\exists(x : X).A(x)$ as the propositional truncation of $\Sigma(x : X).A(x)$. This is related to NuPrI’s squash types [14], Maietti’s mono-types [13], and Awodey–Bauer bracket types in extensional type theory [1]. Here by a proposition we mean a type whose elements are all equal in the sense of the identity type, as in the HoTT book. In a topos with identity types understood as equalizers, the propositions are the truth values (subterminal objects), and the propositional truncation of an object X is its support, namely the image of the unique map $X \rightarrow 1$ to the terminal object. This gives the truth value of the inhabitedness of X , without necessarily revealing an inhabitant of X , and we have that

$$(\exists(x : X).A(x)) = \|\Sigma(x : X).A(x)\|.$$



© Martín Hötzel Escardó and Chuangjie Xu;
licensed under Creative Commons License CC-BY

13th International Conference on Typed Lambda Calculi and Applications (TLCA’15).

Editor: Thorsten Altenkirch; pp. 153–164



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In HoTT, this is taken as the definition of \exists , with truncation taken as a primitive notion. But we don’t (need to) work with the homotopical understanding of type theory or the univalence axiom here.

1.1 The continuity of all functions $\mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}$

In Brouwerian intuitionistic mathematics, all functions $f : \mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}$ on the Baire space $\mathbb{N}^{\mathbb{N}} = (\mathbb{N} \rightarrow \mathbb{N})$ are continuous [2, 19]. This means that, for any sequence $\alpha : \mathbb{N}^{\mathbb{N}}$ of natural numbers, the value $f\alpha$ of the function depends only on a finite prefix of the argument $\alpha : \mathbb{N}^{\mathbb{N}}$. If we write $\alpha =_n \beta$ to mean that the sequences α and β agree at their first n positions, a precise formulation of this continuity principle is

$$\forall(f : \mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}). \forall(\alpha : \mathbb{N}^{\mathbb{N}}). \exists(n : \mathbb{N}). \forall(\beta : \mathbb{N}^{\mathbb{N}}). \alpha =_n \beta \rightarrow f\alpha = f\beta.$$

It is well known that this statement cannot be proved in higher-type Heyting arithmetic (HA^{ω}), but that is consistent and validated by the model of Kleene–Kreisel continuous functionals, and also realizable with Kleene’s second combinatory algebra K_2 [2].

We show that, in intensional Martin–Löf type theory, the Curry–Howard interpretation of the above continuity principle is false: It is a theorem of intensional MLTT, even without universes, that

$$(\Pi(f : \mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}). \Pi(\alpha : \mathbb{N}^{\mathbb{N}}). \Sigma(n : \mathbb{N}). \Pi(\beta : \mathbb{N}^{\mathbb{N}}). \alpha =_n \beta \rightarrow f\alpha = f\beta) \rightarrow 0 = 1.$$

We prove this by adapting Kreisel’s well-known argument that, e.g. in HA^{ω} , extensionality, choice and continuity are together impossible [12][18][2, page 267]. The difference here is that

1. We work in *intensional* type theory.
2. Choice for the Σ interpretation of existence is a theorem of type theory.

So what is left to understand is that extensionality is not needed in Kreisel’s argument when it is rendered in type theory (Section 2).

The above two versions of the notion of continuity can be usefully compared by considering the interpretations of HA^{ω} and MLTT in Johnstone’s *topological topos* [9]. The point of this topos is that it fully embeds a large cartesian closed category of continuous maps of topological spaces, the sequential topological spaces, and the larger locally cartesian closed category of Kuratowski limit spaces [15]. As discussed above, any topos has $\exists, \forall, \Sigma, \Pi$ and therefore models both intuitionistic predicate logic and dependent type theory. We have that

1. The formula

$$\forall(f : \mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}). \forall(\alpha : \mathbb{N}^{\mathbb{N}}). \exists(n : \mathbb{N}). \forall(\beta : \mathbb{N}^{\mathbb{N}}). \alpha =_n \beta \rightarrow f\alpha = f\beta$$

is true in the topological topos.

The informal reading of this is “all functions $\mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}$ are continuous”.

2. There is a function

$$(\Pi(f : \mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}). \Pi(\alpha : \mathbb{N}^{\mathbb{N}}). \Sigma(n : \mathbb{N}). \Pi(\beta : \mathbb{N}^{\mathbb{N}}). \alpha =_n \beta \rightarrow f\alpha = f\beta) \rightarrow 0 = 1$$

in the topological topos, or indeed in any topos whatsoever, by our version of Kreisel’s argument.

The informal reading of this is “not all functions $\mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}$ are continuous”.

But there is no contradiction in the formal versions of the above statements: they simultaneously hold in the same world, the topological topos. From a hypothetical inhabitant of

$$\Pi(f : \mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}). \Pi(\alpha : \mathbb{N}^{\mathbb{N}}). \Sigma(n : \mathbb{N}). \Pi(\beta : \mathbb{N}^{\mathbb{N}}). \alpha =_n \beta \rightarrow f\alpha = f\beta$$

we get a modulus-of-continuity functional

$$M : (\mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}) \times \mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N},$$

by projection (rather than by choice in the topos-logic sense), which gives a modulus of continuity $n = M(f, \alpha)$ of the function $f : \mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}$ at the point $\alpha : \mathbb{N}^{\mathbb{N}}$. Kreisel's argument derives a contradiction from the existence of M . What this shows, then, is that although every function *is* continuous, there is no continuous way of finding a modulus of continuity of a given function f at a given point α . There is no continuous M . Perhaps the difference between the seemingly contradictory statements becomes clearer if we formulate them type theoretically with and without propositional truncation. In the topological topos, the object

$$\Pi(f : \mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}). \Pi(\alpha : \mathbb{N}^{\mathbb{N}}). \|\Sigma(n : \mathbb{N}). \Pi(\beta : \mathbb{N}^{\mathbb{N}}). \alpha =_n \beta \rightarrow f\alpha = f\beta\|$$

is inhabited, but

$$\Pi(f : \mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}). \Pi(\alpha : \mathbb{N}^{\mathbb{N}}). \Sigma(n : \mathbb{N}). \Pi(\beta : \mathbb{N}^{\mathbb{N}}). \alpha =_n \beta \rightarrow f\alpha = f\beta$$

is not.

1.2 The uniform continuity of all functions $\mathfrak{2}^{\mathbb{N}} \rightarrow \mathbb{N}$

The above situation changes radically when we move from the Baire space to the Cantor space, and from continuous functions to uniformly continuous functions. Another Brouwerian continuity principle is that all functions from the Cantor space $\mathfrak{2}^{\mathbb{N}} = (\mathbb{N} \rightarrow \mathfrak{2})$ to the natural numbers are uniformly continuous:

$$\forall(f : \mathfrak{2}^{\mathbb{N}} \rightarrow \mathbb{N}). \exists(n : \mathbb{N}). \forall(\alpha, \beta : \mathfrak{2}^{\mathbb{N}}). \alpha =_n \beta \rightarrow f\alpha = f\beta.$$

Again this is not provable in HA^{ω} but consistent and validated by the model of continuous functionals, by realizability over K_2 , and by the topological topos. We have also constructively developed a model analogous to the topological topos in [20].

By the above discussion, the above principle is equivalent to

$$\Pi(f : \mathfrak{2}^{\mathbb{N}} \rightarrow \mathbb{N}). \|\Sigma(n : \mathbb{N}). \Pi(\alpha, \beta : \mathfrak{2}^{\mathbb{N}}). \alpha =_n \beta \rightarrow f\alpha = f\beta\|.$$

We show that this, in turn, is logically equivalent to its untruncated version

$$\Pi(f : \mathfrak{2}^{\mathbb{N}} \rightarrow \mathbb{N}). \Sigma(n : \mathbb{N}). \Pi(\alpha, \beta : \mathfrak{2}^{\mathbb{N}}). \alpha =_n \beta \rightarrow f\alpha = f\beta.$$

In particular, it follows that this object is inhabited (by a global point) in the topological topos. Each inhabitant gives, by projection, a “fan functional” $(\mathfrak{2}^{\mathbb{N}} \rightarrow \mathbb{N}) \rightarrow \mathbb{N}$ that continuously assigns a modulus of uniform continuity to its argument. There is a canonical one, which assigns the least modulus of uniform continuity.

In order to establish the above logical equivalence, we prove the following general principle for “exiting truncations”: If A is a family of types indexed by natural numbers such that

1. $A(n)$ is a proposition for every $n : \mathbb{N}$, and
2. $A(n)$ implies that $A(m)$ is decidable for every $m < n$,

then

$$\|\Sigma(n : \mathbb{N}). A(n)\| \rightarrow \Sigma(n : \mathbb{N}). A(n).$$

From anonymous existence one gets explicit existence in this case.

1.3 A question regarding Church’s Thesis

Troelstra [18] also shows that extensionality, choice and Church’s Thesis (CT) are together impossible, and Beeson [2, page 268] adapts this argument to conclude that extensional Martin-Löf type theory refutes CT, with existence expressed by Σ . But CT with existence expressed as the truncation of Σ is consistent with MLTT, and validated by Hyland’s *effective topos* [8]. What seems to be open is whether CT formulated with Σ is already refuted by *intensional* MLTT (including the ξ -rule). This question has been popularized by Maria Emilia Maietti.

2 Continuity of functions $\mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}$

We reason informally, but rigorously, in type theory, where, as above, we use the equality sign to denote identity types, unless otherwise indicated. A formal proof, written in Agda [3, 4, 16], is available at [6], but the development here is self-contained and doesn’t show Agda code.

The following says that the Curry–Howard interpretation of “all functions $\mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}$ are continuous” is false.

► **Theorem 1.** *If*

$$\Pi(f: \mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}).\Pi(\alpha: \mathbb{N}^{\mathbb{N}}).\Sigma(n: \mathbb{N}).\Pi(\beta: \mathbb{N}^{\mathbb{N}}). \alpha =_n \beta \rightarrow f\alpha = f\beta$$

then $0 = 1$.

We take the conclusion to be $0 = 1$ rather than the empty type because we are not assuming a universe for the sake of generality. The argument below gives $0 = 1$, and, as is well known, to get to the empty type from $0 = 1$ a universe is needed.

Proof. Let 0^ω denote the infinite sequence of zeros, that is, $\lambda i.0$, and let $0^n k^\omega$ denote the sequence of n many zeros followed by infinitely many k ’s. Then

$$(0^n k^\omega) =_n 0^\omega \quad \text{and} \quad (0^n k^\omega)(n) = k.$$

Assume $\Pi(f: \mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}).\Pi(\alpha: \mathbb{N}^{\mathbb{N}}).\Sigma(n: \mathbb{N}).\Pi(\beta: \mathbb{N}^{\mathbb{N}}).\alpha =_n \beta \rightarrow f\alpha = f\beta$. By projection, with $\alpha = 0^\omega$, this gives a modulus-of-continuity function

$$M: (\mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}) \rightarrow \mathbb{N}$$

such that

$$\Pi(f: \mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}).\Pi(\beta: \mathbb{N}^{\mathbb{N}}).0^\omega =_{Mf} \beta \rightarrow f(0^\omega) = f\beta. \quad (1)$$

We use M to define a function $f: \mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}$ such that $M(f)$ cannot be a modulus of continuity of f and hence get a contradiction. Let

$$m = M(\lambda\alpha.0),$$

and define $f: \mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}$ by

$$f\beta = M(\lambda\alpha.\beta(\alpha m)).$$

The crucial observation is that, by simply expanding the definitions, we have the judgemental equalities

$$f(0^\omega) = M(\lambda\alpha.0^\omega(\alpha m)) = M(\lambda\alpha.0) = m,$$

because $0^\omega(\alpha m) = 0$. By the defining property (1) of M , and the crucial observation,

$$\Pi(\beta : \mathbb{N}^{\mathbb{N}}).0^\omega =_{Mf} \beta \rightarrow m = f\beta. \quad (2)$$

For any $\beta : \mathbb{N}^{\mathbb{N}}$, by the continuity of $\lambda\alpha.\beta(\alpha m)$, by the definition of f , and by the defining property (1) of M , we have that

$$\Pi(\alpha : \mathbb{N}^{\mathbb{N}}).0^\omega =_{f\beta} \alpha \rightarrow \beta 0 = \beta(\alpha m).$$

If we choose $\beta = 0^{Mf+1}1^\omega$, we have $0^\omega =_{Mf+1} \beta$, and so $0^\omega =_{Mf} \beta$, and hence $f(\beta) = m$ by (2). This gives

$$\Pi(\alpha : \mathbb{N}^{\mathbb{N}}).0^\omega =_m \alpha \rightarrow \beta 0 = \beta(\alpha m).$$

Considering $\alpha = 0^m(Mf + 1)^\omega$, we have $0^\omega =_m \alpha$, and therefore

$$0 = \beta 0 = \beta(\alpha m) = \beta(Mf + 1) = 1. \quad \blacktriangleleft$$

► **Remark** (Thomas Streicher, personal communication). The conversion

$$f(0^\omega) = M(\lambda\alpha.0^\omega(\alpha m)) = M(\lambda\alpha.0) = m$$

in the above proof relies on the ξ -rule (reduction under λ), which is not available in a system based on the combinators S and K rather than the λ -calculus. Usually HA^ω is taken in combinatory form, in which case one needs some form of extensionality to conclude that $f(0^\omega) = m$, and this explains how we avoid the extensionality hypothesis in Kreisel's original argument. But notice that the ξ -rule holds in categorical models.

Therefore the argument of the above proof shows that:

► **Theorem 2.** *In HA^ω , the ξ -rule, the axiom of choice, and the continuity of all functions $\mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}$ are together impossible.*

Another observation, offered independently by Thorsten Altenkirch, Thierry Coquand and Per Martin-Löf (personal communication), is that the continuity of a function $\mathbb{N}^{\mathbb{N}} \rightarrow \mathbb{N}$ implies that it is extensional in the sense that it maps pointwise equal arguments to equal values, and so the continuity axiom has some amount of extensionality built into it.

The above formulation and proof of Theorem 1 assumes natural numbers, identity types, Π and Σ types, and no universes. But it uses only the identity type of the natural numbers. If we assume a universe U , this identity type doesn't need to be assumed, because it can be defined by induction. We first define a U -valued equality relation, where \mathbb{O} is the empty type and $\mathbb{1}$ is the unit type with element \star ,

$$(0 = 0) = \mathbb{1}, \quad (m + 1 = 0) = (0 = n + 1) = \mathbb{O}, \quad (m + 1 = n + 1) = (m = n).$$

Then we define $\text{refl} : \Pi(n : \mathbb{N}).n = n$ by induction as

$$\text{refl}(0) = \star, \quad \text{refl}(n + 1) = \text{refl}(n),$$

and $J : \Pi(A : \Pi(m, n).m = n \rightarrow U).(\Pi n.A n n (\text{refl}(n))) \rightarrow \Pi m, n, p.A m n p$ by

$$\begin{aligned} J \ A \ r \ 0 \quad 0 \quad \star &= r\ 0, \\ J \ A \ r \ (m + 1) \ 0 \quad p &= \mathbb{O}\text{-rec}(A(m + 1)\ 0)\ p, \\ J \ A \ r \ 0 \quad (n + 1) \ p &= \mathbb{O}\text{-rec}(A\ 0\ (n + 1))\ p, \\ J \ A \ r \ (m + 1) \ (n + 1) \ p &= J(\lambda mn.A(m + 1)(n + 1))(\lambda n.r(n + 1))\ m\ n\ p. \end{aligned}$$

where $\mathbb{O}\text{-rec} : \Pi(X : U).\mathbb{O} \rightarrow X$ is the recursion combinator of the empty type. The usual computation rule, or judgemental equality, for J when it is given as primitive doesn't hold here, but the above J is enough to define transport (substitution) and hence symmetry, transitivity and application (congruence), which are enough to carry out the above proof formally (and we have checked this in Agda [6]). Hence the theorem and its proof can be expressed in a type theory without a primitive equality type. All is needed to formulate and prove Theorem 1 is a type theory with $\mathbb{O}, \mathbb{1}, \mathbb{N}, \Pi, \Sigma, U$.

3 Propositional truncation and existential quantification

We recall the notion of propositional truncation from the HoTT book and use it to define the quantifiers \exists, \forall , in a slightly different way from that in the HoTT book, so that they satisfy the Lawvere's adjointness conditions that correspond to their intuitionistic introduction and elimination rules. Another difference is that, instead of adding propositional truncations for all types to our type theory, we define what a propositional truncation for a given type is. For some types, their propositional truncation already exist, including the types needed in our discussion of uniform continuity in Section 4.

3.1 Propositional truncation

We adopt the terminology of the HoTT book, which clashes with the terminology of the Curry–Howard interpretation of (syntactical) propositions as types. For us, a *proposition* is a subsingleton, or a type whose elements are all equal, in the sense of the identity type, here written “=” again as in the HoTT book:

$$\text{isProp } X = \Pi(x, y : X).x = y.$$

Perhaps a better terminology, compatible with that of topos theory, would be *truth value*, in order to avoid the clash. But we will stick to the terminology *proposition*, and occasionally use *truth value* synonymously, for emphasis.

A propositional truncation of a type X , if it exists, is a proposition $\|X\|$ together with a map $|-| : X \rightarrow \|X\|$ such that for any proposition P and $f : X \rightarrow P$ we can find $\bar{f} : \|X\| \rightarrow P$. Because P is a proposition, this map \bar{f} is automatically unique up to pointwise equality, and we have $\bar{f}|x| = f(x)$, and hence a propositional truncation is a reflection in the categorical sense, giving a universal map of X into a proposition. This can also be understood as a recursion principle, or elimination rule,

$$\text{isProp } P \rightarrow (X \rightarrow P) \rightarrow \|X\| \rightarrow P,$$

for any types P and X . The induction principle, in this case, can be derived from the recursion principle, but in practice it is seldom needed.

In HoTT, propositional truncations for all types are given as higher-inductive types, with the judgemental equality $\bar{f}|x| = f(x)$. From the existence of the truncation of the two-point type $\mathbb{2}$ with this judgemental equality, one can prove function extensionality (any two pointwise equal functions are equal) [11]. The assumption that $\|X\| \rightarrow X$ for every type X gives a constructive taboo (and also contradicts univalence) [10].

However, for some types X , not only can a propositional truncation $\|X\|$ be constructed in MLTT, but also there is a map $\|X\| \rightarrow X$:

1. If $P = \mathbb{O}$ or $P = \mathbb{1}$, or more generally if P is any proposition, we can take $\|P\| = P$, of course. In particular, if $X \rightarrow \mathbb{O}$, we can take $\|X\| = X$; but also we can take $\|X\| = \mathbb{O}$, even though we can't say $X = \mathbb{O}$ without univalence.

2. If we have an inhabitant of X then we can take $\|X\| = 1$. The map $\|X\| \rightarrow X$ simply picks the given inhabitant.
3. More generally, if X is logically equivalent to a proposition P , then we can take $\|X\| = P$, and we make profitable use of this simple fact.
4. If X is any type and $g : X \rightarrow X$ is a constant map in the sense that any two of its values are equal, we can take $\|X\|$ to be the type $\Sigma(x : X).g(x) = x$ of fixed points of g , together with the function $X \rightarrow \|X\|$ that maps x to $(g(x), p)$, where p is an inhabitant of the type $g(g(x)) = g(x)$ coming from the constancy witness [10]. In this case the first projection gives a map $\|X\| \rightarrow X$. Given a map $f : X \rightarrow P$ into a proposition, we let $\bar{f} : \|X\| \rightarrow P$ be the first projection followed by $f : X \rightarrow P$ (and we don't use the fact that P is a proposition).
5. For any $f : \mathbb{N} \rightarrow \mathbb{N}$, the type $\Sigma(n : \mathbb{N}).f(n) = 0$, which may well be empty, has a constant endomap that sends (n, p) to (n', p') , where we take the least $n' \leq n$ with $p' : f(n') = 0$, using the decidability of equality of \mathbb{N} and bounded search. Hence not only $\|\Sigma(n : \mathbb{N}).f(n) = 0\|$ exists, but also $\|\Sigma(n : \mathbb{N}).f(n) = 0\| \rightarrow \Sigma(n : \mathbb{N}).f(n) = 0$.

3.2 Quantification

For a universe U , let Prop be the type of propositions in U :

$$\text{Prop} = \Sigma(X : U).\text{isProp } X.$$

If we assume that all types in U come with designated propositional truncations, then we have a reflection

$$r : U \rightarrow \text{Prop}$$

that sends $X : U$ to the pair $(\|X\|, p)$ with $p : \text{isProp } \|X\|$ coming from the assumption. In the other direction, we have an embedding

$$s : \text{Prop} \rightarrow U,$$

given by the projection. For $X : U$ we have

$$s(r(X)) = \|X\|.$$

(We also have that s is a section of r if propositional univalence holds.) For a fixed type $X : U$, the type constructors Σ and Π can be regarded as having type

$$\Sigma, \Pi : (X \rightarrow U) \rightarrow U.$$

We define

$$\exists, \forall : (X \rightarrow \text{Prop}) \rightarrow \text{Prop},$$

by, for any $A : X \rightarrow \text{Prop}$,

$$\exists(A) = r(\Sigma(s \circ A)), \quad \forall(A) = r(\Pi(s \circ A)),$$

which we also write more verbosely as

$$\begin{aligned} (\exists(x : X).A(x)) &= r(\Sigma(x : X).s(A(x))), \\ (\forall(x : X).A(x)) &= r(\Pi(x : X).s(A(x))). \end{aligned}$$

This is essentially the same as the definition in the HoTT book, except that we give different types to \exists, \forall . With the type given in the book, \forall gets confused with Π , because, with function extensionality, a product of propositions is a proposition, and so there is no need to distinguish \forall from Π in the book.

The point of choosing the above types is that now it is easy to justify that these quantifiers do satisfy the intuitionistic rules for quantification. It is enough to show that they satisfy Lawvere’s adjointness conditions. For $P, Q : \text{Prop}$, define

$$(P \leq Q) = (s(P) \rightarrow s(Q)).$$

This is a pre-order (and a partial order if propositional univalence holds). Now endow the function type $(X \rightarrow \text{Prop})$ with the pointwise order (using Π to define it). Then the quantifiers $\exists, \forall : (X \rightarrow \text{Prop}) \rightarrow \text{Prop}$ are the left and right adjoints to the exponential transpose $\text{Prop} \rightarrow (X \rightarrow \text{Prop})$ of the projection

$$\text{Prop} \times X \rightarrow \text{Prop},$$

using the universal property of truncation. The exponential transpose maps P to $\lambda x.P$. Hence the adjointness condition for the existential quantifier amounts to

$$\exists(A) \leq P \iff A \leq \lambda x.P.$$

Expanding the definitions, this amounts to

$$\begin{aligned} \|\Sigma(x : X).s(A(x))\| \rightarrow s(P) &\iff \Pi(x : X).s(A(x)) \rightarrow s(P), \\ &\iff (\Sigma(x : X).s(A(x))) \rightarrow s(P). \end{aligned}$$

So we need to check that

$$\|\Sigma(x : X).s(A(x))\| \rightarrow s(P) \iff (\Sigma(x : X).s(A(x))) \rightarrow s(P)$$

holds, but this is the case by the defining property of propositional truncation. For the sake of completeness, we also check the adjointness condition

$$P \leq \forall(A) \iff \lambda x.P \leq A.$$

For this we need function extensionality (which follows from the assumption of truncations supporting the judgemental equality discussed above). By definition, this amounts to

$$s(P) \rightarrow \|\Pi(x : X).s(A(x))\| \iff \Pi(x : X).s(P) \rightarrow s(A(x)).$$

But

$$\Pi(x : X).s(P) \rightarrow s(A(x)) \iff s(P) \rightarrow \Pi(x : X).s(A(x)),$$

and so the above is equivalent to

$$s(P) \rightarrow \|\Pi(x : X).s(A(x))\| \iff s(P) \rightarrow \Pi(x : X).s(A(x)).$$

But this again holds by the defining property of truncation, because, by function extensionality, a product of propositions is a proposition, and each $s(A(x))$ is a proposition. This explains why \forall is identified with Π in the HoTT book.

Having established that the quantifiers $\exists, \forall : (X \rightarrow \text{Prop}) \rightarrow \text{Prop}$ defined from Σ and Π with truncation (via the reflection $r : U \rightarrow \text{Prop}$) do satisfy the adjointness conditions corresponding to the introduction and elimination rules of intuitionistic logic, in practice we prefer to use the notation of the HoTT book, with $\exists(x : X).A(x)$ defined as $\|\Sigma(x : X).A(x)\|$ for propositionally-valued $A : X \rightarrow U$, or even avoid \exists altogether, and just use truncation explicitly, as in the next section.

4 Uniform continuity of functions $\mathcal{2}^{\mathbb{N}} \rightarrow \mathbb{N}$

We now compare the untruncated formulation of the uniform continuity principle

$$\Pi(f: \mathcal{2}^{\mathbb{N}} \rightarrow \mathbb{N}). \Sigma(n: \mathbb{N}). \Pi(\alpha, \beta: \mathcal{2}^{\mathbb{N}}). \alpha =_n \beta \rightarrow f\alpha = f\beta$$

with its truncated version

$$\Pi(f: \mathcal{2}^{\mathbb{N}} \rightarrow \mathbb{N}). \|\Sigma(n: \mathbb{N}). \Pi(\alpha, \beta: \mathcal{2}^{\mathbb{N}}). \alpha =_n \beta \rightarrow f\alpha = f\beta\|.$$

A formal counter-part in Agda of this section is available at [6].

We work in a type theory with $\mathbb{O}, \mathbb{1}, \mathbb{2}, \mathbb{N}, \Pi, \Sigma, \text{Id}$. This time, identity types for types other than \mathbb{N} are needed, but universes are not. But we need more:

1. In principle, we would have to assume the presence of truncations, for example as defined in the HoTT book and explained in the previous section, from which function extensionality follows [11].
2. However, it turns out that function extensionality alone suffices, because it implies the existence of the propositional truncation mentioned above, and hence we can omit propositional truncations from our type theory. (But it doesn't seem to be possible to remove the assumption of function extensionality in the theorem proved here.)

Hence we don't assume propositional truncations in our type theory.

► **Theorem 3.** *Assuming function extensionality, for every $f: \mathcal{2}^{\mathbb{N}} \rightarrow \mathbb{N}$ the type*

$$\Sigma(n: \mathbb{N}). \Pi(\alpha, \beta: \mathcal{2}^{\mathbb{N}}). \alpha =_n \beta \rightarrow f\alpha = f\beta$$

has a propositional truncation, and the proposition

$$\Pi(f: \mathcal{2}^{\mathbb{N}} \rightarrow \mathbb{N}). \|\Sigma(n: \mathbb{N}). \Pi(\alpha, \beta: \mathcal{2}^{\mathbb{N}}). \alpha =_n \beta \rightarrow f\alpha = f\beta\|$$

is logically equivalent to the type

$$\Pi(f: \mathcal{2}^{\mathbb{N}} \rightarrow \mathbb{N}). \Sigma(n: \mathbb{N}). \Pi(\alpha, \beta: \mathcal{2}^{\mathbb{N}}). \alpha =_n \beta \rightarrow f\alpha = f\beta.$$

► **Lemma 4.** *Function extensionality implies that, for any $f: \mathcal{2}^{\mathbb{N}} \rightarrow \mathbb{N}$, the type family*

$$A(n) = \Pi(\alpha, \beta: \mathcal{2}^{\mathbb{N}}). \alpha =_n \beta \rightarrow f\alpha = f\beta$$

satisfies the following conditions:

1. $A(n)$ is a proposition for every $n: \mathbb{N}$, and
2. $A(n)$ implies that $A(m)$ is decidable for every $m < n$,

Proof. By Hedberg's Theorem [7], equality of natural numbers is a proposition. Hence so is $A(n)$, because, by function extensionality, a product of a family of propositions is a proposition. To conclude that for all n , if $A(n)$ holds then $A(m)$ is decidable for all $m < n$, it is enough to show that for all n , (1) $\neg A(n+1)$ implies $\neg A(n)$, and (2) if $A(n+1)$ holds then $A(n)$ is decidable. (1) This follows from $A(n) \rightarrow A(n+1)$, which says that any number bigger than a modulus of uniform continuity is also a modulus, which is immediate. (2): For every n , the type

$$B(n) = \Pi(s: \mathcal{2}^n). f(s0^\omega) = f(s1^\omega),$$

is decidable, because \mathbb{N} has decidable equality and finite products of decidable types are also decidable. Now let $n: \mathbb{N}$ and assume $A(n+1)$. To show that $A(n)$ is decidable, it is

enough to show that $A(n)$ is logically equivalent to $B(n)$, because then $B(n) \rightarrow A(n)$ and $\neg B(n) \rightarrow \neg A(n)$ and hence we can decide $A(n)$ by reduction to deciding $B(n)$.

The implication $A(n) \rightarrow B(n)$ holds without considering the assumption $A(n+1)$. To see this, assume $A(n)$ and let $s : \mathbb{2}^n$. Taking $\alpha = s0^\omega$ and $\beta = s1^\omega$, we conclude from $A(n)$ that $f(s0^\omega) = f(s1^\omega)$, which is the conclusion of $B(n)$.

Now assume $A(n+1)$ and $B(n)$. To establish $A(n)$, let $\alpha, \beta : \mathbb{2}^{\mathbb{N}}$ with $\alpha =_n \beta$. We need to conclude that $f(\alpha) = f(\beta)$. By the decidability of equality of $\mathbb{2}$, either $\alpha(n) = \beta(n)$ or not. If $\alpha(n) = \beta(n)$, then $\alpha =_{n+1} \beta$, and hence $f(\alpha) = f(\beta)$ by the assumption $A(n+1)$. If $\alpha(n) \neq \beta(n)$, we can assume w.l.o.g. that $\alpha(n) = 0$ and $\beta(n) = 1$. Now take the finite sequence $s = \alpha(0), \alpha(1), \dots, \alpha(n-1) (= \beta(0), \beta(1), \dots, \beta(n-1))$. Then $\alpha =_{n+1} s0^\omega$ and $s1^\omega =_{n+1} \beta$, which together with $A(n+1)$ imply $f(\alpha) = f(s0^\omega)$ and $f(s1^\omega) = f(\beta)$. But $f(s0^\omega) = f(s1^\omega)$ by $B(n)$, and hence $f(\alpha) = f(\beta)$ by transitivity. \blacktriangleleft

► **Lemma 5.** *If a type X is logically equivalent to a proposition Q , then*

1. X has the propositional truncation $\|X\| = Q$, and
2. $\|X\| \rightarrow X$.

Proof. We have $X \rightarrow \|X\|$ because this is the assumption $X \rightarrow Q$. If $X \rightarrow P$ for some proposition P , then also $\|X\| \rightarrow P$, because this means $Q \rightarrow P$, which follows from the assumption $Q \rightarrow X$ and transitivity of implication. This shows that our definition of $\|X\|$ has the required property for truncations. And $\|X\| \rightarrow X$ is the assumption that $Q \rightarrow X$. \blacktriangleleft

► **Lemma 6.** *Function extensionality implies that for any family A of types indexed by natural numbers such that*

1. $A(n)$ is a proposition for every $n : \mathbb{N}$, and
 2. $A(n)$ implies that $A(m)$ is decidable for every $m < n$,
- the type $\Sigma(n : \mathbb{N}).A(n)$ is logically equivalent to the proposition

$$P = \Sigma(k : \mathbb{N}).B(k)$$

where

$$B(k) = A(k) \times \Pi(i : \mathbb{N}).A(i) \rightarrow k \leq i.$$

Proof. By function extensionality, the product of a family of propositions is a proposition, and hence the type $\Pi(n : \mathbb{N}).A(n) \rightarrow k \leq n$ is a proposition, because the type $k \leq n$ is a proposition. Because the product of two propositions is a proposition, the type $B(k)$ is a proposition. But now if $B(k)$ and $B(k')$ then, by construction, $k = k'$. Hence any two inhabitants of P are equal, using the fact that $B(k)$ is a proposition, which means that P is indeed a proposition. By projection, $P \rightarrow \Sigma(n : \mathbb{N}).A(n)$. Conversely, if we have $(n, a) : \Sigma(n : \mathbb{N}).A(n)$, then we can find, by the decidability of $A(m)$ for $m < n$, the minimal k such that there is $b : A(k)$, by search bounded by n , and this gives an element $(k, b, \mu) : P$ where $\mu : \Pi(i : \mathbb{N}).A(i) \rightarrow k \leq i$ is the minimality witness. This shows that $\Sigma(n : \mathbb{N}).A(n) \rightarrow P$ and concludes the proof. \blacktriangleleft

► **Remark.** Function extensionality in the above lemma can be avoided using the fact that the type of fixed points of a constant endomap is a proposition [10], where a map is constant if any two of its values are equal. Given $(n, a) : \Sigma(n : \mathbb{N}).A(n)$, we know that $A(m)$ is decidable for all $m < n$ and thus can find the minimal m such that $A(m)$, by search bounded by n , which gives an endomap of $\Sigma(n : \mathbb{N}).A(n)$. This map is constant, because any two minimal witnesses are equal, and because $A(n)$ is a proposition. Then we instead take P to be the type of fixed points of this constant map.

By Lemmas 4, 5, and 6, for any $f : \mathbb{2}^{\mathbb{N}} \rightarrow \mathbb{N}$, the truncation of the type

$$\text{UC}(f) = \Sigma(n : \mathbb{N}). \Pi(\alpha, \beta : \mathbb{2}^{\mathbb{N}}). \alpha =_n \beta \rightarrow f\alpha = f\beta$$

exists and implies $\text{UC}(f)$, which establishes Theorem 3. \blacktriangleleft

Unfolding the above construction of the truncation, the truncated version of uniform continuity says that there is, using Σ to express existence, a minimal modulus of uniform continuity, making this use of Σ into a proposition, and, by function extensionality, the statement of uniform continuity into a proposition too. Then the theorem says that this proposition is logically equivalent to the existence, using Σ again, of some modulus of uniform continuity. This statement is not a proposition, because any number bigger than a modulus of uniform continuity is itself a modulus of uniform continuity.

The situation here is analogous to that of quasi-inverses and equivalences in the sense of the HoTT book. The type expressing that a function has a quasi-inverse is not a proposition in general, but it is equivalent to the type expressing that the function is an equivalence, which is always a proposition. Hence being an equivalence is the propositional truncation of having a quasi-inverse.

Acknowledgements. The material of this paper was presented at the Institute Henri Poincaré in Paris in June 2014 during the research programme *Semantics of proofs and certified mathematics* organized by Pierre-Louis Curien, Hugo Herbelin, Paul-André Mellies, at the workshop *Semantics of proofs and programs* organized by Thomas Ehrhard and Alex Simpson. We are grateful to them for the invitation to take part of such a wonderful programme and workshop, and to the participants, for useful discussions and input. In particular, we mention Thorsten Altenkirch, Andrej Bauer, Thierry Coquand, Nicolai Kraus, Per Martin-Löf, Paulo Oliva, Bas Spitters and Thomas Streicher.

We also thank the anonymous referees for various suggestions for improvement, including a simplification of the proof of Theorem 1 by removing a case distinction (the updated Agda files [6] contain both our original proof and the simplified proof presented here).

References

- 1 Steve Awodey and Andrej Bauer. Propositions as [types]. *Journal of Logic and Computation*, 14(4):447–471, 2004.
- 2 Michael J. Beeson. *Foundations of Constructive Mathematics*. Springer-Verlag, 1985.
- 3 Ana Bove and Peter Dybjer. Dependent types at work. Lecture Notes for the LerNet Summer School, Piriápolis, Uruguay, 2008.
- 4 Ana Bove, Peter Dybjer, and Ulf Norell. A brief overview of Agda—a functional language with dependent types. In *Theorem proving in higher order logics*, volume 5674 of *Lecture Notes in Computer Science*, pages 73–78. Springer, 2009.
- 5 Martín Hötzel Escardó and Chuangjie Xu. A constructive manifestation of the Kleene–Kreisel continuous functionals. To appear in APAL, 2013.
- 6 Martín Hötzel Escardó and Chuangjie Xu. The inconsistency of a Brouwerian continuity principle with the Curry–Howard interpretation (Agda development). School of Computer Science, University of Birmingham, UK. Available at <http://www.cs.bham.ac.uk/~mhe/continuity-false/>, 2015.
- 7 Michael Hedberg. A coherence theorem for Martin-Löf’s type theory. *J. Functional Programming*, pages 413–436, 1998.
- 8 J Martin E Hyland. The effective topos. *Studies in Logic and the Foundations of Mathematics*, 110:165–216, 1982.

- 9 Peter T. Johnstone. On a topological topos. *Proceedings of the London Mathematical Society*, 38(3):237–271, 1979.
- 10 Nicolai Kraus, Martín Escardó, Thierry Coquand, and Thorsten Altenkirch. Generalizations of hedberg’s theorem. In *Typed Lambda Calculi and Applications*, volume 7941 of *Lecture Notes in Computer Science*, pages 173–188. Springer Berlin Heidelberg, 2013.
- 11 Nicolai Kraus, Martín Escardó, Thierry Coquand, and Thorsten Altenkirch. Notions of anonymous existence in Martin-Löf Type Theory. Submitted for publication, 2014.
- 12 G. Kreisel. On weak completeness of intuitionistic predicate logic. *J. Symbolic Logic*, 27:139–158, 1962.
- 13 Maria Emilia Maietti. The internal type theory of a Heyting pretopos. In Eduardo Giménez and Christine Paulin-Mohring, editors, *Types for Proofs and Programs*, volume 1512 of *Lecture Notes in Computer Science*, pages 216–235. Springer Berlin Heidelberg, 1998.
- 14 N. P. Mendler. Quotient types via coequalizers in Martin-Löf type theory. In G. Huet and G. Plotkin, editors, *Informal Proceedings of the First Workshop on Logical Frameworks*, Antibes, May 1990, pages 349–360, 1990, May 1990.
- 15 M. Menni and A. Simpson. Topological and limit-space subcategories of countably-based equilogical spaces. *Math. Struct. Comput. Sci.*, 12(6):739–770, 2002.
- 16 U. Norell. Dependently typed programming in Agda. In *Proceedings of the 4th international workshop on Types in language design and implementation*, TLDI ’09, pages 1–2, 2009.
- 17 The Univalent Foundations Program. *Homotopy Type Theory: Univalent Foundations of Mathematics*. Institute for Advanced Study, 2013.
- 18 A. S. Troelstra. A note on non-extensional operations in connection with continuity and recursiveness. *Indag. Math.*, 39(5):455–462, 1977.
- 19 Anne Sjerp Troelstra and Dirk van Dalen. *Constructivism in Mathematics : An Introduction*. North-Holland, 1988.
- 20 Chuangjie Xu and Martín Hötzel Escardó. A constructive model of uniform continuity. In *Typed Lambda Calculi and Applications*, volume 7941 of *Lecture Notes in Computer Science*, pages 236–249. Springer Berlin Heidelberg, 2013.

Curry-Howard for Sequent Calculus at Last!

José Espírito Santo

Centro de Matemática, Universidade do Minho, Portugal

Abstract

This paper tries to remove what seems to be the remaining stumbling blocks in the way to a full understanding of the Curry-Howard isomorphism for sequent calculus, namely the questions: What do variables in proof terms stand for? What is co-control and a co-continuation? How to define the dual of Parigot's mu-operator so that it is a co-control operator? Answering these questions leads to the interpretation that sequent calculus is a formal vector notation with first-class co-control. But this is just the "internal" interpretation, which has to be developed simultaneously with, and is justified by, an equivalent, "external" interpretation, offered by natural deduction: the sequent calculus corresponds to a bi-directional, agnostic (w.r.t. the call strategy), computational lambda-calculus. Next, the formal duality between control and co-control is studied, in the context of classical logic. The duality cannot be observed in the sequent calculus, but rather in a system unifying sequent calculus and natural deduction.

1998 ACM Subject Classification F.4.1 Mathematical Logic

Keywords and phrases co-control, co-continuation, vector notation, let-expression, formal substitution, context substitution, computational lambda-calculus, classical logic, de Morgan duality

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.165

1 Introduction

Despite the anathema “*From an algorithmic viewpoint, the sequent calculus has no Curry-Howard isomorphism, because of the multitude of ways of writing the same proof*” [9], more than two decades of research have been dedicated to extend the Curry-Howard isomorphism to the sequent calculus. In its purest form, this is the question: if systems of combinators correspond to Hilbert systems, and the ordinary λ -calculus corresponds to natural deduction, what variant of the λ -calculus does correspond to the sequent calculus? Many computational aspects have been shown to be relevant: pattern matching [2, 19], explicit substitutions [1, 18], abstract machines [4]. But we miss a clear-cut answer to the question in its pure form. The textbook [18] says that the sequent calculus corresponds to explicit substitutions, but it immediately admits “this is just the beginning of the story”, as we will see yet again.

We might dismiss the question as closed: maybe the sequent calculus is too complex to admit such a clear-cut computational explanation. But we will not give up, because of the following reason: there are very basic questions, at the bottom of any attempt to understand the sequent calculus computationally, which remain barely uttered and scandalously unanswered; these questions we may now give an answer; and this answer opens the way to the desired clear-cut interpretation of the sequent calculus. The questions are:

- (i) What does a variable stand for in a sequent calculus proof-term?
- (ii) What is the related substitution operation?
- (iii) What is co-control?

What do we mean? Something very simple. Suppose the proof terms L_2 and L_3 represent derivations of $\Gamma \vdash A$ and $\Gamma, y : B \vdash C$ respectively; and suppose we now infer $\Gamma, x : A \supset B \vdash C$



© José Espírito Santo;

licensed under Creative Commons License CC-BY

13th International Conference on Typed Lambda Calculi and Applications (TLCA'15).

Editor: Thorsten Altenkirch; pp. 165–179



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

by left-introduction, building a derivation represented by some term $L(x, L_2, y.L_3)$. What does this x stand for? What can it be substituted for? Certainly not for a proof-term, as $L(L_1, L_2, y.L_3)$ corresponds to no derivation.

Of course we can survive by avoiding (not answering) the questions. We may say $L(x, L_2, y.L_3)$ is a particular form $\langle xL_2/y \rangle L_3$ of “explicit substitution”, the general form of which corresponds to cut, we can write cut-elimination rules with this syntax, many of them consisting of permutations of substitutions, all this even without breaking strong normalization [18]. Left introduction is reified with a cut that will to be eliminated, interpreted as a substitution that will not be executed. After all this, we still don’t know what that x stands for; and we are busy doing explicitly “substitution”, but we do not really know what substitution we mean.

Variables in proof-terms correspond to formulas in the l.h.s. of sequents; and the explicit handling of formulas in the l.h.s. of sequents is typical of the sequent calculus. Now, we have a model of what the handling of formulas in the *r.h.s.* of sequents means computationally: it is the $\lambda\mu$ -calculus, which proves that the *r.h.s.* handling is related to *control operation*. So, by mere formal duality, “co-control” operation has to be involved in the computational interpretation of the sequent calculus. But what is co-control? This is question III.

Toward the system. Of course, our starting point for a co-control operator is the $\tilde{\mu}$ -operator of the $\bar{\lambda}\mu\tilde{\mu}$ -calculus. But in $\bar{\lambda}\mu\tilde{\mu}$: variables are (and stand for) proof-terms – which allows one to represent left-introductions $L(x, L_2, y.L_3)$ as certain cuts that cannot be eliminated; and the operational meaning of $\tilde{\mu}$ is given by a reduction rule that triggers an ordinary term-substitution.

We propose to integrate the $\tilde{\mu}$ -operator in a system keeping the cut= \rightarrow redex paradigm, where the treatment of variables is very much like in the $\bar{\lambda}$ -calculus [10]: variables are not terms, but rather show in a construction that Herbelin writes xl and sees as corresponding to the structural inference of contraction. We prefer to interpret this construction logically as an inference that makes a formula passive on the l.h.s. of the sequent, and computationally as the dual to the “naming” construction aM of the $\lambda\mu$ -calculus. Additionally, in our system the operational meaning of $\tilde{\mu}$ is given by a rule that captures some sort of context, triggering some sort of “structural substitution”, as in $\lambda\mu$. Variables will stand for that sort of context (in the same way as names in $\lambda\mu$ stand for evaluation contexts or continuations), and the new “structural substitution” is the related kind of substitution. But what sort of context and structural substitution? The co-control question sends us back to the questions I and II.

The answers come from recent work about the isomorphism between $\bar{\lambda}\mu\tilde{\mu}$ and natural deduction [17], where a context-like concept named co-context was introduced in the sequent calculus side. Here we adapt the concept to our intuitionistic setting and rename it *co-continuation*. This is the last ingredient of the sequent calculus we propose, named $\bar{\lambda}\tilde{\mu}$. The system is an extension of $\bar{\lambda}$ with first-class co-control.

Computational interpretation. Having obtained the system, what is its computational interpretation? First-class co-control is a part of the interpretation, but not the whole story – it is only the story that goes beyond $\bar{\lambda}$. Fortunately, all that is needed for the rest of the story is already in place. Surprisingly, we can harvest not one but actually two clear-cut interpretations. Amazingly, the alternative can be seen through the little construction xl .

The first interpretation, the *external* one, is through natural deduction. At the basis of it is the idea, going back to [10], that l represents an evaluation context or continuation (a concept derived from the λ -calculus or natural deduction syntax), and that xl represents a

fill instruction: fill x in the hole of the continuation represented by l . This interpretation, if developed along the lines of [7, 17], is the core of the isomorphism Θ between the system $\overline{\lambda\mu}$ and the natural deduction system (named $\underline{\lambda\text{let}}$) that is designed hand-in-hand with it.

Notice the isomorphism holds at the levels of syntax and rewriting. For this reason, Θ is the ultimate term assignment, that completely side-steps the anathema cited above, and the ultimate realization of the idea, going back to Prawitz [15], of sequent calculus as a meta-system for natural deduction. But, more important, Θ is a computational interpretation, because the target system $\underline{\lambda\text{let}}$ – we will see – has a clear computational meaning itself: it is a computational λ -calculus [12, 16] agnostic w.r.t. the CBN vs CBV alternative.

The second interpretation, the *internal* one, is through the “structural substitution” operation of $\overline{\lambda\mu}$: xl is a fill instruction, but l is the stuff to be filled in the hole of the co-continuation that will land at x . Here l is taken literally, as primitive syntax, not as a continuation. So, we need a good word to describe l computationally. It could be “list” [10] or “spine” [3], but the best is “*vector*”, to suggest the (informal) “vector notation” of the λ -calculus [11]. This choice is part of a most needed re-interpretation of a 15-years-old technical result: there is a fragment of $\overline{\lambda}$, here renamed $\overrightarrow{\lambda}$, that is isomorphic to the λ -calculus [6, 3, 5], the isomorphism being essentially the map \mathcal{P} from natural deduction to sequent calculus introduced by Prawitz [15]. But someone has to say loudly that $\overrightarrow{\lambda}$ is a *formal vector notation*. This is the re-interpretation.

Summarizing, we propose two computational interpretations of $\overline{\lambda\mu}$: formal vector notation with first-class co-control, developed in Section 3; and agnostic computational λ -calculus, developed in Section 4. Section 5 looks closely at the duality between control and co-control and extracts unforeseen consequences for logical duality and structural proof theory. Please bear in mind that the words of this introduction just give an approximation of what we want to say. Let the technical development that follows speak for itself.

2 Background

Outside Section 5, we just consider intuitionistic implicational logic. Formulas(=types) are given by: $A, B, C ::= X \mid A \supset B$. In typing systems, contexts Γ are sets of declarations $x : A$ with at most one declaration per variable. In term languages, meta-substitution is denoted with square brackets, as in $[N/x]M$.

2.1 Control operation

Parigot’s $\lambda\mu$ -calculus [13] is our model for the management of formulas and (co-)variables in the r.h.s. of sequents, when the possibility of a distinguished/active/selected formula exists – a model we wish to “dualize” to the l.h.s. of sequents.

Still we diverge from the original. Let $Q := [b]((\mu a.M)N_1 \cdots N_m)$. In the original formulation of $\lambda\mu$, the reduction of Q proceeds by m applications of “structural reduction”, by which μ -abstraction consumes the arguments, one after the other, capturing contexts $[\cdot]N_i$ and triggering a “*structural substitution*”. After m such reduction steps, the resulting term $[b]\mu a.M'$ is reduced by a *renaming* rule, producing $[b/a]M'$. The same effect is obtained with a single, long-step, reduction rule for the μ -operator. Consider the context $\mathcal{C} = [b]([\cdot]N_1 \cdots N_m)$, hence $Q = \mathcal{C}[\mu a.M]$. Now apply the reduction rule $\mathcal{C}[\mu a.M] \rightarrow [\mathcal{C}/a]M$, where $[\mathcal{C}/a]M$ is

(Terms)	$t, u ::= x \mid \lambda x.t \mid \mu a.c$	(β)	$\langle \lambda x.t \mid u :: e \rangle \rightarrow \langle u \mid \tilde{\mu}x.\langle t \mid e \rangle \rangle$
(Co-terms)	$e ::= a \mid u :: e \mid \tilde{\mu}x.c$	$(\tilde{\mu})$	$\langle t \mid \tilde{\mu}x.c \rangle \rightarrow [t/x]c$
(Commands)	$c ::= \langle t \mid e \rangle$	(μ)	$\langle \mu a.c \mid e \rangle \rightarrow [e/a]c$

■ **Figure 1** The $\bar{\lambda}\mu\tilde{\mu}$ -calculus.

context substitution, in whose definition the critical case reads:¹

$$[\mathcal{C}/a](aP) = \mathcal{C}[P'] \text{ where } P' = [\mathcal{C}/a]P . \quad (1)$$

In this formulation, the single reduction rule still says that the μ -operator captures \mathcal{C} , while the definition of context substitution says that aP is a *fill instruction*: fill the hole of the \mathcal{C} that lands here with P .

This style with a single reduction rule is – see [17] – the exact reflection in natural deduction of the reduction rule for the μ -operator in the $\bar{\lambda}\mu\tilde{\mu}$ -calculus. We are calling the latter reduction rule μ – see Fig. 1 – and so it is natural to call μ the corresponding natural deduction rule². Beware that, sometimes, in $\lambda\mu$, μ names solely the “structural reduction” – which we may see as the control operation proper. In the style with a single rule, μ comprehends the whole control operation, and we will seek a single rule $\tilde{\mu}$ to comprehend the whole co-control operation – and find something different from the rule $\tilde{\mu}$ of $\bar{\lambda}\mu\tilde{\mu}$, despite the compelling symmetry of the latter.

2.2 Vector notation

The vector notation for the λ -calculus is the following definition of the λ -terms [11]:

$$M, N, P, Q ::= \lambda x.M \mid x\vec{N} \mid (\lambda x.M)N\vec{N}$$

According to this definition, λ -terms have three forms, which we call *first*, *second* and *third forms*. The advantage of this notation is that the head variable/redex is visible, and the β -normal forms are obtained by omitting the third form of terms in the above grammar.

This notation is informal, since many details are left unspecified. For instance: Are vectors a separate syntactic class, or do the second and third forms correspond to families of rules? How is substitution defined? Notice that, in the second form, it does not make sense to replace x by another term.

Let us introduce the *relaxed* vector notation:

$$M, N, P, Q ::= \lambda x.M \mid x\vec{N} \mid M\vec{N}$$

Some redundancy is now allowed, as the same ordinary λ -term can be represented in many ways. By analyzing the third form, we find four cases. One corresponds to the third form of the original vector notation, the other three may be simplified as follows:

$$\begin{aligned} (\epsilon) \quad & M\vec{} = M \\ (\pi_1) \quad & (x\vec{Q})N\vec{N} = x(\vec{Q}N\vec{N}) \\ (\pi_2) \quad & (P\vec{Q})N\vec{N} = P(\vec{Q}N\vec{N}) \end{aligned}$$

Here $\vec{}$ represents the empty vector. The simplifications were given names that will link them with the reduction rules of sequent calculus.

¹ Structural substitution is the particular case $[a([\cdot]N)/a]_-$ of context substitution.

² In [17] both rules are called σ_μ .

$$\begin{array}{c}
Ax \frac{}{\Gamma | A \vdash [] : A} \quad Cut \frac{\Gamma \vdash t : A \quad \Gamma | A \vdash k : B}{\Gamma \vdash tk : B} \quad Pass \frac{\Gamma, x : A | A \vdash k : B}{\Gamma, x : A \vdash x \hat{k} : B} \\
L \supset \frac{\Gamma \vdash u : A \quad \Gamma | B \vdash k : C}{\Gamma | A \supset B \vdash u :: k : C} \quad R \supset \frac{\Gamma, x : A \vdash t : B}{\Gamma \vdash \lambda x.t : A \supset B} \quad Act \frac{\Gamma, x : B \vdash v : B}{\Gamma | A \vdash \tilde{\mu}x.v : B}
\end{array}$$

■ **Figure 2** Typing rules of the $\bar{\lambda}\tilde{\mu}$ -calculus.

The third form of the original vector notation is a β -redex and hence it means something like “call the head function with the first argument of the vector”. The simplifications that have arisen by relaxing the vector notation can be read as rules for vector bookkeeping: garbage collect an empty vector (ϵ), or append chained vector (π_i). Much more difficult is to recognize in the second form $x\vec{N}$ an instruction for action on the vector \vec{N} : this is visible in a more general system with co-control, like the sequent calculus we are about to introduce, where the base case for vectors is not just $[]$.

3 Sequent calculus

3.1 The $\bar{\lambda}\tilde{\mu}$ -calculus

The proof-expressions of $\bar{\lambda}\tilde{\mu}$ ³ are given by the following grammar:

$$\begin{array}{l}
\text{(Terms)} \quad t, u, v ::= \lambda x.t \mid x \hat{k} \mid tk \\
\text{(Generalized vectors)} \quad k ::= [] \mid \tilde{\mu}x.v \mid u :: k
\end{array}$$

The typing rules are in Fig. 2. They handle two kinds of sequents: $\Gamma \vdash t : A$ and $\Gamma | A \vdash k : B$. The distinguished formula A in the latter is not exactly a “stoup” or a focused formula, because the operator $\tilde{\mu}x.t$ may select an arbitrary formula from the context Γ . The construction $x \hat{k}$ comes from the $\bar{\lambda}$ -calculus, but here it forms a pair with $\tilde{\mu}x.t$: logically, these are a activation/passification pair, in the style of the $\lambda\mu$ -calculus, but acting on the l.h.s. of sequents.

The notation $\tilde{\mu}x.t$ comes from $\bar{\lambda}\mu\tilde{\mu}$ [4]; but here, contrary to what happens in $\bar{\lambda}\mu\tilde{\mu}$, the reduction rule that defines the behavior of $\tilde{\mu}$ does not trigger a term-substitution, but rather a context-substitution, in the style of the above presentation of the μ -operator. The construction $x \hat{k}$ is easily recognized as the accompanying fill-instruction, and what remains is to pin down the right notion of context that $\tilde{\mu}$ will capture. The notion is this:

$$\mathcal{H} ::= x \hat{[\cdot]} \mid t([\cdot]) \mid \mathcal{H}[u :: [\cdot]]$$

These expressions are called *co-continuations*. Later we will argue they are logically dual to continuations.

The reduction rules of $\bar{\lambda}\tilde{\mu}$ are in Fig. 3. Let $\pi := \pi_1 \cup \pi_2$. The co-control rule $\tilde{\mu}$ triggers the context substitution $[\mathcal{H}/x]_{-}$, in whose definition the only non-routine case is:

$$[\mathcal{H}/x](x \hat{k}) = \mathcal{H}[k'] \text{ with } k' = [\mathcal{H}/x]k .$$

This equation gives the meaning of $x \hat{k}$: fill k in the hole of the \mathcal{H} that will substitute x . The π_i -rules employ concatenation of generalized vectors $k@k'$, defined by the obvious equations $[]@k' = k'$ and $(u :: k)@k' = u :: (k@k')$, together with $(\tilde{\mu}x.t)@k' = \tilde{\mu}x.tk'$.

³ We would have adopted the name $\bar{\lambda}\tilde{\mu}$ (to suggest $\bar{\lambda} + \tilde{\mu}$), had it not been already in use [4].

$$\begin{array}{ll}
(\beta) & (\lambda x.t)(u :: k) \rightarrow (u(\tilde{\mu}x.t))k \\
(\tilde{\mu}) & \mathcal{H}[\tilde{\mu}x.t] \rightarrow [\mathcal{H}/x]t \\
(\epsilon) & t[] \rightarrow t \\
(\pi_1) & (x\hat{\ }k)k' \rightarrow x\hat{\ }(k@k') \\
(\pi_2) & (tk)k' \rightarrow t(k@k')
\end{array}$$

■ **Figure 3** Reduction rules of $\overline{\lambda\tilde{\mu}}$.

Rule $\tilde{\mu}$ eliminates all occurrences of the $\tilde{\mu}$ -operator. The remaining rules eliminate all occurrences of cuts tk . So the $\beta\tilde{\mu}\epsilon\pi$ -normal forms correspond to a well-known representation of β -normal λ -terms. There is a critical pair generated by rules $\tilde{\mu}$ and π . This is the *call-by-name vs call-by-value* dilemma [4].

Two particular cases of the reduction rule $\tilde{\mu}$ are:

$$(\rho) \quad y\hat{\ }(\tilde{\mu}x.t) \rightarrow [y\hat{\ }[\cdot]/x]t \qquad (\sigma) \quad u(\tilde{\mu}x.t) \rightarrow [u([\cdot])/x]t$$

We let $\tau := \tilde{\mu} \setminus (\rho \cup \sigma)$. The particular case ρ may be called the *renaming* rule (as sometime does the “dual” rule in the $\lambda\mu$ -calculus). Indeed, the particular case $[y\hat{\ }[\cdot]/x]t$ of context substitution is almost indistinguishable from a substitution operation that renames variables, since the critical case of its definition reads⁴

$$[y\hat{\ }[\cdot]/x](x\hat{\ }k) = y\hat{\ }k' \text{ with } k' = [y\hat{\ }[\cdot]/x]k .$$

If we wanted a set of small step $\tilde{\mu}$ -rules, in the style of the original $\lambda\mu$ -calculus, we would have taken ρ and σ , together with $u :: (\tilde{\mu}x.t) \rightarrow \tilde{\mu}x.[x\hat{\ }(\cdot)/x]t$. The particular case $[x\hat{\ }(\cdot)/x]t$ of context substitution gives a form of “structural substitution” dual to that found in $\lambda\mu$.

3.2 The proof theory of vector notation

We consider notorious fragments of $\overline{\lambda\tilde{\mu}}$. First we do a kind of reconstruction of $\overline{\lambda}$. Next we identify two subsystems of our version of $\overline{\lambda}$ that say something about vector notation.

The $\rho\tau$ -normal-forms of $\overline{\lambda\tilde{\mu}}$ are given by:

$$t, u, v ::= \lambda x.t \mid x\hat{\ }l \mid tk \qquad k ::= l \mid \tilde{\mu}x.t \qquad l ::= [] \mid u :: l$$

These are the $\overline{\lambda}$ -expressions, if we recognize $t(\tilde{\mu}x.v)$ as the “mid-cut”, and if we ignore the other forms of explicit substitution or concatenation that are primitive in the original formulation of $\overline{\lambda}$. In terms of logical derivations, $\rho\tau$ -normalization ends in the focused fragment LJT [10] of the sequent calculus (the fragment corresponding to $\overline{\lambda}$). In this sense $\rho\tau$ -reduction is a focalization process. See the focalization theorem below, saying that $\rho\tau$ -normal forms exist and are unique.

The $\tilde{\mu}$ -rule is now restricted to the σ -rule, but the critical pair with π remains. This is solved by a syntactical trick, that chooses the call-by-name option: erase the case $\tilde{\mu}x.t$ from k 's (so that there is no more a distinction between k 's and l 's, there is a single class of *vectors* ranger over by l), but break the cut tk into the two cases tl (“head-cut”) and $t(\tilde{\mu}x.v)$. A π -redex $(x\hat{\ }l)l'$ or $(tl)l'$ is no longer a σ -redex, and a mid-cut $(y\hat{\ }l)\tilde{\mu}x.v$ or $(ul)\tilde{\mu}x.v$ can never be reduced by π (as it could in $\overline{\lambda\tilde{\mu}}$). This concludes the reconstruction of $\overline{\lambda}$.

This version of $\overline{\lambda}$ is equipped with β , σ , ϵ and π . We now consider the R -normal forms, with $R = \sigma$ or $R = \sigma \cup \epsilon \cup \pi$.

⁴ We say “almost” because, don't forget, variables are not expressions *per se*.

- In the first case, terms have the forms $\lambda x.t$, $x^{\wedge}l$ or tl , equipped with a rule β that corresponds to $\bar{\lambda}$'s β -rule followed by σ -normalization. The rules ϵ and π remain. Let us call this system $\bar{\lambda}$.
- In the second case, terms have the forms $\lambda x.t$, $x^{\wedge}l$ or $(\lambda x.t)(u :: l)$, equipped solely with a rule β that corresponds to $\bar{\lambda}$'s β -rule followed by $\sigma\epsilon\pi$ -normalization. Let us call this system $\bar{\lambda}$.

Around 15 year ago [6, 3, 5], the system $\bar{\lambda}$ was identified and proved isomorphic to the ordinary λ -calculus, with the isomorphism comprising a bijection between the sets of terms and an isomorphism of β -reduction relations. In the author's opinion, this little technical fact has a tremendous importance for the Curry-Howard isomorphism that has never been recognized. First, $\bar{\lambda}$ has a clear and revealing computational interpretation: it is a *formal vector notation*. It is a concrete definition of the notation (it says what vectors are, how substitution is defined, etc.) in perfect correspondence with a logical calculus. Second, we can reconstruct from this interpretation the interpretation of full sequent calculus $\bar{\lambda}\bar{\mu}$ by walking back the path that led from $\bar{\lambda}\bar{\mu}$ to $\bar{\lambda}$. Clearly, $\bar{\lambda}$ is a formalization of the relaxed vector notation we introduced in Section 2; and, since $\bar{\lambda}$ is the fragment of normal forms of $\bar{\lambda}\bar{\mu}$ w.r.t. the co-control rule $\bar{\mu}$, $\bar{\lambda}\bar{\mu}$ can be interpreted as a *formal, relaxed vector notation with first-class co-control*.

4 Natural deduction

The interpretation developed before is an internal and literal one: vectors are vectors; and co-control is understood in a formal way, as dual to control. We now develop a natural deduction system λlet that is isomorphic to $\bar{\lambda}\bar{\mu}$. The isomorphism gives a new, external interpretation, which recovers the view that (some) vectors are “evaluation contexts”[10]; it justifies the design of $\bar{\lambda}\bar{\mu}$, namely its notion of \mathcal{H} and reduction rule $\bar{\mu}$; finally, it allows the transfer of properties among the two systems.

4.1 The λlet -calculus

The proof-expressions of the calculus are given by:

$$\begin{array}{ll} \text{(Terms)} & M, N, P ::= \lambda x.M \mid \text{app}(H) \mid \text{let } x := H \text{ in } P \\ \text{(Heads)} & H ::= x \mid \text{hd}(M) \mid HN \end{array}$$

Notice that variables x and applications HN are not terms. A head $\text{hd}(M)$ is called a *head term*.

The typing rules are in Fig. 4. They handle two kinds of sequents: $\Gamma \vdash M : A$ and $\Gamma \triangleright H : A$. Four rules are standard, with the appropriate kind of sequent determined by the kind of expression being typed. The remaining two rules switch the kind of sequent, and are called *coercions*. However, despite the superficial impression, the two coercions are quite different, one being called weak and the other strong. Normalization will tell them apart radically.

A *normal* derivation is one without occurrences of *Let* and *SCoercion*.

► **Theorem 1** (Subformula property). *Every formula (resp. every formula, including A) occurring in a normal derivation of $\Gamma \vdash M : A$ (resp. $\Gamma \triangleright H : A$) is subformula of A or of some formula in Γ (resp. is a subformula of some formula in Γ).*

$$\begin{array}{c}
\frac{}{\Gamma, x : A \triangleright x : A} \text{Hyp} \quad \frac{\Gamma \triangleright H : A \quad \Gamma, x : A \vdash P : B}{\Gamma \vdash \text{let } x := H \text{ in } P : B} \text{Let} \quad \frac{\Gamma \triangleright H : A}{\Gamma \vdash \text{app}(H) : A} \text{WCoercion} \\
\frac{\Gamma, x : A \vdash M : B}{\Gamma \vdash \lambda x.M : A \supset B} \text{Intro} \quad \frac{\Gamma \triangleright H : A \supset B \quad \Gamma \vdash N : A}{\Gamma \triangleright HN : B} \text{Elim} \quad \frac{\Gamma \vdash M : A}{\Gamma \triangleright \text{hd}(M) : A} \text{SCoercion}
\end{array}$$

■ **Figure 4** Typing rules of λlet .

$$\begin{array}{lcl}
(\text{beta}) & \text{hd}(\lambda x.M)N & \rightarrow \text{hd}(\text{let } x := \text{hd}(N) \text{ in } M) \\
(\text{let}) & \text{let } x := H \text{ in } P & \rightarrow [H/x]P \\
(\text{triv}) & \text{app}(\text{hd}(M)) & \rightarrow M \\
(\text{head}_1) & \text{hd}(\text{app}(H)) & \rightarrow H \\
(\text{head}_2) & \mathcal{K}[\text{hd}(\text{let } x := H \text{ in } P)] & \rightarrow \text{let } x := H \text{ in } \mathcal{K}[\text{hd}(P)]
\end{array}$$

■ **Figure 5** Reduction rules of λlet .

So, the weak coercion loses information regarding the subformula property, while the strong coercion potentially violates that property.

The reduction rules of λlet are in Fig. 5. Rule `let` triggers ordinary substitution $[H/x]P$, while rule `head2` employs certain contexts that we call *continuations*:

$$\mathcal{K} ::= \text{app}([\cdot]) \mid \text{let } x := [\cdot] \text{ in } P \mid \mathcal{K}[[\cdot]N]$$

We single out two particular cases of `let`: `ren`, when $H = x$; and `sub`, when $H = \text{hd}(M)$. We put $\mathfrak{t} := \text{let} \setminus (\text{ren} \cup \text{sub})$. Let $\text{head} := \text{head}_1 \cup \text{head}_2$. Notice that rules `beta` and `head1` are relations on heads. The normal forms w.r.t. all reduction rules are given by:

$$M ::= \lambda x.M \mid \text{app}(H) \quad H ::= x \mid HN$$

That is, these normal forms are characterized by the absence of occurrences of lets and `hd()`. Lets are eliminated by `let` whereas all the other rules concur to eliminate the coercion `hd()`. So, `beta`, `triv`, `let`, `head`-reduction is *normalization*, that is, the reduction to a form corresponding to normal derivations.

4.2 Isomorphism

See Fig. 6 for the map $\Theta : \overline{\lambda\tilde{\mu}} \rightarrow \lambda\text{let}$. There is actually a function $\Theta : \overline{\lambda\tilde{\mu}}\text{-Terms} \rightarrow \lambda\text{let}\text{-Terms}$, together with an auxiliary function $\Theta : \lambda\text{let}\text{-Heads} \times \overline{\lambda\tilde{\mu}}\text{-Vectors} \rightarrow \lambda\text{let}\text{-Terms}$. Let $\Theta(t) = M$, $\Theta(u_i) = N_i$ and $\Theta(v) = P$. The idea is to map, say, $t(u_1 :: u_2 :: \tilde{\mu}x.v)$ to $\text{let } x := \text{hd}(M)N_1N_2 \text{ in } P$, and $x^\wedge(u_1 :: u_2 :: [])$ to $\text{app}(xN_1N_2)$: left-introductions are replaced by applications, inverting the associativity of non-abstractions.

► **Theorem 2 (Isomorphism).** *Map Θ is a sound bijection between the set of $\overline{\lambda\tilde{\mu}}$ -terms and the set of λlet -terms (whose inverse Ψ is shown in Fig. 7). Moreover, let R be rule β (resp. $\tilde{\mu}$, ϵ , π) of $\overline{\lambda\tilde{\mu}}$, and let R' be rule `beta` (resp. `let`, `triv`, `head`) of λlet . Then, $t \rightarrow_R t'$ in $\overline{\lambda\tilde{\mu}}$ iff $\Theta t \rightarrow_{R'} \Theta t'$ in λlet .*

The real action of the isomorphism happens in the translation of non-abstractions. Every non-abstraction λlet -term has the form $\Theta(H, k)$, and $\Psi\Theta(H, k) = \Psi(H, k)$. Every non-abstraction $\overline{\lambda\tilde{\mu}}$ -term has the form $\Psi(H, k)$, and $\Theta\Psi(H, k) = \Theta(H, k)$. So non-abstractions

$$\begin{array}{ll}
\Theta(\lambda x.t) = \lambda x.\Theta t & \Theta(H, []) = \text{app}(H) \\
\Theta(x\hat{k}) = \Theta(x, k) & \Theta(H, \tilde{\mu}x.t) = \text{let } x := H \text{ in } \Theta t \\
\Theta(tk) = \Theta(\text{hd}(\Theta t), k) & \Theta(H, u :: k) = \Theta(H\Theta u, k)
\end{array}$$

■ **Figure 6** Map $\Theta : \overline{\lambda\tilde{\mu}} \rightarrow \underline{\lambda\text{let}}$.

$$\begin{array}{ll}
\Psi(\lambda x.M) = \lambda x.\Psi M & \Psi(x, k) = x\hat{k} \\
\Psi(\text{app}(H)) = \Psi(H, []) & \Psi(\text{hd}(M), k) = (\Psi M)k \\
\Psi(\text{let } x := H \text{ in } P) = \Psi(H, \tilde{\mu}x.\Psi P) & \Psi(HN, k) = \Psi(H, (\Psi N) :: k)
\end{array}$$

■ **Figure 7** Map $\Psi : \underline{\lambda\text{let}} \rightarrow \overline{\lambda\tilde{\mu}}$

have the form $\Theta(H, k)$ (natural deduction) or $\Psi(H, k)$ (sequent calculus), and the isomorphism action between them is just to interchange Θ and Ψ in these expressions.

Ψ can be extended to continuations, establishing a bijection with vectors: $\Psi(\text{app}([\cdot])) = []$, $\Psi(\text{let } x := [\cdot] \text{ in } P) = \tilde{\mu}x.\Psi P$ and $\Psi(\mathcal{K}[[\cdot]N]) = \Psi(N) :: \Psi(\mathcal{K})$. As we knew, continuations are typed “on the left”: the sequent calculus rules for typing vectors are derived typing rules in natural deduction for typing continuations.

Similarly, Θ can be extended to co-continuations, establishing a bijection with heads: $\Theta(x\hat{[\cdot]}) = x$, $\Theta(t([\cdot])) = \text{hd}(\Theta t)$ and $\Theta(\mathcal{H}[u :: [\cdot]]) = \Theta(\mathcal{H})\Theta u$. This tells us how to type co-continuations [17]: the natural deduction rules for typing heads are derived typing rules in sequent calculus for typing co-continuations, and so co-continuations are typed “on the right”.

Let us call Θ the inverse of the bijection between continuations and vectors. Then it is easy to prove that $\Theta(H, k) = \Theta(k)[H]$. This tells us that non-abstractions in $\overline{\lambda\tilde{\mu}}$ are fill instructions, and that Θ executes these instructions. For instance, $\Theta(tk) = \Theta(\text{hd}(\Theta t), k) = \Theta(k)[\text{hd}(\Theta t)]$; so tk means “fill $\text{hd}(M)$ in the hole of continuation \mathcal{K} ”, with $M = \Theta t$ and $\mathcal{K} = \Theta(k)$; and $\Theta(tk)$ is the result of such filling. Similarly, $x\hat{k}$ means “fill x in the hole of \mathcal{K} ”. So Θ realizes again the idea, going back to Prawitz [15], that sequent calculus derivations are instructions for building natural deduction proofs.

4.3 Forgetfulness

The *forgetful map* $|\cdot|$ translates $\underline{\lambda\text{let}}$ -expressions to λ -terms by erasing occurrences of the coercion $\text{hd}(\cdot)$, forgetting the distinction between terms and heads, de-sugaring let-expressions (*i.e.* translating them as β -redexes), and mapping $|\text{app}(H)| = I|H|$, where $I = \lambda x.x$. The following results about $\overline{\lambda\tilde{\mu}}$ are proved through the analysis of this simple translation of the isomorphic calculus $\underline{\lambda\text{let}}$.

► **Theorem 3** (Strong normalization). *Every typable term of $\overline{\lambda\tilde{\mu}}$ (resp. of $\underline{\lambda\text{let}}$) is $\beta\epsilon\tilde{\mu}\pi$ -SN (resp. beta triv let head-SN).*

► **Theorem 4** (Focalization). *Every term of $\overline{\lambda\tilde{\mu}}$ has a unique $\rho\tau$ -normal form, which is a $\overline{\lambda}$ -term (representing a LJT-proof).*

4.4 Computational interpretation

We argue that $\underline{\lambda\text{let}}$ is a *bidirectional, agnostic, computational λ -calculus*. The word “bi-directional” comes from [14], where the organization of a typing system for λ -terms with two kinds of sequents (for synthesis, like $\Gamma \vdash M : A$, and for checking, like $\Gamma \triangleright H : A$) is already

found. The word “agnostic” comes from [19] and means coexistence or superimposition of call-by-name (CBN) and call-by-value (CBV).

Let us go back to Fig. 5. Rule β generates a let-expression, which can be executed by the separate rule **let**. Let-expressions enjoy “associativity”: the particular case of **head**₂ where $\mathcal{K} = \text{let } y := [\cdot] \text{ in } Q$ reads $\text{let } y := \text{hd}(\text{let } x := H \text{ in } P) \text{ in } Q \rightarrow \text{let } x := H \text{ in let } y := \text{hd}(P) \text{ in } Q$. In addition, there is a pair of reduction rules to cancel a sequence of two coercions. So we might view λlet as a sort of computational λ -calculus [12, 16] where rule **let** does not require the computation of a value prior to substitution triggering, and where a pair of trivial reduction rules (**triv** and **head**₁) eliminates odd accumulations of coercions caused by a clumsy syntax.

However, this is not the right view. **head**₁ works together with **head**₂ to reduce every non-abstraction term to one of the forms $\mathcal{K}[x]$ or $\mathcal{K}[\text{hd}(\lambda x.M)]$. This is a hint of what we see next: all reduction rules of λlet have quite clear roles in CBV and CBN computation, and through these roles we will understand how different rules **triv** and **head**₁ are.

Let us make a technical point. Rule **head**₁ is a relation on heads. As with all other reduction rules, **head**₁ generates by compatible closure a relation $\rightarrow_{\text{head}_1}$ on heads and another on terms. The relation $\rightarrow_{\text{head}_1}$ on terms would have been the same, had we taken the rule **head**₁ as the relation on terms $\mathcal{K}[\text{hd}(\text{app}(H))] \rightarrow \mathcal{K}[H]$. A similar remark applies to **beta**. In the discussion of CBN and CBV that follows, we take **head**₁ and **beta** in their alternative formulation, so that it makes sense to speak about **head**₁- or **beta**-reduction at root position of a term.

CBN and CBV are defined through priorities among reduction rules [4]:

- CBV strategy: reduction at root position of a closed, non-abstraction term with priority given to **head**.
- CBN strategy: reduction at root position of a closed, let-free, non-abstraction term with priority given to **triv**.

We will give an alternative characterization of these strategies. For CBN we need the rule $\mathcal{K}[\text{hd}(\lambda x.M)N] \rightarrow \mathcal{K}[\text{hd}([\text{hd}(N)/x]M)]$, which we call *CBN – beta*, and is obtained by **beta** followed by **let**.

► **Theorem 5 (Agnosticism).** *The following is an equivalent description of the CBN and CBV strategies. In this description, “reduction” means root-position reduction of a closed, non-abstraction term. We assume additionally that the initial term is let-free.*

- *CBV. Do head-reduction as long as possible, until the term becomes either a beta, let, or triv redex. In the two first cases, reduce and restart; in the last case, reduce to return the computed abstraction.*
- *CBN. Do head-reduction as long as possible, until the term becomes either a beta or triv redex. In the first case, reduce (with CBN – beta) and restart; in the last case, reduce to return the computed abstraction.*


This description is not in terms of priorities, but rather reveals the shared organization of the computation and the roles of the different reduction rules, which are the same in both strategies – see Fig. 8. The shared organization, in turn, shows how “superimposed” CBV and CBN are in the system, in other words, how agnostic the system is.

4.5 Epilogue

An easy consequence of the isomorphism Θ is that the space of calculi in the sequent calculus format has a mirror image in natural deduction. See Fig. 9 for a roadmap. The λ -calculus is displayed in the “imperfect world” – the terminology is inspired in Remark 2.1. of [4].

		CBV	CBN
iteration of the computation cycle	pre-processing	head ₁ + head ₂	head ₁
	real computation	beta + let	beta; let
return		triv	triv

■ **Figure 8** Agnosticism: the shared organization of CBN and CBV strategies.



■ **Figure 9** The sequent calculus/natural deduction mirror.

The two computational interpretations of sequent calculus are collected in Fig. 10. The external interpretation works like this: in Fig. 10 the shown interpretation is that of λlet ; $\bar{\lambda}\bar{\mu}$ is a language of instructions for λlet (recall Section 4.2); the behavior of $t \in \bar{\lambda}\bar{\mu}$ is the isomorphic behavior of $\Theta t \in \lambda\text{let}$ written in the language of instructions.

Why is co-control almost invisible in natural deduction? Why does it boil down to the low-profile rule let , which is just a substitution triggering rule? The explanation is in the good old associativity of “applicative terms” [10]. In natural deduction, the control operator shows up in the hidden part of an applicative terms, like $(\mu a.M)N_1 \cdots N_m$. Since we want to get to this M , we collect the outer stuff in the context-like structure $\mathcal{K} = [\cdot]N_1 \cdots N_m$ and trigger a context-substitution $[\mathcal{K}/a]M$. But, in natural deduction, the co-control operator is disguised in let -expressions $\text{let } x := HN_1 \cdots N_m \text{ in } P = \Theta(HN_1 \cdots N_m, \tilde{\mu}x.P)$, and so is already at the surface. So we may proceed by ordinary substitution $[HN_1 \cdots N_m/x]P$.

5 Duality

Let us check that co-control, as formulated in $\bar{\lambda}\bar{\mu}$, is dual to control. We will expand $\bar{\lambda}\bar{\mu}$ to a self-dual system where control and co-control are each other’s dual. This is achieved with three steps.

First step: to unify $\bar{\lambda}\bar{\mu}$ and λlet . Some hints are at the end of Section 4.2, the idea comes from [7]. Every non-abstraction term of $\bar{\lambda}\bar{\mu}$ has the form $\Psi(H, k)$. So we unify $x\hat{k}$ and tk as

sequent calculus	$\overline{\lambda\tilde{\mu}}$	internal interpretation	external interpretation
	t	λ -term in formal vector notation with first-class co-control	bi-directional, agnostic, computational λ -term
right intro.	$\lambda x.t$	λ -abstraction	λ -abstraction
contraction	$x\hat{k}$	2nd form of vector notation fill k in the co-continuation x	$\mathcal{K}[x]$
cut	tk	3rd form of vector notation	$\mathcal{K}[\text{hd}(M)]$
	k	generalized vector	continuation \mathcal{K}
axiom	\square	empty vector	$\text{app}([\cdot])$
left selection	$\tilde{\mu}x.t$	co-control operator	$\text{let } x := [\cdot] \text{ in } P$
left intro.	$u :: k$	vector constructor	$\mathcal{K}[[\cdot]N]$
cut elim. + focalization	red. rules		
key-step	β	function call	function call
left-sel. elim.	$\tilde{\mu}$	co-control operation	subst. triggering (let)
focalization	τ	co-control operation proper	subst. triggering (t)
focalization	ρ	subst. triggering (renaming)	subst. triggering (ren)
right-perm.	σ	subst. triggering	subst. triggering (sub)
right-perm.	ϵ	erasure of empty vector	return (triv)
left-perm.	π	append of iterated vectors	re-associate (head)

■ **Figure 10** Curry-Howard for sequent calculus.

$\overline{\Psi}(H, k)$ and allow in $\overline{\lambda\tilde{\mu}}$ a new syntactic class $H ::= x \mid \text{hd}(t)$. Every non-abstraction term of λlet has the form $\Theta(H, k)$. So we unify $\text{app}(H)$ and $\text{let } x := H \text{ in } P$ as $\Theta(H, k)$ and allow in λlet a new syntactic class $k ::= \square \mid \tilde{\mu}x.P$. Next let us unify $\overline{\Psi}(H, K)$ and $\Theta(H, k)$ as $\chi(H|k)$. After this we realize that $\overline{\lambda\tilde{\mu}}$ and λlet are partial views of the same system (the former lacks HN , the latter lacks $u :: k$). So let t and M range over the same set of proof terms. Continuations are internalized and coincide with vectors, co-continuations are internalized and coincide with heads, context-substitution is internalized as ordinary substitution. We work modulo $\chi(HN|k) = \chi(H|N :: k)$, which abstracts the single difference between $\overline{\lambda\tilde{\mu}}$ and λlet .

Second step: to add control. We introduce the class of “commands” $c ::= \chi(H|k)$, and a non-abstraction term is now $\mu a.c$. Continuations are now given by $k ::= a \mid \tilde{\mu}x.c \mid u :: k$. Sequents have full r.h.s.’s: for instance, $\Gamma|k : A \vdash \Delta$. Logically, we moved to classical logic.

Third step: to complete the duality. We add the dual implication $A - B$, and the class of co-terms $r ::= \tilde{\lambda}a.r \mid \tilde{\mu}x.c$. The place left vacant in the grammar of continuations by the move of $\tilde{\mu}x.c$ is occupied by the new construction $\tilde{\text{hd}}(r)$. The full suite of sequents is:

$$\Gamma \vdash t : A \mid \Delta \quad \Gamma|r : A \vdash \Delta \quad \Gamma \triangleright H : A \mid \Delta \quad \Gamma|k : A \triangleright \Delta \quad c : (\Gamma \vdash \Delta)$$

We now easily write the constructors for the inference rules relative to $A - B$, just by dualizing those of implication: the already seen $\tilde{\lambda}a.r$ (left introduction), the co-continuation $H \dot{::} r$ (right introduction), and the continuation $r \sim k$ (elimination, on the left!). The full system is given in Fig. 11. The typing rules are omitted due to space limitations, but writing them down is now just routine.

The classical, de Morgan/Gentzen duality is the duality between hypotheses and conclusions, l.h.s. and r.h.s. of sequents, conjunction and disjunction (if these were present), $A \supset B$ and $B - A$. Gentzen praised LK for its exhibiting of this duality [8]. Let us denote it by $(\tilde{\cdot})$, *justement*. At the level of types $A \supset B = \tilde{B} - \tilde{A}$ and vice-versa. Co-terms are dual

(Terms)	$t, u, M, N ::= \lambda x.t \mid \mu a.c$
(Co-terms)	$r, s ::= \tilde{\lambda} a.r \mid \tilde{\mu} x.c$
(Co-continuations)	$H ::= x \mid \mathbf{hd}(M) \mid HN \mid H \dot{::} r$
(Continuations)	$k ::= a \mid \mathbf{hd}(r) \mid r \tilde{\cdot} k \mid u \dot{::} k$
(Commands)	$c ::= \chi(H k)$
(β)	$\chi(\mathbf{hd}(\lambda x.t) u \dot{::} k) \rightarrow \chi(\mathbf{hd}(u) \tilde{\mathbf{hd}}(\tilde{\mu} x.\chi(\mathbf{hd}(t) k)))$
$(\tilde{\beta})$	$\chi(H \dot{::} s \mathbf{hd}(\tilde{\lambda} a.r)) \rightarrow \chi(\mathbf{hd}(\mu a.\chi(H \mathbf{hd}(r)) \mathbf{hd}(s)))$
(μ)	$\chi(\mathbf{hd}(\mu a.c) k) \rightarrow [k/a]c$
$(\tilde{\mu})$	$\chi(H \mathbf{hd}(\tilde{\mu} x.c)) \rightarrow [H/x]c$
(\cong)	$\chi(HN k) = \chi(H N \dot{::} k)$
(\cong)	$\chi(H \dot{::} r k) = \chi(H r \tilde{\cdot} k)$

■ **Figure 11** The unified calculus.

of terms, and vice-versa. The same for co-continuations and continuations. The notation of constructions and the naming of reduction rules self-explains how $(\tilde{\cdot})$ operates. Commands are self-dual: $\chi(\tilde{H}|k) = \chi(k|\tilde{H})$. The unified system is self-dual, at the level of typing and reduction. For instance, $\Gamma \vdash t : A|\Delta$ iff $\tilde{\Delta}|\tilde{t} : \tilde{A}|\tilde{\Gamma}$, etc.

Notice that the duality between *SC* and *ND* links HN with $r \tilde{\cdot} k$ (and $H \dot{::} r$ with $u \dot{::} k$), whereas the isomorphism between the two systems (internalized as equations in the unified system) links HN with $N \dot{::} k$ (and $H \dot{::} r$ with $r \tilde{\cdot} k$).

Let us take the unified system and forbid four constructions: $\tilde{\lambda} a.r$, HN , $H \dot{::} r$, and $r \tilde{\cdot} k$. The result is a sequent calculus, a kind of classical $\overline{\lambda\mu}$. This system, of course, is not a self-dual system; but even if we enlarge it to a self-dual sequent calculus, *SC* say, by putting back the constructions $\tilde{\lambda} a.r$ and $H \dot{::} r$ for $A - B$, the formal duality between control and co-control is not achieved yet. The formal treatment of continuations requires the presence of $u \dot{::} k$ as much as the formal treatment of co-continuations requires the presence of HN ; and, once HN is in, the last forbidden construction $r \tilde{\cdot} k$ is put back by de Morgan/Gentzen duality.

So classical $\overline{\lambda\mu}$ is not self-dual, but $\overline{\lambda\mu\tilde{\mu}}$ seemingly is [4]. Actually, a critique similar to that of classical $\overline{\lambda\mu}$ applies to $\overline{\lambda\mu\tilde{\mu}}$. Look again at Fig. 1. Despite its compelling symmetry, $\overline{\lambda\mu\tilde{\mu}}$ does not enjoy a duality between control and co-control. Why? Commands in $\overline{\lambda\mu\tilde{\mu}}$, we may say, have the form $\langle H|e \rangle$ with $H ::= \mathbf{hd}(t)$. Hence, the constructors for e 's have no dual in the class of H 's: the class of e 's is fully there, the class of H 's is residually there. We seem to see a duality between terms t and “co-terms” e , but here the word “co-terms” is a misnomer. “Co-terms” e 's are continuations, rightly captured by the μ -operator; then, either we see terms as the “dual” of continuations, and let them be captured by the $\tilde{\mu}$ -operator, but then the latter, although “dual” to the μ -operator, is not a co-control operator; or the $\tilde{\mu}$ -operator, if it is to be a co-control operator, should capture, not terms, but co-continuations, a missing kind of expression, which is also typed “on the right”; and the true co-terms are another missing kind of expressions typed “on the left”. Notice that the distortion in $\overline{\lambda\mu\tilde{\mu}}$ has nothing to do with the fact that the dual of implication is not included in Fig. 1; if it were, one would add one constructor to the grammar of terms and its “dual” to the grammar of “co-terms”, preserving the original “duality” [4], but still failing to achieve true duality, for the same reasons.

What did we learn? There is nothing wrong with classical $\overline{\lambda\mu}$, *SC* or $\overline{\lambda\mu\tilde{\mu}}$. What happens is that the classical sequent calculus, despite its symmetry, is unable to capture the formal

duality between control and co-control, because the latter requires the full extent of the de Morgan/Gentzen duality, which also involves natural deduction, and is captured only in the unified system.

6 Conclusions

Contributions. On a higher-level, this paper has two main contributions. The first is the intended one, about the Curry-Howard isomorphism for sequent calculus. If systems of combinators correspond to Hilbert systems, and the ordinary λ -calculus corresponds to natural deduction, what variant of the λ -calculus does correspond to the sequent calculus? We propose a clear-cut answer, which turns out to be a coin with two faces: sequent calculus corresponds to a formal vector notation with first-class co-control; and to a bi-directional, agnostic, computational λ -calculus.

The second contribution concerns structural proof theory. We knew from our past experience [7, 17] that sequent calculus has to be developed hand-in-hand with natural deduction. And this was confirmed here in many ways. For instance, co-continuations correspond to a primitive syntactic class in natural deduction, from where one can import, say, the typing rules. Also, things look very different in the other side of the sequent-calculus-vs-natural-deduction mirror, different to the point of un-recognizability. For instance, co-control is almost invisible in natural deduction. The surprise came when we moved to classical logic in order to prove the duality between control and co-control. There we learned that the de Morgan/Gentzen duality comprehends the duality between control and co-control, as long as we unify sequent calculus and natural deduction.

On the technical level, the main contribution is the formulation of co-control. The formulation is entirely based on the identification of the concept of co-continuation. This is the entity variables in sequent calculus proof terms stand for, and with which one may formulate the $\bar{\mu}$ -operator as a co-control operator, dualizing the behavior of the μ -operator. We showed the meaning of co-control in natural deduction, and how co-control subsumes a form of focalization. Finally, it is also noteworthy: (i) The analysis contained in the agnosticism theorem of the superimposition of CBN and CBV present in λlet (and $\bar{\lambda}\bar{\mu}$); (ii) The treatment of the logical operation $A - B$ contained in the unified system.

Related and future work. The author apologizes for the title of this paper, if the reader finds it exaggerated. True, the author believes something simultaneously new and very basic was said here about the computational interpretation of the sequent calculus. On the other hand, this contribution corresponds a small step from the wisdom accumulated before. Specifically, our proposal starts from the following ingredients: the $\bar{\lambda}$ -calculus [10], the $\bar{\mu}$ -operator [4], the vector notation [11], the technical result about formal vector notation [6, 3, 5], together with the previous work by the author [7, 17].

It is clear why co-control has been unnoticed in the theory and practice of programming: in a syntax with the natural deduction format, co-control control corresponds to a low-profile substitution triggering rule. Co-control, as such, is only visible in a syntax with the sequent calculus format. Now, such kind of syntax is usually regarded as a machine representation. Therefore, it is natural to ask whether co-control is relevant in the theory and practice of functional languages implementation. This question deserves further investigation.

Acknowledgements. The author was supported by Fundação para a Ciência e Tecnologia through project PEst-OE/MAT/UI0013/2014.

References

- 1 H. Barendregt and S. Ghilezan. Lambda terms for natural deduction, sequent calculus and cut elimination. *Journal of Functional Programming*, 10(1):121–134, 2000.
- 2 S. Cerrito and D. Kesner. Pattern matching as cut elimination. In *Proceedings of 14th annual IEEE Symposium on Logic in Computer Science (LICS'99)*, pages 98–108, 1999.
- 3 I. Cervesato and F. Pfenning. A linear spine calculus. *J. Log. Compt.*, 13(5):639–688, 2003.
- 4 P.-L. Curien and H. Herbelin. The duality of computation. In *Proceedings of the Fifth ACM SIGPLAN International Conference on Functional Programming (ICFP'00), Montreal, Canada, September 18-21, 2000*, SIGPLAN Notices 35(9), pages 233–243. ACM, 2000.
- 5 R. Dyckhoff and C. Urban. Strong normalisation of Herbelin's explicit substitution calculus with substitution propagation. *Journal of Logic and Computation*, 13(5):689–706, 2003.
- 6 J. Espírito Santo. Revisiting the correspondence between cut-elimination and normalisation. In *Proceedings of ICALP'00*, volume 1853 of *Lecture Notes in Computer Science*, pages 600–611. Springer-Verlag, 2000.
- 7 Jo. Espírito Santo. The λ -calculus and the unity of structural proof theory. *Theory of Computing Systems*, 45:963–994, 2009.
- 8 G. Gentzen. Investigations into logical deduction. In M. E. Szabo, editor, *The collected papers of Gerhard Gentzen*, pages 68–131. North Holland, 1969.
- 9 J.-Y. Girard, Y. Lafont, and P. Taylor. *Proofs and Types*. C.U.P., 1989.
- 10 H. Herbelin. A λ -calculus structure isomorphic to a Gentzen-style sequent calculus structure. In L. Pacholski and J. Tiuryn, editors, *Proceedings of CSL'94*, volume 933 of *Lecture Notes in Computer Science*, pages 61–75. Springer-Verlag, 1995.
- 11 F. Joachimski and R. Matthes. Short proofs of normalization for the simply-typed lambda-calculus, permutative conversions and Gödel's T. *Arch. for Math. Logic*, 42:59–87, 2003.
- 12 E. Moggi. Computational lambda-calculus and monads. Technical Report ECS-LFCS-88-86, University of Edinburgh, 1988.
- 13 M. Parigot. $\lambda\mu$ -calculus: an algorithmic interpretation of classic natural deduction. In *Int. Conf. Logic Prog. Automated Reasoning*, volume 624 of *Lecture Notes in Computer Science*, pages 190–201. Springer Verlag, 1992.
- 14 B. Pierce and D. Turner. Local type inference. In *Proc. of POPL'98*, pages 252–265. ACM, 1998.
- 15 D. Prawitz. *Natural Deduction*. Almqvist and Wiksell, Stockholm, 1965.
- 16 A. Sabry and P. Wadler. A reflection on call-by-value. *ACM Transactions on Programming Languages and Systems*, 19(6):916–941, 1997.
- 17 José Espírito Santo. Towards a canonical classical natural deduction system. *Ann. Pure Appl. Logic*, 164(6):618–650, 2013.
- 18 M. Sørensen and P. Urzyczyn. *Lectures on the Curry-Howard Isomorphism*. Elsevier, 2006.
- 19 Noam Zeilberger. On the unity of duality. *Ann. Pure Appl. Logic*, 153(1-3):66–96, 2008.

Dependent Types for Nominal Terms with Atom Substitutions

Elliot Fairweather¹, Maribel Fernández¹, Nora Szasz², and Alvaro Tasistro²

¹ King’s College London, UK

² Universidad ORT Uruguay, Uruguay

Abstract

Nominal terms are an extended first-order language for specifying and verifying properties of syntax with binding. Founded upon the semantics of nominal sets, the success of nominal terms with regard to systems of equational reasoning is already well established. This work first extends the untyped language of nominal terms with a notion of non-capturing atom substitution for object-level names and then proposes a dependent type system for this extended language. Both these contributions are intended to serve as a prelude to a future nominal logical framework based upon nominal equational reasoning and thus an extended example is given to demonstrate that this system is capable of encoding various other formal systems of interest.

1998 ACM Subject Classification F.4.1 Mathematical Logic, lambda calculus and related systems, D.3.3 Language Constructs and Features, data types and structures, frameworks

Keywords and phrases α -equivalence, nominal term, substitution, dependent type

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.180

1 Introduction

There exist many formal systems described by a syntax that makes use of name binding constructs. Nominal terms [35, 15, 14], are, by now, a well-established approach to the specification and verification of properties of such languages and systems. Based upon the name abstraction semantics of nominal sets [23], nominal terms use the properties of permutations of object-level names or ‘atoms’ to provide an explicit formalisation for both the use of side conditions on names and for the axiomatisation of alpha-equivalence between object-level terms. Following the development of efficient algorithms for matching [4] and unification [5, 25], nominal terms have been applied in rewriting [15, 16, 12] and unoriented equational reasoning [22, 10].

In these works, the capture-avoiding substitution used in many systems of interest has, thus far, needed to be encoded by explicit rewrite rules or axioms for the syntax in question. The first contribution of the present work addresses this issue by extending the language of nominal terms with a notion of capture-avoiding atom substitution at the object-level. Definitions of freshness, alpha-equivalence and matching together with informal proofs of decidability are provided for the new syntax.

The second contribution of this paper is the definition of a proposed dependent type system for the extended language. The language of the type system presented in this paper is user-defined; users define an interdependent signature of type- and term-constructors of interest and give their type declarations. It is the responsibility of the user to maintain the adequacy of their encoding and thus to declare types in accordance with the system to be



© Elliot Fairweather, Maribel Fernández, Nora Szasz, and Alvaro Tasistro; licensed under Creative Commons License CC-BY

13th International Conference on Typed Lambda Calculi and Applications (TLCA’15).

Editor: Thorsten Altenkirch; pp. 180–195



Leibniz International Proceedings in Informatics

LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

formalised. In keeping with earlier work on nominal equational reasoning and in contrast with previous work on nominal type theories [8, 9, 32], the system given here remains ‘lambda-free’.

Both these contributions are presented as foundations for the future development of a logical framework based upon nominal equational reasoning. A logical framework is a formal system that provides the facility to define a number of other formal systems, such as programming languages, mathematical structures and logics, by abstraction of their common features [24, 2, 27, 30, 29]. Many logical frameworks are developed as a type system; the minimum expressive power desirable of such a type system is that types be allowed to depend upon terms.

The type system presented here is one of such expressive power and to demonstrate this an extended example is given in the form of a specification for first-order logic for which adequacy is proven. Equality axioms are not yet considered in this prototype ‘framework’ but in the future it is expected that the system be expanded to include a user-defined nominal equational theory, specified as a set of equality axioms or rewriting rules in the style of [20, 15, 16].

2 Syntax

Consider countably infinite, pairwise disjoint sets of **atoms**, $a, b, c, \dots \in \mathbb{A}$, **variables**, $X, Y, Z, \dots \in \mathbb{X}$, **term-formers**, $f, g, \dots \in \mathbb{F}$, and **type-formers**, $\mathbb{C}, \mathbb{C}', \dots \in \mathbb{C}$. The syntax of **permutations**, π , **atom substitutions**, ϑ , **pseudo-terms**, s, t , and **pseudo-types**, σ, τ , is defined by mutual induction and generated by the grammar in Definition 1.

► **Definition 1** (Syntax).

$$\begin{aligned} \pi &::= \text{id} \mid \pi * (a \ b) & s, t &::= a \mid [a : \sigma] t \mid f t \mid (t_1, \dots, t_n) \mid \vartheta \mid \pi \cdot X \\ \vartheta &::= \text{id} \mid [a \mapsto t] * \vartheta & \sigma, \tau &::= [a : \sigma] \tau \mid \mathbb{C} t \mid (\tau_1 \times \dots \times \tau_n) \end{aligned}$$

A **permutation** is a bijection on the set of atoms, \mathbb{A} , represented as a list of swappings, such that $\pi(a) \neq a$ for finitely many atoms, $a \in \mathbb{A}$. $\pi(a)$ is easily computed using swappings (we omit the inductive definition). An **atom substitution** is a mapping from atoms to pseudo-terms, equal to the identity mapping but for finitely many arguments. This mapping is represented as a list of pairs, ϑ , of the form, $[a_i \mapsto t_i]$ such that the atoms, a_i , are *pairwise distinct*. This list is interpreted as a set of *simultaneous* bindings and not as a sequence, and thus the value of an atom substitution is determined directly from the syntactic representation. The final **id** and ‘list cons’ operators in the syntax for both permutations and atom substitutions are commonly omitted. Atom substitutions act upon pseudo-terms and pseudo-types by instantiating atoms and are ‘capture-avoiding’. Atom substitutions suspend upon variables and are applied *after* a suspended permutation.

The constructions for pseudo-terms are called respectively **atom terms**, **abstractions**, **function applications**, **tuples** (where $n \geq 0$) and **moderated variables** and those for pseudo-types, **abstraction types**, **constructed types** and **product types** ($n \geq 0$). Abbreviate $f()$ as f , and $\mathbb{C}()$ as \mathbb{C} . Let M, N, \dots range over elements of the union of the set of pseudo-terms and set of pseudo-types. There is only one kind of well-formed types: **type**.

► **Definition 2** (Permutation Action).

$$\begin{array}{ll}
\pi \cdot \text{id} \triangleq \text{id} & \pi \cdot ([a \mapsto t] * \vartheta) \triangleq [\pi(a) \mapsto \pi \cdot t] * (\pi \cdot \vartheta) \\
\pi \cdot a \triangleq \pi(a) & \pi \cdot (\vartheta | \pi' \cdot X) \triangleq \pi \cdot \vartheta | (\pi @ \pi') \cdot X \\
\pi \cdot [a : \sigma] t \triangleq [\pi(a) : \pi \cdot \sigma] (\pi \cdot t) & \pi \cdot [a : \sigma] \tau \triangleq [\pi \cdot a : \pi \cdot \sigma] (\pi \cdot \tau) \\
\pi \cdot f t \triangleq f (\pi \cdot t) & \pi \cdot \mathbf{C} t \triangleq \mathbf{C} (\pi \cdot t) \\
\pi \cdot (t_1, \dots, t_n) \triangleq (\pi \cdot t_1, \dots, \pi \cdot t_n) & \pi \cdot (\tau_1 \times \dots \times \tau_n) \triangleq (\pi \cdot \tau_1 \times \dots \times \pi \cdot \tau_n)
\end{array}$$

Call $a \# M$ a **freshness constraint**. Let Δ, ∇ , range over sets of freshness constraints of the form $a \# X$; call such sets **freshness contexts**. Write $\Delta \vdash a \# M$ when a derivation exists using the rules given in Definition 3 below; in rule $(atm)^\#$ we assume $a \neq b$.

► **Definition 3** (Freshness Relation).

$$\begin{array}{c}
\frac{}{\Delta \vdash a \# b} (atm)^\# \quad \frac{\text{img}^\#(\Delta, a, \vartheta | \pi \cdot X) \quad a \in \text{dom}(\vartheta)}{\Delta \vdash a \# \vartheta | \pi \cdot X} (var : aa)^\# \\
\frac{\text{img}^\#(\Delta, a, \vartheta | \pi \cdot X) \quad a \notin \text{dom}(\vartheta) \quad \pi^{-1}(a) \# X \in \Delta}{\Delta \vdash a \# \vartheta | \pi \cdot X} (var : ab)^\# \\
\frac{\Delta \vdash a \# \tau}{\Delta \vdash a \# [a : \tau] s} (abs : aa)^\# \quad \frac{\Delta \vdash a \# s \quad \Delta \vdash a \# \tau}{\Delta \vdash a \# [b : \tau] s} (abs : ab)^\# \\
\frac{\Delta \vdash a \# s_1 \dots \Delta \vdash a \# s_n}{\Delta \vdash a \# (s_1, \dots, s_n)} (tpl)^\# \quad \frac{\Delta \vdash a \# s}{\Delta \vdash a \# f s} (app)^\# \\
\frac{\Delta \vdash a \# \sigma}{\Delta \vdash a \# [a : \sigma] \tau} (abt : aa)^\# \quad \frac{\Delta \vdash a \# \tau \quad \Delta \vdash a \# \sigma}{\Delta \vdash a \# [b : \sigma] \tau} (abt : ab)^\# \\
\frac{\Delta \vdash a \# \tau_1 \dots \Delta \vdash a \# \tau_n}{\Delta \vdash a \# (\tau_1 \times \dots \times \tau_n)} (prd)^\# \quad \frac{\Delta \vdash a \# t}{\Delta \vdash a \# \mathbf{C} t} (cns)^\#
\end{array}$$

The main differences with respect to the freshness relation for nominal terms are the introduction of new rules for types, $(abt : aa)^\#$, $(abt : ab)^\#$, $(prd)^\#$ and $(cns)^\#$, and the rules for moderated variables, $(var : aa)^\#$ and $(var : ab)^\#$, which take into account suspended atom substitutions as well as suspended permutations. The notation, $\text{img}^\#(\Delta, a, \vartheta | \pi \cdot X)$, in these rules, defines the conditions necessary for freshness with regard to the suspended atom substitution and is an abbreviation of the following finite set of hypotheses.

$$\{(\Delta \vdash a \# \vartheta(\pi(b))) \vee (b \# X \in \Delta) \mid b \in \mathbb{A}, \pi(b) \in \text{dom}(\vartheta)\}$$

This ensures that the substitution will not introduce the atom a when it is applied to an instance of X . However this disjunction of conditions results in the possibility of multiple derivations $\Delta_i \vdash a \# t$ for a given freshness constraint $a \# t$. If one considers the suspended atom substitution, ϑ , to be id , the conditions upon ϑ are satisfied vacuously and the two rules clearly reduce to that for nominal terms.

Call $M \approx_\alpha N$ an **alpha-equality constraint** and write $\Delta \vdash M \approx_\alpha N$ when a derivation exists using the rules given in Definition 4 below. Note that because alpha-equivalence is defined using freshness, again multiple possible derivations may exist for a given constraint.

► **Definition 4** (Alpha-equivalence Relation).

$$\begin{array}{c}
\frac{}{\Delta \vdash a \approx_\alpha a} (atm)^\alpha \quad \frac{\forall a \in \text{ds}(\vartheta_1|\pi_1, \vartheta_2|\pi_2). a \# X \in \Delta}{\Delta \vdash \vartheta_1|\pi_1 \cdot X \approx_\alpha \vartheta_2|\pi_2 \cdot X} (var)^\alpha \\
\frac{\Delta \vdash s_1 \approx_\alpha t_1 \dots \Delta \vdash s_n \approx_\alpha t_n}{\Delta \vdash (s_1, \dots, s_n) \approx_\alpha (t_1, \dots, t_n)} (tpl)^\alpha \quad \frac{\Delta \vdash \sigma \approx_\alpha \tau \quad \Delta \vdash s \approx_\alpha t}{\Delta \vdash [a: \sigma] s \approx_\alpha [a: \tau] t} (abs : aa)^\alpha \\
\frac{\Delta \vdash s \approx_\alpha t}{\Delta \vdash f s \approx_\alpha f t} (app)^\alpha \quad \frac{\Delta \vdash \sigma \approx_\alpha \tau \quad \Delta \vdash a \# t \quad \Delta \vdash s \approx_\alpha (a b) \cdot t}{\Delta \vdash [a: \sigma] s \approx_\alpha [b: \tau] t} (abs : ab)^\alpha \\
\frac{\Delta \vdash \sigma_1 \approx_\alpha \tau_1 \dots \Delta \vdash \sigma_n \approx_\alpha \tau_n}{\Delta \vdash (\sigma_1 \times \dots \times \sigma_n) \approx_\alpha (\tau_1 \times \dots \times \tau_n)} (prd)^\alpha \quad \frac{\Delta \vdash \sigma_1 \approx_\alpha \tau_1 \quad \Delta \vdash \sigma_2 \approx_\alpha \tau_2}{\Delta \vdash [a: \sigma_1] \sigma_2 \approx_\alpha [a: \tau_1] \tau_2} (abt : aa)^\alpha \\
\frac{\Delta \vdash s \approx_\alpha t}{\Delta \vdash \mathbf{C} s \approx_\alpha \mathbf{C} t} (cns)^\alpha \quad \frac{\Delta \vdash \sigma_1 \approx_\alpha \tau_1 \quad \Delta \vdash a \# \tau_2 \quad \Delta \vdash \sigma_2 \approx_\alpha (a b) \cdot \tau_2}{\Delta \vdash [a: \sigma_1] \sigma_2 \approx_\alpha [b: \tau_1] \tau_2} (abt : ab)^\alpha
\end{array}$$

This presentation of alpha-equivalence is defined by induction on the size of the pair, (M, N) , and is both syntax-directed and decidable when considered as a recursive predicate. It is a generalisation of the notion of alpha-equivalence on nominal terms.

The only case that is not straightforward is again that of a moderated variable, $(var)^\alpha$. Here it is important to remember that both permutations and atom substitutions are finite mappings and that the image of a suspended atom substitution is given as a sub-term of the syntax of the moderated variable. Thus, the disagreement set of two suspensions, $\vartheta_1|\pi_1$ and $\vartheta_2|\pi_2$, written $\text{ds}(\vartheta_1|\pi_1, \vartheta_2|\pi_2)$, is also finite and may be defined as $\{a \mid \vartheta_1(\pi_1(a)) \not\approx_\alpha \vartheta_2(\pi_2(a)), a \in \mathbb{A}\}$, which although recursive, is of decreasing size. The reflexivity, symmetry and transitivity of the alpha-equivalence relation have been proved by adapting the proofs given in [15], themselves simplified from those in [35].

The action of an atom substitution, ϑ , upon a pseudo-term or pseudo-type, M , written, $M \vartheta$, is defined by induction in the presence of a freshness context, Δ , in Definition 5. For the sake of clarity of presentation this freshness context is not explicitly written throughout the definition. Let ϑ^{-a} denote the atom substitution ϑ restricted to the domain, $\text{dom}(\vartheta) \setminus \{a\}$. The composition of two atom substitutions, written $\vartheta_1 \circ \vartheta_2$, is defined as the atom substitution equivalent to applying ϑ_1 followed by ϑ_2 . The syntactic construction of such a composition built from two substitutions represented as sets of bindings can be defined by adapting the algorithm described in [3, 2.1]. Note that this operation itself uses the action of an atom substitution upon pseudo-term and so must be defined simultaneously with Definition 5 and is also parameterised by the freshness context, Δ .

► **Definition 5** (Action of Atom Substitution).

$$\begin{array}{l}
(\vartheta'|\pi \cdot X) \vartheta \triangleq (\vartheta' \circ \vartheta)|\pi \cdot X \quad a \vartheta \triangleq \vartheta(a); a \in \text{dom}(\vartheta) \quad a \vartheta \triangleq a; a \notin \text{dom}(\vartheta) \\
([a: \sigma] s) \vartheta \triangleq [c: \sigma \vartheta] ((a c) \cdot s) \vartheta^{-c}; \Delta \vdash c \# s, c \# \text{img}(\vartheta) \\
(f s) \vartheta \triangleq f(s \vartheta) \quad (t_1, \dots, t_n) \vartheta \triangleq (t_1 \vartheta, \dots, t_n \vartheta) \\
([a: \sigma] \tau) \vartheta \triangleq [c: \sigma \vartheta] ((a c) \cdot \tau) \vartheta^{-c}; \Delta \vdash c \# \tau, c \# \text{img}(\vartheta) \\
(\mathbf{C} s) \vartheta \triangleq \mathbf{C}(s \vartheta) \quad (\tau_1 \times \dots \times \tau_n) \vartheta \triangleq (\tau_1 \vartheta \times \dots \times \tau_n \vartheta)
\end{array}$$

The capture-avoidance of unabstracted atoms is ensured by the fact that when an atom substitution acts upon an abstraction or abstraction type, a suitable alpha-equivalent representative is first chosen with respect to the freshness context, Δ . In practice, this

presentation will result in the creation of freshness constraints for atoms, newly-generated with respect to the system as a whole, and is similar to the approach taken in [16]. A suitable ‘even fresher’ atom always exists, and it is one’s right to add constraints for that atom to the freshness context, a fact which is taken advantage of below in Definition 9. Any implementation of this definition as a recursive function must accommodate a suitable mechanism for the generation of such names; this is most easily achieved by the threading of global state throughout the function or by the use of a global choice function that returns the next available name. Atom substitutions work uniformly on alpha-equivalence classes of pseudo-terms and pseudo-types.

► **Definition 6** (Variable Substitutions). A **variable substitution** is a mapping from variables to pseudo-terms, equal to the identity mapping but for finitely many arguments, and written as a set of bindings $[X_1 \mapsto s_1] \dots [X_n \mapsto s_n]$, such that the variables, X_1, \dots, X_n , are pairwise distinct.

The action of variable substitutions upon atom substitutions, pseudo-terms and pseudo-types is given in Definition 7. A variable substitution, θ , acts upon an atom substitution, ϑ , by instantiating the variables occurring in the pseudo-terms of the image of ϑ , and is written $\vartheta\theta$. Note that the instantiation of a variable requires the application of an atom substitution to a pseudo-term and thus the action of variable substitutions is also parameterised by a freshness context, which again is left implicit in the definition below.

► **Definition 7** (Variable Substitution Action).

$$\begin{aligned} \text{id } \theta &\triangleq \text{id} & ([a \mapsto t] * \vartheta) \theta &\triangleq [a \mapsto t\theta] * (\vartheta\theta) \\ a\theta &\triangleq a \\ (\vartheta|\pi \cdot X)\theta &\triangleq (\pi \cdot \theta(X))(\vartheta\theta); X \in \text{dom}(\theta) & (\vartheta|\pi \cdot X)\theta &\triangleq (\vartheta\theta)|\pi \cdot X; X \notin \text{dom}(\theta) \\ ([a: \sigma]t)\theta &\triangleq [a: \sigma\theta](t\theta) & (ft)\theta &\triangleq f(t\theta) & (t_1, \dots, t_n)\theta &\triangleq (t_1\theta, \dots, t_n\theta) \\ ([a: \sigma]\tau)\theta &\triangleq [a: \sigma\theta](\tau\theta) & (\mathbf{C}t)\theta &\triangleq \mathbf{C}(t\theta) & (\tau_1 \times \dots \times \tau_n)\theta &\triangleq (\tau_1\theta \times \dots \times \tau_n\theta) \end{aligned}$$

The type system introduced in Section 3 requires a formalisation of matching to check that term-formers and type-formers are used in a way that is consistent with their respective type declarations. The concepts of constraint problems and matching are now therefore extended to nominal terms with atom substitutions.

Let \mathcal{C} range over freshness and alpha-equality constraints; a **constraint problem**, \mathcal{C} , is an arbitrary set of such constraints. Extend the above notations for the derivability of constraints element-wise to constraint problems; thus, write $\Delta \vdash \{C_1, \dots, C_n\}$ for $\Delta \vdash C_1, \dots, \Delta \vdash C_n$. Substitution action extends naturally to constraints and constraint problems.

► **Definition 8** (Matching Problem). Given a constraint problem, $\mathcal{C}, \{\dots, a_j \# Q_j, \dots, M_i \approx_\alpha N_i, \dots\}$, a corresponding **matching problem** is defined if $(\bigcup_i \text{vars}(M_i)) \cap (\bigcup_i \text{vars}(N_i)) = \emptyset$ and is written $\{\dots, a_j \# Q_j, \dots, M_i \stackrel{?}{\approx}_\alpha N_i, \dots\}$.

A solution to such a problem, if one exists, is a pair, (Δ, θ) , of a freshness context, Δ , and a variable substitution, θ , such that $\text{dom}(\theta) \subseteq \bigcup \text{vars}(M_i)$ and $\Delta \vdash \mathcal{C}\theta$.

Informally, this says that a matching problem is a constraint problem in which one adds the restriction that the variables in the left-hand sides of alpha-equality constraints are disjoint from the variables in the right-hand sides and that only variables in the left-hand

sides of equality constraints may be instantiated. There may be several Δ_i such that $\Delta_i \vdash \mathcal{C} \theta$. As in the case of nominal terms, one can define an ordering between solutions (Δ_i, θ) and define a most general solution as a least element in the ordering. However, unlike nominal matching, here there is no unique most general solution.

The situation is similar for logical frameworks based on the lambda calculus. The solution there is to restrict the form of matching problems. Inspired by LF, [30], the type system in Section 3 is designed so that one only needs to match against a pattern term, t , such that for any variable, $X \in \text{vars}(t)$, then $\text{id}|\pi \cdot X$ is a sub-term of t . Thus, the value of X is uniquely determined. Note that a solution may only instantiate variables in the pattern and so if an atom substitution occurs in the matched term then it can be treated as a constant sub-term. Using this assumption, if a matching problem has a solution, there is a unique most general one. The algorithm to compute it is similar to that used to check alpha-equivalence, except that when a variable sub-term, $\vartheta|\pi \cdot X$, of the pattern is being matched against, if ϑ is not id then that constraint is postponed until after the constraint for the occurrence of the $\text{id}|\pi \cdot X$ sub-term in the pattern has been solved and a unique variable substitution generated.

In addition, it is also important to note that due to the action of atom substitutions suspended upon variables, the freshness context of a solution may contain constraints for atoms, $a \notin \text{atms}(\mathcal{C})$.

► **Definition 9 (Pattern Matching Problem).** A **pattern matching problem**, consists of two pseudo-terms-in-context or two pseudo-types-in-context, $\nabla \vdash M$ and $\Delta \vdash N$, to be matched, where $\text{vars}(\nabla \vdash M) \cap \text{vars}(\Delta \vdash N) = \emptyset$ and is written $(\nabla \vdash M) \stackrel{?}{\approx}_\alpha (\Delta \vdash N)$.

A solution to such a problem, if one exists, is a variable substitution, θ , such that (∇', θ) is a solution to the matching problem $\nabla \cup \{M \stackrel{?}{\approx}_\alpha N\}$ and there exists a freshness context $\Delta^\#$ (of which each constraint, $a \# X$, is such that $a \notin \text{atms}(M) \cup \text{atms}(N) \cup \text{atms}(\nabla)$), such that $\Delta \cup \Delta^\# \vdash \nabla'$.

A **newly-freshened variant** of a term, t , is a term, written t^n , in which all the atoms and variables have been replaced by newly generated atoms and variables with respect to those occurring in t (and maybe other elements of syntax, always specified.)

Closed terms were introduced in [15] and shown there to be decidable by an algorithm using newly-freshened variants and nominal matching. Intuitively, a closed term has no unabstracted atoms and all occurrences of a variable must appear under the *same* abstracted atoms. Here, closedness can be checked in a similar way using the matching algorithm mentioned above.

3 Type System

This section starts by introducing the syntax of environments, declarations and judgements used in this type system. The validity of environments (Definition 11), validity of sets of declarations (Definition 10), and derivability of typing judgements (Definition 13), are then defined by mutual induction.

A **type association** is a pair of a variable, X , and a type, σ , written $(X : \sigma)$ or an atom and a type, written $(a : \sigma)$. A **pseudo-environment**, Γ , is an ordered list of type associations. A pseudo-environment may contain at most one type association for each variable and atom. Here, let $\Gamma \bowtie (a : \tau)$, denote the result of **appending** $(a : \tau)$ to the end of the list that represents the pseudo-environment, Γ (similarly for a variable association) and let this notation be extended element-wise to lists of associations. The association available for a given atom, a , in Γ is denoted by Γ_a . If there is no type association for a in Γ then Γ_a is

undefined, written \perp (similarly for a variable.) The domain of Γ , denoted $\text{dom}(\Gamma)$, and image of Γ , denoted $\text{img}(\Gamma)$, are defined as usual: $\text{dom}(\Gamma) = \{a \in \mathbb{A} \mid \Gamma_a \neq \perp\} \cup \{X \in \mathbb{X} \mid \Gamma_X \neq \perp\}$ and $\text{img}(\Gamma) = \{\tau \mid \exists a, \Gamma_a = \tau\} \cup \{\tau \mid \exists X, \Gamma_X = \tau\}$. It is important to note that type associations for variables are *never* appended by the typing rules, however the rules for abstractions and abstraction types do append type associations for atoms.

A **pseudo-declaration**, $\Gamma \Vdash \Delta \vdash ft: \langle \sigma \leftrightarrow \tau \rangle$ or $\Gamma \Vdash \Delta \vdash Ct: \langle \sigma \leftrightarrow \text{type} \rangle$, states the type associations and freshness constraints that a term must satisfy in order that an application or constructed type built from that term be well-formed. Thus, informally, this says that if under the type associations in Γ , and the freshness constraints in Δ , t has type σ then ft has type τ , or similarly that Ct is of the kind **type**. In practice, users need not give complete declarations; it is sufficient to write $\Gamma \Vdash \Delta \vdash ft: \tau$ or $\Gamma \Vdash \Delta \vdash Ct: \text{type}$ and the system will infer the complete declaration by computing the type of t .

Pseudo-declarations are given for a term-former or type-former *together* with an argument term in order to allow the use of atoms of that argument in the type of the application or constructed type; for example, see the declarations for all_i and all_e in Section 5.

A **pseudo-judgement**, $\Gamma \Vdash_{\Sigma} \Delta \vdash t: \tau$ or $\Gamma \Vdash_{\Sigma} \Delta \vdash \tau: \text{type}$, specifies that under a given environment, set of declarations and freshness context, either a term has a particular type or a type is well-formed.

A valid set of declarations, Σ , written $\text{validD}(\Sigma)$, is defined inductively as follows.

► **Definition 10** (Valid Set of Declarations).

- The empty set of declarations, \emptyset , is valid; $\text{validD}(\emptyset)$.
- If $\text{validD}(\Sigma)$, $\Gamma \Vdash \Delta \vdash ft: \langle \sigma \leftrightarrow \tau \rangle$ is a valid declaration under the following conditions.
 - $\Gamma \Vdash_{\Sigma} \Delta \vdash t: \sigma$, where $\text{vars}(\sigma) \subseteq \text{vars}(t) \cup \text{vars}(\tau)$
 - $\Gamma \Vdash_{\Sigma} \Delta \vdash \tau: \text{type}$, where τ is not an abstraction type
 - $\Delta \vdash \{t, \sigma, \tau\}$ is closed
 - for any variable, $X \in \text{vars}(t) \cup \text{vars}(\tau)$, then the sub-term, $\text{id}|\pi \cdot X$, occurs in either t or τ .

Then, provided that there is no declaration in Σ for the term-former, f , it holds that $\text{validD}(\Sigma \cup \{\Gamma \Vdash \Delta \vdash ft: \langle \sigma \leftrightarrow \tau \rangle\})$.

- If $\text{validD}(\Sigma)$, then $\Gamma \Vdash \Delta \vdash Ct: \langle \sigma \leftrightarrow \text{type} \rangle$ is a valid declaration under the following conditions.

- $\Gamma \Vdash_{\Sigma} \Delta \vdash t: \sigma$, where $\text{vars}(\sigma) \subseteq \text{vars}(t)$
- $\Delta \vdash \{t, \sigma\}$ is closed
- for any variable, $X \in \text{vars}(t)$ then the sub-term, $\text{id}|\pi \cdot X$, occurs in t

Then, provided that there is no declaration in Σ for the type-former, C , $\text{validD}(\Sigma \cup \{\Gamma \Vdash \Delta \vdash Ct: \langle \sigma \leftrightarrow \text{type} \rangle\})$.

Assuming a freshness context, Δ , and a valid set of declarations, Σ , a valid environment, written $\text{validE}(\Gamma, \Sigma, \Delta)$, is defined as follows.

► **Definition 11** (Valid Environments).

- The empty list of type associations, $-$, is a valid environment; $\text{validE}(-, \Sigma, \Delta)$.
- If $\text{validE}(\Gamma, \Sigma, \Delta)$, then $\text{validE}(\Gamma \bowtie (a: \tau), \Sigma, \Delta \cup \Delta')$, for any Δ' that does not mention atoms in $\text{dom}(\Gamma)$, provided that $a \notin \text{dom}(\Gamma)$, $\Gamma \Vdash_{\Sigma} \Delta \vdash \tau: \text{type}$.
- If $\text{validE}(\Gamma, \Sigma, \Delta)$, then $\text{validE}(\Gamma \bowtie (X: \tau), \Sigma, \Delta \cup \Delta')$, for any Δ' that does not mention atoms in $\text{dom}(\Gamma)$, provided that $X \notin \text{dom}(\Gamma)$, for any atom, a , such that $\Delta \vdash a \# X$ then $\Delta \vdash a \# \tau, a \# \text{img}(\Gamma)$, and $\Gamma \Vdash_{\Sigma} \Delta \vdash \tau: \text{type}$.

Informally, the final condition upon variable associations says that in a valid environment, if an atom cannot occur unabstracted in an instance of a variable, X , then it cannot occur unabstracted in the typing of that variable either.

Typing judgements are derived inductively using the rules given in Definition 13. The declarations provided by the user are required in the rules for term-formers, $(app)^\tau$, and type-formers, $(cns)^\tau$. Declarations are matched to pseudo-terms and pseudo-types using pattern matching and therefore each time an application or constructed type is typed a newly-freshened variant of the declaration for that term-former or type-former is generated; we write $\Theta \Vdash \nabla \vdash f s : \langle \sigma \hookrightarrow \rho \rangle \in \Sigma^n$ (similarly for a constructed type) to emphasise the fact that such a newly-freshened variant of a declaration is being used. In the rule for a variable, $(var)^\tau$, a predicate, typV is used.

► **Definition 12.** Let $\Gamma|_X$, for $X \in \text{dom}(\Gamma)$, be the list of associations in Γ up to but not including the pair $(X : \tau)$. Write $\text{typV}(\Gamma, \Sigma, \Delta, \vartheta|\pi, X)$ if the following conditions hold.

- for any atom $a \in \text{dom}(\Gamma|_X)$, either $\Delta \vdash a \# X$ or $\Gamma \Vdash_\Sigma \Delta \vdash \vartheta(\pi(a)) : (\pi \cdot \Gamma_a) \vartheta$
- for any variable $Y \in \text{dom}(\Gamma|_X)$, $\Gamma \Vdash_\Sigma \Delta \vdash \vartheta|\pi \cdot Y : (\pi \cdot \Gamma_Y) \vartheta$
- $\Gamma \Vdash_\Sigma \Delta \vdash (\pi \cdot \Gamma_X) \vartheta : \text{type}$

This definition indicates that for any atom, a , that can occur unabstracted in an instance of the variable, X , then $\vartheta(\pi(a))$ must be typeable with a type compatible with the type of a . Similarly, variables that may be used in the typing of X (i.e., that occur in Γ before X) also have a type compatible with $\vartheta|\pi$.

In the rule $(abs)^\tau$, the atom b is not in the set of atoms occurring in the judgement, $\Gamma \Vdash_\Sigma \Delta \vdash [a : \sigma] t : [a : \sigma] \tau$ and the freshness context, $\Delta^\#$, is such that $\Delta^\# \vdash b \# t, b \# \sigma, b \# \tau$. Similarly in $(abt)^\tau$, b is not in the set of atoms occurring in $\Gamma \Vdash_\Sigma \Delta \vdash [a : \sigma] \tau : \text{type}$ and $\Delta^\#$, is such that $\Delta^\# \vdash b \# \sigma, b \# \tau$.

In the rule $(app)^\tau$, $\text{sol}(\theta)$ means that the variable substitution, θ , is a most general solution to the pattern matching problem $(\nabla \vdash (\mathbf{C} s \times \rho)) \stackrel{?}{\approx}_\alpha (\Delta \cup \Delta^\# \vdash (\mathbf{C} t \times \tau))$, where $\Delta^\#$ is a freshness context such that $\text{atms}(s) \cup \text{atms}(\rho) \# \text{vars}(t) \cup \text{vars}(\tau)$ and \mathbf{C} , is an arbitrary type-constructor used so that the argument term and application type may both be included within the same pattern matching problem. Similarly in $(cns)^\tau$, θ is a most general solution to $(\nabla \vdash s) \stackrel{?}{\approx}_\alpha (\Delta \cup \Delta^\# \vdash t)$ where $\Delta^\#$ is $\text{atms}(s) \# \text{vars}(t)$.

► **Definition 13** (Typing rules).

$$\begin{array}{c}
\frac{\text{validE}(\Gamma, \Sigma, \Delta) \quad (\pi \cdot \Gamma_X) \vartheta = \tau \quad \text{typV}(\Gamma, \Sigma, \Delta, \vartheta|\pi, X)}{\Gamma \Vdash_\Sigma \Delta \vdash \vartheta|\pi \cdot X : \tau} (var)^\tau \\
\\
\frac{\text{validE}(\Gamma, \Sigma, \Delta) \quad \Gamma_a = \tau}{\Gamma \Vdash_\Sigma \Delta \vdash a : \tau} (atm)^\tau \quad \frac{\Gamma \bowtie (b : \sigma) \Vdash_\Sigma \Delta \cup \Delta^\# \vdash (a b) \cdot t : (a b) \cdot \tau}{\Gamma \Vdash_\Sigma \Delta \vdash [a : \sigma] t : [a : \sigma] \tau} (abs)^\tau \\
\\
\frac{\text{validE}(\Gamma, \Sigma, \Delta)}{\Gamma \Vdash_\Sigma \Delta \vdash () : ()} (tpl : 0)^\tau \quad \frac{\Gamma \Vdash_\Sigma \Delta \vdash t_1 : \tau_1 \dots \Gamma \Vdash_\Sigma \Delta \vdash t_n : \tau_n}{\Gamma \Vdash_\Sigma \Delta \vdash (t_1, \dots, t_n) : (\tau_1 \times \dots \times \tau_n)} (tpl : n)^\tau \\
\\
\frac{\Gamma \Vdash_\Sigma \Delta \vdash t : \sigma \theta \quad \Gamma \Vdash_\Sigma \Delta \vdash \tau : \text{type}}{\Gamma \Vdash_\Sigma \Delta \vdash f t : \tau} (app)^\tau \quad (\Theta \Vdash \nabla \vdash f s : \langle \sigma \hookrightarrow \rho \rangle \in \Sigma^n, \text{sol}(\theta)) \\
\\
\frac{\Gamma \bowtie (b : \sigma) \Vdash_\Sigma \Delta \cup \Delta^\# \vdash (a b) \cdot \tau : \text{type}}{\Gamma \Vdash_\Sigma \Delta \vdash [a : \sigma] \tau : \text{type}} (abt)^\tau \\
\\
\frac{\text{validE}(\Gamma, \Sigma, \Delta)}{\Gamma \Vdash_\Sigma \Delta \vdash () : \text{type}} (prd : 0)^\tau \quad \frac{\Gamma \Vdash_\Sigma \Delta \vdash \tau_1 : \text{type} \dots \Gamma \Vdash_\Sigma \Delta \vdash \tau_n : \text{type}}{\Gamma \Vdash_\Sigma \Delta \vdash (\tau_1 \times \dots \times \tau_n) : \text{type}} (prd : n)^\tau
\end{array}$$

$$\frac{\Gamma \Vdash_{\Sigma} \Delta \vdash t : \sigma \theta}{\Gamma \Vdash_{\Sigma} \Delta \vdash \mathcal{C}t : \text{type}} \text{ (cns)}^{\tau} \quad (\Theta \Vdash \nabla \vdash \mathcal{C}s : \langle \sigma \leftrightarrow \text{type} \rangle \in \Sigma^{\mathfrak{a}}, \text{sol}(\theta))$$

$$\frac{\Gamma \Vdash_{\Sigma} \Delta \vdash t : \tau \quad \Delta \vdash \sigma \approx_{\alpha} \tau}{\Gamma \Vdash_{\Sigma} \Delta \vdash t : \sigma} \text{ (}\alpha\text{)}^{\tau}$$

Note that in the rules, $(\text{atm})^{\tau}$, $(\text{var})^{\tau}$, $(\text{tpl} : 0)^{\tau}$ and $(\text{prd} : 0)^{\tau}$, the validity of the environment is needed as a premise because the environment, Γ , is not assumed to be valid and also, that in the rule, $(\text{var})^{\tau}$, the suspension, $\vartheta|\pi$, must be applied to the type, Γ_X .

4 Meta-theory

The type system presented works uniformly on α -equivalence classes of terms and types (see Theorem 18 below.) In order to prove this property, some standard properties (weakening, strengthening and validity of typing environments) are first stated. The section is concluded with the substitution theorems.

- **Theorem 14 (Type Strengthening).** 1. *If $\Gamma \Vdash_{\Sigma} \Delta \cup \Delta' \vdash t : \tau$ (resp. $\Gamma \Vdash_{\Sigma} \Delta \cup \Delta' \vdash \tau : \text{type}$) and Δ' mentions atoms that are not in $\text{dom}(\Gamma)$ then $\Gamma \Vdash_{\Sigma} \Delta \vdash t : \tau$ (resp. $\Gamma \Vdash_{\Sigma} \Delta \vdash \tau : \text{type}$).*
2. *If $\Gamma \bowtie (a_1 : \sigma_1) \dots (a_n : \sigma_n) \Vdash_{\Sigma} \Delta \vdash t : \tau$ (resp. $\Gamma \bowtie (a_1 : \sigma_1) \dots (a_n : \sigma_n) \Vdash_{\Sigma} \Delta \vdash \tau : \text{type}$) and $\Delta \vdash a_i \# t$, τ ($1 \leq i \leq n$) then $\Gamma \Vdash_{\Sigma} \Delta \vdash t : \tau$ (resp. $\Gamma \Vdash_{\Sigma} \Delta \vdash \tau : \text{type}$).*

Proof. Both parts are proved by induction on the type derivation. ◀

► **Theorem 15 (Type Weakening).**

1. *If $\Gamma \Vdash_{\Sigma} \Delta \vdash t : \tau$ (resp. $\Gamma \Vdash_{\Sigma} \Delta \vdash \tau : \text{type}$), then $\Gamma' \Vdash_{\Sigma} \Delta' \vdash t : \tau$ (resp. $\Gamma' \Vdash_{\Sigma} \Delta' \vdash \tau : \text{type}$), for any Γ' , Δ' such that $\Gamma' = \Gamma \bowtie \Gamma_1$, $\Delta' \supseteq \Delta$ and $\text{validE}(\Gamma', \Sigma, \Delta')$.*
2. *If $\Gamma \bowtie (a_1 : \sigma_1) \bowtie (a_2 : \sigma_2) \bowtie \Gamma' \Vdash_{\Sigma} \Delta \vdash t : \tau$ (resp. $\Gamma \bowtie (a_1 : \sigma_1) \bowtie (a_2 : \sigma_2) \bowtie \Gamma' \Vdash_{\Sigma} \Delta \vdash \tau : \text{type}$), then $\Gamma \bowtie (a_2 : \sigma_2) \bowtie (a_1 : \sigma_1) \bowtie \Gamma' \Vdash_{\Sigma} \Delta \vdash t : \tau$ (resp. $\Gamma \bowtie (a_2 : \sigma_2) \bowtie (a_1 : \sigma_1) \bowtie \Gamma' \Vdash_{\Sigma} \Delta \vdash \tau : \text{type}$) provided $\Delta \vdash a_1 \# \sigma_2$.*

Proof. By simultaneous induction on the type derivation. ◀

► **Theorem 16 (Validity of Typing Environments).** *For any given $\text{validD}(\Sigma)$:*

- If $\Gamma \Vdash_{\Sigma} \Delta \vdash t : \tau$ then $\text{validE}(\Gamma, \Sigma, \Delta)$ and $\Gamma \Vdash_{\Sigma} \Delta \vdash \tau : \text{type}$.*
If $\Gamma \Vdash_{\Sigma} \Delta \vdash \tau : \text{type}$ then $\text{validE}(\Gamma, \Sigma, \Delta)$.

Proof. Both parts are proved by induction on the type derivation, using weakening (in the case where the last rule used was $(\text{atm})^{\tau}$) and strengthening in the cases where the last rule used is an abstraction rule (for types or terms). ◀

► **Lemma 17.** *If $\Delta \vdash \text{ds}(\vartheta_1|\pi_1, \vartheta_2|\pi_2) \# X$ then $\Delta \vdash \text{ds}(\vartheta_1|\pi_1, \vartheta_2|\pi_2) \# \Gamma_X$.*

For any M (pseudo-term or pseudo-type), if $\Delta \vdash \text{ds}(\vartheta_1|\pi_1, \vartheta_2|\pi_2) \# M$, then $\Delta \vdash (\pi_1 \cdot M) \vartheta_1 \approx_{\alpha} (\pi_2 \cdot M) \vartheta_2$.

Proof. The first part is a consequence of the definition of valid environment. The second part is proved by induction on the definition of \approx_{α} . ◀

► **Theorem 18 (Unicity of Types).** *If $\Gamma \Vdash_{\Sigma} \Delta \vdash t : \tau_1$ and $\Gamma \Vdash_{\Sigma} \Delta \vdash t' : \tau_2$, where $\Delta \vdash t \approx_{\alpha} t'$, then $\Delta \vdash \tau_1 \approx_{\alpha} \tau_2$.*

Proof. 1. If $\Gamma \Vdash_{\Sigma} \Delta \vdash t : \tau_1$ and $\Gamma \Vdash_{\Sigma} \Delta \vdash t : \tau_2$, then $\Delta \vdash \tau_1 \approx_{\alpha} \tau_2$.

2. If $\Gamma \Vdash_{\Sigma} \Delta \vdash t : \tau$, $\Delta \vdash t \approx_{\alpha} t'$, $\Delta \vdash \Gamma \approx_{\alpha} \Gamma'$ then $\Gamma' \Vdash_{\Sigma} \Delta \vdash t' : \tau$, where $\Delta \vdash \Gamma \approx_{\alpha} \Gamma'$ indicates that for each mapping $(a : \sigma)$ (resp. $(X : \sigma)$) in Γ , Γ' contains a mapping $(a : \sigma')$ (resp. $(X : \sigma')$) such that $\Delta \vdash \sigma \approx_{\alpha} \sigma'$. Similarly for types: if $\Gamma \Vdash_{\Sigma} \Delta \vdash \tau : \text{type}$, $\Delta \vdash \tau \approx_{\alpha} \tau'$, $\Delta \vdash \Gamma \approx_{\alpha} \Gamma'$ then $\Gamma' \Vdash_{\Sigma} \Delta \vdash \tau' : \text{type}$.

The first part is proved by induction on the type derivation for $\Gamma \Vdash_{\Sigma} \Delta \vdash t : \tau_1$. The only case that is not straightforward is the case in which the derivation concludes using the rule $(app)^{\tau}$. In this case, the unicity of the type is a consequence of the fact that declarations do not overlap (that is, there is a unique declaration that matches the term considered) and matching problems have unique most general solutions under the assumptions for declarations.

The second part is proved by induction on the type derivation. In the case of a variable, Lemma 17 is needed to derive $\text{typV}(\Gamma', \Sigma, \Delta, \vartheta | \pi', X)$. ◀

► **Theorem 19** (Preservation of Types by Atom Substitution). *If $\Gamma \Vdash_{\Sigma} \Delta \vdash t : \tau$ (resp. $\Gamma \Vdash_{\Sigma} \Delta \vdash \tau : \text{type}$) and $\pi, \vartheta, \Gamma', \Delta'$ are such that:*

- $\forall X \in \text{dom}(\Gamma), \Gamma' \Vdash_{\Sigma} \Delta' \vdash \vartheta | \pi \cdot X : (\pi \cdot \Gamma_X) \vartheta$
 - $\forall a \in \text{dom}(\Gamma)$, either $\Delta \vdash a \# t$ or $\Gamma' \Vdash_{\Sigma} \Delta' \vdash \vartheta(\pi(a)) : (\pi \cdot \Gamma_a) \vartheta$
- then $\Gamma' \Vdash_{\Sigma} \Delta' \vdash (\pi \cdot t) \vartheta : (\pi \cdot \tau) \vartheta$ (resp. $\Gamma' \Vdash_{\Sigma} \Delta' \vdash (\pi \cdot \tau) \vartheta : \text{type}$)

Proof. By induction on the type derivation. In the case where the last rule applied is $(app)^{\tau}$ or $(cns)^{\tau}$, one relies on the fact that declarations are closed (that is, there are no unabstracted atoms.) The cases of abstraction rules (for terms or types) follow by induction, since atom substitutions are capture-avoiding. ◀

► **Theorem 20** (Preservation of Types by Variable Substitution). *If $\Gamma \Vdash_{\Sigma} \Delta \vdash t : \tau$ (resp. $\Gamma \Vdash_{\Sigma} \Delta \vdash \tau : \text{type}$) and θ is a variable substitution such that:*

- $\forall X \in \text{dom}(\Gamma), \Gamma' \Vdash_{\Sigma} \Delta' \vdash \theta(X) : \Gamma_X \theta$ and $\Delta' \vdash \Delta \theta$
 - $\forall a \in \text{dom}(\Gamma)$, either $\Delta \vdash a \# t$ or $\Gamma' \Vdash_{\Sigma} \Delta' \vdash a : \Gamma_a \theta$
- then $\Gamma' \Vdash_{\Sigma} \Delta' \vdash t \theta : \tau \theta$ (resp. $\Gamma' \Vdash_{\Sigma} \Delta' \vdash \tau \theta : \text{type}$)

Proof. By induction on the type derivation. In the case where the term is of the form $\theta | \pi \cdot X$ and the last rule applied is $(var)^{\tau}$ Theorem 19 for atom substitutions is used. ◀

5 Extended Example

First-order logic is a proto-typical system with binding. We consider the language of Arithmetic, and start the specification by defining type-formers for natural numbers, propositions and proofs, and term-formers to build numbers, 0 (zero) and s (successor) and propositions, bot (\perp), imp (\Rightarrow) and all (\forall).

$$\begin{aligned}
 & - \Vdash \emptyset \vdash \mathbb{N} : \text{type} \\
 & - \Vdash \emptyset \vdash \text{Prop} : \text{type} \\
 & P : \text{Prop} \Vdash \emptyset \vdash \text{Proof } P : \text{type} \\
 & - \Vdash \emptyset \vdash 0 : \mathbb{N} \\
 & X : \mathbb{N} \Vdash \emptyset \vdash s X : \mathbb{N} \\
 & - \Vdash \emptyset \vdash \text{bot} : \text{Prop} \\
 & P_1 : \text{Prop}, P_2 : \text{Prop} \Vdash \emptyset \vdash \text{imp}(P_1, P_2) : \text{Prop} \\
 & P : \text{Prop} \Vdash \emptyset \vdash \text{all}[x : \mathbb{N}] P : \text{Prop}
 \end{aligned}$$

Now, define declarations for the predicates used to build proofs; the introduction and elimination of imp and all , imp_i , imp_e , all_i and all_e and the elimination of bot , bot_e .

$$\begin{aligned}
& P_1 : \text{Prop}, P_2 : \text{Prop}, Q : \text{Proof } P_2 \\
& \Vdash x \# P_1, x \# P_2 \vdash \text{imp}_i [x : \text{Proof } P_1] Q : \text{Proof } \text{imp} (P_1, P_2) \\
& P_1 : \text{Prop}, P_2 : \text{Prop}, Q : \text{Proof } \text{imp} (P_1, P_2), Q_1 : \text{Proof } P_1 \\
& \quad \Vdash \emptyset \vdash \text{imp}_e (Q, Q_1) : \text{Proof } P_2 \\
& P : \text{Prop}, Q : \text{Proof } \text{bot} \\
& \quad \Vdash \emptyset \vdash \text{bot}_e (P, Q) : \text{Proof } P \\
& P : \text{Prop}, Q : [x : \mathbb{N}] \text{Proof } P \\
& \quad \Vdash \emptyset \vdash \text{all}_i Q : \text{Proof } \text{all} [x : \mathbb{N}] P \\
& P : \text{Prop}, Q : \text{Proof } \text{all} [x : \mathbb{N}] P, N : \mathbb{N} \\
& \quad \Vdash x \# N \vdash \text{all}_e (Q, N) : \text{Proof } [x \mapsto N] \cdot P
\end{aligned}$$

Note that in the declaration for bot_e one must use variables, P and Q , as arguments because of the restriction that all variables in the type of bot_e should occur in its arguments or in its types. Notice also that in the declaration for all_i above, the variable, Q , of type $[x : \mathbb{N}] \text{Proof } P$, is used, in other words, n is not unabstracted, as expected. The full declaration, which can be inferred is $P : \text{Prop}, Q : [x : \mathbb{N}] \text{Proof } P \Vdash \emptyset \vdash \text{all}_i Q : \langle [x : \mathbb{N}] \text{Proof } P \hookrightarrow \text{Proof } \text{all} ([x : \mathbb{N}] P) \rangle$. When this declaration is used to type terms built with all_i , pattern matching is used to obtain the values of P and Q .

An induction principle over the natural numbers could be defined as follows.

$$\begin{aligned}
& P : \text{Prop}, Q_0 : \text{Proof } [x \mapsto 0] \cdot P, Q_1 : [n : \mathbb{N}] [p : \text{Proof } [x \mapsto n] \cdot P] \text{Proof } [x \mapsto sn] \cdot P \\
& \quad \Vdash n \# P \vdash \text{ind} (Q_0, Q_1) : \text{Proof } \text{all} [x : \mathbb{N}] P
\end{aligned}$$

An encoding of a system in a logical framework is adequate if it faithfully reflects the properties of the encoded system. For instance, in the case of an encoding of first-order logic, one needs to show that the terms used in the dependent type system represent first-order terms, that formulae and proofs correspond to their standardly acknowledged notions, and that only provable propositions have a proof in the system. The goal is to prove that there is a bijection between proofs in first-order logic and the corresponding terms in this system. A formal specification of first-order logic terms, formulae and proofs is given and then it is shown that these are encoded by terms of the correct type and that encoded terms represent only well-formed terms, formulae and proofs. The theorems and proofs presented here follow closely those given in [24] to prove the adequacy of LF but are much simpler due to the fact that here lambda calculus β -reduction is not involved and therefore all terms are of canonical form.

The following grammars define the syntax of the sets of terms (Trm) and formulae (Frm) of first-order logic.

► **Definition 21** (First-order Logic Terms and Formulae).

$$\text{T}, \text{T}' ::= 0 \mid s(\text{T}) \mid x \quad \text{F}, \text{F}' ::= \perp \mid \text{F} \rightarrow \text{F}' \mid \forall x. \text{F}$$

Let $\text{fvars}(\text{T})$ and $\text{fvars}(\text{F})$ denote respectively the set of free variables in the term, T , and the formula, F . Extend this notation element-wise to sets of formulae.

A translation function $\llbracket \cdot \rrbracket$ is defined by induction on \mathbb{T} , from Trm to terms in the system using the term-formers, 0 and s , for which, see Section 5. Note that the free variables of \mathbb{T} are encoded as unabstracted atoms in $\llbracket \mathbb{T} \rrbracket$. A corresponding translation function is defined over the elements of Frm .

$$\begin{aligned} \llbracket 0 \rrbracket &= 0 & \llbracket \perp \rrbracket &= \text{bot} \\ \llbracket s(\mathbb{T}) \rrbracket &= s \llbracket \mathbb{T} \rrbracket & \llbracket F_1 \rightarrow F_2 \rrbracket &= \text{imp}(\llbracket F_1 \rrbracket, \llbracket F_2 \rrbracket) \\ \llbracket x \rrbracket &= x & \llbracket \forall x. F \rrbracket &= \text{all}[x: \mathbb{N}] \llbracket F \rrbracket \end{aligned}$$

Using the declarations given above, it can now be proved both that translated terms and formulae are typeable terms of the system and that typeable encoded terms correspond exactly to well-formed first-order logic terms and formulae.

► **Theorem 22.** *For any term, $\mathbb{T} \in \text{Trm}$, such that $\text{fvars}(\mathbb{T}) = x_1, \dots, x_n$ and freshness context, Δ , if $x_1: \mathbb{N}, \dots, x_n: \mathbb{N} \Vdash_{\Sigma} \Delta \vdash \llbracket \mathbb{T} \rrbracket: \mathbb{N}$.*

Similarly for any $F \in \text{Frm}$, such that $\text{fvars}(F) = x_1, \dots, x_n$, and freshness context, Δ , $x_1: \mathbb{N}, \dots, x_n: \mathbb{N} \Vdash_{\Sigma} \Delta \vdash \llbracket F \rrbracket: \text{Prop}$,

Proof. By structural induction on the syntax of elements of Trm and Frm . ◀

► **Theorem 23.** *If $\Gamma \Vdash_{\Sigma} \Delta \vdash t: \mathbb{N}$ is a derivable typing judgement and the environment, Γ , contains only type associations for unabstracted atoms of the form $(a: \mathbb{N})$, then $t \equiv \llbracket \mathbb{T} \rrbracket$ for some term $\mathbb{T} \in \text{Trm}$.*

Similarly, if $\Gamma \Vdash_{\Sigma} \Delta \vdash t: \text{Prop}$ is a derivable typing judgement and the environment, Γ , contains only type associations for unabstracted atoms of the form $(a: \mathbb{N})$, then $t \equiv \llbracket F \rrbracket$ for some formula $F \in \text{Frm}$.

Proof. By induction on typing judgement derivations. The only applicable cases are when the first step of the derivation is by one of the rules, (atm) , (app) or (α) . ◀

In order to show the adequacy of the encoding of proofs of first-order formulae, first, a natural deduction presentation is given for first-order logic, inspired by the one used in [24] to prove the adequacy of the encoding in LF.

Let judgements have the form $\mathcal{E} \vdash_{\mathcal{V}} P: F$, indicating that there is a proof P of the formula F , using the list of hypotheses, \mathcal{E} , and the set of free variables, \mathcal{V} , where $(\text{fvars}(\mathcal{E}) \subseteq \mathcal{V})$. The introduction rules for implication and universal quantification are shown below.

$$\frac{\mathcal{E}, (v_{F_1}: F_1) \vdash_{\mathcal{V}} P: F_2}{\mathcal{E} \vdash_{\mathcal{V} \rightarrow_i} (P \setminus v_{F_1}): F_1 \rightarrow F_2} \quad \frac{\mathcal{E} \vdash_{\mathcal{V} \cup \{x\}} P: F}{\mathcal{E} \vdash_{\mathcal{V}} \forall_i(x.P): \forall x.F}$$

Here, v_{F_1} is the variable name of a proof of F_1 , the notation $(P \setminus v_{F_1})$ denotes the proof P where v_{F_1} is discharged, and in the rule for \forall_i , the condition $\text{fvars}(\mathcal{E}) \subseteq \mathcal{V}$ on judgements implies that x is not used in \mathcal{E} .

A natural deduction judgement, J , of the form $\mathcal{E} \vdash_{\mathcal{V}} P: F$ is translated to a typing judgement of the system, $\llbracket J \rrbracket$, as follows.

$$\llbracket \mathcal{E} \vdash_{\mathcal{V}} P: F \rrbracket = \llbracket \mathcal{V} \rrbracket \times \llbracket \mathcal{E} \rrbracket \Vdash_{\Sigma} \llbracket P \rrbracket_{\text{fvars}(\llbracket J \rrbracket)}^{\Delta} \vdash \llbracket P \rrbracket: \text{Proof} \llbracket F \rrbracket$$

Here, if \mathcal{V} is $\{x_1, \dots, x_n\}$ then $\llbracket \mathcal{V} \rrbracket = x_1: \mathbb{N}, \dots, x_n: \mathbb{N}$ and $\llbracket \mathcal{E} \rrbracket$ contains $v_{F_i}: \text{Proof} \llbracket F_i \rrbracket$ for each $(v_{F_i}: F_i)$ in Γ . The translation function from proofs to terms, $\llbracket P \rrbracket$, is defined inductively; two cases are given.

$$\llbracket \forall_i(x.P) \rrbracket = \text{all}[x: \mathbb{N}] \llbracket P \rrbracket \quad \llbracket \rightarrow_i(P \setminus v_{F_1}) \rrbracket = \text{imp}_i[v_{F_1}: \text{Proof} \llbracket F_1 \rrbracket] \llbracket P \rrbracket$$

A second translation function from a proof, P , to a freshness context, parameterised by a set of variables, \mathcal{X} , and written $\llbracket P \rrbracket_{\mathcal{X}}^{\Delta}$ is also required; again two cases are given.

$$\llbracket \forall_i(x.P) \rrbracket_{\mathcal{X}}^{\Delta} = \llbracket P \rrbracket_{\mathcal{X}}^{\Delta} \quad \llbracket \rightarrow_i(P \setminus v_{F_1}) \rrbracket_{\mathcal{X}}^{\Delta} = \{v_{F_1} \# X \mid X \in \mathcal{X}\} \cup \llbracket P \rrbracket_{\mathcal{X}}^{\Delta}$$

One can now prove the following property relating natural deduction proofs in first-order logic and their encoding in the system presented here.

► **Theorem 24.** *If $\mathcal{E} \vdash_{\mathcal{V}} P : F$ is a derivable judgement, J , of natural deduction then $\llbracket \mathcal{V} \rrbracket \bowtie \llbracket \mathcal{E} \rrbracket \Vdash_{\Sigma} \llbracket P \rrbracket_{\text{vars}(\llbracket J \rrbracket)}^{\Delta} \vdash \llbracket P \rrbracket : \text{Proof } \llbracket F \rrbracket$ is a derivable typing judgement.*

Proof. By induction on the syntax of the proof, P . In the case where P is of the form $\rightarrow_i(P' \setminus v_{F_1})$ of these, the use of the rule (α) follows from the fact that $\llbracket P \rrbracket_{\text{vars}(\llbracket J \rrbracket)}^{\Delta} \vdash v_{F_1} \# \llbracket F_1 \rrbracket, v_{F_1} \# \llbracket F_2 \rrbracket$. ◀

Finally, in order to complete the adequacy proof, it is shown that only provable first-order formulae are encoded by terms of type `Proof` t .

► **Theorem 25.** *If $\Gamma \Vdash_{\Sigma} \Delta \vdash s : \text{Proof } t$ where the environment, Γ , contains only type associations for atoms either of the form $(a : N)$ or $(v_{F_1} : \text{Proof } t_i)$ where $t_i \equiv \llbracket F_i \rrbracket$ is an encoding of some formula F_i , then $s \equiv \llbracket P \rrbracket$ where P is a proof by natural deduction of some formula F such that $t \equiv \llbracket F \rrbracket$.*

Proof. By induction on typing judgement derivations. In the derivation for `imp`, $\Delta \vdash v_{F_1} \# t_1, v_{F_1} \# t_2, v_{F_1} \# \text{imp}_i[v_{F_1} : \text{Proof } t_1]s'$ and so v_{F_1} cannot be a free variable of P or F . ◀

6 Related Work

Nominal sets have been used to give semantics to systems based on nominal abstract syntax (see, for instance, [31, 18, 7, 11]) and proof theories for nominal logic have also been considered [19, 6]). Atom substitutions and their properties have been defined as systems of equational rules in [17, 21, 15]. Nominal equational theories have been investigated in [22, 10, 15] amongst others, and type systems for nominal terms and equational theories, using rank-1 polymorphic types, are defined in [14, 13, 12]. However, although proofs play an important role in all of these works, none of these systems deal explicitly with proof terms, and do not yield directly a nominal type theory. Nominal type theory has been investigated by Schöpp and Stark [34], using categorical models of nominal logic. The nominal dependent type theories developed following this approach are very expressive, but it is not clear whether their computational properties make them suitable for use in a logical framework. Nominal type theory as a basis for logical frameworks has been investigated by Cheney [8, 9], as extensions of a typed λ -calculus with names, name-abstraction and concretion operators, and name-abstraction types. A system combining λ -calculus and nominal features is also investigated by Pitts [32] to define a nominal version of Gödel's System T. A key difficulty encountered when following the approach of combining λ -calculus and nominal syntax is the interaction between name abstraction and functional abstraction (see [8] for a detailed discussion.) Westbrook [36] extends the Calculus of Inductive Constructions with a name-abstraction construct in the style of [8].

One of the best known examples of logical frameworks is LF [24], based on a typed λ -calculus with dependent types. The system presented here has similar expressive power, however there is no primitive notion of functional abstraction, instead there are term- and type- constructors in the user-defined signature. Other differences with LF include the

distinction between atoms (which can be abstracted or unabstracted), variables (which cannot be abstracted but can be instantiated, with non-capture-avoiding substitutions), and the use of name swappings (or more generally, permutations), to axiomatise α -equivalence.

Compared with previous approaches to the definition of λ -free logical frameworks [26, 1, 33], abstraction is a first-class ingredient in the syntax presented here and can be used in arguments for term- or type-constructors (as in [26, 1, 33]), or on its own (unlike [26, 1]) although user-defined constructors of abstraction type are not allowed. Also, as in [26, 1, 33], instantiation is a primitive notion in the system; it is used instead of the application operation used in λ -calculus based logical frameworks. However the approach here does not rely on explicit lists of arities and η -long normal forms as in [26, 1]. The triggering of suspended atom substitutions by instantiating variables is similar to the hereditary substitution mechanism of DMBEL [33].

There may exist similarities between the work here and that of ‘contextual modal type theory’ [28], and the dependent type system of this paper may benefit from a study of the handling of type environments and substitutions therein.

7 Conclusions and Future Work

This paper has presented a dependent type system for nominal terms with atom substitutions. A definition of matching over this syntax have been given together with an algorithm for solving such problems. This algorithm has been implemented but the complexity of problems has not been analysed; this is left for future work. A set of axioms and rules was then defined for determining the typeability of pseudo-terms and pseudo-types in this system in the presence of user-defined declarations for term-formers and type-formers. An extended example for first-order logic was presented and its adequacy proven. The type system itself has not been implemented. In its present form, the inclusion of the rule $(\alpha)^\tau$ means that the inference of derivations is not completely syntax-directed. This property, that derivable typing judgements hold for alpha-equivalent classes of types, may be derivable and if so should help in the development of a type inference algorithm for the system. Further benefits may also be gained from the inclusion of some of the more sophisticated ideas used for other type systems considered in [12]. Although we have not included computation rules in our language, dynamic features, such as reduction in the λ -calculus, or proof normalisation for a logic, may be represented using relations between terms. However, a more direct definition using equality axioms, such as in [26], would be easier to use. An extension of the logical framework to include a user-defined nominal equational theory, specified as a set of equality axioms or rewriting rules, also remains as future work.

References

- 1 Robin Adams. Lambda-free logical frameworks. *CoRR*, abs/0804.1879, 2008.
- 2 Stuart F. Allen, Robert L. Constable, Richard Eaton, Christoph Kreitz, and Lori Lorigo. The Nuprl open logical environment. In *Automated Deduction – CADE-17*, 2000.
- 3 Franz Baader and Wayne Snyder. Unification theory. In *Handbook of Automated Reasoning*. Elsevier, 2001.
- 4 Christophe Calvès and Maribel Fernández. Matching and alpha-equivalence check for nominal terms. *Journal of Computer System Sciences*, 76, 2010.
- 5 Christophe Calvès and Maribel Fernández. The first-order nominal link. In *Logic-Based Program Synthesis and Transformation – 20th International Symposium, LOPSTR 2010*,

- Hagenberg, Austria, July 23-25, 2010, Revised Selected Papers, volume 6564 of *Lecture Notes in Computer Science*, pages 234–248. Springer, 2011.
- 6 James Cheney. A simpler proof theory for nominal logic. In *FoSSaCS*, 2005.
 - 7 James Cheney. Completeness and Herbrand theorems for nominal logic. *Journal of Symbolic Logic*, 71, 2006.
 - 8 James Cheney. A simple nominal type theory. *Electronic Notes in Theoretical Computer Science*, 228, 2009.
 - 9 James Cheney. A dependent nominal type theory. *Logical Methods in Computer Science*, 8, 2012.
 - 10 Ranald A. Clouston. *Equational Logic for Names and Binding*. PhD thesis, University of Cambridge, 2010.
 - 11 Roy L. Crole and Frank Nebel. Nominal lambda calculus: An internal language for fincartesian closed categories. *Electr. Notes Theor. Comput. Sci.*, 298:93–117, 2013.
 - 12 Elliot Fairweather. *Type Systems for Nominal Terms*. PhD thesis, King’s College London, 2014.
 - 13 Elliot Fairweather, Maribel Fernández, and Murdoch J. Gabbay. Principal types for nominal theories. In *Proceedings of the 18th International Symposium on Fundamentals of Computation Theory (FCT 2011)*, 2011.
 - 14 Maribel Fernández and Murdoch J. Gabbay. Curry-style types for nominal terms. In *Types for Proofs and Programs (TYPES’06)*. Springer, 2007.
 - 15 Maribel Fernández and Murdoch J. Gabbay. Nominal rewriting. *Information and Computation*, 205, 2007.
 - 16 Maribel Fernández and Murdoch J. Gabbay. Closed nominal rewriting and efficiently computable nominal algebra equality. In *Proceedings of the 5th International Workshop on Logical Frameworks and Meta-Languages (LFMTP 2010)*, 2010.
 - 17 Murdoch J. Gabbay. A study of substitution, using nominal techniques and Fraenkel-Mostowski sets. *Theoretical Computer Science*, 410, 2009.
 - 18 Murdoch J. Gabbay. Two-level nominal sets and semantic nominal terms: an extension of nominal set theory for handling meta-variables. *Mathematical Structures in Computer Science*, 21, 2011.
 - 19 Murdoch J. Gabbay and James Cheney. A sequent calculus for nominal logic. In *Proceedings of the 19th IEEE Symposium on Logic in Computer Science (LICS 2004)*, 2004.
 - 20 Murdoch J. Gabbay and Aad Mathijssen. Nominal algebra. In *18th Nordic Workshop on Programming Theory*, 2006.
 - 21 Murdoch J. Gabbay and Aad Mathijssen. Capture-avoiding substitution as a nominal algebra. *Formal Aspects of Computing*, 20, 2008.
 - 22 Murdoch J. Gabbay and Aad Mathijssen. Nominal universal algebra: Equational logic with names and binding. *Journal of Logic and Computation*, 19, 2009.
 - 23 Murdoch J. Gabbay and Andrew M. Pitts. A new approach to abstract syntax involving binders. In *Proceedings of the 14th Annual Symposium on Logic in Computer Science (LICS 1999)*, pages 214–224. IEEE Computer Society Press, July 1999.
 - 24 Robert Harper, Furio Honsell, and Gordon Plotkin. A framework for defining logics. In *Proceedings of the 2nd IEEE Symposium on Logic in Computer Science (LICS 1987)*. IEEE Computer Society Press, 1987.
 - 25 Jordi Levy and Mateu Villaret. An efficient nominal unification algorithm. In *Proceedings of the 21st International Conference on Rewriting Techniques and Applications (RTA 2010)*, 2010.
 - 26 Zhaohui Luo. PAL⁺: a lambda-free logical framework. *Journal of Functional Programming*, 13, 2003.

- 27 Lena Magnusson and Bengt Nordström. The ALF proof editor and its proof engine. In *Types for Proofs and Programs, (TYPES'93)*, 1994.
- 28 Aleksandar Nanevski, Frank Pfenning, and Brigitte Pientka. Contextual modal type theory. *ACM Transactions on Computational Logic*, 9, 2008.
- 29 Bengt Nordström, Kent Petersson, and Jan Smith. *Programming in Martin-Löf's Type Theory*. Oxford University Press, 1990.
- 30 Frank Pfenning and Carsten Schürmann. System description: Twelf – a meta-logical framework for deductive systems. In *Automated Deduction – CADE-16*, 1999.
- 31 Andrew M. Pitts. Nominal logic: A first order theory of names and binding. In *Proceedings of the 4th International Symposium on Theoretical Aspects of Computer Software (STACS 2001)*, 2001.
- 32 Andrew M. Pitts. Nominal system T. In *Proceedings of the 37th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages (POPL 2010)*, 2010.
- 33 Gordon Plotkin. An algebraic framework for logics and type theories, 2006. Talk given at LFMTTP'06.
- 34 Ulrich Schöpp and Ian Stark. A Dependent Type Theory with Names and Binding. In *CSL*, 2004.
- 35 Christian Urban, Andrew M. Pitts, and Murdoch J. Gabbay. Nominal unification. *Theoretical Computer Science*, 323, 2004.
- 36 Edwin Westbrook. *Higher-order Encodings with Constructors*. PhD thesis, Washington University in St. Louis, 2008.

Realizability Toposes from Specifications*

Jonas Frey

Department of Computer Science
University of Copenhagen, Denmark
jofr@di.ku.dk

Abstract

We investigate a framework of Krivine realizability with I/O effects, and present a method of associating realizability models to *specifications* on the I/O behavior of processes, by using adequate interpretations of the central concepts of *pole* and *proof-like term*. This method does in particular allow to associate realizability models to computable functions.

Following recent work of Streicher and others we show how these models give rise to *triposes* and *toposes*.

1998 ACM Subject Classification F.4.1 Mathematical Logic

Keywords and phrases Krivine machine, classical realizability, realizability topos, bisimulation

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.196

1 Introduction

Krivine realizability with side effects has been introduced by Miquel in [14]. In this article we demonstrate how an instance of Miquel’s framework including I/O instructions allows to associate *realizability toposes* to *specifications*, i.e. sets of requirements imposed on the I/O behavior of programs. Since the requirement to compute a specific function f can be viewed as a specification, we do in particular obtain a way to *associate toposes to computable functions*.

These toposes are different from traditional ‘Kleene’ realizability toposes such as the *effective topos* [6] in that we associate toposes to *individual* computable functions, whereas the effective topos incorporates *all* recursive functions on equal footing. Another difference to the toposes based on Kleene realizability is that the internal logic of the latter is *constructive*, whereas the present approach is based on Krivine’s realizability interpretation [10], which validates classical logic.

To represent specifications we make use of the fact that Krivine’s realizability interpretation is parametric over a set of processes called the *pole*. The central observation (Lemma 27 and Theorem 28) is that non-trivial specifications on program behavior give rise to poles leading to consistent (i.e. non-degenerate) interpretations.

To give a categorical account of Krivine realizability we follow recent work of Streicher [18] and others [17, 20, 2], which demonstrates how Krivine realizability models give rise to *triposes*. Toposes are then obtained via the tripos-to-topos construction [7].

Our basic formalism is an extension of the Krivine machine (2) that gives an operational semantics to I/O instructions for single bits. We give two formulations of the operational semantics – one (3) in terms of a transition relation on processes including a *state* (which is adequate for reasoning about function computation), and one (4) in terms of a labeled

* This work is supported by the Danish Council for Independent Research Sapere Aude grant “Complexity through Logic and Algebra” (COLA).



© Jonas Frey;

licensed under Creative Commons License CC-BY

13th International Conference on Typed Lambda Calculi and Applications (TLCA’15).

Editor: Thorsten Altenkirch; pp. 196–210



Leibniz International Proceedings in Informatics

LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

transition system admitting to reason about program equivalence in terms of bisimulation. The two operational semantics are related by Corollary 7, which we use to prove a Turing completeness result in Theorem 12.

1.1 Related work

The idea of adding instructions with new evaluation rules to the machine plays a central role in Krivine’s writings, as a means to realize non-logical axioms. Citing from [11]:

“Indeed, when we realize usual axioms of mathematics, we need to introduce, one after the other, the very standard tools in system programming: for the law of Peirce, these are continuations (particularly useful for exceptions); for the axiom of dependent choice, these are the clock and the process numbering; for the ultrafilter axiom and the well ordering of \mathbb{R} , these are no less than I/O instructions on a global memory, in other words assignment.”

Although features like exceptions and memory are often called *effects*, it is arguable whether they should be called *side effects*, since they do not interact with the outside world.

The idea to add instructions for *side effects* which are influenced by – and influence – the outside world, has already been investigated by Miquel [14, Section 2.2], and our execution relation (3) can be viewed as an instance of his framework.

What sets the present approach apart is that Miquel views the state of the world (represented by a forcing condition) as being *part* of a process and requires poles to be saturated w.r.t. all (including effectful) reductions, whereas for us poles are sets of ‘bare’ processes without state, which are saturated only w.r.t. reduction free of side-effects.

This difference is crucial in that it enables the construction of poles from specifications.

2 Syntax and machine

In this section we recall Krivine’s abstract machine with continuations as described in [10]. We then go on to describe an extension of the syntax by I/O instructions, and describe an operational semantics as a transition relation on triples (p, ι, o) of process, input, and output.

2.1 Krivine’s machine

We recall the underlying syntax and machine of Krivine’s classical realizability from [10]. The syntax consists of three syntactic classes called *terms*, *stacks*, and *processes*.

$$\begin{array}{lll}
 \text{Terms:} & t ::= x \mid \lambda x.t \mid tt \mid \mathbf{c} \mid \mathbf{k}_\pi & \\
 \text{Stacks:} & \pi ::= \pi_0 \mid t \cdot \pi & t \text{ closed, } \pi_0 \in \Pi_0 \\
 \text{Processes:} & p ::= t \star \pi & t \text{ closed}
 \end{array} \tag{1}$$

Thus, the terms are the terms of the λ -calculus, augmented by a constant \mathbf{c} for call/cc, and continuation terms \mathbf{k}_π for any stack π . A stack, in turn, is a list of closed terms terminated by an element π_0 of a designated set Π_0 of *stack constants*. A process is a pair $t \star \pi$ of a closed term and a stack. The set of closed terms is denoted by Λ , the set of stacks is Π , and the set of processes is $\Lambda \star \Pi$.

Krivine’s machine is now defined by a transition relation \succ on processes called *evaluation*.

$$\begin{array}{lll}
 \text{(push)} & tu \star \pi & \succ \quad t \star u \cdot \pi \\
 \text{(pop)} & (\lambda x.t[x]) \star u \cdot \pi & \succ \quad t[u] \star \pi \\
 \text{(save)} & \mathbf{c} \star t \cdot \pi & \succ \quad t \star \mathbf{k}_\pi \cdot \pi \\
 \text{(restore)} & \mathbf{k}_\pi \star t \cdot \rho & \succ \quad t \star \pi
 \end{array} \tag{2}$$

The first two rules implement *weak head reduction* of λ -terms, and the third and fourth rule capture and restore continuations.

2.2 The machine with I/O

To incorporate I/O we modify the syntax as follows:

$$\begin{array}{ll} \text{Terms:} & t ::= x \mid \lambda x.t \mid tt \mid \mathfrak{c} \mid k_\pi \mid r \mid \mathfrak{w}0 \mid \mathfrak{w}1 \mid \text{end} \\ \text{Stacks:} & \pi ::= \varepsilon \mid t \cdot \pi \qquad t \text{ closed} \\ \text{Processes:} & p ::= t \star \pi \mid \top \qquad t \text{ closed} \end{array}$$

The grammar for terms is extended by constants $r, \mathfrak{w}0, \mathfrak{w}1, \text{end}$ for reading, writing and termination, and in exchange the stack constants are omitted – ε is the empty stack. Finally there is a *process constant* \top also representing termination – the presence of both end and \top will be important in Section 3.

We write Λ_e and Π_e for the sets of terms and stacks of the syntax with I/O, and P for the set of processes. Furthermore, we denote by Λ_p the set of *pure* terms, i.e. terms not containing any of $r, \mathfrak{w}0, \mathfrak{w}1, \text{end}$.

The operational semantics of the extended syntax is given in terms of *execution contexts*, which are triples (p, ι, o) of a process p , and a pair $\iota, o \in \{0, 1\}^*$ of binary strings representing input and output. On these execution contexts, we define the *execution relation* \rightsquigarrow as follows:

$$\begin{array}{ll} (\tau) & (t \star \pi, \iota, o) \rightsquigarrow (u \star \rho, \iota, o) \quad \text{whenever } t \star \pi \succ u \star \rho \\ (\mathfrak{r}0) & (r \star t \cdot u \cdot v \cdot \pi, 0\iota, o) \rightsquigarrow (t \star \pi, \iota, o) \\ (\mathfrak{r}1) & (r \star t \cdot u \cdot v \cdot \pi, 1\iota, o) \rightsquigarrow (u \star \pi, \iota, o) \\ (\mathfrak{r}\varepsilon) & (r \star t \cdot u \cdot v \cdot \pi, \varepsilon, o) \rightsquigarrow (v \star \pi, \varepsilon, o) \\ (\mathfrak{w}0) & (\mathfrak{w}0 \star t \cdot \pi, \iota, o) \rightsquigarrow (t \star \pi, \iota, 0o) \\ (\mathfrak{w}1) & (\mathfrak{w}1 \star t \cdot \pi, \iota, o) \rightsquigarrow (t \star \pi, \iota, 1o) \\ (\mathfrak{e}) & (\text{end} \star \pi, \iota, o) \rightsquigarrow (\top, \iota, o) \end{array} \quad (3)$$

Thus, if there is neither of $r, \mathfrak{w}0, \mathfrak{w}1, \text{end}$ in head position, the process is reduced as in (2) without changing ι and o . If r is in head position, the computation selects one of the first three arguments depending on whether the input starts with a 0, a 1, or is empty. $\mathfrak{w}0$ and $\mathfrak{w}1$ write out 0 and 1, and end discards the stack and yields \top , which represents successful termination.

We observe that the execution relation is *deterministic*, i.e. for every execution context there is at most one transition possible, which is determined by the term in head position, and in case of r also by the input.

2.3 Representing functions

We view the above formalism as a model of computation that explicitly includes reading of input, and writing of output.

Consequently, when thinking about expressivity we are not so much interested in the ability of the machine to transform abstract representations of data like ‘Church numerals’, but rather in the functions on binary strings that processes can compute by reading their argument from the input, and writing the result to the output.

► **Definition 1.** For $n \in \mathbb{N}$, $\text{bin}(n) \in \{0, 1\}^*$ is the base 2 representation of n . 0 is represented by the empty string, thus we have e.g. $\text{bin}(0) = \varepsilon$, $\text{bin}(1) = 1$, $\text{bin}(2) = 10$, $\text{bin}(3) = 11$, ...

A process p is said to *implement* a partial function $f : \mathbb{N} \rightarrow \mathbb{N}$, if $(p, \text{bin}(n), \varepsilon) \rightsquigarrow^* (\top, \varepsilon, \text{bin}(f(n)))$ for all $n \in \text{dom}(f)$.

► **Remark.** There is a stronger version of the previous definition which requires $(p, \text{bin}(n), \varepsilon)$ to diverge or block for $n \notin \text{dom}(f)$, and a completeness result like Thm. 12 can be shown for the strengthened definition as well.

We use the weaker version, since we expect the poles \perp_f defined in Section 5.2.1 to be better behaved this way.

2.4 β -reduction

To talk about contraction of single β -redexes which are not necessarily in head position in a process p , we define *contexts* – which are terms/stacks/processes with a single designated hole $[\cdot]$ in term position – by the following grammar:

$$\begin{array}{lll} \text{Term contexts:} & t[\cdot] ::= [\cdot] \mid \lambda x.t[\cdot] \mid t[\cdot]t \mid tt[\cdot] \mid \mathbf{k}_{\pi[\cdot]} \\ \text{Stack contexts:} & \pi[\cdot] ::= t \cdot \pi[\cdot] \mid t[\cdot] \cdot \pi & t, t[\cdot] \text{ closed} \\ \text{Process contexts:} & p[\cdot] ::= t[\cdot] \star \pi \mid t \star \pi[\cdot] \end{array}$$

Contexts are used to talk about substitution that allows capturing of variables – as described in [1, 2.1.18], given a context $t[\cdot]/\pi[\cdot]/p[\cdot]$ and a term u , $t[u]/\pi[u]/p[u]$ is the result of replacing the hole $[\cdot]$ in $t[\cdot]/\pi[\cdot]/p[\cdot]$ by u , allowing potential free variables in u to be captured. We say that u is *admissible* for $t[\cdot]/\pi[\cdot]/p[\cdot]$, if $t[u]/\pi[u]/p[u]$ is a valid term/stack/process conforming to the closedness condition for terms making up stacks.

Now we can express β -reduction as the action of contracting a single redex: given a redex $(\lambda x.u)v$ which is admissible for a context $t[\cdot]/\pi[\cdot]/p[\cdot]$, we have

$$t[(\lambda x.u)v] \rightarrow_{\beta} t[u[v/x]] \quad \pi[(\lambda x.u)v] \rightarrow_{\beta} \pi[u[v/x]] \quad p[(\lambda x.u)v] \rightarrow_{\beta} p[u[v/x]],$$

and any single β -reduction can uniquely be written this way. β -equivalence \simeq_{β} is the equivalence relation generated by β -reduction.

3 Bisimulation and \top -equivalence

To reason efficiently about execution of processes with side effects – in particular to show Turing completeness in Section 4 – we want to show that although the computation model imposes a deterministic reduction strategy, we can perform β -reduction anywhere in a process without changing its I/O behavior.

The natural choice of concept to capture ‘equivalence of I/O behavior’ is *weak bisimilarity* (see [13, Section 4.2]), and in order to make this applicable to processes we have to reformulate the operational semantics as a *labeled transition system* (LTS).

We use the set $\mathcal{L} = \{r0, r1, r\varepsilon, w0, w1, e\}$ of labels, where $r0, r1$ represent reading of a 0 or 1, respectively, and $w0, w1$ represent writing of bits. $r\varepsilon$ represents the unsuccessful attempt of reading on empty input, and e represents successful termination. The set $Act = \mathcal{L} \cup \{\tau\}$ of *actions* contains the labels as well as the symbol τ representing a ‘silent’ transition, that is used to represent effect-free evaluation.

The transition system on processes is now given as follows.

$$\begin{array}{lll} (\lambda x.t[x]) \star t \cdot \pi & \xrightarrow{\tau} & t[u] \star \pi & r \star t \cdot u \cdot v \cdot \pi & \xrightarrow{r0} & t \star \pi & w0 \star t \cdot \pi & \xrightarrow{w0} & t \star \pi \\ tu \star \pi & \xrightarrow{\tau} & t \star u \cdot \pi & r \star t \cdot u \cdot v \cdot \pi & \xrightarrow{r1} & u \star \pi & w1 \star t \cdot \pi & \xrightarrow{w1} & t \star \pi \\ \mathbf{c} \star t \cdot \pi & \xrightarrow{\tau} & t \star \mathbf{k}_{\pi} \cdot \pi & r \star t \cdot u \cdot v \cdot \pi & \xrightarrow{r\varepsilon} & v \star \pi & \text{end} \star \pi & \xrightarrow{e} & \top \\ \mathbf{k}_{\pi} \star t \cdot \rho & \xrightarrow{\tau} & t \star \pi & & & & & & \end{array} \quad (4)$$

Observe that the τ -transitions are in correspondence with the transitions of the evaluation relation (2), and the labeled transitions correspond to the remaining transitions of the execution relation (3).

We now recall the definition of *weak bisimulation relation* from [13, Section 4.2].

► **Definition 2.**

- For processes p, q we write $p \xrightarrow{\tau} q$ for $p \xrightarrow{\tau}^* q$, and for $\alpha \neq \tau$ we write $p \xrightarrow{\alpha} q$ for $\exists p', q'. p \xrightarrow{\tau} p' \xrightarrow{\alpha} q' \xrightarrow{\tau} q$.
- A *weak bisimulation* on \mathbf{P} is a binary relation $R \subseteq \mathbf{P}^2$ such that for all $\alpha \in \text{Act}$ and $(p, q) \in R$ we have

$$\begin{aligned} p \xrightarrow{\alpha} p' &\Rightarrow \exists q'. q \xrightarrow{\alpha} q' \wedge (p', q') \in R \quad \text{and} \\ q \xrightarrow{\alpha} q' &\Rightarrow \exists p'. p \xrightarrow{\alpha} p' \wedge (p', q') \in R. \end{aligned} \tag{5}$$

- Two processes p, q are called *weakly bisimilar* (written $p \approx q$), if there exists a weak bisimulation relation R with $(p, q) \in R$.

We recall the following important properties of the weak bisimilarity relation \approx .

► **Lemma 3.** *Weak bisimilarity is itself a weak bisimulation, and furthermore it is an equivalence relation.*

Proof. [13, Proposition 4.2.7] ◀

To show that β -equivalent processes are bisimilar, we have to find a bisimulation relation containing β -equivalence. The following relation does the job.

► **Definition 4** (γ -equivalence). γ -equivalence (written $p \simeq_{\gamma} q$) is the equivalence relation on processes that is generated by β -reduction and τ -transitions.

► **Lemma 5.** *γ -equivalence of processes is a weak bisimulation.*

Proof. It is sufficient to verify conditions (5) on the generators of γ -equivalence, i.e. one-step β -reductions and τ -transitions. Therefore we show the following:

1. if $p \xrightarrow{\tau} q$ and $p \xrightarrow{\alpha} p'$ then there exists q' with $q \xrightarrow{\alpha} q'$ and $p' \simeq_{\gamma} q'$
2. if $p \xrightarrow{\tau} q$ and $q \xrightarrow{\alpha} q'$ then there exists p' with $p \xrightarrow{\alpha} p'$ and $p' \simeq_{\gamma} q'$
3. if $p \rightarrow_{\beta} q$ and $p \xrightarrow{\alpha} p'$ then there exists q' with $q \xrightarrow{\alpha} q'$ and $p' \simeq_{\gamma} q'$
4. if $p \rightarrow_{\beta} q$ and $q \xrightarrow{\alpha} q'$ then there exists p' with $p \xrightarrow{\alpha} p'$ and $p' \simeq_{\gamma} q'$

In the first case, the fact that the LTS can only branch if r is in head position, and this does not involve τ -transitions, implies that $\alpha = \tau$ and $p' = q$, and we can choose $q' = q$ as well. In the second case we have $p \xrightarrow{\alpha} q'$ and thus can choose $p' = q'$.

For cases 3 and 4, which we treat simultaneously, we have to analyze the structure of p and q , which are of the form $r[(\lambda x. s)t]$ and $r[s[t/x]]$ for some context $r[\cdot]$ (see Section 2.4). The proof proceeds by cases on the structure of $r[\cdot]$.

If $r[\cdot]$ is of either of the forms $(st \star \pi)[\cdot]$ ¹, $\mathbf{w0} \star \pi[\cdot]$, $\mathbf{w1} \star \pi[\cdot]$, $\mathbf{c} \star \pi[\cdot]$, $(\mathbf{k}_{\pi} \star \rho)[\cdot]$, or $\mathbf{end} \star \pi[\cdot]$, then it is immediately evident that p and q can perform the same unique transition (if any), and the results will again be β -equivalent (possibly trivially, since the redex can get deleted in the transition).

If $r[\cdot]$ is of the form $((\lambda y. u) \star \pi)[\cdot]$ then this is true as well, regardless of whether the hole is in u or in π (here the redex can be duplicated, if the hole is in the first term in π).

¹ The notation is meant to convey that we don't care if the hole is in s , t , or π .

If $r[\cdot]$ is of the form $r \star \pi[\cdot]$ then several transitions may be possible, but any transition taken by either of p or q can also be taken by the other, and the results will again be β -equivalent.

It remains to consider $r[\cdot]$ of the form $[\cdot] \star \pi$. In this case, $p = (\lambda x. s) t \star \pi$ and $q = s[t/x] \star \pi$. Here p can perform the transition $p \xrightarrow{\tau} (\lambda x. s) \star t \cdot \pi$ which can be matched by $q \xrightarrow{\tau} q$ where we have $(\lambda x. s) \star t \cdot \pi \simeq_{\gamma} p \simeq_{\gamma} q$. In the other direction we have $p \xrightarrow{\alpha} q'$ for every $q \xrightarrow{\alpha} q'$ since $p \xrightarrow{\tau} q$. \blacktriangleleft

The following definition and corollary makes the link between the execution relation (3) and the LTS (4).

► **Definition 6.** Two execution contexts (p, ι, o) , (q, ι', o') are called \top -equivalent (written $(p, \iota, o) \sim_{\top} (q, \iota', o')$), if for all $\iota'', o'' \in \{0, 1\}^*$ we have

$$(p, \iota, o) \rightsquigarrow^* (\top, \iota'', o'') \quad \text{iff} \quad (q, \iota', o') \rightsquigarrow^* (\top, \iota'', o'').$$

► **Corollary 7.**

1. $p \approx q$ implies $(p, \iota, o) \sim_{\top} (q, \iota, o)$ for all $\iota, o \in \{0, 1\}^*$.
2. $(p, \iota, o) \rightsquigarrow^* (q, \iota', o')$ implies $(p, \iota, o) \sim_{\top} (q, \iota', o')$.
3. $(p, \iota, o) \sim_{\top} (\top, \iota', o')$ implies $(p, \iota, o) \rightsquigarrow^* (\top, \iota', o')$.

Proof. For the first claim we show that

$$p \approx q, \quad (p, \iota, o) \rightsquigarrow^* (\top, \iota', o') \quad \text{implies} \quad (q, \iota, o) \rightsquigarrow^* (\top, \iota', o')$$

by induction on the length of $(p, \iota, o) \rightsquigarrow^* (\top, \iota', o')$. The base case is clear. For the induction step assume that $(p, \iota, o) \rightsquigarrow (p^*, \iota^*, o^*) \rightsquigarrow^* (\top, \iota', o')$. If the initial transition is a (τ) in the execution relation (3), then have $p^* \cong q$, $\iota^* = \iota$ and $o^* = o$, and we can apply the induction hypothesis. If the initial transition corresponds to another clause in (3), then there is a corresponding transition $p \xrightarrow{\alpha} q$ with $\alpha \in \mathcal{L}$ in the LTS (4), and by bisimilarity there exists a q^* with $q \xrightarrow{\alpha} q^*$ and $p^* \approx q^*$. Now the induction hypothesis implies $(q^*, \iota^*, o^*) \rightsquigarrow^* (\top, \iota', o')$, and from $q \xrightarrow{\alpha} q^*$ we can deduce $(q, \iota, o) \rightsquigarrow^* (q^*, \iota^*, o^*)$ by cases on α .

The second claim follows since \rightsquigarrow is deterministic.

The third claim follows since (\top, ι', o') can not perform any more transitions. \blacktriangleleft

4 Expressivity

In this section we show that the machine with I/O is Turing complete, i.e. that every computable $f : \mathbb{N} \rightarrow \mathbb{N}$ can be implemented in the sense of Def. 1 by a process p .

Roughly speaking, given f , we define a process p that reads the input, transforms it into a Church numeral, applies a term t that computes f on the level of Church numerals, and then writes the result out.

To decompose the task we define terms R and W for reading and writing, with the properties that $(R \star \pi, \text{bin}(n), o) \sim_{\top} (\bar{n} \star \pi, \varepsilon, o)$ (\bar{n} is the n -th Church numeral), and $(W \bar{n} \star \pi, \iota, \varepsilon) \sim_{\top} (\top, \iota, \text{bin}(n))$ for all $n \in \mathbb{N}$.

Now the naive first attempt to combine R and W with the term t computing the function would be something like $W(tR)$, but this would only work if the operational semantics was call by value. The solution is to use Krivine's *storage operators* [9] which were devised precisely to simulate call by value in call by name, and we use a variation of them.

The following definition introduces the terms R and W , after giving some auxiliary definitions.

► **Definition 8.**

E, Z, B, C, H, Y, S are λ -terms satisfying

$$\begin{array}{lll} B\bar{n} \simeq_\beta \overline{2n} & E\overline{(2n)}st \simeq_\beta s & Yt \simeq_\beta t(Yt) \\ C\bar{n} \simeq_\beta \overline{2n+1} & E\overline{(2n+1)}st \simeq_\beta t & \\ H\bar{n} \simeq_\beta \overline{\text{floor}(n/2)} & Z\overline{(0)}st \simeq_\beta s & \\ S\bar{n} \simeq_\beta \overline{n+1} & Z\overline{(n+1)}st \simeq_\beta t & \end{array}$$

for all terms s, t and $n \in \mathbb{N}$, where \bar{n} is the Church numeral $\lambda fx. f^n x$.²

The terms F, R, W are defined as follows:

$$\begin{array}{l} F = \lambda hy. h(Sy)^3 \\ R = YQ\bar{0} \quad \text{where} \quad Q = \lambda xn. r(x(Bn))(x(Cn))n \\ W = YV \quad \text{where} \quad V = \lambda xn. Zn \text{end}(En(\mathbf{w0}x(Hn))(\mathbf{w1}x(Hn))) \end{array}$$

The next three lemmas explain the roles of the terms R, F , and W .

► **Lemma 9.** For all $n \in \mathbb{N}$, $\pi \in \Pi$ and $o \in \{0, 1\}^*$ we have $(R \star \pi, \text{bin}(n), o) \sim_\top (\bar{n} \star \pi, \varepsilon, o)$.

Proof. For all $n \in \mathbb{N}$ we have $YQ\bar{n} \simeq_\beta Q(YQ)\bar{n} \simeq_\beta r(YQ\overline{(2n)})(YQ\overline{(2n+1)})\bar{n}$, and thus

$$\begin{array}{l} (YQ\bar{n} \star \pi, \varepsilon, o) \sim_\top (\bar{n} \star \pi, \varepsilon, o) \\ (YQ\bar{n} \star \pi, 0\iota, o) \sim_\top (YQ\overline{(2n)} \star \pi, \iota, o) \\ (YQ\bar{n} \star \pi, 1\iota, o) \sim_\top (YQ\overline{(2n+1)} \star \pi, \iota, o) \end{array}$$

The claim follows by induction on the length of $\text{bin}(n)$, since $\text{bin}(2n) = \text{bin}(n)0$ for $n > 0$, and $\text{bin}(2n+1) = \text{bin}(n)1$. ◀

► **Lemma 10.** For $n \in \mathbb{N}$ and t any closed term, we have $\bar{n}Ft\bar{0} \simeq_\beta t\bar{n}$.

Proof. This is because $\bar{n}Ft\bar{0} \simeq_\beta F^n t\bar{0} \simeq_\beta t(S^n\bar{0}) \simeq_\beta t\bar{n}$, where the second step can be shown by induction on n . ◀

► **Lemma 11.** For all $n \in \mathbb{N}$, $\pi \in \Pi$ and $\iota \in \{0, 1\}^*$ we have $(W\bar{n} \star \pi, \iota, \varepsilon) \sim_\top (\top, \iota, \text{bin}(n))$.

Proof. We have $W\bar{n} \simeq_\beta VW\bar{n} \simeq_\beta Z\bar{n} \text{end}(E\bar{n}(\mathbf{w0}W(H\bar{n}))(\mathbf{w1}W(H\bar{n})))$, and therefore

$$\begin{array}{lll} (W\bar{0} \star \pi, \iota, o) & \sim_\top (\text{end} \star \pi, \iota, o) & \sim_\top (\top, \iota, o) \\ (W\overline{(2n)} \star \pi, \iota, o) & \sim_\top (\mathbf{w0}W(H\overline{(2n)}) \star \pi, \iota, o) & \sim_\top (W(\bar{n}) \star \pi, \iota, 0o) \quad \text{for } (n > 0) \\ (W\overline{(2n+1)} \star \pi, \iota, o) & \sim_\top (\mathbf{w1}W(H\overline{(2n+1)}) \star \pi, \iota, o) & \sim_\top (W(\bar{n}) \star \pi, \iota, 1o). \end{array}$$

The claim follows again by induction on the length of $\text{bin}(n)$. ◀

► **Theorem 12.** Every computable function $f : \mathbb{N} \rightarrow \mathbb{N}$ can be implemented by a process p .

Proof. From [5, Thm. 4.23] we know that there exists a term t with $t\bar{n} \simeq_\beta \overline{f(n)}$ for $n \in \text{dom}(f)$. The process p is given by $R \star F \cdot t \cdot \bar{0} \cdot F \cdot W \cdot \bar{0}$. Indeed, for $n \in \text{dom}(f)$ we have

$$\begin{array}{lll} (R \star F \cdot t \cdot \bar{0} \cdot F \cdot W \cdot \bar{0}, \text{bin}(n), \varepsilon) & \sim_\top (\bar{n} \star F \cdot t \cdot \bar{0} \cdot F \cdot W \cdot \bar{0}, \varepsilon, \varepsilon) & \sim_\top (\bar{n} Ft\bar{0} \star F \cdot W \cdot \bar{0}, \varepsilon, \varepsilon) \\ & \sim_\top (t\bar{n} \star F \cdot W \cdot \bar{0}, \varepsilon, \varepsilon) & \sim_\top (\overline{f(n)} \star F \cdot W \cdot \bar{0}, \varepsilon, \varepsilon) \\ & \sim_\top (W\overline{f(n)} \star \varepsilon, \varepsilon, \varepsilon) & \sim_\top (\top, \varepsilon, \text{bin}(f(n))) \end{array}$$

and we deduce $(R \star F \cdot t \cdot \bar{0} \cdot F \cdot W \cdot \bar{0}, \text{bin}(n), \varepsilon) \rightsquigarrow^* (\top, \varepsilon, \text{bin}(f(n)))$ by Corollary 7-3. ◀

² Such terms exist by elementary λ -calculus, see e.g. [5, Chapters 3,4]. In particular, Y is known as *fixed point operator*.

³ This is (part of) a *storage operator* for Church numerals [9].

5 Realizability and Tripeses

The aim of this section is to describe how the presence of I/O instructions allows to define new realizability models, which we do in the categorical language of tripeses and toposes[21].

In Subsection 5.1 we give a categorical reading of Krivine’s realizability interpretation as described in [10] and show how it gives rise to tripeses. In Subsection 5.2 we show how the definitions can be adapted to the syntax and machine with I/O, and how this allows us to define new realizability models from specifications.

The interpretation of Krivine realizability in terms of tripeses is due to Streicher [18], and has further been explored in [2]. However, the presentation here is more straightforward since the constructions and proofs do not rely on *ordered combinatory algebras*, but directly rephrase Krivine’s constructions categorically.

5.1 Krivine’s classical realizability

Throughout this subsection we work with the syntax (1) without I/O instructions but with stack constants.

Krivine’s realizability interpretation is always given relative to a set of processes called a ‘pole’ – the choice of pole determines the interpretation.

► **Definition 13.** A *pole* is a set $\perp \subseteq \Lambda \star \Pi$ of processes which is *saturated*, in the sense that $p \in \perp$ and $p' \succ p$ implies $p' \in \perp$.

As Miquel [15] demonstrated, the pole can be seen as playing the role of the parameter R in Friedman’s negative translation [3]. In the following we assume that a pole \perp is fixed.

A *truth value* is by definition a set $S \subseteq \Pi$ of stacks. Given a truth value S and a term t , we write $t \Vdash S$ – and say ‘ t realizes S ’ – if $\forall \pi \in S. t \star \pi \in \perp$. We write $S^\perp = \{t \in \Lambda \mid t \Vdash S\}$ for the set of realizers of Π . So unlike in Kleene realizability the elements of a truth value are not its realizers – they should rather be seen as ‘refutations’, and indeed larger subsets of Π represent ‘falsier’ truth values⁴; in particular falsity is defined as

$$\perp = \Pi.$$

Given truth values $S, T \subseteq \Pi$, we define the implication $S \Rightarrow T$ as follows.

$$S \Rightarrow T = S^\perp \cdot T = \{s \cdot \pi \mid s \Vdash S, \pi \in T\}$$

With these definitions we can formulate the following lemma, which relates refutations of a truth value S with realizers of its negation.

► **Lemma 14.** *Given $\pi \in S \subseteq \Pi$, we have $k_\pi \Vdash S \Rightarrow \perp$.*

Proof. We have to show that $k_\pi \star t \cdot \rho \in \perp$ for all $t \Vdash S$ and $\rho \in \Pi$. This is because $k_\pi \star t \cdot \rho \succ t \star \pi$, where $\pi \in S$ and $t \Vdash S$. ◀

A (*semantic*) *predicate* on a set I is a function $\varphi : I \rightarrow P(\Pi)$ from I to truth values. On semantic predicates we define the basic logical operations of falsity, implication, universal

⁴ For this reason, Miquel [15, 16] calls the elements of $P(\Pi)$ *falsity values*.

quantification, and reindexing by

$$\begin{aligned}
 \perp(i) &= \perp && \text{(falsity)} \\
 (\varphi \Rightarrow \psi)(i) &= \varphi(i) \Rightarrow \psi(i) = \varphi(i)^\perp \cdot \psi(i) && \text{(implication)} \\
 \forall_f(\theta)(i) &= \bigcup_{f(j)=i} \theta(j) && \text{(universal quantification)} \\
 f^*\varphi &= \varphi \circ f && \text{(reindexing)}
 \end{aligned} \tag{6}$$

for $\varphi, \psi : I \rightarrow P(\Pi)$, $\theta : J \rightarrow P(\Pi)$ and $f : J \rightarrow I$. Thus, for any function $f : J \rightarrow I$, the function \forall_f (called ‘universal quantification along f ’) maps predicates on J to predicates on I ⁵, and the function f^* (called ‘reindexing along f ’) maps predicates on I to predicate on J . We write \forall_I for universal quantification along the terminal projection $I \rightarrow 1$.

Next, we come to the concept of ‘truth/validity’ of the interpretation. We can not simply call a truth value ‘true’ if it has a realizer – this would lead to inconsistency as soon as the pole \perp is nonempty, since $\mathbf{k}_\pi t \Vdash \perp$ for any process $t \star \pi \in \perp$. The solution is to single out a set PL of ‘well-behaved’ realizers called ‘proof-like terms’. We recall the definition from [10].

► **Definition 15.** The set $\text{PL} \subseteq \Lambda$ of *proof-like terms* is the set of terms t that do not contain any continuations \mathbf{k}_π .

As Krivine [10, pg. 2] points out, t is a proof-like term if and only if it does not contain any stack constant $\pi_0 \in \Pi_0$ (since continuation terms \mathbf{k}_π necessarily contain a stack constant at the end of π , and conversely stacks can only occur as continuations in a term).

Proof-like terms give us a concept of logical validity – a truth value S is called *valid*, if there exists a proof-like term t with $t \Vdash S$.

With this notion, we are ready to define the centerpiece of the realizability model, which is the *entailment relation* on predicates.

► **Definition 16.** For any set I and integer n , the $(n + 1)$ -ary entailment relation (\vdash_I) on predicates on I is defined by

$$\varphi_1 \dots \varphi_n \vdash_I \psi \quad \text{if and only if} \quad \exists t \in \text{PL} . t \Vdash \forall_I(\varphi_1 \Rightarrow \dots \Rightarrow \varphi_n \Rightarrow \psi).$$

If the right hand side proposition holds, we call t a *realizer* of $\varphi_1 \dots \varphi_n \vdash_I \psi$.

Thus, $\varphi_1 \dots \varphi_n \vdash_I \psi$ means that the truth value $\forall_I(\varphi_1 \Rightarrow \dots \Rightarrow \varphi_n \Rightarrow \psi)$ is valid. More explicitly this can be written out as

$$\exists t \in \text{PL} \forall i \in I, u_1 \in \varphi_1(i)^\perp, \dots, u_n \in \varphi_n(i)^\perp, \pi \in \psi(i) . t \star u_1 \dots u_n \cdot \pi \in \perp.$$

With the aim to show that the semantic predicates form a tripos in Theorem 22, we now prove that the entailment ordering models the logical rules in Table (1). The first eight rules form a standard natural deduction system for (the \perp, \Rightarrow fragment of) classical propositional logic, but for universal quantification we give categorically inspired rules that bring us quicker to where we want, and in particular avoid having to deal with variables.

► **Lemma 17.** *The rules displayed in Table 1 are admissible for the entailment relation, in the sense that if the hypotheses hold then so does the conclusion.*

⁵ The usual $\forall x : A$ from predicate logic corresponds to taking f to be a projection map $\pi_1 : \Gamma \times A \rightarrow \Gamma$, see e.g. [8, Chapter 4].

■ **Table 1** Admissible rules for the entailment relation.

$\frac{}{\varphi \vdash_I \varphi} \text{ (Ax)}$	$\frac{\Gamma \vdash_I \perp}{\Gamma \vdash_I \psi} \text{ (\perp E)}$
$\frac{\Gamma, \varphi \vdash_I \psi}{\Gamma \vdash_I \varphi \Rightarrow \psi} \text{ (\Rightarrow I)}$	$\frac{\Gamma \vdash_I \psi \quad \Delta \vdash_I \psi \Rightarrow \theta}{\Gamma, \Delta \vdash_I \theta} \text{ (\Rightarrow E)}$
$\frac{\Gamma \vdash_I \psi}{\sigma(\Gamma) \vdash_I \psi} \text{ (S)}$	$\frac{}{\Delta \mid \Gamma \vdash_I ((\psi \Rightarrow \perp) \Rightarrow \psi) \Rightarrow \psi} \text{ (PeL)}$
$\frac{\Gamma \vdash_I \psi}{A, \Gamma \vdash_I \psi} \text{ (W)}$	$\frac{A, A, \Gamma \vdash_I \psi}{A, \Gamma \vdash_I \psi} \text{ (C)}$
$\frac{f^* \Gamma \vdash_J \xi}{\Gamma \vdash_I \forall_f \xi} \text{ (\forall I)}$	$\frac{\Gamma \vdash_I \forall_f \xi}{f^* \Gamma \vdash_J \xi} \text{ (\forall E)}$

φ, ψ, θ are predicates on I , i.e. functions $I \rightarrow P(\Pi)$, and $\Gamma \equiv \varphi_1 \dots \varphi_n$ and $\Delta \equiv \psi_1 \dots \psi_m$ are lists of such predicates. ξ is a predicate on J , and $f : J \rightarrow I$ is a function. σ is a permutation of $\{1, \dots, n\}$. $f^* \Gamma$ is an abbreviation for $f^* \varphi_1 \dots f^* \varphi_n$, and $\sigma(\Gamma)$ is an abbreviation for $\varphi_{\sigma(1)} \dots \varphi_{\sigma(n)}$.

Proof. (Ax) rule: The conclusion is realized by $\lambda x . x$.

(\perp E) rule: every realizer of the hypothesis is also a realizer of the conclusion, since $\psi(i) \subseteq \perp(i) = \Pi$ for all $i \in I$.

(\Rightarrow I) rule: the hypothesis and the conclusion have precisely the same realizers.

(\Rightarrow E) rule: if t realizes $\Delta \vdash_I \psi \Rightarrow \theta$ and u realizes $\Gamma \vdash_i \psi$ then $\Gamma, \Delta \vdash_i \theta$ is realized by $\lambda x_1 \dots x_n y_1 \dots y_m . t y_1 \dots y_m (u x_1 \dots x_n)$.

(PeL) rule ('Peirce's law'): the conclusion is realized by \mathbf{c} . To see this, let $i \in I$, $t \Vdash (\psi(i) \Rightarrow \perp) \Rightarrow \psi(i)$, and $\pi \in \psi(i)$. Then we have $\mathbf{c} \star t \cdot \pi \succ t \star \mathbf{k}_\pi \cdot \pi$, which is in \perp since $\mathbf{k}_\pi \cdot \pi \in (\psi(i) \Rightarrow \perp) \Rightarrow \psi(i)$ by Lemma 14 and the definition (6) of implication.

(W) rule: if t realizes $\Gamma \vdash_I \psi$, then $\lambda x . t$ realizes $A, \Gamma \vdash_I \psi$.

(C) rule: if t realizes $A, A, \Gamma \vdash_I \psi$, then $\lambda x . t x x$ realizes $A, \Gamma \vdash_I \psi$.

(S) rule: if t realizes $\Gamma \vdash_I \psi$, then $\lambda x_{\sigma(1)} \dots x_{\sigma(n)} . t x_1 \dots x_n$ realizes $\sigma(\Gamma) \vdash_I \psi$.

(\forall I) and (\forall E) rules: $\Gamma \vdash_I \forall_f \xi$ and $f^* \Gamma \vdash_J \xi$ have exactly the same realizers. Indeed, a realizer of $f^* \Gamma \vdash_J \xi$ is a term t satisfying

$$\forall j \in J, u_1 \in \varphi_1(f(j))^\perp, \dots, u_n \in \varphi_n(f(j))^\perp, \pi \in \xi(j) . t \star u_1 \cdot \dots \cdot u_n \cdot \pi \in \perp,$$

and a realizer of $\Gamma \vdash_I \forall_f \xi$ is a term t satisfying

$$\forall i \in I, u_1 \in \varphi_1(i)^\perp, \dots, u_n \in \varphi_n(i)^\perp, \pi \in \bigcup_{f(j)=i} \xi(j) . t \star u_1 \cdot \dots \cdot u_n \cdot \pi \in \perp,$$

and both statements can be rephrased as a quantification over pairs (i, j) with $f(j) = i$. ◀

We only defined the propositional connectives \perp, \Rightarrow , since \top, \wedge, \vee, \neg can be encoded as follows:

$$\begin{aligned} \top &\equiv \perp \Rightarrow \perp & \neg \varphi &\equiv \varphi \Rightarrow \perp \\ \varphi \wedge \psi &\equiv (\varphi \Rightarrow (\psi \Rightarrow \perp)) \Rightarrow \perp & \varphi \vee \psi &\equiv (\varphi \Rightarrow \perp) \Rightarrow \psi \end{aligned} \quad (7)$$

With these encodings it is routine to show the following.

► **Lemma 18.** *With the connectives \top, \wedge, \vee, \neg encoded as in (7), the rules of propositional classical natural deduction (e.g. system Nc in [19, Section 2.1.8]) are derivable from the rules in Table 1.*

With this we can show that for any set I , the binary part of the entailment relation makes $P(\Pi)^I$ into a *Boolean prealgebra*.

► **Definition 19.** A *Boolean prealgebra* is a preorder (B, \leq) which

1. has binary *joins* and *meets* – denoted by $x \vee y$ and $x \wedge y$ for $x, y \in B$,
2. has a *least element* \perp and a *greatest element* \top ,
3. is *distributive* in the sense that $x \wedge (y \vee z) \cong (x \wedge y) \vee (x \wedge z)$ for all $x, y, z \in B$, and
4. is *complemented*, i.e. for every $x \in B$ there exists a $\neg x$ with $x \wedge \neg x \cong \perp$ and $x \vee \neg x \cong \top$.

► **Lemma 20.** *Writing $\varphi \leq \psi$ for $\varphi \vdash_I \psi$, $(P(\Pi)^I, \leq)$ is a Boolean prealgebra.*

Proof. The (Ax) rule implies that \leq is reflexive, and transitivity follows from the derivation

$$\frac{\varphi \vdash_I \psi \quad \frac{\psi \vdash_I \theta}{\vdash_I \psi \Rightarrow \theta}}{\varphi \vdash_I \theta}$$

Thus, \leq is a preorder on $P(\Pi)^I$.

The joins, meets, complements, and least and greatest element are given by the corresponding logical operations as defined in (6) and (7).

The required properties all follow from derivability of corresponding entailments and rules in classical natural deduction – for example, $\varphi \wedge \psi$ is a binary meet of φ and ψ since

$$(*) \text{ the entailments } \varphi \wedge \psi \vdash_I \varphi \text{ and } \varphi \wedge \psi \vdash_I \psi \text{ and the rule } \frac{\theta \vdash_I \varphi \quad \theta \vdash_I \psi}{\theta \vdash_I \varphi \wedge \psi}$$

are derivable.

Distributivity follows from derivability of the entailments $\varphi \wedge (\psi \vee \theta) \vdash_I (\varphi \wedge \psi) \vee (\varphi \wedge \theta)$ and $(\varphi \wedge \psi) \vee (\varphi \wedge \theta) \vdash_I \varphi \wedge (\psi \vee \theta)$. ◀

We now come to *triposes*, which are a kind categorical model for higher order logic. We use a ‘strictified’ version of the original definition [7, Def. 1.2] since this bypasses some subtleties and is sufficient for our purposes. Furthermore, we are only interested modeling classical logic here, and thus can restrict attention to triposes whose fibers are Boolean (instead of Heyting) prealgebras.

► **Definition 21.** A *strict⁶ Boolean tripos* is a contravariant functor $\mathcal{P} : \mathbf{Set}^{\text{op}} \rightarrow \mathbf{Ord}$ from the category of sets to the category of preorders such that

- for every set I , the preorder $\mathcal{P}(I)$ is a Boolean prealgebra, and for any function $f : J \rightarrow I$, the induced monotone map $\mathcal{P}(f) : \mathcal{P}(I) \rightarrow \mathcal{P}(J)$ preserves all Boolean prealgebra structure.

⁶ ‘Strict’ refers to the facts that (i) \mathcal{P} is a *functor*, not merely a pseudofunctor (ii) the Boolean prealgebra structure is preserved ‘on the nose’ by the monotone maps $\mathcal{P}(f)$ (iii) the Beck-Chevalley condition is required up to equality, not merely isomorphism, (iv) we require equality and uniqueness in the last condition. Every strict tripos is a tripos in the usual sense, and conversely it can be shown that any tripos is equivalent to a strict one.

- for any $f : J \rightarrow I$, $\mathcal{P}(f)$ has left and right adjoints⁷ $\exists_f \dashv \mathcal{P}(f) \dashv \forall_f$ such that

$$\text{for any pullback square}^8 \quad \begin{array}{ccc} L & \xrightarrow{q} & K \\ p \downarrow & & \downarrow g \\ J & \xrightarrow{f} & I \end{array} \quad (8)$$

we have $\mathcal{P}(g) \circ \forall_f = \forall_q \circ \mathcal{P}(p)$ (this is the *Beck-Chevalley condition*), and

- there exists a *generic predicate*, i.e. a set \mathbf{Prop} and an element $\text{tr} \in \mathcal{P}(\mathbf{Prop})$ such that for every set I and $\varphi \in \mathcal{P}(I)$ there exists a unique function $f : I \rightarrow \mathbf{Prop}$ with $\mathcal{P}(f)(\text{tr}) = \varphi$.

The assignment $I \mapsto (P(\Pi)^I, \leq)$ extends to a functor $\mathcal{P}_{\perp} : \mathbf{Set}^{\text{op}} \rightarrow \mathbf{Ord}$ by letting $\mathcal{P}_{\perp}(f) = f^*$, i.e. mapping every function $f : J \rightarrow I$ to the reindexing function along f , which is monotone since every realizer of $\varphi \vdash_I \psi$ is also a realizer of $\varphi \circ f \vdash_J \psi \circ f$.

► **Theorem 22.** \mathcal{P}_{\perp} is a strict Boolean tripos.

Proof. We have shown in Lemma 20 that the preorders $(P(\Pi)^I, \leq)$ are Boolean prealgebras. It is immediate from (6) that the reindexing functions f^* preserve \perp and \Rightarrow , and the other Boolean operations are preserved since they are given by encodings.

The identity function $\text{id} : P(\Pi) \rightarrow P(\Pi)$ is a generic predicate for \mathcal{P}_{\perp} .

The $(\forall I)$ and $(\forall E)$ rules together imply that the operation $\forall_f : P(\Pi)^I \rightarrow P(\Pi)^J$ is right adjoint to f^* for any $f : J \rightarrow I$. Existential quantification along $f : J \rightarrow I$ is given by $\exists_f = \neg \circ \forall_f \circ \neg$, which is left adjoint to f^* since

$$\neg \forall_f \neg \varphi \vdash_I \psi \quad \text{iff} \quad \neg \psi \vdash_I \forall_f \neg \varphi \quad \text{iff} \quad f^* \neg \psi \vdash_J \neg \varphi \quad \text{iff} \quad \neg f^* \psi \vdash_J \neg \varphi \quad \text{iff} \quad \varphi \vdash_J f^* \psi$$

for all $\varphi : J \rightarrow P(\Pi)$ and $\psi : I \rightarrow P(\Pi)$.

It remains to verify the Beck-Chevalley condition. Given a square as in (8) we have

$$g^* \forall_f (\varphi(k)) = \bigcup \{ \varphi(j) \mid f(j) = g(k) \} \quad \text{and} \quad \forall_q (p^*(k)) = \bigcup \{ \varphi(j) \mid \exists l. pl = j \wedge ql = k \},$$

and the two terms are equal since the square is a pullback. ◀

Thus we obtain a tripos \mathcal{P}_{\perp} for each pole \perp . As Hyland, Johnstone, and Pitts showed in [7], every tripos \mathcal{P} gives rise to a topos $\mathbf{Set}[\mathcal{P}]$ via the *tripos-to-topos construction*. Since the fibers of the triposes \mathcal{P}_{\perp} are Boolean prealgebras, the toposes $\mathbf{Set}[\mathcal{P}_{\perp}]$ are Boolean as well, which means that their internal logic is classical.

5.1.1 Consistency

Triposes of the form \mathcal{P}_{\perp} can be degenerate in two ways: if \perp is empty then $\mathcal{P}_{\perp}(I) \simeq (P(I), \subseteq)$ for every set I , and the topos $\mathbf{Set}[\mathcal{P}_{\perp}]$ is equivalent to the category \mathbf{Set} .

If, in the other extreme, the pole is so big that there exists a proof-like t realizing \perp , i.e. falsity is valid in the model, then we have $\mathcal{P}_{\perp}(I) \simeq 1$ for all I (since t realizes every entailment $\varphi \vdash_I \psi$), and the topos $\mathbf{Set}[\mathcal{P}_{\perp}]$ is equivalent to the terminal category.

By *consistency* we mean that falsity is *not* valid, or equivalently that

$$\forall t \in \text{PL} \exists \pi \in \Pi. t \star \pi \notin \perp. \quad (9)$$

⁷ ‘Adjoint’ in the sense of ‘adjoint functor’, where monotone maps are viewed as functors between degenerate categories.

⁸ The square being a pullback means that $f \circ p = g \circ q$ and $\forall jk. f(j) = g(k) \Rightarrow \exists l. pl = j \wedge ql = k$.

The ‘canonical’ (according to Krivine [12]) non-trivial consistent pole is the *thread model*, which is given by postulating a stack constant π_t for each proof-like term t , and defining $\perp = \{p \in \Lambda \star \Pi \mid \neg \exists t \in \text{PL} . t \star \pi_t \rightsquigarrow^* p\}$. Then the processes $t \star \pi_t$ are not in \perp for any proof-like t , which ensures the validity of condition (9).

In the next section we show how the presence of side effects allows to define a variety of new, ‘meaningful’ consistent poles.

5.2 Krivine realizability with I/O

The developments of the previous section generalize pretty much directly to the syntax with I/O. Concretely, we carry over the definitions of *pole*, *truth value*, *realizer*, *predicate*, and of the basic logical operations $\perp, \Rightarrow, \forall$, by replacing Λ with Λ_e , Π with Π_e , and $\Lambda \star \Pi$ with P .

We point out that in presence of effects, Definition 13 only means that \perp is saturated w.r.t. *effect-free* evaluation, in contrast to Miquel’s approach [14] where a pole is a set of (what we call) execution contexts, closed under the entire execution relation.

The concept of proof-like term deserves some reexamination. It turns out that the appropriate concept of *proof-like term* is ‘term not containing any side effects’. This is consistent with Definition 15 if we read ‘free of side effects’ as ‘free of non-logical constructs’, which are the stack constants in Krivine’s case. Continuation terms, on the other hand, can be considered proof-like. We redefine therefore:

► **Definition 23.** The set $\text{PL} \subseteq \Lambda_e$ of *proof-like terms* is the set of terms not containing any of the constants $r, \mathbf{w0}, \mathbf{w1}, \text{end}$.

With this rephrased definition of proof-like term, we can define the entailment relation on the extended predicates in the same way:

► **Definition 24.** For any set I and integer n , the $(n + 1)$ -ary entailment relation (\vdash_I) on the set $P(\Pi_e)^I$ of extended predicates on I is defined by

$$\varphi_1 \dots \varphi_n \vdash_I \psi \quad \text{if and only if} \quad \exists t \in \text{PL} . t \Vdash \forall_I (\varphi_1 \Rightarrow \dots \Rightarrow \varphi_n \Rightarrow \psi).$$

As a special case, the ordering on extended predicates is defined by

$$\varphi \leq \psi \quad \text{if and only if} \quad \exists t \in \text{PL} . t \Vdash \forall_I (\varphi \Rightarrow \psi).$$

With these definitions, we can state analogues of Lemma 20 and Theorem 22:

► **Theorem 25.**

- For each set I , the order $(P(\Pi_e)^I, \leq)$ of extended predicates is a Boolean prealgebra.
- The assignment $I \mapsto (P(\Pi_e)^I, \leq)$ gives rise to a strict Boolean tripos $\mathcal{P}_\perp : \mathbf{Set}^{\text{op}} \rightarrow \mathbf{Ord}$.

Proof. This follows from the arguments in Section 5.1, since the proofs of Lemmas 14,17, 18,20, and of Theorem 22 are not obstructed in any way by the new constants, nor do they rely on stack constants. The redefinition of ‘proof-like term’ does not cause any problems either, since we never relied on proof-like terms not containing continuation terms. ◀

The above rephrasing of the definition of proof-like term admits an intuitive reformulation of the consistency criterion (9):

► **Lemma 26.** A pole \perp is consistent if and only if every $p \in \perp \setminus \{\top\}$ contains a non-logical constant, i.e. one of $r, \mathbf{w0}, \mathbf{w1}, \text{end}$.

Proof. If every element of $p \in \perp \setminus \{\top\}$ contains a non-logical constant, then $t \star \varepsilon$ is not in \perp for any proof-like t , which implies (9).

On the other hand, if $t \star \pi \in \perp$ does not contain any non-logical constant then $k_\pi t$ is a proof-like term which realizes \perp , since for any $\rho \in \Pi_e$ we have $k_\pi t \star \rho \succ k_\pi t \star \rho \succ t \star \pi \in \perp$. ◀

5.2.1 Poles from specifications

The connection between poles and specifications is established by the following lemma.

► **Lemma 27.** *Every set \perp of processes that is closed under weak bisimilarity is a pole.*

Proof. This is because $p \approx q$ whenever $p \succ q$, which follows from Lemma 5. ◀

Since we can assume that for any reasonable specification the processes implementing it are closed under weak bisimilarity, we can thus conclude that for any specification, the set of processes implementing it is a pole. For example:

- \perp_{cp} is the set of processes that read the input, copy every bit immediately to the output, and terminate when the input is empty. We have $Y \star (\lambda x. r(\mathbf{w}0x)(\mathbf{w}1x)\mathbf{end}) \in \perp_{cp}$.
- $\perp_{cp'}$ contains the processes that first read the entire input, and then write out the same string and terminate. We have $R \star F \cdot W \cdot \bar{0} \in \perp_{cp'}$ with the notations of Section 4.
- For any partial function $f : \mathbb{N} \rightarrow \mathbb{N}$, the pole \perp_f consists of those processes that implement f in the sense of Definition 1.
- Since poles are closed under unions, we can define the pole $\perp_F = \bigcup_{f \in F} \perp_f$ for any set $F \subseteq (\mathbb{N} \rightarrow \mathbb{N})$ of partial functions.

5.2.2 Toposes from computable functions

We are particularly interested in the poles \perp_f associated to computable functions f , and we want to use the associated triposes $\mathcal{P}_f = \mathcal{P}_{\perp_f}$ and toposes $\mathbf{Set}[\mathcal{P}_f]$ to study these functions.

The following theorem provides a first ‘sanity check’, in showing that the associated models are non-degenerate.

► **Theorem 28.** *Let $f : \mathbb{N} \rightarrow \mathbb{N}$.*

- \perp_f is consistent if and only if f is not totally undefined.
- \perp_f is non-empty if and only if f is computable.

Proof. The first claim follows from Lemma 26. If $n \in \text{dom}(f)$ and $t \star \pi$ implements f , then $(t \star \pi, \text{bin}(n), \varepsilon)$ must terminate and thus $t \star \pi$ must contain an **end** instruction. The totally undefined function, on the other hand, is by definition implemented by every process.

For the second claim, we have shown in Theorem 12 that every computable f is implemented by some process. Conversely, every implementable function is computable since the Krivine machine with I/O is an effective model of computation. ◀

5.3 Discussion and future work

The structure and properties of the toposes $\mathbf{Set}[\mathcal{P}_f]$ remain mysterious for the moment, and in future work we want to explore which kind of properties of f are reflected in $\mathbf{Set}[\mathcal{P}_f]$. In the spirit of Grothendieck [4] we want to view the toposes $\mathbf{Set}[\mathcal{P}_f]$ as *geometric* rather than logical objects, the guiding intuition being that $\mathbf{Set}[\mathcal{P}_f]$ can be seen as representation of ‘*the space of solutions to the algorithmic problem of computing f* ’, encoding e.g. information on how algorithms computing f can be decomposed into simpler parts.

Evident problems to investigate are to understand the lattice of truth values in $\mathbf{Set}[\mathcal{P}_f]$, and to determine for which pairs f, g of functions the associated toposes are equivalent, and which functions can be separated.

A more audacious goal is to explore whether $\mathbf{Set}[\mathcal{P}_f]$ can teach us something about the complexity of a computable function f . The Krivine machine with I/O seems to be a model of computation that is fine grained enough to recognize and differentiate time complexity

of different implementations of f , but it remains to be seen in how far this information is reflected in the ‘geometry’ of $\mathbf{Set}[\mathcal{P}_f]$.

Acknowledgements. Thanks to Jakob Grue Simonsen and Thomas Streicher for many discussions.

References

- 1 H.P. Barendregt. *The lambda calculus, Its syntax and semantics*, volume 103 of *Studies in Logic and the Foundations of Mathematics*. North-Holland Publishing Co., Amsterdam, revised edition, 1984.
- 2 W. Ferrer, J. Frey, M. Guillermo, O. Malherbe, and A. Miquel. Ordered combinatory algebras and realizability. *arXiv preprint arXiv:1410.5034*, 2014.
- 3 H. Friedman. Classically and intuitionistically provably recursive functions. In *Higher set theory (Proc. Conf., Math. Forschungsinst., Oberwolfach, 1977)*, volume 669 of *Lecture Notes in Math.*, pages 21–27. Springer, Berlin, 1978.
- 4 A. Grothendieck, M. Artin, and J.L. Verdier. Théorie des topos et cohomologie étale des schémas. *Lecture Notes in Mathematics*, 269, 1972.
- 5 J. Roger Hindley and Jonathan P. Seldin. *Lambda-calculus and combinators, an introduction*. Cambridge University Press, Cambridge, 2008.
- 6 J.M.E. Hyland. The effective topos. In *The L.E.J. Brouwer Centenary Symposium (Noordwijkerhout, 1981)*, volume 110 of *Stud. Logic Foundations Math.*, pages 165–216. North-Holland, Amsterdam, 1982.
- 7 J.M.E. Hyland, P.T. Johnstone, and A.M. Pitts. Tripos theory. *Math. Proc. Cambridge Philos. Soc.*, 88(2):205–231, 1980.
- 8 B. Jacobs. *Categorical logic and type theory*. Elsevier Science Ltd, 2001.
- 9 J.L. Krivine. Lambda-calcul, évaluation paresseuse et mise en mémoire. *RAIRO Inform. Théor. Appl.*, 25(1):67–84, 1991.
- 10 J.L. Krivine. Realizability in classical logic. *Panoramas et synthèses*, 27:197–229, 2009.
- 11 J.L. Krivine. Realizability algebras: a program to well order \mathbb{R} . *Log. Methods Comput. Sci.*, 7(3):3:02, 47, 2011.
- 12 J.L. Krivine. Realizability algebras II: New models of ZF + DC. *Log. Methods Comput. Sci.*, 8(1):1:10, 28, 2012.
- 13 R. Milner. Operational and algebraic semantics of concurrent processes. In *Handbook of theoretical computer science, Vol. B*, pages 1201–1242. Elsevier, Amsterdam, 1990.
- 14 A. Miquel. Classical modal realizability and side effects. *preprint*, 2009.
- 15 A. Miquel. Existential witness extraction in classical realizability and via a negative translation. *Log. Methods Comput. Sci.*, 7(2):2:2, 47, 2011.
- 16 A. Miquel. Forcing as a program transformation. In *26th Annual IEEE Symposium on Logic in Computer Science—LICS 2011*, pages 197–206. IEEE Computer Soc., Los Alamitos, CA, 2011.
- 17 W.P. Stekelenburg. *Realizability Categories*. PhD thesis, Utrecht University, 2013.
- 18 T. Streicher. Krivine’s classical realisability from a categorical perspective. *Mathematical Structures in Computer Science*, 23(06):1234–1256, 2013.
- 19 A.S. Troelstra and H. Schwichtenberg. *Basic proof theory*, volume 43 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, Cambridge, 1996.
- 20 J. van Oosten. *Classical Realizability*. Invited talk at “Cambridge Category Theory Seminar”, slides at <http://www.staff.science.uu.nl/~ooste110/talks/cambr060312.pdf>.
- 21 J. van Oosten. *Realizability: An Introduction to its Categorical Side*. Elsevier Science Ltd, 2008.

Standardization of a Call-By-Value Lambda-Calculus*

Giulio Guerrieri¹, Luca Paolini², and Simona Ronchi Della Rocca²

- 1** Laboratoire PPS, UMR 7126, Université Paris Diderot, Sorbonne Paris Cité
75205 Paris, France
giulio.guerrieri@pps.univ-paris-diderot.fr
- 2** Dipartimento di Informatica, Università degli Studi di Torino
Corso Svizzera 185, Torino, Italy
{paolini,ronchi}@di.unito.it

Abstract

We study an extension of Plotkin’s call-by-value lambda-calculus by means of two commutation rules (sigma-reductions). Recently, it has been proved that this extended calculus provides elegant characterizations of many semantic properties, as for example solvability. We prove a standardization theorem for this calculus by generalizing Takahashi’s approach of parallel reductions. The standardization property allows us to prove that our calculus is conservative with respect to the Plotkin’s one. In particular, we show that the notion of solvability for this calculus coincides with that for Plotkin’s call-by-value lambda-calculus.

1998 ACM Subject Classification D.3.1 Formal Definitions and Theory, F.3.2 Semantics of Programming Language, F.4.1 Mathematical Logic

Keywords and phrases standardization, sequentialization, lambda-calculus, sigma-reduction, parallel reduction, call-by-value, head reduction, internal reduction, solvability, potential valuability, observational equivalence

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.211

1 Introduction

The λ_v -calculus (λ_v for short) has been introduced by Plotkin in [15], in order to give a formal account of the call-by-value evaluation, which is the most commonly used parameter passing policy for programming languages. λ_v shares the syntax with the classical, call-by-name, λ -calculus (λ for short), but its reduction rule, β_v , is a restriction of β , firing only in case the argument is a value (i.e., a variable or an abstraction). While β_v is enough for evaluation, it turned out to be too weak to study operational properties of terms. For example, in λ , the β -reduction is sufficient to characterize solvability and (using extensionality) separability, but, in order to characterize similar properties for λ_v , it has been necessary to introduce different notions of reduction unsuitable for a correct call-by-value evaluation (see [13, 14]): this is disappointing and requires complex reasoning. In this paper we study λ_v^σ , the extension of λ_v proposed in [3]. It keeps the λ_v (and λ) syntax and it adds to the β_v -reduction two commutation rules, called σ_1 and σ_3 , which unblock β_v -redexes that are hidden by the “hyper-sequential structure” of terms. It is well-known (see [14, 1]) that in λ_v there are normal forms that are unsolvable, e.g. $(\lambda yx.xx)(zz)(\lambda x.xx)$. The more evident benefit of λ_v^σ

* This work was partially supported by LINTEL TO_Call1_2012_0085, i.e. a Research Project funded by the “Compagnia di San Paolo”.



is that the commutation rules make all normal forms solvable (indeed $(\lambda yx.xx)(zz)(\lambda x.xx)$ is not a λ_v^σ normal form). More generally, the so obtained language, allows us to characterize operational properties, like solvability and potential valuability, in an internal and elegant way (see [3]). In this paper we prove a standardization property in λ_v^σ , and some of its consequences, namely its soundness with respect to the semantics of λ_v .

Let us recall the notion of standardization, which has been first studied in the ordinary λ -calculus (see, for example [5, 8, 2]). A reduction sequence is standard if its redexes are ordered in a given way, and the corresponding standardization theorem establishes that every reduction sequence can constructively be transformed into a standard one. Standardization is a key tool to grasp the way in which reductions work, that sheds some light on redexes relationships and their dependencies. It is useful for characterization of semantic properties through reduction strategies (the proof of operational semantics adequacy is a typical use).

In the λ_v setting standardization theorems have been proved by Plotkin [15], Paolini and Ronchi Della Rocca [14, 12] and Crary [4]. The definition of standard sequence of reductions considered by Plotkin and Crary coincides, and it imposes a partial order on redexes, while Paolini and Ronchi Della Rocca define a total order on them. All these proofs are developed by using the notion of parallel reduction introduced by Tait and Martin-Löf (see Takahashi [17] for details and interesting technical improvements). We emphasize that this method does not involve the notion of residual of a redex, on which many classical proofs for the λ -calculus are based (see for example [8, 2]). As in [15, 17, 14, 4], we use a suitable notion of parallel reduction for developing our standardization theorem for λ_v^σ . In particular we consider two groups of redexes, head β_v -redexes and head σ -redexes (putting together σ_1 and σ_3), and we induce a total order on head redexes of the two groups, without imposing any order on head σ -redexes themselves. More precisely, when σ -redexes are missing, this notion of standardization coincides with that presented in [14]. Moreover, we show that it is not possible to strengthen our standardization by (locally) ordering σ_1 -reduction to σ_3 -reduction (or viceversa).

As usual, our standardization proof is based on a sequentialization result: inner reductions can always be postponed after the head ones, for a non-standard definition of head reduction. Sequentialization has interesting consequences: it allows us to prove that fundamental operational properties in λ_v^σ , like observational equivalence, potential valuability and solvability, are conservative with respect to the corresponding notions of λ_v . This fully justifies the project in [3] where λ_v^σ has been introduced as a tool for studying the operational behaviour of λ_v .

Other variants of λ_v have been introduced in the literature for modeling the call-by-value computation. We would like to cite here at least the contributions of Moggi [10], Felleisen and Sabry [16], Maraist et al. [9], Herbelin and Zimmerman [7], Accattoli and Paolini [1]. All these proposals are based on the introduction of new constructs to the syntax of λ_v , so the comparison between them is not easy with respect to syntactical properties (some detailed comparison is given in [1]). We point out that the calculi introduced in [10, 16, 9, 7] present some variants of our σ_1 and/or σ_3 rules, often in a setting with explicit substitutions.

Outline. In Section 2 we introduce the language λ_v^σ and its operational behaviour; in Section 3 the sequentialization property is proved; Section 4 contains the main result, i.e., standardization ; in Section 5 some conservativity results with respect to Plotkin's λ_v -calculus are proved. Section 6 concludes the paper, with some hints for future work.

2 The call-by-value lambda calculus with sigma-rules

In this section we present λ_v^σ , a call-by-value λ -calculus introduced in [3] that adds two σ -reduction rules to pure (i.e. without constants) call-by-value λ -calculus defined by Plotkin in [15].

The syntax of terms of λ_v^σ [3] is the same as the one of ordinary λ -calculus and Plotkin's call-by-value λ -calculus λ_v [15] (without constants). Given a countable set \mathcal{V} of *variables* (denoted by x, y, z, \dots), the sets Λ of *terms* and Λ_v of *values* are defined by mutual induction:

$$\begin{array}{lll} (\Lambda_v) & V, U ::= x \mid \lambda x.M & \text{values} \\ (\Lambda) & M, N, L ::= V \mid MN & \text{terms} \end{array}$$

Clearly, $\Lambda_v \subseteq \Lambda$. All terms are considered up to α -conversion. The set of free variables of a term M is denoted by $\text{fv}(M)$. Given $V_1, \dots, V_n \in \Lambda_v$ and pairwise distinct variables x_1, \dots, x_n , $M\{V_1/x_1, \dots, V_n/x_n\}$ denotes the term obtained by the *capture-avoiding simultaneous substitution* of V_i for each free occurrence of x_i in the term M (for all $1 \leq i \leq n$). Note that, for all $V, V_1, \dots, V_n \in \Lambda_v$ and pairwise distinct variables x_1, \dots, x_n , $V\{V_1/x_1, \dots, V_n/x_n\} \in \Lambda_v$.

Contexts (with exactly one hole (\cdot)), denoted by \mathbf{C} , are defined as usual via the grammar:

$$\mathbf{C} ::= (\cdot) \mid \lambda x.\mathbf{C} \mid \mathbf{C}M \mid M\mathbf{C}.$$

We use $\mathbf{C}(M)$ for the term obtained by the capture-allowing substitution of the term M for the hole (\cdot) in the context \mathbf{C} .

► **Notation.** From now on, we set $I = \lambda x.x$ and $\Delta = \lambda x.xx$.

The reduction rules of λ_v^σ consist of Plotkin's β_v -reduction rule, introduced in [15], and two simple commutation rules called σ_1 and σ_3 , studied in [3].

► **Definition 1** (Reduction rules). We define the following binary relations on Λ (for any $M, N, L \in \Lambda$ and any $V \in \Lambda_v$):

$$\begin{array}{l} (\lambda x.M)V \mapsto_{\beta_v} M\{V/x\} \\ (\lambda x.M)NL \mapsto_{\sigma_1} (\lambda x.ML)N \quad \text{with } x \notin \text{fv}(L) \\ V((\lambda x.L)N) \mapsto_{\sigma_3} (\lambda x.VL)N \quad \text{with } x \notin \text{fv}(V). \end{array}$$

For any $r \in \{\beta_v, \sigma_1, \sigma_3\}$, if $M \mapsto_r M'$ then M is a *r-redex* and M' is its *r-contractum*. In this sense, a term of the shape $(\lambda x.M)N$ (for any $M, N \in \Lambda$) is a *β -redex*.

We set $\mapsto_\sigma = \mapsto_{\sigma_1} \cup \mapsto_{\sigma_3}$ and $\mapsto_v = \mapsto_{\beta_v} \cup \mapsto_\sigma$.

The side conditions on \mapsto_{σ_1} and \mapsto_{σ_3} in Definition 1 can be always fulfilled by α -renaming.

Obviously, any β_v -redex is a β -redex but the converse does not hold: $(\lambda x.z)(yI)$ is a β -redex but not a β_v -redex. Redexes of different kind may overlap: for example, the term $\Delta I \Delta$ is a σ_1 -redex and it contains the β_v -redex ΔI ; the term $\Delta(I\Delta)(xI)$ is a σ_1 -redex and it contains the σ_3 -redex $\Delta(I\Delta)$, which contains in turn the β_v -redex $I\Delta$.

According to the Girard's call-by-value "boring" translation $(\cdot)^v$ of terms into Intuitionistic Multiplicative Exponential Linear Logic proof-nets, defined by $(A \Rightarrow B)^v = !A^v \multimap !B^v$ (see [6]), the images under $(\cdot)^v$ of a σ_1 -redex (resp. σ_3 -redex) and its contractum are equal modulo some "bureaucratic" steps of cut-elimination.

► **Notation.** Let R be a binary relation on Λ . We denote by R^* (resp. R^+ ; R^\equiv) the reflexive-transitive (resp. transitive; reflexive) closure of R .

► **Definition 2** (Reductions). Let $r \in \{\beta_v, \sigma_1, \sigma_3, \sigma, \nu\}$.

The r -reduction \rightarrow_r is the contextual closure of \mapsto_r , i.e. $M \rightarrow_r M'$ iff there is a context \mathcal{C} and $N, N' \in \Lambda$ such that $M = \mathcal{C}(N)$, $M' = \mathcal{C}(N')$ and $N \mapsto_r N'$.

The r -equivalence $=_r$ is the reflexive-transitive and symmetric closure of \rightarrow_r .

Let M be a term: M is r -normal if there is no term N such that $M \rightarrow_r N$; M is r -normalizable if there is a r -normal term N such that $M \rightarrow_r^* N$; M is strongly r -normalizing if there is no sequence $(N_i)_{i \in \mathbb{N}}$ such that $M = N_0$ and $N_i \rightarrow_r N_{i+1}$ for any $i \in \mathbb{N}$. Finally, \rightarrow_r is strongly normalizing if every $N \in \Lambda$ is strongly r -normalizing.

Patently, $\rightarrow_\sigma \subsetneq \rightarrow_\nu$ and $\rightarrow_{\beta_v} \subsetneq \rightarrow_\nu$.

► **Remark 3.** For any $r \in \{\beta_v, \sigma_1, \sigma_3, \sigma, \nu\}$ (resp. $r \in \{\sigma_1, \sigma_3, \sigma\}$), values are closed under r -reduction (resp. r -expansion): for any $V \in \Lambda_v$, if $V \rightarrow_r M$ (resp. $M \rightarrow_r V$) then $M \in \Lambda_v$; more precisely, $V = \lambda x.N$ and $M = \lambda x.N'$ for some $N, N' \in \Lambda$ with $N \rightarrow_r N'$ (resp. $N' \rightarrow_r N$).

► **Proposition 4** (See [3]). *The σ -reduction is confluent and strongly normalizing. The ν -reduction is confluent.*

The λ_v^σ -calculus, λ_v^σ for short, is the set Λ of terms endowed with the ν -reduction \rightarrow_ν . The set Λ endowed with the β_v -reduction \rightarrow_{β_v} is the λ_v -calculus (λ_v for short), i.e. the Plotkin's call-by-value λ -calculus [15] (without constants), which is thus a sub-calculus of λ_v^σ .

► **Example 5.** $M = (\lambda y.\Delta)(xI)\Delta \rightarrow_{\sigma_1} (\lambda y.\Delta\Delta)(xI) \rightarrow_{\beta_v} (\lambda y.\Delta\Delta)(xI) \rightarrow_{\beta_v} \dots$ and $N = \Delta((\lambda y.\Delta)(xI)) \rightarrow_{\sigma_3} (\lambda y.\Delta\Delta)(xI) \rightarrow_{\beta_v} (\lambda y.\Delta\Delta)(xI) \rightarrow_{\beta_v} \dots$ are the only possible ν -reduction paths from M and N respectively: M and N are not ν -normalizable, and $M =_\nu N$. Meanwhile, M and N are β_v -normal and different, hence $M \neq_{\beta_v} N$ (by confluence of \rightarrow_{β_v} , see [15]).

Informally, σ -rules unblock β_v -redexes which are hidden by the “hyper-sequential structure” of terms. This approach is alternative to the one in [1] where hidden β_v -redexes are reduced using rules acting at a distance (through explicit substitutions). It can be shown that the call-by-value λ -calculus with explicit substitution introduced in [1] can be embedded in λ_v^σ .

3 Sequentialization

In this section we aim to prove a sequentialization theorem (Theorem 22) for the λ_v^σ -calculus by adapting Takahashi's method [17, 4] based on parallel reductions.

► **Notation.** From now on, we always assume that $V, V' \in \Lambda_v$.

Note that the generic form of a term is $VM_1 \dots M_m$ for some $m \in \mathbb{N}$ (in particular, values are obtained when $m = 0$). The sequentialization result is based on a partitioning of ν -reduction between head and internal reduction.

► **Definition 6** (Head β_v -reduction). We define inductively the head β_v -reduction $\xrightarrow{h}_{\beta_v}$ by the following rules ($m \in \mathbb{N}$ in both rules):

$$\frac{}{(\lambda x.M)VM_1 \dots M_m \xrightarrow{h}_{\beta_v} M\{V/x\}M_1 \dots M_m} \beta_v \quad \frac{N \xrightarrow{h}_{\beta_v} N'}{VNM_1 \dots M_m \xrightarrow{h}_{\beta_v} VN'M_1 \dots M_m} \text{right}$$

The head β_v -reduction $\xrightarrow{h}_{\beta_v}$ reduces exactly the same redexes (see also [13]) as the “left reduction” defined in [15, p. 136] for λ_v and called “evaluation” in [16, 4]. If $N \xrightarrow{h}_{\beta_v} N'$ then N' is obtained from N by reducing the leftmost-outermost β_v -redex, not in the scope of a λ : thus, the head β_v -reduction is deterministic (i.e., it is a partial function from Λ to Λ) and does not reduce values.

► **Definition 7** (Head σ -reduction). We define inductively the *head σ -reduction* \xrightarrow{h}_σ by the following rules ($m \in \mathbb{N}$ in all the rules, $x \notin \text{fv}(L)$ in the rule σ_1 , $x \notin \text{fv}(V)$ in the rule σ_3):

$$\frac{}{(\lambda x.M)NLM_1 \dots M_m \xrightarrow{h}_\sigma (\lambda x.ML)NLM_1 \dots M_m} \sigma_1 \quad \frac{N \xrightarrow{h}_\sigma N'}{VNM_1 \dots M_m \xrightarrow{h}_\sigma VN'M_1 \dots M_m} \text{right}$$

$$\frac{}{V((\lambda x.L)N)M_1 \dots M_m \xrightarrow{h}_\sigma (\lambda x.VL)NM_1 \dots M_m} \sigma_3$$

The *head (v)-reduction* is $\xrightarrow{h}_v = \xrightarrow{h}_{\beta_v} \cup \xrightarrow{h}_\sigma$. The *internal (v)-reduction* is $\xrightarrow{int}_v = \rightarrow_v \setminus \xrightarrow{h}_v$.

Notice that $\mapsto_{\beta_v} \subsetneq \xrightarrow{h}_{\beta_v} \subsetneq \rightarrow_{\beta_v}$ and $\mapsto_\sigma \subsetneq \xrightarrow{h}_\sigma \subsetneq \rightarrow_\sigma$ and $\mapsto_v \subsetneq \xrightarrow{h}_v \subsetneq \rightarrow_v$. Values are normal forms for the head reduction, but the converse does not hold: $xI \notin \Lambda_v$ is head-normal.

Informally, if $N \xrightarrow{h}_\sigma N'$ then N' is obtained from N by reducing “one of the leftmost” σ_1 - or σ_3 -redexes, not in the scope of a λ : in general, a term may contain several head σ_1 - and σ_3 -redexes. Indeed, differently from $\xrightarrow{h}_{\beta_v}$, the head σ -reduction \xrightarrow{h}_σ is not deterministic, for example the leftmost-outermost σ_1 - and σ_3 -redexes may overlap: if $M = (\lambda y.y')(\Delta(xI))I$ then $M \xrightarrow{h}_\sigma (\lambda y.y'I)(\Delta(xI)) = N_1$ by applying the rule σ_1 and $M \xrightarrow{h}_\sigma (\lambda z.(\lambda y.y')(zz))(xI)I = N_2$ by applying the rule σ_3 . Note that N_1 contains only a head σ_3 -redex and $N_1 \xrightarrow{h}_\sigma (\lambda z.(\lambda y.y'I)(zz))(xI) = N$ which is normal for \xrightarrow{h}_v ; meanwhile N_2 contains only a head σ_1 -redex and $N_2 \xrightarrow{h}_\sigma (\lambda z.(\lambda y.y')(zz)I)(xI) = N'$ which is normal for \xrightarrow{h}_v : $N \neq N'$, hence the head reduction \xrightarrow{h}_v is not confluent and a term may have several head-normal forms (this example does not contradict the confluence of σ -reduction because $N' \rightarrow_\sigma N$ but by performing an internal reduction step). Later, in Corollary 26.2 we show that if a term M has a head normal form $N \in \Lambda_v$ then N is the unique head normal form of M .

► **Definition 8** (Parallel reduction). We define inductively the *parallel reduction* \Rightarrow by the following rules ($x \notin \text{fv}(L)$ in the rule σ_1 , $x \notin \text{fv}(V)$ in the rule σ_3):

$$\frac{V \Rightarrow V' \quad M_i \Rightarrow M'_i \quad (m \in \mathbb{N}, 0 \leq i \leq m)}{(\lambda x.M_0)VM_1 \dots M_m \Rightarrow M'_0\{V'/x\}M'_1 \dots M'_m} \beta_v \quad \frac{N \Rightarrow N' \quad L \Rightarrow L' \quad M_i \Rightarrow M'_i \quad (m \in \mathbb{N}, 0 \leq i \leq m)}{(\lambda x.M_0)NLM_1 \dots M_m \Rightarrow (\lambda x.M'_0L')N'M'_1 \dots M'_m} \sigma_1$$

$$\frac{V \Rightarrow V' \quad N \Rightarrow N' \quad L \Rightarrow L' \quad M_i \Rightarrow M'_i \quad (m \in \mathbb{N}, 1 \leq i \leq m)}{V((\lambda x.L)N)M_1 \dots M_m \Rightarrow (\lambda x.V'L')N'M'_1 \dots M'_m} \sigma_3$$

$$\frac{M_i \Rightarrow M'_i \quad (m \in \mathbb{N}, 0 \leq i \leq m)}{(\lambda x.M_0)M_1 \dots M_m \Rightarrow (\lambda x.M'_0)M'_1 \dots M'_m} \lambda \quad \frac{M_i \Rightarrow M'_i \quad (m \in \mathbb{N}, 1 \leq i \leq m)}{xM_1 \dots M_m \Rightarrow xM'_1 \dots M'_m} \text{var}$$

In Definition 8 the rule *var* has no premises when $m = 0$: this is the base case of the inductive definition of \Rightarrow . The rules σ_1 and σ_3 have exactly three premises when $m = 0$.

Intuitively, $M \Rightarrow M'$ means that M' is obtained from M by reducing a number of β_v -, σ_1 - and σ_3 -redexes (existing in M) simultaneously.

► **Definition 9** (Internal and strong parallel reduction). We define inductively the *internal parallel reduction* \xRightarrow{int} by the following rules:

$$\frac{N \Rightarrow N'}{\lambda x.N \xRightarrow{int} \lambda x.N'} \lambda \quad \frac{}{x \xRightarrow{int} x} \text{var} \quad \frac{V \Rightarrow V' \quad N \xRightarrow{int} N' \quad M_i \Rightarrow M'_i \quad (m \in \mathbb{N}, 1 \leq i \leq m)}{VNM_1 \dots M_m \xRightarrow{int} V'N'M'_1 \dots M'_m} \text{right}$$

The *strong parallel reduction* \Rightarrow is defined by: $M \Rightarrow N$ iff $M \Rightarrow N$ and there exist $M', M'' \in \Lambda$ such that $M \xrightarrow{h}_{\beta_v} M' \xrightarrow{h}_\sigma M'' \xRightarrow{int} N$.

Notice that the rule *right* for \xRightarrow{int} has exactly two premises when $m = 0$.

► **Remark 10.** The relations \Rightarrow , \Rightarrow and $\xRightarrow{\text{int}}$ are reflexive. The reflexivity of \Rightarrow follows immediately from the reflexivity of \Rightarrow and $\xRightarrow{\text{int}}$. The proofs of reflexivity of \Rightarrow and $\xRightarrow{\text{int}}$ are both by structural induction on a term: in the case of \Rightarrow , recall that every term is of the form $(\lambda x.N)M_1 \dots M_m$ or $xM_1 \dots M_m$ for some $m \in \mathbb{N}$ and then apply the rule λ or var respectively, together with the inductive hypothesis; in the case of $\xRightarrow{\text{int}}$, recall that every term is of the form $\lambda x.M$ or x or $VNM_1 \dots M_m$ for some $m \in \mathbb{N}$ and then apply the rule λ (together with the reflexivity of \Rightarrow) or var or right (together with the reflexivity of \Rightarrow and the inductive hypothesis) respectively.

One has $\xRightarrow{\text{int}} \subsetneq \Rightarrow \subseteq \xRightarrow{\text{int}}$ (first, prove that $\xRightarrow{\text{int}} \subseteq \Rightarrow$ by induction on the derivation of $M \xRightarrow{\text{int}} M'$, the other inclusions follow from the definition of \Rightarrow) and, since \Rightarrow is reflexive (Remark 10), $\xrightarrow{\beta_v} \subsetneq \Rightarrow$ and $\xrightarrow{\sigma} \subsetneq \Rightarrow$. Observe that $\Delta\Delta \text{ R } \Delta\Delta$ for any $\text{R} \in \{\xrightarrow{\beta_v}, \xrightarrow{\sigma}, \Rightarrow, \xRightarrow{\text{int}}, \Rightarrow\}$, even if for different reasons: for example, $\Delta\Delta \xRightarrow{\text{int}} \Delta\Delta$ by reflexivity of $\xRightarrow{\text{int}}$ (Remark 10), whereas $\Delta\Delta \xrightarrow{\beta_v} \Delta\Delta$ by reducing the (leftmost-outermost) β_v -redex.

Next two further remarks collect many minor properties that can be easily proved.

- **Remark 11. 1.** The head β_v -reduction $\xrightarrow{\beta_v}$ does not reduce a value (in particular, does not reduce under λ 's), i.e., for any $M \in \Lambda$ and any $V \in \Lambda_v$, one has $V \not\xrightarrow{\beta_v} M$.
2. The head σ -reduction $\xrightarrow{\sigma}$ does neither reduce a value nor reduce to a value, i.e., for any $M \in \Lambda$ and any $V \in \Lambda_v$, one has $V \not\xrightarrow{\sigma} M$ and $M \not\xrightarrow{\sigma} V$.
 3. Variables and abstractions are preserved under $\xRightarrow{\text{int}}$ ($\xRightarrow{\text{int}}$ -expansion), i.e., if $M \xRightarrow{\text{int}} x$ (resp. $M \xRightarrow{\text{int}} \lambda x.N'$) then $M = x$ (resp. $M = \lambda x.N$ for some $N \in \Lambda$ such that $N \Rightarrow N'$).
 4. If $M \Rightarrow M'$ then $\lambda x.M \text{ R } \lambda x.M'$ for any $\text{R} \in \{\Rightarrow, \xRightarrow{\text{int}}, \Rightarrow\}$. Indeed, for $\text{R} \in \{\Rightarrow, \xRightarrow{\text{int}}\}$ apply the rule λ to conclude, then $\lambda x.M \Rightarrow \lambda x.M'$ according to the definition of \Rightarrow .
 5. For any $V, V' \in \Lambda_v$, $V \xRightarrow{\text{int}} V'$ iff $V \Rightarrow V'$. The left-to-right direction holds because $\xRightarrow{\text{int}} \subseteq \Rightarrow$; conversely, assume $V \Rightarrow V'$: if V is a variable then necessarily $V = V'$ and hence $V \xRightarrow{\text{int}} V'$ by applying the rule var for $\xRightarrow{\text{int}}$; otherwise $V = \lambda x.N$ for some $N \in \Lambda$, and then necessarily $V' = \lambda x.N'$ with $N \Rightarrow N'$, so $V \xRightarrow{\text{int}} V'$ by applying the rule λ for $\xRightarrow{\text{int}}$.

- **Remark 12. 1.** If $M \Rightarrow M'$ and $N \Rightarrow N'$ then $MN \Rightarrow M'N'$. For the proof, it is sufficient to consider the last rule of the derivation of $M \Rightarrow M'$.
2. If $\text{R} \in \{\xrightarrow{\beta_v}, \xrightarrow{\sigma}\}$ and $M \text{ R } M'$, then $MN \text{ R } M'N$ for any $N \in \Lambda$. For the proof, it is sufficient to consider the last rule of the derivation of $M \text{ R } M'$, for any $\text{R} \in \{\xrightarrow{\beta_v}, \xrightarrow{\sigma}\}$.
 3. If $M \xRightarrow{\text{int}} M'$ and $N \Rightarrow N'$ where $M' \notin \Lambda_v$, then $MN \xRightarrow{\text{int}} M'N'$: indeed, the last rule in the derivation of $M \xRightarrow{\text{int}} M'$ can be neither λ nor var because $M' \notin \Lambda_v$. The hypothesis $M' \notin \Lambda_v$ is crucial: for example, $x \xRightarrow{\text{int}} x$ and $I\Delta \Rightarrow \Delta$ but $I\Delta \not\xRightarrow{\text{int}} \Delta$ and thus $x(I\Delta) \not\xRightarrow{\text{int}} x\Delta$.
 4. $\rightarrow_v \subseteq \Rightarrow \subseteq \rightarrow_v^*$. As a consequence, $\Rightarrow^* = \rightarrow_v^*$ and (by Proposition 4) \Rightarrow is confluent.
 5. $\xrightarrow{\beta_v} \subseteq \xRightarrow{\text{int}} \subseteq \xrightarrow{\beta_v}^*$, so $\xRightarrow{\text{int}}^* = \xrightarrow{\beta_v}^*$. Thus, by Remark 11.3, variables and abstractions are preserved under $\xrightarrow{\beta_v}^*$ -expansion, i.e., if $M \xrightarrow{\beta_v}^* x$ (resp. $M \xrightarrow{\beta_v}^* \lambda x.N'$) then $M = x$ (resp. $M = \lambda x.N$ with $N \rightarrow_v^* N'$).
 6. For any $\text{R} \in \{\xrightarrow{\beta_v}, \xrightarrow{\sigma}\}$, if $M \text{ R } M'$ then $M\{V/x\} \text{ R } M'\{V/x\}$ for any $V \in \Lambda_v$. The proof is by straightforward induction on the derivation of $M \text{ R } M'$ for any $\text{R} \in \{\xrightarrow{\beta_v}, \xrightarrow{\sigma}\}$.

As expected, a basic property of parallel reduction \Rightarrow is the following:

► **Lemma 13** (Substitution lemma for \Rightarrow). *If $M \Rightarrow M'$ and $V \Rightarrow V'$ then $M\{V/x\} \Rightarrow M'\{V'/x\}$.*

Proof. By straightforward induction on the derivation of $M \Rightarrow M'$. ◀

The following lemma will play a crucial role in the proof of Lemmas 18-19 and shows that the head σ -reduction $\xrightarrow{\sigma}$ can be postponed after the head β_v -reduction $\xrightarrow{\beta_v}$.

► **Lemma 14** (Commutation of head reductions).

1. If $M \xrightarrow{h}_\sigma L \xrightarrow{h}_{\beta_v} N$ then there exists $L' \in \Lambda$ such that $M \xrightarrow{h}_{\beta_v} L' \xrightarrow{h}_\sigma N$.
2. If $M \xrightarrow{h}_\sigma^* L \xrightarrow{h}_{\beta_v} N$ then there exists $L' \in \Lambda$ such that $M \xrightarrow{h}_{\beta_v}^* L' \xrightarrow{h}_\sigma^* N$.
3. If $M \xrightarrow{h}_v^* M'$ then there exists $N \in \Lambda$ such that $M \xrightarrow{h}_{\beta_v}^* N \xrightarrow{h}_\sigma^* M'$.

Proof. 1. By induction on the derivation of $M \xrightarrow{h}_\sigma L$. Let us consider its last rule r .

If $r = \sigma_1$ then $M = (\lambda x.M_0)N_0L_0M_1 \dots M_m$ and $L = (\lambda x.M_0L_0)N_0M_1 \dots M_m$ where $m \in \mathbb{N}$ and $x \notin \text{fv}(L_0)$. Since $L \xrightarrow{h}_{\beta_v} N$, there are only two cases:

- either $N_0 \xrightarrow{h}_{\beta_v} N'_0$ and $N = (\lambda x.M_0L_0)N'_0M_1 \dots M_m$ (according to the rule *right* for $\xrightarrow{h}_{\beta_v}$), then $M \xrightarrow{h}_{\beta_v} (\lambda x.M_0)N'_0L_0M_1 \dots M_m \xrightarrow{h}_\sigma N$;
- or $N_0 \in \Lambda_v$ and $N = M_0\{N_0/x\}L_0M_1 \dots M_m$ (by the rule β_v , as $x \notin \text{fv}(L_0)$), so $M \xrightarrow{h}_{\beta_v} N$.

If $r = \sigma_3$ then $M = V((\lambda x.L_0)N_0)M_1 \dots M_m$ and $L = (\lambda x.VL_0)N_0M_1 \dots M_m$ with $m \in \mathbb{N}$ and $x \notin \text{fv}(V)$. Since $L \xrightarrow{h}_{\beta_v} N$, there are only two cases:

- either $N_0 \xrightarrow{h}_{\beta_v} N'_0$ and $N = (\lambda x.VL_0)N'_0M_1 \dots M_m$ (according to the rule *right* for $\xrightarrow{h}_{\beta_v}$), then $M \xrightarrow{h}_{\beta_v} V((\lambda x.L_0)N'_0)M_1 \dots M_m \xrightarrow{h}_\sigma N$;
- or $N_0 \in \Lambda_v$ and $N = VL_0\{N_0/x\}M_1 \dots M_m$ (by the rule β_v , as $x \notin \text{fv}(V)$), so $M \xrightarrow{h}_{\beta_v} N$.

Finally, if $r = \text{right}$ then $M = VN_0M_1 \dots M_m$ and $L = VN'_0M_1 \dots M_m$ with $m \in \mathbb{N}$ and $N_0 \xrightarrow{h}_\sigma N'_0$. By Remark 11.2, $N'_0 \notin \Lambda_v$ and thus, since $L \xrightarrow{h}_{\beta_v} N$, the only possibility is that $N'_0 \xrightarrow{h}_{\beta_v} N''_0$ and $N = VN''_0M_1 \dots M_m$ (according to the rule *right* for $\xrightarrow{h}_{\beta_v}$). By induction hypothesis, there exists $N'''_0 \in \Lambda$ such that $N_0 \xrightarrow{h}_{\beta_v} N'''_0 \xrightarrow{h}_\sigma N''_0$. Therefore, $M \xrightarrow{h}_{\beta_v} VN'''_0M_1 \dots M_m \xrightarrow{h}_\sigma N$.

2. Immediate consequence of Lemma 14.1, using standard techniques of rewriting theory.
3. Immediate consequence of Lemma 14.2, using standard techniques of rewriting theory. ◀

We are now able to travel over again Takahashi's method [17, 4] in our setting with β_v - and σ -reduction. The next four lemmas govern the strong parallel reduction and will be used to prove Lemma 19.

► **Lemma 15.** *If $M \Rightarrow M'$ and $N \Rightarrow N'$ and $M' \notin \Lambda_v$, then $MN \Rightarrow M'N'$.*

Proof. One has $MN \Rightarrow M'N'$ by Remark 12.1 and since $M \Rightarrow M'$. By hypothesis, there exist $m, n \in \mathbb{N}$ and $M_0, \dots, M_m, N_0, \dots, N_n$ such that $M = M_0$, $M_m = N_0$, $N_n \xrightarrow{\text{int}} M'$, $M_i \xrightarrow{h}_{\beta_v} M_{i+1}$ for any $0 \leq i < m$ and $N_j \xrightarrow{h}_\sigma N_{j+1}$ for any $0 \leq j < n$; by Remark 12.2, $M_iN \xrightarrow{h}_{\beta_v} M_{i+1}N$ for any $0 \leq i < m$ and $N_jN \xrightarrow{h}_\sigma N_{j+1}N$ for any $0 \leq j < n$. As $M' \notin \Lambda_v$, one has $N_nN \xrightarrow{\text{int}} M'N'$ by Remark 12.3. Therefore, $MN \Rightarrow M'N'$. ◀

► **Lemma 16.** *If $M \Rightarrow M'$ and $N \Rightarrow N'$ then $MN \Rightarrow M'N'$.*

Proof. If $M' \notin \Lambda_v$ then $MN \Rightarrow M'N'$ by Lemma 15 and since $N \Rightarrow N'$.

Assume $M' \in \Lambda_v$: $MN \Rightarrow M'N'$ by Remark 12.1, as $M \Rightarrow M'$ and $N \Rightarrow N'$. By hypothesis, there are $m, m', n, n' \in \mathbb{N}$ and $M_0, \dots, M_m, M'_0, \dots, M'_{m'}, N_0, \dots, N_n, N'_0, \dots, N'_{n'}$ such that:

- $M = M_0$, $M_m = M'_0$, $M'_{m'} \xrightarrow{\text{int}} M'$, $M_i \xrightarrow{h}_{\beta_v} M_{i+1}$ for any $0 \leq i < m$, and $M'_{i'} \xrightarrow{h}_\sigma M'_{i'+1}$ for any $0 \leq i' < m'$,
- $N = N_0$, $N_n = N'_0$, $N'_{n'} \xrightarrow{\text{int}} N'$, $N_j \xrightarrow{h}_{\beta_v} N_{j+1}$ for any $0 \leq j < n$ and $N'_{j'} \xrightarrow{h}_\sigma N'_{j'+1}$ for any $0 \leq j' < n'$.

By Remark 11.3, $M'_{m'} \in \Lambda_v$ since $M' \in \Lambda_v$, therefore $m' = 0$ by Remark 11.2, and thus $M_m = M'_0 \xrightarrow{\text{int}} M'$ (and $M_m \Rightarrow M'$ since $\xrightarrow{\text{int}} \subseteq \Rightarrow$) and $M_m \in \Lambda_v$. Using the rules *right* for $\xrightarrow{h}_{\beta_v}$ and \xrightarrow{h}_σ , one has $M_mN_j \xrightarrow{h}_{\beta_v} M_mN_{j+1}$ for any $0 \leq j < n$, and $M_mN'_{j'} \xrightarrow{h}_\sigma M_mN'_{j'+1}$ for any $0 \leq j' < n'$. By Remark 12.2, $M_iN_0 \xrightarrow{h}_{\beta_v} M_{i+1}N_0$ for any $0 \leq i < m$. By applying

the rule *right* for $\overset{int}{\Rightarrow}$, one has $M_m N'_n \overset{int}{\Rightarrow} M' N'$. Therefore, $MN = M_0 N_0 \xrightarrow{h}_{\beta_v}^* M_m N_0 \xrightarrow{h}_{\beta_v}^* M_m N_n = M_m N'_0 \xrightarrow{h}_{\sigma}^* M_m N'_n \overset{int}{\Rightarrow} M' N'$ and hence $MN \Rightarrow M' N'$. \blacktriangleleft

► **Lemma 17.** *If $M \overset{int}{\Rightarrow} M'$ and $V \Rightarrow V'$, then $M\{V/x\} \Rightarrow M'\{V'/x\}$.*

Proof. By Lemma 13, one has $M\{V/x\} \Rightarrow M'\{V'/x\}$ since $M \Rightarrow M'$ and $V \Rightarrow V'$. We proceed by induction on $M \in \Lambda$. Let us consider the last rule r of the derivation of $M \overset{int}{\Rightarrow} M'$.

If $r = var$ then there are two cases: either $M = x$ and then $M\{V/x\} = V \Rightarrow V' = M'\{V'/x\}$; or $M = y \neq x$ and then $M\{V/x\} = y = M'\{V'/x\}$, so $M\{V/x\} \Rightarrow M'\{V'/x\}$ by Remark 10.

If $r = \lambda$ then $M = \lambda y.N$ and $M' = \lambda y.N'$ with $N \Rightarrow N'$; we can suppose without loss of generality that $y \notin \text{fv}(V) \cup \{x\}$. One has $N\{V/x\} \Rightarrow N'\{V'/x\}$ according to Lemma 13. By applying the rule λ for $\overset{int}{\Rightarrow}$, one has $M\{V/x\} = \lambda y.N\{V/x\} \overset{int}{\Rightarrow} \lambda y.N'\{V'/x\} = M'\{V'/x\}$ and thus $M\{V/x\} \Rightarrow M'\{V'/x\}$.

Finally, if $r = right$ then $M = U N M_1 \dots M_m$ and $M' = U' N' M'_1 \dots M'_m$ for some $m \in \mathbb{N}$ with $U, U' \in \Lambda_v$, $U \Rightarrow U'$, $N \overset{int}{\Rightarrow} N'$ and $M_i \Rightarrow M'_i$ for any $1 \leq i \leq m$. By induction hypothesis, $U\{V/x\} \Rightarrow U'\{V'/x\}$ (indeed $U \overset{int}{\Rightarrow} U'$ according to Remark 11.5) and $N\{V/x\} \Rightarrow N'\{V'/x\}$. By Lemma 13, $M_i\{V/x\} \Rightarrow M'_i\{V'/x\}$ for any $1 \leq i \leq m$. By Lemma 16, $U\{V/x\} N\{V/x\} \Rightarrow U'\{V'/x\} N'\{V'/x\}$ and hence, by applying Lemma 15 m times since $U'\{V'/x\} N'\{V'/x\} \notin \Lambda_v$, one has $M\{V/x\} = U\{V/x\} N\{V/x\} M_1\{V/x\} \dots M_m\{V/x\} \Rightarrow U'\{V'/x\} N'\{V'/x\} M'_1\{V'/x\} \dots M'_m\{V'/x\} = M'\{V'/x\}$. \blacktriangleleft

In the proof of the two next lemmas, as well as in the proof of Corollary 21 and Theorem 22, our Lemma 14 plays a crucial role: indeed, since the head σ -reduction well interact with the head β_v -reduction, Takahashi's method [17, 4] is still working when adding the reduction rules σ_1 and σ_3 to Plotkin's β_v -reduction.

► **Lemma 18** (Substitution lemma for \Rightarrow). *If $M \Rightarrow M'$ and $V \Rightarrow V'$ then $M\{V/x\} \Rightarrow M'\{V'/x\}$.*

Proof. By Lemma 13, one has $M\{V/x\} \Rightarrow M'\{V'/x\}$ since $M \Rightarrow M'$ and $V \Rightarrow V'$. By hypothesis, there exist $m, n \in \mathbb{N}$ and $M_0, \dots, M_m, N_0, \dots, N_n$ such that $M = M_0$, $M_m = N_0$, $N_n \overset{int}{\Rightarrow} M'$, $M_i \xrightarrow{h}_{\beta_v} M_{i+1}$ for any $0 \leq i < m$ and $N_j \xrightarrow{h}_{\sigma} N_{j+1}$ for any $0 \leq j < n$; by Remark 12.6, $M_i\{V/x\} \xrightarrow{h}_{\beta_v} M_{i+1}\{V/x\}$ for any $0 \leq i < m$, and $N_j\{V/x\} \xrightarrow{h}_{\sigma} N_{j+1}\{V/x\}$ for any $0 \leq j < n$. By Lemma 17, one has $N_n\{V/x\} \Rightarrow M'\{V'/x\}$, thus there exist $L, N \in \Lambda$ such that $M\{V/x\} \xrightarrow{h}_{\beta_v}^* N_0\{V/x\} \xrightarrow{h}_{\sigma}^* N_n\{V/x\} \xrightarrow{h}_{\beta_v}^* N \xrightarrow{h}_{\sigma}^* L \overset{int}{\Rightarrow} M'\{V'/x\}$. By Lemma 14.2, there exists $N' \in \Lambda$ such that $M\{V/x\} \xrightarrow{h}_{\beta_v}^* N_0\{V/x\} \xrightarrow{h}_{\beta_v}^* N' \xrightarrow{h}_{\sigma}^* N \xrightarrow{h}_{\sigma}^* L \overset{int}{\Rightarrow} M'\{V'/x\}$, therefore $M\{V/x\} \Rightarrow M'\{V'/x\}$. \blacktriangleleft

Now we are ready to prove a key lemma, which states that parallel reduction \Rightarrow coincides with strong parallel reduction \Rightarrow (the inclusion $\Rightarrow \subseteq \Rightarrow$ is trivial).

► **Lemma 19** (Key Lemma). *If $M \Rightarrow M'$ then $M \Rightarrow M'$.*

Proof. By induction on the derivation of $M \Rightarrow M'$. Let us consider its last rule r .

If $r = var$ then $M = x M_1 \dots M_m$ and $M' = x M'_1 \dots M'_m$ where $m \in \mathbb{N}$ and $M_i \Rightarrow M'_i$ for any $1 \leq i \leq m$. By reflexivity of \Rightarrow (Remark 10), $x \Rightarrow x$. By induction hypothesis, $M_i \Rightarrow M'_i$ for any $1 \leq i \leq m$. Therefore, $M \Rightarrow M'$ by applying Lemma 16 m times.

If $r = \lambda$ then $M = (\lambda x.M_0) M_1 \dots M_m$ and $M' = (\lambda x.M'_0) M'_1 \dots M'_m$ where $m \in \mathbb{N}$ and $M_i \Rightarrow M'_i$ for any $0 \leq i \leq m$. By induction hypothesis, $M_i \Rightarrow M'_i$ for any $1 \leq i \leq m$. According to Remark 11.4, $\lambda x.M_0 \Rightarrow \lambda x.M'_0$. So $M \Rightarrow M'$ by applying Lemma 16 m times.

If $r = \beta_v$ then $M = (\lambda x.M_0) V M_1 \dots M_m$ and $M' = M'_0\{V'/x\} M'_1 \dots M'_m$ where $m \in \mathbb{N}$, $V \Rightarrow V'$ and $M_i \Rightarrow M'_i$ for any $0 \leq i \leq m$. By induction hypothesis, $V \Rightarrow V'$ and $M_i \Rightarrow M'_i$

for any $0 \leq i \leq m$. By applying the rule β_v for $\xrightarrow{\beta_v}$, one has $M \xrightarrow{\beta_v} M_0\{V/x\}M_1 \dots M_m$; moreover $M_0\{V/x\}M_1 \dots M_m \equiv M'$ by Lemma 18 and by applying Lemma 16 m times, thus there are $L, N \in \Lambda$ such that $M \xrightarrow{\beta_v} M_0\{V/x\}M_1 \dots M_m \xrightarrow{\beta_v^*} L \xrightarrow{\sigma^*} N \xrightarrow{\text{int}} M'$. So $M \equiv M'$.

If $r = \sigma_1$ then $M = (\lambda x.M_0)N_0L_0M_1 \dots M_m$ and $M' = (\lambda x.M'_0L'_0)N'_0M'_1 \dots M'_m$ where $m \in \mathbb{N}$, $L_0 \Rightarrow L'_0$, $N_0 \Rightarrow N'_0$ and $M_i \Rightarrow M'_i$ for any $0 \leq i \leq m$. By induction hypothesis, $N_0 \Rightarrow N'_0$ and $M_i \Rightarrow M'_i$ for any $1 \leq i \leq m$. By applying the rule σ_1 for $\xrightarrow{\sigma}$, one has $M \xrightarrow{\sigma} (\lambda x.M_0L_0)N_0M_1 \dots M_m$. By Remark 12.1, $M_0L_0 \Rightarrow M'_0L'_0$ and thus $\lambda x.M_0L_0 \Rightarrow \lambda x.M'_0L'_0$ according to Remark 11.4. So $(\lambda x.M_0L_0)N_0M_1 \dots M_m \equiv M'$ by applying Lemma 16 $m+1$ times, hence there are $L, N \in \Lambda$ such that $M \xrightarrow{\sigma} (\lambda x.M_0L_0)N_0M_1 \dots M_m \xrightarrow{\beta_v^*} L \xrightarrow{\sigma^*} N \xrightarrow{\text{int}} M'$. By Lemma 14.2, there is $L' \in \Lambda$ such that $M \xrightarrow{\beta_v^*} L' \xrightarrow{\sigma^*} L \xrightarrow{\sigma^*} N \xrightarrow{\text{int}} M'$ and thus $M \equiv M'$.

Finally, if $r = \sigma_3$ then $M = V((\lambda x.L_0)N_0)M_1 \dots M_m$ and $M' = (\lambda x.V'L'_0)N'_0M'_1 \dots M'_m$ with $m \in \mathbb{N}$, $V \Rightarrow V'$, $L_0 \Rightarrow L'_0$, $N_0 \Rightarrow N'_0$ and $M_i \Rightarrow M'_i$ for any $1 \leq i \leq m$. By induction hypothesis, $N_0 \Rightarrow N'_0$ and $M_i \Rightarrow M'_i$ for any $1 \leq i \leq m$. By the rule σ_3 for $\xrightarrow{\sigma}$, one has $M \xrightarrow{\sigma} (\lambda x.VL_0)N_0M_1 \dots M_m$. By Remark 12.1, $VL_0 \Rightarrow V'L'_0$ and thus $\lambda x.VL_0 \Rightarrow \lambda x.V'L'_0$ according to Remark 11.4. So $(\lambda x.VL_0)N_0M_1 \dots M_m \equiv M'$ by applying Lemma 16 $m+1$ times, hence there are $L, N \in \Lambda$ such that $M \xrightarrow{\sigma} (\lambda x.VL_0)N_0M_1 \dots M_m \xrightarrow{\beta_v^*} L \xrightarrow{\sigma^*} N \xrightarrow{\text{int}} M'$. By Lemma 14.2, there is $L' \in \Lambda$ such that $M \xrightarrow{\beta_v^*} L' \xrightarrow{\sigma^*} L \xrightarrow{\sigma^*} N \xrightarrow{\text{int}} M'$, so $M \equiv M'$. \blacktriangleleft

Next Lemma 20 and Corollary 21 show that internal parallel reduction can be shifted after head reductions.

► **Lemma 20 (Postponement).** *If $M \xrightarrow{\text{int}} L$ and $L \xrightarrow{\beta_v} N$ (resp. $L \xrightarrow{\sigma} N$) then there exists $L' \in \Lambda$ such that $M \xrightarrow{\beta_v} L'$ (resp. $M \xrightarrow{\sigma} L'$) and $L' \Rightarrow N$.*

Proof. By induction on the derivation of $M \xrightarrow{\text{int}} L$. Let us consider its last rule r .

If $r = \text{var}$, then $M = x = L$ which contradicts $L \xrightarrow{\beta_v} N$ and $L \xrightarrow{\sigma} N$ by Remarks 11.1-2.

If $r = \lambda$ then $L = \lambda x.L'$ for some $L' \in \Lambda$, which contradicts $L \xrightarrow{\beta_v} N$ and $L \xrightarrow{\sigma} N$ by Remarks 11.1-2.

Finally, if $r = \text{right}$ then $M = VM_0M_1 \dots M_m$ and $L = V'L_0L_1 \dots L_m$ where $m \in \mathbb{N}$, $V \Rightarrow V'$ (so $V \xrightarrow{\text{int}} V'$ by Remark 11.5), $M_0 \xrightarrow{\text{int}} L_0$ (thus $M_0 \Rightarrow L_0$ since $\xrightarrow{\text{int}} \subseteq \Rightarrow$) and $M_i \Rightarrow L_i$ for any $1 \leq i \leq m$.

- If $L \xrightarrow{\beta_v} N$ then there are two cases, depending on the last rule r' of the derivation of $L \xrightarrow{\beta_v} N$.
 - If $r' = \beta_v$ then $V' = \lambda x.N'_0$, $L_0 \in \Lambda_v$ and $N = N'_0\{L_0/x\}L_1 \dots L_m$, thus $M_0 \in \Lambda_v$ and $V = \lambda x.N_0$ with $N_0 \Rightarrow N'_0$ by Remark 11.3. By Lemma 13, one has $N_0\{M_0/x\} \Rightarrow N'_0\{L_0/x\}$. Let $L' = N_0\{M_0/x\}M_1 \dots M_m$: so $M = (\lambda x.N_0)M_0M_1 \dots M_m \xrightarrow{\beta_v} L'$ (apply the rule β_v for $\xrightarrow{\beta_v}$) and $L' \Rightarrow N$ by applying Remark 12.1 m times.
 - If $r' = \text{right}$ then $N = V'N_0L_1 \dots L_m$ with $L_0 \xrightarrow{\beta_v} N_0$. By induction hypothesis, there exists $L'_0 \in \Lambda$ such that $M_0 \xrightarrow{\beta_v} L'_0 \Rightarrow N_0$. Let $L' = VL'_0M_1 \dots M_m$: so $M \xrightarrow{\beta_v} L'$ (apply the rule right for $\xrightarrow{\beta_v}$) and $L' \Rightarrow N$ by applying Remark 12.1 $m+1$ times.
- If $L \xrightarrow{\sigma} N$ then there are three cases, depending on the last rule r' of the derivation of $L \xrightarrow{\sigma} N$.
 - If $r' = \sigma_1$ then $m > 0$, $V' = \lambda x.N'_0$ and $N = (\lambda x.N'_0L_1)L_0L_2 \dots L_m$, thus $V = \lambda x.N_0$ with $N_0 \Rightarrow N'_0$ by Remark 11.3. Using Remarks 12.1 and 11.4, one has $\lambda x.N_0M_1 \Rightarrow \lambda x.N'_0L_1$. Let $L' = (\lambda x.N_0M_1)M_0M_2 \dots M_m$: so $M = (\lambda x.N_0)M_0M_1 \dots M_m \xrightarrow{\sigma} L'$ (apply the rule σ_1 for $\xrightarrow{\sigma}$) and $L' \Rightarrow N$ by applying Remark 12.1 m times.
 - If $r' = \sigma_3$ then $L_0 = (\lambda x.L_{01})L_{02}$ and $N = (\lambda x.V'L_{01})L_{02}L_1 \dots L_m$. Since $M_0 \xrightarrow{\text{int}} (\lambda x.L_{01})L_{02}$, necessarily $M_0 = (\lambda x.M_{01})M_{02}$ with $M_{01} \Rightarrow L_{01}$ and $M_{02} \xrightarrow{\text{int}} L_{02}$ (so $M_{02} \Rightarrow L_{02}$). Using Remarks 12.1 and 11.4, one has $\lambda x.VM_{01} \Rightarrow \lambda x.V'L_{01}$. Let

$L' = (\lambda x.VM_{01})M_{02}M_1 \dots M_m$: therefore $M = V((\lambda x.M_{01})M_{02})M_1 \dots M_m \xrightarrow{h}_\sigma L'$ (apply the rule σ_3 for \xrightarrow{h}_σ) and $L' \Rightarrow N$ by applying Remark 12.1 $m + 1$ times.

- If $r' = \text{right}$ then $N = V'N_0L_1 \dots L_m$ with $L_0 \xrightarrow{h}_\sigma N_0$. By induction hypothesis, there exists $L'_0 \in \Lambda$ such that $M_0 \xrightarrow{h}_\sigma L'_0 \Rightarrow N_0$. Let $L' = VL'_0M_1 \dots M_m$: so $M \xrightarrow{h}_\sigma L'$ (apply the rule right for \xrightarrow{h}_σ) and $L' \Rightarrow N$ by applying Remark 12.1 $m + 1$ times. ◀

► **Corollary 21.** *If $M \xrightarrow{\text{int}} L$ and $L \xrightarrow{h}_{\beta_v} N$ (resp. $L \xrightarrow{h}_\sigma N$), then there exist $L', L'' \in \Lambda$ such that $M \xrightarrow{h}_{\beta_v} L' \xrightarrow{h}_\sigma L'' \xrightarrow{\text{int}} N$ (resp. $M \xrightarrow{h}_{\beta_v} L' \xrightarrow{h}_\sigma L'' \xrightarrow{\text{int}} N$).*

Proof. Immediate by Lemma 20 and Lemma 19, applying Lemma 14.2 if $L \xrightarrow{h}_\sigma N$. ◀

Now we obtain our first main result (Theorem 22): any v -reduction sequence can be sequentialized into a head β_v -reduction sequence followed by a head σ -reduction sequence, followed by an internal reduction sequence. In ordinary λ -calculus, the well-known result corresponding to our Theorem 22 says that a β -reduction sequence can be factorized in a head reduction sequence followed by an internal reduction sequence (see for example [17, Corollary 2.6]).

► **Theorem 22 (Sequentialization).** *If $M \rightarrow_v^* M'$ then there exist $L, N \in \Lambda$ such that $M \xrightarrow{h}_{\beta_v} L \xrightarrow{h}_\sigma N \xrightarrow{\text{int}}_v M'$.*

Proof. By Remark 12.4, $M \Rightarrow^* M'$ and thus there are $m \in \mathbb{N}$ and $M_0, \dots, M_m \in \Lambda$ such that $M = M_0$, $M_m = M'$ and $M_i \Rightarrow M_{i+1}$ for any $0 \leq i < m$. We prove by induction on $m \in \mathbb{N}$ that there are $L, N \in \Lambda$ such that $M \xrightarrow{h}_{\beta_v} L \xrightarrow{h}_\sigma N \xrightarrow{\text{int}}^* M'$, so $N \xrightarrow{\text{int}}_v M'$ by Remark 12.5.

If $m = 0$ then $M = M_0 = M'$ and hence we conclude by taking $L = M' = N$.

Finally, suppose $m > 0$. By induction hypothesis applied to $M_1 \Rightarrow^* M'$, there exist $L', N' \in \Lambda$ such that $M_1 \xrightarrow{h}_{\beta_v} L' \xrightarrow{h}_\sigma N' \xrightarrow{\text{int}}^* M'$. By applying Lemma 19 to M , there exist $L_0, N_0 \in \Lambda$ such that $M \xrightarrow{h}_{\beta_v} L_0 \xrightarrow{h}_\sigma N_0 \xrightarrow{\text{int}} M_1$. By applying Corollary 21 repeatedly, there exists $N \in \Lambda$ such that $N_0 \xrightarrow{h}_\sigma N \xrightarrow{\text{int}} N'$ and hence $M \xrightarrow{h}_{\beta_v} N \xrightarrow{\text{int}}^* M'$. According to Lemma 14.3, there exists $L \in \Lambda$ such that $M \xrightarrow{h}_{\beta_v} L \xrightarrow{h}_\sigma N \xrightarrow{\text{int}}^* M'$. ◀

It is worth noticing that in Definition 7 there is no distinction between head σ_1 - and head σ_3 -reduction steps, and, according to it, the sequentialization of Theorem 22 imposes no order between head σ -reductions. We denote by $\xrightarrow{h}_{\sigma_1}$ and $\xrightarrow{h}_{\sigma_3}$ respectively the reduction relations $\rightarrow_{\sigma_1} \cap \xrightarrow{h}_\sigma$ and $\rightarrow_{\sigma_3} \cap \xrightarrow{h}_\sigma$. So, a natural question arises: is it possible to sequentialize them? The answer is negative, as proved by the next two counterexamples.

- Let $M = x((\lambda y.z')(zI))\Delta$ and $N = (\lambda y.xz'\Delta)(zI)$: $M \xrightarrow{h}_{\sigma_3} (\lambda y.xz')\Delta \xrightarrow{h}_{\sigma_1} N$, but there exists no L such that $M \xrightarrow{h}_{\sigma_1} L \xrightarrow{h}_{\sigma_3} N$. In fact M contains only a head σ_3 -redex and $(\lambda y.xz')\Delta$ contains only a head σ_1 -redex
- Let $M = x((\lambda y.z')(zI)\Delta)$ and $N = (\lambda y.x(z'\Delta))(zI)$: $M \xrightarrow{h}_{\sigma_1} x((\lambda y.z'\Delta)(zI)) \xrightarrow{h}_{\sigma_3} N$ but there is no L such that $M \xrightarrow{h}_{\sigma_3} L \xrightarrow{h}_{\sigma_1} N$. In fact M contains only a head σ_1 -redex and $x((\lambda y.z'\Delta)(zI))$ contains only a head σ_3 -redex.

So, the impossibility of sequentializing a head σ -reduction sequence is due to the fact that a head σ_1 -reduction step can create a head σ_3 -redex, and viceversa. This is not a problem, since head σ -reduction is strongly normalizing (by Proposition 4 and since $\xrightarrow{h}_\sigma \subseteq \rightarrow_\sigma$). Our approach does not force a strict order of head σ -reductions.

4 Standardization

Now we are able to prove the main result of this paper, i.e., a standardization theorem for λ_v^σ (Theorem 25). In particular we provide a notion of standard reduction sequence that avoids any auxiliary notion of residual redexes, by closely following the definition given in [15].

► **Notation.** For any $k, m \in \mathbb{N}$ with $k \leq m$, we denote by $[M_0, \dots, M_k, \dots, M_m]^{head}$ a sequence of terms such that $M_i \xrightarrow{h}_{\beta_v} M_{i+1}$ when $0 \leq i < k$, and $M_i \xrightarrow{h}_{\sigma} M_{i+1}$ when $k \leq i < m$.

It is easy to check that $[M]^{head}$ for any $M \in \Lambda$. The notion of standard sequence of terms is defined by using the previous notion of *head*-sequence. Our notion of standard reduction sequence is mutually defined together with the notion of inner-sequence of terms (Definition 23). This definition allows us to avoid non-deterministic cases remarked in [7] (we provide more details at the end of this section). We denote by $[M_0, \dots, M_m]^{std}$ (resp. $[M_0, \dots, M_m]^{in}$) a standard (resp. inner) sequence of terms.

► **Definition 23** (Standard and inner sequences). *Standard* and *inner sequences* of terms are defined by mutual induction as follows:

1. if $[M_0, \dots, M_m]^{head}$ and $[M_m, \dots, M_{m+n}]^{in}$ then $[M_0, \dots, M_{m+n}]^{std}$, where $m, n \in \mathbb{N}$;
2. $[M]^{in}$, for any $M \in \Lambda$;
3. if $[M_0, \dots, M_m]^{std}$ then $[\lambda z.M_0, \dots, \lambda z.M_m]^{in}$, where $m \in \mathbb{N}$;
4. if $[V_0, \dots, V_h]^{std}$ and $[N_0, \dots, N_n]^{in}$ then $[V_0 N_0, \dots, V_h N_n]^{in}$, where $h, n \in \mathbb{N}$;
5. if $[N_0, \dots, N_n]^{in}$, $[M_0, \dots, M_m]^{std}$ and $N_0 \notin \Lambda_v$, then $[N_0 M_0, \dots, N_n M_0, \dots, N_n M_m]^{in}$, where $m, n \in \mathbb{N}$.

For instance, let $M = (\lambda y.Ix)(z(\Delta I))(II)$: $M \rightarrow_v (\lambda y.Ix)(z(\Delta I))I \rightarrow_v (\lambda y.x)(z(\Delta I))I$ and $M \rightarrow_v (\lambda y.Ix(II))(z(\Delta I)) \rightarrow_v (\lambda y.Ix(II))(z(II))$ are not standard sequences; $M \rightarrow_v (\lambda y.Ix)(z(\Delta I))I$ and $M \rightarrow_v (\lambda y.Ix)(z(II))(II) \rightarrow_v (\lambda y.Ix)(zI)(II) \rightarrow_v (\lambda y.Ix(II))(zI) \rightarrow_v (\lambda y.x(II))(zI) \rightarrow_v (\lambda y.xI)(zI)$ are standard sequences.

► **Remark 24.** For any $n \in \mathbb{N}$, if $[N_0, \dots, N_n]^{in}$ (resp. $[N_0, \dots, N_n]^{head}$) then $[N_0, \dots, N_n]^{std}$. Indeed, $[N_0]^{head}$ (resp. $[N_n]^{in}$ by Definition 23.2), so $[N_0, \dots, N_n]^{std}$ by Definition 23.1.

In particular, $[N]^{std}$ for any $N \in \Lambda$: apply Definition 23.2 and Remark 24 for $n = 0$.

► **Theorem 25** (Standardization).

1. If $M \rightarrow_v^* M'$ then there is a sequence $[M, \dots, M']^{std}$.
2. If $M \xrightarrow{int}_{\beta_v}^* M'$ then there is a sequence $[M, \dots, M']^{in}$.

Proof. Both statements are proved simultaneously by induction on $M' \in \Lambda$.

1. - If $M' = z$ then, by Theorem 22, $M \xrightarrow{h}_{\beta_v}^* L \xrightarrow{h}_{\sigma}^* N \xrightarrow{int}_{\beta_v}^* z$ for some $L, N \in \Lambda$. By Remarks 12.5 and 11.2, $L = N = z$; therefore $M \xrightarrow{h}_{\beta_v}^* z$ and hence there is a sequence $[M, \dots, z]^{head}$. Thus, $[M, \dots, z]^{std}$ by Remark 24.
- If $M' = \lambda z.N'$ then, by Theorem 22, $M \xrightarrow{h}_{\beta_v}^* L \xrightarrow{h}_{\sigma}^* L' \xrightarrow{int}_{\beta_v}^* \lambda z.N'$ for some $L, L' \in \Lambda$. By Remarks 12.5 and 11.2, $L = L' = \lambda z.N$ with $N \rightarrow_v^* N'$. So $M \xrightarrow{h}_{\beta_v}^* \lambda z.N$ and hence there is a sequence $[M, \dots, \lambda z.N]^{head}$. By induction on (1), there is a sequence $[N, \dots, N']^{std}$, thus $[\lambda z.N, \dots, \lambda z.N']^{in}$ by Definition 23.3. Therefore $[M, \dots, \lambda z.N, \dots, \lambda z.N']^{std}$ by Definition 23.1.
- If $M' = N'L'$ then, by Theorem 22, $M \xrightarrow{h}_{\beta_v}^* M'' \xrightarrow{h}_{\sigma}^* M_0 \xrightarrow{int}_{\beta_v}^* N'L'$ for some $M'', M_0 \in \Lambda$. By Remark 3, $M_0 = NL$ for some $N, L \in \Lambda$, since $\xrightarrow{int}_{\beta_v}^* \subseteq \rightarrow_v^*$ and $M' \notin \Lambda_v$. Thus there is a sequence $[M, \dots, M'', \dots, NL]^{head}$. By Remark 12.5, $NL \xrightarrow{int}_{\beta_v}^* N'L'$; clearly, each step of $\xrightarrow{int}_{\beta_v}^*$ is an instance of the rule *right* of Definition 9. There are two sub-cases.

- If $N \in \Lambda_v$ then $N \Rightarrow^* N'$ and $L \xrightarrow{\text{int}^*} L'$, so $N \rightarrow_v^* N'$ and $L \xrightarrow{\text{int}^*}_v L'$ by Remarks 12.4-5. By induction respectively on (1) and (2), there are sequences $[N, \dots, N']^{std}$ and $[L, \dots, L']^{in}$, thus $[NL, \dots, NL', \dots, N'L']^{in}$ by Definition 23.4. Therefore $[M, \dots, M'', \dots, NL, \dots, NL', \dots, N'L']^{std}$ by Definition 23.1.
 - If $N \notin \Lambda_v$ (i.e., $N = VM_1 \dots M_m$ with $m > 0$) then $N \xrightarrow{\text{int}^*} N'$ and $L \Rightarrow^* L'$, so $N \xrightarrow{\text{int}^*}_v N'$ and $L \rightarrow_v^* L'$ by Remarks 12.4-5. By induction respectively on (2) and (1), there are sequences $[N, \dots, N']^{in}$ and $[L, \dots, L']^{std}$. Hence $[NL, \dots, N'L, \dots, N'L']^{in}$ by Definition 23.5. Thus $[M, \dots, M'', \dots, NL, \dots, N'L, \dots, N'L']^{std}$ by Definition 23.1.
2. ■ If $M' = z$ then $M = z$ by Remark 12.5, hence $[z]^{in}$ by Definition 23.2.
- If $M' = \lambda z.L'$ then $M = \lambda z.L$ and $L \rightarrow_v^* L'$ by Remark 12.5. Hence there is a sequence $[L, \dots, L']^{std}$ by induction on (1). By Definition 23.3, $[\lambda z.L, \dots, \lambda z.L']^{in}$.
 - If $M' = N'L'$ then $M = NL$ for some $N, L \in \Lambda$ by Remark 3, since $\xrightarrow{\text{int}^*}_v \subseteq \rightarrow_v^*$ and $M' \notin \Lambda_v$. By Remark 12.5, $NL \xrightarrow{\text{int}^*} N'L'$; clearly, each step of $\xrightarrow{\text{int}^*}$ is an instance of the rule *right* of Definition 9. There are two sub-cases.
 - If $N \in \Lambda_v$ then $N \Rightarrow^* N'$ and $L \xrightarrow{\text{int}^*} L'$, so $N \rightarrow_v^* N'$ and $L \xrightarrow{\text{int}^*}_v L'$ by Remarks 12.4-5. Thus there are sequences $[N, \dots, N']^{std}$ and $[L, \dots, L']^{in}$ by induction respectively on (1) and (2). Therefore, by Definition 23.4, $[NL, \dots, NL', \dots, N'L']^{in}$.
 - If $N \notin \Lambda_v$ (i.e. $N = VM_1 \dots M_m$ with $m > 0$) then $N \xrightarrow{\text{int}^*} N'$ and $L \Rightarrow^* L'$, thus $N \xrightarrow{\text{int}^*}_v N'$ and $L \rightarrow_v^* L'$ by Remarks 12.4-5. By induction respectively on (2) and (1), there are sequences $[N, \dots, N']^{in}$ and $[L, \dots, L']^{std}$. So $[NL, \dots, N'L, \dots, N'L']^{in}$ by Definition 23.5. \blacktriangleleft

Due to non-sequentialization of head σ_1 - and head σ_3 -reductions, several standard sequences may have the same starting term and ending term: for instance, if $M = I(\Delta I)I$ and $N = (\lambda z.(\lambda x.xI)(zz))I$ then $M \rightarrow_v (\lambda x.xI)(\Delta I) \rightarrow_v N$ and $M \rightarrow_v (\lambda z.I(zz))II \rightarrow_v (\lambda z.I(zz)I)I \rightarrow_v N$ are both standard sequences from M to N .

Finally, we can compare our notion of standardization with that given in [15]. To make the comparison possible we avoid σ -reductions and we recall that $\xrightarrow{h}_{\beta_v}$ is exactly the Plotkin's left-reduction [15, p. 136]. As remarked in [7, §1.5 p. 149], both sequences $(\lambda z.II)(II) \rightarrow_v (\lambda z.I)(II) \rightarrow_v (\lambda z.I)I$ and $(\lambda z.II)(II) \rightarrow_v (\lambda z.II)I \rightarrow_v (\lambda z.I)I$ are standard according to [15]. On the other hand, only the second sequence is standard in our sense. It is easy to check that collapsing the two notions of inner and standard sequence given in Definition 23, we get a notion of standard sequence that accept both the above sequences.

5 Some conservativity results

The sequentialization result (Theorem 22) has some interesting semantic consequences. It allows us to prove that (Corollary 29) the λ_v^σ -calculus is sound with respect to the call-by-value observational equivalence introduced by Plotkin in [15] for λ_v . Moreover we can prove that some notions, like that of potential valuability and solvability, introduced in [13] for λ_v , coincide with the respective notions for λ_v^σ (Theorem 31). This justifies the idea that λ_v^σ is a useful tool for studying the properties of λ_v . Our starting point is the following corollary.

► Corollary 26.

1. If $M \rightarrow_v^* V \in \Lambda_v$ then there exists $V' \in \Lambda_v$ such that $M \xrightarrow{h}_{\beta_v} V' \xrightarrow{\text{int}^*}_v V$.
2. For every $V \in \Lambda_v$, $M \xrightarrow{h}_{\beta_v} V$ if and only if $M \xrightarrow{h}_v V$.

Proof. The first point is proved by observing that, by Theorem 22, there are $N, L \in \Lambda$ such that $M \xrightarrow{h}_{\beta_v} L \xrightarrow{h}_\sigma N \xrightarrow{\text{int}^*}_v V$. By Remark 12.5, $N \in \Lambda_v$ and thus $L = N$ according to

Remark 11.2. Concerning the second point, the right-to-left direction is a consequence of Lemma 14.3 and Remark 11.2; the left-to-right direction follows from $\xrightarrow{h}_{\beta_v} \subseteq \xrightarrow{h}_{\nu}$. ◀

Let us recall the notion of observational equivalence defined by Plotkin [15] for λ_v .

► **Definition 27** (Halting, observational equivalence). Let $M \in \Lambda$. We say that (*the evaluation of*) M *halts* if there exists $V \in \Lambda_v$ such that $M \xrightarrow{h}_{\beta_v}^* V$.

The (*call-by-value*) *observational equivalence* is an equivalence relation \cong on Λ defined by: $M \cong N$ if, for every context C , one has that $C(M)$ halts iff $C(N)$ halts.¹

Clearly, similar notions can be defined for λ_v^σ using \xrightarrow{h}_{ν} instead of $\xrightarrow{h}_{\beta_v}$. Head σ -reduction plays no role neither in deciding the halting problem for evaluation (Corollary 26.1), nor in reaching a particular value (Corollary 26.2). So, we can conclude that the notions of halting and observational equivalence in λ_v^σ coincide with the ones in λ_v , respectively.

Now we compare the equational theory of λ_v^σ with Plotkin's observational equivalence.

► **Theorem 28** (Adequacy of ν -reduction). *If $M \rightarrow_v^* M'$ then: M halts iff M' halts.*

Proof. If $M' \xrightarrow{h}_{\beta_v}^* V \in \Lambda_v$ then $M \rightarrow_v^* M' \rightarrow_v^* V$ since $\xrightarrow{h}_{\beta_v} \subseteq \rightarrow_v$. By Corollary 26.1, there exists $V' \in \Lambda_v$ such that $M \xrightarrow{h}_{\beta_v}^* V'$. Thus M halts.

Conversely, if $M \xrightarrow{h}_{\beta_v}^* V \in \Lambda_v$ then $M \rightarrow_v^* V$ since $\xrightarrow{h}_{\beta_v} \subseteq \rightarrow_v$. By confluence of \rightarrow_v (Proposition 4, since $M \rightarrow_v^* M'$) and Remark 3 (since $V \in \Lambda_v$), there is $V' \in \Lambda_v$ such that $V \rightarrow_v^* V'$ and $M' \rightarrow_v^* V'$. By Corollary 26.1, there is $V'' \in \Lambda_v$ such that $M' \xrightarrow{h}_{\beta_v}^* V''$. So M' halts. ◀

► **Corollary 29** (Soundness). *If $M =_v N$ then $M \cong N$.*

Proof. Let C be a context. By confluence of \rightarrow_v (Proposition 4), $M =_v N$ implies that there exists $L \in \Lambda$ such that $M \rightarrow_v^* L$ and $N \rightarrow_v^* L$, hence $C(M) \rightarrow_v^* C(L)$ and $C(N) \rightarrow_v^* C(L)$. By Theorem 28, $C(M)$ halts iff $C(L)$ halts iff $C(N)$ halts. Therefore, $M \cong N$. ◀

Plotkin [15, Theorem 5] has already proved that $M =_{\beta_v} N$ implies $M \cong N$, but our Corollary 29 is not obvious since our λ_v^σ -calculus equates more than Plotkin's λ_v -calculus ($=_{\beta_v} \subseteq =_v$ since $\rightarrow_{\beta_v} \subseteq \rightarrow_v$, and Example 5 shows that this inclusion is strict).

The converse of Corollary 29 does not hold since $\lambda x.x(\lambda y.xy) \cong \Delta$ but $\lambda x.x(\lambda y.xy)$ and Δ are different ν -normal forms and hence $\lambda x.x(\lambda y.xy) \neq_v \Delta$ by confluence of \rightarrow_v (Proposition 4).

A further remarkable consequence of Corollary 26.1 is that the notions of potential valuability and solvability for λ_v^σ -calculus (studied in [3]) can be shown to coincide with the ones for Plotkin's λ_v -calculus (studied in [13, 14]), respectively. Let us recall their definition.

► **Definition 30** (Potential valuability, solvability). Let M be a term:

- M is ν -potentially valuable (resp. β_v -potentially valuable) if there are $m \in \mathbb{N}$, pairwise distinct variables x_1, \dots, x_m and $V, V_1, \dots, V_m \in \Lambda_v$ such that $M\{V_1/x_1, \dots, V_m/x_m\} \rightarrow_v^* V$ (resp. $M\{V_1/x_1, \dots, V_m/x_m\} \rightarrow_{\beta_v}^* V$);
- M is ν -solvable (resp. β_v -solvable) if there are $n, m \in \mathbb{N}$, variables x_1, \dots, x_m and $N_1, \dots, N_n \in \Lambda$ such that $(\lambda x_1 \dots x_m.M)N_1 \dots N_n \rightarrow_v^* I$ (resp. $(\lambda x_1 \dots x_m.M)N_1 \dots N_n \rightarrow_{\beta_v}^* I$).

► **Theorem 31.** *Let M be a term:*

1. M is ν -potentially valuable if and only if M is β_v -potentially valuable;
2. M is ν -solvable if and only if M is β_v -solvable.

¹ Original Plotkin's definition of call-by-value observational equivalence (see [15]) also requires that $C(M)$ and $C(N)$ are closed terms, according to the tradition identifying programs with closed terms.

Proof. In both points, the implication from right to left is trivial since $\rightarrow_{\beta_v} \subseteq \rightarrow_v$. Let us prove the other direction.

1. Since M is v -potentially valuable, there are variables x_1, \dots, x_m and $V, V_1, \dots, V_m \in \Lambda_v$ (with $m \geq 0$) such that $M\{V_1/x_1, \dots, V_m/x_m\} \rightarrow_v^* V$; then, there exists $V' \in \Lambda_v$ such that $M\{V_1/x_1, \dots, V_m/x_m\} \rightarrow_{\beta_v}^* V'$ by Corollary 26.1 and because $\xrightarrow{h}_{\beta_v} \subseteq \rightarrow_{\beta_v}$, therefore M is β_v -potentially valuable.
2. Since M is v -solvable, there exist variables x_1, \dots, x_m and terms N_1, \dots, N_n (for some $n, m \geq 0$) such that $(\lambda x_1 \dots x_m.M)N_1 \dots N_n \rightarrow_v^* I$; then, there exists $V \in \Lambda_v$ such that $(\lambda x_1 \dots x_m.M)N_1 \dots N_n \rightarrow_{\beta_v}^* V \xrightarrow{int}_{\beta_v}^* I$ by Corollary 26.1 and because $\xrightarrow{h}_{\beta_v} \subseteq \rightarrow_{\beta_v}$. According to Remark 12.5, $V = \lambda x.N$ for some $N \in \Lambda$ such that $N \rightarrow_v^* x$. By Corollary 26.1, there is $V' \in \Lambda_v$ such that $N \xrightarrow{h}_{\beta_v}^* V' \xrightarrow{int}_{\beta_v}^* x$, hence $V' = x$ by Remark 12.5 again. Since $\xrightarrow{h}_{\beta_v} \subseteq \rightarrow_{\beta_v}$, $N \rightarrow_{\beta_v}^* x$ and thus $V = \lambda x.N \rightarrow_{\beta_v}^* I$, therefore M is β_v -solvable. \blacktriangleleft

So, due to Theorem 31, the semantic (via a relational model) and operational (via two sub-reductions of \rightarrow_v) characterization of v -potential valuability and v -solvability given in [3, Theorems 24-25] is also a semantic and operational characterization of β_v -potential valuability and β_v -solvability. The difference is that in λ_v^σ these notions can be studied operationally inside the calculus, while it has been proved in [13, 14] that the β_v -reduction is too weak to characterize them: an operational characterization of β_v -potential valuability and β_v -solvability cannot be given inside λ_v . Hence, λ_v^σ is a useful, conservative and “complete” tool for studying semantic properties of λ_v .

6 Conclusions

In this paper we have proved a standardization theorem for the λ_v^σ -calculus introduced in [3]. The used technique is a notion of parallel reduction. Let us recall that parallel reduction in λ -calculus has been defined by Tait and Martin-Löf in order to prove confluence of the β -reduction, without referring to the difficult notion of residuals. Takahashi in [17] has simplified this technique and showed that it can be successfully applied to standardization. We would like to remark that our parallel reduction cannot be used to prove confluence of \rightarrow_v . Indeed, take $M = (\lambda x.L)((\lambda y.N)((\lambda z.N')N''))L'$, $M_1 = (\lambda x.LL')((\lambda y.N)((\lambda z.N')N''))$ and $M_2 = ((\lambda y.(\lambda x.L)N)((\lambda z.N')N''))L'$: then $M \Rightarrow M_1$ and $M \Rightarrow M_2$ but there is no term M' such that $M_1 \Rightarrow M'$ and $M_2 \Rightarrow M'$. To sum up, \Rightarrow does not enjoy the Diamond Property.

The standardization result allows us to formally verify the correctness of λ_v^σ with respect to the semantics of λ_v , so we can use λ_v^σ as a tool for studying properties of λ_v . This is a remarkable result: in fact some properties, like potential valuability and solvability, cannot be characterized in λ_v by means of β_v -reduction (as proved in [13, 14]), but they have a natural operational characterization in λ_v^σ (via two sub-reductions of \rightarrow_v).

We plan to continue to explore the call-by-value computation, using λ_v^σ . As a first step, we would like to revisit and improve the Separability Theorem given in [11] for λ_v . Still the issue is more complex than in the call-by-name, indeed in ordinary λ -calculus different $\beta\eta$ -normal forms can be separated (by the Böhm Theorem), while in λ_v there are different normal forms that cannot be separated, but which are only semi-separable (e.g. I and $\lambda z.(\lambda u.z)(zz)$). We hope to completely characterize separable and semi-separable normal forms in λ_v^σ . This should be a first step aimed to define a semantically meaningful notion of approximants. Then, we should be able to provide a new insight on the denotational analysis of the call-by-value, maybe overcoming limitations as that of the absence of fully abstract filter models [14, Theorem 12.1.25]. Last but not least, an unexplored but challenging

research direction is the use of commutation rules to improve the call-by-value evaluation. We do not have concrete evidence supporting such possibility, but since λ_v^σ is strongly related to the calculi presented in [7, 1], which are endowed with explicit substitutions, we are confident that a sharp use of commutations can have a relevant impact in the evaluation.

References

- 1 Beniamino Accattoli and Luca Paolini. Call-by-Value Solvability, Revisited. In Tom Schrijvers and Peter Thiemann, editors, *Functional and Logic Programming*, volume 7294 of *Lecture Notes in Computer Science*, pages 4–16. Springer-Verlag, 2012.
- 2 Henk Barendregt. *The Lambda Calculus: Its Syntax and Semantics*, volume 103 of *Studies in logic and the foundation of mathematics*. North Holland, 1984.
- 3 Alberto Carraro and Giulio Guerrieri. A Semantical and Operational Account of Call-by-Value Solvability. In Anca Muscholl, editor, *Foundations of Software Science and Computation Structures*, volume 8412 of *Lecture Notes in Computer Science*, pages 103–118. Springer-Verlag, 2014.
- 4 Karl Crary. A Simple Proof of Call-by-Value Standardization. Technical Report CMU-CS-09-137, Carnegie Mellon University, 2009.
- 5 Haskell B. Curry and Robert Feys. *Combinatory Logic*, volume 1. North Holland, 1958.
- 6 Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50(1):1–102, 1987.
- 7 Hugo Herbelin and Stéphane Zimmermann. An Operational Account of Call-by-Value Minimal and Classical lambda-Calculus in “Natural Deduction” Form. In Pierre-Louis Curien, editor, *Typed Lambda Calculi and Applications*, volume 5608 of *Lecture Notes in Computer Science*, pages 142–156. Springer-Verlag, 2009.
- 8 Roger Hindley. Standard and normal reductions. *Transactions of the American Mathematical Society*, pages 253–271, 1978.
- 9 John Maraist, Martin Odersky, David N. Turner, and Philip Wadler. Call-by-name, call-by-value, call-by-need and the linear lambda calculus. *Theoretical Computer Science*, 228(1–2):175–210, 1999.
- 10 Eugenio Moggi. Computational Lambda-Calculus and Monads. In *Logic in Computer Science*, pages 14–23. IEEE Computer Society, 1989.
- 11 Luca Paolini. Call-by-Value Separability and Computability. In Antonio Restivo, Simona Ronchi Della Rocca, and Luca Roversi, editors, *Italian Conference in Theoretical Computer Science*, volume 2202 of *Lecture Notes in Computer Science*, pages 74–89. Springer-Verlag, 2002.
- 12 Luca Paolini and Simona Ronchi Della Rocca. Parametric parameter passing lambda-calculus. *Information and Computation*, 189(1):87–106, 2004.
- 13 Luca Paolini and Simona Ronchi Della Rocca. Call-by-value Solvability. *Theoretical Informatics and Applications*, 33(6):507–534, 1999. RAIRO Series, EDP-Sciences.
- 14 Luca Paolini and Simona Ronchi Della Rocca. *The Parametric λ -Calculus: a Metamodel for Computation*. Texts in Theoretical Computer Science: An EATCS Series. Springer-Verlag, 2004.
- 15 Gordon D. Plotkin. Call-by-name, call-by-value and the lambda-calculus. *Theoretical Computer Science*, 1(2):125–159, 1975.
- 16 Amr Sabry and Matthias Felleisen. Reasoning about programs in continuation-passing style. *Lisp and symbolic computation*, 6(3-4):289–360, 1993.
- 17 Masako Takahashi. Parallel reductions in lambda-calculus. *Information and Computation*, 118(1):120–127, 1995.

Wild ω -Categories for the Homotopy Hypothesis in Type Theory

André Hirschowitz¹, Tom Hirschowitz², and Nicolas Tabareau³

- 1 CNRS, Université de Nice Sophia-Antipolis, France
- 2 CNRS, Université Savoie Mont-Blanc, France
- 3 Inria, Nantes, France

Abstract

In classical homotopy theory, the *homotopy hypothesis* asserts that the fundamental ω -groupoid construction induces an equivalence between topological spaces and weak ω -groupoids. In the light of Voevodsky’s *univalent foundations* program, which puts forward an interpretation of types as topological spaces, we consider the question of transposing the homotopy hypothesis to type theory. Indeed such a transposition could stand as a new approach to specifying higher inductive types. Since the formalisation of general weak ω -groupoids in type theory is a difficult task, we only take a first step towards this goal, which consists in exploring a shortcut through *strict* ω -categories.

The first outcome is a satisfactory type-theoretic notion of strict ω -category, which has hsets of cells in all dimensions. For this notion, defining the ‘fundamental strict ω -category’ of a type seems out of reach. The second outcome is an ‘incoherently strict’ notion of type-theoretic ω -category, which admits arbitrary types of cells in all dimensions. These are the ‘wild’ ω -categories of the title. They allow the definition of a ‘fundamental wild ω -category’ map, which leads to our (partial) homotopy hypothesis for type theory (stating an adjunction, not an equivalence).

All of our results have been formalised in the Coq proof assistant. Our formalisation makes systematic use of the machinery of coinductive types.

1998 ACM Subject Classification F.4.1 Mathematical logic

Keywords and phrases Homotopy Type theory, Omega-categories, Coinduction, Homotopy hypothesis

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.226

1 Introduction

Martin-Löf type theory offers an alternative to the classical set-theoretic approach to mathematics. The univalent foundations program [23] advocates understanding sets as a very special kind of types (so-called *0-types*, or *hsets*). Conversely, types should be understood in set-theoretic terms as some kind of topological spaces, hsets corresponding to those whose connected components are contractible. And, indeed, types have been interpreted as ω -groupoids of various flavours, most notably simplicial [13] and globular [3, 16, 24]. Now, on the set-theoretic side, the *homotopy hypothesis* states that topological spaces and ω -groupoids are equivalent [22]. This work addresses the question of coining a type-theoretic counterpart of this homotopy hypothesis.

We here propose a preliminary and partial version which barely expresses that the ‘fundamental ω -groupoid’ functor has a kind of left adjoint (instead of being an equivalence). For our proposal, it is sufficient to choose an appropriate target category \mathbb{T} of ‘ ω -groupoids’, together with an appropriate ‘fundamental ω -groupoid’ functor $\pi: \text{Type} \rightarrow \mathbb{T}$. Indeed, we may then express our partial homotopy hypothesis as follows:



© André Hirschowitz, Tom Hirschowitz, and Nicolas Tabareau;
licensed under Creative Commons License CC-BY

13th International Conference on Typed Lambda Calculi and Applications (TLCA’15).

Editor: Thorsten Altenkirch; pp. 226–240



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

► **Hypothesis 1.** *There exist $L: \mathbb{T} \rightarrow \text{Type}$ and $\eta: \forall C, \mathbb{T}(C, \pi(L(C)))$ such that for each $C: \mathbb{T}$ and $T: \text{Type}$, the map*

$$\eta(C)^*: \mathbb{T}(\pi(L(C)), \pi(T)) \rightarrow \mathbb{T}(C, \pi(T))$$

given by precomposition with $\eta(C)$ is an equivalence.

The idea of a ‘fundamental ω -category’ map goes back to van den Berg/Garner and Lumsdaine [3, 16]. There, the ‘fundamental’ ω -categories extracted from types are definitely weak, because expressions like $f \circ (g \circ h)$ and $(f \circ g) \circ h$ (for suitable cells f, g, h and higher categorical composition \circ) may differ definitionally. However, they are always equal propositionally, so when reasoning inside type theory as we do, the distinction becomes invisible. Based on this observation, we here explore the shortcut of a ‘fundamental *strict* ω -category’ map.

Thus, our first task is to transpose to type theory the classical notion of strict ω -category. This is the subject of Section 3 (after a brief recap on set-theoretic strict ω -categories in Section 2). There we face a crucial choice concerning the type of cells in all dimensions. If we take these types to be hsets, we get an honest notion of *strict* ω -category. While if we allow these types to be arbitrary, we get an ‘incoherently strict’ notion, which we call *wild*. In wild ω -categories, some coherence diagrams usually showing up in definitions of *weak* ω -categories make sense, but are not required to commute – even weakly, hence the name ‘wild’.

Now the crucial point is that we are able to define a ‘fundamental wild ω -category’ map, while a strict one seems out of reach – a difficulty previously observed by Altenkirch et al. [2] in a similar context.

Let us mention another crucial choice faced when defining both notions. Indeed, we have to express equations as commuting diagrams of morphisms between globular types. Because such morphisms form a coinductive type, there are two standard choices [8] for their equality (identity types and bisimilarity). However, we weren’t able to define our ‘fundamental ω -category’ map using either of them, so we work with a third, coarser one (Definition 8). (Of course, using identity types or bisimilarity yields other, perfectly sensible definitions.)

Altogether this yields in Section 5 a first version of our hypothesis with \mathbb{T} the type for wild ω -categories. Our hypothesis essentially asserts that new types may be constructed from the homotopical information carried by any wild ω -category. This is clearly akin to the assumption of existence for *higher inductive types* [23], as well as to the Rezk completion for precategories in [1] and we briefly discuss the relationship.

In Section 6, we get back to ω -groupoids, as opposed to ω -categories. We briefly discuss problems and solutions concerning the definition of wild ω -groupoids and a perhaps more primitive formulation of our homotopy hypothesis involving them. We finally conclude and sketch further directions in Section 7.

A note on the formalisation

This paper presents informally a mathematical development based not on set theory but on Martin-Löf type theory enriched with coinduction. All our definitions and statements have been formalised in the Coq proof assistant [21]. The formalisation is available as [9]. The code is composed of 2750 lines of definitions/theorems and less than 800 lines of proofs (as given by `coqwc`). Definitions and theorems constitute 75% of the formalisation, which is a lot. This is because `coqwc` only counts as proof what is coded using the tactic language, whereas most of our proofs are coded directly in Gallina. The reason for this is that the

computational content of many of our proofs turned out to be crucial for their usability at later stages. This phenomenon has also been observed, e.g., of the Coq/HoTT library (<https://github.com/HoTT/HoTT>).

2 Set-theoretic strict ω -categories

Before presenting our definition of wild and strict ω -categories in type theory, we briefly sketch the traditional set-theoretic definition of strict ω -categories. We refer the reader to Lafont et al. [15] for further detail.

The base for strict ω -categories is the notion of globular sets, which are generally defined as presheaves over the so-called *globular category* **Glob**:

$$[0] \begin{array}{c} \xrightarrow{s_0} \\ \xleftarrow{t_0} \end{array} [1] \begin{array}{c} \xrightarrow{s_1} \\ \xleftarrow{t_1} \end{array} [2] \quad \dots \quad [n] \begin{array}{c} \xrightarrow{s_n} \\ \xleftarrow{t_n} \end{array} [n+1] \quad \dots$$

where for all n , $s_{n+1} \circ s_n = t_{n+1} \circ s_n$ and $s_{n+1} \circ t_n = t_{n+1} \circ t_n$.

► **Notation 1.** We denote by s_n^p the composite $s_{p-1} \circ \dots \circ s_n$, which returns the n -dimensional source of a p -cell. Similarly, we use t_n^p .

The idea is just that a globular set X has *objects*, the elements of $X[0]$, *1-cells* between them, and so on, each $(n + 1)$ -cell having *parallel* n -cells as source and target. E.g., a typical 2-cell looks like this:

$$x \begin{array}{c} \xrightarrow{\alpha \cdot s_1} \\ \Downarrow \alpha \\ \xrightarrow{\alpha \cdot t_1} \end{array} y,$$

where we use the standard shorthand notation, e.g., $\alpha \cdot s_1$ for $X(s_1)(\alpha)$. In the picture, $x = \alpha \cdot t_1 \cdot s_0 = \alpha \cdot s_1 \cdot s_0$ and $y = \alpha \cdot t_1 \cdot t_0 = \alpha \cdot s_1 \cdot t_0$.

We first concentrate on the data for composition and the so-called *interchange law*. A strict ω -category is a globular set X equipped (among other data), for all $n < p$, with a partial, binary *composition* operation on p -cells, defined on a pair (β, α) when the iterated source $\beta \cdot s_n^p$ of β matches the iterated target $\alpha \cdot t_n^p$ (recall Notation 1). The result is denoted by $\beta \circ_n \alpha$. Each such composition operation is required to be associative on the nose.

The source and target of such a composition are given by obvious globular intuition, generally not even spelled out. When $p = n + 1$, composition is like categorical composition, i.e., we have $(\beta \circ_n \alpha) \cdot s_n = \alpha \cdot s_n$ and similarly $(\beta \circ_n \alpha) \cdot t_n = \beta \cdot t_n$. When $p > n + 1$, composition is more like horizontal composition of 2-cells in a 2-category, as in

$$x \begin{array}{c} \xrightarrow{\alpha \cdot s_1} \\ \Downarrow \alpha \\ \xrightarrow{\alpha \cdot t_1} \end{array} y \begin{array}{c} \xrightarrow{\beta \cdot s_1} \\ \Downarrow \beta \\ \xrightarrow{\beta \cdot t_1} \end{array} z,$$

so we have


$$(\beta \circ_n \alpha) \cdot s_{p-1} = (\beta \cdot s_{p-1}) \circ_n (\alpha \cdot s_{p-1})$$

and similarly for t_{p-1} .

The crux of the definition of strict ω -category is the *interchange law*, which generalises the perhaps more well-known 2-categorical interchange law. It says that whenever $n < p < q$ and $\alpha, \alpha', \beta, \beta'$ are adequately composable q -cells, we have

$$(\beta' \circ_p \beta) \circ_n (\alpha' \circ_p \alpha) = (\beta' \circ_n \alpha') \circ_p (\beta \circ_n \alpha).$$

Graphically, for $n = 0$, $p = 1$, $q = 2$, both ways of composing the diagram below coincide:



Finally, we require identities $id_x: x \rightarrow x$ for all n -cells $x \in X[n]$. Let id_x^p denote the iteration $id_{id_{id_{\dots id_x}}}$, p times (with $id_x^0 = x$). Identities should satisfy

$$id_x^p \circ_n \alpha = \alpha \quad \beta \circ_n id_x^p = \beta \quad id_{\beta \circ_n \alpha} = id_\beta \circ_n id_\alpha, \tag{1}$$

for all $x \in X[n]$, $p \geq 1$, α , and β such that $\alpha \cdot t_n^{n+p} = x$ and $\beta \cdot s_n^{n+p} = x$.

While this presentation is perfectly sensible when working in set theory, it is less convenient in type theory because many properties of the structure are imposed *by equations* rather than obtained *by construction*. For instance, compatibility between source and target maps in globular sets relies on equations. Using equations introduces a lot of non-definitional equalities that are very hard to deal with. On the contrary, we will see in the next section that our use of coinductive definitions allows us to define structures more computationally, thus avoiding the common pitfall of using equations.

3 Wild and strict ω -categories

In this section, we present our definition of wild and strict ω -categories. Our heavy use of coinduction is inspired by Cheng and Leinster’s account [7] of Trimble’s ω -categories, as well as by Lafont et al.’s presentation [15] of, e.g., weak equivalences. We start with globular types, a type-theoretic counterpart of globular sets. Essentially, a wild ω -category is a globular type equipped with identities and compositions in all dimensions, satisfying unitality and associativity axioms, and such that, at each dimension, composition preserves higher-dimensional identities and compositions. A strict ω -category is a wild ω -category whose underlying globular type is a globular hset, i.e., consists of hsets in all dimensions. The main goal of the present section is to give a precise meaning to the previous sketch of definition. Indeed, this is not completely straightforward, and we will explicitly discuss our design choices.

3.1 Globular types

The coinductive presentation of globular sets, which we’ll here call globular types, is well-known and extremely simple:

► **Definition 1.** A *globular type* X consists of a type $|X|$, plus for all $x, y \in |X|$, a globular type $X[x, y]$.

Many of our definitions will, like the previous one, be coinductive or corecursive, and we will systematically omit to mention this feature. For instance, a morphism $X \dashrightarrow Y$ of globular types consists of

- a function $|f|: |X| \rightarrow |Y|$ between object types, and
- for all $x, x' \in |X|$, a morphism $f_{x,x'}: X[x, x'] \dashrightarrow Y[|f|(x), |f|(x')]$ of globular types.

Of course, morphisms of globular types compose.

We also need a definition of the cartesian product of globular types:

► **Definition 2.** The *product* $X \times Y$ of two globular types X and Y is defined by

- $|X \times Y| = |X| \times |Y|$ and
- for all $x, x' \in |X|$ and $y, y' \in |Y|$, $(X \times Y)[(x, y), (x', y')] = X[x, x'] \times Y[y, y']$.

3.2 Data for composition and identities

Let us now start our definition of wild ω -categories.

In the following definition of composition for globular types, we choose to tie together infinitely many elementary compositions into a morphism of globular types:

- **Definition 3.** To equip a globular type X with ω -categorical composition is to give:
- for all $x, y, z \in |X|$, a morphism $\text{comp}_X(x, y, z): X[x, y] \times X[y, z] \rightarrow X[x, z]$ of globular types,
 - and, for all $x, y \in |X|$, ω -categorical composition on $X[x, y]$.

In the terminology of Section 2, the function $|\text{comp}_X|$ between object types represents composition for $p = n + 1$, whereas compositions for $p > n + 1$ are packed in higher-dimensional components $(\text{comp}_X)_{(f,g),(f',g')}$. Similarly:

- **Definition 4.** To equip a globular type X with ω -categorical identities is to give:
- for all $x \in |X|$, an element $\text{id}_x^X \in |X[x, x]|$, and
 - for all $x, y \in |X|$, ω -categorical identities on $X[x, y]$.

► **Definition 5.** An ω -precategory is a globular type equipped with ω -categorical composition and identities.

► **Remark.** Our use of ‘precategory’ conflicts with [1], where it is used for ‘not-necessarily-univalent’ categories. However, it coincides with Cheng’s [6].

Most notions defined so far on globular types lift to ω -precategories, namely the object type $|X|$, the hom- ω -precategory $X[x, y]$, and cartesian product $X \times Y$. We lift notations accordingly. However, lifting the notion of morphism of globular types requires some work, which we now do, yielding the notion of ω -functor.

3.3 Omega-functors

It is not straightforward to define what it means for a morphism $F: X \rightarrow Y$ of globular types between ω -precategories to preserve composition.

Let us treat the object level first, and consider any $x, x', x'' \in |X|$. The idea is that F preserves composition at x, x', x'' iff the diagram

$$\begin{array}{ccc}
 X[x, x'] \times X[x', x''] & \xrightarrow{\text{comp}_X(x, x', x'')} & X[x, x''] \\
 \downarrow F_{x, x'} \times F_{x', x''} & & \downarrow F_{x, x''} \\
 Y[|F|x, |F|x'] \times Y[|F|x', |F|x''] & \xrightarrow{\text{comp}_Y(|F|x, |F|x', |F|x'')} & Y[|F|x, |F|x'']
 \end{array} \tag{2}$$

commutes.

To give a meaning to this commutation, we need to choose a notion of equality of globular morphisms. Because globular morphisms form a coinductive type, identity types are expected to be too fine as a notion of equality between them [8]. The standard choice for this is bisimilarity:

► **Definition 6.** Two globular morphisms $F, G: X \rightarrow Y$ are *bisimilar* iff

- for all $x \in |X|$, $|f|(x) = |g|(x)$, and
- for all $x, x' \in |X|$, $f_{x, x'}$ and $(p_x)_*((p_{x'})_*(g_{x, x'}))$ are bisimilar,

where $p_*(-)$ denotes transport [23] along p and p_x is the given equality $|f|(x) = |g|(x)$.

Transport is necessary here in order to compare $f_{x,x'}: X[x, x'] \rightarrow Y[|f|(x), |f|(x')]$ and $g_{x,x'}: X[x, x'] \rightarrow Y[|g|(x), |g|(x')]$.

Bisimilarity would be a very reasonable notion of equality between globular morphisms. However, we were not able to construct the ‘fundamental wild ω -category’ of a type for the resulting notion of wild ω -category. So we here decide to use the following extensional and inductive definition (based on globular cells).

► **Definition 7.** We define the type of globular n -cells of an ω -precategory A , noted $cell_n A$, by induction on the natural number n as

- $cell_0 A := |A|$, and
- $cell_{n+1} A := \sum_{a,a':|A|} cell_n(A[a, a'])$.

Given a globular morphism $F: A \rightarrow B$ and $c: cell_n A$, the globular cell Fc of B obtained by applying F to c can be defined inductively by

- $|F|c$ when $n = 0$, and
- $(|F|a, |F|a', F_{a,a'}c')$ when $c = (a, a', c')$.

► **Definition 8.** Two globular morphisms $F, G: A \rightarrow B$ are *extensionally equal* iff

$$\text{for all } n \in \mathbb{N} \text{ and } c \in cell_n A, Fc = Gc.$$

We may now settle the following (obviously proof-relevant) definition:

► **Definition 9.** A morphism $F: X \rightarrow Y$ of globular types between ω -precategories *preserves composition* iff

- for all $x, x', x'' \in |X|$, the square (2) commutes extensionally, and
- for all $x, x' \in |X|$, $F_{x,x'}$ preserves composition.

Preservation of composition will be used below in Definition 12 to express the interchange law. E.g., consider any ω -precategory X , and $x, y, z \in |X|$. Viewing $\mathbf{comp}_X(x, y, z)$ as a morphism of globular types between ω -precategories, saying that it preserves composition amounts to stating the interchange law of Section 2, specialised to $n = 0$. For instance, on objects, it means that, taking X in (2) to be $X[x, y] \times X[y, z]$, for all $f, f', f'' \in X[x, y]$, $g, g', g'' \in X[y, z]$, $a \in X[x, y][f, f']$, $a' \in X[x, y][f', f'']$, $b \in X[y, z][g, g']$, and $b' \in X[y, z][g', g'']$, we have $(b' \bullet b) \circ (a' \bullet a) = (b' \circ a') \bullet (b \circ a)$ (using some hopefully clear notation).

We may treat preservation of identities in a similar way:

► **Definition 10.** A morphism $F: X \rightarrow Y$ of globular types between ω -precategories *preserves identities* iff

- for all $x \in |X|$, $|F_{x,x}|(id_x^X) = id_{|F|(x)}^Y$, and
- for all $x, y \in |X|$, $F_{x,y}$ preserves identities.

This will again be used in Definition 12 to enforce the third law of (1). Indeed, for any ω -precategory X , and $x, y, z \in |X|$, saying that $\mathbf{comp}_X(x, y, z)$ preserves identities entails, e.g., that the identity on any pair $(f, g) \in |X[x, y] \times X[y, z]|$ should be mapped by

$$X[x, y][f, f] \times X[y, z][g, g] \xrightarrow{(\mathbf{comp}_X(x, y, z))_{(f, g), (f, g)}} X[x, z][g \circ f, g \circ f]$$

to the identity on $g \circ f$ (abbreviating $|\mathbf{comp}_X(x, y, z)|(f, g)$ to $g \circ f$).

► **Definition 11.** A morphism of globular types between ω -precategories is an *ω -functor* iff it preserves composition and identities.

3.4 Wild ω -categories

It is now routine to extend the previous techniques to define *associativity* as extensional commutation in all dimensions of all diagrams

$$\begin{array}{ccc}
 X[x, y] \times X[y, z] \times X[z, t] & \xrightarrow{\text{comp}_X(x, y, z) \times X[z, t]} & X[x, z] \times X[z, t] \\
 \downarrow \text{comp}_X(y, z, t) & & \downarrow \text{comp}_X(x, z, t) \\
 X[x, y] \times X[y, t] & \xrightarrow{\text{comp}_X(x, y, t)} & X[x, t]
 \end{array}$$

and *unitality* as extensional commutation in all dimensions of all diagrams

$$\begin{array}{ccc}
 X[x, y] & \xrightarrow{\langle \ulcorner id_x \urcorner \circ !, X[x, y] \rangle} & X[x, x] \times X[x, y] \\
 \downarrow \langle X[x, y], \ulcorner id_y \urcorner \circ ! \rangle & \searrow & \downarrow \text{comp}_X(x, x, y) \\
 X[x, y] \times X[y, y] & \xrightarrow{\text{comp}_X(x, y, y)} & X[x, y]
 \end{array}$$

where $! : A \rightarrow 1$ denotes the unique morphism to the terminal globular type with 1 at all stages (for all A), and $\ulcorner id_x \urcorner : 1 \rightarrow X[x, x]$ maps the unique element of $|1|$ to id_x , the unique endo 1-cell over it to id_{id_x} , and so on.

We may at last define the type $\omega\text{-wCat}$ of wild ω -categories:

► **Definition 12.** A *wild ω -category* is an ω -precategory, satisfying associativity and unitality, whose compositions are ω -functors in all dimensions. Morphisms $\omega\text{-wCat}(C, D)$ between wild ω -categories C and D are simply ω -functors between the underlying ω -precategories.

The complete formal definition of wild ω -categories is given in the file `omega_categories.v` [9].

3.5 Strict ω -categories

The definition of wild ω -categories is not satisfactory as a type-theoretic account of strict ω -categories. As a matter of fact, wild ω -categories are not even weak ω -categories. Indeed, they appear to lack some coherence conditions. For instance, for any wild ω -category X , $f \in |X[x, y]|$ and $g \in |X[y, z]|$, there are two proofs of $g \circ (id_y \circ f) = g \circ f$ (the less trivial one goes to $(g \circ id_y) \circ f$ and then simplifies). These proofs induce by transport two 2-cells, say l and r . In weak higher categories, one imposes that l and r are related by a ‘coherence’ 3-cell. This is not the case in our wild ω -categories, which may thus be viewed as ‘incoherently’ weak ω -categories.

One perhaps reassuring perspective is that wild ω -categories can be restricted to ω -categories where all higher coherences are trivially satisfied. This is the case when the involved types are all hsets, which leads to:

► **Definition 13.** A *strict ω -category* is a wild ω -category X such that $|X|$ is an hset and for all $x, y \in |X|$, $X[x, y]$ is a strict ω -category. We call $\omega\text{-sCat}$ the type of strict ω -categories.

Strict ω -categories are intuitively close to set-theoretic strict ω -categories, and, as suggested by a referee, we expect their interpretation in standard models such as the simplicial model [13] to coincide with set-theoretic ω -categories. However, as we have seen, our definition relies on extensional equality (Definition 8), where identity types or bisimilarity could have been used. We thus in fact have three notions of strict ω -categories, and it is not entirely clear that all three are interpreted in the simplicial model as set-theoretic strict ω -categories.

As previously observed in a similar context [2], it seems impossible to define any satisfactory ‘fundamental strict ω -category’. That is why we devote the rest of the paper to the definition of the ‘fundamental wild ω -category’ and use it to propose (partial) homotopy hypotheses.

4 Fundamental wild ω -category

We have defined wild ω -categories, and now turn to the definition of the ‘fundamental wild ω -category’ map $\pi: \text{Type} \rightarrow \omega\text{-wCat}$. This is essentially an internal variant of van den Berg and Garner’s [3] and Lumsdaine’s [16] constructions, with wild ω -categories instead of weak ω -categories. For any type T , we first easily define the globular type underlying $\pi(T)$. We then explain the definition of composition, which is a good example of how we had to generalise some of our statements in order to be able to tie the coinduction loop. We refer the reader to the formalisation for complete definitions and proofs. The part on composition is technical and may safely be skipped: the sequel only makes use of Theorem 16.

We start by defining the globular type underlying the fundamental wild ω -category.

► **Definition 14.** For any type T , the underlying globular type of $\pi(T)$ (still denoted by $\pi(T)$ by abuse of notation) has T itself as its type of objects, and $(\pi(T))[x, y] = \pi(x = y)$ for all $x, y \in T$.

We now turn to defining composition on $\pi(T)$. The definition of composition seems to introduce a lot of choices. Indeed, let us start by fixing $x, y, z \in T$ and look at what $\pi(x = y) \times \pi(y = z) \rightarrow \pi(x = z)$ is on objects (i.e., on 1-cells in $\pi(T)$). The obvious choice is concatenation of equality proofs, which we denote by $(a: x = y), (b: y = z) \mapsto (a \cdot b)$ [23]. But actually, here we need to choose between two different definitions of concatenation, depending on whether a or b is eliminated first. Fortunately, it is well known that both definitions are equal, so the choice does not really matter.

Actually, this situation occurs at every dimension: the definition of composition is not unique, but all potential candidates are equal. This claim is justified by the work of Lumsdaine [16], where he constructs the operad P_{ML}^{Id} of all definable composition laws over a (generic) type and shows that this operad is contractible. Contractibility means that all possible choices of composition are equal. The difficulty is to show that our particular choice of composition gives rise to a wild ω -category. For lack of space, we will only sketch the definition of our compositions, referring to the formalisation [9] for details.

Let us start with the definition of $\pi(x = y) \times \pi(y = z) \rightarrow \pi(x = z)$ in low dimensions, for $x, y, z \in T$. On objects, we have seen that the obvious choice is concatenation. On 1-cells, consider $a, a': x = y, b, b': y = z$, together with $e: a = a'$ and $f: b = b'$. How to compose e and f into a proof of $a \cdot b = a' \cdot b'$? We consider, for all types A, B, C and map $\varphi: A \rightarrow B \rightarrow C$, the obvious function $\text{ap2}_{A,B,C,\varphi}$ of type

$$\forall a, a' \in A, b, b' \in B, e \in (a = a'), f \in (b = b'), \varphi a b = \varphi a' b'.$$

Applying this with $A = (x = y), B = (y = z), C = (x = z)$, and $\varphi = \cdot$, we indeed get $\text{ap2 } a a' b b' e f$ of type $a \cdot b = a' \cdot b'$ (omitting the subscript of ap2 for readability). To now deal with 2-cells, considering $e': a = a', f': b = b', u: e = e',$ and $v: f = f'$, we again apply ap2 , with $A = (e = e'), B = (f = f'), C = (\varphi' e f = \varphi' e' f')$, with $\varphi' e f = \text{ap2 } a a' b b' e f$ for all e, f . In the next dimension, we’ll need a different φ'' with one more layer of ap2 .

It would be obvious how to formalise this process coinductively if it weren’t for the first level, where \cdot is used. The trick is thus to abstract over this. Here, things become slightly

more verbose, and we apologise to the reader: for all types A, B, C , functions $\varphi: A \rightarrow B \rightarrow C$, and elements $a, a' \in A$ and $b, b' \in B$, we coinductively define a morphism of globular types

$$\text{comp2}_{A,B,C,\varphi}(a, a', b, b'): \pi(a = a') \times \pi(b = b') \rightarrow \pi(\varphi a b = \varphi a' b')$$

(the ‘2’ refers to `ap2`) by

- mapping $e: a = a'$ and $f: b = b'$ to `ap2 φ e f: $\varphi a b = \varphi a' b'$` ,
- and then defining `($\text{comp2}_{A,B,C,\varphi}(a, a', b, b')$)(e,f),(e',f')` to be

$$\text{comp2}_{A[\text{ap2}],B[b,b'],C[\varphi a b, \varphi a' b'], \text{ap2 } \varphi}(e, e', f, f'). \quad (3)$$

This is well-defined, since `($\text{comp2}_{A,B,C,\varphi}(a, a', b, b')$)(e,f),(e',f')` should have type

$$(\pi(a = a') \times \pi(b = b'))[(e, f), (e', f')] \rightarrow \pi(\varphi a b = \varphi a' b')[\text{ap2 } \varphi e f, \text{ap2 } \varphi e' f'],$$

i.e.,

$$\pi(a = a')[e, e'] \times \pi(b = b')[f, f'] \rightarrow \pi(\varphi a b = \varphi a' b')[\text{ap2 } \varphi e f, \text{ap2 } \varphi e' f'],$$

or equivalently

$$\pi(e = e') \times \pi(f = f') \rightarrow \pi(\text{ap2 } \varphi e f = \text{ap2 } \varphi e' f'),$$

which is indeed the type of (3). It is now routine to define, for all types A and elements $a, a', a'' \in A$,

$$\text{hcomp}_A(a, a', a''): \pi(a = a') \times \pi(a' = a'') \rightarrow \pi(a = a'')$$

by

- mapping $e: a = a'$ and $f: a' = a''$ to $e \cdot f$, with
- `$\text{hcomp}_A(a, a', a'')$ (e,f),(e',f') = $\text{comp2}_{(a=a'),(a'=a''),(a=a''),(\lambda e.\lambda f.e \cdot f)}$ (e, e', f, f')`.

This is again well-defined, because the latter has type

$$\pi(e = e') \times \pi(f = f') \rightarrow \pi(e \cdot f = e' \cdot f'),$$

i.e.,

$$\pi(a = a')[e, e'] \times \pi(a' = a'')[f, f'] \rightarrow \pi(a = a'')[e \cdot f, e' \cdot f'],$$

as expected.

► **Definition 15.** This yields composition structure on $\pi(T)$, for any T :

- for all $x, y, z \in T$, we let `$\text{comp}_{\pi(T)}(x, y, z)$ = $\text{hcomp}_T(x, y, z)$` ;
- for all $x, y \in T$, we get composition structure on $\pi(x = y)$ by coinduction hypothesis.

One may similarly define identity structure and show:

► **Theorem 16.** *This ω -precategory structure makes $\pi(T)$ into a wild ω -category.*

Proof. Let us say a few words about the proof. It resorts to a high level of generalisation to tie the coinduction loop for each law of ω -categories. In particular, a more explicit coinductive definition of the interchange law is developed and proved equivalent to the compact version that composition preserves composition. This explicit interchange law, plus a proof that composition in all dimensions sends pairs of proofs by reflexivity to some proof by reflexivity, enables us to prove the interchange law by coinduction, using elimination of identity types, a.k.a. path induction. The other laws are dealt with similarly. The complete proof is given in the file `type_to_omega_cat.v` [9]. ◀

5 Partial Homotopy Hypothesis

The classical homotopy hypothesis states an equivalence between spaces and ω -groupoids. It can be formulated either at the level of so-called homotopy categories [12], or at the level of model categories [11] or even at the level of $(\infty, 1)$ -categories [22]. On the type-theoretic side, we propose our first partial homotopy hypothesis:

► **Hypothesis 1.** *There exist $L: \omega\text{-wCat} \rightarrow \text{Type}$ and $\eta: \forall C, \omega\text{-wCat}(C, \pi(L(C)))$ such that for each wild ω -category C and type T , the map*

$$\eta(C)^*: \omega\text{-wCat}(\pi(L(C)), \pi(T)) \rightarrow \omega\text{-wCat}(C, \pi(T))$$

given by precomposition with $\eta(C)$ is an equivalence.

We conjecture that Hypothesis 1 is consistent, and more precisely that it holds in the groupoid model [10]. Indeed, in this model, small types are small discrete groupoids and (small) wild ω -categories are the set-theoretic, strict ω -categories of Section 2. The fundamental wild ω -category $\pi(T)$ of any small discrete groupoid T is a discrete globular set, equipped with the only possible additional structure. Thus, morphisms from any strict ω -category C into $\pi(T)$ are just maps from $C[0]$ to T compatible with (the equivalence relation induced by) $C[1]$. Hence we can define $L(C)$ to be the quotient $C[0]/C[1]$, $\eta(C)$ being induced by the quotient map.

This indicates in particular that Hypothesis 1 is consistent with the Univalence Axiom. But it also shows that the asserted adjunction may be far from an equivalence. An easy way of strengthening our hypothesis is to require $\eta(C)$ to be a weak equivalence, in the sense of [15]:

► **Definition 17.** An ω -functor $F: C \rightarrow D$ is a *weak equivalence* iff

- for all $d \in |D|$, there exists $c \in |C|$ such that $|D|[F](c), d|$ is inhabited, and
- for all $c, c' \in |C|$, $F_{c,c'}$ is a weak equivalence.

Another possibility is to state a proper adjunction between L and π and ask both its unit and counit to be weak equivalences in the appropriate sense. We haven't yet investigated such strengthened hypotheses.

For now, let us modestly check how our hypothesis implies the existence of a type corresponding to the standard type-theoretic circle. For this purpose we introduce the wild ω -category S^1 as follows: it has a single object \star , and $S^1[\star, \star]$ is the discrete wild ω -category on \mathbb{N} -many objects with composition given by addition. Of course, we could have worked with \mathbb{Z} instead. Assuming our hypothesis, we prove the expected, non-dependent induction principle for $L(S^1)$, together with (propositional versions of) computational rules.

► **Theorem 18.** *There exists a term inhabiting the non-dependent recursion principle of the circle*

$$L(S^1)_{\text{rec}}: \forall T \in \text{Type}, b \in T, l \in (b = b), L(S^1) \rightarrow T,$$

satisfying (propositionally) the expected computational laws:

$$\begin{aligned} L(S^1)_{\beta, \star}: \forall T, b, l, L(S^1)_{\text{rec}} T b l (\eta \star) &= b, \\ L(S^1)_{\beta, 1}: \forall T, b, l, p_*((\text{ap } (L(S^1)_{\text{rec}} T b l) (\eta_{\star, \star} 1))) &= l, \end{aligned}$$

where $p_(-)$ denotes transport [23] along $p = L(S^1)_{\beta, \star} T b l$.*

Proof. Any ω -functor from S^1 is determined by the images of \star and $1 \in S^1[\star, \star]$. Using the inverse of the equivalence given by Hypothesis 1, this induces an essentially unique ω -functor $\pi(L(S^1)) \rightarrow \pi(T)$, from which we extract the underlying map $L(S^1) \rightarrow T$. Using the retraction part of the equivalence given by Hypothesis 1, we deduce the computational laws. The complete proof¹ is given in the file `homotopy_hypothesis.v` [9]. ◀

This suggests that Hypothesis 1 may imply existence of certain *higher inductive types* [23]. The basic idea of higher inductive types is to generate not only the elements of an inductive type, but also equality proofs between them, and so on. Lumsdaine and Shulman [17] propose a semantics for them (in particular) in $(\infty, 1)$ -toposes, using so-called strictly Reedy-functorial path objects. Closer to implementation, Sojakova [19] proposes an operational definition of higher inductive types as so-called homotopy-initial algebras. Both accounts fix a particular syntax for higher inductive types.

If, along the lines of Theorem 18, we could show that the partial homotopy hypothesis entails adequate induction principles, it could be understood as *specifying* higher inductive types, in a syntax-independent way. Of course, this would only be a definition from the internal point of view, i.e., the corresponding computational behaviour would not be accounted for.

So, given any candidate syntax for (or combinatorial description of) higher inductive types, this opens the option of describing ‘the corresponding’ wild ω -category, from which Hypothesis 1 would yield the desired type and reasoning principles.

► **Remark.** It is well-known that set-theoretic strict ω -categories cannot represent all homotopy types. E.g., they do not model the homotopy type of the 2-sphere [18]. Depending on the ambient type theory, wild ω -categories may be much more expressive than strict ω -categories. Nevertheless, we suspect that even in such cases, they may not adequately represent all types.

Let us consider a few example syntaxes.

To start with, Ahrens et al.’s categories [1] straightforwardly embed into wild ω -categories, and are enough to specify (a groupoidal version of) S^1 as above (but not S^2). Their categories may express ‘non-freeness’ properties of composition. E.g., we may consider the category obtained by quotienting S^1 under $1 + 1 = 1$, i.e., replace \mathbb{N} by booleans and addition with sup. Or similarly quotient under $1 + 1 = 0$, i.e., work with $\mathbb{Z}/2\mathbb{Z}$. Also, it seems plausible to extend their Yoneda-based *Rezk completion* procedure—which (in their terms) constructs a (univalent) category from a precategory—to a proof of our hypothesis for categories. Please note, however, that they work in homotopy type theory, and their construction uses univalence (because the category of sets needs to be univalent) and higher inductive types (through propositional truncation).

A different ‘syntax’ is offered by globular types themselves, and we may hope for a type-theoretic analogue of the standard adjunction computing the free strict ω -category associated to any globular set. In contrast this syntax does not allow to express non-freeness properties. E.g., we may express S^1 , but none of the quotients of S^1 evoked above. Also, we cannot express the higher inductive type for S^2 with one base point b and an equality proof on refl_b , because we cannot talk about identities. However, we can perfectly consider the globular set with two base points 0 and 1, two 1-cells $s, t: 0 \rightarrow 1$, and two 2-cells $s \rightarrow t$.

¹ The proof that the constructed ω -functor from S^1 actually preserves composition remains incomplete at the time of writing.

Finally, the most expressive such syntax would probably be offered by *computads* [20], a.k.a. *polygraphs* [5, 15]. A computad is essentially a graph, together with a set of 2-cells between parallel paths, together with a set of 3-cells between parallel paths of 2-cells, and so on. In particular, one may talk about identities and composition in all dimensions. And indeed, any strict ω -category is weakly equivalent to some strict ω -category which is free on a computad [15].

6 Towards ω -groupoids

In the fundamental wild ω -category $\pi(T)$ associated to any type T , all cells are invertible. We will see below that defining general ω -groupoids is not straightforward. Nevertheless, following Ahrens et al. [1], we have an economical, yet a bit restrictive definition:

- **Definition 19.** A wild ω -category X is a *univalent ω -groupoid* when for all $x, y \in |X|$,
 - the map $(x = y) \rightarrow |X[x, y]|$ (induced by transporting id_x) is an equivalence (of types) [23],
 - and $X[x, y]$ is a univalent ω -groupoid.

This univalence automatically entails existence of inverses in a very strong sense. In particular, for any wild ω -category X and $x, y, z \in |X|$, the diagram

$$\begin{array}{ccc} (x = y) \times (y = z) & \longrightarrow & (x = z) \\ \downarrow & & \downarrow \\ |X[x, y]| \times |X[y, z]| & \longrightarrow & |X[x, z]| \end{array}$$

(extensionally) commutes in **Type**, so that, if X is univalent, inverses in the ω -groupoidal sense have to be inverses in the identity type sense. The above definition has to be about *ω -groupoids*, as it implies that all cells are invertible. We prove:

- **Proposition 1.** *The ω -category $\pi(T)$ is a univalent ω -groupoid, for all types T .*
- **Proposition 2.** *For all univalent ω -groupoids G , there is a type T and an extensional equivalence $\pi(T) \simeq G$.*

By extensional equivalence we here mean an equivalence $e: |\pi(T)| \simeq |G|$ of types, such that for all $x, y \in T$ the map $\pi(x = y) \rightarrow |G[e\ x, e\ y]|$ induced by e is an extensional equivalence.

Proof. Take $T = |G|$. ◀

We conclude this section by considering perhaps more primitive versions of our hypothesis using some notion of *ω -groupoid* rather than ω -categories. Of course, univalent ω -groupoids are not interesting for this purpose, so we seek a definition of what it means for some (possibly higher) cell in a wild ω -category to be invertible. We first review possible notions of ω -groupoids and transpose a result of Cheng [6] showing that two of them are equivalent. We then get back to ω -groupoidal statements of our hypothesis.

Brown et al. [4] use *strict inverses*, i.e., for them an ω -groupoid is an ω -category in which for every n -cell $f: x \rightarrow y$ there is an n -cell $g: y \rightarrow x$ such that $g \circ_{n-1} f = id_x$ and $f \circ_{n-1} g = id_y$ (for $n > 0$). Others [20, 12, 18] consider *weak inverses*, in several apparently different ways. A recent preprint [14] shows that two such definitions coincide, namely those of Street [20], and Kapranov and Voevodsky [12]. The latter had previously been shown by Simpson [18] to be equivalent to an apparently stronger definition. Finally, Cheng [6] shows that these definitions are further equivalent to a seemingly weaker definition *when required of the whole ω -category*. A bit more precisely, Cheng defines the notion of a *dual*

to an n -cell in an ω -category, which is in general weaker than that of a *weak inverse* in the sense of [20, 12, 14]. But she shows that for a given ω -category, having all duals is equivalent to having all weak inverses. We now show how to recover this equivalence in our setting.

It is easy to define what it means for a wild ω -category to have all duals.

► **Definition 20.** In a wild ω -category X , for all $x, y \in |X|$ and $f \in |X[x, y]|$, a *dual* for f is a 1-cell $g \in |X[y, x]|$ such that there exist 2-cells inhabiting

$$|X[y, y][f \circ g, id_y]| \quad \text{and} \quad |X[x, x][g \circ f, id_x]|.$$

We then say that X *has all duals* iff for all $x, y \in |X|$ and $f \in |X[x, y]|$, f has a dual, and for all $x, y \in |X|$, $X[x, y]$ has all duals.

The notion of weak inverse is a bit harder. We follow the coinductive presentation of [15] (a definition considered ‘unsound’ by Cheng, but which Coq readily accepts!).

► **Definition 21.** We pose the following mutually coinductive definitions:

- two objects x and y of a wild ω -category X are *equivalent*, notation $x \sim y$, iff there exists a reversible 1-cell $f \in |X[x, y]|$;
- a 1-cell $f \in |X[x, y]|$ is *reversible* when it has a weak inverse;
- a *weak inverse* for a 1-cell $f \in |X[x, y]|$ is a 1-cell $g \in |X[y, x]|$ such that $g \circ f \sim id_x$ and $f \circ g \sim id_y$.

► **Definition 22.** A wild ω -category *has all weak inverses* when for all $x, y \in |X|$, any $f \in |X[x, y]|$ is reversible, and $X[x, y]$ has all weak inverses.

► **Proposition 3.** *Any wild ω -category has all duals iff it has all weak inverses.*

This allows us to comfortably state a first definition of ω -groupoid:

► **Definition 23.** A *wild ω -groupoid* is a wild ω -category with all duals.

One possible problem with this definition is that, ideally, being an ω -groupoid should be a mere property of ω -categories, i.e., for all X , the type ‘ X has all weak inverses’ should be a mere proposition [23]. The very definition of univalence [23] suggests that having all duals might not suffice and that one may have to resort to some sensible notion of ω -adjunction in this context, but there does not seem to be any commonly accepted such notion in the literature, and we leave the question open.

There is one possible solution if we wish to delve into HoTT — except for this paragraph, our whole development remains within Martin-Löf type theory with coinduction. Namely, we could *truncate* the naive definition above. Indeed, according to the naive definition above, a wild ω -groupoid is a wild ω -category, equipped with a choice of duals for all cells in all dimensions. So we may define *brutal ω -groupoids* to denote wild ω -categories for which there exists, in the mere propositional sense, duals for all cells in all dimensions. It would however be preferable to have a mere property without resorting to truncation, as is done in the standard treatment of the univalence axiom.

We conclude this section by stating a groupoidal variant of our hypothesis. Since we have several candidate definitions for general ω -groupoids, we state our hypothesis taking this as a parameter. So let $\omega\text{-Gpd}$ denote some type of ω -groupoids.

► **Hypothesis 2.** *There exist $L: \omega\text{-Gpd} \rightarrow \text{Type}$ and $\eta: \forall G \in \omega\text{-Gpd}, \omega\text{-wCat}(G, \pi(L(G)))$ such that for each ω -groupoid G and type T , the following map is an equivalence:*

$$\eta(G)^*: \omega\text{-wCat}(\pi(L(G)), \pi(T)) \rightarrow \omega\text{-wCat}(G, \pi(T)).$$

► **Remark.** We here use ω -functors as morphisms of ω -groupoids. An alternative would be to require morphisms of ω -groupoids to preserve weak inverses. If we could refine the type of inverses to a given morphism into a mere proposition, then both possibilities would coincide.

7 Conclusion and Future Work

We have defined wild and strict ω -categories, as well as wild and univalent ω -groupoids. We have constructed a ‘fundamental ω -groupoid’ map from types into all of these notions but strict ω -categories. We have stated a few variants of our partial homotopy hypothesis which postulates some correspondence between types and the given ω -dimensional structure. We have loosely related such hypotheses to the existence of higher inductive types and to the Rezk completion for one-dimensional categories.

The main remaining issue in our development is that we have used a rather coarse notion of equality between globular morphisms (Definition 8) in order to be able to define our ‘fundamental wild ω -category’ map. We wonder whether this could be done using a more standard notion like bisimilarity.

As explained in the introduction, this paper grew out of an attempt to use strict, as opposed to weak, ω -groupoids in the statement of a type-theoretic homotopy hypothesis. Beyond the issues raised by the definition of type-theoretic ω -groupoids (Section 6), it now seems likely that the lack of coherence of wild ω -categories disqualifies them for the full homotopy hypothesis.

Our main future challenge is thus clearly to propose a type-theoretic notion of weak ω -category allowing the definition of a ‘fundamental weak ω -category’ map. Assuming that we succeed in defining weak ω -categories, stating a full version of the homotopy hypothesis would first require us to define weak ω -groupoids properly. Thus, we expect the discussion of Section 6 about weak inverses to also be relevant in the weak case. Namely, we will need to investigate whether ‘having all duals’ is a mere property of weak ω -categories, and, if not, how to refine it into one. Finally, a full version of the homotopy hypothesis would essentially assert the existence of an infinite-dimensional generalisation of the Rezk completion. We wonder whether the construction of [1] could be adapted to this setting.

Acknowledgements. We thank the anonymous referees for insightful comments and suggestions. T. Hirschowitz and N. Tabareau acknowledge support from R  cr   ANR-11-BS02-0010.

References

- 1 B. Ahrens, K. Kapulkin, and M. Shulman. Univalent categories and the Rezk completion. *Mathematical Structures in Computer Science*, 2015.
- 2 T. Altenkirch, N. Li, and O. Ryp   ek. Some constructions on ω -groupoids. In *LFMTP ’14*, pages 4:1–4:8. ACM, 2014.
- 3 B. van den Berg and R. Garner. Types are weak ω -groupoids. *Proceedings of the London Mathematical Society*, 102(2):370–394, 2011.
- 4 R. Brown and P. Higgins. The equivalence of ∞ -groupoids and crossed complexes. *Cahiers de Topologie et G  om  trie Diff  rentielle Cat  goriques*, 22, 1981.
- 5 A. Burroni. Higher-dimensional word problems with applications to equational logic. *Theoretical Computer Science*, 115(1):43–62, 1993.
- 6 E. Cheng. An ω -category with all duals is an ω -groupoid. *Applied Categorical Structures*, 15:439–453, 2007.

- 7 E. Cheng and T. Leinster. Weak ω -categories via terminal coalgebras. *ArXiv e-prints*, December 2012.
- 8 A. Chlipala. *Certified Programming with Dependent Types*. MIT Press, 2014.
- 9 A. Hirschowitz, T. Hirschowitz, and N. Tabareau. Formalisation of strict omega-categories and the homotopy hypothesis in type theory, using coinduction. Coq code, https://github.com/tabareau/omega_categories, 2015.
- 10 M. Hofmann and T. Streicher. The groupoid model refutes uniqueness of identity proofs. In *LICS*, pages 208–212. IEEE Computer Society, 1994.
- 11 M. Hovey. *Model Categories*, volume 63 of *Mathematical Surveys and Monographs*. American Mathematical Society, 1999.
- 12 M. M. Kapranov and V. A. Voevodsky. ∞ -groupoids and homotopy types. *Cahiers de Topologie et Géométrie Différentielle Catégoriques*, 32(1):29–46, 1991.
- 13 C. Kapulkin, P. LeFanu Lumsdaine, and V. Voevodsky. The simplicial model of univalent foundations. *ArXiv e-prints*, November 2012.
- 14 S. Kasangian, G. Metere, and E. Vitale. Weak inverses for strict n -categories. Preprint, 2009.
- 15 Y. Lafont, F. Métayer, and K. Worytkiewicz. A folk model structure on omega-cat. *Advances in Mathematics*, 224(3):1183 – 1231, 2010.
- 16 P. LeFanu Lumsdaine. Weak omega-categories from intensional type theory. *Logical Methods in Computer Science*, 6(3), 2010.
- 17 P. LeFanu Lumsdaine and M. Shulman. Semantics of higher inductive types. Preprint, 2012.
- 18 C. Simpson. Homotopy types of strict 3-groupoids. *ArXiv e-prints*, 1998.
- 19 K. Sojakova. Higher inductive types as homotopy-initial algebras. In *POPL*, pages 31–42. ACM, 2015.
- 20 R. Street. The algebra of oriented simplexes. *Journal of Pure and Applied Algebra*, 49(3):283–335, 1987.
- 21 The Coq development team. *Coq 8.4 Reference Manual*. Inria, 2012.
- 22 The NLab. Homotopy hypothesis. Accessed April 2015.
- 23 The Univalent Foundations Program. *Homotopy Type Theory: Univalent Foundations of Mathematics*. <http://homotopytypetheory.org/book>, Institute for Advanced Study, 2013.
- 24 M. A. Warren. The strict ω -groupoid interpretation of type theory. *CRM Proceedings and Lecture Notes*, 53:291–340, 2011.

Multivariate Amortised Resource Analysis for Term Rewrite Systems*

Martin Hofmann¹ and Georg Moser²

- 1 Institute of Computer Science
LMU Munich, Germany
hofmann@ifi.lmu.de
- 2 Institute of Computer Science
University of Innsbruck, Austria
georg.moser@uibk.ac.at

Abstract

We study amortised resource analysis in the context of term rewrite systems. We introduce a novel amortised analysis based on the potential method. The method is represented in an inference system akin to a type system and gives rise to polynomial bounds on the innermost runtime complexity of the analysed rewrite system. The crucial feature of the inference system is the admittance of multivariate bounds in the context of arbitrary data structures in a completely uniform way. This extends our earlier univariate resource analysis of typed term rewrite systems and continues our program of applying automated amortised resource analysis to rewriting.

1998 ACM Subject Classification F.3.2 Program Analysis

Keywords and phrases program analysis, amortised analysis, term rewriting, multivariate bounds

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.241

1 Introduction

Amortised resource analysis was pioneered by Sleator and Tarjan who used it to analyse the performance of new data structures that sometimes need to perform costly operations that pay off later on, e.g. rebalancing operations on a search tree.

Briefly, one assigns to the participating datastructures a nonnegative real-value, the *potential* in an a priori arbitrary fashion. One then defines the amortised cost of an operation as its actual cost, e.g. runtime plus the difference in potential of all datastructures before and after the operations. In this way, the amortised cost of a costly operation may be small if it results in a big decrease of potential. On the other hand some cheap operations that increase the potential will be overcharged. In this way, one can “save money” now to pay for costly operations later. By a simple telescoping argument the sum of all amortised costs in a sequence of operations plus the potential of the initial input data structure is also an upper bound on the *actual* cost of that sequence. In this way, amortised analysis yields rigorous bounds on actual resource usage and not just approximate or average bounds. If the potential functions are chosen well then the amortised costs of operations are either constant or exhibit merely a very simple dependency on the maximum size of all intermediate results which considerably facilitates a compositional analysis: the costs of running composite expressions

* This research is partially supported by FWF (Austrian Science Fund) project P25781.



can be calculated as the sum of the individual costs. If costs are highly input-dependent, on the other hand, one must get bounds on the sizes or shapes of intermediate results which can be very difficult. This compositional aspect of amortised analysis makes it attractive for syntax-directed automation.

Of course, the crux of the matter is the choice of the correct potential functions. A simple concrete example is the implementation of a queue by two stacks, an in-tray and an out-tray. Incoming elements are added to the in-tray, outgoing elements are taken from the top of the out-tray. Only if the out-tray becomes empty the entire in-tray is reverse-copied into the out-tray. In this case, the length of the in-tray is clearly a suitable potential function. The costly operation of copying can entirely be paid from the big decrease in potential it causes.

In a nutshell automated amortised analysis works as follows. One selects a collection of basic potential functions, called basic resource functions, and assumes that all potential functions are linear combinations of these basic resource functions. One then performs a symbolic amortised cost analysis where the coefficients of these linear combinations as well as the amortised costs of operations (assumed constant) are unknowns. This yields a system of linear constraints for these unknowns whose solution provides the desired amortised analysis from which actual cost bounds in the form of functions of the size of the initial input can be easily read off.

Most automated amortised analyses are univariate in the sense that the joint potential of several arguments to an operation (a “context”) is calculated as the sum of the individual potentials. This, however, proved unsatisfactory in the analysis of nested data structures such as lists of lists or trees and led to the development of multivariate analysis where also products and sums of products of the individual potentials may be used [9].

While automated amortised analysis has hitherto mostly been applied to functional and to a lesser extent to imperative programs; we are here interested in its application to term rewriting understood as a generalisation of functional programming to arbitrary constructor-defined datatypes. After a first step in this direction [12] which was based on univariate analysis, we now generalise to multivariate analysis and indeed subsume and further extend the entire system from [9].

On the one hand this gives a more general treatment of algebraic datatypes which are now untyped and merely defined by their constructor symbols. On the other hand, this necessitates a more general approach to basic resource functions which also streamlines the existing format in [9] or for that matter [11]. In [9] a basic resource function for a list type $L(A)$ is given by a finite list $[p_1, \dots, p_k]$ of basic resource functions for the underlying type of entries A . The interpretation of such a list as a nonnegative \mathbb{R} -valued potential function was then given by the formula

$$[p_1, \dots, p_k]([v_1, \dots, v_n]) := \sum_{1 \leq i_1 < \dots < i_k \leq n} p_1(v_{i_1}) \cdots p_k(v_{i_k}).$$

By treating tree types as a list of entries in depth-first order this same format could then be applied to trees as well. While these formats provided a smooth interaction with the typing rules and allow a very precise analysis of many examples they still look somewhat arbitrary and unjustified.

In the present paper, these formats are subsumed under a very general pattern that is on the one hand simpler and on the other hand permits an even smoother interaction with the typing rules for constructors and matching.

Namely, for us, a basic resource function is defined simply by a bottom-up tree automaton

\mathcal{A} which acts on values by

$$p_{\mathcal{A}}(v) := \text{number of accepting runs of } \mathcal{A} \text{ on } v .$$

If c is a binary constructor, we have

$$p_{\mathcal{A}}(c(v_1, v_2)) = \sum_{c(\beta_1, \beta_2) \rightarrow \alpha \in \mathcal{A}} p_{(\mathcal{A}, \beta_1)}(v_1) \cdot p_{(\mathcal{A}, \beta_2)}(v_2) .$$

where α is the final state of \mathcal{A} and $c(\beta_1, \beta_2) \rightarrow \alpha$ represents a transition in \mathcal{A} and (\mathcal{A}, β_i) denotes \mathcal{A} with the final state set to β_i . Using this formula, the above formats for potentials on lists and trees are readily derived and obviously much more general potential functions can be defined which for example perform a rudimentary kind of type checking by simply ignoring certain constructors or, more interestingly, insisting on certain local patterns.

We note that the expression $p_{\mathcal{A}}(v)$ is known as the *ambiguity* of \mathcal{A} [13] and has been extensively studied. In particular, the above recursive expansion of $p_{\mathcal{A}}(v)$ is known and attributed to Kuich, cf. [13]. However, these previous studies focused mainly on bounding $\max_v p_{\mathcal{A}}(v)$ as a function of the number of states (in the case where this quantity is at all finite) and has to our knowledge never been applied to complexity analysis of tree-like data structures.

Let us now look at a concrete example. Consider the following TRS $\mathcal{R}_{\text{dyade}}$, encoding vector multiplication. The example forms a direct translation of the `dyade.raml` program discussed in [9].

$$\begin{array}{ll} 0 + y \rightarrow y & \mathfrak{s}(x) + y \rightarrow \mathfrak{s}(x + y) \\ 0 \times y \rightarrow 0 & \mathfrak{s}(x) \times y \rightarrow y + (x \times y) \\ \text{mult}(n, []) \rightarrow [] & \text{mult}(n, x :: xs) \rightarrow (n \times x) :: \text{mult}(n, xs) \\ \text{dyade}([], ls) \rightarrow [] & \text{dyade}(x :: xs, ls) \rightarrow \text{mult}(x, ls) :: \text{dyade}(xs, ls) . \end{array}$$

Consider a call to `dyade`(ls_1, ls_2), where ls_1 and ls_2 are lists. It is easy to see that the runtime complexity of $\mathcal{R}_{\text{dyade}}$ crucially depends on the sum of the entries of ls_1 times the sum of the entries of ls_2 , that is an optimal (automated) analysis should provide us with the certificate $O(|ls_1| \cdot |ls_2|)$. However, state-of-the-art complexity tools, like AProVE [7], or TCT [2] overestimate the actual resource usage. For example TCT will provide a polynomial interpretation of degree 2, which is quadratic in $|ls_1|$, even if the monotonicity conditions are weakened suitably, cf. [8]. Also our earlier amortised resource analysis of typed TRS [12] can only provide the non-optimal bound $O(|ls_1|^2 + |ls_2|^2)$. On the other hand the automated analysis of the RaML prototype is more to the point; the analysis with respect to `dyade.raml`, just overestimates the optimal bound by a linear factor and provides unnecessary big constants. The multivariate amortised analysis provided in this paper allows to lift this analysis to the above example and provides essentially optimal bounds (see the example on page 254).

This paper is structured as follows. In the next section we cover basics. In Section 3 we introduce resource functions as generalisations of resource polynomials to arbitrary constructor-defined datatypes. In Section 4 we present our type system and establish its soundness. Finally, we conclude in Section 5, where we also present related work.

2 Term Rewrite Systems and Tree Automata

We assume familiarity with term rewriting [4, 16] and tree automata [6] but briefly review basic concepts and notations.

$$\begin{array}{c}
\frac{x\sigma = v}{\sigma \stackrel{0}{|} x \Rightarrow v} \qquad \frac{c \in \mathcal{C} \quad x_1\sigma = v_1 \quad \cdots \quad x_n\sigma = v_n}{\sigma \stackrel{0}{|} c(x_1, \dots, x_n) \Rightarrow c(v_1, \dots, v_n)} \\
\\
\frac{f(l_1, \dots, l_n) \rightarrow r \in \mathcal{R} \quad \exists \tau \forall i: x_i\sigma = l_i\tau \quad \sigma \uplus \tau \stackrel{m}{|} r \Rightarrow v}{\sigma \stackrel{m+1}{|} f(x_1, \dots, x_n) \Rightarrow v} \\
\\
\text{all } x_i \text{ are fresh} \\
\frac{\sigma \uplus \rho \stackrel{m_0}{|} f(x_1, \dots, x_n) \Rightarrow v \quad \sigma \stackrel{m_1}{|} t_1 \Rightarrow v_1 \quad \cdots \quad \sigma \stackrel{m_n}{|} t_n \Rightarrow v_n \quad m = \sum_{i=0}^n m_i}{\sigma \stackrel{m}{|} f(t_1, \dots, t_n) \Rightarrow v} \\
\\
\text{Here } \rho := \{x_1 \mapsto v_1, \dots, x_n \mapsto v_n\}.
\end{array}$$

■ **Figure 1** Operational Big-Step Semantics.

Let \mathcal{V} denote a countably infinite set of variables and \mathcal{F} a signature, such that \mathcal{F} contains at least one constant. The set of terms over \mathcal{F} and \mathcal{V} is denoted by $\mathcal{T}(\mathcal{F}, \mathcal{V})$. We write $\text{Var}(t)$ to denote the set of variables occurring in term t . The *size* $|t|$ of a term is defined as the number of symbols in t . We suppose $\mathcal{F} = \mathcal{C} \uplus \mathcal{D}$, where \mathcal{C} denotes a finite, non-empty set of *constructor symbols*, \mathcal{D} is a finite set of *defined function symbols*, and \uplus denotes disjoint union. The set of ground constructor terms is denoted as $\mathcal{T}(\mathcal{C})$, ground constructor terms are also called *values*. A *substitution* σ is a mapping from variables to terms. Substitutions are conceived as sets of assignments: $\sigma = \{x_1 \mapsto t_1, \dots, x_n \mapsto t_n\}$. We write $\text{dom}(\sigma)$ ($\text{rg}(\sigma)$) to denote the domain (range) of σ . Let σ, τ be substitutions such that $\text{dom}(\sigma) \cap \text{dom}(\tau) = \emptyset$. Then we denote the (disjoint) union of σ and τ as $\sigma \uplus \tau$. We call a substitution σ *normalised* if all terms in the range of σ are values.

A *rewrite rule* is a pair $l \rightarrow r$ of terms, such that (i) the root symbol of l is defined, and (ii) $\text{Var}(l) \supseteq \text{Var}(r)$. A rule $l \rightarrow r$ is called *left-linear*, if l is linear. A *term rewrite system* (TRS for short) over \mathcal{F} is a finite set of rewrite rules. In the sequel, \mathcal{R} always denotes a TRS. The rewrite relation is denoted as $\rightarrow_{\mathcal{R}}$ and we use the standard notations for its transitive and reflexive closure. We simply write \rightarrow for $\rightarrow_{\mathcal{R}}$ if \mathcal{R} is clear from context. Let s and t be terms. If exactly n steps are performed to rewrite s to t , we write $s \rightarrow^n t$. In the sequel we are concerned with *innermost* rewriting, that is, an eager evaluation strategy. The *innermost rewrite relation* $\overset{\cdot}{\rightarrow}_{\mathcal{R}}$ of a TRS \mathcal{R} is defined on terms as follows: $s \overset{\cdot}{\rightarrow}_{\mathcal{R}} t$ if there exists a rewrite rule $l \rightarrow r \in \mathcal{R}$, a context C , and a substitution σ such that $s = C[l\sigma]$, $t = C[r\sigma]$, and all proper subterms of $l\sigma$ are normal forms of \mathcal{R} .

A TRS is *left-linear* if all rules are left-linear, it is *non-overlapping* if there are no critical pairs, that is, no ambiguity exists in applying rules. A TRS is *orthogonal* if it is left-linear and non-overlapping. A TRS is *completely defined* if all ground normal-forms are values. Note that an orthogonal TRS is confluent. Let s and t be terms, such that t is in normal-form. Then a *derivation* $D: s \rightarrow_{\mathcal{R}}^* t$ with respect to a TRS \mathcal{R} is a finite sequence of rewrite steps. The *derivation height* of a term s with respect to a well-founded, finitely branching relation \rightarrow is defined as: $\text{dh}(s, \rightarrow) = \max\{n \mid \exists t \ s \rightarrow^n t\}$. A term $t = f(t_1, \dots, t_k)$ is called *basic* if f is defined, and all $t_i \in \mathcal{T}(\mathcal{C})$. We define the (*innermost*) *runtime complexity* (with respect to \mathcal{R}): $\text{rc}_{\mathcal{R}}(n) := \max\{\text{dh}(t, \overset{\cdot}{\rightarrow}_{\mathcal{R}}) \mid t \text{ is basic and } |t| \leq n\}$.

We study *constructor* TRSs \mathcal{R} , that is, for each rule $f(l_1, \dots, l_n) \rightarrow r$ we have that the arguments l_i are constructor terms. Furthermore, we restrict to *completely defined* and *orthogonal* systems. These restrictions are natural in the context of functional programming as orthogonal TRSs correspond to first-order function programs with pattern matching. Let \mathcal{F} denote the signature underlying \mathcal{R} .

As \mathcal{R} is completely defined, any derivation ends in a value. In connection with innermost rewriting this yields a *call-by-value* strategy. Furthermore, as \mathcal{R} is non-overlapping any innermost derivation is determined modulo the order in which parallel redexes are contracted. This allows us to recast innermost rewriting into an operational big-step semantics instrumented with resource counters, cf. Figure 1. The semantics resembles similar definitions given in the literature on amortised resource analysis.

► **Proposition 1.** *Let f be a defined function symbol of arity n and σ a normalised substitution. Then $\sigma \Vdash^m f(x_1, \dots, x_n) \Rightarrow v$ holds iff $\text{dh}(f(x_1\sigma, \dots, x_n\sigma), \rightarrow_{\mathcal{R}}) = m$.*

We suit the standard definition of bottom-up tree automata to our context. A *tree automaton* is a quadruple $\mathcal{A} = (\mathcal{F}, \mathcal{Q}, \alpha, \Delta)$, consisting of a finite signature \mathcal{F} , a finite non-empty set of states \mathcal{Q} (disjoint from \mathcal{F}), a unique final state α , and a set of non-empty transitions Δ . Every rule in Δ has the following form $f(\alpha_1, \dots, \alpha_n) \rightarrow \beta$ with $f \in \mathcal{F}$, $\alpha_1, \dots, \alpha_n, \beta \in \mathcal{Q}$.

Note that we only consider tree automata that consist of at least one state and feature a non-empty transition relation. As we will only be concerned with tree automata, we drop the qualifier “tree” and simply speak of an automaton. Observe that an automaton $\mathcal{A} = (\mathcal{F}, \mathcal{Q}, \alpha, \Delta)$ is conceivable as a finite ground TRS Δ over the signature $\mathcal{F} \cup \mathcal{Q}$, where the shape of the rewrite rules is restricted. The induced rewrite relation on $\mathcal{T}(\mathcal{F} \cup \mathcal{Q})$ is denoted as $\rightarrow_{\mathcal{A}}$. A ground term t is *accepted* by \mathcal{A} if $t \rightarrow_{\mathcal{A}}^* \alpha$; we set $L(\mathcal{A}) := \{t \mid t \rightarrow_{\mathcal{A}}^* \alpha\}$. Two automata \mathcal{A} and \mathcal{B} are *equivalent*, if $L(\mathcal{A}) = L(\mathcal{B})$. We use the notation (\mathcal{A}, β) to refer to the automaton $(\mathcal{F}, \mathcal{Q}, \beta, \Delta)$ where $\beta \in \mathcal{Q}$. Note that $(\mathcal{A}, \alpha) = \mathcal{A}$ and we sometimes use the succinct notation instead of the expanded one.

In the sequel, \mathcal{A} will always denote an automaton. Henceforth, \mathcal{R} and \mathcal{F} , as well as the defined symbols \mathcal{D} and constructors \mathcal{C} are kept fixed. Furthermore, all considered substitutions are normalised.

3 Resource Functions

We define a set \mathcal{BF} of *basic functions*, that map terms to natural numbers. Basic functions are indexed by a pair consisting of an automaton \mathcal{A} and a state α . Resource functions will then be defined as nonnegative rational linear combinations of basic functions.

► **Definition 2.** Let $\mathcal{A} = (\mathcal{C}, \mathcal{Q}, \alpha, \Delta)$. For $v \in \mathcal{T}(\mathcal{C})$ we define the *basic function* $p_{\mathcal{A}}$, whose value $p_{\mathcal{A}}(v)$ is the number of accepting runs of \mathcal{A} on v . The set of basic functions is denoted as \mathcal{BF} .

For any set of constructors \mathcal{C} , there exists an automaton \mathcal{A} with $p_{\mathcal{A}}(v) = 1$ for all values v . Moreover \mathcal{A} is unique upto renaming of the single state α . We call \mathcal{A} the *canonical* automaton, denoted by \emptyset , whose unique state is denoted by \emptyset . As mentioned in the introduction $p_{\mathcal{A}}(v)$ is called *ambiguity* in the literature (see for example [13, 14]). In particular it is known that the finiteness of the *degree of ambiguity* $\sup_v p_{\mathcal{A}}(v)$ is polytime decidable [13]. The following is direct from the definition.

► **Proposition 3.** *Let $v \in \mathcal{T}(\mathcal{C})$ be a value such that $v = c(v_1, \dots, v_n)$ and let $\alpha \in \mathcal{Q}$. We then have:*

$$p_{(\mathcal{A}, \alpha)}(c(v_1, \dots, v_n)) := \sum_{c(\alpha_1, \dots, \alpha_n) \rightarrow \alpha \in \Delta} p_{(\mathcal{A}, \alpha_1)}(v_1) \cdots p_{(\mathcal{A}, \alpha_n)}(v_n).$$

Note that $p_{(\mathcal{A}, \alpha)}(c) = \sum_{c \rightarrow \alpha \in \Delta} 1$, as the empty product equals 1.

We could alternatively have used the latter as a recursive definition of $p_{(\mathcal{A},\alpha)}(v)$. The advantage of the definition based on runs is that we can immediately read off an exponential upper bound:

► **Proposition 4.** *Let q be the number of states of \mathcal{A} and n the size of $v \in \mathcal{T}(\mathcal{C})$. Then $p_{\mathcal{A}}(v) \leq q^n$.*

It is also easy to see that this bound is actually taken on so that unlike the basic functions in [9] ours are not in general polynomials. If desired, it is however easy to impose syntactic restrictions that ensure polynomial growth. For example, we can use a ranking function on states and require that for each transition $c(\alpha_1, \dots, \alpha_n) \rightarrow \alpha$ the ranks of the α_i are not bigger than that of α and that they strictly decrease for all but one transition with symbol c and result state α . All the concrete basic functions we use in examples satisfy this restriction and thus exhibit polynomial growth. We believe, however, that in some cases also exponential bounding functions may prove useful.

Let $\mathcal{C} = \{0, s, [], ::\}$ and consider the following automaton \mathcal{A} with final state β_k .

$$\begin{array}{ccccccc} 0 \rightarrow \alpha_0 & s(\alpha_0) \rightarrow \alpha_1 & \cdots & s(\alpha_{\ell-1}) \rightarrow \alpha_\ell & & & \\ & s(\alpha_0) \rightarrow \alpha_0 & \cdots & s(\alpha_{\ell-1}) \rightarrow \alpha_{\ell-1} & s(\alpha_\ell) \rightarrow \alpha_\ell & & \\ [] \rightarrow \beta_0 & \alpha_{i_k} :: \beta_0 \rightarrow \beta_1 & \cdots & \alpha_{i_1} :: \beta_{k-1} \rightarrow \beta_k & & & \\ & \alpha_0 :: \beta_0 \rightarrow \beta_0 & \cdots & \alpha_0 :: \beta_{k-1} \rightarrow \beta_{k-1} & \alpha_0 :: \beta_k \rightarrow \beta_k & & \end{array}$$

First, by a simple inductive argument we see that $p_{(\mathcal{A},\alpha_i)}(s^n(0)) = \binom{n}{i}$ for $n \geq 0$ and $i = 0, \dots, \ell$. Based on this, we conclude by induction on m :

$$p_{(\mathcal{A},\beta_k)}([\mathbf{n}_1, \dots, \mathbf{n}_m]) = \sum_{1 \leq j_1 < \dots < j_k \leq m} \binom{n_{j_1}}{i_1} \cdots \binom{n_{j_k}}{i_k},$$

where we denote the numeral $s^n(0)$ by \mathbf{n} , $[\mathbf{n}_1, \dots, \mathbf{n}_m]$ abbreviates the corresponding cons-list, and $1 \leq i_j \leq \ell$ for all $j = 1, \dots, k$.

Consider the set, denoted as \mathfrak{A} , of all non-equivalent (and non-empty) automata over \mathcal{C} . In the following we will frequently appeal to an enumeration of \mathfrak{A} , referring to a suitable chosen, but inessential encoding of automata. Note that in effect we will only work with a small, in particular finite subset of \mathfrak{A} .

► **Definition 5.** A *resource function* $r: \mathcal{T}(\mathcal{C}) \rightarrow \mathbb{Q}^+$ is a non-negative rational linear combination of basic functions, that is,

$$r(t) := \sum_{\mathcal{A} \in \mathfrak{A}} q_{\mathcal{A}} \cdot p_{\mathcal{A}}(t),$$

where $p_{\mathcal{A}} \in \mathcal{BF}$. The set of resource functions is denoted as \mathcal{RF} .

The example above hints at the fact that the expressivity of our basic functions exceeds the expressivity of the *base polynomials* considered by Hoffmann et al. [10, 9]. The next proposition makes this fact precise.

► **Lemma 6.** *All the resource polynomials from [9] are also resource functions in the present automata-based setting.*

Proof. This is proved by induction on the definition of resource polynomials [9]. We do not need to recall their definition here; the inductive cases we establish reveal enough detail.

If p, q are basic resource polynomials for types A, B respectively then $\lambda ab(p(a) \cdot q(b))$ is a base resource polynomial for the product type $A \times B$. In rewriting, we can simulate product types by introducing a binary constructor $\text{pair}(x, y)$. Now, if, inductively $p = p_{(\mathcal{A}, \alpha)}$ and $q = p_{(\mathcal{B}, \beta)}$, and w.l.o.g. the two automata have disjoint state sets, then we can build an automaton \mathcal{C} whose states are the union of the states of \mathcal{A} and \mathcal{B} together with a new state γ and a transition $\text{pair}(\alpha, \beta) \rightarrow \gamma$ in addition to the transitions from \mathcal{A} and \mathcal{B} . We then have $p_{(\mathcal{C}, \gamma)}(a, b) = p_{(\mathcal{A}, \alpha)}(a) \cdot p_{(\mathcal{B}, \beta)}(b) = p(a) \cdot q(b)$ as required. If p_1, \dots, p_k are base polynomials for type A then the base polynomial $[p_1, \dots, p_k]$ given by

$$[p_1, \dots, p_k]([a_1, \dots, a_m]) := \sum_{1 \leq i_1 < i_2 < \dots < i_k \leq m} p_1(a_{i_1}) \cdots p_k(a_{i_k}),$$

is the generic base polynomial for lists over type A . Assuming that lists are constructed with the symbols $[]$ and $::$ and that, inductively, $p_i = p_{(\mathcal{A}_i, \alpha_i)}$, we can build an automaton \mathcal{B} as the disjoint union of the \mathcal{A}_i together with $k + 1$ states $\beta_0, \beta_1, \dots, \beta_k$, and transitions:

$$\emptyset :: \beta_i \rightarrow \beta_i \quad \alpha_{i+1} :: \beta_{i+1} \rightarrow \beta_i \quad [] \rightarrow \emptyset,$$

where $i = 0, \dots, k - 1$. (Recall that \emptyset denotes the final state of the canonical automaton \emptyset , whose inclusion we here tacitly assume.) We obtain $[p_1, \dots, p_k]([a_1, \dots, a_n]) = p_{(\mathcal{B}, \beta_k)}(a_1 :: a_2 \cdots a_n :: [])$. Finally, the generic resource polynomial for A -labelled trees takes the form

$$[p_1, \dots, p_k](t) := [p_1, \dots, p_k](l_t),$$

where l_t is the list of entries of t (in the leaves) in depth-first order. Note that we have

$$ls(t) := \begin{cases} \sum_{i=0}^k ([p_1, \dots, p_i](t_1) \cdot & \text{if } ls = [p_1, \dots, p_k], k > 1, \text{ and } t = \text{node}(t_1, t_2) \\ \quad \cdot [p_{i+1}, \dots, p_k](t_2)) & \\ p(a) & \text{if } ls = [p] \text{ and } t = \text{leaf}(a) \\ 1 & \text{if } ls = [] \text{ and } t = \text{leaf}(a) \\ 0 & \text{otherwise.} \end{cases}$$

Thus, assuming the automata \mathcal{A}_i ($i = 1, \dots, k$) as above, we can construct a new automaton \mathcal{B} for $[p_1, \dots, p_k]$ as the disjoint union of the \mathcal{A}_i together with new states $\beta_{i,j}$ for $1 \leq i \leq j \leq k$ and the following transitions:

$$\text{leaf}(\alpha_i) \rightarrow \beta_{i,i+1} \quad \text{leaf}(\emptyset) \rightarrow \beta_{i,i} \quad \text{node}(\beta_{i,t}, \beta_{t,j}) \rightarrow \beta_{i,j},$$

where $t = i, \dots, j$. As above, we obtain $[p_1, \dots, p_k](t) = p_{(\mathcal{B}, \beta_{1,k})}(l_t)$. Thus the lemma follows. \blacktriangleleft

► **Lemma 7.** *If r and r' are resource functions, then $r + r', r \cdot r' \in \mathcal{RF}$.*

Proof. First, resource functions are obviously closed under addition. With respect to multiplication we employ the linearity of resource functions to see that it suffices to prove the claim for basic functions. Here we argue similar to the proof of Lemma 6 by using a product automata construction. \blacktriangleleft

► **Definition 8.** We define the set of *multivariate basic functions*, denoted as $\mathcal{BF}(n)$:

$$\mathcal{BF}(n) := \{v_1, \dots, v_n \mapsto p_{\mathcal{A}_1}(v_1) \cdots p_{\mathcal{A}_n}(v_n) \mid \text{for all } i: v_i \in \mathcal{T}(\mathcal{C}) \text{ and } p_{\mathcal{A}_i} \in \mathcal{BF}\}.$$

We enumerate the set $\mathcal{BF}(n)$ by sequences of automata, such that $p_{(\mathcal{A}_1, \dots, \mathcal{A}_n)}(v_1, \dots, v_n) = p_{\mathcal{A}_1}(v_1) \cdots p_{\mathcal{A}_n}(v_n)$.

In the following sequences of automata (like $(\mathcal{A}_1, \dots, \mathcal{A}_n)$) are sometimes abbreviated as $\vec{\mathcal{A}}$, in particular we set $\vec{\emptyset} = (\emptyset, \dots, \emptyset)$. Note that $\mathcal{BF}(1) = \mathcal{BF}$; in the following we use \mathcal{BF} as unique denotation.

► **Lemma 9.** *For all $p \in \mathcal{BF}(2)$, there exists $p' \in \mathcal{BF}$, such that $p(v, v) = p'(v)$ for all values v .*

Proof. By definition there exists automata \mathcal{A}, \mathcal{B} such that $p(v, v) = p_{\mathcal{A}}(v) \cdot p_{\mathcal{B}}(v)$. By the previous lemma there exists an automaton \mathcal{C} such that $p_{\mathcal{C}}(v) = p_{\mathcal{A}}(v) \cdot p_{\mathcal{B}}(v)$. Thus we set $p' := p_{\mathcal{C}}$ to conclude the lemma. ◀

Let \mathcal{C} denote a set of constructor symbols. A *resource annotation over \mathcal{C}* , or simply *annotation*, is a family $Q = (q_{\mathcal{A}})_{\mathcal{A} \in \mathfrak{A}}$ with $q_{\mathcal{A}} \in \mathbb{Q}^+$ with all but finitely many of the coefficients $q_{\mathcal{A}}$ equal to 0. It represents a (finite) linear combination of basic resource functions. We generalise annotation for sequences of terms. An annotation for a sequence of length n is a family $Q = (q_{(\mathcal{A}_1, \dots, \mathcal{A}_n)})_{\mathcal{A}_i \in \mathfrak{A}}$ again vanishing almost everywhere. We denote annotations with upper-case letters from the end of the alphabet and use the convention that the corresponding lower-case letter denotes the elements of the family.

► **Definition 10.** The *potential* of a value v with respect to an annotation Q (of length 1), that is, $Q = (q_{\mathcal{A}})_{\mathcal{A} \in \mathfrak{A}}$, is defined as:

$$\Phi(v; Q) := \sum_{\mathcal{A} \in \mathfrak{A}} q_{\mathcal{A}} \cdot p_{\mathcal{A}}(v),$$

where $p_{\mathcal{A}} \in \mathcal{BF}$. We generalise this to the potential of a term sequence with respect to an annotation $\bar{Q} = (q_{(\mathcal{A}_1, \dots, \mathcal{A}_n)})_{\mathcal{A}_i \in \mathfrak{A}}$: $\Phi(v_1, \dots, v_n; \bar{Q}) := \sum_{\mathcal{A}_1, \dots, \mathcal{A}_n \in \mathfrak{A}} q_{(\mathcal{A}_1, \dots, \mathcal{A}_n)} \cdot p_{(\mathcal{A}_1, \dots, \mathcal{A}_n)}(v_1, \dots, v_n)$, where $p_{(\mathcal{A}_1, \dots, \mathcal{A}_n)} \in \mathcal{BF}(n)$.

We are ready to generalise the notion of additive shift studied in [10, 9, 12].

► **Definition 11.** Suppose $Q = (q_{(\mathcal{A}_1, \dots, \mathcal{A}_m, \mathcal{B})})_{\mathcal{A}_i, \mathcal{B} \in \mathfrak{A}}$ denotes a resource annotation of length $m + 1$. Let $c \in \mathcal{C}$ be a constructor symbol of arity n . The *additive shift for c* of Q is an annotation $\triangleleft_c(Q) = (q'_{(\mathcal{A}_1, \dots, \mathcal{A}_m, \mathcal{B}_1, \dots, \mathcal{B}_n)})_{\mathcal{A}_i, \mathcal{B}_j \in \mathfrak{A}}$ for a sequence of length $n + m$, where $q'_{(\mathcal{A}_1, \dots, \mathcal{A}_m, \mathcal{B}_1, \dots, \mathcal{B}_n)}$ is defined as follows:

$$q'_{(\mathcal{A}_1, \dots, \mathcal{A}_m, \mathcal{B}_1, \dots, \mathcal{B}_n)} := \sum_{\substack{\mathcal{B} = (\mathcal{C}, \mathcal{Q}, \beta, \Delta) \in \mathfrak{A} \\ \text{with } c(\beta_1, \dots, \beta_n) \rightarrow \beta \in \Delta}} q_{(\mathcal{A}_1, \dots, \mathcal{A}_m, \mathcal{B})}, \quad (1)$$

Here $\mathcal{B}_i = (\mathcal{C}, \mathcal{Q}_{\mathcal{B}_i}, \beta_i, \Delta_{\mathcal{B}_i})$ for each $i = 1, \dots, n$.

The correctness of the additive shift operation follows from the next lemma.

► **Lemma 12.** *Let $v = c(v_1, \dots, v_n)$ be a value and let $Q = (q_{(\mathcal{A}_1, \dots, \mathcal{A}_m, \mathcal{B})})_{\mathcal{A}_1, \dots, \mathcal{A}_m, \mathcal{B} \in \mathfrak{A}}$ be an annotation of length $m + 1$. Then $\Phi(w_1, \dots, w_m, v; Q) = \Phi(w_1, \dots, w_m, v_1, \dots, v_n; \triangleleft_c(Q))$.*

Proof. In proof, we restrict to the case $m = 0$. Consider the value $v = c(v_1, \dots, v_n)$:

$$\begin{aligned}
\Phi(v : Q) &= \sum_{\mathcal{B} \in \mathfrak{A}} q_{\mathcal{B}} \cdot p_{\mathcal{B}}(c(v_1, \dots, v_n)) \\
&= \sum_{\mathcal{B} \in \mathfrak{A}} q_{\mathcal{B}} \cdot \left(\sum_{c(\beta_1, \dots, \beta_n) \rightarrow \beta \in \Delta} p_{(\mathcal{B}, \beta_1)}(v_1) \cdots p_{(\mathcal{B}, \beta_n)}(v_n) \right) \\
&= \sum_{\mathcal{B}_1, \dots, \mathcal{B}_n \in \mathfrak{A}} q'_{(\mathcal{B}_1, \dots, \mathcal{B}_n)} \cdot (p_{\mathcal{B}_1}(v_1) \cdots p_{\mathcal{B}_n}(v_n)) \\
&= \sum_{\mathcal{B}_1, \dots, \mathcal{B}_n \in \mathfrak{A}} q'_{(\mathcal{B}_1, \dots, \mathcal{B}_n)} \cdot p_{(\mathcal{B}_1, \dots, \mathcal{B}_n)}(v_1, \dots, v_n) = \Phi(v_1, \dots, v_n : \triangleleft_c(Q)).
\end{aligned}$$

Here is straightforward to check that $q'_{(\mathcal{B}_1, \dots, \mathcal{B}_n)}$ is as in (1). Thus the last equation (and the lemma) follows. \blacktriangleleft

Let $Q = (q_{(\mathcal{A}_1, \dots, \mathcal{A}_n)})_{\mathcal{A}_i \in \mathfrak{A}}$ denote a resource annotation of length n . Let $\vec{B} = (\mathcal{B}_1, \dots, \mathcal{B}_m)$; we define the *projection* of Q with respect to \vec{B} to an annotation of length $\ell < n$. The projection is denoted as $\pi_{\vec{B}}^{\ell}(Q)$. We set

$$\pi_{\vec{B}}^{\ell}(Q) := (q'_{(\mathcal{A}_1, \dots, \mathcal{A}_{\ell})})_{\mathcal{A}_i \in \mathfrak{A}},$$

where $q'_{(\mathcal{A}_1, \dots, \mathcal{A}_{\ell})} = q_{(\mathcal{A}_1, \dots, \mathcal{A}_{\ell}, \mathcal{B}_1, \dots, \mathcal{B}_m)}$ and $n = \ell + m$. Suppose $\Gamma, v_1, v_2 : Q$ denotes an annotated sequence of length $m + 2$. Suppose $v_1 = v_2$ and we want to *share* the values. Then we make use of the operator $\Upsilon(Q)$ that adapts the potential suitably. The operator is also called *sharing operator*.

► **Lemma 13.** *Let $\Gamma, v_1, v_2 : Q$ denote an annotated sequence of length $m + 2$. Then there exists a resource annotation $\Upsilon(Q)$ such that $\Phi(\Gamma, v_1, v_2 : Q) = \Phi(\Gamma, v : \Upsilon(Q))$, if $v_1 = v_2 = v$.*

Proof. This follows from Lemma 9. \blacktriangleleft

Let Q be an annotation over the set of constructor symbols \mathcal{C} and let $K \in \mathbb{Q}^+$. Then we define $Q' := Q + K$ as follows: $Q' = (q'_{\mathcal{A}})_{\mathcal{A} \in \mathfrak{A}}$, where $q'_{\emptyset} := q_{\emptyset} + K$ and for all $\mathcal{A} \neq \emptyset$, $q'_{\mathcal{A}} := q_{\mathcal{A}}$. We define the comparison \leq of two annotations Q, Q' over \mathcal{C} pointwise: $Q \leq Q'$ if for all $\mathcal{A} \in \mathfrak{A}$: $q_{\mathcal{A}} \leq q'_{\mathcal{A}}$.

4 Amortised Cost Analysis

The rest of the paper is essentially an adaptation of the type systems given in [10, 9, 12]. There are no essential surprises but care must be taken with the rule for the evaluation of non-constructor terms which is essentially a combination of the let rule and the function application rule from [9].

Let Γ denote a sequence of variables x_1, \dots, x_n ; Γ is called a *variable context* or simply a *context*. The *length* n of Γ is denoted as $|\Gamma|$. A *resource annotation for Γ* or simply *annotation* is an annotation of length $|\Gamma|$. Now let Q be a resource annotation for Γ , let Q' be a resource annotation, let t be a term and let v denote its normalform. Then the typing judgement $\Gamma : Q \vdash t : Q'$ expresses that for *bounded* TRS (see Definition 15) the potential of Γ with respect to Q is sufficient to pay for the total cost m of the evaluation $\sigma \stackrel{m}{\mapsto} t \Rightarrow v$ (σ a substitution), plus the potential of the value v with respect to Q' .

We comment on the inference rules in Figure 2. For the majority of the typing rules their definition is straightforward. Consider exemplary the typing rule for constructors

$c \in \mathcal{C}$. Due to Lemma 6 the potential of a value $c(x_1\sigma, \dots, x_n\sigma)$ equals the potential of the variable context modulo an additive shift with respect to c . This is precisely expressed in the corresponding type rule. On the other hand the composition rule is more involved. Consider the judgement $\Gamma_1, \dots, \Gamma_n : Q \vdash f(t_1, \dots, t_n) : Q'$. Observe that on the basis of the sharing rule we can assume that $t = f(t_1, \dots, t_n)$ is linear and thus the variable context splits into its parts Γ_i ($i = 1, \dots, n$). The intuition of the composition is that the potential of $\Gamma_1, \dots, \Gamma_n$ (represented through the annotation Q) should be distributed over the evaluations of all arguments t_i and the evaluation of the function f , in a way such that the interdependency of the bounds is preserved. This is achieved by type checking each of the arguments and the context individually and verifying that all issuing annotations are consistent with each other.

In order to see how this works, we detail some of the constraints. First, consider $\pi_{\Gamma_1}^{\vec{j}_1}(Q) = P_1(\vec{j}_1)$. Here \vec{j}_1 denotes an index for $\Gamma_2, \dots, \Gamma_n$. The projection asserts that the resources annotated in Q are projected to Γ_1 , so that the typing $\Gamma_1 : P_1(\vec{j}_1) \vdash^{(\text{cf})} t_1 : R_1(\vec{j}_1)$ is realisable for every index \vec{j}_1 . The constraint $p_\ell^{i+1, \vec{j}_{i+1}} = r_k^{i, \vec{j}_i}$, where $1 \leq i < n$, ℓ an index for Γ_i and $k \in I$, guarantees that the annotations for the remaining contexts Γ_i , $i > 1$ are consistent with each other. Finally, constraint $\pi_{x_n}^{\vec{j}_n}(S) = R_n(\vec{j}_n)$ links the potential after the evaluation of the last argument t_n consistently with the annotation S for the variable context x_1, \dots, x_n .

The type system given requires the use of *cost-free* judgements. Here the rules of the given TRS are considered as weak rules that are not taken into account in the complexity evaluation (see Definition 15).

► **Definition 14.** An *annotated signature* $\overline{\mathcal{F}}$ is a mapping from \mathcal{D} to sets of pairs of resource annotations:

$$\overline{\mathcal{F}}(f) := \left\{ Q \rightarrow Q' \mid \begin{array}{l} \text{if the arity of } f \text{ is } n, Q \text{ is an annotation of length } \\ n \text{ and } Q' \text{ a resource annotation} \end{array} \right\}.$$

Usually, we confuse the signature and the annotated signature and denote the latter simply as \mathcal{F} .

The set of indices of an annotation for a context Γ is defined as follows:

$$I(\Gamma) := \{(\mathcal{A}_1, \dots, \mathcal{A}_n) \mid \text{if } \mathcal{A} \in \mathfrak{A} \text{ and } |\Gamma| = n\}.$$

We also set $I := \mathfrak{A}$. Let Q be a resource annotation of length n , let $\Gamma = x_1, \dots, x_\ell$ be a variable context, and let $\vec{i} = \mathcal{B}_1, \dots, \mathcal{B}_m$ be an index of length m . We define the *projection with respect to* Γ : $\pi_{\Gamma}^{\vec{i}}(Q) := \pi_{|\Gamma|}^{\vec{i}}(Q)$.

Recall that any rewrite rule $l \rightarrow r \in \mathcal{R}$ can be written as $f(l_1, \dots, l_n) \rightarrow r$ with $l_i \in \mathcal{T}(\mathcal{C}, \mathcal{V})$ and that no variable occurs twice in l . Similar to the notion of well-typed TRSs in [12], we introduce *bounded* TRSs.

► **Definition 15.** Let $f(l_1, \dots, l_n) \rightarrow r$ be a rewrite rule in \mathcal{R} and let $\text{Var}(f(l_1, \dots, l_n)) = \{x_1, \dots, x_\ell\}$. Suppose f is defined and let $Q \rightarrow Q' \in \mathcal{F}(f)$. Suppose further, for any such annotation $Q \rightarrow Q'$ we can derive

$$x_1, \dots, x_\ell : P \vdash r : Q', \tag{2}$$

where $P + 1 = R$ and R is obtained by iteratively applying shift operations with respect to the constructors occurring in $\bigcup_{i=1}^n l_i$ to Q . Then we say f is *bounded* wrt. \mathcal{F} . On the other hand f is *weakly bounded* if the rewrite step is not counted, that is, the judgement $x_1, \dots, x_\ell : R \vdash r : Q'$ is asserted instead of (2). A TRS \mathcal{R} over \mathcal{F} is *bounded* (*weakly bounded*) if any defined f is bounded (weakly bounded).

$$\begin{array}{c}
\frac{}{x:Q \vdash x:Q} \qquad \frac{\Gamma: \pi_{\Gamma}^{\emptyset}(Q) \vdash t:Q'}{\Gamma, x:Q \vdash t:Q'} \\
\frac{f \in \mathcal{D} \quad Q \rightarrow Q' \in \mathcal{F}(f) \quad f \text{ is } n\text{-ary}}{x_1, \dots, x_n:Q \vdash f(x_1, \dots, x_n):Q'} \qquad \frac{c \in \mathcal{C} \quad c \text{ is } n\text{-ary}}{x_1, \dots, x_n:\triangleleft_c(Q) \vdash c(x_1, \dots, x_n):Q} \\
\forall \vec{j}_1, \dots, \vec{j}_n \quad \vec{j}_i \in I^{i-1} \times I(\Gamma_{i+1}) \times \dots \times I(\Gamma_n) \\
P_i(\vec{j}_i) = (p^{\ell, \vec{j}_i})_{\ell \in I(\Gamma_i)} \quad R_i(\vec{j}_i) = (r^{k, \vec{j}_i})_{k \in I} \\
p_{\ell}^{i+1, \vec{j}_{i+1}} = r_{k}^{i, \vec{j}_i} \text{ if } \vec{j}_i = a @ \ell @ b \text{ and } \vec{j}_{i+1} = a @ k @ b \\
\frac{\pi_{\Gamma_1}^{\vec{j}_1}(Q) = P_1(\vec{j}_1) \quad \pi_{x_n}^{\vec{j}_n}(S) = R_n(\vec{j}_n) \quad x_1, \dots, x_n:S \vdash f(x_1, \dots, x_n):Q' \quad \Gamma_1:P_1(\vec{j}_1) \vdash^{(\text{cf})} t_1:R_1(\vec{j}_1) \quad \dots \quad \Gamma_n:P_n(\vec{j}_n) \vdash^{(\text{cf})} t_n:R_n(\vec{j}_n)}{\Gamma_1, \dots, \Gamma_n:Q \vdash f(t_1, \dots, t_n):Q'} \\
\frac{\Gamma, x, y:Q \vdash t[x, y]:Q' \quad x, y \text{ are fresh}}{\Gamma, z:\Upsilon(Q) \vdash t[z, z]:Q'} \qquad \frac{\Gamma, x:P \vdash t:P' \quad P \leq Q \quad P' \geq Q'}{\Gamma, x:Q \vdash t:Q'}
\end{array}$$

In the composition rule a and b denote suitable chosen indices, where $@$ denotes concatenation of indices. Further, the judgement $\Gamma_i:P_i(\vec{j}_i) \vdash^{(\text{cf})} t_i:R_i(\vec{j}_i)$ abbreviates that $\Gamma_i:P_i(\vec{\emptyset}) \vdash t_i:R_i(\vec{\emptyset})$ and $\Gamma_i:P_i(\vec{j}_i) \vdash^{(\text{cf})} t_i:R_i(\vec{j}_i)$ for all $\vec{j}_i \neq \vec{\emptyset}$ and $i = 1, \dots, n$.

■ **Figure 2** Multivariate Analysis of Term Rewrite Systems.

$$\frac{\frac{y_1:P_1(j) \vdash^{(\text{cf})} y_1:R_1(j) \quad \frac{\times:P_2(i) \rightarrow R_2(i)}{x, y_2:P_2(i) \vdash^{(\text{cf})} x \times y_2:R_2(i)} \quad \frac{+:S \rightarrow S'}{u, v:S \vdash u + v:S'}}{y_1, x, y_2:T \vdash y_1 + (x \times y_2):M'} \quad \frac{}{x, y:\bar{M}_2 \vdash y + (x \times y):M'}}{}$$

Here $i \in I$, and $j \in I(x, y_2)$.

■ **Figure 3** Derivation of $x, y:\bar{M}_2 \vdash y + (x \times y):M'$.

Let σ denote a substitution, $\Gamma = x_1, \dots, x_n$ a context and Q a resource annotation. Then we define the *potential* of $\Gamma:Q$ with respect to σ :

$$\Phi(\sigma, \Gamma:Q) := \Phi(x_1\sigma, \dots, x_n\sigma:Q).$$

Note that the above definition employs the shift operator in a similar way as in [9], where this is part of the type system in the case for pattern matching.

Before we state our main result, we exemplify the use of the type system on a simple (but clarifying) example. Consider the following TRS \mathcal{R}_{\times} , restricting our motivating example (see page 243).

$$\begin{array}{ll}
0 + y \rightarrow y & \mathfrak{s}(x) + y \rightarrow \mathfrak{s}(x + y) \\
0 \times y \rightarrow 0 & \mathfrak{s}(x) \times y \rightarrow y + (x \times y).
\end{array}$$

We consider the canonical automaton \emptyset for the constructor symbols $\{0, \mathfrak{s}\}$ together with the automaton \mathcal{A} defined as follows:

$$\mathcal{A}: \quad 0 \rightarrow \emptyset \qquad \mathfrak{s}(\emptyset) \rightarrow \alpha \qquad \mathfrak{s}(\emptyset) \rightarrow \emptyset \qquad \mathfrak{s}(\alpha) \rightarrow \alpha$$

We show that \mathcal{R}_\times is bounded with respect to the following annotated signatures:

$$\begin{aligned} +: & \{ (p_{(\emptyset,\emptyset)}, p_{(\mathcal{A},\emptyset)}) \rightarrow (p'_{\emptyset}) \mid p_{(\emptyset,\emptyset)} \geq 1, p_{(\mathcal{A},\emptyset)} \geq 1, p_{(\emptyset,\emptyset)} - 1 \geq p'_{\emptyset} \} \\ \times: & \left\{ \begin{pmatrix} m_{(\emptyset,\emptyset)} & m_{(\mathcal{A},\emptyset)} \\ 0 & m_{(\mathcal{A},\mathcal{A})} \end{pmatrix} \rightarrow (m'_{\emptyset}, m'_{\mathcal{A}}) \mid \begin{array}{l} m_{(\emptyset,\emptyset)} \geq 1, m_{(\mathcal{A},\emptyset)} \geq 2, \\ m_{(\mathcal{A},\mathcal{A})} \geq 1, m_{(\emptyset,\emptyset)} - 1 \geq m'_{\emptyset} \end{array} \right\}. \end{aligned}$$

Each of the four rules induced one of the following demands, cf. (2). (Recall that we denote annotations with upper-case letters and the elements of the families with the corresponding lower-case letters.)

$$y: \bar{P}_1 \vdash y: P' \qquad \bar{P}_1 + 1 = \triangleleft_0(P) \qquad (3)$$

$$x, y: \bar{P}_2 \vdash s(x + y): P' \qquad \bar{P}_2 + 1 = \triangleleft_s(P) \qquad (4)$$

$$y: \bar{M}_1 \vdash 0: M' \qquad \bar{M}_1 + 1 = \triangleleft_0(M) \qquad (5)$$

$$x, y: \bar{M}_2 \vdash y + (x \times y): M' \qquad \bar{M}_2 + 1 = \triangleleft_s(M) \qquad (6)$$

It is not difficult to see that (3) and (4) induce the constraints $p_{(\emptyset,\emptyset)} \geq 1$, $p_{(\mathcal{A},\emptyset)} \geq 1$, and $p_{(\emptyset,\emptyset)} - 1 \geq p'_{\emptyset}$, and to ease the presentation we set $p'_{\mathcal{A}} = 0$. Furthermore, the typing judgement (5) yields the constraint $m_{(\emptyset,\emptyset)} - 1 \geq m'_{\emptyset}$.

Of more interest is the precise derivation of (6). In this derivation we make use of the weakly boundedness of \mathcal{R}_\times with respect to the cost-free signature $+: (p_{(\emptyset,\emptyset)}) \rightarrow p'_{\emptyset}$ and $\times: (m_{(\emptyset,\emptyset)}) \rightarrow m'_{\emptyset}$, where it suffices to demand that $p_{(\emptyset,\emptyset)} \geq p'_{\emptyset}$ and $m_{(\emptyset,\emptyset)} \geq m'_{\emptyset}$. We obtain the following derivation in Figure 3. This derivation induces the constraint $\Upsilon(T) = \bar{M}_2$, as first (reading bottom-up) we employ a sharing rule. The composition rule yields the constraints $\pi_{y_1}^j(T) = P_1(j)$ ($j \in I(x, y_2)$), $\pi_v^i(S) = R_2(i)$ ($i \in I$), and $r_i^{1,j} = p_j^{2,i}$, where $R_1(j) = (r^{1,j})_{i \in I}$ and $P_2(i) = (p^{2,i})_{j \in I(x, y_2)}$. Finally, the axioms yield the constraints $P_1(j) \geq R_i(j)$ for $j \in I(x, y_2)$ as well as the indicated conditions on the signature. It is tedious, but straightforward to check that these constraints can be met for the given annotations. Thus, \times can be typed with the annotation $m_{(\emptyset,\emptyset)} = 1$, $m_{(\mathcal{A},\emptyset)} = 2$, and $m_{(\mathcal{A},\mathcal{A})} = 1$ which yields the bound $\text{dh}(\mathbf{m} \times \mathbf{n}, \rightarrow_{\mathcal{R}_\times}) \leq m \cdot n + 2m + 1$.

It is instructive to depict the annotation \bar{M}_2 in matrix format, which allows a simple expression of the \triangleleft_s -operator.

$$\bar{M}_2 = \begin{pmatrix} \bar{m}_{\emptyset,\emptyset} & \bar{m}_{\mathcal{A},\emptyset} \\ \bar{m}_{\emptyset,\mathcal{A}} & \bar{m}_{\mathcal{A},\mathcal{A}} \end{pmatrix} = \begin{pmatrix} m_{\emptyset,\emptyset} + m_{\mathcal{A},\emptyset} - 1 & m_{\mathcal{A},\emptyset} \\ m_{\mathcal{A},\mathcal{A}} & m_{\mathcal{A},\mathcal{A}} \end{pmatrix} = \triangleleft_s(M).$$

Then it becomes apparent that the annotation \bar{M}_2 is the result of adding the auxiliary annotation

$$\begin{pmatrix} m_{\mathcal{A},\emptyset} - 1 & 0 \\ m_{\mathcal{A},\mathcal{A}} & 0 \end{pmatrix}, \qquad (7)$$

to the annotation for multiplication. Intuitively the type system asserts that we can split the annotation \bar{M}_2 into an annotation M that pays for the recursive call and the annotation (7) that pays for the call to addition.

► **Theorem 16.** *Let \mathcal{R} be bounded. Suppose $\Gamma: Q \vdash t: Q'$ and $\sigma \stackrel{m}{\vdash} t \Rightarrow v$. Then $\Phi(\sigma, \Gamma: Q) - \Phi(v: Q') \geq m$.*

Proof. Let Π be the proof deriving $\sigma \stackrel{m}{\vdash} t \Rightarrow v$ and let Ξ be the proof of $\Gamma: Q \vdash t: Q'$. The proof of the theorem proceeds by main-induction on the length of Π and by side-induction

on the length of Ξ . We consider the case for composition. Suppose the last rule in Π has the form

$$\frac{\sigma \uplus \rho \mid^{m_0} f(x_1, \dots, x_n) \Rightarrow v \quad \sigma \mid^{m_i} t_i \Rightarrow v_i \quad i = 1, \dots, n \quad m = \sum_{i=0}^n m_i}{\sigma \mid^m t \Rightarrow v} .$$

We can assume that $t = f(x_1, \dots, x_n)$ is linear, due the presence of the share operator. In proof we restrict to the case where $n = 2$. Hence the last rule in the type inference Ξ is of the following form.

$$\frac{\Gamma_1 : P(j) \vdash^{(\text{cf})} t_1 : \bar{P}(j) \quad \Gamma_2 : R(i) \vdash^{(\text{cf})} t_2 : \bar{R}(i) \quad x, y : S \vdash f(x, y) : Q'}{\Gamma_1, \Gamma_2 : Q \vdash f(t_1, t_2) : Q'}$$

The following conditions hold, where we use the notations $P(j) = (p_i^j)_{i \in I(\Gamma_1)}$, $\bar{P}(j) = (\bar{p}_i^j)_{i \in I(\Gamma_1)}$, $R(i) = (r_j^i)_{j \in I(\Gamma_2)}$, and $\bar{R}(i) = (\bar{r}_j^i)_{j \in I(\Gamma_2)}$

$$\forall j \in I(\Gamma_2) \quad \pi_{\Gamma_1}^j(Q) = P(j) \quad \forall i \in I \quad \pi_y^i(S) = \bar{R}(i) \quad \forall i, j \quad \bar{p}_i^j = r_j^i . \quad (8)$$

By induction hypothesis, we have (i) $\Phi(\sigma \uplus \rho, x, y : R) - \Phi(v : Q') \geq m_0$, (ii) for all $j \in I(\Gamma_2)$: $\Phi(\sigma, \Gamma_1 : P(j)) - \Phi(v_1 : \bar{P}(j)) \geq m_1$, and (iii) for all $i \in I$: $\Phi(\sigma, \Gamma_2 : R(i)) - \Phi(v_2 : \bar{R}(i)) \geq m_2$. Let $\vec{x} := x_1, \dots, x_n$, where $\text{Var}(t_1) = \{x_1, \dots, x_n\}$ and let $\vec{y} := y_1, \dots, y_n$ with $\text{Var}(t_2) = \{y_1, \dots, y_n\}$. The theorem follows by a straightforward calculation:

$$\begin{aligned} \Phi(\sigma, \Gamma_1, \Gamma_2 : Q) &= \sum_{i \in I(\Gamma_1), j \in I(\Gamma_2)} q_{(i,j)} \cdot p_i(\vec{x}\sigma) \cdot p_j(\vec{y}\sigma) \\ &= \sum_{j \in I(\Gamma_2)} p_j(\vec{y}\sigma) \cdot \left(\sum_{i \in I(\Gamma_1)} p_i^j \cdot p_i(\vec{x}\sigma) \right) \\ &\geq \sum_{j \in I(\Gamma_2)} p_j(\vec{y}\sigma) \cdot \left(\sum_{i \in I} \bar{p}_i^j \cdot p_i(v_1) \right) + m_1 \\ &= \sum_{i \in I} p_i(v_1) \cdot \left(\sum_{j \in I(\Gamma_2)} r_j^i \cdot p_j(\vec{y}\sigma) \right) + m_1 \\ &\geq \sum_{i \in I} p_i(v_1) \cdot \left(\sum_{j \in I} \bar{r}_j^i \cdot p_j(v_2) \right) + m_1 + m_2 \\ &\geq \sum_{i \in I} q'_i \cdot p_i(v) + \sum_{i=0}^2 m_i = \Phi(v : Q') + \sum_{i=0}^2 m_i . \end{aligned}$$

Here we tacitly employ the conditions (8) together with the induction hypothesis which is employed in line 3, 5, and 6. \blacktriangleleft

The following corollary to the theorem is immediate.

► **Corollary 17.** *Assume the conditions of the theorem. If additionally for all values v and annotations Q , $\Phi(v : Q) \in \mathcal{O}(n^k)$, where $n = |v|$, then $\text{rc}_{\mathcal{R}}(n) \in \mathcal{O}(n^k)$.*

► **Remark.** Recall that we restrict to completely defined, orthogonal constructor TRSs. It is not difficult to see that Theorem 16 (and its corollary) generalise to the case where completely definedness is dropped. While completely definedness is essential for the correctness of the

big-step semantics presented in Figure 1, the proof of Theorem 16 extends with relative ease. The induction on the length of $\Pi: \sigma \stackrel{m}{\vdash} t \Rightarrow v$ is replaced by an induction on the length of an innermost derivation $D: t\sigma \rightarrow_{\mathcal{R}}^+ v$.

Finally, we consider the motivating TRS $\mathcal{R}_{\text{dyade}}$. Based on the above example (page 251) it remains to consider the remaining four rules of $\mathcal{R}_{\text{dyade}}$:

$$\begin{aligned} \text{mult}(n, []) &\rightarrow [] & \text{mult}(n, x :: xs) &\rightarrow (n \times x) :: \text{mult}(n, xs) \\ \text{dyade}([], ls) &\rightarrow [] & \text{dyade}(x :: xs, ls) &\rightarrow \text{mult}(x, ls) :: \text{dyade}(xs, ls) . \end{aligned}$$

We consider the following automata \emptyset , \mathcal{A} , \mathcal{B} and \mathcal{C} , where \emptyset denotes the canonical automata for $\{0, s, [], ::\}$ and \mathcal{A} is defined as on page 251. The definition of \mathcal{B} and \mathcal{C} is given below.

$$\begin{aligned} \mathcal{B}: \quad 0 &\rightarrow \emptyset & s(\emptyset) &\rightarrow \emptyset & [] &\rightarrow \emptyset & \emptyset :: \emptyset &\rightarrow \beta & \emptyset :: \beta &\rightarrow \beta \\ \mathcal{C}: \quad 0 &\rightarrow \emptyset & s(\emptyset) &\rightarrow \emptyset & s(\emptyset) &\rightarrow \alpha & s(\alpha) &\rightarrow \alpha & [] &\rightarrow \emptyset \\ & & \alpha :: \emptyset &\rightarrow \gamma & \emptyset &:: \gamma &\rightarrow \gamma . \end{aligned}$$

Note that $p_{\emptyset}(v) = 1$, $p_{\mathcal{A}}(\mathbf{n}) = n$, $p_{\mathcal{B}}(l) = |l|$, and $p_{\mathcal{C}}(l) = \sum_{i=1}^m n_i$, where $l = [\mathbf{n}_1, \dots, \mathbf{n}_m]$.

We make use of a similar denotation of the annotations as in the example on page 251 and set $\text{mult}: M \rightarrow (m'_{\emptyset, \emptyset})$, where $M = (m_{\mathcal{A}_1, \mathcal{A}_2})_{\mathcal{A}_1 \in \{\emptyset, \mathcal{A}\}, \mathcal{A}_2 \in \{\emptyset, \mathcal{B}, \mathcal{C}\}}$ and similarly for $\text{dyade}: D \rightarrow (d'_{\emptyset, \emptyset})$, where $D = (d_{\mathcal{A}_1, \mathcal{A}_2})_{\mathcal{A}_1 \in \{\emptyset, \mathcal{B}, \mathcal{C}\}, \mathcal{A}_2 \in \{\emptyset, \mathcal{B}, \mathcal{C}\}}$. Thus we can assert the signature of mult and dyade as follows:

$$\text{mult}: \begin{pmatrix} 1 & 0 \\ 2 & 1 \\ 0 & 1 \end{pmatrix} \rightarrow (0) \quad \text{dyade}: \begin{pmatrix} 1 & 2 & 0 \\ 0 & 2 & 1 \\ 0 & 0 & 1 \end{pmatrix} \rightarrow (0) .$$

Considering $\text{dyade}(x :: xs, ls) \rightarrow \text{mult}(x, ls) :: \text{dyade}(xs, ls)$, we study the effects of the additive shift on D . Let $\bar{D} := \triangleleft_{::}(D) - 1$ such that $\bar{D} = (\bar{d}_{\mathcal{A}_1, \mathcal{A}_2, \mathcal{A}_3})_{\mathcal{A}_1 \in \{\emptyset, \mathcal{A}\}, \mathcal{A}_2, \mathcal{A}_3 \in \{\emptyset, \mathcal{B}, \mathcal{C}\}}$. For $\mathcal{A}_1 = \mathcal{A}$, then $\bar{d}_{\mathcal{A}, \mathcal{A}_2, \mathcal{A}_3}$ vanishes almost everywhere, but $\bar{d}_{\mathcal{A}, \emptyset, \mathcal{B}} = d_{\mathcal{C}, \mathcal{B}} = 1$ and $\bar{d}_{\mathcal{A}, \emptyset, \mathcal{C}} = d_{\mathcal{C}, \mathcal{C}} = 1$. To ease the presentation, we ignore these positive annotations and only consider the restricted annotation $(\bar{d}_{\emptyset, \mathcal{A}_2, \mathcal{A}_3})_{\mathcal{A}_2, \mathcal{A}_3 \in \{\emptyset, \mathcal{B}, \mathcal{C}\}}$. This annotation is again representable as a matrix and typability of the rule follows, as we can decompose the matrix suitably:

$$\begin{pmatrix} 2 & 2 & 0 \\ 2 & 3 & 1 \\ 0 & 1 & 1 \end{pmatrix} = \begin{pmatrix} 1 & 2 & 0 \\ 0 & 2 & 1 \\ 0 & 0 & 1 \end{pmatrix} + \begin{pmatrix} 1 & 0 & 0 \\ 2 & 1 & 0 \\ 0 & 1 & 0 \end{pmatrix} .$$

In order to estimate $\Phi(\sigma, xs, ys: D)$ for arbitrary σ we analyse the base functions of the considered automata. Thus the analysis yields the essentially optimal bound of the execution of $\text{dyade}(ls_1, ls_2)$ as a multiplicative bound in the sum of the values of the first and second list $ls_1 (= xs\sigma)$ and $ls_2 (= ys\sigma)$, respectively, together with a linear factor.

5 Conclusion

We have presented a novel amortised resource analysis in the context of term rewrite systems. The method is represented in an inference system akin to a type system and can give rise to polynomial bounds on the innermost runtime complexity of the analysed rewrite system. The crucial feature of the inference system is the admittance of multivariate bounds in the

context of arbitrary data structures in a completely uniform way. This extends our earlier univariate resource analysis of typed term rewrite systems and continues our program of applying automated amortised resource analysis to rewriting.

We already briefly commented on the differences of the here presented study to our earlier work [12] in the introduction. As far as we can tell this and the present result are currently the only attempts to lift amortised cost analysis to rewriting or provide such a study in the context of arbitrary constructor-defined datastructures. Hoffmann and Shao provide in [11] a multivariate amortised analysis of integers and arrays that extend upon [10]. These language extensions are also provided in RaML. However, the treatment still appears to be ad-hoc and does not provide a similar uniform framework than ours. Further we mention some general work on automated resource analysis. Albert et al. [1] underlies COSTA, an automated tool for the resource analysis of Java programs. Sinn et al. provide in [15] related approaches for the runtime complexity analysis of C programs, incorporated into LOOPUS. Very recently Brockschmidt et al. [5] have provided a runtime complexity analysis of integer programs, taking also size considerations into account. Basic steps for a modular complexity framework for rewrite systems have been established in [3]. Finally, the RaML prototype [10] provides an automated potential-based resource analysis for various resource bounds of functional programs. For term rewriting AProVE [7] and TCT [2] are the most powerful tools for complexity analysis of rewrite systems as witnessed during last year's termination competition.¹

In future work we will clarify the automatability of the method. We expect that by restricting the number of states and the format of those tree automata \mathcal{A} , whose annotation do not vanish, we can reduce inference of annotations to linear constraint solving in much the same way as in [10]. More challenging would be a combination of linear programming and combinatorial constraint solving to infer the best possible structure of automata.

References

- 1 E. Albert, P. Arenas, S. Genaim, G. Puebla, and G. Román-Díez. Conditional termination of loops over heap-allocated data. *Sci. Comput. Program.*, 92:2–24, 2014.
- 2 M. Avanzini and G. Moser. Tyrolean complexity tool: Features and usage. In *Proc. 24th RTA*, volume 21 of *LIPICs*, pages 71–80, 2013.
- 3 M. Avanzini and G. Moser. A combination framework for complexity. *IC*, 2015. To appear.
- 4 F. Baader and T. Nipkow. *Term Rewriting and All That*. Cambridge University Press, 1998.
- 5 M. Brockschmidt, F. Emmes, S. Falke, C. Fuhs, and J. Giesl. Alternating runtime and size complexity analysis of integer programs. In *Proc. 20th TACAS*, volume 8413 of *LNCS*, pages 140–155, 2014.
- 6 H. Comon, M. Dauchet, R. Gilleron, C. Löding, F. Jacquemard, D. Lugiez, S. Tison, and M. Tommasi. *Tree automata techniques and applications*, 2007.
- 7 J. Giesl, M. Brockschmidt, F. Emmes, F. Frohn, C. Fuhs, C. Otto, M. Plücker, P. Schneider-Kamp, T. Ströder, S. Swiderski, and R. Thiemann. Proving termination of programs automatically with AProVE. In *Proc. 7th IJCAR*, volume 8562 of *LNCS*, pages 184–191, 2014.
- 8 N. Hirokawa and G. Moser. Automated complexity analysis based on context-sensitive rewriting. In *Proc. of Joint 25th RTA and 12th TLCA*, volume 8560 of *LNCS*, pages 257–271, 2014.

¹ See <http://nfa.imn.htwk-leipzig.de/termcomp/competition/20>.

- 9 J. Hoffmann, K. Aehlig, and M. Hofmann. Multivariate amortized resource analysis. *TOPLAS*, 34(3):14, 2012.
- 10 J. Hoffmann, K. Aehlig, and M. Hofmann. Resource aware ML. In *Proc. 24th CAV*, volume 7358 of *LNCS*, pages 781–786, 2012.
- 11 J. Hoffmann and Z. Shao. Type-based amortized resource analysis with integers and arrays. In *Proc. 12th FLOPS*, volume 8475 of *LNCS*, pages 152–168, 2014.
- 12 M. Hofmann and G. Moser. Amortised resource analysis and typed polynomial interpretations. In *Proc. of Joint 25th RTA and 12th TLCA*, volume 8560 of *LNCS*, pages 272–286, 2014.
- 13 H. Seidl. On the finite degree of ambiguity of finite tree automata. *Acta Informatica*, 26(6):527–542, 1989.
- 14 H. Seidl. Ambiguity, Validity, and Costs. Technical report, TU München, 1992. Habilitation Thesis.
- 15 M. Sinn, F. Zuleger, and H. Veith. A simple and scalable static analysis for bound analysis and amortized complexity analysis. In *Proc. 26th CAV*, volume 8559 of *LNCS*, pages 745–761, 2014.
- 16 TeReSe. *Term Rewriting Systems*, volume 55 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2003.

Termination of Dependently Typed Rewrite Rules

Jean-Pierre Jouannaud^{1,2} and Jianqi Li²

1 LIX, École Polytechnique, INRIA, and Université Paris-Sud, France

2 School of Software, Tsinghua University, NLIST, Beijing, China

Abstract

Our interest is in automated termination proofs of higher-order rewrite rules in presence of dependent types modulo a theory T on base types. We first describe an original transformation to a type discipline without type dependencies which preserves non-termination. Since the user must reason on expressions of the transformed language, we then introduce an extension of the computability path ordering CPO for comparing dependently typed expressions named DCPO. Using the previous result, we show that DCPO is a well-founded order, behaving well in practice.

1998 ACM Subject Classification F.4.1 Mathematical Logic, F.4.2 Other Rewriting Systems

Keywords and phrases rewriting, dependent types, strong normalization, path orderings

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.257

1 Introduction

This paper addresses the problem of (semi-)automating termination proofs for typed higher-order calculi defined by rewrite rules. Since many automated techniques exist for showing termination of simply typed higher-order rewrite rules, our first approach is to reduce the former to the latter.

To this end, we introduce a non-termination preserving transformation from dependently typed algebraic λ -terms to simply typed algebraic λ -terms. Unlike the transformation used for showing strong normalization of LF [13], the present one uses algebraic symbols and type constructors in an essential way. Dependently typed rewrite rules can then be shown terminating via the transformation. The user can therefore benefit from all existing tools allowing to check termination of higher-order rewrite rules. The drawback is that these tools will operate on the transformed rules.

Among all termination proof techniques, we favour the one reducing termination proofs to ordering comparisons between lefthand and righthand sides of rules. These comparisons require well-founded orders on typed algebraic λ -terms which are stable by context application and substitution instance. CPO is such an order on *simply typed* algebraic λ -terms, defined recursively on the structure of the compared terms [9]. CPO is indeed well-founded on *weakly polymorphic* λ -terms, the familiar ML-discipline for which quantifiers on types can only occur in prefix position. A recent extension of core CPO to appear in LMCS handles inductive types, constructors possibly taking functional arguments, and function symbols smaller than application and abstraction.

We formulate here a new extension DCPO of CPO for dependently typed algebraic λ -terms. DCPO is then viewed as an infinite set of dependently typed rewrite rules which are shown terminating by checking the transformed rules with CPO. It follows that DCPO is a well-founded order of the set of dependently typed λ -terms, whose syntax-directed comparisons require little input from the user, in the form of a precedence on the algebraic symbols used in the rules. DCPO is our answer for practice.



© Jean-Pierre Jouannaud and Jianqi Li;
licensed under Creative Commons License CC-BY

13th International Conference on Typed Lambda Calculi and Applications (TLCA'15).

Editor: Thorsten Altenkirch; pp. 257–272



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Variables: $\frac{x:\sigma \in \Gamma}{\Gamma \vdash_{\Sigma} x:\sigma}$	Abstraction: $\frac{\Gamma, x:\sigma \vdash_{\Sigma} t:\tau}{\Gamma \vdash_{\Sigma} (\lambda x:\sigma.t):\sigma \rightarrow \tau}$	Application: $\frac{\Gamma \vdash_{\Sigma} s:\sigma \rightarrow \tau \quad \Gamma \vdash_{\Sigma} t:\sigma}{\Gamma \vdash_{\Sigma} @(s,t):\tau}$	Functions: $\frac{f^n:\bar{\sigma} \rightarrow \sigma \in \mathcal{F} \quad \Gamma \vdash_{\Sigma} \bar{t}:\bar{\sigma}}{\Gamma \vdash_{\Sigma} f(\bar{t}):\sigma}$
--	--	--	---

■ **Figure 1** Type system for monomorphic higher-order algebras.

Dependent programming has become a major trend in recent years [23, 16]. In practice, many types depend on natural numbers. Typing dependent definitions requires then a convertibility relation T including arithmetic laws [20], see Example 3.4. Our results allow for *dependent types modulo T* .

Sections 4 and 5 describe the non-termination preserving transformation and DCPO.

2 Higher-Order Algebras $\lambda_{\Sigma}^{\rightarrow}$

We assume a *signature* $\Sigma = \mathcal{S} \uplus \mathcal{F}$ of *sort symbols* in \mathcal{S} and *function symbols* in \mathcal{F} . The set $\mathcal{T}_{\Sigma}^{\rightarrow}$ of *simple types* (in short, *types*) is generated by the grammar $\sigma, \tau := a \in \mathcal{S} \mid \sigma \rightarrow \tau$. The (arrow) type constructor \rightarrow associates to the right. The output sort of a type σ is itself if $\sigma \in \mathcal{S}$ and the output sort of τ if $\sigma = \nu \rightarrow \tau$. We use σ, τ, μ, ν for simple types.

Function symbols are meant to be algebraic operators upper-indexed by their fixed *arity* n . *Function declarations* are written $f^n : \sigma_1 \rightarrow \dots \rightarrow \sigma_n \rightarrow \sigma$ (in short, $f : \bar{\sigma} \rightarrow \sigma$), where $\bar{\sigma}$ and σ are the *input* and *output types* of f^n . We use f^n, g^m for function symbols, possibly omitting m, n . $\lambda x : \sigma.s$, $@(s, t)$ and $f^n(t_1, \dots, t_n)$ (or $f(\bar{t})$) are an *abstraction*, an *application*, and a *pre-algebraic* raw term. $f^0()$ is identified with f . We use x, y, z for variables, s, t, u, v, w, l, r for raw terms, $\text{FV}(s)$ for the set of free variables of s and $|s|$ for the *size* of s .

Raw terms are seen as finite labeled trees by considering $\lambda x : \sigma.s$, for each $x : \sigma$, as a unary abstraction operator taking s as argument to construct the raw term $\lambda x : \sigma.s$. We abbreviate abstraction operators by λ . *Positions* are strings of strictly positive integers. We use i, j for positive integers, p, q for arbitrary positions. The empty string Λ is the *root* or *head* position and \cdot is string concatenation. $\text{Pos}(t)$ is the set of positions of t .

Given a raw term s , $s|_p$ and $s|_p$ denote respectively the *symbol* and *subterm* of s at position p . For example, $(\lambda x : \sigma.u)|_1 = u$. The result of replacing the subterm $s|_p$ by the term t is written $s[t]_p$. A context term $s[x]_p$, in short $s[\]_p$ or even $s[\]$, is a term s in which x is a fresh variable, called *hole*, occurring at position p . All these notions extend as expected to a set P of disjoint positions, writing $s[t]_P$ for replacement of all terms in $s|_P$ by a single term t , and $s[\bar{x}]_P$ for a context with many holes.

An *environment* Γ is a finite set of pairs $\{x_1 : \sigma_1, \dots, x_n : \sigma_n\}$ where x_i is a variable, σ_i is a type, and $x_i \neq x_j$ for $i \neq j$. $\text{FV}(\Gamma) = \{x_1, \dots, x_n\}$ is the set of variables of Γ . Our typing judgements are written as $\Gamma \vdash_{\Sigma} s : \sigma$. A raw term s has type σ in the environment Γ if the judgement $\Gamma \vdash_{\Sigma} s : \sigma$ is provable in the inference system given in Figure 1. Typable raw terms are called *algebraic λ -terms* (in short, *terms* or *objects*), and their set is denoted by $\lambda_{\mathcal{F}}$ (or $\mathcal{O}_{\mathcal{F}}^{\rightarrow}$). Objects have a unique type in a given, possibly omitted, environment.

Substitutions are type-preserving homomorphisms avoiding captures, see for example [2], written here in postfix form, using the notation $\{x_i \mapsto s_i\}_i$. The congruence on terms generated by renaming the free variable x in s by the *fresh* variable $z \notin \text{FV}(\lambda x.s)$ to yield the term $\lambda z.s\{x \mapsto z\}$ is called α -*conversion*, denoted by $=_{\alpha}$. We use γ, θ for substitutions.

A higher-order rewrite rule is a quadruple $\Delta \vdash_{\Sigma} l \rightarrow r : \sigma$ made of lefthand and righthand side terms l, r , and possibly omitted environment Δ and type σ . Rules β and η are particular

rewrite rule schemas. Reductions are defined as usual, and are a particular case of reductions in presence of dependent types, see Definition 3.3. A *higher-order reduction ordering* \succeq is a quasi-order on terms satisfying: (i) its strict part \succ is well-founded ; (ii) its equivalence is a congruence ; (iii) *monotonicity*: $s \succ t$ implies $u[s]_p \succ u[t]_p$ (assuming typability); (iv) *stability*: $s \succ t$ implies $s\gamma \succ t\gamma$ for all substitutions γ ; and (v) *functionality*: $\longrightarrow_\beta \cup \longrightarrow_\eta \subseteq \succ$.

Given a set E the notation \bar{s} shall be used for a list, multiset, or set of elements of E . Given a binary relation \succ on E , we use $\bar{s} \succ_{lex} \bar{t}$ and $\bar{s} \succ_{mul} \bar{t}$ for its lexicographic and multiset extensions respectively. We use $s \succ \bar{t}$ for $(\forall t \in \bar{t}) s \succ t$, and $\bar{s} \succ t$ for $(\exists s \in \bar{s}) s \succ t$.

A rewrite relation generated by a set of rules $R \cup \{\beta, \eta\}$ can be proved terminating by checking whether $l \succ r$ for all rules in R with some higher-order reduction ordering \succ [14]. CPO is such a higher-order reduction ordering based on three ingredients [9]:

- an order \geq^\rightarrow on simple types, whose strict part $>^\rightarrow$ satisfies
 - (i) *well-foundedness*: $>^\rightarrow \cup \{(\sigma \rightarrow \tau, \sigma) \mid \sigma, \tau \in \mathcal{T}_{\mathcal{S}^\rightarrow}\}$ is well-founded ;
 - (ii) *right arrow subterm*: $\sigma \rightarrow \tau >^\rightarrow \tau$;
 - (iii) *preservation*: $\sigma \rightarrow \tau \Rightarrow \nu$ iff $\nu = \sigma' \rightarrow \tau'$, $\sigma \Rightarrow \sigma'$, $\tau \Rightarrow \tau'$ and
 - (iv) *decreasingness*: $\sigma \rightarrow \tau >^\rightarrow \nu$ implies $\tau \geq^\rightarrow \nu$ or $\nu = \sigma \rightarrow \mu$ and $\tau >^\rightarrow \mu$.
- a quasi-order $\geq_{\mathcal{F}}$ on $\mathcal{F} \cup \{\@, \lambda\} \cup \mathcal{X}$ called *precedence* s.t.: (i) its strict part $>_{\mathcal{F}}$ restricts to $\mathcal{F} \cup \{\@, \lambda\}$, is well-founded, and satisfies $(\forall f \in \mathcal{F}) f >_{\mathcal{F}} \@ >_{\mathcal{F}} \lambda$; (ii) its equivalence $=_{\mathcal{F}}$ contains pairs in $\mathcal{F} \times \mathcal{F}$ and all pairs $\{(x, x) \mid x \in \mathcal{X}\}$.
- $(\forall f \in \mathcal{F} \cup \{\@, \lambda\})$, a status operator $_f \in \{lex, mul\}$ in postfix index position such that $_f \@ = mul$. Equivalent symbols have the same status.

The following auxiliary relations are used to define CPO [9] :

- $s \geq^X t$ iff $s =_\alpha t$ or $s >^X t$, for a set of variables X disjoint from $FV(s)$, is the *main order*;
- $s >^X \bar{t}$ and $\bar{s} >^X t$ defined respectively as $(\forall v \in \bar{t}) s >^X v$ and $(\exists u \in \bar{s}) u >^X t$;
- $s : \sigma >^X t : \tau$ (resp., $s : \sigma \geq^X t : \tau$) for $s >^X t$ (resp., $s \geq^X t$) and $\sigma \geq^\rightarrow \tau$;
- we are interested in the *typed order* $s : \sigma >^\emptyset t : \tau$, written $s : \sigma > t : \tau$.

► **Definition 2.1.** Given $\Gamma \vdash_{\mathcal{F}} s : \sigma$ and $\Gamma \vdash_{\mathcal{F}} t : \tau$, $s >^X t$ iff either

$t \in X$ and $s \notin \mathcal{X}$	VAR
$s = \lambda x : \mu. u$ and $u\{x \mapsto z\} : \nu \geq^X t : \tau$	SUBT λ
$s = f(\bar{s})$, $t = \lambda x : \mu. v$ and $s >^{X \cup \{z : \mu\}} v\{x \mapsto z\}$	$\mathcal{F}\lambda$
$s = \@ (u, w)$, $t = \lambda x : \mu. v$, $x \notin FV(v)$ and $s >^X v$	$\@ \lambda$
$s = \lambda x : \mu. u$, $t = \lambda y : \mu. v$ and $u\{x \mapsto z\} >^X v\{y \mapsto z\}$	$\lambda \lambda$
$s = \@ (\lambda x : \mu. u, w)$ and $u\{x \mapsto w\} \geq^X t$	BETA
$s = \lambda x : \mu. \@ (u, x)$, $x \notin FV(u)$ and $u \geq^X t$	ETA
otherwise $s = f(\bar{s})$, $t = g(\bar{t})$ with $f, g \in \mathcal{F} \cup \{\@, \lambda\} \cup \mathcal{X}$, and either	
SUBT $\bar{s} : \bar{\sigma} \geq t : \tau$ PREC $f >_{\mathcal{F}} g$ and $s >^X \bar{t}$ STAT $f =_{\mathcal{F}} g$, $s >^X \bar{t}$ and $\bar{s} : \bar{\sigma} >_f \bar{t} : \bar{\tau}$	

This definition of CPO is organized differently from [9] to be more compact. The last three cases originating in Dershowitz' recursive path ordering [12] describe the normal behaviour of head symbols, whether or not in \mathcal{F} . Note here that $>_f$ is the status extension (lexicographic or multiset) of the order $>$. The first 7 cases describe other behaviours, either specific (var, beta, eta), or using explicit α -conversion for the others. In its recursive call, case $\mathcal{F}\lambda$ increases the set X of upper variables, while other recursive calls either keep it unchanged or reset it to \emptyset . Relaxations of these recursive calls are indeed ill-founded.

► **Theorem 2.2** ([9]). $>^+$ is a higher-order reduction ordering.

3 Dependent algebras λ_{Σ}^{Π}

We move to a calculus with dependent types inspired by Edinburgh's LF [13], an extension of the simply typed λ -calculus which can be seen as a formal basis for dependently typed programming languages and their formal study as done with Elf [18].

In higher-order algebras, types are typable by a single, usually omitted constant TYPE. In dependent algebras, as in other type theories, types (called *type families*) are typed by *kinds*, TYPE being one of them, which describe their functional structure. Let \mathcal{S}, \mathcal{F} and \mathcal{V} be pairwise disjoint sets of respectively *type symbols*, *algebraic function symbols*, and *variables*. Algebraic function symbols in \mathcal{F} may carry an arity upper-indexing their name. Type symbols in \mathcal{S} are curried constants which kind may be functional. We use respectively f and a for a typical function or type symbol. Raw terms are given by the following grammar:

$$\begin{aligned} \text{Kinds } K & := \text{TYPE} \mid \Pi x : A. K \mid \text{Types } A, B & := a \mid \Pi x : A. B \mid \lambda x : A. B \mid A \ N \\ \text{Objects } M, N & := x \mid \lambda x : A. M \mid M \ N \mid f^n(M_1, \dots, M_n) \end{aligned}$$

λ and Π are LF's binders. Notation $=_{\alpha}$ stands for similarity. Here, λ and Π are binary operators, whose arguments are a type $((\lambda x : A. u)|_1 = A)$, and the body of the abstraction or product $((\lambda x : A. u)|_2 = u)$. Both may originate computations. We write f for $f^0()$.

3.1 Typing judgements

All LF expressions are typed, objects by types, types by kinds, kinds by a special untyped constant KIND (the *universe*) asserting their well-formedness, we also call them *valid*. Five kinds of judgement are recursively defined by the LF type system of Figure 2.

$\vdash_{sig} \Sigma$	Σ is a valid signature
$\Sigma \vdash_{\mathcal{C}} \Gamma$	Γ is a valid context assuming $\vdash_{sig} \Sigma$
$\Sigma; \Gamma \vdash_{\mathcal{K}} K : \text{KIND}$	K is a valid kind assuming $\Sigma \vdash_{\mathcal{C}} \Gamma$
$\Sigma; \Gamma \vdash_{\mathcal{T}} A : K$	type A has kind K assuming $\Sigma; \Gamma \vdash_{\mathcal{K}} K : \text{KIND}$
$\Sigma; \Gamma \vdash_{\mathcal{O}} M : A$	object M has type A assuming $\Sigma; \Gamma \vdash_{\mathcal{T}} A : K$
$\Sigma; \Gamma \vdash_{\mathcal{T} \vee \mathcal{O}} C : D$ and $\Sigma; \Gamma \vdash_{\mathcal{T} \vee \mathcal{K}} C : D$	are self explanatory

Environments pair up a signature and a context. Signatures assign kinds to type symbols and product types to function symbols. Contexts assign product types to variables.

$Env \ \Theta := \Sigma; \Gamma$	where nil is the empty set, $\Pi\{x_i : A_i\}_n.A$ is $\Pi x_1 : A_1. (\dots (\Pi x_n : A_n. A) \dots)$
$Sig \ \Sigma := \text{nil} \mid \Sigma, a : K \mid \Sigma, f^n : \Pi\{x_i : A_i\}_n.A$	
$Con \ \Gamma := \text{nil} \mid \Gamma, x : A$	

In dependent calculi, the order of constants or variables in the environment is determined by their types. This impacts the expression of the so-called substitution lemma:

► **Lemma 3.1.** *Let $\Sigma; \Gamma \vdash_{\mathcal{O}} M : A$ and $\Sigma; \Gamma, x : A, \Gamma' \vdash_{\mathcal{T} \vee \mathcal{O}} s : \sigma$. Then, $\Sigma; \Gamma, \Gamma'\{x \mapsto M\} \vdash_{\mathcal{T} \vee \mathcal{O}} s\{x \mapsto M\} : \sigma\{x \mapsto M\}$.*

Applying the substitution lemma several times introduces an order on the application of elementary substitutions. We use the notation $M \circ^n \{x_i \mapsto N_i\}_i$ to denote the sequential application to M of the *elementary* substitutions $\{x_1 \mapsto N_1\}, \dots, \{x_n \mapsto N_n\}$ in this order. We use the word *dependent substitution* to stress this sequential behaviour of substitutions.

Given a valid signature Σ and a context Γ valid in Σ , the set λ_{Σ} of valid expressions, called *terms* is the (disjoint) union of the sets $\mathcal{K}_{\Sigma}^{\Pi}$ of valid *kinds*, $\mathcal{T}_{\Sigma}^{\Pi}$ of valid *types* and $\mathcal{O}_{\Sigma}^{\Pi}$ of

Signatures

$$\begin{array}{c}
\text{[EMPTY]} \frac{}{\vdash_{sig} \mathbf{nil}} \quad \text{[TCONST]} \frac{\Sigma; \mathbf{nil} \vdash_{\mathcal{K}} K : \text{KIND}}{\vdash_{sig} \Sigma, a : K} \quad a \notin \text{dom}(\Sigma) \\
\text{[CONST]} \frac{\vdash_{sig} \Sigma \quad \Sigma; \{x_i : A_i\}_k \vdash_{\mathcal{T}} A_{k+1} (k = 0..n) : \text{TYPE}}{\vdash_{sig} \Sigma, f^n : \Pi\{x_i : A_i\}_n. A_{n+1}} \quad f \notin \text{dom}(\Sigma)
\end{array}$$

Contexts

$$\begin{array}{c}
\text{[EMPTY]} \frac{\vdash_{sig} \Sigma}{\Sigma \vdash_{\mathcal{C}} \mathbf{nil}} \quad \text{[VAR]} \frac{\Sigma \vdash_{\mathcal{C}} \Gamma \quad \Sigma; \Gamma \vdash_{\mathcal{T}} A : \text{TYPE}}{\Sigma \vdash_{\mathcal{C}} \Gamma, x : A} \quad x \notin \text{dom}(\Gamma)
\end{array}$$

Kinds

$$\begin{array}{c}
\text{[UNIV]} \frac{\Sigma \vdash_{\mathcal{C}} \Gamma}{\Sigma; \Gamma \vdash_{\mathcal{K}} \text{TYPE} : \text{KIND}} \quad \text{[PROD]} \frac{\Sigma; \Gamma, x : A \vdash_{\mathcal{K}} K : \text{KIND}}{\Sigma; \Gamma \vdash_{\mathcal{K}} \Pi x : A. K : \text{KIND}}
\end{array}$$

Types

$$\begin{array}{c}
\text{[AX]} \frac{\Sigma \vdash_{\mathcal{C}} \Gamma}{\Sigma; \Gamma \vdash_{\mathcal{T}} a : K} \quad a : K \in \Sigma \quad \text{[ABS]} \frac{\Sigma; \Gamma, x : A \vdash_{\mathcal{T}} B : K}{\Sigma; \Gamma \vdash_{\mathcal{T}} \lambda x : A. B : \Pi x : A. K} \\
\text{[APP]} \frac{\Sigma; \Gamma \vdash_{\mathcal{T}} A : \Pi x : B. K \quad \Sigma; \Gamma \vdash_{\mathcal{O}} M : B}{\Sigma; \Gamma \vdash_{\mathcal{T}} A M : K\{x \mapsto M\}} \quad \text{[PROD]} \frac{\Sigma; \Gamma, x : A \vdash_{\mathcal{T}} B : \text{TYPE}}{\Sigma; \Gamma \vdash_{\mathcal{T}} \Pi x : A. B : \text{TYPE}} \\
\text{[CONV]} \frac{\Sigma; \Gamma \vdash_{\mathcal{T}} A : K \quad \Sigma; \Gamma \vdash_{\mathcal{K}} K' : \text{KIND}}{\Sigma; \Gamma \vdash_{\mathcal{T}} A : K'} \quad K \equiv K'
\end{array}$$

Objects

$$\begin{array}{c}
\text{[AX]} \frac{\Sigma \vdash_{\mathcal{C}} \Gamma \quad x : A \in \Gamma}{\Sigma; \Gamma \vdash_{\mathcal{O}} x : A} \quad \text{[FUN]} \frac{\Sigma; \Gamma \vdash_{\mathcal{O}} M_i : A_i \{x_1 \mapsto M_1, \dots, x_{i-1} \mapsto M_{i-1}\}}{\Sigma; \Gamma \vdash_{\mathcal{O}} f^n(M_1, \dots, M_n) : A \circ^n \{x_i \mapsto M_i\}_i} \\
\text{where } f^n : \Pi\{x_i : A_i\}_n. A \in \Sigma \\
\text{[ABS]} \frac{\Sigma; \Gamma, x : A \vdash_{\mathcal{O}} M : B}{\Sigma; \Gamma \vdash_{\mathcal{O}} \lambda x : A. M : \Pi x : A. B} \quad \text{[APP]} \frac{\Sigma; \Gamma \vdash_{\mathcal{O}} M : \Pi x : A. B \quad \Sigma; \Gamma \vdash_{\mathcal{O}} N : A}{\Sigma; \Gamma \vdash_{\mathcal{O}} M N : B\{x \mapsto N\}} \\
\text{[CONV]} \frac{\Sigma; \Gamma \vdash_{\mathcal{T} \vee \mathcal{O}} s : \sigma \quad \Sigma; \Gamma \vdash_{\mathcal{T} \vee \mathcal{K}} \sigma' : \text{TYPE/KIND}}{\Sigma; \Gamma \vdash_{\mathcal{T} \vee \mathcal{O}} s : \sigma'} \quad \sigma \equiv \sigma'
\end{array}$$

■ Figure 2 LF Typing rules.

valid *objects* (possibly abbreviated as $\mathcal{K}, \mathcal{T}, \mathcal{O}$). Our presentation of LF performs classically the necessary sanitary checks when forming the signatures and contexts and only those.

\equiv denoting the *convertibility relation*, a congruence discussed next in more details, that is generated by $\beta\eta$ -conversion on similar terms on the one hand and on the other hand by an additional arbitrary congruence between object terms stable by substitution (possibly identifying them all [4]). By its definition as a congruence, convertibility respects our syntactic categories, objects, types and kinds.

Our dependently typed calculus is referred to as λ_{Σ}^{Π} to stress the signature Σ , or simply LF.

Lexicography: we use K for kinds, A,B,C,D for types, M,N for objects, s,t,u,v,w,l,r for (objects or types), σ, τ, μ, ν for (type or kinds), and γ, θ for substitutions.

3.2 Dependent rewriting and convertibility

In LF, the usual rules of the λ -calculus apply at the object and type levels, making four different rules generating a congruence called *$\beta\eta$ -convertibility*:

$$\begin{array}{l} \text{beta: } (\lambda x:A.B) M \rightarrow_{\beta\tau} B\{x \mapsto M\} \quad \left| \quad (\lambda x:A.M) N \rightarrow_{\beta\mathcal{O}} M\{x \mapsto N\} \right. \\ \text{eta: } \lambda x:A.(B x) \rightarrow_{\eta\tau} B \text{ IF } x \notin \text{FV}(B) \quad \left| \quad \lambda x:A.(M x) \rightarrow_{\eta\mathcal{O}} M \text{ IF } x \notin \text{FV}(M) \right. \end{array}$$

Convertibility plays a key role for typing via CONV. In LF, convertibility is defined as the congruence generated by $\beta\eta$ -reductions. In reality, convertibility must be strengthened on object level terms in order to type most examples. This problem has been considered in the framework of the calculus of constructions with the Calculus on Inductive Constructions [17], the Calculus of Algebraic Constructions [7] and the Calculus of Constructions Modulo Theory [20, 3], for which convertibility includes respectively: primitive recursion at higher type generated by the user's inductive types; the user's higher-order rules; and a decidable object-level first-order theory like Presburger arithmetic. These frameworks can be restricted to the LF type system seen as a particular case of the calculus of constructions. In the context of LF, it *relates* to the liquid types discipline [19], which shares similar objectives.

We now introduce dependent rewrite rules and rewriting.

► **Definition 3.2.** Given a valid signature Σ , a *plain dependent rewriting system* is a set $\{\Delta_i \vdash l_i : \sigma_i \rightarrow r_i : \tau_i\}_i$ of quintuples made, for every index i , of a context Δ_i , lefthand and righthand side terms l_i, r_i , and terms σ_i, τ_i , s.t. $\text{FV}(r_i) \subseteq \text{FV}(l_i)$, $\Sigma; \Delta_i \vdash_{\mathcal{T}\vee\mathcal{O}} l_i : \sigma_i$ and $\Sigma; \Delta_i \vdash_{\mathcal{T}\vee\mathcal{O}} r_i : \tau_i$ with $\sigma_i \equiv \tau_i$. Δ_i, σ_i and τ_i may be omitted.

► **Definition 3.3 (Dependent rewriting).** Given a rewriting system R , one step rewriting is a relation over terms, written $\Sigma; \Gamma \vdash_R s \longrightarrow^p t$ (in short, $s \longrightarrow_R t$) defined as:

- s and t are both types or objects which are typable under $\Sigma; \Gamma$,
- $\Delta \vdash l \rightarrow r : A \in R$, where $\text{FV}(\Delta) \cap \text{FV}(\Gamma) = \emptyset$.
- $s = s[l \circ \gamma]_p$ and $t = s[r \circ \gamma]_p$, where γ is a dependent substitution wrt to Δ .

A major semantic property expected from rewrite rules is that rewriting preserves types. In presence of dependencies, types are usually preserved up to some congruence defined by the rules themselves, as in the Calculus of Inductive Constructions or the Calculus of Algebraic Constructions [7]. Here, preservation of typing by rewriting, up to *type erasures*, follows from Lemma 4.7.

► **Example 3.4.** Here is a simple example with dependent lists of elements of a given type A . We allow ourselves with some OBJ-like mixfix syntax, using “ $_$ ” for arguments' positions:

$nat, A : \text{TYPE}; List : \Pi m : nat. \text{TYPE};$
 $0 : nat; _ + 1 : \Pi n : nat. nat; _ + _ : \Pi \{m, n : nat\}. nat$
 $cons : \Pi \{n : nat, a : A, l : List\ n\}. List (n + 1)$
 $app : \Pi \{m, n : nat, k : List\ m, l : List\ n\}. List (m + n)$
 $0 + n \rightarrow n \quad app(0, n, nil, l) \rightarrow l$
 $(m + 1) + n \rightarrow (m + n) + 1 \quad app(m + 1, n, cons(m, a, k), l) \rightarrow cons(m + n, a, app(m, n, k, l))$

Using LF's typing rules given in Figure 2, we get:

$$\begin{array}{l}
\{m, n : nat\} \vdash n, 0 + n, (m + 1) + n, (m + n) + 1 : nat \\
\{n : nat, l : List\ n\} \vdash app(0, n, nil, l) : List (0 + n) \\
\{m, n : nat, k : List\ m, l : List\ n\} \vdash app(m + 1, n, cons(m, a, k), l) : List ((m + 1) + n) \\
\{m, n : nat, k : List\ m, l : List\ n\} \vdash cons(m + n, a, app(m, n, k, l)) : List ((m + n) + 1)
\end{array}$$

Typing these rules requires a conversion relation extending $\beta\eta$ -conversion on objects with Presburger arithmetic to identify $List (0 + n)$, $List\ n$ and $List ((m + 1) + n)$, $List ((m + n) + 1)$.

4 Encoding LF in higher-order algebras

We define here a transformation from the source language to a target language which preserves non-termination of arbitrary reductions, not just β -reductions as in LF. Our target language is the simply-typed λ -calculus enriched with function symbols and type constants of Section 2, a choice which has three important advantages: the target vocabulary can be as close as possible from the source vocabulary; the transformation preserves termination as well as non-termination; the transformed rules can be checked for termination by CPO.

The higher-order algebra encoding λ_{Σ}^{Π} will be $\lambda_{\Sigma^{flat}}^{\rightarrow}$, where $\Sigma^{flat} = \mathcal{S}_{flat} \uplus \Sigma_{flat}$ is a higher-order (non-dependent) signature whose two pieces are described next. The set of types in $\lambda_{\Sigma^{flat}}^{\rightarrow}$ is denoted by $\mathcal{T}_{\Sigma^{flat}}^{\rightarrow}$. Terms in $\lambda_{\Sigma^{flat}}^{\rightarrow}$, whose set is denoted by $\lambda_{\Sigma^{flat}}$, are meant to encode LF objects and types in a way which mimics dependently typed computations. We do not encode LF kinds, since computations in kinds are indeed computations on types or objects. Indeed, dependent types will be encoded in $\lambda_{\Sigma^{flat}}^{\rightarrow}$ both as types and as terms.

4.1 Type erasures

Types. Types in $\lambda_{\Sigma^{flat}}^{\rightarrow}$ are arrow types built from the set of sorts $\mathcal{S}_{flat} = \{*\} \cup \{a \mid a : K \in \Sigma\}$.

The new sort $*$ will serve to encode LF product types as terms of sort $*$ in $\lambda_{\Sigma^{flat}}^{\rightarrow}$. LF objects will be encoded as terms whose types are erasures of LF types, as defined next.

Type erasures. The classical *erasing* transformation from families and kinds in λ_{Σ}^{Π} to types in $\lambda_{\Sigma^{flat}}^{\rightarrow}$ eliminates dependencies from objects by replacing product types by arrow types:

► **Definition 4.1.** The *erasing* map $|\cdot| : \mathcal{T} \cup \mathcal{K} \rightarrow \mathcal{T}_{\Sigma^{flat}}^{\rightarrow}$ is defined as:

$$\begin{array}{lll}
(1) |a| = a & (2) |\Pi x : A. B| = |A| \rightarrow |B| & (3) |\lambda x : A. B| = |B| \\
(4) |TYPE| = * & (5) |\Pi x : A. K| = |A| \rightarrow |K| & (6) |A\ N| = |A|
\end{array}$$

Sorts being constants, an induction shows that variables in types are eliminated in rule (6):

► **Lemma 4.2.** *Let $A \in \mathcal{T} \cup \mathcal{K}$. Then $FV(|A|) = \emptyset$.*

► **Corollary 4.3.** *For any $E \in \mathcal{K} \cup \mathcal{T}$, $y \in \mathcal{V}$, and $s \in \mathcal{O}_{\Sigma^{flat}}^{\rightarrow}$, $|E\{y \mapsto s\}| = |E|\{y \mapsto |s|\} = |E|$.*

► **Lemma 4.4** (Conversion equality). *Let $D, E \in \mathcal{T}$ such that $D \equiv E$. Then, $|D| = |E|$.*

Proof. Conversion is a congruence generated by $\beta\eta$ -rewriting and an equivalence $=_{\mathcal{O}}$ on object terms. We show that the property is true of both relations by induction on D .

1. Case $D = a : K$. Then $E = D$.
2. Case $D = \Pi x : A.B : \text{TYPE}$. Then $E = \Pi x : A'.B' : \text{TYPE}$ with $A \equiv A'$ and $B \equiv B'$. By induction hypothesis (and definition of the erasing map).
3. Case $D = \lambda x : A.B : \Pi x : A.K$. There are two cases:
 - (a) $E = \lambda x : A'.B' : \Pi x : A'.K$, with $A \equiv A'$ and $B \equiv B'$. By induction hypothesis.
 - (b) $D = \lambda x : A.(E x) \rightarrow_{\eta} E$ where $x \notin \text{FV}(E)$. Then $|D| = |\lambda x : A.(E x)| = |E x| = |E|$.
3. Case $D = A M : K$, where M is an object. Again two cases:
 - (a) $E = A' M'$, $A \equiv A'$ and $M \equiv M'$. Then, $|D| = |A|$ and $|E| = |A'|$. By induction.
 - (b) $D = (\lambda x : F.G) M \rightarrow_{\beta} G\{x \mapsto M\} = E$. By Corol. 4.3, $|D| = |G| = |G\{x \mapsto M\}| = |E|$.

◀

Erasing is extended to environments $\Sigma; \Gamma$ by: $|\Gamma| = \{x : |A| \mid x : A \in \Gamma\}$ and $|\Sigma| = \{a : |K| \mid a : K \in \Sigma\} \cup \{f^n : |A_1| \rightarrow \dots \rightarrow |A_n| \rightarrow |A| \text{ where } f^n : \Pi\{x_i : A_i\}_n. A \in \Sigma\}$.

Note here that a constructor like $\text{cons} : \Pi n : \text{Nat}, x : \text{Nat}, l : \text{List}(n). \text{List}(n+1) \in \Sigma$ becomes $\text{cons} : \text{Nat} \rightarrow \text{Nat} \rightarrow \text{List} \rightarrow \text{List} \in |\Sigma|$. Eliminating the first (type) argument from cons would easily allow writing rules for which termination is not preserved by the transformation. This does not mean, however, that writing such rules is actually impossible.

4.2 Term flattening

Besides types and function symbols in $|\Sigma|$, the signature Σ_{flat} will contain algebraic symbols, called *flattening constructors*, used to mimic LF's abstraction and product.

$$\Sigma_{\text{flat}} = |\Sigma| \cup \{lo_{\sigma}^2 : * \rightarrow \sigma \rightarrow \sigma, \quad lf_{\sigma}^2 : * \rightarrow \sigma \rightarrow \sigma, \quad \text{pif}_{\sigma}^2 : * \rightarrow (\sigma \rightarrow *) \rightarrow * \mid \sigma \in \mathcal{T}_{\mathcal{S}_{\text{flat}}}^{\rightarrow}\}$$

We may omit subscripts σ, τ in constructor's names, and use lof^2 for $(lo_{\sigma}^2/lf_{\sigma}^2)$. The first argument of the flattening constructors, of type $*$, is the flattening of some dependent type A . The second argument of lo_{σ}^2 is the interpretation of some object $M : A$, hence $\sigma = |A|$, while that of lf_{σ}^2 is the interpretation of a type $A : K$, hence $\sigma = |K|$. Since signatures are monomorphic and σ is arbitrary, the flattening constructors must be indexed by these types. We now define a *flattening* transformation for expressions of LF which are typed in some environment left unspecified to expressions of $\lambda_{\Sigma_{\text{flat}}}^{\rightarrow}$:

► **Definition 4.5.** The *flattening* function $\|_{_}$ from λ_{Σ} to $\lambda_{\Sigma_{\text{flat}}}$ (the context in which the input term is valid is omitted) is defined as

$$\begin{array}{l|l} \|x : A = x : |A| & \|\Pi x : A.B : \text{TYPE} = \text{pif}_{|A|}^2(\|A, \lambda x : |A|. \|B) \\ \|M N : A = @(\|M, \|N) & \|\lambda x : A.M : \Pi x : A.B = \lambda x : |A|. \text{lo}_{|B|}^2(\|A, \|M) \\ \|A N : K = @(\|A, \|N) & \|\lambda x : A.B : \Pi x : A.K = \lambda x : |A|. \text{lf}_{|K|}^2(\|A, \|B) \\ \|a : K = a & \|f^n(M_1, \dots, M_n) : A = f^n(\|M_1, \dots, \|M_n) \end{array}$$

Since flattening is not surjective, we denote by $\|\lambda_{\Sigma} \subset \lambda_{\Sigma_{\text{flat}}}$ its target.

Note that type symbols in \mathcal{S} become both sorts in $\mathcal{S}_{\text{flat}}$ and function symbols in Σ_{flat} , with the same (overloaded) name. This definition obeys the following principles: (i) flattening is a homomorphism, hence commutes with substitutions; (ii) because types may depend on objects, the type information associated with bound variables must be recorded by the encoding to preserve non-termination; (iii) the types of the flattening constructors are compatible with the erasing transformation, which allows to trace the syntactic categories in the transformed world; (iv) the encoding of abstractions and products preserves their variable's binding, but two different encodings are used. The encoding of abstractions is an

abstraction, in order to preserve beta-redexes via the transformation. Nothing like that is needed for product types which cannot be applied, and can therefore be transformed into terms of sort $|\text{TYPE}| = *$. In that case, the abstraction is encapsulated in the flattening constructor. This allows to single out easily products' encodings in the flattened world. Now,

- an LF object $M : A$ is translated as a term $\|M$ of type $|A|$,
- an LF type $A : K$ is translated as both a type $|A|$ and a term $\|A$ of type $|K|$,
- an LF kind $K : \text{KIND}$ is translated as a type $|K|$.

► **Lemma 4.6.** *The following properties hold:*

1. *soundness: let $\Sigma; \Gamma \vdash_{\mathcal{T}} A : K$. Then $\Sigma_{flat}; |\Gamma| \vdash \|A : |K|$;*
2. *soundness: let $\Sigma; \Gamma \vdash_{\mathcal{O}} M : A$. Then $\Sigma_{flat}; |\Gamma| \vdash \|M : |A|$;*
3. *preservation: let $\Sigma; \Gamma \vdash_{\mathcal{T}\vee\mathcal{O}} s : \sigma$. Then $FV(s) = FV(\|s)$;*
4. *stability: let $\Sigma; \Gamma \vdash_{\mathcal{T}\vee\mathcal{O}} s, t : \sigma, \tau$ and $x : \tau \in \Gamma$. Then $\|s\{x \mapsto t\} = \|s\{x \mapsto \|t\}$.*

Proof. The first three are proved by induction on the typing derivations of A , M and s respectively, using Lemma 4.4 for the third (translation of products) and for the conversion rule. Stability follows by induction on the typing derivation of s . ◀

4.3 Preservation of reductions by flattening

Our goal now is to show that the reductions on objects and types in λ_{Σ}^{Π} are mimicked in $\lambda_{\Sigma_{flat}}^{\rightarrow}$. Since rewriting a term cannot not increase its set of free variables, rewriting commutes with the η -rule. A consequence is that, given a dependent rewrite system R , $\rightarrow_{\beta\eta R}$ terminates iff $\rightarrow_{\beta R}$ terminates: encoding the η -rule will not be necessary. To ease the reading, we use \rightarrow for our rewriting symbol in $\lambda_{\Sigma_{flat}}^{\rightarrow}$, and decorate subterms in flattened rules by their type.

$$\begin{aligned} [\beta_{\mathcal{O}}] \quad & @(\lambda x : \sigma.lof_{\tau}^2(A : *, M : \tau), N : \sigma) \rightarrow M\{x \mapsto N\} \\ [\beta_{\mathcal{T}}] \quad & @(\lambda x : \sigma.lf_{\tau}^2(A : *, B : \tau), N : \sigma) \rightarrow B\{x \mapsto N\} \end{aligned}$$

In these rules, A and σ are the term flattening and type erasure of the same dependent type D , a relationship that cannot be expressed in $\lambda_{\Sigma_{flat}}^{\rightarrow}$, hence is not kept in the transformed rules. For example, assuming $A, s, t \in \|\lambda_{\Sigma}^{\rightarrow}\|$ and $\|A^{-1}\| \neq \sigma$, then $u = @(\lambda x : \sigma.lof^2(A, s), t)$ is a redex which has no counter-part in λ_{Σ}^{Π} . There are indeed new rewrites in the flattened world, making preservation of non-termination a weaker property.

We use $[R]$, $[\beta]$, $[\beta R]$ and $\beta[\beta R]$ for $\{\|l \rightarrow \|r \mid l \rightarrow r \in R\}$, $\{[\beta_{\mathcal{T}}], [\beta_{\mathcal{O}}]\}$, $[\beta] \cup [R]$, and $\{\beta\} \cup [\beta R]$.

As with A, M, N, σ, τ used in $[\beta]$, variables of R that denote expressions in λ_{Σ}^{Π} keep their name in $[R]$, denoting now expressions of $\lambda_{\Sigma_{flat}}^{\rightarrow}$ belonging to the same syntactic categories as in the dependent world. Then, an instance of a $[\beta]$ -rule may not be the encoding of an instance of the β -rule in the dependent world if M/B match flattened terms which are no encoding of dependent expressions. Despite these approximations, reductions are preserved:

► **Lemma 4.7.** *Let $\Sigma; \Gamma \vdash_{\mathcal{T}\vee\mathcal{O}} s : \sigma$, $s \rightarrow_{\beta R} t$ and $\Sigma; \Gamma \vdash_{\mathcal{T}\vee\mathcal{O}} t : \tau$. Then, $|\sigma| = |\tau|$ and $\|s \rightarrow_{[\beta R]} \|t$.*

Proof. The proof is by induction on the typing derivation of s . All cases are by induction except those where the βR redex is at the top. We carry out four typical cases:

1.
$$[\text{APP}] \frac{\Sigma; \Gamma \vdash_{\mathcal{T}} A : \mathbb{I}x : B.K \quad \Sigma; \Gamma \vdash_{\mathcal{O}} M : B}{\Sigma; \Gamma \vdash_{\mathcal{T}} s = AM : K\{x \mapsto M\}}$$
 and $M \rightarrow_{\beta R} M'$, hence $AM \rightarrow_{\beta R} AM' = t$. By induction hypothesis, $\Sigma; \Gamma \vdash_{\mathcal{T}\vee\mathcal{O}} M' : B'$, $|B'| = |B|$ and $\|M \rightarrow_{[\beta R]} \|M'$. By definition of flattening, $\|AM = @(\|A, \|M) : |K\{x \mapsto M\}| = |K|$ by Corollary 4.3. By

definition of erasing, $|\Pi x : B.K| = |B| \rightarrow |K|$. By Lemma 4.6 (soundness) $\|A : |B| \rightarrow |K|$ and $\|M' : |B'|\$. Finally, $\|t = @(\|A, \|M') : |K|$, we are done.

2.
$$\frac{[\text{CONV}] \quad \Sigma; \Gamma \vdash_{\mathcal{T}\vee\mathcal{O}} s : \sigma' \quad \Sigma; \Gamma \vdash_{\mathcal{T}\vee\mathcal{K}} \sigma : \text{TYPE/KIND}}{\Sigma; \Gamma \vdash_{\mathcal{O}} s : \sigma} \sigma \equiv \sigma'.$$
 By induction hypothesis, $\Sigma; \Gamma \vdash_{\mathcal{T}\vee\mathcal{O}} t : \tau$, $|\tau| = |\sigma'|$ and $\|s \rightarrow_{[\beta R]} \|t$. By Lemma 4.6, $\|t : |\tau|$. Lemma 4.4 concludes.
3.
$$\frac{[\text{APP}] \quad \Sigma; \Gamma \vdash_{\mathcal{O}} \lambda x : A.u : \Pi x : A.B \quad \Sigma; \Gamma \vdash_{\mathcal{O}} M : A}{\Sigma; \Gamma \vdash_{\mathcal{O}} s = \lambda x : A.u M : B\{x \mapsto M\}} \text{ and } t = u\{x \mapsto M\}.$$
 By inversion, $\Sigma; \Gamma \vdash_{\mathcal{O}} u : B$, thus $t : B\{x \mapsto M\}$ by Lemma 3.1. $\|\lambda x : A.u M = @(\lambda x : A.lo_{|B|}^2(\|A, \|u), \|M) \rightarrow_{[\beta]} \|u\{x \mapsto \|M\}$ by definition of $[\beta_{\mathcal{O}}]$. Lemma 4.6 and Corollary 4.3 conclude.
4. $\Sigma; \Gamma \vdash_{\mathcal{T}\vee\mathcal{O}} s : \sigma$, $s = l \circ \gamma$ and $t = r \circ \gamma$ for some $l \rightarrow r \in R$ such that l, r have convertible types μ, ν in their environment. By Lemma 3.1 applied repeatedly, s, t have types $\sigma = \mu \circ \gamma$ and $\tau = \nu \circ \gamma$, and since \equiv is a congruence, $\sigma \equiv \tau$. By Lemma 4.6(stability), $\|l\gamma = \|l\|\gamma$ and $\|r\gamma = \|r\|\gamma$ hence $\|s \rightarrow_{[R]} \|t$ by definition of $\rightarrow_{[R]}$. \blacktriangleleft

This Lemma contains the analog of the type-preservation property of non-dependent rewriting, equivalence by conversion being here equivalence modulo type erasures: subject reduction holds for dependent rewriting modulo type erasures.

Thus λ_{Σ}^{Π} is terminating for an empty set R , implying strong normalisation of $\beta\eta$ -reductions at both object and type level, for any convertibility relation \equiv containing $\beta\eta$ -convertibility. In particular, \equiv can contain Presburger arithmetic, an important known extension of LF.

► **Example 4.8.** Here are the transformed signature and rules for our example on Lists:

$$\begin{array}{l} \text{nat}, A : * \quad \left\| \begin{array}{l} 0 : \text{nat} \\ \text{List} : \text{nat} \rightarrow * \end{array} \right\| \left\| \begin{array}{l} _ + 1 : \text{nat} \rightarrow \text{nat} \\ + : \text{nat} \rightarrow \text{nat} \rightarrow \text{nat} \end{array} \right\| \\ \text{cons} : \text{nat} \rightarrow A \rightarrow \text{List} \rightarrow \text{List} \quad \left\| \begin{array}{l} \text{app} : \text{nat} \rightarrow \text{nat} \rightarrow \text{List} \rightarrow \text{List} \rightarrow \text{List} \\ 0 + n \rightarrow n \end{array} \right\| \left\| \begin{array}{l} (m + 1) + n \rightarrow (m + n) + 1 \\ \text{app}(0, n, \text{nil}, l) \rightarrow l \end{array} \right\| \\ \text{app}(m + 1, n, \text{cons}(m, a, k), l) \rightarrow \text{cons}(m + n, a, \text{app}(m, n, k, l)) \end{array}$$

These dependently typed rewrite rules being algebraic, their encoding is the identity. CPO proves their termination with $\text{app} >_{\mathcal{F}} \text{cons} >_{\mathcal{F}} \text{nil} >_{\mathcal{F}} + >_{\mathcal{F}} _ + 1 >_{\mathcal{F}} 0$, and $\text{List} >^{\rightarrow} \{\text{nat}, A\}$.

► **Theorem 4.9.** *Given a signature Σ , a dependent term rewriting system $\beta\eta R$ is terminating in λ_{Σ}^{Π} if its flattening $\beta\|\beta R$ is terminating in $\lambda_{\Sigma}^{\rightarrow \text{flat}}$.*

Proof. By using Lemma 4.7 and a commutation argument for η . \blacktriangleleft

4.4 Inverse encoding

Lemma 4.7 justifies our method for checking strong normalization by a transformation to a higher-order algebra where we can use standard techniques including CPO. But the flattened world is richer than the dependent world, there are more terms and rewrites. Nonetheless, we can also show that termination is partly preserved by defining an inverse transformation such that composing both is the identity. First, we define the inverse transformation on the subset $\|\lambda_{\Sigma} \subseteq \lambda_{\Sigma}^{\text{flat}}$, hence allowing us to show that flattening is injective.

► **Definition 4.10.** The inverse $\|_^{-1} : \|\lambda_{\Sigma} \rightarrow \lambda_{\Sigma}$ is defined as:

$$\begin{array}{l} \|x^{-1} = x \quad \left\| \quad \|a^{-1} = a \quad \left\| \quad \|f^n(s_1, \dots, s_n)^{-1} = f^n(\|s_1^{-1}, \dots, \|s_n^{-1}) \right. \\ \|\@ (s, t)^{-1} = \|s^{-1} \|t^{-1} \quad \left\| \quad \|\text{pif}^2(s, \lambda x : \sigma.t)^{-1} = \Pi x : \|s^{-1}. \|t^{-1} \\ \|\lambda x : \sigma.\text{lf}^2(s, t)^{-1} = \lambda x : \|s^{-1}. \|t^{-1} \quad \left\| \quad \|\lambda x : \sigma.\text{lo}_{\tau}^2(s, t)^{-1} = \lambda x : \|s^{-1}. \|t^{-1} \end{array}$$

► **Lemma 4.11** (Reversibility). *Assume $\Sigma; \Gamma \vdash s : \sigma$ in λ_{Σ}^{Π} . Then $\Sigma; \Gamma \vdash \|(\|s)\|^{-1} = s : \sigma$.*

Proof. By induction on s . Variables, type constants, applications and pre-algebraic terms are clear. We carry out one remaining case, the others being similar. $\|(\|\Pi x : A.B)\|^{-1} = \|\mathit{pif}_{|A|}^2(\|A, \lambda x : |A|. \|B)\|^{-1} = \Pi x : \|(\|A)\|^{-1}. \|(\|B)\|^{-1}$. We conclude by induction. ◀

► **Corollary 4.12.** *Let $\sigma; \Gamma \vdash_{\mathcal{T}\vee\mathcal{O}} s : \sigma$ and $\|s\|_p \in \|\lambda_{\Sigma}$ for some p . Then $\|s\|_p = \|s\|_q$ for some q .*

Flattening being an injection from λ_{Σ} to $\lambda_{\Sigma^{flat}}$, hence a bijection between λ_{Σ} and its target $\|\lambda_{\Sigma}$, provides strong evidence that λ_{Σ}^{Π} is faithfully encoded by flattening. It follows that, if $\|s\| = \|t\|$ for $s, t \in \lambda_{\Sigma}$, then $s = t$ since $s = \|\|s^{-1}\| = \|\|t^{-1}\| = t$. We apply to Lemma 4.6:

► **Corollary 4.13** (inverse stability). *Assume $u\{x \mapsto v\} \in \|\lambda_{\Sigma}$. Then, $\|u\{x \mapsto v\}^{-1} = \|u^{-1}\{x \mapsto \|v^{-1}\}$.*

Therefore, rewrites in $\|\lambda_{\Sigma}$ can be mapped back to λ_{Σ} , showing both that $\|\lambda_{\Sigma}$ is closed under rewriting by $[\beta R]$, and that termination in λ_{Σ} implies termination in $\|\lambda_{\Sigma}$:

► **Lemma 4.14.** *Let $\Sigma; \Gamma \vdash_{\mathcal{T}\vee\mathcal{O}} s : \sigma$ and $\|s \rightarrow_{[\beta R]} u$. Then $u = \|t$ for some t st $s \rightarrow_{\beta R} t$.*

Proof. By assumption $\|s\|_p = \|l\gamma \rightarrow_{[\beta R]} u\|_p$ for some $p \in \mathcal{Pos}(\|s)$, where $l \rightarrow r \in \{\beta\} \cup R$, γ a substitution in $\|\lambda_{\Sigma}$ and $u\|_p = \|r\gamma$. By Corollary 4.12, $\|s\|_p = \|s\|_q$ for some q . By Lemma 4.11, $s\|_q = \|\|l\gamma^{-1}\| = \|\|l^{-1}\|\gamma^{-1}$ (by Corollary 4.13) = $l\theta$ with $\|\theta = \gamma$ by Lemma 4.11. Hence $s \rightarrow_{\beta R} s[r\theta]_q = t$. Now, since flattening is a homomorphism, $\|t = \|s\|\|r\theta\|_p$ (by definition of q) = $\|s\|\|r\gamma\|_p$ (by definition of θ) = u . ◀

This does not prove, however, that termination in λ_{Σ}^{Π} implies termination in $\lambda_{\Sigma^{flat}}^{\rightarrow}$. The problem is that flattening is non-surjective, since many terms, like those headed by lof^2 , are no flattening of a dependent term. Further, some seemingly good-looking terms, like $\mathit{pif}_{\sigma}^2(s, \lambda x : \sigma.t)$ may not be either, even assuming that s, t are themselves flattening of dependent terms. This is the case because the flattened signature checks that term s has type $*$, hence is the encoding of some dependent type A , but cannot check whether $|A| = \sigma$ as it should. The same happens with applications and pre-algebraic terms: (s, t) may not be typable in λ_{Σ}^{Π} when s, t are typable in λ_{Σ}^{Π} and $@(\|s, \|t)$ is typable in $\lambda_{\Sigma^{flat}}^{\rightarrow}$. Indeed, termination in λ_{Σ}^{Π} does not imply termination in $\lambda_{\Sigma^{flat}}^{\rightarrow}$ as shown by this example: let $o : \text{TYPE}; List : \Pi x : o. \text{TYPE}; a, b : o; la : List a; lb : List b; f^2 : \Pi m : o, l : List m. o;$ and $\{f(a, la) \rightarrow f(b, lb); b \rightarrow a, lb \rightarrow la\}$, whose dependent derivations are all finite since b cannot rewrite to a in $f(b, lb)$, a non-typable term. In the flattened signature, o and $List$ are sorts, and $a, b : o; la, lb : List; f^2 : o \rightarrow List \rightarrow o;$ and $\{f(a, la) \rightarrow f(b, lb); b \rightarrow a; lb \rightarrow la\}$. We now have the following infinite derivation: $f(a, la) \rightarrow f(b, lb) \rightarrow f(a, lb) \rightarrow f(a, la) \dots$. Restricting rewrites on parameters could be a solution.

5 The Dependent Computability Path Ordering

DCPO is an extension of CPO obtained by adding new cases for products. All notations for DCPO are simply and systematically obtained by replacing the CPO ordering notation $>$ by \succ . For example, the precedence will be denoted by \succ_{Σ} . Type families are compared alternatively as types with \succeq^{Π} and as terms with \succ . The basic ingredients of DCPO are:

- a precedence \succeq_{Σ} on $\mathcal{F} \cup \mathcal{S} \cup \{\@, \lambda, \Pi\}$ s.t. \succ_{Σ} is well-founded, and $(\forall f \in \mathcal{F}) f \succ_{\Sigma} \@ \succ_{\Sigma} \{\lambda, \Pi\}$.
- a status $f \in \{Mul, Lex\}$ for every symbol $f \in \mathcal{F} \cup \{\@\}$ with $\@ \in Mul$.
- a quasi-order \succeq^{Π} on $\mathcal{K} \cup \mathcal{T}$, whose strict part \succ^{Π} and equivalence $=^{\Pi}$ satisfy:

- (i) *compatibility*: $\equiv \subseteq =^{\Pi}$;
- (ii) *well-foundedness*: $\succ^{\Pi} \cup \{(\Pi x : A.\sigma, A) \mid A \in \mathcal{T}, \sigma \in \mathcal{T} \cup \mathcal{K}\}$ is well-founded;
- (iii) *product body subterm*: $(\forall \sigma \in \mathcal{T} \cup \mathcal{K}) \Pi x : A.\sigma \succ^{\Pi} \sigma$;
- (iv) *product preservation*: $|\Pi x : A.\sigma| =^{\Pi} |\tau|$ implies $\tau = \Pi x : B.\mu$ for some B, μ , such that $|A| =^{\Pi} |B|, |\sigma| =^{\Pi} |\mu|$;
- (v) *decreasingness*: $\Pi x : A.\sigma \succeq^{\Pi} \tau$ implies $\sigma \succeq^{\Pi} \tau$ or else $\tau = \Pi x : B.\nu, A =^{\Pi} B$ and $\tau \succeq^{\Pi} \nu$.

Building quasi-orders \succeq^{Π} on types and kinds with a non-trivial equivalence $=^{\Pi}$ is not hard [14].

We now define a first version of DCPO which can be justified by using CPO [9]. We will then discuss briefly an enhanced version justified by a more elaborated version of CPO [10]. In the following definition, z denotes a fresh variable of type A/B .

► **Definition 5.1 (DCPO)**. Given $\Gamma \vdash_{\Sigma} s : \sigma$ and $\Gamma \vdash_{\Sigma} t : \tau$, then $s \succ^X t$ iff either:

1. $t \in X$ and $s \notin \mathcal{X}$ (var)
2. $s = (u N)$ and $u \succeq^X t$ or $N : B \succeq^X t : \tau$ (subt@)
3. $s = \lambda / \Pi x : A.u$ and $u\{x \mapsto z\} : \mu \succeq^X t : \tau$ (subt $\lambda\Pi$)
4. $s = f(\bar{s}), f \in \Sigma, t = (\lambda / \Pi y : B.v) : \Pi y : B.K$, and $s \succ^X B$ and $s \succ^{X \cup \{z\}} v\{y \mapsto z\}$ (Σ prec $\lambda\Pi$)
5. $s = \lambda x : A.u : \Pi x : A.C, t = \lambda y : B.v : \Pi y : B.D, |A| = |B|, A \succ^X B$ and $u\{x \mapsto z\} : C \succ^X v\{y \mapsto z\} : D$ (stat λ)
6. $s = \Pi x : A.u, t = \Pi y : B.v, |A| = |B|, A \succ^X B$ and $u\{x \mapsto z\} \succ^X v\{y \mapsto z\}$ (stat Π)
7. $s = ((\lambda x : A.u) : (\Pi x : A.C) w : A)$, and $u\{x \mapsto w\} \succeq t$ (beta)
8. otherwise $s = f(\bar{s}), t = g(\bar{t})$ with $f, g \in \mathcal{F} \cup \mathcal{S} \cup \{\text{@}\} \cup \mathcal{X}$, and either of (rpo)
 - $\bar{s} : \bar{\sigma} \succeq t : \tau$ (subt) $f \succ_{\mathcal{F}} g$ and $s \succ^X \bar{t}$ (prec) $f =_{\Sigma} g, s \succ^X \bar{t}$ and $\bar{s} : \bar{\sigma} (\succ)_f \bar{t} : \bar{\tau}$ (stat)

All terms built by DCPO are well-typed under the assumption that both starting terms s, t are well-typed. Note (i) the importance of all our assumptions on \succeq^{Π} , see for example Case stat $\lambda\Pi$; (ii) the order may recursively compare objects with types, even when the input comparison operates on expressions in $\mathcal{T}^2 \cup \mathcal{O}^2$; (iii) compared products cannot be kinds.

5.1 Example

We now consider our list example, skipping the rules on natural numbers which do not have dependencies. We shall use the (user defined) precedence $app \succ_{\Sigma} \{+, cons, nil\}$, multiset status for app , and a quasi-order on types in which, for all $n : nat$, then $List n \succeq^{\Pi} A$ and $List m =^{\Pi} List n$ for all m, n of type nat . Such a type order is easy to get by using a restricted RPO on type erasures, which equivalence therefore contains Presburger arithmetic. See [14]. The goal $app(0, n, nil, l) : List 0 + n \succ l : List n$ is easy, although already requiring identification of $0 + n$ in $List 0 + n$ with n in $List n$. We proceed with the second goal: $app(m+1, n, cons(m, a, k), l) : List (m+1) + n \succ cons(m+n, a, app(m, n, k, l)) : List (m+n)+1$. The type comparison $List (m+1) + n \succeq^{\Pi} List (m+n) + 1$ succeeds by using our type ordering which indeed equates both types, and we are left with the term comparison: $app(m+1, n, cons(m, a, k), l) \succ cons(m+n, a, app(m, n, k, l))$. Using (prec), we get three subgoals, which are processed in turn, using indentation to identify the dependencies between recursive calls, the used Case being indicated between parentheses:

$$\begin{array}{l}
app(m+1, n, cons(m, a, k), l) \succ m+n \quad (prec) \\
app(m+1, n, cons(m, a, k), l) \succ m \quad (subt) \\
m+1 : nat \succeq m : nat \text{ which yields} \\
nat \succeq^{\Pi} nat \text{ which succeeds and } m+1 \succeq m \text{ which succeeds by } (subt) \\
app(m+1, n, cons(m, a, k), l) \succ n \quad \text{which succeeds by } (subt) \\
app(m+1, n, cons(m, a, k), l) \succ a \quad (subt) \\
cons(m, a, k) : List(m+1) \succ a : A \text{ which yields} \\
List(m+1) \succeq^{\Pi} A \text{ which succeeds and } cons(m, a, k) \succ a \text{ which succeeds by } (subt) \\
app(m+1, n, cons(m, a, k), l) \succ app(m, n, k, l) \quad (stat) \\
\{m+1 : nat, n : nat, cons(m, a, k) : List\ m+1, l : List\ n\} \\
\quad \succ_{mul} \{m : nat, n : nat, k : List\ m, l : List\ n\} \text{ which yields} \\
m+1 : nat \succ m : nat \text{ and } cons(m, a, k) : List\ m+1 \succ k : List\ m, \text{ for the reader.}
\end{array}$$

5.2 Properties of DCPO

► **Lemma 5.2** (Monotonicity, stability). *Let $s, t \in \lambda_{\Sigma}$ $st : \sigma \succ t : \sigma$, $C(x : \sigma)$ a context, and γ a substitution. Then $C\{x \mapsto s\gamma\} \succ C\{x \mapsto t\gamma\}$.*

Proof. By induction on the definition of $s \succ t$ and stability of the type order for monotonicity. By induction on the context, and use of the status rules for stability. ◀

As anticipated, we now reduce the well-foundedness of DCPO to the well-foundedness of CPO by using Theorem 4.9, termination in the target higher-order algebra being checked by CPO. To this end, we need to show that, whenever $s : \sigma \succ t : \tau$, then $\|s : |\sigma| > \|t : |\tau|$.

We start showing that an order on types and kinds of LF satisfying the requirements for DCPO becomes naturally an order on types of $\lambda_{\Sigma^{flat}}^{\rightarrow}$ satisfying the requirements for CPO.

► **Definition 5.3.** Given $\sigma, \tau \in \mathcal{T}_{\Sigma^{flat}}^{\rightarrow}$, let $\sigma \succeq_{\Sigma^{flat}}^{\rightarrow} \tau$ iff $\exists \mu, \nu \in \mathcal{T} \cup \mathcal{K}$, $|\mu| = \sigma$, $|\nu| = \tau$, and $\mu \succeq^{\Pi} \nu$.

Transitivity is clear. Of course, different choices of μ, ν may sometimes lead to contradictory orderings for σ, τ , hence the equality $\simeq_{\Sigma^{flat}}^{\rightarrow}$ may be strictly larger than the corresponding equality \simeq^{Π} . This has indeed no negative impact since the strict part \succ^{Π} is never used in comparisons. Further, the properties of \succeq^{Π} transfer naturally to $\succeq_{\Sigma^{flat}}^{\rightarrow}$. Soundness follows:

► **Lemma 5.4.** *Assume \succeq^{Π} is a DCPO type order. Then $\succeq_{\Sigma^{flat}}^{\rightarrow}$ is a CPO type order.*

The set Σ^{flat} of function symbols of Σ^{flat} is the union of \mathcal{S} , \mathcal{F} and the flattening constructors. The precedence $\succ_{\Sigma^{flat}}$ is obtained by letting the flattening constructors be equivalent minimal symbols. The strict part of $\succ_{\Sigma^{flat}}$ is clearly well-founded, while its equivalence is increased by the equality of the flattening constructors. We take status *Mul* for lo^2, lf^2, pif^2 .

► **Lemma 5.5.** *Let $s : \sigma, t : \tau$ and $X \subseteq \mathcal{X}$ such that $s \succ^X t$ by any DCPO rule. Then $\|s >^X \|t$.*

Proof Sketch. The proof is by induction on the definition of DCPO, assuming by induction hypothesis that the property holds at every recursive call. Note that if $s : \sigma, t : \tau, \sigma \succeq^{\Pi} \tau$ and $\|s >^X \|t$, then $\|s : |\sigma| >^X \|t : |\tau|$ by Definition 5.3. ◀

We then obtain the second main result of the paper as a corollary:

► **Theorem 5.6** (Well-foundedness of DCPO). *\succ^{Π} is well-founded.*

Proof. By theorem 4.9. Note that we use the full strength of that result, including the need for type-level rules instances of the various DCPO rules instances which compare types. This is possible since we only need preservation of non-termination by the flattening transformation, that is, the if direction of Theorem 4.9. ◀

It follows that $\succ^{\Pi} \cup \eta$ is well-founded. One may wonder why we did not include η -rules in the definition of DCPO, since it is in the definition of CPO. The reason is the comparison of lefthand and righthand sides of $[\eta]$, $\lambda x:|A|.lof^2(\|A, @(\|B, x)) > \|B$ which does not go through: the type comparison of the subgoal $lof^2(\|A, @(\|B, x)):\sigma > \|B:|A| \rightarrow \sigma$ fails.

5.3 A realistic example

We consider here a more complex, non-algebraic, higher-order example. Given a list l of natural numbers x_1, \dots, x_m , and a natural number y , a higher-order variable g which is meant to be instantiated by $+$, we define a higher-order function $foldr$ such that $(foldr(m, l, y) g)$ calculates $g(x_1, g(x_1, \dots, g(x_m, y)))$, while $(map(m, l) f)$ calculates the list $f(x_1), \dots, f(x_m)$.

```

nat : TYPE | 0 : nat | + 1 :  $\Pi x : nat.nat$ 
List :  $\Pi m : nat.TYPE$  | nil : List 0 | cons :  $\Pi\{m : nat, x : nat, l : List\}m.List\ m + 1$ 
map :  $\Pi\{m : nat, l : List\}m.( $\Pi f : ( $\Pi x : nat.nat$ ).List\ m$ )
foldr :  $\Pi\{m : nat, l : list\}m, y : nat.\Pi g : ( $\Pi\{x_1 : nat, x_2 : nat\}.nat$ ).nat$ 
map(0, l)  $\rightarrow \lambda f : ( $\Pi x : nat.nat$ ).nil$  | foldr(0, l, y)  $\rightarrow \lambda g : \Pi\{x_1 : nat, x_2 : nat\}.nat.y$ 
map(m + 1, cons(m, x, l))  $\rightarrow \lambda f : \Pi\{x : nat\}.nat.cons((f\ x), map(m, l))$ 
foldr(m + 1, cons(m, z, l), y)  $\rightarrow \lambda g : \Pi\{x_1 : nat, x_2 : nat\}.nat.(g\ z\ (foldr(m, l, y)\ g))$$ 
```

To carry out the example, we let $List >_{\mathcal{T}} nat$, $foldr > map > cons > nil > lf^2 > lo^2 > List > nat$ and $(\forall \sigma >_{\mathcal{T}} \tau) lf_{\sigma}^2 > lf_{\tau}^2$ and $lo_{\sigma}^2 > lo_{\tau}^2$. We consider the last rule only.

CPO comparison. It generates the following goals (omitting the type comparisons):

```

 $\|l > \lambda g : nat \rightarrow nat \rightarrow nat.lo_{nat}^2$ 
 $(lf_{nat}^2(nat, \lambda x_1 : nat.lf_{nat}^2(nat, \lambda x_2 : nat.nat)), @(@(g, z), @(foldr(m, l, y) g)))$   $\mathcal{F}\lambda$ 
 $\|l > \{g\} lo_{nat}^2(lf_{nat}^2(nat, \lambda x_1 : nat.lf_{nat}^2(nat, \lambda x_2 : nat.nat)), @(@(g, z), @(foldr(m, l, y) g)))$ 
 $\|l > \{g\} lf_{nat}^2(nat, \lambda x_1 : nat.lf_{nat}^2(nat, \lambda x_2 : nat.nat))$   $\text{PREC}$ 
 $\|l > \{g\} nat$   $\text{succeeds by PREC}$ 
 $\|l > \{g\} \lambda x_1 : nat.lf_{nat}^2(nat, \lambda x_2 : nat.nat)$   $\mathcal{F}\lambda$ 
 $\|l > \{g, x_1\} lf_{nat}^2(nat, \lambda x_2 : nat.nat)$   $\text{PREC}$ 
 $\|l > \{g, x_1\} nat$   $\text{succeeds by PREC}$ 
 $\|l > \{g, x_1\} \lambda x_2 : nat.nat$   $\mathcal{F}\lambda$ 
 $\|l > \{g, x_1, x_2\} nat$   $\text{succeeds by PREC}$ 
 $\|l > \{g\} @(@(g, z), @(foldr(m, l, y) g))$   $\text{PREC}$ 
 $\|l > \{g\} @(g, z)$   $\text{PREC}$ 
 $\|l > \{g\} g$   $\text{succeeds by VAR}$ 
 $\|l = foldr(m + 1, cons(m, z, l), y) > \{g\} z$   $\text{SUBT}$ 
 $cons(m, z, l) : List(m + 1) > \{g\} z : nat$   $\text{which succeeds by SUBT}$ 
 $\|l > \{g\} @(foldr(m, l, y) g)$   $\text{PREC}$ 
 $\|l = foldr(m + 1, cons(m, z, l), y) > \{g\} foldr(m, l, y)$   $\text{STAT}$ 
 $\{m + 1, cons(m, z, l), y\} > \{g\}_{mul}\{m, l, y\}$   $\text{which succeeds by repeated SUBT}$ 
 $\|l > \{g\} g$   $\text{which succeeds by VAR, therefore ending the computation successfully.}$ 

```

DCPO comparison. We now carry out the same computation with DCPO directly:

```

foldr(m + 1, cons(m, z, l), y)  $\succ \lambda g : \Pi\{x_1 : nat, x_2 : nat\}.nat.(g\ z\ (foldr(m, l, y)\ g))$  ( $\Sigma\text{prec}\lambda\Pi$ )
 $l \succ \Pi\{x_1 : nat, x_2 : nat\}.nat$  ( $\Sigma\text{prec}\lambda\Pi$ )
 $l \succ nat$   $\text{which succeeds by (prec)}$ 
 $l \succ \{x_1\} \Pi x_2 : nat.nat$  ( $\Sigma\text{prec}\lambda\Pi$ )
 $l \succ \{x_1\} nat$   $\text{which succeeds by (prec)}$ 

```

$l \succ^{\{x_1, x_2\}} nat$	which succeeds by (prec)
$l \succ^{\{g\}} (g z (foldr(m, l, y) g))$	-we keep the @ operator implicit this time- (prec)
$l \succ^{\{g\}} g$	(var)
$l \succ^{\{g\}} z$	(subt)
$cons(m, z, l) : List(m + 1) \succ^{\{g\}} z : nat$	succeeds by (subt)
$l \succ^{\{g\}} (foldr(m, l, y) g)$	(prec)
$l \succ^{\{g\}} foldr(m, l, y)$	(stat)
$\{m + 1, cons(m, z, l), y\} \succ^{\{g\}} \{m, l, y\}$	which succeed by repeated (subt)
$l \succ^{\{g\}} g$	succeeds by (var), therefore ending the computation successfully.

6 Conclusion

The amount of research work targeting automatic termination is vast. Among the most popular techniques are dependency pairs, introduced by Aart and Giesl and the size-changing principle, pioneered by Neil Jones, which have been generalized to dependently-typed rules [5]. Dependent types can also be *used* to store annotations useful for proving termination [22]. Despite these proposals that recent prototypes try to combine, techniques used in Coq and Agda are still poor, as acknowledged by the authors on their websites. Using our techniques would improve this situation.

Our first contribution is a new transformation for eliminating type dependencies using a framework richer than the simply-typed λ -calculus, which provides with a natural encoding, and allows us to consider arbitrary rewrite rules, not only β - and η -reductions. Furthermore, these results hold for a practical dependent type system made richer than LF's via a convertibility relation possibly stronger than $\beta\eta$ -convertibility, hence allowing us to type many more terms. This easily implementable transformation allow us using existing implementations targeting termination of rules in presence of simple types. The transformation also allows us to show well-foundedness of DCPO, a version of CPO applying to dependently typed terms directly. This is done by considering pairs ordered by DCPO as dependently typed rewrite rules to which the transformation applies. Note that DCPO will naturally benefit from improvements of CPO, without changing the proof technique. In particular, we could easily accommodate size interpretations by using them as a precedence as in [11], or type level rules such as $s = (A u)$, $t = \Pi y : B.v$, with $s \succ^X B$ and $s \succ^{X \cup \{z\}} v\{y \mapsto z\}$ flattened as $\|s = @(\|A, \|u)$, and $\|t = pik_{|B|}(\|B, \lambda y : |B|. \|v)$, with $\|s \succ^X \|B$ and $\|s \succ^{X \cup \{z\}} \|v\{y \mapsto z\}$, whose justification requires the extension of CPO with *small* function symbols, here $pik_{|B|}$, which behave as if they were smaller than application and abstraction [10]. We can then break the main goal into the above solvable subgoals.

Our main interest is indeed to prove termination directly at the dependent type level. Using DCPO allows the programmer, in case of failure, to get an error message in her/his own dependently typed syntax, rather than in the transformed syntax as would be the case when using CPO on the transformed rules. To our knowledge, this is the very first general – Coq and Agda's techniques are very limited –, purely syntactic method that allows one to show termination of a set of dependently typed rewrite rules via computations taking place on the user's dependently typed rules.

Acknowledgements. Work supported by NSFC grants 61272002, 91218302, 973 Program 2010CB328003 and Nat. Key Tech. R&D Program SQ2012BAJY4052 of China. This work was done in part during a visit of the first author supported by Tsinghua University.

References

- 1 T. Arts and J. Giesl. Termination of term rewriting using dependency pairs. *Theor. Comput. Sci.*, 236(1-2):133–178, 2000.
- 2 H. P. Barendregt. Lambda calculus with types. In *Handbook of Logic in Computer Science, Volume 2*, pages 117–309. Oxford Univ. Press, 1993.
- 3 B. Barras, J.-P. Jouannaud, P.-Y. Strub, and Q. Wang. CoqMTU: A higher-order type theory with a predicative hierarchy of universes parametrized by a decidable first-order theory. In *LICS*, pages 143–151. IEEE Computer Society, 2011.
- 4 G. Barthe. The relevance of proof-irrelevance. In *ICALP*, volume 1443 of *LNCS*, pages 755–768. Springer, 1998.
- 5 F. Blanqui. A type-based termination criterion for dependently-typed higher-order rewrite systems. In *RTA*, volume 3091 of *LNCS*, pages 24–39. Springer, 2004.
- 6 F. Blanqui. Definitions by rewriting in the calculus of constructions. *CoRR*, /abs/cs/0610065, 2006.
- 7 F. Blanqui, J.-P. Jouannaud, and M. Okada. The calculus of algebraic constructions. In *RTA*, volume 1631 of *LNCS*, pages 301–316. Springer, 1999.
- 8 F. Blanqui, J.-P. Jouannaud, and A. Rubio. HORPO with computability closure : A reconstruction. In *LPAR*, volume 4790 of *LNCS*. Springer, 2007.
- 9 F. Blanqui, J.-P. Jouannaud, and A. Rubio. The computability path ordering: The end of a quest. In *CSL, LNCS* 5213. 2008.
- 10 F. Blanqui, J.-P. Jouannaud, and A. Rubio. The computability path ordering. To appear in *LMCS*.
- 11 C. Borralleras and A. Rubio. A monotonic higher-order semantic path ordering. In *LPAR*, volume 2250 of *LNCS*, pages 531–547. Springer, 2001.
- 12 Nachum Dershowitz. Orderings for term-rewriting systems. *TCS*, 17:279–301, 1982.
- 13 R. Harper, F. Honsell, and G. Plotkin. A framework for defining logics. *Journal of the ACM*, 40(1):143–184, 1993.
- 14 J.-P. Jouannaud and A. Rubio. Polymorphic higher-order recursive path orderings. *J. ACM*, 54(1), 2007.
- 15 K. Kusakari, Y. Isogai, M. Sakai, and F. Blanqui. Static dependency pair method based on strong computability for higher-order rewrite systems. *CoRR*, abs/1109.5468, 2011.
- 16 U. Norell. Dependently typed programming in Agda. In *TLDI*, pages 1–2. ACM, 2009.
- 17 C. Paulin-Mohring. Inductive definitions in the system Coq - rules and properties. In *TLCA*, volume 664 of *LNCS*, pages 328–345. Springer, 1993.
- 18 F. Pfenning. Elf: A language for logic definition and verified metaprogramming. In *In Fourth Annual IEEE Symposium on Logic in Computer Science*, pages 313–322. 1989.
- 19 P. M. Rondon, M. Kawaguchi, and R. Jhala. Liquid types. In *PLDI*, pages 159–169. ACM, 2008.
- 20 P.-Y. Strub. Coq modulo theory. In *CSL*, volume 6247 of *LNCS*, pages 529–543. Springer, 2010.
- 21 S. Suzuki, K. Kusakari, and F. Blanqui. Argument filterings and usable rules in higher-order rewrite systems. *CoRR*, abs/1109.4357, 2011.
- 22 H. Xi. Dependent ML an approach to practical programming with dependent types. *J. Funct. Program.*, 17(2):215–286, 2007.
- 23 H. Xi and F. Pfenning. Dependent types in practical programming. In *POPL*, pages 214–227. ACM, 1999.

Well-Founded Recursion over Contextual Objects

Brigitte Pientka¹ and Andreas Abel²

1 School of Computer Science
McGill University, Montreal, Canada
bpientka@cs.mcgill.ca

2 Department of Computer Science and Engineering, Gothenburg University
Göteborg, Sweden
andreas.abel@gu.se

Abstract

We present a core programming language that supports writing well-founded structurally recursive functions using simultaneous pattern matching on contextual LF objects and contexts. The main technical tool is a coverage checking algorithm that also generates valid recursive calls. To establish consistency, we define a call-by-value small-step semantics and prove that every well-typed program terminates using a reducibility semantics. Based on the presented methodology we have implemented a totality checker as part of the programming and proof environment Beluga where it can be used to establish that a total Beluga program corresponds to a proof.

1998 ACM Subject Classification D.3.1[Programming Languages] Formal Definitions and Languages. F.3.1[Logics and Meaning of Programs] Specifying and Verifying and Reasoning about Programs

Keywords and phrases Type systems, Dependent Types, Logical Frameworks

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.273

1 Introduction

Mechanizing formal systems and their proofs play an important role in establishing trust in formal developments. A key question in this endeavor is how to represent variables and assumptions to which the logical framework LF [8], a dependently typed lambda-calculus, provides an elegant and simple answer: both can be represented uniformly using LF's function space, modelling binders in the object language using binders in LF. This kind of encoding is typically referred to as *higher-order abstract syntax* (HOAS) and provides a general uniform treatment of syntax, rules and proofs.

While the elegance of higher-order abstract syntax encodings is widely acknowledged, it has been challenging to reason inductively about LF specifications and formulate well-founded recursion principles. HOAS specifications are not inductive in the standard sense. As we recursively traverse higher-order abstract syntax trees, we extend our context of assumptions, and our LF object does not remain closed. To tackle this problem, Pientka and collaborators [11, 4] propose to pair LF objects together with the context in which they are meaningful. This notion is then internalized as a contextual type $[\Psi.A]$ which is inhabited by terms M of type A in the context Ψ [9]. Contextual objects are then embedded into a computation language which supports general recursion and pattern matching on contexts and contextual objects. Beluga, a programming environment based on these ideas [13], facilitates the use of HOAS for non-trivial applications such as normalization-by-evaluation [4] and a type-preserving compiler including closure conversion and hoisting [3]. However, nothing in this work enforces or guarantees that a given program is total.



© Brigitte Pientka and Andreas Abel;
licensed under Creative Commons License CC-BY

13th International Conference on Typed Lambda Calculi and Applications (TLCA'15).

Editor: Thorsten Altenkirch; pp. 273–287



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In this paper, we develop a core functional language for reasoning inductively about context and contextual objects. One can think of this core language as the target of a Beluga program: elaboration may use type reconstruction to infer implicit indices [6] and generate valid well-founded recursive calls that can be made in the body of the function. Type checking will guarantee that we are manipulating well-typed objects and, in addition, that a given set of cases is covering and the given recursive calls are well-founded. To establish consistency, we define a call-by-value small-step semantics for our core language and prove that every well-typed program terminates, using Tait’s method of logical relations. Thus, we justify the interpretation of well-founded recursive programs in our core language as inductive proofs. Based on our theoretical work, we have implemented a totality checker for Beluga.

Our approach is however more general: our core language can be viewed as a language for first-order logic proofs by structural induction over a given domain. The domain must only provide answers to three domain-specific questions: (1) how to unify objects in the domain, (2) how to split on a domain object and (3) how to justify that a domain object is smaller according to some measure. The answer to the first and second question allows us to justify that the given program is covering, while the third allows us to guarantee termination. For the domain of contextual LF presented in this paper, we rely on higher-order unification [2] for (1), and our splitting algorithm (2) and subterm ordering (3) builds on previous work [5, 10]. As a consequence, our work highlights that reasoning about HOAS representations via contextual types can be easily accommodated in a first-order theory. In fact, it is a rather straightforward extension of how we reason inductively about simple domains such as natural numbers or lists.

The remainder of the paper is organized as follows. We first present the general idea of writing and verifying programs to be total in Sec. 2 and then describe in more detail the foundation of our core programming language which includes well-founded recursion principles and simultaneous pattern matching in Sec. 3. The operational semantics together with basic properties such as type safety is given in Sec. 4. In Sec. 5, we review contextual LF [4], define a well-founded measure on contextual objects and contexts, and define the splitting algorithm. Subsequently we describe the generation of valid well-founded recursive calls generically, and prove normalization (Sec. 7). We conclude with a discussion of related work, current status and future research directions. Due to space constraints, proofs have been omitted.

2 General Idea

2.1 Example 1: Equality on Natural Numbers

To explain the basic idea of how we write inductive proofs as recursive programs, we consider first a very simple example: reasoning about structural equality on natural numbers (see Listing 1). We encode natural numbers and equality on them in the logical framework LF.

■ **Listing 1** Encoding of an Inductive Proof as a Recursive Function.

```

nat : type.
z   : nat.
s   : nat → nat.
.
ref : Π M:nat. [eq M M] = Λ M ⇒ rec-case M of
|   ref z           ⇒ [eq_z]
| M':nat ; ref M':[eq M' M']. ref (s M') ⇒ let D = ref M' in [eq_s M' M' D];

```

The free variables M and N in the definition of `eq_s` are implicitly quantified at the outside. Program `ref` proves reflexivity of `eq`: for all $M:\text{nat}$ we can derive `eq M M`. Following type-theoretic notation, we write Π for universal quantification; we embed LF objects which denote base predicates via `[]`. Abstraction over LF object M is written $\Lambda M \Rightarrow$ in our language. Using `rec-case`, we prove inductively that for all M there is a derivation for `[eq M M]`. There are two cases to consider: `ref z` describes the base case where M is zero and the goal refines to `[eq z z]`. In this case, the proof is simply `[eq_z]`. In the step case, written as `ref (s M')`, we also list explicitly the other assumptions: the type of M' and the induction hypothesis written as `ref M':[eq M' M']`. To establish that `[eq (s M') (s M')]`, we first obtain a derivation D of `eq M' M'` by induction hypothesis and then extend it to a derivation `[eq_s M' M' D]` of `[eq (s M') (s M')]`. We highlight in green redundant information which can be inferred automatically. In the pattern, it is the typing (here: $M':\text{nat}$) of the pattern variables [12, 6] and the listing of the induction hypotheses. The dot “.” separates these assumptions from the main pattern. For clarity, we choose to write the pattern as a simultaneous pattern match and make the name of the function explicit; in practice, we only write the main pattern which is left in black, and all other arguments are inferred.

2.2 Example 2: Intrinsically Typed Terms

Next, we encode intrinsically typed λ -terms. This example does exploit the power of LF.

```

tp   : type.
bool : tp.
arr  : tp → tp → tp.

tm   : tp → type.
lam  : (tm A → tm B) → tm (arr A B).
app  : tm (arr A B) → tm A → tm B.

```

We define base types such as `bool` and function types, written as `arr T S`, and represent simply-typed lambda-terms using the constructors `lam` and `app`. In particular, we model the binding in the lambda-calculus (our object language) via HOAS, using the LF function space. For example, the identity function is represented as `lam λx.x` and function composition as `lam λg. lam λf. lam λx. app (f (app g x))`. As we traverse λ -abstractions we record the variables we are encountering in a context $\phi : \text{cxt}$. Its shape is given by a *schema declaration* `schema ctx = tm A` stating that it contains only variable bindings of type `tm A` for some A . To reason about typing derivations, we package the term (or type) together with its context, forming a contextual object (or contextual type, resp.). For example, we write $\phi \vdash \text{tm } A$ for an object of type `tm A` in the context ϕ . Such *contextual types* are embedded into logical statements as `[\phi \vdash tm A]`. When the context ϕ is empty, we may drop the turnstile and simply write `[tm A]`.

Counting constructors: Induction on (contextual) LF object

As an example, we consider counting constructors in a term. This corresponds to defining the overall size of a typing derivation. We recursively analyze terms M of type `tm A` in the context ϕ . In the variable case, written as `count φ B (φ ⊢ p...)`, we simply return zero. The pattern variable p stands for a variable from the context ϕ . We explicitly associate it with the identity substitution, written as `...`, to use p which has declared type $\phi \vdash \text{tm } B$ in the context ϕ . Not writing the identity substitution would enforce that the pattern variable does not depend on ϕ and forces the type of p to be `[⊢ tm B]`. While it is certainly legitimate to use p in the context ϕ , since the empty substitution maps variables from the empty context to ϕ , the type of p is empty; since the context is empty, there are no variables of the type `[⊢ tm B]`. Hence writing `(φ ⊢ p)` would describe an empty pattern. In contrast, types described by

■ **Listing 2** Counting constructors.

```

count:  $\Pi \phi: \text{ctx}. \Pi A: \text{tp}. \Pi M: (\phi \vdash \text{tm } A) . [\text{nat}] =$ 
 $\Lambda \phi \Rightarrow \Lambda A \Rightarrow \Lambda M \Rightarrow \text{rec-case } M \text{ of}$ 
|  $B: \text{tp}, p: (\phi \vdash \text{tm } B); . \text{count } \phi B (\phi \vdash p \dots) \Rightarrow [z]$  % Variable Case
|  $B: \text{tp}, C: \text{tp}, M: (\phi, x: \text{tm } B \vdash \text{tm } C);$  % Abstraction Case
   $\text{count } (\phi, x: \text{tm } B) C (\phi, x: \text{tm } B \vdash M \dots x) : [\text{nat}].$  % IH
   $\text{count } \phi (\text{arr } B C) (\phi \vdash \text{lam } B C \lambda x. M \dots x) \Rightarrow$ 
   $\text{let } X = \text{count } (\phi, x: \text{tm } B) C (\phi, x: \text{tm } B \vdash M \dots x) \text{ in } [s X]$ 
|  $B: \text{tp}, C: \text{tp}, M: (\phi \vdash \text{tm } (\text{arr } B C)), N: (\phi \vdash \text{tm } B);$  % Application Case
   $\text{count } \phi (\text{arr } B C) (\phi \vdash M \dots) : [\text{nat}],$  % IH1
   $\text{count } \phi B (\phi \vdash N \dots) : [\text{nat}].$  % IH2
   $\text{count } \phi C (\phi \vdash \text{app } B C (M \dots) (N \dots)) \Rightarrow$ 
   $\text{let } X = \text{count } \phi (\text{arr } B C) (\phi \vdash M \dots) \text{ in}$ 
   $\text{let } Y = \text{count } \phi B (\phi \vdash N \dots) \text{ in add } (s X) Y$ 

```

■ **Listing 3** Computing length of a context.

```

length =  $\Pi \phi: \text{ctx}. [\text{nat}] =$ 
 $\Lambda \phi \Rightarrow \text{rec-case } \phi \text{ of}$ 
|  $. \text{count } \emptyset \Rightarrow [z]$ 
|  $\psi: \text{ctx}, A: \text{tp}; \text{count } \psi : [\text{nat}] . \text{count } (\psi, x: \text{tm } A) \Rightarrow \text{let } X = \text{count } \psi \text{ in } [s X]$ 

```

meta-variables A or B , for example, are always closed and can be instantiated with any closed object of type tp and we do not associate them with an identity substitution.

In the case for lambda-abstractions, $\text{count } \phi (\text{arr } B C) (\phi \vdash \text{lam } \lambda x. M \dots x)$, we not only list the type of each of the variables occurring in the pattern, but also the induction hypothesis, $\text{count } (\phi, x: \text{tm } B) C (\phi, x: \text{tm } B \vdash M \dots x) : [\text{nat}]$. Although the context grows, the term itself is smaller. In the body of the case, we use the induction hypothesis to determine the size X of $M \dots x$ in the context $\phi, x: \text{tm } B$ and then increment it.

The case for application, $\text{count } \phi C (\phi \vdash \text{app } B C (M \dots) (N \dots))$, is similar. We again list all the types of variables occurring in the pattern as well as the two induction hypotheses. In the body, we determine the size X of $(\phi \vdash M \dots)$ and the size Y of $(\phi \vdash N \dots)$ and then add them.

Computing the length of a context: Induction on the context

As we have the power to abstract and manipulate contexts as first-class objects, we also can reason inductively about them. Contexts are similar to lists and we distinguish between the empty context, written here as \emptyset , and a context consisting of at least one element, written as $\psi, x: \text{tm } A$. In the latter case, we can appeal to the induction hypothesis on ψ (see Listing 3).

3 Core language with well-founded recursion

In this section, we present the core of Beluga's computational language which allows the manipulation of contextual LF objects by means of higher-order functions and primitive recursion. In our presentation of the computation language we keep however our domain abstract simply referring to U , the type of a domain object, and C , the object of a given domain. In fact, our computational language is parametric in the actual domain. To guarantee totality of a program, the domain needs to provide answers for two main questions: 1) how to split on a domain type U and 2) how to determine whether a domain object C is smaller according to some domain-specific measure. We also need to know how to unify two terms and determine when two terms in our domain are equal. In terms of proof-theoretical strength, the language is comparable to Gödel's T or Heyting Arithmetic where the objects

of study are natural numbers. However in our case, U will stand for a (contextual) LF type and C describes a (contextual) LF object.

Types	$\mathcal{I}, \tau ::= [U] \mid \tau_1 \rightarrow \tau_2 \mid \Pi X:U.\tau$
Expressions	$e ::= y \mid [C] \mid \text{fn } y:\tau \Rightarrow e \mid e_1 e_2 \mid \Lambda X:U \Rightarrow e \mid e C$ $\mid \text{let } X = e_1 \text{ in } e_2 \mid \text{rec-case}^{\mathcal{I}} C \text{ of } \vec{b}$
Branches	$b ::= \Delta; \vec{r} . r \Rightarrow e$
Assumptions	$r ::= f \vec{C} C$
Contexts	$\Gamma ::= \cdot \mid \Gamma, y:\tau \mid \Gamma, r : \tau$
Meta Context	$\Delta ::= \cdot \mid \Delta, X:U$

We distinguish between computation variables, simply referred to as variables and written using lower-case letter y ; variables that are bound by Π -types and Λ -abstraction are referred to as meta-variables and written using upper-case letter X . Meta-variables occur inside a domain object. For example we saw earlier the object $(\psi \vdash \text{app } \mathbf{B} \ \mathbf{C} \ (\mathbf{M} \dots) \ (\mathbf{N} \dots))$. Here, ψ , \mathbf{B} , \mathbf{C} , \mathbf{M} , and \mathbf{N} are referred to as meta-variables.

There are three forms of computation-level types τ . The base type $[U]$ is introduced by wrapping a contextual object C inside a box; an object of type $[U]$ is eliminated by a let-expression effectively unboxing a domain object. The non-dependent function space $\tau_1 \rightarrow \tau_2$ is introduced by function abstraction $\text{fn } y:\tau_1 \Rightarrow e$ and eliminated by application $e_1 e_2$; finally, the dependent function type $\Pi X:U.\tau$ which corresponds to universal quantification in predicate logic is introduced by abstraction $\Lambda X:U \Rightarrow e$ over meta-variables X and eliminated by application to a meta objects C written as $e C$. The type annotations on both abstractions ensure that every expression has a unique type. Note that we can index computation-level types τ only by meta objects (but this includes LF contexts!), not by arbitrary computation-level objects. Thus, the resulting logic is just first-order, although the proofs we can write correspond to higher-order functional programs manipulating HOAS objects.

Our language supports pattern matching on a meta-object C using `rec-case`-expressions. Note that one cannot match on a computational object e directly; instead one can bind an expression of type $[U]$ to a meta variable X using `let` and then match on X . We annotate the recursor `rec-case` with the type of the inductive invariant $\Pi \Delta_0.\tau_0$ which the recursion satisfies. Since we are working in a dependently-typed setting, it is not sufficient to simply state the type U of the scrutinee. Instead, we generalize over the index variables occurring in the scrutinee, since they may be refined during pattern matching. Hence, Δ_0 is $\Delta_1, X_0:U_0$ where Δ_1 exactly describes the free meta-variables occurring in U_0 . The intention is that we induct on objects of type U_0 which may depend on Δ_1 . Δ_0 must therefore contain at least one declaration. We also give the return type τ_0 of the recursor, since it might also depend on Δ_0 and might be refined during pattern matching. This is analogous to Coq's `match_as_in_return_with_end` construct.

One might ask whether this form of inductive invariant is too restrictive, since it seems not to capture, e.g., $\Pi \Delta_0.(\tau \rightarrow \Pi X:U_0.\tau')$. While allowing more general invariants does not pose any fundamental issues, we simply note here that the above type is isomorphic to $\Pi(\Delta_0, X:U_0).\tau \rightarrow \tau'$ which is treated by our calculus. Forcing all quantifiers at the outside simplifies our theoretical development; however, our implementation is more flexible.

A branch b_i is expressed as $\Delta_i; \vec{r}_i . r_{i0} \Rightarrow e_i$. As shown in the examples, we explicitly list all pattern variables (i.e. meta-variables) occurring in the pattern in Δ_i . In practice, they often can be inferred (see for example [12]). We also list all valid well-founded recursive calls \vec{r}_i , i.e. r_{ik}, \dots, r_{i1} , for pattern r_{i0} . In practice, they can be also derived dynamically while we check that a given pattern r_{i0} is covering and we give an algorithm in Section 6.

$$\begin{array}{c}
\boxed{\Delta; \Gamma \vdash e : \tau} : \text{Computation } e \text{ has type } \tau \quad \frac{\Gamma(y) = \tau}{\Delta; \Gamma \vdash y : \tau} \quad \frac{\Gamma(r) = \tau \quad r = r'}{\Delta; \Gamma \vdash r' : \tau} \\
\\
\frac{\Delta \vdash C : U}{\Delta; \Gamma \vdash [C] : [U]} \quad \frac{\Delta; \Gamma \vdash e_1 : \tau_2 \rightarrow \tau \quad \Delta; \Gamma \vdash e_2 : \tau_2}{\Delta; \Gamma \vdash e_1 e_2 : \tau} \quad \frac{\Delta; \Gamma \vdash e : \Pi X:U.\tau \quad \Delta \vdash C : U}{\Delta; \Gamma \vdash e C : \llbracket C/X \rrbracket \tau} \\
\\
\frac{\Delta; \Gamma, y:\tau_1 \vdash e : \tau_2}{\Delta; \Gamma \vdash \text{fn } y:\tau_1 \Rightarrow e : \tau_1 \rightarrow \tau_2} \quad \frac{\Delta, X:U; \Gamma \vdash e : \tau}{\Delta; \Gamma \vdash \Lambda X:U \Rightarrow e : \Pi X:U.\tau} \quad \frac{\Delta; \Gamma \vdash e_1 : [U] \quad \Delta, X:U; \Gamma \vdash e_2 : \tau}{\Delta; \Gamma \vdash \text{let } X = e_1 \text{ in } e_2 : \tau} \\
\\
\frac{\mathcal{I} = \Pi \Delta_1. \Pi X_0:U_0. \tau_0 \quad \Delta \vdash C : \llbracket \theta \rrbracket U_0 \quad \Delta \vdash \theta : \Delta_1 \quad b_i : \mathcal{I} \text{ (for all } i) \quad \vec{b} \text{ covers } \mathcal{I}}{\Delta; \Gamma \vdash \text{rec-case}^{\mathcal{I}} C \text{ of } \vec{b} : \llbracket \theta, C/X \rrbracket \tau_0} \\
\\
\boxed{b : \mathcal{I}} : \text{Branch } b \text{ satisfies the invariant } \mathcal{I} \\
\\
\frac{\text{for all } 0 \leq j \leq k. \Delta \vdash_{\mathcal{I}} r_j : \tau_j \quad \Delta; r_k:\tau_k, \dots, r_1:\tau_1 \vdash e : \tau_0}{(\Delta; r_k \dots r_1 . r_0 \Rightarrow e) : \mathcal{I}} \\
\\
\boxed{\Delta \vdash_{\mathcal{I}} r : \tau'} \text{ and } \boxed{\Delta \vdash \vec{C} : \mathcal{I} > \tau'} : \text{Assumption } r/\text{pattern spine } \vec{C} \text{ has type } \tau' \text{ given } \mathcal{I} \\
\\
\frac{\Delta \vdash \vec{C} : \mathcal{I} > \tau'}{\Delta \vdash_{\mathcal{I}} f \vec{C} : \tau'} \quad \frac{\Delta \vdash C : U \quad \Delta \vdash \vec{C} : \llbracket C/X \rrbracket \tau > \tau'}{\Delta \vdash C \vec{C} : \Pi X:U.\tau > \tau'} \quad \frac{\Delta \vdash C : U}{\Delta \vdash C : \Pi Y:U.\tau > \llbracket C/Y \rrbracket \tau}
\end{array}$$

■ **Figure 1** Type system for dependently-typed functional computation language.

The identifier f in assumptions r denotes the local function that is essentially introduced by `rec-case`; this notation is inspired by primitive recursion in *Tutch* [1]. Currently, it just improves the readability of call patterns; however, it is vital for extensions to nested recursion.

3.1 Computation-level Type System

In the typing judgement (Fig. 1), we distinguish between the context Δ for meta-variables from our index domain and the context Γ which includes declarations of computation-level variables. We will tacitly rename bound variables, and maintain that contexts declare no variable more than once. Moreover, we require the usual conditions on bound variables. For example in the rule for Λ -abstraction the meta-variable X must be new and cannot already occur in the context Δ . This can always be achieved via α -renaming. Similarly, in the rule for function abstraction, the variable y must be new and cannot already occur in Γ . We have two variable rules to look up a computation-level variable y and an induction hypothesis r . To verify that the induction hypothesis r' has type τ and its use is valid, we simply check whether there exists $r : \tau$ in Γ where $r = r'$. For now it suffices to think of $=$ as syntactically equivalent.

The most interesting rule is the one for recursion: given the invariant $\mathcal{I} = \Pi \Delta_1. \Pi X_0:U_0. \tau_0$ the expression `rec-case` ^{\mathcal{I}} C of \vec{b} is well-typed under three conditions: First, the meta-object C we are recursing over has some type U and moreover, U is an instance of the type specified in the invariant, i.e. $\Delta_0 = \Delta_1, X_0:U_0$ and $U = \llbracket \theta \rrbracket U_0$ for some meta-substitution θ with domain Δ_1 . Secondly, all branches b_i are well-typed with respect to the given invariant $\Pi \Delta_0. \tau_0$. Finally, \vec{b} must cover the meta-context Δ_0 , i.e., it must be a complete, non-redundant set of patterns covering Δ_0 , and all recursive calls are well-founded. Since the coverage check is domain specific, we leave it abstract for now and return to it when we consider (contextual) LF as one possible domain (see Sec. 5).

Note that we drop the meta-context Δ and the computation context Γ when we proceed to check that all branches satisfy the specified invariant. Dropping Δ is fine, since we require

$$\frac{}{(\text{fn } x:\tau \Rightarrow e) v \longrightarrow [v/x]e} \quad \frac{}{(\Lambda X:U \Rightarrow e) C \longrightarrow \llbracket C/X \rrbracket e} \quad \frac{}{\text{let } X = [C] \text{ in } e \longrightarrow \llbracket C/X \rrbracket e}$$

$$\frac{\exists \text{ unique } (\Delta.r_k, \dots, r_1.r_0 \Rightarrow e) \in \vec{b} \text{ where } r_j = f \vec{C}_j C_{j_0} \text{ such that } \vdash C \doteq C_{00}/\theta}{\text{rec-case}^{\mathcal{I}} C \text{ of } \vec{b} \longrightarrow \llbracket \theta \rrbracket [(\text{rec-case}^{\mathcal{I}} C_{k_0} \text{ of } \vec{b})/r_k, \dots, (\text{rec-case}^{\mathcal{I}} C_{1_0} \text{ of } \vec{b})/r_1]e}$$

■ **Figure 2** Small-step semantics $e \longrightarrow e'$.

the invariant $\Pi\Delta_0.\tau_0$ to be closed. One might object to dropping Γ ; indeed this could be generalized to keeping those assumptions from Γ which do not depend on Δ and generalizing the allowed type of inductive invariant (see our earlier remark).

For a branch $b = \Delta; \vec{r}.r_0 \Rightarrow e$ to be well-typed with respect to a given invariant \mathcal{I} , we check the call pattern r_0 and each recursive call r_j against the invariant and synthesize target types τ_j ($j \geq 0$). We then continue checking the body e against τ_0 , i.e., the target type of the call pattern r_0 , populating the computation context with the recursive calls \vec{r} at their types $\vec{\tau}$. A pattern / recursive call $r_j = f \vec{C}_j$ intuitively corresponds to the given inductive invariant $\mathcal{I} = \Pi\Delta_1.\Pi X_0:U_0.\tau_0$, if the spine \vec{C} matches the specified types in $\Delta_1, X_0:U_0$ and it has intuitively the type $\llbracket C_{j_n}/X_n, \dots, C_{j_0}/X_0 \rrbracket \tau_0$ which we denote with τ'_j .

More generally, we write $\Delta \vdash \theta : \Delta_0$ for a well-typed simultaneous substitution where Δ_0 is the domain and Δ is the range of the substitution. It can be inductively defined (see below) and the standard substitution lemmas hold (see for example [4]).

$$\frac{}{\Delta \vdash \cdot : \cdot} \quad \frac{\Delta \vdash \theta : \Delta_0 \quad \Delta \vdash C : \llbracket \theta \rrbracket U}{\Delta \vdash \theta, C/X : \Delta_0, X : U}$$

4 Operational Semantics

Fig. 2 specifies the call-by-value (cbv) one-step reduction relation $e \longrightarrow e'$; we have omitted the usual congruence rules for cbv. Reduction is deterministic and does not get stuck on closed terms, due to completeness of pattern matching in rec-case . To reduce $(\text{rec-case}^{\mathcal{I}} C \text{ of } \vec{b})$ we find the branch $(\Delta.r_k, \dots, r_1.r_0 \Rightarrow e) \in \vec{b}$ such that the principal argument C_{00} of its clause head $r_0 = f \vec{C}_0 C_{00}$ matches C under meta substitution θ . The reduct is the body e under θ where we additionally replace each place holder r_j of a recursive call by the actual recursive invocation $(\text{rec-case}^{\mathcal{I}} \llbracket \theta \rrbracket C_{j_0} \text{ of } \vec{b})$. The object C_{j_0} in fact just denotes the meta-variable on which we are recursing. We also apply θ to the body e . In the rule, we have lifted out θ . *Values* v in our language are boxed meta objects $[C]$, functions $\text{fn } x:\tau \Rightarrow e$, and $\Lambda X:U.e$.

► **Theorem 1** (Subject reduction). *If $\vdash e : \tau$ and $e \longrightarrow e'$, then $\vdash e' : \tau$.*

Proof. By induction on $e \longrightarrow e'$. ◀

► **Theorem 2** (Progress). *If $\vdash e : \tau$ then either e is a value or $e \longrightarrow e'$.*

Proof. By induction on $\vdash e : \tau$. ◀

5 Contextual LF: Background, Measure, Splitting

If we choose as our domain natural numbers or lists, it may be obvious how to define splitting together with a measure that describes when an object is smaller. Our interest however is to

use the contextual logical framework LF [9] as a general domain language. Contextual LF extends the logical framework LF [8] by packaging an LF objects M of type A together with the context Ψ in which it is meaningful. This allows us to represent rich syntactic structures such as open terms and derivation trees that depend on hypotheses. The core language introduced in Sec. 3 then allows us to implement well-founded recursive programs over these rich abstract syntax trees that correspond to proofs by structural induction.

5.1 Contextual LF

We briefly review contextual LF here. As usual we consider only objects in η -long β -normal form, since these are the only meaningful objects in LF. Further, we concentrate on characterizing well-typed terms; spelling out kinds and kinding rules for types is straightforward.

LF Base Types	P, Q	$::= c \cdot S$
LF Types	A, B	$::= P \mid \Pi x:A.B$
Heads	H	$::= c \mid x \mid p[\sigma]$
Neutral Terms	R	$::= H \cdot S \mid u[\sigma]$
Spines	S	$::= \text{nil} \mid M S$
Normal Terms	M, N	$::= R \mid \lambda x. M$
Substitutions	σ	$::= \cdot \mid \text{id}_\psi \mid \sigma, M \mid \sigma; H$
Variable Substitutions	π	$::= \cdot \mid \text{id}_\psi \mid \sigma; x$
LF Contexts	Ψ, Φ	$::= \cdot \mid \psi \mid \Psi, x:A$

Normal terms are either lambda-abstractions or neutral terms which are defined using a spine representation to give us direct access to the head of a neutral term. Normal objects may contain *ordinary bound variables* x which are used to represent object-level binders and are bound by λ -abstraction or in a context Ψ . Contextual LF extends LF by allowing two kinds of *contextual variables*: the meta-variable u has type $(\Psi.P)$ and stands for a general LF object that has type P and may use the variables declared in Ψ ; the parameter variable p has type $\#(\Psi.A)$ and stands for an LF variable object of type A in the context Ψ .

Contextual variables are associated with a postponed substitution σ which is applied as soon as we instantiate it. More precisely, a meta-variable u stands for a contextual object $\hat{\Psi}.R$ where $\hat{\Psi}$ describes the ordinary bound variables which may occur in R . This allows us to rename the free variables occurring in R when necessary. The parameter variable p stands for a contextual object $\hat{\Psi}.H$ where H must be either an ordinary bound variable from $\hat{\Psi}$ or another parameter variable.

In the simultaneous substitutions σ , we do not make the domain explicit. Rather we think of a substitution together with its domain Ψ and the i -th element in σ corresponds to the i -th declaration in Ψ . We have two different ways of building a substitution entry: either by using a normal term M or a variable x . Note that a variable x is only a normal term M if it is of base type. However, as we push a substitution σ through a λ -abstraction $\lambda x.M$, we need to extend σ with x . The resulting substitution σ, x may not be well-formed, since x may not be of base type and in fact we do not know its type. Hence, we allow substitutions not only to be extended with normal terms M but also with variables x ; in the latter case we write $\sigma; x$. Expression id_ψ denotes the identity substitution with domain ψ while \cdot describes the empty substitution.

Application of a substitution σ to an LF normal form B , written as $[\sigma]B$, is *hereditary* [20] and produces in turn a normal form by removing generated redexes on the fly, possibly triggering further hereditary substitutions.

$$\begin{array}{c}
\boxed{\Delta; \Psi \vdash H \Rightarrow A} \text{ Synthesize type } A \text{ for head } H \\
\frac{\Psi(x) = A \quad \Sigma(c) = A \quad \Delta(p) = \#\Phi.A \quad \Delta; \Psi \vdash \sigma \Leftarrow \Phi}{\Delta; \Psi \vdash x \Rightarrow A \quad \Delta; \Psi \vdash c \Rightarrow A \quad \Delta; \Psi \vdash p[\sigma] \Rightarrow [\sigma]A} \\
\\
\boxed{\Delta; \Psi \vdash S : A > P} \text{ Check spine } S \text{ against } A \text{ with target } P \\
\frac{}{\Delta; \Psi \vdash \text{nil} : P > P} \quad \frac{\Delta; \Psi \vdash M \Leftarrow A \quad \Delta; \Psi \vdash S : [M/x]B > P}{\Delta; \Psi \vdash M S : \Pi x:A.B > P} \\
\\
\boxed{\Delta; \Psi \vdash M \Leftarrow A} \text{ Check normal object } M \text{ against type } A \\
\frac{\Delta; \Psi, x:A \vdash M \Leftarrow B \quad \Delta(u) = \Phi.P \quad \Delta; \Psi \vdash \sigma \Leftarrow \Phi \quad Q = [\sigma]P}{\Delta; \Psi \vdash \lambda x.M \Leftarrow A \rightarrow B \quad \Delta; \Psi \vdash u[\sigma] \Leftarrow Q} \\
\frac{\Delta; \Psi \vdash H \Rightarrow A \quad \Delta; \Psi \vdash S : A > P}{\Delta; \Psi \vdash H \cdot S \Leftarrow P} \\
\\
\boxed{\Delta; \Psi \vdash \sigma \Leftarrow \Phi} \text{ Check substitution } \sigma \text{ against domain } \Phi \\
\frac{}{\Delta; \Psi \vdash \cdot \Leftarrow \cdot} \quad \frac{}{\Delta; (\psi, \Psi^0) \vdash \text{id}_\psi \Leftarrow \psi} \\
\frac{\Delta; \Psi \vdash \sigma \Leftarrow \Phi \quad \Delta; \Psi \vdash M \Leftarrow [\sigma]A}{\Delta; \Psi \vdash (\sigma, M) \Leftarrow (\Phi, x:A)} \quad \frac{\Delta; \Psi \vdash \sigma \Leftarrow \Phi \quad \Delta; \Psi \vdash H \Rightarrow B \quad B = [\sigma]A}{\Delta; \Psi \vdash (\sigma; H) \Leftarrow (\Phi, x:A)}
\end{array}$$

■ **Figure 3** Bi-directional typing for contextual LF.

An LF context Ψ is either a list of bound variable declarations $\overrightarrow{x:A}$ or a context variable ψ followed by such a list. We write Ψ^0 for contexts that do not start with a context variable. We write Ψ, Φ^0 or sometimes Ψ, Φ for the extension of context Ψ by the variable declarations of Φ^0 or Φ , resp. The operation $\text{id}(\Psi)$ that generates an identity substitution for a given context Ψ is defined inductively as follows: $\text{id}(\cdot) = \cdot$, $\text{id}(\psi) = \text{id}_\psi$, and $\text{id}(\Psi, x:A) = \text{id}(\Psi); x$.

We summarize the bi-directional type system for contextual LF in Fig. 3. LF objects may depend on variables declared in the context Ψ and a fixed meta-context Δ which contains contextual variables such as meta-variables u , parameter variables p , and context variables ψ . All typing judgments have access to both contexts and a fixed well-typed signature Σ where we store constants c together with their types and kinds.

5.2 Meta-level Terms and Typing Rules

We lift contextual LF objects to meta-objects to have a uniform definition of all meta-objects. We also define context schemas G that classify contexts.

$$\begin{array}{l}
\text{Context Schemas } G \quad ::= \exists \Phi^0.B \mid G + \exists \Phi^0.B \\
\text{Meta Types } \quad U, V \quad ::= \Psi.P \mid G \mid \#\Psi.A \quad \text{Meta Objects } C, D \quad ::= \hat{\Psi}.R \mid \Psi
\end{array}$$

A consequence of the uniform treatment of meta-terms is that the design of the computation language is modular and parametrized over meta-terms and meta-types. This has two main advantages: First, we can in principle easily extend meta-terms and meta-types without affecting the computation language; second, it will be key to a modular, clean design.

The above definition gives rise to a compact treatment of meta-context Δ . A meta-variable X can denote a meta-variable u , a parameter variable p , or a context variable ψ .

Meta substitution C/X can represent $\hat{\Psi}.R/u$, or Ψ/ψ , or $\hat{\Psi}.x/p$, or $\hat{\Psi}.p'[\pi]/p$ (where π is a variable substitution so that $p[\pi]$ always produces a variable). A meta declaration $X:U$ can stand for $u : \Psi.P$, or $p : \#\Psi.A$, or $\psi : G$. Intuitively, as soon as we replace u with $\hat{\Psi}.R$ in $u[\sigma]$, we apply the substitution σ to R hereditarily. The simultaneous meta-substitution, written as $[\theta]$, is a straightforward extension of the single substitution. For a full definition of meta-substitutions, see [9, 4]. We summarize the typing rules for meta-objects below.

$$\boxed{\Delta \vdash C : U} \text{ Check meta-object } C \text{ against meta-type } U \quad \frac{\Delta; \Psi \vdash R \Leftarrow P}{\Delta \vdash \hat{\Psi}.R : \Psi.P}$$

$$\frac{}{\Delta \vdash \cdot : G} \quad \frac{\Delta(\psi) = G}{\Delta \vdash \psi : G} \quad \frac{\Delta \vdash \Psi : G \quad \exists \Phi^0. B \in G \quad \Delta; \Psi \vdash \sigma \Leftarrow \Phi^0 \quad [\sigma]B = B'}{\Delta \vdash \Psi, x:B' : G}$$

We write $\hat{\Psi}$ for a list of variables obtained by erasing the types from the context Ψ . We have omitted the rules for parameter types $\#\Psi.A$ because they are not important for the further development. Intuitively an object R has type $\#\Psi.A$ if R is either a concrete variable x of type A in the context Ψ or a parameter variable p of type A in the context Ψ . This can be generalized to account for re-ordering of variables allowing the parameter variable p to have some type A' in the context Ψ' s.t. there exists a permutation substitution π on the variables such that $\Psi \vdash \pi : \Psi'$ and $A = [\pi]A'$.

5.3 Well-founded Structural Subterm Order

There are two key ingredients to guarantee that a given function is total: we need to ensure that all the recursive calls are on smaller arguments according to a well-founded order and the function covers all possible cases. We define here a well-founded structural subterm order on contexts and contextual objects similar to the subterm relations for LF objects[10]. For simplicity, we only consider here non-mutual recursive type families; those can be incorporated using the notion of subordination [19].

We first define an ordering on contexts: $\boxed{\Psi \preceq \Phi}$, read as “context Ψ is a subcontext of Φ ”, shall hold if all declarations of Ψ are also present in the context Φ , i.e., $\Psi \subseteq \Phi$. The strict relation $\boxed{\Psi \prec \Phi}$, read as “context Ψ is strictly smaller than context Φ ” holds if $\Psi \preceq \Phi$ but Ψ is strictly shorter than Φ .

Further, we define three relations on contextual objects $\hat{\Psi}.M$: a strict subterm relation \prec , an equivalence relation \equiv , and an auxiliary relation \preceq .

$$\frac{\hat{\Psi} \subseteq \hat{\Phi} \text{ or } \hat{\Phi} \subseteq \hat{\Psi} \quad \pi \text{ is a variable subst. s.t. } M = [\pi]N}{\hat{\Psi}.M \equiv \hat{\Phi}.N}$$

$$\frac{\hat{\Psi}.M \preceq \hat{\Phi}.N_i \text{ for some } 1 \leq i \leq n}{\hat{\Psi}.M \prec \hat{\Phi}.h \cdot N_1 \dots N_n \text{ nil}} \quad \frac{\hat{\Psi}.M \prec \hat{\Phi}.N}{\hat{\Psi}.M \prec \hat{\Phi}.N} \quad \frac{\hat{\Psi}.M \equiv \hat{\Phi}.N}{\hat{\Psi}.M \preceq \hat{\Phi}.N} \quad \frac{\hat{\Psi}.M \preceq \hat{\Phi}.x.N}{\hat{\Psi}.M \preceq \hat{\Phi}.\lambda x.N}$$

$\hat{\Psi}.M$ is a strict subterm of $\hat{\Phi}.N$ if M is a proper subterm of N modulo α -renaming and weakening. Two terms $\hat{\Psi}.M$ and $\hat{\Phi}.N$ are structurally equivalent, if they describe the same term modulo α -renaming and possible weakening. To allow mutual recursive definitions and richer subterm relationships, we can incorporate subordination information and generalize the variable substitution π (see for example [10] for such a generalization). Using the defined subterm order, we can easily verify that the recursive calls in the examples are structurally smaller.

The given subterm relation is well-founded. We define the measure $||\Psi||$ of a ground context Ψ^0 or its erasure $\hat{\Psi}^0$ as its length $|\Psi|$. The measure $||\hat{\Psi}.M||$ of a contextual object

$\hat{\Psi}.M$, is the measure $\|M\|$ of M . The latter is defined inductively by:

$$\begin{aligned} \|h \cdot M_1 \dots M_n \text{ nil}\| &= 1 + \max(\|M_1\|, \dots, \|M_n\|) \\ \|\lambda x.M\| &= \|M\| \end{aligned}$$

► **Theorem 3** (Order on contextual objects is well-founded). *Let θ be a grounding meta-substitution.*

1. If $C \prec C'$ then $\|\llbracket\theta\rrbracket C\| < \|\llbracket\theta\rrbracket C'\|$.
2. If $C \equiv C'$ then $\|\llbracket\theta\rrbracket C\| = \|\llbracket\theta\rrbracket C'\|$.
3. If $C \preceq C'$ then $\|\llbracket\theta\rrbracket C\| \leq \|\llbracket\theta\rrbracket C'\|$.

5.4 Case Splitting

Our language allows pattern matching and recursion over contextual objects. For well-formed recursors ($\text{rec-case}^{\mathcal{I}} C$ of \vec{b}) with invariant $\mathcal{I} = \Pi\Delta.\Pi X:U.\tau$, branches \vec{b} need to cover all different cases for the argument C of type U . We only take the shape of U into account and generate the unique complete set $\mathcal{U}_{\Delta \vdash U}$ of non-overlapping shallow patterns by splitting meta-variable X of type U .

If $U = \Psi.P$ is a base type, then intuitively the set $\mathcal{U}_{\Delta \vdash U}$ contains all neutral terms $R = H \cdot S$ where H is a constructor c , a concrete variable x from Ψ or a parameter variable $p[\text{id}_\psi]$ denoting a variable from the context variable ψ , and S is a most general spine s.t. the type of R is an instance of P in the context Ψ . We note that when considering only closed terms it suffices to consider only terms with $H = c$. However, when considering terms with respect to a context Ψ , we must generate additional cases covering the scenario where H is a variable – either a concrete variable x if $x:A$ is in Ψ or a parameter variable if the context is described abstractly using a context variable ψ .

If U denotes a context schema G , we generate all shallow context patterns of type G . This includes the empty context and a context extended with a declaration formed by $\Psi, x:A$.

From $\mathcal{U}_{\Delta \vdash U}$ we generate the complete minimal set $\mathcal{C} = \{\Delta_i; r_{ik}, \dots, r_{i1}.r_{i0} \mid 1 \leq i \leq n\}$ of possible, non-overlapping cases where the i -th branch shall have the well-founded recursive calls r_{ik}, \dots, r_{i1} for the case r_{i0} . For the given branches \vec{b} to be covering, each element in \mathcal{C} must correspond to one branch b_i .

Splitting on a Contextual Type

Following [5, 17], the patterns R of type $\Psi.P$ are computed by brute force: We first synthesize a set $\mathcal{H}_{\Delta, \Psi}$ of *all* possible heads together with their type: constants $c \in \Sigma$, variables $x \in \Psi$, and parameter variables if Ψ starts with a context variable ψ .

$$\begin{aligned} \mathcal{H}_{\Delta, \Psi} = & \{(\Delta; \Psi \vdash c : A) \mid (c:A) \in \Sigma\} \\ & \cup \{(\Delta; \Psi \vdash x : A) \mid (x:A) \in \Psi\} \\ & \cup \{(\Delta, \overrightarrow{X:\vec{U}}, p:\#(\psi.B'); \Psi \vdash p[\text{id}_\psi] : B') \mid \Psi = \psi, \Psi^0 \text{ and } \psi:G \in \Delta \text{ and } \exists x:\vec{A}.B \in G \\ & \text{and } \text{genMV}(\psi.A_i) = (X_i:U_i, M_i) \text{ for all } i, \text{ and } B' = [\overrightarrow{M/x}]B \} \end{aligned}$$

See Fig. 4. Using a head H of type A from the set $\mathcal{H}_{\Delta, \Psi}$, we then generate, if possible, the most general pattern $H \cdot S$ whose target type is unifiable with P in the context Ψ . We describe unification using the judgment $\boxed{\Delta; \Psi \vdash Q \doteq P / (\Delta', \theta)}$. If unification succeeds then $\llbracket\theta\rrbracket Q = \llbracket\theta\rrbracket P$ and $\Delta' \vdash \theta : \Delta$.

$$\begin{array}{c}
\boxed{\text{genMV}(\Psi.A) = (X:U, M)} \quad \text{Generation of a lowered meta variable} \\
\text{genMV}(\Psi.\overrightarrow{\Pi x:\dot{A}.P}) = (u : (\Psi, \overrightarrow{x:\dot{A}.P}), \lambda \vec{x}.u[\text{id}(\Psi, \overrightarrow{x:\dot{A}})]) \text{ for a fresh meta variable } u \\
\boxed{\Delta; \Psi \vdash R : A \Leftarrow P / (\Delta', \theta, R_0)} \quad \text{Extending } R:A \text{ to most general normal term } R_0 : \llbracket \theta \rrbracket P. \\
\frac{\Delta; \Psi \vdash Q \doteq P / (\Delta_0, \theta)}{\Delta; \Psi \vdash R : Q \Leftarrow P / (\Delta_0, \theta, \llbracket \theta \rrbracket R)} \\
\frac{\text{genMV}(\Psi.A) = (X:U, M) \quad \Delta, X:U; \Psi \vdash R \quad M : [M/x]B \Leftarrow P / (\Delta_0, \theta, R')}{\Delta; \Psi \vdash R : \Pi x:A.B \Leftarrow P / (\Delta_0, \theta, R')}
\end{array}$$

■ **Figure 4** Generation of most general normal objects and call patterns.

$\boxed{\Delta; \Psi \vdash R : A \Leftarrow P / (\Delta', \theta, R_0)}$ describes the generation of a normal pattern where all the elements on the left side of $/$ are inputs and the right side is the output, which satisfies $\Delta' \vdash \theta : \Delta$ and $\Delta'; \llbracket \theta \rrbracket \Psi \vdash R \Rightarrow \llbracket \theta \rrbracket A$ and $\Delta'; \llbracket \theta \rrbracket \Psi \vdash R_0 \Leftarrow \llbracket \theta \rrbracket P$. To generate a normal term R_0 of the expected base type, we start with head $H : A$. As we recursively analyze A , we generate all the arguments H is applied to until we reach an atomic type Q . If Q unifies with the expected type P , then generating a most general neutral term with head H succeeds.

$$\begin{aligned}
\mathcal{U}_{\Delta \vdash \Psi, P} = \{ (\Delta'' \vdash \hat{\Phi}.R : \Phi.Q) \mid (\Delta'; \Psi \vdash H : A) \in \mathcal{H}_{\Delta, \Psi} \text{ and} \\
\Delta'; \Psi \vdash H : A \Leftarrow P / (\Delta'', \theta, R) \text{ and } \Phi = \llbracket \theta \rrbracket \Psi \text{ and } Q = \llbracket \theta \rrbracket P \}
\end{aligned}$$

Splitting on a Context Schema

Spitting a context variable of schema G generates the empty context and the non-empty contexts $(\phi, x:B')$ for each possible form of context entry $\exists \Phi^0.B \in G$.

$$\begin{aligned}
\mathcal{U}_{\Delta \vdash G} = \{ (\Delta \vdash \cdot : G) \} \\
\cup \{ (\Delta, \phi:G, \overrightarrow{X:\dot{U}} \vdash (\phi, x:[M/x]B) : G) \mid \phi \text{ a fresh context variable and} \\
\text{for any } \exists \overrightarrow{x:\dot{A}.B} \in G. \text{genMV}(\psi.A_i) = (X_i:U_i, M_i) \text{ for all } i \}
\end{aligned}$$

► **Theorem 4** (Splitting on meta-types). *The set $\mathcal{U}_{\Delta \vdash U}$ of meta-objects generated is non-redundant and complete.*

Proof. $\mathcal{U}_{\Delta \vdash G}$ is obviously non-redundant. $\mathcal{U}_{\Delta \vdash \Psi, P}$ is non-redundant since all generated neutral terms have distinct heads. Completeness is proven by cases. ◀

6 Generation of Call Patterns and Coverage

Next, we explain the generation of call patterns, i.e. well-founded recursive calls as well as the actual call pattern being considered.

► **Definition 5** (Generation of call patterns). Given the invariant $\mathcal{I} = \Pi(\Delta_1, X_0:U_0).\tau_0$ where $\Delta_1 = X_n:U_n, \dots, X_1:U_1$, the set \mathcal{C} of call patterns $(\Delta_i ; r_{ik}:\tau_{ik}, \dots, r_{i1}:\tau_{i1} \cdot r_{i0})$ is generated as follows:

- For each meta-object $\Delta_i \vdash C_{i0} : V_i$ in $\mathcal{U}_{\Delta_0 \vdash U_0}$, we generate using unification, if possible, a call pattern $r_{i0} = f \ C_{in} \dots C_{i1} \ C_{i0}$ s.t. $\tau_{i0} = \llbracket C_{in}/X_n, \dots, C_{i1}/X_1, C_{i0}/X_0 \rrbracket \tau_0$ and

$\Delta_i \vdash r_{i0} : \tau_{i0}$. This may fail if V_i is not an instance of the scrutinee type U_0 ; then, the case C_{i0} is impossible.

- Further, for all $1 \leq j \leq k$, $\Delta_i = Y_k:V_k, \dots, Y_1:V_1$, we generate a recursive call $r_{ij} = f C_{jn} \dots C_{j1} Y_j$ s.t. $\tau_{ij} = \llbracket C_{jn}/X_n, \dots, C_{j1}/X_1, Y_j/X_0 \rrbracket \tau_0$ and $\Delta_i \vdash r_{ij} : \tau_{ij}$, if $Y_j \prec C_{i0}$. This may also fail, if V_i is not an instance with U_0 ; in this case V_i does not give rise to recursive call.

► **Theorem 6** (Pattern generation). *The set \mathcal{C} of call patterns generated is non-redundant and complete and the recursive calls are well-founded.*

Proof. Using Theorem 4 and the properties of unification. ◀

► **Definition 7** (Coverage). We say \vec{b} covers \mathcal{I} iff for every $\Delta_i ; \vec{r}_i : \vec{\tau}_i . r_{i0} \in \mathcal{C}$ where \mathcal{C} is the set of call patterns given \mathcal{I} , we have one corresponding $\Delta_i ; \vec{r}_i : \vec{\tau}_i . r_{i0} \Rightarrow e_i \in \vec{b}$ and vice versa.

7 Termination

We now prove that every well-typed closed program e terminates (halts) by a standard reducibility argument; closely related is [21]. The set \mathcal{R}_τ of reducible closed programs $\cdot ; \vdash e : \tau$ is defined by induction on the size of τ .

Contextual Type	$\mathcal{R}_{[U]}$	$= \{e \mid \cdot ; \vdash e : [U] \text{ and } e \text{ halts}\}$
Function Type	$\mathcal{R}_{\tau' \rightarrow \tau}$	$= \{e \mid \cdot ; \vdash e : \tau' \rightarrow \tau \text{ and } e \text{ halts and } \forall e' \in \mathcal{R}_{\tau'} . e e' \in \mathcal{R}_\tau\}$
Dependent Type	$\mathcal{R}_{\Pi X:U.\tau}$	$= \{e \mid \cdot ; \vdash e : \Pi X:U.\tau \text{ and } e \text{ halts and } \forall C.s.t. \vdash C : U . e C \in \mathcal{R}_{[C/X]\tau}\}$
Context	\mathcal{R}_Γ	$= \{\eta \mid \cdot ; \vdash \eta : \Gamma \text{ and } \eta(x) \in \mathcal{R}_\tau \text{ for all } (x:\tau) \in \Gamma\}$

For the size of τ all meta types U shall be disregarded, thus, the size is invariant under meta substitution C/X . We also note that since reduction $e \longrightarrow e'$ is deterministic, e halts if and only if e' halts.

► **Lemma 8** (Expansion closure).

1. If $\cdot ; \vdash e : \tau$ and $e \longrightarrow e'$ and $e' \in \mathcal{R}_\tau$, then $e \in \mathcal{R}_\tau$.
2. If $\cdot ; \vdash e : \tau$ and $e \longrightarrow^* e'$ and $e' \in \mathcal{R}_\tau$, then $e \in \mathcal{R}_\tau$.

Proof. The first statement, by induction on the size of type τ . The second statement, inductively on \longrightarrow^* . ◀

► **Lemma 9** (Fundamental Lemma).

If $\Delta; \Gamma \vdash e : \tau$ and grounding substitution θ s.t. $\cdot \vdash \theta : \Delta$ and $\eta \in \mathcal{R}_{[\theta]\Gamma}$ then $[\eta][\theta]e \in \mathcal{R}_{[\theta]\tau}$.

Proof. By induction on $\Delta; \Gamma \vdash e : \tau$. In the interesting case of recursion rec–case, we make essential use of coverage and structural descent in the recursive calls. ◀

► **Theorem 10** (Termination). *If $\cdot ; \vdash e : \tau$ then e halts.*

Proof. Taking the empty meta-context Δ and empty computation-level context Γ , we obtain $e \in \mathcal{R}_\tau$ by the fundamental lemma, which implies that e halts by definition of τ . ◀

8 Related Work

Our work is most closely related to [16] where the authors propose a modal lambda-calculus with iteration to reason about closed HOAS objects. In their work the modal type \Box describes closed LF objects. Our work extends this line to allow open LF objects and define functions by pattern matching and well-founded recursion.

Similar to our approach, Schürmann [15] presents a meta-logic \mathcal{M}^2 for reasoning about LF specifications and describes the generation of splits and well-formed recursive calls. However, \mathcal{M}^2 does not support higher-order computations. Moreover, the foundation lacks first-class contexts, but all assumptions live in an ambient context. This makes it less direct to justify reasoning with assumptions, but maybe more importantly complicates establishing meta-theoretic results such as proving normalization.

Establishing well-founded induction principles to support reasoning about higher-order abstract syntax specifications has been challenging and a number of alternative approaches have been proposed. These approaches have led to new reasoning logics – either based on nominal set theory [14] or on nominal proof theory [7]. Proving consistency of these theories in the presence of induction is often significantly more complicated [18]. Our work shows that reasoning about HOAS representations can be supported using first-order logic by modelling HOAS objects as contextual objects. As a consequence, we can directly take advantage of established proof and mechanization techniques. This also opens up the possibility of supporting contextual reasoning as a domain in other systems.

9 Conclusion

We have developed a core language with structural recursion for implementing total functions about LF specification. We describe a sound coverage algorithm which, in addition to verifying that there exists a branch for all possible contexts and contextual objects, also generates and verifies valid primitive recursive calls. To establish consistency of our core language we prove termination using reducibility semantics.

Our framework can be extended to handle mutual recursive functions: By annotating a given `rec`-case-expression with a list of invariants using the subordination relation, we can generate well-founded recursive calls matching each of the invariants. Based on these ideas, we have implemented a totality checker in Beluga. We also added reasoning principles for inductive types [4] that follow well-trodden paths; we must ensure that our inductive type satisfies the positivity restriction and define generation of patterns for them.

Our language not only serves as a core programming language but can be interpreted by the Curry-Howard isomorphism as a proof language for interactively developing proofs about LF specifications. In the future, we plan to implement and design such a proof engine and to generalize our work to allow lexicographic orderings and general well-founded recursion.

Acknowledgements. We thank Sherry Shanshan Ruan for her work during her Summer Undergraduate Research Internship in 2013 at the beginning of this project.

References

- 1 Andreas Abel. *Tutch User's Guide*. Carnegie-Mellon University, Pittsburgh, PA, 2002. Section 7.1: Proof terms for structural recursion.
- 2 Andreas Abel and Brigitte Pientka. Higher-order dynamic pattern unification for dependent types and records. In TLCA'11, volume 6690 of *LNCS*, pages 10–26. Springer, 2011.

- 3 Olivier Savary Belanger, Stefan Monnier, and Brigitte Pientka. Programming type-safe transformations using higher-order abstract syntax. In *CPP'13*, volume 8307 of *LNCS*, pages 243–258. Springer, 2013.
- 4 Andrew Cave and Brigitte Pientka. Programming with binders and indexed data-types. In *POPL'12*, pages 413–424. ACM, 2012.
- 5 Joshua Dunfield and Brigitte Pientka. Case analysis of higher-order data. *ENTCS*, 228:69–84, 2009.
- 6 Francisco Ferreira and Brigitte Pientka. Bidirectional elaboration of dependently typed languages. In *PPDP'14*. ACM, 2014.
- 7 Andrew Gacek, Dale Miller, and Gopalan Nadathur. Combining generic judgments with recursive definitions. In *LICS'08*, pages 33–44. IEEE CS Press, 2008.
- 8 Robert Harper, Furio Honsell, and Gordon Plotkin. A framework for defining logics. *JACM*, 40(1):143–184, 1993.
- 9 Aleksandar Nanevski, Frank Pfenning, and Brigitte Pientka. Contextual modal type theory. *ACM TOCL*, 9(3):1–49, 2008.
- 10 Brigitte Pientka. Verifying termination and reduction properties about higher-order logic programs. *JAR*, 34(2):179–207, 2005.
- 11 Brigitte Pientka. A type-theoretic foundation for programming with higher-order abstract syntax and first-class substitutions. In *POPL'08*, pages 371–382. ACM, 2008.
- 12 Brigitte Pientka. An insider's look at LF type reconstruction: Everything you (n)ever wanted to know. *JFP*, 1(1–37), 2013.
- 13 Brigitte Pientka and Joshua Dunfield. Beluga: A framework for programming and reasoning with deductive systems (system description). In *IJCAR'10*, volume 6173 of *LNCS*, pages 15–21. Springer, 2010.
- 14 Andrew Pitts. Nominal logic, a first order theory of names and binding. *Inf. Comput.*, 186(2):165–193, 2003.
- 15 Carsten Schürmann. *Automating the Meta Theory of Deductive Systems*. PhD thesis, Department of Computer Science, Carnegie Mellon University, 2000. CMU-CS-00-146.
- 16 Carsten Schürmann, Joëlle Despeyroux, and Frank Pfenning. Primitive recursion for higher-order abstract syntax. *TCS*, 266(1-2):1–57, 2001.
- 17 Carsten Schürmann and Frank Pfenning. A coverage checking algorithm for LF. In *TPHOLS'03*, volume 2758 of *LNCS*, pages 120–135, Rome, Italy, 2003. Springer.
- 18 Alwen Tiu and Alberto Momigliano. Cut elimination for a logic with induction and co-induction. *J. Applied Logic*, 10(4):330–367, 2012.
- 19 Roberto Virga. *Higher-Order Rewriting with Dependent Types*. PhD thesis, Department of Mathematical Sciences, Carnegie Mellon University, 1999. CMU-CS-99-167.
- 20 Kevin Watkins, Iliano Cervesato, Frank Pfenning, and David Walker. A concurrent logical framework I: Judgements and properties. Technical report, School of Computer Science, Carnegie Mellon University, Pittsburgh, 2003.
- 21 Hongwei Xi. Dependent types for program termination verification. *HOSC*, 15(1):91–131, 2002.

Polynomial Time in the Parametric Lambda Calculus

Brian F. Redmond

Department of Computing, Mathematics and Statistical Sciences
Grande Prairie Regional College
10726 – 106 Avenue, Grande Prairie, AB, T8V 4C4, Canada
bredmond@GPRC.ab.ca

Abstract

In this paper we investigate Implicit Computational Complexity via the parametric lambda calculus of Ronchi Della Rocca and Paolini [13]. We show that a particular instantiation of the set of input values leads to a characterization of polynomial time computations in a similar way to Lafont’s Soft Linear Logic [9]. This characterization is manifestly type-free and does not require any ad hoc extensions to the pure lambda calculus. Moreover, there is a natural extension to nondeterminism with the addition of explicit products.

1998 ACM Subject Classification F.4.1 Mathematical Logic

Keywords and phrases Parametric Lambda Calculus, Polynomial Time Complexity, Combinators, Nondeterminism and Explicit Products, Implicit Computational Complexity

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.288

1 Introduction

There is an inherent problem in studying computational complexity within the lambda calculus in that a single beta reduction step can effectively square the size of a term, thus leading to an exponential size term after only a linear number of steps. It therefore seems unreasonable to simply count the number of beta reduction steps as a measure of the computational complexity of reduction. Indeed, there have been studies, most recently in [1], which give more reasonable measures of the complexity of a given λ -term using explicit substitutions and notions of sharing. These delicate issues are entirely avoided here as all beta reduction steps, except for a constant number (which does not depend on the size of the input), do not increase the size of the lambda term. Therefore, we feel justified in taking a very simplistic approach to measuring the complexity of reduction of a given lambda term. We define a simple “by-value” operational semantics and define the complexity of reduction as the size of its corresponding evaluation tree, which we show is polynomial in the size of the input binary word. We claim that any such reduction can be simulated on a Turing machine with polynomial overhead.

Most studies of complexity within the (pure) lambda calculus rely on typing restrictions to ensure that terms are strongly normalizing. For example, it is well known that terms in the simply typed lambda calculus are strongly normalizing and, moreover, that the class of representable numerical functions is precisely the class of *extended polynomials* [15]. Adding an (impredicative) operation of abstraction on types, as in Girard’s system **F** [6], greatly increases the class of representable functions to the class of functions provably total in second order Peano arithmetic [7]. Nevertheless, the system remains strongly normalizing. A system somewhere in the middle of these two extremes is obtained by stratifying type abstraction into a finite number of levels. Indeed, in this case, the class of representable functions



© Brian F. Redmond;

licensed under Creative Commons License CC-BY

13th International Conference on Typed Lambda Calculi and Applications (TLCA’15).

Editor: Thorsten Altenkirch; pp. 288–301



Leibniz International Proceedings in Informatics

LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

is precisely the class of super-elementary functions, i.e. the class of \mathcal{E}_4 in Grzegorzczuk's subrecursive hierarchy [10]. Finally, there are also well known studies on the complexity of beta equivalence in the simply typed lambda calculus [16], as well as related results when restricting to low functional orders [14].

Another fruitful approach to investigating complexity in the lambda calculus is through linear logic and related type systems. In these studies, it is not typically type abstraction that is limited, but instead duplication is controlled via the modified exponentials ! and § (see [5] and [2], for example). These type systems were derived from their corresponding systems of (light) linear logic and proof nets, and illustrate the computational power of duplication in the lambda calculus. Unfortunately, type checking and type inference are undecidable in the presence of unrestricted polymorphism [4].

In this paper we take an entirely different approach to studying complexity within the lambda calculus via the parametric lambda calculus [13]. With a frugal choice of so-called input values, we show that strong normalization is guaranteed, yet the system remains expressive enough to capture polynomial time computations. Moreover, there is a natural extension to nondeterminism with the addition of explicit products. In contrast to the above studies, however, this approach does not rely on typing restrictions as we work in an entirely type-free setting. Nevertheless, we believe a system of intersection types can be introduced post hoc if desired.

This work is closely related to the author's work on Bounded Combinatory Logic [12]. In that work, the usual Curry combinators B, C, K, W are introduced, but the duplication combinator W has one of its arguments restricted to a proper subset of combinators, namely the BCK -combinators. This ensures that only *affine linear* terms are duplicated, and leads to a simple characterization of polynomial time computations. The "moral" analogue in the lambda calculus corresponds to a particular instantiation of the parametric lambda calculus, which is the focus of the current paper. However, the systems are not equivalent and in fact use completely different reduction strategies and encodings. For this reason, we chose to present this work independently and investigate the relationship between the two system in future work.

2 An Instance of the Parametric Lambda Calculus

One of the aims of the parametric lambda calculus is to study in a uniform way various systems of the pure lambda calculus, in particular its call-by-name and call-by-value versions [13]. This is done by restricting beta reduction to subsets of lambda terms, called input values, that satisfy certain closure conditions. These closure conditions guarantee important properties like confluence are satisfied. In this paper we study various instantiations of the parametric lambda calculus in the context of Implicit Computational Complexity (ICC). We refer the reader to [13] for much of the notation used as well as some of the basic definitions in the lambda calculus. However, all nonstandard notation, terminology, and definitions will be explicitly stated.

The set of lambda terms, Λ , is defined in the usual manner. We assume a countably infinite set of variables, Var , and define the set of lambda terms Λ as follows:

$$M, N ::= x \mid (MN) \mid (\lambda x.M)$$

where $x \in \text{Var}$. For notational convenience we tacitly assume the standard conventions regarding parentheses and for contracting multiple lambda abstractions. We use the symbol \equiv to denote the syntactical identity of terms up to α -congruence, and $M[N/x]$ denotes the

capture-free substitution of N for x in M . Finally, we need the notion of a context $C[\cdot]$ as defined in [13], Def. 1.1.11.

Let M be a generic lambda term. For each subterm of M of the form $\lambda x.P$, let $\#(\lambda x.P)$ denote the number of times x occurs free in P . The **size**¹ of a lambda term M , denoted $s(M)$, is defined by induction on M as follows: $s(x) = 1$ if x is a variable, $s(\lambda x.P) = 1 + s(P)$ and $s(PQ) = s(P) + s(Q)$. Recall from [13] that a subset Δ of Λ is called a set of **input values** if it satisfies the following three conditions:

1. $\text{Var} \subseteq \Delta$
2. If P and Q are in Δ , then so is $P[Q/x]$, for each $x \in \text{Var}$
3. If $M \in \Delta$ and $M \rightarrow_{\Delta} N$, then $N \in \Delta$

In condition 3, the symbol \rightarrow_{Δ} denotes a one-step Δ -reduction [13].

In this paper we introduce a set of input values, denoted Φ , consisting of “pseudo-affine” terms. This is made precise in the following definition: Let Φ denote the smallest subset of Λ satisfying the following closure properties:

- $x \in \text{Var}$ implies $x \in \Phi$,
- If $M, N \in \Phi$, then $(MN) \in \Phi$,
- If $M \in \Phi$ and $x \in \text{Var}$ and x occurs free at most once in M , then $(\lambda x.M) \in \Phi$.

We shall often refer to Φ as the set of **player** terms and to $\Lambda \setminus \Phi$ as the set of **opponent** terms. For example, the closed terms $I \equiv \lambda x.x$, $B \equiv \lambda xyz.x(yz)$, $C \equiv \lambda xyz.xzy$, and $K \equiv \lambda xy.x$ are all player terms, but $D \equiv \lambda x.xx$ is an opponent term. Note that *free* variables may appear more than once in a player term. This follows terminology introduced in [12].

The one-step Φ -reduction \rightarrow_{Φ} is defined as the contextual closure of the following rule:

$$(\lambda x.M)N \rightarrow_{\Phi} M[N/x] \quad \text{if and only if} \quad N \in \Phi$$

Let \rightarrow_{Φ}^* denote the reflexive and transitive closure of \rightarrow_{Φ} . A redex $R \equiv (\lambda x.P)Q$ with $(\lambda x.P)$ an opponent term and $Q \in \Phi$ is called an **opponent redex**. A redex $R \equiv (\lambda x.P)Q \in \Phi$ is called a **player redex**. This terminology extends to their respective reductions as well.

Define the **degree** of a term M , denoted $d(M)$, as follows: $d(M) = 0$ if $M \in \Phi$ and $d(MN) = d(M) + d(N)$ if $MN \in \Lambda \setminus \Phi$ and $d(\lambda x.P) = 1 + d(P)$ if $\lambda x.P \in \Lambda \setminus \Phi$. For example, the degree of the term $\lambda x.(\lambda y.yy)x$ is 2.

► **Theorem 1** (confluence). *The subset Φ as defined above is a set of input values. Therefore, the reduction relation \rightarrow_{Φ}^* is confluent by Theorem 1.2.5 in [13].²*

Proof. Condition 1 is obvious. For condition 2, we use induction on the structure of P :

- If $P \equiv y$, then $P[Q/x]$ is either $Q \in \Phi$ if $y \equiv x$, or $y \in \Phi$ if $y \neq x$.
- If $P \equiv P_1P_2 \in \Phi$, then $P_1 \in \Phi$ and $P_2 \in \Phi$. By the induction hypothesis, we have $P_1[Q/x], P_2[Q/x] \in \Phi$. Thus, $P[Q/x] \equiv P_1[Q/x]P_2[Q/x] \in \Phi$.
- If $P \equiv \lambda z.M \in \Phi$, then $M \in \Phi$ and z occurs free at most once in M . If $z \equiv x$, then $P[Q/x] \equiv P \in \Phi$. Otherwise, by α -congruence, we may assume that $z \notin \text{FV}(Q)$, so z occurs free at most once in $M[Q/x]$. By the induction hypothesis we have $M[Q/x] \in \Phi$. Thus, $P[Q/x] \equiv \lambda z.M[Q/x] \in \Phi$.

Finally, for condition 3, we proceed by induction on the structure of M :

- If M is a variable, then the result is vacuously true.

¹ This is called *length* in [8].

² The set $\text{Var} \cup \Phi^0$, where Φ^0 denotes the set of closed terms in Φ , also forms a set of input values. However, as we shall see in Section 3.1, pseudo-affine terms provide more flexibility when defining programs.

- If $M \equiv P_1 P_2 \in \Phi$ and the Φ -redex/reduction is entirely in P_i ($i \in \{1, 2\}$), then $P_i \rightarrow_{\Phi} P'_i$. By the induction hypothesis $P'_i \in \Phi$, so either $N \equiv P'_1 P_2 \in \Phi$ or $N \equiv P_1 P'_2 \in \Phi$. On the other hand, if $M \equiv (\lambda x.P)Q$ and $(\lambda x.P)Q \rightarrow_{\Phi} P[Q/x] \equiv N$, then $N \in \Phi$ by condition 2.
 - If $M \equiv \lambda x.P \in \Phi$, then x occurs free in $P \in \Phi$ at most once, and $P \rightarrow_{\Phi} P'$. By the induction hypothesis, $P' \in \Phi$. Moreover, since the redex contracted in $P \in \Phi$ must be a player redex, the number of times x occurs free in P' is no more than 1. Thus, $\lambda x.P' \in \Phi$.
- Therefore, Φ is a valid set of input values. ◀

Note that any term in the $\lambda\Phi$ -calculus has the form: $\lambda x_1 \cdots x_n. \zeta M_1 \cdots M_m$ ($n, m \geq 0$) where ζ is either a variable or a Φ -redex or a head block. (A head block is a term $(\lambda x.P)Q$, where $Q \notin \Phi$.) We say that a term is in Φ -normal form if it has the form $\lambda x_1 \cdots x_n. \zeta M_1 \cdots M_m$, where M_i is in Φ -normal form ($1 \leq i \leq m$) and ζ is either a variable or a head block $(\lambda x.P)Q$, where both P and Q are in Φ -normal form. Note that a lambda term is in Φ -normal form iff it contains no Φ -redexes.

► **Lemma 2.** *If P is any term and Q is a player term, then $d(P) = d(P[Q/x])$.*

Proof. First note that if P is a player term, then so is $P[Q/x]$ and both have degree 0. So assume P is an opponent term and argue by induction on P . If $P \equiv P_1 P_2$, then by the induction hypothesis, $d(P_1[Q/x]) = d(P_1)$ and $d(P_2[Q/x]) = d(P_2)$. Thus, $d(P[Q/x]) = d(P_1[Q/x]P_2[Q/x]) = d(P_1[Q/x]) + d(P_2[Q/x]) = d(P_1) + d(P_2) = d(P)$. If $P \equiv \lambda z.P_1$, then $d(P) = 1 + d(P_1)$. If $z \equiv x$, then $P[Q/x] \equiv P$ and the result follows. If $z \not\equiv x$, then $P[Q/x] \equiv \lambda z.P_1[Q/x]$ where we may assume by α -congruence that $z \notin \text{FV}(Q)$. Then by the induction hypothesis, we have $d(P_1) = d(P_1[Q/x])$. Thus, $d(P) = 1 + d(P_1) = 1 + d(P_1[Q/x]) = d(P[Q/x])$. ◀

► **Theorem 3 (strong normalization).** *The $\lambda\Phi$ -calculus is strongly normalizing.*

Proof. Let M be an arbitrary term. Let $R \equiv (\lambda x.P)Q$ be a redex in M and let M change to M' by contracting R . If R is an opponent redex, then by induction on M we show that $d(M') < d(M)$:

- If $M \equiv N_1 N_2$, and R is contained in N_i , then let N_i change to N'_i by contracting R . By the induction hypothesis, $d(N'_i) < d(N_i)$. Thus, $d(M') < d(M)$. If $M \equiv R$, then $M' \equiv P[Q/x]$. Then $d(M') = d(P) < 1 + d(P) = d(M)$, where the first equality follows from Lemma 2.
- If $M \equiv \lambda z.N$, then R must be contained in N . Let N change to N' by contracting R , so $M' \equiv \lambda z.N'$. By the induction hypothesis, we have $d(N') < d(N)$. Thus, $d(M') \leq 1 + d(N') < 1 + d(N) = d(M)$.

On the other hand, if R is a player redex, then $s(M') < s(M)$ and $d(M') \leq d(M)$. Therefore, in between each of the at most $d(M)$ opponent reductions, there can be at most a finite number of player reductions, and reduction always terminates. ◀

It is useful to have a rough estimate on the complexity of reduction. Note that since each opponent term can at most square the size of the term, the size of any reduct is bounded by $s(M)^{2^d}$, where $d = d(M)$. Therefore, in between each of the at most $d = d(M)$ opponent reductions, there can be at most $s(M)^{2^d}$ player reductions. This leads to normalization in *double exponential time*, $s(M)^{2^{O(d)}}$.

There is some flexibility in what to take as the set of **output values**, Θ (see [13]). At the very least, the set of Φ -normal forms, denoted Φ -NF, should be contained in Θ . For the purposes of this paper, it suffices to assume that $\Theta = \Phi$ -NF. In the next section we shall represent polynomial time algorithms by programs in the $\lambda\Phi$ -calculus.

3 The Parametric Lambda Calculus and Polynomial Time

To begin we must represent some basic data structures like booleans and boolean strings in the $\lambda\Phi$ -calculus, but we shall not require them to be encoded as player terms, so they are not necessarily input values³. Other data structures will be introduced as needed.

Booleans

Booleans are represented by the set $\{True, False\}$ and the conditional $Cond$, where:

$$True \equiv \lambda xy.x, \quad False \equiv \lambda xy.y, \quad Cond \equiv I$$

Note that $CondTrueMN \rightarrow_{\Phi}^* M$ and $CondFalseMN \rightarrow_{\Phi}^* N$ for any $M, N \in \Phi$, and that the terms $True, False, Cond$ belong to $\Phi \cap \Theta$.

Boolean Strings

Boolean strings $w = b_1b_2 \cdots b_n \in \{0, 1\}^*$, are represented using a Church-style encoding as follows:

$$W \equiv \lambda fx.f a_1(f a_2(\cdots f a_{n-1}(f a_n x) \cdots)), \quad a_i \equiv \begin{cases} zero \equiv \lambda xyz.y, & b_i = 0 \\ one \equiv \lambda xyz.z, & b_i = 1 \end{cases}$$

For example, the boolean string $1011 \in \{0, 1\}^*$ is encoded by the opponent term:

$$\lambda fx.f one(f zero(f one(f one x)))$$

Note that $d(W) \leq 1$ for all string encodings W ; this fact will be important in the proof of Theorem 5. Note that boolean strings are in Φ -normal form, but they are not necessarily input values. Thus, boolean strings are not in general duplicable in our setting. For this reason a slightly more general notion of representability is required:

► **Definition 4.** A predicate $A \subseteq \{0, 1\}^*$ is representable in the $\lambda\Phi$ -calculus if there exists a context $C[\cdot]$ such that, for all $w \in \{0, 1\}^*$, $C[W] \rightarrow_{\Phi}^* Bool$, where W is the encoding of w (as defined above) and $Bool \equiv True$ if $w \in A$ and $Bool \equiv False$ if $w \notin A$. In this case, the context $C[\cdot]$ is said to represent the predicate in the $\lambda\Phi$ -calculus.

We now prove one of the main results of this paper:

► **Theorem 5 (soundness).** *Let $C[\cdot]$ represent a predicate in the $\lambda\Phi$ -calculus. Then, for all $w \in \{0, 1\}^*$, the term $C[W]$, where W is the encoding of w , reduces to $True$ or $False$ in time polynomial in the size/length, denoted $|w|$, of w .*

Proof. We define a simple call-by-value operational semantics that proves judgements of the

³ An alternative system in which all of the required basic data structures are input values will be sketched in the conclusion.

form $M \Downarrow_{\mathbf{V}} N$, where M is any lambda term and N is a term in Φ -normal form:

$$\frac{(M_i \Downarrow_{\mathbf{V}} N_i)_{i \leq m}}{xM_1 \cdots M_m \Downarrow_{\mathbf{V}} xN_1 \cdots N_m} \text{ (var)}$$

$$\frac{M \Downarrow_{\mathbf{V}} N}{\lambda x.M \Downarrow_{\mathbf{V}} \lambda x.N} \text{ (abs)}$$

$$\frac{Q \Downarrow_{\mathbf{V}} Q' \quad Q' \in \Phi \quad P[Q'/x]M_1 \cdots M_m \Downarrow_{\mathbf{V}} N}{(\lambda x.P)QM_1 \cdots M_m \Downarrow_{\mathbf{V}} N} \text{ (head)}$$

$$\frac{Q \Downarrow_{\mathbf{V}} Q' \quad Q' \notin \Phi \quad P \Downarrow_{\mathbf{V}} P' \quad (M_i \Downarrow_{\mathbf{V}} N_i)_{i \leq m}}{(\lambda x.P)QM_1 \cdots M_m \Downarrow_{\mathbf{V}} (\lambda x.P')Q'N_1 \cdots N_m} \text{ (block)}$$

An easy induction on the evaluation tree shows that if $M \Downarrow_{\mathbf{V}} N$, then $M \rightarrow_{\Phi}^* N$, where N is in Φ -normal form, and the length of this reduction sequence is bounded by the size of the evaluation tree.

Our runtime bound is therefore obtained by bounding the size of the (rooted) evaluation tree – i.e. the total number of rules in the evaluation tree for the judgement $M \Downarrow_{\mathbf{V}} N$. Let n denote the total number of (*head*)-rules contained in such a tree. We show by induction on the height of the tree that the size of the tree is bounded by $(n+1)h$, where h denotes the maximum size of any reduct of M . To this end, we define a partial order on the vertices (proof rules) of the evaluation tree such that $u \leq v$ iff the unique path from the root to v passes through u . Consider the subtree obtained by removing the *right* branch of each of the $p \geq 0$ occurrences of minimal (with respect to the above partial order) (*head*)-rules. This subtree has size bounded by $s(M) \leq h$ as each rule in the subtree decreases the size of the term. By the induction hypothesis, each of the removed branches, which end with judgements of the form $P[Q'/x]M_1 \cdots M_m \Downarrow_{\mathbf{V}} N$, has size bounded by $(n_i+1)h$, where n_i ($1 \leq i \leq p$) denotes the number of (*head*)-rules in branch i . Therefore, the total size of the evaluation tree is bounded by $h + \sum_{i \leq p} (n_i+1)h = (1 + \sum_{i \leq p} n_i + p)h = (1+n)h$, as claimed.

Therefore, it follows by the discussion immediately following the proof of Theorem 3 that the size of the evaluation tree is bounded by $s(C[W])^{2^{O(d)}}$, where d is the degree of $C[W]$, as both n and h are so bounded. And since $s(C[W]) = O(|w|)$ and d remains fixed, the bound is in fact polynomial in $|w|$ (albeit with possibly large degree). Therefore, our reduction machine, as defined by the operational semantics above, runs in polynomial time. ◀

3.1 Polynomial Time Completeness

In this section we shall use the notation $M^n N \equiv N$, if $n = 0$, and $M^{n+1} N \equiv M(M^n N)$, if $n > 0$.

We begin by defining a context $Iter_{P,M}[\cdot]$ which is used to iterate a player term $M \in \Phi$ a polynomial $P(n)$ number of times. More precisely, we claim:

$$Iter_{P,M}[W] \rightarrow_{\Phi}^* \lambda y.M^{P(n)}y \in \Phi \quad (1)$$

where $n = |w|$. Recall that any polynomial with natural number coefficients can be represented in *Horner normal form*. For example, the polynomial $2n^3 + 4n^2 + 3n + 5$ is represented in Horner normal form as $((((2)n + 4)n + 3)n + 5)$. Given a polynomial $P(n)$, which is either a constant a_0 or has the form $P(n) = P_1(n)n + a_0$, where $P_1(n)$ is in Horner normal form, the context $Iter_{P,M}[\cdot]$ is defined inductively on the structure of P as follows:

$$Iter_{a_0,M}[\cdot] \equiv \lambda y.M^{a_0}y$$

$$Iter_{P,M}[\cdot] \equiv \lambda y.M^{a_0}([\cdot](K Iter_{P_1,M}[\cdot])y) \quad (K \equiv \lambda xy.x)$$

We argue by induction on the structure of the polynomial that this context satisfies reduction (1). Indeed, the base case is clear. Otherwise we suppose $P(n) = (P_1(n))n + a_0$. By the induction hypothesis we have that $Iter_{P_1, M}[W] \rightarrow_{\Phi}^* \lambda y. M^{P_1(n)} y \in \Phi$, where $n = |w|$. Then:

$$\begin{aligned}
Iter_{P, M}[W] &\equiv \lambda x. M^{a_0} (W(K Iter_{P_1, M}[W])x) \\
&\rightarrow_{\Phi}^* \lambda x. M^{a_0} (W(K \lambda y. M^{P_1(n)} y)x) \\
&\rightarrow_{\Phi}^* \lambda x. M^{a_0} ((\lambda y. M^{P_1(n)} y)^n x) \\
&\rightarrow_{\Phi}^* \lambda x. M^{a_0} (M^{(P_1(n))n} x) \\
&\equiv \lambda x. M^{(P_1(n))n + a_0} x \\
&\equiv \lambda x. M^{P(n)} x \in \Phi
\end{aligned}$$

This finishes the induction. Therefore, for any player term N , we have:

$$Iter_{P, M}[W]N \rightarrow_{\Phi}^* M^{P(n)}N \in \Phi$$

Note, in particular, that $\lambda f. Iter_{P, f}[W] \rightarrow_{\Phi}^* \lambda f x. f^{P(n)}x$, which is the Church representation of $P(n)$. We shall use this iteration combinator to iterate the transition function of a space-bounded Turing machine a polynomial number of times. But first we need the following preliminaries.

Affine Tensor Products

Given player terms $N_1, \dots, N_m \in \Phi$, we write $N_1 \otimes \dots \otimes N_m$ for the term $\lambda x. xN_1 \dots N_m$ and Prj_i for the term $\lambda f. f(\lambda x_1 \dots x_m. x_i)$ for each $1 \leq i \leq m$. Note that $N_1 \otimes \dots \otimes N_m, Prj_i \in \Phi$ and satisfy: $Prj_i N_1 \otimes \dots \otimes N_m \rightarrow_{\Phi}^* N_i$ for each $1 \leq i \leq m$. Occasionally we shall also use the notation $\lambda x_1 \otimes \dots \otimes x_m. M$ for the term $\lambda f. f(\lambda x_1 \dots x_m. M)$. This satisfies $(\lambda x_1 \otimes \dots \otimes x_m. M)(N_1 \otimes \dots \otimes N_m) \rightarrow_{\Phi}^* M[N_1/x_1] \dots [N_m/x_m]$ for any terms $M, N_1, \dots, N_m \in \Phi$. For example, Prj_i could be written instead as $\lambda x_1 \otimes \dots \otimes x_m. x_i$ for each $1 \leq i \leq m$.

Lists

Lists are encoded using player terms for the constructors *nil* and *cons* as follows:

$$nil \equiv \lambda xy. y \quad H :: T \equiv \lambda x. xHT \equiv H \otimes T$$

Note that $nil \in \Phi$ and $H :: T \in \Phi$ iff $H, T \in \Phi$. We shall assume that “ $::$ ” associates to the right.

Space-Bounded Turing Machines

Suppose we are given a Turing machine with k states, tape alphabet $\{\sqcup, 0, 1\}$, and input alphabet $\{0, 1\}$, where \sqcup is a special symbol for “blank”. These three symbols are encoded, respectively, by the terms $blank \equiv \lambda xyz. x$, $zero \equiv \lambda xyz. y$ and $one \equiv \lambda xyz. z$. For convenience, we shall assume that the set of states always contains (distinct) special *accepting* and *rejecting* states, and that the machine cannot change states once it reaches one of these two terminating states. Moreover, we assume that the tape is sufficiently large so that either the accepting or rejecting state is always reached before the machine encounters either end of the tape.

A configuration of the Turing machine is encoded as an affine triple tensor product, $S \otimes L \otimes R$, where S encodes the current state, L encodes the left part of the tape (in reverse

order), and R encodes the right part of the tape. The head of the Turing machine is always assumed to be positioned on the head of R . The left and right parts of the tape are encoded as lists of the tape alphabet. The k states S_i are encoded as follows. Suppose, for example, in state i , the machine's instructions, upon reading the symbol on the head, are:

$$\begin{aligned} \sqcup &\mapsto (s_{j_1}, 0, \text{move_right}) \\ 0 &\mapsto (s_{j_2}, 1, \text{move_left}) \\ 1 &\mapsto (s_{j_3}, \sqcup, \text{move_left}) \end{aligned}$$

Then state S_i is encoded as follows:

$$S_i \equiv \lambda x_1 \dots x_k, h_1 \otimes t_1, h_2 \otimes t_2. h_2(F_1)(F_2)(F_3)x_1 \dots x_k h_1 t_1 t_2$$

$$F_1 \equiv \lambda x_1 \dots x_k h_1 t_1 t_2. x_{j_1} \otimes (\text{zero} :: h_1 :: t_1) \otimes t_2$$

$$F_2 \equiv \lambda x_1 \dots x_k h_1 t_1 t_2. x_{j_2} \otimes t_1 \otimes (h_1 :: \text{one} :: t_2)$$

$$F_3 \equiv \lambda x_1 \dots x_k h_1 t_1 t_2. x_{j_3} \otimes t_1 \otimes (h_1 :: \text{blank} :: t_2)$$

The transition function is then defined as $T \equiv \lambda z \otimes l \otimes r. z S_1 \dots S_k l r$ and satisfies $T(S_i \otimes L \otimes R) \rightarrow_{\Phi}^* S_j \otimes L' \otimes R'$, where S_j is the new current state and L' and R' are encodings of the updated left and right parts of the tape after one iteration of the machine. The special accepting and rejecting states simply remain in the same state and write *one* or *zero*, respectively, on the head of the tape. Finally, observe that there is a term *out* (encoded via appropriate projections) that returns the head of the right tape from a given configuration of the machine. Note that all the terms in this encoding belong to Φ .

► **Theorem 6 (completeness).** *If a predicate is computable by a Turing machine in polynomial time $P(n)$ and polynomial space $Q(n)$ ⁴, then it is representable in the $\lambda\Phi$ -calculus by a context $C[\cdot]$.*

Proof. Observe that there is a player term, denoted *pad*, which satisfies $\text{pad}(S \otimes L \otimes R) \rightarrow_{\Phi}^* S \otimes (\text{blank} :: L) \otimes (\text{blank} :: R)$. Then $\text{Iter}_{Q, \text{pad}}[W](S_1 \otimes \text{nil} \otimes \text{nil})$ pads out the tape sufficiently for the full computation of the machine and puts it in the initial state S_1 . Next, apply a player context *write* such that $\text{write}[W](S \otimes L \otimes R) \rightarrow_{\Phi}^* S \otimes L \otimes (W(\lambda xy.(x :: y))R)$, which writes the binary string w onto the right part of the tape. Now apply $\text{Iter}_{P, T}[W]$, where T is the encoding of the transition function of the machine (as described above), to iterate the transition function $P(|w|)$ times, which suffices for the machine to reach a terminating state. Finally, use the player term *out* (mentioned above) to return the head of the right tape R and apply it to the arguments $I, \text{False}, \text{True}$ to get a boolean value which indicates whether the machine is in an accepting or a rejecting state. ◀

4 Various Extensions

In this section we investigate a few natural extensions of the $\lambda\Phi$ -calculus. The first extends the set of input values to include the so-called Φ -valuable terms (defined below). This larger set of input values allows for further flexibility in defining programs, but at the expense of strong normalization. The second extension adds explicit products to the language for the purpose of characterizing nondeterministic polynomial computations.

⁴ Of course, the explicit space bound $Q(n)$ is not necessary here since one may simply take $Q(n) = P(n)$. The explicit accounting of both time and space resources in the encoding of TMs is similar to that given in [9].

4.1 Φ -valuability

In order to reduce a redex $(\lambda x.P)Q$ in the $\lambda\Phi$ -calculus, it is first necessary to reduce Q to an input value. However, suppose we relax this condition as follows. Recall that a term M is called Φ -valuable if there is an $N \in \Phi$ such that $M \rightarrow_{\Phi}^* N$. Let Φ_v denote the set of Φ -valuable terms.

► **Theorem 7.** Φ_v forms a set of input values such that $\Phi \subset \Phi_v$. Therefore, the reduction relation $\rightarrow_{\Phi_v}^*$ is confluent by Theorem 1.2.5 in [13].

Proof. Condition 1 is immediate. For condition 2, let $P, Q \in \Phi_v$, so there are terms $P', Q' \in \Phi$ such that $P \rightarrow_{\Phi}^* P'$ and $Q \rightarrow_{\Phi}^* Q'$. Then $P[Q/x] \rightarrow_{\Phi}^* P'[Q'/x]$ (by Lemmas 1.2.21 and 1.2.22 in [13]). But then $P'[Q'/x]$ is in Φ by the substitution closure of Φ . Hence $P[Q/x] \in \Phi_v$. For condition 3, let $M \equiv C[(\lambda x.P)Q]$ and $N \equiv C[P[Q/x]]$. Then $Q \in \Phi_v$ in order for this reduction to happen. So there is a term $Q' \in \Phi$ such that $Q \rightarrow_{\Phi}^* Q'$. Then $M \rightarrow_{\Phi}^* C[(\lambda x.P)Q'] \rightarrow_{\Phi} C[P[Q'/x]]$. But since $M \in \Phi_v$, there is a term $M' \in \Phi$ such that $M \rightarrow_{\Phi}^* M'$. So by the confluence of \rightarrow_{Φ}^* there is a term $M'' \in \Phi$ such that $C[P[Q'/x]] \rightarrow_{\Phi}^* M''$. But then $N \rightarrow_{\Phi}^* C[P[Q'/x]] \rightarrow_{\Phi}^* M''$, which shows that $N \in \Phi_v$. ◀

The $\lambda\Phi_v$ -calculus is *not* strongly normalizing. For example, consider the term $D \equiv \lambda y.KI(yy)I$, which is Φ -valuable since it reduces to $\lambda yx.x \in \Phi$. However, $DD \rightarrow_{\Phi_v} KI(DD)I \rightarrow_{\Phi_v} KI(KI(DD)I)I \rightarrow_{\Phi_v} \dots$, which leads to an infinite reduction sequence. Nevertheless, every term in the $\lambda\Phi_v$ -calculus has a (unique) normal form:

► **Theorem 8 (weak normalization).** Every term in the $\lambda\Phi_v$ -calculus has a unique normal form.

Proof. Let M be a term. If there are no Φ_v -redexes, then M is a Φ_v -nf and we are done. Otherwise, let $(\lambda x.P)Q$ be a Φ_v -redex in M . Then $Q \in \Phi_v$, and thus $Q \rightarrow_{\Phi}^* Q'$, where $Q' \in \Phi$. Consider the Φ -reduction of M : $M \equiv C[(\lambda x.P)Q] \rightarrow_{\Phi}^* C[(\lambda x.P)Q'] \rightarrow_{\Phi} C[P[Q'/x]]$. If $C[P[Q'/x]]$ is a Φ_v -nf, then we are done. Otherwise, repeat this procedure and extend the Φ -reduction of M . This process must terminate since there are no infinite Φ -reductions. Therefore, M is weakly normalizable. Uniqueness comes from the confluence property (Theorem 7). ◀

Of course, by including the Φ -valuable terms in the set of input values, we have not obtained any new complexity results. However, Theorem 7 is an interesting general result about the parametric lambda calculus that is true of any set of input values and may have future applications. Theorem 8 is less applicable as it requires the additional fact that the $\lambda\Phi$ -calculus is strongly normalizing.

4.2 Explicit Products and Nondeterminism

In this section we study nondeterminism in the parametric lambda calculus with the use of explicit products. Of course, one could simply add a new term constructor $M + N$ (*sum*) together with nondeterministic projections, but the resulting system would not be confluent (by construction). Here we present an alternative approach based on the idea of a polynomial verifier.

Let Λ_{\times} denote the set of lambda calculus with explicit products:

$$M, N ::= x \mid MN \mid \lambda x.M \mid \langle M, N \rangle \mid \pi_1 M \mid \pi_2 M$$

where $x \in \text{Var}$. The notions of substitution and α -equivalence are extended in the obvious manner. Here, the term $\langle M, N \rangle$ is called an **explicit product** and $\pi_i M$ are called **projections**. For a given set of input values $\Delta \subseteq \Lambda_\times$, we extend \rightarrow_Δ as the contextual closure of the following rules:

$$\begin{aligned} (\lambda x.M)N &\rightarrow_\Delta M[N/x] \quad \text{iff } N \in \Delta \\ \pi_1 \langle M, N \rangle &\rightarrow_\Delta M \\ \pi_2 \langle M, N \rangle &\rightarrow_\Delta N \end{aligned}$$

As before, we let \rightarrow_Δ^* denote the reflexive and transitive closure of \rightarrow_Δ . The following is a straightforward generalization of Theorem 1.2.5 in [13].

► **Theorem 9.** *The $\lambda\Delta$ -calculus with explicit products and \rightarrow_Δ^* as defined above is confluent.*

Proof. The definitions of deterministic parallel reduction \hookrightarrow_Δ and nondeterministic parallel reduction \Rightarrow_Δ , Definition 1.2.19 in [13], are extended with the following clauses:

5. $M \hookrightarrow_\Delta M', N \hookrightarrow_\Delta N'$ imply $\langle M, N \rangle \hookrightarrow_\Delta \langle M', N' \rangle$;
6. $M \hookrightarrow_\Delta M'$ and $M \not\equiv \langle M_1, M_2 \rangle$ imply $\pi_i M \hookrightarrow_\Delta \pi_i M'$, for $i \in \{1, 2\}$;
7. $M_1 \hookrightarrow_\Delta M'_1, M_2 \hookrightarrow_\Delta M'_2$ imply $\pi_i \langle M_1, M_2 \rangle \hookrightarrow_\Delta M'_i$, for $i \in \{1, 2\}$.

5. $M \Rightarrow_\Delta M', N \Rightarrow_\Delta N'$ imply $\langle M, N \rangle \Rightarrow_\Delta \langle M', N' \rangle$;
6. $M \Rightarrow_\Delta M'$ implies $\pi_i M \Rightarrow_\Delta \pi_i M'$, for $i \in \{1, 2\}$;
7. $M_1 \Rightarrow_\Delta M'_1, M_2 \Rightarrow_\Delta M'_2$ imply $\pi_i \langle M_1, M_2 \rangle \Rightarrow_\Delta M'_i$, for $i \in \{1, 2\}$.

Then Lemmas 1.2.21, 1.2.22, Property 1.2.23 and Lemma 1.2.24 in [13] all have straightforward generalizations by checking the extra cases. The details are left to the reader. Finally, the proof of Lemma 1.2.25 and the rest of the proof of Theorem 1.2.5 are unchanged. ◀

We expand the set of input/player values to include affine linear terms with explicit products: Let Φ_\times be defined the smallest subset of Λ_\times satisfying the following closure properties:

- $x \in \text{Var}$ implies $x \in \Phi_\times$,
- If $M, N \in \Phi_\times$, then $(MN) \in \Phi_\times$,
- If $M \in \Phi_\times$, then $\pi_i M \in \Phi_\times$ for $i \in \{1, 2\}$,
- If $M, N \in \Phi_\times$, then $\langle M, N \rangle \in \Phi_\times$,
- If $M \in \Phi_\times$ and $x \in \text{Var}$ and x occurs free at most once in M , then $(\lambda x.M) \in \Phi_\times$.

Note that if $M \in \Phi_\times$ and is closed, then M is affine linear even with (additive) explicit products. One could define a more general notion of input values based on the notion of *slice* (as defined in [11], for example) which includes terms like $\lambda x.\langle x, x \rangle$ and is strongly normalizing. However, such a system would require a lazy reduction strategy for explicit products as well as pointers to avoid an exponential explosion in the size of a term (cf. [12]). We don't believe this added complication is necessary here.⁵

The following two theorems are straightforward generalizations of Theorems 1 and 3, so the proofs have been omitted.

► **Theorem 10.** Φ_\times forms a set of input values such that $\Phi \subset \Phi_\times$. Therefore, the reduction relation $\rightarrow_{\Phi_\times}^*$ is confluent by Theorem 9.

► **Theorem 11.** The $\lambda\Phi_\times$ -calculus is strongly normalizing.

⁵ However, we do believe this more general set of input values is the starting point for a characterization of PSPACE (see [12]).

We extend the simple call-by-value operational semantics introduced in the proof of Theorem 5 with the following rules:

$$\frac{P_1 \Downarrow_{\mathbf{V}} P'_1 \quad P_2 \Downarrow_{\mathbf{V}} P'_2 \quad (M_i \Downarrow_{\mathbf{V}} N_i)_{i \leq m}}{\langle P_1, P_2 \rangle M_1 \cdots M_m \Downarrow_{\mathbf{V}} \langle P'_1, P'_2 \rangle N_1 \cdots N_m} \text{ (pair)}$$

$$\frac{Q \Downarrow_{\mathbf{V}} Q' \quad Q' \equiv \langle Q'_1, Q'_2 \rangle \quad Q'_i M_1 \cdots M_m \Downarrow_{\mathbf{V}} N}{\pi_i Q M_1 \cdots M_m \Downarrow_{\mathbf{V}} N} \text{ (proj}_i\text{)}$$

$$\frac{Q \Downarrow_{\mathbf{V}} Q' \quad Q' \not\equiv \langle Q'_1, Q'_2 \rangle \quad (M_i \Downarrow_{\mathbf{V}} N_i)_{i \leq m}}{\pi_i Q M_1 \cdots M_m \Downarrow_{\mathbf{V}} \pi_i Q' N_1 \cdots N_m} \text{ (block}_i\text{)}$$

Once again, an easy induction on the evaluation tree shows that if $M \Downarrow_{\mathbf{V}} N$, then $M \rightarrow_{\Phi_{\times}}^* N$, where N is in Φ_{\times} -normal form, and the length of this reduction sequence is bounded by the size of the evaluation tree. On the other hand, by Theorem 10, if $M \rightarrow_{\Phi_{\times}}^* N$, where N is in Φ_{\times} -normal form, then M evaluates to N according to the operational semantics defined above.

Additive Booleans

Explicit products allow for an alternative definition of booleans and conditional:

$$Proj_1 \equiv \lambda f. \pi_1 f \quad Proj_2 \equiv \lambda f. \pi_2 f \quad \text{if } b \text{ then } M \text{ else } N \equiv \lambda b. b \langle M, N \rangle$$

A term M is called **eventually true** if there exists a sequence of additive booleans $Proj_{i_1}, \dots, Proj_{i_k}$ such that $M Proj_{i_1} \cdots Proj_{i_k} \rightarrow_{\Phi_{\times}}^* True$.

► **Definition 12.** A predicate $A \subseteq \{0, 1\}^*$ is representable in the $\lambda\Phi_{\times}$ -calculus if there is a context $C[\cdot]$ such that, for all $w \in \{0, 1\}^*$, $w \in A$ iff $C[W]$ is eventually true.

We have the following result:

► **Theorem 13.** A predicate $A \subseteq \{0, 1\}^*$ is representable in the $\lambda\Phi_{\times}$ -calculus by a context $C[\cdot]$ iff it is computable in nondeterministic polynomial time (NP).

Proof. (\Rightarrow) Suppose a predicate A is representable by a context $C[\cdot]$ in the $\lambda\Phi_{\times}$ -calculus. If $w \in A$, then a straightforward generalization of Theorem 5 shows that, for any choice of projections $Proj_{i_1}, Proj_{i_2}, \dots, Proj_{i_k}$, the term $C[W] Proj_{i_1} Proj_{i_2} \cdots Proj_{i_k}$ reduces in time bounded by a polynomial in $s(C[W] Proj_{i_1} Proj_{i_2} \cdots Proj_{i_k})$ to $True$.⁶ Moreover, note that k must be bounded by a polynomial in $|w|$. Indeed, each projection input requires a head lambda abstraction. This head lambda abstraction cannot itself be a projection term because otherwise the normal form would have a (projection) block. And there can only be a polynomial in $|w|$ such head reductions. Therefore, the entire reduction is polynomial time in $|w|$ only.

(\Leftarrow) Conversely, let A be a predicate computable on a nondeterministic Turing machine in polynomial time $P(n)$ and polynomial space $Q(n)$ (i.e. the maximum time and space used by any computational branch). The encoding of nondeterministic Turing machines is based on the encoding of deterministic Turing machines found in Section 3.1. However, for a nondeterministic machine, we assume a pair of transition functions T_l and T_r instead of just

⁶ As noted above, we may assume, by Theorem 10, that this reduction sequence is determined by the operational semantics.

one. The following player term can be iterated $P(n)$ times using the iteration combinator from Section 3.1:

$$\text{Branch} \equiv \lambda f z b.b\langle \lambda xy.x(T_l y), \lambda xy.x(T_r y) \rangle f z$$

Let w be any binary word and let $n = |w|$. Let $\text{Initial}[W] \rightarrow_{\Phi_x}^* \text{Config}$ initialize the machine by padding the tape out to size $Q(n)$, writing w on the initial segment of the tape, and putting it in the start state. Finally, let out be a term that reduces to True if a given configuration is accepting and reduces to False otherwise.

Let Z_k denote the normal form of $\text{Branch}^k \text{out}$, which has the form:

$$\begin{aligned} Z_k &\equiv \lambda z b.b\langle \lambda xy.x(T_l y), \lambda xy.x(T_r y) \rangle Z_{k-1} z, & k > 0 \\ Z_0 &\equiv \text{out} \end{aligned}$$

If $w \in A$, then there exists a sequence of i_1, \dots, i_k , with $i_j \in \{1, 2\}$ and $k = P(n)$, specifying a path down the nondeterministic evaluation tree to an accepting leaf. This path is encoded by the series of projections $\text{Proj}_{i_1}, \dots, \text{Proj}_{i_k}$ and verified as follows:

$$\begin{aligned} & \text{Iter}_{P, \text{Branch}}[W] \text{out Config Proj}_{i_1} \cdots \text{Proj}_{i_k} \\ \rightarrow_{\Phi_x}^* & \text{Branch}^k \text{out Config Proj}_{i_1} \cdots \text{Proj}_{i_k} \\ \rightarrow_{\Phi_x}^* & (\lambda b.b\langle \lambda xy.x(T_l y), \lambda xy.x(T_r y) \rangle Z_{k-1} \text{Config}) \text{Proj}_{i_1} \cdots \text{Proj}_{i_k} \\ \rightarrow_{\Phi_x}^* & \text{Proj}_{i_1} \langle \lambda xy.x(T_l y), \lambda xy.x(T_r y) \rangle Z_{k-1} \text{Config Proj}_{i_2} \cdots \text{Proj}_{i_k} \\ \rightarrow_{\Phi_x}^* & Z_{k-1} \text{Config}_1 \text{Proj}_{i_2} \cdots \text{Proj}_{i_k} \\ \rightarrow_{\Phi_x}^* & \cdots \\ \rightarrow_{\Phi_x}^* & Z_1 \text{Config}_{k-1} \text{Proj}_{i_k} \\ \rightarrow_{\Phi_x}^* & \text{out Config}_k \\ \rightarrow_{\Phi_x}^* & \text{True} \end{aligned}$$

Moreover, this reduction proceeds according to the operational semantics defined above. Thus, $(\text{Iter}_{P, \text{Branch}}[W] \text{out Config}) \text{Proj}_{i_1} \cdots \text{Proj}_{i_k} \Downarrow_{\mathbf{V}} \text{True}$. On the other hand, if $w \notin A$, then all such reductions reduce to False . ◀

5 Conclusion

In this paper we have demonstrated that a characterization of polynomial time computations can be obtained in the lambda calculus without requiring any typing information and/or ad hoc extensions to the language. Indeed, the characterization is obtained simply by restricting the set of input values to the so-called pseudo-affine terms. Moreover, a characterization of nondeterministic polynomial time is obtained with the addition of explicit products.

It would be interesting to investigate other (decidable) instantiations of the parametric lambda calculus in the context of Implicit Computational Complexity. For example, the choice to allow weakening in the language was made simply because it made the encoding of polynomial time TMs much more natural. However, if we change the final clause in the definition of Φ to specify that x occurs *exactly* once in P , we conjecture that this smaller set of input values also characterizes polynomial time. In this case, the encoding of polynomial bounded TMs might follow that in [11].

Finally, one unfortunate aspect of our characterization is that binary words (and Church numerals) are not in general input values. For this reason our definitions of representability use contexts instead, which is not standard. We consider two possibilities for dealing with

this situation. First, one could use a player encoding of binary strings instead (in the style of the Barendregt numerals [3, 13]), together with a more natural definition of representability. However, it is not difficult to show that this leads to a system for linear time computations only. A second possibility is to use a system of *abstract binary numerals* (see [8]) instead of pure terms to represent binary strings. In this case, four new atomic constants, denoted $\epsilon, \sigma_1, \sigma_2$ and Z , are added to the system such that ϵ, σ_1 and σ_2 belong to the set of input values, but Z does not. Then any binary string $w \in \{0, 1\}^*$ can be represented by an input value \widehat{w} in the obvious way. Furthermore, we add the contextual closure of the following reduction rule:

$$Z\widehat{w} \rightarrow_{\Phi} \bar{w}$$

where \bar{w} is the Church-style representation of w described in Section 3.⁷ This *arithmetical extension* [8] of the $\lambda\Phi$ -calculus leads to an *applied* system for polynomial time computations, similar to the pure system presented in this paper. It is an alternative if a standard notion of representability is desired.

Acknowledgements.

The author would like to thank Simona Ronchi Della Rocca for reading a preliminary draft of this paper and offering some useful suggestions for improvement. The author would also like to thank the anonymous referees for their valuable comments and suggestions regarding the presentation of the results of this paper.

References

- 1 Beniamino Accattoli and Ugo Dal Lago. Beta reduction is invariant, indeed. In Thomas A. Henzinger and Dale Miller, editors, *Joint Meeting of the Twenty-Third EACSL Annual Conference on Computer Science Logic (CSL) and the Twenty-Ninth Annual ACM/IEEE Symposium on Logic in Computer Science (LICS), CSL-LICS'14, Vienna, Austria, July 14–18, 2014*, page 8. ACM, 2014.
- 2 Patrick Baillot and Kazushige Terui. Light types for polynomial time computation in lambda-calculus. In *19th IEEE Symposium on Logic in Computer Science (LICS 2004), 14–17 July 2004, Turku, Finland, Proceedings*, pages 266–275. IEEE Computer Society, 2004.
- 3 Hendrik Pieter Barendregt. *The lambda calculus: its syntax and semantics*. Studies in logic and the foundations of mathematics. North-Holland, Amsterdam, New-York, Oxford, 1981.
- 4 J. Chrzęszcz and A. Schubert. The role of polymorphism in the characterisation of complexity by soft types. (To appear).
- 5 Marco Gaboardi and Simona Ronchi Della Rocca. A soft type assignment system for lambda-calculus. In Jacques Duparc and Thomas A. Henzinger, editors, *Computer Science Logic, 21st International Workshop, CSL 2007, 16th Annual Conference of the EACSL, Lausanne, Switzerland, September 11–15, 2007, Proceedings*, volume 4646 of *Lecture Notes in Computer Science*, pages 253–267. Springer, 2007.
- 6 Jean-Yves Girard. *Interprétation fonctionnelle et élimination des coupures de l'arithmétique d'ordre supérieur*. Thèse d'état, Université Paris 7, June 1972.

⁷ Note that if Z is permitted to be an input value, then \bar{w} must be as well. Otherwise, the system would not satisfy reduction closure (i.e. condition 3 in the definition of input values). This leads to (at least) elementary time complexity.

- 7 Jean-Yves Girard, Yves Lafont, and Paul Taylor. *Proofs and Types*. Number 7 in Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 1989.
- 8 J. Roger Hindley and Jonathan P. Seldin. *Lambda-Calculus and Combinators: An Introduction*. Cambridge University Press, New York, NY, USA, 2 edition, 2008.
- 9 Yves Lafont. Soft linear logic and polynomial time. *Theor. Comput. Sci.*, 318(1-2):163–180, 2004.
- 10 Daniel Leivant. Finitely stratified polymorphism. *Inf. Comput.*, 93(1):93–113, 1991.
- 11 Harry G. Mairson and Kazushige Terui. On the computational complexity of cut-elimination in linear logic. In Carlo Blundo and Cosimo Laneve, editors, *Theoretical Computer Science, 8th Italian Conference, ICTCS 2003, Bertinoro, Italy, October 13-15, 2003, Proceedings*, volume 2841 of *Lecture Notes in Computer Science*, pages 23–36. Springer, 2003.
- 12 B. Redmond. Bounded combinatory logic and lower complexity. (To appear).
- 13 Simona Ronchi Della Rocca and Luca Paolini. *The parametric lambda calculus: a meta-model for computation*. Texts in theoretical computer science. Springer-Verlag, New York, 2004.
- 14 Aleksy Schubert. The complexity of beta-reduction in low orders. In *TLCA*, pages 400–414, 2001.
- 15 H. Schwichtenberg. Definierbare Funktionen im λ -Kalkül mit Typen. *Archiv für mathematische Logik und Grundlagenforschung*, 17:113–114, 1976.
- 16 R. Statman. The typed lambda calculus is not elementary recursive. *Theoretical Computer Science*, 9:73–82, 1979.

Fibrations of Tree Automata

Colin Riba

ENS de Lyon, Université de Lyon, LIP*, France
colin.riba@ens-lyon.fr

Abstract

We propose a notion of morphisms between tree automata based on game semantics. Morphisms are winning strategies on a synchronous restriction of the linear implication between acceptance games. This leads to split indexed categories, with substitution based on a suitable notion of synchronous tree function. By restricting to tree functions issued from maps on alphabets, this gives a fibration of tree automata. We then discuss the (fibrewise) monoidal structure issued from the synchronous product of automata. We also discuss how a variant of the usual projection operation on automata leads to an existential quantification in the fibered sense. Our notion of morphism is correct (it respects language inclusion), and in a weaker sense also complete.

1998 ACM Subject Classification F.3.2 Semantics of Programming Languages, F.4.1 Mathematical Logic, F.4.2 Formal Languages

Keywords and phrases Tree automata, Game semantics, Categorical logic

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.302

1 Introduction

This paper proposes a notion of morphism between tree automata based on game semantics. We follow the Curry-Howard-like slogan: *Automata as objects, Executions as morphisms*.

We consider general alternating automata on infinite ranked trees. These automata encompass Monadic Second-Order Logic (MSO) and thus most of the logics used in verification [6]. Tree automata are traditionally viewed as positive objects: one is primarily interested in satisfaction or satisfiability, and the primitive notion of quantification is existential. In contrast, Curry-Howard approaches tend to favor proof-theoretic oriented and negative approaches, *i.e.* approaches in which the predominant logical connective is the implication, and where the predominant form of quantification is universal.

We consider full infinite ranked trees, built from a non-empty finite set of directions D and labeled in non-empty finite alphabets Σ . The base category **Tree** has alphabets as objects and morphisms from Σ to Γ are $(\Sigma \rightarrow \Gamma)$ -labeled D -ary trees.

The fibre categories are based on a generalization of the usual acceptance games, where for an automaton \mathcal{A} on alphabet Γ (denoted $\Gamma \vdash \mathcal{A}$), input characters can be precomposed with a tree morphism $M \in \mathbf{Tree}[\Sigma, \Gamma]$, leading to substituted acceptance games of type $\Sigma \vdash \mathcal{G}(\mathcal{A}, M)$. Usual acceptance games, which correspond to the evaluation of $\Sigma \vdash \mathcal{A}$ on a Σ -labeled input tree, are substituted acceptance games $\mathbf{1} \vdash \mathcal{G}(\mathcal{A}, t)$ with $t \in \mathbf{Tree}[\mathbf{1}, \Sigma]$. Games of the form $\Sigma \vdash \mathcal{G}(\mathcal{A}, M)$ are the objects of the fibre category over Σ .

For morphisms, we introduce a notion of “synchronous” simple game between acceptance games. We rely on Hyland & Schalk’s functor (denoted HS) from simple games to **Rel** [9]. A synchronous strategy $\Sigma \vdash \sigma : \mathcal{G}(\mathcal{A}, M) \multimap \mathcal{G}(\mathcal{B}, N)$ is a strategy in the simple game $\mathcal{G}(\mathcal{A}, M) \multimap \mathcal{G}(\mathcal{B}, N)$ required to satisfy (in **Set**) a diagram of the form of (1) below,

* LIP is associated with UMR CNRS – ENS Lyon – UCB Lyon 1 – INRIA 5668.



© Colin Riba;

licensed under Creative Commons License CC-BY

13th International Conference on Typed Lambda Calculi and Applications (TLCA’15).

Editor: Thorsten Altenkirch; pp. 302–316



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

expressing that \mathcal{A} and \mathcal{B} are evaluated along the same path of the tree and read the same input characters:

$$\begin{array}{ccc} \text{HS}(\sigma) & \longrightarrow & \mathcal{G}(\mathcal{B}, N) \\ \downarrow & & \downarrow \\ \mathcal{G}(\mathcal{A}, M) & \longrightarrow & (D + \Sigma)^* \end{array} \quad (1)$$

This gives a split fibration **game** of tree automata and acceptance games. When restricting the base to *alphabet* morphisms (*i.e.* functions $\Sigma \rightarrow \Gamma$), substitution can be internalized in automata. By change-of-base of fibrations, this leads to a split fibration **aut**. In the fibers of **aut**, the substituted acceptance games have finite-state winning strategies, whose existence can be checked by trivial adaptation of usual algorithms.

Each of these fibrations is monoidal in the sense of [16], by using a natural synchronous product of tree automata. We also investigate a linear negation, as well as existential quantifications, obtained by adapting the usual projection operation on non-deterministic automata to make it a left-adjoint to weakening, the adjunction satisfying the usual Beck-Chevalley condition.

Our linear implication of acceptance games seems to provide a natural notion of prenex universal quantification on automata not investigated before. As expected, if there is a synchronous winning strategy $\sigma \Vdash \mathcal{A} \multimap \mathcal{B}$, then $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(\mathcal{B})$ (*i.e.* each input tree accepted by \mathcal{A} is also accepted by \mathcal{B}). Under some assumptions on \mathcal{A} and \mathcal{B} the converse holds: $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(\mathcal{B})$ implies $\sigma \Vdash \mathcal{A} \multimap \mathcal{B}$ for some σ .

At the categorical level, thanks to (1), the constructions mimic relations in slices categories $\mathbf{Set}/(D + \Sigma)^*$ of the co-domain fibration: substitution is given by a (well chosen) pullback, and the monoidal product of automata is issued from the Cartesian product of plays in $\mathbf{Set}/(D + \Sigma)^*$ (*i.e.* also by a well chosen pullback).

The paper is organized as follows. Section 2 presents notations for trees and tree automata. Our notions of substituted acceptance games and synchronous arrow games are then discussed in Sect. 3. Substitution functors and the corresponding fibrations are presented in Sect. 4, and Section 5 overviews the monoidal structure. We then state our main correctness results in Sect. 6. Section 7 presents existential quantifications and quickly discusses non-deterministic automata.

2 Preliminaries

Fix a singleton set $\mathbf{1} = \{\bullet\}$ and a finite non-empty set D of (tree) *directions*.

Alphabets and Trees. We write Σ, Γ, \dots for *alphabets*, *i.e.* finite non-empty sets. We let **Alph** be the category whose objects are alphabets and whose morphisms $\beta \in \mathbf{Alph}[\Sigma, \Gamma]$ are functions $\beta : \Sigma \rightarrow \Gamma$.

We let **Tree** $[\Sigma]$ be the set of Σ -labeled full D -ary trees, *i.e.* the set of maps $T : D^* \rightarrow \Sigma$. Let **Tree** be the category with *alphabets* as objects and with morphisms $\mathbf{Tree}[\Sigma, \Gamma] := \mathbf{Tree}[(\Sigma \rightarrow \Gamma)]$, *i.e.* $(\Sigma \rightarrow \Gamma)$ -labeled trees. Maps $M \in \mathbf{Tree}[\Sigma, \Gamma]$ and $L \in \mathbf{Tree}[\Gamma, \Delta]$ are composed as $L \circ M : p \in D^* \mapsto (a \in \Sigma \mapsto L(p)(M(p)(a)))$ and the identity $\text{Id}_\Sigma \in \mathbf{Tree}[\Sigma, \Sigma]$ is defined as $\text{Id}_\Sigma(p)(a) := a$. Note that $\mathbf{Tree}[\mathbf{1}, \Sigma]$ is in bijection with $\mathbf{Tree}[\Sigma]$.

There is a faithful functor from **Alph** to **Tree**, mapping $\beta \in \mathbf{Alph}[\Sigma, \Gamma]$ to the constant tree morphism $(_ \mapsto \beta) \in \mathbf{Tree}[\Sigma, \Gamma]$ that we simply write β .

Tree Automata. Alternating tree automata [14] are finite state automata running on full infinite Σ -labeled D -ary trees. Their distinctive feature is that transitions are given by *positive Boolean* formulas with atoms pairs (q, d) of a state q and a tree direction $d \in D$ ($((q, d)$ means that one copy of the automaton should start in state q from the d -th son of the current tree position).

Acceptance for alternating tree automata can be defined either *via* run trees or *via* the existence of winning strategies in *acceptance games* [14]. In both cases, we can w.l.o.g. restrict to transitions given by formulas in (irredundant) disjunctive normal form [15]. In our setting, it is quite convenient to follow the presentation of [18], in which disjunctive normal forms with atoms in $Q \times D$ are represented as elements of $\mathcal{P}(\mathcal{P}(Q \times D))$.

An alternating tree automaton \mathcal{A} on alphabet Σ has the form (Q, q^i, δ, Ω) where Q is the finite set of states, $q^i \in Q$ is the *initial state*, the *acceptance condition* is $\Omega \subseteq Q^\omega$ and following [18], the *transition function* δ has the form

$$\delta : Q \times \Sigma \longrightarrow \mathcal{P}(\mathcal{P}(Q \times D))$$

We write $\Sigma \vdash \mathcal{A}$ if \mathcal{A} is a tree automaton on Σ . Usual acceptance games are described in Sect. 3.1. It is customary to put restrictions on the acceptance condition $\Omega \subseteq Q^\omega$, typically by assuming it is generated from a *Muller family* $\mathcal{F} \in \mathcal{P}(\mathcal{P}(Q))$ as the set of $\pi \in Q^\omega$ such that $\text{Inf}(\pi) \in \mathcal{F}$. We call such automata *regular*¹. They have decidable emptiness checking and the same expressive power as MSO on D -ary trees (see e.g. the survey [17]).

3 Categories of Acceptance Games and Automata

We present in this Section the categories $\mathbf{SAG}_\Sigma^{(W)}$ of *substituted acceptance games*. Their objects will be *substituted acceptance games* (to be presented in Sect. 3.1) and their morphisms will be strategies in corresponding *synchronous arrow games* (to be presented in Sect. 3.2). Substituted acceptance games and synchronous arrow games are the two main notions we introduce in this paper. Our categories of $\mathbf{Aut}_\Sigma^{(W)}$ of automata will be full subcategories of $\mathbf{SAG}_\Sigma^{(W)}$, while $\mathbf{SAG}_\Sigma^{(W)}$ and $\mathbf{Aut}_\Sigma^{(W)}$ will be the total categories of our fibrations

$$\text{game}^{(W)} : \mathbf{SAG}^{(W)} \longrightarrow \mathbf{Tree} \qquad \text{aut}^{(W)} : \mathbf{Aut}^{(W)} \longrightarrow \mathbf{Alph}$$

to be presented in Sect. 4.

3.1 Substituted Acceptance Games

Consider a tree automaton $\mathcal{A} = (Q, q^i, \delta, \Omega)$ on Γ and a morphism $M \in \mathbf{Tree}[\Sigma, \Gamma]$. The *substituted acceptance game* $\Sigma \vdash \mathcal{G}(\mathcal{A}, M)$ is the positive game

$$\mathcal{G}(\mathcal{A}, \vec{M}) := (D^* \times (A_P + A_O), E, *, \lambda, \xi, \mathcal{W})$$

whose positions are given by $A_P := Q$ and $A_O := \Sigma \times \mathcal{P}(Q \times D)$, whose polarized root is $* := (\varepsilon, q^i)$ with $\xi(*) = P$, whose polarized moves (E, λ) are given by

$$\begin{aligned} \text{from } (D^* \times A_P) \text{ to } (D^* \times A_O) : & \quad (p, q) \xrightarrow{P} (p, a, \gamma) \quad \text{iff } \gamma \in \delta(q, \vec{M}(p)(a)) \\ \text{from } (D^* \times A_O) \text{ to } (D^* \times A_P) : & \quad (p, a, \gamma) \xrightarrow{O} (p.d, q) \quad \text{iff } (q, d) \in \gamma \end{aligned}$$

¹ By adding states to \mathcal{A} if necessary, one can describe Ω by an equivalent *parity* condition.

and whose winning condition is given by

$$(\varepsilon, q_0) \cdot (\varepsilon, a_0, \gamma_0) \cdot (p_1, q_1) \cdot \dots \cdot (p_n, q_n) \cdot (p_n, a_n, \gamma_n) \cdot \dots \in \mathcal{W} \quad \text{iff} \quad (q_i)_{i \in \mathbb{N}} \in \Omega$$

The input alphabet of $\Gamma \vdash \mathcal{A}$ is Γ , and we use the tree morphism $M \in \mathbf{Tree}[\Sigma, \Gamma]$ in a contravariant way to obtain a game with “input alphabet” Σ , that we emphasize by writing $\Sigma \vdash \mathcal{G}(\mathcal{A}, \vec{M})$. Input characters $a \in \Sigma$ are chosen by P, directions $d \in D$ are chosen by O.

Write $\Sigma \vdash \sigma \Vdash \mathcal{G}(\mathcal{A}, \vec{M})$ if σ is a winning P-strategy on $\Sigma \vdash \mathcal{G}(\mathcal{A}, \vec{M})$, and $\Sigma \Vdash \mathcal{G}(\mathcal{A}, M)$ if $\Sigma \vdash \sigma \Vdash \mathcal{G}(\mathcal{A}, M)$ for some σ .

Correspondence with usual Acceptance Games. Usual acceptance games model the evaluation of automata $\Sigma \vdash \mathcal{A}$ on input trees $t \in \mathbf{Tree}[\Sigma]$. They correspond to games of the form $\mathbf{1} \vdash \mathcal{G}(\mathcal{A}, \dot{t})$, where $\dot{t} \in \mathbf{Tree}[\mathbf{1}, \Sigma]$ is the tree morphism corresponding to $t \in \mathbf{Tree}[\Sigma]$.

Note that in these cases, A_O is of the form $\mathbf{1} \times \mathcal{P}(Q \times D) \simeq \mathcal{P}(Q \times D)$, so that the games $\mathbf{1} \vdash \mathcal{G}(\mathcal{A}, \dot{t})$ are isomorphic to the acceptance games of [18].

► **Definition 3.1.** Let $\Sigma \vdash \mathcal{A}$.

- (i) \mathcal{A} *accepts* the tree $t \in \mathbf{Tree}[\Sigma]$ if there is a strategy σ such that $\mathbf{1} \vdash \sigma \Vdash \mathcal{G}(\mathcal{A}, \dot{t})$.
- (ii) Let $\mathcal{L}(\mathcal{A}) \subseteq \mathbf{Tree}[\Sigma]$, the *language* of \mathcal{A} , be the set of trees accepted by \mathcal{A} .

3.2 Synchronous Arrow Games

Consider games $\Sigma \vdash \mathcal{G}(\mathcal{A}, M)$ and $\Sigma \vdash \mathcal{G}(\mathcal{B}, N)$ with $\mathcal{A} = (Q_{\mathcal{A}}, q_{\mathcal{A}}^i, \delta_{\mathcal{A}}, \Omega_{\mathcal{A}})$ and $\mathcal{B} = (Q_{\mathcal{B}}, q_{\mathcal{B}}^i, \delta_{\mathcal{B}}, \Omega_{\mathcal{B}})$. Similarly as in Sect. 3.1 above, write

$$A_P := Q_{\mathcal{A}} \quad A_O := \Sigma \times \mathcal{P}(Q_{\mathcal{A}} \times D) \quad B_P := Q_{\mathcal{B}} \quad B_O := \Sigma \times \mathcal{P}(Q_{\mathcal{B}} \times D)$$

We define the *synchronous arrow game*

$$\Sigma \vdash \mathcal{G}(\mathcal{A}, \vec{M}) \text{ } \text{---} \otimes \text{ } \mathcal{G}(\mathcal{B}, \vec{N})$$

as the negative game $(V, E, *, \lambda, \xi, \mathcal{W})$ whose positions are given by

$$V := (D^* \times A_P) \times (D^* \times B_P) + (D^* \times A_O) \times (D^* \times B_P) + (D^* \times A_O) \times (D^* \times B_O)$$

whose polarized root is $*$:= $((\varepsilon, q_{\mathcal{A}}^i), (\varepsilon, q_{\mathcal{B}}^i))$ with $\xi(*) := O$, whole polarized edges (E, λ) are given in Table 1, and whose winning condition is given by

$$\begin{aligned} ((\varepsilon, q_{\mathcal{A}}^0), (\varepsilon, q_{\mathcal{B}}^0)) \cdot \dots \cdot ((\varepsilon, q_{\mathcal{A}}^n), (\varepsilon, q_{\mathcal{B}}^n)) \cdot \dots \in \mathcal{W} \\ \text{iff} \quad ((q_{\mathcal{A}}^i)_{i \in \mathbb{N}} \in \Omega_{\mathcal{A}} \implies (q_{\mathcal{B}}^i)_{i \in \mathbb{N}} \in \Omega_{\mathcal{B}}) \end{aligned}$$

Note that P-plays end in positions of the form

$$\begin{aligned} ((p, q_{\mathcal{A}}), (p, q_{\mathcal{B}})) &\in (D^* \times A_P) \times (D^* \times B_P) \\ \text{and} \quad ((p, a, \gamma_{\mathcal{A}}), (p, a, \gamma_{\mathcal{B}})) &\in (D^* \times A_O) \times (D^* \times B_O) \end{aligned}$$

Each of these position is of homogeneous type, and moreover in each case the D^* and Σ components coincide. On the other hand, O-plays end in positions of the form

$$\begin{aligned} ((p, a, \gamma_{\mathcal{A}}), (p, q_{\mathcal{A}})) &\in (D^* \times A_O) \times (D^* \times B_P) \\ \text{and} \quad ((p, a, \gamma_{\mathcal{A}}), (p \cdot d, q_{\mathcal{B}})) &\in (D^* \times A_O) \times (D^* \times B_P) \end{aligned}$$

Each of these intermediate position is of heterogeneous type, and in the second one, the D^* components do not coincide.

λ	$\mathcal{G}(\mathcal{A}, \vec{M})$	\multimap	$\mathcal{G}(\mathcal{B}, \vec{N})$	
	$((p, q_{\mathcal{A}})$,	$(p, q_{\mathcal{B}}))$	
O	\downarrow			
	$((p, a, \gamma_{\mathcal{A}})$,	$(p, q_{\mathcal{B}}))$	if $\gamma_{\mathcal{A}} \in \delta_{\mathcal{A}}(q_{\mathcal{A}}, \vec{M}(p)(a))$
P			\downarrow	
	$((p, a, \gamma_{\mathcal{A}})$,	$(p, a, \gamma_{\mathcal{B}}))$	if $\gamma_{\mathcal{B}} \in \delta_{\mathcal{B}}(q_{\mathcal{B}}, \vec{N}(p)(a))$
O			\downarrow	
	$((p, a, \gamma_{\mathcal{A}})$,	$(p.d, q'_{\mathcal{B}}))$	if $(q'_{\mathcal{B}}, d) \in \gamma_{\mathcal{B}}$
P	\downarrow			
	$((p.d, q'_{\mathcal{A}})$,	$(p.d, q'_{\mathcal{B}}))$	if $(q'_{\mathcal{A}}, d) \in \gamma_{\mathcal{A}}$

■ **Figure 1** Moves of $\mathcal{G}(\mathcal{A}, \vec{M}) \multimap \mathcal{G}(\mathcal{B}, \vec{N})$.

We write $\Sigma \vdash \sigma : \mathcal{G}(\mathcal{A}, M) \multimap \mathcal{G}(\mathcal{B}, N)$ if σ is a P-strategy on $\mathcal{G}(\mathcal{A}, M) \multimap \mathcal{G}(\mathcal{B}, N)$, and $\Sigma \vdash \sigma \Vdash \mathcal{G}(\mathcal{A}, M) \multimap \mathcal{G}(\mathcal{B}, N)$ if σ is moreover winning. Finally, we write

$$\Sigma \Vdash \mathcal{G}(\mathcal{A}, M) \multimap \mathcal{G}(\mathcal{B}, N)$$

if there is a winning P-strategy σ on $\mathcal{G}(\mathcal{A}, M) \multimap \mathcal{G}(\mathcal{B}, N)$.

► **Remark.** Recall that if $\Omega_{\mathcal{A}}$ and $\Omega_{\mathcal{B}}$ are Borel sets, then \mathcal{W} is a Borel set and by Martin's Theorem [12], either P or O has a winning strategy. Moreover, if the automata \mathcal{A} and \mathcal{B} are regular (in the sense of Sect. 2), then \mathcal{W} is an ω -regular language. If in addition the trees M and N are regular (in the usual sense), then the game is equivalent to a finite regular game. By Büchi-Landweber Theorem, the existence of a winning strategy for a given player is decidable, and the winning player has *finite state* winning strategies (see e.g. [17]).

3.3 Characterization of the Synchronous Arrow Games

We now give a characterization of synchronous arrow games in traditional games semantics. Our characterization involve relations in slices categories \mathbf{Set}/J , that will give rise to a strong analogy between our fibrations $\mathbf{game}^{(W)}$ and $\mathbf{aut}^{(W)}$ and substitution (a.k.a *change-of-base*) in the codomain fibration $\mathbf{cod} : \mathbf{Set}^{\rightarrow} \rightarrow \mathbf{Set}$.

Simple Games. Recall the usual notion of *simple games* (see e.g. [1, 7]). Simple games are usually negative, but given positive games A and B , their *negative* linear arrow $A \multimap B$ can still be defined. Moreover, simple games, with linear arrows $A \multimap B$ between games A and B of the same polarity, form a category that we write \mathbf{SGG} . When equipped with winning conditions, winning strategies compose, giving rise to a category that we write \mathbf{SGG}^W .

A P-strategy $\Sigma \vdash \sigma : \mathcal{G}(\mathcal{A}, M) \multimap \mathcal{G}(\mathcal{B}, N)$ is a morphism of \mathbf{SGG} from the substituted acceptance game $\mathcal{G}(\mathcal{A}, M)$ to the substituted acceptance game $\mathcal{G}(\mathcal{B}, N)$. If σ is moreover winning, then it is a morphism of \mathbf{SGG}^W .

The Hyland & Schalk Functor. Hyland & Schalk have presented in [9] a faithful functor, that we denote HS, from simple games to the category \mathbf{Rel} of sets and relations. This functor can easily be extended to a functor $\mathbf{HS} : \mathbf{SGG}^{(W)} \rightarrow \mathbf{Rel}$.

Given a play $s \in \wp(A \multimap B)$ we let $s|A \in \wp(A)$ be its projection on A and similarly for

B ,² so that $\text{HS}(s) := (s|A, s|B)$. Given a P-strategy $\sigma : A \multimap B$ we have $\sigma \subseteq \wp^{\text{P}}(A \multimap B)$ and thus

$$\text{HS}(\sigma) := \{\text{HS}(s) \mid s \in \sigma\} \subseteq \wp(A) \times \wp(B)$$

We write $\wp_{\Sigma}(\mathcal{A}, M)$ for the plays of the substituted acceptance game $\Sigma \vdash \mathcal{G}(\mathcal{A}, M)$. Given $\Sigma \vdash \sigma : \mathcal{G}(\mathcal{A}, M) \multimap \mathcal{G}(\mathcal{B}, N)$, we thus have

$$\text{HS}(\sigma) \subseteq \wp_{\Sigma}(\mathcal{A}, M) \times \wp_{\Sigma}(\mathcal{B}, N)$$

Synchronous Relations. We will now see that P-strategies on a synchronous arrow game can be seen as relations in slice categories \mathbf{Set}/J . We call such relations *synchronous*.

Given a set J , define the category $\mathbf{Rel}(\mathbf{Set}/J)$ as follows:

Objects are indexed sets $A \xrightarrow{g} J$, written simply A when g is understood from the context.

Morphisms from $A \xrightarrow{g} J$ to $B \xrightarrow{h} J$ are given by relations $\overset{\circ}{R} : A \dashrightarrow B$ such that the following commutes:

$$\begin{array}{ccc} & \overset{\circ}{R} & \\ \pi_1 \swarrow & & \searrow \pi_2 \\ A & & B \\ g \searrow & & \swarrow h \\ & J & \end{array}$$

Traces. For the synchronous arrow games, synchronization is performed using the following notion of *trace*. Given $\Gamma \vdash \mathcal{A}$ and $M \in \mathbf{Tree}[\Gamma, \Sigma]$, define

$$\text{tr} : \wp_{\Sigma}(\mathcal{A}, M) \longrightarrow (D + \Sigma)^*$$

inductively as follows

$$\text{tr}(\varepsilon) := \varepsilon \quad \text{tr}(s \rightarrow (p, a, \gamma)) := \text{tr}(s) \cdot a \quad \text{tr}(s \rightarrow (p \cdot d, q)) := \text{tr}(s) \cdot d$$

The image of tr is the set $\text{Tr}_{\Sigma} := (\Sigma \cdot D)^* + (\Sigma \cdot D)^* \cdot \Sigma$.

Characterization of the Synchronous Arrow. We can now characterize the synchronous arrow games.

► **Proposition 3.2.** *Strategies on the synchronous arrow game $\mathcal{G}(\mathcal{A}, M) \multimap \mathcal{G}(\mathcal{B}, N)$ are exactly the strategies $\sigma : \mathcal{G}(\mathcal{A}, M) \multimap \mathcal{G}(\mathcal{B}, N)$ such that*

$$\begin{array}{ccc} \text{HS}(\sigma) & \longrightarrow & \wp_{\Sigma}(\mathcal{B}, \vec{N}) \\ \downarrow & & \downarrow \text{tr} \\ \wp_{\Sigma}(\mathcal{A}, \vec{M}) & \xrightarrow{\text{tr}} & \text{Tr}_{\Sigma} \end{array} \quad (2)$$

► **Proposition 3.3.** *Let $\Sigma \vdash \mathcal{G}(\mathcal{A}, M)$ and $\Sigma \vdash \mathcal{G}(\mathcal{B}, N)$. The following is a pullback in \mathbf{Set} :*

$$\begin{array}{ccc} \wp_{\Sigma}^{\text{P}}(\mathcal{G}(\mathcal{A}, M) \multimap \mathcal{G}(\mathcal{B}, N)) & \xrightarrow{(-)|\mathcal{G}(\mathcal{B}, N)} & \wp_{\Sigma}(\mathcal{B}, N) \\ \downarrow (-)|\mathcal{G}(\mathcal{A}, M) & \lrcorner & \downarrow \text{tr} \\ \wp_{\Sigma}(\mathcal{A}, M) & \xrightarrow{\text{tr}} & \text{Tr}_{\Sigma} \end{array}$$

We write tr^{\multimap} for any of two equal maps $\text{tr} \circ (-)|\mathcal{G}(\mathcal{A}, M)$, $\text{tr} \circ (-)|\mathcal{G}(\mathcal{B}, N)$.

² We write $\wp(A)$ for the set of plays on A , and $\wp^{\text{P}}(A)$ for the set of P-plays.

3.4 Categories of Substituted Acceptance Games and Automata

We now define our categories $\mathbf{SAG}_\Sigma^{(W)}$ of substituted acceptance games and their full subcategories $\mathbf{Aut}_\Sigma^{(W)}$ of tree automata. That they indeed form categories follows from the characterization Prop. 3.2, together with the fact that $\mathbf{Rel}(\mathbf{Set}/J)$ and $\mathbf{SGG}^{(W)}$ are categories, and the fact that the identity strategies $\text{id} : \mathcal{G}(\mathcal{A}, M) \multimap \mathcal{G}(\mathcal{B}, N)$ are synchronous.

Objects of \mathbf{SAG}_Σ and \mathbf{SAG}_Σ^W are games $\Sigma \vdash \mathcal{G}(\mathcal{A}, M)$,

Morphisms of \mathbf{SAG}_Σ are synchronous strategies $\Sigma \vdash \sigma : \mathcal{G}(\mathcal{A}, M) \multimap \mathcal{G}(\mathcal{B}, N)$,

Morphisms of \mathbf{SAG}_Σ^W are synchronous winning strategies $\Sigma \vdash \sigma \Vdash \mathcal{G}(\mathcal{A}, M) \multimap \mathcal{G}(\mathcal{B}, N)$.

Objects of \mathbf{Aut}_Σ and \mathbf{Aut}_Σ^W are automata $\Sigma \vdash \mathcal{A}$,

Morphisms of \mathbf{Aut}_Σ are synchronous strategies $\Sigma \vdash \sigma : \mathcal{G}(\mathcal{A}, \text{Id}_\Sigma) \multimap \mathcal{G}(\mathcal{B}, \text{Id}_\Sigma)$,

Morphisms of \mathbf{Aut}_Σ^W are synchronous winning strategies $\Sigma \vdash \sigma \Vdash \mathcal{G}(\mathcal{A}, \text{Id}_\Sigma) \multimap \mathcal{G}(\mathcal{B}, \text{Id}_\Sigma)$.

A Lifting Property. Among the useful consequences of Prop. 3.3, we state the following lifting property.

► **Proposition 3.4.** *Consider $\Sigma \vdash \mathcal{G}(\mathcal{A}, M)$ and $\Sigma \vdash \mathcal{G}(\mathcal{B}, N)$. Assume that, in $\mathbf{Rel}(\mathbf{Set}/\text{Tr}_\Sigma)$ we have an isomorphism $\mathring{R} : (\wp_\Sigma(\mathcal{A}, M) \xrightarrow{\text{tr}} \text{Tr}_\Sigma) \dashv\dashv_{/\text{Tr}_\Sigma} (\wp_\Sigma(\mathcal{B}, N) \xrightarrow{\text{tr}} \text{Tr}_\Sigma)$.*

There is a (unique, total) isomorphism $\sigma : \mathcal{G}(\mathcal{A}, M) \xrightarrow{\mathbf{SAG}_\Sigma} \mathcal{G}(\mathcal{B}, N)$ s.t. $\text{HS}(\sigma) = R$.

In general we can not ask σ to be winning, and in particular to be a morphism of \mathbf{SAG}_Σ^W .

4 Fibrations of Acceptance Games and Automata

A tree morphism $L \in \mathbf{Tree}[\Sigma, \Gamma]$ defines a map L^* from the objects of \mathbf{SAG}_Γ to the objects of \mathbf{SAG}_Σ : we let $L^*(\Gamma \vdash \mathcal{G}(\mathcal{A}, M)) := \Sigma \vdash \mathcal{G}(\mathcal{A}, M \circ L)$.

In this Section, we show that L^* extends to functors $L^* : \mathbf{SAG}_\Gamma^{(W)} \rightarrow \mathbf{SAG}_\Sigma^{(W)}$ and that the operation $(-)^*$ is itself functorial and thus leads to split indexed categories $(-)^* : \mathbf{Tree}^{\text{op}} \rightarrow \mathbf{Cat}$. By applying Groethendieck completion, we obtain our split fibrations of acceptance games $\mathbf{game}^{(W)} : \mathbf{SAG}^{(W)} \rightarrow \mathbf{Tree}$.

On the other hand, by restricting substitution to tree morphisms generated by alphabet morphisms $\beta \in \mathbf{Alph}[\Sigma, \Gamma]$, we obtain functors $\beta^* : \mathbf{Aut}_\Gamma^{(W)} \rightarrow \mathbf{Aut}_\Sigma^{(W)}$ giving rise to split fibrations of tree automata $\mathbf{aut}^{(W)} : \mathbf{Aut}^{(W)} \rightarrow \mathbf{Alph}$.

Our substitution functors L^* are build in strong analogy with change-of-base functors $\mathbf{Set}/\text{Tr}_\Gamma \rightarrow \mathbf{Set}/\text{Tr}_\Sigma$ of the codomain fibration $\text{cod} : \mathbf{Set}^{\rightarrow} \rightarrow \mathbf{Set}$. We refer to [10] for basic material about fibrations.

4.1 Substitution Functors

Change-of-Base in $\mathbf{Set}^{\rightarrow}$. A morphism $L \in \mathbf{Tree}[\Sigma, \Gamma]$ induces a map $\text{Tr}(L) : \text{Tr}_\Sigma \rightarrow \text{Tr}_\Gamma$ inductively defined as follows (where $(-)_D$ is the obvious projection $\text{Tr}_\Sigma \rightarrow D^*$):

$$\text{Tr}(L)(\varepsilon) := \varepsilon \quad \text{Tr}(L)(w \cdot a) := \text{Tr}(L)(w) \cdot L(w_D)(a) \quad \text{Tr}(L)(w \cdot d) := \text{Tr}(L)(w) \cdot d$$

The map $\text{Tr}(L)$ gives rise to the usual change-of-base functor $L^\bullet : \mathbf{Set}/\text{Tr}_\Gamma \rightarrow \mathbf{Set}/\text{Tr}_\Sigma$, defined using chosen pullbacks in \mathbf{Set} :

$$\begin{array}{ccc} L^\bullet(\wp_\Gamma(\mathcal{A}, M)) & \longrightarrow & \wp_\Gamma(\mathcal{A}, M) \\ \downarrow L^\bullet(\text{tr}) & \lrcorner & \downarrow \text{tr} \\ \text{Tr}_\Sigma & \xrightarrow{\text{Tr}(L)} & \text{Tr}_\Gamma \end{array}$$

Substitution on Plays. The action of the substitution L^* on plays can be described, similarly as the action of L^\bullet on objects of $\mathbf{Set}/\mathbf{Tr}_\Gamma$, by a pullback property.

Consider $\Gamma \vdash \mathcal{G}(\mathcal{A}, M)$, so that $\Sigma \vdash \mathcal{G}(\mathcal{A}, M \circ L)$. A position $(p, a, \gamma_{\mathcal{A}})$ of the game $\Sigma \vdash \mathcal{G}(\mathcal{A}, M \circ L)$ can be mapped to the position $(p, L(p)(a), \gamma_{\mathcal{A}})$ of the game $\Gamma \vdash \mathcal{G}(\mathcal{A}, M)$. Moreover, since $\delta_{\mathcal{A}}(q_{\mathcal{A}}, (M \circ L)(p)(a)) = \delta_{\mathcal{A}}(q_{\mathcal{A}}, M(p)(L(p)(a)))$, we have

$$(p, q_{\mathcal{A}}) \rightarrow (p, a, \gamma_{\mathcal{A}}) \quad \text{if and only if} \quad (p, q_{\mathcal{A}}) \rightarrow (p, L(p)(a), \gamma_{\mathcal{A}})$$

This gives a map

$$\wp(L) : \wp_{\Sigma}(\mathcal{A}, M \circ L) \longrightarrow \wp_{\Gamma}(\mathcal{A}, M)$$

If we are also given $\Gamma \vdash \mathcal{G}(\mathcal{B}, N)$, then we similarly obtain

$$\wp(L)_{-\otimes} : \wp_{\Sigma}(\mathcal{G}(\mathcal{A}, M \circ L) -\otimes \mathcal{G}(\mathcal{B}, N \circ L)) \longrightarrow \wp_{\Gamma}((\mathcal{G}(\mathcal{A}, M) -\otimes \mathcal{G}(\mathcal{B}, N)))$$

These two maps are related *via* HS as expected: $\text{HS} \circ \wp(L)_{-\otimes} = (\wp(L) \times \wp(L)) \circ \text{HS}$. Moreover,

► **Proposition 4.1.** *We have, in \mathbf{Set} :*

$$\begin{array}{ccc} \wp_{\Sigma}(\mathcal{A}, M \circ L) & \xrightarrow{\wp(L)} & \wp_{\Gamma}(\mathcal{A}, M) & \wp_{\Sigma}^{\mathbb{P}}(\mathcal{G}(\mathcal{A}, M \circ L) -\otimes \mathcal{G}(\mathcal{B}, N \circ L)) \\ \downarrow \text{tr} & \lrcorner & \downarrow \text{tr} & \downarrow \text{tr}^{-\otimes} \\ \text{Tr}_{\Sigma} & \xrightarrow{\text{Tr}(L)} & \text{Tr}_{\Gamma} & \text{Tr}_{\Sigma} \end{array} \quad \begin{array}{ccc} & & \wp_{\Gamma}^{\mathbb{P}}(\mathcal{G}(\mathcal{A}, M) -\otimes \mathcal{G}(\mathcal{B}, N)) \\ & \searrow \wp(L)_{-\otimes} & \downarrow \text{tr}^{-\otimes} \\ & & \text{Tr}_{\Gamma} \end{array}$$

Substitution on Strategies. The action of L^* on strategies is defined using Prop. 4.1: Given $\Gamma \vdash \sigma : \mathcal{G}(\mathcal{A}, M) -\otimes \mathcal{G}(\mathcal{B}, N)$, so that $\sigma \subseteq \wp_{\Gamma}^{\mathbb{P}}(\mathcal{G}(\mathcal{A}, M) -\otimes \mathcal{G}(\mathcal{B}, N))$, we define

$$L^*(\sigma) := \wp(L)_{-\otimes}^{-1}(\sigma) \subseteq \wp_{\Sigma}^{\mathbb{P}}(\mathcal{G}(\mathcal{A}, M \circ L) -\otimes \mathcal{G}(\mathcal{B}, N \circ L))$$

► **Proposition 4.2.** *$L^*(\sigma)$ is a strategy. If moreover σ is winning, then $L^*(\sigma)$ is also winning.*

Functoriality of Substitution. Proposition 4.1 can be formulated by saying that the maps $\langle \text{tr}, \wp(L) \rangle$ and $\langle \text{tr}^{-\otimes}, \wp(L)_{-\otimes} \rangle$ are bijections, respectively:

$$\begin{array}{ccc} \wp_{\Sigma}(\mathcal{A}, M \circ L) & \xrightarrow{\cong} & \text{Tr}_{\Sigma} \times_{\text{Tr}_{\Sigma}} \wp_{\Gamma}(\mathcal{A}, M) \\ \wp_{\Sigma}^{\mathbb{P}}(\mathcal{G}(\mathcal{A}, M \circ L) -\otimes \mathcal{G}(\mathcal{B}, N \circ L)) & \xrightarrow{\cong} & \text{Tr}_{\Sigma} \times_{\text{Tr}_{\Sigma}} \wp_{\Gamma}^{\mathbb{P}}(\mathcal{G}(\mathcal{A}, M) -\otimes \mathcal{G}(\mathcal{B}, N)) \end{array}$$

These bijections are crucial to prove that

► **Proposition 4.3.** *L^* is a functor from $\mathbf{SAG}_{\Gamma}^{(W)}$ to $\mathbf{SAG}_{\Sigma}^{(W)}$.*

4.2 Fibrations of Acceptance Games

Consider now $L \in \mathbf{Tree}[\Sigma, \Gamma]$ and $K \in \mathbf{Tree}[\Gamma, \Delta]$. Since $\text{Tr}(K \circ L) = \text{Tr}(K) \circ \text{Tr}(L)$ and $\wp(K \circ L)_{(-\otimes)} = \wp(K)_{(-\otimes)} \circ \wp(L)_{(-\otimes)}$ we immediately get

► **Proposition 4.4.** *The operations $(-)^* : \mathbf{Tree}^{\text{op}} \rightarrow \mathbf{Cat}$, mapping Σ to $\mathbf{SAG}_{\Sigma}^{(W)}$, and mapping $L \in \mathbf{Tree}[\Sigma, \Gamma]$ to $L^* : \mathbf{SAG}_{\Gamma}^{(W)} \rightarrow \mathbf{SAG}_{\Sigma}^{(W)}$ are functors.*

By using Grothendieck completion (see e.g. [10, §1.10]), this gives us split fibrations of acceptance games $\text{game}^{(W)} : \mathbf{SAG}^{(W)} \rightarrow \mathbf{Tree}$ that we do not detail here by lack of space.

4.3 Fibrations of Automata

In order to obtain fibrations of automata, we restrict substitution to tree morphisms generated by alphabet morphisms $\beta \in \mathbf{Alph}[\Sigma, \Gamma]$. The crucial point is that these restricted substitutions can be internalized in automata.

Given $\Gamma \vdash \mathcal{A}$ with $\mathcal{A} = (Q, q^t, \delta, \Omega)$, and $\beta \in \mathbf{Alph}[\Sigma, \Gamma]$, define the automaton $\Sigma \vdash \mathcal{A}[\beta]$ as $\mathcal{A}[\beta] := (Q, q^t, \delta_\beta, \Omega)$ where $\delta_\beta(q, a) := \delta(q, \beta(a))$.

► **Proposition 4.5.** $\Sigma \vdash \mathcal{G}(\mathcal{A}[\beta], \text{Id}_\Sigma) = \Sigma \vdash \mathcal{G}(\mathcal{A}, \beta)$.

It is easy to see that $(-)^*$ restricts to a functor from $\mathbf{Alph}^{\text{op}}$ to \mathbf{Cat} , so that we get fibrations

$$\text{aut}^{(W)} : \mathbf{Aut}^{(W)} \longrightarrow \mathbf{Alph} \quad \text{with, in Cat:} \quad \begin{array}{ccc} \mathbf{Aut}^{(W)} & \longrightarrow & \mathbf{SAG}^{(W)} \\ \text{aut}^{(W)} \downarrow & \lrcorner & \downarrow \text{game}^{(W)} \\ \mathbf{Alph} & \longrightarrow & \mathbf{Tree} \end{array}$$

5 Symmetric Monoidal Structure

We now consider a synchronous product of automata. When working on *complete* automata (to be defined in Sect. 5.1 below), it gives rise to split symmetric monoidal fibrations, in the sense of [16].

According to [16, Thm. 12.7], split symmetric monoidal fibrations can equivalently be obtained from split symmetric monoidal indexed categories. In our context, this means that the functors $(-)^*$ extend to

$$(-)^* : \mathbf{Tree}^{\text{op}} \longrightarrow \mathbf{SymMonCat} \quad (-)^* : \mathbf{Alph}^{\text{op}} \longrightarrow \mathbf{SymMonCat}$$

where $\mathbf{SymMonCat}$ is the category of symmetric monoidal categories and strong monoidal functors. Hence, we equip our categories of (complete) acceptance games and automata with a symmetric monoidal structure. Substitution turns out to be *strict* symmetric monoidal.

We refer to [13] for background on symmetric monoidal categories.

5.1 Complete Tree Automata

An automaton \mathcal{A} is *complete* if for every $(q, a) \in Q \times \Sigma$, the set $\delta(q, a)$ is not empty and moreover for every $\gamma \in \delta(q, a)$ and every $d \in D$, we have $(q', d) \in \gamma$ for some $q' \in Q$.

Given an automaton $\mathcal{A} = (Q, q^t, \delta, \Omega)$ its *completion* is the automaton $\widehat{\mathcal{A}} := (\widehat{Q}, q^t, \widehat{\delta}, \widehat{\Omega})$ with states $\widehat{Q} := Q + \{\text{true}, \text{false}\}$, with acceptance condition $\widehat{\Omega} := \Omega + Q^* \cdot \text{true} \cdot \widehat{Q}^\omega$, and with transition function $\widehat{\delta}$ defined as

$$\begin{aligned} \widehat{\delta}(\text{true}, q) &:= \{ \{ (\text{true}, d) \mid d \in D \} \} & \widehat{\delta}(\text{false}, q) &:= \{ \{ (\text{false}, d) \mid d \in D \} \} \\ \widehat{\delta}(q, a) &:= \{ \{ (\text{false}, d) \mid d \in D \} \} & & \text{if } q \in Q \text{ and } \delta(q, a) = \emptyset \\ \widehat{\delta}(q, a) &:= \{ \widehat{\gamma} \mid \gamma \in \delta(q, a) \} & & \text{otherwise} \end{aligned}$$

where, given $\gamma \in \delta(q, a)$, we let $\widehat{\gamma} := \gamma \cup \{ (\text{true}, d) \mid \text{there is no } q \in Q \text{ s.t. } (q, d) \in \gamma \}$.

► **Proposition 5.1.** $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\widehat{\mathcal{A}})$.

Restricting to complete automata gives rise to full subcategories $\widehat{\mathbf{SAG}}_\Sigma^{(W)}$ and $\widehat{\mathbf{Aut}}_\Sigma^{(W)}$ of resp. $\mathbf{SAG}_\Sigma^{(W)}$ and $\mathbf{Aut}_\Sigma^{(W)}$, and thus induces fibrations

$$\widehat{\text{game}} : \widehat{\mathbf{SAG}}_\Sigma^{(W)} \longrightarrow \mathbf{Tree} \quad \widehat{\text{aut}} : \widehat{\mathbf{Aut}}_\Sigma^{(W)} \longrightarrow \mathbf{Alph}$$

5.2 The Synchronous Product

Assume given complete automata $\Sigma \vdash \mathcal{A}$ and $\Sigma \vdash \mathcal{B}$. Define $\Sigma \vdash \mathcal{A} \otimes \mathcal{B}$ as

$$\mathcal{A} \otimes \mathcal{B} := (Q_{\mathcal{A}} \times Q_{\mathcal{B}}, (q_{\mathcal{A}}^i, q_{\mathcal{B}}^i), \delta_{\mathcal{A} \otimes \mathcal{B}}, \Omega_{\mathcal{A} \otimes \mathcal{B}})$$

where $(q_{\mathcal{A}}^n, q_{\mathcal{B}}^n)_{n \in \mathbb{N}} \in \Omega_{\mathcal{A} \otimes \mathcal{B}}$ iff $((q_{\mathcal{A}}^n)_{n \in \mathbb{N}} \in \Omega_{\mathcal{A}}$ and $(q_{\mathcal{B}}^n)_{n \in \mathbb{N}} \in \Omega_{\mathcal{B}})$, and where we let $\delta_{\mathcal{A} \otimes \mathcal{B}}((q_{\mathcal{A}}, q_{\mathcal{B}}), a)$ be the set of all the $\gamma_{\mathcal{A}} \otimes \gamma_{\mathcal{B}}$ for $\gamma_{\mathcal{A}} \in \delta_{\mathcal{A}}(q_{\mathcal{A}}, a)$ and $\gamma_{\mathcal{B}} \in \delta_{\mathcal{B}}(q_{\mathcal{B}}, a)$, with $\gamma_{\mathcal{A}} \otimes \gamma_{\mathcal{B}} := \{((q'_{\mathcal{A}}, q'_{\mathcal{B}}), d) \mid d \in D \text{ and } (q'_{\mathcal{A}}, d) \in \gamma_{\mathcal{A}} \text{ and } (q'_{\mathcal{B}}, d) \in \gamma_{\mathcal{B}}\}$.

Note that since \mathcal{A} and \mathcal{B} are complete, each $\gamma_{\mathcal{A} \otimes \mathcal{B}} \in \delta_{\mathcal{A} \otimes \mathcal{B}}((q_{\mathcal{A}}, q_{\mathcal{B}}), a)$ uniquely decomposes as $\gamma_{\mathcal{A} \otimes \mathcal{B}} = \gamma_{\mathcal{A}} \otimes \gamma_{\mathcal{B}}$.

Action on Plays. The unique decomposition property of $\gamma_{\mathcal{A} \otimes \mathcal{B}}$ allows to define projections

$$\begin{aligned} \varpi_i & : & \varphi_{\Sigma}(\mathcal{A}_1 \otimes \mathcal{A}_2, M) & \longrightarrow & \varphi_{\Sigma}(\mathcal{A}_i, M) \\ \varpi_i^{-\otimes} & : & \varphi_{\Sigma}(\mathcal{G}(\mathcal{A}_1 \otimes \mathcal{B}_1, M) -\otimes \mathcal{G}(\mathcal{A}_2 \otimes \mathcal{B}_2, N)) & \longrightarrow & \varphi_{\Sigma}(\mathcal{G}(\mathcal{A}_i, M) -\otimes \mathcal{G}(\mathcal{B}_i, N)) \end{aligned}$$

We write $\text{SP} := \langle \varpi_1, \varpi_2 \rangle$ and $\text{SP}_{-\otimes} := \langle \varpi_1^{-\otimes}, \varpi_2^{-\otimes} \rangle$.

► **Proposition 5.2.** *We have, in Set:*

$$\begin{array}{ccc} \varphi_{\Sigma}(\mathcal{A} \otimes \mathcal{B}, M) & \xrightarrow{\varpi_2} & \varphi_{\Sigma}(\mathcal{B}, M) & \varphi_{\Sigma}^{\text{P}}(\mathcal{G}(\mathcal{A} \otimes \mathcal{B}, M) -\otimes \mathcal{G}(\mathcal{C} \otimes \mathcal{D}, N)) \\ \varpi_1 \downarrow \lrcorner & & \downarrow \text{tr} & \downarrow \lrcorner \\ \varphi_{\Sigma}(\mathcal{A}, M) & \xrightarrow{\text{tr}} & \text{Tr}_{\Sigma} & \varphi_{\Sigma}^{\text{P}}(\mathcal{G}(\mathcal{B}, M) -\otimes \mathcal{G}(\mathcal{D}, N)) \\ & & & \downarrow \text{tr} \\ & & & \varphi_{\Sigma}^{\text{P}}(\mathcal{G}(\mathcal{A}, M) -\otimes \mathcal{G}(\mathcal{C}, N)) \xrightarrow{\text{tr}} \text{Tr}_{\Sigma} \end{array}$$

Action on Synchronous Games. The action of \otimes on the objects of $\widehat{\text{SAG}}_{\Sigma}^{(\text{W})}$ is given by

$$(\Sigma \vdash \mathcal{G}(\mathcal{A}, M)) \otimes (\Sigma \vdash \mathcal{G}(\mathcal{B}, N)) := \Sigma \vdash \mathcal{G}(\mathcal{A}[\pi] \otimes \mathcal{B}[\pi'], \langle M, N \rangle)$$

where π and π' are suitable projections. For morphisms, let $\Sigma \vdash \sigma : \mathcal{G}(\mathcal{A}_0, M_0) -\otimes \mathcal{G}(\mathcal{A}_1, M_1)$ and $\Sigma \vdash \tau : \mathcal{G}(\mathcal{B}_0, N_0) -\otimes \mathcal{G}(\mathcal{B}_1, N_1)$. Then since $\Sigma \vdash \mathcal{G}(\mathcal{A}_i[\pi_i], \langle M_i, N_i \rangle) = \Sigma \vdash \mathcal{G}(\mathcal{A}_i, M_i)$ and $\Sigma \vdash \mathcal{G}(\mathcal{B}_i[\pi'_i], \langle M_i, N_i \rangle) = \Sigma \vdash \mathcal{G}(\mathcal{B}_i, N_i)$, thanks to Prop. 5.2 we can simply let $\sigma \otimes \tau := \text{SP}_{-\otimes}^{-1}(\sigma, \tau)$.

► **Proposition 5.3.** *The product $-\otimes-$ gives functors $\widehat{\text{SAG}}_{\Sigma}^{(\text{W})} \times \widehat{\text{SAG}}_{\Sigma}^{(\text{W})} \longrightarrow \widehat{\text{SAG}}_{\Sigma}^{(\text{W})}$.*

5.3 Symmetric Monoidal Structure

Thanks to Prop. 5.2 and Prop. 3.4 the symmetric monoidal structure of \otimes in $\widehat{\text{SAG}}_{\Sigma}^{(\text{W})}$ can be directly obtained from the symmetric monoidal structure of the tensorial product of $\text{Rel}(\text{Set}/\text{Tr}_{\Sigma})$.

Symmetric Monoidal Structure in $\text{Rel}(\text{Set}/J)$. We define a product \otimes in $\text{Rel}(\text{Set}/J)$:

On Objects: for (A, g) and (B, h) objects in $\text{Rel}(\text{Set}/J)$ the product $A \otimes B$ is $A \times_J B$ with the corresponding map, that is

$$A \otimes B := \{(a, b) \in A \times B \mid g(a) = h(b)\} \xrightarrow{g \circ \pi_1 = h \circ \pi_2} J$$

On Morphisms: given $R \in \mathbf{Rel}(\mathbf{Set}/J)[A, C]$ and $P \in \mathbf{Rel}(\mathbf{Set}/J)[B, D]$, we define $R \otimes P \in \mathbf{Rel}(\mathbf{Set}/J)[A \otimes B, C \otimes D]$ as

$$R \otimes P := \{((a, b), (c, d)) \in (A \otimes B) \times (C \otimes D) \mid (a, c) \in R \text{ and } (b, d) \in P\}$$

For the unit, we *choose* some $\mathbf{I} = (\mathbf{j} : I \xrightarrow{\simeq} J)$. Note that \mathbf{j} is required to be a bijection. The natural isomorphisms are given by:

$$\begin{aligned} \hat{\alpha}_{A,B,C} &:= \{(((a, b), c), (a, (b, c))) \mid g_A(a) = g_B(b) = g_C(c)\} \\ \hat{\lambda}_A &:= \{((e, a), a) \mid \mathbf{j}(e) = g_A(a)\} \\ \hat{\rho}_A &:= \{((a, e), a) \mid g_A(a) = \mathbf{j}(e)\} \\ \hat{\gamma}_{A,B} &:= \{((a, b), (b, a)) \mid g_A(a) = g_B(b)\} \end{aligned}$$

► **Proposition 5.4.** *The category $\mathbf{Rel}(\mathbf{Set}/J)$, equipped with the above data, is symmetric monoidal.*

Unit Automata. The requirement that the monoidal unit $\mathbf{j} : I \rightarrow J$ of $\mathbf{Rel}(\mathbf{Set}/J)$ should be a bijection leads us to the following unit automata. We let $\mathcal{I} := (Q_{\mathcal{I}}, q_{\mathcal{I}}, \delta_{\mathcal{I}}, \Omega_{\mathcal{I}})$ where $Q_{\mathcal{I}} := \mathbf{1}$, $q_{\mathcal{I}} := \bullet$, $\Omega_{\mathcal{I}} = Q_{\mathcal{I}}^{\omega}$ and $\delta_{\mathcal{I}}(q_{\mathcal{I}}, a) := \{\{(q_{\mathcal{I}}, d) \mid d \in D\}\}$.

Note that since $\delta_{\mathcal{I}}$ is constant, we have $\Sigma \vdash \mathcal{G}(\mathcal{I}, M) = \Sigma \vdash \mathcal{G}(\mathcal{I}, \text{Id})$. Moreover,

► **Proposition 5.5.** *Given $M \in \mathbf{Tree}[\Sigma, \Gamma]$, we have, in \mathbf{Set} , a bijection*

$$\text{tr} : \wp_{\Sigma}(\mathcal{I}, M) \xrightarrow{\simeq} \text{Tr}_{\Sigma}$$

Symmetric Monoidal Structure. Using Prop. 3.4, the structure isos of $\mathbf{Rel}(\mathbf{Set}/\text{Tr}_{\Sigma})$ can be lifted to $\widehat{\mathbf{SAG}}_{\Sigma}^{(W)}$ (winning is trivial). Moreover, the required equations (naturality and coherence) follows from Prop. 3.3, Prop 5.2, and the fact that $((\text{SP} \times \text{SP}) \circ \text{HS})(\sigma \otimes \tau) = \text{HS}(\sigma) \otimes \text{HS}(\tau)$ (where composition on the left is in \mathbf{Set} , and the expression denotes the actions of the resulting function on the set of plays $(\sigma \otimes \tau)$).

► **Proposition 5.6.** *The categories $\widehat{\mathbf{SAG}}_{\Sigma}^{(W)}$ and $\widehat{\mathbf{Aut}}_{\Sigma}^{(W)}$ equipped with the above data, are symmetric monoidal.*

5.4 Symmetric Monoidal Fibrations

In order to obtain symmetric monoidal fibrations, by [16, Thm. 12.7], it remains to check that substitution is strong monoidal. It is actually *strict* monoidal: it directly commutes with \otimes and preserves the unit, as well as all the structure maps.

► **Proposition 5.7.**

- (i) *Given $L \in \mathbf{Tree}[\Sigma, \Gamma]$, the functors $L^* : \widehat{\mathbf{SAG}}_{\Gamma}^{(W)} \rightarrow \widehat{\mathbf{SAG}}_{\Sigma}^{(W)}$ are strict monoidal.*
- (ii) *Given $\beta \in \mathbf{Alph}[\Sigma, \Gamma]$, the functors $\beta^* : \widehat{\mathbf{Aut}}_{\Gamma}^{(W)} \rightarrow \widehat{\mathbf{Aut}}_{\Sigma}^{(W)}$ are strict monoidal.*

6 Correctness w.r.t. Language Operations

This Section gathers several properties stating the correctness of our constructions w.r.t. operations on recognized languages. We begin in Sect. 6.1 by properties on the symmetric monoidal structure, the most important one being that the synchronous arrow is *correct*, in the sense that $\Sigma \vdash \mathcal{A} \text{---}\otimes\text{---} \mathcal{B}$ implies $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(\mathcal{B})$. Then, in Sect. 6.2, we discuss complementation of automata, and its relation with the synchronous arrow.

6.1 Correctness of the Symmetric Monoidal Structure

We begin by a formal correspondence between acceptance games and synchronous games of a specific form. This allows to show that the synchronous arrow is *correct*, in the sense that $\Sigma \vdash \mathcal{A} \multimap \mathcal{B}$ implies $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(\mathcal{B})$. We then briefly discuss the correctness of the synchronous product w.r.t. language intersection.

► **Proposition 6.1.** *Given $\Sigma \vdash \mathcal{A}$ and $t \in \mathbf{Tree}[\Sigma]$, there is a bijection:*

$$\{\sigma \mid \mathbf{1} \vdash \sigma \Vdash \mathcal{G}(\mathcal{A}, t)\} \quad \simeq \quad \{\theta \mid \mathbf{1} \vdash \theta \Vdash \mathcal{G}(\mathcal{I}, \text{Id}_{\mathbf{1}}) \multimap \mathcal{G}(\mathcal{A}, t)\}$$

► **Remark.** The above correspondence is only possible for acceptance games over $\mathbf{1}$:

- In $\Sigma \vdash \sigma \Vdash \mathcal{G}(\mathcal{A}, M)$, σ is a positive P-strategy, hence chooses the input characters in Σ .
- In $\Sigma \vdash \theta \Vdash \mathcal{G}(\mathcal{I}_{\Sigma}, \text{Id}_{\Sigma}) \multimap \mathcal{G}(\mathcal{A}, M)$, the strategy θ is a negative. It plays positively in $\Sigma \vdash \mathcal{G}(\mathcal{A}, M)$, but must follow the input characters chosen by \mathbf{O} in $\Sigma \vdash \mathcal{G}(\mathcal{I}_{\Sigma}, \text{Id}_{\Sigma})$.

We now check that the arrow $\mathcal{G}(\mathcal{A}, M) \multimap \mathcal{G}(\mathcal{B}, N)$ is correct w.r.t. language inclusion:

► **Proposition 6.2 (Correctness of the Arrow).** *Assume given $\Sigma \vdash \sigma \Vdash \mathcal{G}(\mathcal{A}, M) \multimap \mathcal{G}(\mathcal{B}, N)$.*

- (i) *For all $t \in \mathbf{Tree}[\Sigma]$, we have $t^*(\sigma) \Vdash \mathcal{G}(\mathcal{A}, M \circ t) \multimap \mathcal{G}(\mathcal{B}, N \circ t)$.*
- (ii) *If $\mathbf{1} \Vdash \mathcal{G}(\mathcal{A}, M \circ t)$ then $\mathbf{1} \Vdash \mathcal{G}(\mathcal{B}, N \circ t)$.*
- (iii) *For all tree $t \in \mathbf{Tree}[\Sigma]$, if $M(t) \in \mathcal{L}(\mathcal{A})$ then $N(t) \in \mathcal{L}(\mathcal{B})$.*

The converse property will be discussed in Sect. 7. We finally check that the synchronous product is correct.

► **Proposition 6.3.** $\mathcal{L}(\widehat{\mathcal{A}} \otimes \widehat{\mathcal{B}}) = \mathcal{L}(\mathcal{A}) \cap \mathcal{L}(\mathcal{B})$.

6.2 Complementation and Falsity

Complementation. Given an automaton $\mathcal{A} = (Q, q^i, \delta, \Omega)$, following [18], we let its complement be $\sim \mathcal{A} := (Q, q^i, \delta_{\sim \mathcal{A}}, \Omega_{\sim \mathcal{A}})$, where $\Omega_{\sim \mathcal{A}} := Q^{\omega} \setminus \Omega$ and

$$\delta_{\sim \mathcal{A}}(q, a) := \{\gamma_{\sim} \in \mathcal{P}(Q \times D) \mid \forall \gamma \in \delta(q, a), \gamma_{\sim} \cap \gamma \neq \emptyset\}$$

The idea is that P on $\sim \mathcal{A}$ simulates O on \mathcal{A} , so that the correctness of $\sim \mathcal{A}$ relies on determinacy of acceptance games. In particular, thanks to Borel determinacy [12], we have:

► **Proposition 6.4 ([18]).** *Given \mathcal{A} with $\Omega_{\mathcal{A}}$ a Borel set, we have $\mathcal{L}(\sim \mathcal{A}) = \mathbf{Tree}[\Sigma] \setminus \mathcal{L}(\mathcal{A})$.*

Note that if \mathcal{A} is complete, then $\sim \mathcal{A}$ is not necessarily complete, but $\delta_{\sim \mathcal{A}}$ is always not empty and so are the γ 's in its image.

The Falsity Automaton \perp . We let $\perp := (Q_{\perp}, q_{\perp}, \delta_{\perp}, \Omega_{\perp})$ where $Q_{\perp} := \mathbf{1}$, $q_{\perp} := \bullet$, $\Omega_{\perp} = \emptyset$ and $\delta_{\perp}(q_{\perp}, a) := \{\{(q_{\perp}, d)\} \mid d \in D\}$. Note that $\mathcal{I} = \sim \perp$. In particular, it is actually P who guides the evaluation of \perp , by choosing the tree directions.

► **Proposition 6.5.** *Let \mathcal{A} and \mathcal{B} be complete. Then $\Sigma \Vdash \mathcal{A} \otimes \mathcal{B} \multimap \widehat{\perp}$ iff $\Sigma \Vdash \mathcal{A} \multimap \widehat{\sim \mathcal{B}}$.*

► **Corollary 6.6.** *Let \mathcal{A} be a complete automaton on Σ . Then $\mathbf{1} \Vdash \widehat{\sim \mathcal{A}}$ iff $\mathbf{1} \Vdash \mathcal{A} \multimap \widehat{\perp}$.*

7 Projection and Fibred Simple Coproducts

We now check that automata can be equipped with existential quantifications in the fibred sense. Namely, given a projection $\pi \in \mathbf{Alph}[\Sigma \times \Gamma, \Sigma]$, the induced weakening functor $\pi^* : \widehat{\mathbf{Aut}}_{\Sigma}^{(W)} \rightarrow \widehat{\mathbf{Aut}}_{\Sigma \times \Gamma}^{(W)}$ has a left-adjoint $\Pi_{\Sigma, \Gamma}$, and moreover this structure is preserved by substitution, in the sense of the Beck-Chevalley condition (see e.g. [10]). This will lead to a (weak) completeness property of the synchronous arrow on *non-deterministic* automata, to be discussed below.

Recall from [11, Thm. IV.1.2.(ii)] that an adjunction $\Pi_{\Sigma, \Gamma} \dashv \pi^*$, with π^* a functor, is completely determined by the following data: To each object $\Sigma \times \Gamma \vdash \mathcal{A}$, an object $\Sigma \vdash \Pi_{\Sigma, \Gamma} \mathcal{A}$, and a map $\eta_{\mathcal{A}} : \Sigma \times \Gamma \vdash \mathcal{A} \rightarrow \Sigma \times \Gamma \vdash (\Pi_{\Sigma, \Gamma} \mathcal{A})[\pi]$ satisfying the following universal lifting property:

$$\begin{array}{l} \text{For every} \\ \sigma : \Sigma \times \Gamma \vdash \mathcal{A} \longrightarrow \Sigma \times \Gamma \vdash \mathcal{B}[\pi] \\ \text{there is a unique} \\ \tau : \Sigma \vdash \Pi_{\Sigma, \Gamma} \mathcal{A} \longrightarrow \Sigma \vdash \mathcal{B} \end{array} \quad \text{s.t.} \quad \begin{array}{ccc} \mathcal{A} & \xrightarrow{\eta_{\mathcal{A}}} & (\Pi_{\Sigma, \Gamma} \mathcal{A})[\pi] \\ & \searrow \sigma & \downarrow \pi^*(\tau) \\ & & \mathcal{B}[\pi] \end{array} \quad (3)$$

In our context, the Beck-Chevalley condition amounts to the equalities

$$\Delta \vdash (\Pi_{\Sigma, \Gamma} \mathcal{A})[\beta] = \Delta \vdash \Pi_{\Delta, \Gamma}(\mathcal{A}[\beta \times \text{Id}_{\Gamma}]) \quad \eta_{\mathcal{A}[\beta \times \text{Id}_{\Gamma}]} = (\beta \times \text{Id}_{\Gamma})^*(\eta_{\mathcal{A}}) \quad (4)$$

It turns out that the usual projection operation on automata (see e.g. [18]) is not functorial. Surprisingly, this is independent from whether automata are non-deterministic or not³. We devise a *lifted* projection operation, which indeed leads to a fibred existential quantification, and which is correct, on non-deterministic automata, w.r.t. the recognized languages.

The Lifted Projection. Consider $\Sigma \times \Gamma \vdash \mathcal{A}$ with $\mathcal{A} = (Q, q^i, \delta, \Omega)$. Define $\Sigma \vdash \Pi_{\Sigma, \Gamma} \mathcal{A}$ as $\Pi_{\Sigma, \Gamma} \mathcal{A} := (Q \times \Gamma + \{q^i\}, q^i, \delta_{\Pi \mathcal{A}}, \Omega_{\Pi \mathcal{A}})$ where

$$\begin{aligned} \delta_{\Pi \mathcal{A}}(q^i, a) &:= \bigcup_{b \in \Gamma} \{\gamma^{+b} \mid \gamma \in \delta(q^i, (a, b))\} \\ \delta_{\Pi \mathcal{A}}((q, _), a) &:= \bigcup_{b \in \Gamma} \{\gamma^{+b} \mid \gamma \in \delta(q, (a, b))\} \end{aligned}$$

and, given $\gamma \in \mathcal{P}(Q \times D)$ and $b \in \Gamma$, we let $\gamma^{+b} := \{(q^{+b}, d) \mid (q, d) \in \gamma\}$ with $q^{+b} := (q, b)$.

For the acceptance condition, we let $q^i \cdot (q_0, b_0) \cdot \dots \cdot (q_n, b_n) \cdot \dots$ in $\Omega_{\Pi \mathcal{A}}$ iff $q^i \cdot q_0 \cdot \dots \cdot q_n \cdot \dots \in \Omega$.

Action on Plays of The Lifted Projection. The action on plays of $\Pi_{\Sigma, \Gamma}$ is characterized by the map $\wp(\Pi) : \wp_{\Sigma \times \Gamma}(\mathcal{A}) \rightarrow \wp_{\Sigma}(\Pi_{\Sigma, \Gamma} \mathcal{A})$ inductively defined as $\wp(\Pi)(\varepsilon, q^i) := (\varepsilon, q^i)$ and

$$\begin{aligned} \wp(\Pi)((\varepsilon, q^i) \rightarrow^* (p, q) \rightarrow (p, (a, b), \gamma)) &:= \wp(\Pi)((\varepsilon, q^i) \rightarrow^* (p, q) \rightarrow (p, a, \gamma^{+b})) \\ \wp(\Pi)((\varepsilon, q^i) \rightarrow^* (p, (a, b), \gamma) \rightarrow (p.d, q)) &:= \wp(\Pi)((\varepsilon, q^i) \rightarrow^* (p, (a, b), \gamma) \rightarrow (p.d, q^{+b})) \end{aligned}$$

► **Proposition 7.1.** *If \mathcal{A} is a complete automaton, then $\wp(\Pi)$ is a bijection.*

³ It is well-known that the projection operation is correct w.r.t. the recognized languages only on *non-deterministic automata*.

The Unit Maps $\eta_{(-)}$. Consider the injection $\iota_{\Sigma, \Gamma} : \wp_{\Sigma}(\Pi_{\Sigma, \Gamma} \mathcal{A}) \longrightarrow \wp_{\Sigma \times \Gamma}((\Pi_{\Sigma, \Gamma} \mathcal{A})[\pi])$ inductively defined as $\iota_{\Sigma, \Gamma}((\varepsilon, q_{\mathcal{A}}^t)) := (\varepsilon, q_{\mathcal{A}}^t)$ and $\iota_{\Sigma, \Gamma}(s \rightarrow (p, q^{+b})) := \iota_{\Sigma, \Gamma}(s) \rightarrow (p, q^{+b})$ and $\iota_{\Sigma, \Gamma}(s \rightarrow (p, a, \gamma^{+b})) := \iota_{\Sigma, \Gamma}(s) \rightarrow (p, (a, b), \gamma^{+b})$.

If $\Sigma \times \Gamma \vdash \mathcal{A}$ is complete, we let the unit $\eta_{\mathcal{A}}$ be the unique strategy of $\widehat{\mathbf{SAG}}_{\Sigma \times \Gamma}^{\mathbf{W}}$ such that $\text{HS}(\eta_{\mathcal{A}}) = \{(t, \iota_{\Sigma, \Gamma} \circ \wp(\Pi)(t)) \mid t \in \wp_{\Sigma \times \Gamma}(\mathcal{A})\}$. We do not detail the B.-C. condition (4).

The Unique Lifting Property (3). Consider some $\Sigma \times \Gamma \vdash \sigma : \mathcal{A} \multimap \mathcal{B}[\pi]$ with \mathcal{A} complete. We let τ be the unique strategy such that $\text{HS}(\tau) = \{(\wp(\Pi)(s), \wp(\pi)(t)) \mid (s, t) \in \text{HS}(\sigma)\}$. It is easy to see that τ is winning whenever σ is winning. Moreover

► **Lemma 7.2.** $\sigma = \pi^*(\tau) \circ \eta_{\mathcal{A}}$.

For the unicity part of the lifting property of $\eta_{\mathcal{A}}$, it is sufficient to check:

► **Lemma 7.3.** *If $\pi^*(\theta) \circ \eta_{\mathcal{A}} = \pi^*(\theta') \circ \eta_{\mathcal{A}}$ then $\theta = \theta'$.*

Non-Deterministic Tree Automata. An automaton \mathcal{A} is *non-deterministic* if for every γ in the image of δ and every direction $d \in D$, there is at most one state q such that $(q, d) \in \gamma$.

► Remark. If \mathcal{A} and \mathcal{B} are non-deterministic, then so are $\mathcal{A} \otimes \mathcal{B}$ and $\Pi(\mathcal{A})$.

► **Proposition 7.4** ([4, 15, 18]). *For each regular automaton $\Sigma \vdash \mathcal{A}$ there is a complete non-deterministic automaton $\Sigma \vdash \text{ND}(\mathcal{A})$ such that $\mathcal{L}(\mathcal{A}) = \mathcal{L}(\text{ND}(\mathcal{A}))$.*

► **Proposition 7.5.** *If $\Sigma \times \Gamma \vdash \mathcal{A}$ is non-deterministic and complete, then $\mathcal{L}(\Pi_{\Sigma, \Gamma} \mathcal{A}) = \pi_{\Sigma, \Gamma}(\mathcal{L}(\mathcal{A}))$ where $\pi_{\Sigma, \Gamma} \in \mathbf{Alph}[\Sigma \times \Gamma, \Sigma]$ is the first projection.*

► **Proposition 7.6.** *Consider complete regular automata $\Sigma \vdash \mathcal{A}$ and $\Sigma \vdash \mathcal{B}$. If $\mathcal{L}(\mathcal{A}) \subseteq \mathcal{L}(\mathcal{B})$ then $\Sigma \Vdash \text{ND}(\mathcal{A}) \multimap \widehat{\sim} \mathcal{C}$ with $\mathcal{C} := \text{ND}(\sim \mathcal{B})$.*

8 Conclusion

We presented monoidal fibrations of tree automata and acceptance games, in which the fibre categories are based on a synchronous restriction of linear simple games.

For technical simplicity, we did not yet consider monoidal closure, but strongly expect that it holds. One of the main question is whether suitable restrictions of these categories are Cartesian closed, so as to interpret proofs from intuitionistic variants of MSO. Among other questions are the status of non-determinization (*i.e.* whether it can be made functorial, or even co-monadic), as well as relation with the Dialectica interpretation (in the vein of e.g. [8]). Our result of weak completeness (Prop. 7.6) suggests strong connections with the notion of *guidable* non-deterministic automata of [2]. On a similar vein, connections with *game automata* [3, 5] might be relevant to investigate.

Acknowledgments. This work benefited from numerous discussions with Pierre Clairambault and Thomas Colcombet.

References

- 1 S. Abramsky. Semantics of Interaction. In A. M. Pitts and P. Dybjer, editors, *Semantics and Logics of Computation*, volume 14 of *Publications of the Newton Institute*, page 1. Cambridge University Press, 1997.

- 2 T. Colcombet and C. Löding. The Non-deterministic Mostowski Hierarchy and Distance-Parity Automata. In *ICALP 2008*, volume 5126 of *Lecture Notes in Computer Science*, pages 398–409. Springer, 2008.
- 3 J. Duparc, F. Facchini, and F. Murlak. Definable Operations On Weakly Recognizable Sets of Trees. In *FSTTCS*, volume 13 of *LIPICs*, pages 363–374. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2011.
- 4 E. A. Emerson and C. S. Jutla. Tree Automata, Mu-Calculus and Determinacy (Extended Abstract). In *FOCS*, pages 368–377. IEEE Computer Society, 1991.
- 5 A. Facchini, F. Murlak, and M. Skrzypczak. Rabin-Mostowski Index Problem: A Step beyond Deterministic Automata. In *LICS*, pages 499–508. IEEE Computer Society, 2013.
- 6 E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *LNCS*. Springer, 2002.
- 7 J. M. E. Hyland. Game Semantics. In A. M. Pitts and P. Dybjer, editors, *Semantics and Logics of Computation*, volume 14 of *Publications of the Newton Institute*, page 131. Cambridge University Press, 1997.
- 8 J. M. E. Hyland. Proof theory in the abstract. *Ann. Pure Appl. Logic*, 114(1-3):43–78, 2002.
- 9 J. M. E. Hyland and A. Schalk. Abstract Games for Linear Logic. *Electr. Notes Theor. Comput. Sci.*, 29:127–150, 1999.
- 10 B. Jacobs. *Categorical Logic and Type Theory*. Studies in logic and the foundations of mathematics. Elsevier, 2001.
- 11 S. Mac Lane. *Categories for the Working Mathematician*. Springer, 2 edition, 1998.
- 12 D. A. Martin. Borel Determinacy. *The Annals of Mathematics, Second Series*, 102(2):363–371, 1975.
- 13 P.-A. Mellies. Categorical semantics of linear logic. In *Interactive models of computation and program behaviour*, volume 27 of *Panoramas et Synthèses*. SMF, 2009.
- 14 D. E. Muller and P. E. Schupp. Alternating Automata on Infinite Trees. *Theor. Comput. Sci.*, 54:267–276, 1987.
- 15 D. E. Muller and P. E. Schupp. Simulating Alternating Tree Automata by Nondeterministic Automata: New Results and New Proofs of the Theorems of Rabin, McNaughton and Safra. *Theor. Comput. Sci.*, 141(1&2):69–107, 1995.
- 16 M. Shulman. Framed bicategories and monoidal fibrations. *Theory and Applications of Categories*, 20(18):650–738, 2008.
- 17 W. Thomas. Languages, Automata, and Logic. In G. Rozenberg and A. Salomaa, editors, *Handbook of Formal Languages*, volume III, pages 389–455. Springer, 1997.
- 18 I. Walukiewicz. Monadic second-order logic on tree-like structures. *Theor. Comput. Sci.*, 275(1-2):311–346, 2002.

Multi-Focusing on Extensional Rewriting with Sums

Gabriel Scherer

Gallium, INRIA, France

Abstract

We propose a logical justification for the rewriting-based equivalence procedure for simply-typed lambda-terms with sums of Lindley [8]. It relies on maximally multi-focused proofs, a notion of canonical derivations introduced for linear logic. Lindley's rewriting closely corresponds to *preemptive rewriting* [5], a technical device used in the meta-theory of maximal multi-focus.

1998 ACM Subject Classification F.4.1 Lambda calculus and related systems

Keywords and phrases Maximal multi-focusing, extensional sums, rewriting, natural deduction

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.317

1 Introduction

Deciding observational equality of pure typed lambda-terms in presence of sum types is a difficult problem. After several solutions based on complex syntactic [6] or semantic [1, 2] techniques, Sam Lindley presented a surprisingly simple rewriting solution [8]. While the underlying intuition (extrude contexts to move pattern-matchings as high as possible in the term) makes sense, the algorithm is still mysterious in many aspects: even though they were synthesized from the previous highly-principled approach, the rewriting rules may feel strangely ad-hoc.

In this paper, we will propose a *logical* justification of this algorithm. It is based on recent developments in proof search, maximally multi-focused proofs [5]. The notion of *preemptive rewriting* was introduced in the meta-theory of multi-focusing as a purely technical device; we claim that it is in fact strongly related to Lindley's rewriting, and formally establish the correspondence.

The reference work on multi-focused systems [5] has been carried in a sequent calculus for linear logic. We will first establish the meta-theory of maximal multi-focusing for intuitionistic logic (Section 2). We start from a sequent calculus presentation, which is closest to the original system. Our first contribution is to propose an equivalent multi-focusing system in natural deduction 2.2. We then define preemptive rewriting in this natural deduction 2.4 and establish canonicity of maximally multi-focused proofs 2.6.

In Section 3, we transpose the preemptive rewrite rules into a relation on proof terms. We can then formally study the correspondence between rewriting a multi-focused proof into a canonical maximally multi-focused one, and Lindley's γ -reduction on lambda-terms. We demonstrate that they compute the same normal forms, modulo a form of redundancy elimination that is missing in the multi-focused system.

We finally introduce redundancy-elimination rewriting and equivalence for the proof terms of the multi-focused natural deduction (Section 4). The resulting notion of canonical proofs, *simplified maximal proofs*, precisely corresponds to normal forms of Lindley's rewriting relation. The natural notion of local equivalence between simplified maximal proofs therefore captures extensional equality.



© Gabriel Scherer;

licensed under Creative Commons License CC-BY

13th International Conference on Typed Lambda Calculi and Applications (TLCA'15).

Editor: Thorsten Altenkirch; pp. 317–331

Leibniz International Proceedings in Informatics



LIPIC Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

2 Intuitionistic multi-focusing

The space of proofs in sequent calculus or natural deduction exhibits a lot of redundancy: many proofs that are syntactically distinct really encode the same semantics. In particular, it is often possible to permute two inference rules in a way that preserves the validity of proofs, but also the reduction semantics of the corresponding proof terms. If a permutation transforms a proof with rule A applied above rule B into a proof with rule B applied above rule A, we say that it is an A/B permutation (A is above the slash, as in the source proof).

Focusing is a general discipline that can be imposed upon proof system, based on the separation of inference rules into two classes. *Invertible* rules (called as such because their inverse is admissible) always preserve provability, and can thus be applied as early as possible. *Non-invertible* rules may result in dead ends if they are applied too early (consider proving $A+B \vdash A+B$ by first introducing the sum on the right-hand side)). In focusing calculi, derivations are structured in “sequences” or “phases”, that either only apply invertible rules or only non-invertible rules. Focusing imposes that phases be as long as possible. During invertible phases, one must apply any valid invertible rule. During non-invertible phases, one *focuses* on a set of formulas, and applies non-invertible operations on those formulas as long as possible – if the phase is started too early, this may result in a dead end.

Invertibility determines a notion of polarity of logical connectives: we call *positive* those whose right-hand-side rule is non-invertible (they are “only interesting in positive position”), and *negative* those whose left-hand-side rule is non-invertible. In single-succedent intuitionistic logic, (\rightarrow) is negative, $(+)$ is positive, and the product (\times) may actually be assigned either polarity.

In single-sided calculi, non-invertible rules are those that introduce positive connectives, and are called “positive”. For continuity of vocabulary, we will also call non-invertible rules *positive*, and invertible rules *negative*. In particular, a permutation that moves a non-invertible rule below an invertible rule is a “pos/neg permutation”.

2.1 Multi-focused sequent calculus

Multi-focusing ([9, 5]) is an extension of focusing calculi where, instead of focusing on a single formula of the sequent (either on the left or on the right), we allow to simultaneously focus on several formulas at once. The multiple foci do not interact during the focusing phase, and this allows to express the fact that several focusing sequences are in fact independent and can be performed in parallel, condensing several distinct focused proofs into a single multi-focused derivation.

We start with a multi-focused variant of the intuitionistic sequent calculus, presented in Fig. 1. We denote focus using brackets: the rules with no brackets are invertible. This notation will change in natural deduction calculi.

In particular, we write A_n or Δ_p for formula or contexts that must be all negative or positive, and X , Y or Z for atoms. We write B_{pa} and Γ_{na} when either a positive (resp. negative) or an atom is allowed. For readability reasons, we only add polarity annotations when necessary; if we consider only derivations whose end conclusion is unfocused, then the invariant holds that the unfocused left-hand-side context is always all-negative, while the unfocused right-hand-side formula is always positive.

Our intuitionistic calculi are, as is most frequent, single-succedent. The notation $A \mid B$ on the right does not denote a real disjunction but a single formula, one of the two variables being empty. The focusing rule SEQ-FOCUS with conclusion $\Gamma, \Delta \vdash A \mid B$ can be instantiated in two ways, one when A is empty, and the premise is $\Gamma, [\Delta] \vdash [B]$ (the succedent is part of

$$\begin{array}{c}
\text{SEQ-ATOM} \\
\frac{X \text{ atomic}}{\Gamma_n, X \vdash X} \\
\\
\text{SEQ-INV-SUM-L} \\
\frac{\Gamma, A \vdash C \quad \Gamma, B \vdash C}{\Gamma, A + B \vdash C} \\
\\
\text{SEQ-INV-PROD-R} \\
\frac{\Gamma \vdash A \quad \Gamma \vdash B}{\Gamma \vdash A \times B} \\
\\
\text{SEQ-INV-ARR-R} \\
\frac{\Gamma, A \vdash B}{\Gamma \vdash A \rightarrow B} \\
\\
\text{SEQ-FOCUS} \\
\frac{\Gamma_{na}, [\Delta_n] \vdash A_{pa} \mid [B_{pa}]}{\Gamma_{na}, \Delta_n \vdash A_{pa} \mid B_{pa}} \\
\\
\text{SEQ-RELEASE} \\
\frac{\Gamma, \Delta_{pa} \vdash A \mid B_{na}}{\Gamma, [\Delta_{pa}] \vdash A \mid [B_{na}]} \\
\\
\text{SEQ-FOC-ARR-L} \\
\frac{\Gamma \vdash [A] \quad \Gamma, [\Delta, B] \vdash C \mid [D]}{\Gamma, [\Delta, A \rightarrow B] \vdash C \mid [D]} \\
\\
\text{SEQ-FOC-PROD-L} \\
\frac{\Gamma, [\Delta, A_i] \vdash B \mid [C]}{\Gamma, [\Delta, A_1 \times A_2] \vdash B \mid [C]} \\
\\
\text{SEQ-FOC-SUM-L} \\
\frac{\Gamma, [\Delta] \vdash [A_i]}{\Gamma, [\Delta] \vdash [A_1 + A_2]}
\end{array}$$

■ **Figure 1** Multifocused sequent calculus for intuitionistic logic.

the multi-focus), and one when B is empty, and the premise is $\Gamma, [\Delta] \vdash A$ (the succedent is not part of the multi-focus). Note that Δ is a set and may be empty, in which case the focus only happens on the right.

As a minor presentation difference to the reference work on multi-focusing [5], our contexts are unordered multi-sets, and all the formulas under focus are released at once – by SEQ-RELEASE, which releases positives (resp. negatives) or atoms.

This multi-focused calculus proves exactly the same formulas as the singly-focused sequent calculus. The latter is trivially included in the former, and conversely one can turn a multi-focus into an arbitrarily ordered sequence of single foci. As a corollary, relying on non-trivial proofs from the literature (e.g., [11]), it is equivalent in provability to the (non-focused) sequent calculus for intuitionistic logic.

2.2 Multi-focused natural deduction

While the multi-focusing sequent calculus closely corresponds to existing focused presentations, its natural deduction presentation in Fig. 2 is new. We took inspiration from the presentation of focused linear logic in natural deduction of [3], in particular the \uparrow and \downarrow notations coming from intercalation calculi.

There are three main judgments. $\Gamma \vdash A$ is the unfocused judgment with the invertible rules. $\Gamma; A \downarrow B$ is the “elimination-focused” judgment, and $A \uparrow B$ is the “introduction-focused” judgment (focused on A). $\Gamma; A \downarrow B$ means that the assertion B can be produced from the hypothesis A by non-invertible elimination rules; the context Γ is used in any non-focused subgoal. $A \uparrow B$ means that proving the goal A can be reduced, by applying non-invertible introduction rules, to proving the goal B . Those two judgments do not come separately, they are introduced by the focusing rule NAT-FOCUS.

In Fig. 2, we used auxiliary rules (NAT-START-INTRO, NAT-START-NO-INTRO, NAT-START-ELIM) to present the focusing compactly (this is important when rewriting proofs); those rules can only happen immediately above NAT-FOCUS, and can thus be considered definitional syntactic sugar – we used a double bar to reflect this. If we inlined these auxiliary rules, the focusing rule would read (equivalently):

$$\frac{(\overline{A_n^i})^{i \in I} \subseteq \Gamma_{na} \quad (\Gamma_{na}; A_n^i \downarrow A_{pa}^i)^{i \in I} \quad (B_p \uparrow B'_{na} \mid B = B') \quad \Gamma_{na}, (A_{pa}^i)^{i \in I} \vdash B'}{\Gamma_{na} \vdash B_{pa}}$$

This rule can only be used when all invertible rules have been performed: the context must be negative or atomic, and the goal positive or atomic. It selects set of foci on the

$$\begin{array}{c}
\text{NAT-ATOM} \\
\frac{X \text{ atomic}}{\Gamma_{na}, X \vdash X} \\
\\
\text{NAT-INV-SUM-L} \\
\frac{\Gamma, A \vdash C \quad \Gamma, B \vdash C}{\Gamma, A + B \vdash C} \\
\\
\text{NAT-INV-PROD-R} \\
\frac{\Gamma \vdash A \quad \Gamma \vdash B}{\Gamma \vdash A \times B} \\
\\
\text{NAT-INV-ARR-R} \\
\frac{\Gamma, A \vdash B}{\Gamma \vdash A \rightarrow B} \\
\\
\text{NAT-FOCUS} \\
\frac{\Gamma_{na} \Downarrow \Gamma' \quad A_{pa} \Uparrow^? A' \quad \Gamma_{na}, \Gamma' \vdash A'}{\Gamma_{na} \vdash A_{pa}} \\
\\
\text{NAT-END-ELIM} \\
\frac{}{\Gamma; A_{na} \Downarrow A_{na}} \\
\\
\text{NAT-END-INTRO} \\
\frac{}{A_{na} \Uparrow A_{na}} \\
\\
\text{NAT-ELIM-ARR} \\
\frac{\Gamma; A \Downarrow B \rightarrow C \quad B \Uparrow B' \quad \Gamma \vdash B'}{\Gamma; A \Downarrow C} \\
\\
\text{NAT-ELIM-PROD} \\
\frac{\Gamma; A \Downarrow B_1 \times B_2}{\Gamma; A \Downarrow B_i} \\
\\
\text{NAT-INTRO-SUM} \\
\frac{A_i \Uparrow B}{A_1 + A_2 \Uparrow B} \\
\\
\text{NAT-START-NO-INTRO} \\
\frac{}{A \Uparrow^? A} \\
\\
\text{NAT-START-INTRO} \\
\frac{A_p \Uparrow B_{na}}{A_p \Uparrow^? B_{na}} \\
\\
\text{NAT-START-ELIM} \\
\frac{(A_n^i)^{i \in I} \subseteq \Gamma \quad (\Gamma; A_n^i \Downarrow A_{pa}^i)^{i \in I}}{\Gamma \Downarrow (A_{pa}^i)^{i \in I}}
\end{array}$$

■ **Figure 2** Multifocused natural deduction for intuitionistic logic.

left, the family of strictly negative assumptions $(A_n^i)^{i \in I}$ (we consistently use the superscript notation for family indices), and optionally a focus on the right; if the goal is focused it must be strictly positive. All foci must be as long as possible: elimination foci go from a variable down to a positive or atomic A_{pa}^i , and the introduction focus goes up until it encounters a negative or atomic B_{na}^i .

In comparison to the sequent calculus, the positive or atomic formulas $(A_{pa}^i)^{i \in I}$ appearing at the *start* of the elimination-focus correspond to the formulas *released* at the *end* of a multi-focus in a sequent proof; natural deduction, when compared to the sequent calculus, has elimination rules “upside down”. Also characteristic of natural deduction is the horizontal parallelism between eliminations and introductions; for example, the following two partial derivations correspond to the same natural deduction:

$$\begin{array}{c}
\frac{\frac{\frac{A_{pa} \times B, A_{pa} \vdash C_{na}}{A_{pa} \times B, [A_{pa}] \vdash [C_{na}]}{A_{pa} \times B, [A_{pa}] \vdash [C_{na} + D]}}{A_{pa} \times B, [A_{pa} \times B] \vdash [C_{na} + D]}}{A_{pa} \times B \vdash C_{na} + D} \\
\\
\frac{\frac{\frac{A_{pa} \times B, A_{pa} \vdash C_{na}}{A_{pa} \times B, [A_{pa}] \vdash [C_{na}]}{A_{pa} \times B, [A_{pa} \times B] \vdash [C_{na} + D]}}{A_{pa} \times B, [A_{pa} \times B] \vdash [C_{na} + D]}}{A_{pa} \times B \vdash C_{na} + D} \\
\\
\frac{\frac{A_{pa} \times B; A_{pa} \times B \Downarrow A_{pa} \times B}{A_{pa} \times B; A_{pa} \times B \Downarrow A_{pa}}}{A_{pa} \times B \vdash C_{na} + D} \quad \frac{\frac{C_{na} \Uparrow C_{na}}{C_{na} + D \Uparrow C_{na}}}{A_{pa} \times B, A_{pa} \vdash C_{na}}
\end{array}$$

On the other hand, we kept the less important invertible rules in sequent style: the sum elimination is a left introduction. Invertible rules being morally “automatically” applied, the sequent-style left introduction, which is directed by the type of its conclusion, is more natural in this context. Ironically, this brings us rather close to the sequent calculus of Krishnaswami [7] which, for presentation purposes, preserved a function-elimination rule in natural deduction style.

► **Lemma 1.** *The multi-focused natural deduction system proves exactly the same non-focused judgments as the multi-focused sequent calculus.*

$$\begin{array}{c}
\text{PREEMPT-FOCUS} \\
\frac{\Gamma_{npa} \Downarrow \Delta'_{pa} \quad B_p \Uparrow^? B'_{na} \quad \Gamma_{npa}, \Delta'_{pa} \vdash A_{npa} \mid B'_{na}}{\Gamma_{npa} \vdash A_{npa} \mid B_p}
\end{array}
\qquad
\begin{array}{c}
\text{PREEMPT-ELIM} \\
\frac{\Gamma_{npa} \Downarrow \Delta'_{pa} \quad \Gamma_{npa}, \Delta'_{pa}; A \Downarrow A'}{\Gamma_{npa}; A \Downarrow A'}
\end{array}$$

■ **Figure 3** Preemptive rules for intuitionistic multifocused natural deduction.

2.3 A preemptive variant of multi-focused natural deduction

Multi-focusing was introduced to express the idea of *parallelism* between non-invertible rules on several independent foci. A proof has more parallelism than another if two sequential foci of the latter are merged (through rule permutations) in a single multi-focus in the former. A natural question is whether there exists “maximally parallel proofs”. To answer it (affirmatively), the original article on multi-focusing ([5]) introduced a rewriting relation that permutes non-invertible phases down in proof derivations, until they cannot go any further without losing provability – neighboring phases can then be merged into a maximally focused proof.

In the process of moving down, a non-invertible phase will traverse invertible phases below. The intermediary states of this reduction sequence may break the invariant that invertible rules must be applied as early as possible; we say that the non-invertible phase *preempts* (a part of) the invertible phase. As this intermediary state is not a valid proof in off-the-shelf multi-focusing systems, the original article introduced a relaxed variant called a preemptive system, in which the phase-sinking transformation, called *preemptive rewriting*, can be defined following [5].

We present in Fig. 3 a preemptive variant of multi-focused natural deduction, except for the invertible and focused-introduction rules that are strictly unchanged from the previous multi-focusing rules in Fig. 2. There are two important differences:

- Preemption of invertible phases. To allow the start of a focusing phase when some invertible rules could still be applied, we lifted the polarity constraints for starting focusing. In the rule PREEMPT-FOCUS, the goal $\Gamma_{npa} \vdash A_{npa}$ may be of any polarity. We use a tautological Γ_{npa} annotation to emphasize this change.
- Preemption of non-invertible phases. This is expressed by the rule PREEMPT-ELIM, where an ongoing focus on A is preempted by a complete focus on Δ'_{pa} . Note that stored contexts are *not* available during the current elimination phase (they are unused in NAT-END-ELIM); they are only available to non-focused phases that appear as subgoals (in the arrow elimination rule). This preserves the central idea that the simultaneous foci of a single focusing rule are *independent*.

2.4 Preemptive rewriting

We can then define in Fig. 4 the rewriting relation on the preemptive calculus, that lets any non-invertible phase move as far as possible down the derivation tree. Maximally multi-focused proofs, which can be characterized on permutation-equivalence classes of multi-focused proofs, correspond to normal forms of this rewriting relation.

A focused phase cannot move below an inference rule if some of the foci depend on this inference rule. Instead of expressing the non-dependency requirement by implicit absence of the foci, we have explicitly canceled out the foci that *must be absent* to improve readability. In the first rule for example, $\Gamma, \cancel{A} \Downarrow \Delta$ means that the A hypothesis must be weakened (not used) in the derivation of $\Gamma \Downarrow \Delta$, or else it cannot move below the introduction of A .

$$\begin{array}{c}
\frac{\frac{\Gamma, \cancel{A} \Downarrow \Delta \quad \cancel{B} \Uparrow \cancel{B}' \quad \Gamma, A, \Delta \vdash B}{\Gamma, A \vdash B}}{\Gamma \vdash A \rightarrow B} \quad \rightarrow \quad \frac{\frac{\Gamma, \Delta, A \vdash B}{\Gamma \Downarrow \Delta \quad \Gamma, \Delta \vdash A \rightarrow B}}{\Gamma \vdash A \rightarrow B} \\
\\
\frac{\frac{\Gamma \Downarrow \Delta \quad \cancel{A} \Uparrow \cancel{A}' \quad \Gamma, \Delta \vdash A}{\Gamma \vdash A} \quad \Gamma \vdash B}{\Gamma \vdash A \times B} \quad \rightarrow \quad \frac{\frac{\Gamma, \Delta \vdash A \quad \Gamma \vdash B}{\Gamma, \Delta \vdash A \times B}}{\Gamma \Downarrow \Delta \quad \Gamma \vdash A \times B} \\
\\
\left(\frac{\frac{\Gamma, \cancel{A} \Downarrow \Delta \quad C \Uparrow^? D \quad \Gamma, A, \Delta \vdash D}{\Gamma, A \vdash C}}{\Gamma, B \vdash C} \right) \quad \rightarrow \quad \frac{\frac{\Gamma, A, \Delta \vdash D \quad \Gamma, B, \Delta \vdash D}{\Gamma, A + B, \Delta \vdash D}}{\Gamma \Downarrow \Delta \quad C \Uparrow^? D} \\
\left(\frac{\frac{\Gamma, \cancel{B} \Downarrow \Delta \quad C \Uparrow^? D \quad \Gamma, B, \Delta \vdash D}{\Gamma, B \vdash C}}{\Gamma, A + B \vdash C} \right) \quad \rightarrow \quad \frac{\Gamma, A + B \vdash C, D}{\Gamma, A + B \vdash C, D} \\
\\
\frac{\frac{\Gamma \Downarrow \Gamma' \quad \Gamma, \Gamma'; A \Downarrow B_1 \times B_2}{\Gamma; A \Downarrow B_1 \times B_2}}{\Gamma, \Gamma'; A \Downarrow B_i} \quad \rightarrow \quad \frac{\frac{\Gamma, \Gamma'; A \Downarrow B_1 \times B_2}{\Gamma, \Gamma'; A \Downarrow B_i}}{\Gamma \Downarrow \Gamma' \quad \Gamma; A \Downarrow B_i} \\
\\
\frac{\frac{\Gamma \Downarrow \Gamma' \quad \Gamma, \Gamma'; A \Downarrow B \rightarrow C}{\Gamma; A \Downarrow B \rightarrow C} \quad B \Uparrow B' \quad \Gamma \vdash B'}{\Gamma; A \Downarrow C} \quad \rightarrow \quad \frac{\frac{\Gamma, \Gamma'; A \Downarrow B \rightarrow C \quad B \Uparrow B' \quad \Gamma \vdash B'}{\Gamma, \Gamma' \Downarrow B \rightarrow C}}{\Gamma \Downarrow \Gamma' \quad \Gamma; A \Downarrow C} \\
\\
\frac{\Gamma; A \Downarrow B \rightarrow C \quad B \Uparrow B'_{na} \quad \frac{\Gamma \Downarrow \Gamma' \quad \Gamma, \Gamma' \vdash B'_{na}}{\Gamma \vdash B'_{na}}}{\Gamma; A \Downarrow C} \quad \rightarrow \quad \frac{\frac{\Gamma; A \Downarrow B \rightarrow C \quad B \Uparrow B'_{na} \quad \Gamma, \Gamma' \vdash B'_{na}}{\Gamma, \Gamma'; A \Downarrow C}}{\Gamma \Downarrow \Gamma' \quad \Gamma; A \Downarrow C} \\
\\
\frac{\frac{\frac{\frac{\Gamma \Downarrow \Gamma' \quad A_n \in \Gamma \quad \frac{\Gamma \Downarrow \Delta \quad \Gamma, \Delta; A_n \Downarrow A'}{\Gamma; A_n \Downarrow A'}}{\Gamma \Downarrow \Gamma', A'}}{\Gamma \Downarrow \Gamma', A'} \quad B \Uparrow^? B' \quad \Gamma, \Gamma', A' \vdash B'}{\Gamma \vdash B} \\
\rightarrow \quad \frac{\frac{\frac{\frac{\Gamma \Downarrow \Gamma' \quad A_n \in \Gamma \quad \Gamma, \Delta; A_n \Downarrow A'}{\Gamma, \Delta \Downarrow \Gamma', A'}}{\Gamma, \Delta \Downarrow \Gamma', A'} \quad B \Uparrow^? B' \quad \Gamma, \Gamma', A' \vdash B'}{\Gamma \Downarrow \Delta \quad \Gamma, \Delta \vdash B}}{\Gamma \vdash B} \\
\\
\frac{\Gamma \Downarrow \Delta \quad A \Uparrow^? B \quad \frac{\Gamma \Downarrow \Delta' \quad B \Uparrow^? C \quad \Gamma, \Delta, \Delta' \vdash C}{\Gamma, \Delta \vdash B}}{\Gamma \vdash A} \quad \leftrightarrow \quad \frac{\Gamma \Downarrow \Delta, \Delta' \quad A \Uparrow^? C \quad \Gamma, \Delta, \Delta' \vdash C}{\Gamma \vdash A}
\end{array}$$

■ **Figure 4** Preemptive rewriting for multifocused natural deduction for intuitionistic logic.

In this situation, it may be the case that other parts of the multi-focus do not depend on the rule below, and those should not be blocked. To allow rewriting to continue, the last rewrite of our system is bidirectional. It allows to separate the foci of a multi-focus, in particular separate the foci that depend on the rule below from those that do not – and can thus permute again. This corresponds to the first rule of the original preemptive rewriting system [5], which splits a multi-focus in two. We only need to apply this rule when the result can make one more unidirectional rewrite step – this strategy ensures termination.

In the left-to-right direction, this rule relies on the possibility of merging together two elimination-focused derivations, or two optional introduction-focused derivations, with the implicit requirement that at least one of them is empty.

2.5 Reinversion

After the preemptive rewriting rules have been applied, the result is not, in general, a valid derivation in the non-preemptive system. Consider for example the following rewriting process:

$$\begin{bmatrix} \nu_3 \\ \pi_3 \\ \nu_2 \\ \pi_2 \\ \nu_1 \\ \pi_1 \end{bmatrix} \rightarrow^* \begin{bmatrix} \nu_3 \\ \nu_2 & \pi_3; \nu_2 \\ \pi_2 \\ \nu_1 \\ \pi_1 \end{bmatrix} \rightarrow^* \begin{bmatrix} \nu_2 & \nu_3 \\ \pi_2 & \pi_3 \\ \nu_1 \\ \pi_1 \end{bmatrix} \rightarrow^* \begin{bmatrix} \nu_2 & \nu_3 \\ \pi_2; \nu_3 & \\ \nu_1 & \pi_3; \nu_1 \\ \pi_1 \end{bmatrix} \rightarrow^* \begin{bmatrix} \nu_2 & \nu_3 \\ \pi_2; \nu_3 & \\ \nu_1 & \\ \pi_1 & \pi_3 \end{bmatrix}$$

We are here representing derivations from a high-level point of view, by naming complete sequences of rules of the same polarity. Sequences of positive (non-invertible) are named π_n , and sequences of negative (invertible) rules ν_m . We use horizontal position to denote parallelism, or dependencies between phases: each dipole (π_k, ν_k) is vertically aligned as the invertibles of ν_k have been produced by the foci of π_k , but we furthermore assume that the second dipole depends on formulas released by the first, while the third dipole is independent.

The third dipole is independent from the others, and its foci in π_3 move downward in the derivation as expected in the preemptive system. After the first step, its negative phase has preempted the invertible phase ν_2 , and it is thus written $\pi_3; \nu_2$ to emphasize that any rule of this sequence will have all the invertible formulas of ν_2 in non-focused positions (positives in the hypotheses, and negatives in the succedent). It can then be merged with the foci of π_2 , in which case it does not see the invertibles of ν_2 anymore. When it moves further down, the invertible formulas in its topmost sequent, those consumed by ν_3 , are present/preempted by all the non-invertible rules of π_2 . It is eventually merged with π_1 .

The normal form of this rewrite sequence could be considered a maximally multi-focused proof, in the sense that the foci happen as soon as possible in the derivation – which was not the case in the initial proof, where π_3 was delayed. However, while the initial proof is a valid proof in the non-preemptive system, the last derivation is not: the invertible formulas produced by π_3 are not consumed as early as possible, but only at the very end of the derivation, and the foci of π_2 therefore happen while there are still invertible rules to be applied.

We introduce a *reinversion* relation between proofs, written $\mathcal{D} \triangleright \mathcal{E}$, that turns the proof \mathcal{D} with possible preemption into a proof \mathcal{E} valid in the non-preemptive system, by doing the inversions where they are required, without changing the structure of the negative phases –

the foci are exactly the same. In our example, we have:

$$\left[\begin{array}{cc} \nu_2 & \nu_3 \\ \pi_2; \nu_3 & \\ \nu_1 & \\ \pi_1 & \pi_3 \end{array} \right] \triangleright \left[\begin{array}{cc} \nu_2 & \\ \pi_2 & \\ \nu_1 & \nu_3 \\ \pi_1 & \pi_3 \end{array} \right]$$

► **Definition 2** (Rewriting relation). If \mathcal{D} and \mathcal{E} are proofs of the non-preemptive system, we write $\mathcal{D} \Rightarrow \mathcal{E}$ if there exists a \mathcal{E}' such that $\mathcal{D} \rightarrow^* \mathcal{E}' \triangleright \mathcal{E}$.

Reinversion was not discussed directly in the original multi-focusing work [5], but it plays an important role and can be described and understood in several fairly different ways. For lack of space, we omit this discussion from this short article, and will only formally define reinversion as a relation on the (more concise) proof terms in Section 3.1, Definition 4.

2.6 Maximal multi-focusing and canonicity

Now that we have defined the focusing-lowering rewrite (\Rightarrow) between non-preemptive proof, we can define the notion of *maximal* multi-focusing and its meta-theory. It is defined by looking at the width of multi-focus phases in equivalence classes of rule permutations; but it can also be characterized as the normal forms of the (\Rightarrow) relation.

For lack of space, we have omitted this development (which is a mere adaptation of the previous work [5]) from this short article. The central result is summarized below.

► **Definition.** We say that two proofs \mathcal{D} and \mathcal{E} are locally equivalent, or iso-polar, written $\mathcal{D} \approx_{loc} \mathcal{E}$, if one can be rewritten into the other using only local positive/positive and negative/negative permutations, preserving their initial sequents.

► **Definition.** We say that two proofs \mathcal{D} and \mathcal{E} are globally equivalent, or iso-initial, written $\mathcal{D} \simeq_{glob} \mathcal{E}$, when one can be rewritten into the other using local permutations of any polarity (so when seen as proofs in a non-focused system), preserving their initial sequents.

► **Fact.** Two multi-focused proofs are globally equivalent if and only if they are rewritten by (\Rightarrow) in locally equivalent maximal proofs.

3 On the side of proof terms

3.1 Preemption and reinversion as term rewriting

Now that we have a notion of maximally multi-focused proofs in natural deduction, we can cross the second bridge between multi-focusing and Lindley's work by moving to a term system. We define in Figure 5 a term syntax for multi-focused derivations in natural deduction.

As the distinction between the preemptive and the non-preemptive systems are mostly about invariants of the focusing rule, the same term calculus is applicable to both. The only syntactic difference is that preemptive terms allow a multi-focusing $f[n]$ to preempt an ambient elimination focus n' .

Structural constraints on the multi-focusing system (preemptive or not) guarantee that strong typing invariants are verified. In particular, in a focused term ($\text{let } \bar{x} = \bar{n} \text{ in } p^?t$), the \bar{n} are typed by the formulas in Δ at the end of a $\Gamma \Downarrow \Delta$ elimination phase: by our release discipline they have a positive or atomic type, so the let -introduced \bar{x} are always bound to positive types. The rewriting rules corresponding to the preemptive rewriting relation are defined in Figure 6.

$$\begin{array}{l}
t ::= \\
| x, y, z \quad \text{variable} \\
| \lambda(x) t \quad \text{lambda} \\
| (t, t) \quad \text{pair} \\
| \delta(x, x.t, x.t) \quad \text{case} \\
| f[t] \quad \text{focusing}
\end{array}
\quad
\begin{array}{l}
\text{terms} \\
\text{variable} \\
\text{lambda} \\
\text{pair} \\
\text{case} \\
\text{focusing}
\end{array}
\quad
\frac{X \text{ atomic}}{\Gamma_{na}, x : X \vdash x : X}$$

$$\frac{\Gamma, x : A \vdash t : C \quad \Gamma, x : B \vdash u : C}{\Gamma, x : A + B \vdash \delta(x, x.t, x.u) : C} \quad
\frac{\Gamma \vdash t : A \quad \Gamma \vdash u : B}{\Gamma \vdash (t, u) : A \times B} \quad
\frac{\Gamma, x : A \vdash t : B}{\Gamma \vdash \lambda(x) t : A \rightarrow B}$$

$$\begin{array}{l}
f[\square] ::= \text{let } \bar{x} = \bar{n} \text{ in } p^? \square \quad \text{multi-focusing} \\
p^? ::= \text{optional introduction focus} \\
| \emptyset \quad \text{no introduction} \\
| p \quad \text{introduction focus}
\end{array}$$

$$\frac{\Gamma_{na} \Downarrow \text{let } \bar{x} = \bar{n} : \Gamma' \quad A_{pa} \Uparrow^? p^? : A' \quad \Gamma_{na}, \Gamma' \vdash t : A'}{\Gamma_{na} \vdash \text{let } \bar{x} = \bar{n} \text{ in } p^? t : A_{pa}}$$

$$\frac{A \Uparrow^? \emptyset : A \quad A_p \Uparrow^? p : B_{na} \quad (\Gamma; x^i : A_n^i)^{i \in I} \subseteq \Gamma \quad (\Gamma; x^i : A_n^i \Downarrow n^i : A_{pa}^i)^{i \in I}}{\Gamma \Downarrow \text{let } (x^i)^{i \in I} = (n^i)^{i \in I} : (A_{pa}^i)^{i \in I}}$$

$$\begin{array}{l}
n ::= \\
| x, y, z \quad \text{variable} \\
| \pi_i n \quad \text{pair projection} \\
| n p(t) \quad \text{function application} \\
| \text{let } \bar{x} = \bar{n} \text{ in } n \quad \text{focusing (only in the preemptive calculus)}
\end{array}
\quad
\frac{}{\Gamma; x : A_{na} \Downarrow x : A_{na}}$$

$$\frac{\Gamma; A \Downarrow n : B \rightarrow C \quad B \Uparrow p : B' \quad \Gamma \vdash t : B'}{\Gamma; A \Downarrow n p(t) : C} \quad
\frac{\Gamma; A \Downarrow n : B_1 \times B_2}{\Gamma; A \Downarrow \pi_i n : B_i} \quad i \in \{1, 2\}$$

$$\frac{\Gamma_{npa} \Downarrow \text{let } (x^i)^{i \in I} = (n^i)^{i \in I} : \Delta'_{pa} \quad \Gamma_{npa}, \Delta'_{pa}; A \Downarrow n' : A'}{\Gamma_{npa}; A \Downarrow \text{let } (x^i)^{i \in I} = (n^i)^{i \in I} \text{ in } n' : A'}$$

$$\begin{array}{l}
p ::= \\
| \star \quad \text{identity} \\
| \sigma_i p \quad \text{sum injection}
\end{array}
\quad
\frac{}{A_{na} \Uparrow \star : A_{na}} \quad
\frac{A_i \Uparrow p : B}{A_1 + A_2 \Uparrow \sigma_i p B} \quad i \in \{1, 2\}$$

■ **Figure 5** Preemptive term calculus.

► **Lemma 3.** *If t is a proof term for the preemptive derivation \mathcal{D} , then $t \rightarrow u$ if and only if u is a proof term for a preemptive derivation \mathcal{E} with $\mathcal{D} \rightarrow \mathcal{E}$.*

The reinversion relation also has a corresponding term-rewriting interpretation. To perform each invertible rule as early as it should be, it suffices to let any invertible rule skip over a non-invertible phase it does not depend on. Depending on the order of the invertible rules after this phase, the invertible rule we want to move may be after a series of invertible rules that cannot be moved.

We “skip” over invertible contexts, we reduce invertible rules happening inside contexts of

$$\begin{array}{lcl}
\lambda(y) \text{ let } \bar{x} = \bar{n} \text{ in } t & \xrightarrow{y \notin \bar{n}} & \text{let } \bar{x} = \bar{n} \text{ in } \lambda(y) t \\
((\text{let } \bar{x} = \bar{n} \text{ in } t_1), t_2) & \rightarrow & \text{let } \bar{x} = \bar{n} \text{ in } (t_1, t_2) \\
(t_1, (\text{let } \bar{x} = \bar{n} \text{ in } t_2)) & \rightarrow & \text{let } \bar{x} = \bar{n} \text{ in } (t_1, t_2) \\
\delta(y, y_1. (\text{let } \bar{x} = \bar{n} \text{ in } p^? t_1), y_2. (\text{let } \bar{x} = \bar{n} \text{ in } p^? t_2)) & \xrightarrow{y_1, y_2 \notin \bar{n}} & \text{let } \bar{x} = \bar{n} \text{ in } p^? \delta(y, y_1.t_1, y_2.t_2) \\
\pi_i (\text{let } \bar{x} = \bar{n} \text{ in } n') & \rightarrow & \text{let } \bar{x} = \bar{n} \text{ in } \pi_i n' \\
(\text{let } \bar{x} = \bar{n} \text{ in } n') t & \rightarrow & \text{let } \bar{x} = \bar{n} \text{ in } n' t \\
n' p(\text{let } \bar{x} = \bar{n} \text{ in } t) & \rightarrow & \text{let } \bar{x} = \bar{n} \text{ in } n' p(t) \\
\text{let } y = (\text{let } \bar{x} = \bar{n} \text{ in } n') \text{ in } p^? t & \rightarrow & \text{let } \bar{x} = \bar{n} \text{ in } \text{let } y = n' \text{ in } p^? t \\
\text{let } \bar{x} = \bar{n} \text{ in } p^? (\text{let } \bar{y} = \bar{n}' \text{ in } q^? t) & \xleftrightarrow{\bar{x} \notin \bar{n}'} & \text{let } \bar{x}, \bar{y} = \bar{n}, \bar{n}' \text{ in } (p^?.q^?)t \\
\\
p^?.\emptyset = p^? & \star.q = q & \\
\emptyset.q^? = q^? & (\sigma_i p).q = \sigma_i (p.q) &
\end{array}$$

■ **Figure 6** Preemptive rewriting on proof terms.

$$\begin{array}{lcl}
F_i[\square] ::= \lambda(x) \square & & C_{neg}[\square] ::= n p(\square) \\
| \delta(x, x_1.\square, x_2.t) & & | C_{neg}[\square] p(t) \\
| \delta(x, x_1.t, x_2.\square) & & | \pi_i C_{neg}[\square] \\
| (t, \square) & & | C_{ni}[C_{neg}[\square]] \\
| (\square, t) & & \\
\\
C_i[\square] ::= \square \mid F_i[C_i[\square]] & & C_{ni}[\square] ::= \text{let } \bar{x} = \bar{n} \text{ in } p^? \square \\
& & | \text{let } \bar{x}, y = \bar{n}, C_{neg}[\square] \text{ in } p^? t
\end{array}$$

■ **Figure 7** Invertible frames and contexts, non-invertible contexts and elimination contexts.

the form $C_{ni}[C_i[\]]$, where $C_i[t]$ is a notation for invertible contexts (defined using invertible frames $F_i[t]$), and $C_{ni}[t]$ for non-invertible contexts. Defining the latter requires describing negative/elimination contexts $C_{neg}[t]$, with holes where a term may appear in a series of elimination-focused terms.

► **Definition 4.** Reversion can be precisely defined as the transitive congruence closure of the rewrite rules listed in Figure 8.

The rewrite conditions are expressed in terms of a $C[\square] \prec c$ relation (read “context C blocks term-constructor c ”) that indicates a dependency of an invertible construction c on a given context $C[\square]$. For example, it would make no sense to extrude a λ in argument position in a destructor, or move a sum-elimination $\delta(x)$ across the frame that defined the variable x . This blocking relation is defined in Figure 9 – $(A \mid B)$ in a rule means that the rule holds with either A or B in place of $(A \mid B)$.

It may at first seem surprising that reversion rules have instances that are the *opposite* of some of the preemptive rewriting rules – those about pos/neg permutations. But that is precisely one of the purposes of reversion: after preemptive rewriting rules have been fully applied, we undo those that have gone “too far”, in the sense that they let a non-invertible phase preempt a portion of an invertible phase below, but were blocked by dependencies without reaching the next non-invertible phase. This blocked phase does not increase the parallelism of multi-focusing in the proof, but stops the derivation from being valid in the original multi-focusing system, so reversion undoes its preemption.

$$\begin{aligned}
& C_{ni}[C_i[\lambda(x) t]] \stackrel{C_{ni}[C_i[\square]] \not\prec \lambda}{\triangleright} \lambda(x) C_{ni}[C_i[t]] \\
& C_{ni}[C_i[(t_1, t_2)]] \stackrel{C_{ni}[C_i[\square]] \not\prec (\cdot, \cdot)}{\triangleright} (C_{ni}[C_i[t_1]], C_{ni}[C_i[t_2]]) \\
& C_{ni}[C_i[\delta(x, x_1.t_1, x_2.t_2)]] \stackrel{C_{ni}[C_i[\square]] \not\prec \delta(x)}{\triangleright} \delta(x, x_1.C_{ni}[C_i[t_1]], x_2.C_{ni}[C_i[t_2]])
\end{aligned}$$

■ **Figure 8** Reversion rewrite rules.

$$\begin{aligned}
c & ::= (\cdot) \mid \lambda \mid \delta(x) \\
\frac{y \in \bar{x}}{\text{let } \bar{x} = \bar{n} \text{ in } p^? \square \prec \delta(y)} & \quad \frac{p \neq \emptyset}{\text{let } \bar{x} = \bar{n} \text{ in } p \square \prec (\cdot) \mid \lambda} & \quad \frac{C_{neg}[\square] \prec c}{\text{let } \bar{x}, y = \bar{n}, C_{neg}[\square] \text{ in } p^? t \prec c} \\
((\square, t) \mid (t, \square) \mid \lambda(x) \square) \prec ((\cdot) \mid \lambda) & \quad \lambda(x) \square \prec \delta(x) & \quad (\delta(x, y. \square, z.t) \mid \delta(x, z.t, y. \square)) \prec \delta(y) \\
n p(\square) \prec (\cdot) \mid \lambda & \quad \frac{C_{neg}[\square] \prec c}{C_{neg}[\square] p(t) \mid \pi_i C_{neg}[\square] \mid C_{ni}[C_{neg}[\square]] \prec c} \\
\frac{C_{ni}[\square] \prec c \mid C_{neg}[\square] \prec c}{C_{ni}[C_{neg}[\square]] \prec c} & \quad \frac{F_i[\square] \prec c \mid C_i[\square] \prec c}{F_i[C_i[\square]] \prec c} & \quad \frac{C_{ni}[\square] \prec c \mid C_i[\square] \prec c}{C_{ni}[C_i[\square]] \prec c}
\end{aligned}$$

■ **Figure 9** Reversion blocking relation.

Remark, in relation to this situation, that preemptive rewriting cannot be easily defined on equivalence classes of neg/neg permutations (or other presentations of focusing that crush the invertible phase in one not-so-interesting step, such as higher-order focusing), as the order of the invertible rules in a single phase may determine where a non-invertible phase stops its preemption and is blocked in the middle of the invertible phase. Reversion restores this independence on invertible ordering. This explains why the meta-theory of maximal multi-focusing was conducted in the non-preemptive system, using the relation between proofs that always applies reversion after preemptive rewriting.

The other interesting case is a non-invertible phase π_0 having traversed a family of non-invertible phases $(\pi'_i)_{i \in I}$, before merging into some non-invertible phase π_1 . Reversion will move its negative phase ν_0 , reverting the preemption of the (π'_i) on the invertible formulas introduced by π_0 . But the important preemptions that happened, namely the traversal by π_0 of each of the invertible phases $(\nu'_i)_{i \in I}$, are not reverted: each ν'_i is blocked by the π'_i below and thus cannot be reverted below π_0 . As π_0 traversed *both* the ν'_i and the π'_i , it does not have the corresponding invertible formulas in its context anymore, and is well-positioned even in a non-preemptive proof.

► **Lemma 5.** *If t is the proof term of the preemptive derivation $\mathcal{D} : \Gamma \vdash A$, and u is such that $t \triangleright u$, then u is a valid (preemptive) proof term for $\Gamma \vdash A$.*

► **Lemma 6.** *If u is a valid proof term in the preemptive system, and a normal form of the relation (\triangleright) , then u is also a valid proof term for the non-preemptive system.*

► **Theorem 7.** *If t is a proof term for \mathcal{D} and u for \mathcal{E} , then $\mathcal{D} \Rightarrow \mathcal{E}$ if and only there is a u' such that $t \rightarrow^* u' \triangleright u$, and u is a normal form for (\triangleright) .*

3.2 Multi-focused terms as lambda-terms

There is a natural embedding $[t]$ of a multi-focused term t into the standard lambda-calculus, generated by the following transformation, where $t[\bar{x} := \bar{u}]$ represents simultaneous substitution:

$$\begin{aligned} [\text{let } \bar{x} = \bar{n} \text{ in } p^\circ t] &:= [p^\circ]([t][\bar{x} := [\bar{n}]])) \\ [\emptyset](t) &:= t & [\star](t) &:= t & [\sigma_i p](t) &:= \sigma_i [p](t) \end{aligned}$$

The substitutions break the invariant that the scrutinee of a sum-elimination construct is always a variable. However, as only negative terms are substituted, sum-elimination scrutinees are always neutrals – embedding of negative terms. In particular, this embedding does not create any β -redex. Proof terms coming from non-preemptive multi-focusing are also always in η -long form, and this is preserved by the embedding; with the restriction present in Lindley’s work that only neutral terms (eliminations) are expanded – this avoids issues of commuting conversions. We mean here the *weak* η -long form, determined by the weak equation $(m : A + B) =_{\text{weak-}\eta} \delta(m, x_1.\sigma_1 \ x_1, x_2.\sigma_2 \ x_2)$.

► **Lemma 8.** *If $\Gamma \vdash t : A$ in the preemptive multi-focused system, then $\Gamma \vdash [t] : A$ in simply-typed lambda-calculus, and $[t]$ is in β -normal form. If t is valid in the non-preemptive system, then the pure neutral subterms of $[t]$ are also in weak η -long form.*

3.3 Lindley’s rewriting relation

The strong η -equivalence for sums makes lambda-term equivalence a difficult notion. For any term $m : A + B$ and well-typed context $C[\square]$, it dictates that $C[m] \approx \delta(m, x_1.C[x_1], x_2.C[x_2])$. In his article [8], Sam Lindley breaks it down in four simpler equations, including in particular the “weak”, non-local η -rule (where F represents a frame, that is a context of term-size exactly 1):

$$\begin{aligned} m &\approx \delta(m, x_1.\sigma_1 \ x_1, x_2.\sigma_2 \ x_2) \quad (+.\eta) \\ F[\delta(p, x_1.t_1, x_2.t_2)] &\approx \delta(p, x_1.F[t_1], x_2.F[t_2]) \quad (\text{move-case}) \\ \delta \left(p, \begin{array}{l} x_1.\delta(p, y_1.t_1, y_2.t_2), \\ x_2.\delta(p, z_1.u_1, z_2.u_2) \end{array} \right) &\approx \delta(p, x_1.t_1[y_1 := x_1], x_2.u_2[z_2 := x_2]) \quad (\text{repeated-guard}) \\ \delta(p, x_1.t, x_2.t) &\stackrel{x_1, x_2 \notin t}{\approx} t \quad (\text{redundant-guard}) \end{aligned}$$

Lindley further refines the *move-case* equivalence into a less-local *hoist-case* rule. Writing D for a frame that is either $\delta(p, x_1.\square, x_2.t)$ or $\delta(p, x_1.t, x_2.\square)$, D^* for an arbitrary (possibly empty) sequence of them, and H any frame that is *not* of this form, *hoist-case* is defined as:

$$H[D^*[\delta(t, x_1.t_1, x_2.t_2)]] \rightarrow \delta(t, x_1.H[D^*[t_1]], x_2.H[D^*[t_2]])$$

Lindley’s equivalence algorithm (Theorem 36, p. 13) proceeds in three steps: rewriting terms in $\beta\eta\gamma_E$ -normal forms (using the weak $(+.\eta)$ on sums), then rewriting them in γ -normal form, and finally using a decidable redundancy-eliminating equivalence relation called \sim . The rewriting relation γ is defined as the closure of *repeated-guard*, *redundant-guard* (when read left-to-right) and *hoist-case*; γ_E is a weak restriction of it defined below. The equivalence \sim is the equivalence closure of the equivalence *repeated-guard*, *redundant-guard*, and *move-case* restricted to D -frames – clauses of a sum elimination.

We discuss redundancy elimination, that is aspects related to *repeated-guard* and *redundant guard*, in Section 4, and focus here on explanation of the other rewriting processes ($\beta\eta\gamma_E$ and *hoist-case*) in logical terms. We show that multi-focused terms in (\Rightarrow) -normal form embed into $\beta\eta\gamma_E\gamma$ -normal forms. As we ignore redundancy elimination, this is modulo \sim .

The β and η rewriting rules are standard – for sums, this is the weak, local η -relation, and not the strong η -equivalence. As explained in the previous subsection, embeddings of proof terms valid in the non-preemptive system – as are (\Rightarrow) -normal forms – are in $\beta\eta$ -normal form. The rewriting γ_E is defined as the extrusion of a sum-elimination out of an elimination context: $\square t \mid \pi_i \square \mid \delta(\square, x_1.t, x_2.t)$.

► **Lemma 9.** *Terms for valid preemptive multi-focusing derivations are in γ_E -normal form.*

This rigid structure of focused proofs is well-known, just as $\beta\eta$ -normality or commuting conversions are not the interesting points of Lindley’s work. The crux of the correspondence is between the transformation to maximal proofs, computed by (\Rightarrow) , and his γ -rewriting relation. There is an interesting dichotomy:

- Preemptive rewriting, which merges non-invertible phases, is where most of the work happens from a logical point of view. Yet this transformation, on the embeddings of the multi-focused proof terms, corresponds to the identity!
- Reinverson, which is obvious logically as it only concerns invertible rules which commute easily, corresponds to γ -rewriting on the embeddings.

Of course, preemptive rewriting is in fact crucial for γ -rewriting. It is the one that determines upto where negative terms can move in the derivation, and in particular the scrutinees of sum eliminations. Reinverson would not work without the first preemptive rewriting step, and applying reinverson on a proof term that is not in preemptive-normal form may not give a γ -normal embedding. Note that the proof of the last theorem in this section makes essential use of the confluence of γ -rewriting, one of Lindley’s key results.

► **Lemma 10.** *If $t \rightarrow u$, then $\lfloor t \rfloor =_\alpha \lfloor u \rfloor$.*

► **Lemma 11.** *If $t \triangleright u$, then $\lfloor t \rfloor \rightarrow_\gamma^* \lfloor u \rfloor$.*

► **Lemma 12.** *If u is in (\Rightarrow) -normal form, then for some $u' \approx_{loc} u$, $\lfloor u' \rfloor$ is in γ -normal form modulo \sim .*

► **Theorem 13** (γ -normal forms are embeddings of maximally-focused proofs). *If $\lfloor t \rfloor \rightarrow_\gamma^* n$ and n is γ -normal, then there are $u \approx_{loc} u'$ such that $t \Rightarrow u$ and $\lfloor u' \rfloor \sim n$. In particular, u is maximally multi-focused.*

4 Redundancy elimination

In the previous section, we have glossed over the fact that Lindley’s γ -reduction also simplifies redundant and duplicated sum-eliminations. Those simplifications are *not* implied by multi-focusing – they are not justified by proof theory alone. Our understanding is that they correspond to purity assumptions that are stronger than the natural equational theory of focused proofs. On the other hand, starting from maximally multi-focused forms is essential to being able to define those extra simplifications. We do so in this section, to obtain a system that is completely equivalent to Lindley’s.

We simply have to add the following simplifications on proof terms:

$$\begin{array}{ll}
\text{REDUNDANT-FOCUS} & \text{REDUNDANT-GUARD} \\
\text{let } \bar{x}, y, z = \bar{n}, n', n' \text{ in } p^? t \rightarrow_s \text{let } \bar{x}, y = \bar{n}, n' \text{ in } p^? t[z := y] & \delta(x, x_1.t, x_2.t) \stackrel{x_1, x_2 \notin t}{\approx_{\text{loc}}} t \\
\\
\text{REPEATED-CASE-1} & \\
\delta(x, x_1.\delta(x, y_1.u_1, y_2.u_2), x_2.t_2) \approx_{\text{loc}} \delta(x, x_1.u_1[y_1 := x_1], x_2.t_2) & \\
\\
\text{REPEATED-CASE-2} & \\
\delta(x, x_1.t_1, x_2.\delta(x, y_1.u_1, y_2.u_2)) \approx_{\text{loc}} \delta(x, x_1.y_1, x_2.u_2[y_2 := x_2]) &
\end{array}$$

While those rules are not implied by focusing, they are reasonable in a focused setting, as they respect the phase separation. As the redundancy-elimination rules test for equality of subterms, they have an unpleasant non-atomic aspect (repeated cases only test variables), but this seems unavoidable to handle sum equivalence (Lindley [8], or Balat, Di Cosmo and Fiore [2], have a similar test in their normal form judgments), and have also been used previously in the multi-focusing literature, for other purposes; in Alexis Saurin's PhD thesis [10], an equality test is used to give a convenient $\otimes/\&$ permutation rule (p. 231).

► **Definition 14.** We define the relation $t \Rightarrow_s u$ between proof terms of the (preemptive) multi-focusing calculus as follows, where t_1 is a preemptive normal form, t_2 is a redundant-foci normal form, and u_0 is a (\triangleright)-normal form: $t \rightarrow^* t_1 \rightarrow_s^* t_2 \triangleright^* u_0 \approx_{\text{loc}} u$

► **Definition 15.** We call the u in the target of the (\Rightarrow_s) relation *simplified maximal forms*.

► **Theorem 16** (Simplified maximal forms are γ -normal). *Given a multi-focused term t , there exists some u such that $t \Rightarrow_s u$, $[t] \rightarrow_\gamma^* [u]$, and $[u]$ is in γ -normal form. This u is unique modulo local equivalence.*

► **Corollary 17.** *Two multi-focused proof terms are extensionally equivalent if their maximally multi-focused normal forms are locally equivalent (modulo redundancy elimination).*

Related and Future work

Maximally multi-focused proofs were previously used to bridge the gap between sequent calculus, as a rather versatile way of defining proof systems, and specialized proof structures designed to minimize redundancy for a fixed logic. The original paper on multi-focusing [5] demonstrated an isomorphism between maximal proofs and proof nets for a subset of linear logic. In recent work [4], maximally multi-focused proof of a sequent calculus for first-order logic have been shown isomorphic to *expansion proofs*, a compact description of first-order classical proofs.

There are some recognized design choices in the land of equivalence-checking presentation that can now be linked to design choices of focused system. For example, Altenkirch et al. [1] proposed to make the syntax more canonical with respect to redundancy-elimination by using a n -ary sum elimination construct, while Lindley prefers to quotient over local reorderings of unary sum-eliminations. This sounds similar to the choice between higher-order focusing ([12]), where all invertible rules are applied at once, or quotienting of concrete proofs by neg/neg permutations as used here.

When we started this work, we planned to also study the proof-term presentation of preemptive rewriting, in a term language for sequent calculus. We have been collaborating with Guillaume Munch-Maccagnoni to study the normal forms of an intuitionistic restriction of System L, with sums. In this untyped calculus, syntactic phases appear that closely

resemble a focusing discipline, and equivalence relations can be defined in a more uniform way, thanks to the symmetric status of the (non)-invertible rules that “change the type of the result” (terms, values) and those that only manipulate the environments (co-terms, stacks).

Conclusion

We propose a multi-focused calculus for intuitionistic logic in natural deduction, and establish the canonicity of maximally multi-focused proofs by transposing the preemptive rewriting technique [5] in our intuitionistic, natural deduction setting. By studying the computational effect of preemptive rewriting on proof terms, we demonstrate the close correspondence with the rewriting on lambda-terms with sums proposed by Lindley [8] to compute extensional equivalence. Adding a notion of redundancy elimination to our multi-focused system makes preemptive rewriting precisely equivalent to Lindley’s γ -rules. In particular, the resulting canonical forms, *simplified maximal proofs*, capture extensional equality.

Acknowledgements. The author thanks his advisor, Didier Rémy for the freedom of making an overly long detour through the proof search literature; Alexis Saurin for his advice, and Pierre-Évariste Dagand, Sam Lindley, and anonymous reviewers for their helpful comments.

References

- 1 Thorsten Altenkirch, Peter Dybjer, Martin Hofmann, and Philip J. Scott. Normalization by evaluation for typed lambda calculus with coproducts. In *LICS*, 2001.
- 2 Vincent Balat, Roberto Di Cosmo, and Marcelo P. Fiore. Extensional normalisation and type-directed partial evaluation for typed lambda calculus with sums. In *POPL*, 2004.
- 3 Taus Brock-Nannestad and Carsten Schürmann. Focused natural deduction. In *LPAR (Yogyakarta)*, 2010.
- 4 Kaustuv Chaudhuri, Stefan Hetzl, and Dale Miller. A Systematic Approach to Canonicity in the Classical Sequent Calculus. In *21st EACSL Annual Conference on Computer Science Logic*, volume 16, September 2012.
- 5 Kaustuv Chaudhuri, Dale Miller, and Alexis Saurin. Canonical sequent proofs via multi-focusing. In *IFIP TCS*, 2008.
- 6 N Ghani. Beta-eta equality for coproducts. In *Proceedings of TLCA'95*, number 902 in Lecture Notes in Computer Science. Springer-Verlag, 1995.
- 7 Neelakantan R. Krishnaswami. Focusing on pattern matching. In *POPL*, 2009.
- 8 Sam Lindley. Extensional rewriting with sums. In *TLCA*, 2007.
- 9 Dale Miller and Alexis Saurin. From proofs to focused proofs: A modular proof of focalization in linear logic. In *CSL*, 2007.
- 10 Alexis Saurin. *Une étude logique du contrôle (appliquée à la programmation fonctionnelle et logique)*. PhD thesis, École Polytechnique, 2008.
- 11 Robert J. Simmons. Structural focalization. *CoRR*, abs/1109.6273, 2011.
- 12 Noam Zeilberger. *The Logical Basis of Evaluation Order and Pattern-Matching*. PhD thesis, Carnegie Mellon University, 2009.

A Proof-theoretic Characterization of Independence in Type Theory

Yuting Wang¹ and Kaustuv Chaudhuri²

1 University of Minnesota, USA
yuting@cs.umn.edu

2 Inria & LIX/École polytechnique, France
kaustuv.chaudhuri@inria.fr

Abstract

For λ -terms constructed freely from a type signature in a type theory such as LF, there is a simple inductive *subordination* relation that is used to control type-formation. There is a related – but not precisely complementary – notion of *independence* that asserts that the inhabitants of the function space $\tau_1 \rightarrow \tau_2$ depend vacuously on their arguments. Independence has many practical reasoning applications in logical frameworks, such as pruning variable dependencies or transporting theorems and proofs between type signatures. However, independence is usually not given a formal interpretation. Instead, it is generally implemented in an *ad hoc* and uncertified fashion. We propose a formal definition of independence and give a proof-theoretic characterization of it by: (1) representing the inference rules of a given type theory and a closed type signature as a theory of intuitionistic predicate logic, (2) showing that typing derivations in this signature are adequately represented by a focused sequent calculus for this logic, and (3) defining independence in terms of *strengthening* for intuitionistic sequents. This scheme is then formalized in a meta-logic, called \mathcal{G} , that can represent the sequent calculus as an inductive definition, so the relevant strengthening lemmas can be given explicit inductive proofs. We present an algorithm for automatically deriving the strengthening lemmas and their proofs in \mathcal{G} .

1998 ACM Subject Classification F.4.2. Mathematical logic: proof theory

Keywords and phrases subordination, independence, sequent calculus, focusing, strengthening

Digital Object Identifier 10.4230/LIPIcs.TLCA.2015.332

1 Introduction

In logical frameworks such as LF [6] or \mathcal{G} [4] that are designed to reason about typed λ -terms *qua* data, one notion that appears again and again is *dependency*: when can the structure of one group of λ -terms depend essentially on that of another group of λ -terms? The most widely studied general notion of dependency is *subordination* [15, 17, 7], which is best explained using an example. Consider λ -terms built out of the following *type signature* of constants, where **nat** and **bt** respectively denote natural numbers and binary trees with natural numbers in the leaves.

$$\mathbf{z} : \mathbf{nat}. \quad \mathbf{s} : \mathbf{nat} \rightarrow \mathbf{nat}. \quad \mathbf{leaf} : \mathbf{nat} \rightarrow \mathbf{bt}. \quad \mathbf{node} : \mathbf{bt} \rightarrow \mathbf{bt} \rightarrow \mathbf{bt}.$$

From this signature, it is immediately evident that a closed β -normal term of type **bt** can – indeed, *must* – contain a subterm of type **nat**, so we say that **nat** is *subordinate* to **bt**. The subordination relation \leq on types can be derived from a type signature as follows (adapting [7, Definition 2.14]). For any type τ , write $\mathcal{H}(\tau)$ for its *head*, which is the basic type that occurs rightmost in the chain of \rightarrow s (or dependent products in the case of dependent types) in



© Yuting Wang and Kaustuv Chaudhuri;
licensed under Creative Commons License CC-BY

13th International Conference on Typed Lambda Calculi and Applications (TLCA'15).

Editor: Thorsten Altenkirch; pp. 332–346



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

τ . Then, for every type $\tau_1 \rightarrow \cdots \rightarrow \tau_n \rightarrow a$ (here, a is the head) that occurs anywhere in the signature, set $\mathcal{H}(\tau_i) \leq a$ for every $i \in 1..n$. Finally, define $\tau_1 \leq \tau_2$ generally as $\mathcal{H}(\tau_1) \leq \mathcal{H}(\tau_2)$ and close \leq under reflexivity and transitivity.

With this definition, and the above signature, we have that $\mathbf{nat} \leq \mathbf{bt}$ but $\mathbf{bt} \not\leq \mathbf{nat}$. This notion of subordination, when strictly enforced such as in canonical LF [17, 7], enables a kind of modularity of reasoning: inductive theorems about the λ -terms of a given type can be proved in the *smallest* relevant signature and imported into larger signatures that do not contain subordinate types. For instance, meta-theorems about \mathbf{nats} , proved in a context of \mathbf{nat} assumptions, can be transported to contexts of \mathbf{bt} assumptions since $\mathbf{bt} \not\leq \mathbf{nat}$. It is indeed this complement of subordination that is most useful in reasoning: intuitively, $\mathbf{bt} \not\leq \mathbf{nat}$ means that the inhabitants of $\mathbf{bt} \rightarrow \mathbf{nat}$ are functions whose arguments cannot occur in their bodies in the β -normal form. This negative reading of subordination can be used to prune dependencies during unification, which may bring an unsolvable higher-order problem into a solvable pattern problem [9], or to prevent *raising* a variable over non-subordinate arguments, producing more concise proofs [3].

However, this pleasingly simple notion of subordination has a somewhat obscure formal interpretation: the definition is independent of the typing rules and it is unclear how they are related. We set out to formalize such a relation in terms of an inductive characterization of the (β -normal) inhabitants of types, and in the process we discovered a curious aspect of the above definition of subordination that manifests for higher-order constructors. Take, for instance, the alternative type $(\mathbf{nat} \rightarrow \mathbf{bt}) \rightarrow \mathbf{bt}$ for \mathbf{leaf} . Nothing changes as far as the \leq relation is concerned: $\mathbf{nat} \rightarrow \mathbf{bt}$ still occurs in the signature, so $\mathbf{nat} \leq \mathbf{bt}$ is still true.¹ Yet, if we look at all β -normal terms of type \mathbf{bt} now, there can be no subterms of type \mathbf{nat} since there does not exist a constructor for injecting them into terms of type \mathbf{bt} in the base case. The definition of \leq above clearly over-approximates possible dependencies, for reasons that are not at all obvious, so its complement is too conservative.

In this paper we propose an alternative view of dependency that is not based on the subordination relation. We directly characterize when one type is *independent* of another with a *proof-theoretic* view of independence: for every claimed independence, we establish, by means of an induction over all typing derivations based on a given signature, that indeed a certain dependency is not strict. This view has several benefits:

- First, our notion of independence is larger than non-subordination, which means that it can be used for more aggressive pruning of dependencies.
- More important is that we have strong confidence in the correctness of our definition of independence, since it is now a derived property of the type system. Indeed, we propose an algorithm that extracts formal inductive proofs of independence that can be verified without needing a built-in notion of subordination or independence. This changes independence from a trusted framework-level procedure to a certifying procedure.
- Finally, we use only standard and simple proof-theoretic machinery in our definition. We require neither rich type systems nor sophisticated meta-theoretic tools.

Our view of independence has the following outline, which is also the outline of the paper.

1. We start (Section 2) by defining independence as a property of a given type theory.
2. We then (Section 3) describe a *specification language* built around the logic of higher-order hereditary Harrop formulas (HH) [10]. This is a fragment of first-order intuitionistic logic

¹ In canonical LF [7, Definition 2.14], well-formedness of signatures requires subordination of argument types of all dependent products.

with a particularly simple focused proof system [16]. It can also be seen as the minimal logical core of every logical framework.

3. We then (Section 4) use the language to give an *adequate* encoding of the inference system that underlies a given type system. To keep the presentation simple, we have chosen to use the simply typed λ -calculus, but the technique generalizes at least to LF [13, 14]. In terms of this encoding, we characterize independence as a particular kind of *strengthening*.
4. We then (Section 5) take the focused HH sequent calculus as an object logic in the reasoning logic \mathcal{G} that adds the remaining crucial ingredients: a *closed world* reading, induction, $\beta\eta$ -equality of λ -terms, and generic reasoning. This is the *two-level logic approach* [5] that underlies the Abella theorem prover [18].
5. Lastly (Section 6) we show how strengthening for typing derivations is formalized in \mathcal{G} , and give an algorithm for automatically deriving these lemmas from a given signature. We show an application of the formalization to pruning unnecessary dependence between variables in \mathcal{G} .

The Abella development of examples in this paper is available at:

<http://abella-prover.org/independence>.

2 Independence in Type Theory

Intuitively, in a given type theory, τ_2 is independent of τ_1 if and only if the type $\tau_1 \rightarrow \tau_2$ is only inhabited by abstractions whose bodies in the β -normal form do not contain the arguments. We can write this as a property on typing derivations:

► **Definition 1** (Independence). Let $\Gamma \vdash t : \tau$ be the typing judgment in the given type theory. The type τ_2 is *independent* of τ_1 in Γ if whenever $\Gamma, x:\tau_1 \vdash t : \tau_2$ holds for some t , the β -normal form of t does not contain x , *i.e.*, $\Gamma \vdash t : \tau_2$ holds. ◀

A straightforward way to prove such a property is to perform inductive reasoning on the first typing derivation. For this, we need to know not only what are the possible ways to prove a typing judgment, but also that they are the only ways. This is the *closed-world assumption* that underlies reasoning about a fixed type signature. However, even with this assumption, the inductive proofs for independence can be hard to establish: since the target type τ_2 is fixed, the inductive hypothesis will not be applicable to new types encountered during the induction. To see this, suppose we are working with the simply-typed λ -calculus (STLC). Typing rules for STLC derive judgments of the form $\Gamma \vdash t : \tau$ where Γ is a context that assigns unique types to distinct variables, t is a λ -term and τ is its type. The typing rules are standard:

$$\frac{x:\tau \in \Gamma}{\Gamma \vdash x : \tau} \text{t-bas} \quad \frac{\Gamma \vdash t_1 : \tau_1 \rightarrow \tau \quad \Gamma \vdash t_2 : \tau_1}{\Gamma \vdash t_1 t_2 : \tau} \text{t-app} \quad \frac{\Gamma, x:\tau_1 \vdash t : \tau \quad (x \notin \Gamma)}{\Gamma \vdash \lambda x:\tau_1. t : \tau_1 \rightarrow \tau} \text{t-abs}$$

A direct induction on $\Gamma \vdash t : \tau$ using the typing rules produces three cases. The case when t is an application results in two premises where the premise for the argument has a new target type τ_1 (as shown in **t-app**). Since τ_1 can be any arbitrary type that is not necessarily related to τ , it is not possible to appeal to the inductive hypothesis.

The key observation is that the signature must be fixed for the dependence between types to be fully characterized. We propose an approach to formalize independence by (1) giving an encoding of a given type theory and a closed signature in a specification logic which finitely characterizes the dependence between types, (2) proving that the encoding is faithful to the

$$\begin{array}{c}
\frac{\Sigma, \bar{x}:\bar{\tau}; \Gamma, F_1, \dots, F_n \vdash A}{\Sigma; \Gamma \vdash \Pi \bar{x}:\bar{\tau}. F_1 \Rightarrow \dots \Rightarrow F_n \Rightarrow A} \text{ reduce} \quad \frac{(F \in \Gamma) \quad \Sigma; \Gamma, [F] \vdash A}{\Sigma; \Gamma \vdash A} \text{ focus} \\
\\
\frac{\Sigma \vdash \theta : \bar{\tau} \quad A'[\theta] = A \quad \{\Sigma; \Gamma \vdash F_i[\theta]\} \text{ for } i \in 1..n}{\Sigma; \Gamma, [\Pi x_1:\tau_1, \dots, x_n:\tau_n. F_1 \Rightarrow \dots \Rightarrow F_n \Rightarrow A'] \vdash A} \text{ backchain}
\end{array}$$

■ **Figure 1** Inference rules for HH. For *reduce*, we assume that $\bar{x} \notin \Sigma$. In the *backchain* rule, θ stands for a *substitution* $[t_1/x_1, \dots, t_n/x_n]$; we write $\Sigma \vdash \theta : \bar{\tau}$ to mean $\Sigma \vdash t_i : \tau_i$ for each $i \in 1..n$.

original type theory and signature, and (3) formally stating and proving the independence as lemmas in a reasoning logic that gives an inductive reading of the encoding. The first two tasks are covered by Section 3 and 4. The last task is the topic of Section 5 and 6. We use STLC as the example throughout the paper; extending to other type theories such as LF is left for future work.

3 The Specification Language

Let us begin with a sketch of a specification language for encoding rule-based systems. This language will be used to encode type theories in later sections.

3.1 The HH Proof System

The logic of *higher-order hereditary Harrop formulas* (HH) is an intuitionistic and predicative fragment of Church's simple theory of types. Expressions in HH are simply typed λ -terms. Types are freely formed from built-in or user-defined atomic types containing at least the built-in type \circ (of formulas) and the arrow type constructor \rightarrow (right-associative). The *head type* of τ is the atomic type that occurs rightmost in a chain of \rightarrow s. Terms are constructed from a *signature* (written Σ); we write $\Sigma \vdash t : \tau$ if a λ -term t has type τ in Σ with the usual rules. Logic is built into HH by means of logical constants including $\Rightarrow : \circ \rightarrow \circ \rightarrow \circ$ (written infix and right-associative) for implications and $\Pi_\tau : (\tau \rightarrow \circ) \rightarrow \circ$ for universal quantification over values of a type τ that does not contain \circ . Predicates are constants in the signature with head type \circ . For readability, we will often write $t_1 \Leftarrow t_2$ for $t_2 \Rightarrow t_1$, and abbreviate $\Pi_\tau(\lambda x:\tau. t)$ as $\Pi x:\tau. t$ and $\Pi x_1:\tau_1 \dots \Pi x_n:\tau_n. t$ as $\Pi \bar{x}:\bar{\tau}. t$ where $\bar{x} = x_1, \dots, x_n$ and $\bar{\tau} = \tau_1, \dots, \tau_n$. We also omit type subscripts when they are obvious from context.

An *atomic formula* (written A) is a term of type \circ that head-normalizes to a term that is not an application of one of the logical constants \Rightarrow or Π_τ . Every HH formula is equivalent to a *normalized formula* (written F) of the form $\Pi \bar{x}:\bar{\tau}. F_1 \Rightarrow \dots \Rightarrow F_n \Rightarrow A$. In the rest of this paper we will assume that all formulas are normalized unless explicitly excepted. Given a normalized formula $F = \Pi \bar{x}:\bar{\tau}. F_1 \Rightarrow \dots \Rightarrow F_n \Rightarrow A$, we write: $\mathcal{H}(F)$, the *head of* F , for A ; $\mathcal{L}(F)$, the *body of* F , for the multiset $\{F_1, \dots, F_n\}$; and $\mathcal{B}(F)$, the *binders of* F , for the typing context $\bar{x}:\bar{\tau}$.

The inference system for HH is a negatively polarized focused sequent calculus, also known as a uniform proof system, that contains two kinds of *sequents*: the *goal reduction* sequent $\Sigma; \Gamma \vdash F$ and the *backchaining* sequent $\Sigma; \Gamma, [F] \vdash A$ with F *under focus*. In either case, Γ is a multiset of formulas called *program clauses* and the right hand side formula of \vdash is called the *goal formula*. Figure 1 contains the inference rules.

An HH proof usually starts (reading conclusion upwards) with a goal-reduction sequent $\Sigma; \Gamma \vdash F$. If F is not atomic, then the *reduce* rule is applied to make it atomic; this rule extends the signature with $\mathcal{B}(F)$ and the context with $\mathcal{L}(F)$. Then the *focus* rule is applied which selects a formula from the context Γ . This formula is then decomposed in the *backchain*

rule that produces fresh goal reduction sequents for (instances of the) body, assuming that the instantiated head of the focused formula matches the atomic goal formula. This process repeats unless the focused formula has no antecedents.

The three rules of HH can be combined together into one *synthetic* rule for goal reduction sequents that have an atomic goal formula.

$$\frac{F \in \Gamma \quad \Sigma \vdash \theta : \bar{\tau} \quad \mathcal{H}(F)[\theta] = A \quad \{\Sigma, \mathcal{B}(G[\theta]); \Gamma, \mathcal{L}(G[\theta]) \vdash \mathcal{H}(G[\theta])\} \text{ for } G \in \mathcal{L}(F)}{\Sigma; \Gamma \vdash A} \text{ bcred}$$

Every premise of this rule is a goal reduction sequent with an atomic goal formula. In the rest of this paper we will limit our attention to this fragment of HH.

3.2 Encoding Rule-based Systems in HH

Because the expressions of HH are λ -terms, we can use the *λ -tree approach to syntax* (λ TS), sometimes known as *higher-order abstract syntax* (HOAS), to represent the rules of deductive systems involving binding using the binding structure of λ -terms. Binding in object-language syntax is represented explicitly by meta-language λ -abstraction, and recursion over such structures is realized by introducing fresh new constants using universal goals and recording auxiliary properties for such constants via hypothetical goals. This kind of encoding is concise and, as we shall see in later sections, has logical properties that we can use in reasoning.

We present the encoding of typing rules for STLC as described in Section 2 as a concrete example of specifying in HH. Two basic types, **ty** and **tm**, are used for classifying types and terms in STLC. We then use the following signature defining the type and term constructors:

$$\mathbf{b} : \mathbf{ty}. \quad \mathbf{arr} : \mathbf{ty} \rightarrow \mathbf{ty} \rightarrow \mathbf{ty}. \quad \mathbf{app} : \mathbf{tm} \rightarrow \mathbf{tm} \rightarrow \mathbf{tm}. \quad \mathbf{abs} : \mathbf{ty} \rightarrow (\mathbf{tm} \rightarrow \mathbf{tm}) \rightarrow \mathbf{tm}.$$

The type of the **abs** uses λ TS to encode binding. To illustrate, $(\lambda y:b \rightarrow b. \lambda x:b. y x)$ is encoded as **abs** (**arr** **b** **b**) ($\lambda y. \mathbf{abs} \mathbf{b} (\lambda x. \mathbf{app} y x)$).

We define a predicate **of** : **tm** \rightarrow **ty** \rightarrow **o** to encode the typing judgments in STLC, with the context implicitly represented by the context of HH sequents. The typing rules are encoded by the following program clauses (where the outermost universal quantifiers are omitted):

$$\begin{aligned} \mathbf{of} (\mathbf{app} M_1 M_2) T &\Leftarrow \mathbf{of} M_1 (\mathbf{arr} T_1 T) \Leftarrow \mathbf{of} M_2 T_1 \\ \mathbf{of} (\mathbf{abs} T_1 R) (\mathbf{arr} T_1 T) &\Leftarrow (\Pi x. \mathbf{of} (R x) T \Leftarrow \mathbf{of} x T_1) \end{aligned}$$

Here, we are following the usual logic programming convention of writing universally quantified variables using upper-case identifiers.

To see that these clauses accurately represent the typing rules of STLC, consider deriving a HH sequent $\Sigma; \Gamma \vdash \mathbf{of} M T$, where Γ contains the clauses above and the possible assignments of types to variables introduced by the second clause corresponding to the abstraction rule **t-abs**. The only way to proceed is by focusing and backchaining on one of the clauses. Backchaining on the first clause unifies the goal with the head formula and produces two premises that corresponds to the premises of **t-app**. Backchaining on the second clause followed by goal reduction results in $\Sigma, x : \mathbf{tm}; \Gamma, \mathbf{of} x T_1 \vdash \mathbf{of} (R x) T$. Note that x is a fresh variable introduced by reducing the universal goal and $\mathbf{of} x T_1$ is the typing assignment of x from reducing the hypothetical goal, which exactly captures the side condition of **t-abs** that x must be fresh for Γ . The rule **t-bas** is modeled by backchaining on assumptions in Γ that assigns types to variables introduced by **t-abs**.

$$\begin{aligned} \hat{\text{ty}} \mathbf{b}. \quad \hat{\text{ty}} (\text{arr } T_1 T_2) &\Leftarrow \hat{\text{ty}} T_1 \Leftarrow \hat{\text{ty}} T_2. & \hat{\text{tm}} (\text{app } M_1 M_2) &\Leftarrow \hat{\text{tm}} M_1 \Leftarrow \hat{\text{tm}} M_2. \\ \hat{\text{tm}} (\text{abs } T R) &\Leftarrow (\Pi x:\text{tm}. \hat{\text{tm}} (R x) \Leftarrow \hat{\text{tm}} x) \Leftarrow \hat{\text{ty}} T. \end{aligned}$$

■ **Figure 2** An example encoding of an STLC signature into HH clauses

4 Independence via Strengthening

This section presents an encoding of STLC in HH that finitely captures the dependence between types in a closed signature. Then the independence property can be stated as strengthening lemmas and proved by induction.

4.1 An Adequate Encoding of STLC

The encoding is based on the *types-as-predicates* principle: every type is interpreted as a predicate that is true of its inhabitants. The atomic types and constants of STLC are imported directly into the HH signature. For every atomic type b , we define a predicate $\hat{b} : b \rightarrow \circ$. We then define a mapping $\llbracket - \rrbracket$ from STLC types τ to a function $\tau \rightarrow \circ$ as follows:

$$\llbracket b \rrbracket = \lambda t. \hat{b} t \quad \text{if } b \text{ is an atomic type.} \quad \llbracket \tau_1 \rightarrow \tau_2 \rrbracket = \lambda t. \Pi x:\tau_1. \llbracket \tau_1 \rrbracket x \Rightarrow \llbracket \tau_2 \rrbracket (t x)$$

The mapping $\llbracket - \rrbracket$ is extended to typing contexts: for $\Gamma = x_1:\tau_1, \dots, x_n:\tau_n$, we write $\llbracket \Gamma \rrbracket$ for the multiset of HH formulas $\{\llbracket \tau_1 \rrbracket x_1, \dots, \llbracket \tau_n \rrbracket x_n\}$. A typing judgment $\Gamma \vdash t : \tau$ is encoded as an HH sequent $\Gamma; \llbracket \Gamma \rrbracket \vdash \llbracket \tau \rrbracket t$.

As an example, consider the signature with STLC as an object language as described in Section 3.2 (not to be confused with STLC we are encoding). We have two predicates $\hat{\text{tm}} : \text{tm} \rightarrow \circ$ and $\hat{\text{ty}} : \text{ty} \rightarrow \circ$. Let Γ be the STLC signature containing \mathbf{b} , arr , app and abs . It is translated into $\llbracket \Gamma \rrbracket$ containing the program clauses in Figure 2 (in normalized form and with outermost quantifiers elided): The typing judgment $\Gamma, y : \text{tm} \rightarrow \text{tm} \vdash \text{abs } \mathbf{b} y : \text{tm}$ which is provable in STLC is encoded as $\Gamma, y:\text{tm} \rightarrow \text{tm}; \llbracket \Gamma \rrbracket, (\Pi x:\text{tm}. \hat{\text{tm}} x \Rightarrow \hat{\text{tm}} (y x)) \vdash \hat{\text{tm}} (\text{abs } \mathbf{b} y)$ which is also provable in HH.

This encoding generalizes to richer type systems than STLC. An encoding of LF into HH was presented in [2, 13, 14], which is essentially a superset of the encoding we are doing. The soundness and completeness of the STLC encoding follows easily from the results in [13].

► **Theorem 2.** *Let Γ be a well-formed context and τ be a type in STLC. If $\Gamma \vdash t : \tau$ has a derivation for a $\beta\eta$ -normal term t , then there is a derivation of $\Gamma; \llbracket \Gamma \rrbracket \vdash \llbracket \tau \rrbracket t$ in HH. Conversely, if $\Gamma; \llbracket \Gamma \rrbracket \vdash \llbracket \tau \rrbracket t$ for any term t that is well-typed in HH, then $\Gamma \vdash t' : \tau$ where t' is the canonical (β -normal η -long) form of t .*

Proof. This theorem is a special case of [13, Theorem 1]. The proof is almost the same. ◀

4.2 Independence as Strengthening Lemmas

By Definition 1 and Theorem 2, τ_2 is independent from τ_1 in Γ iff the following *strengthening* lemma is true: if $\Gamma, x:\tau_1; \llbracket \Gamma \rrbracket, \llbracket \tau_1 \rrbracket x \vdash \llbracket \tau_2 \rrbracket t$ holds for some t , then the β -normal form of t does not contain x and $\Gamma; \llbracket \Gamma \rrbracket \vdash \llbracket \tau_2 \rrbracket t$ holds. Because the typing information in formulas generated by $\llbracket - \rrbracket$ is statically determined by predicates and the translated program is finite, it is possible to determine the dependence between types finitely. Thus, the independence argument can be proved by induction.

Using the previous encoding example, let's see how to prove that ty is independent of tm in Γ which is the signature with STLC in Section 3.2. We need to show that $\Gamma, x:\text{tm} \vdash t : \text{ty}$

implies $\Gamma \vdash t : \hat{\tau}y$, which by Theorem 2 is equivalent to $\Gamma, x:\hat{\tau}m; [\Gamma], \hat{\tau}m \ x \vdash \hat{\tau}y \ t$ implying $\Gamma; [\Gamma] \vdash \hat{\tau}y \ t$. The proof proceeds by induction on the first assumption. Apparently, $\hat{\tau}y \ t$ can only be proved by backchaining on clauses whose heads start with $\hat{\tau}y$. The case when $\hat{\tau}y \ t$ is proved by backchaining on $\hat{\tau}y \ b$ is immediate. For the case when $t = \text{arr } t_1 \ t_2$, backchaining produces two premises $\Gamma, x:\hat{\tau}m; [\Gamma], \hat{\tau}m \ x \vdash \hat{\tau}y \ t_i$ ($i \in \{1, 2\}$). Applying the inductive hypothesis we get $\Gamma; [\Gamma] \vdash \hat{\tau}y \ t_i$, from which the conclusion $\Gamma; [\Gamma] \vdash \hat{\tau}y \ (\text{arr } t_1 \ t_2)$ follows easily.

5 The Two-Level Logic Approach

5.1 The Reasoning Logic \mathcal{G}

When a relation is described as an inductive inference system, the rules are usually understood as fully characterizing the relation. When interpreting the relation as a computation, we can just give the inference rules a *positive* interpretation. However, when reasoning about the properties of the inference system, the inductive definition must be seen in a *negative* interpretation, *i.e.*, as a prescription of the *only* possible ways to establish the encoded property. Concretely, given the rules for typing λ -terms in STLC, we not only want to identify types with typable terms, but also to argue that a term such as $\lambda x. x \ x$ does not have a type. We sketch the logic \mathcal{G} that supports this complete reading of rule-based specifications by means of *fixed-point definitions* [4].

To keep things simple, \mathcal{G} uses the same term language as HH and is also based on Church's simple theory of types. At the type level, the only difference is that \mathcal{G} formulas have type **prop** instead of **o**. The non-atomic formulas of \mathcal{G} include formulas of ordinary first-order intuitionistic logic, built using the constants $\top, \perp : \mathbf{prop}$, $\wedge, \vee, \supset : \mathbf{prop} \rightarrow \mathbf{prop} \rightarrow \mathbf{prop}$ (written infix), and $\forall_\tau, \exists_\tau : (\tau \rightarrow \mathbf{prop}) \rightarrow \mathbf{prop}$ for types τ that do not contain **prop**. To this, \mathcal{G} adds intensional equality at all types, $=_\tau : \tau \rightarrow \tau \rightarrow \mathbf{prop}$, which is interpreted as $\beta\eta$ -convertibility. Like with HH earlier, we will drop the explicit types for these constants and write quantifiers using more familiar notation. The proof system of \mathcal{G} is a sequent calculus with the standard derivation rules for logical constants [5], which we will not repeat here.

The next crucial ingredient in \mathcal{G} is a mechanism for *defining* predicates by means of *fixed-point definitions*. Such a definition is characterized by a collection of *definitional clauses* each of which has form $\forall \bar{x}. A \triangleq B$ where A is an atomic formula all of whose free variables are bound by \bar{x} and B is a formula whose free variables must occur in A . A is called the head of such a clause and B is called its body. (For a fixed-point definition to be well-formed, it must satisfy certain stratification conditions [8] that we do not elaborate on here.) To illustrate definitions, let **olist** be a new basic type for lists of HH formulas, built from the type signature **nil** : **olist** and $(::) : \mathbf{o} \rightarrow \mathbf{olist} \rightarrow \mathbf{olist}$. Then, list membership (**member**) and list concatenation (**append**) may be specified in \mathcal{G} using the following definitional clauses:

$$\begin{array}{ll} \text{member } X \ (X :: L) \triangleq \top. & \text{member } X \ (Y :: L) \triangleq \text{member } X \ L. \\ \text{append nil } L \ L \triangleq \top. & \text{append } (X :: L_1) \ L_2 \ (X :: L_3) \triangleq \text{append } L_1 \ L_2 \ L_3. \end{array}$$

As before, we use the convention of indicating universally closed variables with upper-case identifiers. Read positively, these clauses can be used in the standard backchaining style to derive atomic formulas: the goal must match with the head of a clause and the backchaining step reduces the goal to deriving the instances of corresponding body. Read negatively, they are used to do case analysis on an assumption: if an atomic formula holds, then it must be the case that it unifies with the head of some clause defining it and the body of the clause is derivable. It therefore suffices to show that the conclusion follows from each such possibility.

Fixed-point definitions can also be interpreted inductively or coinductively, leading to corresponding reasoning principles. We use an annotated style of reasoning to illustrate how induction works in \mathcal{G} . Its interpretation into the proof theory of \mathcal{G} is described in [3]. When proving a formula $(\forall \bar{x}. D_1 \supset \dots \supset A \supset \dots \supset D_n \supset G)$ by induction on the atom A , the proof reduces to proving G with the inductive hypothesis $(\forall \bar{x}. D_1 \supset \dots \supset A^* \supset \dots \supset D_n \supset G)$ and assumptions $D_1, \dots, A^\circ, \dots, D_n$. Here, A^* and A° are simply annotated versions of A standing for *strictly smaller* and *equal sized* measures respectively. When A° is *unfolded* by using a definitional clause, the predicates in the body of the corresponding clause are given the $*$ annotation. Thus, the inductive hypothesis can only be used on A^* that results from unfolding A° at least once.

The negative reading of fixed-point definitions in \mathcal{G} requires some care. In the negative view, universally quantified variables are interpreted *extensionally*, *i.e.*, as standing for all their possible instances. To illustrate, if \mathbf{nat} were defined by $\mathbf{nat} \ z \triangleq \top$ and $\mathbf{nat} \ (s \ X) \triangleq \mathbf{nat} \ X$, then one can derive $\forall x. \mathbf{nat} \ x \supset G$ by case-analysis of the assumption $\mathbf{nat} \ x$, which amounts to proving $[z/x]G$ and $\forall y. \mathbf{nat} \ y \supset [s \ y/x]G$. This extensional view of universal variables is not appropriate when reasoning about binding structures viewed as syntax, where the syntactic variables are not stand-ins for all terms but rather for *names*. To see this clearly, consider the formula $\forall w. (\lambda x. w) = (\lambda x. x) \supset \perp$. If equality of λ -terms were interpreted extensionally with \forall , we would be left with $\forall w. (\forall x. w = x) \supset \perp$ which is not provable. Yet, the λ -terms $(\lambda x. w)$ and $(\lambda x. x)$ denote the constant and identity functions, respectively, and are therefore intensionally different – neither is $\beta\eta$ -convertible to the other.

To achieve this intensional view, we come to the final ingredient of \mathcal{G} : *generic reasoning*. Every type of \mathcal{G} is endowed with an infinite number of *nominal constants*, and there is a quantifier $\nabla_\tau : (\tau \rightarrow \mathbf{prop}) \rightarrow \mathbf{prop}$ (for τ not containing \mathbf{prop}) to abstract over such constants. The critical features of nominal constants are: (1) they are not $\beta\eta$ -convertible, so $\nabla x. \nabla y. x = y \supset \perp$ is derivable; and (2) formulas on the left and right of a \mathcal{G} sequent interact up to *equivariance*, which allows the nominal constants in one formula to be systematically permuted to match those of the other. This latter property allows the ∇ quantifier to permute with all other connectives, including disjunction. The rules for introducing ∇ quantified formulas both as assumption and conclusion are similar: a formula $\nabla x. A$ is reduced to $[c/x]A$ where c is a nominal constant that is not already present in A .

The general form of a \mathcal{G} definitional clause therefore allows the ∇ quantifier in heads: $\forall \bar{x}. (\nabla \bar{z}. A) \triangleq B$. An instance of the clause must have \bar{z} replaced by distinct nominal constants that are not already present in the formula. Since \bar{x} is quantified outside of \bar{z} , their instantiations cannot contain \bar{z} , which is used to encode structural properties of terms in the definitions. For example, $(\nabla x. \mathbf{name} \ x) \triangleq \top$ defines \mathbf{name} such that $\mathbf{name} \ X$ holds only if X is a nominal constant. Another example is $\forall T. (\nabla x. \mathbf{fresh} \ x \ T) \triangleq \top$ which defines \mathbf{fresh} such that $\mathbf{fresh} \ X \ T$ holds only if X is a nominal constant that does not occur in T .

5.2 An Encoding of HH in \mathcal{G}

The HH proof system can be encoded as a fixed-point definition in \mathcal{G} . The logical constants, signatures, and terms are imported into \mathcal{G} transparently from HH – this is possible since \mathcal{G} and HH use the same simple type system. The sequents of HH are then encoded in \mathcal{G} using the predicates $\mathbf{seq} : \mathbf{olist} \rightarrow \circ \rightarrow \mathbf{prop}$ and $\mathbf{bch} : \mathbf{olist} \rightarrow \circ \rightarrow \circ \rightarrow \mathbf{prop}$. The encodings and their abbreviated notations are as follows:

HH	\mathcal{G}	notation
$\Sigma; \Gamma \vdash F$	$\mathbf{seq} \ \Gamma \ F$	$\{\Gamma \vdash F\}$
$\Sigma; \Gamma, [F] \vdash A$	$\mathbf{bch} \ \Gamma \ F \ A$	$\{\Gamma, [F] \vdash A\}$

$$\begin{array}{ll}
\mathbf{seq} L (F \Rightarrow G) \triangleq \mathbf{seq} (F :: L) G. & \mathbf{bch} L (G \Rightarrow F) A \triangleq \mathbf{seq} L G \wedge \mathbf{bch} L F A. \\
\mathbf{seq} L (\Pi x:\tau. F x) \triangleq \nabla x:\tau. \mathbf{seq} L (F x). & \mathbf{bch} L (\Pi x:\tau. F x) A \triangleq \exists t:\tau. \mathbf{bch} L (F t) A. \\
\mathbf{seq} L A \triangleq \mathbf{atom} A \wedge \mathbf{member} F L \wedge \mathbf{bch} L F A. & \mathbf{bch} L A A \triangleq \top.
\end{array}$$

■ **Figure 3** Encoding of HH rules as inductive definitions in \mathcal{G} .

The definitional clause for \mathbf{seq} and \mathbf{bch} are listed in Figure 3. Note that we use a list of formulas to represent the multiset context in HH. This is not a problem since we can prove that the ordering of clauses in the list in \mathbf{seq} and \mathbf{bch} formulas does not affect their derivability as theorems about \mathbf{seq} and \mathbf{bch} in \mathcal{G} .

The encoding of HH in \mathcal{G} is adequate. The first two clauses defining \mathbf{seq} exactly capture the *reduce* rule. Note that in the second clause we use ∇ to encode Π in HH, since Π introduces fresh and unanalyzable eigenvariables into the HH signature. The third clause encodes the *focus* rule, where \mathbf{atom} is defined as

$$\mathbf{atom} F \triangleq (\forall G. (F = \Pi x:\tau. G x) \supset \perp) \wedge (\forall G_1, G_2. (F = (G_1 \Rightarrow G_2)) \supset \perp).$$

The clauses defining \mathbf{bch} exactly capture the *backchain* rule. Note that the two rules for introducing Π actually represents a collection of rules for each instance of τ . In other words, these rule are actually schematic rules. This is possible since the proof theory allows inductive definitions to have infinitely many clauses. We will often write elements in \mathbf{olist} in reverse order separated by commas; the exact meaning of the comma depends on its context. For example, given $L_1, L_2 : \mathbf{olist}$ and $A : \mathbf{o}$, $\{L_1, L_2, A \vdash G\}$ stands for $\mathbf{seq} (A :: L) G$ for some $L : \mathbf{olist}$ such that $\mathbf{append} L_1 L_2 L$ holds.

Theorems of HH specifications can be proved through this encoding in \mathcal{G} . As an example, consider proving the uniqueness of the encoding of typing in STLC shown in Section 3.2. The theorem can be stated as follows:

$$\forall L, M, T, T'. \{\Gamma, L \vdash \mathbf{of} M T\} \supset \{\Gamma, L \vdash \mathbf{of} M T'\} \supset T = T'.$$

where Γ represents the program clauses defining \mathbf{of} in Section 3.2.²

6 Formalizing Independence

This section first describes the formalization of independence in terms of *strengthening* for HH sequents of a certain shape in \mathcal{G} . Then a general algorithm for automatically deriving such strengthening lemmas is presented. Lastly, an application of the formalization for pruning variable dependencies in \mathcal{G} is described.

6.1 Independence as Strengthening Lemmas in \mathcal{G}

A strengthening lemma has the following form in \mathcal{G} : $\forall \bar{x}. \{\Gamma, F \vdash G\} \supset \{\Gamma \vdash G\}$. As we have discussed in Section 3, an HH derivation always starts with applying *reduce* to turn the right hand side to atomic form. Thus, it suffices to consider instances of strengthening lemmas where G is atomic. This lemma is usually proved by induction on the only assumption. The proof proceeds by inductively checking that a derivation of $\{\Gamma', F \vdash A\}$ cannot contain any application of *focus* rule on F . In Section 3 we have shown that an HH derivation starting

² This is not precisely correct, since the typing context L also needs to be characterized by some inductive property; a complete exposition on this encoding can be found in the Abella tutorial [1].

with an atomic goal can be seen as repeatedly applying `bcred`. At every point in such a proof, we check whether the atomic goal A matches the head of F . If such a match never occurs, we can drop the assumption F from every HH sequent.

We formalize this intuition to establish independence for STLC in terms of a kind of strengthening lemma in \mathcal{G} . Concretely, by the adequacy of the encoding of STLC in HH and the adequacy of the encoding of HH in \mathcal{G} , τ_2 is independent of τ_1 in Γ if the following strengthening lemma is provable in \mathcal{G} :

$$\forall t. \nabla x. \{ \llbracket \Gamma \rrbracket, \llbracket \tau_1 \rrbracket x \vdash \llbracket \tau_2 \rrbracket (t x) \} \supset \exists t'. t = (\lambda y. t') \wedge \{ \llbracket \Gamma \rrbracket \vdash \llbracket \tau_2 \rrbracket t' \}.$$

Note that the variables in the context Γ become nominal constants of appropriate types which are absorbed into formulas. The term $(t x)$ indicates the *possible* dependence of t on x . The conclusion $t = (\lambda y. t')$ asserts that t is $\beta\eta$ -convertible to a term with a vacuous abstraction; this is indicated by the fact that t' is bound outside the scope of λ . To prove any instance of this lemma for particular types τ_1 and τ_2 , we proceed by induction on the only assumption. The conclusion $\exists t'. t = (\lambda y. t')$ is immediately satisfied when the atomic goals in the derivation do not match $\llbracket \tau_1 \rrbracket x$, since this is the only hypothesis where x occurs.

As an example, \mathbf{ty} is independent of \mathbf{tm} in $\llbracket \Gamma \rrbracket$ which contains the clauses in Figure 2 is formalized as the following lemma:

$$\forall t. \nabla x. \{ \llbracket \Gamma \rrbracket, \hat{\mathbf{tm}} x \vdash \hat{\mathbf{ty}} (t x) \} \supset \exists t'. t = (\lambda y. t') \wedge \{ \llbracket \Gamma \rrbracket \vdash \hat{\mathbf{ty}} t' \}$$

By induction on the assumption and introducing the goals, we get an inductive hypothesis

$$\forall t. \nabla x. \{ \llbracket \Gamma \rrbracket, \hat{\mathbf{tm}} x \vdash \hat{\mathbf{ty}} (t x) \}^* \supset \exists t'. t = (\lambda y. t') \wedge \{ \llbracket \Gamma \rrbracket \vdash \hat{\mathbf{ty}} t' \}$$

and a new hypothesis $\{ \llbracket \Gamma \rrbracket, \hat{\mathbf{tm}} x \vdash \hat{\mathbf{ty}} (t x) \}^\circ$. Unfolding this hypothesis amounts to analyzing all possible ways to derive this using the definitional clauses of Figure 3. Since $\hat{\mathbf{ty}} (t x)$ cannot match $\hat{\mathbf{tm}} x$, the selected focus cannot be $\hat{\mathbf{tm}} x$, so the only options are clauses selected from $\llbracket \Gamma \rrbracket$, of which only two clauses have heads compatible with $\hat{\mathbf{ty}} (t x)$. In the first case, for the clause $\hat{\mathbf{ty}} \mathbf{b}$, we unite $t x$ with \mathbf{b} , instantiating the eigenvariable t with $\lambda y. \mathbf{b}$; this in turn gives us the witness for t' to finish the proof. In the other case, $t x$ is united with $\mathbf{arr} t_1 t_2$ (for fresh eigenvariables t_1 and t_2). Here, t_1 and t_2 are first *raised* over x to make the two terms have the same nominal constants, and `bcred` then reduces the hypothesis to the pair of hypotheses, $\{ \llbracket \Gamma \rrbracket, \hat{\mathbf{tm}} x \vdash \hat{\mathbf{ty}} (t_1 x) \}^*$ and $\{ \llbracket \Gamma \rrbracket, \hat{\mathbf{tm}} x \vdash \hat{\mathbf{ty}} (t_2 x) \}^*$. The inductive hypothesis applies to both of these, so we conclude that $\{ \llbracket \Gamma \rrbracket \vdash \hat{\mathbf{ty}} t'_1 \}^*$ and $\{ \llbracket \Gamma \rrbracket \vdash \hat{\mathbf{ty}} t'_2 \}^*$ for suitable t'_1 and t'_2 that are independent of x . We can then finish the proof by forward-chaining on the definitional clauses for HH. Observe that this proof follows the informal proof described in Section 4.2

As another example, let's see why \mathbf{tm} is not independent of \mathbf{ty} , since abstractions contain their argument types. The relevant lemma would be:

$$\forall t. \nabla x. \{ \llbracket \Gamma \rrbracket, \hat{\mathbf{ty}} x \vdash \hat{\mathbf{tm}} (t x) \} \supset \exists t'. t = (\lambda y. t') \wedge \{ \llbracket \Gamma \rrbracket \vdash \hat{\mathbf{tm}} t' \}.$$

Here, a direct induction on the assumption would not work because we can now focus on the fourth clause of Figure 2 that would extend the context of the HH sequent with a fresh assumption of the form $\hat{\mathbf{tm}} x'$. The inductive hypothesis is prevented from being used because the context has *grown*. This is a standard feature of HH derivations, and we therefore need to give an inductive characterization of the structure of the *dynamically growing* context. We use the technique of defining these dynamic extensions in terms of *context definitions*; for example, for the $\hat{\mathbf{tm}}$ predicate, this definition has the following clauses:

$$\mathbf{ctx} \mathbf{nil} \triangleq \top; \quad (\nabla x. \mathbf{ctx} (\hat{\mathbf{tm}} x :: L)) \triangleq \mathbf{ctx} L.$$

Then we generalize the lemmas as follows:

$$\forall t, L. \nabla x. \text{ctx } L \supset \{[\Gamma], L, \hat{\text{ty}} x \vdash \hat{\text{tm}}(t x)\} \supset \exists t'. t = (\lambda y. t') \wedge \{[\Gamma], L \vdash \hat{\text{tm}} t'\}.$$

Now, by induction, when the fourth clause of Figure 2 is selected for focus, one of the hypotheses of the resulting `bcred` is $\{[\Gamma], L, \hat{\text{ty}} x \vdash \hat{\text{ty}}(t_1 x)\}$ (for a new eigenvariable t_1). It is entirely possible that $\hat{\text{ty}}(t_1 x)$ is proved by selecting $\hat{\text{ty}} x$. Thus we cannot conclude the original assumption does not depend on $\hat{\text{ty}} x$ – so `tm` is not independent of `ty`.

6.2 Automatically Deriving Strengthening Lemmas

The above illustrations show that these strengthening lemmas have a predictable form and proof structure. We will now give an algorithm that extracts these proofs automatically. The key insight is that the strengthening lemmas can be provable because of a *failure* to match the heads of the encoded clauses against the right hand sides of the HH sequent. Therefore, we simply need to enumerate the possible forms of the right hand sides and generate a mutually inductive lemma to cover all possibilities. This intuition is not entirely trivial to implement, since the HH contexts can potentially grow on every `bcred` step, which must then be accounted for.

For $F = (\Pi \bar{x}. F_1 \Rightarrow \dots \Rightarrow F_n \Rightarrow A)$, let $\mathcal{H}_p(F)$ stand for the head predicate in A and $\mathcal{L}_p(F)$ for $\{\mathcal{H}_p(F_i) \mid 1 \in i..n\}$. To prove $\forall \bar{x}. \{\Gamma, F \vdash A\} \supset \{\Gamma \vdash A\}$, we proceed as follows:

1. For every predicate a in Γ , we compute the possible dynamic contexts that arise in proofs of atomic formulas of head a .
2. For every predicate a in Γ and A , we compute a collection of predicates $S(a)$ which contains the head predicates of atomic formulas that may occur as the goal in a derivation starting with an atomic goal formula A' for which $\mathcal{H}_p(A') = a$. That is, the provability of any goal A' of head a only depends on formulas whose heads are in $S(a)$.
3. If $\mathcal{H}_p(A) = a$ and $\mathcal{H}_p(F) \notin S(a)$, then for every predicate $a' \in S(a)$ and any atomic goal A' s.t. $\mathcal{H}_p(A') = a'$, $\forall \bar{x}. \{\Gamma, F \vdash A'\} \supset \{\Gamma \vdash A'\}$ is provable. The proof proceeds by a simultaneous induction on all these formulas.
4. Since $a \in S(a)$, the required lemma is just one of the cases of the simultaneous induction.

Before we elaborate on these steps in the following subsections, note that our algorithm is sound and terminating, but not complete. The existence of a decision procedure for strengthening lemmas is outside the scope of this paper.

6.2.1 Calculating the dynamic contexts

Let Γ be a context that contains only finite distinct clauses, A be an atomic goal and F be some program clause. We would like to prove the strengthening lemma $\forall \bar{x}. \{\Gamma, F \vdash A\} \supset \{\Gamma \vdash A\}$ in \mathcal{G} . Let Δ be the set of predicates occurring in Γ . For every predicate $a \in \Delta$, let $C(a)$ denote a set of formulas that can possibly occur in the dynamic contexts of atomic formulas of head a . The only way formulas can be introduced into dynamic contexts is by applying `bcred`. Given a program clause $\Pi \bar{x}. G_1 \Rightarrow \dots \Rightarrow G_n \Rightarrow A$, the dynamic context of $\mathcal{H}(G_i)$ is obtained by extending the context of A with program clauses in $\mathcal{L}(G_i)$ for $1 \leq i \leq n$. Algorithm 1 derives a set of constraints \mathcal{C} on C based on this observation. It traverses all the program clauses and their sub-program clauses in formulas in Γ and collects a set \mathcal{C} of constraints for dynamic contexts which must be satisfied by derivations starting with Γ as the context.

To compute a set of dynamic contexts that satisfies \mathcal{C} , we start with $C(a) = \emptyset$ for all $a \in \Delta$ and iteratively apply the constraint equations until the constraint is satisfied. It is

Algorithm 1 Collecting constraints on dynamic contexts

Let Γ' be a finite set equal to Γ and $\mathcal{C} \leftarrow \emptyset$
while $\Gamma' \neq \emptyset$ **do**
 pick some $D = (\Pi \bar{x}. G_1 \Rightarrow \dots \Rightarrow G_n \Rightarrow A)$ from Γ'
 add equations $\{C(\mathcal{H}_p(G_i)) = C(\mathcal{H}_p(A)) \cup C(\mathcal{H}_p(A)) \cup \mathcal{L}(G_i) \mid i \in 1..n\}$ to \mathcal{C}
 remove D from Γ' and add clauses in $\bigcup_{i \in 1..n} \mathcal{L}(G_i)$ to Γ'
end while

Algorithm 2 Collecting constraints on the dependency relation

let Γ' be a finite set equal to Γ and $\mathcal{S} \leftarrow \emptyset$
for all $a \in \Delta$ **do**
 for all $D \in \Gamma' \cup C(a)$ where $D = (\Pi \bar{x}. G_1 \Rightarrow \dots \Rightarrow G_n \Rightarrow A)$ and $\mathcal{H}_p(A) = a$ **do**
 add $(S(a) = S(a) \cup \bigcup_{i \in 1..n} S(\mathcal{H}_p(G_i)))$ to \mathcal{S}
 end for
end for

easy to see that this algorithm terminates: since the iterations never shrink $C(a)$ for $a \in \Delta$, the only way the algorithm goes on forever is to keep adding new clauses in every iteration. This is impossible since there are only finitely many distinct program clauses. In the end, we get a finite set of dynamic clauses $C(a)$ for every $a \in \Delta$ that satisfies constraints in \mathcal{C} . Suppose $C(a) = \{D_1, \dots, D_n\}$ for $a \in \Delta$. We define the context relation ctx_a as follows:

$$\text{ctx}_a \text{ nil} \triangleq \top; \quad \text{ctx}_a (D_1 :: L) \triangleq \text{ctx}_a L; \quad \dots \quad \text{ctx}_a (D_n :: L) \triangleq \text{ctx}_a L.$$

6.2.2 Generating the dependency relation between predicates

By the `bcred` rule, given a program clause $\Pi \bar{x}. G_1 \Rightarrow \dots \Rightarrow G_n \Rightarrow A$, the provability of A will depend on the provability of $\mathcal{H}(G_i)$ for $1 \in 1..n$. For any $a \in \Delta \cup \{\mathcal{H}_p(A)\}$, a can be proved by backchaining on either some clause in Γ or in the dynamic contexts $C(a)$. Since both Γ and $C(a)$ are known and finite, we can derive a set $S(a)$ containing predicates that a depends on.

The steps for computing S are similar to that for C . First, we derive a set of constraints \mathcal{S} , by Algorithm 2. To generate the dependency relations, we start with $S(a) = \{a\}$ for all $a \in \Delta \cup \{\mathcal{H}_p(A)\}$ and iteratively apply the constraint equations until the constraints in \mathcal{S} are satisfied. The algorithm terminates by an analysis similar to the previous one.

6.2.3 Constructing proofs for the strengthening lemmas

Now we are in a position to prove the strengthening theorem $\forall \bar{x}. \{\Gamma, F \vdash A\} \supset \{\Gamma \vdash A\}$ in \mathcal{G} . Since the proof of A may depend on formulas with different heads, we generalize the strengthening lemma to take into account of related predicates.

► **Theorem 3.** *Let $S(\mathcal{H}_p(A)) = \{a_1, \dots, a_n\}$, The generalized strengthening lemma is*

$$(\forall \Gamma', \bar{x}_1. \text{ctx}_{a_1} \Gamma' \supset \{\Gamma, \Gamma', F \vdash a_1 \bar{x}_1\} \supset \{\Gamma, \Gamma' \vdash a_1 \bar{x}_1\}) \wedge \dots \wedge$$

$$(\forall \Gamma', \bar{x}_n. \text{ctx}_{a_n} \Gamma' \supset \{\Gamma, \Gamma', F \vdash a_n \bar{x}_n\} \supset \{\Gamma, \Gamma' \vdash a_n \bar{x}_n\})$$

If $\mathcal{H}_p(F) \notin S(\mathcal{H}_p(A))$, then this lemma has a proof in \mathcal{G} .

Proof. By simultaneous induction on $\{\Gamma, \Gamma', F \vdash a_i \bar{x}_i\}$ for $1 \leq i \leq n$. ◀

The original strengthening lemma is an immediate corollary of Theorem 3:

► **Corollary 4.** *If $\mathcal{H}_p(F) \notin S(\mathcal{H}_p(A))$, then $\forall \bar{x}. \{\Gamma, F \vdash A\} \supset \{\Gamma \vdash A\}$ has a proof in \mathcal{G} .* ◀

Since the proofs for Theorem 3 is constructive, we obtain a certifying algorithm to state and prove strengthening lemmas in \mathcal{G} on demand.

6.3 Application

This section describes an application of the formalization of independence: to prune unnecessary dependence on variables in \mathcal{G} . When analyzing the terms containing nominal constants in \mathcal{G} , it is often necessary to introduce variable dependencies on such constants. As an example, the \forall introduction rule creates a new eigenvariable that is raised over the nominal constants in the principal formula. However, if we can show that the type of the eigenvariable is independent from the types of the nominal constants, then we can suppress this dependency. In Abella, the theorem prover based on \mathcal{G} , an ad-hoc algorithm based on checking of subordination relations is used to prune such dependencies, but the exact logical basis of the algorithm has never been adequately formulated.

Now that we can derive independence lemmas, we can recast this pruning of nominal constants in terms of the closed-world reading of types. To give it a formal treatment in the logic, we reflect the closed-world reading of types into the proofs. This can be done by encoding the type theory of HH (which is STLC) into an HH specification as described in Section 4 and then to use the derived strengthening lemmas directly.

As an example, suppose we want to prove the following theorem

$$\forall X, T. \text{name } X \supset \{\text{of } X T\} \supset \dots$$

where X has type \mathfrak{tm} and T has type \mathfrak{ty} . By introducing the assumptions and case analysis on $\text{name } X$, X will be unified with a nominal constant n and the dependence of T on n will be introduced, resulting in the hypothesis $\{\text{of } n (T n)\}$. This dependence of T on n is vacuous as we have already seen – types cannot depend on terms in STLC. However, to formally establish the independence, we require T to be a well-formed type, so we need to change the theorem to

$$\forall X, T, L. \text{ctx } L \supset \{\Gamma, L \vdash \hat{\mathfrak{ty}} T\} \supset \text{name } X \supset \{\text{of } X T\} \supset \dots$$

where Γ contains the program clauses described in Figure 2. For the purpose of demonstration, we assume the context definition ctx contains infinitely many nominal constants of type $\hat{\mathfrak{tm}}$, which is defined as follows:

$$\text{ctx nil} \triangleq \top; \quad (\forall x. \text{ctx } (\hat{\mathfrak{tm}} x :: L)) \triangleq \text{ctx } L.$$

We perform the same introduction and case analysis on $\text{name } X$ and get hypotheses $\text{ctx } (L n)$, $\{\Gamma, L n \vdash \hat{\mathfrak{ty}} (T n)\}$ and $\{\text{of } n (T n)\}$. By the definition of ctx , when treated as a multiset, $L n$ must be equivalent to $(\hat{\mathfrak{tm}} n :: L')$ where L' does not contain n . Thus $\{\Gamma, L', \hat{\mathfrak{tm}} n \vdash \hat{\mathfrak{ty}} (T n)\}$ holds. At this point, we can use the algorithm in Section 6.2 to derive and prove the following strengthening lemma

$$\forall T. \forall x. \{\Gamma, L', \hat{\mathfrak{tm}} x \vdash \hat{\mathfrak{ty}} (T x)\} \supset \exists T'. T = (\lambda y. T') \wedge \{\Gamma, L' \vdash \hat{\mathfrak{ty}} T'\}.$$

By applying it to $\{\Gamma, L', \hat{\mathfrak{tm}} n \vdash \hat{\mathfrak{ty}} (T n)\}$, we get $T = \lambda y. T'$ for some T' not containing y . Hence $\{\text{of } n (T n)\}$ becomes $\{\text{of } n T'\}$. Note that we choose a particular definition of ctx for demonstration; the exact context of well-formed terms can vary in practice, based on the signature. Nevertheless, pruning of nominal constants can always be expressed by proving and applying the strengthening lemmas derived from independence.

7 Related Work

The earliest formulation of dependency that we were able to find is by Miller [9, Definition 11.3], where it is defined in terms of derivability of a certain fragment of strict logic for use in a unification procedure. Although we do not have a proof of this, this notion appears to coincide with the negation of our notion of independence. It is unclear why this definition was never adopted in subsequent work on logical frameworks, but we can speculate that one reason is its inherent cost, since it requires proving an arbitrary theorem of relevant logic. Indeed, using independence in the core unification engine is probably inadvisable if the unification engine is to be trusted.

It is more popular to define subordination relations independently of proof-theory, as has been done for many variants of LF. In [15], Virga proposed a *dependence relation* between types and type families in LF to constraint higher-order rewritings to well-behaved expressions. Later, this relation was popularized as *subordination* and used in the type theory of canonical LF to show that canonical terms of one type τ is not affected by introduction of terms of another type that is not subordinate to τ [17, 7]. The subordination relations in these cases are defined to be strong enough so that the type theory only deals with well-formed instances.

For reasoning applications, in order to move from one context to another, it is often necessary to check if a term of type τ_1 can occur in the normal form of another type τ_2 . Traditionally, the complement of subordination has been thought to be the right interpretation of independence. However, it is unclear how exactly it can be translated into evidence in the theory that supports the reasoning. Thus, *ad hoc* algorithms have been developed in systems like Twelf [11], Beluga [12] and Abella [18], which all lack formal definitions.

8 Conclusion and Future Work

We proposed a proof-theoretic characterization of independence in a two-level logic framework, and gave an example of such characterization by encoding the type theory of STLC in the logic HH and interpreting the independence relation as strengthening lemmas in a reasoning logic \mathcal{G} . We developed an algorithm to automatically establish the independence relation and strengthening lemmas and showed its application to pruning variable dependence in \mathcal{G} .

Interpreting independence as strengthening should be realizable in other logical frameworks that support inductive reasoning. We chose the two-level logic framework because it provides a first-class treatment of contexts, which makes proofs of strengthening lemmas easy. It would be worth investigating a similar formal development in logical frameworks such as Twelf and Beluga where contexts are either implicit or built into the type system.

The characterization of independence can be extended to more sophisticated type theories. Recently, it was shown that the encoding of LF in \mathcal{G} [13] can be used transparently to perform inductive reasoning over LF specifications [14]. We plan to develop a characterization of independence of LF based on this approach, which will formalize the important concept of *world subsumption* for migrating LF meta-theorems between different LF context schema.

Finally, one benefit of a logical characterization that is almost too obvious to state is that it opens up independence to both external validation and user-guidance. In LF, where inhabitation is undecidable, the notion of independence proposed in this paper will generally not be automatically derivable. Presenting the user with unsolved independence obligations may be an interesting interaction mode worth investigating.

Acknowledgements. This work has been partially supported by the NSF Grant CCF-0917140 and the ERC grant ProofCert. The first author has been partially supported by the

Doctoral Dissertation Fellowship from the University of Minnesota, Twin Cities. Opinions, findings, and conclusions or recommendations expressed in this paper are those of the authors and do not necessarily reflect the views of the National Science Foundation, the University of Minnesota, or Inria.

References

- 1 David Baelde, Kaustuv Chaudhuri, Andrew Gacek, Dale Miller, Gopalan Nadathur, Alwen Tiu, and Yuting Wang. Abella: A system for reasoning about relational specifications. *Journal of Formalized Reasoning*, 7(2), 2014.
- 2 Amy Felty and Dale Miller. Encoding a dependent-type λ -calculus in a logic programming language. In *Conference on Automated Deduction*, volume 449 of *LNAI*, pages 221–235. Springer, 1990.
- 3 Andrew Gacek. *A Framework for Specifying, Prototyping, and Reasoning about Computational Systems*. PhD thesis, University of Minnesota, 2009.
- 4 Andrew Gacek, Dale Miller, and Gopalan Nadathur. Nominal abstraction. *Information and Computation*, 209(1):48–73, 2011.
- 5 Andrew Gacek, Dale Miller, and Gopalan Nadathur. A two-level logic approach to reasoning about computations. *J. of Automated Reasoning*, 49(2):241–273, 2012.
- 6 Robert Harper, Furio Honsell, and Gordon Plotkin. A framework for defining logics. *Journal of the ACM*, 40(1):143–184, 1993.
- 7 Robert Harper and Daniel R. Licata. Mechanizing metatheory in a logical framework. *Journal of Functional Programming*, 17(4–5):613–673, July 2007.
- 8 Raymond McDowell and Dale Miller. Cut-elimination for a logic with definitions and induction. *Theoretical Computer Science*, 232:91–119, 2000.
- 9 Dale Miller. Unification under a mixed prefix. *J. of Symbolic Comp.*, 14(4):321–358, 1992.
- 10 Dale Miller and Gopalan Nadathur. *Programming with Higher-Order Logic*. Cambridge University Press, June 2012.
- 11 Frank Pfenning and Carsten Schürmann. System description: Twelf — A meta-logical framework for deductive systems. In *16th Conf. on Automated Deduction (CADE)*, number 1632 in *LNAI*, pages 202–206. Springer, 1999.
- 12 Brigitte Pientka and Joshua Dunfield. Beluga: A framework for programming and reasoning with deductive systems (system description). In *Fifth International Joint Conference on Automated Reasoning*, number 6173 in *LNCS*, pages 15–21, 2010.
- 13 Zachary Snow, David Baelde, and Gopalan Nadathur. A meta-programming approach to realizing dependently typed logic programming. In *ACM SIGPLAN Conference on Principles and Practice of Declarative Programming (PPDP)*, pages 187–198, 2010.
- 14 Mary Southern and Kaustuv Chaudhuri. A two-level logic approach to reasoning about typed specification languages. In *34th Int. Conf. on Foundations of Software Technology and Theoretical Computer Science (FSTTCS)*, volume 29 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 557–569, December 2014.
- 15 Roberto Virga. *Higher-order Rewriting with Dependent Types*. PhD thesis, Carnegie Mellon University, 1999.
- 16 Yuting Wang, Kaustuv Chaudhuri, Andrew Gacek, and Gopalan Nadathur. Reasoning about higher-order relational specifications. In *15th International Symposium on Principles and Practice of Declarative Programming (PPDP)*, pages 157–168, September 2013.
- 17 Kevin Watkins, Iliano Cervesato, Frank Pfenning, and David Walker. A concurrent logical framework I: Judgments and properties. Technical Report CMU-CS-02-101, Carnegie Mellon University, 2003. Revised, May 2003.
- 18 The Abella web-site. <http://abella-prover.org/>, 2015.