DAGSTUHL
REPORTS

Volume 7, Issue 6, June 2017

ISSN 2192-5283

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/2192-5283

Publication date
January, 2018

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`.

Digital Object Identifier: 10.4230/DagRep.7.6.i

Aims and Scope
The periodical *Dagstuhl Reports* documents the program and the results of Dagstuhl Seminars and Dagstuhl Perspectives Workshops.
In principal, for each Dagstuhl Seminar or Dagstuhl Perspectives Workshop a report is published that contains the following:

- an executive summary of the seminar program and the fundamental results,
- an overview of the talks given during the seminar (summarized as talk abstracts), and
- summaries from working groups (if applicable).

This basic framework can be extended by suitable contributions that are related to the program of the seminar, e. g. summaries from panel discussions or open problem sessions.

Contact
Schloss Dagstuhl – Leibniz-Zentrum für Informatik
Dagstuhl Reports, Editorial Office
Oktavie-Allee, 66687 Wadern, Germany
reports@dagstuhl.de
http://www.dagstuhl.de/dagrep

Report from Dagstuhl Seminar 17231

Epistemic Planning

Edited by

Chitta Baral[1], Thomas Bolander[2], Hans van Ditmarsch[3], and Sheila McIlraith[4]

1 **Arizona State University – Tempe, US,** chitta@asu.edu
2 **Technical University of Denmark – Lyngby, DK,** tobo@dtu.dk
3 **LORIA, CNRS – Univ. of Lorraine – Nancy, FR,** hans.van-ditmarsch@loria.fr
4 **University of Toronto, CA,** sheila@cs.toronto.edu

Abstract

The seminar Epistemic Planning brought together the research communities of *Dynamic Epistemic Logic*, *Knowledge Representation and Reasoning*, and *Automated Planning* to address fundamental problems on the topic of epistemic planning. In the context of this seminar, dynamic epistemic logic investigates the formal semantics of communication and communicative actions, knowledge representation and reasoning focuses on theories of action and change, and automated planning investigates computational techniques and tools to generate plans. The original goals of the seminar were to develop benchmarks for epistemic planning, to explore the relationship between knowledge and belief in multi-agent epistemic planning, to develop models of agency and capability in epistemic planning and to explore action types and their representations (these originally separate goals were merged during the seminar), and finally to identify practical tools and resources. An additional goal explored during the workshop was the correspondence between planning problems and games.

Seminar June 5–9, 2017 – http://www.dagstuhl.de/17231
1998 ACM Subject Classification F.4.1 Mathematical Logic, I.2.4 Knowledge Representation Formalisms and Methods, I.2.8 Problem Solving, Control Methods, and Search, I.2.11 Distributed Artificial Intelligence
Keywords and phrases Automated Planning, Knowledge Representation and Reasoning, Reasoning About Actions, Dynamic Epistemic Logic, Multi-Agent Systems
Digital Object Identifier 10.4230/DagRep.7.6.1

1 Executive Summary

Chitta Baral
Thomas Bolander
Hans van Ditmarsch
Sheila McIlraith

This seminar brought together three largely independent research communities: Dynamic Epistemic Logic (DEL), Knowledge Representation and Reasoning (KR&R, subsequently KR) and Automated Planning. All three communities have a tradition of investigating the interaction between dynamical systems and epistemic states, but with a different focus and by different means. In the context of this seminar, despite occasional overlap, DEL has mainly investigated the formal semantics of communication and communicative actions,

KR&R has mainly focussed on the theories of action and change, and AP has mainly focussed on computational techniques to automatically generate plans. This seminar aimed to encourage and nurture increasing synergies between these three strong and largely independent research communities leading to frameworks for Epistemic Planning: planning with epistemic states, actions, and goals. The seminar succeeded in strengthening and broadening the cross-disciplinary research community that emerged after a prior Dagstuhl Seminar on the subject, entitled "Planning with Epistemic Goals", that was held in January 2014 (seminar number 14032). This follow-up seminar led to a better understanding and articulation of commonalities, synergies, and deficits between the DEL, KR and Planning communities.

The main components of the seminar were tutorial talks and group work. There were four tutorials, by Andreas Herzig, Bernhard Nebel, Tran Cao Son and Ramaswamy Ramanujam. The four tutorial talks presented epistemic planning from the perspective of four different research communities. Andreas Herzig presented epistemic planning from a knowledge representation perspective, Bernhard Nebel from a classical planning perspective, Tran Cao Son from a theories of action and change perspective, and Ramaswamy Ramanujam from a distributed systems and temporal logic perspective. Abstracts for the four tutorials are included in Section 3.

The group work was performed in five separate working groups, each addressing a distinct important objective in epistemic planning of shared interest to the DEL, KR and Planning communities. Specifically, the five groups worked on (1) *Developing Benchmarks for Epistemic Planning*, (2) *Exploring Action Types and their Representations*, (3) *Exploring the Relations Between Knowledge and Belief in Multi-Agent Epistemic Planning*, (4) *Practical Tools, Resources and Computational Techniques*, and (5) *Correspondence Between Planning Problems and Games*. These group themes are briefly described below, and the outcome of the work in each group is documented in Section 4.

1. **Developing Benchmarks for Epistemic Planning.** In the planning community benchmarks are common, but in the epistemic planning community they are not. The overall goal of the group working on this theme was to formulate a list of ten benchmark epistemic planning problems, in view of setting goals that can evolve into competitions. In particular, the focus was on targeting planning problems that are truly epistemic, meaning problems in which the epistemic dimension – knowledge and ignorance – cannot easily be disregarded.

   Guiding questions for the work in this group were: What problems help define and circumscribe what we are studying? What are specific tasks that motivate this area of study, e.g. epistemic planning, protocol synthesis, automated diagnosis, verification, and communication? What problems can help drive future research in developing formalisms and implementing systems of epistemic planning? What are the features of the relevant planning problems in terms of knowledge vs. belief, single- vs. multi-agent, communicative vs. sensing vs. ontic actions, deterministic vs. non-deterministic actions, etc. How do we evaluate "hardness" of problems, e.g. in terms of level of nesting of belief/knowledge, types of actions, size of problem, scalability, and quality of solution. Do benchmarks for some of these problems already exist?
2. **Exploring Action Types and their Representations.** It is important to identify and broaden the list of action types relevant to epistemic planning. In terms of communicative actions we can for instance at least distinguish between announcements, questions, requests and instructions. How these actions are best represented is also an important issue. It should be explored whether formalizations in dynamic epistemic logic are appropriate for planning or whether a more simplified way of representing multi-agent actions is needed.

The overall goal of the group working on this theme was to identify, classify and possibly broaden the list of action types relevant to epistemic planning, as well as to explore formalisms for representing these action types.

Guiding questions for the work in this group were: What are the relevant distinctions between types of actions, e.g. epistemic vs. ontic, deterministic vs. non-deterministic vs. probabilistic, instantaneous vs. durative, sensing vs. announcements, degree of observability (public, private, semi-private), etc. What formalisms can support expressing and distinguishing between these action types (e.g. DEL, Situation Calculus, Knowledge-based Programs)? How are multi-agent actions represented, in particular how are conflicts between concurrently occurring actions specified, and how is observability of concurrent actions specified in terms of the observability of the constituting actions?

3. **Exploring the Relations Between Knowledge and Belief in Multi-Agent Epistemic Planning.** In multi-agent planning an important problem is that knowledge may turn into false belief for some agents after a partially observable action has taken place. This is a problem for several formalisms for epistemic planning, e.g. dynamic epistemic logic, since agents might not be able to recover from false beliefs. It relates to the general issues of devising appropriate formalisms for doxastic planning (treating beliefs instead of knowledge) and how to deal with belief revision in such settings. The overall goal of the work in this group was to identify the theoretical and computational challenges in planning with knowledge vs. planning with belief; when one or the other is appropriate, or both are needed.

   Guiding questions for the work in this group were: How are knowledge and beliefs represented and distinguished from representations of the actual world? How do we formally handle that knowledge may turn into false belief after a partially observable action has occurred? What are the relevant formalisms for planning with knowledge and/or belief and what are their theoretical and computational properties? How do we deal with belief revision in planning? Are there specific types of interesting goals for epistemic planning? For example, in planning with beliefs, goals can be about making some agents have false beliefs. This necessitates formalizing false-belief tasks, lying and deception.
4. **Practical Tools, Resources and Computational Techniques.** The goal of the group working on this theme was to identify practical tools and resources that facilitate the development and experimental evaluation of automated techniques for epistemic planning.

   Guiding questions for the work in this group were: What tools and computational techniques already exist in epistemic planning? What are the models and formulas used? What are the shortcomings and challenges of these tools and computational techniques? Are there tools or computational techniques from other communities that we are not availing ourselves of to the fullest extent (e.g., for the planning people model checking)? What are the trade-offs between different tools? What are the computational complexities of the different approaches and under different assumptions.
5. **Correspondence Between Planning Problems and Games.** This working group was introduced as an additional discussion topic during the seminar, since several participants found it very relevant and important to epistemic planning. The issue is that (epistemic) game theory studies many of the same problems as (epistemic) planning, but mainly by separate research communities using separate vocabularies. The goal of the group working on this theme was to establish formal connections between the area of automated planning and the area of game theory.

Guiding questions for the work in this group were: What are the formalisms, tools and results from game theory relevant to automated planning? Symmetrically, what are the formalisms, tools and results from automated planning relevant to game theory? Can problems formulated in one of the settings easily be translated into the other? What is gained and lost in such translations?

2 Table of Contents

3 Overview of Talks

3.1 Epistemic Planning: A Knowledge Representation Perspective

Andreas Herzig (Paul Sabatier University, Toulouse, FR)

I start by a brief historical review of knowledge representation approaches to planning. I then argue that the integration of reasoning about knowledge requires a thorough analysis of the main concepts and models that are involved, viz. initial situation, goal, and action laws. While epistemic logic provides a satisfactory account of the former two, the representation of action laws deserves a closer look. I focus on the dynamic epistemic logic account in terms of event models and point out some conceptual and practical difficulties, including a mismatch between action tokens as modelled in the former on the one hand, and action types as needed for action laws on the other.

3.2 Epistemic Planning Based on DEL for Researchers Working on Classical Planning

Bernhard Nebel (Universität Freiburg, DE)

Joint work of Bernhard Nebel, Thorsten Engesser, Robert Mattmüller

After a quick introduction of what states and actions in DEL are, we show how classical planning, fully observable non-deterministic (FOND) planning, conformant planning, and partially observable non-deterministic (POND) planning can be implemented using DEL. On top of POND planning, we introduce the notions of multi-agent planning, non-uniform observability, and implicitly coordinated plans, and demonstrate how these notions are modelled in DEL. The final part of the tutorial highlights strength and weaknesses of DEL planning by presenting three small examples: Nano-Hanabi, multi-agent path finding with destination uncertainty, and secretly travelling to Dagstuhl.

3.3 KRR and Epistemic Planning: Specification and Implementation Issues

Tran Cao Son (New Mexico State University, US)

In this tutorial, we discuss specification and implementation issues of epistemic planning. We start with a short motivation of the epistemic planning problem. We then discuss the issues and how they have been addressed in single agent environments. Specifically, we describe an action language formalism and present possible approaches for planning with incomplete information and sensing actions. We highlight the advantage of approximation-based approaches to epistemic planning in single agent domains. We continue with a discussion of various issues of the problem formulation in multi agent domains. We present

an action language, $m\mathcal{A}^+$, and the notion of finitary S5-theory whose combination with $m\mathcal{A}^+$ could be used for the specification of epistemic planning problems. Additionally, we show that a large class of $m\mathcal{A}^+$ theories maintains the KD45-property of the initial epistemic state, a desirable feature when one needs to distinguish between knowledge and beliefs. We conclude with a list of research issues related to the formalization and implementation of epistemic planning systems.

3.4 Distributed Presentations of Multi-Agent Systems

Ramaswamy Ramanujam (Institute of Mathematical Sciences, Chennai, IN)

In Dynamic Epistemic Logics (DEL) and Epistemic Temporal Logics (ETL) the modelling of multi-agent systems tends to be 'from above', describing epistemic ascriptions to agents by the designer. In distributed protocols we looks for 'inside' or local presentations of agents, where agents have partial views of system evolution and acquire knowledge about others by explicit mechanisms such as shared variables or message passing. We present a framework of systems of epistemic finite state automata that explicitly acquire and update knowledge tokens specified by a separate (pre-ordered) alphabet. These tokens can be thought of as finite representations of sets of knowledge formulas (ordered by implication). We also present an ETL built on local propositions where structural composition ensures that epistemic assertions by agents about others refer to knowledge constructed by agents based on their partial views. For every consistent formula we construct a system of automata that realizes the models of the formula (as a set of partial orders). There are refinement methodologies and theorem prover based synthesis of such behaviours as well. In general, given a global presentation of an epistemic multi-agent system, decomposing it into a system of communicating agents seems to be interesting and challenging.

4 Working groups

4.1 Developing Benchmarks for Epistemic Planning

Malvin Gattinger, Fillippos Kominis, Jérôme Lang, Robert Mattmüller, Tim Miller (coordinator), Ron Petrick, François Schwarzentruber, Bruno Zanuttini

We distinguish between *long term applications*, which can hardly be expressed by simplified examples, *examples of problems*, that can more easily be formalized and are good candidates for developing benchmarks, and *benchmarks* stricto sensu, that are parameterized families of simplified examples (such as puzzles) and that are primarily meant to test and compare the different approaches for verifying and generating plans.

4.1.1 Applications

We classify here potential long-term applications of epistemic planning.

One major application is the design of *cooperative agents*, who share the same utility function and that are generally programmed in a centralized fashion (their plans are designed off-line, but of course the online execution of these plans will be distributed and perhaps asynchronous). As an example: multirobot exploration of a nuclear plant (modelled by decentralized POMDP ([20])).

At the other extremity of this spectrum, we have *purely adversarial games*, where the utilities of the players sum to zero. This is typically the field of game playing; knowledge plays a role as soon as incomplete knowledge of the state of the game and/or communicative actions are present in the game. Two paradigmatic examples: *poker* and *bridge*.

Between these two extreme cases, we have classes of problems which are neither purely cooperative nor purely adversarial. A class of applications of crucial importance is the class of *assistive agent problems*, where an autonomous agent has to reason about the mental state of a human agent and assist them in performing various tasks, such as: robots helping disabled or elderly persons, tutorial systems and serious games, dialogue/interaction management (including interactive generation of explanations). The problems in this rich class of applications are "almost cooperative", in the sense that the utility of the helper should generally increase with the utility of the helped agent, but there are exceptions, especially when the helped agent is not fully rational (e.g., patients with Alzheimer). Another class of applications is where the environment has agents that can be partitioned to groups and one group of agents may collaboratively try to achieve something while keeping agents in the other groups in the dark or even making them have false beliefs.

Another important class of applications of epistemic planning is *security games*, where an autonomous agent or a collection of autonomous agents have to reason about the mental states of an adversary and to prevent it from performing malevolent actions, such as organization of airport security, phishing detection, reasoning about a malware that could detect a sandbox (e.g. WannaCry), etc.

The difference with games in the 'game playing' category is that there is a safety/maintenance goal, agents are not symmetric, the airport's actions are only checking properties, etc.

4.1.2 Problems as Benchmarks

In this section, we provide a list of examples that could be potentially transformed into benchmarks. Indeed, the applications presented in the previous section are too informal to lead to benchmarks. However we keep the same classification.

Purely cooperative problems

Let us start with a list of purely cooperative problems, that is, in which there is no adversary.

- Tiger [49]
- Hanabi (recently proven to be NP-hard [6])
- Active Muddy Children (one active child, asking questions with the goal of knowing its own state) [57, p. 152]
- Multirobot exploration of a nuclear plant (modeled by decentralized POMDP [20])
- 100 Prisoners and a Lightbulb [108]
- Word Rooms (see [57])
- Corridor (see [57])
- Gossip (see below and [46])
- Grapevine (combination of Corridor and Gossip, see [74])
- Dining Cryptographers [30]

Security games

In the following examples, several agents cooperate for an adversary not to perform bad actions and not to obtain certain knowledge.

- Russian cards
- Scotland Yard
- Sandbox and detection thereof (e.g. WannaCry)
- Organization of airport security

There are key differences with general games. First, the goal is essentially a safety/maintenance condition. Second, agents are not symmetric in the sense that they do not have the same role at all. For instance, in the organization of an airport security, the airport's actions are checking properties only while other agents have physical actions.

Adversary game playing

The following list provides examples where agents are programmed to react against other agents in a non-cooperative imperfect information game.

- Clue (Cluedo)
- Mobile battleships
- Military strategy
- Magic tricks [86]

Multi-agent deception

This scenario can be illustrated with respect to a simple example[1] taken from the storybook, *Winnie the Pooh*. At one point in the story, Rabbit, Pooh and Piglet are forming a joint plan to distract Kanga's attention so that they can substitute Piglet for Roo so that Kanga will not notice that Roo is gone so that Rabbit can run off with Roo.

In the above example, the goal of Rabbit, Pooh and Piglet is not only to substitute Piglet for Roo but also to keep Kanga in the dark whereby Kanga has false beliefs.

Features to evaluate

Each problem is characterized by the following features:

- *Knowledge vs. belief*: Is the problem one of true knowledge, of (possibly incorrect) belief, or both? Typically, knowledge involves indistinguishability between possible situations modeled as equivalence relations over possible worlds, one for each agent. On the contrary, (possibly incorrect) belief is modeled by means of so-called KD45 relation (see [31]). Note also that belief may require complex mechanisms such as belief revision [42].
- *Single vs. multi-agent*: Is the problem a multi- or single-agent problem?
- *Communication vs. sensing vs. ontic*: Are actions ontic (effect only the environment), sensing (only sense the environment and have no effect), or complex form of communication (effect the mental states of other, etc. as for action models of dynamic epistemic logic [9]) actions, or a mix of these?
- *Determinism/reliability of communication*: Are actions deterministic or non-deterministic, in which the latter includes concepts such as unreliable sending and communication?
- *Epistemic goals*: Are goals epistemic, ontic, or both?

[1] A version of this example was presented by Ernie Davis in his invited talk at the ACS 2017 conference.

- *Cooperative/competitive*: Is the problem cooperative, adversarial/competitive, or both; or potentially neither (other agents as 'noise' in the environment)?
- *Benchmarks*: Are there already existing benchmarks that can be accessed?

4.1.3 Application-Inspired Benchmarks

4.1.3.1 Gossip

Description: The Gossip problem is a communication problem in which agents must share secrets with each other; thus goals are epistemic. The actions available to each agent are to share a secret with another agent, as in a telephone call. In the simplest version, the goal is that all agents know each others' secrets. However, more sophisticated goals are possible:

- Everybody knows all secrets
- All secrets are common knowledge
- Deception (of some subset of agents): some agents believe the opposite of some secrets
- Privacy-preserving: Some agents know all secrets, while some can only know part of the secrets. For example, a single source agent knows a set of secrets and wants a destination agent to know all of these, but can communicate only via other agents, to whom the secret should not be revealed. Thus, it must communicate some secrets via some paths and other secrets via other paths.

Similarly, the *nature* of the secrets can be more expressive than simple propositions, such as:

- s_i (just the secret)
- $K_j s_i$ (secret + epistemic info)
- histories (secret + when I got it + from whom)
- how much does the agent say per turn? (one secret, all secrets, some secrets, everything that agent knows)
- unilateral or not?
- negotiation of secrets? values of secrets?

Figure 1 Gossip problem (Taken from Faustine Maffre's thesis.) [67]

Different variations of the gossip protocols have been investigated [109, 3, 106].

Motivation: The Gossip problem has the nice properties that one would like for a running example in a paper, but can be easily made more challenging by increasing the number of agents, number of secrets, nature of the goals, nature of the secrets, and the structure of the network in which the secrets are shared. Thus, it forms an excellent benchmark problem.

Further, variants of the problem, parameters, and complexity have been well studied [34, 50, 107].

Features:

- *Knowledge vs. belief*: Either.
- *Single vs. multi-agent*: Multi-agent.
- *Communication vs. sensing vs. ontic*: Communication.
- *Determinism/reliability of communication*: Reliable.
- *Epistemic Goals*: Yes.
- *Cooperative/competitive*: Cooperative, but adversarial versions possible.
- *(De)Centralized*: Centralized.
- *Benchmarks*: Different implementations available, such as at Toulouse/IRIT[2] and Melbourne/Toronto[3].
- *Evaluation of hardness*: Cooper et al. [34] for optimality, levels of nesting of belief and knowledge, size of problems, and scalability.

4.1.3.2 Security and patrolling games

Description: In security games, the aim is to find a strategy for a defender to deter attackers from reaching their targets. Usually, the defender's strategy is a mixed strategy which is a best response to what the defender believes to be the attacker's strategy. Patrolling games is a specific subfamily of security games where actions consist in performing some security checks at some points at each time step. Depending on the type of game, the nature of the attackers and the cost involved, the techniques used and the rationality assumptions about the attacker largely differ. Some examples of high-stake security games are protecting airports against terrorist attacks [95] ; some middle-stake games include "green security games", such as protecting the environment against smuggling and poaching; low-stake domains include deterring fare evasion in public transport.

Motivation: Programming a defender's strategy can be seen as an adversarial epistemic planning problem. The defender has to reason about the opponent's intended actions and about his higher-order beliefs (his beliefs about the defender's strategy and about the defender's beliefs about his own strategy). The base version consists in casting the problem as a Stackelberg game: a defender must defend a set of targets using a limited amount of resources whereas the attacker can observe and learn the (mixed) strategy of the defender. In this base model, the payoffs of the two agents depend only on the target attacked and whether or not it was protected by the defender. More sophisticated versions have been developed, including versions where the attacker can only partially observe the attacker's actions and/or where his rationality is bounded.

Features:

- *Knowledge vs. belief*: Belief (with deception).
- *Single vs. multi-agent*: Multi-agent.
- *Communication vs. sensing vs. ontic*: Sensing and ontic.
- *Determinism/reliability of communication*: Either.
- *Epistemic Goals*: Usually not, but possible (maintenance goal to always know the position of the attacker, for instance).
- *(De)Centralized*: Centralized.

[2] See https://github.com/FaustineMaffre/GossipProblem-PDDL-generator.
[3] See https://bitbucket.org/haz/pdkb-planning.

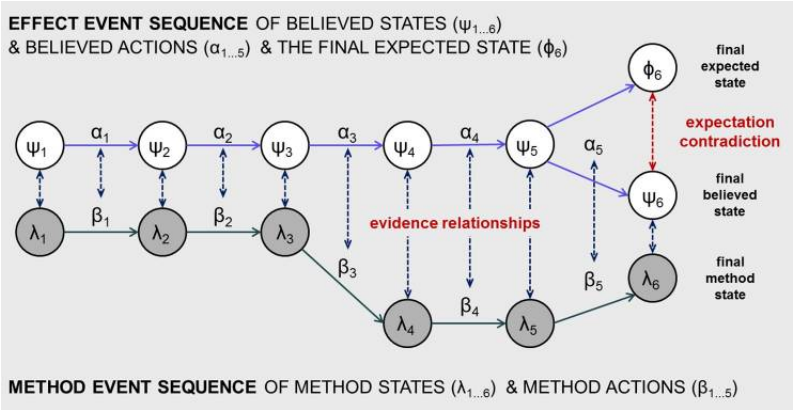

Figure 2 Parallel processes involved in conjuring/magic.

- *Cooperative/competitive*: Competitive (adversarial).
- *Benchmarks*: Existing problem description from Tambe [95]; see also follow-up work of *SimPatrol: Establishing a Testbed for Multi-agent Patrolling*).

4.1.3.3 Conjuring Magic Tricks

Description: The aim of these problems is to generate an element of 'surprise' in a group observers. A magic trick consists of a set of actions, in which the state of the world is partially observable to some or all of the observers. The aim of the magician is to transition the real state of the world into some state that is different from that expected by some of the observers. Figure 2 presents an abstraction of this, in which two parallel processes occur: the effect sequence, which is what the observer believes has occurred, and the method sequence, which is what the magician knows has occurred. At the end, there are three states: the *expected* state of the observer, the *believed* state of the observer, and the *method* state, which is the real state of the affairs only known by the magician.

As a first step, we could aim to discover the correct sequences for existing magic tricks, by specifying action models describing these, and using an epistemic planner to try to discover the correct set of moves. As a second step, epistemic planning could be used to find new tricks using existing or modified actions models.

Smith et al. [86] define a formalism of conjuring using propositional dynamic logic (PDL), as well as examples of some tricks using this formalism.

Motivation: Being able to derive magic tricks, even in a limited setting, is a somewhat creative task, and something that explicitly requires a Theory of Mind of the observers. In addition, the goals are *temporal epistemic goals* – the true state of the world must remain hidden until such time that the observer's perspective is incorrect.

Higher-order beliefs also play a part in some tricks; for example, in tricks involving more than two people, in which the magician has a helper for the trick, and thus needs to track not only the perspectives that the helper and 'victim' have of the world, but also of each other and the magician.

Further to this, generating surprise is key to many tactical scenarios, such as game playing, military planning, and story telling.

These are interesting and fun problems to solve, so this may be a good way to attract masters and internship students to work on epistemic planning.

Figure 3 Cover of Hanabi, © Asmodée éditions.

Features:

- *Knowledge vs. belief*: Belief (with deception).
- *Single vs. multi-agent*: Multi-agent.
- *Communication vs. sensing vs. ontic*: All three.
- *Determinism*: Non-deterministic, or probabilistic (uncertain whether participants will infer certain beliefs).
- *(De)Centralized*: Decentralized.
- *Epistemic Goals*: Temporal epistemic goals.
- *Collaborative/competitive*: Both (collaborative helper, competitive).
- *Benchmarks*: Existing model of conjuring from Smith et al. [86].

4.1.3.4 Hanabi

Description: Hanabi is a fully collaborative card game[4]. The game consists of cards each with a given colour and number. Each player has a hand drawn from the deck, which they cannot see, while they can see the others' hand. The cooperative goal is to stack up cards of each colour in sequential order. For this, on their turn, players can either (i) play a card from their hand (and draw a new one), (ii) discard a card and draw a new one, or (iii) give a clue to another player about what cards they have. A clue is the indication of all cards of a given colour, or of all cards of a given number (e.g., "the blue cards in your hand are the first and fourth ones".

When a card is played, it must be either the first of its colour played by any player, or the one numbered immediately after the last one played. Otherwise, the move is a mistake; if a predefined number of mistakes is reached, the players lose the game. Moreover, a maximal number of clues to be given by anyone is predefined. Giving a clue uses one, and discarding a card gives one back. When the deck is empty (and players have not lost the game before), the score is computed as the sum of the highest number in each colour.

Motivation: Hanabi requires the player to maintain knowledge about their own cards (e.g., "I know that my second card is a 2", "I know that I don't have any red card", etc.) and to reason about what information is needed by their mates, so as to inform them while using as few tokens as possible. When playing it for real, players most often agree on a strategy about which information they will give to each other (e.g., "Always give information about what next card to play"), which can be seen as centralized planning (with decentralized execution of the strategy agreed upon). Of course, it also makes sense to plan in a decentralized or distributed fashion, which corresponds in real life to players starting a game without knowing each other.

[4] Created by Antoine Bauza and published in 2010.

An interesting feature of Hanabi is that it has a formal and simple description, easily amenable to description in an input language for epistemic planners. Moreover, by adjusting the numbers of players, colors, and numbers in each, it can be used for analyzing the scalability of algorithms. Other features can also be controlled; for instance, loosely speaking, the number of clue tokens is related to the number of communication actions which can be taken in a row, and the number of mistakes allowed is related to whether we need "cautious" plans.

Finally, the goal can be controlled by framing it as an optimization problem (maximize the score while not losing) or one of reachability (empty the deck before losing). Orthogonally, the goal can be defined taking into account the probabilities involved (maximize the expected score, maximize the expected time-to-defeat, etc.) or looking for a winning strategy in the adversarial sense.

Features:

- *Knowledge vs. belief*: Knowledge in the first place ("what am I/is my teammate sure of?"), but probabilistic knowledge ("I know that there is a 33 % chance that my second card is blue") or beliefs ("since all but one red card have been played, let me bet that I have none") make sense.
- *Single vs. multi-agent*: Multi-agent.[5]
- *Communication vs. sensing vs. ontic*: Communication (mostly) and ontic (but these are mostly dictated by the information which the player has).
- *Determinism*: Stochastic because the order of the deck is unknown (and random), otherwise deterministic.
- *(De)Centralized*: Decentralized for the execution of the plan; naturally centralized for computation of a strategy but may be solved in decentralized or distributed fashion.
- *Epistemic Goals*: Mo, but reasoning with epistemic subgoals is clearly needed for playing optimally (first-order knowledge should be enough).
- *Collaborative/competitive*: Collaborative.
- *Benchmarks*: Game rules are already formal, and easy to encode; can be made small, tiny, nano, or huge by adjusting the number of colours, cards and players. AI implementations exist, see for instance https://github.com/Quuxplusone/Hanabi and https://github.com/WuTheFWasThat/hanabi.rs

4.1.3.5 Clue (Cluedo)

Description: It is an adversarial board game in which the goal of a player is to guess rightly the value of each of three variables: who committed the crime; where; and how. Players iteratively ask other players to privately tell them if they possess one card of a specified triple of values (one for each variable). The value of variable 'where' must correspond to the place where the player is located at the request time; players can move from one place to another subject to some constraints.

Motivation: Clue is perhaps the most well-known game with epistemic goals, and it possesses almost all the possible features of an epistemic planning problem. In order to rightly guess the three variables, a player exploits information from the cards that other agents show them,

[5] Interestingly, deciding whether there is a winning sequence of plays even with full information *and* a single player is already NP-complete [6].

Figure 4 Scotland Yard.

from the fact that an agent denies having anyone of a triple of cards, or from the fact that an agent shows a card to another agent, but also from what other agents ask. A good player should maintain probabilistic beliefs and choose questions maximizing expected information gain, but should also try to reason about the inferences that the other agents will draw from his questions.

Features:

- *Knowledge vs. belief*: Knowledge (acquired from non-noisy observations) but the players may perform some probabilistic reasoning.
- *Single vs. multi-agent*: Multi-agent.
- *Communication vs. sensing vs. ontic*: Communication and ontic.
- *Determinism/reliability of communication*: Movement resources and card drawing probabilistic, communication reliable.
- *Epistemic Goals*: Yes, and only that (guess the right triple).
- *Collaborative/competitive*: Competitive.
- *Benchmarks*: Do probably not exist, but have been discussed several times in the epistemic planning community. Parameters that can vary is the number of players, the number of variables, and the number of values per variable.

This problem and its relationship to epistemic reasoning has been discussed by van Ditsmarsch [105].

4.1.3.6 Scotland Yard

This board game is played on a map of London. A group of detectives is hunting a criminal using public transport. The position of the criminal is only revealed occasionally and detectives have to reason about possible paths the criminal could have taken. Additionally, resources for movement (bus tickets etc.) are limited for all players.

References and existing Implementations:

- Wikipedia: https://en.wikipedia.org/wiki/Scotland_Yard_(board_game)
- Spiel des Jahres 1983: http://spieldesjahres.com/de/scotland-yard.

- Scotland Yard is PSPACE: http://www.illc.uva.nl/Research/Publications/Reports/PP-2006-18.text.pdf.
- Monte-Carlo Tree Search for Scotland Yard: [76].
- Cooperative planning: https://eldorado.tu-dortmund.de/bitstream/2003/2653/1/86.pdf.

Features:

- *Knowledge vs. belief*: Both, there is hard (e.g. criminal is at certain position) and probabilistic (e.g. bus ticket was used) information.
- *Single vs. multi-agent*: Multi-agent, though this depends on model and implementation. If detectives use centralized planning then it is essentially a two-player game. For decentralized planning each detective should work on their own and communication between them might be limited, in particular to public announcements.
- *Communication vs. sensing vs. ontic*: Seen as a two-player game there are only ontic actions, if detectives are multiple agents then a range of communication can be allowed. The only sensing actions are when one can sense whether a criminal is in their sensing agent's current location.
- *Determinism and reliability of communication*: Probabilistic uncertainty is in the initial state only, all communications are reliable successful public announcements.
- *Epistemic Goals*: The goal is not explicitly epistemic but just to be at the same position. But the implied and intermediate goal is to know where the criminal is.
- *Collaborative/competitive*: both: 1 vs n players.
- *Benchmarks*: Different planning algorithms and their implementations could play against each other. Alternatively, a tool could assist the detectives.

4.1.3.7 Bitstring

Description: *Bitstring* is a contingent problem of communication where we have a string of bits, two agents that have the same uncertainty concerning the actual truth values of the bits in the string and each agent can sense only a subset of the bits, disjoint between them.

The simple version of the problem is the one where the goal is for both agents to know the actual truth values of all the bits in the string.

Formally, we have two agents $A = \{A_1, A_2\}$, we have a set B of bits, $B = \{ b_1, b_2, ..., b_n \}$, which denotes a set S of possible states, where $|S| = 2^n$ and a hidden, true state s^*. It is common knowledge among the two agents that only a set S' is actually possible, where $S' \subset S$. Allowed actions are communication actions of the knowledge of each agent concerning the truth value of each bit ($K_i b_j / \neg b_j$, for $i \in A$ and $j = 1...n$). The goal then is $G = \bigwedge_{j=1...n}(K_i b_j \vee K_i \neg b_j)$, for $i \in A$.

The simple version of the problem allows for simple plans where, for example, we can have an agent deriving that he does not need to communicate the truth value of a certain bit since he knows that based on the knowledge the other agent has acquired through his observations, he already knows the truth value of the specific bit.

A more interesting view on the *Bitstring* is its extended version: Assume a third agent A_3 who is not aware of S' but only of S – at the same time he is able to listen to all communication between the two agents A_1 and A_2 and he knows what A_1 and A_2 have sensed. Now, the question becomes: is there a plan which achieves the goal $G_s = \bigwedge_{j=1...n}(K_i b_j \vee K_i \neg b_j) \wedge \bigwedge_{j=1...n}(\neg K_{A_3} b_j \wedge \neg K_{A_3} \neg b_j)$. In other words, we are checking whether there is a plan where the the first two agents know the truth values of all bits, while the third one learns nothing.

This problems allows for a number of approaches:

- Checking for existence of such a plan.
- Treat the second part of the goal as a soft goal: the first two agents must learn the truth value of all bits and, at the same time, try to minimize the knowledge the third agent acquires.
- Increase the number of agents and/or restrict the set of agents that each one can communicate with.
- Is there a way to derive a set of states S' such that (i) $S' \subset S$, and (ii) there exists a plan that satisfies G_s for some $s^* \in S'$.

Motivation: It allows for different planning approaches. We can focus on existence of a plan, minimizing leaked information, maximizing size of S' and more. It easily allows decrease/increase in difficulty either by size of state space, or number of agents or restriction of communication. It is related to other examples in dynamic epistemic logic (russian cards, gossip protocols etc).

Features:

- *Knowledge vs. belief*: Knowledge.
- *Single vs. multi-agent*: Multi-agent.
- *Communication vs. sensing vs. ontic*: Communication with, possibly, sensing.
- *Determinism/reliability of communication*: Deterministic.
- *(De)Centralized*: Possibly both.
- *Epistemic Goals*: Yes.
- *Collaborative/competitive*: Collaborative, since the third agent does not act.
- *Benchmarks*: None known.

See also work on communication complexity about related problems (A very short introduction to this and references can be found in Capelli's thesis [27].

4.1.3.8 BW4T – BlocksWorld For Teams

Description: Blocks World for Teams (BW4T) [54] is a multi-agent simulation platform that is an abstraction of several different applications and scenarios. In this problem, we have a number of rooms and a drop zone. Each room contains a number of colored blocks (or it is empty), and, initially, none of the agents know where the blocks are. Each agent can carry at most one block at a time, and agents can communicate with each other concerning their observations.

The goal is for the agents to deliver a number of colored blocks, in a particular order, to the drop zone. Figure 5 shows an example scenario, taken from the BW4T simulation environment.

Motivation: This problem is an abstraction of several important application domains, such as disaster management and search & rescue. For example, blocks represent survivors of a disaster, the drop zone represents a medical tent, agents represent autonomous vehicles and humans with differing capabilities, such as searchers, medical officers, or aid delivery vehicles.

From an epistemic point of view, we can define a BW4T with private actions, restrict communication between subgroups of agents, and, if in a decentralized approach, reason about the other agents' intentions and plan accordingly. In any case, agents need to reason in terms of *who* to ask about *what*.

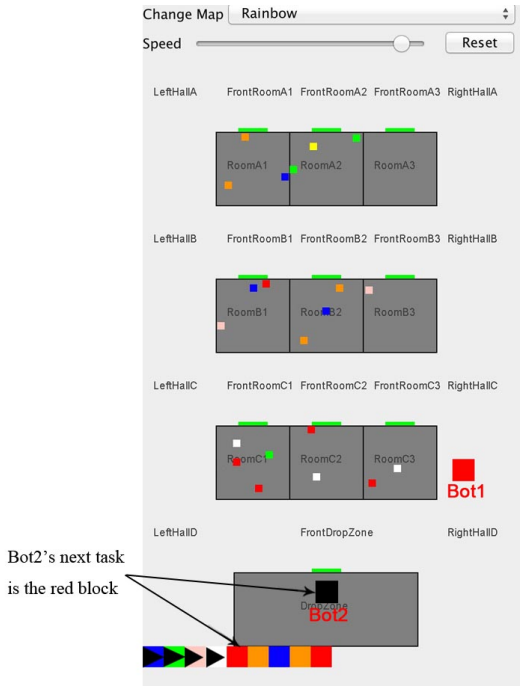

Figure 5 Blocks World for Teams simulation interface.

In addition, the BW4T simulator is designed as a testbed for human-agent interaction experiments, making it an interesting framework to study human interaction in epistemic terms: bounded reasoning about high-order beliefs, belief change, interaction between belief, intention, and action.

The problem is easy to scale by modifying the number of rooms, blocks, agents, depth of epistemic goals, and area to be searched.

While the original problem is not particularly rich with respect to epistemic reasoning problems, there are many extensions that one can make that are seen in domains that BW4T abstracts:

1. Joint activities: Certain coloured blocks (e.g. blue) are 'heavy' blocks, and need two agents to carry them simultaneously. This requires agents to participate in and coordinate this joint activity; something that can be done only via communication or by reasoning about the other agents' perspectives and intentions. In the scenario of no communication, it explicitly requires reasoning about higher-order beliefs and intentions.
2. Heterogenous agents: Different agents may have different capabilities; for example, some agents can just search and find blocks, while others must retrieve them, and certain rooms may only be accessible to certain agents.
3. Knowledge goals & Communication: Another scenario is one in which a 'commander' of the task must be told of the location of all blocks once found. Only some agents can communicate back to the commander, so the problem is deciding what and when to communicate block locations.

Features:

- *Knowledge vs. belief*: Possibly both.
- *Single vs. multi-agent*: Multi-agent.
- *Communication vs. sensing vs. ontic*: All three.

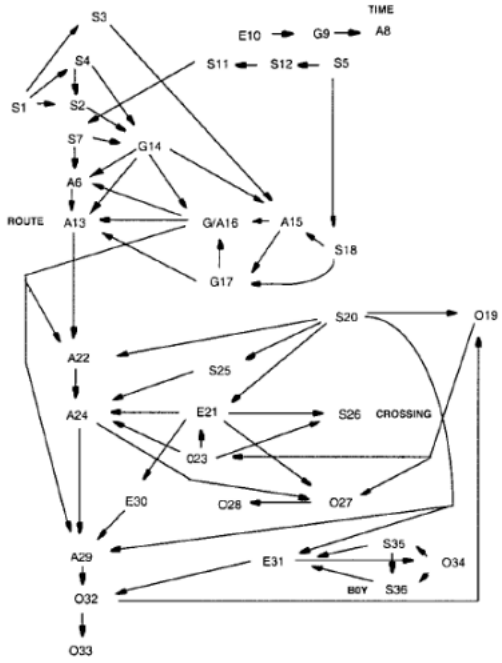

Figure 6 Example causal chain.

- *Determinism*: Yes, but non-determinism is possible.
- *(De)Centralized*: Both.
- *Epistemic Goals*: Yes.
- *Collaborative/competitive*: Both, if two teams try to get the same blocks while their number is not sufficent for both teams.
- *Benchmarks*: A simulation tool for BW4T: http://eishub.github.io/BW4T/

4.1.3.9 Social Explanation Dialogue in AI

Description: Consider an agent that makes a decision, and a human observer asks why that decision was made. The aim of a social explanation can be described as follows. Given a causal chain of 'events', such as the one in Figure 6, and a request to provide a socially-acceptable explanation for a specific event in that chain, taking into account what the receiver already knows, to extract an explanation that allows the person who posed the query to understand why the event occurred.

The explainer may have a model of the explainee's (partial) understanding of the causal chain. The explainee can communicate values of the variables in the chain, or relations between variables. Halpern and Pearl offer a formal model for the problem using structural equations [45].

This problem can be extended to a dialogue. At each step, the explainee provides only one part of the explanation, selecting/prioritising the 'best' explanation, measured by e.g. simplicity, generality, conformance with prior beliefs of explainee.

Motivation: The ability to explain a decision to a human observer is key to providing transparency and trust to human operators, and is currently one of the most challenging and important problems in artificial intelligence.

Further, such social explanations require a Theory of Mind to avoid explaining redundant information. In richer settings, explanations can be given to multiple observers at once. The challenge in this setting is to provide the explanation with the most explanatory power to the most observers.

Features:

- *Knowledge vs. belief*: Both.
- *Single vs. multi-agent*: Multi-agent.
- *Communication vs. sensing vs. ontic*: Communication.
- *Determinism/reliability of communication*: Deterministic.
- *(De)Centralized*: Decentralized.
- *Epistemic Goals*: Yes (perhaps only first-order).
- *Collaborative/competitive*: Collaborative (but deceptive versions could exist).
- *Benchmarks*: None known.

4.1.4 Modelling

Currently there is no standardized formal language which can express all or even just a substantial subset of these epistemic planning problems and benchmarks. For the future it would be good to fix a notation for DEL and generalize (or pick one of the existing generalizations of) PDDL.

Further, it would be fruitful to build a centralized website that collects the benchmarks (as http://www.satcompetition.org/ for SAT and http://planning.domains/ for planning).

4.2 Exploring Action Types and their Representations

Yves Lespérance (coordinator), Louwe B. Kuijer (coordinator), Nina Gierasimczuk, Barteld Kooi, Michael Thielscher, Ivan José Varzinczak, Thomas Bolander, Andrés Occhipinti Liberman, Tran Cao Son, Yongmei Liu, Hans van Ditmarsch

4.2.1 The Need of Pure Action Types in DEL

One of the topics discussed was the distinction between action tokens and action types. In particular, we focused on the question how the action models in DEL fit this distinction. Can DEL action models be seen as action types? It turns out that some aspects of actions which could be seen as properties of tokens are by design included in action models.

As an example, consider an action model corresponding to a private announcement. Such an action model contains information about which agents learn the content of the information, which agents learn that there has been an announcement but do not learn the actual content, and which agents learn nothing. The problem here is that which agents are "paying attention" should be independent of the action of announcing, but rather follow from the situation in which the announcement is made.

In other words, who does or does not learn the content of an announcement should be determined not by the action model, but by the state model to which the action model is applied. What is needed, then, is an alternative formulation of such an event where the

"observability" information is encoded in the state model. In the context of high level action languages such an approach was suggested in [13] and elaborated in [14, 15, 16].

One can also do this in DEL itself, but there are also other (equivalent) formalisms where this comes a bit more natural, such as Generalized Arrow Update Logic [58] and Edge-Conditioned Event Models [22]. The (epistemic) Game Description Language [97] provides an alternative action representation formalism for epistemic domains that also comes with a clear separation between an action itself and information about which agents are observing it; see below.

One can introduce *observation atoms*, combined with *agentive actions* and *observability change*. Agent i announcing φ can then be the action $i : \varphi$, represented by the following event model.

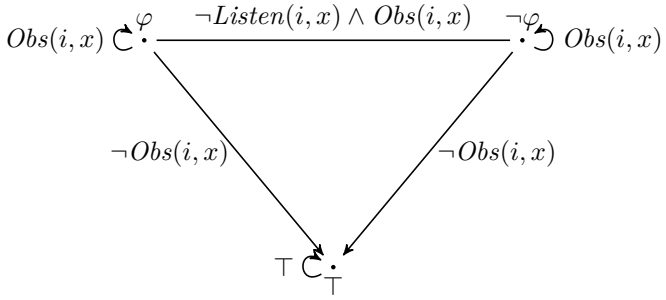

Here $Listen(i, x)$ means that agent x is listening to agent i, and therefore x will learn the content of i's announcement. The atom $Obs(i, x)$, meanwhile, means that x observes i making the announcement, but does not necessarily listen to the content of the message.

Importantly, in addition to the preconditions on the worlds of the event model, there are also preconditions on the arrows in the event model. In the notation used above, there would be an arrow for agent x from event e to event f if $\psi(x)$ is true at e, where ψ is the label on the arrow from e to f.[6]

As a consequence of these additional preconditions on arrows, one does not need to put any information about who observes the announcement in the event model. Agent i takes an action of the type "announcing", whether this is a public or a private announcement depends on the state model.

In this framework, if i definitely wants to make a *public* announcement of φ, this means that she has to perform two actions. First, she has to take an action which has the result that $Listen(i, x)$ becomes common knowledge for all x, then she takes the action of announcing φ.

There is still one significant restriction on this framework: the agents that observe but do not listen to the announcement $i : \varphi$ will still learn that it was the truth or falsity of φ that was announced. There does not seem to be an elegant way around this in a DEL-like system like the one presented; an agent who observes an announcement, but has no idea what the announcement is about would, in theory, need to be represented by an exponential (or even infinite) model with one copy for each possible announcement.

It seems the single "announcement" action type can model public, semi-private, and private announcements. It is conjectured that similar representations can be found for every action type that can be represented in DEL.

[6] See [22] for details. This formalization is very similar, but not identical, to that of Generalized Arrow Update Logic [58].

Unfortunately, one cannot give this conjecture a technical formulation, because it is unclear what, on a technical level, the objective is. For every type of action, which is represented by multiple action models in DEL, one would like to have a single event model in the language that that was just described. However, it is not clear what, on a technical level, it means for different action models to represent actions of the same type.

As a final note, it should be noted that lunch is usually not free. In the above formalism, information about which agents observe which parts of an action need not be encoded in the event models. In some sense, this makes them simpler. However, one pays for this by having to include information about observability in the state model. So complexity is not removed, but moved to a different part of the problem.

It is believed that in some cases it is useful to move this complexity from event models to state models. The first advantage of doing this is the motivation for this system, namely that it allows a clearer representation of action types. The second advantage is that it is often considered harder to intuitively understand event models than state models, so keeping event models as simple as possible may aid compehensibility. But one should keep in mind that in some cases the complexity shift may be harmful, and that in those cases one is probably better off using DEL.

Action Types in the Epistemic Game Description Language

The (epistemic) game description language (GDL-III, [97]) may provide an alternative to DEL for representing actions and knowledge especially when the aim is to avoid complex event models. A special feature of this language is that only what agents can see and do needs to be specified. It is then implicit in the semantics how actions and observations affect the knowledge of agents. The scenario from above, for example, would be axiomatized thus (using a first-order variation of the actual syntax):

$$\mathbf{Sees}(x, i_a) \;\Leftrightarrow\; \mathbf{Does}(i, \mathit{announce}(z)) \wedge \mathit{Obs}(i, x)$$
$$\mathbf{Sees}(y, \varphi) \;\Leftrightarrow\; \mathbf{Does}(i, \mathit{announce}(\varphi)) \wedge \mathit{Obs}(i, y) \wedge \mathit{Listen}(i, y) \wedge \varphi$$

Both **Sees** and **Does** are keywords that loosely correspond to modalities. The two axioms can be understood as follows: If agent i makes some announcement z then all agents x observing i will receive the "information token" i_a. Any agent y that observes and also listens to i will get to see the announcement φ itself. From the semantics of GDL-III it follows that agents who only see i_a will know that i has made an announcement but without learning the content. Agents who see φ, however, will know that φ must be true. Moreover, if for example an agent x only observes agent i but at the same time knows that another agent y listens to i, then the semantics entails that x will know that y knows the content of the announcement afterwards.

In the Game Description Language, actions such as $\mathit{announce}(z)$ can be seen as action types. Their effects, especially on the knowledge of other agents, depends on the state in which they are executed. Our example action can thus model public, semi-public, and private announcements. Form this perspective GDL-III might be a viable alternative to DEL in cases where the complexity shift from the event model to the state representation is desirable. It has been shown that many of the standard DEL-examples can be expressed in GDL-III [97]; however, a precise characterization of the expressivity of the two languages in comparison has yet to be given.

4.2.2 Concurrent Actions

The issue of finding a suitable representation for *concurrent actions* that can be used in multiagent epistemic planning was also discussed. Amongst other things, concurrency has attracted attention in multiagent planning because parallelizing actions can reduce overall plan length and lead to a significant speed-up in plan execution time.

Concurrent actions can have interaction effects; their parallel execution sometimes alters the effects that the actions would have if executed individually. The following example from [24] illustrates this. Suppose agents x and y need to move a large set of blocks from *room 1* to *room 2*. While x and y could each pick up and transport single blocks repeatedly, they can do better by using an existing table as follows. First, the agents put all the blocks on the table. Then, they each lift one side of the table. Lifting the table simultaneously from both sides allows the agents to jointly carry all blocks to *room 2* in one trip. However, if only one side of the table is lifted, all the blocks will fall to the floor. That is, the *table lifting* action has different effects when applied individually or jointly, and its joint execution allows the agents to achieve their goal faster.

In the reasoning about actions community several formalisms and approaches, such as in [63, 17], have been proposed to represent and reason abut concurrent actions. More recently, Boutilier [24] and Knoblock [56] extended the STRIPS representation of actions to allow for concurrent actions. Reiter proposed axioms for concurrent planning in the situation calculus framework [81]. There has also been some work to model concurrent actions in the DEL framework [110], as well as in the GDL framework [96]. It would be useful to provide a comparison of these different approaches to concurrency as well as their applicability in epistemic planning.

4.3 Exploring the Relations Between Knowledge and Belief in Multi-Agent Epistemic Planning

Chitta Baral, Hans van Ditmarsch, Jan van Eijck, Esra Erdem, Andreas Herzig, Mathias Justesen (coordinator), Yongmei Liu, Richard Scherl (coordinator), Tran Cao Son

In multi-agent domains actions can happen without all agents being aware of their occurrence. This leads to situations where some of the agents' beliefs may not match with reality. On the other hand agents' may have incomplete information about the world and there may be actions such as sensing actions that provide additional knowledge about the world to the agent that does the sensing. Agents may have knowledge goals that are about knowing part of the world and agents may have goals that are about decieving other agents in terms of mismatch between the beliefs of the other agents and reality. Thus the notions of knowledge and belief play crucial roles in a multi-agent epistepic planning domain and several important questions about them need to be addressed. Some of those questions are: How do we identify the theoretical and computational challenges in planning with knowledge vs. planning with belief? When is one or the other appropriate and when are both needed?

Every approach has to have a way of distinguishing the representation of the knowledge and beliefs of different agents from the representation of the actual world. Even if knowledge is being used, the knowledge may turn into false belief after a partially observable action has occurred. This necessitates addressing questions such as:

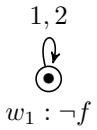

Figure 7 Initial state s_0.

- What are the different formalisms for planning in knowledge and/or belief and what are their theoretical and computational properties?
- How is revision dealt with in planning?
- Are there specific types of interesting goals for epistemic planning? (Note that, in planning with beliefs, goals can be false beliefs with relevance to formalizing false-belief tasks such as lying and deception.)

4.3.1 Knowledge vs. Belief

In planning we sometimes want to achieve something with certainty and therefore knowledge is then needed. However, the certainty required for knowledge is not always achievable and then we should work with belief. From an external perspective the planner has a goal involving knowledge, but from an internal perspective each agent executing its part of the plan only needs belief. One approach is to define knowledge in terms of belief as true belief. There are also a variety of ways of defining belief in terms of knowledge [4].

In the planning based on dynamic epistemic logic (DEL) work in [10], states are modeled using Kripke structures and depending on the accesibility relations among the worlds, different axioms of knowledge and belief are obtained. If the accessibility relations satisfy that all relations are reflexive, it corresponds to the axiom of truthfulness, T: $K_a\varphi \rightarrow \varphi$. When modelling the planning problems that satisfy axiom T (e.g. S5), the agents cannot have false beliefs. However, this is often not a suitable model of the problem, so for illustrating the difference between knowledge and belief, KD45 is used. (In this context, in the work [88] it is shown that the KD45$_n$ property is maintained for a large class of actions.)

The rest of this subsection is organzied as follows. First, we motivate the use of planning with beliefs with an example based on DEL. Continuing with the example, we proceed to show that with DEL, agents cannot recover from a false belief. Using various formalisms, we then explore different options for incorporating belief revision in epistemic planning.

4.3.2 Planning with beliefs

Consider the following scenario. There are two agents. Initially, there is common knowledge of $\neg f$. The goal is to have agent 1 correctly believe f and agent 2 to incorrectly believe $\neg f$. So agent 1 has the plan of changing f to true while agent 2 is not looking.

- Initial State: $C_{\{1,2\}}\neg f$
- Goal: $K_1 f \wedge K_1 K_2 \neg f$
- Plan: agent 1 makes f true while agent 2 is not looking

Consider the example in DEL. We first have the initial state, where there is common knowledge of $\neg f$ depicted in Figure 7.

The plan can be modelled by an action a depicted in Figure 8, i.e., agent 1 makes make f true while agent 2 is not looking.

The result of executing action a is $s_1 = s_0 \otimes a$ as depicted in Figure 9, where $\otimes$ is the usual product-update operator. Observe that s_1 satisfies the goal, $s_1 \models K_1 f \wedge K_1 K_2 \neg f$.

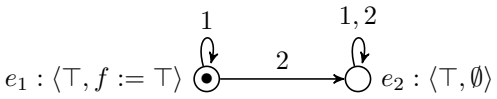

Figure 8 Event model for action a.

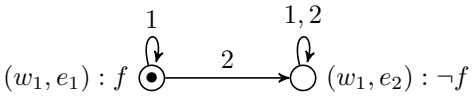

Figure 9 Resulting state $s_1 = s_0 \otimes a$.

Figure 10 Event model for action a_2.

Now let us imagine that there is a public announcement of f. The action model for this action a_2 is given in Figure 10.

The result in computing $s_2 = s_1 \otimes a_2$ is that agent 2 now believes everything (including contradictions), as shown in Figure 11, since there are no accessibility relations for agent 2. This appears counter-intuitive. We would expect agent 2 to revise their beliefs, and now believe f. Hence, we need some sort of belief revision to handle such scenarios where an agent has a false belief.

4.3.3 Belief revision

The topic of belief revision in general, and with regard to DEL in particular [38, 113, 12, 102], has been studied thoroughly in the literature. We look at several approaches of incorporating belief revision:

- Plausibility models [11]
- Syntactic belief revision [53]
- Recovery [112]
- Modifying the product update of DEL (a recent work by Chitta Baral and Tran Cao Son).

Plausibility models

Here there is a plausibility ordering on the worlds considered possible by a given agent. The belief of the agent is captured by the most plausible worlds. If worlds are lost due to application of an action, then previously less plausible worlds become the most plausible. It is a requirement of plausibility orderings that they are reflexive, hence, we get the a action in Figure 12, and the resulting state s_1 (Figure 13) now also has a reflexive edge for agent 2 in the designated world. Still, in s_1 agent 2 believes $\neg f$ more strongly than f as indicated by the plausibility orderings ($\leq_i$) on the worlds, where the minimum element is considered most plausible.

Now, we apply the public announcement action of f and the result is given in Figure 14. Note that we no longer lose the indistinguishibility relations of agent 2, as their belief was revised by rejecting the world that was previously believed to be most plausible.

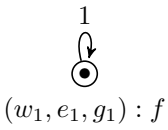

Figure 11 State s_2 after the public announcement of f.

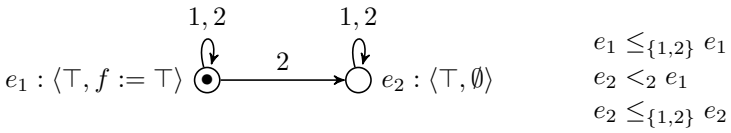

Figure 12 Plausibility event model for action a.

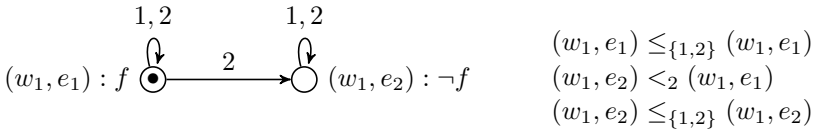

Figure 13 Resulting state $s_1 = s_0 \otimes a$.

Syntactic revision

In [53] formulas are kept in a normal form called Alternating Cover Disjunctive Formula (ACDF). There is an algorithm for syntactically manipulating the form for update and revision so that the result is still in ACDF. The running example can be done with an initial state of

$$\neg f \wedge K_1 \neg f \wedge K_2 \neg f.$$

The action a with effect $f \wedge K_1 f$ is used with the update algorithm to yield

$$f \wedge K_1 f \wedge K_2 \neg f.$$

Next the action a_2 with effect $K_1 f \wedge K_2 f$ is used along with the revision algorithm to yield

$$f \wedge K_1 f \wedge K_2 f.$$

Recovery

Another approach is to do a recovering action before the public announcement [112]. The recovering action is defined as $R \cup (R; R^u) \cup R^u$, where R is the original set of accessibility relations for an agent, R^u is the inverse relation, and $R; R^u$ is the composition of the relation and its inverse. Recovery can be seen as a form of epiphany, where agents suddenly expand their beliefs and consider worlds that they previously did not. Recovery is applied to the state of Figure 9 to yield Figure 15. Then when the public announcement is applied, the desired result is obtained.

Modifying the DEL product operator

An alternative approach, suggested and being worked on by Baral and Son, is based on modifying the DEL product-update operator. In the above discussion of the product update,

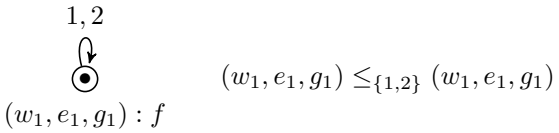

Figure 14 State after public announcement of f is $s_2 = s_1 \otimes a_2$.

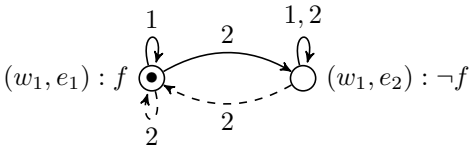

Figure 15 State after recovery. Dashed edges are the edges added by recovery.

we saw that when executing an action a in the epistemic state s, there is an edge with label l between nodes n_1 and n_2 in the resulting epistemic state, if there is a similar edge in both a and s. In the modified operator, this requirement is relaxed.

More precisely, if node n_1 is a composition of node ns_1 (of s) and na_1 (of a) and node n_2 is a composition of node ns_2 (of s) and na_2 (of a) then the traditional DEL operator will add an edge with label l from n_1 and n_2 if there is an edge both from ns_1 to ns_2 and from na_1 to na_2 with label l. Baral and Son suggest that, when one can conclude that the action a requires revising s, then a *relaxed* product operator can be used where a gets preference over s. For example, with respect to a simple relaxed operator, there will be an edge with label l from n_1 and n_2, if there is an edge with label l from na_1 to na_2. This addresses the example in question. Finding the right relaxed product operator, and conditions when they should be used is a topic of ongoing work.

Update and revision for proper epistemic knowledge bases

Miller and Muise [70] present a belief update and revision approach for proper epistemic knowledge bases, which underlies the approach for compiling epistemic planning to classical planning discussed in Section 4.4.1. The syntactical restrictions on proper epistemic knowledge bases allow entailment, update, revision, and erasure operations that have quadratic time and space complexity, and thus enable efficient search over epistemic states.

4.3.4 Story understanding

The distinction between knowledge and belief is important in many scenarios, e.g., story understanding. As a story progresses, the characters can obtain false beliefs. As is common in TV serials, the truth is revealed at the end and all characters need to revise their beliefs.

This is related to the Sally-Anne test. A short skit is enacted. Sally takes a marble and hides it in her basket. She then leaves the room and Anne takes the marble out of Sally's basket and puts it in her own basket. The question of where Sally will look for her marble is asked. Autistic children tend to reply that the marble is in Anne's basket. The test shows the inability of people on the autism spectrum to distinguish between the actual world and a person's belief about the world.

Story understanding relates more directly to plan verification than to plan synthesis: Answering questions about the beliefs of agents at the end of a story corresponds to checking whether certain epistemic or doxastic formulas ("goal formula") hold after the execution of an action sequence (where the action sequence represents the story).

4.4 Practical Tools, Resources and Computational Techniques

Tristan Charrier (coordinator), Sophie Pinchinat, Vaishak Belle, Thorsten Engesser, Torsten Schaub, Malte Helmert, Bastien Maubert, Jens Claßen, Ioannis Kokkinis, Sunil Easaw Simon, Robert Mattmüller

We survey existing techniques to deal with epistemic planning. For each of these techniques, we explain what are the models and formulas considered, and what are the advantages and the drawbacks of such techniques. We do not restrict the survey to already implemented epistemic planners, but also detail theoretical techniques to deal with epistemic planning.

4.4.1 Compilation to Classical Planning: Syntactic Approaches

An approach to epistemic planning that recently gained attention is the idea of resorting to a suitably restricted fragment of DEL and use a STRIPS-style encoding of action operators over epistemic formulas instead of event models and product updates. This allows to compile the given problem into an instance of classical planning, and subsequently leverage the performance of state-of-the-art PDDL planners for solving it efficiently. The proposed methods differ in

1. the restrictions that are applied,
2. whether a centralised or first-person view is adopted,
3. the type of knowledge that is assumed.

Muise *et al.* [74] present a planner for multi-agent systems where agents are assumed to have so-called *proper epistemic knowledge bases*, which effectively rule out disjunctions. With the additional restriction to a finite depth of nesting of modalities, the compilation is exponential in the depth of the maximum length of modal strings. The expressiveness of the action model is independent of the approach, but depends only on that of the underlying planner. Both $KD45_n$ and $S5_n$ are supported. Planning is from the view of a single agent, including a "perspective shift" where one agent reasons about the belief of another from that agent's point of view.

Kominis and Geffner [57] compile a multi-agent epistemic planning problem into a classical one using techniques from single-agent conformant planning. The compilation is similarly exponential in the maximal depth of nesting of modal operators. They assume $S5_n$ type knowledge and adopt a centralised perspective with command-and-control first-person planning. Preconditions and effects of actions are expressed by Boolean combinations of propositional literals with modality prefixes, including disjunctions.

Cooper *et al.* [33] use an encoding based on "seeing" variables for atomic propositions, which represent the agent "knowing whether" the corresponding variable is true. Knowing a variable p can then be defined as knowing whether p holds and p being actually true at the same time. The approach supports both $S5_n$ and K_n knowledge types and plans from a centralized perspective.

A related, earlier approach is Petrick and Bacchus' PKS ("planning with knowledge and sensing") system [78] for single-agent epistemic planning. While they do not compile into classical planning, they also adopt a syntactic representation for a planning domain where an agent's knowledge is assumed to be expressed by three different types of formulas: the knowledge base is partitioned into one part for "knowing that", one for "knowing whether",

and one for "knowing what". Their planner then generates conditional plans over the space of knowledge states expressible by means of such knowledge bases, relying on an efficient, yet incomplete, inference algorithm.

4.4.2 Forward Search in Space of Epistemic Models

Forward search in the space of epistemic models is possibly one of the most obvious approaches to solving epistemic planning tasks. The first definitions of (centralized) epistemic plans over epistemic states and DEL action models were given by Bolander and Andersen [23] and independently by Löwe *et al.* [66]. The plans are sequences of DEL action models applicable in the initial epistemic state that guarantee to reach an epistemic state satisfying the goal formula. Those were later generalized to branching plans [2] where branching is essential, since epistemic action models allow non-deterministic action outcomes, possibly necessitating different follow-up actions.

While the above-mentioned definitions of plans may not, strictly speaking, imply a computation based on forward search in the space of epistemic states, they at least suggest such a search as an obvious first approach to explore. To our knowledge, so far only simple uninformed explicit-state forward search algorithms such as breadth-first search (BFS) have been implemented.

Whereas the first definitions assumed a centralized perspective, the definitions were later adapted to a more decentralized view. Engesser *et al.* [40] introduced the notion of implicitly-coordinated (linear and branching) plans. While these plans are still computed from a single-agent perspective, they consider explicitly the perspectives of the owner of each subsequent action (*i.e.*, the agent that will execute it). The intention is to facilitate cooperative behavior by ensuring that agents can come up with those actions when planning on their own. The approach heavily relies on the framework's capacity to arbitrarily switch perspective between the agents. Since the necessary depth of knowledge nesting is dependent on the plan length and therefore not known beforehand, compilation into classical planning is not possible.

4.4.3 Efficient progression for epistemic states

One way to handle epistemic planning is to adapt a standard forward search (progression) or backward search (regression) algorithm to the epistemic setting. Compared to a classical (non-epistemic) setting, a major question in this approach is how to represent the search states. In the classical setting, search states are either propositional models (in the case of planning with complete information) or representations of sets of propositional models, such as propositional logic formulas (in the case of planning with incomplete information, for example conformant planning). This latter representation can be generalized to the epistemic setting by using epistemic formulas instead of classical ones. An algorithmic challenge, then, is how to efficiently perform the basic operations on search states that are required by standard search algorithms like A^* or greedy best-first search, such as testing whether a goal state has been reached, testing whether a given action is applicable, and computing the (epistemic) successor state for a given (epistemic) search state and (epistemic) action. The paper by Bienvenu *et al.* [21] addresses these issues in the context of the (single-agent) epistemic logic S5. They present a fragment of $S5$ formulas, called $S5DNF_{DNF,CNF}$, which efficiently supports the progression operation: given an epistemic formula describing the possible (epistemic) states of the world and a description of an epistemic or classical action in a certain formula representation, compute a formula representing the possible (epistemic)

successor states after applying the action. The operation is polynomial-time in the description size of the inputs (epistemic state and action), but presumably can be exponential in the sequence length if sequences of actions are applied, as each application can increase the size of the resulting formula. Miller and Muise [70] present a similar approach for proper epistemic knowledge bases, which are epistemic knowledge bases that forbid disjunction. Progression of such knowledge bases is quadratic in the number and depth of epistemic formulae. Presumably, the same techniques can be used for regression instead of progression, *i.e.*, for computing possible predecessor states given formulas representing current states and actions. More generally, the paper is framed in the context of knowledge compilation pioneered by Darwiche and Marquis [36]. It emphasizes the strengths and weaknesses of different kinds of representations of knowledge (usually in terms of propositional models, here extended to an epistemic setting) by focusing on their closure properties and the set of operations (such as logical connectives, logical implication, existential and universal abstraction etc.) that can or cannot be performed efficiently for a given representation. Beyond the immediate contribution of the paper by Bienvenu *et al.*, the knowledge compilation perspective might be a useful angle for coming up with search state representations for epistemic planning.

4.4.4 Symbolic model checking of strategy logics with knowledge (AETL)

Logics for strategic reasoning have been developed in the last two decades, first with Alternating-time Temporal Logic (ATL) and its many variants, and more recently with Strategy Logic (SL). ATL, introduced in [1], extends the branching-time logic Computation Tree Logic (CTL) [7] by replacing CTL temporal operators with strategic cooperation modalities expressing what state of affairs a coalition of agents can bring about in a system, irrespective of the actions of the other agents. It can thus naturally express the existence of a strategy, or plan, to reach a given objective. So can SL, as it is strictly more expressive than ATL [71].

MCMAS [65] is a model checker for logics for strategic reasoning that handles several variants of ATL and SL, including AETL, an extension of ATL with knowledge operators [104] that can talk about epistemic objectives. Models are concurrent game structures, which are essentially finite state automata where states represent positions of the game, or states of the system, and transitions are labelled by tuples of actions, one for each agent. In addition, equivalence relations on the set of states represent agents' observation of the system. The specification of the system is completed by providing MCMAS with a set of initial states, an optional set of fairness conditions, and the set of formulas to be verified.

Given an epistemic formula φ, the AETL formula $\langle A\rangle F\varphi$ means that "there exists a strategy for coalition A such that every outcome of this strategy hits a state where goal φ holds, irrespective of the strategies of the other agents". The model checker MCMAS can therefore be used to solve a form of epistemic planning problem, even in an adversarial setting. Since MCMAS, as other standard model checkers, can also generate witnesses of strategies when they exist, it is possible to synthesize an (adversarial) plan when the answer to the epistemic planning problem is positive. Notice that AETL can express more than mere reachability. This is an advantage if one wants to solve variants of the epistemic planning, *e.g.* with extra constraints on the searched plan. For example, one may seek a plan that reaches the goal without ever taking the system into some predefined set of bad/risky states.

However, there are some strong limitations when using the MCMAS model checker. First, it does not allow reasoning about general beliefs, as epistemic relations are $S5$. Second, it is not possible to reason about systems with perfect recall: an agent does not distinguish between two executions of the system as long as their last states are equivalent to her (*i.e.*

memoryless knowledge). Third, unlike other frameworks such as DEL, the entire system needs to be described as a finite state automaton (or concurrent game structure). Therefore, only regular systems can be considered, as opposed to those described by arbitrary DEL presentations.

We also mention MCK, a model checker for logics of knowledge and time (that supports both symbolic and explicit model checking); DEMO, an explicit model checker for DEL; and SMCDEL, a symbolic model checker for DEL [103]. A comparison of MCMAS, MCK and DEMO can be found in [111].

4.4.5 Model checking via automata

Automata theory provides very powerful techniques to study a wide range of logics of programs, such as the linear-time temporal logic LTL, the branching-time temporal logic CTL or the mu-calculus (see, e.g, [99] and [114]). Concerning the interplay between time and knowledge, little work has been done using automata. The reason may be that adding the "transversal" dimension of knowledge in trees representing the temporal evolution transform the usual tree models into graphs, for which no good model of automata is known.

A few years ago, Bozzelli *et al.* [25] showed that if the transversal, or epistemic, relation is representable by a two-tape automaton (*i.e.* it is a rational relation), then we can model-check branching-time temporal logic with epistemic operators. This result was then applied to solve the propositional DEL-based epistemic planning problem: the automata constructions presented in [5] can, given a propositional DEL presentable domain of the epistemic planning problem, produce a finite word automaton that accepts the set of successful plans. Indeed, in this restricted class of domains, the epistemic relations of agents with perfect recall fit in the framework considered in [5]: they are rational relations, and in fact quite simple ones as they can be accepted by transducers with only one state.

The first main advantage of this approach is that it generates a finite representation of the whole set of solutions to the planning problem, which may be useful for instance if the addition of some new constraint reduces the set of solutions: it may be more efficient to cut some transitions in the automaton to obtain the new set of solutions, instead of computing from scratch new solutions to the over-constrained problem.

A second strength of this approach is that it can actually solve, without additional effort, a much more general problem, which is called Epistemic Protocol Synthesis in [5]. In this problem, the input is again a DEL-based epistemic planning problem, but instead of an epistemic goal formula φ that has to be achieved, the problem takes an arbitrary formula from CTL^*K_n, the full branching-time temporal logic with knowledge, and produces, if possible, one possible infinite tree of events that satisfies it. The tree may represent a non-deterministic plan, a set of plans, or a strategy for instance. In CTL^*K_n one can express, for instance, the existence of a path that eventually satisfies φ, which corresponds to the classic epistemic planning problem. But one can also for example require instead that in the solution tree there is a path where φ_1 holds infinitely often, another path in which after some point φ_2 always holds, and in addition that in all paths of the solution tree, some fairness property holds.

The main drawbacks of this approach are that, because of the perfect recall semantics resulting from the DEL setting, it only works for propositional events, *i.e.*, events with propositional preconditions and effects/postconditions. However the problem does not come from the automata approach: the DEL-based epistemic planning problem is known to be undecidable for alternation depth one in the pre/post-conditions. However it is still a seemingly difficult question whether it is decidable for alternation depth 0 and even modal depth 1.

Finally, this approach gives an algorithm in $(k+1)$-EXPTIME, where k is the modal depth of the goal formula. Note that it is the same complexity for the epistemic planning problem and its generalization, epistemic protocol synthesis. Also, for DEL-based epistemic planning it is the best known upper bound, but matching lower bounds are an open question.

4.4.6 Reduction of Epistemic Reasoning to Theorem Proving

One technique proposed for dealing with epistemic formulas in a Situation-Calculus context is a reduction to theorem proving (*e.g.* see [59] for a modal variant of the epistemic Situation Calculus). In the single agent case (see [18] for an extension to the multi-agent case), the idea is that the agent's belief is represented by a knowledge base in the form of a collection of (non-epistemic) first-order formulas about the current state of the world, and it is assumed that this knowledge base is all that is known by the agent. An epistemic subformula such as $K(Open(window))$ can then be evaluated by checking whether $Open(window)$ is entailed by the knowledge base, and subsequently substituting it by either true or false, depending on the outcome of the entailment check. Furthermore, open formulas like $K(OnTable(x))$ will be substituted by a representation of the known instances with respect to the knowledge base in terms of equalities, *e.g.* $(x = mug \vee x = laptop)$, and nested K's can be dealt with by means of recursion. Together with regression (a syntactical method for rewriting a formula about the situation after the execution of a sequence of actions to a formula talking only about the initial state of the world), this corresponds to a reduction of the projection problem to standard first-order theorem proving as follows: In order to decide whether formula $[Action_1]...[Action_k]K(\varphi)$ holds, use regression to eliminate the action modalities $[Action_i]$, turning the input into an equivalent formula (with respect to axioms describing action effects and sensing outcomes in the knowledge base), use reduction (which involves theorem proving in the form of entailment checks) to eliminate K modalities, and finally check the resulting formula against the knowledge base (another entailment check). Projection is of course only a subtask needed for planning. To create a planner, one has to perform a search over possible action sequences, where in each case projection is used to decide whether it constitutes a plan achieving the given goal formula. One option to constrain the search space and thus speed up planning is by incorporating procedural domain knowledge in the form of knowledge-based agent programs (e.g. [83, 69, 82, 61, 32]). Apart from imperative constructs such as while-loops and if-conditionals, such programs may also contain nondeterministic choices. A program thus can be understood as a plan sketch, where certain parts are left open and have to be resolved by the system. As the general Situation Calculus allows for full first-order expressivity, in practice one needs to integrate further restrictions to obtain decidability or tractability. Apart from the usual syntactical restrictions on the language, another interesting approach is resorting to a limited (sound, but incomplete) form of reasoning [60].

4.4.7 Language $m\mathcal{A}$

The language $m\mathcal{A}$ is described in [16]. Some of its earlier versions appeared in [14, 15]. It could be seen as an extension of the high-level action languages to reasoning about actions and changes in single agent environment [43] to multi-agent environment.

In $m\mathcal{A}$, a multi-agent environment is characterized by a set of agents $\mathcal{AG}$, a set of fluents $\mathcal{F}$, and a set of actions $\mathcal{A}$.

A *belief formula* over $\mathcal{AG}$ and $\mathcal{F}$ is of the following form

$$\varphi \mid \neg\varphi \mid \varphi \wedge \psi \mid \varphi \vee \psi \mid \varphi \rightarrow \psi \mid B_i\varphi \mid C\varphi \mid E_\alpha\varphi$$

where φ and ψ can be a propositional formulas over $\mathcal{F}$ or belief formulas, $i \in \mathcal{AG}$, and $\alpha \subseteq \mathcal{AG}$.

Syntax

A *dynamic domain* over $\mathcal{AG}$, $\mathcal{F}$, and $\mathcal{A}$ consists of a set D of statements describing the effects of actions in $\mathcal{A}$, and a set O of statements describing the observability of agents when actions are executed. Specifically,

- Each statement in D is of a following form

$$\textbf{executable } a \textbf{ if } \psi \tag{1}$$
$$a \textbf{ causes } \ell \textbf{ if } \psi \tag{2}$$
$$a \textbf{ determines } f \tag{3}$$
$$a \textbf{ announces } \ell \tag{4}$$

  where $a \in \mathcal{A}$, f is a fluent $\mathcal{F}$, ℓ is a fluent $f \in \mathcal{F}$ or its negation $\neg f$, and ψ is a belief formula. (1) encodes the executability condition of action a, *i.e.*, the execution of a is possible only if ψ is true. (2) represents the effects of action a. Intuitively, if the real state of the world and of the beliefs match the condition described by ψ, then the real state of the world is affected by the change that makes the literal ℓ true after the execution of a. (3) encodes a sensing action a which enables the agent(s) to learn the value of the fluent f. (4) states that action a announces the truth value of ℓ.
- Each statement in O is of a following form

$$\alpha \textbf{ observes } a \textbf{ if } \varphi \tag{5}$$
$$\alpha \textbf{ aware\_of } a \textbf{ if } \varphi \tag{6}$$

  where $\alpha \subseteq \mathcal{AG}$ is a set of agents, $a \in \mathcal{A}$, and φ is a fluent formula.
  (5) (resp. (6)) states that if φ is true in the real state of the world then agents in α are fully (resp. partially) aware of the occurrence of the action a.

4.4.7.1 Semantics

Let $\mathcal{S}$ be the set of pointed Kripke structures over $\mathcal{AG}$ and $\mathcal{F}$. The semantics of a dynamic domain (D, O) over $\mathcal{AG}$, $\mathcal{F}$, and $\mathcal{A}$ is defined by a function $\Phi_{(D,O)}$ that maps pairs of actions and pointed Kripke structures in $\mathcal{S}$ to subsets of $\mathcal{S}$, *i.e.*, given an action a and a pointed Kripke structure M, $\Phi_{(D,O)}(a, M)$ is a set of pointed Kripke structures that results from the execution of a in M. It does so by

- constructing the update model of the action occurrence a given M, denoted by $\mathcal{U}(a, M)$;
- defining $\Phi_{(D,O)}(a, M)$ as the update of M by $\mathcal{U}(a, M)$, $M \otimes \mathcal{U}(a, M)$, following *e.g.* [8].

4.4.7.2 Relationship to Epistemic Planning

The language $m\mathcal{A}$ provides a practical tool as well as the foundation for generalizing contemporary planning techniques in single agent domains to epistemic planning in presence of multiple agents. The features of $m\mathcal{A}$ that are suitable for the development of epistemic planning systems are:

- *The separation between the description of effects of actions and the observability of action occurrences of the agents*: In $m\mathcal{A}$, action types and action tokens are clearly separated. Statements of the forms (1) – (4) specify action types. Action types are specialized with respect to a pointed Kripke structure and statements of the forms (1) – (4) to create action tokens. In this sense, $m\mathcal{A}$ addresses the issues raised in the discussion for the need of pure action types in DEL (Sub-section 4.2.1). Given the similarity between action languages for single-agent domains and PDDL and the fact that statements of the forms (1) – (3) have been a part of PDDL, the development of $m\mathcal{A}$ shows that PDDL can be easily extended to specifying multi-agent planning domains. To do so, PDDL needs to be extended with (*i*) the specification of agents; (*ii*) consideration of belief formulas and an additional type of actions (announcement actions); and (*iii*) the specification of observability.
- *Definition for the transition function*: The function $\Phi_{(D,O)}$ can easily be realized and utilized by heuristic search algorithms. Indeed, the function $\Phi_{(D,O)}$ can be implemented in heuristic search planner similar to the implementation of the `result` function of PDDL domains. This paves way for the studying and development of heuristics in epistemic planning.
- *Conditions for the finiteness of the set of pointed Kripke structures satisfying a belief formula and algorithms for computing this set*: The initial state of an epistemic planning problem can be represented by a belief formula. Planners employing pointed Kripke structures as states need to create the initial state by computing the set of pointed Kripke structures satisfying a belief formula. This set, however, might be infinite. As such, conditions for the finiteness of the set of initial pointed Kripke structures and algorithms for computing this set are important for the development of epistemic planners. Such conditions are provided in [87].

4.4.8 Multi-agent planning in general (with epistemic side-effects)

In multi-agent planning, one of the settings usually studied is cooperative planning involving agents with heterogeneous capabilities. An even more specific setting that gained attention recently is privacy-preserving multi-agent planning [77]. While the problem is not epistemic per se (*i.e.*, there are no knowledge goals and preconditions), it certainly possesses some epistemic characteristics. First, some dimensions of the state space can be completely unknown to all but some agents. Private variables can then be changed by private actions that are only known to their owners. The intention is for the agents to come up with a joint plan to achieve their goal without having to give away their private information (variables, actions) in the process.

There already exists a variety of search algorithms [101, 26] as well as different kind of heuristics [93, 68, 91, 92]. An overview of state-of-the-art planners can be found in [94].

4.4.9 Classical planning techniques

For approaches based on compilation to classical planning, or perhaps also to draw inspiration from classical planning algorithms, it is useful to consider the state of the art in classical planning. The three prevalent approaches in the classical planning literature at the moment are heuristic forward search, planning as satisfiability, and planning as symbolic model checking. In planning competitions, for more than a decade almost all successful planning systems in the sequential planning tracks have been based on one of these paradigms or are portfolios of planning systems based on these paradigms. For an overview of planning tech-

niques used in state-of-the-art planning systems, the planner descriptions in the competition booklets of the International Planining Competitions (IPC) are a good starting point (see the "Deterministic Track" entries on http://icaps-conference.org/index.php/Main/Competitions). Most current planning systems based on heuristic forward search are built on top of the Fast Downward planner [47], and some others are based on FF [51]. It should be pointed out that the paper describing the Fast Downward planner is more than a decade old, and the current version of Fast Downward contains many features (in particular more heuristics) not described in the paper. A good place for discussion and practical information about planning as heuristic forward search is the Fast Downward mailing list, linked from the planner homepage (http://www.fast-downward.org). Generally speaking, there is a wide difference in scalability between planning systems that guarantee optimality of solutions (optimal planners) and ones that do not (satisficing planners). Most optimal heuristic search planners use A^* with an admissible heuristic, in some cases enhanced with pruning techniques based on symmetries and partial order reduction. The paper by Helmert and Domshlak [48] is a good starting point to get an overview of heuristics for classical planning, although it has become a bit dated. For satisficing search algorithms, the design space and hence the list of available techniques is larger. In recent work, much emphasis has been placed on guaranteeing sufficient exploration, *i.e.*, making sure to visit diverse parts of the search space, as opposed to exploitation guided by a heuristic. Exploration can be based on randomness or based on systematically favouring states with high novelty. The latter approach has been very successful, and many relevant papers have been published recently by Lipovetzky and Geffner [64]. Planners based on compilations to satisfiability have also been very successful over the last ten years, thanks mainly to the work by Jussi Rintanen [84, 85]. They are a very strong alternative to heuristic search planners in the case of satisficing planning; for optimal planning, no competitive SAT-based planners exist at the moment. SAT-based planners and planners based on heuristic search tend to have different strengths and weaknesses, so both approaches are very worth trying out in the context of epistemic planning. Similar remarks apply to planning approaches based on symbolic search, *i.e.*, BDD-based exploration. Symbolic search planners have been very successful for optimal planning, where they are as strong as the best heuristic planners (again, as with SAT-based planners, the different approaches excel in different domains). A very good reference for BDD-based search techniques is Torralba's PhD thesis [100].

4.4.10 Generalized planning

Generalized planning [19, 52, 62, 89, 90] is an extension of classical planning where the goal is, instead of solving a single planning task, to solve a general class of planning tasks (e.g., all blocks world planning tasks) by essentially synthesizing a program with loops. Semantically, one posits a set of possible worlds, where each world is one particular plan instance. A generalized plan, then, is one which is correct for all world states: that is, in each world state, the plan enables the goal and then terminates (by reaching a distinguished stop state).

Typically, as input it accepts a standard planning problem with one exception: a particular parameter for the initial state is left open. For example, in a blocks world setting, the number of blocks on the table may be left open, in which the possible worlds are all possible instantiations of this parameter. The output can be a sequential plan, a conditional plan or a loopy plan. The obvious advantage of plans created by generalized planners is that they have a broad domain coverage. However in order use a generalized planner for a planning problem instance, an applicability test and then (if the test is successful) a plan instantiation are necessary. These procedures, if they have high complexity, may add computational effort to the planner.

4.5 Correspondence Between Planning Problems and Games

Guillaume Aucher (coordinator), Tim French, Kai Li, Sheila McIlraith, Bernhard Nebel, Ramaswamy Ramanujam (coordinator), Yanjing Wang

The research communities dealing with games and planning problems both deal with similar issues: a number of agents can perform actions in a given context and aim at achieving an objective. However, these two research communities use different vocabularies. For example, a *goal* in a planning problem corresponds to an *objective* in a game; a *plan* or *policy* in a planning problem corresponds to a *strategy* in a game and a *planning domain* corresponds to an *arena* in a game.

Setting rigorous and formal connections between both parties has a number of benefits for both of them. For example, meta-theorems of game theory are of interest to the planning community and heuristic for planning synthesis could be used to synthesize strategies more efficiently. As it turns out, some connections have already been set [37].

The rest of this subsection is organized as follows. First, in Section 4.5.1, we recall the formal setting of games based on graphs. Then, in Section 4.5.2, we recall the *Fully Observable and Non-Deterministic* planning problem (FOND) as it is usually defined in the planning community and we show how it can be reformulated in the setting of games based on graphs. In Section 4.5.3, we discuss what is gained and lost in this translation. Finally, in Section 4.5.4, we discuss which novel perspectives and questions arise in this confrontation between planning problems and games.

4.5.1 Games Based on Graphs

Games are natural models for modeling interaction between multiple agents. In this report, we consider infinite games played on finite graphs [28, 39]. These games are supposed to model reactive systems like an operating system evolving in an arbitrary or hostile environment. The first player represents a program (or "control", here identified with Player 1), and the second player represents the environment (or "disturbance", here identified with Player 2). More precisely, we consider zero-sum two-player turn-based games of infinite duration played on finite graphs. This means that when a player wins a game, the other player loses it ("zero-sum"), that players play in turn ("turn-based"), that the game lasts forever ("infinite duration"), and that plays are determined by a finite graph. The graph describes the possible interactions of the system. The game is of infinite duration, because reactive systems are usually not expected to terminate. As recalled by [29], zero-sum graph games have been widely used in the synthesis (or control) of reactive systems, as well as for defining and checking the realizability of specifications, the compatibility of interfaces, simulation relations between transition systems, and for generating test cases.

We introduce the definitions and notations that we will use in the sequel. They are (almost) identical to the definitions and notations of [28, p. 290-292].

Game Graphs

A *game arena* (of imperfect information) is a tuple $G = (L, l_0, \Sigma, \Delta, \mathcal{O}, \gamma)$ where L is the finite set of states, $l_0 \in L$ is the initial state, $\Delta \subseteq L \times \Sigma \times L$ is a set of labeled transitions,

$\mathcal{O}$ is a finite set of observations and $\gamma : \mathcal{O} \to 2^L - \{\emptyset\}$ maps each observation to the set of states that it represents. We require the following two properties on G:

- *Total*: for all $l \in L$ and all $\sigma \in \Sigma$, there exists $l' \in L$ such that $(l, \sigma, l') \in \Delta$;
- *Partition*: the set $\{\gamma(o) : o \in \mathcal{O}\}$ partitions L.

When the partition is composed of singletons only, the game is a *game of perfect information*. The non-determinism of the transition function is very important since it allows players to play in turn.

Plays and Strategies

In a game arena, at each turn, Player 1 chooses a letter in Σ, and Player 2 resolves nondeterminism by choosing the successor state. A *play* in G is an infinite sequence $\pi = l_0\sigma_0 l_1 \ldots l_n\sigma_n l_{n+1} \ldots$ such that for all $i \geq 0$, $(l_i, \sigma_i, l_{i+1}) \in \Delta$. The *prefix up to* l_n of the play π is denoted $\pi(0, n)$; its length is $|\pi(0, n)| = n + 1$; and its last element is $\mathsf{Last}(\pi(0, n)) = l_n$. The *observation sequence* of π is the unique infinite sequence $\gamma^-(\pi) = o_0 o_1 \ldots o_{n-1} o_n \ldots$ such that for all $i \geq 0$, we have $l_i \in \gamma(o_i)$. Similarly, the *observation sequence* $\pi(0, n)$ is the prefix up to o_n of $\gamma^-(\pi)$. The set of infinite plays in G is denoted $\mathsf{Plays}(G)$, and the set of corresponding finite prefixes is denoted $\mathsf{Prefs}(G)$.

A *deterministic strategy* in G for Player 1 is a function $\alpha : \mathsf{Prefs}(G) \to \Sigma$. A strategy α for Player 1 is *observation-based* if for all prefixes $\rho, \rho' \in \mathsf{Prefs}(G)$, if $\gamma^-(\rho) = \gamma^-(\rho')$, then $\alpha(\rho) = \alpha(\rho')$. A strategy for Player 1 is *memory-less* if for all prefixes $\rho, \rho' \in \mathsf{Prefs}(G)$ such that $\mathsf{Last}(\rho) = \mathsf{Last}(\rho')$, $\alpha(\rho) = \alpha(\rho')$. A *deterministic strategy* for Player 2 is a function $\beta : \mathsf{Prefs}(G) \times \Sigma \to L$ such that for all $\rho \in \mathsf{Prefs}(G)$ and all $\sigma \in \Sigma$, we have $(\mathsf{Last}(\rho), \sigma, \beta(\rho, \sigma)) \in \Delta$. We denote by $\mathcal{A}_G$ the set of all strategies for Player 1.

The *outcome* of two deterministic strategies α (for Player 1) and β (for Player 2) in G is the play $\pi = l_0\sigma_0 l_1 \ldots \sigma_{n-1} l_n \sigma_n \ldots \in \mathsf{Plays}(G)$ such that for all $i \geq 0$, we have $\sigma_i = \alpha(\pi(0, i))$ and $l_{i+1} = \beta(\pi(0, i), \sigma_i)$. This play is denoted $\mathsf{outcome}(G, \alpha, \beta)$. The *outcome set* of the deterministic strategy α for Player 1 in G is the set $\mathsf{Outcome}_1(G, \alpha)^\omega$ of plays π such that there exists a deterministic strategy β for Player 2 with $\pi = \mathsf{outcome}(G, \alpha, \beta)$. The set $\mathsf{Outcome}_1(G, \alpha)^*$ denotes the set of finite prefixes of $\mathsf{Outcome}_1(G, \alpha)^\omega$.

Objectives

An *objective* for a game of imperfect information is a set $\mathcal{K} \subseteq (\mathcal{O} \times \Sigma)^\omega$ which is assumed to be Borel measurable: a Borel objective is a Borel set in the Cantor topology on $(\mathcal{O} \times E)^\omega$ [55]. Reachability, safety, Büchi, coBüchi, and parity objectives are all Borel measurable. The parity objectives are a canonical form to express all ω-regular objectives [98]. An objective is induced by a set of target observations $\mathcal{T} \subseteq \mathcal{O}$. Here are some examples of objectives:

- $\mathsf{Reach}^\omega_\mathcal{T} := \{l_0\sigma_0 l_1\sigma_1 \ldots \in \mathsf{Plays}(G) : \text{there is } k \geq 0, \text{there is } o \in \mathcal{T} \text{ such that } k \in \gamma(o)\}$. The *reachability objective* $\mathsf{Reach}^\omega_\mathcal{T}$ requires that an observation in $\mathcal{T}$ be visited at least once;
- $\mathsf{Safe}^\omega_\mathcal{T} := \{l_0\sigma_0 l_1\sigma_1 \ldots \in \mathsf{Plays}(G) : \gamma^-(l_k) \in \mathcal{T} \text{ for all } k \geq 0\}$. The *safety objective* $\mathsf{Safe}^\omega_\mathcal{T}$ requires that only observations in $\mathcal{T}$ be visited;
- $\mathsf{Buchi}^\omega_\mathcal{T} := \{\pi : \mathsf{Inf}(\pi) \cap \mathcal{T} \neq \emptyset\}$. The *Büchi objective* $\mathsf{Buchi}^\omega_\mathcal{T}$ requires that an observation in $\mathcal{T}$ be visited infinitely often;
- $\mathsf{Parity}(p)^\omega := \{\pi : min\{p(o) : o \in \mathsf{Inf}(\pi)\} \text{ is even}\}$, where $p : \mathcal{O} \to \{0, 1, \ldots, d\}$ is a *priority function* that maps each observation to a nonnegative integer priority (where $d \in \mathbb{N}$). The *parity objective* $\mathsf{Parity}(p)^\omega$ requires that the minimum priority that appears infinitely often be even.

In this report, we consider only the *reachability objective* $\mathsf{Reach}^{\omega}_{\mathcal{T}}$. A *game* $\mathcal{G} = \{G, \mathcal{T}\}$ is a pair consisting of a game arena and a set of target observations $\mathcal{T}$. Then, we say that a strategy $\alpha \in \mathcal{A}_G$ is *winning* for $\mathsf{Reach}^{\omega}_{\mathcal{T}}$ when $\mathsf{Outcome}_1(G, \alpha)^{\omega} \subseteq \mathsf{Reach}^{\omega}_{\mathcal{T}}$.

4.5.2 Translation from Planning Problems to Games

Planning problems in the most general sense are specified over some logic (often propositional logic), where the initial situation and the goals are given as formulas and the dynamics of the world is described by action models. These action models induce a transition system that is similar to a game arena. Depending on the form of the initial situation description and the action description language, you end up with different planning problems. In this report, we only consider Fully Observable Non-Deterministic (FOND) planning problems, but we conjecture that the result that we obtain can be extended to other forms of planning problems such as *Partially Observable Non-Deterministic* planning problems (POND) and classical planning problems.

FOND Planning

Following [44], a *Fully Observable Non-Deterministic* (FOND) planning problem consists of a tuple $\langle \mathcal{F}, \mathcal{I}, \mathcal{G}, \mathcal{A} \rangle$, where $\mathcal{F}$ is a set of propositions that we call *fluents*, $\mathcal{I} \subseteq \mathcal{F}$ characterizes what holds in the initial state; $\mathcal{G} \subseteq \mathcal{F}$ characterizes what must hold for the goal to be achieved. Finally $\mathcal{A}$ is the set of actions. The set of literals of $\mathcal{F}$ is $\mathit{Lits}(\mathcal{F}) = \mathcal{F} \cup \{\neg f \mid f \in \mathcal{F}\}$.

Each action $a \in \mathcal{A}$ is associated with $\langle \mathit{Pre}_a, \mathit{Eff}_a \rangle$, where $\mathit{Pre}_a \subseteq \mathit{Lits}(\mathcal{F})$ is the precondition

and Eff_a is a set of outcomes of a.

Each outcome $e \in \mathit{Eff}_a$ is a set of conditional effects, each of the form $(C \to \ell)$, where $C \subseteq \mathit{Lits}(\mathcal{F})$ and $\ell \in \mathit{Lits}(\mathcal{F})$. Given a planning state $s \subseteq \mathcal{F}$ and a fluent $f \in \mathcal{F}$, we say that s satisfies f, denoted $s \models f$ iff $f \in s$. In addition $s \models \neg f$ if $f \notin s$, and $s \models L$ for a set of literals L, if $s \models \ell$ for every $\ell \in L$.

Action a is *applicable* in state s if $s \models \mathit{Pre}_a$. We say s' *is a result of applying* a *in* s iff, for one outcome

e in Eff_a, s' is equal to $s \setminus \{p \mid (C \to \neg p) \in e, s \models C\} \cup \{p \mid (C \to p) \in e, s \models C\}$.

Unlike classical planning, where a solution is a sequence of actions, a solution to a FOND planning task is a policy that maps a state to an appropriate action (so that the agent eventually reaches the goal). A state s is said to be *reachable* by a policy if the policy can lead the agent to s, and a policy is *closed* if it returns an action for every non-goal state that a policy reaches. When the agent executes a non-deterministic action the effect is randomly chosen, so a closed policy must handle every possible outcome of an action it returns. There are three primary variations of plans for a FOND planning problem [35]: *weak*, *strong*, and *strong cyclic*.

▶ **Definition 1** (Weak Plan [35])**.** A *weak plan* is a policy that achieves the goal with at least one possible set of outcomes for the non-deterministic actions.

A weak plan may be as simple as a sequence of actions that achieves the goal with a particular setting of non-deterministic action outcomes. The policy for a weak plan need not be closed.

▶ **Definition 2** (Strong Plan [35])**.** A *strong plan* is a closed policy that achieves the goal and never visits the same state twice.

A strong plan provides a guarantee on the maximum number of steps to achieve the goal but is often too restrictive. For example, a strong plan cannot contain an action that can fail with no effect on the state.

▶ **Definition 3** (Strong Cyclic Plan [35]). A *strong cyclic plan* is a closed policy where the goal is reachable from every state reachable using the policy.

A strong cyclic plan guarantees that the agent eventually reaches the goal, but does not guarantee that the agent can do so in a fixed number of steps. Researchers in the planning community are primarily interested in computing strong cyclic plans, as they are guaranteed to achieve the goal under an assumption of fairness – every outcome of a non-deterministic action will occur infinitely often if the action is executed infinitely often.

▶ Remark. The text in this section is taken near-verbatim from the following two sources: [72, 73].

Translation Theorem

▶ **Theorem 4.** *Given a FOND planning problem* $\langle \mathcal{F}, \mathcal{I}, \mathcal{G}, \mathcal{A} \rangle$*, one can define an associated game graph* $G = (L, l_0, \Sigma, \Delta, \mathcal{O}, \gamma)$ *of perfect information and a target observation* $\mathcal{T}$ *such that:*

- $\Sigma = \mathcal{A}$*;*
- *there exists a bijective mapping* f *between sets of fluents and states of* L *such that* $f(\mathcal{F}) = l_0$ *and* $f(\mathcal{G}) = \bigcup_{t \in \mathcal{T}} \gamma^-(t)$*;*
- *every policy* π *can be mapped to a (memory-less) strategy in the associated game* G*, denoted* $g(\pi)$*, such that for all reachable states* $s \subseteq \mathcal{F}$*,* $f(\pi(s)) = g(\pi)(f(s))$*;*
- *there exists a strong plan* π *in the original FOND planning problem if, and only if, Player 1 has a winning strategy* $g(\pi)$ *in the associated game of perfect information* G *with reachability objective* $\mathsf{Reach}^{\omega}_{\mathcal{T}}$*.*

We conjecture that this theorem can be extended to other forms of planning, such as planning with *partial observability* and even *implicitly coordinated* planning [41].

Muise et al. [75] have informally identified the same relationship between multi-agent games and FOND planning. They modified an existing FOND planner to find strong and weak solutions for first-person multi-agent planning problems. They treated the outcome of other players moves as non-deterministic outcomes of the actions from the first-person perspective, but used standard planning heuristics to prioritize which states to explore given a fixed time budget. Calculating a strong solution for this is equivalent to calculating a winning strategy for the game.

4.5.3 Lost and Gained in Translation

Translating planning problems into games has a number of benefits, but also some shortcomings. On the one hand, the succinctness of the description of a planning problem and in particular of actions is lost when it is translated into a game problem. Moreover, the planning approach allows to provide a more fine–grained and rich description of actions than in the game–theoretical reformulation of the problem, where actions are simply represented as (labeled) transitions between states. On the other hand, one inherits from the game–theoretical formulation some meta–theorems such as determinacy and formal methods to synthesize the plans/winning strategies of the planning problem/game. Moreover, we can also consider other objectives than mere reachability objectives that deal with the overall desired behavior of the system when a plan/strategy is executed.

We should also point out some important differences between planning problems and games. Games are often non-cooperative while in planning the agents (most often) cooperate to achieve a common objective. Moreover, the game arena is usually assumed to be common knowledge among agents, while in planning transition systems are often not explicitly given.

4.5.4 New Perspectives

While the idea of mapping between multi-player co-ordination games and multi-agent planning is to make use of techniques and results from one domain in the other, it is also to instructive to examine ideas from one to widen perspectives in the other, in the process generating new interesting questions for study in that domain. Below we list some instances of such possibilities.

Distributed Games

In the context of coordination games against an environment, one interesting question that arises is the existence of not merely a winning strategy for the coalition, but a *distributed* one: these are strategies played out by players locally, based on their partial views of game position. These are imperfect information games and strategy existence is in general undecidable [80, 79] for synchronous executions. The winning conditions in these games are naturally definable in epistemic temporal logics and correspond to distributed plans for multi-agent systems.

In synchronous play, the strategy question is decidable for broadcast environments and hierarchical systems of agents (where if $i \leq j$ and $s \sim_i s' \rightarrow s \sim_j s'$; i is j's boss, so whatever information is available to j is also available to i). It is also decidable when the goal formulas have local structure. Mapping the decidable subclasses is an interesting exercise.

When players move asynchronously, the strategy question has been shown to be decidable for several subclasses of communication architectures, but the general question is open.

All these questions have an obvious reformulation in the world of multi-agent epistemic planning, and can be answered by mapping planning into games. In the process, we may have new subclasses inspired by plans enriching game theory as well.

Large Games

Games with a large number of players can be called large games. In such games, payoffs are associated not with strategy profiles but with choice distributions. Repeated normal forms of this kind can lead to herd behavior etc. Studying planning in the context of multi-agent systems with a large number of agents might be interesting as well.

Goals, Preferences

In planning, it is customary to distinguish between hard goals (that need to be met definitely) and soft ones (where we have a preference for achieving them over otherwise, but non-achievement is OK). Notions of rationality may be relevant in such goal setting. In game theory, an important question is where utilities arise from and the structure of outcomes. This again suggests a close relationship between plans and games.

Another interesting question in the multi-agent case is coordination by separating 'social' goals from individual agent goals: as long as the social goals are achieved, agents might behave competitively. Such behavior has been studied in game theory and it may be relevant for some planning problems.

References

1. Rajeev Alur, Thomas A Henzinger, and Orna Kupferman. Alternating-time temporal logic. *Journal of the ACM (JACM)*, 49(5):672–713, 2002.
2. Mikkel Andersen, Thomas Bolander, and Martin Jensen. Conditional epistemic planning. *Logics in Artificial Intelligence*, pages 94–106, 2012.
3. Maduka Attamah, Hans van Ditmarsch, Davide Grossi, and Wiebe van der Hoek. Knowledge and gossip. In *ECAI 2014 - 21st European Conference on Artificial Intelligence, 18-22 August 2014, Prague, Czech Republic - Including Prestigious Applications of Intelligent Systems (PAIS 2014)*, pages 21–26, 2014.
4. Guillaume Aucher. Intricate Axioms as Interaction Axioms. *Studia Logica*, pages 1035–1062, 2015.
5. Guillaume Aucher, Bastien Maubert, and Sophie Pinchinat. Automata techniques for epistemic protocol synthesis. *arXiv preprint arXiv:1404.0844*, 2014.
6. Jean-François Baffier, Man-Kwun Chiu, Yago Diez, Matias Korman, Valia Mitsou, André van Renssen, Marcel Roeloffzen, and Yushi Uno. Hanabi is np-hard, even for cheaters who look at their cards. *Theor. Comput. Sci.*, 675:43–55, 2017.
7. Christel Baier, Joost-Pieter Katoen, and Kim Guldstrand Larsen. *Principles of model checking*. MIT press, 2008.
8. A. Baltag and L. Moss. Logics for epistemic programs. *Synthese*, 2004.
9. Alexandru Baltag, Lawrence S. Moss, and Slawomir Solecki. The logic of public announcements and common knowledge and private suspicions. In *Proceedings of the 7th Conference on Theoretical Aspects of Rationality and Knowledge (TARK-98), Evanston, IL, USA, July 22-24, 1998*, pages 43–56, 1998.
10. Alexandru Baltag and Bryan Renne. Dynamic epistemic logic. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, winter 2016 edition, 2016.
11. Alexandru Baltag and Sonja Smets. Probabilistic dynamic belief revision. *Synthese*, 165(2):179, 2008.
12. Alexandru Baltag and Sonja Smets. *A Qualitative Theory of Dynamic Interactive Belief Revision*, pages 813–858. Springer International Publishing, Cham, 2016.
13. Chitta Baral and Gregory Gelfond. On representing actions in multi-agent domains. In M. Balduccini and T. Son, editors, *Logic programming, knowledge representation, and nonmonotonic reasoning*, pages 213–232. Springer, 2011.
14. Chitta Baral, Gregory Gelfond, Enrico Pontelli, and Tran Cao Son. An action language for reasoning about beliefs in multi-agent domains. In *Proceedings of NMR*, 2012.
15. Chitta Baral, Gregory Gelfond, Enrico Pontelli, and Tran Cao Son. Reasoning about the beliefs of agents in multi-agent domains in the presence of state constraints: The action language mal. In J. Leite, T. C. Son, P. Torroni, L. van deer Torre, and S. Woltran, editors, *Proceedings of the 14th International Workshop, Computational Logic in Multi-Agent Systems, CLIMA VIX, Coruna, Spain, September 16-18, 2013*, volume 8143 of *Lecture Notes in Computer Science*, pages 290–306. Springer, 2013.
16. Chitta Baral, Gregory Gelfond, Enrico Pontelli, and Tran Cao Son. An Action Language for Mutli-Agent Domains: Foundations. *arXiv.org*, page http://arxiv.org/abs/1511.01960, 2015.
17. Chitta Baral and Michael Gelfond. Representing concurrent actions in extended logic programming. In *IJCAI*, pages 866–873, 1993.
18. Vaishak Belle and Gerhard Lakemeyer. Multiagent only knowing in dynamic systems. *Journal of Artificial Intelligence Research*, 49:363–402, 2014.
19. Vaishak Belle and Hector J Levesque. Foundations for generalized planning in unbounded stochastic domains. In *KR*, pages 380–389, 2016.

20 Daniel S. Bernstein, Robert Givan, Neil Immerman, and Shlomo Zilberstein. The complexity of decentralized control of markov decision processes. *Math. Oper. Res.*, 27(4):819–840, 2002.

21 Meghyn Bienvenu, Hélene Fargier, and Pierre Marquis. Knowledge compilation in the modal logic s5. In *AAAI*, 2010.

22 Thomas Bolander. Seeing is believing: Formalising false-belief tasks in dynamic epistemic logic. In Andreas Herzig and Emiliano Lorini, editors, *Proceedings of the European Conference on Social Intelligence (ECSI-2014)*, volume 1283 of *CEUR Workshop Proceedings*, pages 87–107. CEUR-WS.org, 2014.

23 Thomas Bolander and Mikkel Birkegaard Andersen. Epistemic planning for singleand multi-agent systems. *Journal of Applied Non-Classical Logics*, 2011.

24 Craig Boutilier, Ronen I Brafman, et al. Planning with concurrent interacting actions. In *AAAI/IAAI*, pages 720–726, 1997.

25 Laura Bozzelli, Bastien Maubert, and Sophie Pinchinat. Uniform strategies, rational relations and jumping automata. *Information and Computation*, 242:80–107, 2015.

26 Ronen I Brafman and Israel Be'er Sheva. A privacy preserving algorithm for multi-agent planning and search. In *IJCAI*, pages 1530–1536, 2015.

27 Florent Capelli. *Structural restriction of CNF-formulas: Application to model counting and knowledge compilation.* PhD thesis, Université Paris Diderot, 2016.

28 Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger, and Jean-François Raskin. Algorithms for omega-regular games with imperfect information. In *CSL*, pages 287–302, 2006.

29 Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. Strategy logic. *Inf. Comput.*, 208(6):677–693, 2010.

30 David Chaum. The dining cryptographers problem: Unconditional sender and recipient untraceability. *J. Cryptology*, 1(1):65–75, 1988.

31 Brian F Chellas. *Modal logic: an introduction.* Cambridge university press, 1980.

32 Jens Claßen and Gerhard Lakemeyer. Foundations for knowledge-based programs using $\mathcal{ES}$. In *KR*, pages 318–328, 2006.

33 Martin C Cooper, Andreas Herzig, Faustine Maffre, Frédéric Maris, and Pierre Régnier. A simple account of multi-agent epistemic planning. In *ECAI*, pages 193–201, 2016.

34 Martin C. Cooper, Andreas Herzig, Faustine Maffre, Frédéric Maris, and Pierre Régnier. Simple epistemic planning: generalised gossiping. *CoRR*, abs/1606.03244, 2016.

35 Marco Daniele, Paolo Traverso, and Moshe Y. Vardi. Strong cyclic planning revisited. In *5th European Conference on Planning*, pages 35–48. Springer, 2000.

36 Adnan Darwiche and Pierre Marquis. A knowledge compilation map. *Journal of Artificial Intelligence Research*, 17(1):229–264, 2002.

37 Giuseppe De Giacomo, Aniello Murano, Sasha Rubin, and Antonio Di Stasio. Imperfect-information games and generalized planning. In Subbarao Kambhampati, editor, *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 1037–1043. IJCAI/AAAI Press, 2016.

38 Hans van Ditmarsch. Prolegomena to dynamic logic for belief revision. *Synthese*, 147(2):229–275, 2005.

39 Laurent Doyen and Jean-François Raskin. Games with imperfect information: Theory and algorithms. *Lecture Notes in Game Theory for Computer Scientists*, pages 185–212, 2011.

40 Thorsten Engesser, Thomas Bolander, Robert Mattmüller, and Bernhard Nebel. Cooperative epistemic multi-agent planning for implicit coordination. *arXiv preprint arXiv:1703.02196*, 2017.

41 Thorsten Engesser, Thomas Bolander, Robert Mattmüller, and Bernhard Nebel. Cooperative epistemic multi-agent planning for implicit coordination. In Sujata Ghosh and

R. Ramanujam, editors, *Proceedings of the Ninth Workshop on Methods for Modalities, M4M@ICLA 2017, Indian Institute of Technology, Kanpur, India, 8th to 10th January 2017.*, volume 243 of *EPTCS*, pages 75–90, 2017.

42 Peter Gärdenfors and Hans Rott. Belief revision. In Dov M. Gabbay, C. J. Hogger, and J. A. Robinson, editors, *Handbook of Logic in Artificial Intelligence and Logic Programming (Vol. 4)*, pages 35–132. Oxford University Press, Oxford, UK, 1995.

43 M. Gelfond and V. Lifschitz. Representing actions and change by logic programs. *Journal of Logic Programming*, 17(2,3,4):301–323, 1993.

44 Malik Ghallab, Dana Nau, and Paolo Traverso. *Automated planning: theory & practice*. Elsevier, 2004.

45 Joseph Y Halpern and Judea Pearl. Causes and explanations: A structural-model approach. part ii: Explanations. *The British journal for the philosophy of science*, 56(4):889–911, 2005.

46 Sandra Mitchell Hedetniemi, Stephen T. Hedetniemi, and Arthur L. Liestman. A survey of gossiping and broadcasting in communication networks. *Networks*, 18(4):319–349, 1988.

47 Malte Helmert. The fast downward planning system. *Journal of Artificial Intelligence Research*, 26:191–246, 2006.

48 Malte Helmert and Carmel Domshlak. Landmarks, critical paths and abstractions: What's the difference anyway? In *Proceedings of the 19th International Conference on Automated Planning and Scheduling (ICAPS)*, 2009.

49 Andreas Herzig, Jérôme Lang, Dominique Longin, and Thomas Polacsek. A logic for planning under partial observability. In *AAAI/IAAI*, pages 768–773, 2000.

50 Andreas Herzig and Faustine Maffre. How to share knowledge by gossiping. *AI Communications*, 30(1):1–17, 2017.

51 Jörg Hoffmann and Bernhard Nebel. The ff planning system: Fast plan generation through heuristic search. *Journal of Artificial Intelligence Research*, 14:253–302, 2001.

52 Yuxiao Hu and Giuseppe De Giacomo. A generic technique for synthesizing bounded finite-state controllers. In *ICAPS*, 2013.

53 Xiao Hung, Biqing Gang, Hai Wan, and Yongmei Liu. A general multi-agent epistemic planner based on higher-order belief change. In *IJCAI 2017, Proceedings of the 26th International Joint Conference on Artificial Intelligence, Melbourne, Australia, Sydney, 2017*, 2017.

54 Matthew Johnson, Catholijn M Jonker, M Birna van Riemsdijk, Paul J Feltovich, and Jeffrey M Bradshaw. Joint activity testbed: Blocks world for teams (bw4t). In *ESAW*, volume 9, pages 254–256. Springer, 2009.

55 Alexander S Kechris. *Classical descriptive set theory*, volume 156. Springer-Verlag, New York, 1995.

56 Craig A Knoblock. Generating parallel execution plans with a partial order planner. In *AIPS*, volume 94, pages 98–103, 1994.

57 Filippos Kominis and Hector Geffner. Beliefs in multiagent planning: From one agent to many. In *Proceedings of the Twenty-Fifth International Conference on Automated Planning and Scheduling, ICAPS 2015, Jerusalem, Israel, June 7-11, 2015.*, pages 147–155, 2015.

58 Barteld Kooi and Bryan Renne. Generalized arrow update logic. In *Proceedings of the 13th Conference on Theoretical Aspects of Rationality and Knowledge*, pages 205–211. ACM, 2011.

59 Gerhard Lakemeyer and Hector J Levesque. Situations, si! situation terms, no! In *KR*, pages 516–526, 2004.

60 Gerhard Lakemeyer and Hector J Levesque. Decidable reasoning in a fragment of the epistemic situation calculus. In *KR*, 2014.

61 Noël Laverny and Jérôme Lang. From knowledge-based programs to graded belief-based programs, part I: on-line reasoning*. *Synthese*, 147(2):277–321, 2005.

62 Hector Levesque. What is planning in the presence of sensing? In *Proceedings of the 13th National Conference on Artificial Intelligence (AAAI)*, pages 1139–1146, Portland, Oregon, 1996.

63 Fangzhen Lin and Yoav Shoham. Concurrent actions in the situation calculus. In *AAAI*, volume 92, pages 590–595, 1992.

64 Nir Lipovetzky and Hector Geffner. Best-first width search: Exploration and exploitation in classical planning. In *AAAI*, pages 3590–3596, 2017.

65 Alessio Lomuscio, Hongyang Qu, and Franco Raimondi. Mcmas: An open-source model checker for the verification of multi-agent systems. *International Journal on Software Tools for Technology Transfer*, 19(1):9–30, 2017.

66 Benedikt Löwe, Eric Pacuit, Andreas Witzel, et al. Del planning and some tractable cases. *LORI*, 6953:179–192, 2011.

67 Faustine Maffre. *Ignorance is bliss : observability-based dynamic epistemic logics and their applications*. PhD thesis, Université Paul Sabatier - Toulouse III, September 2016.

68 Shlomi Maliah, Guy Shani, and Roni Stern. Privacy preserving landmark detection. In *Proceedings of the Twenty-first European Conference on Artificial Intelligence*, pages 597–602. IOS Press, 2014.

69 Sheila A. McIlraith and Richard B. Scherl. What sensing tells us: Towards a formal theory of testing for dynamical systems. In *Proceedings of the Seventeenth National Conference on Artificial Intelligence (AAAI)*, pages 483–490, 2000.

70 Tim Miller and Christian J. Muise. Belief update for proper epistemic knowledge bases. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 1209–1215, 2016.

71 Fabio Mogavero, Aniello Murano, Giuseppe Perelli, and Moshe Y Vardi. Reasoning about strategies: On the model-checking problem. *ACM Transactions on Computational Logic (TOCL)*, 15(4):34, 2014.

72 Christian Muise. *Exploiting Relevance to Improve Robustness and Flexibility in Plan Generation and Execution*. PhD thesis, University of Toronto, 2014.

73 Christian Muise, Sheila A. McIlraith, and J. Christopher Beck. Improved Non-deterministic Planning by Exploiting State Relevance. In *The 22nd International Conference on Automated Planning and Scheduling*, The 22nd International Conference on Automated Planning and Scheduling, 2012.

74 Christian J Muise, Vaishak Belle, Paolo Felli, Sheila A McIlraith, Tim Miller, Adrian R Pearce, and Liz Sonenberg. Planning over multi-agent epistemic states: A classical planning approach. In *AAAI*, pages 3327–3334, 2015.

75 Christian J. Muise, Paolo Felli, Tim Miller, Adrian R. Pearce, and Liz Sonenberg. Planning for a single agent in a multi-agent environment using FOND. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 3206–3212, 2016.

76 J. A. M. Nijssen and M. H. M. Winands. Monte-carlo tree search for the game of scotland yard. In *2011 IEEE Conference on Computational Intelligence and Games (CIG'11)*, pages 158–165, Aug 2011.

77 Raz Nissim and Ronen I Brafman. Distributed heuristic forward search for multi-agent planning. *J. Artif. Intell. Res.(JAIR)*, 51:293–332, 2014.

78 Ronald PA Petrick and Fahiem Bacchus. Extending the knowledge-based approach to planning with incomplete information and sensing. In *ICAPS*, pages 2–11, 2004.

79 Amir Pnueli and Roni Rosner. On the synthesis of a reactive module. In *POPL*, pages 179–190, 1989.

80 Amir Pnueli and Roni Rosner. On the synthesis of an asynchronous reactive module. In *Automata, Languages and Programming*, pages 652–671. Springer, 1989.

81 Raymond Reiter. Natural actions, concurrency and continuous time in the situation calculus. *KR*, 96:2–13, 1996.

82 Raymond Reiter. *Knowledge in Action: Logical Foundations for Specifying and Implementing Dynamical Systems*. MIT Press, Cambridge, MA, 2001.

83 Raymond Reiter. On knowledge-based programming with sensing in the situation calculus. *ACM Trans. Comput. Log.*, 2(4):433–457, 2001.

84 Jussi Rintanen. Planning as satisfiability: Heuristics. *Artificial Intelligence*, 193:45–86, 2012.

85 Jussi Rintanen, Keijo Heljanko, and Ilkka Niemelä. Planning as satisfiability: parallel plans and algorithms for plan search. *Artificial Intelligence*, 170(12-13):1031–1080, 2006.

86 Wally Smith, Frank Dignum, and Liz Sonenberg. The construction of impossibility: a logic-based analysis of conjuring tricks. *Frontiers in Psychology*, 7:748, 2016.

87 Tran Cao Son, Enrico Pontelli, Chitta Baral, and Gregory Gelfond. Finitary s5-theories. In *European Workshop on Logics in Artificial Intelligence*, pages 239–252. Springer, 2014.

88 Tran Cao Son, Enrico Pontelli, Chitta Baral, and Gregory Gelfond. Exploring the kd45 property of a kripke model after the execution of an action sequence. In *AAAI*, pages 1604–1610, 2015.

89 Siddharth Srivastava. Foundations and applications of generalized planning. *AI Communications*, 24(4):349–351, 2011.

90 Siddharth Srivastava, Neil Immerman, and Shlomo Zilberstein. Computing applicability conditions for plans with loops. In *ICAPS*, pages 161–168, 2010.

91 Michal Štolba, Daniel Fišer, and Antonın Komenda. Admissible landmark heuristic for multi-agent planning. In *International Conference on Automated Planning and Scheduling (ICAPS)*, 2015.

92 Michal Štolba, Daniel Fišer, and Antonın Komenda. Potential heuristics for multi-agent planning. In *Proceedings of the Twenty-Sixth International Conference on International Conference on Automated Planning and Scheduling*, pages 308–316. AAAI Press, 2016.

93 Michal Stolba and Antonín Komenda. Relaxation heuristics for multiagent planning. In *ICAPS*, 2014.

94 Michal Štolba, Antonın Komenda, and Daniel L Kovacs. Competition of distributed and multiagent planners (codmap). *The International Planning Competition (WIPC-15)*, 24, 2015.

95 Milind Tambe. *Security and game theory: algorithms, deployed systems, lessons learned*. Cambridge University Press, 2011.

96 Michael Thielscher. A general game description language for incomplete information games. In *Twenty-Fourth AAAI Conference on Artificial Intelligence*, pages 994–999, 2010.

97 Michael Thielscher. Gdl-iii: A description language for epistemic general game playing. In *Proceedings of the International Joint Conference on Artificial Intelligence*, Melbourne, August 2017. AAAI Press.

98 Wolfgang Thomas. *Handbook of formal languages, vol. 3: beyond words*, chapter Languages, automata, and logic. Springer-Verlag, New York, 1997.

99 Wolfgang Thomas et al. *Automata, logics, and infinite games: a guide to current research*, volume 2500. Springer Science & Business Media, 2002.

100 Alvaro Torralba. Symbolic search and abstraction heuristics for cost-optimal planning. *Universidad Carlos III de Madrid*, 2015.

101 Alejandro Torreno, Eva Onaindia, and Oscar Sapena. Fmap: Distributed cooperative multi-agent planning. *Applied Intelligence*, 41(2):606–626, 2014.

102 Johan van Benthem. Dynamic logic for belief revision. *Journal of Applied Non-Classical Logics*, 17(2):129–155, 2007.

103 Johan Van Benthem, Jan Van Eijck, Malvin Gattinger, and Kaile Su. Symbolic model checking for dynamic epistemic logic. In *International Workshop on Logic, Rationality and Interaction*, pages 366–378. Springer, 2015.

104 Wiebe Van der Hoek and Michael Wooldridge. Cooperation, knowledge, and time: Alternating-time temporal epistemic logic and its applications. *Studia Logica*, 75(1):125–157, 2003.

105 Hans van Ditmarsch. *Knowledge games*. PhD thesis, Groningen University, 2000.

106 Hans van Ditmarsch. Epistemic gossip protocols. In *Proceedings of the 1st Chinese Conference on Logic and Argumentation (CLAR 2016), Hangzhou, China, April 2-3, 2016.*, page 1, 2016.

107 Hans van Ditmarsch, Davide Grossi, Andreas Herzig, Wiebe van der Hoek, and Louwe B Kuijer. Parameters for epistemic gossip problems. *Proc. LOFT 2016*, 2016.

108 Hans van Ditmarsch and Barteld Kooi. *One Hundred Prisoners and a Light Bulb*. Springer, 2015.

109 Hans van Ditmarsch, Jan van Eijck, Pere Pardo, Rahim Ramezanian, and François Schwarzentruber. Epistemic protocols for dynamic gossip. *J. Applied Logic*, 20:1–31, 2017.

110 Hans P van Ditmarsch, Wiebe van der Hoek, and Barteld P Kooi. Concurrent dynamic epistemic logic for mas. In *Proceedings of the second international joint conference on Autonomous agents and multiagent systems*, pages 201–208. ACM, 2003.

111 Hans P van Ditmarsch, Wiebe Van Der Hoek, Ron Van Der Meyden, and Ji Ruan. Model checking russian cards. *Electronic Notes in Theoretical Computer Science*, 149(2):105–123, 2006.

112 J. van Eijck and K. Li. Conditional belief, knowledge and probability. In J. Lang, editor, *Proc. of 16th TARK*, 2017. to appear.

113 Jan van Eijck and Yanjing Wang. Propositional dynamic logic as a logic of belief revision. In *Logic, Language, Information and Computation, 15th International Workshop, WoLLIC 2008, Edinburgh, UK, July 1-4, 2008, Proceedings*, pages 136–148, 2008.

114 Moshe Y Vardi. Alternating automata and program verification. In *Computer Science Today*, pages 471–485. Springer, 1995.

Participants

- Guillaume Aucher
IRISA – Rennes, FR
- Chitta Baral
Arizona State University – Tempe, US
- Vaishak Belle
University of Edinburgh, GB
- Thomas Bolander
Technical University of Denmark – Lyngby, DK
- Tran Cao Son
New Mexico State University, US
- Tristan Charrier
INRIA – Rennes, FR
- Jens Claßen
RWTH Aachen, DE
- Thorsten Engesser
Universität Freiburg, DE
- Esra Erdem
Sabanci University – Istanbul, TR
- Tim French
The Univ. of Western Australia, AU
- Malvin Gattinger
University of Amsterdam, NL
- Nina Gierasimczuk
Technical University of Denmark – Lyngby, DK
- Malte Helmert
Universität Basel, CH
- Andreas Herzig
Paul Sabatier University – Toulouse, FR
- Mathias Justesen
Technical University of Denmark – Lyngby, DK
- Ioannis Kokkinis
LORIA – Nancy, FR
- Filippos Kominis
UPF – Barcelona, ES
- Barteld Kooi
University of Groningen, NL
- Louwe B. Kuijer
University of Liverpool, GB
- Jérôme Lang
University Paris-Dauphine, FR
- Yves Lesperance
York University – Toronto, CA
- Kai Li
Peking University, CN
- Yongmei Liu
Sun Yat-sen University – Guangzhou, CN
- Robert Mattmüller
Universität Freiburg, DE
- Bastien Maubert
University of Napoli, IT
- Sheila McIlraith
University of Toronto, CA
- Timothy Miller
The University of Melbourne, AU
- Bernhard Nebel
Universität Freiburg, DE
- Andrés Occhipinti Liberman
Technical University of Denmark – Lyngby, DK
- Ron Petrick
Heriot-Watt University – Edinburgh, GB
- Sophie Pinchinat
IRISA – Rennes, FR
- Ramaswamy Ramanujam
Institute of Mathematical Sciences – Chennai, IN
- Torsten Schaub
Universität Potsdam, DE
- Richard Scherl
Monmouth Univ. – West Long Branch, US
- Francois Schwarzentruber
IRISA – ENS Rennes, FR
- Maayan Shvo
Utrecht University, NL
- Sunil Easaw Simon
Indian Institute of Technology Kanpur, IN
- Michael Thielscher
UNSW – Sydney, AU
- Hans Van Ditmarsch
LORIA – Nancy, FR
- Jan van Eijck
CWI – Amsterdam, NL
- Ivan José Varzinczak
CRIL, University Artois & CNRS, FR
- Yanjing Wang
Peking University – Beijing, CN
- Bruno Zanuttini
Caen University, FR

Report from Dagstuhl Seminar 17232

Computational Interactivity

Edited by

Xiaojun Bi[1], Otmar Hilliges[2], Takeo Igarashi[3], and Antti Oulasvirta[4]

1 **Stony Brook University , US,** xiaojun@cs.stonybrook.edu
2 **ETH Zürich, CH,** otmar.hilliges@inf.ethz.ch
3 **University of Tokyo, JP,** takeo@acm.org
4 **Aalto University, FI,** antti.oulasvirta@aalto.fi

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar on Computational Interaction organized in June 2017. The seminar focused on the use of computational methods to represent and enhance human-computer interaction. This topic is gaining traction, but efforts have been diluted over multiple research areas ranging from HCI to computer graphics and design. The main objective of the seminar was to get an overview and, moreover, discuss shared fundaments, such as what computational interaction is, formally and in practice. The seminar invitees were 22 researchers from areas such as Human-Computer Interaction, Computer Graphics, Operations Research, and more. The seminar consisted of three days of events, with emphasis on presentations, panels, and group discussions. The following summarizes the main outcomes.

Seminar June 5–8, 2017 – http://www.dagstuhl.de/17232
1998 ACM Subject Classification H5.m
Keywords and phrases crowd-computing, graphics, HCI, Machine learning, optimization, simulation
Digital Object Identifier 10.4230/DagRep.7.6.48
Edited in cooperation with Anna Maria Feit

1 Executive Summary

Xiaojun Bi
Otmar Hilliges
Takeo Igarashi
Antti Oulasvirta

The field of Human Computer Interaction (HCI) as a whole has been tremendously successful in the past both in terms of growth and impact of the premier academic conferences and in reshaping the IT industry. However, as we enter the post-PC era new technologies emerge and bring new challenges along that the traditional user-centered-design approach is not well equipped to meet. For example, artificial intelligence, wearable computing, augmented and virtual reality and custom interactive devices enabled by emerging digital fabrication technologies pose increasingly wicked challenges for interaction design, where designers must consider the entire stack from low-level hardware, through software all the way to the human

factors, implying that it is no longer feasible to abstract away technology and hence design spaces that explode in their complexity.

In June 2016 we assembled a group of 25 researchers to discuss aspects relating to a computational view of interactions. Through a series of talks, breakout discussions and panel discussions we established a broad consensus that HCI can and should be approached via a computational lense. We also identified several areas in which computational models are already being used to answer HCI research questions and to move the field forward. However, it became clear that the area is in its infancy and that much work is necessary to turn computational HCI into one of the mainstream approaches in the larger research community. Primarily, researchers and students need to begin thinking in computational terms (abstraction, modelling, automation) and need to learn how to incorporate such thinking into the typically more design driven thinking prevalent in current research. Furthermore, it was also discussed at length how state-of-the-art methods in numerical optimization and machine learning can advance HCI research and likewise how HCI research can identify and refine research requirements in these adjacent research communities.

In terms of concrete outcomes, many of the present researchers agreed to contribute to a forthcoming book on "computational interaction" and to write a joint overview article further refining the discussions and outcomes of the Dagstuhl seminar. In summary, a very fruitful and productive seminar led to interesting and in-depth discussions and provided starting points for much collaborative and community-driven future work. We also identified the need for further community building work including establishing of recurring workshops, symposia, and similar outlets, outreach via summer-schools, tutorials and other educational efforts as well as establishment of a sub-committee at ACM SIGCHI the premier venue for HCI research.

2 Table of Contents

3 Overview of Talks

3.1 Usor Economicus: Modelling Interaction with Economic Models

Leif Azzopardi (University of Strathclyde – Glasgow, GB)

Joint work of Guido Zuccon, Diane Kelly, Kathy Brennan, Leif Azzopardi
Main reference Leif Azzopardi: "Modelling interaction with economic models of search", in Proc. of the 37th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '14, pp. 3–12, ACM, 2014.
URL http://dx.doi.org/10.1145/2600428.2609574

In trying to understand how people interact with search systems I have developed several economic models of search and search behaviour. These models assume an Economic User (i.e. Usor Economicus), one who inevitably does that by which s/he may obtain the greatest amount of information and knowledge, with the smallest quantity of effort. In my talk, I first provided an overview of the typically Interactive Information Retrieval process. Then I introduced an economic model of search, which is derived from production theory. I showed how the model enables us to generate compelling, intuitive and crucially testable hypotheses about the search behaviour of users. They provide insights into how we can manipulate the system and the interface in order to change the behaviour of users. In a series of user experiments, I showed how well the models characterise, predict and explain observed behaviours (and where they fall down). I believe the models not only provide a concise and compact representation of search and search behaviour, but also provide a strong theoretical basis for future research into Interactive Information Retrieval. Furthermore, these economic models can be developed for all sorts of human computer interactions, and so are likely to provide many more insights into how people use systems and how we should design such systems.

3.2 Enabling Human-Data Supported Interfaces Through Computational Models of Human Behavior

Nikola Banovic (Carnegie Mellon University – Pittsburgh, US)

Joint work of Jennifer Mankoff, Anind K. Dey, Nikola Banovic
Main reference Nikola Banovic, Tofi Buzali, Fanny Chevalier, Jennifer Mankoff, Anind K. Dey: "Modeling and Understanding Human Routine Behavior", in Proc. of the 2016 CHI Conference on Human Factors in Computing Systems, pp. 248–260, ACM, 2016.
URL http://dx.doi.org/10.1145/2858036.2858557

User interfaces that learn about people's behaviors by observing them and interacting with them enable a future in which technology helps people to be productive, comfortable, healthy, and safe. However, despite current advances in Artificial Intelligence that make such interfaces possible, it remains challenging to ensure that models that power those interfaces are free of biases that may negatively impact people. Thus, to make user interfaces that have a positive impact on people requires technology that can accurately model people's behaviors. In this body of work, we focus on behaviors in the domain of human routines that people enact as sequences of actions they perform in specific situations, which we call behavior instances. We propose a probabilistic, generative model of human routine behaviors, that can describe, reason about, and act in response to people's behaviors. We ground our model in a holistic

definition of human routines to constrain the patterns it extracts from the data to those that match routine behaviors. We train the model by estimating the likelihood that people will perform certain actions in different situations in a way that matches their demonstrated preference for those actions and situations in behavior logs. We leverage this computational model of routines to create visual analytics tools to aid stakeholders, such as domain experts and end users, in exploring, making sense of, and generating new insights about human behavior stored in large behavior logs in a principled way.

3.3 Rig Animation with a Tangible and Modular Input Device

Daniele Panozzo

Joint work of Alec Jacobson, Alex (Wan-Chun) Ma, Olga Sorkine-Hornung, Oliver Glauser, Otmar Hilliges, Daniele Panozzo
Main reference Oliver Glauser, Wan-Chun Ma, Alec Jacobson, Daniele Panozzo, Otmar Hilliges, Olga Sorkine-Hornung: "Rig animation with a tangible and modular input device", in ACM Trans. Graph., Vol. 35(4), pp. 144:1–144:11, 2016.
URL http://dx.doi.org/10.1145/2897824.2925909

We propose a novel approach to digital character animation, combining the benefits of tangible input devices and sophisticated rig animation algorithms. A symbiotic software and hardware approach facilitates the animation process for novice and expert users alike. We overcome limitations inherent to all previous tangible devices by allowing users to directly control complex rigs using only a small set (5-10) of physical controls. This avoids oversimplification of the pose space and excessively bulky device configurations. Our algorithm derives a small device configuration from complex character rigs, often containing hundreds of degrees of freedom, and a set of sparse sample poses. Importantly, only the most influential degrees of freedom are controlled directly, yet detailed motion is preserved based on a pose interpolation technique. We designed a modular collection of joints and splitters, which can be assembled to represent a wide variety of skeletons. Each joint piece combines a universal joint and two twisting elements, allowing to accurately sense its configuration. The mechanical design provides a smooth inverse kinematics-like user experience and is not prone to gimbal locking. We integrate our method with the professional 3D software Autodesk Maya® and discuss a variety of results created with characters available online. Comparative user experiments show significant improvements over the closest state-of-the-art in terms of accuracy and time in a keyframe posing task.

3.4 Optimization of Text Input

Anna Maria Feit (Aalto University, FI)

Joint work of Srinath Sridhar, Antti Oulasvirta, Andreas Karrenbauer, Daryl Weir, Mathieu Nancel, Christian Theobalt, Anna Maria Feit
Main reference Anna Maria Feit: "Computational Design of Input Methods", in Proc. of the 2017 CHI Conference on Human Factors in Computing Systems, Extended Abstracts., pp. 274–279, ACM, 2017.
URL http://dx.doi.org/10.1145/3027063.3027134

Designing a user interface or input method requires evaluating and trading-off many criteria. The corresponding design spaces are huge, making it impossible to manually build and test

every potential design. For example, if we want to design a method to enter letters in mid-air via finger gestures, there are $10^3 3$ possibilities to assign 27 characters to 32 hand gestures. Therefore, my work focuses on using optimization methods for the design of (text input) systems. The use of optimization methods allows us to efficiently and rigorously search very large design spaces, it gives quantitative guarantees on the goodness of a design and helps us to explicitly formulate and trade-off different criteria and constraints. Using optimization methods to design an input method or user interface requires 3 steps.

The first step is how to formulate the design problem. This requires to explicitly state the decisions and constraints in a mathematical way and helps to understand the characteristics of the problem. Second, we need to formulate an objective function that can be used to evaluate and compare different designs. The input data and models we use for evaluation determines the outcome of the optimization. Third, in order to solve the optimization problem, we need computational search methods. This can be mathematical solvers which guarantee to find the optimum, or it can be approximation algorithms.

I present several projects in which the focus is to develop more plausible, empirically valid formulations and objectives for more realistic optimization approaches. Among others, we modeled the performance and anatomical constraints of the hand to computationally optimize multi-finger gestures for mid-air input [1] and studied how people type on physical keyboards, in order to understand the performance of two-hand typing [2]. Most recently, we used Integer Programming to optimize the special character layout of the French keyboard to facilitate typing of correct French. Therefore, we reformulated the commonly used letter assignment problem to make it applicable to large real-world cases with over 120 to-be-mapped characters and quantified the performance, ergonomics, and ease-of-use of a keyboard layout. While my work focuses on text entry, the same problem formulations and optimization methods can be applied to the design of many other input methods and user interfaces.

References

1 Sridhar, S., Feit, A. M., Theobalt, C. and Oulasvirta, A. *Investigating the dexterity of multi-finger input for mid-air text entry*. In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems, ACM, 2015.

2 Feit, A. M., Weir, D. and Oulasvirta, A. *How we type: Movement strategies and performance in everyday typing*. In Proceedings of the 34th Annual ACM Conference on Human Factors in Computing Systems, ACM, 2016.

3.5 From n=15 to n=15,000 – Recruiting Unpaid Volunteers For Computational Design Research

Krzysztof Gajos (Harvard University, US)

Joint work of Katharina Reinecke, Krzysztof Z. Gajos
Main reference Katharina Reinecke, Krzysztof Z. Gajos: "LabintheWild: Conducting Large-Scale Online Experiments With Uncompensated Samples", in Proc. of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing, CSCW 2015, pp. 1364–1378, ACM, 2015.
URL http://dx.doi.org/10.1145/2675133.2675246

Many Computational Design projects rely on large quantities of user data. Currently, most researchers obtain such data from public data sets, collaborations with external organizations, recruitment of paid online participants or simulation. We offer an alternative: directly recruiting unpaid, intrinsically-motivated online volunteers. For over three years now, we

have run LabintheWild.org, an online platform for conducting behavioral studies with unpaid online volunteers. Our participants are motivated by the promise of finding out something about themselves and the ability to compare their performance or preferences with others. The results of our validation studies demonstrate that results of studies performed with such curiosity-motivated participants match those obtained in conventional laboratory studies as long as a few best practices are followed [1]. Depending on topic and design, individual LabintheWild studies are completed by between one hundred and several thousand participants weekly. LabintheWild was developed without any external funding and without leveraging the brand of our university (Harvard). In other words, we believe this methodology is accessible to all researchers regardless of budget or affiliation. For researchers interested in launching their own studies on LabintheWild, we offer guidelines, a tutorial and code templates at http://labinthewild.org/researchers.php.

References

1 Reinecke, K., and Gajos, K. Z. LabintheWild: conducting large-scale online experiments with uncompensated samples. In *Proceedings of the 18th ACM Conference on Computer Supported Cooperative Work & Social Computing*, CSCW '15, ACM (New York, NY, USA, 2015), 1364–1378.

3.6 Human-Math Interaction

Alec Jacobson (University of Toronto, CA)

Human-computer interaction (HCI) both studies how people interact with computers and designs new ways for people and computers to interact. In this talk, I propose human-*math* interaction (HMI), a new field of research that would study how people interact with math and design new ways to interact with math. I purposefully treat computers and math as distinct — if only as a thought experiment. HCI has been so successful for computing, I ask whether we can return the favor to applied math. As a motivating example, I first look back on how different notations for calculus provide better or worse affordances for theorem proofs or mathematical derivations. Then I turn to a ripe opportunity for HMI: the interaction with partial differential equations (PDEs) common to computer graphics and engineering. Traditional interactions are via raw boundary conditions. I show an example in the context of computer animation of how analysis of the PDE can lead to homogenization of boundary conditions relieving the burden of their explicit specification from the human "user". In another prototypical scenario of smoothing data over a bounded domain, I show how a small change can lead to dramatically simple natural boundary conditions, allowing the user to focus on the primary smoothing parameters without also managing boundary behavior. I conclude with thoughts on future directions for HMI research.

3.7 Functionality Selection in Interaction Design via Discrete Optimization

Andreas Karrenbauer (MPI für Informatik – Saarbrücken, DE)

Joint work of Antti Oulasvirta, Anna Feit, Perttu Lähteenlahti, Andreas Karrenbauer
Main reference Antti Oulasvirta, Anna Maria Feit, Perttu Lähteenlahti, Andreas Karrenbauer: "Computational Support for Functionality Selection in Interaction Design", in ACM Trans. Comput.-Hum. Interact., Vol. 24(5), pp. 34:1–34:30, 2017.
URL http://dx.doi.org/10.1145/3131608

Designing interactive technology entails several objectives, one of which is identifying and selecting appropriate functionality. Given candidate functionalities such as "Print", "Bookmark", and "Share", a designer has to choose which functionalities to include and which to leave out. Such choices critically affect the acceptability, productivity, usability, and experience of the design. However, designers may overlook reasonable designs, because there is an exponential number of functionality sets and multiple factors to consider. To tackle this problem, we use discrete optimization techniques and propose algorithmic methods to support designers to explore alternative functionality sets in early stage design. Based on interviews with professional designers, we mathematically define the task of identifying functionality sets that strike the best balance among four objectives: usefulness, satisfaction, ease of use, and profitability. We develop an integer linear programming solution that can quickly solve very large instances (set size over 1,300) on a regular computer. Further, we build on techniques of robust optimization to search for diverse and surprising functionality designs. Empirical results from a controlled study and field deployment are encouraging. Most designers rated computationally created sets to be of the comparable or superior quality than their own. Designers reported gaining a better understanding of available functionalities and the design space.

3.8 The Future of Programming and Data Sciences

Jun Kato (AIST – Ibaraki, JP)

Joint work of Jun Kato, Takeo Igarashi, Masataka Goto
Main reference Jun Kato, Takeo Igarashi, Masataka Goto: "Programming with Examples to Develop Data-Intensive User Interfaces", in IEEE Computer, Vol. 49(7), pp. 34–42, 2016.
URL http://dx.doi.org/10.1109/MC.2016.217

In this talk, I discuss three gaps between development and runtime environments of the programs that prevent a fluid programming experience. The first gap is the lack of intuitive representation of complex big data in the integrated development environment (IDE) and can be addressed by adding IDE support for the "programming with examples" workflow [1], such as integrated graphical representations in IDEs. The second gap is the so-called gulf of execution, which forces the programmer to imagine the runtime behavior of the program from the abstract source code. This gulf can be addressed by "live programming" that makes programs editable while they virtually keep running. The third gap is the high threshold for the program users to modify it to match their needs. To fill the gap, the live programming technique can be expanded to provide the end-users partial but intuitive user interfaces to tune the program behavior during its runtime [2].

These research on programming experience has deepened the understanding of, and to some extent, deconstructed the definition of "programming". I foresee that the current programmers' role will be shared among more people with diverse technical backgrounds. Those who make most of such programming deconstruction include creators since the process of creating media content becomes more and more "computational" in the future, as the seminar title implies. In particular, I consider that the process of scientific research – the process of creating scientific knowledge – can be greatly enhanced. I believe that research in computational science, when combined with the programming experience research, will make science more reproducible, enable sharing the process rather than the results, and support reuse and mashups of scientific knowledge.

References

1 Jun Kato, Takeo Igarashi, Masataka Goto. *Programming with Examples to Develop Data-Intensive User Interfaces*. IEEE Computer 49(7), pp.34–42. 2016. http://junkato.jp/programming-with-examples

2 Jun Kato, Masataka Goto. *Live Tuning: Expanding Live Programming Benefits to Non-Programmers*. In Proceedings of the Second Workshop on Live Programming Systems. 2016. http://junkato.jp/live-tuning

3.9 Crowdsourcing Visual Design

Yuki Koyama (AIST – Ibaraki, JP)

Joint work of Issei Sato, Daisuke Sakamoto, Takeo Igarashi, Yuki Koyama
Main reference Yuki Koyama, Issei Sato, Daisuke Sakamoto, Takeo Igarashi: "Sequential line search for efficient visual design optimization by crowds", in ACM Trans. Graph., Vol. 36(4), pp. 48:1–48:11, 2017.
URL http://dx.doi.org/10.1145/3072959.3073598

Parameter tweaking is one of the most fundamental tasks in many design domains. The goal is to maximize the quality of designed objects based on some criteria (i.e., objective functions). In visual design domains, aesthetic preference (i.e., how aesthetically preferable the designed object looks) often plays the role of the objective function. However, as aesthetic preference is closely tied with human perception, it is difficult to mathematically quantify this criterion using only simple rules.

To support this preference-driven parametric design tasks, we have investigated the use of crowdsourced human computation. By asking crowd workers to perform perceptual microtasks in the form of function calls, it is possible to handle human preference in a computational manner. In this talk, I introduced the following two methods. The first method is to construct a computational preference model by analyzing data from crowds and then support interactive design exploration by a designer based on the constructed model [2]. The second method is to solve a numerical optimization problem in a crowds-in-the-loop manner so that the system can efficiently find the design that maximizes the aesthetic quality [1].

References

1 Yuki Koyama, Issei Sato, Daisuke Sakamoto, Takeo Igarashi. Sequential line search for efficient visual design optimization by crowds. *ACM Trans. Graph.*, Vol. 36(4), pp. 48:1–48:11, 2017.

2 Yuki Koyama, Daisuke Sakamoto, and Takeo Igarashi. Crowd-powered parameter analysis for visual design exploration. In *Proceedings of the 27th Annual ACM Symposium on User Interface Software and Technology*, UIST '14, pages 65–74, 2014.

3.10 Probabilistic Models for Text Input and Beyond

Per Ola Kristensson (University of Cambridge, GB)

Main reference Per Ola Kristensson: "Next-Generation Text Entry", in IEEE Computer, Vol. 48(7), pp. 84–87, 2015.
URL http://dx.doi.org/10.1109/MC.2015.185

Next-generation text entry systems will need to satisfy multiple objectives, such as a high entry rate, low error rate, minimal training requirements, etc. However, via an analysis of the design space of mobile text entry methods we reveal that successful mainstream text entry methods must carry two traits in particular: 1) they must exhibit high immediate efficacy; and 2) they must have an effective entry rate which is at least as high as existing commercial alternatives. We demonstrate that insisting that a new design of a text entry method must carry these two traits leads to a very narrow design space with few viable alternative designs. However, we further demonstrate that there exist at least seven solution principles that can increase the design space: 1) from closed to open-loop; 2) continuous novice-to-expert transition; 3) path dependency; 4) flexibility; 5) probabilistic error correction; 6) fluid regulation of uncertainty; 7) efficiency.

3.11 A BIG (Bayesian Information Gain) Approach for Human-Computer Interaction

Wanyu Liu (ENST – Paris, FR)

Joint work of Rafael Lucas D'Oliveira, Michel Beaudouin-Lafon, Olivier Rioul, Wanyu Liu
Main reference Wanyu Liu, Rafael Lucas D'Oliveira, Michel Beaudouin-Lafon, Olivier Rioul: "BIGnav: Bayesian Information Gain for Guiding Multiscale Navigation", in Proc. of the 2017 CHI Conference on Human Factors in Computing Systems, pp. 5869–5880, ACM, 2017.
URL http://dx.doi.org/10.1145/3025453.3025524

We propose a novel approach BIG (Bayesian Information Gain), which is based on Bayesian Experimental Design using the criterion of mutual information from Information Theory. It captures the uncertainty the computer has about the user's goal as well as the information carried in the user input expressing what the user has in mind in the interaction. By maximizing the notion of the expected information gain, the computer can play a more active role, instead of simply executing user commands. We illustrated this idea with an instance in multiscale navigation – BIGnav. The controlled experiment demonstrated that BIGnav is significantly faster than conventional pan and zoom and requires fewer commands for distant targets, especially in non-uniform information spaces. Though being more efficient, BIGnav incurs higher cognitive load for the users, which leads us to consider more balanced interaction and shared control by leveraging the expected information gain.

3.12 Computational Interaction from a control theory perspective

Roderick Murray-Smith (University of Glasgow, GB)

As part of the Panel: "What is Computational Interaction?" I reviewed the role of control theory and dynamic systems theory in understanding common interaction techniques. Most existing techniques can potentially be described in a canonical differential equation form. I discussed how control can be seen to be at the foundations of Human–Computer Interaction and might be essential for making progress in novel forms of interface.

Traditional views of the human–machine control loop have focussed on the controlled application of external power sources. Here, a framework was proposed for considering how the human–computer interaction loop could be augmented with computational power. I considered how this could reduce mental effort, or empower disabled users to achieve more in their everyday lives.

However, the application of computational power has ethical consequences, depending on where the information supplied by the computational unit comes from.

I also discussed the limitations of control theory for the design of human–computer control loops compared to traditional engineering approaches. In contrast to traditional control theory, you cannot just swap out blocks in your control loop. Given the human's ability to predict the behaviour of other elements of the loop, and apply feed-forward compensation, we often find that the overall closed-loop behaviour does not change in as predictable a manner as in an automatic engineering system.

3.13 Personalized Interactive Devices

Jürgen Steimle (Universität des Saarlandes, DE)

Joint work of Jürgen Steimle, Daniel Groeger, Martin Weigel, Martin Schmitz, Aditya Shekhar Nittala, Alex Olwal, Tong Lu, Gilles Bailly, Antti Oulasvirta, Carmel Majidi, Jochen Huber, Niloofar Dezfuli, Max Mühlhäuser, Elena Chong Loo

Main reference Jürgen Steimle: "Skin-The Next User Interface", in IEEE Computer, Vol. 49(4), pp. 83–87, 2016.

URL http://dx.doi.org/10.1109/MC.2016.93

Advanced interactive devices for physical interaction are deployed in the user's body or embedded in the physical environment. This demands for personalized form factors and functionalities. Printable electronics, paired with computational design and fabrication, is an enabling technology for such personalized interfaces.

I gave an overview of our recent work on printed sensor and display surfaces. I presented techniques for realizing interactive temporary tattoos, which sense touch and deformation and offer visual output. These tattoos conform to fine wrinkles and highly curved body locations and hence turn the human skin into an input and output surface. They can be personalized to individual users and individual body locations through digital design.

Moreover, I presente novel ways of integrating sensing and output in 3D printed objects. These make use of material composites that comprise rigid, flexible, temperature-sensitive, and conductive 3D-printable materials. By 3D printing specific geometric patterns, the desired functional behavior is integrated. Localized touch sensors, deformation sensors, and shape-change capabilities can be realized at the desired locations within a 3D printed object.

Together, these approaches demonstrate the potential of digital fabrication and printed electronics for realizing sensors and output in new form factors, which are compatible with highly individual geometries and custom user preferences.

3.14 A Formal Methods Perspective on Computational Interaction

Harold Thimbleby (Swansea University, GB)

Best practice in HCI is to use iterative design (see figure 1), because getting complex human systems right the first time is hard, and the best solution is to implement a prototype, evaluate it, and improve it. Then repeat until the design meets the usability criteria (and other criteria such as performance and dependability).

In practice this approach does not work, since by the time a system is ready to be evaluated it is also probably ready to be sold and used — and users will anyway start getting used to and even start to rely on the suboptimal design features; further delay in release represents an opportunity cost. It is, moreover, hard work improving a complex interactive system program, particularly if evaluation suggests 'orthogonal' improvements the developers never originally anticipated. Why refactor when the first version is already in use? (This creates technical debt.)

Formal methods can help produce a more dependable and flexible initial implementation, for instance, that satisfies important principles like predictability. However, formal methods are often seen as arcane and even at odds with the 'user experience,' and in any case, for most developers, it is easier to get a program (seemingly) working than using best software engineering practice. The idea that formal methods somehow conflicts with user experience is common but is, nevertheless, a profound misunderstanding beyond the scope of this abstract to address.

A solution is to use computational interaction. Here, a separate module of the interactive program performs a restricted form of interaction $\mathcal{C}$, for instance, machine learning, signal processing, information theory, optimisation, inference, control theory or formal methods, as the case may be. This module now embodies mathematical reasoning, and the reasoning is parameterised (for instance, with data from user evaluations). This means that iterating the user interface with respect to $\mathcal{C}$ is now easy and reliable, for $\mathcal{C}$ preserves all the properties of machine learning or whatever.

Now the iteration of the ISO iterative design cycle iterates $\mathcal{C}$; this is fast and easy, if not trivial, and by appropriate choice of one or more $\mathcal{C}_i$ the most important features of the user interface quickly converge onto what the evaluation and data collection suggests. This is a radical improvement over conventional user interface design, provided only that the right data is collected.

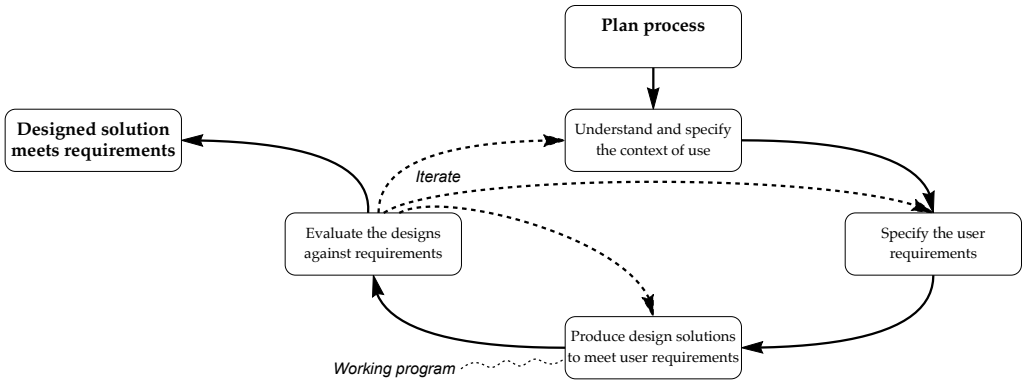

Figure 1 The international standards ISO 9241:210 iterative design cycle. Note that a working program is available before even one iteration is complete, and in conventional program development this may dissuade developers from embarking on the expensive iterative cycles. In contrast, computational interaction identifies a module that is easy to parameterise and easy to update from user evaluation data.

3.15 A Proposed Definition of Computational Interaction

John Williamson (University of Glasgow, GB)

Computational interaction is a distinctive approach to human-computer interaction, which emphasises the role of mathematical modeling of user-system behaviour and computation in augmenting the design and operation of interfaces. We propose some defining characteristics which computational interaction would typically be expected to have at least some of:

- an explicit mathematical model of user-system behavior;
- a way of updating that model with observed data from users;
- an algorithmic element that, using this model, can directly synthesise or adapt the design;
- a way of automating and instrumenting the modeling and design process;
- the ability to simulate or synthesise elements of the expected user-system behavior.

Potential advantages of computational approaches include:

- reduced design time of interfaces, as automation, data and models can supplant hand-tweaking;
- more robust and efficient interfaces, as better models can better predict how interactions evolve;
- better tailored interfaces: to users, contexts, devices, as structure can be learned rather than dictated
- new technologies can be harnessed quickly as fundamental, repeatable processes generalise to new contexts
- the evaluation burden can be reduced, as strong, executable models can predict much of expected user behavior
- our ability to reason rather than rely on experimentation increases as HCI problems can be defined formally
- algorithmic design can support designers in tedious tasks and help designers focus on creative aspects of design

4 Working groups

4.1 Modeling for Computational HCI

Leif Azzopardi (University of Strathclyde – Glasgow, GB), Nikola Banovic (Carnegie Mellon University – Pittsburgh, US), Andreas Karrenbauer (MPI für Informatik – Saarbrücken, DE), and Wanyu Liu (ENST – Paris, FR)

Computational models of human behavior, interactive systems, and the interaction between human and the system are the foundation of computational interactivity. However, despite availability of such models, user experience (UX) designers still mostly depend on intuition, supported by existing knowledge and experience, when designing user interfaces. This likely follows from lack of awareness about the existing computational models, lack of information about the benefits of such models, and the difficulty that designers face when integrating such models into their work. Our attempt to provide a brief taxonomy of such models is a first step in bringing these models closer to the UX design community. We propose the following dimensions that describe the space of computational models: 1) granularity – the level at which humans, systems, and interactions are being modeled (e.g., cognitive, activity, task); 2) realism – how closely to the real world the model attempts to represent information; 3) flexibility – the ability of the model to change in response to new information; 4) dynamicity – whether the model is static (offline) or active (online); 5) explainability – if the model is predictive or explanatory; and 6) origin – data driven or theory driven. These dimensions offer the ability to compare and contrast the existing models and provide an overview of the current state of the field. It is also a first step towards defining metrics to evaluate such computational models to ensure they provide scientific, principled foundation for computational interactivity. These dimensions will lead to the development of new and the improvement of existing computational models and offer a promise of a unified, overarching model of human behavior which will drive the future of computational interactivity.

4.2 Challenges and Problems in Computational Interaction

Alec Jacobson (University of Toronto, CA), Anna Maria Feit (Aalto University, FI), Krzysztof Gajos (Harvard University, US), and Harold Thimbleby (Swansea University, GB)

Our working group undertook the daunting task of defining "grand challenges for computational interaction". After brainstorming, the group separated proposed challenges into two categories: A) ways computational interaction can make the world a better place and B) technical challenges within computational interaction research. The working definition of "computational interaction" was based on the notion that computational interaction is interaction that improves with improved computational power. We identified two modalities on the spectrum of computational interaction. I) Interactions that are created as the result of intensive computation. For example, using machine learning to determine optimal placement for buttons or knobs. II) Interactions that involve intensive computation at run time. For example, a gesture based controller that learns to better recognize a particular user's motions after continued use. With these possibilities in mind, we enumerate promising and potentially impactful directions for the field.

Ways computational interaction can make the world a better place:

1. Automatically generate personalized curricula for people given their knowledge goals, current knowledge, and cognitive abilities
2. Automatically plan meeting activities given goals, time budget and personalities of participants
3. Interactively model solution spaces for designers or design teams and identify under-explored regions of this space by breaking implicit, human assumptions
4. Invent computational tools for facilitating arguments and mediating conflicts
5. Aid understanding of deep neural networks by designing human-understandable networks or providing tools to determine whether a network is understandable
6. Invent information extraction tools for life-long datasets such as photos, GPS data or files
7. Improve user experiences in mixed initiative interactions by modeling predictability (user's guess about the system) and uncertainty (system's guesses about the user)

Technical challenges for computational interaction

1. Create benchamarks for quantitative comparisons for fundamental computational interaction problems à la computer vision.
   a. Utilize labinthewild.org for open-source, large-scale testing
   b. For example, X minutes to build a model of the user, then quantify improved performance of some UI Y.
   c. Info-vis is especially ripe for such testing. Fix dataset and measure time and how well a visualization conveys information.
   d. Other immediately testable tasks are text input, GUIs, and pointing.
   e. Unlike fully automatic benchmarks in computer vision, extra care will be needed to prevent bias, noise, and foul play.
2. Foster a community of open-source, shared, reusable toolboxes. Upstart time for experiments could be dramatically reduced by collectively creating plugins for common methods and models in the field. Reproducibility should increase as well.

Solving these challenges would mark great strides for the emerging field of computational interaction.

4.3 What is Computational HCI? Definition, Types, and Scope

Per Ola Kristensson (University of Cambridge, GB), Andrew Howes (University of Birmingham, GB), Jun Kato (AIST – Ibaraki, JP), and John Williamson (University of Glasgow, GB)

In this working group, we discussed a definition of Computational HCI. The outcome of the discussion was the following definition:

Computational HCI applies computational thinking (abstraction, automation, analysis) to explain and enhance the interaction between a user (or users) and a system. It is underpinned by modelling which admits formal reasoning and involves at least one of:

- a way of updating a model with observed data from users;
- an algorithmic element that, using a model, can directly synthesise or adapt the design;
- a way of automating and instrumenting the modeling and design process;
- the ability to simulate or synthesise elements of the expected user-system behavior.

Computational HCI is founded on fundamental theoretical considerations but is constructive in its approach rather than descriptive. It is empirical and quantitative but focuses on how to use computationally powered models to do what could not be done before rather than describing what has been done before. It emphasises generating motor themes in HCI, and robust, replicable and durable approaches which go beyond point sampling of the interaction space.

4.4 Machine Learning for and in HCI

Roderick Murray-Smith (University of Glasgow, GB), Emre Aksan (ETH Zürich, CH), Otmar Hilliges (ETH Zürich, CH), Takeo Igarashi (University of Tokyo, JP), and Jerry Talton (Slack Technologies – New York, US)

This working group discussed the relevance of machine learning (ML) with human computer interaction (HCI) and raised questions in order to better address problems of HCI research. Our discussion consisted of four main topics:

1. Why has HCI often not seen the 'ratchet effect' seen in other fields where people productively build on the progress of others?
   - How can we change this (motor themes)?
   - Would having an HCI Kaggle/OpenAI Gym like combo be beneficial to encourage sharing of data, models, and problems?
   - How do you generate, adapt and curate data in HCI for ML tools?

   We anticipate triple benefits for the community:
   - Defining problems and providing datasets generates citations.
   - Beating established baselines generates progress (and papers).
   - Solving hard problems requires the development of new methods (and creation of insights) that may become useful outside of academic research.
2. What is different about the nature of HCI compared to typical problems in ML?
   - Data is more expensive (sometimes).
   - Users (and hence data) adapt over time which requires 'living' datasets (see information retrieval approaches).
   - Conventional metrics in ML are not representative enough. For example, optimizing (e.g.) recognition accuracy is a poor proxy for usability.
   - Users learn and adapt, which means that cost functions and data changes.
   - Variability across users (behavior, preferences) is high (high variance problems typically require a lot of data).
3. Should we approach ML as a tool or invest resources in ML research for HCI problems?
   - Should HCI researchers be more involved with ML and seek solutions for HCI related issues?
   - How can you calculate usability gradients through the user? Can we optimise usability as a cost function?
   - How do you optimise usability when the human is adapting continuously to the new system?
   - How do we enable and ensure a smooth co-evolution of the user and system to a state of improved interaction?

- How do you design for the individual (en masse, via adapting to each end user)?
- Many recent ML approaches learn from large datasets and then become rigid. Adaptive algorithms such as one-shot learning should be favored.

4. What can HCI contribute to ML?
   - Can we enable nonML-trained end users to productively use machine learning?
   - Can we 'un-black-box' ML models for naive users?
   - ML is making certain interaction scenarios worth investigating again from an HCI perspective. (e.g. sudden use in speech-based interaction in home settings with Alexa)
   - Has HCI then lost 'control'? (e.g. much modern commercial VR/AR activity ignored the older AR/VR research.)

4.5 Computational Design

Antti Oulasvirta (Aalto University, FI), Yuki Koyama (AIST – Ibaraki, JP), Ranjitha Kumar (University of Illinois – Urbana-Champaign, US), Daniele Panozzo (New York University, US), and Jürgen Steimle (Universität des Saarlandes, DE)

This workgroup looked at the intersection of computational design and computational interaction. This intersection is highly relevant for HCI research because although interactive computational design methods are used across several fields, there seems to be no principled understanding of how to design efficient user interfaces for these tools. Computational design is researched, among others, in graphics, visualization, computer-aided design, engineering design, and architecture, and there are multiple fields that come close, such as recommendation systems and computational creativity. To begin with, the group defined computational design as (1) influencing the design of an artefact, (2) supporting any part of the design process from user research to generation and evaluation (although it noted that most work focuses on generation), (3) represents the problem mathematically as the problem of finding the best design among some options, and (4) solves it using algorithmic solution methods that broadly fall within the scope of mathematical optimization.

There are several practical and scientific benefits to this approach, which were observed to stem from two sources: (1) new ways of representing design problems (a new way to think and imagine design); (2) increased problem-solving capacity (e.g., solving problems in design that were previously unsolvable). The group made a somewhat startling conclusion: at the core of it computational design is applied math, and in particular mathematical optimization. Interactive design tools in computational design are essentially about interfacing mathematics. The group concluded that this poses a grand challenge to computational HCI: what is the best way to interface mathematics in a given design project? Can something foundational be said about the success or performance of a designer using a particular interaction style in a particular project, such as about success probability, time cost, or mental load? The challenge for HCI research is to identify ways to interface objective functions with their constraints and parameters. This invites research on how to obtain, edit, analyze, visualize, validate, and interpret objective functions as part of a design process.

The group suggested that while an essential part of this multi-disciplinary field has looked at the generation or refinement of designs (construction), there is a significant opportunity space in other parts of the design process, such as user research, sketching, prototyping,

deployment, and evaluation. Another significant contribution HCI research can make in this space is systematic reporting of the needs, requirements, and practices of professional and non-professional designers. The group concluded with five objectives for future work: (1) Formulate the research problems of (interactive) computational design; (2) Define success metrics of interactive computational design; (3) Review methods explored and progress achieved this far; (4) Articulate the changing roles, requirements, and practices of design; and (5) To position computational design in user-centered design, its position to current understanding of design thinking must be articulated.

Participants

- Emre Aksan
ETH Zürich, CH
- Leif Azzopardi
University of Strathclyde – Glasgow, GB
- Nikola Banovic
Carnegie Mellon University – Pittsburgh, US
- Anna Maria Feit
Aalto University, FI
- Krzysztof Gajos
Harvard University, US
- Otmar Hilliges
ETH Zürich, CH
- Andrew Howes
University of Birmingham, GB
- Takeo Igarashi
University of Tokyo, JP
- Alec Jacobson
University of Toronto, CA
- Andreas Karrenbauer
MPI für Informatik – Saarbrücken, DE
- Jun Kato
AIST – Ibaraki, JP
- Yuki Koyama
AIST – Ibaraki, JP
- Per Ola Kristensson
University of Cambridge, GB
- Ranjitha Kumar
University of Illinois – Urbana-Champaign, US
- Wanyu Liu
ENST – Paris, FR
- Roderick Murray-Smith
University of Glasgow, GB
- Antti Oulasvirta
Aalto University, FI
- Daniele Panozzo
New York University, US
- Jürgen Steimle
Universität des Saarlandes, DE
- Jerry Talton
Slack Technologies – New York, US
- Harold Thimbleby
Swansea University, GB
- John Williamson
University of Glasgow, GB

Report from Dagstuhl Seminar 17251

Game Theory Meets Computational Learning Theory

Edited by

Paul W. Goldberg[1], Yishay Mansour[2], and Paul Dütting[3]

1 **University of Oxford, GB,** paul.goldberg@cs.ox.ac.uk
2 **Tel Aviv University, IL,** mansour@tau.ac.il
3 **London School of Economics, GB,** p.d.duetting@lse.ac.uk

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 17251 "Game Theory Meets Computational Learning Theory". While there have been many Dagstuhl seminars on various aspects of Algorithmic Game Theory, this was the first one to focus on the emerging field of its intersection with computational learning theory.

Seminar June 18–23, 2017 – http://www.dagstuhl.de/17251
1998 ACM Subject Classification I.2 Artificial Intelligence, I.2.6 Learning (K.3.2), F. Theory of Computation, J.4 Social and Behavioral Sciences (economics)
Keywords and phrases Algorithmic Game Theory, Computational Learning Theory, Economics
Digital Object Identifier 10.4230/DagRep.7.6.68

1 Executive Summary

Paul W. Goldberg
Yishay Mansour
Paul Dütting

Algorithmic Game Theory (AGT) has been an identifiable research field for about 20 years by now. It emerged as an important research community in the 1990s, with the ACM-EC conference starting in 1999, and the conferences WINE and SAGT also support this community; in addition, the field is also represented in the main CS theory and AI conferences. Among former Dagstuhl seminars on topics in AGT, there have been a sequence of Dagstuhl seminars on Equilibrium Computation, and another sequence of seminars on Computational Social Choice, and also on Electronic Markets and Auctions.

Machine learning has of course become very pervasive, with the vast accessibility of "big data" and has the motivation to develop new methodologies for harvesting the vast amounts of data, improve our ability to automatically carry out many tasks, from classifying documents and pictures, to identifying normal trends and anomalies.

It is perhaps not surprising that it is timely to investigate the connections between economics and big data, more specifically the interface between game theory and machine learning. Much of econometrics is about handing data and deriving understanding from data.

In the AGT context, this would seem to apply most readily to data emanating from "economic" sources, and of course there are plenty of examples. The most notable of these is learning user preferences from examples.

There have been workshops at the AGT/Machine Learning interface in the ACM-EC conference (the 2017 EC held the 3rd workshop on Algorithmic Game Theory and Data

Science) but there is clearly space for more meetings on this topic, and this is the first one to take place at Dagstuhl. As such, it contributed to the development of the European community in this research area (noting that ACM-EC is usually held in the USA).

It was pleasing that the seminar attracted quite a high proportion of participants who were visiting Dagstuhl for this first time, alongside others who have made multiple visits. There was a good balance amongst representatives of Algorithmic Game Theory, Machine Learning, and Economics.

One can classify the AGT/Machine Learning topics as follows.

- Usage of ML ideas (reinforcement learning, multi-arm bandits, etc.) into decision making under uncertainty (and the search for game-theoretic solution notions such as equilibria)
- usage of game-theoretic tools into machine learning approaches (as in Generative Adversarial Nets).
- A basic test case is learning user valuations from historical data. For example, given the outcome of previous auctions to learn the distribution of the users' valuations and the goal is to define near optimal mechanisms. (This is also an aspect in learning "revealed preferences".)
- Query complexity of solution concepts of games, aspects of which are applicable to learning adversary preferences in the context of security/patrolling games.

The seminar was structured around longer invited/tutorial talks (typically lasting 1.5-2 hours), one or two such talks taking place each day. These were followed by shorter contributed talks.

We thank Argy Deligkas for serving as collector of the abstracts.

Keynote/tutorial talks

Monday Sven Seuken, University of Zurich Design of Machine Learning-Based Mechanisms; Yaron Singer, Harvard University: Learning, Optimization, and Noise

Tuesday Claudio Gentile, Universita dell'Insubria: No Regret and Sequential Prediction

Wednesday Denis Nekipelov, University of Virginia: Robust Inference for Non-Robust Models

Thursday Jamie Morgenstern, University of Pennsylvania: The Sample Complexity of Single-Parameter Auction Design

Friday Yakov Babichenko, Technion: Informational Bounds on Equilibria, and its Relation to Learning

Related topics not covered

There is ongoing work at the intersection of macroeconomics and machine-learning techniques, that was out of scope of this meeting but may be of interest later. For example, ongoing work on Vector Autoregressive models in the context of multivariate time series modelling, which may later lead to interesting problem in computational learning theory. Another topic that was only touched-on is agent-based models of macroeconomics.

2 Table of Contents

3 Overview of Talks

3.1 Forecast Aggregation

Yakov Babichenko (Technion – Haifa, IL)

Joint work of Itai Arieli, Yakov Babichenko, Smorodinsky Rann

Bayesian experts with a common prior that are exposed to different types of evidence possibly make contradicting probabilistic forecasts. A policy maker who receives the forecasts must aggregate them in the best way possible. This is a challenge whenever the policy maker is not familiar with the prior nor the model and evidence available to the experts. We propose a model of non-Bayesian forecast aggregation and adapt the notion of regret as a means for evaluating the policy maker's performance. Whenever experts are Blackwell ordered taking a weighted average of the two forecasts, the weight of which is proportional to its precision (the reciprocal of the variance), is optimal. The resulting minimal regret is equal to $(5^{1.5} - 11)/8 \simeq 0.0225425$, which is 3 to 4 times better than naive approaches such as choosing one expert at random or taking the non-weighted average.

3.2 Aggregating Earnings-Per-Share Predictions

Amir Ban (Tel Aviv University, IL)

Joint work of Amir Ban, Yishay Mansour

We investigate whether stock-market analysts possess differential expertise, and if so, whether such an observation makes it possible to aggregate multiple analyst estimates into a result that is significantly more accurate than their consensus average. We do a retrospective study using historical quarterly earnings-per-share forecasts and actual results for large publicly traded companies, setting a goal of getting the most accurate aggregated forecast possible. Using this goal as criterion, we find that analysts possess both an individual forecast bias as well as individual expertise. The individual bias is significant, while the individual expertise is of lower significance. Together, they enable a $20\% - 30\%$ accuracy improvement over consensus average.

3.3 The optimal strategy for a linear regression game

Peter L. Bartlett (University of California – Berkeley, US)

Joint work of Peter Bartlett, Alan Malek

We consider a linear regression game: at each round, an adversary reveals a covariate vector, the learner predicts a real value, the adversary reveals a label, and the learner incurs the squared prediction error. The aim is to minimize the difference between the cumulative loss and that of the linear predictor that is best in hindsight. We present the minimax optimal

strategy for this game, and show that it can be efficiently computed, for natural constraints on the adversarially chosen covariate sequence that prevent the adversary from misrepresenting the scale of the problem. This strategy is horizon-independent, that is, it incurs no more regret than the optimal strategy that knows in advance the number of rounds of the game. We also provide an interpretation of the minimax algorithm as a follow-the-regularized-leader strategy with a data-dependent regularizer.

3.4 Bayesian Methods for Market Clearing

Gianluca Brero (Universität Zürich, CH) and Sébastien Lahaie (Google – New York, US)

We cast the problem of combinatorial auction design in a Bayesian framework in order to incorporate prior information into the auction process and minimize the number of rounds. We develop a generative model of agent valuations and market prices such that clearing prices become maximum a posteriori estimates given observed agent valuations. This generative model then forms the basis of an auction process which alternates between refining estimates of agent valuations and computing candidate clearing prices. We provide an implementation of the auction using assumed density filtering to estimate valuations and expectation maximization to compute prices. An empirical evaluation over a range of valuation domains demonstrates that our Bayesian auction mechanism is very competitive against a conventional combinatorial clock auction, even under the most favorable choices of price increment for this baseline.

3.5 Distributed Methods for Computing Approximate Equilibria

Argyris Deligkas (Technion – Haifa, IL), John Fearnley (University of Liverpool, GB), and Rahul Savani (University of Liverpool, GB)

Joint work of Artur Czumaj, Argyrios Deligkas, Michail Fasoulakis, John Fearnley, Marcin Jurdzinski, Rahul Savani
Main reference Artur Czumaj, Argyrios Deligkas, Michail Fasoulakis, John Fearnley, Marcin Jurdzinski, Rahul Savani: "Distributed Methods for Computing Approximate Equilibria", in Proc. of the Web and Internet Economics - 12th International Conference, WINE 2016, Montreal, Canada, December 11–14, 2016, Proceedings, LNCS, Vol. 10123, pp. 15–28, Springer, 2016.
URL http://dx.doi.org/10.1007/978-3-662-54110-4_2

We present a new, distributed method to compute approximate Nash equilibria in bimatrix games. In contrast to previous approaches that analyze the two payoff matrices at the same time (for example, by solving a single LP that combines the two players' payoffs), our algorithm first solves two independent LPs, each of which is derived from one of the two payoff matrices, and then computes an approximate Nash equilibrium using only limited communication between the players. Our method gives improved bounds on the complexity of computing approximate Nash equilibria in a number of different settings. Firstly, it gives a polynomial-time algorithm for computing *approximate well supported Nash equilibria (WSNE)* that always finds a 0.6528-WSNE, beating the previous best guarantee of 0.6608. Secondly, since our algorithm solves the two LPs separately, it can be applied to give an improved

bound in the limited communication setting, giving a randomized expected-polynomial-time algorithm that uses poly-logarithmic communication and finds a 0.6528-WSNE, which beats the previous best known guarantee of 0.732. It can also be applied to the case of *approximate Nash equilibria*, where we obtain a randomized expected-polynomial-time algorithm that uses poly-logarithmic communication and always finds a 0.382-approximate Nash equilibrium, which improves the previous best guarantee of 0.438. Finally, the method can also be applied in the query complexity setting to give an algorithm that makes $O(n \log n)$ payoff queries and always finds a 0.6528-WSNE, which improves the previous best known guarantee of 2/3.

3.6 Optimal Auctions through Deep Learning

Paul Dütting (London School of Economics, GB), Zhe Feng, Harikrishna Narasimhan, and David C. Parkes

Main reference Paul Dütting, Zhe Feng, Harikrishna Narasimhan, David C. Parkes: "Optimal Auctions through Deep Learning", CoRR, Vol. abs/1706.03459, 2017.
URL http://arxiv.org/abs/1706.03459

Designing an auction that maximizes expected revenue is an intricate task. Indeed, as of today–despite major efforts and impressive progress over the past few years–only the single-item case is fully understood. In this work, we initiate the exploration of the use of tools from deep learning on this topic. The design objective is revenue optimal, dominant-strategy incentive compatible auctions. We show that multi-layer neural networks can learn almost-optimal auctions for settings for which there are analytical solutions, such as Myerson's auction for a single item, Manelli and Vincent's mechanism for a single bidder with additive preferences over two items, or Yao's auction for two additive bidders with binary support distributions and multiple items, even if no prior knowledge about the form of optimal auctions is encoded in the network and the only feedback during training is revenue and regret. We further show how characterization results, even rather implicit ones such as Rochet's characterization through induced utilities and their gradients, can be leveraged to obtain more precise fits to the optimal design. We conclude by demonstrating the potential of deep learning for deriving optimal auctions with high revenue for poorly understood problems.

3.7 Optimal Auctions for Correlated Bidders with Sampling

Hu Fu (University of British Columbia – Vancouver, CA), Jason D. Hartline (Northwestern University – Evanston, US), Nima Haghpanah, and Robert Kleinberg

Main reference Hu Fu, Nima Haghpanah, Jason D. Hartline, Robert Kleinberg: "Optimal Auctions for Correlated Buyers with Sampling", CoRR, Vol. abs/1406.1571, 2014.
URL http://arxiv.org/abs/1406.1571

Correlation among buyers' valuations enables a revenue maximizing seller to fully extract the social surplus, with no money left to the buyers. This was shown by a classical result by Cremer and McLean. The model has been criticized for allowing arbitrary dependence of the mechanism on the prior: any uncertainty on the prior disrupts the mechanism. We examine

this criticism from a learning point of view. We allow uncertainty on the prior but grant the seller sample access from the true prior, and study the number of samples that suffice for surplus extraction. We give precise bounds on the number of samples needed, which show that surplus extraction needs much less information than learning the prior itself. In a sense, this is because the buyers "collaborate" in the learning, driven by their incentives. Our upper bound on the number of samples is by an algebraic argument.

3.8 No Regret and Sequential Prediction

Claudio Gentile (University of Insubria – Varese, IT)

A 90 minute tutorial on classical and recent research in the context of sequential prediction algorithms (expert, bandit, and variants thereof), as well as in online convex optimization and online learning of Lipschitz policies.

3.9 Non-revelation Mechanism Design

Jason D. Hartline (Northwestern University – Evanston, US) and Samuel Taggart

Main reference Jason D. Hartline, Samuel Taggart: "Non-Revelation Mechanism Design", CoRR, Vol. abs/1608.01875, 2016.
URL http://arxiv.org/abs/1608.01875

We consider mechanism design and redesign for markets like Internet advertising where many frequent, small transactions are organized by a principal. Mechanisms for these markets rarely have truthtelling equilibria. We identify a family of winner-pays-bid mechanisms for such markets that exhibit three properties. First, equilibria in these mechanisms are simple. Second, the mechanisms' parameters are easily reoptimized from the bid data that the mechanism generates. Third, the performance of mechanisms in the family is near the optimal performance possible by any mechanism (not necessarily within the family). Our mechanisms are based on batching across multiple iterations of an auction environment, and our approximation bound is asymptotically optimal, with loss inversely proportional to the cube root of the number of iterations batched. Our analysis methods are of broader interest in mechanism design and, for example, we also use them to give new sample complexity bounds for mechanism design in general single-dimensional agent environments.

3.10 Optimizing Worst-case Benchmarks

Aleck Johnsen (Northwestern University – Evanston, US) and Jason D. Hartline (Northwestern University – Evanston, US)

Joint work of Sam Taggart, Aleck Johnsen, Jason D. Hartline

Three common information settings under which algorithms with incomplete information are studied are Bayesian- known distribution, Bayesian- unknown distribution, and worst-case. Whereas Bayesian settings automatically embed optimization problems, worst-case requires benchmarks as a freely chosen parameter, with algorithms subsequently designed to compete pointwise with the benchmark. This talk studies a property of worst-case benchmarks (normalization) that would necessarily allow algorithms approximating the benchmark to extend to give guarantees in the setting of an unknown Bayesian distribution; and another property (resolution) to allow comparison of the efficacy of "normalized" benchmarks, turning benchmark design into an optimization question. Two disparate algorithmic settings of incomplete information are used for the analysis, (online) Expert Learning, and (private values) Auctions.

3.11 Best-Response Dynamics in Combinatorial Auctions with Item Bidding

Thomas Kesselheim (TU Dortmund, DE) and Paul Dütting (London School of Economics, GB)

Main reference Paul Dütting, Thomas Kesselheim: "Best-Response Dynamics in Combinatorial Auctions with Item Bidding", in Proc. of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2017, Barcelona, Spain, Hotel Porta Fira, January 16–19, pp. 521–533, SIAM, 2017.
URL http://dx.doi.org/10.1137/1.9781611974782.33

In a combinatorial auction with item bidding, agents participate in multiple single-item second-price auctions at once. As some items might be substitutes, agents need to strategize in order to maximize their utilities. A number of results indicate that high welfare can be achieved this way, giving bounds on the welfare at equilibrium. Recently, however, criticism has been raised that equilibria are hard to compute and therefore unlikely to be attained.

In this paper, we take a different perspective. We study simple best-response dynamics. That is, agents are activated one after the other and each activated agent updates his strategy myopically to a best response against the other agents' current strategies. Often these dynamics may take exponentially long before they converge or they may not converge at all. However, as we show, convergence is not even necessary for good welfare guarantees. Given that agents' bid updates are aggressive enough but not too aggressive, the game will remain in states of good welfare after each agent has updated his bid at least once.

In more detail, we show that if agents have fractionally subadditive valuations, natural dynamics reach and remain in a state that provides a 1/3 approximation to the optimal welfare after each agent has updated his bid at least once. For subadditive valuations, we can guarantee an $\Omega(1/\log m)$ approximation in case of m items that applies after each agent has updated his bid at least once and at any point after that. The latter bound is complemented

by a negative result, showing that no kind of best-response dynamics can guarantee more than an $o(\log\log m/\log m)$ fraction of the optimal social welfare.

3.12 Bayesian Methods for Clearing Markets

Sébastien Lahaie (Google – New York, US) and Gianluca Brero (Universität Zürich, CH)

We cast the problem of combinatorial auction design in a Bayesian framework in order to incorporate prior information into the auction process and minimize the number of rounds. We develop a generative model of agent valuations and market prices such that clearing prices become maximum a posteriori estimates given observed agent valuations. This generative model then forms the basis of an auction process which alternates between refining estimates of agent valuations and computing candidate clearing prices. We provide an implementation of the auction using assumed density filtering to estimate valuations and expectation maximization to compute prices. An empirical evaluation over a range of valuation domains demonstrates that our Bayesian auction mechanism is very competitive against a conventional combinatorial clock auction, even under the most favorable choices of price increment for this baseline.

3.13 Submultiplicative Glivenko-Cantelli and Uniform Convergence of Revenues

Yishay Mansour (Tel Aviv University, IL)

Joint work of Noga Alon, Moshe Babaioff, Yannai A. Gonczarowski, Yishay Mansour, Shay Moran, Amir Yehudayoff
Main reference Noga Alon, Moshe Babaioff, Yannai A. Gonczarowski, Yishay Mansour, Shay Moran, Amir Yehudayoff: "Submultiplicative Glivenko-Cantelli and Uniform Convergence of Revenues", CoRR, Vol. abs/1705.08430, 2017.
URL http://arxiv.org/abs/1705.08430

In this work we derive a variant of the classic Glivenko-Cantelli Theorem, which asserts uniform convergence of the empirical Cumulative Distribution Function (CDF) to the CDF of the underlying distribution. Our variant allows for tighter convergence bounds for extreme values of the CDF.

We apply our bound in the context of revenue learning, which is a well-studied problem in economics and algorithmic game theory. We derive sample-complexity bounds on the uniform convergence rate of the empirical revenues to the true revenues, assuming a bound on the k-th moment of the valuations, for any (possibly fractional) $k > 1$.

For uniform convergence in the limit, we give a complete characterization and a zero-one law: if the first moment of the valuations is finite, then uniform convergence almost surely occurs; conversely, if the first moment is infinite, then uniform convergence almost never occurs.

3.14 The Sample Complexity of Single-Parameter Auction Design

Jamie Morgenstern (University of Pennsylvania – Philadelphia, US)

Joint work of Jamie Morgenstern, Tim Roughgarden

This tutorial will overview recent literature on the sample complexity of learning (nearly) revenue-optimal auctions for selling to buyers in the single item (and more generally single parameter) setting.

3.15 Revenue Optimization with Approximate Bid Predictions

Andrés Muñoz Medina (Google – New York, US)

Joint work of Andrés Muñoz Medina, Sergei Vassilvitskii
Main reference Andrés Muñoz Medina, Sergei Vassilvitskii: "Revenue Optimization with Approximate Bid Predictions", CoRR, Vol. abs/1706.04732, 2017.
URL http://arxiv.org/abs/1706.04732

In the context of advertising auctions, finding good reserve prices is a notoriously challenging learning problem. This is due to the heterogeneity of ad opportunity types and the non-convexity of the objective function. In this work, we show how to reduce reserve price optimization to the standard setting of prediction under squared loss, a well understood problem in the learning community. We further bound the gap between the expected bid and revenue in terms of the average loss of the predictor. This is the first result that formally relates the revenue gained to the quality of a standard machine learned model.

3.16 Robust Inference for Games via Theoretical Guarantees

Denis Nekipelov (University of Virginia – Charlottesville, US)

Joint work of Darrell Hoy, Denis Nekipelov, Vasilis Syrgkanis

In the past decade the Economics literature has developed a unified approach to inference and prediction of strategic environments. This approach starts with a full theoretical model that characterizes the preferences of agents and the mechanism of interaction between them.

The Econometrician then infers the components of the theoretical model from the data and provides prediction for the new settings by computing a new equilibrium of the strategic model. The issue with this approach is that the initial step where the Econometrician recovers the preferences of the agents from the data boils down to an inversion of a nonlinear mapping that can be discontinuous and even set-valued. In the talk I will discuss the properties of this mapping for classes of simple games and demonstrate that even in cases where this mapping is invertible, the recovered pre-image is very sensitive to the specification of the model. I will also discuss that such poor behavior of the solution will be preserved even when one uses strong conditions to "regularize" this mapping. All these problems will be further amplified in the prediction.

The approach of set inference provides a robust alternative to traditional inference. In this approach the Econometrician recovers an entire set of preferences of the agents in the interactions that are compatible with many possible specifications of the theoretical model. However, computation of such sets is difficult even in simple games. In my talk I discuss a new approach to inference that is based on the idea of the price of anarchy in Koutsoupias and Papadimitriou (1999). The idea of the approach is to bypass the set inference for the primitives of the model and instead directly infer the outcomes of interest such as welfare or revenue in the game. However, unlike the standard price of anarchy which is based on the "worst case scenario"-based prediction for the outcomes, we propose to consider the bounds that are informed by the distribution of the data. I talk about the new notion of the empirical price of anarchy that yields the price of anarchy over all preferences of agents that could have generated the observable distribution of the data. I then discuss some connections between our notion of the empirical price of anarchy and Economic literature on set inference.

3.17 Weighted Voting Via No-Regret Learning

Ariel D. Procaccia (Carnegie Mellon University – Pittsburgh, US)

Joint work of Nika Haghtalab, Ritesh Noothigattu, Ariel D. Procaccia
Main reference Nika Haghtalab, Ritesh Noothigattu, Ariel D. Procaccia: "Weighted Voting Via No-Regret Learning", CoRR, Vol. abs/1703.04756, 2017.
URL http://arxiv.org/abs/1703.04756

Voting systems typically treat all voters equally. We argue that perhaps they should not: Voters who have supported good choices in the past should be given higher weight than voters who have supported bad ones. To develop a formal framework for desirable weighting schemes, we draw on no-regret learning. Specifically, given a voting rule, we wish to design a weighting scheme such that applying the voting rule, with voters weighted by the scheme, leads to choices that are almost as good as those endorsed by the best voter in hindsight. We derive possibility and impossibility results for the existence of such weighting schemes, depending on whether the voting rule and the weighting scheme are deterministic or randomized, as well as on the social choice axioms satisfied by the voting rule.

3.18 Design of Machine Learning-based Mechanisms

Sven Seuken (Universität Zürich, CH)

Joint work of Gianluca Brero, Benjamin Lubin, Sven Seuken
Main reference Gianluca Brero, Benjamin Lubin, Sven Seuken: "Probably Approximately Efficient Combinatorial Auctions via Machine Learning", in Proc. of the Thirty-First AAAI Conference on Artificial Intelligence, February 4–9, 2017, San Francisco, California, USA., pp. 397–405, AAAI Press, 2017.
URL http://aaai.org/ocs/index.php/AAAI/AAAI17/paper/view/15040

In this talk, we present a new paradigm we call "designing machine learning-based mechanisms." In contrast to most prior work, our paradigm uses machine learning (ML) directly on the agents' reports, not to optimize some future mechanism, but to immediately make use of the learning outcome (for the current instance). We instantiate this new idea via combinatorial auctions (CAs), and show how using ML inside CAs can substantially simplify the interaction

with the bidders. In our CAs, the bidders report their values (bids) to a proxy agent by answering a small number of value queries. The proxy agent then uses an ML algorithm to generalize from those bids to the whole value space, and the efficient allocation is computed based on the generalized valuations. We discuss that this new design leads to new challenges regarding allocative efficiency, individual rationality, and incentives. However, we show that an iterative auction design and an epsilon-expressive ML algorithm address these challenges. To instantiate our design, we use support vector regression (SVR) as the ML algorithm, which enables us to formulate the winner determination problem as a succinct integer program. Finally, we present some experimental results for two stylized spectrum auction domains. Our results demonstrate that even with a small number of bids, our ML-based auctions achieve high allocative efficiency.

3.19 Learning, Optimization, and Noise

Yaron Singer (Harvard University – Cambridge, US)

Joint work of Eric Balkanski, Avinatan Hassidim, Thibaut Horel, Nicole Immorlica, Harikrishna Narasimhan, David, Parkes, Aviad Rubinstein, Jan Vondrak, Yaron Singer

We will discuss a body of work that revolves around a simple question: what can we optimize from data? This question is relevant to several areas of interest of the workshop including mechanism design and auctions, social networks, decision theory, and online learning. In many cases we wish to optimize a function that is not known, but rather learned from data. In combinatorial auctions, for example, an agent may not know her valuation but rather learns it by observing data. In influence maximization in social networks we do not know the influence model we optimize over but rather learn it from data. In a recent line of work we began exploring these questions, as discussed below.

Optimization under Noise

A natural approach to optimization from data is to first learn a surrogate function that approximates the true function generating the data, and then optimize the surrogate. The problem however, is that even when the function is convex or submodular and the learned function well approximates the true function, the optimization problem may be inapproximable, as shown in the papers below. On the positive side, in work currently under submission we show that when noise is stochastic, the optimal guarantees are achievable for canonical submodular optimization problems.

- Optimization from Noisy Preferences with Avinatan Hassidim Working paper
- Optimal Guarantees for Maximizing Noisy Submodular Functions with Avinatan Hassidim In Submission
- Submodular Optimization under Noise with Avinatan Hassidim Conference on Learning Theory (COLT) 2017
- Robust Guarantees for Stochastic Greedy Algorithms with Avinatan Hassidim International Conference on Machine Learning (ICML) 2017
- Maximizing Approximately Submodular Functions with Thibaut Horel Annual Conference on Neural Information Processing Systems (NIPS) 2016
- Information-theoretic Lower Bounds for Convex Optimization with Erroneous Oracles with Jan Vondrak Annual Conference on Neural Information Processing Systems (NIPS) 2015

Optimization from Samples

Recently, my group has been focused on the question of optimization from data: how much training data do we need in order to optimize a function? In a paper published at NIPS 2015 we show that various submodular functions that arise in the context of diffusion in networks are statistically learnable from data. In recent work we show a sharp impossibility result. We show that various important submodular and convex functions that are statistically learnable such as diffusion functions cannot be optimized from sampled data. The moral is that there are models (functions) that may be statistically learnable and amenable to optimization (submodular or convex), though would still require exponentially-many samples for any algorithm to optimize. On the positive side, in work published at NIPS 2016 we give optimal algorithms for functions with bounded curvature and show they can be optimized well from samples.

- Maximizing the Spread of Influence From Training Data with Eric Balkanski and Nicole Immorlica In Submission
- Minimizing a Submodular Function from Samples with Eric Balkanski In Submission
- The Sample Complexity of Optimizing a Convex Function with Eric Balkanski Conference on Learning Theory (COLT) 2017
- The Limitations of Optimization from Samples with Eric Balkanski and Aviad Rubinstein Symposium on Theory of Computation (STOC) 2017
- The Power of Optimization from Samples with Eric Balkanski and Aviad Rubinstein Annual Conference on Neural Information Processing Systems (NIPS) 2016
- Learnability of Influence in Networks with Harikrishna Narasimhan and David Parkes Annual Conference on Neural Information Processing Systems (NIPS) 2015

3.20 Learning from Untrusted Data

Gregory Valiant (Stanford University, US)

We consider the problems of estimation, learning, and optimization over a large dataset of which a subset consists of points drawn from the distribution of interest, and we make no assumptions on the remaining points–they could be well-behaved, extremely biased, or adversarially generated. We investigate this question via two models for studying robust estimation, learning, and optimization. One of these models, which we term the "semi-verified" model, assumes access to a second much smaller (typically constant-sized) set of "verified" or trusted data points that have been drawn from the distribution of interest. The key question in this model is how a tiny, but trusted dataset can allow for the accurate extraction of the information contained in the large, but untrusted dataset. The second model, "list-decodable learning", considers the task of returning a small list of proposed answers. Underlying this model is the question of whether the structure of "good" datapoints can be overwhelmed by the remaining data–surprisingly, the answer is often "no". We present several strong algorithmic results for these models, for a large class of mathematically clean and practically relevant robust estimation and learning tasks.

The talk is based on several joint works with Jacob Steinhardt and Moses Charikar, and with Michela Meister.

3.21 Prediction with a Short Memory

Gregory Valiant (Stanford University, US)

We consider the problem of predicting the next observation given a sequence of past observations. We show that for any distribution over sequences of observations, if the mutual information between past observations and future observations is upper bounded by J, then a simple Markov model over the most recent J/ϵ observations obtains expected KL error ϵ–and hence L_1 error $\sqrt{\epsilon}$–with respect to the optimal predictor that has access to the entire past. For a Hidden Markov Model with n states, J is bounded by $\log n$, a quantity that does not depend on the mixing time. We also establish that this result cannot be improved upon, in the following senses: First, a window length of J/ϵ is information-theoretically necessary for expected KL error epsilon or L_1 error $\sqrt{\epsilon}$. Second, the $d^{(J/\epsilon)}$ samples required to accurately estimate the Markov model when observations are drawn from an alphabet of size d is necessary for any computationally tractable learning/prediction algorithm, assuming the hardness of strongly refuting a certain class of CSPs.

3.22 Sample Complexity of Multi-Item Profit Maximization

Ellen Vitercik (Carnegie Mellon University – Pittsburgh, US)

We study the design of pricing mechanisms and auctions when the mechanism designer does not know the distribution of buyers' values. Instead, the mechanism designer receives a set of samples from this distribution and his goal is to use the sample to design a pricing mechanism or auction with high expected profit. We provide generalization guarantees which bound the difference between average profit on the sample and expected profit over the distribution. These bounds are directly proportional to the intrinsic complexity of the mechanism class the designer is optimizing over. We present a single, general theorem that uses empirical Rademacher complexity to measure the intrinsic complexity of a variety of widely-studied single- and multi-item auction classes, including affine maximizer auctions, mixed-bundling auctions, and second-price item auctions. This theorem also applies to multi- and single-item pricing mechanisms in both multi- and single-unit settings, such as linear and non-linear pricing mechanisms. Despite the applicability of our main theorem, we match or improve over the best-known generalization guarantees for many mechanism classes. Finally, our central theorem allows us to easily derive generalization guarantees for every class in several finely grained hierarchies of auction and pricing mechanism classes. We demonstrate how to determine the precise level in a hierarchy with the optimal tradeoff between profit and generalization using structural profit maximization. The mechanism classes we study are significantly different from well-understood function classes typically found in machine learning, so bounding their complexity requires a sharp understanding of the interplay between mechanism parameters and buyer valuations.

3.23 Learning Mastermind – Some Ideas and Challenges

Bernhard von Stengel (London School of Economics, GB)

In the puzzle game "Mastermind" (available under the name "Guess" at http://hewgill.com/puzzles/), a hidden random code has to be guessed, which is a sequence of 4 pegs, each of which has 6 possible colors. Partial feedback is given on each guess in the form of black pegs (one for each correct colour and place, but never where) and white pegs (one for each other correct colour in the wrong place). The 14 possible feedback combinations (all are possible except 3 black 1 white) partition the remaining codes. Minimizing the maximal partition class allows to find the code with at most 5 guesses (Knuth 1976), and 4.34 expected guesses can be achieved optimally. These may require inconsistent guesses which cannot win immediately.

However, these strategies are not suitable for humans. The main cognitive problem is to anticipate the possible feedback information in order to find a good guess. Instead, a human will use simple rules such as "break symmetry", "use two colors", "stay consistent from the third guess onwards". This puzzle should be a suitable challenge for learning algorithms in order to identify complexity-reducing rules that are obvious to humans, such as symmetry of a guessing strategy in colors and position, and how the rules of the game are defined. Methodologically, one needs to be clear if these rules will be assumed or are to be learned. Another issue, of possible general interest, is the use of introspection (of how a human plays) to validate a learning method.

3.24 Peer Prediction Mechanisms and their Connections to Machine Learning

Jens Witkowski (ETH Zürich, CH)

Joint work of David C. Parkes, Rafael Frongillo, Jens Witkowski
Main reference Rafael M. Frongillo, Jens Witkowski: "A Geometric Perspective on Minimal Peer Prediction", ACM Trans. Economics and Comput., Vol. 5(3), pp. 17:1–17:27, 2017.
URL http://dx.doi.org/10.1145/3070903

Peer prediction mechanisms truthfully elicit private information, such as opinions, experiences, or ratings, from self-interested participants. For example, peer prediction can be used to elicit truthful responses to questions such as "Does this blog contain offensive content?", "Would you recommend this hotel to a friend?", or "Would you consider this article fake news?" Importantly, peer prediction mechanisms elicit truthful responses to these questions without ever observing ground truth at any point, e.g., whether a blog does indeed contain offensive content. In this talk, I will give a brief introduction to the peer prediction problem followed by some of the recent work that is exploring connections between peer prediction mechanisms and machine learning.

3.25 Learning Cooperative Solution Concepts

Yair Zick (National University of Singapore, SG)

Joint work of Maria-Florina Balcan, Ariel D. Procaccia, Yair Zick

Solution concepts in cooperative games have usually been considered in the setting where agent preferences are fully known. In recent work, we consider the setting where agent preferences are unknown; we lay the theoretical foundations for studying the interplay between coalitional stability and (PAC) learning in hedonic games and cooperative games. We introduce the notion of PAC stability – the equivalent of core stability under uncertainty – and examine the PAC stabilizability and learnability of several popular classes of hedonic and cooperative games. We show that all classic cooperative games with transferable utility can be PAC stabilized; however, not all hedonic games can be PAC stabilized.

Participants

- Yakov Babichenko
Technion – Haifa, IL
- Amir Ban
Tel Aviv University, IL
- Siddharth Barman
Indian Institute of Science – Bangalore, IN
- Peter L. Bartlett
University of California – Berkeley, US
- Gianluca Brero
Universität Zürich, CH
- Rachel Cummings
California Institute of Technology – Pasadena, US
- Argyris Deligkas
Technion – Haifa, IL
- Paul Dütting
London School of Economics, GB
- John Fearnley
University of Liverpool, GB
- Felix Fischer
University of Glasgow, GB
- Hu Fu
University of British Columbia – Vancouver, CA
- Claudio Gentile
University of Insubria – Varese, IT
- Paul W. Goldberg
University of Oxford, GB
- Jason D. Hartline
Northwestern University – Evanston, US
- Martin Hoefer
Goethe-Universität – Frankfurt a. M., DE
- Aleck Johnsen
Northwestern University – Evanston, US
- Thomas Kesselheim
TU Dortmund, DE
- Max Klimm
HU Berlin, DE
- Sébastien Lahaie
Google – New York, US
- Yishay Mansour
Tel Aviv University, IL
- Jamie Morgenstern
University of Pennsylvania – Philadelphia, US
- Andrés Muñoz Medina
Google – New York, US
- Paresh Nakhe
Goethe-Universität – Frankfurt a. M., DE
- Denis Nekipelov
University of Virginia – Charlottesville, US
- Ariel D. Procaccia
Carnegie Mellon University – Pittsburgh, US
- Rahul Savani
University of Liverpool, GB
- Sven Seuken
Universität Zürich, CH
- Yaron Singer
Harvard University – Cambridge, US
- Alexander Skopalik
Universität Paderborn, DE
- Gregory Valiant
Stanford University, US
- Ellen Vitercik
Carnegie Mellon University – Pittsburgh, US
- Bernhard von Stengel
London School of Economics, GB
- Jens Witkowski
ETH Zürich, CH
- Yair Zick
National University of Singapore, SG

Report from Dagstuhl Seminar 17252

Computational Challenges in RNA-Based Gene Regulation: Protein-RNA Recognition, Regulation and Prediction

Edited by

Rolf Backofen[1], Yael Mandel-Gutfreund[2], Uwe Ohler[3], and Gabriele Varani[4]

1 **Universität Freiburg, DE,** backofen@informatik.uni-freiburg.de
2 **Technion – Haifa, IL,** yaelmg@tx.technion.ac.il
3 **Max-Delbrück-Centrum – Berlin, DE,** uwe.ohler@mdc-berlin.de
4 **University of Washington – Seattle, US,** varani@chem.washington.edu

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 17252 "Computational Challenges in RNA-Based Gene Regulation: Protein-RNA Recognition, Regulation and Prediction".

Seminar June 18–21, 2017 – http://www.dagstuhl.de/17252
1998 ACM Subject Classification I.5 Pattern Recognition, I.5.1 Models, Statistical, I.5.2 Design Methodology, Feature evaluation and selection, I.5.3 Clustering, J.2 Physical Sciences and Engineering, Chemistry, J.3 Life and Medical Sciences, Biology and genetics, Health
Keywords and phrases Machine learning, algorithms, genomics analysis, gene expression networks, big data analysis, quantitative prediction, proteins, RNA
Digital Object Identifier 10.4230/DagRep.7.6.86

1 Executive Summary

Rolf Backofen
Yael Mandel-Gutfreund
Uwe Ohler
Gabriele Varani

All living organism must be able to differentially regulate the expression of genes encoded in their genome. Genes are first transcribed into RNA, which are either translated to proteins or functionally active as non-coding RNAs. Beside the direct regulation of the transcription of DNA into RNA, an important additional layer is the direct regulation of RNAs by RNA binding proteins (RBPs). This layer of regulation controls cellular decisions as part of gene expression networks composed of both proteins and RNAs. While being a dark matter of the cell for a long time, recent years have shown the development of sophisticated high throughput experimental technologies that greatly increased our understanding of protein-RNA recognition and regulation. Nevertheless, the quantitative molecular understanding of the transcriptome-level processes remains very limited. Especially complexity **(both in the form of data as in the required computational approaches)** limits the exploitation of these advances towards a quantitative understanding of post-transcriptional regulation.

Computational Challenges in RNA-Based Gene Regulation: Protein-RNA Recognition, Regulation and Prediction, *Dagstuhl Reports*, Vol. 7, Issue 06, pp. 86–108
Editors: Rolf Backofen, Yael Mandel-Gutfreund, Uwe Ohler, and Gabriele Varani
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The objective of the seminar to discuss urgently needed computational approaches allowing to exploit the wealth of new data. More specifically, the seminar focused on

- addressing major computational challenges in this field
- mining the extensive genomic information on RNA and associated proteins
- investigation of RNA-protein interactions on an atomic level
- quantitative prediction of cellular regulatory networks and their dynamics.

2 Table of Contents

3 Introduction

3.1 Seminar Format

This seminar arose from the organizer's joint experience that an integration of new experimental tools with new computational approaches is needed to understand RNA-based gene regulation at the level of sophistication and depth required to answer fundamental biological and biomedical questions. It is not sufficient for structural biologists or biochemists to only understand in-depth what other structural biologists are doing: the greatest opportunity for progress lies in the joint analysis of data from different experimental approaches by new computational tools. Experimentalists often communicate with each other in different venues, but it is rare for computational biologists and experimentalists to meet.

In the seminar, we were able to bring together structural and genome biologists, who have developed powerful experimental methods to investigate RNA-protein interaction across genomes, with computational biologists who seek to model and develop predictive tools based on the confluence of these experimental advances. Individual sessions covered different aspects of RNA-based regulation and usually consisted of a mixture of presentations from experimental and computational groups. Discussions where very vivid, fostering the exchange of ideas between the experimental and computational biologists. Discussions also catalyzed the development of new and improved technologies to merge experimental analysis with new computational techniques. The seminar was attended by many leading scientists in the field, further stimulating discussions on how to better interact (possibly on the level of funding opportunities) for a better understanding of the rules of protein-RNA recognition.

3.2 Advances and challenges in studying protein-RNA interactions

The most significant advance in the field of RNA-protein interactions in the last 10 years has been the development of genome-wide approaches to determine the RNA population targeted by an RNA-binding protein (RBP) and to establish its specificity using computational approaches. These approaches, termed CLIP-seq, have now been introduced in several variants, requiring specific computational analysis pipelines. In this session, the participants reported on the one hand about new and large-scale experimental approaches to determine binding sites of RNA-binding proteins and their dynamic behavior. On the other hand, sophisticated computational tools for analyzing CLIP-seq data were shown.

3.3 Identifying new RBPs

Great advances have recently been made in the development of high-throughput screens to identify novel RBPs in cultured cell lines and in tissues. While the technology has contributed dramatically to the field of RNA-protein interactions, increasing the number of identified RBPs and suggesting novel cellular mechanisms for these proteins, these high throughput technologies generate many false positive and false negative results. The participants discussed new approaches for the purification of RNPs, which are complexes of proteins and RNA, as well as specific binding properties. Here, computational approaches allow to investigate multiple functions for these RBPs, showing the flexibility of this layer of regulation.

3.4 Integrative analysis of protein-RNA data

The genome-wide CLIP-methods and their analysis as discussed in session 1 provides the raw data for investigating RNA-centered gene regulation at the level of individual binding sites. However, only an integrative analysis of these data can determine the modes of regulation in more detail. The major topic was how to combine the information of several single CLIP-experiments for exploring RBP-based regulation. The participants discussed how an integrated map of RBP-based regulation based on several CLIP experiments can be constructed, and how advanced machine learning approaches such as matrix factorization can be used for an integrative analysis of multiple CLIP data sets.

3.5 Exploring the world of non-coding RNAs

The majority of the human genome is transcribed into non-coding RNAs, which are not further translated into proteins. Several non-coding RNAs play an important role as assembly platform for proteins. This process is mediated via RNA-protein interactions. For that reason, we decided to invite participants that report about novel findings on non-coding RNA, showing how RNA structure evolves and which effect single mutations (SNP) can have on the structure of non-coding RNAs. In addition, we discussed non-coding RNAs with atypical features.

3.6 Inferring RNA binding specificity

The genome-wide approaches reported in Session 1 (CLIP-seq) determine *in vivo* interactions. These experiments, however, do not determine the biophysical affinity of binding. Here, *in vitro* approaches such as SELEX or RNAcompete are the method of choice, which require their own type of computational analysis. The participants discussed progress both on the experimental side, as well as on the analysis of these types of data. An interesting discussion evolved about how to compare the results of these *in vitro* approaches with the results of *in vivo* CLIP experiments.

3.7 RBPs and small RNAs

In the preceding sessions, we talked in great detail about RNA-protein interactions. However, another mode of regulation is given by direct interactions between RNAs themselves. Very often, these RNA-RNA interactions are mediated by RNA-binding proteins, a well-known example being microRNAs. The interaction of microRNAs with its target RNA is mediated by Argonaute, which is an RNA-binding protein. The participants discussed new large-scale approaches for determining RNA-RNA interactions, and how one can model the combined effect of RNA-RNA and RNA-protein interactions.

3.8 Conclusions

The community around genome-wide determination protein-RNA interaction was very interdisciplinary from the beginning as the whole area started with the availability of CLIP

data, which was challenging both from the experimental as well as from the computational side. However, as there are only few venues where computational and experimental people meet on a specific topic, this workshop was actually the second event where this community could meet. For that reason, it was well appreciated by all participants and had initiated very interesting interdisciplinary collaborations that would not have been formed without this meeting.

4 Overview of Talks

4.1 The solution structure of FUS bound to RNA reveal a bipartite mode of RNA recognition with both sequence and shape specificities.

Frédéric Allain (ETH Zürich, CH)

Joint work of Fionna Loughlin, Peter Lukavsky, Sebastien Campagne, Stefan Reber, Martino Colombo, Eva-Maria Hock , Oliver Muhlemann, Magdalini Polymenidou, Marc-David Ruepp, Frédéric Allain

FUsed in Sarcoma (FUS) is a multifunctional hnRNP which regulates transcription, pre-mRNA splicing, miRNA biogenesis, DNA damage and can drive subcellular phase transitions in RNP granule formation. Two neurodegenerative diseases Amyotrophic Lateral Sclerosis (ALS) and Fronto-Temporal Lobar Degeneration (FTLD) show neuropathological protein aggregates containing FUS and it is hypothesized that mis-regulation of RNA processing could play a major role in these diseases.

FUS consists of a N-terminal prion-like region and an C-terminal RNA binding region including several RNA binding domains (RBD), an RNA Recognition Motif (RRM), and a zinc-finger (ZnF) domain interspersed between RGG repeats. FUS associates with a large variety of RNAs including pre-mRNA, pri-miRNA, lncRNA from CCND1 promoter and TERRA. Results from a dozen in vivo (CLIP) and in vitro (SELEX, binding assays) experiments suggest that RNA binding by FUS is very complex potentially combining structural, sequence and nonspecific interactions with additional influences from post-translational modifications and protein partners.

We present the solution structures of FUS RRM and ZnF domains bound to RNA. The ZnF shows a sequence-specific recognition for a single-stranded NGGU motif and this interaction accounts for the preference for GU-rich motifs found in several CLIP based experiments. The FUS RRM structure was solved bound to a stem-loop RNA and revealed an unusual shape-specific binding mode. The RRM binds the 3' side of the RNA loop using the β-sheet and C-terminal helix with very limited sequence-specificity and the RNA stem using a α_1-β_2 extension unique to FUS, EWS and TAF15 that inserts in the major groove. Furthermore, the RGG repeats between the two RBDs significantly increase the general affinity and unfold the 5'part of the RNA loop. The RRM and ZnF can coordinate in binding a bi-partite RNA site and mutation of the RNA binding surfaces of both the ZnF and RRM domains are sufficient to abolish FUS function in splicing. These results revealed why RNA recognition by FUS has been so difficult to decipher.

4.2 RNA recognition in Syncrip-mediated exosomal loading of miRNAs

Andre Dallmann (Institute of Structural and Molecular Biology, University College London, London, UK and Humboldt University, Berlin, DE)

Joint work of Andres Ramos, Evangelos Christodoulou, Fruzsina Hobor, Ian Taylor, Marco Tripodi, Neil Ball, Roksane Ogrodowicz, Andre Dallmann
Main reference unpublished

Exosomal-mediated miRNA transfer is an important mechanism of inter-cellular communication that regulates gene expression in a broad range of cells and tissues, and has been implicated in cancer, neurological diseases and cardiomyopathies. Syncrip/hnRNPQ is a highly conserved RNA binding protein that regulates mRNA metabolism, transport and translation. It was recently reported that Syncrip plays an essential role in the exosomal partitioning of a group of miRNAs. Here we analyse the molecular basis of the selectivity of his mechanism. We show that Syncrip contains a cryptic RNA binding domain which recognises a short sequence in the miRNA targets. We also discuss how this domain cooperates with the other RNA binding domains of the protein and how the RNA-binding mode of Syncrip allows the protein to recognise selectively a group of miRNA with different sequences. Finally, we discuss other examples where the combined use of bioinformatics, structural biology and functional assays was used to dissect the specificity of multi-domain proteins.

4.3 PTex – a novel method for unbiased purification of crosslinked RNPs

Benedikt Beckman (Humboldt University, Berlin, DE)

Joint work of Erika Urdaneta, Carlos Vieira, Matthias Selbach, Benedikt Beckman

RNA-binding proteins (RBPs) play important roles in post-transcriptional regulation of gene expression, e.g. as key proteins in splicing, RNA-transport or translation. In the past years, discovery of RNA-protein complexes (RNPs) has been driven by high-throughput techniques such as CLIP-Seq or mRNA interactome capture, leading to deeper insights into RNA targets bound by specific proteins or the discovery of hundreds of novel mRNA-binding proteins, respectively. Common to all of these methods is the utilisation of UV light to induce covalent crosslinks between RNA and proteins prior to purification.

Here, we present PTex (Phenol-Toulol EXtraction), a novel technique for the unbiased purification of in vivo UV-crosslinked RNPs. Our method consists of a fast extraction using different organic compounds thereby removing non-crosslinked RNA and proteins. Purification with PTex relies solely on physicochemical differences between crosslinked and "free" RNA/protein and thus enables analysis of non-poly(A) RNA interacting proteins, e.g. for interactors of ribosomal or tRNA. PTex can be performed in 2.5–3 hrs.

We tested PTex in human HEK293 cells and analysed the purified proteins by mass spectrometry. PTex-purified proteins (up to 3000, FDR 0.01) are highly enriched in known RBPs from mRNA interactome capture experiments as well as ribosomal proteins and tRNA-interacting RBPs.

4.4 Towards fully flexible modeling of protein-RNA complexes

Janusz Bujnicki (Int. Inst. of Molecular and Cell Biology – Warsaw, PL) and Michal Boniecki

Joint work of Janusz Bujnicki, Michal Boniecki
URL http://genesilico.pl

Macromolecular complexes composed of proteins and nucleic acids play fundamental roles in many biological processes, such as the regulation of gene expression, RNA splicing and protein synthesis. Our group has developed computational tools for protein and RNA 3D structure prediction, which covered approaches for template-based and template-free modeling. These tools, including the GeneSilico metaserver and Frankenstein Monster approach for proteins, and ModeRNA and SimRNA for RNA have been validated in CASP and RNA Puzzles experiments, respectively. Recently, we combined our approaches for protein and RNA modeling, and developed SimRNP, a method for modeling of protein-RNA complex structures.

4.5 Matrix factorization-based integrative analysis of multiple protein-RNA data sets

Tomaž Curk (University of Ljubljana, SI)

Joint work of Martin Stražar, Marinka Žitnik, Blaž Zupan, Jernej Ule, Tomaž Curk
Main reference Martin Strazar, Marinka Zitnik, Blaz Zupan, Jernej Ule, Tomaz Curk: "Orthogonal matrix factorization enables integrative analysis of multiple RNA binding proteins", in Bioinformatics, Vol. 32(10), pp. 1527–1535, 2016.
URL http://dx.doi.org/10.1093/bioinformatics/btw003

We have presented an integrative orthogonality-regularized nonnegative matrix factorization (iONMF) to integrate multiple data sources and discover non-overlapping, class-specific RNA binding patterns of varying strengths. The orthogonality constraint halves the factor model and outperforms other NMF models in predicting RBP interaction sites on RNA. We have integrated a large data compendium, which includes 31 CLIP experiments on 19 RBPs involved in splicing (such as hnRNPs, U2AF2, ELAVL1, TDP-43 and FUS) and processing of 3'UTR (Ago, IGF2BP). The integration of multiple data sources improves the predictive accuracy of retrieval of RNA binding sites. The key predictive factors of protein–RNA interactions were the position of RNA structure and sequence motifs, RBP co-binding and gene region type. We report on a number of protein-specific patterns, many of which are consistent with experimentally determined properties of RBPs.

4.6 Identification of transcriptomic regulatory elements from CLIP-Seq data

Philipp Drewe (Max-Delbrück-Centrum – Berlin, DE)

Joint work of Philipp Drewe-Boss, Hans-Hermann Wessels, Uwe Ohler
Main reference Philipp Drewe-Boss, Hans-Hermann Wessels, Uwe Ohler: "omniCLIP: Bayesian identification of protein-RNA interactions from CLIP-Seq data", in bioRxiv, Cold Spring Harbor Laboratory, 2017.
URL http://dx.doi.org/10.1101/161877

High-throughput immunoprecipitation methods to analyze RNA binding protein-RNA interactions and modifications have great potential to further the understanding of post-transcriptional gene regulation. Due to the differences between individual approaches, each of a diverse number of computational methods can typically be applied to only one specific sequencing protocol. Here, we present a Bayesian model called omniCLIP that can be applied to data from all protocols to detect regulatory elements in RNAs. omniCLIP greatly simplifies the data analysis, increases the reliability of results and paves the way for integrative studies based on data from different sources.

4.7 RNA binding properties of a metabolic enzyme

André Gerber (University of Surrey – Guildford, GB)

Joint work of Ana M. Matia-González, Valentina Iadevaia, Emma E. Laing, André P. Gerber

RNA binding proteins (RBPs) are essential for post-transcriptional regulation of gene expression with widespread implications in development and disease. We recently described the repertoire of hundreds of RBPs that interact with polyadenylated mRNAs in the unicellular yeast Saccharomyces cerevisiae and in the multicellular nematode Caenorhabditis elegans, respectively [1]. Thereby, we found that most proteins comprising the mRNA-binding proteomes (mRBPomes) were evolutionarily conserved, including components of early metabolic pathways such as glycolysis, which indicated an ancient origin of the mRBPomes. To further investigate the RNA-binding properties of metabolic enzymes, we are currently focusing on phosphofructokinase (Pfk), which is a key enzyme of glycolysis. In the yeast S. cerevisiae, the enzyme is a hetero-octamer comprised of Pfk1 and Pfk2 subunits. We found that both paralogous subunits can independently interact with mRNAs. Furthermore, Pfk2p specifically interacts with PFK2 mRNA; and one point of interaction could be located to a conserved sequence element in the 3'UTR of PFK2 mRNA. Furthermore, we showed that enzyme-RNA interactions depend on the metabolic state of the cells, indicating cross-talk with cellular metabolism. Furthermore, we found that interaction of PFK with its own mRNA is conserved between yeast and human. We speculate that conserved enzyme-mRNA interactions could relate to an ancient mechanism for the post-transcriptional coordination of metabolic pathways, integrating metabolic activity into dynamic gene regulatory responses.

References

1 Matia-González, A. M., Laing, E. E, Gerber, A. P. (2015) Conserved mRNA-binding proteomes in eukaryotic organisms. Nat. Struct. Mol. Biol. 22(12), 1027-33.

4.8 Large-scale elucidation of RNase III targets and cleavage patterns

Hanah Margalit (The Hebrew University of Jerusalem, IL)

Joint work of Hanah Margalit, Yael Altuvia, Liron Argaman, Amir Bar, Ehud Karavani

The family of Ribonuclease III (RNase III) enzymes spans all kingdoms of life, with its most famous representatives the eukaryotic enzymes Drosha and Dicer. A founding member of the family is RNase III of Escherichia coli, which was discovered about 50 years ago and extensively studied. RNase III cleaves double-stranded RNA, typically generating two cleavages in the two sides of a stem structure while leaving 2-nt 3' overhangs. At present, 23 target genes of RNase III are known in E. coli, all of them were discovered in specific experiments focusing on particular RNAs. We reasoned that the availability of E. coli genome and the opportunities that the large-scale RNA-seq technique offers open the way for transcriptome-wide mapping of RNase III targets in E. coli. To this end we developed Cleave-seq, a method for global mapping of the cleavage sites of any endoribonuclease of interest at a nucleotide resolution, taking advantage of the chemical group at the 5' end of the cleavage product. Exploiting the characteristic 5'-phosphomonoester termini of RNase III cleavage products, we constructed sequencing libraries deigned in a way that the start position of a RNA sequenced segment should be a 5'P position, and hence, generated by an RNase that leaves 5'P ends. To identify 5'P ends generated by RNAse III we applied the protocol to RNA extracted from E. coli MG1655 WT strain and a mutant strain lacking active RNase III, and mapped the sequenced fragments to the genome. Positions where the start of sequenced fragments mapped in the WT but not in the mutant ($p \leq 0.05$ by DESeq2) were identified as RNase III cleavage sites. We re-discovered 20 out of the 23 known targets, and extended the set of cleavage sites in E. coli to $\tilde{4}00$ putative sites, residing in RNAs derived from a variety of genomic entities (CDSs, 5' and 3' UTRs, intergenic regions and antisense to annotated transcripts). 30 % of the targets had a pair of cleavage sites that followed the known cleavage rules, while the rest presented only one cleavage site. We are now investigating these single sites in search of new, additional, cleavage rules that might dictate RNase III cleavage of these targets. Especially intriguing is the possibility that base-pairing of regulatory small RNAs with these target RNAs generates substrates for RNAse III cleavage.

4.9 Kinetics of RNA-protein interactions in cells

Eckhard Jankowsky (Case Western Reserve University – Cleveland, US)

RNA-binding proteins (RBPs) often interact with many different RNAs at sometimes large numbers of binding sites. It is widely accepted that the kinetics by with a given RBP interacts with each of its binding sites is critical for the impact of the RBP on RNA metabolism. While kinetics of RBP binding to RNA can be readily measured in vitro, it has not yet been possible to determine kinetic parameters of RBP interactions with individual RNA sites in cells. Here, we report an approach to measure kinetic parameters by which the RNA binding protein Dazl from Mouse interacts with thousands of RNA binding sites in cells. We describe

how a combination of time-resolved UV-laser crosslinking, next generation sequencing and kinetic modeling can be utilized to determine on- and off rate constants, as well as fractional occupancy of Dazl to thousands of RNA binding sites in cells. Our data enable quantitative, biophysical descriptions of RNA-protein interactions in cells.

4.10 Modeling the combined effect of RBPs and miRNAs in post-transcriptional regulation

Hilal Kazan (Antalya International University, TR)

Joint work of Saber HafezQorani, Atefeh Lafzi, Ruben G de Bruin, Anton Jan van Zonneveld, Eric P van der Veer, Yesim Aydin Son, Hilal Kazan
Main reference Saber HafezQorani, Atefeh Lafzi, Ruben G. de Bruin, Anton Jan van Zonneveld, Eric P. van der Veer, Yeşim Aydın Son, Hilal Kazan: "Modeling the combined effect of RNA-binding proteins and microRNAs in post-transcriptional regulation", in Nucleic Acids Research, Vol. 44(9), p. e83, 2016.
URL http://dx.doi.org/10.1093/nar/gkw048

Recent studies support that RNA-binding proteins (RBPs) and microRNAs (miRNAs) function in coordination with each other to control post-transcriptional regulation (PTR). However, the majority of research to date has focused on the regulatory effect of individual RBPs or miRNAs. Here, we mapped both RBP and miRNA binding sites on human 3'UTRs and utilized this collection to better understand PTR. We show that the transcripts that lack competition for HuR binding are destabilized more after HuR depletion. We also validate this finding for PUM1(2) by knocking down PUM1 and PUM2 in HEK293 cells and measuring genome-wide gene expression changes. Next, to find potential cooperative interactions, we identified the pairs of factors whose sites co-localize more often than expected by random chance. We show that transcripts where the sites of PUM1(2) and its interacting miRNA form a stem-loop are stabilized more upon PUM1(2) depletion. Finally, using dinucleotide frequency and counts of regulatory sites as features in a regression model, we achieved an AU- ROC of 0.86 in predicting mRNA half-life in BEAS-2B cells. Our results suggest that combining the effects of RBPs and miRNAs lead to more accurate models of PTR.

4.11 RNA genotype-phenotype mapping

Grzegorz Kudla (University of Edinburgh, GB)

Recent advances in synthetic biology and next-generation sequencing allow experimental investigation of long-standing, fundamental questions in molecular biology and evolution, such as: Which mutations influence function? How do mutations influence function? How do the effects of mutations depend on location within molecular structures? How do the effects depend on environmental conditions and genetic background? We are studying these questions using yeast U3 snoRNA as a model system. U3 is an abundant, evolutionarily conserved noncoding RNA, which plays an essential role in ribosome biogenesis. By measuring the effects of 60,000 mutated variants of U3 on yeast growth, we found that the effects of individual mutations were correlated with evolutionary conservation and structural stability. Many

mutations had no measurable effect in an otherwise wild-type background, but were deleterious in combination with additional mutations in U3. We also found pairs of compensatory mutations, and used these to predict the secondary structure of the RNA. The effects of mutations depended on environmental conditions, and the effects of destabilizing mutations were increased at high temperature. In parallel with mutational studies, we use a high-throughput method for mapping RNA-RNA interactions, called CLASH, to investigate the folding of U3 and preribosomal RNA. By integrating these data with recent cryo-electron microscopy structural studies of preribosomes, we aim to systematically elucidate the relations between genotype, structure, molecular function, and evolutionary fitness.

4.12 Exploring global changes in protein-mRNA interactions

Markus Landthaler (Max-Delbrück-Centrum – Berlin, DE)

Joint work of Miha Milek, Koshi Imami, Neelanjan Mukherjee, Francesca De Bortoli, Ulrike Zinnall, Orsalia Hazapis, Christian Trahan, Marlene Oeffinger, Florian Heyd, Uwe Ohler, Matthias Selbach, Markus Landthaler
Main reference Miha Milek, Koshi Imami, Neelanjan Mukherjee, Francesca De Bortoli, Ulrike Zinnall, Orsalia Hazapis, Christian Trahan, Marlene Oeffinger, Florian Heyd, Uwe Ohler, Matthias Selbach, Markus Landthaler: "DDX54 regulates transcriptome dynamics during DNA damage response", in Genome Research, 10.1101/gr.218438.116, 2017
URL http://dx.doi.org/10.1101/gr.218438.116

The cellular response to genotoxic stress is mediated by a characterized network of DNA surveillance pathways. The contribution of posttranscriptional gene regulatory networks to the DNA damage response has been studied poorly. In an attempt to examine global changes in protein-mRNA interactions, we systematically identified RNA-binding proteins differentially bound to polyadenylated transcripts upon exposure of MCF-7 human breast carcinoma cells to ionizing radiation (IR). Surprisingly, about 260 proteins including many nucleolar proteins showed increased binding to mRNA in IR-treated cells. The analysis of DDX54, an RNA helicase, revealed that this protein is an immediate-to-early DDR regulator required for the splicing efficacy of its target IR-induced pre-mRNAs. Upon IR exposure, DDX54 acts by increased interaction with a well-defined class of pre-mRNAs which harbor introns with weak 3' acceptor splice sites, as well as by protein-protein contacts within components of U2 snRNP and spliceosomal B complex, resulting in reduced intron retention and higher processing rates of its target pre-mRNA transcripts. Since DDX54 promotes survival after exposure to IR, its expression may impact DDR-related pathologies. This work indicates the relevance of many uncharacterized RNA-binding proteins involved in the cellular response to DNA damage.

4.13 iCLIP determines the binding landscape of the clock-regulated RNA-binding protein AtGRP7

Martin Lewinski (Universität Bielefeld, DE)

Joint work of Katja Meyer, Tino Köster, Christine Nolte, Claus Weinholdt, Martin Lewinski, Ivo Grosse, Dorothee Staiger

Background: A key function for RNA-binding proteins in orchestrating plant development and environmental responses is well established. However, the lack of a genome-wide view on their in vivo binding targets and binding landscapes represents a gap in understanding the mode of action of plant RNA-binding proteins. Here, we adapt individual nucleotide resolution crosslinking immunoprecipitation (iCLIP) for genome-wide determining the binding repertoire of the circadian clock-regulated Arabidopsis thaliana glycine-rich RNA-binding protein AtGRP7. Results: iCLIP identified 858 transcripts with significantly enriched crosslink sites in plants expressing AtGRP7-GFP and absent in plants expressing an RNA-binding-dead AtGRP7 variant or GFP alone. To independently validate the targets, we performed RNA immunoprecipitation (RIP)-sequencing of AtGRP7-GFP plants subjected to formaldehyde fixation. 452 of the iCLIP targets were also identified by RIP-seq, thus representing a set of high-confidence binders. AtGRP7 can bind to all transcript regions with a preference for 3´untranslated regions. In the vicinity of crosslink sites, UC-rich motifs were overrepresented. Cross-referencing the targets against transcriptome changes in AtGRP7 loss-of-function mutants or AtGRP7-overexpressing plants revealed a predominantly negative effect of AtGRP7 on its targets. In particular, elevated AtGRP7 levels lead to damping of circadian oscillations of transcripts including DORMANCY/AUXIN ASSOCIATED FAMILY PROTEIN2 and CCR-LIKE. Furthermore, several targets show changes in alternative splicing or polyadenylation in response to altered AtGRP7 levels. Conclusion: We have established iCLIP for plants to identify target transcripts of the RNA-binding protein AtGRP7. This paves the way to investigate the dynamics of posttranscriptional networks in response to exogenous and endogenous cues.

4.14 Challenges in exploring the world of RNA binding proteins

Yael Mandel-Gutfreund (Technion – Haifa, IL)

Joint work of Shlomi Dvir, Chen Lesnik, Inbal Paz, Michal Amit, Yael Mandel-Gutfreund

DNA binding proteins (DBPs) and RNA binding proteins (RBPs) are the main regulators of gene expression in the cell. In the context of tight interactions between transcriptional and post-transcriptional regulation, proteins that bind both DNA and RNA, namely DNA and RNA Binding Proteins (DRBPs), are highly likely to be key players in mediating the cross talk between the different processes of gene expression pathway. Human Embryonic Stem Cells (hESCs) provide a biologically valuable and experimentally tractable model system to study the unappreciated group of dual nucleic acids binding proteins. As a first step to study the role of DRBPs in pluripotency we employed an RNA-interactome experiment on hESCs, identifying over 800 high-confidence RBPs, among them we identified 25% DRBPs. Among the proteins detected as RNA binding in our experiment more 120 proteins

were not previously shown to interact with either RNA or DNA in mammalian cells. We further differentiated the hESCs to embryoid bodies (EBs) using a spontaneous differentiation protocol and followed the RNA and protein expression levels the process. We found that a large fraction of the detected RBPs and DRBPs in the hESC interactome were significantly down regulated upon differentiation, suggesting that dual DNA and RNA binding function may play an important role in shaping the stem cell state in human cells.

4.15 Accurate identification of RBP binding sites and RNA sequence-structure motifs from CLIP-seq data

Annalisa Marsico (MPI für Molekulare Genetik – Berlin, DE)

Joint work of Sabrina Krakau, Hugues Richard, Annalisa Marsico

Interactions between RNAs and RNA binding proteins (RBPs) play essential roles in both transcriptional and post-transcriptional gene regulation. iCLIP and eCLIP techniques allow the detection of protein-RNA interaction sites at high resolution, based on diagnostic events at crosslink sites. Previous methods do not explicitly model the specfics of iCLIP and eCLIP truncation patterns and possible biases. We introduce PureCLIP, a hidden Markov model based approach, which simultaneously performs peak calling and individual crosslink site detection. It explicitly incorporates RNA abundances and, for the first time, non-specfic sequence biases. On both simulated and real data, PureCLIP is more accurate in calling crosslink sites than other state-of-the-art methods and has a higher agreement across replicates.

4.16 CLIPing STAR proteins: target specificity via compartmentalisation

Daniel Maticzka (Universität Freiburg, DE)

Joint work of Daniel Maticzka, Ibrahim Avsar Ilik, Tugce Aktas, Rolf Backofen, Asifa Akhtar
Main reference Daniel Maticzka, Ibrahim Avsar Ilik, Tugce Aktas, Rolf Backofen, Asifa Akhtar: "uvCLAP: a fast, non-radioactive method to identify in vivo targets of RNA-binding proteins", in bioRxiv, Cold Spring Harbor Laboratory, 2017.
URL http://dx.doi.org/10.1101/158410

We have developed uvCLAP (UV crosslinking and affinity purification), a method for measuring RNA-protein interactions in vivo. To test its performance and applicability we investigated binding of 15 RBPs from fly, mouse and human cells. Our results show that despite their different subcellular localizations, STAR proteins (KHDRBS1-3, QKI) bind to a similar RNA motif in vivo. Consistently a point mutation (KHDRBS1Y440F) or a natural splice isoform (QKI-6) that changes the respective RBP subcellular localization, dramatically alters target selection without changing the targeted RNA motif. Combined with the knowledge that RBPs can compete and cooperate for binding sites, our data shows that compartmentalization of RBPs can be used as an elegant means to generate RNA target specificity.

4.17 Alternative RNA structure expression and its functional roles

Irmtraud Meyer (Max-Delbrück-Centrum – Berlin, DE)

Joint work of Alborz Mazloomian, Irmtraud M. Meyer
Main reference Alborz Mazloomian, Irmtraud M. Meyer: "Genome-wide identification and characterization of tissue-specific RNA editing events in D. melanogaster and their potential role in regulating alternative splicing", in RNA Biology, Vol. 12(12), pp. 1391–1401, Taylor & Francis, 2015.
URL http://dx.doi.org/10.1080/15476286.2015.1107703

Many studies of protein-RNA interactions aim to identify new protein-RNA interaction sites, e.g. by utilising experimentally derived CLIP-data as input for a computational prediction pipeline. One important additional aim of these computational studies is often to also deduce the yet unknown characteristic binding features for new proteins. This often requires the computational modelling of the potential RNA secondary-structure context of the protein-RNA binding sites.

We here present results of a transcriptome-wide, computational study of A-to-I editing sites for several, tissue-specific RNA-seq data sets in the fruit fly *Drosophila melanogaster*. By designing a dedicated computational prediction pipeline which employs probabilistic algorithms and which utilises only transcriptome and genome sequencing data as input (rather than experimental data on protein-RNA interaction sites), we manage to reliably identify a considerable number of new protein-RNA binding sites for this particular protein. This is only possible because our prediction pipeline explicitly captures several key features of the corresponding RNA-binding protein ADAR which is responsive for these A-to-I edits: the requirement of double-stranded RNA for potential ADAR-binding sites and the typical accumulation of several editing sites in the same stretch of double-stranded RNA. Using this dedicated probabilistic prediction strategy for one particular RNA-binding protein (ADAR) in combination with a comparative approach allows us to gain considerable biological insights into the protein's potential functional role in alternative splicing and the underlying molecular mechanisms *in vivo* without having to resort to the typical assumptions made in the field.

Our main insight gained is that ADAR-binding to double-stranded RNA can result in several A-to-I edits which can induce RNA-structure changes which can in turn induce change the splicing pattern of the underlying pre-mRNA. This turns out to be of particular importance for cells of the central nervous system, not only in the fruit fly, but also the mouse and human, and proposes a dedicated RNA-structure-based molecular mechanism for tissue-specific changes of splicing patterns.

To summarize, our study shows that

- protein-specific computational studies have the potential to gain significant biological insight,
- protein-RNA binding may require a very specific RNA secondary-structure context,
- protein-RNA binding may change the RNA secondary-structure context upon binding/acting and
- modelling the RNA secondary-structure context at it is *in vivo* can be achieved by employing a comparative approach.

References

1 I. M. Meyer, In silico methods for co-transcriptional RNA secondary structure prediction and for investigating alternative RNA structure expression, Methods 120 (2017) 3–16.
2 D. Lai, I. M. Meyer, e-RNA: a collection of web servers for comparative RNA structure prediction and visualisation, Nucleic Acids Research 42 (2014) W373–6.

4.18 Inferring RNA motifs from millions of binding sites using billions of features

Quaid Morris (University of Toronto, CA) and Tim Hughes (University of Toronto, CA)

Joint work of Kate B. Cook, Shankar Vembu, Kevin Ch. Ha, Hong Zheng, Kaitlin U. Laverty, Deb Ray, Quaid Morris, Tim Hughes
Main reference Kate B. Cook, Shankar Vembu, Kevin C.H. Ha, Hong Zheng, Kaitlin U. Laverty, Timothy R. Hughes, Debashish Ray, Quaid D. Morris: "RNAcompete-S: Combined RNA sequence/structure preferences for RNA binding proteins derived from a single-step in vitro selection", in Methods, Vol. 126(Supplement C), pp. 18–28, 2017.
URL http://dx.doi.org/10.1016/j.ymeth.2017.06.024

RNA-binding proteins (RBPs) recognize RNA sequences and structures, but there are few systematic and accurate method to derive large (>12base) motifs de novo that capture an RBPs sequence and structural binding preferences. My talk will introduce RNAcompete-S, a new method that couples a single-step competitive binding reaction with an excess of random RNA 40-mers to a new computational pipeline to derive SSMs (Sequence and Structure Models) from the millions of RNA oligos in selected and background fractions of oligos.

We use Vowpal Wabbit to identify sequence and structure k-mers that are predictive of an RNA oligo being in the selected fraction and then use Gaussian mixture model clustering to group k-mers into motifs. We then combine frequently co-occurring motifs into a single, larger motifs.

RNAcompete-S confirms that HuR, QKI, and SRSF1 prefer binding sites that are single stranded, and recapitulates known 8-10bp sequence and structure preferences for Vts1p and RBMY. We also derive an 18-base long SSM for Drosophila SLBP, which has not been previously determined by selections from pure random sequence, and accurately discriminates human replication-dependent histone mRNAs. Thus, RNAcompete-S enables accurate identification of large, intrinsic sequence-structure specificities with a uniform assay.

4.19 An integrated map of RNA-binding protein mediated gene regulation

Uwe Ohler (Max-Delbrück-Centrum – Berlin, DE), Mahsa Ghanbari (Max-Delbrück-Centrum – Berlin, DE)

Joint work of Uwe Ohler, Mahsa Ghanbari, Svetlana Lebedeva, Neelanjan Mukherjee, Alina Munteanu, Hans-Hermann Wessels, Aitor Garcia, Thomas Tuschl

RNA-binding proteins (RBPs) control and coordinate each stage in the life cycle of RNAs. Recent sequencing protocols, particularly variants of in vivo RBP-RNA crosslinking and immunoprecipitation (CLIP), have in principle enabled the identification of target sites for many of these RBPs. Yet, there has been little effort to systematically reanalyze previously published data generated by different labs, both to allow for meaningful benchmarking of new algorithms e.g. for site and motif identification and for integrative analyses of these data.

Here, we examined 118 transcriptome-wide data sets representing 67 RBPs that were generated by the PAR-CLIP protocol, to characterize target RNA class binding preference, sequence and regional specificity. Furthermore, by integrating RNA metabolism measurements, we identified regulatory modules defined by subsets of mRNAs bound by specific subsets of RBPs, which may represent post-transcriptional RNA operons. This study represents

the most comprehensive investigation of empirical RBP-RNA interaction evidence and their regulatory function in a human cell line to date, and the data has potential to be a useful gold standard for future algorithmic developments.

4.20 AptaTRACE – a novel method to discover RNA sequence-structure binding motifs

Teresa Przytycka (National Center for Biotechnology – Bethesda, US)

Joint work of Phuong Dao, Jan Hoinka, Mayumi Takahashi, Jiehua Zhou, Michelle Ho, Yijie Wang, Fabrizio Costa, John J. Rossi, Rolf Backofen, John Burnett, Teresa M. Przytycka
Main reference Phuong Dao, Jan Hoinka, Mayumi Takahashi, Jiehua Zhou, Michelle Ho, Yijie Wang, Fabrizio Costa, John J. Rossi, Rolf Backofen, John Burnett, Teresa M. Przytycka: "AptaTRACE Elucidates RNA Sequence-Structure Motifs from Selection Trends in HT-SELEX Experiments", in Cell Systems, Vol. 3(1), pp. 62–70, 2016.
URL http://dx.doi.org/10.1016/j.cels.2016.07.003

RNA aptamers are short RNA molecules capable of binding, with high affinity and specificity, a specific target molecule via sequence and structure features that are complementary to the biochemical characteristics of the target's surface. While the specifics vary depending on the target, aptamers are typically identified through the Systematic Evolution of Ligands by Exponential Enrichment (SELEX) protocol. A key reason for the resurgence of interest in aptamer research relates to the emergence the HT-SELEX – a procedure which utilizes affordable next-generation sequencing technologies along with traditional SELEX. HT-SELEX calls for new scalable analytic tools for identifying sequence-structure biding motifs from HT-SELEX data. In this talk I will present AptaTRACE, a computational approach that leverages the experimental design of the HT-SELEX protocol, RNA secondary structure, and the potential presence of many secondary motifs to identify sequence-structure motifs that selected for in the HT-SELEX experiment.

AptaTRACE is not limited to the detection of a single motif but capable of elucidating an arbitrary number of binding sites along with their corresponding structural preferences. Unlike previous methods, AptaTRACE does not rely on aptamer frequency or its derivative – cycle-to-cycle enrichment. Instead, our method builds on tracing the dynamics of the SELEX process itself to uncover motif-induced selection trends. Specifically, we expect that in the initial pool the structural contexts of each k-mer are distributed according to a background distribution that can be determined from the data. However, for sequence motifs involved in binding, in later selection cycles, this distribution becomes biased towards the structural context favored by the binding interaction with the target site. Consequently, AptaTRACE aims at identifying sequence motifs whose tendency of residing in a hairpin, bulge loop, inner loop, multiple loop, dangling end, or of being paired converges to a specific structural context throughout the selection.

We applied AptaTRACE to identify nine motifs in C-C chemokine receptor type 7 targeted by aptamers in an in vitro cell-SELEX experiment. We experimentally validated two aptamers whose binding required both sequence and structural features. AptaTRACE can identify low-abundance motifs, and we show through simulations that because of this it could lower HT-SELEX cost and time by reducing the number of selection cycles required. AptaTRACE is available for download at http://www.ncbi.nlm.nih.gov/CBBresearch/Przytycka/index.cgi#aptatools.

In the second part of the talk I will describe our recently developed method, AptaBlocks, to design Aptamer – drug delivery systems.

4.21 Determination of RNA-protein interactions

Rolf Backofen (Universität Freiburg, DE)

Main reference Tuğçe Aktaş, İbrahim Avşar Ilık, Daniel Maticzka, Vivek Bhardwaj, Cecilia Pessoa Rodrigues, Gerhard Mittler, Thomas Manke, Rolf Backofen, Asifa Akhtar: "DHX9 suppresses RNA processing defects originating from the Alu invasion of the human genome", Nature, Vol. 544, pp. 115–119, Springer Nature, 2017.
URL http://dx.doi.org/10.1038/nature21715
Main reference Daniel Maticzka, Sita J. Lange, Fabrizio Costa, Rolf Backofen: "GraphProt: modeling binding preferences of RNA-binding proteins", Genome Biology, Vol. 15(1), p. R17, 2014.
URL http://dx.doi.org/10.1186/gb-2014-15-1-r17

It is becoming increasingly clear that a RNA-binding proteins are key elements in regulating the cell's transcriptome. Thus, unraveling the interaction network of the RNA-binding proteins by determining their binding sites is becoming an increasingly important topic. CLIP-seq is one of the major tools to determine binding sites but suffers from high false negative rate due its expression dependency. This critical hinders the use of public CLIP-data. We will show in several examples how use of raw public CLIPp data can lead to false biological reasoning and how advanced advanced machine learning approach can overcome this problem. I will further discuss our results from our new Nature paper, showing that the human RNA helicase DHX9 predominantly binds to IRAlu elements and such suppresses the negative effect of Alu inflation in transcripts.

4.22 Modelling the RNA lifecycle with χ-CRAC data

Guido Sanguinetti (University of Edinburgh, GB)

Gene expression responses to cellular stress involve a wide repertoire of RNA dynamic profiles. The regulation of RNA abundance is a balance of production, processing and decay, yet how these processes combine to shape RNA profiles is only partly understood. In this talk, I describe how χ-CRAC, a new technology to assay dynamic protein-RNA interactions, can be used to start deconvolving the various components of the RNA life cycle. Evidence from our recent work [1] suggests that degradation factors such as the yeast protein Nab3 respond to stress by dynamically altering their binding. We are in the process of collecting data about other processing pathways; in the talk, I describe how we are planning to integrate such data sets within a modelling framework.

References

1 van Nues, Rob and Schweikert, Gabriele and de Leau, Erica and Selega, Alina and Langford, Andrew and Franklin, Ryan and Iosub, Ira and Wadsworth, Peter and Sanguinetti, Guido and Granneman, Sander. *Kinetic CRAC uncovers a role for Nab3 in determining gene expression profiles during stress.* Nature Communications 8, 2017

4.23 Decoding regulatory protein-RNA interactions in gene regulation using integrated structural biology

Michael Sattler (TU München, DE)

Main reference Andreas Schlundt, Jan-Niklas Tants, Michael Sattler: "Integrated structural biology to unravel molecular mechanisms of protein-RNA recognition", Methods, Vol. 118-119(Supplement C), pp. 119–136, 2017.
URL http://dx.doi.org/10.1016/j.ymeth.2017.03.015

RNA plays essential roles in virtually all aspects of gene regulation. This involves the recognition of cis regulatory RNA sequences by RNA binding proteins (RBPs). Most eukaryotic RBPs are multi-domain proteins that comprise multiple structural domains to mediate protein-RNA or protein-protein interactions. Large scale biochemical approaches to map RBP-RNA interactions in vivo at a genome-wide scale combined with computational analysis provide important information about RNA sites recognized. Also, large scale approaches such RNACompete identify RNA sequence motifs that are recognized by a given RBP in vitro. However, the roles of individual RNA binding domains and other structural modules in these multidomain RNA binding proteins are not revealed by these approaches. We employ integrative structural biology to unravel the molecular mechanisms involving these regulatory RNP (ribonucleoprotein) complexes. For these studies, solution NMR-spectroscopy, SAXS/SANS and FRET experiments provide unique information on functionally important dynamics and are combined with X-ray crystallography and electron microscopy to elucidate the structural mechanisms and dynamics of regulatory RNPs. In collaboration with the Koenig group (IMB Mainz) we show how the combination and comparative bioinformatics analysis of in vitro and in vivo CLIP experiments identifies novel regulators of RBP-RNA interactions and is a useful approach to link large scale approaches in vivo with mechanistic analysis of RNA binding by RBPs using structural biology. For the heterodimeric constitutive splicing factor U2AF we found that the tandem RNA recognition motif (RRM) domains of U2AF65 adopt an equilibrium of open and closed domain arrangements in solution. RNA binding shifts this equilibrium towards the open, active domain arrangement depending on the overall RNA binding affinity. The population shift towards the open conformation induced by RNA binding quantitatively correlates with the efficiency of spliceosome assembly on corresponding pre-mRNAs (Mackereth et al Nature, 2011). Notably, the addition of the small U2AF35 subunit preshifts this equilibrium and thereby enhances RNA binding of U2AF65 also to weaker RNA ligands (von Voithenberg et al, PNAS 2016). We also found that a linker preceding one of the RRM domains in U2AF65 has an autoinhibitory role to proofread RNA binding. For a set of paralogous multidomain alternative splicing factor RBPs we show that distinct contributions and dynamic domain arrangements mediate RNA binding specificity and enagement of different cellular pre-RNAs. These results indicate that distinct molecular mechanisms of dynamic RNA recognition by multidomain RNA binding proteins represent important features for the regulation of gene expression.

References

1 Voith von Voithenberg L, Sanchez-Rico C, Kang HS, Madl T, Zanier K, Barth A, Warner LR, Sattler M, and Lamb DC Recognition of the 3' splice site RNA by the U2AF heterodimer involves a dynamic population shift. (2016) Proc Natl Acad Sci U S A 113, E7169-E7175. doi: 10.1073/pnas.1605873113

2 Huang JR, Warner LR, Sanchez C, Gabel F, Madl T, Mackereth CD, Sattler M, Blackledge M. Transient Electrostatic Interactions Dominate the Conformational Equilibrium Sampled

by Multidomain Splicing Factor U2AF65: A Combined NMR and SAXS Study. (2014) J Am Chem Soc. 136:7068-76. doi:10.1021/ja502030n

3 Mackereth CD, Madl T, Bonnal S, Simon B, Zanier K, Gasch A, Rybin V, Valcarcel J, and Sattler M. Multi-domain conformational selection underlies pre-mRNA splicing regulation by U2AF. (2011) Nature 475, 408-11. doi:10.1038/nature10171

4.24 ANRIL and STAIR18 – two long non-coding RNAs with atypical features

Peter F. Stadler (Universität Leipzig, DE)

Joint work of the groups of Daniel Teupser, Friedemann Horn, Steve Hoffmann, Peter F. Stadler
Main reference Lesca M. Holdt, Steve Hoffmann, Kristina Sass, David Langenberger, Markus Scholz, Knut Krohn, Knut Finstermeier, Anika Stahringer, Wolfgang Wilfert, Frank Beutner, Stephan Gielen, Gerhard Schuler, Gabor Gäbel, Hendrik Bergert, Ingo Bechmann, Peter F. Stadler, Joachim Thiery, Daniel Teupser: "Alu Elements in ANRIL Non-Coding RNA at Chromosome 9p21 Modulate Atherogenic Cell Functions through Trans-Regulation of Gene Networks", PLOS Genetics, Vol. 9(7), pp. 1–12, Public Library of Science, 2013.
URL http://dx.doi.org/10.1371/journal.pgen.1003588
Main reference Stefanie Binder, Nadine Hösler, Diana Riedel, Ivonne Zipfel, Tilo Buschmann, Christoph Kämpf, Kristin Reiche, Renate Burger, Martin Gramatzki, Jörg Hackermüller, Peter F. Stadler, Friedemann Horn: "STAT3-induced long noncoding RNAs in multiple myeloma cells display different properties in cancer", Scientific Reports, Vol. 7(1), p. 7976, Springer Nature, 2017.
URL http://dx.doi.org/10.1038/s41598-017-08348-5

ANRIL has been discovered by GWAS studies and is known to function via polycomb-based gene silencing. The "gene" comprises a plethora of isoforms with discernible effects on a large number of targets. Interestingly, a Alu-DEIN element is associated in the targetting. An interesting facet of the story is that the standard analysis of the ChIP-seq data completely miss the targeting specificity, which requires the use multi-mapping reads to account for the repetitive element right in the functionally interacting site. STAR18 is a very long, inefficiently spliced transcript located in the largest human-specific duplicated region of the genome. Dependent on the the transcription factor STAT3 for its transcription, it interacts with STAT3 in manifold ways in an autoregulatory loop.

4.25 Structure, evolution and targeting of long non coding RNAs

Gabriele Varani (University of Washington – Seattle, US)

The molecular basis for activity of long non coding RNAs and their structure-function relationship remain to be established. A commonly stated but unproven hypothesis is that their secondary structures are conserved and functional, despite low levels of primary sequence conservation. We have discovered a complex secondary structure in the functional core of Cyrano, a rare lincRNAs conserved over significant evolutionary distances, at the center of which is a strikingly conserved cloverleaf structure maintained over >420 million years of evolution. This structure provides protein interaction sites and is recognized by miR-7 in a non-canonical fashion. Structures within ncRNAs are functional, as we demonstrated for the promoter associated transcript controlling expression of the tumor suppressor E-cadherin,

where a single SNP that alters RNA structure affects differential recruitment of epigenetic enzymes to the promoter and leads to different outcome in cancer patients. These RNA structures provide unexploited targets for intervention using peptidic and small molecule chemistry, as we illustrate with targeting of the microRNA precursor coding for the oncogenic miR-21.

4.26 Large-scale prediction and analysis of RBPs

Gene Yeo (UC – San Diego, US)

I will present my lab's recent efforts in expanding the repertoire of RNA binding proteins (RBPs) using available affinity precipitation followed by mass spectrometry datasets. Our algorithm SONAR can be applied to yeast, fly and human protein-protein interaction datasets to identify novel RBPs. I will also discuss our improvements to CLIP technologies (enhanced CLIP) and large-scale CRISPR-tagging of RBPs in human cell lines.

Participants

- Frédéric Allain
ETH Zürich, CH
- Rolf Backofen
Universität Freiburg, DE
- Benedikt Beckman
HU Berlin, DE
- Janusz Bujnicki
Int. Inst. of Molecular & Cell Biology – Warsaw, PL
- Tomaž Curk
University of Ljubljana, SI
- Philipp Drewe
Max-Delbrück-Centrum – Berlin, DE
- André Gerber
University of Surrey – Guildford, GB
- Mahsa Ghanbari
Max-Delbrück-Centrum – Berlin, DE
- Tim Hughes
University of Toronto, CA
- Eckhard Jankowsky
Case Western Reserve University – Cleveland, US
- Hilal Kazan
Antalya International University, TR
- Grzegorz Kudla
University of Edinburgh, GB
- Markus Landthaler
Max-Delbrück-Centrum – Berlin, DE
- Martin Lewinski
Universität Bielefeld, DE
- Yael Mandel-Gutfreund
Technion – Haifa, IL
- Hannah Margalit
The Hebrew University of Jerusalem, IL
- Annalisa Marsico
MPI für Molekulare Genetik – Berlin, DE
- Daniel Maticzka
Universität Freiburg, DE
- Irmtraud Meyer
Max-Delbrück-Centrum – Berlin, DE
- Quaid Morris
University of Toronto, CA
- Uwe Ohler
Max-Delbrück-Centrum – Berlin, DE
- Teresa Przytycka
National Center for Biotechnology – Bethesda, US
- Andres Ramos
University College London, GB
- Guido Sanguinetti
University of Edinburgh, GB
- Michael Sattler
TU München, DE
- Peter F. Stadler
Universität Leipzig, DE
- Gabriele Varani
University of Washington – Seattle, US
- Gene Yeo
UC – San Diego, US

Report from Dagstuhl Seminar 17261

Voting: Beyond Simple Majorities and Single-Winner Elections

Edited by

Dorothea Baumeister[1], Piotr Faliszewski[2], Annick Laruelle[3], and Toby Walsh[4]

1 **Heinrich-Heine-Universität Düsseldorf, DE,** baumeister@cs.uni-duesseldorf.de
2 **AGH University of Science and Technology – Krakow, PL,** faliszew@agh.edu.pl
3 **University of the Basque Country – Bilbao, ES,** annick.laruelle@ehu.es
4 **TU Berlin, DE,** Toby.Walsh@data61.csiro.au

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 17261 "Voting: Beyond simple majorities and single-winner elections". The seminar featured five survey talks, a series of classic scientific presentations, working group discussions, open problems sessions (with the first one used to establish working groups and the last one to present their results). The seminar was mostly focused on multiwinner elections (from discussions of their algorithmic properties to political-science considerations), but the topics of real-life voting experiments and strategic behavior received attention as well.

Seminar June 25–30, 2017 – http://www.dagstuhl.de/17261
1998 ACM Subject Classification I.2.11 Distributed Artificial Intelligence
Keywords and phrases artificial intelligence, collective decision making, computational social choice, multi agent systems, preference aggregation, preference elicitation and preference learning, voting
Digital Object Identifier 10.4230/DagRep.7.6.109
Edited in cooperation with Ann-Kathrin Selker, Lisa Rey

1 Executive Summary

Dorothea Baumeister
Piotr Faliszewski
Annick Laruelle

Computational social choice is an interdisciplinary field of research, focused on the issue of aggregating preferences of agents–perhaps self-interested and strategic – and providing them with joint decisions. Computational social choice combines the tools and approaches of social choice theory, computer science (with particular focus on artificial intelligence and theoretical computer science), economics, political science, and operations research. The distinctive feature of computational social choice–as opposed to the classic social choice theory – is that computational considerations (e.g., efficiency of computing outcomes of the preference aggregation processes) are given significant attention. Further, researchers

working on computational social choice often study virtual elections where either people vote through electronic means (such as in Doodle polls for scheduling of meetings) or the elections are used as a tool and the votes are derived automatically in some way (e.g., voting can be used as a selection procedure in a genetic algorithm). Nonetheless, the two research areas are deeply connected and there is significant interaction between them.

One of the most classic problems studied within (computational) social choice regards conducting a single-winner election. For example, consider the situation where members of some society wish to choose their president. They collect the set of candidates (who usually have to register well ahead of time, typically by gathering necessary popular support), then the candidates run their campaigns, argue who would be the best president, air commercials, etc. Eventually, the voters form their preferences (either they simply decide who would be the best president, or they form rankings of the candidates, or they decide on a set of acceptable presidents, depending on the voting rule used) and, on the election day, they cast their votes. In the end, an electoral commission gathers the votes and applies an agreed upon voting rule to decide who would be the next president.

Due to the fantastic progress in social choice (since the middle of the twentieth century) and in computational social choice (over the last fifteen years or so), essentially all the stages of the above-described process are quite well understood. However, In the modern world – especially in the era of ubiquitous use of social media – it appears that there is a great range of preference aggregation settings where the classic approach falls short. For example, consider a situation where a company wants to hire a team of specialists. There may be quite a large number of possible candidates to employ (as opposed to the few candidates in a typical presidential election), each of the candidates may have quite different skills and abilities, various employees may either complement each other or be counter-productive if teamed up (as opposed to our presidential election, where we pick a single person for the whole task). Finally, the recruiting committee typically consists of just a few people (few voters, as opposed to the millions of people voting for presidents), but their preferences might have a very involved structure (for example, if we hire a specialist in X then we also need a specialist in Y, but otherwise a specialist in Z would suffice; on top of that, each member of the recruitment committee may judge candidates' abilities differently).

The goal of the seminar was to bring together researchers who work on various aspects of aggregation problems that go beyond the classic high-stake, rarely conducted, single-winner elections, and to discuss the following issues:

1. Scenarios with multiple winners and/or settings where each aggregation outcome may consist of separate entities (e.g., multi-winner elections and single-winner elections in combinatorial domains).
2. Various non-typical ways of expressing preferences, going from (variants of) non-binary preferences to settings where the agents can express complex statements, including conditional ones (e.g., CP-nets).

The seminar also dealt with some real-world applications, such the question of drawing constituency boundaries in the United States of America or the choice of the voting rule in EU Council of Ministers.

The seminar brought together 42 researchers from 14 countries, working in artificial intelligence, theoretical computer science, mathematics, economics, social choice, and political science. Discussions regarding the new challenges in the area of preference aggregation should fertilize research in all these areas.

The technical program of the seminar was structured over five working groups. On the first day the participants were invited to give 5-minute presentations of research topics

that they found interesting and, based on those presentations, the organizers suggested five working groups:

- Working Group 1: Voting Experiments.
- Working Group 2: Understanding Diversity in Multiwinner Elections.
- Working Group 3: Aggregation Procedures with Nonstandard Input and Output Types.
- Working Group 4: Voting in Larger Contexts.
- Working Group 5: Proportionality in Multiwinner Elections.

The seminar attendees accepted these groups and each chose one to participate in the discussions. During the seminar each group met twice for extended discussions. Also, one afternoon was free for unstructured discussions (of which many were used for in-depth discussions of topics initiated during the working group discussions). Finally, on the last day of the seminar representatives of each group presented the results of their discussions (which ranged from making actual technical contributions to presenting research agendas for future work[1]). The working groups were supported by 21 regular scientific presentations, and extended presentation of the real-life experiments regarding French presidential elections, and 6 survey talks.

The seminar acknowledged that there are new, exciting research topics regarding computational social choice. In particular, multiwinner elections were represented very prominently and it became clear that work on them has only started. We hope and believe that one of the effects of the seminar was convincing many more people (also outside of the core computational social choice community) that studying more general forms of elections (such as multiwinner elections, or elections with nonstandard input formats) is important and promising.

Given the personal feedback we received, we believe that the participants were very happy with the setting of working groups (from the process of forming them based on surveyed interest, through actual meetings, to presentation of results). We also receive very strong, positive feedback regarding leaving one afternoon unstructured, for the participants to use as they preferred. We have seen many ad-hoc discussions, meetings of coauthors, and problem-solving sessions. Indeed, it seems that for seminars with well-established communities, such an unstructured afternoon is far more useful than the traditional excursion (which, on the other hand, seems to be very effective for not as well-established communities).

We are very grateful to all the participants for their contributions, ideas, and discussions, which made this seminar truly enjoyable. We would also like to thank the Schloss Dagstuhl team for their support and excellent organization and patience.

[1] The organizers are already aware of three independent research projects that follow the agendas presented at the seminar.

2 Table of Contents

Working groups

3 Overview of Talks

3.1 Proportional Representation in Approval-based Committee Voting and Beyond

Haris Aziz (Data61 – Sydney, AU)

Proportional representation (PR) is one of the central principles in voting. Elegant rules with compelling PR axiomatic properties have the potential to be adopted for several important collective decision making settings. I survey some recent ideas and results on axioms and rules for proportional representation in committee voting.

References

1 H. Aziz and S. Huang. *Computational complexity of testing proportional justified representation.* Technical Report arXiv:1612.06476. arXiv.org, 2016.
2 H. Aziz and S. Huang. *A polynomial-time algorithm to achieve extended justified representation.* Technical Report arXiv:1703.10415. arXiv.org, 2017.
3 H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. *Justified representation in approval-based committee voting.* In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, pp. 784–790. AAAI Press, 2015.
4 H. Aziz, S. Gaspers, J. Gudmundsson, S. Mackenzie, N. Mattei, and T. Walsh. *Computational aspects of multi-winner approval voting.* In *Proceedings of the 14th International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS)*, pp. 107–115. IFAAMAS, 2015.
5 H. Aziz, M. Brill, V. Conitzer, E. Elkind, R. Freeman, and T. Walsh. *Justified representation in approval-based committee voting. Social Choice and Welfare* 48(2):461–485, 2017.
6 H. Aziz, E. Elkind, P. Faliszewski, M. Lackner, and P. Skowron. *The Condorcet principle for multiwinner elections: From shortlisting to proportionality.* Technical Report arXiv:1701.08023. arXiv.org, 2017.
7 M. Balinski and H. P. Young. *Fair Representation: Meeting the Ideal of One Man, One Vote.* Brookings Institution Press, 2nd edition, 2001.
8 M. Brill, R. Freeman, S. Janson, and M. Lackner. *Phragmén's voting methods and justified representation.* In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pp. 406–413. AAAI Press, 2017.
9 M. Brill, J.-F. Laslier, and P. Skowron. *Multiwinner approval rules as apportionment methods.* In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pp. 414–420. AAAI Press, 2017.
10 P. Faliszewski, P. Skowron, A. Slinko, and N. Talmon. *Multiwinner voting: A new challenge for social choice theory.* In *Trends in Computational Social Choice* (U. Endriss, ed), Chapter 2. 2017. To appear.
11 S. Janson. *Phragmén's and Thiele's election methods.* Technical Report arXiv:1611.08826 [math.HO]. arXiv.org, 2016.
12 D. M. Kilgour. *Approval balloting for multi-winner elections.* In *Handbook on Approval Voting* (J.-F. Laslier and M. R. Sanver, eds), Chapter 6, pp. 105–124. Springer, 2010.
13 J-L. Petit and E. Térouanne. *A theory of proportional representation. SIAM Journal on Discrete Mathematics* 3(1):116–139, 1990.
14 F. Pukelsheim. *Proportional Representation: Apportionment Methods and Their Applications.* Springer, 2014.

15 L. Sánchez-Fernández, N. Fernández, and L. A. Fisteus. *Fully open extensions to the D'Hondt method.* Technical Report arXiv:1609.05370 [cs.GT]. arXiv.org, 2016.

16 L. Sánchez-Fernández, E. Elkind, and M. Lackner. *Committees providing EJR can be computed efficiently.* Technical Report arXiv:1704.00356. arXiv.org, 2017.

17 L. Sánchez-Fernández, E. Elkind, M. Lackner, N. Fernández, J. A. Fisteus, P. Basanta Val, and P. Skowron. *Proportional justified representation.* In *Proceedings of the 31st AAAI Conference on Artificial Intelligence (AAAI)*, pp. 670–676. AAAI Press, 2017.

18 P. Skowron, M. Lackner, E. Elkind, and L. Sánchez-Fernández. *Optimal average satisfaction and extended justified representation in polynomial time.* Technical Report arXiv:1704.00293. arXiv.org, 2017.

19 P. Skowron, P. Faliszewski, and J. Lang. *Finding a collective set of items: From proportional multirepresentation to group recommendation. Artificial Intelligence* 241:191–216, 2016.

20 T. N. Thiele. *Om flerfoldsvalg. Oversigt over det Kongelige Danske Videnskabernes Selskabs Forhandlinger*, pp. 415–441. 1895.

3.2 Voter Dissatisfaction in Committee Elections

Dorothea Baumeister (Heinrich-Heine-Universität Düsseldorf, DE)

Joint work of Toni Böhnlein, Lisa Rey, Oliver Schaudt, Ann-Kathrin Selker
Main reference Dorothea Baumeister, Toni Böhnlein, Lisa Rey, Oliver Schaudt, Ann-Kathrin Selker: "Minisum and Minimax Committee Election Rules for General Preference Types", in Proc. of the ECAI 2016 – 22nd European Conference on Artificial Intelligence, 29 August–2 September 2016, The Hague, The Netherlands, Frontiers in Artificial Intelligence and Applications, Vol. 285, pp. 1656–1657, IOS Press, 2016.
URL http://dx.doi.org/10.3233/978-1-61499-672-9-1656

In committee elections it is often assumed that voters only (dis)approve of each candidate or that they rank all candidates, as it is common for single-winner elections. We suggest an intermediate approach, where the voters rank the candidates into a fixed number of groups. This allows more diverse votes than approval votes, but leaves more freedom than in a linear order. A committee is then elected by applying the minisum or minimax approach to minimize the voters' dissatisfaction. We study the axiomatic properties of these committee election rules as well as the complexity of winner determination and show fixed-parameter tractability for our minimax rules.

3.3 Multiwinner Approval Voting: An Apportionment Approach

Steven J. Brams (New York University, US)

Joint work of Steven J. Brams, D. Marc Kilgour, Richard F. Potthoff
Main reference Steven J. Brams, D. Marc Kilgour, Richard F. Potthoff: "Multiwinner Approval Voting: An Apportionment Approach", Preprint, 2017.

We extend approval voting so as to elect multiple candidates, who may be either individuals or members of a political party, in rough proportion to their approval in the electorate. We analyze two divisor methods of apportionment, first proposed by Jefferson and Webster, that iteratively depreciate the approval votes of voters who have one or more of their approved candidates already elected. We compare the usual sequential version of these methods with

a nonsequential version, which is computationally complex but feasible for many elections. Whereas Webster apportionments tend to be more representative of the electorate than those of Jefferson, the latter, whose equally spaced vote thresholds for winning seats duplicate those of cumulative voting in 2-party elections, is more even-handed or balanced.

3.4 On the Tradeoff Between Efficiency and Strategyproofness in Probabilistic Social Choice

Felix Brandt (TU München, DE)

Joint work of Haris Aziz, Florian Brandl, Felix Brandt, Markus Brill, Christian Geist
Main reference Haris Aziz, Florian Brandl, Felix Brandt, Markus Brill: "On the tradeoff between efficiency and strategyproofness", Working paper, 2017.
URL http://dss.in.tum.de/files/brandt-research/stratlott_journal.pdf

Two fundamental notions in microeconomic theory are efficiency–no agent can be made better off without making another one worse off–and strategyproofness–no agent can obtain a more preferred outcome by misrepresenting his preferences. The conflict between these two notions is already apparent in Gibbard and Satterthwaite's seminal theorem, which states that the only single-valued social choice functions that satisfy non-imposition–a weakening of efficiency–and strategyproofness are dictatorships. This talk will be concerned with efficiency and strategyproofness in the context of social decision schemes, i.e., functions that map a preference profile to a probability distribution (or lottery) over a fixed set of alternatives. Depending on how preferences over alternatives are extended to preferences over lotteries, there are varying degrees of efficiency and strategyproofness. I will dicuss positive results for random serial dictatorship and maximal lotteries as well as a number of impossibility theorems, one of which was recently shown using computer-aided solving techniques.

3.5 On Coalitional Manipulation for Multiwinner Elections: Shortlisting

Robert Bredereck (TU Berlin, DE)

Joint work of Robert Bredereck, Andrzej Kaczmarczyk, Rolf Niedermeier
Main reference Robert Bredereck, Andrzej Kaczmarczyk, Rolf Niedermeier: "On Coalitional Manipulation for Multiwinner Elections: Shortlisting", in Proc. of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017, pp. 887–893, ijcai.org, 2017.
URL http://dx.doi.org/10.24963/ijcai.2017/123

Shortlisting of candidates – selecting a group of "best" candidates – is a special case of multiwinner elections. We provide the first in-depth study of the computational complexity of strategic voting for shortlisting based on the most natural and simple voting rule in this scenario, ℓ-Bloc (every voter approves ℓ candidates). In particular, we investigate the influence of several tie-breaking mechanisms (e.g. pessimistic versus optimistic) and group evaluation functions (e.g. egalitarian versus utilitarian) and conclude that in an egalitarian setting strategic voting may indeed be computationally intractable regardless of the tie-breaking rule. We provide a fairly comprehensive picture of the computational complexity landscape of this neglected scenario.

3.6 Lars Edvard Phragmén: The Man, the Myth, the Legend

Markus Brill (TU Berlin, DE)

I survey some contributions of Swedish mathematician Lars Edvard Phragmén (1863–1937) to the theory of multi-winner elections. Examples are taken from Svante Janson's arXiv paper titled "Phragmén's and Thiele's election methods" (arXiv:1611.08826 [math.HO]).

3.7 Proportional Rankings

Edith Elkind (University of Oxford, GB)

Joint work of Edith Elkind, Piotr Skowron, Martin Lackner, Markus Brill, Dominik Peters

We extend the principle of proportional representation to rankings: given approval preferences, we aim to generate aggregate rankings so that cohesive groups of voters are represented proportionally in each initial segment of the ranking. Such rankings are desirable in situations where initial segments of different lengths may be relevant, e.g., in recommender systems, for hiring decisions, or for the presentation of competing proposals on a liquid democracy platform. We define what it means for rankings to be proportional, provide bounds for well-known aggregation rules, and experimentally evaluate the performance of these rules.

3.8 How Should We Model Incomplete Information in Strategic Voting?

Ulle Endriss (University of Amsterdam, NL)

The Gibbard-Satterthwaite Theorem suggests that every reasonable voting rule is subject to strategic manipulation. But this result relies on the arguably unrealistic assumption that the manipulator has full information about the voting intentions of everybody else. Maybe, if we drop this assumption, better results are attainable? In this talk, I have reviewed a small number of contributions to the literature that are relevant to this kind of research agenda [1, 2, 3, 4, 5, 6]. I have also asked (but not answered conclusively) the question of what constitutes a good model of incomplete information in this context.

References

1 M. J. Osborne and A. Rubinstein. *Sampling Equilibrium, with an Application to Strategic Voting. Games and Economic Behavior* 45(2):434–441, 2003.
2 S. Chopra, E. Pacuit, and R. Parikh. *Knowledge-Theoretic Properties of Strategic Voting.* In *Proceedings of the 9th European Conference on Logics in Artificial Intelligence (JELIA)*, pp. 18–30. 2004.

3 V. Conitzer, T. Walsh, and L. Xia. *Dominating Manipulations in Voting with Partial Information.* In *Proceedings of the 25th AAAI Conference on Artificial Intelligence*, pp. 638–643. 2011.

4 A. Reijngoud and U. Endriss. *Voter Response to Iterated Poll Information.* In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, pp. 635–644. 2012.

5 R. Meir, O. Lev, and J. S. Rosenschein. *A Local-Dominance Theory of Voting Equilibria.* In *Proceedings of the 15th ACM Conference on Economics and Computation (EC)*, pp. 313–330. 2014.

6 U. Endriss, S. Obraztsova, M. Polukarov, and J. S. Rosenschein. *Strategic Voting with Incomplete Information.* In *Proceedings of the 25th International Joint Conference on Artificial Intelligence (IJCAI)*, pp. 236–244. 2016.

3.9 Robustness Among Multiwinner Voting Rules

Piotr Faliszewski (AGH University of Science and Technology – Krakow, PL)

Joint work of Piotr Faliszewski, Andrzej Kaczmarczyk, Robert Bredereck, Rolf Niedermeier, Piotr Skowron, Nimrod Talmon

We investigate how robust are results of committee elections to small changes in the input reference orders, depending on the voting rules used. We find that for typical rules the effect of making a single swap of adjacent candidates in a single preference order is either that (1) at most one committee member can be replaced, or (2) it is possible that the whole committee can be replaced. We also show that the problem of computing the smallest number of swaps that lead to changing the election outcome is typically NP-hard, but there are natural FPT algorithms. Finally, for a number of rules we assess experimentally the average number of random swaps necessary to change the election result.

3.10 The Complexity of Campaigning

Judy Goldsmith (University of Kentucky – Lexington, US)

Joint work of Judy Goldsmith, Cory Siler, Luke Harold Miles
Main reference Cory Siler, Luke Harold Miles, Judy Goldsmith: "The Complexity of Campaigning", in Proc. of the Algorithmic Decision Theory - 5th International Conference, ADT 2017, Luxembourg, Luxembourg, October 25-27, 2017, Proceedings, Lecture Notes in Computer Science, Vol. 10576, pp. 153–165, Springer, 2017.
URL http://dx.doi.org/10.1007/978-3-319-67504-6_11

In "The Logic of Campaigning" [1], Dean and Parikh consider a candidate making campaign statements to appeal to the voters. They model these statements as Boolean formulas over variables that represent stances on the issues, and study optimal candidate strategies under three proposed models of voter preferences based on the assignments that satisfy these formulas. We prove that voter utility evaluation is computationally hard under these preference models (in one case, $\#P$-hard), along with certain problems related to candidate strategic reasoning. Our results raise questions about the desirable characteristics of a voter preference model and to what extent a polynomial-time-evaluable function can capture them.

References

1 Walter Dean and Rohit Parikh. *The logic of campaigning*. In *Proceedings of the 4th Indian Conference on Logic and its Applications*, pp. 38–49. Springer, 2011.

3.11 Apportionment as a Sequential Portfolio Allocation Method: Application to Northern Ireland

Bernard Grofman (University of California – Irvine, US)

Joint work of Brendan O'Leary, Jorgen Elklit, Bernard Grofman

There are three elements of the usual cabinet coalition bargaining game that need to be determined:

1. Which parties will be in government?
2. What (perhaps minimal/incomplete) policy platform can these parties agree upon?
3. How will cabinet portfolios be allocated? (The set of (junior) ministries may be flexible)

The process in Northern Ireland is quite different: on the one hand, incorporating both a diversity requirement (re Unionist and Nationalist representation in the governing coalition) and, on the other hand, making use of a proportional allocation algorithm to allocate ministerial portfolios in a fashion that does not actually require the agreement of the parties on which parties get which portfolio in any given instance, but only an agreement in advance on a quite abstract but always determinate rule as to the sequence of choice, with ministries are basically fixed in advance.

1. The leading unionist and the leading nationalist party will always be in government as the executive and deputy executive.
2. Cabinet portfolios will be allocated in accord with party vote shares to those parties that agree to be in government using the D'Hondt apportionment rule as a sequential allocation process, i.e., the largest party gets to be First Minister, the next largest Deputy First Minister and the next ministry is chosen by whichever party has the largest unallocated divisor quotient. Etc.
3. There is no overall agreement on government policy; roughly speaking, the party that controls a ministry controls the policy within the jurisdiction of that ministry.

We consider the properties of such a sequential allocation rule

3.12 Graph Theory Perspectives on Committee Representativeness

Bernard Grofman (University of California – Irvine, US)

Joint work of Scott L. Feld, Bernard Grofman

After briefly reviewing several centuries of work on theories of representation and offering a conceptual synthesis, we offer a new approach to determining the overall "representativeness"of a committee or legislature that combines representativeness with diversity. The basic ideas, presented for the first time, we believe, in Grofman and Feld (1988, unpublished) are similar

in spirit to that of Chamberlin and Courant (1978), but draw on two seminal ideas in graph theory, the idea of a k-cover and the idea of a k-basis (Harary, Cartwright and Norman, 1965). Like Chamberlin and Courant, we abstract away from the mechanism by which representative are chosen to look at voter preferences about members of possible committees so as to build up evaluations of committees as a whole from voter evaluations of the individual committee members. We then look at a minimal k-cover which is the equivalent of a minimax version of Chamberlin and Courant. Because this approach yields too many feasible committees, we propose to narrow the selection by imposing a diversity requirement on the set of min-max committees, in terms of the idea of a minimal k-basis, a set whose members are maximally distinct from one another.

3.13 Mathematical and Computational Aspects of Constituency Boundary Drawing in the United States

Bernard Grofman (University of California – Irvine, US)

We consider mathematical and computational aspects of redistricting (constituency boundary drawing) in the United States, with particular attention to issues of partisan and racial gerrymandering and legal challenges to such gerrymandering involving expert witness testimony. We pay particular attention to issues tied to computer analysis, e.g., analysis of districting features such as compactness and computer simulation of all possible plans that satisfy certain legal criteria. We briefly discuss a pending partisan gerrymandering legal challenge to Wisconsin congressional districts that will be heard by the U.S. Supreme Court in Fall, 2017, and we offer a five-pronged test for unconstitutional partisan gerrymandering. We also look at how U.S. redistricting practices differ from those elsewhere in the democratic world.

3.14 Multi-Winner Elections: A Review

Marc Kilgour (Wilfrid Laurier University, CA)

A Multi-Winner Election is an election that is intended to select a subset of the candidates. What properties should that subset possess? In most multi-winner elections, it is desirable that the voters support strongly not only each individual elected candidate, but also the elected subset as a whole. In multi-winner elections, the desiderata of level of support and breadth of support are often in conflict, amply illustrated by the well-known "Tyranny of the Majority" problem.

Multi-winner elections can be classified according to whether the number of winners is fixed in advance (FNW), or variable – where the voters determine not only who wins, but also how many win (VNW). Common multi-winner elections are reviewed in light of this classification and the often-conflicting objectives of individual vs. group support.

Prominent among the reasons why it is difficult to find a good multi-winner election procedure is that ballots describe only individual candidates, and do not record any assessment

of subsets of candidates. A voter simply cannot express nuances of preference, such as candidates who work well, or poorly, together.

One strategy to ensure diversity in the winning subset is to specify which subsets are admissible, or permitted to win the election. Another method is to select a procedure for counting the ballots that measures the level of support for individual candidates while penalizing candidates whose support lacks breadth.

Ballots available for multi-winner elections are reviewed, and the issue of strategy raised. Even for a sincere voter with well-defined utilities and linear preferences over subsets, it is not obvious how to fill out some ballots. In addition, even for an FNW election with sincere voters, the winning subset can depend on the particular ballot form chosen.

Some simple observations about the less-known category of VNW (Variable Number of Winners) elections are offered, including the work of Duddy and Piggins (and Zwicker) on aggregation. A new property is proposed that eliminates the possibility of ties. It is appropriate for a VNW election, including elevation to a hall of fame and short-listing, where the only objective is to identify the candidates with sufficient individual support.

Multi-winner elections are an important area of rapid scientific progress. Moreover, improving electronic capabilities have made multi-winner elections more common. If they are matched by improvements in procedures, there will be new options for efficient opinion-gathering and accurate expression and understanding of voters' opinions.

References

1 D. M. Kilgour. "Approval Balloting for Multi-Winner Elections", In Handbook on Approval Voting (J.-F. Laslier and M. R. Sanver, eds), Chapter 6, pp. 105–124. Springer, Heidelberg, 2010.
2 D. M. Kilgour and E. Marshall. "pproval Balloting for Fixed-Size Committees", In Electoral Systems: Paradoxes, Assumptions, and Procedures (M. Machover and D. Felsenthal, eds), pp. 305–326. Springer, Heidelberg, 2012.
3 S. J. Brams and D. M. Kilgour. "Satisfaction Approval Voting", In Voting Power and Procedures: Essays in Honour of Dan Felsenthal and Moshe Machover (R. V. Fara, D. Leech, and M. Salles, eds), pp. 323–346. Springer, Cham, 2014.
4 D. M. Kilgour. "Approval Elections with a Variable Number of Winners", Theory and Decision, 81(2):199–211, 2016.
5 S. J. Brams, D. M. Kilgour, and R. Potthoff. "Multiwinner Approval Voting: An Apportionment Approach", 2017. Preprint.

3.15 Consistent Approval-Based Multi-Winner Rules

Martin Lackner (University of Oxford, GB)

Joint work of Martin Lackner, Piotr Skowron
Main reference Martin Lackner, Piotr Skowron: "Consistent Approval-Based Multi-Winner Rules", in CoRR, Vol. abs/1704.02453, 2017.
URL http://arxiv.org/abs/1704.02453

The goal of this talk is to provide an overview of recent work on approval-based multi-winner rules and discuss research directions. Approval-based multi-winner rules are voting rules that select a fixed-size group of candidates based on approval ballots. Such rules are applicable to a wide range of scenarios concerning group decision making and consequently may aim for widely diverging objectives.

In this talk, I introduce the class of counting rules and discuss axiomatic characterizations of rules in this class. In particular, I present axiomatic characterizations of three import consistent multi-winner rules: Proportional Approval Voting, Multi-Winner Approval Voting and Approval Chamberlin–Courant. These results demonstrate the variety of multi-winner rules and the different, orthogonal goals that multi-winner voting rules may pursue.

3.16 Voting in Combinatorial Domains

Jérôme Lang (University Paris-Dauphine, FR)

This talk addresses preference aggregation and voting on domains which are the Cartesian product (or sometimes, a subset of the Cartesian product) of finite domain values, each corresponding to an issue, a variable, or an attribute. Practical examples of voting on such combinatorial domains are multiple referenda, group configuration, and committee elections. The talk gives a structured survey of classes of methods for voting in combinatorial domains and relates them to classes of multiwinner elections rules. The talk mostly follows the handbook chapter J. Lang and L. Xia, "Voting in Combinatorial Domains", Chapter 9, Handbook of Computational Social Choice.

3.17 Voting Experiments: First Lessons from the 2017 French Presidential Election

Annick Laruelle (University of the Basque Country – Bilbao, ES), Sylvain Bouveret (LIG – Grenoble, FR), and Vincent Merlin (Caen University, FR)

Joint work of Annick Laruelle, Sylvain Bouveret, Vincent Merlin, Antoinette Baujard, Renaud Blanch, Herrade Igersheim, Jean-François Laslier, Isabelle Lebon

During the French presidential election in April 2017, a voting experiment was carried online (on the website https://vote.imag.fr/) and at five polling. Participants were proposed to test alternative voting procedure such as approval voting, evaluative voting, STV. The aims of this talk are to the reasons for realizing in situ experiments, to describe the experimental protocol and to give a few preliminary results and feedback. Above all, the aim is to raise discussions about what voting theory can learn from this kind of experiments.

3.18 Multiwinner Voting Rules in Practice: an Efficiency – Inequality Dilemma

Jean-Francois Laslier (Paris School of Economics, FR)

Multiwinner voting rules take as input preferences over candidates and return sets of candidates. These rules could be used for parliamentary election if electoral law would allow panachage across parties. Using survey data, we study six such rules that have been proposed in the literature and illustrate what they could produce using 73 political elections, each of whom has between 6 and 9 parties. We show in particular how these rules differentiate the ones from the others according to Proportionality, Efficiency, and Inequality.

3.19 Strategyproof Peer Selection using Randomization, Partitioning, and Apportionment

Nicholas Mattei (IBM TJ Watson Research Center – Yorktown Heights, US)

Joint work of Nicholas Mattei, Haris Aziz, Omer Lev, Jeffrey S. Rosenschein, Toby Walsh
Main reference Haris Aziz, Omer Lev, Nicholas Mattei, Jeffrey S. Rosenschein, Toby Walsh: "Strategyproof Peer Selection", in CoRR, Vol. abs/1604.03632, 2016.
URL http://arxiv.org/abs/1604.03632

Peer review, evaluation, and selection is a fundamental aspect of modern science. Funding bodies the world over employ experts to review and select the best proposals of those submitted for funding. The problem of peer selection, however, is much more universal: a professional society may want to give a subset of its members awards based on the opinions of all members; an instructor for a MOOC or online course may want to crowdsource grading; or a marketing company may select ideas from group brainstorming sessions based on peer evaluation. We make three fundamental contributions to the study of procedures or mechanisms for peer selection, a specific type of group decision making problem studied in computer science, economics, and political science. First, we propose a novel mechanism that is strategyproof, i.e., agents cannot benefit by reporting insincere valuations. Second, we demonstrate the effectiveness of our mechanism by a comprehensive simulation based comparison with a suite of mechanisms found in the literature. Finally, our mechanism employs a randomized rounding technique that is of independent interest, as it solves the apportionment problem that arises in various settings where discrete resources such as parliamentary representation slots need to be divided proportionally.

3.20 The Optimal Voting Mechanism in the Council of the a Federal Union

Vincent Merlin (Caen University, FR)

Joint work of Michel Le Breton, Dominique Lepelley, Antonin Macé, Vincent Merlin
Main reference M. Le Breton, D. Lepelley, A. Macé, and V. Merlin. "Le Mécanisme Optimal de Vote au Sein du Conseil des Représentant d'un Système Fédéral, L'actualité Economique," 2017. To appear.

The main objective of this survey is to review the literature on the selection of a Council or a Committee in a Federal Union. We will assume that several autonomous authorities (cities, regions, states) have decided to form a federal union. We do not discuss here the reasons why the federation exists, nor the extent of its competencies. We posit that each authority is represented at the federal level by a unique representative. The example we have in mind is the Council of Minister of the EU, though it can easily be applied to other contexts (eg the US Electoral College, a parliament, etc.). In most of the paper, we also assume that the federation has to take only binary decisions. A representative will vote either 'yes' or 'no' on each issue (abstention is not allowed). In this context, we face a mechanism design problem: What is the best voting mechanism, in order to fulfill a given objective? Said differently, given the natural boundaries of the districts/cities/region/states, and the impossibility to organize a direct decision mechanism, what is the best two-tiers voting rule? More precisely, which weight should we award to each representative, and which quota should we use for the decision? As we will see, the answer to this question depends upon the criteria we adopt to judge upon the fairness of the decision, and upon the assumptions we retain to model the voting behavior of the citizens in each jurisdiction.

3.21 The Power, Success and Size of Voter Groups: Borda vs. Condorcet vs. Plurality Rule

Stefan Napel (Universität Bayreuth, DE)

The talk considers collective decisions between more than two options by a fixed number of shareholders or homogenous voter groups. The respective voting weights or sizes of the groups differ. Collective outcomes are then known to be highly sensitive to the adopted collective choice rule. How could one clarify if – and ideally quantify the extent to which – adoption of a particular rule benefits a given group a priori, assuming that voting weights and the number of alternatives are known while the specific alternatives and hence individual preference rankings are still open? Possible answers are complicated by big differences in the numbers of structurally distinct weight configurations across decision rules. For instance, with three groups and three alternatives, there are 51 distinct mappings from the 216 possible strict preference profiles to a single winner which can be supported by weighted Borda voting, but only 6 such equivalence classes of weights for plurality voting and 4 under Copeland's method; or 5255 vs. 34 vs. 9 equivalence classes of weight distributions with four groups. Investigation of the geometry of these equivalence classes is a novel and fascinating subject in itself.

3.22 The k-Kemeny Rule, and an Incompatibility between Proportionality and Strategyproofness

Dominik Peters (University of Oxford, GB)

In this talk, I discuss two topics. First, I describe some interesting complexity questions about the k-Kemeny rule recently introduced by Bill Zwicker; these mainly concern structured preferences and parameterized questions. The k-Kemeny rule selects a k-chotomous ordering that best represents the input preferences. For $k \geq 3$, this rule is NP-hard to evaluate. However, this problem becomes easy if the input preferences are single-peaked, or if they are purely cocyclic. An interesting challenge for future research would be to see whether these tractability results can be extended: is evaluation also easy for single-crossing preferences, or for other notions of structure? Is evaluation easy for almost cocyclic preferences? Further, many natural approximation and fixed-parameter questions are open for this interesting rule.

Second, I presented an impossibility theorem for approval-based multi-winner rules. In this setting, Approval Voting (AV), which returns the k candidates with highest approval scores, is strategyproof in a strong sense. However, AV is not proportional, and if a majority of voters agrees on k candidates, then the preferences of the minority are completely ignored by AV. On the other hand, Proportional Approval Voting (PAV), due to Thiele, satisfies a strong notion of proportionality (namely extended justified representation (EJR), as well as being a d'Hondt extension). One may wonder whether both of these virtues can be satisfied together: is there an approval-based multi-winner rule that is both proportional and strategyproof? Using a technique based on SAT solvers, I show that the answer is no. The incompatibility holds even for very weak notions of proportionality and of strategyproofness.

3.23 Evaluationwise Strategy-proofness

M. Remzi Sanver (University Paris-Dauphine, FR)

Joint work of Bora Erdamar, Shin Sato, M. Remzi Sanver

We consider manipulation of collective decision making rules in a framework where voters not only rank candidates but also evaluate them as "acceptable" or "unacceptable". In this richer informational setting, we adopt a new notion of strategy-proofness, called evaluationwise strategy-proofness, where incentives of manipulation exist if and only if a voter can replace an outcome which he finds unacceptable with an acceptable one. Evaluationwise strategy-proofness is weaker than strategy-proofness. However, we establish the prevalence of a logical incompatibility between evaluationwise strategy-proofness, anonymity and efficiency. On the other hand, we show possibility results when either anonymity or efficiency is weakened.

3.24 Practical Algorithms for Computing STV and Other Multi-Round Voting Rules

Sujoy Sikdar (Rensselaer Polytechnic Institute – Troy, US)

Joint work of Sujoy Sikdar, Chunheng Jiang, Hejun Wang, Lirong Xia, Zhibing Zhao
Main reference Chunheng Jiang, Sujoy Sikdar, Hejun Wang, Lirong Xia, Zhibing Zhao: "Practical Algorithms for Computing STV and Other Multi-Round Voting Rules", in Proc. of 4th Workshop on Exploring Beyond the Worst Case in Computational Social Choice, EXPLORE-2017, 2017.
URL http://www.cs.rpi.edu/ sikdas/papers/practical.pdf

STV is one of the most commonly-used voting rules for group decision-making, especially for political elections. However, the literature is vague about which tie-breaking mechanism should be used to eliminate alternatives. We propose anytime algorithms for computing co-winners under STV, each of which corresponds to the winner under some tie-breaking mechanism. This problem is known as parallel-universes-tiebreaking (PUT)-STV, which is known to be NP-complete to compute. We conduct experiments on synthetic data and Preflib data, and show that standard search algorithms work much better than ILP. We explore improvements to the search algorithm with various features including pruning, reduction, caching and sampling. We use deep learning models to develop priority functions that further improve the performance of the search algorithms.

3.25 Proportionality under Multiwinner Rules

Piotr Skowron (TU Berlin, DE)

Joint work of Piotr Faliszewski, Nimrod Talmon, Piotr Skowron

We consider multiwinner election rules and discuss some of their properties. In particular we discuss the concept of proportional representation of multiwinner election rules. Informally, proportional representation requires that the extent to which a particular preference or opinion is represented in the outcome of elections should be proportional to the frequency with which this preference or opinion occurs within the population. We discuss several approaches which allow to better understand the nature of multiwinner rules with respect to proportional representation.

3.26 Gibbard-Satterthwaite Games for k-Approval Voting Rules

Arkadii Slinko (University of Auckland, NZ)

Joint work of Daniel Hughes, Francesca Rossi, Umberto Grandi, Arkadii Slinko

The Gibbard-Satterthwaite theorem implies the existence of voters, called manipulators, who can change the election outcome in their favour by voting strategically. When a given preference profile admits several such manipulators, voting becomes a game played by these voters, who have to reason strategically about each others' actions. To complicate the game even further, counter-manipulators may try to counteract potential actions of manipulators.

Previously, voting manipulating games have been studied mostly for Plurality rule. We extend this to k-Approval voting rules rules. However, unlike previous studies, we assume that voters are boundedly rational and do not think beyond manipulating or countermanipulating. In this paper we look for conditions on strategy sets of voting manipulation games that guarantee the existence of Nash equilibrium in pure strategies.

3.27 Relation Between Multiwinner Elections and k-Median Problems.

Krzysztof Sornat (University of Wroclaw, PL)

Joint work of Jarosław Byrka, Piotr Skowron, Krzysztof Sornat
Main reference Jaroslaw Byrka, Piotr Skowron, Krzysztof Sornat: "Proportional Approval Voting, Harmonic k-median, and Negative Association", in CoRR, Vol. abs/1704.02183, 2017.
URL http://arxiv.org/abs/1704.02183

In multiwinner elections we want to choose k winners that they somehow represents voters. In k-median problems we want to open k facilities that they somehow serve all clients. Voters has dissatisfaction from choosing particular candidate. Clients has cost of being served by particular facility. It is all the same! I will show examples of voting rules and k-median type of problems that are closely related. Then I will focus on approximation algorithms for such problems. Widely used technique for constructing approximation algorithms is linear programming relaxation with proper rounding procedures. I will present dependent rounding [1] that helps us in opening exactly k facilites (choose exactly k winners) keeping marginal probabilities of opening for each facility. Also I will show our last results on approximating the minimization version of Proportional Approval Voting (PAV) [2].

References

1 A. Srinivasan. *Distributions on level-sets with applications to approximation algorithms.* In *Proceedings of the 42nd IEEE Symposium on Foundations of Computer Science*, pp. 588–597, 2001.
2 J. Byrka, P. Skowron, K. Sornat. *Proportional Approval Voting, Harmonic k-median, and Negative Association.* CoRR abs/1704.02183, 2017.

3.28 Constructing a Democratically-Optimal Budget

Nimrod Talmon (Weizmann Institute – Rehovot, IL)

Joint work of Ehud Shapiro, Nimrod Talmon

The budget is the key means for effecting policy in democracies, yet its preparation is typically an opaque and arcane process, at the end of which the governing body is presented with a take-it-or-leave-it budget proposal. Participatory budgeting is making inroads in municipalities, but is typically limited to a small fraction of the total budget. Current participatory budgeting methods do not scale to entire budgets as they cannot handle quantitative budget items nor hierarchical budget construction. Here we apply the Condorcet principle to participatory budgeting and introduce the concept of a democratically-optimal budget – a budget within a given limit for which there is no other budget within this limit

that is preferred by a majority of the voters. While a democratically-optimal budget does not always exist, we show an algorithm that, given a budget proposal, a budget limit, and a ranking of its items by voters, produces a democratically-optimal budget. Our method can handle quantitative budget items and supports hierarchical budget construction and thus may be applicable to entire budgets.

3.29 Aggregation and Orthogonal Decomposition

William S. Zwicker (Union College – Schenectady, US)

In voting, preferences (submitted as ballots, in any of a variety of formats) are aggregated into an election outcome; in cluster analysis, similarities and differences are aggregated into a partition of objects by type; in judgment aggregation, discordant views as to which logical statements are true and which false are aggregated into a logically consistent collective judgment; in short, there exist a variety of settings wherein information from multiple sources is aggregated into a collective decision. Barthélemy and Monjardet [1981] suggest the context of binary relations as a natural home for a general theory of aggregation, and recommend the median procedure, defined via a metric on relations, as a flexible method that, when restricted to various relational classes, yields a number of useful specific aggregation rules. We reformulate the median procedure in terms of inner product, and argue that two key orthogonal decompositions:

- Explain why certain interesting aggregation rules can be obtained as restrictions of the median procedure, while other, closely related rules cannot
- Suggest modifications to the median procedure to obtain related general procedures that yield some of these "missing" rules
- Identify a key source of disagreement among specific rules: components of information that play an active role in aggregation via some rules, but are irrelevant for others.
- Explain why polynomial time aggregation algorithms exist for some rules, while the aggregation problem is NP-complete for others
- Reveal, for large sub-family of aggregation rules that have been proposed over centuries, that these rules are less ad hoc, and more closely related to one another, than may first appear.

4 Working groups

4.1 Voting Experiments

Sylvain Bouveret (LIG – Grenoble, FR), Christian Klamler (Universität Graz, AT), Annick Laruelle (University of the Basque Country – Bilbao, ES), Jean-Francois Laslier (Paris School of Economics, FR), and Vincent Merlin (Caen University, FR)

Christian Klamler reported the experiment realised during regional elections in Austria. He explained the administrative difficulties met and the findings, in particular on the voting

behavior. This experiment was compared to the French one realized by the other participants. Possible classifications of the candidates were discussed: popular/unpopular/polarized/medium versus exclusive/inclusive. Issues concerning the treatment of the data were raised, in particular the question of the bias, the robustness of the results. Tests to be realized were discussed, for instance the single-peakness of voters' preferences. We end up the discussion with lines for further research: how to do experiments in the future? In which types of elections? In which countries?

4.2 Understanding Diversity in Multiwinner Elections

Robert Bredereck (TU Berlin, DE), William Bailey (University of Kentucky – Lexington, US), Dorothea Baumeister (Heinrich-Heine-Universität Düsseldorf, DE), Ulle Endriss (University of Amsterdam, NL), Piotr Faliszewski (AGH University of Science and Technology – Krakow, PL), Bernard Grofman (University of California – Irvine, US), Martin Lackner (University of Oxford, GB), Sujoy Sikdar (Rensselaer Polytechnic Institute – Troy, US), Arkadii Slinko (University of Auckland, NZ), and Nimrod Talmon (Weizmann Institute – Rehovot, IL)

Diversity in voting is a multi-faceted issue which is widely neglected from the theoretical perspective in the area of Computational Social Choice. The discussion group identified several interesting questions and research directions in which issues of diversity play a role or may arise:

Diversity with respect to preferences. Are preferences in a profile diverse in the sense that they reflect a wide range of opinions? In this context, the main idea is to identify functions that measure the diversity of a preference profile. Natural approaches are based on the Kendall tau rank distance or on the positions of committee members in the voters' preferences with respect to natural multi-winner voting rules. This question has been studied, e.g., by Hashemi and Endriss [3].

Diversity with respect to external attributes. Assuming that committee members have additional attributes, is a chosen committee diverse with respect to these attributes? There are various natural applications including gender balanced committees or program selection processes aiming for a specific ratio of senior and junior researchers. External attributes, however, require the development of new formal models [2].

Diversity of (multi-winner) rules. Does a multi-winner rule translate diverse preferences into a diverse committee? Towards developing formal properties of diverse rules, our discussion group agreed that diverse rules should, to some extent, resist against "cloned candidates". The minimax version of Chamberlin Courant together with utilitarian tie-breaking is a rule which intuitively respects diversity quite well so that characterizations of that rules seem very interesting. Note that "blind rules cannot distinguish colors", i.e., if a rule is not aware of external attributes of candidates, it cannot take these into account *even* if voters' preferences are diverse with respect to these attributes. Furthermore, diversity of voting rules can be visualized in the two-dimensional euclidean domain with histogram plots as introduced by Elkind et al. [1]. Intuitively, diverse voting rules should reflect the general shape of distributions but ignore intensities.

References

1 E. Elkind, P. Faliszewski, J.-F. Laslier, P. Skowron, A. Slinko, and N. Talmon. *What Do Multiwinner Voting Rules Do? An Experiment Over the Two-Dimensional Euclidean Domain*. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence*, pp. 494-501. AAAI Press, 2017.

2 J. Lang and P. Skowron. *Multi-Attribute Proportional Representation*. In *Proceedings of the 30th AAAI Conference on Artificial Intelligence*, pp. 530-536. AAAI Press, 2016.

3 V. Hashemi and U. Endriss. *Measuring Diversity of Preferences in a Group*. In *Proceedings of the 21st European Conference on Artificial Intelligence*, pp. 423-428. IOS Press, 2014.

4.3 Aggregation Procedures with Nonstandard Input and Output Types

Dominik Peters (University of Oxford, GB), Sirin Botan (University of Amsterdam, NL), Jiehua Chen (Ben Gurion University of the Negev – Beer Sheva, IL), Marcin Dziubinski (University of Warsaw, PL), Jérôme Lang (University Paris-Dauphine, FR), Nicholas Mattei (IBM TJ Watson Research Center – Yorktown Heights, US), Jörg Rothe (Heinrich-Heine-Universität Düsseldorf, DE), M. Remzi Sanver (University Paris-Dauphine, FR), Gerhard J. Woeginger (RWTH Aachen, DE), and William S. Zwicker (Union College – Schenectady, US)

In this working group, we considered the design and analysis of voting rules that do not fall into the standard framework of taking as input either approval ballots or linear orders and returning as output either a set of winners or a linear order.

In particular, we made some progress in studying an open problem by Bill Zwicker, but did not settle it. This problem concerns the aggregation of a profile of equivalence relations into a consensus equivalence relation using the median procedure. This problem is known to be NP-hard. The question is whether hardness holds even if the input consists of equivalence relations each of which has at most 2 equivalence classes. An intermediate problem we considered is which weighted graphs can be constructed using profiles of such equivalence relations; here, we found a non-universality: some weighted graphs cannot be constructed using 2 equivalence classes only.

We also discussed some other problem settings. One setting that seemed very natural to us is the following: Consider the problem of selecting a set of several winners, for example papers to admit to a conference, students to admit to a PhD program etc. In these cases, there are several spots available, but the number of winners is not fixed, but endogenously determined. In such cases, a rule should identify a class of "obvious winners", a class of "obvious losers", and a class of candidates which belong to neither group. This latter, middle class should ideally be ranked in order of desirability, so that the deciding body can then choose where to draw the line. The design of good rules of this type is an interesting question for future work.

We also discussed some sensible rules that aggregate linear orders with an additional "bar" that identifies some initial segment of each voter's preference as their approval set. Several ways of using this information seem sensible, and a closer analysis should form an interesting starting point for further analysis.

4.4 Voting in Larger Contexts

Lisa Rey (Heinrich-Heine-Universität Düsseldorf, DE), Judy Goldsmith (University of Kentucky – Lexington, US), Marc Kilgour (Wilfrid Laurier University, CA), Stefan Napel (Universität Bayreuth, DE), Ann-Kathrin Selker (Heinrich-Heine-Universität Düsseldorf, DE), and Kristen Brent Venable (Tulane University – New Orleans, US)

In this group two questions were discussed: "Voting on ethical theories", proposed by K. Brent Venable and "Emotion manipulation", proposed by Judy Goldsmith. The former is based on the fact that a human beings' attitude to morality originates from various influences such as a local or religious law and follows different mentalities such as deantology or utilitariarism. In computer science the question arises how to consider these influences on decision making in for example self-driving cars or automated processes of granting credits or solving problems. While this question includes the fields of multi-criteria optimisation, negotiation, and machine learning, we considered decision making in the context of voting and discussed which properties a voting system should fulfil. On the one hand the voting system should be unanimous and abstain from electing the Condorcet loser. On the other hand the winner does not have to be at least one voter's favourite candidate, which is a desirable property in many other settings. The question can be extended to different context (Should the decision be different if there are children in the car?) or by including weights on the types of consequences. The latter, emotion manipulation, is based on the hypothesis that people vote differently under the influence of emotions. Angry voters tend to act while frightened voters tend to do careful research first. This observation is connected with observations of people with addictions. Amongst others, ideas were to formalise voting over time in this context, to consider to which extend the prevention of such behaviour is still democratic, or how this behaviour can be recognised in the voter's preferences. As an aim we considered a continuity of the voting rule with respect to its input, comparable to Condorcet-consistency, taking statistical models into account. This model should measure the impact of occurring changes. A theoretical analysis of changes on the election outcome could interdependent with a psychological study of the kind of the occurring changes in votes under different moods.

4.5 Proportionality in Multiwinner Elections

Piotr Skowron (TU Berlin, DE), Haris Aziz (Data61 – Sydney, AU), Steven J. Brams (New York University, US), Markus Brill (TU Berlin, DE), Edith Elkind (University of Oxford, GB), Svetlana Obraztsova (Nanyang TU – Singapore, SG), and Krzysztof Sornat (University of Wroclaw, PL)

In a multiwinner election the goal is to select a given-size committee, i.e., a subset of the set of candidates, based on the preferences of the voters over the candidates. Selecting a group of representatives for a given society is a prime example of multiwinner elections. In such cases, one often expects the election rule to follow the principle of proportionality, i.e., that sufficiently large homogeneous groups of voters are represented in a committee and that the numbers of representatives that such groups get in a committee are proportional to their

sizes. The concept of proportionality is well-understood in the context of apportionment, i.e., when each voter votes for a single political party, and when parliamentary seats need to be distributed among political parties. Informally speaking, in the apportionment model the proportionality means that the number of seats assigned to a political party should be roughly proportional to the size of the electorate that voted for such a party [1].

The concept of proportionality of apportionment is helpful for understanding how proportional or disproportional are certain multiwinner rules [2]. Indeed, an instance of an apportionment can be also viewed as an instance of an approval-based multiwinner election, by assuming that a vote for a party is a vote approving all the members of the respective party. When viewed from this perspective, each multiwinner rule induces a certain apportionment method, and looking at which rules induce which apportionment methods can give us some insights into how proportional are these rules.

An interesting model which gives more flexibility to the voters than the classic apportionment setting, but less than the model of multiwinner elections, is when we allow the voters to approve multiple parties. Then, the concept of proportionality can be interpreted in a number of ways which we illustrate using the following example.

▶ **Example 1.** Assume there are four parties, A, B, C, and D, that we are looking for a committee of size $k = 10$, and that there are the following votes:

20 votes: $\{A\}$

60 votes: $\{B, C\}$

20 votes: $\{C, D\}$

The above notation means that 20 voters approve party A, 60 voters approve two parties: B and C, and 20 voters – parties C and D.

One approach to interpret a vote $\{B, C\}$ in the above example is that a voter who casts such a vote is perfectly happy if his vote supports party B or party C. With this approach, we can interpret the above instance as an approval-based multiwinner election, by assuming that a vote for $\{B, C\}$ corresponds to approving all members of B and all members of C. Next, we can apply one of the proportional approval-based multiwinner rules to such obtained instance. For instance, by applying Proportional Approval Voting (PAV) to the instance obtained from the profile from Example 1, we see that a solution which allocates 2 seats to party A, 6 seats to party B and 2 seats to party D is proportional.

One drawback of the previous interpretation is that party C can get no seats even though it is approved by 80% of voters. Thus, this interpretation does not provide fairness to the parties. This motivates another approach: a vote for $\{B, C\}$ shows an intention of distributing the voter's support equally between the two parties. If this is the case, we can say that the total support that parties A, B, C, and D received is equal to, respectively, 20, 30, 40, and 10. Thus, an allocation where these four parties get, respectively, 2, 3, 4, and 1 seat would be proportional.

Summarizing, an exploration of different approaches to designing and understanding systems of apportionment where voters can approve multiple political parties is an interesting line of research.

References

1 M. Balinski and H. P. Young. *Fair Representation: Meeting the Ideal of One Man, One Vote*. Yale University Press, 1982. (2nd Edition [with identical pagination], Brookings Institution Press, 2001).

2 M. Brill, J. Laslier, and P. Skowron. *Multiwinner approval rules as apportionment methods*. In ţextlsProceedings of the 31st AAAI Conference on Artificial Intelligence, pp. 414-420. AAAI Press, 2017.

Participants

- Haris Aziz
Data61 – Sydney, AU
- William Bailey
University of Kentucky – Lexington, US
- Dorothea Baumeister
Heinrich-Heine-Universität Düsseldorf, DE
- Sirin Botan
University of Amsterdam, NL
- Sylvain Bouveret
LIG – Grenoble, FR
- Steven J. Brams
New York University, US
- Felix Brandt
TU München, DE
- Robert Bredereck
TU Berlin, DE
- Markus Brill
TU Berlin, DE
- Jiehua Chen
Ben Gurion University of the Negev – Beer Sheva, IL
- Marcin Dziubinski
University of Warsaw, PL
- Edith Elkind
University of Oxford, GB
- Ulle Endriss
University of Amsterdam, NL
- Piotr Faliszewski
AGH University of Science and Technology – Krakow, PL
- Judy Goldsmith
University of Kentucky – Lexington, US
- Bernard Grofman
University of California – Irvine, US
- Marc Kilgour
Wilfrid Laurier University, CA
- Werner Kirsch
FernUniversität in Hagen, DE
- Christian Klamler
Universität Graz, AT
- Martin Lackner
University of Oxford, GB
- Jérôme Lang
University Paris-Dauphine, FR
- Annick Laruelle
University of the Basque Country – Bilbao, ES
- Jean-Francois Laslier
Paris School of Economics, FR
- Andrea Loreggia
University of Padova, IT
- Nicholas Mattei
IBM TJ Watson Research Center – Yorktown Heights, US
- Vincent Merlin
Caen University, FR
- Stefan Napel
Universität Bayreuth, DE
- Svetlana Obraztsova
Nanyang TU – Singapore, SG
- Dominik Peters
University of Oxford, GB
- Clemens Puppe
KIT – Karlsruher Institut für Technologie, DE
- Lisa Rey
Heinrich-Heine-Universität Düsseldorf, DE
- Jörg Rothe
Heinrich-Heine-Universität Düsseldorf, DE
- M. Remzi Sanver
University Paris-Dauphine, FR
- Ann-Kathrin Selker
Heinrich-Heine-Universität Düsseldorf, DE
- Sujoy Sikdar
Rensselaer Polytechnic Institute – Troy, US
- Piotr Skowron
TU Berlin, DE
- Arkadii Slinko
University of Auckland, NZ
- Krzysztof Sornat
University of Wroclaw, PL
- Nimrod Talmon
Weizmann Institute – Rehovot, IL
- Kristen Brent Venable
Tulane University – New Orleans, US
- Gerhard J. Woeginger
RWTH Aachen, DE
- William S. Zwicker
Union College – Schenectady, US

Report from Dagstuhl Seminar 17262

Federated Semantic Data Management

Edited by

Olaf Hartig[1], Maria-Esther Vidal[2], and Johann-Christoph Freytag[3]

1 **Linköping University, SE,** olaf.hartig@liu.se
2 **Universidad S. Bolivar, Caracas, VE,** mvidal@usb.ve
3 **Humboldt-Universität zu Berlin, DE,** freytag@informatik.hu-berlin.de

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 17262 "Federated Semantic Data Management" (FSDM). The purpose of the seminar was to gather experts from the Semantic Web and Database communities, together with experts from application areas, to discuss in-depth open issues that have impeded FSDM approaches to be used on a large scale. The discussions were centered around the following four themes, each of which was the focus of a separate working group: i) graph data models, ii) federated query processing, iii) access control and privacy, and iv) use cases and applications. The main outcome of the seminar is a deeper understanding of the state of the art and of the open challenges of FSDM.

Seminar June 25–30, 2017 – http://www.dagstuhl.de/17262
1998 ACM Subject Classification H.1 Models and Principles, H.2 Database Management, H.3 Information Storage and Retrieval
Keywords and phrases Linked Data, Query Processing, RDF, SPARQL
Digital Object Identifier 10.4230/DagRep.7.6.135

1 Executive Summary

Olaf Hartig
Maria-Esther Vidal
Johann-Christoph Freytag

The Semantic Web is an extension of the World Wide Web in which *structured data and its meaning* is represented in a form that can be readily accessed and exploited by machines. The foundation of this representation is a graph-based data model defined by the Resource Description Framework (RDF). This framework allows for data management approaches that focus on manipulating and using data in terms of its meaning. We refer to this type of data management as *semantic data management*.

In addition to centralized access to RDF datasets, Web-based protocols such as the SPARQL protocol enable software clients to access or to query RDF datasets made available by remote servers. By integrating such remote data sources as members of a *federated system*, software clients may answer cross-dataset queries without having to retrieve various datasets into a single repository. Given such a federation, the complexity of problems of query processing and semantic data management increases due to additional parameters such as variable data transfer delays, a changing availability of federation members, the size of the federation, and distribution criteria followed to place and semantically link data in different datasets of the federation. Moreover, whenever data is replicated across federations,

Federated Semantic Data Management, *Dagstuhl Reports*, Vol. 7, Issue 06, pp. 135–167
Editors: Olaf Hartig, Maria-Esther Vidal, and Johann-Christoph Freytag
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

synchronization is required to ensure that all changes are propagated and the semantics of data is preserved. Despite a large number of technologies developed by the Semantic Web and Database communities to address problems of semantic data management, we still observe a significant lack of efficient and effective solutions to the problems of federated semantic data management (FSDM), which prevents the development of real-world applications on top of Semantic Web technologies. Additionally, existing proposals to evaluate such solutions do not sufficiently cover the large number of parameters that affect FSDM and the complexity of tradeoffs. More specifically, variables and configurations that considerably affect the federated semantic data management problems are not sufficiently defined or even considered in state-of-the-art testbeds (e.g., network latency, data fragmentation and replication, query properties, or frequency of updates).

The aim of the Dagstuhl seminar was to gather experts from the Semantic Web and Database communities, together with experts from application areas, to discuss in-depth open issues that have impeded FSDM approaches to be used on a large scale.

The following crucial questions were posed as a basis for the discussions during the seminar:

Q1 *Can traditional techniques developed for federations of relational databases be enriched with RDF semantics, and thus provide effective and efficient solutions to problems of FSDM?*

Q2 *What problems of FSDM present new research challenges that require the definition of novel techniques?*

Q3 *What is the role of RDF semantics in the definition of the problems of FSDM?*

To discuss these questions the participants of the seminar were grouped according to their areas of expertise and interests. In particular, the seminar focused on four main topic areas (see below). The results of the group discussions were presented in plenary sessions and will be compiled into manuscripts with which the seminar outcomes will be disseminated. As a basis of the group work, and to establish a common understanding of key concepts and terminology, the seminar included a few short, survey-style talks on a number of related topics. In particular, these talks covered:

- "Graph data models and graph databases" (given by Olaf Hartig),
- "RDF and semantics" (by Claudio Gutierrez),
- "Policies and access control" (by Sabrina Kirrane and Piero Andrea Bonatti),
- "Database privacy" (by Johann-Christoph Freytag),
- "Distributed database systems" (by Katja Hose), and
- "Federated query processing" (by Maria-Esther Vidal).

In addition to these survey talks, every participant was given the chance to briefly highlight their research as relevant for the seminar. Moreover, in a demo session, some of the participants showcased their FSDM-related systems and tools, which gave interested attendees of this session an opportunity to play with and better understand these systems and tools. The systems and tools demonstrated in this session were the following:

- *Triple Pattern Fragments client* that runs in a browser and executes queries over a federation of Triple Pattern Fragment (TPF) interfaces (demonstrated by Joachim Van Herwegen),
- *Network of Linked Data Eddies* (*nLDE*), an efficient client-side SPARQL query engine for querying server-side data that can be accessed via a TPF interface (demonstrated by Maribel Acosta),
- *Ladda*, a framework for delegating TPF-based query executions among multiple browsers (demonstrated by Hala Skaf-Molli),

- *Quartz*, a system for querying replicated Triple Pattern Fragments (demonstrated by Hala Skaf-Molli),
- *Ontario*, a federated SPARQL query engine for heterogeneous sources represented in different raw formats (demonstrated by Maria-Esther Vidal),
- *UltraWrap*, a framework for integrating relational databases using SPARQL federation (demonstrated by Juan Sequeda),
- *Ephedra*, a SPARQL federation engine that combines SPARQL services with other services (demonstrated by Peter Haase),
- *JedAI*, an entity resolution toolkit (demonstrated by Themis Palpanas), and
- *Exemplar Queries*, a framework for query answering using knowledge graphs (demonstrated by Themis Palpanas).

As mentioned before, besides the short survey talks, the demos, and the participants' presentations, the major focus of the seminar was on discussions in four working groups, where each of these groups addressed a different topic area. The remainder of this section provides a brief overview of the four topic areas covered by the groups and the respective results. More detailed summaries provided by each of the four groups can be found in a separate section of this report.

Graph Data Models. Graph data models such as the RDF data model allow for a representation of both data and metadata using graphs of nodes that represent entities, and edges that model connections between entities. Graph data management encompasses techniques for managing, querying, and analyzing graph data by utilizing graph-oriented operations. SQL-like query languages have been defined for evaluating declarative queries over graph data; additionally, well-known algorithms are utilized for computing graph invariants (e.g., triangle counting or degree centrality) and for solving typical graph problems (e.g., finding shortest paths, traversals, or dense subgraphs). Furthermore, several real-world applications have been built on top of existing graph-based tools (e.g., community detection, centrality analysis, and link prediction). Graphs naturally represent a wide variety of domains (e.g., social networks, biological networks) in which data, interconnectivity, and data topology all are first-class citizens, with RDF data being one example of graph data.

During the Dagstuhl seminar, a working group was formed to discuss whether tools for graph data management are sufficient to model and to manage the semantics in RDF data, taking into account that characteristics of the RDF data model (e.g., blank nodes and SPARQL operators) may affect tractability of the graph-based tasks in a federation of RDF graphs. As a first result of this discussion, the working group made the following observation. In contrast to other graph data models and query languages, the RDF data model is a "universal" data model in the sense that it is designed for sharing data and knowledge in an unbounded space such as the Web. To continue the discussion, the group introduced a definition of the notion of FSDM and identified five principles that characterize FSDM: universality, unboundedness, dynamicity, network protocols, and semantics. Based on further discussion that took into account these principles, the group made two conjectures that they plan to elaborate on in a future publication and that can be summarized as follows. First, it is impossible to build a FSDM system that fully achieves universality, unboundedness, and dynamicity, all at the same time. Second, the concepts of federation and semantics are interdependent and must be tackled together to develop effective and efficient solutions for building FSDM systems.

Federated Query Processing. A vast number of approaches have been developed to provide a unified interface for querying federations of data sources. In the context of federations of

RDF datasets, existing approaches focus on two problems: the problem of selecting the RDF datasets required to execute a federated query, and the problem of executing the resulting sub-queries efficiently against the selected data sources. Although federated query processing has been studied extensively, a number of important problems are still open, and more challenges are likely to come up as the complexity of federations increases (e.g., by increasing numbers of federation members, by replication and fragmentation of RDF data, and by federation members that update their RDF data autonomously).

During the Dagstuhl seminar, a working group was formed to discuss the problem of federated query processing over RDF data sources. Challenges imposed by the semi-structured nature of RDF, unpredictable behavior and dynamicity of Web-accessible RDF sources, and the role of the entailment regimes guided the group discussions and allowed for enumerating the main differences with the problem of federated query processing against relational databases. The group focused on the formal definition of the problem, as well as on the formalization of the subproblems of source selection, query decomposition, and query execution. As a first result, the group identified that the entailment regimes to be performed over a federation of RDF sources, as well as data replication and dynamicity, access control policies, and SPARQL query capabilities, play a crucial role in source selection, query decomposition, and query execution. State-of-the-art techniques implemented by existing approaches (e.g., FedX, ANAPSID, or Linked Data Fragments) were discussed and compared based on this formalization; the group concluded that none of existing approaches takes into account all these characteristics of RDF data sources, being required further analysis and work to empower them to solve the formalized problems. Finally, the impact of these characteristics on the performance of SPARQL operators (e.g., join, union, or optional) was discussed. The group concluded that although physical operators implemented by existing approaches are capable of adjusting query execution schedulers to RDF source availability, they are unable to adapt their execution to other RDF source characteristics, e.g., supported entailment regimes or data evolution. These issues remain open as well, and require further study from the semantic data management community.

Access Control and Privacy. Solutions to the problem of modeling access control policies for Web resources have been benefited from Semantic Web technologies. Existing rule-based logic languages rely on ontology-based reasoning tasks to represent reactive policies for access control, and to enforce and to propagate trusted and policy-compliant interactions across resources in RDF datasets. For instance, the Open Digital Rights Language (ODRL) is a rule-based approach that allows for a description of policies to access and to exchange Web resources. Nevertheless, as per the Linked Data publishing principles, RDF properties associated with any resource can be accessed by de-referencing their corresponding URL. In applications of domains of FSDM such as personalized medicine or finances, only authorized and privacy-respecting access is allowed. Thus, novel approaches are required to bridge the gap between access-control models and unrestricted access to RDF resources.

A working group with a focus on access control and privacy discussed the following open issues: a) formalisms to specify access-control and privacy policies of federation resources and to reason over the meaning of these resources; and b) techniques that enable systems to enforce privacy-aware and security-aware policies whenever a resource is accessed. After concluding that there are too many open challenges to be all solved immediately, the group agreed to focus on access control. Next, the group discussed conceptual access control models and achieved a better understanding of requirements of a conceptual framework to analyze policy-aware federated Semantic Web architectures. Finally, the group defined such a framework and made plans for a publication about it.

Use Cases and Applications. In addition to the first three working groups that focused on various more technical aspects of building FSDM systems, a fourth working group looked into applications of FSDM and use cases in which adopting FSDM would be beneficial. Specifying such use cases, as well as documenting the usage of FSDM systems in existing applications, is important to better understand the requirements and the challenges of FSDM and to derive realistic testbeds for approaches to build FSDM systems.

A key observation of the work group was that approaches to apply FSDM can be categorized into two classes depending on whether they focus i) on explorative, open-domain querying or ii) on controlled, close-domain querying. Then, the working group identified a broad set of general use cases of FSDM. Thereafter, the group defined a framework for developing specific use cases. This framework introduces a set of requirements for the specification of a use case. Finally, the group applied their framework to develop a number of example use cases.

2 Table of Contents

3 Overview of Talks

This section contains brief summaries by all participants of their research related to the topic of the seminar. During the seminar, all participants gave a lightning talk to highlight their work. Additionally, some of this work has been presented at a demo session during the seminar, others has been the topic of discussions during the seminar.

3.1 Query Processing over Graph-structured Data on the Web

Maribel Acosta (KIT – Karlsruher Institut für Technologie, DE)

Linked Data initiatives have encouraged the publication of large datasets on the Web. As a result, a huge dataspace of graph data has emerged, where data is represented using the RDF data model and can be queried using the SPARQL language. Despite these developments, the Web-like characteristics of Linked Data sources pose fundamental challenges on the efficiency and effectiveness of query processing engines over autonomous Linked Data sources. To address these challenges, this thesis focuses on the definition of flexible query processing strategies over RDF graphs on the Web.

Regarding efficient query processing, the lack of statistics about the data and unpredictable data transfer delays can negatively impact the performance of engines that consume Linked Data. This problem is mainly generated because existing engines execute fixed plans following the traditional optimize-then-execute paradigm. To tackle this problem, this thesis presents an adaptive SPARQL engine tailored to execute queries against remote Linked Data sources. Our solution comprises query optimization techniques to devise effective plans. The plans are executed following an adaptive strategy to change execution schedulers according to current conditions and reduce query runtime. The results of our empirical studies indicate that our solution outperforms static query schedulers. Our results also provide novel insights about the tradeoffs of different adaptive strategies when evaluating selective and non-selective queries.

An orthogonal but equally important aspect of querying Linked Data is the quality of the retrieved data. Executing SPARQL queries against graphs with quality issues leads to low-quality results. To tackle this problem, we propose a novel hybrid engine that integrates humans into query processing to enhance the quality of SPARQL query answers. Our solution relies on the graph structure of RDF data to decide on-the-fly which parts of a query should be crowdsourced. Experimental results show that our engine is able to enhance the completeness of SPARQL queries.

3.2 SPARQL Query Processing with Apache Spark

Bernd Amann (University Pierre & Marie Curie – Paris, FR)

Joint work of Hubert Naacke, Bernd Amann, Olivier Curé
Main reference Hubert Naacke, Bernd Amann, Olivier Curé: "SPARQL Graph Pattern Processing with Apache Spark", in Proc. of the Fifth International Workshop on Graph Data-management Experiences & Systems, GRADES@SIGMOD/PODS 2017, Chicago, IL, USA, May 14–19, 2017, pp. 1:1–1:7, ACM, 2017.
URL http://dx.doi.org/10.1145/3078447.3078448

For guaranteeing scalability, high availability and fault tolerance, RDF store implementations are rarely built from scratch but rather designed on top of a existing data processing engines. Following this line of work, we propose and compare five SPARQL query processing approaches using standard hash-join and broadcast join implementations on top of Apache Spark. Our experimentations on real-world and synthetic data sets emphasize that hybrid join plans using both broadcast or hash-join join operators simultaneously are outperforming plans using only one kind of operator.

3.3 Linked Data Containers – Shipping Linked Data and Data Management Capabilities to Consumers

Sören Auer (Universität Bonn, DE)

The amount of Linked Data both open, made available on the Web, and private, exchanged across companies and organizations, have been increasing in recent years. Maintaining and making this data available is mainly in the responsibility of data providers. Moreover, building applications on top of Linked Data in order to provide, for instance, analytics, data access control, and privacy is left to the end user or data consumers. However, many resources in terms of development costs and equipment are required by both data providers and consumers, thus impeding the development of real-world applications over Linked Data. We propose to encapsulate Linked Data and data processing functionalities in a client-side system called Linked Data Container, intended to be used by data consumers. Linked Data Containers can be deployed on the data consumer environments, ranging from Big Data to light-weight platforms.

As we learned in numerous workshops with more than 50 partner companies of the Industrial Data Space Association, keeping some level of control over the data – called data sovereignty – is a key requirement in industrial data sharing scenarios and currently the main obstacle for establishing data value chains in the industry. In many cases, cooperation partners should only gain access to a well defined fragment or usage access regime of the data.

For example, a cooperation partner in a customer bonus program, should be enabled to access information about a specific customer (e.g., identified by name or member id), but not be allowed to retrieve all customers with email and mailing addresses. In physical value chains, containers play a key role in material, component, half product, and product exchange. Containers in most cases fulfill the function to secure, condition (e.g., cool/warm), observe, or provide access to their containment.

We propose the concept of Linked Data Containers, which is a key element in the Industrial Data Space reference architecture. In order to realize the concept, we advocate for bundling data, security, access, and data processing functionality in a single artifact – the Linked Data Container. The approach can based on the recently emerging light-weight virtualization techniques and load balancing for scalable, high-performance query execution. As a result, LDC represents a novel data sharing and access paradigm, which balances costs and efforts differently between data provider and consumer than prior solutions (such as dumps, SPARQL endpoints, or TPFs). It enables controlling data access even after data shipping, thus it contributes to increased data sovereignty and consequently, in summary, it better fulfills the requirements of industrial data value chains.

3.4 Decentralizing the Semantic Web: Who will pay to realize it?

Abraham Bernstein (Universität Zürich, CH)

Joint work of Tobias Grubenmann, Daniele Dell'Aglio, Dmitry Moor, Sven Seuken, Abraham Bernstein
Main reference Tobias Grubenmann, Daniele Dell'Aglio, Abraham Bernstein, Dmitry Moor, Sven Seuken: "Decentralizing the Semantic Web: Who Will Pay to Realize It?", in Proc. of the Workshop on Decentralizing the Semantic Web 2017 co-located with 16th International Semantic Web Conference (ISWC 2017), CEUR Workshop Proceedings, Vol. 1934, CEUR-WS.org, 2017.
URL http://ceur-ws.org/Vol-1934/contribution-01.pdf

Fueled by enthusiasm of volunteers, government subsidies, and open data legislation, the Web of Data (WoD) has enjoyed a phenomenal growth. Commercial data, however, has been stuck in proprietary silos, as the monetization strategy for sharing data in the WoD is unclear. This is in contrast to the traditional web where advertisement fueled a lot of the growth. This raises the question how the WoD can (i) maintain its success when government subsidies disappear and (ii) convince commercial entities to share their wealth of data.

In this talk based on a paper [1], we propose a marketplace for decentralized data following basic WoD principles. Our approach allows a customer to buy data from different, decentralized providers in a transparent way. As such, our marketplace presents a first step towards an economically viable WoD beyond subsidies.

References

1 Tobias Grubenmann, Daniele Dell' Aglio, Abraham Bernstein, Dmitry Moor, and Sven Seuken. Decentralizing the Semantic Web: Who will pay to realize it?. In *Proceedings of the ISWC2017 workshop on Decentralizing the Semantic Web*, 2017.

3.5 Security, Privacy, and Semantics: Challenges and Opportunities

Piero Andrea Bonatti (University of Naples, IT)

After the initial focus on fully open data, the research on semantic data management is now facing the lack of support to access control and privacy enforcement. The knowledge-based nature of semantic (meta)data and the size of policies and policy-related information introduce further difficulties in the enforcement mechanisms, including anonymization, inference control

etc. There is an urgent need of collecting requirements both from federated query processing and from security/privacy enforcement, and assembling a framework for secure and privacy-enhancing federated, semantic query processing. Some of the hard challenges are:

1. Finding an optimal tradeoff between the expressivity of policy languages and the complexity of reasoning about policies.
2. Choosing a suitable confidentiality criterion to protect knowledge from attacks based on inference and metaknowledge. Such criterion should take into account also the probabilistic inferences that can be made with the help of machine learning algorithms.
3. According to the forthcoming General Data Protection Regulation, none of the anonymization methods known today produces data that can be regarded as anonymous in a legal sense. Consequently, it is crucial to manage data-subjects' consent to data processing (which makes it legal when data are not ideally anonymous and the processing does not belong to a short list of special cases of public interest).
   Here semantic languages and technologies can solve a number of problems related to expressiveness, flexibility, and interoperability. This is the approach taken, for instance, in the H2020 project SPECIAL.

References

1 Piero A. Bonatti: Datalog for Security, Privacy and Trust. Datalog 2010: 21–36
2 Piero A. Bonatti, Luigi Sauro: A Confidentiality Model for Ontologies. International Semantic Web Conference (1) 2013: 17–32
3 Joachim Biskup, Piero A. Bonatti, Clemente Galdi, Luigi Sauro: Optimality and Complexity of Inference-Proof Data Filtering and CQE. ESORICS 2014: 165–181
4 Piero A. Bonatti, Sabrina Kirrane, Axel Polleres, Rigo Wenning: Transparent Personal Data Processing: The Road Ahead. SAFECOMP Workshops 2017: 337–349

3.6 Framework for Allowing Secure and Private Access over a Federation of SPARQL Endpoints

Carlos Buil-Aranda (TU Federico Santa María – Valparaíso, CL)

Joint work of Sabrina Kirrane, Johann-Christoph Freytag, Piero Andrea Bonatti, Katja Hose, Jorge Lobo, Stasinos Konstantopoulos, Carlos Buil-Aranda

Semantic Federated Query Processing has been focused so far in improving the access to a set of SPARQL endpoints (RDF databases) and in selecting to which databases send the SPARQL queries in the main federated query. However all these improvement assume that all data is distributed across open and free to access RDF databases and none of the existing systems assume that these data may have restricted access or security policies to access the exposed data. To solve this problem we envisaged an abstract model for enabling policies in a federated data environment, security management and enforce nodes enforce in the federation engine to use a security and access framework. This model presents a framework in which a Semantic Federated Query Processing System should accommodate for effectively implementing security and privacy over the data it is being federated.

3.7 ACQUA: Approximate Continuous QUery Answering over Streams and Dynamic Linked Data Sets

Emanuele Della Valle (Polytechnic University of Milan, IT)

Joint work of Abraham Bernstein, Soheila Dehghanzadeh, Daniele Dell'Aglio, Shen Gao, Alessandra Mileo, Shima Zahmatkesh, Emanuele Della Valle

Web applications that federate dynamic data stream with distributed background data are getting a growing attention in recent years. Answering in a timely fashion, i.e., reactiveness, is one of the most important performance indicators for those applications.

The Semantic Web community showed that RDF Stream Processing (RSP) [1] is an adequate framework to develop this type of applications. However, RSP engines may lose their reactiveness due to the time necessary to access the background data when it is distributed over the Web. State-of-the-art RSP engines remain reactive using a local replica of the background data, but it progressively becomes stale if not updated to reflect the changes in the remote background data. For this reason, in the last two years, we investigated maintenance policies of the local replica that guarantee reactiveness while maximizing the freshness of the replica. They are collectively named ACQUA: Approximate Continuous QUery Answering over streams and Dynamic Linked Data sets.

In the early work [2], we focused on a continuous join operator between background data (accessed using a SPARQL 1.1 service clause) and stream data (accessed using an RSPQL WINDOW clause) assuming a 1:1 correspondence between the mappings returned on the window clause and those returned by the service clause. Then, we extended it in three directions: 1) we allowed an N:M relationship in the join [3], 2) we showed it is possible to dynamically adjust the policy [3] and 3) we added a filter clause to the service clause [4]. More recently, we showed that the opinion of multiple policies can be combined using rank aggregation[5].

References

1 Daniele Dell'Aglio, Emanuele Della Valle, Frank van Harmelen and Abraham Bernstein. Stream Reasoning: a Survey and Outlook: A summary of ten years of research and a vision for the next decade. In *Data Science Journal*, 1, 2017.
2 Soheila Dehghanzadeh, Daniele Dell'Aglio, Shen Gao, Emanuele Della Valle, Alessandra Mileo, and Abraham Bernstein: Approximate Continuous Query Answering over Streams and Dynamic Linked Data Sets. In *ICWE* 2015: 307–325
3 Shen Gao, Daniele Dell'Aglio, Soheila Dehghanzadeh, Abraham Bernstein, Emanuele Della Valle, Alessandra Mileo: Stream-Driven Linked-Data Access Under Update-Budget Constraints. In *ISWC*(1) 2016: 252–270
4 Shima Zahmatkesh, Emanuele Della Valle, Daniele Dell'Aglio: When a FILTER Makes the Difference in Continuously Answering SPARQL Queries on Streaming and Quasi-Static Linked Data. In *ICWE* 2016: 299–316
5 Shima Zahmatkesh, Emanuele Della Valle, Daniele Dell'Aglio: Using Rank Aggregation in Continuously Answering SPARQL Queries on Streaming and Quasi-static Linked Data. In *DEBS* 2017: 170–179

3.8 Why Federated Semantic Data Management Must Be FAIR

Michel Dumontier (Maastricht University, NL)

Main reference Mark D. Wilkinson, Michel Dumontier et al.: "The FAIR Guiding Principles for scientific data management and stewardship", Scientific Data 3, Vol. 3, 2016.
URL http://dx.doi.org/10.1038/sdata.2016.18

New infrastructure is needed to make digital content findable, accessible, interoperable, and reusable, as defined by the FAIR (Findability, Accessibility, Interoperability, and Reuse) principles [1]. The FAIR principles articulate a new direction for the management of digital content: that the use of globally unique, persistent identifiers to denote and retrieve structured data and metadata that meet the expectations of their communities and are expressed using global standards for semantic knowledge representation. These are all crucial aspects of Federated Semantic Data Management (FSDM). However, for FSDM to be truly realized on a global scale, new efforts must be made to create a social-technical infrastructure. Critically, we believe that efforts must be made to build out capacity, in sustainable manner, to publish, find, and reuse FDSM components including identifiers, descriptions, mappings, queries, formats, procedures, analytics, visualizations, etc. Having such components available will usher a next generation of the semantic web that people can truly embrace.

References

1 Mark D. Wilkinson, Michel Dumontier, IJsbrand Jan Aalbersberg, Gabrielle Appleton, Myles Axton, Arie Baak, Niklas Blomberg, Jan-Willem Boiten, Luiz Bonino da Silva Santos, Philip E. Bourne, Jildau Bouwman, Anthony J. Brookes, Tim Clark, Mercè Crosas, Ingrid Dillo, Olivier Dumon, Scott Edmunds, Chris T. Evelo, Richard Finkers, Alejandra Gonzalez-Beltran, Alasdair J.G. Gray, Paul Groth, Carole Goble, Jeffrey S. Grethe, Jaap Heringa, Peter A.C 't Hoen, Rob Hooft, Tobias Kuhn, Ruben Kok, Joost Kok, Scott J. Lusher, Maryann E. Martone, Albert Mons, Abel L. Packer, Bengt Persson, Philippe Rocca-Serra, Marco Roos, Rene van Schaik, Susanna-Assunta Sansone, Erik Schultes, Thierry Sengstag, Ted Slater, George Strawn, Morris A. Swertz, Mark Thompson, Johan van der Lei, Erik van Mulligen, Jan Velterop, Andra Waagmeester, Peter Wittenburg, Katherine Wolstencroft, Jun Zhao, and Barend Mons. The FAIR Guiding Principles for Scientific Data Management and Stewardship. *Scientific Data 3*, 2016.

3.9 Privacy in the Context of Federated Semantic Data Management (FSDM) Systems

Johann-Christoph Freytag (Humboldt-Universität zu Berlin, DE)

Although privacy and its protection is quite well understood in the context of tabular data there has been little to no work how to take the concepts based on k-anonymity and differential privacy into the federated semantic web world. Based on the existing concepts for FSDM systems this workshop should give an understanding what the requirements and the challenges are to introduce privacy into this world.

For this workshop I presented our work on how to detect breaches of privacy when executing a sequence of queries over a database table that stays unchanged. I show how to transform this problem into a (bipartite) graph problem and outline the challenges of how to perform inference on a set of graphs that represent the anonymized query results.

3.10 Semantics of RDF and SPARQL: Some Considerations

Claudio Gutierrez (University of Chile – Santiago de Chile, CL)

Main reference Renzo Angles, Claudio Gutierrez: "The Multiset Semantics of SPARQL Patterns", in Proc. of the The Semantic Web - ISWC 2016 - 15th International Semantic Web Conference, Kobe, Japan, October 17-21, 2016, Proceedings, Part I, Lecture Notes in Computer Science, Vol. 9981, pp. 20–36, 2016.
URL http://dx.doi.org/10.1007/978-3-319-46523-4_2

The semantics of RDF and SPARQL combine opportunities and challenges for dealing with federation in the open world of the Web. On one hand, it they allow (real) distributed creation and management of resources and vocabulary and distributed population and linking of distributed data. On the other hand, addressing incomplete information in these specifications is both complex and obscure and the logic of RDF combined with that of SPARQL is complex (even if we restrict to RDFS). In one sentence, RDF and SPARQL offer rich opportunities to deal with federation at Web level.

Under the above border conditions, my suggestion is to address, to start building, federation with basic RDF and the relational core of SPARQL (Select, Filter, And, Union, and Except, that is, the Select, Where, Natural Join, Union All and Except of SQL, see [1]), frameworks that offer all the securities and background of SQL.

Once this basic floor is firmly established, one could think of extending in the several possible directions that this core offers, namely, Bags, Incomplete information (blanks and unbound), Paths, Subqueries and Aggregation combined with some of the others, and most important, delegation features (the From Named, Graph, and Service features). Of course there are more possible extension, that I consider at this point –due to the state of the art in the previous levels– theoretical exercises: depth Logical reasoning, the interplay between CWA and OWA, several protocols (like update), etc.

Summarizing, SPARQL is an extremely complex language, and we do not know yet how the semantics of each extension interacts with other parts of the specification. This void is extremely dangerous when developing federation (that assumes each party will trust other pieces of the system). Thus my suggestion to start studying and developing federation using the simple and trusted core indicated.

References

1 Renzo Angles, Claudio Gutierrez: The Multiset Semantics of SPARQL Patterns. International Semantic Web Conference (1) 2016: 20–36

3.11 Ephedra: Extending SPARQL Federation for Efficient Combination of RDF Data and Services

Peter Haase (Metaphacts GmbH, DE)

Joint work of Andriy Nikolov, Peter Haase

Knowledge graph management use cases often require addressing hybrid information needs that involve multitude of data sources, multitude of data modalities (e.g., structured, keyword, geospatial search), and availability of computation services (e.g., machine learning and graph analytics algorithms). Although SPARQL queries provide a convenient way of expressing

information needs over RDF knowledge graphs, the level of support for hybrid information needs is limited: existing query engines usually focus on retrieving RDF data and only support a set of hard-coded built-in services. In this paper we describe representative use cases of metaphacts in the cultural heritage and pharmacy domains and the hybrid information needs arising in them. To address these needs, we present Ephedra: a SPARQL federation engine aimed at processing hybrid queries. Ephedra provides a flexible declarative mechanism for including hybrid services into a SPARQL federation and implements a number of static and runtime query optimization techniques for improving the hybrid SPARQL queries performance.

3.12 Integration and Interoperability of Graph-Data Systems

Olaf Hartig (Linköping University, SE)

My current research agenda focuses on establishing and on studying the notion of a federation of graph data systems. More precisely, I will investigate approaches (i) to integrate graph data across different systems that manage and process such data, and (ii) to integrate such systems as members of a federated system; this federated system will be able to perform workloads of queries and analysis algorithms transparently on the data that is distributed over the federation members. As an initial step towards such an integration I am investigating approaches to reconcile RDF and Property Graphs, which are the two prevalent graph data models used in many graph data systems. My current effort to achieve such a reconciliation [1] focuses on extending the RDF data model and its query language SPARQL to allow users to capture and query statement-level metadata [2].

References

1 Olaf Hartig. *Reconciliation of RDF* and Property Graphs*. In CoRR abs/1409.3288, 2014
2 Olaf Hartig. *Foundations of RDF* and SPARQL* – An Alternative Approach to Statement-Level Metadata in RDF*. In Proceedings of the 11th Alberto Mendelzon International Workshop on Foundations of Data Management (AMW), 2017

3.13 Federated Linked Data in Libraries

Jana Hentschke (Deutsche Nationalbibliothek – Frankfurt am Main, DE)

URL http://www.dnb.de/EN/lds

As a provider of RDF data the German National Library is tracking requirements to data modeling as well as infrastructure to support the intended usage of its data on the web: to be queried along with related and linked datasets. In order to be able to offer a sensible and reliable service we need to forecast user needs by monitoring the latest developments in the relevant research disciplines. Library metadata management is currently seeing a shift in data model – semantically as well as technically. This is a worldwide process which the German National Library actively involves in.

3.14 Linked Open Data, Federations, and beyond

Katja Hose (Aalborg University, DK)

In the past couple of years numerous approaches and techniques for query processing in federations of SPARQL endpoints and over Linked Open Data have been proposed by the research community. These techniques cover various subproblems, such as indexing, join processing, reasoning, query optimization, knowledge extraction, quality and completeness of knowledge bases as well as data integration, semantic data warehouses, and many more. As discussed in this talk, we are still far from having reached a point where we can conclude that we have found sufficiently good solutions for query processing in this setup. Apart from finding an overall solution that combines solutions to all these subproblems [1, 2], we are even still missing solutions for seemingly small problems, such as encoding metadata.

References

1 L. Galárraga, K. Hose, S. Razniewski. *Enabling Completeness-aware Querying in SPARQL*. In Proceedings of the 20th Int. Workshop on the Web and Databases (WebDB), 2017.
2 Jacobo Rouces, Gerard de Melo, Katja Hose. *Heuristics for Connecting Heterogeneous Knowledge via FrameBase*. In The Semantic Web. Latest Advances and New Domains. (ESWC), 2016.

3.15 The Next Generation Internet of Autonomy

Sabrina Kirrane (Wirtschaftsuniversität Wien, AT)

The Next Generation Internet should be about autonomy. Both organisations and individuals need to be able to publish both structured and unstructured data in a manner that allows them to control who can access, what data, under which constraints. Such a vision will require the adoption of existing and the development of new security and privacy mechanisms for control, transparency and compliance checking. Data consumers will naturally need to deal with diversity in the data and the query mechanism, this becomes much more complicated when there is a need to query distributed data sources. The focus of our work is to understand how existing federated query engines can be enhanced in such a way that both open and closed data can be queried in a manner that is capable of dealing with access and usage policies that are attached to the data, taking into consideration various robustness requirements.

3.16 Intelligent Data Management

Stasinos Konstantopoulos (Demokritos – Athens, GR)

In general, looking into the intersection of artificial intelligence with various subjects, including robot perception [1], computational linguistics [2], health data processing, and in particular when it comes to Dagstuhl, federated query execution planning and optimization [3, 4].

The access control and privacy group attracted my attention, as it puts forward challenging requirements for federated query processors: to offer the SPARQL programmer a seamless and transparent integrated view of a system of endpoints that comprises public endpoints and endpoints that impose complex and heterogeneous restrictions on data access and processing.

References

1 Giannakopoulos T., Konstantopoulos S., Siantikos G., and Karkaletsis V.: Design for a system of multimodal interconnected ADL recognition services. Components and Services for IoT Platforms. Springer, September 2016.
2 Konstantopoulos S.: Looking for meaning in names. From Semantics to Dialectometry. College Publications, January 2017.
3 Charalambidis, A., Troumpoukis, A., Konstantopoulos, S.: Semagrow: Optimizing federated SPARQL queries. In: Proceedings 11th Int'l Conference on Semantic Systems (SEMANTiCS 2015), Vienna, Austria, September 2015.
4 Zamani, K., Charalambidis, A., Konstantopoulos, S., Zoulis, N., and Mavroudi, E.: Workload-aware self-tuning histograms for the Semantic Web. Transactions on Large-Scale Data- and Knowledge-Centered Systems 28, Springer, 2016.

3.17 Privacy and Security in the Semantic Web

Jorge Lobo (UPF – Barcelona, ES)

Although much work has been done in the semantic web, issues of privacy and security have not been explored. Nevertheless, the field and the existing federation systems have reached the maturity to require a more systematic study of these issues. I have been working in policy-based management for distributed systems for more than a decade and I have come to the meeting to better understand the particularities of working on the open web that might need to be considered when developing a security and privacy policy management framework for the semantic web. It is foreseen that such a framework will touch upon several lines of work in policy management, from authoring to refinement, composition and analysis.

3.18 Semantic Web in the Fog of Browsers

Pascal Molli (University of Nantes, FR)

URL http://pagesperso.lina.univ-nantes.fr/ molli-p/pmwiki/pmwiki.php

Imagine connecting thousands of web browsers with browser-to-browser connections, sharing storage, bandwidth, and CPU. This builds a fog of browsers where end-user devices are ready to collaborate. Imagine semantic fog applications running in fogs of browsers, querying the linked data servers hosted in the cloud and data hosted in the fog. Fogs of browsers running semantic fog applications create a new massively decentralized infrastructure where RDF data and SPARQL query processing are available both on web servers and on browsers. I explore new opportunities and research challenges opened by a fog of browsers for the semantic web.

3.19 Query Optimization against Federations of SPARQL Endpoints

Gabriela Montoya (Aalborg University, DK)

Joint work of Katja Hose, Hala Skaf-Molli, Pascal Molli, Maria-Esther Vidal, Gabriela Montoya
Main reference Gabriela Montoya, Hala Skaf-Molli, Katja Hose: "The Odyssey Approach for Optimizing Federated SPARQL Queries", in CoRR, Vol. abs/1705.06135, 2017.
URL http://arxiv.org/abs/1705.06135

Optimization of SPARQL queries against federations of SPARQL endpoints includes: i) source selection: identifying relevant sources for each triple pattern; ii) query decomposition: combining triple patterns into subqueries to be evaluated at the endpoints; iii) join ordering: identifying the best order to evaluate the subqueries and the best ways to combine their results.

Query decomposition in the context of federations with replicated data can exploit knowledge about how data have been replicated to decompose the queries into subqueries that reduce the amount of data transfer by sending more selective subqueries to endpoints. These subqueries exploit data locality present at the endpoints to improve their availability.

If some knowledge about the data exposed by the endpoints is available, it could be exploited to obtain good estimations of cardinality that lead to generate good plans. These plans have less subqueries and during execution they require the transfer of less data and exhibit fast execution time.

References

1 Gabriela Montoya, Hala Skaf-Molli, Pascal Molli, Maria-Esther Vidal: "Decomposing federated queries in presence of replicated fragments". J. Web Sem. 42: 1-18 (2017)

3.20 Online Query Answering Using Knowledge Graphs, and Entity Resolution for Very Large and Highly Heterogeneous Data

Themis Palpanas (Paris Descartes University, FR)

Search engines are continuously employing advanced techniques that aim to capture user intentions and provide results that go beyond the data that simply satisfy the query conditions. Examples include the personalized results, related searches, similarity search, popular and relaxed queries. In this work we introduce a novel query paradigm that considers a user query as an example of the data in which the user is interested. We call these queries "exemplar queries". and claim that they can play an important role in dealing with the information deluge. We provide a formal specification of the semantics of such queries and show that they are fundamentally different from notions like queries by example, approximate and related queries. We provide an implementation of these semantics for graph-based data and present an exact solution with a number of optimizations that improve performance without compromising the quality of the answers. We study two different similarity functions, isomorphism and strong simulation, for retrieving the answers to an exemplar query, and we provide solutions for both. We also provide an approximate solution that prunes the search space and achieves considerably better time-performance with minimal or no impact on effectiveness. We experimentally evaluate the effectiveness and efficiency of these solutions with synthetic and real datasets, and illustrate the usefulness of exemplar queries in practice.

In addition, we present JedAI, a toolkit for Entity Resolution that can be used in three different ways: as an open-source Java library that implements numerous state-of-the-art, domain-independent methods, as a workbench that facilitates the evaluation of their relative performance and as a desktop application that offers out-of-the-box ER solutions. JedAI bridges the gap between the database and the Semantic Web communities, offering solutions that are applicable to both relational and RDF data. It also conveys a modular architecture that facilitates its extension with more methods and with more comprehensive workflows.

3.21 FOWLA: A Federated Architecture for Ontologies

Ana Maria Roxin (Universite de Bourgogne, FR)

Joint work of Ana Maria Roxin, Christophe Nicolle, Tarcisio Mendes de Farias
Main reference Tarcisio M. Farias, Ana Roxin, Christophe Nicolle: "FOWLA, A Federated Architecture for Ontologies", in Proc. of the Rule Technologies: Foundations, Tools, and Applications - 9th International Symposium, RuleML 2015, Berlin, Germany, August 2-5, 2015, Proceedings, Lecture Notes in Computer Science, Vol. 9202, pp. 97–111, Springer, 2015.
URL http://dx.doi.org/10.1007/978-3-319-21542-6_7

The progress of information and communication technologies has greatly increased the quantity of data to process. Thus managing data heterogeneity is a prob-lem nowadays. In the 1980s, the concept of a Federated Database Architecture (FDBA) was introduced as a collection of components that, by means of loosely coupled federation, share and exchange information. Semantic web technologies mitigate the data heterogeneity problem, however due to the data structure heterogeneity the integration of several ontologies is still a complex task. For tackling this problem, I have worked on the definition of a loosely coupled

federated ontology architecture (FOWLA). This approach allows the coexistence of various ontologies sharing common data dynamically at query execution through Horn-like rules and inference. It is also at query time that the data access policies for the federated ontologies are checked. The implementation of the FOWLA architecture comes with several advantages for interoperating several ontologies, as it allows: (1) inferring new ontology alignments; (2) avoiding data redundancy; (3) modularizing the maintainability, thought preserving the autonomy among the considered ontology-based systems, (4) addressing queries composed of vocabulary terms issued from different ontologies and (5) improving query execution time through a selection of the rules pertaining to a given query.

3.22 DREAM: Distributed RDF Engine with Adaptive Query Planner and Minimal Communication

Sherif Sakr (KSAU – Riyadh, SA)

Joint work of Mohammad Hammoud, Dania Abed Rabbou,Seyed Mohammad Reza Nouri, Seyed Mehdi Reza Beheshti, Sherif Sakr
Main reference Mohammad Hammoud, Dania Abed Rabbou, Reza Nouri, Seyed-Mehdi-Reza Beheshti, Sherif Sakr: "DREAM: Distributed RDF Engine with Adaptive Query Planner and Minimal Communication", in PVLDB, Vol. 8(6), pp. 654–665, 2015.
URL http://www.vldb.org/pvldb/vol8/p654-Hammoud.pdf

DREAM is a distributed and adaptive RDF system. To the contrary of all existing RDF systems, DREAM partitions SPARQL queries instead of partitioning RDF datasets. By not partitioning datasets, DREAM offers a general paradigm for all types of queries, and entirely averts intermediate data shutting (only meta-data are transferred). On the other hand, by partitioning queries, DREAM presents an adaptive scheme, which automatically runs queries on different numbers of machines depending on their complexities. This is achieved via employing a novel graph-based, rule-oriented query planner and a new cost model. As a result, DREAM combines the advantages of the state-of-the-art centralized and distributed RDF systems, where data communication is avoided and cluster resources are aggregated, and precludes their disadvantages, where system resources are limited and communication overhead is typically hindering.

3.23 Federated Semantic Data Management Systems in Practice

Juan F. Sequeda (Capsenta Inc. – Austin, US)

We are observing the rise and deployment of real world data integration systems based on federation and semantic technologies. The common setup we see in practice consists of the following elements: a set of source relational databases; a target ontology which provides a global semantic description of the domain, independent of the sources; and a set of mappings from the databases to the ontology. The goal is to answer queries in terms of the target ontology in a federated manner. From a practical point of view, this begs the question: where does the target ontology and the mappings come from?

We are investigating and developing methodologies and tools that can help non-experts to design the main components of a federated semantic data management system. For example in one project, we propose a pay-as-you-go methodology to design the target ontology and mapping driven by the expected questions that the semantic federated system should answer. The goal is to create the target ontology and mappings in an incremental manner, thus provide answers to questions as early as possible.

3.24 Data Availability and Efficient Query Processing for the Semantic Web

Hala Skaf-Molli (University of Nantes, FR)

Data availability and efficient query processing are challenging problems for the Semantic Web. My current research is to build decentralized and federated infrastructures to reach these objectives. More precisely, data replication improves data availability but degrades federated query processing performances, how to handle replicated data during federated query processing? Cache at client-side reduces the overhead on the server but clients do not share their caches, how to build a decentralized cooperative cache so clients can share caches? Parallel query processing and SPARQL query processing in the Fog could be improve query execution time. How to execute SPARQL queries in the Fog?

3.25 Adaptive Decentralized Control in Distributed Web Applications

Rudi Studer (KIT – Karlsruher Institut für Technologie, DE)

Joint work of Felix Keppmann, Andreas Harth

Currently, we are witnessing the rise of new technology-driven trends such as the Internet of Things, Web of Things, and Factories of the Future that are accompanied by an increasingly heterogeneous landscape of small, embedded, and highly modularized devices and applications, multitudes of manufactures and developers, and pervasion of network-accessible "things" within all areas of life. At the same time, we can observe increasing complexity of the task of integrating subsets of heterogeneous components into applications that fulfil certain needs by providing value-added functionality beyond the pure sum of their components. Enabling integration in these multi-stakeholder scenarios requires new architectural approaches for adapting components, while building on existing technologies and thus ensuring broader acceptance.

To this end, we discuss current integration-related challenges, present our approach for automated component adaptation, and describe our integration architecture that enables decentralized control.

References

1 Felix Leif Keppmann, Maria Maleshkova, Andreas Harth. *Semantic Technologies for Realising Decentralised Applications for the Web of Things*. 21st International Conference on Engineering of Complex Computer Systems, Nov 6-8, 2016, Dubai, UAE

3.26 Federated Querying on the Web

Joachim Van Herwegen (Ghent University, BE)

Joint work of Ruben Verborgh, Miel Vander Sande, Olaf Hartig, Joachim Van Herwegen, Laurens De Vocht, Ben De Meester, Gerald Haesendonck, Pieter Colpaert
Main reference Ruben Verborgh, Miel Vander Sande, Olaf Hartig, Joachim Van Herwegen, Laurens De Vocht, Ben De Meester, Gerald Haesendonck, Pieter Colpaert: "Triple Pattern Fragments: A low-cost knowledge graph interface for the Web", in J. Web Sem., Vol. 37-38, pp. 184–206, 2016.
URL http://dx.doi.org/10.1016/j.websem.2016.03.003

Triple Pattern Fragments (TPF) is a lightweight interface, allowing for SPARQL queries to be evaluated by moving some of the workload from the server to the client. TPF endpoints can be used to evaluate SPARQL queries by requesting individual pattern information through many HTTP requests and joining the results locally. Due to this querying process, these can easily be queried in a federated way by sending pattern requests to all endpoints at the same time and ignoring an endpoint if it can not answer a pattern. Interestingly, despite the substantially lighter server-side interface, the completeness and execution time of the FedBench benchmark of a TPF setup is comparable to that of a SPARQL endpoint setup [1].

To fully support querying on the Web, many other problems still have to be overcome. These include investigating:

- how to handle heterogeneous interfaces, all with their own restrictions, at the same time,
- the existing benchmarks and whether they are sufficient,
- multiple metrics besides response time, and
- how to discover sources on the web.

References

1 Verborgh, Ruben and Vander Sande, Miel and Hartig, Olaf and Van Herwegen, Joachim and De Vocht, Laurens and De Meester, Ben and Haesendonck, Gerald and Colpaert, Pieter. *Triple Pattern Fragments: a Low-cost Knowledge Graph Interface for the Web*. Journal of Web Semantics, 2016.

3.27 Federated Query Processing over RDF Data

Maria-Esther Vidal (Universidad S. Bolivar – Caracas, VE)

Main reference Maria-Esther Vidal, Simón Castillo, Maribel Acosta, Gabriela Montoya, Guillermo Palma: "On the Selection of SPARQL Endpoints to Efficiently Execute Federated SPARQL Queries", in Trans. Large-Scale Data- and Knowledge-Centered Systems, Vol. 25, pp. 109–149, 2016.
URL http://dx.doi.org/10.1007/978-3-662-49534-6_4

The increasing number of RDF data sources that allow for querying Linked Data via Web services form the basis for federated query processing over Web-accessible RDF data sources. Federated SPARQL query engines provide a unified view of a federation of RDF data sources and rely on different components to exploit the semantics encoded in RDF data during query execution. The problem of federation query processing has been extensively studied by Database and Semantic Web communities; however, these technologies have not been used in large scale yet. Additionally, there is no standard and formal definition of the problem of query processing over a federation of RDF data sources, impeding a formal evaluation of properties of the state-of-the-art approaches. During this seminar, we analyze different

scenarios of federations of RDF data sources, e.g., entailment regimes to be considered during query processing, source query capabilities, or access control, and propose a formal definition of the problem of query processing over a federation of RDF data sources. State-of-the-art approaches are evaluated in terms of the proposed formalization. We hypothesize that this characterization will allow for a better understanding of the state-of-the-art, as well as for uncovering the limitations of the current technologies that have impeded the use of existing approaches on a large scale.

References

1 Maribel Acosta, Maria-Esther Vidal, Tomas Lampo, Julio Castillo, and Edna Ruckhaus. ANAPSID: An Adaptive Query Processing Engine for SPARQL Endpoints. In *The Semantic Web – ISWC 2011 – 10th International Semantic Web Conference, Bonn, Germany, October 23-27, 2011, Proceedings, Part I.*

4 Working groups

In this section, each of the four working groups provides a summary of their discussions and the results of their work during the seminar.

4.1 Foundations of Federated Semantic Data Management on the Web

Bernd Amann (University Pierre & Marie Curie – Paris, FR), Emanuele Della Valle (Polytechnic University of Milan, IT), Claudio Gutierrez (University of Chile – Santiago de Chile, CL), Olaf Hartig (Linköping University, SE), Themis Palpanas (Paris Descartes University, FR), and Rudi Studer (KIT – Karlsruher Institut für Technologie, DE)

4.1.1 Introduction

The Semantic Web vision introduced by Tim Berners-Lee almost 20 years ago has attracted a considerable attention from various computer science domains such as databases (research tracks on RDF data management in the main database conferences like VLDB[1] and SIGMOD[2]), Artificial Intelligence (AI Magazine[3]), Web (WWW conference), and Information Retrieval [11]. The corresponding communities developed solutions for generating [10], analyzing [8], storing, and querying [4] large RDF / knowledge graphs [3, 7] which are used in many uses cases and applications[4]. SPARQL query federation engines and Linked Open Data infrastructures are initial steps towards building such applications at the Web level. However, the vision of a universal and open space for meaningfully sharing data on the Web is still not fully achieved[5].

[1] http://vldb2016.persistent.com/VLDB2016-FullProgram.html#TueF1115T1245R2
[2] http://www.sigmod2015.org/toc_sigmod.shtml
[3] https://www.aaai.org/ojs/index.php/aimagazine/article/view/2161
[4] https://www.w3.org/2001/sw/sweo/public/UseCases/
[5] https://cacm.acm.org/magazines/2016/9/206254-a-new-look-at-the-semantic-web/fulltext

During the week of the Dagstuhl seminar, the "RDF and Graph DB" working group spent several group sessions discussing the current and future issues in federated semantic data management. During the first session we mainly discussed how standard, RDF-based Semantic Web technologies compare to other data and knowledge graph systems and to existing data federation infrastructures. We made two main observations:

1. The RDF data model is a "universal" data model in the sense that it is designed for sharing data and knowledge in an unbounded space such as the Web [2]. The universality and the unbounded nature of this new scenario, for which RDF/SPARQL were originally thought of, presents challenges that are fundamentally different from other data/knowledge graph data models and query languages which are implemented in "closed" systems (regarding data and users) for complex graph queries and analytics.
2. There exist several federated data management frameworks (relational, P2P, Web services) including advanced declarative approaches for building overlay networks [6], and for distributed query processing and reasoning [1, 5]. These approaches address many fundamental concepts of federated semantic web management but they still require a further development and integration effort in order to move from the research prototype stage to standardized open federation frameworks for building applications.

Based on these observations, we decided to revisit the initial Semantic Web vision by taking into account the current scientific and technological state of the art related to Federated Semantic Data Management (FSDM). The goal was to define the fundamental characteristics of FSDM, where the purpose of this definition was to provide a basis for analyzing the opportunities and the limitations of existing semantic and/or federated data management solutions and for preparing a scientific roadmap in FSDM research.

4.1.2 Foundations and Characteristics

Our first step was to define the notion of federated semantic data management and the main abstract principles that distinguish FSDM from other frameworks like federated RDBMSs, P2P data management, graph databases, etc.

Definition (initial version; inspired by the definition of LOD[6]): Federated Semantic Data Management (FSDM) refers to universally and meaningfully publishing, connecting, and processing data in an unbounded space through a network of autonomous data sources exposed on the Web. The word "meaningfully" in this context refers to the transparent use of knowledge that is made available in an autonomous way by the data sources that participate in the federation.

Based on this definition, we have identified five principles that characterize FSDM: universality, unboundedness, dynamicity, network protocols, and semantics. Universality and unboundedness are two main principles related to the Web. Universality[7] denotes the possibility for any federation member (data source or client) to publish, connect, and consume data "anywhere on the Web." In the context of RDF, the universality principle is mainly represented by the notion of URI. Unboundedness reflects the possibility to build graphs of unbounded size where the notion of graph may refer to any of the following: raw and structured data, knowledge (vocabulary, schema, ontology), and data sources connected through a network. As a consequence of these two principles, universality and unboundedness, the complete set of all federation members cannot be assumed to be known in advance, and

[6] http://linkeddata.org/
[7] https://www.w3.org/1999/04/WebData#gloloc

neither can the exact content or the size of data and data sources. Dynamicity reflects the temporal evolution of these graphs; this evolution may be fast (e.g., RDF data streams) or slow (e.g., ontology evolution, network topology). The notion of federation in FSDM is mainly represented by the principle of network protocols. This principle stresses the application of a system of rules that allow two or more entities of a communication system to transmit information via any kind of variation of a physical quantity[8]. These rules are a fundamental part in the definition of distributed data and knowledge processing algorithms, cost models, and optimization techniques. In the RDF context, this principle is mainly represented by the SPARQL protocol recommendation, a means of conveying SPARQL queries and updates from clients to SPARQL processors[9]. Finally, the notion of semantics in FSDM represents the capacity to define "intensional" data and knowledge which can be made explicit through different inference mechanisms (RDF entailment regimes).

4.1.3 Systems

Based on the previous analysis, our second goal was to understand the impact that the identified characteristics and principles have on current and future FSDM systems. To measure this impact we started to study existing models and systems in the literature and in practice. This first study led us to the formulation of the following two hypotheses.

1. First, we argue that it is impossible to build a "perfect" FSDM system that fully achieves universality, unboundedness, and dynamicity, all at the same time [9]. As a consequence, we see the need to define new concepts and new metrics that will play, in this space of FSDM, the role played by soundness and completeness play in logic, or by precision and recall in information retrieval.
2. Second, we conjecture that the two principles of federation and semantics are interdependent, and must be tackled together. In particular, we believe that, for building effective and efficient solutions, it is not sufficient to “simply” extend a federated data management system with semantics or, vice versa, extend a semantic data management system with the notion of federation.

These two hypotheses raise a number of new challenges for current and future FSDM systems including:

- the formalization of notions of federated semantic queries,
- the definition of effective cost models and optimization techniques for federated query engines,
- the definition and the implementation of benchmarks for evaluating and comparing FSDM systems,
- the elaboration of guidelines for choosing solutions and building applications (which might have different levels of constraints for various characteristics).

4.1.4 Next Steps

The immediate next step for our working group is twofold: We aim to survey the state of the art of the aforementioned existing frameworks of federated data management and highlight their relationship (or the lack thereof) to the five principles of FSDM that we have identified, and we want to document in detail the discussions that we had in Dagstuhl. The purpose

[8] https://en.wikipedia.org/wiki/Communications_protocol
[9] https://www.w3.org/TR/sparql11-protocol/

of this work will be to provide a detailed justification and rationale for the aforementioned observations, hypotheses, and challenges. Thereafter, and based on this work, we aim to provide recommendations on research topics and problems that need to be addressed in order to build FSDM systems. We are planning to bring together the results of this work in a publication co-authored by all members of the working group.

References

1. Serge Abiteboul, Pierre Senellart, and Victor Vianu. The erc webdam on foundations of web data management. In *Proceedings of the 21st International Conference on World Wide Web*, pages 211–214. ACM, 2012.
2. Tim Berners-Lee. Www: Past, present, and future. *Computer*, 29(10):69–77, 1996.
3. Christian Bizer, Jens Lehmann, Georgi Kobilarov, Sören Auer, Christian Becker, Richard Cyganiak, and Sebastian Hellmann. Dbpedia – a crystallization point for the web of data. *Web Semant.*, 7(3):154–165, September 2009.
4. Zoi Kaoudi and Ioana Manolescu. Rdf in the clouds: A survey. *The VLDB Journal*, 24(1):67–91, February 2015.
5. Boon Thau Loo, Tyson Condie, Minos Garofalakis, David E Gay, Joseph M Hellerstein, Petros Maniatis, Raghu Ramakrishnan, Timothy Roscoe, and Ion Stoica. Declarative networking: language, execution and optimization. In *Proceedings of the 2006 ACM SIGMOD international conference on Management of data*, pages 97–108. ACM, 2006.
6. Eng Keong Lua, Jon Crowcroft, Marcelo Pias, Ravi Sharma, and Steven Lim. A survey and comparison of peer-to-peer overlay network schemes. *IEEE Communications Surveys & Tutorials*, 7(2):72–93, 2005.
7. Fabian M. Suchanek, Gjergji Kasneci, and Gerhard Weikum. Yago: A core of semantic knowledge. In *Proceedings of the 16th International Conference on World Wide Web*, WWW '07, pages 697–706, New York, NY, USA, 2007. ACM.
8. Radhika Sridhar, Padmashree Ravindra, and Kemafor Anyanwu. Rapid: Enabling scalable ad-hoc analytics on the semantic web. In *Proceedings of the 8th International Semantic Web Conference*, ISWC'09, pages 715–730, Berlin, Heidelberg, 2009. Springer-Verlag.
9. Jürgen Umbrich, Claudio Gutierrez, Aidan Hogan, Marcel Karnstedt, and Josiane Xavier Parreira. The ace theorem for querying the web of data. In *Proceedings of the 22nd International Conference on World Wide Web*, pages 133–134. ACM, 2013.
10. Julius Volz, Christian Bizer, Martin Gaedke, and Georgi Kobilarov. Discovering and maintaining links on the web of data. In *Proceedings of the 8th International Semantic Web Conference*, ISWC'09, pages 650–665, Berlin, Heidelberg, 2009. Springer-Verlag.
11. Gerhard Weikum, Gjergji Kasneci, Maya Ramanath, and Fabian Suchanek. Database and information-retrieval methods for knowledge discovery. *Commun. ACM*, 52(4):56–64, 2009.

4.2 Summary of Federated Query Processing

Juan F. Sequeda (Capsenta Inc. – Austin, US), Maribel Acosta (KIT – Karlsruher Institut für Technologie, DE), Peter Haase (Metaphacts GmbH, DE), Katja Hose (Aalborg University, DK), Gabriela Montoya (Aalborg University, DK), Sherif Sakr (KSAU – Riyadh, SA), Hala Skaf-Molli (University of Nantes, FR), Joachim Van Herwegen (Ghent University, BE), and Maria-Esther Vidal (Universidad S. Bolivar – Caracas, VE)

Our group had extensive discussions on the state of the art in Federated Query Processing from the traditional Relational Databases and Semantic Web perspectives. The goal was to understand the limitations of current approaches in considering ontological knowledge during federated query processing.

We started off by discussing what we understood by the term "semantics" within federated query processing. For some members of the group, it was assumed that federated query processing over RDF already implied "semantics" because it was considering the simple entailment[10] of RDF for reasoning. For others, federated query processing over RDF was rather a change of representation from relational to graph data model for the federated query processing problem. Thus, the differences between traditional federated query processing and federated semantic query processing was not clear from the beginning.

In order to overcome the discrepancies of our assumptions, our goal was to come up with a formal definition of federated semantic query processing where different ontological entailments of reasoning (RDFS, OWL 2 QL, etc) were explicit in the definition. This formal definition was one of the main result of our group.

During our presentation of our formal definition to the entire Dagstuhl group, we received comments that exposed possible research problems in the area of federated semantic data management. These are listed in the answer to Question 2 below.

Our next step is to write a survey paper with the goal of highlighting

1. What is Federated Semantic Data Management (FSDM)
2. How FSDM differs from traditional Federated Data Management (FDM)
3. What are the research problems and open challenges in FSDM.

4.2.1 Results

The main results of our group are the following:

- Formal definition of the Federated Semantic Query Processing problem which includes entailment regimes for reasoning (i.e., RDFS, OWL 2 QL, etc).
- Analysis of state-of-the-art Federated Semantic Data Management tools (FedX, ANAPSIS, Triple Pattern Fragment) with respect to the definition of the Federated Semantic Query Processing problem.
- Definition of the Source Selection and Query Decomposition Problem for Federated Semantic Query Processing based on our previous definition.

[10] https://www.w3.org/TR/rdf11-mt/#simpleentailment

4.2.2 Answers to Seminar Questions

Q1 *Can traditional techniques developed for federations of relational databases be enriched with RDF semantics, and thus provide effective and efficient solutions to problems of federated semantic data management?*

The characteristics of the RDF model impose restrictions on query processing over RDF data sources that impede Relational Databases technologies from providing efficient and effective solutions in general. For example, because datasets are described using binary predicates, for the source selection problem, we start out in the worst possible case scenario. Furthermore, in the Semantic Web, datasets also have "general predicates" e.g., from ontologies such as RDF/S and OWL), which cannot be used efficiently for source selection.

SPARQL operators have been implemented following the computational models of the relational algebra operators, e.g., block nested loop, dependent join, XJoin. Existing SPARQL physical operators implement the simple entailment regime. The question is how/if these operators need to be further extended to support entailments with higher expressivity (i.e., RDFS, OWL 2 QL, etc.)

Q2 *What problems of federated semantic data management present new research challenges that require the definition of novel techniques?*

The following are new research challenges within FSDM:

1. Unboundness: In the traditional federated data management problem, we have as input the set of known sources that we want to federate against. On the Web, this may not be the case: The set of sources may be unknown.
2. Correctness: In the traditional federated data management problem, we have a strict definition of correctness (a federated query is equal to a query over the entire universe of graphs). If the sources are not all known, then we may need a relaxed version of correctness. The tradeoff between soundness and completeness vs. precision and recall needs to be studied.
3. Dynamicity: How to deal with data that may change in different sources during the time of execution?

The following are extensions to existing problems of federated data management.

4. Access Control: Check access control before a query is executed and rewrite the query in order to make sure policies are enforced OR check access control during the execution of the query.
5. Source Selection: There are many more new constraints to consider within federated semantic data management, which makes the problem harder. For example,
   - Timeouts of the sources
   - Max k results
   - Different versions of SPARQL (1.0 vs 1.1)
   - Hardware capability
   - Different semantics of replicated sources (mirrors of sources)
   - Pay for accessing a source (public vs private)
   - Robots.txt
   - Query expressivity of different RDF sources: SPARQL, Linked Data Fragments, Triple Pattern Fragments
6. Heterogeneity: Although we are assuming a common data model (RDF), federated sources that access semantic data may be heterogeneous at the level of:
   - the schema (different ontologies used in different sources),

- different versions of SPARQL having an impact on the type of queries that sources are able to answer (e.g., SPARQL endpoint vs. TPF),
- computational/physical resources of the source,
- type of supported entailment regimes (RDFS, OWL 2 QL, etc)

7. Query Results: Adding provenance to the results to explain where the answer came from.

Q3 *What is the role of RDF semantics in the definition of the problems of federated semantic data management?*

Our simple answer: being able to do 1) reasoning/inferencing over 2) unbounded/unknown sources.

4.3 Privacy and Security Group Summary

Sabrina Kirrane (Wirtschaftsuniversität Wien, AT), Abraham Bernstein (Universität Zürich, CH), Piero Andrea Bonatti (University of Naples, IT), Carlos Buil-Aranda (TU Federico Santa María – Valparaíso, CL), Johann-Christoph Freytag (Humboldt-Universität zu Berlin, DE), Katja Hose (Aalborg University, DK), Stasinos Konstantopoulos (Demokritos – Athens, GR), and Jorge Lobo (UPF – Barcelona, ES)

4.3.1 Day 1 – Knowledge Sharing & Planning

The aim of day one was to discuss the status quo, to identify gaps that need to be addressed and to come up with a working plan for the rest of the week. After discussing existing work by the semantic web community on anonymisation, encryption and access control the group concluded that there are many open research challenges and it is clear that all problems cannot be solved immediately. As such the group agreed to focus on enhancing federated querying with access control. The output of the discussion was a plan for the remaining days with a view to working towards the following objectives:

- Identify a set of requirements that need to be considered (e.g. secure (compliant), soundness, maximal, computational complexity, bandwidth, robust against loss, leakage of meta policies, availability)
- Derive a conceptual framework that can be used to examine the trade-offs between different architectures and implementations
- Propose query and policy evaluation strategies taking into consideration the fact that there will be a tight coupling between access control and query planning
- Define an execution strategy towards optimisation for the identified requirements

4.3.2 Day 2 – Brainstorming

Building on existing work from the database and security communities, the group started by discussing conceptual access control models. This was followed by the mutual sharing of background information in relation to federated querying and policy enforcement. Here the term policy is used in the broader sense, for example constraints, recommendations, access policies, privacy policies, agreements etc. This naturally let to a discussion on the tight coupling between federated query planning and policies. The output of the discussion was a better understanding of:

- The set of requirements that need to be considered
- Initial thoughts on what the conceptual framework might look like

4.3.3 Day 3 – The conceptual framework

Day 3 focused exclusively on formally defining the conceptual framework that can be used to analyse policy aware federated semantic web architectures and their implementations. Key discussion points included:

- Who evaluates the policy, when and where? federation engine, endpoints, both?
- Do we need authentication, and if so who is responsible for authentication?
- What does the user/client send to the federation engine?
- What happens at the federation engine?
- What information is sent to the endpoints?
- What happens at the endpoints?
- What information is sent back from the endpoints?

4.3.4 Day 4 – Bringing it all together in the form of a paper

The final day was dedicated to discussing the shape of the paper, brainstorming about suitable publishing outlets, creating an initial structure for the paper, identifying who will be responsible for what, and deciding on next steps.

4.3.5 Results

The primary output of the group is a conceptual framework that can be used to analyse policy aware federated semantic web architectures and their implementations. In follow up work the framework will be used to examine the design space of possible solutions and discuss the tradeoffs of various architectural choices.

4.3.6 Answers to Seminar Questions

Q1 *Can traditional techniques developed for federations of relational databases be enriched with RDF semantics, and thus provide effective and efficient solutions to problems of federated semantic data management?*

Although it is possible to draw inspiration from databases to a certain extent, these techniques are not always directly applicable. One instance of this is how policies bring additional semantics that can be very naturally captured using Semantic Web technologies. Other topics such as the open nature of the Web which were mentioned by T1 also need to be considered.

Q2 *What problems of federated semantic data management present new research challenges that require the definition of novel techniques?*

Inclusion of policies in the overall architecture has not been addressed. We need to look into semantic specification, modelling, enforcement and inference, and general implications for federation engines e.g. when planning query execution.

Q3 *What is the role of RDF semantics in the definition of the problems of federated semantic data management?*

Policies always have semantics, and to leverage this in the semantic federation engine, RDF specific entailment needs to be explored (hinting at the topic of the second working group). The Semantic Web is also an opportunity for policy interoperability and smooth integration into the query planner (here we mean policy in a broad sense as in it can also be used for planning).

4.4 Federated Semantic Data Management: Use Cases and Applications

Michel Dumontier (Maastricht University, NL), Sören Auer (Universität Bonn, DE), Jana Hentschke (Deutsche Nationalbibliothek – Frankfurt am Main, DE), Pascal Molli (University of Nantes, FR), and Ana Maria Roxin (Universite de Bourgogne, FR)

The main objective of the group was to articulate a vision, benefits, use cases, and current limitations facing federated semantic data management (FDSM). The group consisted of 5 members: Sören Auer (University of Hannover), Ana Maria Roxin (University of Burgundy), Jana Hentschke (Deutsche Nationalbibliothek), Pascal Molli (Nantes University), and Michel Dumontier (Maastricht University).

Our group envisioned that Federated Semantic Data Management would enable access to structured information across heterogeneous, distributed knowledge sources in an accurate, reliable, and performant manner that will usher a new era of research and innovation. FSDM enables two main approaches to querying: (A) Explorative, open domain querying, where users are able to query the accessible web and adapt to its continuous evolution by including new relevant sources, enriching queries with relevant attributes, and suggesting improved queries based on similar federated queries. (B) Controlled, closed domain querying, in which specific data sources are used, queries are optimized for performance, and quality assessment is performed through constraint satisfaction, and that accuracy and recall query results are used in workflows and have real world applications. Our group established a framework for developing use cases. FSDM should enable a broad set of use cases including: i) the automatic discovery and querying of newly published knowledge sources, ii) the ability to answer previously unanswered questions, iii) the automatic, but parameterized gathering of more relevant data to strengthen statistical analyses, iv) the ability to perform real time fact checking, vi) the catalyst for marketplace of queries and their answers, and vii) the discovery of subtle, but important findings obtained through the analysis of massively distributed knowledge sources.

Our group established a framework for further developing specific use cases. This framework requires that use cases addresses aspects of query formulation, query execution, result generation, stakeholders, social, legal, ethical aspects, performance and availability, change management, and quality considerations. We used this framework to develop use cases to illustrate explorative open domain querying as well as controlled, closed domain querying.

We also addressed the 3 main questions of the seminar. We argued that while there may be research in relational database federation that FSDM can learn from, it is open world reasoning and inconsistency management that offer a tantalizing opportunity to move beyond the relational model, although this has yet to be fully explored. We indicate that the main problems of FSDM that require novel techniques include combinations of expressive logics,

synchronicity, and optimizations. Finally, we believe that the role of RDF semantics is to use the open world assumption to reason over an unbounded knowledge graph, and perhaps that this could help develop more advanced artificial intelligence systems.

Participants

- Maribel Acosta
KIT – Karlsruher Institut für Technologie, DE
- Bernd Amann
University Pierre & Marie Curie – Paris, FR
- Sören Auer
Universität Bonn, DE
- Abraham Bernstein
Universität Zürich, CH
- Piero Andrea Bonatti
University of Naples, IT
- Carlos Buil-Aranda
TU Federico Santa María – Valparaíso, CL
- Emanuele Della Valle
Polytechnic University of Milan, IT
- Michel Dumontier
Maastricht University, NL
- Johann-Christoph Freytag
Humboldt-Universität zu Berlin, DE
- Claudio Gutierrez
University of Chile – Santiago de Chile, CL
- Peter Haase
Metaphacts GmbH – Walldorf, DE
- Olaf Hartig
Linköping University, SE
- Jana Hentschke
Deutsche Nationalbibliothek – Frankfurt am Main, DE
- Katja Hose
Aalborg University, DK
- Sabrina Kirrane
Wirtschaftsuniversität Wien, AT
- Stasinos Konstantopoulos
Demokritos – Athens, GR
- Jorge Lobo
UPF – Barcelona, ES
- Pascal Molli
University of Nantes, FR
- Gabriela Montoya
Aalborg University, DK
- Themis Palpanas
Paris Descartes University, FR
- Ana Maria Roxin
Universite de Bourgogne, FR
- Sherif Sakr
KSAU – Riyadh, SA
- Juan F. Sequeda
Capsenta Inc. – Austin, US
- Hala Skaf-Molli
University of Nantes, FR
- Rudi Studer
KIT – Karlsruher Institut für Technologie, DE
- Joachim Van Herwegen
Ghent University, BE
- Maria-Esther Vidal
Universidad S. Bolivar – Caracas, VE