A Cyclic Proof System for Full Computation Tree Logic

Bahareh Afshari

Institute for Logic, Language and Computation, University of Amsterdam, The Netherlands
Department of Philosophy, Linguistics and Theory of Science, University of Gothenburg, Sweden

Graham E. Leigh

Department of Philosophy, Linguistics and Theory of Science, University of Gothenburg, Sweden

Guillermo Menéndez Turata

Institute for Logic, Language and Computation, University of Amsterdam, The Netherlands

Abstract

Full Computation Tree Logic, commonly denoted CTL^*, is the extension of Linear Temporal Logic LTL by path quantification for reasoning about branching time. In contrast to traditional Computation Tree Logic CTL, the path quantifiers are not bound to specific linear modalities, resulting in a more expressive language. We present a sound and complete hypersequent calculus for CTL^*. The proof system is cyclic in the sense that proofs are finite derivation trees with back-edges. A syntactic success condition on non-axiomatic leaves guarantees soundness. Completeness is established by relating cyclic proofs to a natural ill-founded sequent calculus for the logic.

2012 ACM Subject Classification Theory of computation → Modal and temporal logics; Theory of computation → Proof theory

Keywords and phrases Full computation tree logic, Hypersequent calculus, Cyclic proofs

Digital Object Identifier 10.4230/LIPIcs.CSL.2023.5

Funding Supported by the Knut and Alice Wallenberg Foundation [2020.0199]; Swedish Research Council [2017-05111, 2016-03502]; and the Dutch Research Council [OCENW.M20.048].

1 Introduction

Full Computation Tree Logic, CTL^*, is a well studied temporal logic in theoretical computer science. Its roots can be traced to Prior's stance of allowing time-distinctions (temporality) and future-alternatives (branching) in the assessment of truth. As such, it builds on a binary temporal operator *until*, written $\varphi \mathsf{U} \psi$ to express "φ is true until ψ becomes true", and a unary *next* modality, written $\mathsf{X}\varphi$, for denoting "φ is true at the next step". Additionally, path quantification is explicitly available via the *universal* and *existential* operators, $\mathsf{A}\varphi$ and $\mathsf{E}\varphi$ which, respectively, have the intended interpretation of *all* or *some* execution paths satisfy property φ. If these path quantifiers are only allowed as guards of an eventuality, namely in the form of $\mathsf{Q}(\varphi \mathsf{U} \psi)$ for $\mathsf{Q} \in \{\mathsf{A}, \mathsf{E}\}$, a strictly less expressive logic known as Computation Tree Logic, CTL, is realised. Dispensing with the path quantifiers altogether results in the most widely used temporal logic of all, Pnueli's Linear Temporal Logic, LTL.[1]

Much has already been achieved for the proof theory of CTL^* in terms of introducing finitary, infinitary, and cyclic tableaux systems (see e.g. [1, 22, 24, 25, 12, 13]). Most noteworthy is the complete (Hilbert-style) axiomatisation [22] provided by M. Reynolds in 2001. The system was later extended to include past modalities [23], a process which also simplified Reynolds' axioms and the completeness proof, though it remains a highly intricate analysis.

[1] For a comprehensive coverage of temporal logics in computer science we refer the reader to [10].

31st EACSL Annual Conference on Computer Science Logic (CSL 2023).
Editors: Bartek Klin and Elaine Pimentel; Article No. 5; pp. 5:1–5:19
Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

In this article we revisit axiomatisation of CTL^* in the light of recent developments in the area of cyclic proof theory. We introduce a sound and complete cyclic hypersequent calculus for the logic based on an intuitive set of inference rules. Local soundness of the inferences is immediate. Global soundness is achieved by a correctness condition on cyclic proofs similar to the annotated mechanism introduced by Jungteerapanich [14] and Stirling [30] for the modal μ-calculus. There are two notable deviations from the Jungteerapanich–Stirling framework. First, it is sufficient to restrict attention to annotations of length ≤ 1, i.e., annotations that identify at most one temporal operator in a formula. Second, and in strong contrast with the modal μ-calculus, annotations may identify an occurrence of either the release operator (the dual operator to "until", corresponding to the ν quantifier in μ-calculus) *or* the until operator (corresponding to the μ quantifier). Annotations of until operators play a crucial role in ensuring soundness of cycles featuring the existential path quantifier, and allow the correctness condition on proofs to be determined by the simple cycles. In particular, recognising whether a cyclic derivation is a proof requires only linear time. Finally, our system is cut-free, a property which does not come naturally for temporal logics.[2]

Aside from furthering the study of temporal logics, the present work contributes to the development of cyclic proof systems beyond the traditional realm of Gentzen-style sequent calculi where the main focus of work in cyclic proof theory currently resides (see, e.g., [30, 15, 4, 11, 17, 29, 27, 2] for recent contributions to modal and temporal logics). Such a move seems inevitable as ever more structured fixpoint logics are analysed, as witnessed by Rooduijn's cyclic hypersequent calculi for modal logics with the master modality [26], and Das and Girlando's ill-founded hypersequent calculus for Transitive Closure Logic [9].

2 The full computation tree logic CTL^*

Let Prop be a countably infinite set of propositional constants. Formulas of CTL^* are given by the following grammar.

$$\varphi ::= p \mid \neg p \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \mathsf{X}\varphi \mid \varphi \mathsf{U} \varphi \mid \varphi \mathsf{R} \varphi \mid \mathsf{A}\varphi \mid \mathsf{E}\varphi$$

where p ranges over Prop. We use Q to denote either the *universal (path) quantifier* A or the *existential (path) quantifier* E, and O to denote either the *until operator* U or the *release operator* R. A *literal formula* is either a propositional constant or the negation of one. Given a set of formulas Φ, we define $\mathsf{X}\Phi := \{\mathsf{X}\varphi : \varphi \in \Phi\}$. The notion of *subformula* is defined in the usual manner. We write $\psi \leq \varphi$ to denote that ψ is a subformula of φ.

A *labelled transition system (LTS)* is a triple $\mathcal{S} = (S, \to, \lambda)$ where S is a non-empty set of *states*, $\to$ is a binary relation on S, and $\lambda \colon S \to \mathcal{P}(\mathsf{Prop})$ is a *labelling map* which assigns to each state a set of propositional constants. $\mathcal{S}$ is *serial* if for every $s \in S$ there is some $t \in S$ such that $s \to t$. A *path* through a serial LTS $\mathcal{S}$ is an infinite sequence of states $\sigma = s_0 s_1 \cdots$ such that $s_i \to s_{i+1}$ for every $i < \omega$. The i-th state in σ is denoted by $\sigma(i)$, and (σ, i) denotes the path $s_i s_{i+1} \ldots$. Given paths σ, and σ', we write $\sigma \sim \sigma'$ if $\sigma(0) = \sigma'(0)$.

Satisfaction for CTL^* formulas is defined relative to paths through serial labelled transition systems:

- $\mathcal{S}, \sigma \models p$ iff $p \in \lambda(\sigma(0))$;
- $\mathcal{S}, \sigma \models \neg p$ iff $p \notin \lambda(\sigma(0))$;
- $\mathcal{S}, \sigma \models \varphi \wedge \psi$ iff $\mathcal{S}, \sigma \models \varphi$ and $\mathcal{S}, \sigma \models \psi$;

[2] For a discussion on cut-free sequent systems for temporal logic see, e.g., [7].

- $\mathcal{S}, \sigma \models \varphi \vee \psi$ iff $\mathcal{S}, \sigma \models \varphi$ or $\mathcal{S}, \sigma \models \psi$;
- $\mathcal{S}, \sigma \models \mathsf{X}\varphi$ iff $\mathcal{S}, (\sigma, 1) \models \varphi$;
- $\mathcal{S}, \sigma \models \varphi \mathsf{U} \psi$ iff there is a j such that $\mathcal{S}, (\sigma, j) \models \psi$ and $\mathcal{S}, (\sigma, i) \models \varphi$ for every $i < j$;
- $\mathcal{S}, \sigma \models \varphi \mathsf{R} \psi$ iff for every j, either $\mathcal{S}, (\sigma, j) \models \psi$, or there is an $i < j$ such that $\mathcal{S}, (\sigma, i) \models \varphi$;
- $\mathcal{S}, \sigma \models \mathsf{A}\varphi$ iff $\mathcal{S}, \sigma' \models \varphi$ for every $\sigma' \sim \sigma$;
- $\mathcal{S}, \sigma \models \mathsf{E}\varphi$ iff $\mathcal{S}, \sigma' \models \varphi$ for some $\sigma' \sim \sigma$.

Explicit mention of $\mathcal{S}$ may be omitted when no ambiguity arises.

The *negation* of a formula φ, in symbols $\neg\varphi$, is defined inductively via the De Morgan dualities with $\neg\mathsf{X}\varphi := \mathsf{X}\neg\varphi$, $\neg(\varphi\mathsf{U}\psi) := \neg\varphi\mathsf{R}\neg\psi$, $\neg(\varphi\mathsf{R}\psi) := \neg\varphi\mathsf{U}\neg\psi$, $\neg\mathsf{A}\varphi := \mathsf{E}\neg\varphi$, and $\neg\mathsf{E}\varphi := \mathsf{A}\neg\varphi$. A formula φ is *satisfiable* if there is a serial LTS $\mathcal{S}$ and a path σ through $\mathcal{S}$ such that $\mathcal{S}, \sigma \models \varphi$, and *unsatisfiable* otherwise. We say that φ is *valid* if $\neg\varphi$ is unsatisfiable. We write $\varphi \equiv \psi$, and say that φ and ψ are *equivalent*, if for every serial LTS $\mathcal{S}$ and every path σ on $\mathcal{S}$, we have $\mathcal{S}, \sigma \models \varphi$ iff $\mathcal{S}, \sigma \models \psi$.

Note that $\varphi\mathsf{U}\psi \equiv \psi \vee [\varphi \wedge \mathsf{X}(\varphi\mathsf{U}\psi)]$ and dually $\varphi\mathsf{R}\psi \equiv \psi \wedge [\varphi \vee \mathsf{X}(\varphi\mathsf{R}\psi)]$. These equivalences exhibit the fixpoint nature of the until and release operators.

3 Ill-founded proofs

In this section we present a sound and complete, ill-founded, sequent-style proof system for the logic CTL^* inspired by Dam's syntax trees for an embedding of CTL^* into the modal μ-calculus [8]. Proofs are potentially infinite trees whose infinite branches satisfy a syntactic correctness condition that ensures soundness of the calculus.

A *tree* is a pair $(T, \leq_T)$ where T is a non-empty, possibly countably infinite set of *vertices*, and $\leq_T$ is a partial order on T such that $\{t \in T : t \leq_T s\}$ is well-ordered for every $s \in T$ and there is a *root* $r \in T$ satisfying $r \leq_T s$ for every vertex s. We abuse notation and write T in place of $(T, \leq_T)$. For vertices $s, t \in T$, we let $[s, t]_T = s \cdots t$ be the maximal (finite) path from s to t. If $s \not\leq_T t$, $[s, t]_T$ will be the empty sequence.

A *sequent* is an expression of the form $\mathsf{Q}\Phi$, where $\mathsf{Q} \in \{\mathsf{A}, \mathsf{E}\}$ and Φ is a finite set of formulas. We identify the sequent $\mathsf{Q}\{\varphi\}$ with the formula $\mathsf{Q}\varphi$, and write $\mathsf{Q}\{\Phi, \varphi\}$ as shorthand for $\mathsf{Q}(\Phi \cup \{\varphi\})$. To each sequent $\mathsf{Q}\Phi$ we associate a corresponding formula $(\mathsf{Q}\Phi)^*$ by setting $(\mathsf{A}\Phi)^* := \mathsf{A}\bigvee\Phi$ and $(\mathsf{E}\Phi)^* := \mathsf{E}\bigwedge\Phi$. We abuse notation and write $\sigma \models \mathsf{Q}\Phi$ for $\sigma \models (\mathsf{Q}\Phi)^*$. We define $\top := \mathsf{E}\varnothing$ and $\bot := \mathsf{A}\varnothing$. A *literal sequent* is either $\top$, $\bot$, or a sequent of the form $\mathsf{Q}\lambda$ where λ is a literal formula.

A *hypersequent* is a finite set of sequents. Hypersequents are denoted by symbols Γ, Δ, Σ. We extend the translation $(\cdot)^*$ to hypersequents by setting $\Gamma^* := \bigvee\{(\mathsf{Q}\Phi)^* : \mathsf{Q}\Phi \in \Gamma\}$. As with sequents, we abuse notation and write $\sigma \models \Gamma$ for $\sigma \models \Gamma^*$.

▶ **Definition 3.1** (Derivation). *A* CTL^*_∞ derivation *of a formula* φ *is a finite or infinite tree* T *whose vertices are labelled according to the rules in Figure 1 and such that:*

1. *The root of* T *has label* $\mathsf{A}\varphi$.
2. *Every infinite branch of* T *contains infinitely many applications of* AX *or* EX.

A vertex of a derivation is *modal* if it is the conclusion of an instance of AX or EX. Condition 2 in the definition requires that every infinite branch of a derivation contains infinitely many modal vertices. This prevents the construction of infinite branches by "degenerate" applications of rules, for instance applying rule ALit repeatedly to sequent $\mathsf{A}p$.

The rules of CTL^*_∞ are correct in the following sense.

▶ **Proposition 3.2.** *The sequent rules of* CTL^*_∞ *are sound: If* $\dfrac{\Gamma_1 \quad \cdots \quad \Gamma_n}{\Delta}$ *is an instance of a* CTL^*_∞ *rule and each of* $\Gamma_1, \ldots, \Gamma_n$ *is valid, then* Δ *is valid.*

$$\mathsf{ax}_p\ \frac{}{\mathsf{Q}p, \mathsf{Q}'\neg p, \Delta} \qquad\qquad \mathsf{ax}_\top\ \frac{}{\top, \Delta}$$

$$\mathsf{ALit}\ \frac{\mathsf{A}\Phi, \mathsf{A}\lambda, \Delta}{\mathsf{A}\{\Phi, \lambda\}, \Delta} \qquad\qquad \mathsf{ELit}\ \frac{\mathsf{E}\Phi, \Delta \qquad \mathsf{E}\lambda, \Delta}{\mathsf{E}\{\Phi, \lambda\}, \Delta}$$

$$\mathsf{A}\vee\ \frac{\mathsf{A}\{\Phi, \varphi, \psi\}, \Delta}{\mathsf{A}\{\Phi, \varphi \vee \psi\}, \Delta} \qquad\qquad \mathsf{E}\vee\ \frac{\mathsf{E}\{\Phi, \varphi\}, \mathsf{E}\{\Phi, \psi\}, \Delta}{\mathsf{E}\{\Phi, \varphi \vee \psi\}, \Delta}$$

$$\mathsf{A}\wedge\ \frac{\mathsf{A}\{\Phi, \varphi\}, \Delta \qquad \mathsf{A}\{\Phi, \psi\}, \Delta}{\mathsf{A}\{\Phi, \varphi \wedge \psi\}, \Delta} \qquad\qquad \mathsf{E}\wedge\ \frac{\mathsf{E}\{\Phi, \varphi, \psi\}, \Delta}{\mathsf{E}\{\Phi, \varphi \wedge \psi\}, \Delta}$$

$$\mathsf{AA}\ \frac{\mathsf{A}\Phi, \mathsf{A}\{\psi\}, \Delta}{\mathsf{A}\{\Phi, \mathsf{A}\psi\}, \Delta} \qquad\qquad \mathsf{EA}\ \frac{\mathsf{E}\Phi, \Delta \qquad \mathsf{A}\{\psi\}, \Delta}{\mathsf{E}\{\Phi, \mathsf{A}\psi\}, \Delta}$$

$$\mathsf{AE}\ \frac{\mathsf{A}\Phi, \mathsf{E}\{\psi\}, \Delta}{\mathsf{A}\{\Phi, \mathsf{E}\psi\}, \Delta} \qquad\qquad \mathsf{EE}\ \frac{\mathsf{E}\Phi, \Delta \qquad \mathsf{E}\{\psi\}, \Delta}{\mathsf{E}\{\Phi, \mathsf{E}\psi\}, \Delta}$$

$$\mathsf{AX}\ \frac{\mathsf{A}\Phi_i, \mathsf{E}\Psi_1, \ldots, \mathsf{E}\Psi_m}{\mathsf{AX}\Phi_1, \ldots, \mathsf{AX}\Phi_n, \mathsf{EX}\Psi_1, \ldots, \mathsf{EX}\Psi_m, \Sigma} \qquad\qquad \mathsf{EX}\ \frac{\mathsf{E}\Psi_1, \ldots, \mathsf{E}\Psi_m}{\mathsf{EX}\Psi_1, \ldots, \mathsf{EX}\Psi_m, \Sigma}$$

$$\mathsf{AU}\ \frac{\mathsf{A}\{\Phi, \varphi_1, \varphi_2\}, \Delta \qquad \mathsf{A}\{\Phi, \varphi_2, \mathsf{X}(\varphi_1 \mathsf{U} \varphi_2)\}, \Delta}{\mathsf{A}\{\Phi, \varphi_1 \mathsf{U} \varphi_2\}, \Delta} \qquad\qquad \mathsf{EU}\ \frac{\mathsf{E}\{\Phi, \varphi_2\}, \mathsf{E}\{\Phi, \varphi_1, \mathsf{X}(\varphi_1 \mathsf{U} \varphi_2)\}, \Delta}{\mathsf{E}\{\Phi, \varphi_1 \mathsf{U} \varphi_2\}, \Delta}$$

$$\mathsf{AR}\ \frac{\mathsf{A}\{\Phi, \varphi_2\}, \Delta \qquad \mathsf{A}\{\Phi, \varphi_1, \mathsf{X}(\varphi_1 \mathsf{R} \varphi_2)\}, \Delta}{\mathsf{A}\{\Phi, \varphi_1 \mathsf{R} \varphi_2\}, \Delta} \qquad\qquad \mathsf{ER}\ \frac{\mathsf{E}\{\Phi, \varphi_1, \varphi_2\}, \mathsf{E}\{\Phi, \varphi_2, \mathsf{X}(\varphi_1 \mathsf{R} \varphi_2)\}, \Delta}{\mathsf{E}\{\Phi, \varphi_1 \mathsf{R} \varphi_2\}, \Delta}$$

Figure 1 Rules of the system CTL^*_∞. Q and Q' are meta-variables for path quantifiers. In the rules ALit and ELit, λ ranges over literal formulas.

Let T be a CTL^*_∞ derivation. For vertices $s, t \in T$, we write $s \longrightarrow t$ if t is an immediate successor of s in T. We write $s \xrightarrow{\mathrm{R}} t$ to specify the rule R applied at s. A *path* through T is a finite or infinite sequence of vertices $s_0 \longrightarrow s_1 \longrightarrow \cdots$.

It is convenient to refer to the sequents labelling a vertex according to their role in the applied CTL^*_∞ rules. The distinguished sequent(s) in the conclusion of rules in Figure 1, such as the sequent $\mathsf{A}\{\Phi, \varphi \wedge \psi\}$ in $\mathsf{A}\wedge$, are said to be *principal*. This includes all sequents in the conclusion of rules AX and EX. The distinguished sequent(s) in the premises of the rules, for instance $\mathsf{A}\{\Phi, \varphi\}$ and $\mathsf{A}\{\Phi, \psi\}$ in $\mathsf{A}\wedge$, are called *active*. All sequents occurring in the hypersequent Δ are *side* sequents.

Let $\pi = (s_i)_{i<N}$ be a path through T where $N \in \omega \cup \{\omega\}$. For every i, we let $\triangleright_i$ be the sequent trace relation between the labels of s_i and s_{i+1} given by $\mathsf{Q}\Xi \triangleright_i \mathsf{Q}'\Psi$ if $\mathsf{Q}'\Psi$ arises from $\mathsf{Q}\Xi$ in the rule with conclusion s_i. That is, $\mathsf{Q}\Xi \triangleright_i \mathsf{Q}'\Psi$ holds if $\mathsf{Q}\Xi = \mathsf{Q}'\Psi$ is a side sequent or $\mathsf{Q}\Xi$ is principal and $\mathsf{Q}'\Psi$ an active sequent arising from $\mathsf{Q}\Xi$. In particular, if s_i is a modal rule, we require $\Xi = \mathsf{X}\Psi$. A *sequent trace* on π is a sequence $(\mathsf{Q}_i\Xi_i)_{i<N}$ of sequents such that $\mathsf{Q}_i\Xi_i$ is in the label of s_i and $\mathsf{Q}_i\Xi_i \triangleright_i \mathsf{Q}_{i+1}\Xi_{i+1}$. We drop the subscript from $\triangleright_i$ when no ambiguity arises. A *context extraction* is a sequent trace of the form $\mathsf{Q}\{\Xi, \mathsf{Q}'\psi\} \triangleright \mathsf{Q}'\psi$ or $\mathsf{Q}\{\Xi, \lambda\} \triangleright \mathsf{Q}\lambda$ for a literal formula λ. A finite sequent trace $\mathsf{Q}_0\Xi_0 \triangleright \cdots \triangleright \mathsf{Q}_n\Xi_n$ is *stable* if $\mathsf{Q}_0 = \cdots = \mathsf{Q}_n$, and *circular* if $n > 0$ and $\mathsf{Q}_n\Xi_n = \mathsf{Q}_0\Xi_0$. When it is useful to specify the rule R by which $\mathsf{Q}_{i+1}\Xi_{i+1}$ results from $\mathsf{Q}_i\Xi_i$, we write $\mathsf{Q}_i\Xi_i \triangleright^{\mathrm{R}} \mathsf{Q}_{i+1}\Xi_{i+1}$.

As with sequents, it is convenient to have a means of referring to the formulas in principal and active sequents according to their role in the CTL^*_∞ rules. Distinguished formulas in principal sequents in Figure 1 are called *principal*. Thus, $\varphi \wedge \psi$ is the principal formula of the rule $\mathsf{A}\wedge$. Every formula in the conclusion of rules AX and EX is principal. Formulas in the set Φ are *side* formulas. The distinguished formulas in active sequents are called *active*.

Let $(\mathsf{Q}_i\Xi_i)_{i<N}$ be a sequent trace on π. For each i, let $\rhd_i$ denote the corresponding formula trace relation between Ξ_i and Ξ_{i+1}, given by $\varphi \rhd_i \psi$ if formula ψ results from φ in the sense that either φ is principal and ψ an active formula corresponding to φ, or $\varphi = \psi$ is a side formula. A *formula trace* on $(\mathsf{Q}_i\Xi_i)_{i<N}$ is a sequence of formulas $(\varphi_i)_{i<N}$ such that $\varphi_i \in \Xi_i$ and $\varphi_i \rhd_i \varphi_{i+1}$. As with sequent traces, we drop the subscript from $\rhd_i$ when no ambiguity arises. When we want to specify the rule R by means of which φ_{i+1} results from φ_i, we write $\varphi_i \rhd^{\mathrm{R}} \varphi_{i+1}$.

A *(fixpoint) unfolding* is a formula trace of the form $\psi \rhd \mathsf{X}\psi$. Note that unfoldings can only be produced by rules AU, AR, EU, ER, and thus ψ is of the form $\psi_1\mathsf{O}\psi_2$. Due to the presence of fixpoint unfoldings, the system CTL^*_∞ does not satisfy the subformula property: $\varphi \rhd \psi$ does not imply $\psi \leq \varphi$. However, ψ does belong to the *closure* of φ, which is the natural replacement of the notion of subformula in this context:

▶ **Definition 3.3.** *The* closure *of a formula φ is the smallest set of formulas $\mathsf{Clos}(\varphi)$ satisfying:*
1. *$\varphi \in \mathsf{Clos}(\varphi)$;*
2. *If $\psi_1 \star \psi_2 \in \mathsf{Clos}(\varphi)$, for $\star \in \{\wedge, \vee\}$, then $\psi_1, \psi_2 \in \mathsf{Clos}(\varphi)$;*
3. *If $\psi_1\mathsf{O}\psi_2 \in \mathsf{Clos}(\varphi)$, for $\mathsf{O} \in \{\mathsf{U}, \mathsf{R}\}$, then $\psi_1, \psi_2, \mathsf{X}(\psi_1\mathsf{O}\psi_2) \in \mathsf{Clos}(\varphi)$;*
4. *If $\mathsf{X}\psi \in \mathsf{Clos}(\varphi)$, then $\psi \in \mathsf{Clos}(\varphi)$;*
5. *If $\mathsf{Q}\psi \in \mathsf{Clos}(\varphi)$, for $\mathsf{Q} \in \{\mathsf{A}, \mathsf{E}\}$, then $\psi \in \mathsf{Clos}(\varphi)$.*

It is easy to see that $\mathsf{Clos}(\varphi)$ is always finite. And, clearly, $\varphi \rhd \psi$ implies $\psi \in \mathsf{Clos}(\varphi)$. Moreover, since $\mathsf{X}\varphi \rhd \psi$ implies $\psi \in \{\mathsf{X}\varphi, \varphi\}$, we have:

▶ **Lemma 3.4.** *Let $\rho = (\varphi_i)_{i<N}$ be a finite or infinite formula trace. For every $\varphi \in \rho$, either $\varphi \leq \varphi_0$, or $\varphi = \mathsf{X}\psi$ for some $\psi \leq \varphi_0$.*

The following is a fundamental result about infinite formula traces that follows easily from Lemma 3.4.

▶ **Proposition 3.5.** *Let T be a CTL^*_∞ derivation of a formula φ, and let $(\varphi_i)_{i<\omega}$ be an infinite formula trace on an infinite branch of T. There is a formula $\psi = \psi_1\mathsf{O}\psi_2 \in \mathsf{Clos}(\varphi)$ and some $j < \omega$ such that for every $k \geq j$ we have $\varphi_k \in \{\psi, \mathsf{X}\psi\}$. Moreover, both ψ and $\mathsf{X}\psi$ occur infinitely often in $(\varphi_{j+i})_{i<\omega}$.*

Another consequence of Lemma 3.4 is that there cannot be a formula trace of the form $\varphi \rhd \cdots \rhd \mathsf{Q}\varphi$, whence it follows that circular sequent traces must be stable.

▶ **Proposition 3.6.** *Let $(\mathsf{Q}_i\Phi_i)_{i<\omega}$ be an infinite sequent trace. There is some $j < \omega$ such that $\mathsf{Q}_{j+k} = \mathsf{Q}_j$ for every $k < \omega$.*

Given an infinite sequent trace $\tau = (\mathsf{Q}_i\Phi_i)_{i<\omega}$, we say that τ is *of type* A (E) if there is a $j < \omega$ such that $\mathsf{Q}_k = \mathsf{A}$ for all $k \geq j$ (resp., $\mathsf{Q}_k = \mathsf{E}$). Proposition 3.6 then says that every infinite sequent trace is either of type A or of type E. Similarly, an infinite formula trace $\rho = (\varphi_i)_{i<\omega}$ is *of type* U (R) if the operator O given by Proposition 3.5 is the until operator (resp., the release operator). We call $\psi_1\mathsf{O}\psi_2$ the *dominating* formula in ρ.

We are now ready to identify the CTL^*_∞ derivations that constitute *proofs*. Informally, a derivation T is a proof if every leaf of T is axiomatic and every infinite branch of T contains a "good" sequent trace.

▶ **Definition 3.7** (Ill-founded proof). *A* CTL^*_∞ proof *of a formula* φ *is a* CTL^*_∞ *derivation* T *of* φ *satisfying:*

1. *Every leaf of* T *is labelled by an instance of an axiom.*
2. *On every infinite branch of* T *there is an infinite sequent trace* τ *such that either*
   a. τ *is of type* E *and contains no infinite formula trace of type* U*, or*
   b. τ *is of type* A *and contains an infinite formula trace of type* R.

The next section is devoted to showing that φ is valid iff there exists a CTL^*_∞ proof of φ.

4 Soundness and completeness of CTL^*_∞

To prove that the system CTL^*_∞ is sound we work with *signatures*, maps that assign natural numbers to until and release formulas. We follow [8], where signatures are called *indices*.

For every $n < \omega$ we define the *n-th approximation* $\varphi\mathsf{U}^n\psi$ of a formula $\varphi\mathsf{U}\psi$ by setting $\varphi\mathsf{U}^0\psi := \psi$ and $\varphi\mathsf{U}^{n+1}\psi := \psi \lor (\varphi \land \mathsf{X}(\varphi\mathsf{U}^n\psi))$. Dually, we define the *n-th approximation* $\varphi\mathsf{R}^n\psi$ of $\varphi\mathsf{R}\psi$ as $\varphi\mathsf{R}^n\psi := \neg(\neg\varphi\mathsf{U}^n\neg\psi)$. Clearly, $\sigma \models \varphi\mathsf{U}\psi$ iff $\sigma \models \varphi\mathsf{U}^n\psi$ for some $n < \omega$, and dually $\sigma \models \varphi\mathsf{R}\psi$ iff $\sigma \models \varphi\mathsf{R}^n\psi$ for all $n < \omega$.

An occurrence in φ of a subformula $\psi_1\mathsf{O}\psi_2$ is said to be an O-*eventuality* of φ. An O-eventuality of φ is *top-level* if it is not under the scope of U, R, A or E in φ. An O-eventuality of a set of formulas Φ is an O-eventuality of $\bigwedge\Phi$. We borrow this terminology from [8].

▶ **Definition 4.1.** *An* O-signature *of a formula* φ *is a map* ι *associating a natural number to each top-level* O*-eventuality of* φ*. An* O-signature *of a sequent* $\mathsf{Q}\Phi$ *is an* O*-signature of the formula* $\bigwedge\Phi$.

Given an O-signature ι of φ, the O-signature ι^- of φ is defined as $\iota^-(\psi_1\mathsf{O}\psi_2) := \max\{\iota(\psi_1\mathsf{O}\psi_2) - 1, 0\}$ for each top-level O-eventuality $\psi_1\mathsf{O}\psi_2$ of φ. We inductively define *signed formulas* $\varphi[\iota]$, with ι an O-signature of φ:

- $\lambda[\iota] := \lambda$ for every literal formula λ.
- $(\psi_1 \star \psi_2)[\iota] := (\psi_1[\iota]) \star (\psi_2[\iota])$ for $\star \in \{\land, \lor\}$.
- $(\mathsf{Q}\psi)[\iota] := \mathsf{Q}\psi$ for $\mathsf{Q} \in \{\mathsf{A}, \mathsf{E}\}$.
- $(\psi_1\mathsf{O}\psi_2)[\iota] := \psi_1\mathsf{O}^{\iota(\psi_1\mathsf{O}\psi_2)}\psi_2$.
- $(\psi_1\mathsf{O}'\psi_2)[\iota] := \psi_1\mathsf{O}'\psi_2$ for $\mathsf{O}' \neq \mathsf{O}$.
- $(\mathsf{X}\psi)[\iota] := \mathsf{X}(\psi[\iota^-])$.

A *signed sequent* is one of the form $\mathsf{Q}\Phi[\iota] := \mathsf{Q}\{\varphi[\iota] : \varphi \in \Phi\}$ where ι is a signature of $\mathsf{Q}\Phi$.

The following two fundamental results about existence of signatures are immediate.

▶ **Proposition 4.2.** *If* $\sigma \not\models \mathsf{A}\Phi$*, then there is an* R*-signature* ι *of* $\mathsf{A}\Phi$ *such that* $\sigma \not\models \mathsf{A}\Phi[\iota]$.

▶ **Proposition 4.3.** *The following hold, where* ι *is an* R*-signature of the appropriate sequent.*

1. *If* $\sigma \not\models \mathsf{A}\{\Phi, \lambda\}[\iota]$ *and* λ *is a literal formula, then* $\sigma \not\models \mathsf{A}\Phi[\iota]$ *and* $\sigma \not\models \lambda$.
2. *If* $\sigma \not\models \mathsf{A}\{\Phi, \varphi \lor \psi\}[\iota]$*, then* $\sigma \not\models \mathsf{A}\{\Phi, \varphi, \psi\}[\iota]$.
3. *If* $\sigma \not\models \mathsf{A}\{\Phi, \varphi \land \psi\}[\iota]$*, then either* $\sigma \not\models \mathsf{A}\{\Phi, \varphi\}[\iota]$ *or* $\sigma \not\models \mathsf{A}\{\Phi, \psi\}[\iota]$.
4. *If* $\sigma \not\models \mathsf{A}\{\Phi, \mathsf{Q}\psi\}[\iota]$*, for* $\mathsf{Q} \in \{\mathsf{A}, \mathsf{E}\}$*, then* $\sigma \not\models \mathsf{A}\Phi[\iota]$.
5. *If* $\sigma \not\models \mathsf{A}\{\Phi, \varphi\mathsf{R}\psi\}[\iota]$ *then there is an* R*-signature* ι' *which agrees with* ι *on all top-level* R*-eventualities of* $\Phi \cup \{\varphi\mathsf{R}\psi\}$ *and either* $\sigma \not\models \mathsf{A}\{\Phi, \psi\}[\iota']$ *or* $\sigma \not\models \mathsf{A}\{\Phi, \varphi, \mathsf{X}(\varphi\mathsf{R}\psi)\}[\iota']$.
6. *If* $\sigma \not\models \mathsf{A}\{\Phi, \varphi\mathsf{U}\psi\}[\iota]$ *then there is an* R*-signature* ι' *which agrees with* ι *on all top-level* R*-eventualities of* $\Phi \cup \{\varphi\mathsf{U}\psi\}$ *and either* $\sigma \not\models \mathsf{A}\{\Phi, \varphi, \psi\}[\iota']$ *or* $\sigma \not\models \mathsf{A}\{\Phi, \psi, \mathsf{X}(\varphi\mathsf{U}\psi)\}[\iota']$.
7. *If* $\sigma \not\models \mathsf{AX}\Phi[\iota]$*, then there is a* $\sigma' \sim \sigma$ *such that* $(\sigma', 1) \not\models \mathsf{A}\Phi[\iota^-]$.

We are now ready to prove that the system CTL^*_∞ is sound. Some details are omitted for brevity.

▶ **Proposition 4.4** (Soundness). *If there is a* CTL^*_∞ *proof of* φ*, then* φ *is valid.*

Proof. Let T be a CTL^*_∞ proof of φ. Towards a contradiction, suppose φ is not valid. Let $\mathcal{S}$ be an LTS and σ a path on $\mathcal{S}$ such that $\sigma \not\models \mathsf{A}\varphi$. We inductively find an infinite branch $\pi = (s_i)_{i<\omega}$ on T and paths $(\sigma_i)_{i<\omega}$ on $\mathcal{S}$ such that $\sigma_i \not\models \Gamma_i$ where Γ_i is the label of s_i. Propositions 4.2 and 4.3 associate to each sequent $\mathsf{A}\Phi \in \Gamma_i$ an R-signature ι such that $\sigma_i \not\models \mathsf{A}\Phi[\iota]$. Choices at $\{\mathsf{A}\wedge, \mathsf{AU}, \mathsf{AR}\}$-vertices are resolved by Proposition 4.3. Choices at $\{\mathsf{ELit}, \mathsf{EA}, \mathsf{EE}\}$-vertices are resolved by picking the premise with simpler active sequent whenever possible. Moreover, either $\sigma_{i+1} = \sigma_i$, if s_i is not modal, or else $\sigma_{i+1} = (\sigma', 1)$ for some $\sigma' \sim \sigma_i$.

For brevity we only consider the inductive case where the rule AX is applied at vertex s_i. So assume that at s_i we have:

$$\mathsf{AX}\;\frac{\mathsf{A}\Phi_j, \mathsf{E}\Psi_1, \dots, \mathsf{E}\Psi_m}{\mathsf{AX}\Phi_1, \dots, \mathsf{AX}\Phi_n, \mathsf{EX}\Psi_1, \dots, \mathsf{EX}\Psi_m, \Delta}$$

By the inductive hypothesis, σ_i refutes every sequent in the label of s_i. In particular, $\sigma_i \not\models \mathsf{AX}\Phi_j[\iota]$ where ι is the signature given by the inductive hypothesis. By Proposition 4.3, there is a path $\sigma' \sim \sigma_i$ such that $(\sigma', 1) \not\models \mathsf{A}\Phi_j[\iota^-]$. It is easy to see that $(\sigma', 1) \not\models \mathsf{E}\Psi_k$ for any $1 \leq k \leq m$. We then let s_{i+1} be the premise of s_i and $\sigma_{i+1} := (\sigma', 1)$. To the sequent $\mathsf{A}\Phi_j$ assign signature ι^-.

We claim that the existence of the infinite branch π contradicts the fact that T is a proof. The choice of signatures guarantees that there is no infinite sequent trace of type A on π containing an infinite formula trace of type R, as otherwise the signatures assigned to the occurrences of the dominating formula would yield an infinite descending chain of natural numbers.

Therefore, since T is a proof π must contain an infinite sequent trace $\tau = (\mathsf{Q}_i\Phi_i)_{i<\omega}$ of type E in which there is no infinite formula trace of type U. Let then $\rho = (\varphi_i)_{i<\omega}$ be an infinite formula trace through τ of type R (at least one exists by Kőnig's lemma). Let $\psi = \psi_1\mathsf{R}\psi_2$ be the dominating formula in ρ and let $N < \omega$ be such that $\varphi_N = \psi$, $\varphi_{N+1} = \mathsf{X}\psi$, and $\varphi_i \in \{\psi, \mathsf{X}\psi\}$ for all $i \geq N$.

For each $i \geq N$, let $h(i) \geq i$ be least such that rule AX or EX is applied at $s_{h(i)}$. By construction, we have $\sigma_{h(i)+1} = (\sigma', 1)$ for some $\sigma' \sim \sigma_i$, so $\sigma_{h(i)+1}(0)$ is an immediate successor of $\sigma_i(0)$. Let σ^* be the path $\sigma_N(0)\,\sigma_{h(N)+1}(0)\,\sigma_{h(h(N)+1)+1}(0)\cdots$ through $\mathcal{S}$.

We inductively define a function $g\colon \{N, N+1, \dots\} \to \omega$ such that $\sigma_i \sim (\sigma^*, g(i))$ for every $i \geq N$. To that end, let $g(N) = g(N+1) = \cdots = g(h(N)) = 0$ and assume that g is defined on $N, N+1, \dots, h(i)$; set $g(h(i)+1) = \cdots = g(h(h(i)+1)) = g(h(i)) + 1$. The subtrace $(\varphi_i)_{i\geq N}$ then has the following form, where the numbers below the formulas are the indices and the ones below the braces are the g-images of the indices:

$$\underbrace{\underset{N}{\psi} \rhd \underset{N+1}{\mathsf{X}\psi} \rhd \cdots \rhd \underset{h(N)}{\mathsf{X}\psi}}_{0} \rhd \underbrace{\underset{h(N)+1}{\psi} \rhd \cdots \rhd \psi \rhd \mathsf{X}\psi \rhd \cdots \rhd \underset{h(h(N)+1)}{\mathsf{X}\psi}}_{1} \rhd \psi \rhd \cdots$$

Since $\sigma_N \not\models \mathsf{E}\Phi_N$ and $\sigma^* \sim \sigma_N$, we have $\sigma^* \not\models \bigwedge \Phi_N$. To reach a contradiction we show, by induction on χ, that if $\chi \in \Phi_i$, $i \geq N$, then $(\sigma^*, g(i)) \models \chi$. In particular, $\sigma^* \models \bigwedge \Phi_N$.

For brevity we present only the inductive case where $\chi = \chi_1\mathsf{U}\chi_2$. We show the existence of an $n < \omega$ such that $(\sigma^*, g(i) + n) \models \chi_2$ and $(\sigma^*, g(i) + m) \models \chi_1$ for all $m < n$. The fact that there is no infinite formula trace of type U on τ and a context extraction is never encountered along $(\varphi_k)_{k\geq N}$, means that there is a least $j \geq i$ such that $\chi_1 \mathsf{U} \chi_2$ is principal in s_j with unique active formula χ_2 in s_{j+1}. Let n be such that $g(j) = g(j+1) = g(i) + n$. By the inductive hypothesis, $(\sigma^*, g(i)+n) \models \chi_2$. Let $m < n$. Then, $g(i) \leq g(i) + m < g(j)$, so there is a least $i \leq k < j$ such that $g(k) = g(i) + m$. Since $g(k) < g(j)$, there is an instance of AX or EX in between s_k and s_j. So, by the minimality of j, there is a $k \leq k' < h(k)$ such that $\Phi_{k'} \ni \chi \rhd \chi_1, \mathsf{X}\chi \in \Phi_{k'+1}$, whence $(\sigma^*, g(k'+1)) \models \chi_1$ by the inductive hypothesis. And $g(k'+1) = g(k') = g(k) = g(i) + m$, so $(\sigma^*, g(i)+m) \models \chi_1$. ◀

Now we turn to the proof that the system CTL^*_∞ is complete. We argue game-theoretically and appeal to the soundness and completeness result for satisfiability of CTL^* formulas in [12].

▶ **Definition 4.5.** *A* proof-search tree *for a formula φ is a finite or infinite tree T whose vertices are labelled according to the rules on Figure 1 and such that*

1. *The root of T has label $\mathsf{A}\varphi$.*
2. *A vertex of T is a leaf iff it is either axiomatic or a set of literal sequents.*
3. *Every instance of rule EX in T is of the form*

$$\mathsf{EX}\ \frac{\mathsf{E}\Psi_1, \ldots, \mathsf{E}\Psi_m}{\mathsf{EX}\Psi_1, \ldots, \mathsf{EX}\Psi_m, \Lambda}$$

   where Λ is a non-axiomatic set of literal sequents.

4. *In place of rule AX, a branching rule AX_b is used in T:*

$$\mathsf{AX}_b\ \frac{\mathsf{A}\Phi_1, \mathsf{E}\Psi_1, \ldots, \mathsf{E}\Psi_m \qquad \cdots \qquad \mathsf{A}\Phi_n, \mathsf{E}\Psi_1, \ldots, \mathsf{E}\Psi_m}{\mathsf{AX}\Phi_1, \ldots, \mathsf{AX}\Phi_n, \mathsf{EX}\Psi_1, \ldots, \mathsf{EX}\Psi_m, \Lambda}$$

   where again Λ is a non-axiomatic set of literal sequents.

5. *Every infinite path of T contains infinitely many applications of AX or EX.*

Our aim is to show that any proof-search tree for φ contains either a CTL^*_∞ proof of φ or else a *refutation* of φ, a subtree from which the satisfiability of $\neg\varphi$ follows.

▶ **Definition 4.6** (Refutation)**.** *A* refutation *of a formula φ is a subtree T' of a proof-search tree T for φ satisfying*

1. *Every leaf of T' is labelled by a non-axiomatic set of literal sequents.*
2. *If a vertex $s \in T'$ is obtained in T by an application of a rule other than AX_b, then s has exactly one immediate successor in T'.*
3. *If a vertex $s \in T'$ is obtained in T by an application of rule AX_b, then T' contains every immediate successor of s in T.*
4. *On every infinite sequent trace τ in T'*
   a. *if τ is of type A, then every infinite formula trace is of type U;*
   b. *if τ is of type E, then some infinite formula trace is of type U.*

The term *refutation*, which we borrow from [20], is justified by the following proposition, dual to the soundness and completeness result in [12]:

▶ **Proposition 4.7** ([12, Thm. 10])**.** *A formula φ is valid iff there is no refutation of φ.*

To finish the proof of completeness we set up a game for two players whose arena is a proof-search tree in which one of the players looks for a proof and the other one for a refutation. Determinacy of the game then yields completeness of CTL^*_∞.

▶ **Definition 4.8.** *Let T be a proof-search tree for φ. We define the game $\mathcal{G}(\varphi, T)$ with arena T and players* Prov *and* Ref *as follows.*

1. *The starting position is the root of T.*
2. Prov *owns all modal vertices.*
3. Ref *owns every branching non-modal vertex.*
4. *A finite play is won by* Prov *if it is maximal and the last position is labelled by an instance of an axiom; otherwise* Ref *wins.*
5. *An infinite play is won by* Prov *if it contains an infinite sequent trace τ such that either*
   a. *τ is of type* E *and contains no infinite formula trace of type* U*, or*
   b. *τ is of type* A *and contains an infinite formula trace of type* R.

   Otherwise Ref *wins.*

Informally, Prov attempts to find a (branch of a) proof within T, whereas Ref tries to find a (branch of a) refutation.

▶ **Proposition 4.9.** *There is a winning strategy for* Prov *(*Ref*) in $\mathcal{G}(\varphi, T)$ iff T contains a* CTL^*_∞ *proof of φ (resp., a refutation of φ).*

Proof. It is clear from Definition 4.8 that a winning strategy for Prov (Ref) determines a subtree of T which is a CTL^*_∞ proof of φ (resp., a refutation). Conversely, let T' be a CTL^*_∞ proof of φ (resp. a refutation) contained in T. Then, Prov (resp., Ref) has the following winning strategy: move so as to always remain inside T'. ◀

As the winning condition in $\mathcal{G}(\varphi, T)$ is Borel, Martin's theorem [18] implies the games are always determined, whence completeness of CTL^*_∞ follows.

▶ **Proposition 4.10** (Completeness)**.** *If φ is valid, then there is a* CTL^*_∞ *proof of φ.*

Proof. Let T be a proof-search tree for φ (one always exists). Since φ is valid, Proposition 4.9 and Proposition 4.7 ensure that Ref cannot have a winning strategy in $\mathcal{G}(\varphi, T)$. By determinacy, there is a winning strategy for Prov and thus φ is provable by Proposition 4.9. ◀

Combining Proposition 4.4 and Proposition 4.10:

▶ **Theorem 4.11.** *A formula φ is valid iff there is a* CTL^*_∞ *proof of φ.*

CTL^*_∞ was formulated without any internal (i.e., formula level) or external (sequent level) structural rules except for the implicit external weakening that may accompany applications of AX, EX. Contraction (at both levels) is implicit in our choice of sets rather than multisets or sequences. The external and internal weakening rules are, respectively,

$$\mathsf{eW}\,\frac{\mathsf{Q}\Phi, \Delta}{\mathsf{Q}\Phi, \mathsf{Q}'\Psi, \Delta} \qquad\qquad \mathsf{iW}\,\frac{\mathsf{A}\Phi, \Delta}{\mathsf{A}\{\Phi \cup \Psi\}, \Delta}$$

Note, the internal weakening rule is only sound for A sequents. Soundness and completeness confirms both rules are admissible in CTL^*_∞. When working with cyclic proofs, however, it will be convenient to have at our disposal the external weakening as an explicit inference; internal weakening can be included, but is not necessary for completeness.

In the rule eW above, both $\mathsf{Q}\Phi$ and $\mathsf{Q}'\Psi$ are considered principal in the conclusion and $\mathsf{Q}\Phi$ is the only active sequent in the premise. As in the modal rules we impose that $\mathsf{Q}'\Psi \not\rhd \mathsf{Q}''\Xi$ for any sequent $\mathsf{Q}''\Xi$ in the premise.

We denote by $\mathsf{CTL}^*_\infty + \mathsf{eW}$ the proof system resulting by adding eW to CTL^*_∞.

▶ **Proposition 4.12.** *There is a* CTL^*_∞ *proof of φ iff there is a* $\mathsf{CTL}^*_\infty + \mathsf{eW}$ *proof of φ.*

5 Cyclic proofs

We now introduce a cyclic version of the system CTL^*_∞. Formulas are annotated in the style of [14, 30] to keep track of fixpoint unfoldings and determine the existence of "good" traces on cycles. Whereas a CTL^*_∞ derivation is represented as a possibly infinite tree, a $\mathsf{CTL}^*_\circ$ derivation is a finite *tree with back-edges*, that is, a pair $(T, l \mapsto c_l)$ where T is a finite tree and $l \mapsto c_l$ is a partial function defined on a subset Rep_T of the leaves of T such that each $l \in \mathsf{Rep}_T$ is mapped to a vertex $c_l <_T l$. We call c_l the *companion* of l, and leaves in Rep_T are called *repeats*. For repeats l and l', we say that l' is *reachable* from l, in symbols $l \lhd l'$, if $c_l <_T l'$. Given a tree with back-edges T, we denote by T° the result after adding an edge from each repeat $l \in \mathsf{Rep}_T$ to its companion c_l. The sequence of vertices visited by an infinite branch on T° is always of the form

$$[r, l_0]_T + [c_0, l_1]_T + \cdots + [c_n, l_{n+1}]_T + \cdots$$

where r is the root of T, each l_i is a repeat with companion c_i, and "+" denotes sequence concatenation. The following result ensures that a path through T° which visits a repeat l more than once also passes through every vertex in $[c_l, l]_T$.

▶ **Proposition 5.1.** *Let $(T, l \mapsto c_l)$ be a tree with back-edges and π a path through T°. If there are $0 \leq m < n$ such that $\pi(m) = l$ and $\pi(n) = l'$ for repeats $l, l' \in \mathsf{Rep}_T$ such that $l \lhd l'$, then $[c_l, l']_T \subseteq \{\pi(k) : m \leq k \leq n\}$.*

Fix a formula φ. To each eventuality $\psi_1 \mathsf{O} \psi_2$ in φ we associate a unique *identifier* X and write $\psi_1 \mathsf{O}^X \psi_2$. We say that X is an O-*identifier* if the eventuality corresponding to X is an O-formula. For each identifier X we assume a countably infinite set $\mathsf{N}_X = \{x_0, x_1, \dots\}$ of *names* for X. A name $x \in \mathsf{N}_X$ is an O-*name* if X is an O-identifier. In the sequel, eventualities with different identifiers will be considered as different subformulas.

An *annotated formula* is a pair (ψ, u), henceforth written ψ^u, where ψ is a formula and u is either the empty string or a name for an identifier in ψ. We call u an *annotation*. An *annotated sequent* is an expression of the form $\mathsf{Q}\Phi$, where $\mathsf{Q} \in \{\mathsf{A}, \mathsf{E}\}$ and Φ is a finite set of annotated formulas. An *annotated hypersequent* is an expression of the form $\Theta : \Gamma$, where Γ is a finite set of annotated sequents and Θ is a linear ordering of the names occurring in Γ. We call Θ the *control* of $\Theta : \Gamma$.

Given a finite sequence of names Θ, we define the following strict linear order $\prec_\Theta$ on the collection of annotations contained in Θ: $u \prec_\Theta v$ if either $v = \varnothing$ and $u \neq \varnothing$, or both u and v are non-empty and the name in u occurs in Θ strictly before the name in v.

▶ **Definition 5.2** (Cyclic derivation). *A $\mathsf{CTL}^*_\circ$ derivation of a formula φ is a finite tree with back-edges $(T, l \mapsto c_l)$ whose vertices are labelled according to the rules in Figures 2 and 3 and satisfying*

1. *The root of T has label $\mathsf{A}\varphi$.*
2. *Every leaf not labelled by an instance of an axiom is a repeat and has the same label as its companion.*
3. *For every $l \in \mathsf{Rep}_T$, there is an instance of rule AX or EX in $[c_l, l]_T$.*
4. *In rules AR and EU, if $u = \varnothing$ then x is the first name for X not already occurring in Θ. In ER_0, either $u = \varnothing$, in which case $\Theta u = \Theta$, or u is the first name for X not already occurring in Θ.*
5. *Rule iThin has priority over other rules: If a hypersequent in T can be witnessed as the conclusion of an application of iThin, then it is.*

$$\mathsf{ax}_p \; \frac{}{\Theta : \mathsf{Q}p, \mathsf{Q}'\neg p, \Delta} \qquad \mathsf{ax}_\top \; \frac{}{\Theta : \top, \Delta}$$

$$\mathsf{ALit} \; \frac{\Theta : \mathsf{A}\Phi, \mathsf{A}\lambda, \Delta}{\Theta : \mathsf{A}\{\Phi, \lambda\}, \Delta} \qquad \mathsf{ELit} \; \frac{\Theta : \mathsf{E}\Phi, \Delta \qquad \Theta' : \mathsf{E}\lambda, \Delta}{\Theta : \mathsf{E}\{\Phi, \lambda\}, \Delta}$$

$$\mathsf{A}\vee \; \frac{\Theta : \mathsf{A}\{\Phi, \varphi, \psi\}, \Delta}{\Theta : \mathsf{A}\{\Phi, \varphi \vee \psi\}, \Delta} \qquad \mathsf{E}\vee \; \frac{\Theta : \mathsf{E}\{\Phi, \varphi\}, \mathsf{E}\{\Phi, \psi\}, \Delta}{\Theta : \mathsf{E}\{\Phi, \varphi \vee \psi\}, \Delta}$$

$$\mathsf{A}\wedge \; \frac{\Theta : \mathsf{A}\{\Phi, \varphi\}, \Delta \qquad \Theta : \mathsf{A}\{\Phi, \psi\}, \Delta}{\Theta : \mathsf{A}\{\Phi, \varphi \wedge \psi\}, \Delta} \qquad \mathsf{E}\wedge \; \frac{\Theta : \mathsf{E}\{\Phi, \varphi, \psi\}, \Delta}{\Theta : \mathsf{E}\{\Phi, \varphi \wedge \psi\}, \Delta}$$

$$\mathsf{AA} \; \frac{\Theta : \mathsf{A}\Phi, \mathsf{A}\{\psi\}, \Delta}{\Theta : \mathsf{A}\{\Phi, \mathsf{A}\psi\}, \Delta} \qquad \mathsf{EA} \; \frac{\Theta : \mathsf{E}\Phi, \Delta \qquad \Theta' : \mathsf{A}\{\psi\}, \Delta}{\Theta : \mathsf{E}\{\Phi, \mathsf{A}\psi\}, \Delta}$$

$$\mathsf{AE} \; \frac{\Theta : \mathsf{A}\Phi, \mathsf{E}\{\psi\}, \Delta}{\Theta : \mathsf{A}\{\Phi, \mathsf{E}\psi\}, \Delta} \qquad \mathsf{EE} \; \frac{\Theta : \mathsf{E}\Phi, \Delta \qquad \Theta' : \mathsf{E}\{\psi\}, \Delta}{\Theta : \mathsf{E}\{\Phi, \mathsf{E}\psi\}, \Delta}$$

$$\mathsf{AX} \; \frac{\Theta' : \mathsf{A}\Phi_i, \mathsf{E}\Psi_1, \ldots, \mathsf{E}\Psi_m}{\Theta : \mathsf{AX}\Phi_1, \ldots, \mathsf{AX}\Phi_n, \mathsf{EX}\Psi_1, \ldots, \mathsf{EX}\Psi_m, \Sigma} \qquad \mathsf{EX} \; \frac{\Theta' : \mathsf{E}\Psi_1, \ldots, \mathsf{E}\Psi_m}{\Theta : \mathsf{EX}\Psi_1, \ldots, \mathsf{EX}\Psi_m, \Sigma}$$

$$\mathsf{iThin} \; \frac{\Theta' : \mathsf{Q}\{\Phi, \varphi^x\}, \Delta}{\Theta : \mathsf{Q}\{\Phi, \varphi^x, \varphi^v\}, \Delta} \; x \prec_\Theta v \qquad \mathsf{eW} \; \frac{\Theta' : \mathsf{Q}\Phi, \Delta}{\Theta : \mathsf{Q}\Phi, \mathsf{Q}'\Psi, \Delta}$$

Figure 2 Non-fixpoint rules of system $\mathsf{CTL}^*_\circ$. In ALit and ELit, λ is a literal formula. In all rules, Θ' is the result of removing from Θ all names not occurring in the associated hypersequent.

For clarity of presentation the ER inference is split into two rules, depending on whether the principal formula is annotated. We refer to the two rules ER_0 and ER_1 jointly as ER. Sequent traces, as well as principal and active sequents and formulas, follow the definition from the unannotated calculus. In the case of iThin, both φ^x and φ^v are principal and we let $\varphi^x \rhd \varphi^x$ and $\varphi^v \rhd \varphi^x$.

The prioritisation of iThin in the definition above is not necessary and the notion of a cyclic proof given below is sound and complete without this restriction. However, it ensures that all sequents in a derivation are bounded in size. Judicious use of external weakening in the form of an external *thinning* rule can be used to ensure a bound on the size of hypersequents also. We isolate a particular form of the rule eW that will be useful in proving completeness. This is rule

$$\mathsf{eThin} \; \frac{\Theta' : \mathsf{Q}\{\varphi_0^{u_0}, \ldots, \varphi_n^{u_n}\}, \Delta}{\Theta : \mathsf{Q}\{\varphi_0^{u_0}, \ldots, \varphi_n^{u_n}\}, \mathsf{Q}\{\varphi_0^{v_0}, \ldots, \varphi_n^{v_n}\}, \Delta}$$

with the restriction that some name u_i precedes all the names $v_0, \ldots, v_n$, i.e., for some $i \leq n$, we have $u_i \prec_\Theta v_j$ for every $j \leq n$. The condition ensures that the first name in Θ distinguishing the sequent $\mathsf{Q}\{\varphi_0^{u_0}, \ldots, \varphi_n^{u_n}\}$ from $\mathsf{Q}\{\varphi_0^{v_0}, \ldots, \varphi_n^{v_n}\}$ occurs in $\{u_0, \ldots, u_n\}$ and is preserved in the premise.

Prioritising applications of eThin alongside iThin ensures that sequents and hypersequents never grow past a finite bound. More precisely, given a $\mathsf{CTL}^*_\infty + \mathsf{eW}$ derivation T, let $\mathsf{ann}(T)$ be the result after annotating the hypersequents in T according to the rules of $\mathsf{CTL}^*_\circ$ and applying rules iThin and eThin whenever possible. We then have

$$\mathsf{AU}\ \frac{\Theta : \mathsf{A}\{\Phi, \varphi_1, \varphi_2\}, \Delta \qquad \Theta : \mathsf{A}\{\Phi, \varphi_2, \mathsf{X}(\varphi_1 \mathsf{U} \varphi_2)\}, \Delta}{\Theta : \mathsf{A}\{\Phi, \varphi_1 \mathsf{U} \varphi_2\}, \Delta}$$

$$\mathsf{AR}\ \frac{\Theta' : \mathsf{A}\{\Phi, \varphi_2\}, \Delta \qquad \Theta x : \mathsf{A}\{\Phi, \varphi_1, \mathsf{X}(\varphi_1 \mathsf{R}^X \varphi_2)^x\}, \Delta}{\Theta : \mathsf{A}\{\Phi, (\varphi_1 \mathsf{R}^X \varphi_2)^u\}, \Delta}\ u \in \{\varnothing, x\}$$

$$\mathsf{EU}\ \frac{\Theta x : \mathsf{E}\{\Phi, \varphi_2\}, \mathsf{E}\{\Phi, \varphi_1, \mathsf{X}(\varphi_1 \mathsf{U}^X \varphi_2)^x\}, \Delta}{\Theta : \mathsf{E}\{\Phi, (\varphi_1 \mathsf{U}^X \varphi_2)^u\}, \Delta}\ u \in \{\varnothing, x\}$$

$$\mathsf{ER}_0\ \frac{\Theta u : \mathsf{E}\{\Phi, \varphi_1, \varphi_2\}, \mathsf{E}\{\Phi, \varphi_2, \mathsf{X}(\varphi_1 \mathsf{R}^X \varphi_2)^u\}, \Delta}{\Theta :\ \mathsf{E}\{\Phi, \varphi_1 \mathsf{R}^X \varphi_2\}, \Delta}\ u \in \{\varnothing\} \cup \mathsf{N}_X$$

$$\mathsf{ER}_1\ \frac{\Theta' : \mathsf{E}\{\Phi, \varphi_1, \varphi_2\}, \mathsf{E}\{\Phi, \varphi_2, \mathsf{X}(\varphi_1 \mathsf{R}^X \varphi_2)^u\}, \Delta}{\Theta :\ \mathsf{E}\{\Phi, (\varphi_1 \mathsf{R}^X \varphi_2)^x\}, \Delta}\ u \in \{\varnothing, x\}$$

Figure 3 Fixpoint rules of system $\mathsf{CTL}^*_\circ$. In all rules, Θ' denotes the result of removing from Θ all names not occurring in the associated hypersequent. The control Θx denotes Θ if $x \in \Theta$ and the concatenation of Θ and x otherwise. Rule ER_0 is subject to the restriction: u is either $\varnothing$ or a name for X not occurring in Θ.

▶ **Proposition 5.3.** *There are only finitely many distinct annotated hypersequents in* $\mathsf{ann}(T)$.

A name is *fixed* on a path π through a $\mathsf{CTL}^*_\circ$ derivation if it occurs in every control in π. Similarly, a name is *fixed* on a sequent (formula) trace if it occurs in each sequent (resp., formula) on the trace.

▶ **Definition 5.4** (Good trace). *Let T be a $\mathsf{CTL}^*_\circ$ derivation, and let τ be a (finite) stable sequent trace on a path through T. We say that τ is* good *if the following hold.*

1. *There is an R-name fixed on τ.*
2. *No U-name is fixed on τ.*

Otherwise τ is said to be bad.

▶ **Definition 5.5** (Successful repeat). *Let T be a $\mathsf{CTL}^*_\circ$ derivation. A repeat $l \in \mathsf{Rep}_T$ is* successful *if the following hold.*

1. *There is a good sequent trace on $[c_l, l]_T$.*
2. *No R-name is fixed on a bad trace on $[c_l, l]_T$.*

In other words, l is successful if there is a good trace on $[c_l, l]_T$ and, moreover, the path contains no E-trace on which an R-name and a U-name are both fixed. Observe that the rule EU always annotates the active U operator and the only rules eliminating U-names along sequent traces are the thinning rules; on the contrary, the rule ER_1 provides a mechanism for eliminating R-names along a path. Thus, on a successful repeat, if EU affects an E-trace between companion and leaf, then all R-names on the trace must eventually be eliminated (cf. Lemma 6.6 below).

▶ **Definition 5.6** (Cyclic proof). *A $\mathsf{CTL}^*_\circ$* proof *of a formula φ is a $\mathsf{CTL}^*_\circ$ derivation T of φ each of whose leaves is either axiomatic or else a successful repeat.*

Figure 4 contains a CTL^*_∞ proof of the valid formula $(\neg p \mathsf{U} p) \vee (\bot \mathsf{R} \neg p)$ together with its cyclic version. The unique infinite branch in the ill-founded proof corresponds to the unique cycle in the cyclic one.

$$
\begin{prooftree}
\AxiomC{}
\UnaryInfC{$\mathsf{A}\{\mathsf{X}(\neg p\mathsf{U}p), \mathsf{X}(\bot\mathsf{R}\neg p)\}, \mathsf{A}p, \mathsf{A}\neg p$}
\UnaryInfC{$\mathsf{A}\{\mathsf{X}(\neg p\mathsf{U}p), \neg p, \mathsf{X}(\bot\mathsf{R}\neg p)\}, \mathsf{A}p$}
\AxiomC{$\vdots$}
\noLine
\UnaryInfC{$\mathsf{A}\{\neg p\mathsf{U}p, \bot\mathsf{R}\neg p\}$}
\UnaryInfC{$\mathsf{A}\{\mathsf{X}(\neg p\mathsf{U}p), \mathsf{X}(\bot\mathsf{R}\neg p)\}, \mathsf{A}p$}
\UnaryInfC{$\mathsf{A}\{\mathsf{X}(\neg p\mathsf{U}p), p, \mathsf{X}(\bot\mathsf{R}\neg p)\}, \mathsf{A}p$}
\BinaryInfC{$\mathsf{A}\{\mathsf{X}(\neg p\mathsf{U}p), \bot, \mathsf{X}(\bot\mathsf{R}\neg p)\}, \mathsf{A}p$}
\doubleLine
\UnaryInfC{$\mathsf{A}\{\mathsf{X}(\neg p\mathsf{U}p), \bot\mathsf{R}\neg p\}, \mathsf{A}p$}
\UnaryInfC{$\mathsf{A}\{p, \mathsf{X}(\neg p\mathsf{U}p), \bot\mathsf{R}\neg p\}$}
\doubleLine
\UnaryInfC{$\mathsf{A}\{\neg p\mathsf{U}p, \bot\mathsf{R}\neg p\}$}
\UnaryInfC{$\mathsf{A}\{\neg p\mathsf{U}p \vee \bot\mathsf{R}\neg p\}$}
\end{prooftree}
$$

$$
\begin{prooftree}
\AxiomC{}
\UnaryInfC{$x : \mathsf{A}\{\mathsf{X}(\neg p\mathsf{U}p), \mathsf{X}(\bot\mathsf{R}\neg p)^x\}, \mathsf{A}p, \mathsf{A}\neg p$}
\UnaryInfC{$x : \mathsf{A}\{\mathsf{X}(\neg p\mathsf{U}p), \neg p, \mathsf{X}(\bot\mathsf{R}\neg p)^x\}, \mathsf{A}p$}
\AxiomC{$(\dagger) \quad x : \mathsf{A}\{\neg p\mathsf{U}p, (\bot\mathsf{R}\neg p)^x\}$}
\UnaryInfC{$x : \mathsf{A}\{\mathsf{X}(\neg p\mathsf{U}p), \mathsf{X}(\bot\mathsf{R}\neg p)^x\}, \mathsf{A}p$}
\UnaryInfC{$x : \mathsf{A}\{\mathsf{X}(\neg p\mathsf{U}p), p, \mathsf{X}(\bot\mathsf{R}\neg p)^x\}, \mathsf{A}p$}
\BinaryInfC{$x : \mathsf{A}\{\mathsf{X}(\neg p\mathsf{U}p), \bot, \mathsf{X}(\bot\mathsf{R}\neg p)^x\}, \mathsf{A}p$}
\doubleLine
\UnaryInfC{$x : \mathsf{A}\{\mathsf{X}(\neg p\mathsf{U}p), (\bot\mathsf{R}\neg p)^x\}, \mathsf{A}p$}
\UnaryInfC{$x : \mathsf{A}\{p, \mathsf{X}(\neg p\mathsf{U}p), (\bot\mathsf{R}\neg p)^x\}$}
\doubleLine
\UnaryInfC{$(\dagger) \quad x : \mathsf{A}\{\neg p\mathsf{U}p, (\bot\mathsf{R}\neg p)^x\}$}
\UnaryInfC{$x : \mathsf{A}\{\mathsf{X}(\neg p\mathsf{U}p), \mathsf{X}(\bot\mathsf{R}\neg p)^x\}, \mathsf{A}p$}
\doubleLine
\UnaryInfC{$x : \mathsf{A}\{\mathsf{X}(\neg p\mathsf{U}p), \bot, \mathsf{X}(\bot\mathsf{R}\neg p)^x\}, \mathsf{A}p$}
\doubleLine
\UnaryInfC{$\mathsf{A}\{\mathsf{X}(\neg p\mathsf{U}p), \bot\mathsf{R}\neg p\}, \mathsf{A}p$}
\doubleLine
\UnaryInfC{$\mathsf{A}\{\neg p\mathsf{U}p \vee \bot\mathsf{R}\neg p\}$}
\end{prooftree}
$$

Figure 4 Top: A CTL^*_∞ proof of the valid formula $(\neg p\mathsf{U}p) \vee (\bot\mathsf{R}\neg p)$ where $\bot := p \wedge \neg p$. Bottom: A $\mathsf{CTL}^*_\circ$ proof of the same formula. The unique repeat in the cyclic proof and its companion are marked with †. Double line indicates that some vertices are omitted for brevity.

In the next section we prove that φ is valid iff there exists a $\mathsf{CTL}^*_\circ$ proof of φ. We do so by showing that ill-founded proofs can be seen as unravellings of cyclic proofs.

6 Soundness and completeness of $\mathsf{CTL}^*_\circ$

Every successful repeat l in a $\mathsf{CTL}^*_\circ$ derivation T has an associated *invariant*, denoted $\mathsf{inv}(l)$, the shortest sequence of names wx such that x is an R-name witnessing condition 1 of Definition 5.5 and wx is a prefix of every control in $[c_l, l]_T$. The existence of invariants follows from the fact that new names are always appended to the right of the controls.

Invariants induce the following (reflexive) quasi-order on repeats of a proof: $l \preccurlyeq l'$ if $\mathsf{inv}(l)$ is a prefix of $\mathsf{inv}(l')$. The orders $\lhd$ and $\preccurlyeq$ are related in the sense of the following propositions, both of which are easily verified.

▶ **Proposition 6.1.** *For every infinite reachability sequence $l_0 \lhd l_1 \lhd \cdots$ there exists $k \geq 0$ such that $l_k \preccurlyeq l_j$ for all $j \geq k$.*

▶ **Proposition 6.2.** *If $l_0 \lhd l_1 \lhd \cdots \lhd l_m \lhd l_0$ and w is a prefix of $\mathsf{inv}(l_i)$ for each $i \leq m$, then w is a prefix of each control on $[c_{l_m}, l_0]_T$.*

The following lemmas concerning names fixed on paths through cyclic derivations are used to show soundness and completeness for $\mathsf{CTL}^*_\circ$. We omit the proofs for brevity.

▶ **Lemma 6.3.** *Let $\pi = (s_i)_{i\leq n}$ be a finite path through a $\mathsf{CTL}^*_\circ$ derivation. If names $x_1, \dots, x_m$ are fixed on π and there is a sequent $\mathsf{Q}\Phi$ in the label of s_n such that $x_1, \dots, x_m$ all occur in $\mathsf{Q}\Phi$, then there is a sequent trace τ on π such that $x_1, \dots, x_m$ are all fixed on τ.*

▶ **Lemma 6.4.** *Let $\pi = (s_i)_{i<\omega}$ be an infinite path through a $\mathsf{CTL}^*_\circ$ derivation. If names $x_1, \dots, x_m$ are fixed on π and for every $i < \omega$ there is a $j > i$ and a sequent $\mathsf{Q}\Phi$ in the label of s_j such that $x_1, \dots, x_m$ all occur in $\mathsf{Q}\Phi$, then there is an infinite sequent trace τ on π fixing $x_1, \dots, x_m$.*

The next lemmas are immediate consequences of the success condition on $\mathsf{CTL}^*_\circ$ proofs.

▶ **Lemma 6.5.** *Let T be a $\mathsf{CTL}^*_\circ$ proof, and let π be an infinite path through T°. There is an R-name x and an infinite sequent trace on π containing an infinite formula trace on which x is fixed.*

▶ **Lemma 6.6.** *Let T be a $\mathsf{CTL}^*_\circ$ derivation, τ an infinite sequent trace on T° of type E, and ρ an infinite formula trace on τ of type U. There is a U-name eventually fixed on ρ.*

Soundness of $\mathsf{CTL}^*_\circ$ now follows easily.

▶ **Proposition 6.7** (Soundness). *If there is a $\mathsf{CTL}^*_\circ$ proof of φ, then φ is valid.*

Proof. Let T be a $\mathsf{CTL}^*_\circ$ proof of φ. We show that the trace conditions on Definition 3.7 hold for T°. By the priority assigned to rule iThin, this suffices to ensure that the (possibly infinite) tree of paths on T°, once stripped of the annotations, is a $\mathsf{CTL}^*_\infty + \mathsf{eW}$ proof of φ.

Let π be an infinite path on T°. By Lemma 6.5 there is an infinite sequent trace τ on π containing an infinite formula trace ρ on which an R-name x is eventually fixed. If τ is of type A we are done, so suppose τ is of type E. Towards a contradiction, assume that τ contains an infinite formula trace ξ of type U. By Lemma 6.6, there is a U-name y which is eventually fixed on ξ. Then, there is a tail τ' of τ such that both x and y are fixed on τ'. Let π' be the tail of π corresponding to τ', and let l be a repeat encountered infinitely often on π'. By Proposition 5.1, every vertex on $[c_l, l]_T$ occurs infinitely often on π', so x and y are fixed on $[c_l, l]_T$ and some sequent in the label of l contains both x and y. By Lemma 6.3, there is a bad trace on $[c_l, l]_T$ on which x is fixed, contradicting the fact that l is successful. ◀

We now turn to the completeness proof for $\mathsf{CTL}^*_\circ$. An application of rule iThin is *trivial* if, using the same notation as in Figure 2, $v = \varnothing$. Due to the priority given to iThin in derivations, non-trivial instances of iThin satisfy the following.

▶ **Lemma 6.8.** *Let T be a $\mathsf{CTL}^*_\circ$ derivation, and let $s \in T$ be labelled by the premise of a non-trivial instance of* iThin*, say:*

$$\frac{\Theta' : \mathsf{Q}\{\Phi, \varphi^{x_1}\}, \Delta}{\Theta : \mathsf{Q}\{\Phi, \varphi^{x_1}, \varphi^{x_2}\}, \Delta}$$

Then, $\varphi = \mathsf{X}\psi$ for some $\psi = \psi_1 \mathsf{O}^X \psi_2$ and there is a vertex $t <_T s$ such that there is a sequent trace on $[t, s]_T$ of the form

$$\mathsf{Q}\{\Psi, \mathsf{X}\psi^{x_1}, \psi\} \rhd \mathsf{Q}\{\Psi', \mathsf{X}\psi^{x_1}, \mathsf{X}\psi^{x_2}\} \overset{\mathsf{iThin}\ \textit{only}}{\rhd \cdots \rhd} \mathsf{Q}\{\Phi, \mathsf{X}\psi^{x_1}, \mathsf{X}\psi^{x_2}\} \rhd \mathsf{Q}\{\Phi, \mathsf{X}\psi^{x_1}\}.$$

Proof. By the priority given to iThin and the fact that x_1 occurs before x_2 in Θ. ◀

▶ **Proposition 6.9** (Completeness). *If φ is valid, then there is a $\mathsf{CTL}^*_\circ$ proof of φ.*

Proof. Let T be a CTL^*_∞ proof of φ, and let $\mathsf{ann}(T)$ be the result of annotating T in accordance with the rules of $\mathsf{CTL}^*_\circ$, applying iThin and eThin whenever possible and always choosing $u \neq \varnothing$ when applying either ER_0 or ER_1. It suffices to show that every infinite branch on $\mathsf{ann}(T)$ contains a vertex satisfying the requirements of a successful repeat. Proposition 5.3 implies each branch of $\mathsf{ann}(T)$ contains only finitely many distinct annotated hypersequents.

Let $\mathsf{ann}(\pi)$ be an infinite branch on T, where $\pi = (s_i)_{i<\omega}$ is the corresponding branch on T. Let T^+ be the collection of all infinite sequent traces τ on $\mathsf{ann}(\pi)$ such that either τ is of type A and contains an infinite formula trace of type R, or τ is of type E and contains no infinite formula trace of type U. Since T is a proof, $T^+ \neq \varnothing$. Let T^- be the collection of all infinite sequent traces on $\mathsf{ann}(\pi)$ of type E that contain an infinite formula trace of type U.

We modify the annotations on $\mathsf{ann}(\pi)$ by applying the following procedure to every trace $\tau \in T^-$. Let $\tau = (\mathsf{Q}_i\Phi_i)_{i<\omega}$, and let $\rho = (\varphi_i^{u_i})_{i<\omega}$ be an infinite formula trace on τ of type U, say with dominating formula ψ^y. Let $n < \omega$ be such that $\varphi_i^{u_i} \in \{\psi^y, \mathsf{X}\psi^y\}$ for all $i \geq n$. For every application of ER on $\pi_{\geq n} = (s_i)_{i\geq n}$ with principal sequent in τ, we remove the annotation (using instead the rule instance with $u = \varnothing$) and propagate this change upwards appropriately. Let $\tilde{\pi} = (\tilde{s}_i)_{i<\omega}$ be the result after applying this procedure to every $\tau \in T^-$.

We claim that after all these changes the result remains an (annotated) derivation. This is clear for all rules except iThin and eThin. Each application of eThin will either vanish (because the two principal sequents now coincide), or remain an instance of eW. And every trivial instance of iThin will remain so or else vanish. Finally, consider a non-trivial instance of iThin in $\mathsf{ann}(\pi)$:

$$\frac{\Theta' : \mathsf{Q}\{\Psi, \psi^{x_1}\}, \Delta}{\Theta : \mathsf{Q}\{\Psi, \psi^{x_1}, \psi^{x_2}\}, \Delta}\ x_1 \prec_\Theta x_2 \tag{1}$$

The only way for this rule application to cease to be an instance of a $\mathsf{CTL}^*_\circ$ rule after the changes performed is if x_1 has been removed and x_2 remains. We claim this is impossible. Let s be the vertex on $\mathsf{ann}(\pi)$ corresponding to the premise of (1). By Lemma 6.8, $\psi = \mathsf{X}\chi$ for some formula $\chi = \chi_1 \mathsf{R}^X \chi_2$ (since x_1 is an R-name) and there is a vertex t below s such that there is a sequent trace on $[t, s]_{\mathsf{ann}(\pi)}$ of the form

$$\mathsf{Q}\{\Sigma, \mathsf{X}\chi^{x_1}, \chi\} \rhd \mathsf{Q}\{\Sigma', \mathsf{X}\chi^{x_1}, \mathsf{X}\chi^{x_2}\} \overset{\text{iThin only}}{\rhd \cdots \rhd} \mathsf{Q}\{\Psi, \mathsf{X}\chi^{x_1}, \mathsf{X}\chi^{x_2}\} \rhd \mathsf{Q}\{\Psi, \mathsf{X}\chi^{x_1}\}.$$

So by the time x_2 is introduced x_1 has already been removed, whence by construction of $\tilde{\pi}$ we also removed x_2 and thus (1) simply vanishes on $\tilde{\pi}$.

Finally, we show that $\tilde{\pi}$ passes through a successful repeat. Fix a trace $\tau = (\mathsf{Q}_i\Phi_i)_{i<\omega} \in T^+$ (note that $\tau \notin T^-$), and let $\rho = (\varphi_i^{u_i})_{i<\omega}$ be an infinite formula trace on τ of type R, say with dominating formula ψ^x. Let $n < \omega$ be such that

1. $\varphi_i^{u_i} \in \{\psi^x, \mathsf{X}\psi^x\}$ for every $i \geq n$, and
2. every pair $(\Theta_i : \Gamma_i, \mathsf{Q}_i\Phi_i)$ for $i \geq n$ is encountered infinitely often on $\{(\Theta_j : \Gamma_j, \mathsf{Q}_j\Phi_j)\}_{j<\omega}$, where $\Theta_j : \Gamma_j$ is the label of $\tilde{s}_j$.

Let $S = \{(\mathsf{E}\Psi_i, y_i, z_i)\}_{i\in I}$ be the collection of all triples $(\mathsf{E}\Psi_i, y_i, z_i)$ where $\mathsf{E}\Psi_i$ is a sequent in the label of $\tilde{s}_n$, y_i is a U-name occurring in $\mathsf{E}\Psi_i$, and z_i is an R-name occurring in $\mathsf{E}\Psi_i$. Note that S is finite. For every $i \in I$ there is a $j_i \geq n$ such that there is no sequent trace on $[\tilde{s}_n, \tilde{s}_{j_i}]_{\tilde{\pi}}$ starting from $\mathsf{E}\Psi_i$ and where y_i and z_i are both fixed. Otherwise by Lemma 6.4 there would be an infinite sequent trace τ' on $\tilde{\pi}$ of type E such that y_i and z_i are both fixed on τ', contradicting the construction of $\tilde{\pi}$. Let $m > \max\{n, \max\{j_i : i \in I\}\}$ be such that $\tilde{s}_n$ and $\tilde{s}_m$ have identical labels. We claim that $\tilde{s}_m$ is a successful repeat with companion $\tilde{s}_n$.

Let $\tau_n^m := (\mathsf{Q}_i\Phi_i)_{n\leq i\leq m}$. We know that x is fixed on τ_n^m, and by the choice of m no U-name is fixed on τ_n^m, so τ_n^m is a good trace on $[\tilde{s}_n, \tilde{s}_m]_{\tilde{\pi}}$. Towards a contradiction, suppose there is a bad trace ξ on $[\tilde{s}_n, \tilde{s}_m]_{\tilde{\pi}}$ such that an R-name z is fixed on ξ. By definition, there is a U-name y fixed on ξ. Let $\mathsf{E}\Psi$ be the first sequent on ξ. Then, $(\mathsf{E}\Psi, y, z) \in S$, say $(\mathsf{E}\Psi, y, z) = (\mathsf{E}\Psi_k, y_k, z_k)$. So ξ restricted to $[\tilde{s}_n, \tilde{s}_{j_k}]_{\tilde{\pi}}$ is a sequent trace starting from $\mathsf{E}\Psi_k$ and on which both y_k and z_k are fixed, contradicting the choice of j_k. ◀

Combining Proposition 6.9 and Proposition 6.7 we have

▶ **Theorem 6.10.** *A formula φ is valid iff there is a $\mathsf{CTL}^*_\circ$ proof of φ.*

7 A decision procedure for the universal fragment of CTL*

The completeness proofs for the cyclic and ill-founded calculi provide a deterministic proof-search procedure which always yields a proof if the initial formula is valid. The argument leaves open the question of whether validity can be decided via proof-search, i.e., that no proof exists if none has been found within sufficiently many steps. In this section we provide a positive answer for the *universal* fragment of CTL*, that is, for formulas containing no existential quantifier.

Given a universal formula φ, we build a finite proof-search tree for φ, annotate it according to the rules of the cyclic calculus, and show that if it does not contain a $\mathsf{CTL}^*_\circ$ proof of φ then φ is not valid.

▶ **Definition 7.1.** *An* annotated proof-search tree *for a universal formula φ is a finite tree T whose vertices are labelled according to the rules in Figure 2 and Figure 3 and such that*

1. *The root of T has label $\mathsf{A}\varphi$.*
2. *A vertex $u \in T$ is a leaf iff either*
   - **a.** *u is axiomatic,*
   - **b.** *u is labelled by literal sequents only, or*
   - **c.** *there is a vertex $v <_T u$ such that u and v have identical labels.*
3. *In place of rule AX, the branching rule AX^*_b is used in T:*

$$\mathsf{AX}^*_b\ \frac{\Theta_1 : \mathsf{A}\Phi_1 \qquad \cdots \qquad \Theta_n : \mathsf{A}\Phi_n}{\Theta : \mathsf{AX}\Phi_1, \ldots, \mathsf{AX}\Phi_n, \Lambda}$$

   where Λ is a non-axiomatic set of literal sequents and Θ_i is the result of removing from Θ all names not occurring in $\mathsf{A}\Phi_i$.
4. *In every application of rule ALit in T, say with principal formula λ and principal sequent $\mathsf{A}\{\Phi, \lambda\}$, we have $\Phi \setminus \{\lambda\} \neq \varnothing$.*
5. *All instances of eW in T are instances of eThin.*
6. *The thinning rules iThin and eThin are prioritised over the rest in T.*

As a consequence of Proposition 5.3, we have:

▶ **Proposition 7.2.** *For every universal formula φ there exists an annotated proof-search tree for φ.*

In the absence of existential quantifiers, the success condition for repeats drastically simplifies.

▶ **Proposition 7.3.** *A repeat l in an annotated proof-search tree T is successful iff there is an R-name fixed on $[c_l, l]_T$.*

Proof. Since unfoldings of U-formulas under a universal quantifier are never annotated, there are no U-names in T, whence l is successful iff there is a sequent trace on $[c_l, l]_T$ where an R-name is fixed. By Lemma 6.3, this is the case iff there is an R-name fixed on $[c_l, l]_T$. ◀

Finally, we show that if an annotated proof-search tree for φ does not contain a $\mathsf{CTL}^*_\circ$ proof of φ, then φ is not valid.

▶ **Proposition 7.4.** *Let T be an annotated proof-search tree for a universal formula φ. If T cannot be pruned at AX^*_b-vertices down to a $\mathsf{CTL}^*_\circ$ proof of φ, then φ is not valid.*

Proof. Starting bottom-up, let T' be the result of pruning T at $\mathsf{A}\wedge$-, AU- and AR-vertices by keeping the subtree that fails to produce a proof. Note that branching in T' is due solely to AX^*_b, that there are no axiomatic leaves in T', and that every repeat in T' is unsuccessful. So, by Proposition 7.3, for every repeat $l \in T'$ the following holds: there is no R-name fixed on $[c_l, l]_{T'}$. It follows that there cannot be an infinite sequent trace of type A containing an infinite formula trace of type R on $(T')^\circ$, because unfoldings of R-formulas under a universal quantifier are always annotated. Therefore, stripping $(T')^\circ$ of the annotations yields a refutation of φ, whence φ is not valid by Proposition 4.7. ◀

Proposition 7.4 yields a decision procedure for validity of universal formulas: given a universal formula φ, build an annotated proof-search tree for φ and check whether it contains a $\mathsf{CTL}^*_\circ$ proof of φ. This procedure fails in the presence of existential quantifiers. Using the notation from the proof of Proposition 7.4, T' may contain an unsuccessful repeat l with a good sequent trace of type A on $[c_l, l]_{T'}$. Thus, we cannot guarantee that $(T')^\circ$, once stripped of the annotations, is a refutation.

8 Conclusion

We introduce a sound and complete cyclic hypersequent calculus for Full Computation Tree Logic CTL^* and a decision procedure for validity of the universal fragment. Hypersequents – sets of sets of formulas – offer a natural framework for accommodating the existential and universal path quantifiers of the logic. Each "sequent" in a hypersequent is a labelled set of formulas, either $\mathsf{A}\Phi$ or $\mathsf{E}\Phi$, interpreted as *along all paths* $\bigvee \Phi$ and *along some path* $\bigwedge \Phi$ respectively. Through this interpretation, a natural system of ill-founded proofs arises wherein every infinite path of a proof must contain either an infinite sequent trace of type A through which *some* infinite formula trace stabilises (on a release operator), or an infinite trace of type E in which *all* infinite formula traces stabilise. Correctness conditions of the latter kind are rare in ill-founded (and cyclic) proof calculi. Indeed, together with [9], which employs a similar trace condition, these appear to be the only examples of cyclic systems that fall outside the scope of the category theoretic notion of cyclic proof introduced in [5].

In contrast to the ill-founded calculus, correctness of cyclic proofs is determined by the simple cycles only, i.e., the shortest path between leaf and companion. Soundness is ensured by annotating formulas in the cyclic calculus in a manner similar to [2, 3, 14, 17, 30]. As a result, proof-checking a cyclic derivation is linear time, in contrast to the trace condition along all paths which is PSPACE complete in general [21, 5].[3] Even so, the annotation mechanism necessary for CTL^* is significantly simpler than those developed for the μ-calculus and related logics. For example, each formula is annotated by at most one name and vice versa.

[3] It should be noted, however, that the trace condition for regular ill-founded CTL^* proofs in Definition 3.7 does not directly fit within the known PSPACE completeness results.

Future directions include investigating whether an analytic calculus for CTL* can assist with proving completeness of Hilbert-style calculi. Also of interest is the development of robust proof systems for related logics such as hybrid [28], graded [19], memory-full [16], and multi-agent [6] extensions of CTL*.

References

1 Pietro Abate, Rajeev Goré, and Florian Widmann. One-pass tableaux for computation tree logic. In Nachum Dershowitz and Andrei Voronkov, editors, *Logic for Programming, Artificial Intelligence, and Reasoning, 14th International Conference, LPAR 2007, Yerevan, Armenia, October 15-19, 2007, Proceedings*, volume 4790 of *Lecture Notes in Computer Science*, pages 32–46. Springer, 2007. doi:10.1007/978-3-540-75560-9_5.

2 Bahareh Afshari, Sebastian Enqvist, and Graham E Leigh. Cyclic proofs for the first-order μ-calculus. *Logic Journal of the IGPL*, 2022. doi:10.1093/jigpal/jzac053.

3 Bahareh Afshari and Graham E. Leigh. Circular proofs for the modal mu-calculus. *PAMM*, 16(1):893–894, 2016. doi:10.1002/pamm.201610435.

4 Bahareh Afshari and Graham E. Leigh. Cut-free completeness for modal mu-calculus. In *32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, Reykjavik, Iceland, June 20-23, 2017*, pages 1–12. IEEE Computer Society, 2017. doi:10.1109/LICS.2017.8005088.

5 Bahareh Afshari and Dominik Wehr. Abstract cyclic proofs. In Agata Ciabattoni, Elaine Pimentel, and Ruy J. G. B. de Queiroz, editors, *Logic, Language, Information, and Computation - 28th International Workshop, WoLLIC 2022, Iaşi, Romania, September 20-23, 2022, Proceedings*, volume 13468 of *Lecture Notes in Computer Science*, pages 309–325. Springer, 2022. doi:10.1007/978-3-031-15298-6_20.

6 Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. *J. ACM*, 49(5):672–713, 2002. doi:10.1145/585265.585270.

7 Kai Brünnler and Martin Lange. Cut-free sequent systems for temporal logic. *J. Log. Algebraic Methods Program.*, 76(2):216–225, 2008. doi:10.1016/j.jlap.2008.02.004.

8 Mads Dam. Translating CTL* into the modal μ-calculus. *Tech. rep. ECS-LFCS-90-123, University of Edinburgh, Department of Computer Science, Laboratory for Foundations of Computer Science*, 1990.

9 Anupam Das and Marianna Girlando. Cyclic proofs, hypersequents, and transitive closure logic. In Jasmin Blanchette, Laura Kovács, and Dirk Pattinson, editors, *Automated Reasoning - 11th International Joint Conference, IJCAR 2022, Haifa, Israel, August 8-10, 2022, Proceedings*, volume 13385 of *Lecture Notes in Computer Science*, pages 509–528. Springer, 2022. doi:10.1007/978-3-031-10769-6_30.

10 Stéphane Demri, Valentin Goranko, and Martin Lange. *Temporal Logics in Computer Science: Finite-State Systems.* Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2016. doi:10.1017/CBO9781139236119.

11 Sebastian Enqvist. A circular proof system for the hybrid μ-calculus. In Nicola Olivetti, Rineke Verbrugge, Sara Negri, and Gabriel Sandu, editors, *13th Conference on Advances in Modal Logic, AiML 2020, Helsinki, Finland, August 24-28, 2020*, pages 169–188. College Publications, 2020.

12 Oliver Friedmann, Markus Latte, and Martin Lange. A decision procedure for CTL* based on tableaux and automata. In *International Joint Conference on Automated Reasoning*, pages 331–345. Springer, 2010.

13 Dov M. Gabbay and Amir Pnueli. A sound and complete deductive system for CTL* verification. *Logic Journal of the IGPL*, 16(6):499–536, 2008.

14 Natthapong Jungteerapanich. A tableau system for the modal μ-calculus. In Martin Giese and Arild Waaler, editors, *Automated Reasoning with Analytic Tableaux and Related Methods, 18th International Conference, TABLEAUX 2009, Oslo, Norway, July 6-10, 2009. Proceedings*, volume 5607 of *Lecture Notes in Computer Science*, pages 220–234. Springer, 2009.

15 Ioannis Kokkinis and Thomas Studer. Cyclic proofs for linear temporal logic. *Ontos Mathematical Logic*, 6:171–192, 2016.

16 Orna Kupferman and Moshe Y. Vardi. Memoryful branching-time logic. In *21th IEEE Symposium on Logic in Computer Science (LICS 2006), 12-15 August 2006, Seattle, WA, USA, Proceedings*, pages 265–274. IEEE Computer Society, 2006. doi:10.1109/LICS.2006.34.

17 Johannes Marti and Yde Venema. A focus system for the alternation-free μ-calculus. In Anupam Das and Sara Negri, editors, *Automated Reasoning with Analytic Tableaux and Related Methods - 30th International Conference, TABLEAUX 2021, Birmingham, UK, September 6-9, 2021, Proceedings*, volume 12842 of *Lecture Notes in Computer Science*, pages 371–388. Springer, 2021. doi:10.1007/978-3-030-86059-2_22.

18 Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, 1975.

19 Faron Moller and Alexander Moshe Rabinovich. Counting on CTL*: On the expressive power of monadic path logic. *Inf. Comput.*, 184(1):147–159, 2003. doi:10.1016/S0890-5401(03)00104-4.

20 Damian Niwiński and Igor Walukiewicz. Games for the μ-calculus. *Theoretical Computer Science*, 163(1-2):99–116, 1996.

21 Rémi Nollet, Alexis Saurin, and Christine Tasson. PSPACE-Completeness of a thread criterion for circular proofs in linear logic with least and greatest fixed points. In S. Cerrito and A. Popescu, editors, *Automated Reasoning with Analytic Tableaux and Related Methods*, Lecture Notes in Computer Science, pages 317–334. Springer International Publishing, Cham, 2019. doi:10.1007/978-3-030-29026-9_18.

22 Mark Reynolds. An axiomatization of full computation tree logic. *The Journal of Symbolic Logic*, 66(3):1011–1057, 2001.

23 Mark Reynolds. An axiomatization of PCTL*. *Inf. Comput.*, 201(1):72–119, 2005. doi:10.1016/j.ic.2005.03.005.

24 Mark Reynolds. A tableau for CTL*. In Ana Cavalcanti and Dennis R. Dams, editors, *Formal Methods*, pages 403–418. Springer Berlin Heidelberg, 2009.

25 Mark Reynolds. A tableau-based decision procedure for CTL*. *Formal aspects of computing*, 23(6):739–779, 2011.

26 Jan Rooduijn. Cyclic hypersequent calculi for some modal logics with the master modality. In Anupam Das and Sara Negri, editors, *Automated Reasoning with Analytic Tableaux and Related Methods - 30th International Conference, TABLEAUX 2021, Birmingham, UK, September 6-9, 2021, Proceedings*, volume 12842 of *Lecture Notes in Computer Science*, pages 354–370. Springer, 2021. doi:10.1007/978-3-030-86059-2_21.

27 Jan Rooduijn and Lukas Zenger. An analytic proof system for common knowledge logic over S5. In *Proceedings of Advances in Modal Logic (AiML 2022)*, (to appear).

28 Ulrike Sattler and Moshe Y. Vardi. The hybrid μ-calculus. In Rajeev Goré, Alexander Leitsch, and Tobias Nipkow, editors, *Automated Reasoning, First International Joint Conference, IJCAR 2001, Siena, Italy, June 18-23, 2001, Proceedings*, volume 2083 of *Lecture Notes in Computer Science*, pages 76–91. Springer, 2001.

29 Yury Savateev and Daniyar S. Shamkanov. Non-well-founded proofs for the Grzegorczyk modal logic. *Rev. Symb. Log.*, 14(1):22–50, 2021. doi:10.1017/S1755020319000510.

30 Colin Stirling. A tableau proof system with names for modal mu-calculus. In Andrei Voronkov and Margarita V. Korovina, editors, *HOWARD-60: A Festschrift on the Occasion of Howard Barringer's 60th Birthday*, volume 42 of *EPiC Series in Computing*, pages 306–318. EasyChair, 2014.