

Cyclic Proof Systems for Modal Fixpoint Logics

$$\begin{array}{c}
 \frac{p, \Box\varphi \Rightarrow p}{\varphi, \Box\varphi, p \Rightarrow \Box p, p} \quad \frac{\varphi, \Box\varphi, p \Rightarrow \Box p, p}{\varphi, \Box\varphi \Rightarrow p \rightarrow \Box p, p} \quad \frac{\varphi, \Box\varphi, p \Rightarrow p \quad \varphi, \Box\varphi \Rightarrow p}{\varphi, \Box\varphi, p \Rightarrow \Box p} \\
 \frac{\varphi, \Box\varphi, p \Rightarrow \Box p, p}{\varphi, \Box\varphi \Rightarrow p \rightarrow \Box p, p} \quad \frac{\varphi, \Box\varphi, p \Rightarrow \Box p}{\varphi, \Box\varphi \Rightarrow p \rightarrow \Box p} \\
 \frac{p, \Box\varphi \Rightarrow p}{\Box\varphi \Rightarrow \Box(p \rightarrow \Box p), p} \quad \frac{\varphi, \Box\varphi, p \Rightarrow p}{\Box[\underbrace{\Box(p \rightarrow \Box p) \rightarrow p}_{\varphi}] \Rightarrow p} \\
 \leftarrow
 \end{array}$$

Guillermo Menéndez Turata

Cyclic Proof Systems for Modal Fixpoint Logics

Guillermo Menéndez Turata

Cyclic Proof Systems for Modal Fixpoint Logics

ILLC Dissertation Series DS-2023-11



INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION

For further information about ILLC-publications, please contact

Institute for Logic, Language and Computation
Universiteit van Amsterdam
Science Park 107
1098 XG Amsterdam
phone: +31-20-525 6051
e-mail: illc@uva.nl
homepage: <http://www.illc.uva.nl/>

Copyright © 2023 by Guillermo Menéndez Turata

Cover design by Guillermo Menéndez Turata.
Printed and bound by Ipskamp Printing.

ISBN: 978-94-6473-347-1



Cyclic Proof Systems for Modal Fixpoint Logics

ACADEMISCH PROEFSCHRIFT

ter verkrijging van de graad van doctor
aan de Universiteit van Amsterdam
op gezag van de Rector Magnificus
prof. dr. ir. P.P.C.C. Verbeek
ten overstaan van een door het College voor Promoties ingestelde commissie,
in het openbaar te verdedigen in de Agnietenkapel
op dinsdag 30 januari 2024, te 13.00 uur

door Guillermo Menéndez Turata
geboren te Tarragona

Promotiecommissie

Promotor: prof. dr. Y. Venema Universiteit van Amsterdam

Copromotor: dr. B. Afshari Göteborgs universitet

Overige leden: dr. A. Baltag Universiteit van Amsterdam
dr. N. Bezhnashvili Universiteit van Amsterdam
prof. dr. U. Endriss Universiteit van Amsterdam
prof. dr. R. Iemhoff Universiteit Utrecht
prof. dr. ing. R.A.M. van Rooij Universiteit van Amsterdam
dr. A.F.E.L.E.W.M. Saurin Université Paris Cité
dr. T. Studer Universität Bern

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

LOGIC, *n.* The art of thinking and reasoning in strict accordance with the limitations and incapacities of the human misunderstanding.

Ambrose Bierce, *The Devil's Dictionary*

Desenterró la intolerable hipótesis griega de la eterna repetición y procuró educir de esa pesadilla mental una ocasión de júbilo. Buscó la idea más horrible del universo y la propuso a la delectación de los hombres.

Jorge Luis Borges, «La doctrina de los ciclos»

Contents

Acknowledgments	xi
Introduction	1
Main contributions	4
Sources of the material	5
1 Mathematical Preliminaries	7
1.1 Sequences and trees	7
1.2 Trees with back-edges	11
1.3 Sequent calculi	16
1.4 The Borel hierarchy	18
1.5 Gale–Stewart games	20
1.6 Büchi automata	22

Propositional Modal Logics

2 Cyclic Companions of Modal Logics	27
2.1 Propositional normal modal logics	31
2.2 Cycles and infinite chains	36
2.2.1 GL as a cyclic companion of K4	38
2.2.2 S4Grz as a cyclic companion of S4	40
2.3 K4Grz as a cyclic companion of K4	44
2.3.1 The finite model property and decidability	56
2.4 Conclusion	57

Temporal Logics

3	A Cyclic Proof System for CTL^*	63
3.1	Full computation tree logic CTL^*	65
3.2	The ill-founded system CTL^*_∞	68
3.2.1	Soundness of CTL^*_∞	76
3.2.2	Completeness of CTL^*_∞	84
3.3	Regular ill-founded proofs	95
3.4	The cyclic system CTL^*_\circ	98
3.4.1	Soundness of CTL^*_\circ	105
3.4.2	Completeness of CTL^*_\circ	109
3.5	Cycle elimination	114
3.5.1	Inductive acyclic proofs	115
3.5.2	Inductible cyclic proofs	125
3.5.3	Hilbert-style proofs	136
3.6	Conclusion	157
4	A Cyclic Proof System for iLTL	159
4.1	Intuitionistic linear-time temporal logic iLTL	160
4.2	The ill-founded system iLTL_∞	162
4.2.1	Soundness of iLTL_∞	171
4.2.2	Completeness of iLTL_∞	175
4.3	Regular ill-founded proofs	204
4.4	The cyclic system iLTL_\circ	206
4.4.1	Soundness of iLTL_\circ	211
4.4.2	Completeness of iLTL_\circ	213
4.5	Conclusion	217

Cyclic Proof Systems and Interpolation

5	Uniform Interpolants from Cyclic Proofs	221
5.1	The modal μ -calculus	222
5.2	The Jungteerapanich–Stirling proof system	227
5.3	Uniform interpolation	232
5.3.1	Verification	236
5.4	Conclusion	247

Conclusion	249
-------------------	------------

Bibliography	253
---------------------	------------

Index	269
List of Figures	277
List of Tables	279
List of Symbols	281
Samenvatting	285
Abstract	287

Acknowledgments

This thesis is the culmination of four years of doctoral studies. Naturally, many people have contributed to making it possible.

First and foremost, I am grateful to Bahareh Afshari and Yde Venema for their guidance and patience, especially during the last stages of the writing of the thesis. If I am capable of conducting any academic research at all, I owe that to both of you.

On this note, gratitude is also due to Graham E. Leigh for his input on many a topic.

I thank the members of the defence committee for agreeing to read and evaluate the dissertation: Alexandru Baltag, Nick Bezhanishvili, Ulle Endriss, Rosalie Iemhoff, Robert van Rooij, Alexis Saurin, and Thomas Studer.

The ILLC would not be the comfortable place that it is were it not for its administrative staff. I am thankful to all of them for their help.

I thank Nick Bezhanishvili for giving me the opportunity to lecture in the Master of Logic.

Over the last few years I have met many nice people in Amsterdam and abroad. Special thanks to the following, in strictly random order: Aleksandar, Alyssa, Armand, Avgerinos, Claudia, Cristina, Dominik, Drew, Elisa, Iris, Jan, Johannes, Kirill, Lukas, Marianna, Mattias, Noel, Orvar, Pablo, Pedro, Ramón, Raufs, Teresa, Tjeerd, Vlad, Xavier. (Reader, if your name has been unjustly omitted, let me know and I will include it in the second, revised and expanded edition of this bestseller.)

Jinghui, thank you for the time that we spent together, always too short. I remember our discussing Chinese poetry on a French island (including your reaction when you first saw the Côte d'Azur), and traversing endless, perfectly flat Dutch fields.

Lwenn, you have always been there ready to deconstruct me. You managed to make it fun to go to work every day from nine to five: the office — nay, the ILLC — never quite recovered from your departure. I appreciate the advice that you gave me regarding the writing of the thesis. Whenever I think of you, I am reminded of the verses you once taught me: ‘One for sorrow, / Two for joy’; thank you for the joy that you brought to my life.

Giacomo, grazie per la tua amicizia e il tuo senso dell’umorismo. Spero di rivederti presto!

Lide, you once asked me: ‘Are you happy you met me?’ Let this serve as a definitive answer. « L’homme seul est quelque chose d’imparfait ; il faut qu’il trouve un second pour être heureux. Il le cherche le plus souvent dans l’égalité de la condition, à cause que la liberté et que l’occasion de se manifester s’y rencontrent plus aisément. » Vaarwel.

Mi familia se ha encargado de que cuatro años en el extranjero no hayan hecho de mí un extraño. Gracias a ellos puedo decir que, com a casa, enlloc.

Raúl y Borja, os agradezco que os ofrecierais como paraninfos pese a las molestias logísticas que ello conlleva. Echaré de menos nuestras conversaciones. Una vez incluso hablamos de lógica; ¿me pusisteis algo en la bebida? Con vosotros de guardaespaldas, puedo desmayarme tranquilo durante la defensa.

Amsterdam
November, 2023

G. M. T.

Introduction

Cyclic and ill-founded proof theory allow infinity to enter the realm of proofs, the one corner of mathematics in which everyone seems to agree that it should not be welcome. Indeed: the number line may be infinite and infinitely dense; numbers may have infinite, aperiodic decimal expansions; even circles may have infinite radii; but, surely, proofs ought to be finite.

Infinitely long proofs, however, have been part of mathematics since antiquity, in the form of proofs by infinite descent. The term was coined by Fermat (see, e.g., [159, Ch. II, § X]), but the technique is already used in Euclid's proof that every composite number is divided by some prime number (Prop. 31 of Book 7 of the *Elements*). A proof by infinite descent, as the name implies, is infinite. Its 'form', however, allows us to grasp that it is not a vicious circle. For even though we restart the argument after finitely many steps, we do so from a place which occupies a strictly lower position than the previous one (e.g., a smaller natural number). There has been some 'progress', and thus the argument is not so much circular as it is spiral: we move in circles, but ascending (or, rather, descending) towards a sound proof. Moreover, once we 'see' that the limit of the argument is an infinite descending chain of, say, natural numbers, we confidently reject the starting hypothesis without having to continue reasoning ad infinitum.

Ill-founded proof theory, and its more presentable sibling cyclic proof theory, formalise this idea by admitting infinitely long proofs in sequent calculi whose infinite branches satisfy some 'progress condition' ensuring that they yield valid conclusions.

This thesis designs and manipulates cyclic and ill-founded calculi for several modal fixpoint logics. Let us then briefly introduce both concepts, cyclic proofs and modal fixpoint logics, as an appetizer for the chapters that follow.

Cyclic proofs. Proofs by induction are difficult to mechanise because of the need to guess the right inductive hypothesis. It is not uncommon, when proving a statement by induction, to need a strengthening of the ‘natural’ inductive hypothesis. A well known example is provided by the following statement

$$1 + \frac{1}{4} + \frac{1}{9} + \cdots + \frac{1}{n^2} < 2,$$

which clearly does not admit a direct inductive proof. But we can easily prove by induction the *stronger* statement:

$$1 + \frac{1}{4} + \frac{1}{9} + \cdots + \frac{1}{n^2} \leq 2 - \frac{1}{n}.$$

This is not an isolated case, but an instance of a general phenomenon. Inductive constructions and proofs are ultimately based on least fixpoints of (monotone) ‘constructor’ maps. Logic itself provides plenty of inductive definitions in which this is manifest. For example, each time that we define the set of (well-formed) formulas over an alphabet Σ , we characterise it as the *least* collection of strings over Σ closed under certain properties (e.g., if α is a formula, then so is $\neg\alpha$).

Relying on informal notation, the least fixpoint operator of a monotone map f on a complete lattice is characterised by the following axiom and rule, where $\text{LFP } f$ denotes the least fixpoint of f :¹

$$\frac{}{f(\text{LFP } f) \leq \text{LFP } f} \qquad \frac{f(x) \leq x}{\text{LFP } f \leq x}$$

The axiom on the left says that $\text{LFP } f$ is a pre-fixpoint of f . The rule on the right is the fixpoint induction principle of Park’s [105]. Together, they characterise $\text{LFP } f$ as the least pre-fixpoint of f , hence also the least fixpoint by the Knaster–Tarski theorem [150].

Looking at Park’s rule, we can see at once where the problem lies with induction: the rule is not invertible. Consider, for example, the set E of all even numbers. It is the least fixpoint of the map $f : 2^\omega \rightarrow 2^\omega : X \mapsto X \cup \{0\} \cup \{n + 2 \mid n \in X\}$, and $E \subseteq E \cup \{1\}$. However, $E \cup \{1\}$ is not closed under f .

From the point of view of proof-search, right inductive invariants can be as hard to guess as right cut formulas. While several automated theorem provers rely on induction in one form or another, they all encounter difficult challenges (see, e.g., [17, § 1.2]). Cyclic and ill-founded proof theory propose an alternative approach: instead of reasoning by induction, proceed by infinite descent in an analytic calculus, for example by means of fixpoint unfoldings.

¹Existence of $\text{LFP } f$ is guaranteed by the well-known Knaster–Tarski theorem [150], for f is monotone.

It is usually accepted that infinite descent is as strong as induction. Making this claim precise, however, is far from straightforward. The reader may consult [139, 17, 18, 137, 103] for recent approximations to the formalisation and clarification of this question. Abstract approaches to cyclic theory, aiming at a general, logic-agnostic notion of ‘cyclic proof’, are pursued in [17, 2] (see also [84]).

Modal fixpoint logics. At the risk of stating the obvious, there can be no cyclic or ill-founded proofs in a calculus if it lacks the power to create infinitely long branches in a ‘non-trivial’ manner.

Temporal modalities, like ‘eventually’, ‘until’, ‘henceforth’, etc., are a natural source of fixpoint equivalences capable of producing infinite proofs. Consider, for example, the eventually modality Fp , with the meaning: ‘either p holds now, or it will hold at some point in the future’. This is equivalent to saying: ‘either p is the case now, or Fp will be the case at the next instant of time’. In symbols, $Fp \equiv p \vee XFp$, where X is the temporal operator ‘next’. Therefore, in place of an induction rule we may work with the following fixpoint unfolding rule:

$$\frac{\varphi \vee XF\varphi}{F\varphi}$$

By imposing a correctness condition on infinite branches, corresponding to the fact that $F\varphi$ requires that φ eventually be the case, we can ensure that our calculus does not yield invalid conclusions.

This is probably the oldest occurrence of fixpoints in modal logic, if we take this expression to include pre-modern times. Indeed, the ancient Greek logician Diodorus Cronus conceived possibility (and necessity) in temporal terms: he identified the possible with that which is the case or will eventually be the case (so the F operator above). His ideas had great influence in the work of Arthur N. Prior, the father of (modern) temporal logic [53, 110, 111, 104]. In more recent years, cyclic and ill-founded proof-theoretical approaches, based on rules like the one above for F , have met with success in dealing with temporal logics [19, 117, 47, 118].

Another source of fixpoints is provided by logics such as GL and $S4Grz$. This is less obvious, because neither logic includes operators of a distinct fixpoint nature (despite both of them enjoying well-known fixpoint theorems). Nevertheless, both involve infinitary frame conditions on infinite chains that make it possible to design natural cyclic or ill-founded calculi for them.² For recent developments in the cyclic and ill-founded proof theory of propositional modal logics, the reader may consult [134, 128, 129, 130, 69, 120, 122, 136, 35].

There is yet another source of fixpoints in modal logic. Instead of operators which happen to satisfy fixpoint equivalences, one can introduce operators specifically designed to denote fixpoints. The paradigm of this approach is

²We elaborate on this idea in Chapter 2.

Kozen’s modal μ -calculus [81], an extension of propositional (multi-)modal logic with explicit fixpoint quantifiers μ and ν . Syntactically, they behave like quantifiers in predicate logic, in the sense that they bound occurrences of variables in formulas. Semantically, μ and ν denote, respectively, least and greatest fixed points of functions, and are thus a kind of monadic second-order quantifiers. The μ -calculus, informally, has one foot on modal logic and the other on monadic second-order logic. Much success has been achieved in the cyclic and ill-founded proof theory of the μ -calculus and related fixpoint logics (see, e.g., [102, 73, 146, 76, 144, 9, 43, 95, 4, 37, 121, 127]).

Main contributions

Adding cycles to ordinary sequent calculi for **K4** and **S4** yields cyclic proof systems for the Gödel–Löb provability logic (**GL**) and the Grzegorzcyk logic (**S4Grz**), respectively [134, 130]. Rather than isolated contrivances, we argue in Chapter 2 that these systems arise from a natural correspondence between cycles in proofs and infinite chains in frames that enables the former to capture frame conditions involving the latter. We propose to understand *cyclic companionship* by a combination of proof-theoretic and semantical considerations. According to our explanation, one should be able to obtain a cyclic system for the weak Grzegorzcyk logic **K4Grz** by adding cycles to a system for **K4**. We show that this is indeed the case.

In the first part of Chapter 3, we introduce a cut-free, cyclic hypersequent calculus for the full computation tree logic **CTL***. Local soundness of inferences is immediate, and a global correctness condition ensures that cycles yield valid conclusions. Hypersequents offer a natural framework for accommodating the existential (**E**) and universal (**A**) path quantifiers of the logic, as well as their interplay with the next operator **X**. Each ‘sequent’ in a hypersequent is a labelled set of formulas, either **A** Φ or **E** Φ , interpreted as ‘along all paths, $\bigvee \Phi$ is the case’ and ‘along some path, $\bigwedge \Phi$ is the case’, 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.

A simple annotation mechanism on formulas allows us to isolate a finitary condition which suffices to guarantee that a derivation is a proof.

In the second half of Chapter 3, we isolate a class of ‘inductible’ cyclic proofs whose cycles can be transformed into inductive arguments based on the following Park-style characterisation of until:

$$\frac{}{(\beta \vee (\alpha \wedge \mathbf{X}(\alpha \mathbf{U}\beta))) \rightarrow \alpha \mathbf{U}\beta} \qquad \frac{(\beta \vee (\alpha \wedge \mathbf{X}\gamma)) \rightarrow \gamma}{\alpha \mathbf{U}\beta \rightarrow \gamma}$$

In the end, we arrive at a Hilbert-style system and compare it to a fragment of a known axiomatisation for the full logic. Our axiom system is complete for a well-known variant of CTL* obtained by allowing the evaluation of formulas in a bigger class than the standard one.

Chapter 4 introduces a cut-free cyclic proof system for the intuitionistic linear-time temporal logic iLTL, with a fully finitary correctness condition. The calculus uses labelled formulas in order to accommodate the interplay between the ‘temporal dimension’, represented by the modal rule for the next operator \mathbf{X} , and the ‘intuitionistic dimension’, corresponding to the right-implication rule $\mathbf{R}\rightarrow$. Simple annotations on release and until formulas suffice to provide a finitary characterisation of good infinite branches.

Lastly, in Chapter 5 we present a proof of the uniform interpolation theorem for the modal μ -calculus which differs from the original, automata-based one [32]. We build uniform interpolants from cyclic derivations in the system for the μ -calculus due to Jungteerapanich [76] and Stirling [144].

Sources of the material

Most of the content below has been written specifically for this dissertation. The exceptions are:

- Sections 3.2 and 3.4 of Chapter 3, which are based on joint work with B. Afshari and G. E. Leigh [6]. The latter section incorporates substantial revisions, though.
- Chapter 5, which is based on joint work with B. Afshari and G. E. Leigh [3].

Chapter 1

Mathematical Preliminaries

This chapter introduces notation and terminology used in the rest of the dissertation. Of particular importance are Sections 1.1 to 1.3, for the concepts defined therein recur throughout the thesis. Sections 1.4 to 1.6, on the other hand, are only needed in Chapter 3 and may be safely ignored until then.

We work in the standard framework of Zermelo–Fraenkel set theory with the axiom of choice (ZFC). The reader is assumed to be familiar with elementary notions from set theory and order theory, whose definitions will thus not be given here except when we deviate from conventional practice.

1.1 Sequences and trees

Ultimately, this thesis is about building and manipulating derivations in (possibly ill-founded or cyclic) sequent calculi. Unsurprisingly, then, sequences and trees appear in almost every proof. Here we provide formal definitions of these and related notions, and fix some notation.

Sequences. Let X be a set and α an ordinal. An α -sequence on X is a map $s : \alpha \rightarrow X$, usually written $(s_\beta)_{\beta < \alpha}$, where $s_\beta := s(\beta)$. A *finite sequence* on X is an n -sequence on X for some $n < \omega$. We shall only concern ourselves with finite and ω -sequences. Hence, by *infinite sequence* on X we mean an ω -sequence on X . The collection of all finite sequences on X is denoted by $X^{<\omega}$, and X^ω is the collection of all infinite sequences on X .

We identify finite tuples and finite sequences. In addition, we sometimes abuse notation and treat sequences as sets, for example writing $x \in (s_n)_{n < N \leq \omega}$ to mean that $s_i = x$ for some $i < N$.

Given a sequence s on X , we denote by $\text{Inf}(s)$ the (possibly empty) collection of elements of X which occur infinitely often in s . If s is infinite, for every $n < \omega$

we denote by $s_{\geq n}$ ($s_{>n}$) the sequence $(s_i)_{i \geq n}$ (respectively, $(s_i)_{i > n}$). And, if the length of s is at least $N \leq \omega$, for every $n < N$ we denote by $s_{\leq n}$ ($s_{<n}$) the finite sequence (s_0, \dots, s_n) (respectively, (s_0, \dots, s_{n-1})).

If s is a finite sequence on X and t is a finite or infinite sequence on X , we denote by $s \frown t$ the *concatenation* of s and t , defined in the usual manner. We may abuse notation and write $s \frown x$ in place of $s \frown (x)$, for $x \in X$. Additionally, when working with strings of symbols we might denote concatenation simply by juxtaposition, if doing so carries no ambiguity.

Let s and t be sequences. We write $s \sqsubseteq t$, and say that s is a *prefix* of t , if $t = s \frown t'$ for some (possibly empty) sequence t' . As expected, by $s \sqsubseteq t$ we mean $s \sqsubseteq t$ and $s \neq t$. Analogously, we say that s is a *suffix*, or a *tail*, of t if $t = t' \frown s$ for some finite (possibly empty) sequence t' .

Every map $f : X \rightarrow Y$ induces a map $f^* : X^{<\omega} \cup X^\omega \rightarrow Y^{<\omega} \cup Y^\omega$ given by $f^*((s_n)_{n < N \leq \omega}) := (f(s_n))_{n < N}$ for every finite or infinite sequence $(s_n)_{n < N \leq \omega}$ on X . We abuse notation and denote f^* by f .

Let (X_0, \leq_0) and (X_1, \leq_1) be well-ordered sets. We denote by (\leq_1, \leq_2) the lexicographic well-order on $X_0 \times X_1$ induced by \leq_1 and \leq_2 , that is to say:

$$(x, y) (\leq_1, \leq_2) (x', y')$$

if, and only if, either $x <_1 x'$, or $x = x'$ and $y \leq_2 y'$. We extend this notation to arbitrary finite products of well-ordered sets, writing (\leq_1, \dots, \leq_n) .

Trees. There are a few, slightly different notions of tree commonly used in mathematics. For example, in graph theory trees are acyclic connected graphs, while in descriptive set theory a tree is usually a prefix-closed collection of finite sequences (for the latter notion see, e.g., [80]). We follow a third approach, standard in set theory, and define trees as partially ordered sets in which every element has a well-ordered set of predecessors (see [83, 75]).

1.1.1. DEFINITION (Tree). A *tree* is a pair $\mathcal{T} = (T, <_{\mathcal{T}})$, where T is a set of *vertices* and $<_{\mathcal{T}}$ is a strict partial order on T such that $\{v \in T \mid v <_{\mathcal{T}} u\}$ is well-ordered by $<_{\mathcal{T}}$ for every $u \in T$. ◻

Fix a tree $\mathcal{T} = (T, <_{\mathcal{T}})$. A *subtree* of \mathcal{T} is a tree $\mathcal{T}' = (T', <_{\mathcal{T}'})$ such that the following hold:

- (i) $T' \subseteq T$;
- (ii) if $v \in T'$ and $u <_{\mathcal{T}} v$, then $u \in T'$;
- (iii) $<_{\mathcal{T}'} = <_{\mathcal{T}} \upharpoonright (T' \times T')$.

In other words: a subtree of \mathcal{T} is a downwards-closed subset of T with the induced order. Note that, in general, $[u, \rightarrow)_{\mathcal{T}} := \{v \in T \mid u \leq_{\mathcal{T}} v\}$ is not considered to be a subtree of \mathcal{T} . Instead, we call (the tree induced by) $[u, \rightarrow)_{\mathcal{T}}$ a *cone* of \mathcal{T} .

If $u <_{\mathcal{T}} v$, we say that u is a *predecessor* of v , and that v is a *successor* of u . We write $u <_{\mathcal{T}}^0 v$ if v is an *immediate successor* of u or, equivalently, if u is an *immediate predecessor* of v , i.e.: $u <_{\mathcal{T}} v$ and there is no $u' \in T$ such that $u <_{\mathcal{T}} u' <_{\mathcal{T}} v$. We write $u \leq_{\mathcal{T}}^0 v$ if either $u <_{\mathcal{T}}^0 v$ or $u = v$.

A *path* through \mathcal{T} is a finite or infinite sequence of vertices $(u_n)_{n < N \leq \omega}$ such that $u_{n-1} <_{\mathcal{T}}^0 u_n$ for every $0 < n < N$. If $u \leq_{\mathcal{T}} v$, we let $[u, v]_{\mathcal{T}}$ be the unique finite path (u_0, \dots, u_n) on \mathcal{T} , $n < \omega$, with $u_0 = u$ and $u_n = v$. Similarly, $[u, v)_{\mathcal{T}}$ is the unique (possibly empty) finite path on \mathcal{T} such that $[u, v]_{\mathcal{T}} = [u, v)_{\mathcal{T}} \hat{\ } (v)$. If $u \not\leq_{\mathcal{T}} v$, we let both $[u, v]_{\mathcal{T}}$ and $[u, v)_{\mathcal{T}}$ be the empty sequence. The intervals $(u, v]_{\mathcal{T}}$ and $(u, v)_{\mathcal{T}}$ are defined analogously.

A *branch* of \mathcal{T} is a maximal subset of T linearly ordered by $<_{\mathcal{T}}$.

We say that \mathcal{T} is *rooted* if there is a (necessarily unique) vertex $r \in T$ such that $r \leq_{\mathcal{T}} u$ for every $u \in T$. We call r the *root* of \mathcal{T} . If \mathcal{T} is rooted, a path π on \mathcal{T} is said to be *rooted* if π is non-empty and $\pi(0)$ is the root of \mathcal{T} .

The *height* of a vertex $u \in T$ is the order type of $\{v \in T \mid v <_{\mathcal{T}} u\}$ with respect to $<_{\mathcal{T}}$. For every ordinal α , the α -*level* of \mathcal{T} , denoted by $\text{Lev}_{\alpha}(\mathcal{T})$, is the collection of all vertices of \mathcal{T} with height α . The *height* of \mathcal{T} is the least ordinal α such that $\text{Lev}_{\alpha}(T) = \emptyset$.

The *width* of a vertex $u \in T$ is the cardinality of the set $\{v \in T \mid u <_{\mathcal{T}}^0 v\}$. A vertex is *branching* if it has width greater than 1, and *final* if it has width 0. A final vertex is also called a *leaf*, and we let $\text{Leaf}(\mathcal{T})$ denote the collection of all leaves of \mathcal{T} . We say that \mathcal{T} is *finitely branching* if every vertex of \mathcal{T} has finite width. The *width* of \mathcal{T} is the supremum of the widths of the vertices of \mathcal{T} .

The following is a well-known, fundamental result about infinite trees:

1.1.2. LEMMA (König's lemma). *If \mathcal{T} is a tree of height ω and $|\text{Lev}_n(\mathcal{T})| < \aleph_0$ for every $n < \omega$, then \mathcal{T} has an infinite branch.*

The following convention is in line with our use of trees to represent (possibly ill-founded) proofs in sequent calculi:

1.1.3. CONVENTION. Unless stated otherwise, all trees are assumed to be rooted, finitely branching, and containing only vertices of finite height. In particular, every infinite branch has order type ω .

Labelled trees. We shall use trees mostly to represent proofs, hence they will be labelled. All notions defined below are standard except for those of *low* and *thin* tree, which will become important in Chapter 4.

Let Λ be a set. A Λ -labelled tree is a triple $\mathcal{T} = (T, <_T, \lambda_T)$, where $(T, <_T)$ is a tree and $\lambda_T : T \rightarrow 2^\Lambda$ is a labelling map. The set Λ is said to be the labelling set of \mathcal{T} . A labelled tree is a Λ -labelled tree for some unspecified set Λ . A finitely-labelled tree is a labelled tree with finite labelling set.

A Λ -labelled tree $\mathcal{T} = (T, <_T, \lambda_T)$ is *low* if there are no vertices $u, v \in T$ such that $u <_T^0 v$ and $\lambda_T(u) = \lambda_T(v)$.

Let $\mathcal{T} = (T, <_T)$ be an unlabelled tree, and let Λ be a set. A Λ -labelling of \mathcal{T} is a Λ -labelled tree of the form $(T, <_T, \lambda_T)$ for some labelling map $\lambda_T : T \rightarrow 2^\Lambda$.

Let Λ be a non-empty set, and let $\mathcal{T}_1 = (T_1, <_1, \lambda_1)$ and $\mathcal{T}_2 = (T_2, <_2, \lambda_2)$ be Λ -labelled trees. A Λ -isomorphism between \mathcal{T}_1 and \mathcal{T}_2 is a map $f : T_1 \rightarrow T_2$ such that:

- (i) $f : (T_1, <_1) \rightarrow (T_2, <_2)$ is an order isomorphism;
- (ii) $\lambda_1 = \lambda_2 \circ f$.

We write $f : \mathcal{T}_1, u \sim_\Lambda \mathcal{T}_2, v$ if there is a Λ -isomorphism f between $[u, \rightarrow)_{\mathcal{T}_1}$ and $[v, \rightarrow)_{\mathcal{T}_2}$. We may omit f if clear from context.

A Λ -labelled rooted tree \mathcal{T} is *thin* if there are no vertices $u, v_1, v_2 \in T$, with $v_1 \neq v_2$, such that $u <_T^0 v_1, v_2$ and $\mathcal{T}, v_1 \sim_\Lambda \mathcal{T}, v_2$.

1.1.4. LEMMA. *Let Λ be a finite set and $0 < h < \omega$. There are only finitely many thin Λ -labelled rooted trees of height h up to Λ -isomorphism.*

Proof. We proceed by (strong) induction on h . Clearly, there are at most $2^{|\Lambda|}$ -many Λ -labelled rooted trees of height 1, as each consists of exactly one vertex. Assume that the claim holds for $1, \dots, h$, and let $b_1, \dots, b_h < \omega$ be corresponding upper bounds given by the inductive hypothesis.

Let \mathcal{T}_{h+1} be the unlabelled rooted tree of height $h+1$ built as follows. The only vertex of height 0 is the root. For any vertex of height $h' < h$, let it have exactly $b_1 + \dots + b_{h-h'}$ many immediate successors. Vertices of height h have no immediate successors.

1.1.4.1. CLAIM. *For every thin Λ -labelled rooted tree \mathcal{T} of height $h+1$, there exists a Λ -labelling \mathcal{T}^* of \mathcal{T}_{h+1} such that \mathcal{T} embeds into \mathcal{T}^* .*

Proof of claim. It suffices to embed \mathcal{T} into \mathcal{T}_{h+1} , say via a map f built inductively as follows. Send the root of \mathcal{T} to the root of \mathcal{T}_{h+1} . Suppose that $f(u)$ has been defined, say for $u \in T$ of height $h' < h$. Every immediate successor of u determines a thin Λ -labelled rooted cone of height at most $h-h'$. Note that $0 < h-h' \leq h$. By thinness and the inductive hypothesis, it follows that u has at most $b_1 + \dots + b_{h-h'}$ many immediate successors, whence we can map them injectively to the immediate successors of $f(u)$. \square

Since \mathcal{T}_{h+1} and Λ are both finite, by the claim there can only be finitely many thin Λ -labelled rooted trees of height $h + 1$ up to Λ -isomorphism. ■

Note that we have in addition established the following:

1.1.5. COROLLARY. *Let Λ be a finite set and $h < \omega$. Every thin Λ -labelled rooted tree of height h is finite.*

When working with infinite trees we shall be interested in those which, despite being infinite, admit a finitary presentation. These are usually called ‘regular’ in the literature, although some authors prefer the term ‘rational’.

1.1.6. DEFINITION (Regular tree). A labelled tree is *regular* if it contains only finitely many labelled cones up to isomorphism. ◊

So a Λ -labelled tree \mathcal{T} is regular if there are $v_1, \dots, v_n \in T$ such that for every $u \in T$ there is an $1 \leq i \leq n$ such that $[u, \rightarrow)_{\mathcal{T}}$ and $[v_i, \rightarrow)_{\mathcal{T}}$ are Λ -isomorphic.

Regular trees are well known to admit representations as finite graphs [29, 27, 28, 97, 153, 17]. We shall make this assertion more precise in Section 3.3.

1.2 Trees with back-edges

Whereas an ordinary sequent-style proof is represented as a finite (labelled) tree, a cyclic proof is a finite *tree with back-edges*.

1.2.1. DEFINITION (Tree with back-edges). A (*labelled*) *tree with back-edges* is a tuple $\mathcal{T} = (T, \leq_T, \lambda_T, l \mapsto c_l)$, where (T, \leq_T, λ_T) is a finite labelled tree and $l \mapsto c_l$ is a function from a subset $\text{Rep}_{\mathcal{T}} \subseteq \text{Leaf}(\mathcal{T})$ of the leaves of \mathcal{T} to T such that the following hold for every $l \in \text{Rep}_{\mathcal{T}}$:

- (i) $c_l <_T l$;
- (ii) $\lambda_T(l) = \lambda_T(c_l)$.

Elements of $\text{Rep}_{\mathcal{T}}$ are called *repeats*, and c_l is said to be the *companion* of l . The function $l \mapsto c_l$ is the *back-edge map* of \mathcal{T} . ◊

1.2.2. REMARK. Note that trees with back-edges are by definition finite.

Given a tree with back-edges \mathcal{T} , we denote by \mathcal{T}° the graph which results from \mathcal{T} by taking $<_T^0$ as the edge relation and adding an edge from each repeat $l \in \text{Rep}_{\mathcal{T}}$ to its companion c_l . Infinite paths on \mathcal{T}° are thus always of the form

$$[u, l_0]_{\mathcal{T}} \hat{\ } [c_0, l_1]_{\mathcal{T}} \hat{\ } \cdots \hat{\ } [c_n, l_{n+1}]_T \hat{\ } \cdots,$$

where u is the first vertex on the path and each l_i is a repeat with companion c_i .

Similarly, if \mathcal{T} is a tree with back-edges, \mathcal{T}^ω denotes the (possibly infinite) tree of paths on \mathcal{T}° ordered by the prefix ordering \sqsubset . Clearly, \mathcal{T}^ω is infinite if, and only if, $\text{Rep}_{\mathcal{T}} \neq \emptyset$. We call \mathcal{T}^ω the ω -unravelling of \mathcal{T} .

The following is a fundamental result about paths on trees with back-edges. It ensures that a path through \mathcal{T}° which visits a repeat l more than once also passes through every vertex in $[c_l, l]_{\mathcal{T}}$.

1.2.3. PROPOSITION. *Let \mathcal{T} be a tree with back-edges, and π a path through \mathcal{T}° . If there are $0 \leq m < n < |\pi|$ such that $\pi(m) = l$ and $\pi(n) = l'$ for repeats $l, l' \in \text{Rep}_{\mathcal{T}}$ such that $l \triangleleft l'$, then $[c_l, l']_{\mathcal{T}} \subseteq \{\pi(k) \mid m \leq k \leq n\}$.*

Proof. Let $\pi' := (\pi(m), \dots, \pi(n))$, and let $l = l_0, \dots, l_k = l'$ be repeats such that

$$\pi' = l_0 \widehat{\ } [c_0, l_1]_{\mathcal{T}} \widehat{\ } [c_1, l_2]_{\mathcal{T}} \widehat{\ } \cdots \widehat{\ } [c_{k-2}, l_{k-1}]_{\mathcal{T}} \widehat{\ } [c_{k-1}, l_k],$$

where c_i abbreviates c_{l_i} . We proceed by induction on $k \geq 1$. The base case is clear because then $\pi' = l \widehat{\ } [c_l, l']_{\mathcal{T}}$. For the inductive case, let u be $\leq_{\mathcal{T}}$ -least in $[c_l, l']_{\mathcal{T}}$ such that $[u, l']_{\mathcal{T}} \subseteq \pi'$. Such a vertex exists because $[c_l, l']_{\mathcal{T}}$ is well-ordered by $\leq_{\mathcal{T}}$ and $l' \in \pi'$. Moreover, $u <_{\mathcal{T}} l'$ because, since $l \triangleleft l'$ and $|\pi'| > 1$, the immediate predecessor of l' belongs to $[c_l, l']_{\mathcal{T}} \cap \pi'$.

Towards a contradiction, assume that $u \neq c_l$. Then, u has an immediate predecessor $u^- \in [c_l, l']_{\mathcal{T}} \setminus \pi'$. We claim that $u \neq l$: otherwise $l \in [c_l, l']_{\mathcal{T}}$ and thus $u = l = l'$, contradiction. Hence, π' does not start from u , so u must be the companion of some repeat $l'' \in \pi'$ because $u^- \notin \pi'$. Note that $l'' \neq l$, for otherwise $[c_l, l']_{\mathcal{T}} = [u, l']_{\mathcal{T}} \subseteq \pi'$. So $l'' \in \{l_1, \dots, l_k, l_{k+1}\}$. Also, $l'' \in [u, \rightarrow)_{\mathcal{T}} \subseteq [c_l, \rightarrow)$, whence $l \triangleleft l''$. We distinguish two cases.

Case 1: $l'' \neq l'$. Let $1 \leq i \leq k$ be such that $l'' = l_i$, and let

$$\pi'' := l_0 \widehat{\ } [c_0, l_1]_{\mathcal{T}} \widehat{\ } [c_1, l_2]_{\mathcal{T}} \widehat{\ } \cdots \widehat{\ } [c_{i-1}, l_i]_{\mathcal{T}}.$$

By the inductive hypothesis, $[c_l, l'']_{\mathcal{T}} = [c_{l_0}, l_i]_{\mathcal{T}} \subseteq \pi'' \subseteq \pi'$. Since $l \triangleleft l''$ and $c_l <_{\mathcal{T}} u = c_{l''}$, we have $u^- \in [c_l, l'']_{\mathcal{T}}$ and thus $u^- \in \pi'$, contradiction.

Case 2: $l'' = l'$. Note that $l' \neq l$ because $u \neq c_l$. Let $m < j \leq n$ be least such that $v := \pi(j) \in [u, l']_{\mathcal{T}}$. Since $c_l < u$, we have $j - 1 > m$. Also, $v \neq u$ because otherwise, since $u^- \notin \pi'$, we would have $\pi(j - 1) = l'' = l' \in [u, l']_{\mathcal{T}}$, contradicting the minimality of j . Thus, v has an immediate predecessor $v^- \in [u, l']_{\mathcal{T}}$. And, since $\pi(j - 1) \neq v^-$, $\pi(j - 1)$ is a repeat l''' such that $l \triangleleft l'''$ and with companion v . Also, $l''' \neq l, l'$ because $v \neq u, c_l$. We now argue as in the previous case with l''' in place of l'' . ■

Fix a tree with back-edges \mathcal{T} . For repeats $l, l' \in \text{Rep}_{\mathcal{T}}$, we say that l' is *reachable* from l , in symbols $l \triangleleft l'$, if $c_l <_{\mathcal{T}} l'$. A *reachability path* (on \mathcal{T}) is a finite or infinite

sequence $(l_n)_{n < N \leq \omega}$ of repeats of \mathcal{T} such that $l_{n-1} \triangleleft l_n$ for every $0 < n < N$. A finite reachability path $l_0 \triangleleft \cdots \triangleleft l_n$ is *circular* if $l_n = l_0$.

The following are two results about reachability paths invoked below.

1.2.4. LEMMA. *Let $l_0 \triangleleft \cdots \triangleleft l_m$ be a reachability path on \mathcal{T} . If $l_0 \not\triangleleft l_m$, then there is a $0 < k < m$ such that $c_{l_k} <_T c_{l_0} <_T l_k$.*

Proof. By induction on $m < \omega$. The claim holds vacuously if $m \leq 1$. For the inductive case, let $m > 1$ and assume that the claim holds for $m - 1$ and that we have $l_0 \triangleleft \cdots \triangleleft l_m$ and $l_0 \not\triangleleft l_m$. If $c_{l_1} <_T c_{l_0}$ we are done because $l_0 \triangleleft l_1$. Otherwise, since $l_0 \triangleleft l_1 \triangleleft l_2$ we have $l_0 \triangleleft l_2 \triangleleft \cdots \triangleleft l_m$ and the inductive hypothesis yields a $2 \leq k \leq m - 1$ such that $c_{l_k} <_T c_{l_0} <_T l_k$. ■

1.2.5. LEMMA. *Let $l_0 \triangleleft \cdots \triangleleft l_m \triangleleft l_0$ be a circular reachability path on \mathcal{T} , and for every $i \leq m$ let $c_i := c_{l_i}$. There are $n \leq m$ and $0 = i_0 < i_1 < \cdots < i_n = m$ such that:*

- (i) $l_{i_0} \triangleleft l_{i_1} \triangleleft \cdots \triangleleft l_{i_n}$;
- (ii) $c_{i_{j+1}} <_T c_{i_j} <_T l_{i_{j+1}}$ for all $j < n - 1$;
- (iii) $l_{i_j} \not\triangleleft l_m$ for any $j < n - 1$;
- (iv) $l_{i_j} \triangleleft l_0$ for all $j \leq n$.

Proof. By induction on $j < \omega$, we build an infinite sequence $(i_j)_{j < \omega}$ of natural numbers not greater than m as follows. Let $i_0 := 0$. For the inductive case, assume that i_j has been defined. If $l_{i_j} \triangleleft l_m$, we let $i_{j+1} := m$. Otherwise, since $l_{i_j} \triangleleft l_{i_{j+1}} \triangleleft \cdots \triangleleft l_m$, Lemma 1.2.4 yields a $i_j < k < m$ such that $c_k <_T c_{i_j} <_T l_k$. We let $i_{j+1} := k$. Note that in either case we have $l_{i_j} \triangleleft l_{i_{j+1}}$.

Since $i_j \leq m$ for every $j < \omega$, by construction there is an $n < \omega$ such that $i_n = m$ and $i_j < i_{j+1}$ for all $j < n$. It remains to see that i_0, \dots, i_n satisfy conditions (ii)–(iv) above. Let $j < n - 1$. Then, $i_{j+1} \neq m$ and thus by construction we have $c_{i_{j+1}} <_T c_{i_j} <_T l_{i_{j+1}}$ and $l_{i_j} \not\triangleleft l_m$. This establishes (ii) and (iii). To see that (iv) holds, we argue by induction on $j = 0, \dots, n - 1$. The base case $j = 0$ is clear because $l_{i_0} = l_0 \triangleleft l_0$. For the inductive case, suppose that $l_{i_j} \triangleleft l_0$ for some $j < n - 1$. Then, we have $c_{i_{j+1}} <_T c_{i_j} <_T l_0$, so $l_{i_{j+1}} \triangleleft l_0$. Finally, note that we have $l_m \triangleleft l_0$ by assumption. ■

For the remaining of this section, fix a finite set Σ , a tree with back-edges $\mathcal{T} = (T, <_T, \lambda_T, l \mapsto c_l)$, and maps $\Theta : T \rightarrow \Sigma^{<\omega}$ and $\text{inv} : \text{Rep}_{\mathcal{T}} \rightarrow \Sigma^{<\omega}$ such that for every $l \in \text{Rep}_{\mathcal{T}}$ we have: $\text{inv}(l) \sqsubseteq \Theta(u)$ for every $u \in [c_l, l]_{\mathcal{T}}$. We call Θ a *control*

map, and inv an *invariant map*. We prove some general results that will be later used in Chapters 3 to 5.

The invariant map inv induces the following (reflexive) quasi-order \preceq on the repeats of \mathcal{T} : $l \preceq l'$ if, and only if, $\text{inv}(l) \sqsubseteq \text{inv}(l')$. The orders \triangleleft and \preceq are related in the sense of the following results.

1.2.6. LEMMA. *For every $l \in \text{Rep}_{\mathcal{T}}$, the set $\{l' \in \text{Rep}_{\mathcal{T}} \mid l' \preceq l\}$ is linearly ordered by \preceq .*

Proof. If $l_1, l_2 \preceq l$, then $\text{inv}(l_1)$ and $\text{inv}(l_2)$ are both prefixes of $\text{inv}(l)$. ■

1.2.7. LEMMA. *For all $l, l' \in \text{Rep}_{\mathcal{T}}$, if $c_{\nu} \leq_T c_l <_T l'$, then l and l' are \preceq -comparable.*

Proof. Since $c_l \in [c_{\nu}, l']_{\mathcal{T}}$, both $\text{inv}(l)$ and $\text{inv}(l')$ are prefixes of $\Theta(c_l)$. ■

1.2.8. LEMMA. *For all $l, l' \in \text{Rep}_{\mathcal{T}}$, if $l \triangleleft l' \triangleleft l$, then l and l' are \preceq -comparable.*

Proof. From $c_l <_T l'$ and $c_{\nu} <_T l$ it follows that either $c_{\nu} \leq_T c_l <_T l'$ or $c_l \leq_T c_{\nu} <_T l$, whence the claim follows from Lemma 1.2.7. ■

1.2.9. LEMMA. *Let $(l_i)_{i < N \leq \omega}$ be a finite or infinite reachability path on \mathcal{T} such that l_{i-1} and l_i are \preceq -comparable for every $0 < i < N$. Then, there is a $j < N$ such that $l_j \preceq l_i$ for all $i < N$.*

Proof. By induction on $i < N$ we find a $j_i < N$ such that $l_{j_i} \preceq l_0, \dots, l_i$. This establishes the claim because the set $\text{Rep}_{\mathcal{T}}$ is finite. Let $l_{j_0} := l_0$, and for the inductive case let $l_{j_{i+1}}$ be any \preceq -minimum of $\{l_{i+1}, l_{j_i}\}$, which exists by the comparability assumption and Lemma 1.2.6. ■

1.2.10. LEMMA. *For every circular reachability path $l_0 \triangleleft \dots \triangleleft l_m \triangleleft l_0$ on \mathcal{T} , there are $n \leq m$ and $0 = i_0 < \dots < i_n = m$ such that l_{i_j} and $l_{i_{j+1}}$ are \preceq -comparable for every $j < n$.*

Proof. The claim clearly holds if $m = 0$, so assume that $m > 0$. We abbreviate c_{l_i} to c_i for all $i \leq m$.

By Lemma 1.2.5, there are $n \leq m$ and $0 = i_0 < i_1 < \dots < i_n = m$ such that:

- (i) $l_{i_0} \triangleleft l_{i_1} \triangleleft \dots \triangleleft l_{i_n}$;

- (ii) $c_{i_{j+1}} <_T c_{i_j} <_T l_{i_{j+1}}$ for all $j < n - 1$;
- (iii) $l_{i_j} \not\prec l_m$ for any $j < n - 1$;
- (iv) $l_{i_j} \triangleleft l_0$ for all $j \leq n$.

Note that $n > 0$.

Lemma 1.2.7 implies that l_{i_j} and $l_{i_{j+1}}$ are \preceq -comparable for all $j < n - 1$. It remains to see that $l_{i_{n-1}}$ and $l_n = l_m$ are also \preceq -comparable. This is an immediate consequence of Lemma 1.2.8 if $n = 1$, so assume that $n > 1$.

By (ii) and (iii), we have $c_{i_{n-1}} <_T c_{i_{n-2}} <_T l_{i_{n-1}}$ and $c_{i_{n-2}} \not\prec_T l_{i_n}$. And, by (i), $c_{i_{n-1}} <_T l_{i_n}$, whence either $c_{i_n} <_T c_{i_{n-1}} <_T l_{i_n}$, or $c_{i_{n-1}} \leq_T c_{i_n} <_T c_{i_{n-2}} <_T l_{i_{n-1}}$, where in the latter case $c_{i_n} <_T c_{i_{n-2}}$ is given by (iii). In either case, Lemma 1.2.7 implies that l_{i_n} and $l_{i_{n-1}}$ are \preceq -comparable. ■

We arrive finally at the two fundamental results about invariants.

1.2.11. PROPOSITION. *For every infinite reachability path $(l_i)_{i < \omega}$ on \mathcal{T} , there exists a $k < \omega$ such that $l_k \preceq l_j$ for all $j \geq k$.*

Proof. Let $i < \omega$ be such that all the repeats among l_i, l_{i+1}, \dots occur infinitely often on the path. Since the set $\text{Rep}_{\mathcal{T}}$ is finite, it suffices to find, for every $j < \omega$, a $k_j \geq i$ such that $l_{k_j} \preceq l_i, l_{i+1}, \dots, l_{i+j}$.

Let $k_0 := i$. For the inductive case, assume that l_{k_j} has been defined. We know that $l_{i+j} \triangleleft l_{i+j+1}$ and that there is a reachability path from l_{i+j+1} to l_{i+j} because l_{i+j} occurs infinitely often on l_i, l_{i+1}, \dots . Let $n > j + 1$ be such that $l_{i+n} = l_{i+j}$. We then have $l_{i+j+1} \triangleleft \dots \triangleleft l_{i+n} \triangleleft l_{i+j+1}$, so by Lemmas 1.2.9 and 1.2.10 there is some $j + 1 \leq m \leq n$ such that $l_{i+m} \preceq l_{i+j+1}, l_{i+j}$. Let $l_{k_{j+1}}$ be any \preceq -minimum of $\{l_{i+m}, l_{k_j}\}$, which exists by Lemma 1.2.6. ■

1.2.12. PROPOSITION. *Let $l_0 \triangleleft \dots \triangleleft l_m \triangleleft l_0$ be a circular reachability path on \mathcal{T} , and let $w \in \Sigma^{< \omega}$. If $w \sqsubseteq \text{inv}(l_i)$ for each $i \leq m$, then $w \sqsubseteq \Theta(u)$ for every $u \in [c_{i_m}, l_0]_{\mathcal{T}}$.*

Proof. The claim clearly holds if $m = 0$, so assume that $m > 0$. We abbreviate c_i to c_i for all $i \leq m$.

By Lemma 1.2.5, there are $n \leq m$ and $0 = i_0 < i_1 < \dots < i_n = m$ such that:

- (i) $l_{i_0} \triangleleft l_{i_1} \triangleleft \dots \triangleleft l_{i_n}$;
- (ii) $c_{i_{j+1}} <_T c_{i_j} <_T l_{i_{j+1}}$ for all $j < n - 1$;
- (iii) $l_{i_j} \not\prec l_m$ for any $j < n - 1$;
- (iv) $l_{i_j} \triangleleft l_0$ for all $j \leq n$.

Note that $n > 0$.

The claim follows easily if $l_0 \triangleleft l_m$, so assume $l_0 \not\triangleleft l_m$ (so $m > 1$).

By (i), $l_{i_{n-1}} \triangleleft l_m$, so we have either $c_{i_{n-1}} \in [c_m, l_m]_{\mathcal{T}}$, or $c_{i_{n-1}} <_T c_m$. Assume the former. Then, (ii) and (iv) yield:

$$[c_m, l_0]_{\mathcal{T}} = [c_{i_n}, l_{i_0}]_{\mathcal{T}} = [c_{i_n}, c_{i_{n-1}}]_{\mathcal{T}} \widehat{\cdot} \cdots \widehat{\cdot} [c_{i_1}, c_{i_0}]_{\mathcal{T}} \widehat{\cdot} [c_{i_0}, l_{i_0}]_{\mathcal{T}}.$$

By (ii) and the assumption, we have $[c_{i_{j+1}}, c_{i_j}]_{\mathcal{T}} \subseteq [c_{i_{j+1}}, l_{i_{j+1}}]_{\mathcal{T}}$ for every $j < n$, and thus $w \sqsubseteq \Theta(u)$ for every $u \in [c_m, l_0]_{\mathcal{T}}$.

Now assume that $c_{i_{n-1}} <_T c_m$. If $n = 1$, we have $c_{i_0} <_T c_{i_1} <_T l_{i_0}$ and thus

$$[c_m, l_0]_{\mathcal{T}} = [c_{i_1}, l_{i_0}]_{\mathcal{T}} \subseteq [c_{i_0}, l_{i_0}]_{\mathcal{T}},$$

whence $w \sqsubseteq \Theta(u)$ for every $u \in [c_m, l_0]_{\mathcal{T}}$. Finally, suppose that $n > 1$. We then have $c_{i_{n-1}} <_T c_{i_{n-2}} \not\triangleleft_T l_m$ by (ii) and (iii), so $c_m <_T c_{i_{n-2}}$. In this case, (ii) and (iv) yield:

$$[c_m, l_0]_{\mathcal{T}} = [c_{i_n}, l_{i_0}]_{\mathcal{T}} = [c_{i_n}, c_{i_{n-2}}]_{\mathcal{T}} \widehat{\cdot} [c_{i_{n-2}}, c_{i_{n-3}}]_{\mathcal{T}} \widehat{\cdot} \cdots \widehat{\cdot} [c_{i_1}, c_{i_0}]_{\mathcal{T}} \widehat{\cdot} [c_{i_0}, l_{i_0}]_{\mathcal{T}}.$$

By (ii) and the assumption, we have $[c_{i_{j+1}}, c_{i_j}]_{\mathcal{T}} \subseteq [c_{i_{j+1}}, l_{i_{j+1}}]_{\mathcal{T}}$ for every $j < n-2$ and, moreover, $[c_{i_n}, c_{i_{n-2}}]_{\mathcal{T}} \subseteq [c_{i_{n-1}}, l_{i_{n-1}}]_{\mathcal{T}}$. Hence, again we conclude that $w \sqsubseteq \Theta(u)$ for each $u \in [c_m, l_0]_{\mathcal{T}}$. \blacksquare

1.3 Sequent calculi

Most of the proof systems that we consider in this dissertation are sequent calculi.¹ However, we use several different formalisms throughout the thesis, such as hypersequents and labelled sequents, with and without ‘annotations’. Moreover, some of the systems will be finitary, others ill-founded, and still others cyclic. A definition general enough to encompass all these calculi would probably be too abstract to be useful for our purposes. Therefore, here we only define a few general concepts and fix some notation. For each of the systems used in the thesis, the reader will find specific definitions (and refinements of some notions) in due time. We assume that the reader is acquainted with (structural) proof theory, in particular Gentzen systems and variations thereof. Introductions to the subject may be found in [152, 101].

Fix a set \mathcal{S} of (*abstract*) *sequents*. A (*sequent*) *rule* is an expression of the form

$$\text{R} \frac{S_1 \quad \cdots \quad S_n}{S_0},$$

¹Sequent calculi were originally introduced by Gentzen [54, 55] in the 1930s, together with natural deduction calculi.

where each S_i is a sequent and R is the *name* of the rule. We call S_0 the *conclusion* of the rule, and S_1, \dots, S_n the *premises* thereof. Rule R is *axiomatic* if $n = 0$.

A (*sequent*) *calculus*, or (*proof*) *system*, is a collection of sequent rules. Given a calculus G and rules R_1, \dots, R_n , we denote $G \cup \{R_1, \dots, R_n\}$ by $G + R_1 + \dots + R_n$. The rules of a calculus will usually be specified by rule schemata.² We shall often abuse notation and say *instance of a rule* in place of the more rigorous *instance of a rule schema*.

Let G be a sequent calculus. A G *derivation* of a sequent S is a finite or infinite tree \mathcal{T} built according to the rules of G and whose root is labelled by S . We shall not make this definition more precise here for the aforementioned reasons. Note, however, that it would not suffice to simply label the vertices of \mathcal{T} with rule names and sequents, for traces should also be taken into account. Consider, for example, the following (branch of a) conjunction rule, where sequents are disjunctive sets of classical propositional formulas:

$$\frac{\alpha, \alpha \wedge \beta, \alpha \wedge \gamma}{\alpha \wedge \beta, \alpha \wedge \gamma}$$

The implicit use of contraction makes it impossible to know, just by looking at the rule, which of the formulas $\alpha \wedge \beta$ and $\alpha \wedge \gamma$ in the conclusion ‘produced’ the formula α in the premise. This is often irrelevant in ordinary sequent calculi. When working with ill-founded or cyclic calculi, however, it will be necessary to have clear notions of traces. Since the definition of trace depends heavily on the specifics of the calculus, we do not provide a definition here and instead give one in due time for every system that needs it.

Let \mathcal{T} be a G derivation. A vertex $u \in T$ is *axiomatic* if the sequent in its label is the conclusion of an axiomatic rule of G . A vertex $u \in T$ is a *cul-de-sac* if the sequent in its label is not the conclusion of any rule of G .

1.3.1. REMARK. For every calculus, we distinguish between derivations and *proofs*. Derivations, informally, are simply proof-trees, sometimes satisfying some non-degeneracy condition. Proofs, in contrast, are a specific kind of derivation, whose definition varies from system to system. In ordinary sequent calculi, for example, proofs are finite derivations whose leaves are all axiomatic. In ill-founded systems, the notion of proof typically includes a correctness condition on infinite branches of derivations.

Let G be a calculus for which a notion of proof has been defined, and let S be a sequent. We write $G \vdash S$, and say that S is *provable* in G , if there exists a G proof with conclusion S .

²This is relevant only for the logic CTL* considered in Chapter 3, because it is not closed under substitution.

A *cyclic (sequent) calculus*, or *cyclic (sequent) system*, is a pair $G^\circ = (G, \Pi)$, where G is a sequent calculus and Π is a collection of G derivations, representing the proofs of the system.³ Typically, Π is specified by means of correctness conditions imposed on infinite paths through derivations. If Π is characterised finitarily, we say that G° is *finitary*. Not all cyclic systems that we shall consider will be finitary. The *base* of G° , denoted by $\mathcal{B}(G^\circ)$, is the underlying, ‘acyclic’ calculus G .

As is customary, when working with sequents we abbreviate $\Gamma \cup \Delta$ to Γ, Δ and $\Gamma \cup \{\varphi\}$ to Γ, φ . And when describing derivations, the use of a dashed line

$$R \frac{S_1 \text{ --- } \cdots \text{ --- } S_n}{S_0 \text{ ---}}$$

indicates the omission of some vertices. Similarly, a double line

$$R \frac{S_1 \quad \cdots \quad S_n}{S_0}$$

indicates the application of multiple rules in succession.

1.4 The Borel hierarchy

We recall the notion of the Borel hierarchy of a metrizable space, which will be used in conjunction with automata- and game-theoretic techniques to establish the completeness of an ill-founded calculus for the logic CTL^* in Chapter 3. Familiarity with basic topological notions is assumed, in particular: topological spaces and subspaces, bases, continuous maps, homeomorphisms, metrizability. The reader may consult [98] for an introduction to topology. Our main reference for this section is [80, § 11].

Let (X, T) be a topological space. The class of *Borel sets* of (X, T) , denoted by $B(X, T)$, is the smallest family of subsets of X containing all the open sets in T and closed under complements and countable unions (hence also under countable intersections). In other words, $B(X, T)$ is the σ -algebra on X generated by T . It is easy to see that $B(X, T)$ always exists.

When (X, T) is metrizable, its Borel sets form a natural hierarchy as follows. For every $1 \leq \alpha < \omega_1$, we define the subsets $\Sigma_\alpha^0(X)$, $\Pi_\alpha^0(X)$ and $\Delta_\alpha^0(X)$ of X by setting:

- $\Sigma_1^0(X) := T$;
- $\Pi_\alpha^0(X) := \{X \setminus A \mid A \in \Sigma_\alpha^0(X)\}$;

³This is a rough description of a cyclic calculus. As mentioned previously, more polished definitions are given in the chapters to follow. Additionally, the reader may consult [17, 2] for abstract definitions of cyclic proofs.

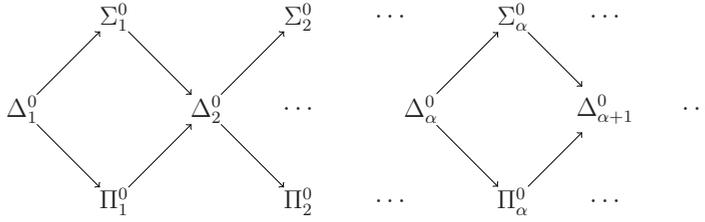


FIGURE 1.1: The Borel hierarchy of a metrizable space (X, T) , where arrows indicate inclusion and explicit mentions of X have been omitted.

- $\Sigma_\alpha^0(X) := \left\{ \bigcup_{n < \omega} A_n \mid A_n \in \Pi_{\beta_n}^0(X), \beta_n < \alpha \right\}$, if $\alpha > 1$;
- $\Delta_\alpha^0(X) := \Sigma_\alpha^0(X) \cap \Pi_\alpha^0(X)$.

A straightforward transfinite induction up to ω_1 shows that

$$\Sigma_\alpha^0(X) \cup \Pi_\alpha^0(X) \subseteq \Delta_{\alpha+1}^0(X)$$

for every $1 \leq \alpha < \omega_1$. Moreover, we have:

$$B(X, T) = \bigcup_{1 \leq \alpha < \omega_1} \Sigma_\alpha^0(X) = \bigcup_{1 \leq \alpha < \omega_1} \Pi_\alpha^0(X) = \bigcup_{1 \leq \alpha < \omega_1} \Delta_\alpha^0(X),$$

whence there emerges the *Borel hierarchy* of (X, T) , depicted in Figure 1.1.

A set $A \subseteq X$ is said to be Σ_α^0 in (X, T) if $A \in \Sigma_\alpha^0(X)$. And analogously for Π_α^0 and Δ_α^0 .

Continuous maps are easily shown to respect the levels of the Borel hierarchy, in the sense that we have:

1.4.1. PROPOSITION. *Let (X, T) and (X', T') be metrizable topological spaces, and $f : X \rightarrow X'$ a continuous map. For every $A \subseteq X'$ and $1 \leq \alpha < \omega_1$, the following hold:*

- (i) *if $A \in \Sigma_\alpha^0(X')$, then $f^{-1}(A) \in \Sigma_\alpha^0(X)$;*
- (ii) *if $A \in \Pi_\alpha^0(X')$, then $f^{-1}(A) \in \Pi_\alpha^0(X)$;*
- (iii) *if $A \in \Delta_\alpha^0(X')$, then $f^{-1}(A) \in \Delta_\alpha^0(X)$.*

We shall be interested in Borel sets of spaces endowed with a ‘prefix’ topology. Let X be any finite or infinite set. The *prefix topology* on X^ω is the topology T_{\sqsubseteq} on X^ω generated by the base $\{B_p \mid p \in X^{<\omega}\}$, where $B_p := \{x \in X^\omega \mid p \sqsubseteq x\}$ for every $p \in X^{<\omega}$. It is well known that $(X^\omega, T_{\sqsubseteq})$ is metrizable [106, Prop. 3.2].

When the starting set X is finite, the prefix topology on X^ω is often called the *Cantor topology* on X^ω (see, e.g., [151, 102, 140]). This is justified because, if X is finite, then $(X^\omega, T_{\sqsubseteq})$ is homeomorphic to the Cantor space [106, Prop. 3.14].

1.5 Gale–Stewart games

Two-player infinite games with perfect information, or Gale–Stewart games, were introduced in [52]. Following a now standard approach that goes back at least to [102], in Chapter 3 we shall derive completeness for an ill-founded calculus by playing an infinite game on a proof-search tree. Here we provide all necessary definitions.

Our main reference for this section is Gale and Stewart’s original article [52]. In particular, we do not require players to take turns alternately.

1.5.1. DEFINITION (Gale–Stewart game). A (*Gale–Stewart*) *game* is an octuple $\mathcal{G} = (x^*, X_{\text{I}}, X_{\text{II}}, X, f, \Pi, W_{\text{I}}, W_{\text{II}})$ such that:

- (i) X is a non-empty set of *positions*;
- (ii) $x^* \in X$ is the *initial position* of \mathcal{G} ;
- (iii) $X = X_{\text{I}} \uplus X_{\text{II}}$;
- (iv) $f : X \setminus \{x^*\} \rightarrow X$ is the (*immediate*) *predecessor* function of \mathcal{G} ;
- (v) f is surjective;
- (vi) for every $x \in X$ there is an $n < \omega$ such that $f^n(x) = x^*$;
- (vii) $\Pi := \{(x_n)_{n < \omega} \in X^\omega \mid x_0 = x^* \text{ and } x_i = f(x_{i+1}) \text{ for all } i < \omega\}$ is the *space* of \mathcal{G} ;
- (viii) $\Pi = W_{\text{I}} \uplus W_{\text{II}}$. –

In their original article [52], Gale and Stewart refer to what we (following standard practice) call Gale–Stewart games as *win-lose games*, which are a particular instance of their more general definition of *game*. Most of their attention, nevertheless, is devoted to win-lose games.

Fix a game $\mathcal{G} = (x^*, X_{\text{I}}, X_{\text{II}}, X, f, \Pi, W_{\text{I}}, W_{\text{II}})$. A (*total*) *play* of \mathcal{G} is an element of Π . A *partial play* of \mathcal{G} is a finite prefix of a total play of \mathcal{G} .⁴ The set W_{I} (W_{II}) is the *winning set for player I* (respectively, *II*). A play p of \mathcal{G} is a *win for player I* (player *II*) if $p \in W_{\text{I}}$ (respectively, $p \in W_{\text{II}}$).

The fact that the predecessor map of a game is a function (and only defined for non-initial positions) implies that, for every position $x \in X$, if x_0, \dots, x_n and y_0, \dots, y_m are partial plays such that $x_n = y_m = x$, then $n = m$ and $x_i = y_i$ for every $i \leq n$. So, informally, the positions of the game are ‘history-aware’. In particular, no position may appear twice in a play:

⁴It would be more natural to define finite games ‘from below’, i.e., as finite sequences of positions satisfying the requirements imposed on total plays. The two approaches coincide due to the surjectivity of the predecessor map.

1.5.2. LEMMA. *Let \mathcal{G} be a game, and p a play of \mathcal{G} . Every position x of \mathcal{G} occurs at most once in p .*

Proof. Let $p = (x_n)_{n < \omega}$, and let $i \leq j < \omega$ be such that $x_i = x_j$. By Definition 1.5.1(vii) we have $x_{i-k} = x_{j-k}$ for all $k = 0, \dots, i$, so $x_{j-i} = x_0 = x^*$ and thus $j - i = 0$ because x^* is not in the domain of f . ■

A \mathcal{G} -strategy for player I (player II) is a function $\sigma : X_I \rightarrow X$ (respectively, $\sigma : X_{II} \rightarrow X$) such that $\sigma(x) \in f^{-1}(x)$ for every $x \in X_I$ (respectively, $x \in X_{II}$). We denote by $\Sigma_I(\mathcal{G})$ ($\Sigma_{II}(\mathcal{G})$) the collection of all \mathcal{G} -strategies for player I (respectively, player II).

Let σ be a \mathcal{G} -strategy for player I (player II). A play $p \in \Pi$ is σ -consistent if for every $n < \omega$, if $p(n) \in X_I$ (respectively, $p(n) \in X_{II}$), then $p(n+1) = \sigma(p(n))$. Given \mathcal{G} -strategies σ and τ for players I and II, respectively, we denote by $\langle \sigma, \tau \rangle$ the unique game which is both σ - and τ -consistent. As a consequence of Lemma 1.5.2, every game is of that form.

1.5.3. PROPOSITION. *Let \mathcal{G} be a game. Every play of \mathcal{G} is of the form $\langle \sigma, \tau \rangle$ for some \mathcal{G} -strategies σ and τ for players I and II, respectively.*

Proof. Let $\mathcal{G} = (x^*, X_I, X_{II}, X, f, \Pi, W_I, W_{II})$, and let $p = (x_n)_{n < \omega}$ be a play of \mathcal{G} . We define the \mathcal{G} -strategies $\sigma : X_I \rightarrow X$ and $\tau : X_{II} \rightarrow X$ as follows. Let $x \in X_I$. If x does not occur in p , we let $\sigma(x) := x'$ for any $x' \in X$ such that $f(x') = x$ (at least one such x' exists by the surjectivity of f). Suppose now that x occurs in p . By Lemma 1.5.2, there is a unique $n < \omega$ such that $p(n) = x$. We then let $\sigma(x) := p(n+1)$. The strategy τ is defined analogously. ■

A \mathcal{G} -strategy σ for player I (player II) is *winning* if for every σ -consistent play $p \in \Pi$ we have $p \in W_I$ (respectively, $p \in W_{II}$). That this definition is equivalent to the one in [52] is given by the following proposition:

1.5.4. PROPOSITION. *Let \mathcal{G} be a game. A \mathcal{G} -strategy σ for player I (player II) is winning if, and only if, for every \mathcal{G} -strategy τ for player II (respectively, player I) we have $\langle \sigma, \tau \rangle \in W_I$ (respectively, $\langle \tau, \sigma \rangle \in W_{II}$).*

Proof. We prove only the claim for player I. The left-to-right direction is clear. For the converse, let p be a σ -consistent play of \mathcal{G} . By the proof of Proposition 1.5.3, p is of the form $p = \langle \sigma, \tau \rangle$ for some \mathcal{G} -strategy τ for player II, whence $p \in W_I$ by assumption. ■

A game \mathcal{G} is *determined* if there exists a winning \mathcal{G} -strategy for one of the players. Clearly, if \mathcal{G} is determined then there is one and only one player for which

a winning \mathcal{G} -strategy exists. Gale and Stewart showed in [52] that not every game is determined, but they managed to prove determinacy under simple topological requirements.

Following [96], we endow Π with the topology $T_{\mathcal{G}}$ generated by taking as a base the collection $\{B_s \mid s \in X^{<\omega}\}$, where for every $s \in X^{<\omega}$ we let

$$B_s := \{p \in \Pi \mid s \sqsubseteq p\}.$$

It is straightforward to see that this base is the same as the one used originally in [52], and that the topology it generates is exactly the subspace topology that $\Pi \subseteq X^\omega$ inherits from $(X^\omega, T_{\sqsubseteq})$, where T_{\sqsubseteq} is the prefix topology on X^ω defined at the end of Section 1.4.

The game $\mathcal{G} = (x^*, X_I, X_{II}, X, f, \Pi, W_I, W_{II})$ is *Borel* if W_I is a Borel set in $(\Pi, T_{\mathcal{G}})$. *Open* and *closed* games are defined analogously, as well as Σ_α^0 , Π_α^0 and Δ_α^0 games, for every $1 \leq \alpha < \omega_1$.

Gale and Stewart [52] proved determinacy of games which are open or closed. This result was gradually improved upon by several authors, culminating in Martin's celebrated proof that all Borel games are determined [96]. For our purposes in Chapter 3, determinacy of Δ_3^0 games, established earlier in [36], will suffice.

1.5.5. THEOREM ([36]). *If a game \mathcal{G} is Σ_3^0 or Π_3^0 , then \mathcal{G} is determined.*

1.6 Büchi automata

Büchi automata were introduced by Büchi [20, 21] to establish the decidability of the monadic second-order theory of one successor (S1S). They are a natural generalisation of finite-state automata to infinite words. Here we provide only the definitions and results needed in Chapter 3. The reader may consult [151] for an overview of Büchi automata.

An *alphabet* is a finite set, whose elements are called *symbols* or *letters*. An ω -word over an alphabet Σ is an element of Σ^ω . An ω -*language* over Σ is a subset of Σ^ω . Since we restrict ourselves to ω -words, unless stated otherwise by *word* and *language* we mean ω -word and ω -language, respectively.

It will be convenient to work with several automata sharing the same underlying transition graph. With this in mind, we first introduce the following notion:

1.6.1. DEFINITION (Automaton skeleton). An *automaton skeleton* is a triple $S = (Q, \Sigma, \Delta)$, where Q is a non-empty finite set of *states*, Σ is an alphabet, and $\Delta \subseteq Q \times \Sigma \times Q$ is the *transition relation* of S .

A *run* of S from a state $q \in Q$ on an ω -word $(\sigma_n)_{n < \omega} \in \Sigma^\omega$ is an infinite sequence of states $(q_n)_{n < \omega} \in Q^\omega$ such that $q_0 = q$ and $(q_n, \sigma_n, q_{n+1}) \in \Delta$ for every $n < \omega$. ⊣

1.6.2. DEFINITION (Büchi automaton). A (*non-deterministic*) *Büchi automaton* (NBA) over an alphabet Σ is a tuple $\mathcal{A} = (Q, \Sigma, \Delta, q^*, F)$, where (Q, Σ, Δ) is an automaton skeleton, $q^* \in Q$ is the *initial state* of \mathcal{A} , and $F \subseteq Q$ is the set of *final states* of \mathcal{A} .

A *run* of \mathcal{A} on an ω -word $w \in \Sigma^\omega$ is a run of (Q, Σ, Δ) from q^* on w . A run ρ is *accepting* if $\text{Inf}(\rho) \cap F \neq \emptyset$. The automaton \mathcal{A} *accepts* a word $w \in \Sigma^\omega$ if there is an accepting run of \mathcal{A} on w . We let $\mathcal{L}(\mathcal{A}) := \{w \in \Sigma^\omega \mid \mathcal{A} \text{ accepts } w\}$. \dashv

A language $L \subseteq \Sigma^\omega$ is said to be ω -*regular*, or *Büchi-recognisable*, if there is an NBA \mathcal{A} such that $\mathcal{L}(\mathcal{A}) = L$. It is well known that the collection of ω -regular languages over an alphabet Σ is closed under union, intersection and complementation:

1.6.3. THEOREM ([20, 151]). *Let Σ be an alphabet. If $L_1, L_2 \subseteq \Sigma^\omega$ are ω -regular, then so are $L_1 \cup L_2$, $L_1 \cap L_2$, and $\Sigma^\omega \setminus L_1$.*

Whereas closure under union and intersection is not difficult to see, closure of ω -regular languages under complementation is far from trivial. In contrast to the case of finite-state automata, a reduction of NBA's to deterministic Büchi automata (DBA, defined by requiring the transition relation to be functional) is not possible because the latter are strictly less expressive than the former (see, e.g., [151, § 4]). Büchi's original proof in [20] relied on Ramsey's theorem. Another well-known proof, based on transforming NBA's into deterministic Rabin automata, was later obtained by Safra [123].

Topology turns out to shed light on the study of ω -languages (see, e.g., [67, 151, 140, 106]). In particular, every ω -regular language over an alphabet Σ is Δ_3^0 in Σ^ω with the prefix topology (recall the definitions given in Section 1.4 above).

1.6.4. THEOREM ([67, 151]). *Let Σ be an alphabet, and \mathcal{A} an NBA over Σ . Then, $\mathcal{L}(\mathcal{A}) \in \Delta_3^0(\Sigma^\omega, T_{\square})$, where T_{\square} is the prefix topology on Σ^ω .*

Propositional Modal Logics

Chapter 2

Cyclic Companions of Modal Logics

Modal logic traces back at least to Aristotle’s modal syllogistic (see, e.g., [31, 119]). As a modern, mathematical discipline, however, its origins are much more recent: it is usually dated back to C. I. Lewis’s work, in the early years of the 20th century, on the paradoxes of material implication and his proposal of a system of strict implication to capture the usual, intensional meaning of the word ‘implies’. This work culminated in the 1932 book *Symbolic Logic* [87], cowritten by C. I. Lewis and C. H. Langford, in which the five axiomatic systems of modal logic **S1–S5** are introduced. For a historical overview of modern modal logic, we refer the reader to [13, § 1.7] and [60].

The language of propositional modal logic extends that of propositional logic by adding two unary operators, typically denoted \Box (‘box’) and \Diamond (‘diamond’), which express modalities of truth such as being necessary, possible, known, believed, provable... Modal logics are then obtained as extensions of classical or intuitionistic propositional logic by means of axioms and rules of inference involving the operators \Box and \Diamond . We restrict our attention to modal logics extending classical propositional logic (CPC). In this setting, the operators \Box and \Diamond are dual to one another and thus interdefinable via the equivalence $\Diamond p \equiv \neg\Box\neg p$. For example, if $\Box p$ is read as ‘it is necessarily the case that p ’, then both $\Diamond p$ and $\neg\Box\neg p$ express that ‘it is possibly the case that p ’. This duality is analogous to the one between the universal (\forall) and existential (\exists) quantifiers in predicate logic, with \Box corresponding to \forall and \Diamond to \exists .

We assume that the reader is familiar with both classical logic and modal logic, and in particular the systems **K**, **K4**, **S4**, **GL** (*Gödel–Löb provability logic*), and **S4Grz** (*Grzegorzcyk logic*), though we formally define all of them below. An introduction to CPC may be found in [24, Ch. 1], and for mathematically-oriented introductions to modal logic the reader may consult [24, 13].

Ordinary sequent calculi have proved to be insufficient for modal logic, partly due to the difficulties involved in obtaining ‘good’ rules for the introduction of the modal operators, that is, rules yielding calculi with nice structural properties, such as the admissibility of cut-elimination (see [126, 158, 99, 100]). In his survey on the technical and philosophical limitations of ordinary sequent calculi for modal logic, Wansing writes:

[M]any normal modal and temporal logics are presentable as ordinary Gentzen calculi [...]. However, no uniform way of presenting only the most important normal modal and temporal propositional logics as ordinary Gentzen calculi is known. Further, the standard approach fails to be modular: in general it is not the case that a single axiom schema is captured by a single sequent rule (or a finite set of such rules). [158, p. 68]

Consider, for example, the logic **K4**, which results by adding the transitivity axiom $\Box p \rightarrow \Box \Box p$ to the basic modal logic **K**. The most natural attempt at obtaining a sequent calculus for **K4** is adding the rule

$$\Box_{\mathbf{K4}}^{\times} \frac{\Gamma \Rightarrow \Box \varphi, \Delta}{\Gamma \Rightarrow \Box \Box \varphi, \Delta}$$

to a system for CPC. The resulting system is indeed sound and complete with respect to **K4**, but unfortunately it does not enjoy cut-elimination [85, 126]. Similar difficulties arise when considering other logics. In the words of Sambin and Valentini:

It is usually not difficult to choose suitable rules for each modal logic if one is content with completeness of rules. The real problem however is to find a set of rules also satisfying the subformula property. [126, p. 316]

Several generalisations of ordinary sequent calculi have therefore been proposed to better deal with modal logic (and non-classical logics), such as display calculi, hypersequents, and labelled calculi. We refer the reader to [158, 100] for surveys of such methods.¹

Of particular interest to us is the case of **GL**. A sequent system for the logic enjoying cut-elimination was first obtained in [125, 85] by the addition of the rule

$$\Box_{\mathbf{GL}} \frac{\Gamma, \Box \Gamma, \Box \varphi \Rightarrow \varphi}{\Box \Gamma \Rightarrow \Box \varphi}$$

¹See also Chapters 3 and 4 below, where we introduce a hypersequent calculus for the temporal logic **CTL*** and a labelled calculus for the intuitionistic temporal logic **iLTL**, respectively.

to a system for CPC.² It is of course not obvious why this is the ‘right’ rule; the reader may consult [85, 126] for an analysis of various natural but unsuccessful attempts at obtaining a sequent calculus for GL admitting cut-elimination.

Proof systems for GL with better structural properties have since been proposed, for example the labelled calculus of [99] and the tree-hypersequent calculus of [109].

More recently, Shamkanov [134] introduced a sequent calculus for GL obtained, not by enriching the syntax or finding another suitable modal rule, but by admitting cyclic proofs in a standard system for K4. These are finite derivations in which a leaf may be non-axiomatic provided that it be labelled by a sequent appearing earlier on the path from the root to the leaf. This work is generalised in [69] to sequent calculi extending standard systems for classical or intuitionistic logic with modal rules satisfying certain constraints.

Following a similar line of research, a labelled calculus for GL is obtained in [35] by allowing ill-founded proofs in a labelled system for K4. Moreover, restricting the system for GL to sequents having at most one formula in the consequent yields a system for intuitionistic GL.

From his cyclic system, Shamkanov was able to prove the Lyndon interpolation property for GL syntactically, a result that he had first shown in [133] by semantic means. The other finitary proof systems that we have mentioned do not seem to be suitable for this purpose because all of them contain modal inference rules in which, as in \Box_{GL} , the principal formula undergoes a polarity change. Shamkanov’s system can also be used to build uniform Craig interpolants for GL, as explained in the conclusion of Chapter 5 below.

We are also interested in the (cyclic) proof theory of S4Grz. The first ordinary sequent calculus for the logic, with a semantical proof of cut-elimination, was given in [8], and a cut-free, labelled one can be found in [40].

Similarly to the work done in [133] for GL, Savateev and Shamkanov [130] proposed a cyclic system for S4Grz whose underlying, acyclic system is easily seen to be sound and complete with respect to S4. Whereas any cycle is acceptable in the cyclic system for GL, this is not the case for S4Grz. Instead, a (finitary) correctness condition is imposed on the cycles in proofs.³

We summarise the finding of these two cyclic systems by saying that GL and S4Grz are *cyclic companions* of K4 and S4, respectively (see Definition 2.2.2 below).

²The cut-elimination proof in [85], however, contained a mistake that was identified and corrected in [154, 8]. See also [64], where the cut-elimination proof from [154] is clarified and adapted to the multiset formulation of the calculus.

³As we shall see, the cyclic system for GL also comes with a correctness condition, yet an implicit one that every cycle satisfies. The condition, however, must be made explicit if one adds weakening and contraction rules to the system.

This terminology is inspired by the one used in [69], where the notion of a *circular companion* of a calculus is defined.

Despite the simplicity of the cyclic systems in [134, 130], it is not *a priori* clear why adding cycles to a calculus for **K4** or **S4** should yield one for **GL** or **S4Grz**. Soundness and completeness are established in [134, 130] by means of translations between ordinary sequent calculi for the logics and the cyclic systems or their ill-founded variants, thus not shedding much light on the question.

This chapter aims to elucidate the relation between a modal logic and its cyclic companion(s). In contrast to the proof-theoretic analysis of [69], we propose to understand cyclic companions by a combination of proof-theoretic and semantical considerations. We argue that cyclic systems like the ones in [134, 130] arise from natural correspondences between cycles in derivations and infinite chains in frames.

Consider **GL**, for instance, which is characterised by the class of transitive and conversely well-founded frames [14, 132]. The former condition is first-order definable and only involves three states of the frame at a time, whereas the latter concerns infinite chains and is known not to be expressible as a first-order property ([14, Ch. 4], [13, § 3.2]). We shall see that cycles capture in a natural way frame conditions involving infinite ascending chains, such as converse well-foundedness and its weak version. This explains why a calculus for **K4**, when coupled with a suitable notion of cycle, captures exactly the theorems of **GL**. An analogous explanation may be given for **S4Grz** and **S4**. In other words: the ‘finitary’, first-order frame conditions are captured by the underlying, acyclic system, while the cycles take care of the ‘infinitary’, second-order ones.

The logic **GL**, however, is not the only one characterised by a class of transitive frames satisfying some second-order condition on infinite chains. Indeed, the logic **K4Grz**, which may be informally described as **S4Grz** minus reflexivity, is characterised by transitive and weakly conversely well-founded frames [7, 94]. So **K4Grz** shares its first-order frame condition with **GL** and its second-order one with **S4Grz** (see Table 2.2 below). According to the view of cyclic companionship that we have just outlined, then, it should be possible to obtain a cyclic system for **K4Grz** by adding cycles to an ordinary sequent calculus for **K4**, not in the way followed in [134] for **GL**, but rather with cycles analogous to the ones in [130] for **S4Grz**. We show in Section 2.3 that this is the case, thus establishing that **K4Grz** is another cyclic companion of **K4**.

Due to the close correspondence between cycles and infinite chains, the frame conditions that one can capture by adding cycles to acyclic systems are probably limited to variations of converse well-foundedness. Therefore, we do not consider the cyclic approach to the proof theory of modal logic to be as fruitful as other methods, such as hypersequent or labelled calculi. Nevertheless, we believe that the cyclic systems studied in this chapter serve as a good introduction to cyclic

proof theory in general and to the more complicated cyclic systems of the next chapters in particular.

Outline of the chapter. Section 2.1 introduces the modal logics that we are interested in, namely: **K4**, **S4**, **GL**, **S4Grz**, and **K4Grz**. Section 2.2 analyses the cyclic systems from [133, 130] and shows the close correspondence between cycles in proofs and infinite chains in frames that enables the construction of such systems. Section 2.3 tests the ideas presented in Section 2.2 by adding cycles to a system for **K4** in order to obtain one for **K4Grz**, thus showing that **K4Grz**, like **GL**, is a cyclic companion of **K4**. Section 2.4 concludes the chapter and discusses some further lines of research based on the material therein.

2.1 Propositional normal modal logics

As mentioned in the introduction, we restrict our attention to modal logics extending classical propositional logic (CPC). Familiarity with both classical logic and modal logic is assumed. The reader is referred to [24, Ch. 1] for an introduction to classical logic, and to [24, 13] for mathematically-oriented introductions to modal logic.

The *language of (propositional) modal logic*, denoted by \mathcal{L}_\square , consists of the following: countably many *propositional letters* drawn from a set **Prop**; the *constants* \perp (falsum) and \top (verum); the *Boolean connectives* \wedge (conjunction), \vee (disjunction), $\bar{\cdot}$ (negation) and \rightarrow (implication); and the *modal operators* \square (box) and \diamond (diamond). The *formulas* of modal logic are given by the following grammar:

$$\varphi ::= \perp \mid \top \mid p \mid \bar{p} \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi) \mid (\varphi \rightarrow \varphi) \mid (\square\varphi) \mid (\diamond\varphi),$$

where p ranges over **Prop**. Formulas are denoted by small Greek letters $\alpha, \beta, \varphi, \dots$, and sets or multisets of formulas by capital Greek letters $\Gamma, \Delta, \Sigma, \dots$. The collection of all modal logic formulas is denoted by Form_\square , and the set of all *subformulas* of a formula φ , defined as usual, is denoted by $\text{Sub}(\varphi)$. A *literal* is a formula of the form \top , \perp , p or \bar{p} . Given a set or multiset of formulas Γ , we let $\circ\Gamma := \{\circ\gamma \mid \gamma \in \Gamma\}$, for $\circ \in \{\square, \diamond\}$.

If no ambiguity arises, we drop the outer parenthesis and stipulate that \square and \diamond bind more strongly than \wedge and \vee , and that these, in turn, bind more strongly than \rightarrow .

It is well known that some of the constants, connectives and operators that we have chosen as primitive are definable in terms of others. For example, \perp , \rightarrow and \square suffice to define \top , $\bar{\cdot}$, \wedge , \vee and \diamond by the equivalences $\top \equiv \neg\perp$, $\bar{p} \equiv \neg p$, $\varphi \wedge \psi \equiv \neg(\varphi \rightarrow \neg\psi)$, $\varphi \vee \psi \equiv \neg\varphi \rightarrow \psi$ and $\diamond\varphi \equiv \neg\square\neg\varphi$, where $\neg\varphi := \varphi \rightarrow \perp$ is

the *negation* of φ . Depending on the context, we shall restrict our language to a subset of \mathcal{L}_\square and rely on these equivalences.

The *dual* of a formula φ , in symbols φ^∂ , is inductively defined as follows:

$$\begin{array}{ll} \perp^\partial := \top & \top^\partial := \perp \\ p^\partial := \bar{p} & \bar{p}^\partial := p \\ (\varphi \wedge \psi)^\partial := \varphi^\partial \vee \psi^\partial & (\varphi \vee \psi)^\partial := \varphi^\partial \wedge \psi^\partial \\ (\Box\varphi)^\partial := \Diamond\varphi^\partial & (\Diamond\varphi)^\partial := \Box\varphi^\partial \\ (\varphi \rightarrow \psi)^\partial := \varphi \wedge \psi^\partial & \end{array}$$

A *substitution* is a map $\sigma : \mathbf{Prop} \rightarrow \mathbf{Form}_\square$. Every substitution σ induces a map $\sigma^* : \mathbf{Form}_\square \rightarrow \mathbf{Form}_\square$ given by:

- (i) $\sigma^*(\perp) := \perp$ and $\sigma^*(\top) := \top$;
- (ii) $\sigma^*(p) := \sigma(p)$ and $\sigma^*(\bar{p}) := \sigma(p)^\partial$, for every $p \in \mathbf{Prop}$;
- (iii) $\sigma^*(\varphi \star \psi) := \sigma^*(\varphi) \star \sigma^*(\psi)$, for $\star \in \{\wedge, \vee, \rightarrow\}$;
- (iv) $\sigma^*(\circ\varphi) := \circ\sigma^*(\varphi)$, for $\circ \in \{\Box, \Diamond\}$.

We abuse notation and denote σ^* by σ .

As is often the case when studying modal logics, we define them as sets of formulas containing certain axioms and closed under the inference rules of *modus ponens*, substitution, and necessitation.

2.1.1. DEFINITION (Modal logic). A (*propositional normal*) modal logic \mathbf{L} is a collection of modal logic formulas containing:

- $(p \rightarrow q) \rightarrow ((q \rightarrow r) \rightarrow (p \rightarrow r))$;
- $(\neg p \rightarrow p) \rightarrow p$;
- $p \rightarrow (\neg p \rightarrow q)$;
- the *normality* axiom \mathbf{k} : $\Box(p \rightarrow q) \rightarrow (\Box p \rightarrow \Box q)$,

and which is closed under the following rules:

- *modus ponens*: if $\alpha, \alpha \rightarrow \beta \in \mathbf{L}$, then $\beta \in \mathbf{L}$;
- *substitution*: if $\alpha \in \mathbf{L}$, then $\sigma(\alpha) \in \mathbf{L}$ for every substitution σ ;
- *necessitation*: if $\alpha \in \mathbf{L}$, then $\Box\alpha \in \mathbf{L}$.

†

The first three of the formulas listed in Definition 2.1.1 are well-known axioms for CPC, originally due to Łukasiewicz [90, 72].

Given a modal logic L and a formula φ , we write $L \vdash \varphi$ if $\varphi \in L$, and we denote by $L \oplus \varphi$ the smallest modal logic containing $L \cup \{\varphi\}$. Following standard notation (see, e.g., [24]), we denote the smallest modal logic by K .

Let L be a modal logic and G a sequent calculus. We say that G is *sound* for L if $G \vdash \varphi$ implies $L \vdash \varphi$ for every formula φ . And G is *complete* for L if $L \vdash \varphi$ implies $G \vdash \varphi$ for every formula φ .

Modal logics admit different semantics: relational (or Kripke), algebraic, topological (see [13]). We restrict our attention to relational semantics, for we are interested in relating cycles in proofs to infinite chains on Kripke frames.

A (*Kripke*) *frame* is a pair $\mathcal{F} = (W, R)$, where W is a non-empty set of *worlds* or *states* and $R \subseteq W \times W$ is a binary relation on W , called the *accessibility relation* of \mathcal{F} . For every $s \in W$, we define $R[s] := \{t \in W \mid s R t\}$. An *infinite (ascending) R-chain* in \mathcal{F} is an infinite sequence of states $(s_n)_{n < \omega}$ such that $s_n R s_{n+1}$ for every $n < \omega$.

A *valuation* on a frame $\mathcal{F} = (W, R)$ is a map $V : \mathbf{Prop} \rightarrow 2^W$. A (*Kripke*) *model* is a triple $\mathcal{M} = (W, R, V)$, where $\mathcal{F} := (W, R)$ is a frame and V a valuation on \mathcal{F} . We say that the model \mathcal{M} is *based on* the frame \mathcal{F} . Given a model $\mathcal{M} = (W, R, V)$, we inductively define a *satisfaction* or *forcing* relation \Vdash between states of \mathcal{M} and formulas in the usual manner:

- $\mathcal{M}, s \not\Vdash \perp$ and $\mathcal{M}, s \Vdash \top$;
- $\mathcal{M}, s \Vdash p$ if, and only if, $s \in V(p)$, for every $p \in \mathbf{Prop}$;
- $\mathcal{M}, s \Vdash \bar{p}$ if, and only if, $\mathcal{M}, s \not\Vdash p$, for every $p \in \mathbf{Prop}$;
- $\mathcal{M}, s \Vdash \varphi \wedge \psi$ if, and only if, $\mathcal{M}, s \Vdash \varphi$ and $\mathcal{M}, s \Vdash \psi$;
- $\mathcal{M}, s \Vdash \varphi \vee \psi$ if, and only if, $\mathcal{M}, s \Vdash \varphi$ or $\mathcal{M}, s \Vdash \psi$;
- $\mathcal{M}, s \Vdash \varphi \rightarrow \psi$ if, and only if, $\mathcal{M}, s \Vdash \varphi$ implies $\mathcal{M}, s \Vdash \psi$;
- $\mathcal{M}, s \Vdash \Box \varphi$ if, and only if, $\mathcal{M}, t \Vdash \varphi$ for every $t \in R[s]$;
- $\mathcal{M}, s \Vdash \Diamond \varphi$ if, and only if, $\mathcal{M}, t \Vdash \varphi$ for some $t \in R[s]$.

If $\mathcal{M}, s \Vdash \varphi$, we say that φ is *true* in s , or that s *satisfies* φ . We write $\mathcal{M} \models \varphi$ if $\mathcal{M}, s \Vdash \varphi$ for every $s \in W$. Given a frame $\mathcal{F} = (W, R)$, we write $\mathcal{F} \models \varphi$, and say that φ is *valid* in \mathcal{F} , if $(W, R, V) \models \varphi$ for every valuation V on \mathcal{F} . Given a class of frames \mathcal{K} , we write $\mathcal{K} \models \varphi$, and say that φ is *\mathcal{K} -valid*, if $\mathcal{F} \models \varphi$ for every $\mathcal{F} \in \mathcal{K}$. Finally, given any formulas φ and ψ , we write $\varphi \equiv \psi$, and say that φ and ψ are *equivalent*, if for every model $\mathcal{M} = (W, R, V)$ and every $s \in W$, we have $\mathcal{M}, s \Vdash \varphi$ if, and only if, $\mathcal{M}, s \Vdash \psi$.

Let \mathcal{K} be a class of frames and L a modal logic. We say that L is *sound* with respect to \mathcal{K} if $L \vdash \varphi$ implies $\mathcal{K} \models \varphi$; that L is *complete* with respect to \mathcal{K} if $\mathcal{K} \models \varphi$ implies $L \vdash \varphi$; and that L is *characterised* by \mathcal{K} if L is both sound and complete with respect to \mathcal{K} .

Let $\mathcal{F} = (W, R)$ be a Kripke frame. We say that \mathcal{F} is *reflexive* (*irreflexive*, *transitive*, *antisymmetric*, *partially ordered*) if R is reflexive (respectively, irreflexive, transitive, antisymmetric, a partial order on W). We say that \mathcal{F} is *conversely well-founded*, or *Noetherian*, if there are no infinite R -chains in \mathcal{F} . Finally, we say that \mathcal{F} is *weakly conversely well-founded*, or *weakly Noetherian*, if every infinite R -chain in \mathcal{F} is eventually constant, i.e., for every infinite R -chain $(s_n)_{n < \omega}$ in \mathcal{F} there is an $i < \omega$ such that $s_i = s_j$ for all $j \geq i$. Clearly, being Noetherian is equivalent to being weakly Noetherian and irreflexive.

Weak converse well-foundedness is the natural formulation of converse well-foundedness in the absence of irreflexivity. It is easy to see that a frame \mathcal{F} is weakly conversely well-founded if, and only if, \mathcal{F} satisfies the following *strong chain condition (SCC)*: for every infinite R -chain $(s_n)_{n < \omega}$ in \mathcal{F} there is an $i < \omega$ such that $s_i = s_{i+1}$.

Table 2.2 lists several well-known modal logics defined in terms of the axioms in Table 2.1, together with characteristic classes of frames for them. We assume that the reader is familiar with the axioms **k**, **4** and **t**, as well as with the corresponding logics **K**, **K4** and **S4** characterised, respectively, by all frames, by all transitive frames, and by all reflexive and transitive frames [24].

The logic **GL** is the well-known *Gödel–Löb provability logic*, in which the operator \Box denotes provability in some arithmetical theory, typically Peano arithmetic. We shall not concern ourselves with **GL** besides the analysis of the cyclic system for the logic given in [134]. The interested reader may consult [15] for a historical account of the emergence of provability logic, as well as Boolos’s classic book [14].

The origins of the *Grzegorzcyk logic S4Grz* can be traced back to Grzegorzcyk’s investigation of the connections between intuitionistic and modal logic [65, 14, 92]. The logic was first axiomatised by adding to **S4** the axiom

$$\Box[(\Box(q \rightarrow \Box p) \rightarrow \Box p) \wedge (\Box(\neg q \rightarrow \Box p) \rightarrow \Box p)] \rightarrow \Box p,$$

due to Grzegorzcyk [65]. Subsequently, Segerberg [132] found the simpler, current axiomatisation of **S4Grz** in terms of the axiom **grz**, first studied by Sobociński [138].

A common, alternative name for **S4Grz** is **Grz** (see, e.g., [8, 24, 92, 94, 130]). This is justified by the fact that **S4Grz** = **K** \oplus **grz** [63, 14, 24]. Moreover, all of the following hold:⁴

$$\mathbf{K} \oplus \mathbf{grz} = \mathbf{K4} \oplus \mathbf{grz} = \mathbf{S4} \oplus \mathbf{grz} = \mathbf{S4} \oplus \mathbf{grz}_1, \quad (2.1)$$

⁴The last equality was first proved in [132], and the first one in [12].

Axiom	Definition
k (<i>normality axiom</i>)	$\Box(p \rightarrow q) \rightarrow (\Box p \rightarrow \Box q)$
4 (<i>transitivity axiom</i>)	$\Box p \rightarrow \Box \Box p$
t (<i>reflexivity axiom</i>)	$\Box p \rightarrow p$
gl (<i>Löb axiom</i>)	$\Box(\Box p \rightarrow p) \rightarrow \Box p$
grz (<i>Grzegorzczuk axiom</i>)	$\Box(\Box(p \rightarrow \Box p) \rightarrow p) \rightarrow p$
grz ₁	$\Box(\Box(p \rightarrow \Box p) \rightarrow p) \rightarrow \Box p$

TABLE 2.1: Some modal logic axioms.

but note that

$$\mathbf{K4} \oplus \text{grz} \neq \mathbf{K4} \oplus \text{grz}_1$$

because $\mathbf{S4Grz} \neq \mathbf{K4Grz}$ (see below). We follow [132, 14] and denote $\mathbf{S4} \oplus \text{grz}$ by $\mathbf{S4Grz}$, in order to highlight the similarities and differences between $\mathbf{S4Grz}$ and $\mathbf{K4Grz}$.

As with \mathbf{GL} , our interest in $\mathbf{S4Grz}$ is restricted to the cyclic system for the logic introduced in [130]. For an overview of $\mathbf{S4Grz}$, its history, and its close connection to \mathbf{GL} , we refer the reader to [92].

The logic $\mathbf{K4Grz}$ has been called the *weak Grzegorzczuk logic* [88]. According to [61], the name $\mathbf{K4Grz}$ originates in [51] and [44], but the logic had been studied earlier under different names (e.g., in [113, 63] it is referred to as \mathbf{Go} , and in [7] it is called \mathbf{G}_0). This logic is also known as $\mathbf{Grz.1}$ [94]. We avoid this notation because it misleadingly suggests that $\mathbf{K4Grz}$ is an extension of \mathbf{Grz} , when in fact this is not the case: the reflexivity axiom $\Box p \rightarrow p$ belongs by definition to $\mathbf{S4Grz}$, but it cannot belong to $\mathbf{K4Grz}$ because there are frames which are transitive and weakly conversely well-founded but not reflexive (e.g., the frame consisting of a single irreflexive state). For an overview of results about $\mathbf{K4Grz}$, the reader may consult [88].

Just as the names $\mathbf{S4Grz}$ and $\mathbf{K4Grz}$ suggest, from the point of view of Kripke semantics $\mathbf{K4Grz}$ is the non-reflexive version of $\mathbf{S4Grz}$, in the sense that we have:

$$\mathbf{K4Grz} \subsetneq \mathbf{K4Grz} \oplus \mathbf{t} = \mathbf{S4Grz},$$

where the equality follows from (2.1) and the proper inclusion from the fact that, as we have observed, $\mathbf{K4Grz} \not\vdash \mathbf{t}$.

Converse well-foundedness is well known not to be first-order definable (see, e.g., [14, Ch. 4] or [13, § 3.2]). Since irreflexivity is a first-order property, it follows that weak converse well-foundedness is not first-order definable either. Thus, the frame conditions given in Table 2.2 for \mathbf{GL} , $\mathbf{S4Grz}$ and $\mathbf{K4Grz}$ can be split up into two groups: ‘finitary’, first-order definable conditions (reflexivity and transitivity); and ‘infinitary’ conditions which are not expressible as first-order properties (converse

Logic	Definition	Characteristic frames
K4	$K \oplus 4$	transitive [24]
S4	$K4 \oplus t$	reflexive and transitive [24]
GL	$K4 \oplus gl$	transitive and Noetherian [14, 132]
S4Grz	$S4 \oplus grz$	reflexive, transitive and weakly Noetherian [14, 132]
K4Grz	$K4 \oplus grz_1$	transitive and weakly Noetherian [7, 94]

TABLE 2.2: Some modal logics and characteristic frames thereof.

well-foundedness and its weak version). As we shall see in the next section, it is natural to try to find proof systems for these logics by adding cyclicity to acyclic sequent calculi for logics characterised by frame conditions of the first kind only, such as K4 and S4.

2.2 Cycles and infinite chains

Our study of cyclic companions begins with an analysis of the cyclic systems given in [133, 130] for GL and S4Grz. Both systems work with multiset-based sequents, so we adopt the following:

2.2.1. DEFINITION (Modal sequent). A (*modal*) *one-sided sequent* is a finite multiset of formulas. And a (*modal*) *two-sided sequent* is a pair (Γ, Δ) , henceforth written $\Gamma \Rightarrow \Delta$, where Γ and Δ are finite multisets of formulas. \dashv

The *antecedent* of the two-sided sequent $\Gamma \Rightarrow \Delta$ is the multiset Γ , and Δ is the *consequent* or *succedent* of the sequent. As usual, the *interpretation* of a one-sided sequent Γ is the formula $\Gamma^\sharp := \bigvee \Gamma$, and the *interpretation* of a two-sided sequent $\Gamma \Rightarrow \Delta$ is $(\Gamma \Rightarrow \Delta)^\sharp := \bigwedge \Gamma \rightarrow \bigvee \Delta$. A sequent S is *valid* if S^\sharp is valid. And we write $\mathcal{M}, s \Vdash S$ if $\mathcal{M}, s \Vdash S^\sharp$.

The difference between working with set-based sequents or multiset-based ones might seem insignificant. However, the availability of some form of contraction is known to be essential to obtaining complete calculi for modal logics characterised by classes of reflexive frames. Therefore, in a multiset setting without an explicit contraction rule it might be necessary to build contraction into some of the rules. This is the case, for example, of the rule *refl* of the system G3S4 for S4, depicted in Figure 2.3. If the formula $\Box\varphi$ were not preserved in the premise, then the formula $\Diamond\Box(\Diamond p \rightarrow \Box\Diamond p)$, valid on all reflexive and transitive frames, would no longer be provable (see [63, § 3.6] and [112, § 3.5]).

From the beginning of this chapter we have been using the expression ‘cyclic companion’ without properly defining it. Recalling the definitions about cyclic derivations given in Section 1.3, we may define it thus:

$$\begin{array}{c}
\text{ax} \frac{}{\Gamma, p \Rightarrow p, \Delta} \\
\text{L} \rightarrow \frac{\Gamma \Rightarrow \varphi, \Delta \quad \Gamma, \psi \Rightarrow \Delta}{\Gamma, \varphi \rightarrow \psi \Rightarrow \Delta}
\end{array}
\qquad
\begin{array}{c}
\text{ax}_{\perp} \frac{}{\Gamma, \perp \Rightarrow \Delta} \\
\text{R} \rightarrow \frac{\Gamma, \varphi \Rightarrow \psi, \Delta}{\Gamma \Rightarrow \varphi \rightarrow \psi, \Delta}
\end{array}$$

FIGURE 2.1: Rules of the system G3.

2.2.2. DEFINITION (Cyclic companion). Let L_1 and L_2 be modal logics. We say that L_2 is a *cyclic companion* of L_1 if there exists a cyclic proof system G such that the following hold for every formula φ :

- (i) $G \vdash \varphi$ if, and only if, $L_2 \vdash \varphi$;
- (ii) $\mathcal{B}(G) \vdash \varphi$ if, and only if, $L_1 \vdash \varphi$.

–

In other words: L_2 is a cyclic companion of L_1 if there is a sound and complete cyclic system G for L_2 such that the underlying, acyclic system $\mathcal{B}(G)$ is sound and complete for L_1 .

2.2.3. REMARK. The reader might have expected the definition of cyclic companionship to be based on adding cycles to an acyclic system rather than removing them from a cyclic one. Whilst the latter operation has a straightforward definition, it is not clear how to adequately define the former to accommodate all the cases that we are interested in: as explained in Section 2.2.2 below, in general the addition of cycles to a calculus amounts to more than simply admitting cyclic derivations, for one has to make sure to exclude viciously circular reasoning.

2.2.4. REMARK. In contrast to the proof-theoretic approach followed in [69], where the notion of *circular companion* is defined for proof systems, our notion of cyclic companionship applies to logics. This is line with our plan to investigate the rôle that cycles play in [134, 130] from a point of view which combines proof-theoretic and semantical considerations.

As we shall see, cyclic companions need not be unique. This is because, informally, there are different ways of adding cycles to a system (more precisely: different cyclic systems might be based on calculi which are sound and complete for the same logic).

It is convenient to fix a sequent calculus for CPC to which we can add modal rules to characterise different modal logics. To this end, let $G3$ be the two-sided sequent calculus with rules given in Figure 2.1. The system $G3$ is the $\{\perp, \rightarrow\}$ -fragment of the well-known calculus $G3cp$ for CPC from [152].

$$\begin{array}{c}
\text{ax} \frac{}{\Gamma, \varphi, \varphi^\theta} \qquad \text{ax}_\top \frac{}{\Gamma, \top} \\
\wedge \frac{\Gamma, \varphi \quad \Gamma, \psi}{\Gamma, \varphi \wedge \psi} \qquad \vee \frac{\Gamma, \varphi, \psi}{\Gamma, \varphi \vee \psi} \qquad \Box_{\text{K4}} \frac{\Gamma, \Diamond \Gamma, \varphi}{\Diamond \Gamma, \Box \varphi, \Delta}
\end{array}$$

FIGURE 2.2: Rules of the system GL_{circ} from [134].

2.2.1 GL as a cyclic companion of K4

Shamkanov [134] obtained a sound and complete cyclic proof system GL_{circ} for GL by admitting cyclic proofs in a standard calculus for K4. The system GL_{circ} works on one-sided sequents and the sublanguage of \mathcal{L}_\Box lacking implication. The rules are given in Figure 2.2.

The creation of cycles in GL_{circ} derivations⁵ is possible due to the preservation of $\Diamond \Gamma$ in the premise of rule \Box_{K4} . For example:

$$\begin{array}{c}
\Box_{\text{K4}} \frac{\perp, \Diamond \Box \perp, \Box \perp}{\perp, \Diamond \Box \perp, \Box \perp} \\
\Box_{\text{K4}} \frac{\perp, \Diamond \Box \perp, \Box \perp}{\Box \perp, \Diamond \Box \perp} \\
\vee \frac{\Box \perp, \Diamond \Box \perp}{\Box \perp \vee \Diamond \Box \perp}
\end{array}$$

2.2.5. DEFINITION (GL_{circ} proof). A GL_{circ} proof of a formula φ is a GL_{circ} derivation \mathcal{T} of φ such that every leaf of \mathcal{T} is either axiomatic or a repeat. \dashv

Observe that any cycle in a GL_{circ} derivation is acceptable, which might give the impression that GL_{circ} proofs allow viciously circular, unsound reasoning. This is not the case because in fact GL_{circ} comes with an implicit correctness condition that any cycle satisfies, namely: for every repeat l in a proof \mathcal{T} , say with companion c , there must be an instance of rule \Box_{K4} in $[c, l]_{\mathcal{T}}$, i.e., there must be a vertex in $[c, l]_{\mathcal{T}}$ at which rule \Box_{K4} is applied (we call such vertices *modal*). This condition is always satisfied because rule \Box_{K4} is the only one that need not decrease the complexity of the sequent to which it is applied. However, if one wishes to add weakening (**wk**) and contraction (**contr**) to the system, either explicitly or by using set-based instead of multiset-based sequents, then the correctness condition must be made explicit to preclude unsound reasoning, such as

$$\begin{array}{c}
\text{wk} \frac{\perp}{\perp} \\
\text{contr} \frac{\perp, \perp}{\perp}
\end{array}$$

⁵Recall the definition of (cyclic) derivation given in Section 1.3.

from being considered valid. The same applies to the intuitionistic version of GL_{circ} [69, 136], as well as to a variant of GL_{circ} considered in the conclusion of Chapter 5 below.

The system GL_{circ} is indeed sound and complete:

2.2.6. PROPOSITION ([134]). *For every formula φ , the following hold:*

- (i) $\text{GL}_{\text{circ}} \vdash \varphi$ if, and only if, $\text{GL} \vdash \varphi$;
- (ii) $\mathcal{B}(\text{GL}_{\text{circ}}) \vdash \varphi$ if, and only if, $\text{K4} \vdash \varphi$.

2.2.7. COROLLARY. *The logic GL is a cyclic companion of K4 .*

Recall that GL is characterised by the class of all transitive, Noetherian frames. By Proposition 2.2.6(ii), the underlying, acyclic system $\mathcal{B}(\text{GL}_{\text{circ}})$ characterises exactly the class of transitive frames, whence it follows that it is the cycles of GL_{circ} that capture converse well-foundedness. To explain how this is so, we now provide a proof of the soundness of GL_{circ} that highlights the correspondence between cycles in GL_{circ} proofs and infinite ascending chains on Kripke frames. This is not the proof given in [134], for there soundness and completeness of GL_{circ} are established via translations among GL_{circ} , the corresponding ill-founded system, and a standard, acyclic sequent calculus for GL .

2.2.8. PROPOSITION (cf. [134, Lem. 3.7]). *For every formula φ , if $\text{GL}_{\text{circ}} \vdash \varphi$, then $\text{GL} \vdash \varphi$.*

Proof. Let \mathcal{T} be a GL_{circ} proof of φ , and suppose towards a contradiction that $\text{GL} \not\vdash \varphi$. Then, there is a model $\mathcal{M} = (W, R, V)$ based on a transitive Noetherian frame and a state $s \in W$ such that $\mathcal{M}, s \not\vdash \varphi$. We inductively build an infinite path $(u_n)_{n < \omega}$ through \mathcal{T}° and an infinite sequence $(s_n)_{n < \omega}$ of states in W such that, for every $i < \omega$:

- (i) $\mathcal{M}, s_i \not\vdash \Gamma_i$, where Γ_i is the label of u_i ;
- (ii) either u_i is not modal and $s_i = s_{i+1}$, or $s_i R s_{i+1}$.

For the base case, let u_0 be the root of \mathcal{T} and $s_0 := s$. For the inductive case, assume that u_n and s_n have been defined. Note that u_n cannot be axiomatic because $\mathcal{M}, s_n \not\vdash \Gamma_n$. If u_n is a repeat, we let $u_{n+1} := c_{u_n}$ and $s_{n+1} := s_n$. Otherwise, we distinguish cases according to the rule R applied at u_n .

Suppose first that $R \in \{\wedge, \vee\}$. Clearly, then, there is an immediate successor v of u_n such that $\mathcal{M}, s_n \not\vdash \Gamma$, where Γ is the label of v . We let $u_{n+1} := v$ and $s_{n+1} := s_i$.

$$\text{refl} \frac{\Gamma, \varphi, \Box\varphi \Rightarrow \Delta}{\Gamma, \Box\varphi \Rightarrow \Delta} \qquad \Box_{S4} \frac{\Box\Pi \Rightarrow \varphi}{\Gamma, \Box\Pi \Rightarrow \Box\varphi, \Delta}$$

FIGURE 2.3: Additional rules for the system G3S4.

Suppose now that u_n is modal, and let

$$\Box_{K4} \frac{\Gamma, \Diamond\Gamma, \psi}{\Diamond\Gamma, \Box\psi, \Delta}$$

be the instance of \Box_{K4} at u_n . By the inductive hypothesis, $\mathcal{M}, s_n \not\models \Diamond\Gamma, \Box\psi, \Delta$, whence there is a state $t \in W$ such that $s_n R t$ and $\mathcal{M}, t \not\models \psi, \Gamma$. Additionally, $\mathcal{M}, t \not\models \Diamond\Gamma$ because otherwise transitivity would yield $\mathcal{M}, s_n \models \Diamond\Gamma$. We thus let u_{n+1} be the unique immediate successor of u_n and $s_{n+1} := t$.

Note that the path $\pi := (u_n)_{n < \omega}$ must be infinite because axioms are labelled by valid sequents. By Proposition 1.2.3, π passes through infinitely many instances of rule \Box_{K4Grz} , so by construction there are infinitely many $i < \omega$ such that $s_i R s_{i+1}$ and thus the sequence $(s_n)_{n < \omega}$ contains a subsequence which is an infinite ascending R -chain, contradicting the converse well-foundedness of (W, R) . ■

Looking at the proof of Proposition 2.2.8, we see that instances of rule \Box_{K4} correspond semantically to transitions through the accessibility relation of a model. So, since every cycle in a GL_{circ} proof contains at least one modal vertex, infinite paths through (the cycles of) a GL_{circ} proof correspond to infinite ascending chains. The assumption that the formula at the root of a proof is not satisfied in a state of a transitive Noetherian model yields one such infinite path, whence soundness of GL_{circ} immediately follows. Informally: something happens on every cycle (a transition) that cannot happen infinitely often. This explains why any cycle should be considered sound.

2.2.2 S4Grz as a cyclic companion of S4

We now turn to S4Grz. Recall (see Table 2.2 above) that S4Grz is characterised by the class of all reflexive, transitive and weakly Noetherian frames. Since S4 is the logic of reflexive and transitive frames, we might be tempted to admit cyclic proofs in a system for S4 to obtain one for S4Grz. This is not the right approach, as we now show.

Let G3S4 be the system obtained by adding the rules in Figure 2.3 to G3. It is well known that G3S4 is sound and complete for S4 [79, 152]. Informally, rule *refl* corresponds to reflexivity, and \Box_{S4} to transitivity. Let us see that admitting cyclic proofs in G3S4 does not yield a system for S4Grz.

$$\begin{array}{c}
\Box_{S4} \frac{(\dagger) \quad \Box(\Box p \rightarrow p) \Rightarrow p}{\Box(\Box p \rightarrow p) \Rightarrow p, \Box p} \quad \text{ax} \frac{}{\Box(\Box p \rightarrow p), p \Rightarrow p} \\
L \rightarrow \frac{}{\Box(\Box p \rightarrow p), \Box p \rightarrow p \Rightarrow p} \\
\text{refl} \frac{}{\Box(\Box p \rightarrow p), \Box p \rightarrow p \Rightarrow p} \\
\Box_{S4} \frac{(\dagger) \quad \Box(\Box p \rightarrow p) \Rightarrow p}{\Box(\Box p \rightarrow p) \Rightarrow \Box p} \\
R \rightarrow \frac{}{\Rightarrow \Box(\Box p \rightarrow p) \rightarrow \Box p}
\end{array}$$

FIGURE 2.4: A proof of the axiom gl in G3S4° . The unique repeat and its companion are marked by the symbol ‘ \dagger ’.

Suppose that, as Shamkanov did for GL , we accept cyclic proofs in G3S4 . More precisely, let G3S4° be the cyclic system with base G3S4 given as follows:

2.2.9. DEFINITION (G3S4° proof). A G3S4° proof of a formula φ is a G3S4° derivation \mathcal{T} of φ such that every non-axiomatic leaf of \mathcal{T} is a repeat. \dashv

As we are about to see, requiring that there be an instance of rule \Box_{S4} or refl on every cycle in a G3S4° proof makes no difference.

That G3S4° is not sound for S4Grz is easy to see. Indeed, as shown in Figure 2.4, we can prove the axiom gl in G3S4° (cf. [69]). Therefore, the least modal logic containing the theorems of G3S4° contains

$$\text{S4} \oplus \text{gl} = \text{GL} \oplus \mathbf{t} = \text{Form}_\Box$$

and is therefore the inconsistent logic. Semantically, the equality $\text{GL} \oplus \mathbf{t} = \text{Form}_\Box$ stems from the fact that axiom gl demands that there be no infinite ascending chains, whereas validating axiom \mathbf{t} implies the existence of at least one such chain in any non-empty frame. Syntactically, $\text{GL} \oplus \mathbf{t}$ contains both $\Box(\Box p \rightarrow p) \rightarrow \Box p$ and $\Box p \rightarrow p$, so by closure under *modus ponens* and necessitation we get $p \in \text{GL} \oplus \mathbf{t}$ and thus $\text{GL} \oplus \mathbf{t} = \text{Form}_\Box$ by closure under substitution.

As our analysis of the system GL_{circ} suggested, admitting any kind of cycle in G3S4 semantically corresponds to imposing that there be no infinite ascending chains. So, since S4Grz does not impose such restriction, not every cycle should be accepted. Let us then see how to distinguish the ‘good’ cycles from the ‘bad’ ones.

In any weakly Noetherian frame (W, R) , every infinite R -chain is eventually constant. Therefore, the ‘good’ cycles are the ones that, when traversed as in the proof of Proposition 2.2.8, cause transitions of the form $s R t \neq s$, for then we are able to obtain a contradiction as we did in said proof. Analogously, cycles that only cause transitions of the form $s R s$ are to be considered ‘bad’. We may capture this distinction proof-theoretically by adding a seemingly superfluous premise to the rule \Box_{S4} so that it becomes:

$$\begin{array}{c}
\text{ax} \frac{}{\Gamma, p \Rightarrow p, \Delta} \\
\text{L} \rightarrow \frac{\Gamma \Rightarrow \varphi, \Delta \quad \Gamma, \psi \Rightarrow \Delta}{\Gamma, \varphi \rightarrow \psi \Rightarrow \Delta} \\
\text{refl} \frac{\Gamma, \varphi, \Box \varphi \Rightarrow \Delta}{\Gamma, \Box \varphi \Rightarrow \Delta}
\end{array}
\qquad
\begin{array}{c}
\text{ax}_{\perp} \frac{}{\Gamma, \perp \Rightarrow \Delta} \\
\text{R} \rightarrow \frac{\Gamma, \varphi \Rightarrow \psi, \Delta}{\Gamma \Rightarrow \varphi \rightarrow \psi, \Delta} \\
\Box_{S4}^{\circ} \frac{\Gamma, \Box \Pi \Rightarrow \varphi, \Delta \quad \Box \Pi \Rightarrow \varphi}{\Gamma, \Box \Pi \Rightarrow \Box \varphi, \Delta}
\end{array}$$

FIGURE 2.5: Rules of the system Grz_{circ} from [130].

$$\Box_{S4}^{\circ} \frac{\Gamma, \Box \Pi \Rightarrow \varphi, \Delta \quad \Box \Pi \Rightarrow \varphi}{\Gamma, \Box \Pi \Rightarrow \Box \varphi, \Delta}$$

The new premise on the left is clearly superfluous in G3S4 because it is a weakening of the right premise. From the point of view of cyclic proof-theory, however, having both premises allows us to tell ‘good’ and ‘bad’ cycles apart: a cycle is to be considered ‘good’ if, and only if, it contains a vertex labelled by the conclusion of an instance of \Box_{S4}° and a right premise thereof. Indeed, arguing as in the proof of Proposition 2.2.8, in the case of \Box_{S4}° we can now distinguish cases according to whether we transition to a new state or not. In the first case we pick the right premise, and in the latter the left one (see the proof of Proposition 2.2.14 below).

In this way we obtain the cyclic calculus Grz_{circ} for S4Grz due to Savateev and Shamkanov [128, 130]. It should be noted, however, that they do not present their system as arising from an acyclic one for S4 in the manner that we have just described.⁶ In particular, their proof of soundness is based on translations from the ill-founded version of Grz_{circ} to a standard sequent calculus for S4Grz .

We therefore prove that $\mathcal{B}(\text{Grz}_{\text{circ}})$ is indeed sound and complete for S4 , and afterwards, as we did for GL , we provide an alternative proof of the soundness of Grz_{circ} analogous to the proof of Proposition 2.2.8, to better show how rule \Box_{S4}° distinguishes between ‘good’ and ‘bad’ cycles.

The rules of Grz_{circ} are given in Figure 2.5. The system works on two-sided sequents and the $\{\perp, \rightarrow, \Box\}$ -fragment of \mathcal{L}_{\Box} . We call the left premise of rule \Box_{S4}° *stationary*, and the right premise *transitional*.

2.2.10. DEFINITION (Grz_{circ} proof). A Grz_{circ} *proof* of a formula φ is a Grz_{circ} derivation \mathcal{T} with conclusion $\emptyset \Rightarrow \varphi$ and such that, if l is a non-axiomatic leaf of \mathcal{T} , then l is a repeat and there is a vertex in $(c_l, l]_{\mathcal{T}}$ labelled by a transitional premise of an instance of rule \Box_{S4}° . \dashv

⁶And, in fact, they are mostly concerned with the ill-founded version of the system Grz_{circ} , to the extent that Grz_{circ} is not given a name in [130].

Let us see that, as we claimed, Grz_{circ} witnesses that S4Grz is a cyclic companion of S4 . This does not follow from the work in [69] because, due to the stationary premise, rule \Box_{S4}° does not meet the requirements imposed in [69] on modal rules.

2.2.11. LEMMA. *The weakening rule*

$$\text{wk} \frac{\Gamma \Rightarrow \Delta}{\Gamma, \Pi \Rightarrow \Sigma, \Delta}$$

is admissible in $\mathcal{B}(\text{Grz}_{\text{circ}})$.

Proof. It follows by a straightforward induction on the height of $\mathcal{B}(\text{Grz}_{\text{circ}}) + \text{wk}$ proofs, relying on the implicit weakening in rules ax and ax_{\perp} . ■

2.2.12. PROPOSITION. *For every formula φ , the following hold:*

- (i) $\text{Grz}_{\text{circ}} \vdash \varphi$ if, and only if, $\text{S4Grz} \vdash \varphi$;
- (ii) $\mathcal{B}(\text{Grz}_{\text{circ}}) \vdash \varphi$ if, and only if, $\text{S4} \vdash \varphi$.

Proof. Item (i) is proved in [130], so we focus on (ii). The left-to-right direction is straightforward: axiomatic leaves of $\mathcal{B}(\text{Grz}_{\text{circ}})$ proofs are labelled by valid sequents, and all the rules of $\mathcal{B}(\text{Grz}_{\text{circ}})$ preserve validity with respect to reflexive and transitive frames.

For the right-to-left direction, assume that $\text{S4} \vdash \varphi$. Then, there is a G3S4 proof of φ . Every rule of G3S4 is a rule of $\mathcal{B}(\text{Grz}_{\text{circ}})$ except for \Box_{S4} , so by Lemma 2.2.11 it suffices to see that every instance of \Box_{S4} can be simulated in $\mathcal{B}(\text{Grz}_{\text{circ}}) + \text{wk}$:

$$\Box_{\text{S4}} \frac{\Box \Pi \Rightarrow \psi}{\Gamma, \Box \Pi \Rightarrow \Box \psi, \Delta} \quad \rightsquigarrow \quad \frac{\text{wk} \frac{\Box \Pi \Rightarrow \psi}{\Gamma, \Box \Pi \Rightarrow \psi, \Delta}}{\Box_{\text{S4}}^{\circ} \frac{\Box \Pi \Rightarrow \psi}{\Gamma, \Box \Pi \Rightarrow \Box \psi, \Delta}}$$

Therefore, every G3S4 proof can be transformed into a $\mathcal{B}(\text{Grz}_{\text{circ}}) + \text{wk}$ proof and we are done. ■

2.2.13. COROLLARY. *The logic S4Grz is a cyclic companion of S4 .*

We now prove the soundness of Grz_{circ} as we did for GL_{circ} , highlighting the correspondence between the cycles in Grz_{circ} proofs and weak converse well-foundedness.

2.2.14. PROPOSITION (cf. [130, Thm. 3.7]). *For every formula φ , if $\text{Grz}_{\text{circ}} \vdash \varphi$, then $\text{S4Grz} \vdash \varphi$.*

Proof. Let \mathcal{T} be a Grz_{circ} proof of φ , and suppose towards a contradiction that $\text{S4Grz} \not\vdash \varphi$. Then, there is a model $\mathcal{M} = (W, R, V)$ based on a reflexive, transitive and weakly Noetherian frame and a state $s \in W$ such that $\mathcal{M}, s \not\vdash \varphi$. We inductively build an infinite path $(u_n)_{n < \omega}$ through \mathcal{T}° and an infinite sequence $(s_n)_{n < \omega}$ of states in W such that, for every $i < \omega$:

- (i) $\mathcal{M}, s_i \not\vdash \Gamma_i \Rightarrow \Delta_i$, where $\Gamma_i \Rightarrow \Delta_i$ is the label of u_i ;
- (ii) either rule \Box_{S4}° is not applied at u_i and $s_i = s_{i+1}$, or $s_i R s_{i+1}$.

The base case, as well as the inductive case where neither refl nor \Box_{S4}° is applied at u_n , are as in the proof of Proposition 2.2.8. Assume, then, that a rule $R \in \{\text{refl}, \Box_{\text{S4}}^\circ\}$ is applied at u_n .

Suppose first that $R = \text{refl}$, and let

$$\text{refl} \frac{\Gamma, \psi, \Box\psi \Rightarrow \Delta}{\Gamma, \Box\psi \Rightarrow \Delta}$$

be the instance of refl at u_n . By the inductive hypothesis, $\mathcal{M}, s_n \not\vdash \Gamma, \Box\psi \Rightarrow \Delta$, so by reflexivity we have $\mathcal{M}, s_n \not\vdash \Gamma, \psi, \Box\psi \Rightarrow \Delta$. We thus let u_{n+1} be the unique immediate successor of u_n and $s_{n+1} := s_n$.

Suppose now that $R = \Box_{\text{S4}}^\circ$, and let

$$\Box_{\text{S4}}^\circ \frac{\Gamma, \Box\Pi \Rightarrow \psi, \Delta \quad \Box\Pi \Rightarrow \psi}{\Gamma, \Box\Pi \Rightarrow \Box\psi, \Delta}$$

be the instance of \Box_{S4}° at u_n . By the inductive hypothesis, $\mathcal{M}, s_n \not\vdash \Gamma, \Box\Pi \Rightarrow \Box\psi, \Delta$, whence there is a state $t \in W$ such that $s_n R t$ and $\mathcal{M}, t \not\vdash \psi$. If $t = s_n$, we let u_{n+1} be the immediate successor of u_n for the stationary premise. Otherwise, transitivity yields $\mathcal{M}, t \not\vdash \Box\Pi \Rightarrow \psi$ and we let u_{n+1} be the immediate successor of u_n for the transitional premise. In either case we let $s_{n+1} := t$.

The path $\pi := (u_n)_{n < \omega}$ must be infinite because axioms are labelled by valid sequents. By Proposition 1.2.3, there are infinitely many $i < \omega$ such that u_i is labelled by the conclusion of an instance of \Box_{S4}° and u_{i+1} is the immediate successor of u_i for the transitional premise. Hence, by construction there are infinitely many $i < \omega$ such that $s_i R s_{i+1} \neq s_i$ and thus the sequence $(s_n)_{n < \omega}$ contains a subsequence which is an infinite ascending R -chain and not eventually constant, contradicting the weak converse well-foundedness of (W, R) . \blacksquare

2.3 K4Grz as a cyclic companion of K4

Looking at the frame conditions in Table 2.2, we can see that K4Grz shares its first-order condition (transitivity) with GL and its second-order one (weak converse

$$\Box_{K4} \frac{\Pi, \Box\Pi \Rightarrow \varphi}{\Gamma, \Box\Pi \Rightarrow \Box\varphi, \Delta}$$

FIGURE 2.6: The modal rule of the system G3K4.

well-foundedness) with S4Grz. According to the view of cyclic companionship presented in Section 2.2, then, we should expect to be able to obtain a cyclic system for K4Grz by adding cycles to an ordinary sequent calculus for K4, not as done for GL in [134], but rather with cycles analogous to the ones in [130] for S4Grz. In this section we show that this is indeed the case.⁷

Let G3K4 be the system obtained by adding the rule in Figure 2.6 to G3. The calculus G3K4 is well known to be sound and complete for K4 [126]. The only difference between \Box_{K4} and \Box_{S4} is that the latter does not include Π in its premise, as it can be obtained by applying rule *refl*. The logic K4, however, does not impose reflexivity, and thus Π must be present in the premise of \Box_{K4} in order for G3K4 to be complete. For example, it is necessary to prove the transitivity axiom $\Box p \rightarrow \Box\Box p$:

$$\begin{array}{c} \text{ax} \frac{}{p, \Box p \Rightarrow p} \\ \Box_{K4} \frac{p, \Box p \Rightarrow p}{p, \Box p \Rightarrow \Box p} \\ \Box_{K4} \frac{p, \Box p \Rightarrow \Box p}{\Box p \Rightarrow \Box\Box p} \\ \text{R} \rightarrow \frac{\Box p \Rightarrow \Box\Box p}{\emptyset \Rightarrow \Box p \rightarrow \Box\Box p} \end{array}$$

Naturally, allowing any kind of cycle in G3K4 proofs yields (a two-sided version of) Shamkanov's cyclic system for GL. We follow the explanation given in Section 2.2.2 to turn G3K4 into a cyclic calculus for K4Grz.⁸ That is, we add a new, seemingly superfluous premise to the rule \Box_{K4} to distinguish between 'good' and 'bad' cycles. The new rule is:

$$\Box_{K4}^{\circ} \frac{\Gamma, \Pi, \Box\Pi \Rightarrow \varphi, \Delta \quad \Pi, \Box\Pi \Rightarrow \varphi}{\Gamma, \Box\Pi \Rightarrow \Box\varphi, \Delta}$$

The left premise of \Box_{K4}° is called *stationary*, and the right one *transitional*. As in the case of S4Grz, the new premise on the left is superfluous in G3K4 because it follows from the right one by an application of *wk*.

We let $K4^{\circ}$ be the cyclic system with base G3 + \Box_{K4}° + *wk* specified as follows.⁹

⁷Like Proposition 2.2.12 above, the fact that K4Grz is a cyclic companion of K4 does not follow from the work in [69] because the rule \Box_{K4}° , defined below, does not satisfy the constraints imposed in [69] on modal rules.

⁸An ordinary, acyclic calculus for K4Grz was first given in [7].

⁹An ill-founded calculus for K4Grz resembling the ill-founded version of our cyclic system was introduced by Savateev and Shamkanov in [129]. As in [130], they establish soundness and

$$\begin{array}{c}
\text{ax} \frac{}{B, \Box B, p \Rightarrow \Box p, p} \quad \text{wk} \frac{(\dagger) \quad \Box B \Rightarrow \Box p}{B, \Box B, p \Rightarrow \Box p} \\
\text{R} \rightarrow \frac{}{B, \Box B \Rightarrow C, p} \quad \text{R} \rightarrow \frac{}{B, \Box B \Rightarrow C} \\
\Box_{\text{K4}}^{\circ} \frac{}{\Box B \Rightarrow \Box C, p} \quad \text{ax} \frac{}{p, \Box B \Rightarrow p} \\
\text{L} \rightarrow \frac{}{B, \Box B \Rightarrow p} \\
\vdots \\
\vdots \quad \vdots \\
\Box_{\text{K4}}^{\circ} \frac{B, \Box B \Rightarrow p \quad B, \Box B \Rightarrow p}{(\dagger) \quad \Box B \Rightarrow \Box p} \\
\text{R} \rightarrow \frac{}{\Rightarrow \Box(\Box(p \rightarrow \Box p) \rightarrow p) \rightarrow \Box p}
\end{array}$$

FIGURE 2.7: A K4° proof of the K4Grz axiom $\text{grz}_1 = \Box(\Box(p \rightarrow \Box p) \rightarrow p) \rightarrow \Box p$, where $C := p \rightarrow \Box p$ and $B := \Box C \rightarrow p$. Repeats and their associated companions are marked by the symbol ‘ \dagger ’.

2.3.1. DEFINITION (K4° derivation). A K4° *derivation* of a formula φ is a finite tree with back-edges built according to the rules in Figure 2.1 plus \Box_{K4}° and wk , and whose root is labelled by $\emptyset \Rightarrow \varphi$. \dashv

A repeat l in a K4° derivation \mathcal{T} , say with companion c , is *successful* if $(c, l]_{\mathcal{T}}$ contains a vertex labelled by a transitional premise of an instance of \Box_{K4}° . A vertex $u \in T$ is *modal* if rule \Box_{K4}° is applied at u .

2.3.2. DEFINITION (K4° proof). A K4° *proof* of a formula φ is a K4° derivation \mathcal{T} of φ such that every non-axiomatic leaf of \mathcal{T} is a successful repeat. \dashv

For the importance (or, rather, lack thereof) of including the explicit weakening rule wk in K4° , see Remark 2.3.20 below.

The distinguished formulas in the conclusions (premises) of K4° rules are said to be *principal* (respectively, *active*). All formulas in $\Gamma \cup \Delta$ are *side* formulas.

2.3.3. EXAMPLE. Figure 2.7 depicts a K4° proof of the axiom K4Grz . From the two premises of the lower instance of \Box_{K4}° we proceed as indicated in the upper derivation. All the leaves are either axiomatic or successful repeats, the latter marked by the symbol ‘ \dagger ’.

Let us first see that $\mathcal{B}(\text{K4}^{\circ})$ is sound and complete for K4 :

completeness by means of proof transformations and a cut-elimination procedure that constitutes the focus of their article.

2.3.4. PROPOSITION. *For every formula φ , we have $\mathcal{B}(\mathbf{K4}^\circ) \vdash \varphi$ if, and only if, $\mathbf{K4} \vdash \varphi$.*

Proof. The left-to-right direction (soundness) is straightforward: axiomatic leaves of $\mathcal{B}(\mathbf{K4}^\circ)$ proofs are labelled by valid sequents, and all the rules of $\mathbf{K4}^\circ$ preserve validity with respect to transitive frames.

For completeness, assume that $\mathbf{K4} \vdash \varphi$. Then, there is a **G3K4** proof of φ . Every rule of **G3K4** is a rule of $\mathcal{B}(\mathbf{K4}^\circ)$ except for $\Box_{\mathbf{K4}}$, so it suffices to see that every instance of $\Box_{\mathbf{K4}}$ can be simulated in $\mathcal{B}(\mathbf{K4}^\circ)$:

$$\Box_{\mathbf{K4}} \frac{\Pi, \Box\Pi \Rightarrow \psi}{\Gamma, \Box\Pi \Rightarrow \Box\psi, \Delta} \quad \rightsquigarrow \quad \frac{\text{wk} \frac{\Pi, \Box\Pi \Rightarrow \psi}{\Gamma, \Pi, \Box\Pi \Rightarrow \psi, \Delta} \quad \Pi, \Box\Pi \Rightarrow \psi}{\Box_{\mathbf{K4}^\circ} \frac{\Gamma, \Box\Pi \Rightarrow \Box\psi, \Delta}{\Gamma, \Box\Pi \Rightarrow \Box\psi, \Delta}}$$

Therefore, every **G3K4** proof can be translated into a $\mathcal{B}(\mathbf{K4}^\circ)$ proof. \blacksquare

Now we prove soundness and completeness of $\mathbf{K4}^\circ$ for **K4Grz**.

The soundness proof for $\mathbf{K4}^\circ$ is analogous to the proofs of Propositions 2.2.8 and 2.2.14 above.

2.3.5. PROPOSITION. *For every formula φ , if $\mathbf{K4}^\circ \vdash \varphi$, then $\mathbf{K4Grz} \vdash \varphi$.*

Proof. Let \mathcal{T} be a $\mathbf{K4}^\circ$ proof of φ , and suppose towards a contradiction that $\mathbf{K4Grz} \not\vdash \varphi$. Then, there is a model $\mathcal{M} = (W, R, V)$ based on a transitive weakly Noetherian frame and a state $s \in W$ such that $\mathcal{M}, s \not\vdash \varphi$. We inductively build an infinite path $(u_n)_{n < \omega}$ through \mathcal{T}° and an infinite sequence $(s_n)_{n < \omega}$ of states in W such that the following hold for every $i < \omega$:

- (i) $\mathcal{M}, s_i \not\vdash \Gamma_i \Rightarrow \Delta_i$, where $\Gamma_i \Rightarrow \Delta_i$ is the label of u_i ;
- (ii) either u_i is not modal and $s_i = s_{i+1}$, or $s_i R s_{i+1}$.

The base case, as well as the inductive case where $\Box_{\mathbf{K4}^\circ}$ is not applied at u_n , are as in the proof of Proposition 2.2.8. Assume, then, that u_n is modal, and let

$$\Box_{\mathbf{K4}^\circ} \frac{\Gamma, \Pi, \Box\Pi \Rightarrow \psi, \Delta \quad \Pi, \Box\Pi \Rightarrow \psi}{\Gamma, \Box\Pi \Rightarrow \Box\psi, \Delta}$$

be the instance of $\Box_{\mathbf{K4}^\circ}$ at u_n . By the inductive hypothesis, $\mathcal{M}, s_n \not\vdash \Gamma, \Box\Pi \Rightarrow \Box\psi, \Delta$, whence there is a state $t \in R[s]$ such that $\mathcal{M}, t \Vdash \Pi$ and $\mathcal{M}, t \not\vdash \psi$. If $t = s_n$, we let u_{n+1} be the immediate successor of u_n for the premise $\Gamma, \Pi, \Box\Pi \Rightarrow \psi, \Delta$. Otherwise, transitivity yields $\mathcal{M}, t \not\vdash \Pi, \Box\Pi \Rightarrow \psi$ and we let u_{n+1} be the immediate successor of u_n for premise $\Pi, \Box\Pi \Rightarrow \psi$. In either case we let $s_{n+1} := t$.

The path $\pi := (u_n)_{n < \omega}$ must be infinite because axioms are labelled by valid sequents. By Proposition 1.2.3, there are infinitely many $i < \omega$ such that u_i

is labelled by the conclusion of an instance of $\Box_{\mathbf{K4}}^\circ$ and u_{i+1} is the immediate successor of u_i for the transitional premise. Hence, by construction there are infinitely many $i < \omega$ such that $s_i R s_{i+1} \neq s_i$ and thus the sequence $(s_n)_{n < \omega}$ contains a subsequence which is an infinite ascending R -chain and not eventually constant, contradicting the weak converse well-foundedness of (W, R) . ■

Completeness of $\mathbf{K4}^\circ$ is established by contraposition: assuming that φ is not provable, we build a countermodel for φ from a failed proof-search for φ .

To simplify some parts of the completeness proof, as well as shed light on the frame conditions that cycles can capture (see Remark 2.3.8 below), it is convenient to give a characterisation of the logic $\mathbf{K4Grz}$ which is slightly different than, but equivalent to, the one in Table 2.2. By Table 2.2, $\mathbf{K4Grz}$ is characterised by transitive, weakly Noetherian frames. Recall (see p. 34) that these are exactly the transitive frames which satisfy the SCC (call such frames *SCC*). Clearly, every SCC frame $\mathcal{F} = (W, R)$ satisfies the following *weak chain condition (WCC)*: for every infinite R -chain $(s_n)_{n < \omega}$ in \mathcal{F} there are $i < j < \omega$ such that $s_i = s_j$. Frames satisfying the WCC are said to be *WCC*. In general, the WCC does not imply the SCC, but in the presence of transitivity we have:

2.3.6. PROPOSITION (cf. [94]). *Let $\mathcal{F} = (W, R)$ be a transitive frame. Then, \mathcal{F} is SCC if, and only if, \mathcal{F} is antisymmetric and WCC.*

Proof. Assume that \mathcal{F} is SCC. Clearly, then, \mathcal{F} is WCC. And if $s, t \in W$ are such that $s R t$ and $t R s$, then (s, t, s, t, \dots) is an infinite R -chain and thus by the SCC we have $s = t$, so \mathcal{F} is antisymmetric.

Conversely, assume that \mathcal{F} is antisymmetric and WCC, and let $(s_n)_{n < \omega}$ be an infinite R -chain. By the WCC, there are $i < j$ such that $s_i = s_j$ and $s_i \neq s_k$ for every $i < k < j$. Towards a contradiction, suppose that $j > i + 1$. Then, by transitivity we have $s_i R s_{j-1} R s_j = s_i$ and thus $s_i = s_{j-1}$ by antisymmetry, contradicting the choice of i and j . Hence, $j = i + 1$ and \mathcal{F} is SCC. ■

2.3.7. COROLLARY. *The logic $\mathbf{K4Grz}$ is characterised by the class of all transitive, antisymmetric and WCC frames.*

2.3.8. REMARK. The characterisation of $\mathbf{K4Grz}$ given in Corollary 2.3.7 is not suitable for finding (cyclic or acyclic) proof systems for the logic. Informally, this is due to the fact that both antisymmetry and the WCC require keeping track of previously visited states (the last visited state for antisymmetry; any prior state for the WCC). More formally, it is well known that antisymmetry is not modally definable (see, e.g., [13, § 4.5]), and a standard argument based on bisimulations and tree-unravellings of frames shows that being WCC is not modally definable

either. One should not expect, therefore, to find proof rules capturing these two properties.

Moreover, since the distinction between ‘good’ and ‘bad’ cycles in a cyclic calculus is ultimately made at the level of the rules of the system (in $K4^\circ$, for example, at the modal rule \Box_{K4}°), we should not expect to be able to capture those conditions by means of cycles either, or at least not by following the approach presented in this chapter. We return to this question in Section 2.4 below.

A formula ψ is *left-duplicated* (*right-duplicated*) in a sequent $\Gamma \Rightarrow \Delta$ if ψ occurs at least twice in the multiset Γ (respectively, Δ). A formula ψ is *duplicated* in a sequent S if ψ is left- or right-duplicated in S . For a multiset Γ , we let $[\Gamma]$ denote the *set* $[\Gamma] := \{\gamma \mid \gamma \in \Gamma\}$.

2.3.9. DEFINITION ($K4^\circ$ proof-search tree). A $K4^\circ$ *proof-search tree* for a formula φ is a finite tree with back-edges \mathcal{T} built according to the rules of $K4^\circ$ and such that:

- (i) The root of \mathcal{T} has label $\emptyset \Rightarrow \varphi$.
- (ii) For every vertex $u \in T$, if there is a duplicated formula in the label of u , then rule wk is applied at u with conclusion $\Gamma \Rightarrow \Delta$ and premise $[\Gamma] \Rightarrow [\Delta]$, where $\Gamma \Rightarrow \Delta$ is the label of u .
- (iii) There are no other instances of wk in \mathcal{T} .
- (iv) A vertex $u \in T$ is a leaf if, and only if, there are no duplicated formulas in the label of u and u is either axiomatic, or a repeat, or a cul-de-sac.
- (v) The following branching version of rule \Box_{K4}° is used in \mathcal{T} in place of \Box_{K4}° :

$$\Box_{K4}^* \frac{\{\Gamma, \Pi, \Box \Pi \Rightarrow \varphi_i, \Box \Phi_i, \Delta \mid i \leq n\} \quad \{\Pi, \Box \Pi \Rightarrow \varphi_i \mid i \leq n\}}{\Gamma, \Box \Pi \Rightarrow \Box \varphi_0, \dots, \Box \varphi_n, \Delta},$$

where $\Phi_i := \{\varphi_0, \dots, \varphi_{i-1}, \varphi_{i+1}, \dots, \varphi_n\}$,¹⁰ $\Gamma \cup \Delta$ contains only literals, $\perp \notin \Gamma$, and the conclusion contains no duplicated formula. –

Let \mathcal{T} be a proof-search tree for a formula φ , and $u \in T$ a vertex such that rule \Box_{K4}^* is applied at u (call such vertices *modal*). Let

$$\Box_{K4}^* \frac{\{\Gamma, \Pi, \Box \Pi \Rightarrow \varphi_i, \Box \Phi_i, \Delta \mid i \leq n\} \quad \{\Pi, \Box \Pi \Rightarrow \varphi_i \mid i \leq n\}}{\Gamma, \Box \Pi \Rightarrow \Box \varphi_0, \dots, \Box \varphi_n, \Delta}$$

¹⁰Each Φ_i is a multiset.

be the instance of $\Box_{\mathbf{K4}}^*$ at u . For every $i \leq n$, let v_i^s be the immediate successor of u for the *stationary* premise $\Gamma, \Pi, \Box\Pi \Rightarrow \varphi_i, \Box\Phi_i, \Delta$, and let v_i^t be the immediate successor of u for the *transitional* premise $\Pi, \Box\Pi \Rightarrow \varphi_i$. We call each (v_i^s, v_i^t) a *modal pair* of u .

The notion of successful repeat is defined for $\mathbf{K4}^\circ$ proof-search trees as expected: a repeat l of a $\mathbf{K4}^\circ$ proof-search tree \mathcal{T} , say with companion c , is *successful* if there is a vertex in $(c, l]_{\mathcal{T}}$ labelled by a transitional premise of an instance of $\Box_{\mathbf{K4}}^*$.

Let \mathcal{T} be a $\mathbf{K4}^\circ$ proof-search tree. A leaf l of \mathcal{T} is *good* if l is either axiomatic or a successful repeat of \mathcal{T} . Otherwise, l is said to be *bad*. So l is bad if l is a cul-de-sac or an unsuccessful repeat.

Every formula occurring in a $\mathbf{K4}^\circ$ proof-search tree for a formula φ is contained in $\text{Sub}(\varphi)$. Hence, conditions (ii) and (iv) in Definition 2.3.9 immediately yield:

2.3.10. PROPOSITION. *For every formula φ , there exists a $\mathbf{K4}^\circ$ proof-search tree for φ .*

Additionally, since rule $\Box_{\mathbf{K4}}^*$ is the only one which need not decrease the complexity of the sequent to which it is applied, we clearly have:

2.3.11. PROPOSITION. *Let \mathcal{T} be a $\mathbf{K4}^\circ$ proof-search tree. For any repeat $l \in \text{Rep}_{\mathcal{T}}$, say with companion c , there is a premise of rule $\Box_{\mathbf{K4}}^*$ in $(c, l]_{\mathcal{T}}$.*

Our goal is to show that every $\mathbf{K4}^\circ$ proof-search tree for φ contains either a $\mathbf{K4}^\circ$ proof of φ , or else a *refutation* of φ , a subtree from which a countermodel for φ can be built.

2.3.12. DEFINITION ($\mathbf{K4}^\circ$ refutation). A *$\mathbf{K4}^\circ$ refutation* of a formula φ is a subtree \mathcal{T}' of a $\mathbf{K4}^\circ$ proof-search tree \mathcal{T} for φ satisfying:

- (i) The root of \mathcal{T}' is the root of \mathcal{T} .
- (ii) Every leaf of \mathcal{T}' is a bad leaf of \mathcal{T} .
- (iii) If a vertex $u \in \mathcal{T}'$ is labelled in \mathcal{T} by the conclusion of an instance of a non-axiomatic rule other than $\Box_{\mathbf{K4}}^*$, then u has exactly one immediate successor in \mathcal{T}' .
- (iv) If a vertex $u \in \mathcal{T}'$ is modal in \mathcal{T} , say with modal pairs $(v_0^s, v_0^t), \dots, (v_n^s, v_n^t)$, then for every $i \leq n$ there is a $v_i \in \{v_i^s, v_i^t\}$ such that the immediate successors of u in \mathcal{T}' are exactly v_0, \dots, v_n .

—

Towards establishing the completeness of $\mathbf{K4}^\circ$, let us first see that every $\mathbf{K4}^\circ$ proof-search tree \mathcal{T} which does not contain a proof can be pruned down to a refutation. Informally, we do this by pruning \mathcal{T} bottom-up as follows: at each

non-modal branching vertex, we keep a premise that fails to produce a proof and discard the other one; similarly, at each modal vertex, for each modal pair we keep a premise that fails to produce a proof and discard the other one. The resulting subtree is then easily shown to be a refutation.

Let \mathcal{T} be a K4° proof-search tree, and $\mathcal{T}' = [u, \rightarrow]_{\mathcal{T}}$ a cone of \mathcal{T} . We denote by $M(\mathcal{T}')$ the collection of all modal vertices of \mathcal{T} contained in \mathcal{T}' . A *modal pruning map* on \mathcal{T}' is a function $f : M(\mathcal{T}') \rightarrow T' \times T'$ such that $f(u)$ is a modal pair of u for each vertex $u \in M(\mathcal{T}')$. The *modal pruning of \mathcal{T}' with respect to f* is the subtree \mathcal{T}'' of \mathcal{T}' given by letting $v \in \mathcal{T}''$ if, and only if, the following holds: for every $w \in M(\mathcal{T}') \cap [u, v]_{\mathcal{T}'}$, the unique immediate successor of w in $[u, v]_{\mathcal{T}'}$ belongs to the modal pair $f(w)$. A *modal pruning of \mathcal{T}'* is the modal pruning of \mathcal{T}' with respect to some unspecified modal pruning map on \mathcal{T}' .

2.3.13. PROPOSITION. *Let \mathcal{T} be a K4° proof-search tree for a formula φ . If no subtree of \mathcal{T} is a K4° proof of φ , then there is a subtree of \mathcal{T} which is a K4° refutation of φ .*

Proof. We inductively build an infinite sequence $((\mathcal{T}_n, P_n))_{n < \omega}$ such that the following hold for every $n < \omega$:

- (i) $\mathcal{T}_0 = \mathcal{T}$;
- (ii) \mathcal{T}_n is a subtree of \mathcal{T}_{n+1} ;
- (iii) $P_n \subseteq P_{n+1}$;
- (iv) if $T_n \setminus P_n \neq \emptyset$, then $|P_{n+1}| = |P_n| + 1$;
- (v) if $T_n \setminus P_n = \emptyset$, then $\mathcal{T}_{n+1} = \mathcal{T}_n$ and $P_{n+1} = P_n$;
- (vi) for every $u \in T_n \setminus P_n$, we have $[u, \rightarrow]_{\mathcal{T}} \subseteq T_n \setminus P_n$;
- (vii) for every $u \in T_n \setminus P_n$ of minimal height in $T_n \setminus P_n$, every modal pruning of $[u, \rightarrow]_{\mathcal{T}_n}$ contains a bad leaf of \mathcal{T} ;
- (viii) for every leaf l of \mathcal{T} , if $l \in P_n$, then l is a bad leaf of \mathcal{T} .

For the base case, let $\mathcal{T}_0 := \mathcal{T}$ and $P_0 := \emptyset$. Note that (vii) holds because \mathcal{T} contains no subtree which is a K4° proof of φ .

For the inductive case, assume that \mathcal{T}_n and P_n have been defined. If $T_n \setminus P_n = \emptyset$, we let $\mathcal{T}_{n+1} := \mathcal{T}_n$ and $P_{n+1} := P_n$. Suppose now that $T_n \setminus P_n \neq \emptyset$, and let $u \in T_n \setminus P_n$ be of minimal height in $T_n \setminus P_n$. We first let $P_{n+1} := P_n \cup \{u\}$. Note that (viii) follows from (vii) and (viii) of the inductive hypothesis. Let us now build \mathcal{T}_{n+1} .

If u is a non-branching vertex of \mathcal{T}_n , we let $\mathcal{T}_{n+1} := \mathcal{T}_n$. If u is branching, we distinguish cases according to the rule $R \in \{\text{L}\rightarrow, \square_{\text{K4}}^*\}$ applied at u .

Suppose first that $R = L \rightarrow$. Let v_0, v_1 be the two immediate successors of u in \mathcal{T}_n . By (vii) of the inductive hypothesis, there is an $i \leq 1$ such that every modal pruning of $[v_i, \rightarrow)_{\mathcal{T}_n}$ contains a bad leaf of \mathcal{T} . We then let \mathcal{T}_{n+1} be the subtree of \mathcal{T}_n given by setting $T_{n+1} := T_n \setminus [v_{1-i}, \rightarrow)_{\mathcal{T}_n}$. Note that $P_{n+1} \subseteq T_{n+1}$ by (iii) and (vi) of the inductive hypothesis.

Suppose now that $R = \Box_{\mathcal{K}4}^*$. Let $(v_0^s, v_0^t), \dots, (v_n^s, v_n^t)$ be the modal pairs of u . By (vii) of the inductive hypothesis, for every $i \leq n$ there is a $v_i \in \{v_i^s, v_i^t\}$ such that every modal pruning of $[v_i, \rightarrow)_{\mathcal{T}_n}$ contains a bad leaf of \mathcal{T} . We then let \mathcal{T}_{n+1} be the subtree of \mathcal{T}_n given by setting

$$T_{n+1} := T_n \setminus \bigcup \{ [v, \rightarrow)_{\mathcal{T}_n} \mid u <_T^0 v \text{ and } v \notin \{v_0, \dots, v_n\} \}.$$

As before, $P_{n+1} \subseteq T_{n+1}$ by (iii) and (vi) of the inductive hypothesis.

Since \mathcal{T} contains only finitely many vertices, by (iv) and (v) there is an $n < \omega$ such that $(\mathcal{T}_{n+k}, P_{n+k}) = (\mathcal{T}_n, P_n)$ for every $k < \omega$. We claim that \mathcal{T}_n is a $\mathbf{K}4^\circ$ refutation of φ . Conditions (i), (iii) and (iv) of Definition 2.3.12 are clear by the construction of \mathcal{T}_n . And that every leaf of \mathcal{T}_n is a bad leaf of \mathcal{T} is an immediate consequence of (viii). ■

2.3.14. COROLLARY. *Let \mathcal{T} be a $\mathbf{K}4^\circ$ proof-search tree for a formula φ . If $\mathbf{K}4^\circ \not\vdash \varphi$, then there is a subtree of \mathcal{T} which is a $\mathbf{K}4^\circ$ refutation of φ .*

As an immediate consequence of Propositions 2.3.5 and 2.3.10 and Corollary 2.3.14, we have:

2.3.15. COROLLARY. *For every formula φ , if $\mathbf{K}4\text{Grz} \not\vdash \varphi$, then there is a $\mathbf{K}4^\circ$ refutation of φ .*

We now establish the converse of Corollary 2.3.15: the existence of a $\mathbf{K}4^\circ$ refutation of φ implies that $\mathbf{K}4\text{Grz} \not\vdash \varphi$. We do so by building a countermodel of φ from the refutation.

Let \mathcal{T} be a $\mathbf{K}4^\circ$ refutation. A path $\pi = (u_n)_{n < N}$ through \mathcal{T}° , with $N \leq \omega$, is *non-modal* if the following holds for every $n < N$: if $\pi(n)$ is a modal vertex, then $N = n + 1$. By Propositions 1.2.3 and 2.3.11, every non-modal path through \mathcal{T}° is finite.

2.3.16. PROPOSITION. *For every formula φ , if there is a $\mathbf{K}4^\circ$ refutation of φ , then $\mathbf{K}4\text{Grz} \not\vdash \varphi$.*

Proof. Let \mathcal{T} be a $\mathbf{K}4^\circ$ refutation of φ . We construct a model $\mathcal{M}_{\mathcal{T}} = (W_{\mathcal{T}}, R_{\mathcal{T}}, V_{\mathcal{T}})$ from \mathcal{T} as follows, and then show it to be a countermodel of φ .

First, let W be the collection of all (finite) non-modal paths π through \mathcal{T}° such that there is no non-modal path $\pi' \neq \pi$ satisfying $\pi' = \pi_0 \hat{\wedge} \pi \hat{\wedge} \pi_1$ for some paths π_0, π_1 on \mathcal{T}° . That is to say: W contains all non-modal, \sqsubset -maximal paths through \mathcal{T}° starting from the root of \mathcal{T} or from an immediate successor of a modal vertex. By Propositions 1.2.3 and 2.3.11, every path in W is finite.

Then, let R^* be the transitive closure of the binary relation R on W defined by setting $s R t$ if, and only if, the last vertex u of s is modal and the first vertex of t is an immediate successor of u in \mathcal{T}° .

Finally, let \sim be the binary relation on W given by setting $s \sim t$ if, and only if, either $s = t$ or both $s R^* t$ and $t R^* s$. It is straightforward to check that \sim is an equivalence relation. For every $s \in W$, let $[s]$ be the \sim -equivalence class of s .

We let $W_{\mathcal{T}} := W/\sim$ and $[s] R_{\mathcal{T}} [t]$ if, and only if, $s R^* t$. It is easy to see that $R_{\mathcal{T}}$ is well-defined. And, clearly, $R_{\mathcal{T}}$ is transitive and antisymmetric. Moreover, $(W_{\mathcal{T}}, R_{\mathcal{T}})$ satisfies the WCC because W is finite. It remains to define the valuation $V_{\mathcal{T}}$. For every $s \in W$ and $p \in \mathbf{Prop}$, we let $[s] \in V_{\mathcal{T}}(p)$ if, and only if, p occurs in the antecedent of the sequent labelling the last vertex of s .

2.3.16.1. CLAIM. *The valuation $V_{\mathcal{T}}$ is well-defined.*

Proof of claim. It suffices to see that, if $[s] \in V_{\mathcal{T}}(p)$ and $[s] = [t]$, then $[t] \in V_{\mathcal{T}}(p)$. Assume, then, that $s \sim t$. If $s = t$ there is nothing to prove, so suppose that $s \neq t$. Then, $s R^* t R^* s$ and thus there are $s_0, \dots, s_m, t_0, \dots, t_{m'} \in W$, with $m, m' > 0$, such that

$$s = s_0 R s_1 R \cdots R s_m = t = t_0 R t_1 R \cdots R t_{m'} = s. \quad (2.2)$$

Let u^s and u^t be the last vertices in s and t , respectively. By (2.2), there is a finite path $\pi = (u_0, \dots, u_n)$ through \mathcal{T}° , with $n > 0$, such that $u_0 = u_n = u^s$ and $u_j = u^t$ for some $j \leq n$.

Since $[s] \in V_{\mathcal{T}}(p)$, the propositional letter p occurs in the antecedent of the sequent labelling u^s . The only rule applicable in \mathcal{T} which, when read bottom-up, might completely remove a propositional letter from the antecedent is \Box_{K4}^* , and only when going from the conclusion to a transitional premise. To establish that $[t] \in V_{\mathcal{T}}(p)$, then, it suffices to show that there is no $i < n - 1$ such that $\pi(i)$ is modal and $\pi(i + 1)$ a transitional premise of $\pi(i)$.

Towards a contradiction, suppose that there exists such an $i < n - 1$. Let $v := \pi(i)$ and $v^+ := \pi(i + 1)$. Also, let $\pi_\omega := \pi_{<n} \hat{\wedge} \pi_{<n} \hat{\wedge} \cdots$. Note that π_ω is an infinite path through \mathcal{T}° because $u_n = u_0$. Since $\pi_\omega(k) = v$ for infinitely many $k < \omega$, there is a repeat $l \in \mathbf{Rep}_{\mathcal{T}} \cap [v^+, \rightarrow)_{\mathcal{T}}$ such that $c_l \leq_{\mathcal{T}} v$. Hence, $v^+ \in (c_l, l]_{\mathcal{T}}$ and thus l is a successful repeat, contradicting Definition 2.3.12(ii). \square

We now show the following:

for every $s \in W$ and every sequent $\Lambda \Rightarrow \Xi$ labelling a vertex in s , (*)
we have $\mathcal{M}_{\mathcal{T}}, [s] \not\vdash \Lambda \Rightarrow \Xi$.

We do this in two stages: first, for every $s \in W$ such that $[s]$ has no $R_{\mathcal{T}}$ -successor (Claim 2.3.16.2); then, inductively, for every other $s \in W$.

2.3.16.2. CLAIM. *If $s \in W$ is such that $([s], [t]) \notin R_{\mathcal{T}}$ for any $t \in W$, then $\mathcal{M}_{\mathcal{T}}, [s] \not\vdash \Lambda \Rightarrow \Xi$ for every sequent $\Lambda \Rightarrow \Xi$ labelling a vertex in s .*

Proof of claim. Let $s = (u_0, \dots, u_n)$. Since $[s]$ has no $R_{\mathcal{T}}$ -successor, u_n is a non-repeat leaf of \mathcal{T} . Fix any $i \leq n$, and let $\Lambda \Rightarrow \Xi$ be the sequent labelling u_i . Towards a contradiction, suppose that $\mathcal{M}_{\mathcal{T}}, [s] \vdash \Lambda \Rightarrow \Xi$. Since $\Box_{\mathbb{K}_4}^*$ is not applied anywhere in s and u_n is a cul-de-sac, the only rules applied in s are $L \rightarrow$, $R \rightarrow$ and \mathbf{wk} , and the last one only to remove duplicated formulas. It is immediate to see that, in any case, if a world of a model satisfies the conclusion of the rule then it also satisfies any of its premises. Therefore, we have $\mathcal{M}_{\mathcal{T}}, [s] \vdash \Pi \Rightarrow \Sigma$, where $\Pi \Rightarrow \Sigma$ is the label of u_n . Given that u_n is a cul-de-sac, $\Pi \cup \Sigma$ contains only literals and we have $\perp \notin \Pi$ and $\Pi \cap \Sigma = \emptyset$. By the definition of $V_{\mathcal{T}}$, then, $\mathcal{M}_{\mathcal{T}}, [s] \not\vdash \Pi \Rightarrow \Sigma$, contradiction. \square

Finally, by structural induction on ψ we show that for every formula ψ , every $s \in W$ and every sequent $\Lambda \Rightarrow \Xi$ labelling a vertex in s , the following hold:

- (i) if $\psi \in \Lambda$, then $\mathcal{M}_{\mathcal{T}}, [s] \vdash \psi$;
- (ii) if $\psi \in \Xi$, then $\mathcal{M}_{\mathcal{T}}, [s] \not\vdash \psi$.

By Claim 2.3.16.2, we may assume that $[s]$ has at least one $R_{\mathcal{T}}$ -successor. Then, the last vertex v of s is modal.

Case $\psi = \perp$. Since no vertex in \mathcal{T} is axiomatic, we can only have $\perp \in \Xi$ and are done.

Case $\psi = p$. Suppose first that $p \in \Lambda$. The only rule applicable in \mathcal{T} which might completely remove a propositional letter when read bottom-up is $\Box_{\mathbb{K}_4}^*$, and only when going from the conclusion to a transitional premise. Hence, p occurs in the antecedent of the sequent labelling v and we thus have $\mathcal{M}_{\mathcal{T}}, [s] \vdash p$ by the definition of $V_{\mathcal{T}}$.

Suppose now that $p \in \Xi$. Arguing as before, we get that p occurs in the consequent of the sequent labelling v . Since no vertex in \mathcal{T} is axiomatic, then, p does not occur in the antecedent of the sequent labelling v , whence $\mathcal{M}_{\mathcal{T}}, [s] \not\vdash p$ by the definition of $V_{\mathcal{T}}$.

Case $\psi = \psi_0 \rightarrow \psi_1$. Suppose first that $\psi \in \Lambda$. By construction, rule $\Box_{\mathbf{K4}}^*$ is only applied in \mathcal{T} when no other rule may be applied, so there is an instance of $\mathbf{L}\rightarrow$ in s with principal formula $\psi_0 \rightarrow \psi_1$. Hence, there is a vertex $w \in s$, say with label $\Pi \Rightarrow \Sigma$, such that either $\psi_0 \in \Sigma$ or $\psi_1 \in \Pi$. In either case, by the inductive hypothesis and 2.3.16.2 we have $\mathcal{M}_{\mathcal{T}}, [s] \Vdash \psi_0 \rightarrow \psi_1$.

Suppose now that $\psi \in \Xi$. Reasoning as before, with $\mathbf{R}\rightarrow$ in place of $\mathbf{L}\rightarrow$, we conclude that there is a vertex $w \in s$, say with label $\Pi \Rightarrow \Sigma$, such that $\psi_0 \in \Pi$ and $\psi_1 \in \Sigma$, whence $\mathcal{M}_{\mathcal{T}}, [s] \not\Vdash \psi_0 \rightarrow \psi_1$ by the inductive hypothesis and Claim 2.3.16.2.

Case $\psi = \Box\psi'$. By construction, ψ remains non-principal until v . Suppose first that $\psi \in \Xi$. Then, there is an immediate successor w of v in \mathcal{T} such that ψ' is active in the consequent of the label of w . Let t be the (unique) path in W starting from w . Then, $s \mathbf{R} t$, so $[s] \mathbf{R}_{\mathcal{T}} [t]$. And $\mathcal{M}_{\mathcal{T}}, [t] \not\Vdash \psi'$ by the inductive hypothesis and Claim 2.3.16.2, whence $\mathcal{M}_{\mathcal{T}}, [s] \not\Vdash \Box\psi'$.

Finally, suppose that $\psi \in \Lambda$. Let $t \in W$ be such that $[s] \mathbf{R}_{\mathcal{T}} [t]$, i.e., $s \mathbf{R}^* t$. So there are $s_0, \dots, s_n \in W$, $n > 0$, such that

$$s = s_0 \mathbf{R} s_1 \mathbf{R} \dots \mathbf{R} s_n = t.$$

By a secondary induction on $1 \leq k \leq n$, we show that both ψ' and $\Box\psi'$ occur in the antecedent of the sequent labelling the first vertex of s_k . By the principal inductive hypothesis, Claim 2.3.16.2, and the arbitrariness of t , this suffices to establish that $\mathcal{M}_{\mathcal{T}}, [s] \Vdash \Box\psi'$.

In the base case, $k = 1$ and $s \mathbf{R} s_1$. Since $\Box\psi'$ occurs in the antecedent of the sequent labelling v , both ψ' and $\Box\psi'$ occur in the antecedent of the first vertex of s_1 . In the inductive case, by the secondary inductive hypothesis both ψ' and $\Box\psi'$ occur in the antecedent of the sequent labelling the first vertex of s_k , where k is such that $k + 1 \leq n$. No rule applied in \mathcal{T} completely removes modal formulas from the antecedent when read bottom-up, so $\Box\psi'$ occurs in the antecedent of the sequent labelling the last vertex of s_k . Thus, both ψ' and $\Box\psi'$ occur in the antecedent of the sequent labelling the first vertex of s_{k+1} and we are done.

This concludes the proof of (*). In particular, $\mathcal{M}_{\mathcal{T}}, [r] \not\Vdash \varphi$ for any $r \in W$ starting from the root of \mathcal{T} . By Corollary 2.3.7, then, $\mathbf{K4Grz} \not\vdash \varphi$. ■

Completeness of $\mathbf{K4}^\circ$ for $\mathbf{K4Grz}$ now follows immediately:

2.3.17. COROLLARY. *For every formula φ , if $\mathbf{K4Grz} \vdash \varphi$, then $\mathbf{K4}^\circ \vdash \varphi$.*

Proof. By contraposition. Suppose that $\mathbf{K4}^\circ \not\vdash \varphi$, and let \mathcal{T} be a $\mathbf{K4}^\circ$ proof-search tree for φ (at least one exists by Proposition 2.3.10). By Proposition 2.3.13, there is a subtree of \mathcal{T} which is a $\mathbf{K4}^\circ$ refutation of φ , whence $\mathbf{K4Grz} \not\vdash \varphi$ by Proposition 2.3.16. ■

In addition, we can now fully justify the use of the term ‘refutation’:

2.3.18. COROLLARY. *For every formula φ , there is a $\mathsf{K4}^\circ$ refutation of φ if, and only if, $\mathsf{K4Grz} \not\vdash \varphi$.*

Proof. The left-to-right implication is Proposition 2.3.16, and the converse follows immediately from Propositions 2.3.5, 2.3.10 and 2.3.13. ■

Bringing Proposition 2.3.5 and Corollary 2.3.17 together yields:

2.3.19. COROLLARY (Soundness and completeness of $\mathsf{K4}^\circ$). *For every formula φ , we have $\mathsf{K4}^\circ \vdash \varphi$ if, and only if, $\mathsf{K4Grz} \vdash \varphi$.*

2.3.20. REMARK. The completeness proofs for $\mathsf{K4}^\circ$ and $\mathcal{B}(\mathsf{K4}^\circ)$ rely on the explicit weakening rule wk . Clearly, wk can be removed from $\mathcal{B}(\mathsf{K4}^\circ)$ and the system remains sound and complete for $\mathsf{K4}$.

By considering the ill-founded version of $\mathsf{K4}^\circ$, in which proofs are possibly infinite trees whose infinite branches encounter infinitely many transitional premises of rule \Box_{K4}° , it is straightforward to see that wk is also superfluous in $\mathsf{K4}^\circ$: every $\mathsf{K4}^\circ$ proof may be unravelled into an ill-founded proof from which instances of wk are easily removable by relying on the weakening implicit in the axiomatic rules and the transitional premise of \Box_{K4}° ; the resulting weakening-free ill-founded proof can readily be folded back down to a weakening-free $\mathsf{K4}^\circ$ proof.

Alternatively, one may leave weakening out of $\mathsf{K4}^\circ$ proof-search by ignoring duplicated formulas and only removing them when passing through a transitional premise of \Box_{K4}° .

2.3.1 The finite model property and decidability

Observe that from Corollary 2.3.15 and the proof of Proposition 2.3.16 we obtain the *finite model property (FMP)* for $\mathsf{K4Grz}$: for every formula φ , if $\mathsf{K4Grz} \not\vdash \varphi$, then there is a finite model \mathcal{M} based on a transitive and weakly Noetherian frame such that $\mathcal{M} \not\models \varphi$.

Moreover, since the size of any $\mathsf{K4}^\circ$ proof-search tree for φ can clearly be bound by a computable function of $|\varphi| := |\text{Sub}(\varphi)|$, it follows that $\mathsf{K4Grz}$ has the *effective FMP* (and is therefore decidable¹¹): there exists a computable function $f : \omega \rightarrow \omega$ such that, for every formula φ , if $\mathsf{K4Grz} \not\vdash \varphi$, then there is a model $\mathcal{M} = (W, R, V)$ based on a transitive and weakly Noetherian frame such that $|W| \leq f(|\varphi|)$ and $\mathcal{M} \not\models \varphi$. Decidability (by purely syntactical means) is also a direct consequence of Propositions 2.3.5, 2.3.10, 2.3.13 and 2.3.16.

¹¹A logic L is *decidable* if there is an algorithm which, for every formula φ , decides in finitely many steps whether $L \vdash \varphi$ or $L \not\vdash \varphi$.

These results were first obtained in [7] by a tableau-based argument. We do not elaborate on this topic because questions about computational complexity fall outside the scope of this dissertation and, besides, our approach does not yield better complexity bounds than the work in [7].

2.4 Conclusion

Adding cycles to ordinary sequent calculi for **K4** and **S4** yields cyclic proof systems for **GL** and **S4Grz**, respectively [134, 130]. Rather than isolated contrivances, we argued that these systems arise from a natural correspondence between cycles in proofs and infinite chains in frames that enables the former to capture frame conditions involving the latter.

Our explanation of *cyclic companionship* suggested that we should be able to obtain a cyclic system for **K4Grz**, which shares its first-order frame condition with **GL** and its second-order one with **S4Grz**, by adding cycles to a system for **K4**. We showed that this is indeed the case. There is nothing special about **K4Grz**, and our approach should also apply to other logics characterised in the style of **GL**, **S4Grz** and **K4Grz**. Natural candidates are **S4.2Grz** and **S4.3Grz**.

The logic $\mathbf{S4.2} := \mathbf{S4} \oplus \diamond\Box p \rightarrow \Box\diamond p$ was introduced in [39] as part of an investigation of modal logics between **S4** and **S5**. It has since found applications in epistemology, where it has been argued to be an adequate logic of knowledge [86, 155, 141]; in temporal logic, where it has been shown to be the Diodorean¹² logic of Minkowski spacetime [59]; and, more recently, in the study of forcing from the point of view of modal logic, where it has been shown to be the logic of forcing in ZFC [66]. We refer the reader to [25] for a survey of these and other results about **S4.2**.

It is well known that **S4.2** is characterised by the class of reflexive and transitive frames (W, R) which are in addition *convergent*: if $s R s_0$ and $s R s_1$, then there is a $t \in W$ such that $s_0 R t$ and $s_1 R t$ (see, e.g., [68, Ch. 7]). Concordantly, the logic $\mathbf{S4.2Grz} := \mathbf{S4.2} \oplus \mathbf{grz}$ is characterised by reflexive, transitive, convergent and weakly Noetherian frames [48, § 2.2.1].¹³

A sound and complete calculus for **S4.2** was first given in [62]. It extends a system for **S4** by adding an analytic cut rule and a modal rule for convergence. Reasoning as we did in Sections 2.2.2 and 2.3, we can add cycles to the system

¹²The *Diodorean* (temporal) reading of $\Box p$, named after the ancient Greek logician Diodorus Cronus, is: ‘ p is the case now and will always remain the case’. To the best of our knowledge, the usage of the term in modern logic originates, under a slightly different spelling, in Prior’s [110].

¹³In the literature, **S4.2Grz** is also known as **Grz.2** (see, e.g., [48]). Contrary to **Grz.1**, **Grz.2** is an extension of **Grz**. We use the alternative name **S4.2Grz** simply for consistency.

from [62] and obtain as a result a cyclic proof system $S4.2^\circ$ for $S4.2Grz$.¹⁴ The cut rule and the additional modal rule make proof-search for $S4.2^\circ$ more involved than for $K4^\circ$, but the difficulties that arise are easy to overcome and, since they do not stem from the addition of cycles, for all intents and purposes they were already encountered and solved in [62]. To ensure convergence when building the counter-model corresponding to a given refutation we can argue as in the completeness proof in [62, § 4.3]. We leave the details to the interested reader.

Another sound and complete sequent calculus for $S4.2$ is given in [149]. Like the one in [62], the system extends a standard calculus for $S4$ with a new modal rule and an analytic cut rule. We conjecture that it is also possible to add cycles to this system to obtain a cyclic calculus for $S4.2Grz$.

The logic $S4.3 := S4 \oplus \Box(\Box p \rightarrow q) \vee \Box(\Box q \rightarrow p)$ was, like $S4.2$, introduced in [39]. It turned out to be the Diodorean logic of linear, continuous time [22, 131], and of course it features in Bull's celebrated result that all normal extensions of $S4.3$ have the FMP [23].

From the point of view of relational semantics, $S4.3$ is characterised by the class of reflexive and transitive frames (W, R) which are moreover (*strongly*) *connected*: if $s R s'$ and $s R s''$, then either $s' R s''$ or $s'' R s'$ (see, e.g., [68, Ch. 7]). Several other characterisations were given in [131], for example: the class of linear orders; the single frame (\mathbb{Q}, \leq) ; the single frame (\mathbb{R}, \leq) . The logic $S4.3Grz := S4.3 \oplus grz$ is characterised both by the class of finite linear orders and by the single frame (ω, \geq) [132].¹⁵

A sound and complete cut-free calculus for $S4.3$ was introduced in [63, § 6.1]. We conjecture that the addition of cycles to the system in [63] as done for $S4Grz$ and $K4Grz$ would yield a cyclic calculus for $S4.3Grz$. Indeed, working in the frame (ω, \geq) and arguing as in the proofs of Propositions 2.2.8, 2.2.14 and 2.3.5, we should be able to build an infinite descending chain of natural numbers and thus establish the soundness of the system. Proof-search in this setting should not be more difficult than for $K4^\circ$, whence completeness would follow.

The frame conditions that we can capture by adding cycles to acyclic systems seem to be limited to variations of converse well-foundedness. Given how close the correspondence between cycles and infinite chains is, this is not surprising. But, as explained in Remark 2.3.8 above, there are common, simple frame conditions about infinite chains, such as the WCC, which we expect to be out of reach of the cyclic approach.

It is well known that converse well-foundedness is not a modally definable property [11, p. 52]. And it is easy to see that the same is true of its weak

¹⁴A cut-free, ordinary sequent calculus for $S4.2Grz$ was first obtained in [113], where $S4.2Grz$ is referred to as $Gr.2$.

¹⁵In the literature, $S4.3Grz$ is also known as $Grz.3$ (see, e.g., [63, 24, 48]).

counterpart: the single, reflexive frame is weakly conversely well-founded, but its tree-unravelling is not. The same example shows the WCC not to be modally definable either. The first two properties, as we have seen, can be characterised by means of cycles. The latter, for the reasons given in Remark 2.3.8, probably cannot.

A natural continuation of our work in this chapter, then, would be to investigate the relation between modal definability of frame conditions about infinite chains and their characterisability by cyclic means.

Temporal Logics

Chapter 3

A Cyclic Proof System for CTL*

We concluded the previous chapter with references to two propositional modal logics, S4.2 and S4.3, having interesting temporal interpretations. Logical considerations about the nature of time trace back at least to Aristotle's analysis of the problem of future contingents, and were common in the works of medieval logicians (see, e.g., [53, 104]). In the 1950s and 1960s, Prior's invention of Tense Logic [111] laid the foundation for modern temporal logic. He was mostly driven by philosophical and theological questions, to which he applied the tools of mathematical logic (especially modal logic). Computer scientists realised the relevance of temporal logic within their discipline in the 1970s. It has since become established as an important tool in the specification and verification of computer programs. The reader may consult [104] for a detailed history of temporal logic from antiquity to the late 20th century. A comprehensive treatment of temporal logics in computer science can be found in [38].

Our work on temporal logic belongs in the computational rather than the philosophical tradition. In this chapter, in particular, we focus on Emerson and Halpern's full computation tree logic CTL* [42]. It assumes a branching model of time and extends the well-known linear-time temporal logic LTL of Pnueli's [108]¹ by means of the *universal* (A) and *existential* (E) temporal path quantifiers of Clarke and Emerson's computation tree logic CTL [26]. In contrast to the latter, it imposes no restriction on the placement of the quantifiers or the temporal operators *next* (X), *until* (U) and *release* (R) of LTL, resulting in a strictly more expressive logic than both LTL and CTL.

Much has already been achieved for the proof theory of CTL* in terms of finitary, infinitary, and cyclic tableau systems (see, e.g., [115, 117, 118, 47, 49]).

¹The until operator U, which had been earlier studied by Kamp [78], was only added to the logic in [50].

Most noteworthy is the complete Hilbert-style axiomatisation of Reynolds’s [115], which was later extended to include past modalities [116].

In the first part of this chapter, we introduce a cyclic hypersequent calculus for the logic, based on an intuitive set of inference rules. Local soundness of inferences is thus immediate, and a global correctness condition ensures that cycles yield valid conclusions. Our system is cut-free, a property which does not come naturally for temporal logics.²

Hypersequents offer a natural framework for accommodating the existential and universal path quantifiers, as well as their interplay with the next operator X . Each ‘sequent’ in a hypersequent is a labelled set of formulas, either $A\Phi$ or $E\Phi$, interpreted as ‘along all paths, $\forall \Phi$ is the case’ and ‘along some path, $\exists \Phi$ is the case’, 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 typical of ill-founded or cyclic tableaux (see, e.g., [102, 47] and [77, pp. 67–8]), but are rare in ill-founded or cyclic proof calculi. Indeed, other than our system and the one in [34], which employs a similar trace condition, there appear to be no other examples of ill-founded or cyclic systems that fall outside the scope of the category-theoretic notion of cyclic proof introduced in [2].

Developing cyclic and ill-founded proof systems beyond the traditional realm of Gentzen-style sequent calculi seems inevitable as more structured fixpoint logics are analysed. Recent examples include Rooduijn’s cyclic hypersequent calculi for modal logics with the master modality [120], and Das and Girlando’s cyclic hypersequent calculus for transitive closure logic [34].

To keep track of fixpoint unfoldings and detect ‘good’ traces, formulas are enriched with annotations similar to – but considerably simpler than – the ones introduced by Jungteerapanich [76] and Stirling [144] for the modal μ -calculus. Specifically, our annotations are strings of symbols of length ≤ 1 . Annotations allow us to isolate a finitary condition that suffices to guarantee that a cyclic derivation is a proof. This is, at the moment, only possible if correctness of every infinite path is witnessed by at least one sequent trace of type A . Further work is required to see whether existential traces admit a finitary characterisation. Indeed, whereas infinite formula traces are easy to detect by finitary means, good E -traces require the *absence* of infinite U -traces. This seems to be beyond the capabilities of our annotations, contrary to the claim made in [6]. We thus decide not to annotate any formula under the existential quantifier and rely on an infinitary correctness condition in the absence of good A -traces.

²For a discussion on cut-free sequent systems for temporal logic, see, e.g., [19].

In the second half of the chapter, we isolate a class of ‘inductible’ cyclic proofs whose cycles can be transformed into inductive arguments based on the following characterisation of until:

$$\frac{}{(\beta \vee (\alpha \wedge X(\alpha U\beta))) \rightarrow \alpha U\beta} \qquad \frac{(\beta \vee (\alpha \wedge X\gamma)) \rightarrow \gamma}{\alpha U\beta \rightarrow \gamma}$$

The axiom expresses that $\alpha U\beta$ is a pre-fixpoint of the ‘function’ $\beta \vee (\alpha \wedge X\cdot)$. The rule is the well-known fixpoint induction principle of Park’s [105] for the until operator. Together, the axiom and the rule say that $\alpha U\beta$ is the least pre-fixpoint of said function, hence the least fixpoint by the Knaster–Tarski theorem [150].

One potential application of cycle elimination, and the one driving our work, is assisting with proving completeness of Hilbert-style calculi. The problem of axiomatising CTL* remained open for several years until Reynolds presented a sound and complete axiom system for it in [115]. Our cyclic system is not fully finitary and it could not therefore yield an axiomatisation of CTL* by removal of cycles. In the end, we arrive at a Hilbert-style system and compare it to a fragment of Reynolds’s axiomatisation. Our system is complete for a well-known variant of CTL* obtained by allowing the evaluation of formulas in a bigger class than the standard one.

Outline of the chapter. Section 3.1 introduces the syntax and semantics of (the standard version of) the logic CTL*. Section 3.2 presents an ill-founded hypersequent calculus for the logic. Soundness and completeness are established, the former by a signature-based argument, the latter by means of a deterministic proof-search procedure and game- and automata-theoretic tools. It is shown that proof-search yields regular proofs, hence admitting a partially finitary presentation. Section 3.3 makes this idea precise by introducing a cyclic version of the ill-founded system, keeping the infinitary correctness condition intact. Section 3.4 introduces an ‘annotated’ cyclic version of the ill-founded system. A simple annotating mechanism keeps track of unfoldings of release formulas under universal quantifiers. A finitary condition is isolated which suffices to guarantee that a derivation is a proof. Section 3.5 isolates a class of cyclic proofs which can be transformed into acyclic, inductive ones. A corresponding, ordinary sequent calculus is presented, together with its Hilbert-style version. The latter is compared with a fragment of a known axiomatisation of full CTL*. Section 3.6 concludes the chapter and discusses some further lines of research based on the material therein.

3.1 Full computation tree logic CTL*

The *language of* CTL*, denoted by $\mathcal{L}_{\text{CTL}^*}$, consists of the following: countably many *propositional letters* drawn from a set **Prop**; the *Boolean connectives* \wedge (con-

junction), \vee (disjunction) and $\bar{}$ (negation); the *temporal operators* X (*next*), U (*until*) and R (*release*); and the (*path*) *quantifiers* A (*universal quantifier*) and E (*existential quantifier*). The *formulas* of CTL* are given by the following grammar:

$$\varphi ::= p \mid \bar{p} \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi) \mid (X\varphi) \mid (\varphi U \varphi) \mid (\varphi R \varphi) \mid (A\varphi) \mid (E\varphi),$$

where p ranges over **Prop**. Formulas are denoted by small Greek letters $\alpha, \beta, \varphi, \dots$, and sets of formulas by capital Greek letters Φ, Ψ, Ξ, \dots . We use O to denote either U or R , and Q to denote either A or E . The collection of all CTL* formulas is denoted by $\text{Form}_{\text{CTL}^*}$, and the set of all *subformulas* of a formula φ , defined as expected, is denoted by $\text{Sub}(\varphi)$.

If no ambiguity arises, we drop the outer parenthesis and stipulate that X , A and E bind more strongly than \wedge , \vee , U and R .

A *literal* is a formula of the form p or \bar{p} . An *X-formula*, or *modal formula*, is a formula of the form $X\varphi$. Similarly, an *O-formula*, for $O \in \{U, R\}$, is a formula of the form $\varphi O \psi$. An *XO-formula* is a formula of the form $X(\varphi O \psi)$. Given a set of formulas Φ , we define $X\Phi := \{X\varphi \mid \varphi \in \Phi\}$.

Observe that we allow negation to be applied to propositional letters solely. This is not necessary and one could instead work with unrestricted negation and a smaller language, but we are of the opinion that restricting negation to propositional letters makes for more readable proof rules and trace conditions.

3.1.1. REMARK. Some authors draw a grammatical distinction between *path* and *state* formulas (see, e.g., [38, § 7.1.4]). Informally, satisfiability of the former depends on the entire ‘computational’ or ‘temporal’ path that one is considering, whilst satisfiability of the latter depends only on the current state on said path. We follow [33, 115] and decide not to make this distinction grammatically. The result is a more succinct presentation of CTL* formulas and, besides, the distinction is not necessary for our purposes.

The *size*, or *complexity*, of a formula φ , denoted by $\langle \varphi \rangle$, is given by:

- $\langle p \rangle := 1$ and $\langle \bar{p} \rangle := 1$, for every $p \in \text{Prop}$;
- $\langle \varphi \star \psi \rangle := \langle \varphi \rangle + \langle \psi \rangle + 1$, for $\star \in \{\wedge, \vee, U, R\}$;
- $\langle \circ \varphi \rangle := \langle \varphi \rangle + 1$, for $\circ \in \{X, A, Q\}$.

We extend $\langle \cdot \rangle$ to finite sets of formulas by setting $\langle \Phi \rangle := \sum \{\langle \varphi \rangle \mid \varphi \in \Phi\}$.

The *dual* of a formula φ , in symbols φ^∂ , is inductively defined as follows:

$$\begin{array}{ll} p^\partial := \bar{p} & \bar{p}^\partial := p \\ (\varphi \wedge \psi)^\partial := \varphi^\partial \vee \psi^\partial & (\varphi \vee \psi)^\partial := \varphi^\partial \wedge \psi^\partial \end{array}$$

$$\begin{aligned}
(\varphi \cup \psi)^\partial &:= \varphi^\partial \mathbf{R} \psi^\partial & (\varphi \mathbf{R} \psi)^\partial &:= \varphi^\partial \cup \psi^\partial \\
(\mathbf{A}\varphi)^\partial &:= \mathbf{E}\varphi^\partial & (\mathbf{E}\varphi)^\partial &:= \mathbf{A}\varphi^\partial \\
(\mathbf{X}\varphi)^\partial &:= \mathbf{X}\varphi^\partial
\end{aligned}$$

Note that $(\cdot)^\partial$ is an involution: $(\varphi^\partial)^\partial = \varphi$. We let $\Phi^\partial := \{\varphi^\partial \mid \varphi \in \Phi\}$ for any set of formulas Φ .

We define implication as $\varphi \rightarrow \psi := \varphi^\partial \vee \psi$. Additionally, we define the constants $\perp := p \wedge \bar{p}$ (falsum) and $\top := \perp^\partial$ (verum), where p is an arbitrary propositional letter; and the temporal operators \mathbf{F} (*eventually*) and \mathbf{G} (*henceforth*), by setting $\mathbf{F}\varphi := \top \cup \varphi$ and $\mathbf{G}\varphi := \perp \mathbf{R} \varphi$. Note that $(\mathbf{F}\varphi)^\partial = \mathbf{G}\varphi^\partial$ and $(\mathbf{G}\varphi)^\partial = \mathbf{F}\varphi^\partial$. Both \mathbf{F} and \mathbf{G} bind as strongly as \mathbf{X} , and \rightarrow binds less strongly than the other connectives and operators.

Formulas of CTL* have been interpreted on different classes of structures, giving rise to different logics (see, e.g., [115, §§ 3–4]). The standard semantics for CTL* formulas, and thus also the standard CTL* logic, is obtained by evaluating formulas on maximal paths through *serial* Kripke models [38, 33, 115],³ i.e., Kripke models (W, R, V) such that $R[s] \neq \emptyset$ for every $s \in W$. A (*maximal*) *path* through a serial Kripke model (W, R, V) is an infinite sequence of states $\sigma = (s_n)_{n < \omega}$ such that $s_n R s_{n+1}$ for every $n < \omega$. We denote by (σ, i) the path $(s_k)_{k \geq i}$, for every $i < \omega$. For paths σ and σ' , we write $\sigma \sim \sigma'$ if $\sigma(0) = \sigma'(0)$.

Given a serial model $\mathcal{M} = (W, R, V)$, we inductively define a *satisfaction* relation \models between paths on \mathcal{M} and formulas in the expected manner:

- $\mathcal{M}, \sigma \models p$ if, and only if, $p \in V(\sigma(0))$;
- $\mathcal{M}, \sigma \models \bar{p}$ if, and only if, $\mathcal{M}, \sigma \not\models p$;
- $\mathcal{M}, \sigma \models \varphi \wedge \psi$ if, and only if, $\mathcal{M}, \sigma \models \varphi$ and $\mathcal{M}, \sigma \models \psi$;
- $\mathcal{M}, \sigma \models \varphi \vee \psi$ if, and only if, $\mathcal{M}, \sigma \models \varphi$ or $\mathcal{M}, \sigma \models \psi$;
- $\mathcal{M}, \sigma \models \mathbf{X}\varphi$ if, and only if, $\mathcal{M}, (\sigma, 1) \models \varphi$;
- $\mathcal{M}, \sigma \models \varphi \cup \psi$ if, and only if, there is a $j < \omega$ such that $\mathcal{M}, (\sigma, j) \models \psi$ and $\mathcal{M}, (\sigma, i) \models \varphi$ for every $i < j$;
- $\mathcal{M}, \sigma \models \varphi \mathbf{R} \psi$ if, and only if, for every $j < \omega$, either $\mathcal{M}, (\sigma, j) \models \psi$, or there is an $i < j$ such that $\mathcal{M}, (\sigma, i) \models \varphi$;
- $\mathcal{M}, \sigma \models \mathbf{A}\varphi$ if, and only if, $\mathcal{M}, \sigma' \models \varphi$ for every $\sigma' \sim \sigma$;
- $\mathcal{M}, \sigma \models \mathbf{E}\varphi$ if, and only if, $\mathcal{M}, \sigma' \models \varphi$ for some $\sigma' \sim \sigma$.

As an immediate consequence, we get:

³Recall the definition of Kripke model given in Section 2.1.

- $\mathcal{M}, \sigma \models \varphi^\partial$ if, and only if, $\mathcal{M}, \sigma \not\models \varphi$;
- $\mathcal{M}, \sigma \models F\varphi$ if, and only if, $\mathcal{M}, (\sigma, n) \models \varphi$ for some $n < \omega$;
- $\mathcal{M}, \sigma \models G\varphi$ if, and only if, $\mathcal{M}, (\sigma, n) \models \varphi$ for every $n < \omega$.

If $\mathcal{M}, \sigma \models \varphi$, we say that σ *satisfies* φ . A formula φ is *satisfiable* if there is a serial model \mathcal{M} and a path σ through \mathcal{M} such that $\mathcal{M}, \sigma \models \varphi$, and *unsatisfiable* otherwise. We write $\text{CTL}^* \models \varphi$, and say that φ is *valid*, if φ^∂ is unsatisfiable. We write $\varphi \equiv \psi$, and say that φ and ψ are *equivalent*, if for every serial model \mathcal{M} and every path σ on \mathcal{M} , we have $\mathcal{M}, \sigma \models \varphi$ if, and only if, $\mathcal{M}, \sigma \models \psi$. Explicit mention of \mathcal{M} may be omitted when no ambiguity arises.

3.1.2. OBSERVATION. It is easy to see that $\varphi U \psi \equiv \psi \vee [\varphi \wedge X(\varphi U \psi)]$ and that, dually, $\varphi R \psi \equiv \psi \wedge [\varphi \vee X(\varphi R \psi)]$. These equivalences exhibit the fixpoint nature of the until and release operators and lie at the core of our ill-founded and cyclic calculi.

At the end of Section 3.5.3 we shall examine another, more generous notion of validity for CTL* formulas.

3.2 The ill-founded system CTL_∞^*

In this section we present a cut-free, ill-founded hypersequent calculus for CTL* and show it to be sound and complete. The system is inspired by Dam's syntax trees for an embedding of CTL* into the modal μ -calculus [33].

Hypersequents — sets of sets of formulas — offer a natural framework for accommodating the existential and universal path quantifiers of the logic, as well as their interplay with the next operator X. The hypersequents that we shall use also feature in the tableau for satisfiability of CTL* formulas given in [47], where they are called ‘quantifier-bound formula blocks’.

A (*one-sided CTL**) *sequent* is a pair (Q, Φ) , henceforth written $Q\Phi$, where $Q \in \{A, E\}$ and Φ is a finite set of formulas. We identify the sequent $Q\{\varphi\}$ with the formula $Q\varphi$, and write $Q\{\Phi, \varphi\}$ as shorthand for $Q(\Phi \cup \{\varphi\})$. A sequent $Q\Phi$ is *literal* if Φ is either empty or a singleton containing a single literal. A sequent $Q\Phi$ is *universal*, or an *A-sequent*, if $Q = A$; otherwise $Q\Phi$ is said to be *existential*, or an *E-sequent*. The *interpretation* of a universal sequent $A\Phi$ is the formula $(A\Phi)^\sharp := A \vee \Phi$, and the *interpretation* of an existential sequent $E\Phi$ is $(E\Phi)^\sharp := E \wedge \Phi$. A sequent $Q\Phi$ is *valid* if $(Q\Phi)^\sharp$ is valid. We abuse notation and write $\mathcal{M}, \sigma \models Q\Phi$ in place of $\mathcal{M}, \sigma \models (Q\Phi)^\sharp$. The *dual* of a sequent $Q\Phi$ is the sequent $(Q\Phi)^\partial := Q^\partial \Phi^\partial$, where $A^\partial := E$ and $E^\partial := A$. The *size*, or *complexity*, of a sequent $Q\Phi$ is $\langle Q\Phi \rangle := 1 + \langle \Phi \rangle$.

A (CTL^*) *hypersequent* is a finite set of sequents. Hypersequents are denoted by capital Greek letters $\Gamma, \Delta, \Sigma, \dots$. When working with hypersequents, we abbreviate $\{\mathbf{Q}_1\Phi_1, \dots, \mathbf{Q}_n\Phi_n\}$ to $\mathbf{Q}_1\Phi_1, \dots, \mathbf{Q}_n\Phi_n$. We extend the interpretation $(\cdot)^\sharp$ to hypersequents by setting $\Gamma^\sharp := \bigvee\{(\mathbf{Q}\Phi)^\sharp \mid \mathbf{Q}\Phi \in \Gamma\}$. A hypersequent Γ is *valid* if Γ^\sharp is valid. We abuse notation and write $\mathcal{M}, \sigma \models \Gamma$ as shorthand for $\mathcal{M}, \sigma \models \Gamma^\sharp$. The *size*, or *complexity*, of a hypersequent Γ is $\langle \Gamma \rangle := \sum\{(\mathbf{Q}\Phi) \mid \mathbf{Q}\Phi \in \Gamma\}$.

A formula α *occurs* in a sequent $\mathbf{Q}\Phi$ if $\alpha \in \Phi$. Analogously, a sequent $\mathbf{Q}\Phi$ *occurs* in a hypersequent Γ if $\mathbf{Q}\Phi \in \Gamma$. Finally, a formula α *occurs* in a hypersequent Γ if α occurs in a sequent occurring in Γ .

The rules of the system CTL_∞^* are given in Figures 3.1 to 3.3. In the rules ALit and ELit , ℓ ranges over literals. The fixpoint rules AU , AR , EU and ER are based on the equivalences in Observation 3.1.2. Contraction, both at the internal and the external level, is implicit in our choice of set-based rather than multiset-based sequents and hypersequents. Note that the internal weakening rule iW is only sound for A -sequents.

3.2.1. DEFINITION (CTL_∞^* derivation). A CTL_∞^* *derivation* of a formula φ is a finite or infinite labelled tree \mathcal{T} built according to the rules in Figures 3.1 to 3.3 and such that:

- (i) The root of \mathcal{T} has label $\text{A}\varphi$.
- (ii) Every infinite branch of \mathcal{T} passes through infinitely many instances of the rule AX or EX .

–

3.2.2. OBSERVATION. Condition (i) in Definition 3.2.1 carries no loss of generality, because, clearly, a formula φ is valid if, and only if, $\text{A}\varphi$ is valid.

A vertex of a derivation is *modal* if it is labelled by the conclusion of an instance of AX or EX . Condition (ii) in Definition 3.2.1 requires that every infinite branch of a derivation contain infinitely many modal vertices. This prevents the construction of infinite branches by ‘degenerate’ applications of the rules, for instance applying rule ALit repeatedly to sequent $\text{A}p$ or invoking implicit contraction.

3.2.3. REMARK. Hypersequents and universal sequents are interpreted disjunctively, while existential sequents are interpreted conjunctively. Concordantly, external commas and commas under A in CTL_∞^* rules stand for disjunctions, and commas under E are treated conjunctively.

This seemingly inconsequential difference will turn out to be the source of endless complications. Informally, E -sequents in CTL^* derivations behave like sequents in a tableau rather than a proof-tree, and thus resist many of the proof-theoretic manipulations that we are interested in. See Remark 3.2.15 below.

$$\begin{array}{l}
\text{ax} \frac{}{Q\rho, Q'\bar{\rho}, \Delta} \\
\text{ALit} \frac{A\Phi, A\ell, \Delta}{A\{\Phi, \ell\}, \Delta} \\
\text{AV} \frac{A\{\Phi, \varphi, \psi\}, \Delta}{A\{\Phi, \varphi \vee \psi\}, \Delta} \\
\text{A}\wedge \frac{A\{\Phi, \varphi\}, \Delta \quad A\{\Phi, \psi\}, \Delta}{A\{\Phi, \varphi \wedge \psi\}, \Delta} \\
\text{AA} \frac{A\Phi, A\{\psi\}, \Delta}{A\{\Phi, A\psi\}, \Delta} \\
\text{AE} \frac{A\Phi, E\{\psi\}, \Delta}{A\{\Phi, E\psi\}, \Delta} \\
\text{AX} \frac{A\Phi, E\Psi_1, \dots, E\Psi_m}{A(X\Phi \cup \Xi), EX\Psi_1, \dots, EX\Psi_m, \Sigma} \\
\text{ax}\top \frac{}{E\emptyset, \Delta} \\
\text{ELit} \frac{E\Phi, \Delta \quad E\ell, \Delta}{E\{\Phi, \ell\}, \Delta} \\
\text{EV} \frac{E\{\Phi, \varphi\}, E\{\Phi, \psi\}, \Delta}{E\{\Phi, \varphi \vee \psi\}, \Delta} \\
\text{E}\wedge \frac{E\{\Phi, \varphi, \psi\}, \Delta}{E\{\Phi, \varphi \wedge \psi\}, \Delta} \\
\text{EA} \frac{E\Phi, \Delta \quad A\{\psi\}, \Delta}{E\{\Phi, A\psi\}, \Delta} \\
\text{EE} \frac{E\Phi, \Delta \quad E\{\psi\}, \Delta}{E\{\Phi, E\psi\}, \Delta} \\
\text{EX} \frac{E\Psi_1, \dots, E\Psi_m}{EX\Psi_1, \dots, EX\Psi_m, \Sigma}
\end{array}$$

FIGURE 3.1: Non-fixpoint, logical rules of the system CTL_∞^* .

The rules of CTL_∞^* are correct in the following sense:

3.2.4. PROPOSITION. *If*

$$R \frac{\Gamma_1 \quad \dots \quad \Gamma_n}{\Delta}$$

is an instance of a CTL_∞^ rule and each of $\Gamma_1, \dots, \Gamma_n$ is valid, then Δ is valid.*

Proof. All cases are straightforward, so we only prove the case $R = \text{AX}$. Let then

$$\text{AX} \frac{A\Phi, E\Psi_1, \dots, E\Psi_m}{A(X\Phi \cup \Xi), EX\Psi_1, \dots, EX\Psi_m, \Sigma}$$

be an instance of AX, and suppose that the conclusion is not valid. Let \mathcal{M} be a serial model and σ a path on \mathcal{M} such that $\sigma \not\models A(X\Phi \cup \Xi), EX\Psi_1, \dots, EX\Psi_m, \Sigma$. Then, there is some $\sigma \sim \sigma'$ such that $\sigma' \not\models X\vee\Phi$, whence $(\sigma', 1) \not\models A\Phi$. It suffices to see that $(\sigma', 1) \not\models E\Psi_i$ for any $1 \leq i \leq m$. Towards a contradiction, suppose that $(\sigma', 1) \models E\Psi_i$. Then, there is some $(\sigma', 1) \sim \sigma''$ such that $\sigma'' \models \wedge \Psi_i$. Let $\sigma^* := (\sigma(0)) \frown \sigma''$. Since $(\sigma^*, 1) = \sigma''$ and $\sigma \sim \sigma^*$, we get $\sigma \models EX \wedge \Psi_i$, contradicting the starting assumption. \blacksquare

$$\begin{array}{c}
\text{AU} \frac{A\{\Phi, \varphi, \psi\}, \Delta \quad A\{\Phi, \psi, X(\varphi U \psi)\}, \Delta}{A\{\Phi, \varphi U \psi\}, \Delta} \quad \text{EU} \frac{E\{\Phi, \psi\}, E\{\Phi, \varphi, X(\varphi U \psi)\}, \Delta}{E\{\Phi, \varphi U \psi\}, \Delta} \\
\text{AR} \frac{A\{\Phi, \psi\}, \Delta \quad A\{\Phi, \varphi, X(\varphi R \psi)\}, \Delta}{A\{\Phi, \varphi R \psi\}, \Delta} \quad \text{ER} \frac{E\{\Phi, \varphi, \psi\}, E\{\Phi, \psi, X(\varphi R \psi)\}, \Delta}{E\{\Phi, \varphi R \psi\}, \Delta}
\end{array}$$

FIGURE 3.2: Fixpoint rules of the system CTL_∞^* .

The distinguished sequents in the conclusions of the rules in Figures 3.1 to 3.3, such as $A\{\Phi, \varphi \wedge \psi\}$ in $A\wedge$, are said to be *principal*. In the special cases of the *modal rules* AX and EX , we consider all sequents in the conclusion to be principal. And in the case of eW , the principal sequents are the ones occurring in Γ . The distinguished sequents in the premises of the rules, for instance $A\{\Phi, \varphi\}$ and $A\{\Phi, \psi\}$ in $A\wedge$, are said to be *active*. In the case of the modal rules, all sequents in the premise are considered active. And in the case of eW , there are no active sequents. All sequents occurring in the hypersequents Δ are *side* sequents.

Analogously, the distinguished formulas in principal sequents, such as $\varphi \wedge \psi$ in rule $A\wedge$, are said to be *principal*. In the case of the modal rules, we consider all formulas in the conclusion to be principal. And in the cases of iW and eW , the principal formulas are the ones occurring in Ψ and Γ , respectively. The distinguished formulas in active sequents, for instance φ and ψ in $A\wedge$, are said to be *active*. In the case of the modal rules, all formulas in the premise are considered active. And in the cases of iW and eW , no formula is active. All formulas occurring in the sets Φ , excluding the case of AX , are *side* formulas.

Since we work with set-based sequents and hypersequents, a side formula or sequent might also be principal. For example, the following is an instance of AV in which the principal formula $\varphi \vee \psi$ is also a side formula:

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

An instance of a non-modal CTL^* rule is *externally preserving* if some principal sequent is also a side sequent, and it is *externally discarding* otherwise. Analogously, an instance of a non-modal rule is *internally preserving* if some principal formula is also a side formula, and it is *internally discarding* otherwise. Finally, an instance of a non-modal rule is *discarding* if it is both externally and internally discarding. Note that, by definition, instances of axiomatic or modal rules are neither preserving nor discarding.

A CTL_∞^* derivation \mathcal{T} is *discarding* if no rule is applied preservingly (internally or externally) in \mathcal{T} .

$$\text{iW} \frac{A\Phi, \Delta}{A(\Phi \cup \Psi), \Delta} \quad \text{eW} \frac{\Delta}{\Gamma, \Delta}$$

FIGURE 3.3: Structural rules of the system CTL_∞^* .

Looking at Figure 3.1, we say that an instance of ALit or ELit is *degenerate* if $\Phi = \emptyset$.

Let \mathcal{T} be a CTL_∞^* derivation, let $u, v \in T$ be such that $u <_T^0 v$, and let Γ and Γ' be the labels of u and v , respectively. We define the *sequent trace* relation \triangleright_u^v between Γ and Γ' by setting $\text{Q}\Phi \triangleright_u^v \text{Q}'\Psi'$ if, and only if, one of the following holds:

- u is non-modal, $\text{Q}\Phi$ and $\text{Q}'\Psi$ are side sequents, and $\text{Q}\Phi = \text{Q}'\Psi$;
- u is non-modal, $\text{Q}\Phi$ is principal, and $\text{Q}'\Psi$ is active;
- u is modal, $\text{Q} = \text{E}$, and $\text{Q}\Phi = \text{Q}'\text{X}\Psi$.
- u is modal, $\text{Q} = \text{A}$, and $\text{Q}\Phi = \text{Q}'(\text{X}\Psi \cup \Xi)$ for some set of formulas Ξ .

So, informally, we have $\text{Q}\Phi \triangleright_u^v \text{Q}'\Psi$ if $\text{Q}'\Psi$ arises from $\text{Q}\Phi$ in the rule with conclusion u . We may drop one or both indices from \triangleright_u^v if no ambiguity arises.

Let $\pi = (u_n)_{n < N \leq \omega}$ be a finite or infinite path through \mathcal{T} . A *sequent trace* on π is a sequence of sequents $(\text{Q}_n\Phi_n)_{n < N}$ such that each $\text{Q}_n\Phi_n$ occurs in the label of u_n and $\text{Q}_n\Phi_n \triangleright_{u_n}^{u_{n+1}} \text{Q}_{n+1}\Phi_{n+1}$ for every $n < n+1 < N$. A *context extraction* is a sequent trace of the form $\text{Q}\{\Phi, \text{Q}'\psi\} \triangleright \text{Q}'\{\psi\}$ or $\text{Q}\{\Phi, \ell\} \triangleright \text{Q}\{\ell\}$, where ℓ is a literal. Note that context extractions are only due to rules ALit, ELit, AA, EA, AE and EE. A finite sequent trace $\text{Q}_0\Phi_0 \triangleright \dots \triangleright \text{Q}_n\Phi_n$ is *stable* if $\text{Q}_0 = \dots = \text{Q}_n$, and *circular* if $n > 0$ and $\text{Q}_n\Phi_n = \text{Q}_0\Phi_0$.

Let again $u, v \in T$ be such that $u <_T^0 v$, and let Γ and Γ' be the labels of u and v , respectively. For all $\text{Q}\Phi \in \Gamma$ and $\text{Q}'\Psi \in \Gamma'$ such that $\text{Q}\Phi \triangleright_u^v \text{Q}'\Psi$, and all formulas $\varphi \in \Phi$ and $\psi \in \Psi$, we write $\text{Q}\Phi, \varphi \triangleright_u^v \text{Q}'\Psi, \psi$ if, and only if, one of the following holds:

- u is non-modal, $\text{Q}\Phi$ and $\text{Q}'\Psi$ are side sequents, and $\varphi = \psi$;
- u is non-modal, $\text{Q}\Phi$ is principal, both φ and ψ are side formulas, and $\varphi = \psi$;
- u is non-modal, $\text{Q}\Phi$ and φ are principal, and ψ is active;
- u is modal, and $\varphi = \text{X}\psi$.

As before, $\text{Q}\Phi, \varphi \triangleright_u^v \text{Q}'\Psi, \psi$ holds, informally, if ψ arises from φ in the rule with conclusion u . We may drop one or both indices from \triangleright_u^v if no ambiguity arises, as well as write $\varphi \triangleright_u^v \psi$ in place of $\text{Q}\Phi, \varphi \triangleright_u^v \text{Q}'\Psi, \psi$ when $\text{Q}\Phi$ and $\text{Q}'\Psi$ are clear from context.

Let $\tau = (\mathbf{Q}_n \Phi_n)_{n < N}$ be a sequent trace on a finite or infinite path $\pi = (u_n)_{n < N \leq \omega}$. A *formula trace* on τ is a sequence of formulas $(\varphi_n)_{n < N}$ such that each $\varphi_n \in \Phi_n$ and $\mathbf{Q}_n \Phi_n, \varphi_n \triangleright_{u_n}^{u_{n+1}} \mathbf{Q}_{n+1} \Phi_{n+1}, \varphi_{n+1}$. An infinite formula trace on an infinite path π of a CTL_∞^* derivation is a formula trace on some (infinite) sequent trace on π .

A (*fixpoint*) *unfolding* is a formula trace of the form $\psi \triangleright \mathbf{X}\psi$. We say that ψ is *unfolded (to $\mathbf{X}\psi$)*. Note that unfoldings can only be produced by rules AU, AR, EU, ER, and thus ψ is an O-formula. Due to the presence of fixpoint unfoldings, the system CTL_∞^* does not satisfy the subformula property: $\varphi \triangleright \psi$ does not imply $\psi \in \text{Sub}(\varphi)$. However, ψ does belong to the *closure* of φ , which is the natural replacement of the notion of subformula in this context:

3.2.5. DEFINITION (CTL^* closure). The (CTL^*) *closure* of a formula φ is the smallest set of formulas $\text{Clos}(\varphi)$ satisfying:

- (i) $\varphi \in \text{Clos}(\varphi)$;
- (ii) if $\bar{p} \in \text{Clos}(\varphi)$, for $p \in \text{Prop}$, then $p \in \text{Clos}(\varphi)$;
- (iii) if $\psi_1 \star \psi_2 \in \text{Clos}(\varphi)$, for $\star \in \{\wedge, \vee\}$, then $\psi_1, \psi_2 \in \text{Clos}(\varphi)$;
- (iv) if $\psi_1 \mathbf{O} \psi_2 \in \text{Clos}(\varphi)$, for $\mathbf{O} \in \{\mathbf{U}, \mathbf{R}\}$, then $\psi_1, \psi_2, \mathbf{X}(\psi_1 \mathbf{O} \psi_2) \in \text{Clos}(\varphi)$;
- (v) if $\mathbf{X}\psi \in \text{Clos}(\varphi)$, then $\psi \in \text{Clos}(\varphi)$;
- (vi) if $\mathbf{Q}\psi \in \text{Clos}(\varphi)$, for $\mathbf{Q} \in \{\mathbf{A}, \mathbf{E}\}$, then $\psi \in \text{Clos}(\varphi)$. ⊣

By a straightforward structural induction, the closure of a formula can be characterised as follows:

3.2.6. LEMMA. *The following hold:*

- (i) $\text{Clos}(p) = \{p\}$ and $\text{Clos}(\bar{p}) = \{p, \bar{p}\}$, for every $p \in \text{Prop}$;
- (ii) $\text{Clos}(\psi_1 \star \psi_2) = \{\psi_1 \star \psi_2\} \cup \text{Clos}(\psi_1) \cup \text{Clos}(\psi_2)$, for $\star \in \{\wedge, \vee\}$;
- (iii) $\text{Clos}(\circ\psi) = \{\circ\psi\} \cup \text{Clos}(\psi)$, for $\circ \in \{\mathbf{X}, \mathbf{A}, \mathbf{E}\}$;
- (iv) $\text{Clos}(\psi_1 \mathbf{O} \psi_2) = \{\psi_1 \mathbf{O} \psi_2, \mathbf{X}(\psi_1 \mathbf{O} \psi_2)\} \cup \text{Clos}(\psi_1) \cup \text{Clos}(\psi_2)$, for $\mathbf{O} \in \{\mathbf{U}, \mathbf{R}\}$.

3.2.7. COROLLARY. *For any CTL^* formula φ , the set $\text{Clos}(\varphi)$ is finite.*

For every formula φ , we let $|\varphi| := |\text{Clos}(\varphi)|$. We also let Seq_φ be the collection of all sequents $\mathbf{Q}\Phi$ with $\Phi \subseteq \text{Clos}(\varphi)$. Finally, we let $\text{HSeq}_\varphi := 2^{\text{Seq}_\varphi}$ be the collection of all hypersequents Γ such that $\Gamma \subseteq \text{Seq}_\varphi$. By Corollary 3.2.7, both Seq_φ and HSeq_φ are finite.

Clearly, $\varphi \triangleright \psi$ implies $\psi \in \text{Clos}(\varphi)$, whence Seq_φ is closed under \triangleright and every hypersequent labelling a vertex of a CTL_∞^* derivation of φ belongs to HSeq_φ . Moreover, since $\mathbf{X}\varphi \triangleright \psi$ implies $\psi \in \{\mathbf{X}\varphi, \varphi\}$, we have the following fundamental lemma about formula traces:

3.2.8. LEMMA. *Let $\rho = (\varphi_n)_{n < N \leq \omega}$ be a finite or infinite formula trace. For every $n < N$, either $\varphi_n \in \text{Sub}(\varphi_0)$, or $\varphi_n = X\psi$ for some $\psi \in \text{Sub}(\varphi_0)$.*

Proof. Let $\pi = (u_n)_{n < N}$ be a path through a CTL* derivation such that ρ is a formula trace on some sequent trace on π . We proceed by induction on $n < N$. The base case $n = 0$ is clear. Assume that the claim holds for an n such that $n + 1 < N$. If $\varphi_{n+1} = \varphi_n$ we are done, so assume otherwise. Then, φ_n must be principal at u_n .

Suppose first that $\varphi_n = X\psi$ for some $\psi \in \text{Sub}(\varphi_0)$. Then, rule X must be applied at u_n because $\varphi_{n+1} \neq \varphi_n$, so by the definition of \triangleright we have $\varphi_n = X\varphi_{n+1}$ and thus $\varphi_{n+1} = \psi \in \text{Sub}(\varphi_0)$.

Suppose now that $\varphi_n \in \text{Sub}(\varphi_0)$. If $\varphi_{n+1} \in \text{Sub}(\varphi_n)$ we are done. Otherwise, φ_n is an O-formula for some $O \in \{U, R\}$ and a rule among AO, EO is applied at u_n with principal formula φ_n and φ_{n+1} an active formula. Since $\varphi_{n+1} \notin \text{Sub}(\varphi_n)$, we have $\varphi_{n+1} = X\varphi_n$ and we are done. ■

We shall mainly concern ourselves with infinite traces. All of them eventually stabilise on some O-formula that is unfolded infinitely often, as the next proposition shows.

3.2.9. PROPOSITION. *Let \mathcal{T} be a CTL*_∞ derivation of a formula φ , and $\rho = (\varphi_i)_{i < \omega}$ an infinite formula trace on an infinite path of \mathcal{T} . There is an O-formula $\psi \in \text{Clos}(\varphi)$, for $O \in \{U, R\}$, and some $N < \omega$ such that for every $i \geq N$ we have $\varphi_i \in \{\psi, X\psi\}$. Moreover, both ψ and $X\psi$ occur infinitely often in $(\varphi_{N+i})_{i < \omega}$.*

Proof. Towards a contradiction, suppose that ρ encounters only finitely many fixpoint unfoldings. Let then $n < \omega$ be such that there are no unfoldings on $\rho_{\geq n}$. Thus, $\langle \varphi_{i+1} \rangle \leq \langle \varphi_i \rangle$ for all $i \geq n$. If $i < \omega$ is such that $\varphi_i = X\varphi_{i+1}$, then $\langle \varphi_{i+1} \rangle < \langle \varphi_i \rangle$. So, since by definition every infinite branch of \mathcal{T} contains infinitely modal vertices, we have $\langle \varphi_{i+1} \rangle < \langle \varphi_i \rangle$ for infinitely many $i \geq n$, which is impossible. We conclude that there are infinitely many unfoldings on ρ .

Since $\text{Clos}(\varphi)$ is finite, there is some O-formula $\psi \in \text{Clos}(\varphi)$, for $O \in \{U, E\}$, such that ψ is unfolded infinitely often in ρ . It remains to see that ρ has a tail which contains only ψ and $X\psi$. Suppose otherwise, and let α be a formula occurring infinitely often on ρ and such that $\alpha \notin \{\psi, X\psi\}$. By Lemma 3.2.8, then, $\psi \in \text{Sub}(\alpha)$ and $\alpha = X\beta$ for some $\beta \in \text{Sub}(\psi)$. As $\psi \neq \alpha$, we have $\psi \in \text{Sub}(\beta)$ and thus $\psi = \beta$, whence $\alpha = X\psi$ and we have reached a contradiction. ■

The formula $\psi_1 O \psi_2$ given by Proposition 3.2.9 is said to be the *dominating* formula in ρ . An infinite formula trace is *of type O*, or an *O-trace*, for $O \in \{U, R\}$, if its dominating formula is an O-formula. Proposition 3.2.9 then says that every infinite formula trace is of type U or of type E.

Another consequence of Lemma 3.2.8 is that there cannot be a formula trace of the form $\varphi \triangleright \cdots \triangleright \mathbf{Q}\varphi$, whence it follows that infinite sequent traces must be stable.

3.2.10. PROPOSITION. *Let $\tau = (\mathbf{Q}_i\Phi_i)_{i<\omega}$ be an infinite sequent trace. There is some $N < \omega$ such that $\mathbf{Q}_{N+k} = \mathbf{Q}_N$ for every $k < \omega$.*

Proof. By König's lemma, there is at least one infinite formula trace $\rho = (\varphi_i)_{i<\omega}$ on τ . Since $\varphi_i \in \text{Clos}(\varphi_0)$ for every $i < \omega$, by Corollary 3.2.7 there is an $N < \omega$ such that φ_{N+i} occurs infinitely often on ρ for every $i < \omega$. Towards a contradiction, suppose that there is an $n \geq N$ such that $\mathbf{Q}_{n+1} \neq \mathbf{Q}_n$. By the definition of \triangleright , we then have $\mathbf{Q}_n\Phi_n = \mathbf{Q}\{\Phi, \mathbf{Q}'\psi\}$ and $\mathbf{Q}_{n+1}\Phi_{n+1} = \mathbf{Q}'\{\psi\}$, so $\varphi_n = \mathbf{Q}'\psi$ and $\varphi_{n+1} = \psi$. Since φ_n occurs infinitely often on ρ , by Lemma 3.2.8 we have $\mathbf{Q}'\psi \in \text{Sub}(\psi)$, which is impossible. Hence, $\mathbf{Q}_{N+k} = \mathbf{Q}_N$ for all $k < \omega$. ■

Given an infinite sequent trace $\tau = (\mathbf{Q}_i\Phi_i)_{i<\omega}$, we say that τ is of *type Q*, or that it is a *Q-trace*, for $\mathbf{Q} \in \{\mathbf{A}, \mathbf{E}\}$, if there is an $N < \omega$ such that $\mathbf{Q}_n = \mathbf{Q}$ for all $n \geq N$. Proposition 3.2.10 then says that every infinite sequent trace is either of type **A** or of type **E**.

We are now ready to identify the CTL_∞^* derivations that constitute proofs. Informally, a derivation \mathcal{T} is a proof if every leaf of \mathcal{T} is axiomatic and every infinite branch of \mathcal{T} contains a 'good' sequent trace.

We say that an infinite sequent trace τ *contains* an infinite formula trace ρ if there is a tail τ' of τ such that ρ is a trace on τ' .

3.2.11. DEFINITION (Good and bad sequent trace). An infinite sequent trace τ is *good* if one of the following hold:

- (i) τ is of type **A** and contains an infinite **R**-trace;
- (ii) τ is of type **E** and contains no infinite **U**-trace.

Otherwise, τ is said to be *bad*. ⊣

A branch π of a CTL_∞^* derivation is *good* if either π is finite and ends at an axiomatic vertex, or π is infinite and there is a good sequent trace on π . Otherwise π is *bad*.

3.2.12. DEFINITION (CTL $_\infty^*$ proof). A *CTL $_\infty^*$ proof* of a formula φ is a CTL_∞^* derivation \mathcal{T} of φ such that every branch of \mathcal{T} is good. ⊣

3.2.13. OBSERVATION. Let \mathcal{T} be a CTL_∞^* derivation, and π an infinite branch of \mathcal{T} . Then, for every tail π' of π and every sequent trace τ' on π' , there is a sequent trace τ on π such that τ' is a tail of τ . Therefore, requiring that π be good is equivalent to requiring that some tail of π be good.

$$\begin{array}{c}
\vdots \\
\text{AX} \frac{(\dagger) \text{A}\{\top \text{U}p, \perp \text{R}\bar{p}\}}{\text{A}\{p, \text{X}(\top \text{U}p), \text{X}(\perp \text{R}\bar{p})\}} \\
\text{AU} \frac{\text{A}\{\top \text{U}p, p, \text{X}(\perp \text{R}\bar{p})\}}{\text{A}\{\top \text{U}p, \perp, \text{X}(\perp \text{R}\bar{p})\}} \\
\text{AR} \frac{(\dagger) \text{A}\{\top \text{U}p, \perp \text{R}\bar{p}\}}{\text{A}\{\top \text{U}p \vee \perp \text{R}\bar{p}\}} \\
\text{AV} \frac{\text{A}\{\top \text{U}p, \perp \text{R}\bar{p}\}}{\text{A}\{\top \text{U}p \vee \perp \text{R}\bar{p}\}}
\end{array}$$

FIGURE 3.4: An infinite branch of a CTL_{∞}^* proof of the formula $(\top \text{U}p) \vee (\perp \text{R}\bar{p})$. The symbol ‘†’ marks roots of identical subproofs. Dashed lines indicate the omission of some vertices.

3.2.14. EXAMPLE. Figure 3.4 depicts (the unique infinite branch of) a CTL_{∞}^* proof of the valid formula $(\top \text{U}p) \vee (\perp \text{R}\bar{p})$. On the branch there is an **A**-trace containing an **R**-trace with dominating formula $\perp \text{R}\bar{p}$.

3.2.15. REMARK. The correctness condition for **E**-traces requires that they contain no infinite **U**-trace. This is line with the ‘satisfiability flavour’ of **E**-sequents pointed out in Remark 3.2.3 above. Correctness conditions of this kind are typical of ill-founded or cyclic tableaux (see, e.g., [102, 47] and [77, pp. 67–8]), but are rare in ill-founded or cyclic proof calculi. Indeed, other than our system and the one in [34], which employs a similar trace condition, there appear to be no other examples of ill-founded or cyclic systems that fall outside the scope of the category-theoretic notion of cyclic proof introduced in [2].

We shall spend the remaining of this section proving the soundness and completeness of CTL_{∞}^* .

3.2.1 Soundness of CTL_{∞}^*

To prove that the system CTL_{∞}^* is sound we follow a standard argument based on *signatures*, maps that assign natural numbers to until and release formulas. This idea goes back at least to the signatures introduced in [145] for the modal μ -calculus.

For every $n < \omega$ we define the *n*-th approximation $\varphi \text{U}^n \psi$ of a formula $\varphi \text{U} \psi$ by setting $\varphi \text{U}^0 \psi := \psi$ and $\varphi \text{U}^{n+1} \psi := \psi \vee (\varphi \wedge \text{X}(\varphi \text{U}^n \psi))$. Dually, we define the *n*-th approximation $\varphi \text{R}^n \psi$ of $\varphi \text{R} \psi$ as $\varphi \text{R}^n \psi := (\varphi^{\partial} \text{U}^n \psi^{\partial})^{\partial}$.

3.2.16. LEMMA. *For any serial model \mathcal{M} , any path σ on \mathcal{M} , any formulas φ, ψ and any $n < \omega$, the following hold:*

- (i) $\mathcal{M}, \sigma \models \varphi \text{U}^n \psi$ if, and only if, there is a $j \leq n$ such that $\mathcal{M}, (\sigma, j) \models \psi$ and, for every $i < j$, $\mathcal{M}, (\sigma, i) \models \varphi$.
- (ii) $\mathcal{M}, \sigma \models \varphi \text{R}^n \psi$ if, and only if, for every $j \leq n$, either $\mathcal{M}, (\sigma, j) \models \psi$, or there is an $i < j$ such that $\mathcal{M}, (\sigma, i) \models \varphi$.

Proof. We only prove the first claim, for the second one can be proved analogously. We proceed by induction on $n < \omega$. The base case $n = 0$ is clear. Assume that the claim holds for n .

Suppose that $\mathcal{M}, \sigma \models \varphi \text{U}^{n+1} \psi$, i.e., $\mathcal{M}, \sigma \models \psi \vee (\varphi \wedge \text{X}(\varphi \text{U}^n \psi))$. If $\mathcal{M}, \sigma \models \psi$ we pick $j := 0$ and are done. Otherwise, $\mathcal{M}, \sigma \models \varphi$ and $\mathcal{M}, (\sigma, 1) \models \varphi \text{U}^n \psi$, so by the inductive hypothesis there is a $j' \leq n$ such that $\mathcal{M}, (\sigma, 1 + j') \models \psi$ and $\mathcal{M}, (\sigma, 1 + i) \models \varphi$ for every $i < j'$. We may thus pick $j := j' + 1$.

Conversely, suppose that there is a $j \leq n + 1$ such that $\mathcal{M}, (\sigma, j) \models \psi$ and $\mathcal{M}, (\sigma, i) \models \varphi$ for every $i < j$. If $j = 0$ we are done. Otherwise, the inductive hypothesis yields $\mathcal{M}, (\sigma, 1) \models \varphi \text{U}^n \psi$, whence $\mathcal{M}, \sigma \models \psi \vee (\varphi \wedge \text{X}(\varphi \text{U}^n \psi))$. \blacksquare

3.2.17. COROLLARY. *The following hold:*

- (i) $\mathcal{M}, \sigma \models \varphi \text{U} \psi$ if, and only if, $\mathcal{M}, \sigma \models \varphi \text{U}^n \psi$ for some $n < \omega$.
- (ii) $\mathcal{M}, \sigma \models \varphi \text{R} \psi$ if, and only if, $\mathcal{M}, \sigma \models \varphi \text{R}^n \psi$ for all $n < \omega$.

An occurrence in φ of a subformula $\psi_1 \text{O} \psi_2$ 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 [33].

3.2.18. REMARK. Note that eventualities are not just subformulas, but occurrences thereof.

3.2.19. DEFINITION (Signature). 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 $\text{Q}\Phi$ is an O-signature of the formula $\bigwedge \Phi$.⁴ \dashv

3.2.20. REMARK. Although we have defined signatures both for U- and R- eventualities, R-signatures suffice for our purposes.

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

⁴This conjunction is unrelated to the interpretation of the sequent $\text{Q}\Phi$; it is simply used to bring the formulas in Φ together.

- $\ell[\iota] := \ell$, for every literal ℓ .
- $(\psi_1 \star \psi_2)[\iota] := (\psi_1[\iota]) \star (\psi_2[\iota])$, for $\star \in \{\wedge, \vee\}$.
- $(Q\psi)[\iota] := Q\psi$, for $Q \in \{A, E\}$.
- $(\psi_1 O\psi_2)[\iota] := \psi_1 O^{\iota(\psi_1 O\psi_2)}\psi_2$.
- $(\psi_1 O'\psi_2)[\iota] := \psi_1 O'\psi_2$, for $O' \neq O$.
- $(X\psi)[\iota] := X(\psi[\iota^-])$.

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

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

3.2.21. PROPOSITION. *Let \mathcal{M} be a serial model and σ a path on \mathcal{M} . For every formula φ , the following hold:*

- (i) *if $\sigma \not\models \varphi$, then there is an R-signature ι of φ such that $\sigma \not\models \varphi[\iota]$.*
- (ii) *If $\sigma \not\models A\Phi$, then there is an R-signature ι of $A\Phi$ such that $\sigma \not\models A\Phi[\iota]$.*

Proof.

- (i) By structural induction on φ . The literal and Boolean cases are clear. And so are the cases of U, A and E because there are no top-level R- eventualities in a formula of the form $\psi_1 U\psi_2$ or $Q\psi$.

Suppose that $\sigma \not\models X\psi$. By the inductive hypothesis, there is an R-signature ι of ψ such that $(\sigma, 1) \not\models \psi[\iota]$. Let ι_+ be the R-signature of ψ given by $\iota_+ := \iota + 1$. Then, ι_+ is an R-signature of $X\psi$ and $(\iota_+)^- = \iota$, so $(\sigma, 1) \not\models \psi[(\iota_+)^-]$ and thus $\sigma \not\models (X\psi)[\iota_+]$.

Finally, suppose that $\sigma \not\models \psi_1 R\psi_2$. Let $n < \omega$ be given by Corollary 3.2.17 such that $\sigma \not\models \psi_1 R^n \psi_2$. Let then ι be the R-signature of $\psi_1 R\psi_2$ given by $\iota(\psi_1 R\psi_2) := n$.

- (ii) Let $\Phi = \{\varphi_1, \dots, \varphi_n\}$ and let $\sigma' \sim \sigma$ be such that $\sigma' \not\models \varphi_1 \vee \dots \vee \varphi_n$. By (i), there is a signature ι of $\bigvee_{1 \leq k \leq n} \varphi_k$ such that $\sigma' \not\models (\varphi_1 \vee \dots \vee \varphi_n)[\iota]$. Then, ι is a signature of $A\Phi$ and $(A\Phi[\iota])^\sharp = A(\varphi_1[\iota] \vee \dots \vee \varphi_n[\iota]) = A((\varphi_1 \vee \dots \vee \varphi_n)[\iota])$, so we are done. ■

The previous result related signatures to formulas and sequents. The next one extends this to hypersequents, and will be used in the soundness proof to guide us along an infinite branch.

3.2.22. PROPOSITION. *Let \mathcal{M} be a serial model and σ a path on \mathcal{M} . The following hold, where in each case ι is an R-signature of the appropriate sequent:*

- (i) If $\sigma \not\models \mathbf{A}\{\Phi, \ell\}[l]$ and ℓ is a literal, then $\sigma \not\models \mathbf{A}\Phi[l]$ and $\sigma \not\models \ell[l]$.
- (ii) If $\sigma \not\models \mathbf{A}\{\Phi, \varphi \vee \psi\}[l]$, then $\sigma \not\models \mathbf{A}\{\Phi, \varphi, \psi\}[l]$.
- (iii) If $\sigma \not\models \mathbf{A}\{\Phi, \varphi \wedge \psi\}[l]$, then either $\sigma \not\models \mathbf{A}\{\Phi, \varphi\}[l]$ or $\sigma \not\models \mathbf{A}\{\Phi, \psi\}[l]$.
- (iv) If $\sigma \not\models \mathbf{A}\{\Phi, \mathbf{Q}\psi\}[l]$, for $\mathbf{Q} \in \{\mathbf{A}, \mathbf{E}\}$, then $\sigma \not\models \mathbf{A}\Phi[l]$.
- (v) If $\sigma \not\models \mathbf{A}\{\Phi, \varphi\mathbf{R}\psi\}[l]$, then one of the following hold:
 - a) there is an R-signature ι' of $\mathbf{A}\{\Phi, \psi\}$ which agrees with ι on all top-level R- eventualities of $\wedge \Phi$, and such that $\sigma \not\models \mathbf{A}\{\Phi, \psi\}[\iota']$;
 - b) there is an R-signature ι' of $\mathbf{A}\{\Phi, \varphi, \mathbf{X}(\varphi\mathbf{R}\psi)\}$ which agrees with ι on all top-level R- eventualities of $\wedge \Phi$ and $\mathbf{X}(\varphi\mathbf{R}\psi)$, and such that $\sigma \not\models \mathbf{A}\{\Phi, \varphi, \mathbf{X}(\varphi\mathbf{R}\psi)\}[\iota']$.
- (vi) If $\sigma \not\models \mathbf{A}\{\Phi, \varphi\mathbf{U}\psi\}[l]$, then one of the following hold:
 - a) there is an R-signature ι' of $\mathbf{A}\{\Phi, \varphi, \psi\}$ which agrees with ι on all top-level R- eventualities of $\wedge \Phi$, and such that $\sigma \not\models \mathbf{A}\{\Phi, \varphi, \psi\}[\iota']$;
 - b) there is an R-signature ι' of $\mathbf{A}\{\Phi, \psi, \mathbf{X}(\varphi\mathbf{U}\psi)\}$ which agrees with ι on all top-level R- eventualities of $\wedge \Phi$ and $\mathbf{X}(\varphi\mathbf{U}\psi)$, and such that $\sigma \not\models \mathbf{A}\{\Phi, \psi, \mathbf{X}(\varphi\mathbf{U}\psi)\}[\iota']$.
- (vii) If $\sigma \not\models \mathbf{A}\mathbf{X}\Phi[l]$, then there is a $\sigma' \sim \sigma$ such that $(\sigma', 1) \not\models \mathbf{A}\Phi[l^-]$.

Proof. The proofs of (i)–(iv) are straightforward, and (v), (vi) follow easily from Proposition 3.2.21(i).

For (vii), assume that $\sigma \not\models \mathbf{A}\mathbf{X}\Phi[l]$, say with $\Phi = \{\varphi_1, \dots, \varphi_n\}$. Then, $\sigma \not\models \mathbf{A}\{\mathbf{X}\varphi_1[l], \dots, \mathbf{X}\varphi_n[l]\}$, i.e., $\sigma \not\models \mathbf{A}\{\mathbf{X}(\varphi_1[l^-]), \dots, \mathbf{X}(\varphi_n[l^-])\}$. Let $\sigma' \sim \sigma$ be such that $\sigma' \not\models \bigvee_{1 \leq k \leq n} \mathbf{X}(\varphi_k[l^-])$. We then have $(\sigma', 1) \not\models \bigvee_{1 \leq k \leq n} \varphi_k[l^-]$, whence $(\sigma', 1) \not\models \mathbf{A}\Phi[l^-]$. \blacksquare

We are now ready to prove that the system CTL_∞^* is sound.

3.2.23. PROPOSITION (Soundness of CTL_∞^*). *For every formula φ , if $\text{CTL}_\infty^* \vdash \varphi$, then $\text{CTL}^* \models \varphi$.*

Proof. Let \mathcal{T} be a CTL_∞^* proof of φ . Towards a contradiction, suppose that φ is not valid. Let \mathcal{M} be a serial model and σ a path on \mathcal{M} such that $\sigma \not\models \mathbf{A}\varphi$. We inductively find an infinite branch $\pi = (u_i)_{i < \omega}$ of \mathcal{T} and paths $(\sigma_i)_{i < \omega}$ on \mathcal{M} such that we have $\sigma_i \not\models \Gamma_i$ for every $i < \omega$, where Γ_i is the label of u_i . Propositions 3.2.21 and 3.2.22 associate to each sequent $\mathbf{A}\Phi \in \Gamma_i$ an R-signature ι such that $\sigma_i \not\models \mathbf{A}\Phi[\iota]$. Choices at $\{\mathbf{A}\wedge, \mathbf{A}\mathbf{U}, \mathbf{A}\mathbf{R}\}$ -vertices are resolved by Proposition 3.2.22. Choices at $\{\mathbf{E}\mathbf{L}\mathbf{i}\mathbf{t}, \mathbf{E}\mathbf{A}, \mathbf{E}\mathbf{E}\}$ -vertices are resolved by a priority mechanism. Moreover, either $\sigma_{i+1} = \sigma_i$, if u_i is not modal, or else $\sigma_{i+1} = (\sigma', 1)$ for some $\sigma' \sim \sigma_i$.

Let u_0 be the root of \mathcal{T} and $\sigma_0 := \sigma$. We assign to $A\varphi$ a signature ι such that $\sigma \not\models A\varphi[\iota]$ given by Proposition 3.2.21. We now proceed inductively from u_i . If no rule is applied at u_i , then u_i is an axiomatic leaf because \mathcal{T} is a proof. But then $\sigma_i \models \Gamma_i$, contradicting the inductive hypothesis $\sigma_i \not\models \Gamma_i$. So this case is not possible and π will be an infinite branch. Assume now that u_i is not a leaf. We distinguish cases according to the rule R applied at u_i :

- Suppose that $R \in \{\text{ALit}, \text{AV}, \text{AA}, \text{AE}, \text{E}\wedge, \text{EV}, \text{EU}, \text{ER}\}$. Let u_{i+1} be the unique immediate successor of u_i , and $\sigma_{i+1} := \sigma_i$. Signatures are assigned to A-sequents in Γ_{i+1} according to Proposition 3.2.22, and also Proposition 3.2.21 in the case of AA.
- Suppose that $R \in \{\text{ELit}, \text{EA}, \text{EE}\}$. By Proposition 3.2.4, there is at least one immediate successor v of u_i such that $\sigma_i \not\models \Gamma$, where Γ is the label of v . We pick as u_{i+1} the immediate successor with simpler label whenever possible, i.e., containing the active literal sequent in the case of ELit, the active A-sequent in the case of EA, and the active sequent $E\{\psi\}$ in the case of EE with principal formula $E\psi$. We let $\sigma_{i+1} := \sigma_i$. In the case of EA we assign to the new A-sequent in the premise a signature given by Proposition 3.2.21.
- Suppose that $R \in \{\text{iW}, \text{eW}\}$. We let u_{i+1} be the unique immediate successor of u_i , and $\sigma_{i+1} := \sigma_i$. We carry to Γ_{i+1} the signatures given by the inductive hypothesis.
- Suppose that $R = A\wedge$, say:

$$A\wedge \frac{A\{\Phi, \varphi_1\}, \Delta \quad A\{\Phi, \varphi_2\}, \Delta}{A\{\Phi, \varphi_1 \wedge \varphi_2\}, \Delta}$$

By the inductive hypothesis, an R-signature ι is assigned to $A\{\Phi, \varphi_1 \wedge \varphi_2\}$ such that $\sigma_i \not\models A\{\Phi, \varphi_1 \wedge \varphi_2\}[\iota]$. By Proposition 3.2.22, there is some $j \in \{1, 2\}$ such that $\sigma_i \not\models A\{\Phi, \varphi_j\}[\iota]$. We let $\sigma_{i+1} := \sigma_i$ and u_{i+1} be the immediate successor of u_i for premise $A\{\Phi, \varphi_j\}, \Delta$. We assign signature ι to $A\{\Phi, \varphi_j\}$.

- Suppose that $R = \text{AU}$, say:

$$\text{AU} \frac{A\{\Phi, \varphi, \psi\}, \Delta \quad A\{\Phi, \psi, X(\varphi \cup \psi)\}, \Delta}{A\{\Phi, \varphi \cup \psi\}, \Delta}$$

By the inductive hypothesis, an R-signature ι is assigned to $A\{\Phi, \varphi_1 \cup \varphi_2\}$ such that $\sigma_i \not\models A\{\Phi, \varphi_1 \cup \varphi_2\}[\iota]$. By Proposition 3.2.22, one of the following is the case:

- a) there is an R-signature ι' of $\mathbf{A}\{\Phi, \varphi, \psi\}$ which agrees with ι on all top-level R-eventualities of $\wedge \Phi$, and such that $\sigma \not\models \mathbf{A}\{\Phi, \varphi, \psi\}[\iota']$;
- b) there is an R-signature ι' of $\mathbf{A}\{\Phi, \psi, \mathbf{X}(\varphi \cup \psi)\}$ which agrees with ι on all top-level R-eventualities of $\wedge \Phi$ and $\mathbf{X}(\varphi \cup \psi)$, and such that $\sigma \not\models \mathbf{A}\{\Phi, \psi, \mathbf{X}(\varphi \cup \psi)\}[\iota']$.

If (a) is the case, we let u_{i+1} be the immediate successor of u_i for the premise $\mathbf{A}\{\Phi, \varphi, \psi\}, \Delta$; otherwise, the one for premise $\mathbf{A}\{\Phi, \psi, \mathbf{X}(\varphi \cup \psi)\}, \Delta$. In either case we let $\sigma_{i+1} := \sigma_i$ and assign ι' to the selected A-sequent.

- The case $\mathbf{R} = \mathbf{AR}$ is analogous to the previous one.
- Suppose that $\mathbf{R} = \mathbf{EX}$, say:

$$\mathbf{EX} \frac{\mathbf{E}\Psi_1, \dots, \mathbf{E}\Psi_m}{\mathbf{EX}\Psi_1, \dots, \mathbf{EX}\Psi_m, \Delta}$$

By the inductive hypothesis, $\sigma_i \not\models \Gamma_i$. It is then easy to see that $(\sigma_i, 1) \not\models \mathbf{E}\Psi_j$ for any $1 \leq j \leq m$, so we let u_{i+1} be the unique immediate successor of u_i and $\sigma_{i+1} := (\sigma_i, 1)$.

- Finally, suppose that $\mathbf{R} = \mathbf{AX}$, say:

$$\mathbf{AX} \frac{\mathbf{A}\Phi, \mathbf{E}\Psi_1, \dots, \mathbf{E}\Psi_m}{\mathbf{A}(\mathbf{X}\Phi \cup \Xi), \mathbf{EX}\Psi_1, \dots, \mathbf{EX}\Psi_m, \Delta}$$

By the inductive hypothesis, $\sigma_i \not\models \Gamma_i$. In particular, $\sigma_i \not\models \mathbf{AX}\Phi[\iota]$, where ι is the signature assigned to $\mathbf{A}(\mathbf{X}\Phi \cup \Xi)$ by the inductive hypothesis. By Proposition 3.2.22, there is a $\sigma' \sim \sigma_i$ such that $(\sigma', 1) \not\models \mathbf{A}\Phi[\iota^-]$. Arguing as in the proof of Proposition 3.2.4, it is easy to see that $(\sigma', 1) \not\models \mathbf{E}\Psi_k$ for any $1 \leq k \leq m$. We then let u_{i+1} be the unique immediate successor of u_i , and $\sigma_{i+1} := (\sigma', 1)$. To $\mathbf{A}\Phi$ we assign signature ι^- .

We now derive a contradiction from the existence of π . Since \mathcal{T} is a proof, π contains a good infinite sequent trace.

Case 1: π contains an infinite A-trace $\tau = (\mathbf{Q}_i\Phi_i)_{i < \omega}$ such that there is an infinite R-trace $\rho = (\varphi_i)_{i < \omega}$ on τ . For every $j < \omega$ with $\mathbf{Q}_j = \mathbf{A}$, let ι_j be the R-signature associated to $\mathbf{Q}_j\Phi_j$. Let $\psi = \psi_1\mathbf{R}\psi_2$ be the dominating formula in ρ , and let $N < \omega$ be such that $\varphi_N = \psi$, $\varphi_{N+1} = \mathbf{X}\psi$, and $\varphi_i \in \{\psi, \mathbf{X}\psi\}$ for all $i \geq N$. Note that rule \mathbf{EX} is not applied anywhere on $(u_i)_{i \geq N}$ because $\mathbf{Q}_i = \mathbf{A}$ for every $i \geq N$. Let $i_0 \geq N$ be least such that u_{i_0} is modal. By construction, $\iota_{i_0}(\psi) = \iota_N(\psi) > 0$ because at every encounter of rule \mathbf{AR} in $(u_i)_{i \geq N}$ with ψ principal we always take the right branch. Therefore, $\iota_{i_0+1}(\psi) = \iota_{i_0}^-(\psi) < \iota_{i_0}(\psi)$. Let now $i' > i_0$ be least

such that rule AR is applied at $u_{i'}$ with ψ active, and $i_1 > i'$ least such that u_{i_1} is modal. As before, $\iota_{i_1}(\psi) > 0$ and we have:

$$\iota_{i_1+1}(\psi) = \iota_{i_1}^-(\psi) = \iota_{i'}^-(\psi) = \iota_{i_0+1}^-(\psi) = \iota_{i_0+1}(\psi) - 1 < \iota_{i_0+1}(\psi).$$

Repeating this argument yields an infinite descending chain of natural numbers.

Case 2: π contains an infinite E-trace $\tau = (\mathbf{Q}_i\Phi_i)_{i<\omega}$ containing no infinite U-trace. Let then $\rho = (\varphi_i)_{i<\omega}$ be an infinite R-trace on τ given by König's lemma. Let $\psi = \psi_1\mathbf{R}\psi_2$ be the dominating formula in ρ , and let $N < \omega$ be such that $\varphi_N = \psi$, $\varphi_{N+1} = \mathbf{X}\psi$, and $\varphi_i \in \{\psi, \mathbf{X}\psi\}$ for all $i \geq N$. Note that no context extraction is ever encountered in $(u_i)_{i \geq N}$ because $(\varphi_i)_{i \geq N}$ contains only ψ and $\mathbf{X}\psi$.

For every $i \geq N$, let $h(i) \geq i$ be least such that $u_{h(i)}$ is modal. By construction, we have $\sigma_{h(i)+1} = (\sigma', 1)$ for some $\sigma' \sim \sigma_i$, and thus $\sigma_{h(i)+1}(0)$ is an immediate successor of $\sigma_i(0)$. Let then σ^* be the path $\sigma_N(0), \sigma_{h(N)+1}(0), \sigma_{h(h(N)+1)+1}(0), \dots$

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

$$\underbrace{\psi \triangleright \mathbf{X}\psi \triangleright \dots \triangleright \mathbf{X}\psi}_{0} \triangleright \underbrace{\psi \triangleright \dots \triangleright \psi \triangleright \mathbf{X}\psi \triangleright \dots \triangleright \mathbf{X}\psi}_{1} \triangleright \psi \triangleright \dots$$

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

Suppose that $\chi = \ell$ for some literal ℓ . Let $j \geq i$ be least such that ℓ is principal in u_j . Note that there must be such a j because we never encounter a context extraction on $(\varphi_k)_{k \geq N}$ and every sequent in a modal vertex is active. As all principal sequents in a modal vertex which have a \triangleright -successor are of the form $\mathbf{Q}\mathbf{X}\Psi$, we have $i \leq j \leq h(i)$ and thus $\sigma_j = \sigma_i$.⁵ Moreover, $j < h(i)$ because otherwise ρ would die out at $h(i)$. Then, at u_j we have:

$$\text{ELit} \frac{\mathbf{E}\Psi, \Delta \quad \mathbf{E}\ell, \Delta}{\mathbf{E}\{\Psi, \ell\}, \Delta}$$

Since in the construction of π we prioritised the right branch but did not take it, we have $\sigma_j \models \mathbf{E}\ell$, whence $(\sigma^*, g(i)) \models \mathbf{E}\ell$ because $(\sigma^*, g(i)) = (\sigma^*, g(j)) \sim \sigma_j$.

Suppose that $\chi = \chi_1 \vee \chi_2$. Let $j \geq i$ be least such that χ is principal in u_j . As before, such a j exists and $i \leq j < h(i)$, so $\sigma_j = \sigma_i$. Let k be such

⁵This observation will be used repeatedly in what follows.

that $\chi \triangleright_{u_j} \chi_k \in \Phi_{j+1}$. By the inductive hypothesis, $(\sigma^*, g(j+1)) \models \chi_k$. And $g(j+1) = g(j) = g(i)$ because u_j is not modal, so $(\sigma^*, g(i)) \models \chi$.

The case $\chi = \chi_1 \wedge \chi_2$ is analogous to the previous one.

Suppose that $\chi = \text{E}\chi'$. Let $j \geq i$ be least such that χ is principal in u_j . As before, such a j exists and $i \leq j < h(i)$, so $\sigma_j = \sigma_i$. Then, at u_j we have:

$$\text{EE} \frac{\text{E}\Psi, \Delta \quad \text{E}\{\chi'\}, \Delta}{\text{E}\{\Psi, \text{E}\chi'\}, \Delta}$$

Since we prioritised the right branch but did not take it, we have $\sigma_i \models \text{E}\chi'$, so there is some $\sigma' \sim \sigma_i$ such that $\sigma' \models \chi'$. Then, since $\sigma_i \sim (\sigma^*, g(i))$, we have $\sigma' \sim (\sigma^*, g(i))$ and thus $(\sigma^*, g(i)) \models \text{E}\chi'$.

The case $\chi = \text{A}\chi'$ is analogous to the previous one.

Suppose that $\chi = \text{X}\chi'$. Let $j \geq i$ be least such that χ is principal in u_j . As before, such a j exists but this time $j = h(i)$. We have $\Phi_{h(i)} \ni \chi \triangleright \chi' \in \Phi_{h(i)+1}$, so by the inductive hypothesis $(\sigma^*, g(h(i)+1)) \models \chi'$. So, since $g(h(i)+1) = g(h(i))+1$, $(\sigma^*, g(h(i))) \models \text{X}\chi'$. And $g(h(i)) = g(i)$, so $(\sigma^*, g(i)) \models \chi$.

Suppose that $\chi = \chi_1 \text{U}\chi_2$. We find an $n < \omega$ such that $(\sigma^*, g(i) + n) \models \chi_2$ and $(\sigma^*, g(i) + m) \models \chi_1$ for all $m < n$. By assumption, τ contains no infinite U -trace. This, together with the facts that we never encounter a context extraction on $(\varphi_k)_{k \geq N}$, that every sequent in a modal vertex is principal, and that ρ is infinite, implies that there is a least $j \geq i$ such that χ is principal in u_j and $\chi \not\triangleright_{u_j} \text{X}\chi$. So $\chi \triangleright_{u_j} \chi_2 \in \Phi_{j+1}$. Let n be such that $g(j+1) = g(i) + n$. 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 a modal vertex in $[u_k, u_j)\tau$. So, by the minimality of j , there is a $k \leq k' < h(k)$ such that $\Phi_{k'} \ni \chi \triangleright \chi_1, \text{X}\chi \in \Phi_{k'+1}$, whence $(\sigma^*, g(k'+1)) \models \chi_1$ by the inductive hypothesis. And $g(k'+1) = g(k') = g(k) = g(i) + m$, so $(\sigma^*, g(i) + m) \models \chi_1$.

Finally, suppose that $\chi = \chi_1 \text{R}\chi_2$. We show that either $(\sigma^*, g(i) + n) \models \chi_2$ for every $n < \omega$, or else there is some $n < \omega$ such that $(\sigma^*, g(i) + n) \models \chi_1$ and $(\sigma^*, g(i) + m) \models \chi_2$ for every $m \leq n$. Let $j \geq i$ be least such that χ is principal in u_j . As before, such a j exists and $i \leq j < h(i)$. Since $\chi_2 \in \Phi_{j+1}$, by the inductive hypothesis $(\sigma^*, g(j+1)) \models \chi_2$. And $g(j+1) = g(j) = g(i)$, so $(\sigma^*, g(i)) \models \chi_2$. If $\chi_1 \in \Phi_{j+1}$, then by the inductive hypothesis $(\sigma^*, g(i)) \models \chi_1$, whence $(\sigma^*, g(i)) \models \chi$ and we are done. If $\chi_1 \notin \Phi_{j+1}$, then $\text{X}\chi \in \Phi_{j+1}$. Since we never encounter a context extraction on $(\varphi_k)_{k \geq N}$, we know that $\text{X}\chi \in \Phi_{h(j+1)}$, so $\chi \in \Phi_{h(j+1)+1}$. Let $k \geq h(j+1) + 1$ be least such that χ is principal in u_k . As before, such a k exists and $h(j+1) + 1 \leq k < h(h(j+1) + 1)$. By the inductive hypothesis, $(\sigma^*, g(k+1)) \models \chi_2$. And

$$g(k+1) = g(k) = g(h(j+1) + 1) = g(h(j+1)) + 1 = g(h(j)) + 1 = g(h(i)) + 1 = g(i) + 1,$$

so $(\sigma^*, g(i) + 1) \models \chi_2$. If $\chi_1 \in \Phi_{k+1}$, then $(\sigma^*, g(k+1)) \models \chi_1$ by the inductive hypothesis and thus $(\sigma^*, g(i) + 1) \models \chi_1$, whence $(\sigma^*, g(i)) \models \chi$ and we are done. Otherwise, $X_\chi \in \Phi_{k+1}$. Repeating this argument yields $(\sigma^*, g(i)) \models \chi$. ■

3.2.2 Completeness of CTL $^*_\infty$

We now turn to the proof of the completeness of CTL $^*_\infty$. We use game- and automata-theoretic techniques and appeal to the soundness and completeness result for satisfiability of CTL* formulas in [47]. Although the approach that we follow is standard (see, e.g., [102]), the correctness condition for E-sequents creates some unusual complications, as in [47]. All topological, game-theoretic and automata-theoretic preliminaries needed for the completeness proof can be found in Sections 1.4 to 1.6.

A *proof-search guide* for a formula φ is a pair (\leq_S, \leq_F) , where \leq_S is a well-ordering of Seq_φ and \leq_F a well-ordering of $\text{Clos}(\varphi)$.

3.2.24. DEFINITION (CTL $^*_\infty$ proof-search tree). Let φ be a CTL* formula, and $\mathcal{G} = (\leq_S, \leq_F)$ a proof-search guide for φ . A *CTL $^*_\infty$ proof-search tree for φ guided by \mathcal{G}* is a finite or infinite labelled tree $\mathcal{T} = (T, <_T, \lambda_T)$ built according to the rules of CTL $^*_\infty$ in such a way that the following hold:

- (i) The root of \mathcal{T} has label $A\varphi$.
- (ii) A vertex of \mathcal{T} is a leaf if, and only if, it is either axiomatic or a cul-de-sac.
- (iii) The weakening rules iW and eW are not applied anywhere in \mathcal{T} .
- (iv) All instances in \mathcal{T} of the rules $ALit$, $ELit$, AV , EV , $A\wedge$, $E\wedge$, AA , EA , AE , EE , AU , AR , EU and ER are discarding.
- (v) For every non-final vertex $u \in T$, say with label Γ , the following conditions hold:
 - a) If there is a sequent in Γ containing a non-literal, non-modal formula, then a CTL $^*_\infty$ rule R other than AX , EX , iW and eW is applied at u with principal sequent $Q\Psi$ and principal formula $\psi \in \Psi$, where $(Q\Psi, \psi)$ is the \mathcal{G} -least pair such that ψ is non-literal and non-modal, $\psi \in \Psi$, and $Q\Psi \in \Gamma$.
 - b) Else, if there is a sequent in Γ containing a literal and some other formula, then a rule among $ALit$ and $ELit$ is applied at u with principal sequent $Q\Psi$ and principal literal ℓ , where $(Q\Psi, \ell)$ is the \mathcal{G} -least pair such that $Q\Psi \in \Gamma$, $\ell \in \Psi$, and $\Psi \setminus \{\ell\} \neq \emptyset$.

c) Else, if there is no A-sequent in Γ , then Γ is of the form

$$\text{EX}\Psi_1, \dots, \text{EX}\Psi_m, \Lambda,$$

where Λ is a (non-axiomatic) set of literal sequents, and rule EX is applied at u with premise $\text{E}\Psi_1, \dots, \text{E}\Psi_m$.

d) In any other case Γ has the form

$$\text{AX}\Phi_1, \dots, \text{AX}\Phi_n, \text{EX}\Psi_1, \dots, \text{EX}\Psi_m, \Lambda,$$

where $n > 0$ and Λ is a (non-axiomatic) set of literal sequents, and the following branching version of rule AX is applied at u :

$$\text{AX}_b \frac{\text{A}\Phi_1, \text{E}\Psi_1, \dots, \text{E}\Psi_m \quad \cdots \quad \text{A}\Phi_n, \text{E}\Psi_1, \dots, \text{E}\Psi_m}{\text{AX}\Phi_1, \dots, \text{AX}\Phi_n, \text{EX}\Psi_1, \dots, \text{EX}\Psi_m, \Lambda}, \quad \dashv$$

A vertex u of a CTL_∞^* proof-search tree is *modal* if rule AX_b or EX is applied at u .

3.2.25. PROPOSITION. *Every infinite branch of a CTL_∞^* proof-search tree \mathcal{T} contains infinitely many modal vertices.*

Proof. Let $\pi = (u_n)_{n < \omega}$ be an infinite branch of \mathcal{T} . For every $n < \omega$, let Γ_n be the hypersequent labelling u_n . For every sequent $\text{Q}\Psi$, we define:

$$\langle \text{Q}\Psi \rangle^- := \sum \{ \langle \psi \rangle \mid \psi \in \Psi \text{ and } \psi \text{ is not an X-formula} \}.$$

Additionally, we let $\langle \Gamma_n \rangle^- := \sum \{ \langle \text{Q}\Psi \rangle^- \mid \text{Q}\Psi \in \Gamma \}$ for every $n < \omega$. By Definition 3.2.24(iv) and the fact that rules ALit and ELit are never applied degenerately on \mathcal{T} , it is straightforward to see that, if u_n is not modal, then $\langle \Gamma_{n+1} \rangle^- < \langle \Gamma_n \rangle^-$. Therefore, π must pass through infinitely many modal vertices. \blacksquare

Observe that Definition 3.2.24 describes a *deterministic* algorithm for building a proof-search tree for φ guided by \mathcal{G} , in the sense that the rule applied at a vertex u is uniquely determined by the label of u . Thus, we have:

3.2.26. PROPOSITION. *Let \mathcal{T} and \mathcal{T}' be CTL_∞^* proof-search trees for φ guided by \mathcal{G} . For any vertices $u \in \mathcal{T}$ and $v \in \mathcal{T}'$, if u and v have identical labels, then the labelled cones $[u, \rightarrow]_{\mathcal{T}}$ and $[v, \rightarrow]_{\mathcal{T}'}$ are isomorphic.*

3.2.27. COROLLARY. *For every formula φ and every proof-search guide \mathcal{G} for φ , there is exactly one CTL_∞^* proof-search tree for φ guided by \mathcal{G} up to isomorphism.*

Proof. Uniqueness follows from Proposition 3.2.26, and existence from the fact that every $\Gamma \in \text{HSeq}_\varphi$ falls under one (and only one) of the cases described in points (ii) and (v) of Definition 3.2.24. ■

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 φ^∂ follows.

3.2.28. DEFINITION (CTL_∞^* refutation). A CTL_∞^* *refutation* of a formula φ is a subtree \mathcal{T}' of a CTL_∞^* proof-search tree \mathcal{T} for φ satisfying:

- (i) The root of \mathcal{T}' is the root of \mathcal{T} .
- (ii) Every leaf of \mathcal{T}' is a cul-de-sac of \mathcal{T} .
- (iii) If a vertex $u \in \mathcal{T}'$ is labelled in \mathcal{T} by the conclusion of an instance of a non-axiomatic rule other than AX_b , then u has exactly one immediate successor in \mathcal{T}' .
- (iv) If a vertex $u \in \mathcal{T}'$ is labelled in \mathcal{T} by the conclusion of an instance of rule AX_b , then \mathcal{T}' contains every immediate successor of u in \mathcal{T} .
- (v) On every infinite sequent trace τ on an infinite path through \mathcal{T}' :
 - a) if τ is of type **A**, then every infinite formula trace on τ is of type **U**;
 - b) if τ is of type **E**, then some infinite formula trace on τ is of type **U**. ⊥

The term *refutation*, which we borrow from [102], is justified by the following proposition, dual to the soundness and completeness result in [47]:

3.2.29. PROPOSITION ([47, Thm. 10]). *A formula φ is valid if, and only if, there is no CTL_∞^* 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, established by automata-theoretic means, yields completeness of CTL_∞^* .

Fix a CTL_∞^* proof-search tree $\mathcal{T} = (T, <_T, \lambda_T)$ for a formula φ guided by $\mathcal{G} = (\leq_S, \leq_F)$. For every $l \in \text{Leaf}(\mathcal{T})$ and $n < \omega$, we let $\odot_l^n := (l, n)$. Since all plays of a Gale–Stewart game must be infinite, we extend \mathcal{T} to a tree $\mathcal{T}_\odot := (T_\odot, <_\odot, \lambda_\odot)$ with no finite branches by setting:

- $T_\odot := T \cup \{\odot_l^n \mid l \in \text{Leaf}(\mathcal{T}) \text{ and } n < \omega\}$;
- $<_\odot := <_T \cup \{(l, \odot_l^0) \mid l \in \text{Leaf}(\mathcal{T})\} \cup \{(\odot_l^n, \odot_l^{n+1}) \mid l \in \text{Leaf}(\mathcal{T}) \text{ and } n < \omega\}$;

- $\lambda_\odot := \lambda_T \cup \{(\odot_l^n, \emptyset) \mid l \in \text{Leaf}(\mathcal{T}) \text{ and } n < \omega\}$.

We let $h : T_\odot \rightarrow T$ be given by setting $h(u) := u$ for every $u \in T \subseteq T_\odot$ and $h(\odot_l^n) := l$ for every $l \in \text{Leaf}(\mathcal{T})$ and $n < \omega$. Clearly, h sends branches of T_\odot to branches of \mathcal{T} . A branch π of T_\odot is *good* (*bad*) if $h(\pi)$ is good (respectively, bad).⁶

The *proof-search game based on \mathcal{T}* is the Gale–Stewart game

$$\mathcal{G}_\mathcal{T} := (r, T_{\text{I}}, T_{\text{II}}, T_\odot, f, \Pi_\omega, W_{\text{I}}, W_{\text{II}}),$$

where:

- (i) r is the root of T_\odot ;
- (ii) T_{II} is the collection of all vertices of T_\odot which are branching and non-modal;
- (iii) $T_{\text{I}} := T_\odot \setminus T_{\text{II}}$;
- (iv) $f : T_\odot \setminus \{r\} \rightarrow T_\odot$ is given by letting $f(u)$ be the unique immediate predecessor of u in T_\odot , for every $u \in T_\odot \setminus \{r\}$;
- (v) Π_ω is the collection of all branches of T_\odot ;
- (vi) W_{I} is the collection of all good branches of T_\odot ;
- (vii) $W_{\text{II}} := \Pi_\omega \setminus W_{\text{I}}$.

We refer to player I (player II) as **Prov** (respectively, **Ref**), to emphasise the fact that player I (player II) is looking for a proof of φ (respectively, a refutation thereof).

We define the *ancestor* function $a : T_\odot \rightarrow T_\odot^{<\omega}$ by letting $a(u)$ be the unique rooted path on T_\odot with last vertex u , for every $u \in T_\odot$. Note that every $a(u)$ is a partial play of $\mathcal{G}_\mathcal{T}$.

Let σ be a $\mathcal{G}_\mathcal{T}$ -strategy for any player. A vertex $u \in T_\odot$ is σ -*visible* if the partial play $a(u)$ is σ -consistent.

As expected, winning strategies for **Prov** correspond to proofs:

3.2.30. PROPOSITION. *There exists a winning $\mathcal{G}_\mathcal{T}$ -strategy for **Prov** if, and only if, \mathcal{T} contains a subtree which is a CTL_∞^* proof of φ .*

Proof. For the left-to-right direction, let $\sigma : T_{\text{I}} \rightarrow T_\odot$ be a winning $\mathcal{G}_\mathcal{T}$ -strategy for **Prov**. Let \mathcal{T}' be the result after removing from T_\odot all vertices which are either not σ -visible, or of the form \odot_l^n for some $l \in \text{Leaf}(\mathcal{T})$ and $n < \omega$. We claim that \mathcal{T}' is a CTL_∞^* proof of φ .

Clearly, the root of T_\odot , which is also that of \mathcal{T} , is σ -visible, so \mathcal{T} and \mathcal{T}' have the same root. And if a vertex $u \in T_\odot$ is σ -visible, then so is any $v \in a(u)$. This shows that \mathcal{T}' is a subtree of \mathcal{T} .

⁶Recall (see Section 1.1) that we write $h(\pi)$ for $(h(\pi(n)))_{n < \omega}$.

Let $u \in T_{\odot}$ be a σ -visible vertex. If u is not branching, then the unique immediate successor of u is also σ -visible. If u is branching and non-modal, then the two immediate successors of u are both σ -visible because $u \in T_{\text{IT}}$. Finally, if u is branching and modal, then $u \in T_{\text{I}}$ and there is exactly one immediate successor of u which is σ -visible, namely, $\sigma(u)$. This establishes that \mathcal{T}' is a CTL_{∞}^* derivation of φ .

Let $l \in \text{Leaf}(\mathcal{T}')$, and let $\pi := a(l) \frown (\odot_l^n)_{n < \omega}$. As l is σ -visible, $\pi \in \Pi_{\omega}$ is a σ -consistent play of $\mathcal{G}_{\mathcal{T}}$, so $\pi \in W_{\text{I}}$ because σ is winning. Hence, $h[\pi]$ is a good finite branch of \mathcal{T} and thus l is axiomatic.

Finally, let $\pi = (u_n)_{n < \omega}$ be an infinite branch of \mathcal{T}' . Since every u_n is σ -visible, π is a σ -consistent play of $\mathcal{G}_{\mathcal{T}}$ and thus $\pi \in W_{\text{I}}$. Hence, $h[\pi]$ is a good infinite branch of \mathcal{T} . And no vertex on π is of the form \odot_l^k , whence $\pi = h[\pi]$ is a good branch of \mathcal{T}' . We have thus shown that \mathcal{T}' is a CTL_{∞}^* proof of φ .

For the converse direction, let \mathcal{T}' be a subtree of \mathcal{T} such that \mathcal{T}' is a CTL_{∞}^* proof of φ . Define $\sigma : T_{\text{I}} \rightarrow T_{\odot}$ as follows, for every $u \in T_{\text{I}}$. If u is not branching, let $\sigma(u)$ be the unique immediate successor of u in T_{\odot} . Suppose that u is branching. Then, u is modal because $u \in T_{\text{I}}$. If $u \in \mathcal{T}'$, we let $\sigma(u)$ be the unique immediate successor of u in \mathcal{T}' , and otherwise we let $\sigma(u)$ be any immediate successor of u . Clearly, σ is a $\mathcal{G}_{\mathcal{T}}$ -strategy for **Prov**. We claim that σ is winning.

Let $\pi = (u_n)_{n < \omega} \in \Pi_{\omega}$ be a σ -consistent play of $\mathcal{G}_{\mathcal{T}}$. By induction on $n < \omega$, we show that $h(u_n) \in \mathcal{T}'$ for every $n < \omega$, whence it follows that $h[\pi]$ is a good branch of \mathcal{T} and thus $\pi \in W_{\text{I}}$. The base case $n = 0$ is clear because u_0 is the root of T_{\odot} and \mathcal{T} , thus also the root of \mathcal{T}' . For the inductive case, assume that $h(u_n) \in \mathcal{T}'$. If $h(u_{n+1}) = h(u_n)$ we are done, so suppose otherwise. Then, u_n is a non-final vertex of \mathcal{T} and thus $h(u_n) = u_n$ and $h(u_{n+1}) = u_{n+1}$. In particular, u_n is not axiomatic by Definition 3.2.24(ii), so u_n is also non-final in \mathcal{T}' . If u_n is not branching in \mathcal{T} , then u_{n+1} is the unique immediate successor of u_n in \mathcal{T} and thus $u_{n+1} \in \mathcal{T}'$ because u_n is non-final in \mathcal{T}' . Suppose now that u_n is branching in \mathcal{T} . If u_n is non-modal, then u_{n+1} is one of the two immediate successors of u_n in \mathcal{T} and we have $u_{n+1} \in \mathcal{T}'$ because \mathcal{T}' is a CTL_{∞}^* derivation and u_n is non-final in \mathcal{T}' . Finally, if u_n is modal the σ -consistency of π yields $u_{n+1} = \sigma(u_n) \in \mathcal{T}'$ and we are done. \blacksquare

A similar argument shows that that winning strategies for **Ref** correspond to refutations:

3.2.31. PROPOSITION. *There exists a winning $\mathcal{G}_{\mathcal{T}}$ -strategy for **Ref** if, and only if, \mathcal{T} contains a subtree which is a CTL_{∞}^* refutation of φ .*

By Proposition 3.2.26, if $u, v \in T$ are labelled by the same hypersequent, then for every immediate successor u' of u there exists an immediate successor v' of v

such that u' and v' have identical labels and, moreover, $\triangleright_u^{u'} = \triangleright_v^{v'}$ and $\triangleright_u^{u'} = \triangleright_v^{v'}$. Hence, for hypersequents $\Gamma, \Gamma' \in \text{HSeq}_\varphi$ we write $\Gamma <_\odot^0 \Gamma'$ if there are $u, v \in T$ labelled by Γ and Γ' , respectively, and such that $u <_\odot^0 v$. In this case, we let $\triangleright_\Gamma^{\Gamma'} := \triangleright_u^v$ and $\triangleright_\Gamma^{\Gamma'} := \triangleright_u^v$.

We endow Π_ω with the topology \mathcal{O} generated by taking as a base the collection $\{B_p \mid p \in T_\omega^{<\omega}\}$, where $B_p := \{\pi \in \Pi_\omega \mid p \sqsubseteq \pi\}$. It is easy to see that $(\Pi_\omega, \mathcal{O})$ is a subspace of $(T_\omega^\omega, T_\square)$, where T_\square is the prefix topology on T_ω^ω .

Our aim is to show that W_I is Borel in $(\Pi_\omega, \mathcal{O})$, for then determinacy of \mathcal{G}_T follows. We do so by building a non-deterministic Büchi automaton (NBA) \mathcal{A} which recognises the good branches of \mathcal{T}_\odot . Observe that \mathcal{A} cannot take as inputs ω -sequences of vertices of \mathcal{T}_\odot because automata require a finite alphabet.⁷ We overcome this difficulty by building \mathcal{A} over the alphabet of hypersequents occurring in \mathcal{T}_\odot . Informally, \mathcal{G}_T and \mathcal{A} will not speak the same language: the game \mathcal{G}_T sees the vertices of \mathcal{T}_\odot , while the automaton \mathcal{A} will only have access to their labels. Topology will bridge this gap by means of a continuous map induced by λ_\odot .

Let $\Lambda_\odot \subseteq \text{HSeq}_\varphi$ be the collection of hypersequents labelling vertices of \mathcal{T}_\odot . We let \mathcal{O}' be the prefix topology on Λ_\odot^ω . We shall later see that the labelling map λ_\odot induces a continuous map from $(\Pi_\omega, \mathcal{O})$ to $(\Lambda_\odot^\omega, \mathcal{O}')$, and that W_I is the preimage of a Δ_3^0 set under said map.

Let $\Sigma := \{(\Gamma, \mathbf{Q}\Psi) \in \Lambda_\odot \times \text{Seq}_\varphi \mid \mathbf{Q}\Psi \in \Gamma\}$. Note that Σ is finite. Finally, let $Q_0 := \{(\Gamma, \mathbf{Q}\Psi, \psi) \in \Lambda_\odot \times \text{Seq}_\varphi \times \text{Clos}(\varphi) \mid (\Gamma, \mathbf{Q}\Psi) \in \Sigma \text{ and } \psi \in \Psi\}$.

For all $(\Gamma, \mathbf{Q}\Psi, \psi), (\Gamma', \mathbf{Q}'\Psi', \psi') \in Q_0$, we write $(\Gamma, \mathbf{Q}\Psi, \psi) \rightsquigarrow (\Gamma', \mathbf{Q}'\Psi', \psi')$ if the following hold:

- (i) $\Gamma <_\odot^0 \Gamma'$;
- (ii) $\mathbf{Q}\Psi \triangleright_\Gamma^{\Gamma'} \mathbf{Q}'\Psi'$;
- (iii) $\mathbf{Q}\Psi, \psi \triangleright_\Gamma^{\Gamma'} \mathbf{Q}'\Psi', \psi'$.

We define an automaton skeleton $S_{\mathcal{A}} = (Q, \Sigma, \Delta)$ as follows. We let $Q := Q_0 \uplus \{q^*\}$, where q^* is a fresh symbol. And, for all $q, q' \in Q$ and $\sigma \in \Sigma$, we let $(q, \sigma, q') \in \Delta$ if, and only if, one of the following holds:

- (i) $q = q^*, \sigma = (\{\mathbf{A}\{\varphi\}\}, \mathbf{A}\{\varphi\})$, and $q' = (\{\mathbf{A}\{\varphi\}\}, \mathbf{A}\{\varphi\}, \varphi)$;
- (ii) $q \neq q^* \neq q', \sigma = (q'(0), q'(1))$, and $q \rightsquigarrow q'$.

Note that no transition leads to q^* , and that $(q, \sigma, q') \in \Delta$ implies $\sigma = (q'(0), q'(1))$.

We say that $(\Gamma_n)_{n < \omega} \in \Lambda_\odot^\omega$ is *branch-like* if there is an infinite branch $\pi = (u_n)_{n < \omega}$ of \mathcal{T} such that $\lambda_T(u_n) = \Gamma_n$ for every $n < \omega$. Similarly, given a word

⁷Natural generalisations of Büchi automata to countably infinite alphabets have been considered in the literature (e.g., [106]). However, it does not seem that we can use them without assuming the topological result that we seek to derive.

$w = ((\Gamma_n, \mathbf{Q}_n \Psi_n))_{n < \omega} \in \Sigma^\omega$, we say that w is *branch-like* if $(\Gamma_n)_{n < \omega}$ is branch-like and $(\mathbf{Q}_n \Psi_n)_{n < \omega}$ is an infinite sequent trace on π .

For every branch-like word $w = ((\Gamma_n, \mathbf{Q}_n \Psi_n))_{n < \omega} \in \Sigma^\omega$, we fix an arbitrary infinite branch $\pi_w = (u_n)_{n < \omega}$ of \mathcal{T} such that $\lambda_{\mathcal{T}}(u_n) = \text{Gamma}_n$ for every $n < \omega$, and let $\tau_w := (\mathbf{Q}_n \Psi_n)_{n < \omega}$. Note that, by Proposition 3.2.26, in place of π_w we could chose any other infinite branch of \mathcal{T} labelled by the same hypersequents as π_w .

As expected, runs of $S_{\mathcal{A}}$ correspond to sequent and formula traces on infinite branches of \mathcal{T} :

3.2.32. PROPOSITION. *For every $w \in \Sigma^\omega$, the following hold:*

- (i) *if w is branch-like, then $q^{\wedge}((\lambda_{\mathcal{T}}(\pi_w(n)), \tau_w(n), \rho(n)))_{n < \omega}$ is a run of $S_{\mathcal{A}}$ from q^* on w , for every infinite formula trace ρ on τ_w ;*
- (ii) *conversely, if there is a run $q^{\wedge}(q_n)_{n < \omega}$ of $S_{\mathcal{A}}$ from q^* on w , then w is branch-like and there exists an infinite formula trace ρ on τ_w such that $q_n = (\lambda_{\mathcal{T}}(\pi_w(n)), \tau_w(n), \rho(n))$ for every $n < \omega$.*

Proof. Fix a word $w = ((\Gamma_n, \mathbf{Q}_n \Psi_n))_{n < \omega} \in \Sigma^\omega$. For (i), assume that w is branch-like and let $(\psi_n)_{n < \omega}$ be an infinite formula trace on τ_w . For every $n < \omega$, let $q_n := (\lambda_{\mathcal{T}}(\pi_w(n)), \tau_w(n), \psi_n)$. We are to show:

- a) $(q^*, (\Gamma_0, \mathbf{Q}_0 \Psi_0), q_0) \in \Delta$;
- b) $(q_n, (\Gamma_{n+1}, \mathbf{Q}_{n+1} \Psi_{n+1}), q_{n+1}) \in \Delta$ for every $n < \omega$.

As $\pi_w(0)$ is the root of \mathcal{T} , we have $\Gamma_0 = \lambda_{\mathcal{T}}(\pi_w(0)) = \{\mathbf{A}\{\varphi\}\}$, $\mathbf{Q}_0 \Psi_0 = \tau_w(0) = \mathbf{A}\{\varphi\}$, and $\psi_0 = \varphi$, whence (a) holds. And for every $n < \omega$ we have $q_n = (\lambda_{\mathcal{T}}(\pi_w(n)), \tau_w(n), \psi_n) = (\Gamma_n, \mathbf{Q}_n \Psi_n, \psi_n)$, so, since π_w is a branch, τ_w a sequent trace on π_w , and $(\psi_n)_{n < \omega}$ a formula trace on τ_w , (b) follows.

For (ii), let $q^{\wedge}(q_n)_{n < \omega}$ be a run of $S_{\mathcal{A}}$ from q^* on w . Thus, every q_n is of the form $q_n = (\Gamma_n, \mathbf{Q}_n \Psi_n, \psi_n)$ for some $\psi_n \in \Psi_n$. By induction on $n < \omega$, we build an infinite sequence $(u_n)_{n < \omega}$ of vertices of \mathcal{T} such that the following hold for every $n < \omega$:

- a) $u_n <_T^0 u_{n+1}$;
- b) $\lambda_{\mathcal{T}}(u_n) = \Gamma_n$;
- c) $\mathbf{Q}_n \Psi_n \triangleright_{u_n}^{u_{n+1}} \mathbf{Q}_{n+1} \Psi_{n+1}$;
- d) $\mathbf{Q}_n \Psi_n, \psi_n \triangleright_{u_n}^{u_{n+1}} \mathbf{Q}_{n+1} \Psi_{n+1}, \psi_{n+1}$.

Clearly, this suffices to establish (ii).

Let u_0 be the root of \mathcal{T} . Since the only transition leaving q^* leads to $q_0 = (\{\mathbf{A}\{\varphi\}\}, \mathbf{A}\{\varphi\}, \varphi)$, we have $\Gamma_0 = \{\mathbf{A}\{\varphi\}\} = \lambda_T(u_0)$. For the inductive case, assume that u_n has been defined. By construction of $S_{\mathcal{A}}$ we have $\Gamma_n <_{\odot}^0 \Gamma_{n+1}$, $\mathbf{Q}_n \Psi_n \triangleright_{\Gamma_n}^{\Gamma_{n+1}} \mathbf{Q}_{n+1} \Psi_{n+1}$, and $\mathbf{Q}_n \Psi_n, \psi_n \triangleright_{\Gamma_n}^{\Gamma_{n+1}} \mathbf{Q}_{n+1} \Psi_{n+1}, \psi_{n+1}$. Hence, by Proposition 3.2.26 there is an immediate successor $v \in T$ of u_n labelled by Γ_{n+1} and such that $\mathbf{Q}_n \Psi_n \triangleright_{u_n}^v \mathbf{Q}_{n+1} \Psi_{n+1}$ and $\mathbf{Q}_n \Psi_n, \psi_n \triangleright_{u_n}^v \mathbf{Q}_{n+1} \Psi_{n+1}, \psi_{n+1}$. We let $u_{n+1} := v$ and are done. \blacksquare

We now define several NBA's with skeleton $S_{\mathcal{A}}$:

- (i) $\mathcal{A}_b := (Q, \Sigma, \Delta, q^*, F^b)$, where $F^b := Q$;
- (ii) $\mathcal{A}_{\text{AR}} := (Q, \Sigma, \Delta, q^*, F^{\text{AR}})$, where

$$F^{\text{AR}} := \{(\Gamma, \mathbf{Q}\Psi, \psi) \in Q_0 \mid \mathbf{Q} = \mathbf{A} \text{ and } \psi \text{ is an R-formula}\};$$

- (iii) $\mathcal{A}_{\text{E}} := (Q, \Sigma, \Delta, q^*, F^{\text{E}})$, where $F^{\text{E}} := \{(\Gamma, \mathbf{Q}\Psi, \psi) \in Q_0 \mid \mathbf{Q} = \mathbf{E}\}$;
- (iv) $\mathcal{A}_{\text{U}} := (Q, \Sigma, \Delta, q^*, F^{\text{U}})$, where $F^{\text{U}} := \{(\Gamma, \mathbf{Q}\Psi, \psi) \in Q_0 \mid \psi \text{ is a U-formula}\}$.

The following propositions ensure that these automata behave as expected.

3.2.33. PROPOSITION. *For every $w \in \Sigma^\omega$, we have $w \in \mathcal{L}(\mathcal{A}_b)$ if, and only if, w is branch-like.*

Proof. As no transition in Δ leads to q^* , we have $w \in \mathcal{L}(\mathcal{A}_b)$ if, and only if, there is a run of \mathcal{A}_b on w . So, since the runs of \mathcal{A}_b are exactly the runs of $S_{\mathcal{A}}$, the claim follows immediately from Proposition 3.2.32 and the fact that every infinite sequent trace contains at least one infinite formula trace, by König's lemma. \blacksquare

3.2.34. PROPOSITION. *For every $w \in \Sigma^\omega$, we have $w \in \mathcal{L}(\mathcal{A}_{\text{AR}})$ if, and only if, w is branch-like, τ_w is of type \mathbf{A} , and there is an infinite formula trace of type \mathbf{R} on τ_w .*

Proof. Let $w = ((\Gamma_n, \mathbf{Q}_n \Psi_n))_{n < \omega}$. For the left-to-right direction, assume that $w \in \mathcal{L}(\mathcal{A}_{\text{AR}})$ and let $q^* \frown (q_n)_{n < \omega}$ be an accepting run of \mathcal{A}_{AR} on w , say with $q_n = (\Gamma_n, \mathbf{Q}_n \Psi_n, \psi_n)$ for every $n < \omega$. By construction, $q^* \frown (q_n)_{n < \omega}$ is also an accepting run of \mathcal{A}_b on w , so by Propositions 3.2.32 and 3.2.33 w is branch-like and $\rho := (\psi_n)_{n < \omega}$ is an infinite formula trace on τ_w . By Propositions 3.2.9 and 3.2.10, there is an $N < \omega$ such that $\mathbf{Q}_{i+1} = \mathbf{Q}_i$ and $\psi_i \in \{\alpha, \mathbf{X}\alpha\}$ for every $i \geq N$, where $\alpha \in \text{Clos}(\varphi)$ is an \mathbf{O} -formula for some $\mathbf{O} \in \{\mathbf{U}, \mathbf{R}\}$. As there are infinitely many $i \geq N$ for which $q_i \in F^{\text{AR}}$, the traces τ_w and ρ are of type \mathbf{A} and \mathbf{R} , respectively.

Conversely, assume that w is branch-like, that τ_w is of type A, and that there is an infinite formula trace ρ on τ_w of type R. By Proposition 3.2.32, $q^*\widehat{(q_n)_{n<\omega}}$, where $q_n := (\lambda_T(\pi_w(n)), \tau_w(n), \rho(n))$, is a run of S_A from q^* on w . Since τ_w and ρ are of type A and R, respectively, by Propositions 3.2.9 and 3.2.10 there is an $N < \omega$ such that for infinitely many $i \geq N$ we have that $\tau_w(n)$ is an A-sequent and $\rho(n)$ is an R-formula. Therefore, $q^*\widehat{(q_n)_{n<\omega}}$ is an accepting run of \mathcal{A}_{AR} on w . ■

3.2.35. PROPOSITION. *For every $w \in \Sigma^\omega$, we have $w \in \mathcal{L}(\mathcal{A}_E)$ if, and only if, w is branch-like and τ_w is of type E.*

Proof. Let $w = ((\Gamma_n, \mathbf{Q}_n \Psi_n))_{n<\omega}$. For the left-to-right direction, assume that $w \in \mathcal{L}(\mathcal{A}_E)$ and let $q^*\widehat{(q_n)_{n<\omega}}$ be an accepting run of \mathcal{A}_E on w , say with $q_n = (\Gamma_n, \mathbf{Q}_n \Psi_n, \psi_n)$ for every $n < \omega$. By construction, $q^*\widehat{(q_n)_{n<\omega}}$ is also an accepting run of \mathcal{A}_b on w , so by Proposition 3.2.33 w is branch-like. By Proposition 3.2.10, there is an $N < \omega$ such that $\mathbf{Q}_{i+1} = \mathbf{Q}_i$ for every $i \geq N$. As there are infinitely many $i \geq N$ for which $q_i \in F^E$, the trace τ_w is of type E.

Conversely, assume that w is branch-like and that τ_w is of type E. Let ρ be an infinite formula trace on τ_w given by König's lemma. By Proposition 3.2.32, $q^*\widehat{(q_n)_{n<\omega}}$, where $q_n := (\lambda_T(\pi_w(n)), \tau_w(n), \rho(n))$, is a run of S_A from q^* on w . Since τ_w is of type E, by Proposition 3.2.10 there is an $N < \omega$ such that $\tau_w(n)$ is an E-sequent for infinitely many $i \geq N$. Therefore, $q^*\widehat{(q_n)_{n<\omega}}$ is an accepting run of \mathcal{A}_E on w . ■

3.2.36. PROPOSITION. *For every $w \in \Sigma^\omega$, we have $w \in \mathcal{L}(\mathcal{A}_U)$ if, and only if, w is branch-like and there is an infinite formula trace on τ_w of type U.*

Proof. Let $w = ((\Gamma_n, \mathbf{Q}_n \Psi_n))_{n<\omega}$. For the left-to-right direction, assume that $w \in \mathcal{L}(\mathcal{A}_U)$ and let $q^*\widehat{(q_n)_{n<\omega}}$ be an accepting run of \mathcal{A}_U on w , say with $q_n = (\Gamma_n, \mathbf{Q}_n \Psi_n, \psi_n)$ for every $n < \omega$. By construction, $q^*\widehat{(q_n)_{n<\omega}}$ is also an accepting run of \mathcal{A}_b on w , so by Propositions 3.2.32 and 3.2.33 w is branch-like and $\rho := (\psi_n)_{n<\omega}$ is an infinite formula trace on τ_w . By Proposition 3.2.9, there is an $N < \omega$ such that $\psi_i \in \{\alpha, \mathbf{X}\alpha\}$ for every $i \geq N$, where $\alpha \in \text{Clos}(\varphi)$ is an O-formula for some $O \in \{\mathbf{U}, \mathbf{R}\}$. As there are infinitely many $i \geq N$ for which $q_i \in F^U$, the trace ρ is of type U.

Conversely, assume that w is branch-like and that there is an infinite formula trace ρ on τ_w of type U. By Proposition 3.2.32, $q^*\widehat{(q_n)_{n<\omega}}$, where $q_n := (\lambda_T(\pi_w(n)), \tau_w(n), \rho(n))$, is a run of S_A from q^* on w . Since ρ is of type U, by Proposition 3.2.9 there is an $N < \omega$ such that $\rho(n)$ is a U-formula for infinitely many $i \geq N$. Therefore, $q^*\widehat{(q_n)_{n<\omega}}$ is an accepting run of \mathcal{A}_U on w . ■

Let $\mathcal{L} := \mathcal{L}(\mathcal{A}_b) \cap [\mathcal{L}(\mathcal{A}_{\text{AR}}) \cup (\mathcal{L}(\mathcal{A}_E) \cap \mathcal{L}(\mathcal{A}_U)^c)]$, where $\mathcal{L}(\mathcal{A}_U)^c := \Sigma^\omega \setminus \mathcal{L}(\mathcal{A}_U)$. By Theorem 1.6.3, there exists an NBA $\mathcal{A}_g = (Q_g, \Sigma, \Delta_g, q_g^*, F_g)$ such that $\mathcal{L}(\mathcal{A}_g) = \mathcal{L}$. The automaton \mathcal{A}_g recognises the ω -words over Σ corresponding to infinite branches and good sequent traces:

3.2.37. PROPOSITION. *For every $w \in \Sigma^\omega$, we have $w \in \mathcal{L}(\mathcal{A}_g)$ if, and only if, w is branch-like and τ_w is a good trace on π_w .*

Proof. Assume that $w \in \mathcal{L}(\mathcal{A}_g)$. Then, w is branch-like by Proposition 3.2.33. Suppose first that $w \in \mathcal{L}(\mathcal{A}_{\text{AR}})$. Then, τ_w is a good sequent trace of type A by Proposition 3.2.34. Suppose now that $w \in \mathcal{L}(\mathcal{A}_E) \cap \mathcal{L}(\mathcal{A}_U)^c$. By Proposition 3.2.35, τ_w is of type E. If τ_w were bad, then there would be an infinite formula trace of type U on τ_w and thus $w \in \mathcal{L}(\mathcal{A}_U)$ by Proposition 3.2.36, contradiction.

Conversely assume that w is branch-like and that τ_w is good. By Proposition 3.2.33, $w \in \mathcal{L}(\mathcal{A}_b)$. Suppose first that τ_w is of type A. Then, there is an infinite formula trace on τ_w of type R, so $w \in \mathcal{L}(\mathcal{A}_{\text{AR}})$ by Proposition 3.2.34 and thus $w \in \mathcal{L}(\mathcal{A}_g)$. Suppose now that τ_w is of type E. Then, $w \in \mathcal{L}(\mathcal{A}_E)$ by Proposition 3.2.35. Towards a contradiction, suppose that $w \notin \mathcal{L}(\mathcal{A}_U)^c$, that is, $w \in \mathcal{L}(\mathcal{A}_U)$. Then, by Proposition 3.2.36 there is an infinite formula trace on τ_w of type U, contradicting the fact that τ_w is good. Hence, $w \in \mathcal{L}(\mathcal{A}_U)^c$ and thus $w \in \mathcal{L}(\mathcal{A}_g)$. ■

Observe that, while good A-traces are easy to recognise, those of type E require (or at least seem to) complementing a Büchi automaton.

To show that W_I is Borel in $(\Pi_\omega, \mathcal{O})$ we build one last automaton, which can be informally described as the ‘projection’ of \mathcal{A}_g onto Λ_\odot .

Let $\mathcal{A} := (Q_g, \Lambda_\odot, \Delta^A, q_g^*, F_g)$, where for all $q, q' \in Q_g$ and $\Gamma \in \Lambda_\odot$ we let $(q, \Gamma, q') \in \Delta^A$ if, and only if, there exists a $\mathbf{Q}\Psi \in \Gamma$ such that $(q, (\Gamma, \mathbf{Q}\Psi), q') \in \Delta_g$.

3.2.38. PROPOSITION. *For every $w = (\Gamma_n)_{n < \omega} \in \Lambda_\odot^\omega$, we have $w \in \mathcal{L}(\mathcal{A})$ if, and only if, there exists a $(\mathbf{Q}_n \Psi_n)_{n < \omega} \in \text{Seq}_\varphi^\omega$ such that $(\Gamma_n, \mathbf{Q}_n \Psi_n)_{n < \omega} \in \mathcal{L}(\mathcal{A}_g)$.*

Proof. For the left-to-right direction, assume that $w \in \mathcal{L}(\mathcal{A})$ and let $(q_n)_{n < \omega}$ be an accepting run of \mathcal{A} on w . By construction of \mathcal{A} , for every $n < \omega$ there is a $\mathbf{Q}_n \Psi_n \in \Gamma_n$ such that $(q_n, (\Gamma_n, \mathbf{Q}_n \Psi_n), q_{n+1}) \in \Delta_g$. Hence, $(q_n)_{n < \omega}$ is an accepting run of \mathcal{A}_g on $((\Gamma_n, \mathbf{Q}_n \Psi_n)_{n < \omega})$.

Conversely, assume that there exists a $(\mathbf{Q}_n \Psi_n)_{n < \omega} \in \text{Seq}_\varphi^\omega$ such that $w' := (\Gamma_n, \mathbf{Q}_n \Psi_n)_{n < \omega} \in \mathcal{L}(\mathcal{A}_g)$. Let $(q_n)_{n < \omega}$ be an accepting run of \mathcal{A}_g on w' . Then, for every $n < \omega$ we have $(q_n, \Gamma_n, q_{n+1}) \in \Delta^A$, whence $(q_n)_{n < \omega}$ is an accepting run of \mathcal{A} on w . ■

As an immediate consequence of Propositions 3.2.37 and 3.2.38, we have:

3.2.39. COROLLARY. *For every $w = (\Gamma_n)_{n < \omega} \in \Lambda_{\odot}^{\omega}$, we have $w \in \mathcal{L}(\mathcal{A})$ if, and only if, w is branch-like and π_w is a good branch of \mathcal{T} .*

We abuse notation and extend the map λ_{\odot} to $\lambda_{\odot} : \Pi_{\omega} \rightarrow \Lambda_{\odot}^{\omega}$ by setting $\lambda_{\odot}((u_n)_{n < \omega}) := (\lambda_{\odot}(u_n))_{n < \omega}$ for every $(u_n)_{n < \omega} \in \Pi_{\omega}$.

3.2.40. PROPOSITION. *The map $\lambda_{\odot} : (\Pi_{\omega}, \mathcal{O}) \rightarrow (\Lambda_{\odot}^{\omega}, \mathcal{O}')$ is continuous.*

Proof. Recall that \mathcal{O} is generated by the base $\{B_p \mid p \in T_{\odot}^{<\omega}\}$, where each $B_p := \{\pi \in \Pi_{\omega} \mid p \sqsubseteq \pi\}$. Let $\{B'_q \mid q \in \Lambda_{\odot}^{<\omega}\}$, with $B'_q := \{s \in \Lambda_{\odot}^{\omega} \mid q \sqsubseteq s\}$, be a base for \mathcal{O}' .

Let $B'_p \in \mathcal{O}'$ be basic open, $p \in \Lambda_{\odot}^{<\omega}$. If p is the empty sequence, then $\lambda_{\odot}^{-1}[B'_p] = \Pi_{\omega} \in \mathcal{O}$. Assume then that p is non-empty, say $p = (\Gamma_0, \dots, \Gamma_k)$ for some $\Gamma_0, \dots, \Gamma_k \in \Lambda_{\odot}$ and $k < \omega$. Let P be the (finite, possibly empty) collection of all rooted paths (u_0, \dots, u_k) on \mathcal{T}_{\odot} of length $k + 1$ such that $\lambda_{\odot}(u_i) = \Gamma_i$ for all $i \leq k$. Note that $P \subseteq T_{\odot}^{<\omega}$. We claim that $\lambda_{\odot}^{-1}[B'_p] = \bigcup\{B_q \mid q \in P\}$.

Let $\pi \in \lambda_{\odot}^{-1}[B'_p]$, say $\pi = (v_n)_{n < \omega}$. Let Δ_n be the label of v_n for each $n < \omega$. Then,

$$(\Gamma_0, \dots, \Gamma_k) = p \sqsubset \lambda_{\odot}(\pi) = (\Delta_0, \Delta_1, \dots),$$

whence $\Delta_i = \Gamma_i$ for every $i \leq k$ and thus $(\pi(0), \dots, \pi(k)) \in P$, from which $\pi \in \bigcup\{B_q \mid q \in P\}$ follows.

Conversely, let $\pi \in \bigcup\{B_q \mid q \in P\}$, and let $q \in P$ be such that $q \sqsubset \pi$, say $q = (u_0, \dots, u_k)$. We then have

$$p = (\Gamma_0, \dots, \Gamma_k) = \lambda_{\odot}(q) \sqsubset \lambda_{\odot}(\pi),$$

so $\lambda_{\odot}(\pi) \in B'_p$ and we are done.

Therefore, $\lambda_{\odot}^{-1}[B'_p] \in \mathcal{O}$ for every basic open set $B'_p \in \mathcal{O}'$, which establishes the continuity of λ_{\odot} . \blacksquare

We are finally ready to show that $W_{\mathbf{I}}$ is Borel:

3.2.41. PROPOSITION. *The set $W_{\mathbf{I}}$ is Δ_3^0 in $(\Pi_{\omega}, \mathcal{O})$.*

Proof. Let

$$X := \bigcup\{B_p \mid p \in T_{\odot}^{<\omega} \text{ and } p(n) = l \text{ for some axiomatic } l \in \text{Leaf}(\mathcal{T}) \text{ and } n < \omega\}.$$

As a union of basic open sets, X is open in $(\Pi_{\omega}, \mathcal{O})$.

By Propositions 1.4.1 and 3.2.40 and Theorem 1.6.4, it suffices to show that $W_{\mathbf{I}} = \lambda_{\odot}^{-1}[\mathcal{L}(\mathcal{A})] \cup X$.

Let $\pi \in W_{\mathbf{I}}$. Suppose first that $h(\pi)$ is finite. Then, $h(\pi)$ is a finite branch of \mathcal{T} ending at an axiomatic leaf l , whence π also passes through l and thus $\pi \in X$.

Suppose now that $h(\pi)$ is infinite. Then, $\pi = h(\pi)$ is an infinite good branch of \mathcal{T} . Let τ be an infinite good sequent trace on π , and let $w := ((\lambda_{\odot}(\pi(n)), \tau(n)))_{n < \omega} \in \Sigma^{\omega}$. We thus have that w is branch-like and $\tau_w = \tau$ is a good trace, so $w \in \mathcal{L}(\mathcal{A}_g)$ by Proposition 3.2.37. Hence, $\lambda_{\odot}(\pi) \in \mathcal{L}(\mathcal{A})$ by Proposition 3.2.38.

Conversely, let $\pi \in \Pi_{\omega}$ be such that $\pi \in \lambda_{\odot}^{-1}[\mathcal{L}(\mathcal{A})] \cup X$. If $\pi \in X$, then $h(\pi)$ is a good finite branch of \mathcal{T} and thus $\pi \in W_{\mathcal{T}}$. Suppose now that $\pi \in \lambda_{\odot}^{-1}[\mathcal{L}(\mathcal{A})]$, i.e., $\lambda_{\odot}(\pi) \in \mathcal{L}(\mathcal{A})$. By Corollary 3.2.39, there exists a good infinite branch π' of \mathcal{T} such that $\lambda_{\mathcal{T}}(\pi'(n)) = \lambda_{\odot}(\pi(n)) = \lambda_{\mathcal{T}}(\pi(n))$. Therefore, by Proposition 3.2.26 $h(\pi) = \pi$ is also a good infinite branch of \mathcal{T} and thus $\pi \in W_{\mathcal{T}}$. ■

From Proposition 3.2.41 and Theorem 1.5.5, we get:

3.2.42. COROLLARY. *The game $\mathcal{G}_{\mathcal{T}}$ is determined.*

3.2.43. COROLLARY. *Let \mathcal{T} be any CTL_{∞}^* proof-search tree for a formula φ . If φ is valid, then \mathcal{T} contains a subtree which is a CTL_{∞}^* proof of φ .*

Proof. By Corollary 3.2.42, there exists a winning $\mathcal{G}_{\mathcal{T}}$ -strategy for either **Prov** or **Ref**. The latter contradicts the validity of φ by Propositions 3.2.29 and 3.2.31, so **Prov** has a winning strategy and thus the claim follows from Proposition 3.2.30. ■

3.2.44. COROLLARY (Completeness of CTL_{∞}^*). *For any formula φ , if $\text{CTL}^* \models \varphi$, then $\text{CTL}_{\infty}^* \vdash \varphi$.*

Proof. Assume that $\text{CTL}^* \models \varphi$, and let \mathcal{T} be a CTL_{∞}^* proof-search tree for φ (at least one exists by Corollary 3.2.27). By Corollary 3.2.43, $\text{CTL}_{\infty}^* \vdash \varphi$. ■

Putting Proposition 3.2.23 and Corollary 3.2.44 together yields:

3.2.45. THEOREM. *A CTL^* formula φ is valid if, and only if, there is a CTL_{∞}^* proof of φ .*

3.3 Regular ill-founded proofs

Before introducing the cyclic version of CTL_{∞}^* we consider *regular* ill-founded proofs, which will assist us in proving that the cyclic system is sound.

Although CTL_{∞}^* proof-search trees are in general infinite, they only contain finitely many pairwise different hypersequents. Moreover, by construction they are deterministic, in the sense that the rule applied at a vertex is completely determined by the hypersequent labelling the vertex (recall Proposition 3.2.26). It stands to reason, then, that proofs obtained via CTL_{∞}^* proof-search should admit

at least a partially finitary presentation. We now formalise this idea. The question of whether the correctness condition imposed on infinite branches of CTL_∞^* proofs can also be finitised will be left open for reasons discussed in Section 3.4 below.

Recall (see Section 1.1) that a labelled tree is regular if it only contains finitely many (labelled) cones up to isomorphism.

3.3.1. DEFINITION (Deterministic CTL_∞^* derivation). A CTL_∞^* derivation \mathcal{T} is *deterministic* if any vertices $u, v \in T$ labelled by the same hypersequent are also labelled by the same instance of the same CTL_∞^* rule. \dashv

Our proof-search procedure yields proofs which are both deterministic and regular:

3.3.2. PROPOSITION. *For every valid formula φ , there exists a CTL_∞^* proof of φ which is regular and deterministic.*

Proof. Let $\langle \cdot \rangle_\varphi$ be any proof-search guide for φ , and let \mathcal{T} be a proof-search tree for φ guided by $\langle \cdot \rangle_\varphi$ (at least one exists by Corollary 3.2.27). Since φ is valid, by Corollary 3.2.43 there is a subtree \mathcal{T}' of \mathcal{T} such that \mathcal{T}' is a CTL_∞^* proof of φ . By Proposition 3.2.26, for any two vertices $u, v \in T'$ labelled by the same hypersequent, the labelled cones $[u, \rightarrow)_{\mathcal{T}'}$ and $[v, \rightarrow)_{\mathcal{T}'}$ are isomorphic. In particular, \mathcal{T}' is deterministic. Only finitely many pairwise different hypersequents occur in \mathcal{T}' , whence \mathcal{T}' contains only finitely many labelled cones up to isomorphism and is thus regular. \blacksquare

Infinite regular proofs may be folded down into (finite) trees with back-edges. To make this formal we introduce the system $\text{CTL}_{\text{reg}}^*$ as a partially finitary version of CTL_∞^* , in the sense that $\text{CTL}_{\text{reg}}^*$ derivations are finite objects but the correctness condition imposed on proofs remains infinitary.

3.3.3. DEFINITION ($\text{CTL}_{\text{reg}}^*$ derivation). A $\text{CTL}_{\text{reg}}^*$ derivation of a formula φ is a labelled tree with back-edges \mathcal{T} built according to the rules in Figures 3.1 to 3.3 and such that:

- (i) the root of \mathcal{T} has label $A\varphi$;
- (ii) for every repeat $l \in \text{Rep}_\mathcal{T}$ there exists a modal vertex in $[c_l, l)_\mathcal{T}$. \dashv

3.3.4. OBSERVATION. Trees with back-edges are by definition finite (recall Definition 1.2.1), so every $\text{CTL}_{\text{reg}}^*$ derivation is a finite object.

3.3.5. DEFINITION ($\text{CTL}_{\text{reg}}^*$ proof). A $\text{CTL}_{\text{reg}}^*$ proof of a formula φ is a $\text{CTL}_{\text{reg}}^*$ derivation \mathcal{T} of φ such that every non-repeat leaf of \mathcal{T} is axiomatic and on every infinite path through \mathcal{T}° there is a good sequent trace. \dashv

3.3.6. REMARK. By Observation 3.2.13, the second condition in Definition 3.3.5 is equivalent to requiring that any infinite rooted path on \mathcal{T}° contain a good trace.

Soundness and completeness of $\text{CTL}_{\text{reg}}^*$ are both immediate consequences of our work on CTL_∞^* :

3.3.7. PROPOSITION (Soundness of $\text{CTL}_{\text{reg}}^*$). *For every formula φ , if $\text{CTL}_{\text{reg}}^* \vdash \varphi$, then $\text{CTL}^* \models \varphi$.*

Proof. Let \mathcal{T} be a $\text{CTL}_{\text{reg}}^*$ proof of φ . We claim that \mathcal{T}^ω is a CTL_∞^* proof of φ . Note that \mathcal{T}^ω is labelled in accordance with the rules of CTL_∞^* because, since every repeat $l \in \text{Rep}_{\mathcal{T}}$ is labelled by a non-empty hypersequent H_l , we may apply eW externally preservingly to H_l to obtain H_l again, which also labels c_l .

For every leaf $l \in T^\omega$ there is a non-repeat leaf $l' \in T$ such that l and l' have identical labels, so l is axiomatic.

Let $\pi = (u_n)_{n < \omega}$ be an infinite branch of \mathcal{T}^ω . Then, there is an infinite path $\pi' = (u'_n)_{n < \omega}$ through \mathcal{T}° such that, for every $n < \omega$, u_n and u'_n are labelled by the same instance of the same CTL_∞^* rule. By Definition 3.3.3(ii) and Proposition 1.2.3, π encounters infinitely many modal vertices. And, by Definition 3.3.5, there is a good sequent trace on π . Therefore, \mathcal{T}^ω is a CTL_∞^* proof of φ and thus $\text{CTL}^* \models \varphi$ by Proposition 3.2.23. ■

3.3.8. REMARK. In the proof of Proposition 3.3.7 we relied on the contraction implicitly built into the rules of CTL_∞^* in order to show that \mathcal{T}^ω itself is a CTL_∞^* derivation. This is clearly unnecessary because we could instead identify each repeat and its companion, but it simplifies the argument by allowing us to work directly with \mathcal{T}^ω .

3.3.9. LEMMA. *Let \mathcal{T} be a deterministic CTL_∞^* derivation. For all vertices $u, v \in T$ such that $u <_{\mathcal{T}} v$, if u and v are labelled by the same hypersequent, then there is a modal vertex in $[u, v)_{\mathcal{T}}$.*

Proof. Suppose otherwise. By determinism, the labelled cones $[u, \rightarrow)_{\mathcal{T}}$ and $[v, \rightarrow)_{\mathcal{T}}$ are isomorphic. Since $u <_{\mathcal{T}} v$ and, by assumption, there is no modal vertex in $[u, v)_{\mathcal{T}}$, there exists an infinite branch in \mathcal{T} which contains only finitely many modal vertices, contradicting the fact that \mathcal{T} is a CTL_∞^* derivation. ■

3.3.10. PROPOSITION (Completeness of $\text{CTL}_{\text{reg}}^*$). *For any formula φ , if $\text{CTL}^* \models \varphi$, then $\text{CTL}_{\text{reg}}^* \vdash \varphi$.*

Proof. Let \mathcal{T} be a regular deterministic CTL $^*_\infty$ proof of φ given by Proposition 3.3.2.

On any infinite branch of \mathcal{T} there are infinitely many vertices v such that there is a $u <_{\mathcal{T}} v$ labelled by the same hypersequent (and thus, since \mathcal{T} is deterministic, also by the same instance of the same CTL $^*_\infty$ rule). Let then \mathcal{T}_r be the result after pruning each infinite branch of \mathcal{T} at the encounter of the first such vertex on the branch (and keeping said vertex as a repeat in \mathcal{T}_r). We claim that \mathcal{T}_r is a CTL $^*_{\text{reg}}$ proof of φ .

By Lemma 3.3.9, for every repeat $l \in \text{Rep}_{\mathcal{T}_r}$ there exists a modal vertex in $(c_l, l)_{\mathcal{T}_r}$. This establishes that \mathcal{T}_r is a CTL $^*_{\text{reg}}$ derivation (of φ).

Clearly, every non-repeat leaf of \mathcal{T}_r is axiomatic, for it is also a leaf of \mathcal{T} . Let $\pi = (u_n)_{n < \omega}$ be an infinite path through \mathcal{T}_r° such that u_0 is the root of \mathcal{T}_r . To see that π contains a good trace it suffices to find an infinite branch $\pi' = (u'_n)_{n < \omega}$ of \mathcal{T} such that u_n and u'_n are labelled by the same instance of the same CTL $^*_\infty$ rule, for every $n < \omega$. We do so by induction on $n < \omega$. Let u'_0 be the root of \mathcal{T} , which has the same label as u_0 . For the inductive case, suppose that u'_n has been defined. Since any two vertices in \mathcal{T} labelled by the same hypersequent are labelled by the same instance of the same CTL $^*_\infty$ rule, the inductive hypothesis yields an immediate successor v of u'_n in \mathcal{T} labelled by the same hypersequent as u_{n+1} and such that $\triangleright_{u_n}^{u_{n+1}} = \triangleright_{u'_n}^v$ and $\triangleright_{u_n}^{u_{n+1}} = \triangleright_{u'_n}^v$. We then let $u'_{n+1} := v$ and are done. By Remark 3.3.6, \mathcal{T}_r is a CTL $^*_{\text{reg}}$ proof of φ . ■

Putting Propositions 3.3.7 and 3.3.10 together, we get:

3.3.11. THEOREM. *For any formula φ , CTL $^*_{\text{reg}} \vdash \varphi$ if, and only if, CTL $^* \models \varphi$.*

3.4 The cyclic system CTL $^*_\circ$

We now introduce a cyclic version of the system CTL $^*_\infty$ which is ‘less infinitary’ than CTL $^*_{\text{reg}}$ but still not finitary. Formulas are enriched with annotations similar to the ones introduced by Jungteerapanich [76] and Stirling [144] for the modal μ -calculus. This allows us to keep track of fixpoint unfoldings and (under certain conditions) detect good cycles. The annotations we utilise, however, are considerably simpler: annotations of length ≤ 1 suffice for our purposes, and only on O- and XO-formulas (so they never permeate down to other subformulas).

Our cyclic system differs from, and corrects, the one in [6]. It is there claimed that the aforementioned annotations yield a finitary calculus. But the completeness proof for the cyclic system in [6] misses a potentially intractable problem caused by purely existential infinite branches (i.e., those containing no good A-trace): after the annotating procedure performed in the completeness proof, purely existential branches may end up unannotated. Due to its universal form, the correctness condition on these branches does not seem to be finitisable by means of

the simple annotations used in [6] and below. Indeed, one needs to detect the absence of certain infinite traces rather than the presence thereof. This is one the ‘complications’ referred to in Remark 3.2.3 (recall also Remark 3.2.15). It may be possible to devise more complex annotations capable of detecting the absence of traces, for example by employing automata-theoretic tools (see Section 3.6).⁸ We do not explore this line of research here and instead decide not to annotate any formula under the existential quantifier and rely on an infinitary correctness condition in the absence of good A-traces. We then isolate a *universality* finitary condition that suffices to ensure that a cyclic derivation in our calculus is a proof.

Additionally, we use an ‘external thinning’ rule which differs from the one in [6] and merges formula traces more ‘gently’. See Remarks 3.4.6 and 3.4.9 below.

We begin by fixing a countably infinite collection $\mathbf{N} = \{x, y, z, \dots\}$ of *names*. An *annotation* is either the empty string ε or a single name. Annotations are denoted by a, b, \dots . We identify a non-empty annotation with the unique name it contains.

An *annotated (CTL^*) formula* is a pair (φ, a) , henceforth written φ^a , where φ is a CTL^* formula and a is an *annotation*. We identify each unannotated formula φ with the formula φ^ε annotated by the empty annotation. A name x *occurs* in a set of annotated formulas Φ if there is a formula φ such that $\varphi^x \in \Phi$.

An *annotated (CTL^*) sequent* is a pair (\mathbf{Q}, Φ) , henceforth written $\mathbf{Q}\Phi$, where $\mathbf{Q} \in \{\mathbf{A}, \mathbf{E}\}$ and Φ is a finite set of annotated formulas. A name x *occurs* in an annotated sequent $\mathbf{Q}\Phi$ if x occurs in Φ .

An *annotated (CTL^*) hypersequent* is a pair (Θ, Γ) , henceforth written $\Theta \dashv \Gamma$, 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 $\Theta \dashv \Gamma$. A name x *occurs* in a hypersequent Γ if there is a sequent $\mathbf{Q}\Phi \in \Gamma$ such that x occurs in $\mathbf{Q}\Phi$.

The *base* of an annotated formula φ^a is $\mathbf{b}(\varphi^a) := \varphi$. Analogously, the base of an annotated sequent $\mathbf{Q}\Phi$ is $\mathbf{b}(\mathbf{Q}\Phi) := \mathbf{Q}\{\mathbf{b}(\varphi) \mid \varphi \in \Phi\}$. We extend the base function \mathbf{b} to sets of sequents and hypersequents by setting $\mathbf{b}(\Gamma) := \{\mathbf{b}(\mathbf{Q}\Phi) \mid \mathbf{Q}\Phi \in \Gamma\}$ and $\mathbf{b}(\Theta \dashv \Gamma) := \mathbf{b}(\Gamma)$.

An annotated formula φ^a is *well-annotated* if either $a = \varepsilon$, or else $a \in \mathbf{N}$ and φ is an R- or an XR-formula.

Given a finite sequence of names Θ , we define the strict linear order \prec_Θ on $\{\varepsilon\} \cup \{x \in \mathbf{N} \mid x \in \Theta\}$ by letting $a \prec_\Theta b$ if, and only if, either $a \neq \varepsilon = b$, or both a and b are non-empty strings and the name in a occurs in Θ strictly before the name in b .

The rules of the cyclic system CTL_\circ^* , defined below, are given in Figures 3.5 to 3.7. In rules AX, EX, iW, eW, iThin, eThin and del, we denote by Θ' the result

⁸Recall that the construction of the Büchi automaton \mathcal{A}_g in Section 3.2.2 involved the complement of the automaton \mathcal{A}_U precisely to detect the absence of infinite U-traces.

$$\begin{array}{c}
\text{ax} \frac{}{\Theta \dashv Qp, Q'\bar{p}, \Delta} \qquad \text{ax}_\top \frac{}{\Theta \dashv E\emptyset, \Delta} \\
\text{ALit} \frac{\Theta \dashv A\Phi, A\ell, \Delta}{\Theta \dashv A\{\Phi, \ell\}, \Delta} \qquad \text{ELit} \frac{\Theta \dashv E\Phi, \Delta \quad \Theta \dashv E\ell, \Delta}{\Theta \dashv E\{\Phi, \ell\}, \Delta} \\
\text{AV} \frac{\Theta \dashv A\{\Phi, \varphi, \psi\}, \Delta}{\Theta \dashv A\{\Phi, \varphi \vee \psi\}, \Delta} \qquad \text{EV} \frac{\Theta \dashv E\{\Phi, \varphi\}, E\{\Phi, \psi\}, \Delta}{\Theta \dashv E\{\Phi, \varphi \vee \psi\}, \Delta} \\
\text{A}\wedge \frac{\Theta \dashv A\{\Phi, \varphi\}, \Delta \quad \Theta \dashv A\{\Phi, \psi\}, \Delta}{\Theta \dashv A\{\Phi, \varphi \wedge \psi\}, \Delta} \qquad \text{E}\wedge \frac{\Theta \dashv E\{\Phi, \varphi, \psi\}, \Delta}{\Theta \dashv E\{\Phi, \varphi \wedge \psi\}, \Delta} \\
\text{AA} \frac{\Theta \dashv A\Phi, A\{\psi\}, \Delta}{\Theta \dashv A\{\Phi, A\psi\}, \Delta} \qquad \text{EA} \frac{\Theta \dashv E\Phi, \Delta \quad \Theta \dashv A\{\psi\}, \Delta}{\Theta \dashv E\{\Phi, A\psi\}, \Delta} \\
\text{AE} \frac{\Theta \dashv A\Phi, E\{\psi\}, \Delta}{\Theta \dashv A\{\Phi, E\psi\}, \Delta} \qquad \text{EE} \frac{\Theta \dashv E\Phi, \Delta \quad \Theta \dashv E\{\psi\}, \Delta}{\Theta \dashv E\{\Phi, E\psi\}, \Delta} \\
\text{AX} \frac{\Theta' \dashv A\Phi, E\Psi_1, \dots, E\Psi_m}{\Theta \dashv A(X\Phi \cup \Xi), EX\Psi_1, \dots, EX\Psi_m, \Sigma} \qquad \text{EX} \frac{\Theta' \dashv E\Psi_1, \dots, E\Psi_m}{\Theta \dashv EX\Psi_1, \dots, EX\Psi_m, \Sigma}
\end{array}$$

FIGURE 3.5: Non-fixpoint, logical rules of the system CTL_\circ^* .

after removing from Θ all names not occurring in the premise with control Θ' . In rule AR_1 , we define Θx as the concatenation $\Theta \hat{\ } x$, if $x \notin \Theta$, and otherwise $\Theta x := \Theta$. In rule eThin , we define $c_i := \min_{\prec_\Theta} \{a_i, b_i\}$ for every $i \leq n$.

3.4.1. DEFINITION (CTL_\circ^* derivation). A CTL_\circ^* derivation of a formula φ is a (finite) labelled tree with back-edges \mathcal{T} built according to the rules in Figures 3.5 to 3.7 and such that:

- (i) the root of \mathcal{T} has label $A\varphi$;
- (ii) for every repeat $l \in \text{Rep}_\mathcal{T}$ there exists a vertex $u \in [c_l, l)_\mathcal{T}$ such that rule AX or EX is applied at u . -

Sequent and formula traces, as well as principal and active sequents and formulas, follow the definition from the system CTL_∞^* in the cases of the rules in Figures 3.5 and 3.6 and rules iW and eW .

$$\begin{array}{c}
\text{AU} \frac{\Theta \vdash \text{A}\{\Phi, \varphi_1, \varphi_2\}, \Delta \quad \Theta \vdash \text{A}\{\Phi, \varphi_2, \text{X}(\varphi_1 \text{U} \varphi_2)\}, \Delta}{\Theta \vdash \text{A}\{\Phi, \varphi_1 \text{U} \varphi_2\}, \Delta} \\
\text{AR}_0 \frac{\Theta \vdash \text{A}\{\Phi, \varphi_2\}, \Delta \quad \Theta \vdash \text{A}\{\Phi, \varphi_1, \text{X}(\varphi_1 \text{R} \varphi_2)^a\}, \Delta}{\Theta \vdash \text{A}\{\Phi, (\varphi_1 \text{R} \varphi_2)^a\}, \Delta} \\
\text{AR}_1 \frac{\Theta \vdash \text{A}\{\Phi, \varphi_2\}, \Delta \quad \Theta x \vdash \text{A}\{\Phi, \varphi_1, \text{X}(\varphi_1 \text{R} \varphi_2)^x\}, \Delta}{\Theta \vdash \text{A}\{\Phi, \varphi_1 \text{R} \varphi_2\}, \Delta} \quad x \in \mathbb{N} \setminus \Theta \\
\text{EU} \frac{\Theta \vdash \text{E}\{\Phi, \varphi_2\}, \text{E}\{\Phi, \varphi_1, \text{X}(\varphi_1 \text{U} \varphi_2)\}, \Delta}{\Theta \vdash \text{E}\{\Phi, \varphi_1 \text{U} \varphi_2\}, \Delta} \\
\text{ER} \frac{\Theta \vdash \text{E}\{\Phi, \varphi_1, \varphi_2\}, \text{E}\{\Phi, \varphi_2, \text{X}(\varphi_1 \text{R} \varphi_2)\}, \Delta}{\Theta \vdash \text{E}\{\Phi, \varphi_1 \text{R} \varphi_2\}, \Delta}
\end{array}$$

FIGURE 3.6: Fixpoint rules of the system CTL_\circ^* .

In rule AR_0 with $a \neq \emptyset$, the unique name in a is principal in the conclusion and active in the premise corresponding to the unfolding. Analogously, in rule AR_1 the name x is principal in the conclusion and active in the premise corresponding to the unfolding.

In the *internal thinning* rule iThin , the name x is *principal* in the conclusion and *active* in the premise, and if $a \neq \varepsilon$ then the unique name in a is *principal* in the conclusion. In the *external thinning* rule eThin , the names in $a_0, \dots, a_n, b_0, \dots, b_n$ are *principal* in the conclusion, and the ones in c_0, \dots, c_n are *active* in the premise.

In the *deletion* rule del , the name x is *principal* in the conclusion.

In the cases of iThin , eThin and del , principal and active sequents and formulas are defined as expected: the distinguished sequents and formulas in the conclusion are principal, the distinguished sequents and formulas in the premise are active, every sequent in Δ is a side sequent, and every formula in Φ is a side formula. Sequent and formula traces are also defined as expected for iThin and del , with the (possible) exception that in iThin we let $\text{Q}\{\Phi, \varphi^x, \varphi^a\}, \varphi^a \triangleright \text{Q}\{\Phi, \varphi^x\}, \varphi^x$ and in eThin we let:

- $\text{Q}\{\varphi_0^{a_0}, \dots, \varphi_n^{a_n}\} \triangleright \text{Q}\{\varphi_0^{c_0}, \dots, \varphi_n^{c_n}\};$
- $\text{Q}\{\varphi_0^{b_0}, \dots, \varphi_n^{b_n}\} \triangleright \text{Q}\{\varphi_0^{c_0}, \dots, \varphi_n^{c_n}\};$
- $\text{Q}\{\varphi_0^{a_0}, \dots, \varphi_n^{a_n}\}, \varphi_i^{a_i} \triangleright \text{Q}\{\varphi_0^{c_0}, \dots, \varphi_n^{c_n}\}, \varphi_i^{c_i}$, for all $i \leq n$;

$$\begin{array}{c}
\text{iW} \frac{\Theta' \dashv \mathbf{A}\Phi, \Delta}{\Theta \dashv \mathbf{A}(\Phi \cup \Psi), \Delta} \qquad \text{eW} \frac{\Theta' \dashv \Delta}{\Theta \dashv \Gamma, \Delta} \\
\text{iThin} \frac{\Theta' \dashv \mathbf{A}\{\Phi, \varphi^x\}, \Delta}{\Theta \dashv \mathbf{A}\{\Phi, \varphi^x, \varphi^a\}, \Delta} x \prec_{\Theta} a \qquad \text{del} \frac{\Theta' \dashv \mathbf{A}\{\Phi, \varphi\}, \Delta}{\Theta \dashv \mathbf{A}\{\Phi, \varphi^x\}, \Delta} \\
\text{eThin} \frac{\Theta' \dashv \mathbf{A}\{\varphi_0^{c_0}, \dots, \varphi_n^{c_n}\}, \Delta}{\Theta \dashv \mathbf{A}\{\varphi_0^{a_0}, \dots, \varphi_n^{a_n}\}, \mathbf{A}\{\varphi_0^{b_0}, \dots, \varphi_n^{b_n}\}, \Delta}
\end{array}$$

FIGURE 3.7: Structural rules of the system CTL_{\circ}^* .

- $\mathbf{Q}\{\varphi_0^{b_0}, \dots, \varphi_n^{b_n}\}, \varphi_i^{b_i} \triangleright \mathbf{Q}\{\varphi_0^{c_0}, \dots, \varphi_n^{c_n}\}, \varphi_i^{c_i}$, for all $i \leq n$.

Observe, then, that even though the hypersequent $\Theta' \dashv \mathbf{A}\{\Phi, \varphi^x\}, \Delta$ may be obtained from $\Theta \dashv \mathbf{A}\{\Phi, \varphi^x, \varphi^a\}, \Delta$ via an application of either **iW** or **iThin**, in the former case we have $\varphi^a \not\triangleright \varphi^x$ but in the latter $\varphi^a \triangleright \varphi^x$.

The merging of sequent and formula traces in the thinning rules, on which our proofs of the soundness and completeness of CTL_{\circ}^* rely, corresponds to the fact that, if the annotations are dropped, then, for example, φ^x and φ^a in $\mathbf{A}\{\varphi^x, \varphi^a\}$ become the same formula and their traces merge. And analogously for **eThin** and sequent traces.

For a comparison of our **eThin** rule and the one in [6], see Remark 3.4.6 below.

Internally and externally discarding and preserving instances of non-modal rules of CTL_{\circ}^* are defined as expected from the definitions given for CTL_{∞}^* rules. We say that a CTL_{\circ}^* derivation \mathcal{T} is *discarding* if no rule is applied preservingly (internally or externally) in \mathcal{T} .

An inspection of the rules of CTL_{\circ}^* immediately yields:

3.4.2. LEMMA. *Let \mathcal{T} be a discarding CTL_{\circ}^* derivation. For every name x and every vertex $u \in T$, say with label $\Theta \dashv \Gamma$, the following hold:*

- (i) x occurs in at most one \mathbf{A} -sequent in Γ ;
- (ii) for every \mathbf{A} -sequent $\mathbf{A}\Phi \in \Gamma$ there is at most one formula $\varphi^a \in \Phi$ with $a = x$.

3.4.3. REMARK. Our attention will mostly be restricted to discarding derivations. As shown in Lemma 3.4.2, the annotations occurring in them are very well-behaved. It is not clear whether our work admits a non-discarding setting (see Remark 3.4.9 below).

When unfolding an unannotated R-formula φ under the universal quantifier, we may mark the unfolding by annotating the active formula $X\varphi$ (rule AR_1), or we may choose to leave $X\varphi$ unannotated (AR_0). If annotated, subsequent unfoldings may only be performed by rule AR_0 and thus the annotation is preserved unless we remove it by applying thinning, weakening or deletion. This degree of freedom corresponds to the fact that, in order to identify the CTL_\circ^* derivations that should count as proofs, we shall look for infinite R-traces, and thus it does not matter when we begin to annotate R-unfoldings on a sequent trace as long as we do so eventually.

A vertex u of a CTL_\circ^* derivation is *modal* if rule AX or EX is applied at u . And u is *thinning* if rule iThin or eThin is applied at u .

A name x is *fixed* on a path π through a CTL_\circ^* derivation if x occurs in the control of every hypersequent labelling a vertex in π . Similarly, x is *fixed* on a sequent (formula) trace if x occurs in each sequent (respectively, formula) on the trace.

Let \mathcal{T} be a CTL_\circ^* derivation. A name x is *eventually fixed* on an infinite path π through \mathcal{T}° (infinite sequent trace τ , infinite formula trace ρ) if there is a tail π' of π (respectively, a tail τ' of τ , a tail ρ' of ρ) such that x is fixed on π' (respectively, τ' , ρ').

We are ready to identify the CTL_\circ^* derivations that should count as proofs.

3.4.4. DEFINITION (CTL_\circ^* proof). A CTL_\circ^* *proof* of a formula φ is a CTL_\circ^* derivation \mathcal{T} of φ satisfying the following:

- (i) \mathcal{T} is discarding;
- (ii) every non-repeat leaf of \mathcal{T} is axiomatic;
- (iii) for every infinite path π through \mathcal{T}° :
 - a) either there is a name eventually fixed on π ;
 - b) or there is a good E-trace on π .

–

Observe that condition (iii) in Definition 3.4.4 is infinitary. The reader may thus wonder whether CTL_\circ^* has any advantage over $\text{CTL}_{\text{reg}}^*$. The answer is that annotations are a first step towards full finitariness. We shall later see (Proposition 3.4.14) that they allow us to isolate finitary conditions that suffice to guarantee that a derivation is a proof. Additionally, in Section 3.5 we shall use annotations to find inductive invariants, thereby allowing a translation of certain cyclic proofs into ordinary, acyclic ones.

3.4.5. EXAMPLE. Figure 3.8 depicts (the unique cycle in) the CTL_\circ^* proof corresponding to the CTL_∞^* proof in Figure 3.4 of the formula $(\top U p) \vee (\perp R \bar{p})$. The

$$\begin{array}{c}
\text{AX} \frac{(\dagger) \quad x \dashv \text{A}\{\top \text{U}p, (\perp \text{R}\bar{p})^x\}}{x \dashv \text{A}\{\text{X}(\top \text{U}p), p, \text{X}(\perp \text{R}\bar{p})^x\}} \\
\text{AU} \frac{x \dashv \text{A}\{\top \text{U}p, p, \text{X}(\perp \text{R}\bar{p})^x\}}{x \dashv \text{A}\{\top \text{U}p, \perp, \text{X}(\perp \text{R}\bar{p})^x\}} \\
\text{A}\wedge \frac{x \dashv \text{A}\{\top \text{U}p, \perp, \text{X}(\perp \text{R}\bar{p})^x\}}{x \dashv \text{A}\{\top \text{U}p, \perp, \text{X}(\perp \text{R}\bar{p})^x\}} \\
\text{AR}_0 \frac{(\dagger) \quad x \dashv \text{A}\{\top \text{U}p, (\perp \text{R}\bar{p})^x\}}{x \dashv \text{A}\{\text{X}(\top \text{U}p), p, \text{X}(\perp \text{R}\bar{p})^x\}} \\
\text{AX} \frac{x \dashv \text{A}\{\text{X}(\top \text{U}p), p, \text{X}(\perp \text{R}\bar{p})^x\}}{x \dashv \text{A}\{\top \text{U}p, p, \text{X}(\perp \text{R}\bar{p})^x\}} \\
\text{AU} \frac{x \dashv \text{A}\{\top \text{U}p, p, \text{X}(\perp \text{R}\bar{p})^x\}}{x \dashv \text{A}\{\top \text{U}p, \perp, \text{X}(\perp \text{R}\bar{p})^x\}} \\
\text{A}\wedge \frac{x \dashv \text{A}\{\top \text{U}p, \perp, \text{X}(\perp \text{R}\bar{p})^x\}}{x \dashv \text{A}\{\top \text{U}p, \perp, \text{X}(\perp \text{R}\bar{p})^x\}} \\
\text{AR}_1 \frac{x \dashv \text{A}\{\top \text{U}p, \perp, \text{X}(\perp \text{R}\bar{p})^x\}}{\text{A}\{\top \text{U}p, \perp \text{R}\bar{p}\}} \\
\text{A}\vee \frac{\text{A}\{\top \text{U}p, \perp \text{R}\bar{p}\}}{\text{A}\{\top \text{U}p \vee \perp \text{R}\bar{p}\}}
\end{array}$$

FIGURE 3.8: A cycle in a CTL_o^* proof of the valid formula $(\top \text{U}p) \vee (\perp \text{R}\bar{p})$, corresponding to the ill-founded proof in Figure 3.4. The unique repeat and its companion are marked with the symbol ‘†’. Dashed lines indicate the omission of some vertices.

infinite branch in the ill-founded proof corresponds to the cycle in the cyclic one. It will become clear in Section 3.5 why we did not stop after applying AR_0 , even though doing so would have resulted in a shorter proof.

3.4.6. REMARK. Our eThin rule differs from the one in [6]. Whereas we keep the \prec_Θ -least annotation for each principal formula, in [6] the active sequent in an instance of eThin is always one of the two principal sequents⁹ and the choice is made by considering a single name only, thus allowing formula traces of the form $\varphi^x \triangleright \varphi^y$ with $x \prec_\Theta y$. This is a potential source of problems when proving completeness of the cyclic system by annotating ill-founded proofs (in particular, in the proof of Lemma 3.4.21 below), for it might give rise to good A -traces on which no name is eventually fixed.

The merging of formula traces in our rule avoids this problem by ensuring that annotations never grow with respect to \prec_Θ . It could not have been used in [6], however, because it would have invalidated the multi-name version of Lemma 3.4.7 below, needed in [6] to accommodate annotations under the existential quantifier of both R - and U -formulas.

We now turn to proving the soundness and completeness of CTL_o^* . Essentially, we do so by showing that ill-founded proofs can be seen as unravellings of cyclic proofs.

⁹We shall consider such instances of (our version of) eThin in Section 3.5.

3.4.1 Soundness of CTL_\circ^*

We prove soundness of CTL_\circ^* by dropping the annotations from a given CTL_\circ^* proof and showing that the resulting tree is a $\text{CTL}_{\text{reg}}^*$ proof. We first need some results relating traces and names in CTL_\circ^* derivations.

3.4.7. LEMMA. *Let \mathcal{T} be a CTL_\circ^* derivation, and $\pi = (u_i)_{i \leq n}$ a finite path through \mathcal{T}° . If a name x is fixed on π , then there is a sequent trace τ on π such that x is fixed on τ .*

Proof. If $n = 0$ there is nothing to prove, so assume otherwise. Let $\Theta_i \dashv \Gamma_i$ be the label of u_i for every $i \leq n$. Since x is fixed on π , there is a sequent $\mathbf{Q}\Phi \in \Gamma_n$ such that x occurs in $\mathbf{Q}\Phi$. It suffices to find a sequent $\mathbf{Q}\Phi' \in \Gamma_{n-1}$ such that x occurs in $\mathbf{Q}\Phi'$ and $\mathbf{Q}\Phi' \triangleright \mathbf{Q}\Phi$, for then repeating the same argument finitely many times yields the desired trace τ .

Towards a contradiction, suppose that there is no such sequent $\mathbf{Q}\Phi' \in \Gamma_{n-1}$. In particular, then, rule \mathbf{eThin} is not applied at u_{n-1} with $\mathbf{Q}\Phi$ active.

Let $\mathbf{Q}'\Psi \in \Gamma_{n-1}$ be such that $\mathbf{Q}'\Psi \triangleright \mathbf{Q}\Phi$. Note that such a sequent $\mathbf{Q}'\Psi$ exists by the definition of \triangleright . By assumption, x does not occur in $\mathbf{Q}'\Psi$. Then, x is introduced in $\mathbf{Q}\Phi$ from $\mathbf{Q}'\Psi$ by an application of rule \mathbf{AR}_1 at u_{n-1} , contradicting the fact that $x \in \Theta_{n-1}$. ■

3.4.8. LEMMA. *Let \mathcal{T} be a discarding CTL_\circ^* derivation, π a finite path through \mathcal{T}° , and τ a sequent trace on π . If a name x is fixed on τ , then there is a formula trace on τ on which x is fixed.*

Proof. Let $\pi = (u_i)_{i \leq n}$ and $\tau = (\mathbf{Q}_i\Phi_i)_{i \leq n}$. If $n = 0$ there is nothing to prove, so assume otherwise. Since x occurs in Φ_n , there is some $\varphi^x \in \Phi_n$. It suffices to find a formula $\psi^x \in \Phi_{n-1}$ such that $\psi^x \triangleright \varphi^x$, for then repeating the same argument finitely many times yields the desired trace ρ . Towards a contradiction, suppose that no such ψ^x exists. Let $\psi^a \in \Phi_{n-1}$ be such that $\psi^a \triangleright \varphi^x$. Note that such a formula ψ^a exists by the definition of \triangleright . By assumption, $a \neq x$, so a rule \mathbf{R} among \mathbf{AR}_1 , \mathbf{iThin} and \mathbf{eThin} is applied at u_{n-1} with ψ^a principal and φ^x active.

The case $\mathbf{R} = \mathbf{AR}_1$ contradicts the fact that x occurs in Φ_{n-1} .

Suppose that $\mathbf{R} = \mathbf{iThin}$. Then, $\Phi_{n-1} \ni \varphi^x \triangleright \varphi^x$, contradicting our assumption.

Finally, suppose that $\mathbf{R} = \mathbf{eThin}$. Since \mathcal{T} is discarding and x occurs in Φ_{n-1} , by Lemma 3.4.2 $\mathbf{Q}_{n-1}\Phi_{n-1}$ is the only sequent in the label of u_{n-1} in which x occurs, whence again $\Phi_{n-1} \ni \varphi^x \triangleright \varphi^x$, contradicting our assumption. ■

$$\begin{array}{c}
\text{eThin} \frac{xy \vdash A\{(pRq)^x, X(pRq)^y\}}{xy \vdash A\{X(pRq), (pRq)^y\}, A\{(pRq)^x, X(pRq)^y\}} \\
\text{iW} \frac{xy \vdash A\{p, X(pRq), (pRq)^y\}, A\{(pRq)^x, X(pRq)^y\}}{xy \vdash A\{pRq, (pRq)^y\}, A\{(pRq)^x, X(pRq)^y\}} \\
\text{AR}_0 \frac{xy \vdash A\{pRq, (pRq)^y\}, A\{(pRq)^x, X(pRq)^y\}}{xy \vdash A\{pRq, (pRq)^x, (pRq)^y\}} \\
\text{iW} \frac{xy \vdash A\{pRq, (pRq)^x, (pRq)^y\}, A\{pRq, (pRq)^x, p, (pRq)^y, X(pRq)^y\}}{xy \vdash A\{pRq, (pRq)^x, (pRq)^y\}} \\
\text{AX} \frac{xy \vdash A\{X(pRq), X(pRq)^x, X(pRq)^y\}}{xy \vdash A\{p, X(pRq), X(pRq)^x, X(pRq)^y\}} \\
\text{iW} \frac{xy \vdash A\{p, X(pRq), X(pRq)^x, X(pRq)^y\}}{xy \vdash A\{p, pRq, X(pRq)^x, X(pRq)^y\}} \\
\text{AR}_0 \frac{xy \vdash A\{p, pRq, X(pRq)^x, X(pRq)^y\}}{xy \vdash A\{p, pRq, (pRq)^x, X(pRq)^y\}} \\
\text{AR}_1 \frac{xy \vdash A\{p, pRq, (pRq)^x, X(pRq)^y\}}{x \vdash A\{pRq, (pRq)^x\}} \\
\text{AX} \frac{x \vdash A\{X(pRq), X(pRq)^x\}}{x \vdash A\{p, X(pRq), X(pRq)^x\}} \\
\text{iW} \frac{x \vdash A\{p, X(pRq), X(pRq)^x\}}{x \vdash A\{p, pRq, X(pRq)^x\}} \\
\text{AR}_0 \frac{x \vdash A\{p, pRq, X(pRq)^x\}}{x \vdash A\{p, pRq, X(pRq)^x\}} \\
\text{AR}_1 \frac{x \vdash A\{p, pRq, X(pRq)^x\}}{A\{pRq\}}
\end{array}$$

FIGURE 3.9: A non-discarding CTL_\circ^* derivation showing that discardingness is a necessary condition in Lemma 3.4.8. Dashed lines indicate the omission of some vertices. The name y is fixed on the sequent trace $A\{X(pRq), (pRq)^y\} \triangleright A\{(pRq)^x, X(pRq)^y\}$ at the top, yet there is no formula trace on the sequent trace on which y is fixed.

3.4.9. REMARK. Figure 3.9 depicts a non-discarding CTL_\circ^* derivation showing that the discardingness assumption in Lemma 3.4.8 cannot be dropped. Persistence of names on a path does not seem to permeate all the way down to the formula level in the presence of (explicit or implicit) contraction, regardless of the rule used to merge sequents (**eThin**, in our case). This is overlooked in [6] and may give rise to an intractable problem due to the use therein of names under existential quantifiers, for which there does not seem to be an analogue of Lemma 3.4.2.

We now obtain, from König's lemma, infinitary versions of the two preceding results.

3.4.10. LEMMA. *Let \mathcal{T} be a CTL_\circ^* derivation, π an infinite path through \mathcal{T}° , and $x \in \mathbb{N}$ any name. If there are infinitely many $n < \omega$ for which there exists a sequent trace on $\pi_{\leq n}$ on which x is fixed, then there is an infinite sequent trace on π on which x is fixed.*

Proof. For every sequent $\mathbf{Q}\Phi$ in the label of $\pi(0)$ and every $n < \omega$, let $S(\mathbf{Q}\Phi, n)$ be the collection of all sequent traces on $\pi_{\leq n}$ starting from $\mathbf{Q}\Phi$ and on which x is

fixed. As the label of $\pi(0)$ is finite, by assumption there is a sequent $\mathbf{Q}\Phi$ in the label of $\pi(0)$ such that $S(\mathbf{Q}\Phi, n) \neq \emptyset$ for every $n < \omega$ (note that every trace in $S(\mathbf{Q}\Phi, n+1)$ yields one in $S(\mathbf{Q}\Phi, n)$). Let $S := \bigcup_{n < \omega} S(\mathbf{Q}\Phi, n)$. Then, (S, \sqsubset) is a tree of height ω each of whose levels is finite, so by König's lemma (S, \sqsubset) has an infinite branch, whence clearly there is an infinite sequent trace on π starting from $\mathbf{Q}\Phi$ and on which x is fixed. ■

3.4.11. LEMMA. *Let \mathcal{T} be a discarding CTL^*_\circ derivation, π an infinite path through \mathcal{T}° , $x \in \mathbf{N}$ any name, and τ an infinite sequent trace on π such that x is fixed on τ . Then, there is an infinite formula trace on τ on which x is fixed.*

Proof. Let $\tau = (\mathbf{Q}_i\Phi_i)_{i < \omega}$. By Lemma 3.4.2, there is a unique formula φ such that $\varphi^x \in \Phi_0$. For every $n < \omega$, let $F(n)$ be the collection of all formula traces on $\tau_{\leq n}$ on which x is fixed. By Lemma 3.4.8, $F(n) \neq \emptyset$ for every $n < \omega$. Let $F := \bigcup_{n < \omega} F(n)$. Then, (F, \sqsubset) is a tree of height ω each of whose levels is finite, so by König's lemma (F, \sqsubset) has an infinite branch, whence clearly there is an infinite formula trace on τ starting from φ^x and on which x is fixed. ■

Soundness of CTL^*_\circ now follows easily. Informally, given a CTL^*_\circ proof \mathcal{T} of a formula φ , we drop the annotations from \mathcal{T} and show that the resulting tree with back-edges is a $\text{CTL}^*_{\text{reg}}$ proof of φ . Applications of rules other than **iThin**, **eThin** and **del** are preserved when dropping the annotations due to the implicit admissibility of contraction. For example, an instance of rule **AV** from CTL^*_\circ of the form

$$\text{AV} \frac{x \dashv \mathbf{A}\{\varphi^x, \alpha, \beta\}, \mathbf{A}\{\varphi, \alpha \vee \beta\}}{x \dashv \mathbf{A}\{\varphi^x, \alpha \vee \beta\}, \mathbf{A}\{\varphi, \alpha \vee \beta\}}$$

becomes the following preserving instance of rule **AV** from $\text{CTL}^*_{\text{reg}}$:

$$\text{AV} \frac{\mathbf{A}\{\varphi, \alpha, \beta\}, \mathbf{A}\{\varphi, \alpha \vee \beta\}}{\mathbf{A}\{\varphi, \alpha \vee \beta\}}$$

3.4.12. PROPOSITION (Soundness of CTL^*_\circ). *For every CTL^* formula φ , if there is a CTL^*_\circ proof of φ , then φ is valid.*

Proof. Let $\mathcal{T} = (T, <_T, \lambda_T, l \mapsto c_l)$ be a CTL^*_\circ proof of φ . Let $\mathcal{T}_r := (T, <_T, \lambda_r, l \mapsto c_l)$ be given by setting $\lambda_r(u) := \mathbf{b}(\lambda_T(u))$ for every $u \in T$. That is to say, \mathcal{T}_r results from replacing every hypersequent H labelling a vertex of \mathcal{T} by its base $\mathbf{b}(H)$. We claim that \mathcal{T}_r is a $\text{CTL}^*_{\text{reg}}$ proof of φ .

Let us first show that \mathcal{T}_r is a $\text{CTL}^*_{\text{reg}}$ derivation of φ . It suffices to see that instances of CTL^*_\circ rules in \mathcal{T} turn into instances of $\text{CTL}^*_{\text{reg}}$ rules in \mathcal{T}_r . Let $u \in T_r$ be a non-final vertex of \mathcal{T}_r . Then, u is non-final in \mathcal{T} as well. Let **R** be the CTL^*_\circ

rule applied at u in \mathcal{T} . If R is among ELit, EV, E \wedge , EA, EE, EU, ER, AX, EX, iW, eW, then R is applied at u in \mathcal{T}_r . Suppose now that R is among ALit, AV, A \wedge , AA, AE, AU, AR₀, AR₁. Let $\mathbf{Q}\Psi$ and ψ^a be the unique principal sequent and formula, respectively, at u . If there is a side sequent $\mathbf{Q}'\Psi' \in \lambda_{\mathcal{T}}(u)$ such that $\mathbf{b}(\mathbf{Q}'\Psi') = \mathbf{b}(\mathbf{Q}\Psi)$ or a side formula $\chi^b \in \Psi$ such that $\mathbf{b}(\chi^b) = \mathbf{b}(\psi^a)$, then R is applied preservingly at u in \mathcal{T}_r . Otherwise, R is applied discardingly at u in \mathcal{T}_r . Finally, instances of iThin, eThin and del in \mathcal{T} become preserving instances of weakening in \mathcal{T}_r . This establishes that \mathcal{T}_r is a $\text{CTL}_{\text{reg}}^*$ derivation. And, since the roots of \mathcal{T} and \mathcal{T}_r have identical labels, \mathcal{T}_r is a $\text{CTL}_{\text{reg}}^*$ derivation of φ .

Every leaf of \mathcal{T}_r which is not a repeat is clearly axiomatic. And for every repeat $l \in \text{Rep}_{\mathcal{T}_r}$ there is a modal vertex in $[c_l, l]_{\mathcal{T}_r}$. It only remains to see that on every infinite path through \mathcal{T}_r° there is a good sequent trace. Let $\pi_r = (u_n)_{n < \omega}$ be an infinite path on \mathcal{T}_r° . Then, $\pi := (u_n)_{n < \omega} = \pi$ is an infinite path on \mathcal{T}° . By Definition 3.4.4(iii), one of the following holds:

- (i) either there is a name x eventually fixed on π ;
- (ii) or there is a good E-trace on π .

Formulas are never annotated in E-sequents, so if the latter is the case we are done. Assume, then, the former. Let π' be a tail of π such that x is fixed on π' . By Lemma 3.4.10, there is a sequent trace τ' on π' such that x is fixed on τ' . In particular, τ' is an A-trace. By Lemma 3.4.11, there is a formula trace ρ' on τ' such that x is fixed on ρ' (note that \mathcal{T} is by definition discarding). In particular, ρ' is an R-trace. Let $\tau = (\mathbf{Q}_n \Phi_n)_{n < \omega}$ be a sequent trace on π such that τ' is a tail of τ . Similarly, let ρ be a formula trace on τ such that ρ' is a tail of $\rho = (\varphi_n^{a_n})_{n < \omega}$. Existence of τ and ρ is immediate from the definitions of \triangleright and \triangleright . Then, $\tau_r := (\mathbf{b}(\mathbf{Q}_n \Phi_n))_{n < \omega}$ is an A-trace on π_r and $\rho_r := (\mathbf{b}(\varphi_n^{a_n}))_{n < \omega}$ is an R-trace on τ , whence π_r contains a good A-trace.

Therefore, \mathcal{T}_r is a $\text{CTL}_{\text{reg}}^*$ proof of φ , so φ is valid by Proposition 3.3.7. \blacksquare

Now that we have established the soundness of the system CTL_\circ^* , let us, before showing completeness, isolate a finitary correctness condition which suffices to guarantee that a CTL_\circ^* derivation is a proof.

3.4.13. DEFINITION (Universal CTL_\circ^* derivation). A CTL_\circ^* derivation \mathcal{T} is *universal* if for every repeat $l \in \text{Rep}_{\mathcal{T}}$ there is a name fixed on $[c_l, l]_{\mathcal{T}}$. \dashv

Every repeat l in a universal CTL_\circ^* derivation \mathcal{T} has an associated *invariant*, denoted by $\text{inv}(l)$, defined as the shortest sequence of names wx such that wx is a prefix of every control in $[c_l, l]_{\mathcal{T}}$. The existence of invariants follows immediately from Definition 3.4.13 and the fact that, when reading CTL_\circ^* rules bottom-up, new names are always appended to the right end of the controls. The invariant map

inv induces the following (reflexive) quasi-order \preceq on the repeats of \mathcal{T} : $l \preceq l'$ if, and only if, $\text{inv}(l) \sqsubseteq \text{inv}(l')$.

Universality suffices as a replacement of the infinitary condition (iii) in Definition 3.4.4, as the next proposition shows.

3.4.14. PROPOSITION. *Let \mathcal{T} be a universal, discarding CTL_\circ^* derivation such that every non-repeat leaf of \mathcal{T} is axiomatic. Then, \mathcal{T} is a CTL_\circ^* proof.*

Proof. It suffices to show that for every infinite path π through \mathcal{T}° there is a name which is eventually fixed on π .

Let $\pi' := (u_i)_{i < \omega}$ be a tail of π such that each u_i occurs infinitely often on π and $u_0 \in \text{Rep}_{\mathcal{T}}$. We may then write

$$\pi' = l_0 \widehat{[c_0, l_1]}_{\mathcal{T}} \widehat{[c_1, l_2]}_{\mathcal{T}} \widehat{\cdots},$$

where each l_i is a repeat with companion c_i . Note that $l_i \triangleleft l_{i+1}$ for every $i < \omega$. By Proposition 1.2.11, there is some $k \geq 0$ such that $l_k \preceq l_i$ for every $i < \omega$. Let $\text{inv}(l_k) = wx$. By Proposition 1.2.12, wx is a prefix of each control on π' , and thus by Lemma 3.4.7 for every $j < \omega$ there is a sequent trace on (u_0, \dots, u_j) on which x is fixed. Therefore, by Lemma 3.4.10 there exists an infinite sequent trace on π' on which x is fixed. ■

Observe that all conditions imposed on \mathcal{T} in Proposition 3.4.14 are finitary, and that so are the requirements for being a CTL_\circ^* derivation. Therefore, Proposition 3.4.14 isolates a fully finitary fragment of CTL_\circ^* (which is, clearly, incomplete). In Section 3.5 we shall use annotations to find inductive invariants¹⁰ on universal CTL_\circ^* proofs.

3.4.2 Completeness of CTL_\circ^*

We establish the completeness of CTL_\circ^* by annotating CTL_∞^* proofs obtained via proof-search and folding them down to trees with back-edges.

For the remainder of this section, fix an arbitrary well-order $\leq_{\mathbb{N}}$ on \mathbb{N} , a valid formula φ , and a CTL_∞^* proof \mathcal{T} of φ given by Corollaries 3.2.27 and 3.2.43. In particular, \mathcal{T} is a subtree of a CTL_∞^* proof-search tree for φ , say guided by $\mathcal{G} = (\leq_S, \leq_F)$.

Let Γ be a set of annotated formulas and $n < \omega$. A non-annotated formula α is *n-annotated* in Γ if there are pairwise different annotations a_1, \dots, a_n such that $\alpha^{a_i} \in \Gamma$ for every $1 \leq i \leq n$. A non-annotated formula is *twice-annotated* in Γ if it is 2-annotated in Γ .

¹⁰No relation to the invariants given by $\text{inv}(\cdot)$.

Let $\mathbf{Q}\Psi$ and $\mathbf{Q}'\Xi$ be sequents. We say that $\mathbf{Q}\Psi$ and $\mathbf{Q}'\Xi$ are *similar* if the multisets $\{\psi \mid \psi^a \in \Psi\}$ and $\{\xi \mid \xi^a \in \Xi\}$ are identical.

We begin by inductively building a (possibly infinite) tree $\tilde{\mathcal{T}}$ according to the rules of CTL_o^* , together with a function $f : \tilde{\mathcal{T}} \rightarrow \mathcal{T}$ such that the following hold for all $u, v \in \tilde{\mathcal{T}}$:

- (i) if $u <_{\tilde{\mathcal{T}}}^0 v$ and u is not thinning, then $f(u) <_{\mathcal{T}}^0 f(v)$;
- (ii) if $u <_{\tilde{\mathcal{T}}}^0 v$ and u is thinning, then $f(u) = f(v)$;
- (iii) if $f(u) <_{\mathcal{T}}^0 f(v)$, then $u <_{\tilde{\mathcal{T}}}^0 v$;
- (iv) if u has label $\Theta \dashv \Gamma$, then $f(u)$ is labelled by $\mathbf{b}(\Gamma)$.

For the base case, let the root of $\tilde{\mathcal{T}}$ be labelled by the same hypersequent as the one labelling the root of \mathcal{T} , and map it via f to the root of \mathcal{T} . For the inductive case, assume that a vertex $u \in \tilde{\mathcal{T}}$ has been defined, say with label $\Theta \dashv \Gamma$, and that $f(u)$ is defined as well. We proceed as follows:

1. If Γ contains a sequent in which some formula is twice-annotated, then we apply rule \mathbf{iThin} at u with principal sequent $\mathbf{Q}\Phi$ and principal formulas α^a and α^b , where $(\mathbf{Q}\Phi, \alpha, a, b)$ is the $(\leq_S, \leq_F, <_{\Theta}, <_{\Theta})$ -least tuple such that $\mathbf{Q}\Phi \in \Gamma$, $\alpha^a, \alpha^b \in \Phi$, and $a \neq b$. We map the unique immediate successor of u via f to $f(u)$.
2. Else, if Γ contains two different, similar sequents, then we apply rule \mathbf{eThin} at u with principal sequents $\mathbf{Q}_0\Phi_0$ and $\mathbf{Q}_1\Phi_1$, where $(\mathbf{Q}_0\Phi_0, \mathbf{Q}_1\Phi_1)$ is the (\leq_S, \leq_S) -least pair such that $\mathbf{Q}_0\Phi_0, \mathbf{Q}_1\Phi_1 \in \Gamma$, $\mathbf{Q}_0\Phi_0 \neq \mathbf{Q}_1\Phi_1$, and $\mathbf{b}(\mathbf{Q}_0\Phi_0) = \mathbf{b}(\mathbf{Q}_1\Phi_1)$. We map the unique immediate successor of u via f to $f(u)$.
3. Otherwise, by the inductive hypothesis the base function \mathbf{b} induces bijections $\mathbf{b} : \Gamma \rightarrow \mathbf{b}(\Gamma)$ and $\mathbf{b}_{\mathbf{Q}\Phi} : \Phi \rightarrow \mathbf{b}(\Phi)$ for every $\mathbf{Q}\Phi \in \Gamma$, and we distinguish cases according to the rule \mathbf{R} applied at $f(u)$:
 - a) Case $\mathbf{R} = \mathbf{AR}$. Let $\mathbf{A}\Phi \in \Gamma$ and $\alpha^a \in \Phi$ be (unique) such that $\mathbf{b}(\mathbf{A}\Phi)$ and $\mathbf{b}_{\mathbf{A}\Phi}(\alpha^a)$ are the principal sequent and formula, respectively, at $f(u)$. If $a = \varepsilon$, then we apply rule \mathbf{AR}_1 at u with principal sequent $\mathbf{A}\Phi$, principal formula α , and active name the \leq_N -least name not occurring in Θ . Otherwise, we apply rule \mathbf{AR}_0 at u with principal sequent $\mathbf{A}\Phi$ and principal formula α^a .
 - b) In any other case we apply rule \mathbf{R} at u with principal and active sequents and formulas the b^{-1} -images of principal and active sequents and formulas, respectively, at $f(u)$.

In both cases (a) and (b), we map each immediate successor of u via f to the corresponding immediate successor of $f(u)$.

Since the thinning rules are only applied in $\tilde{\mathcal{T}}$ to hypersequents containing pairs of different similar sequents, and to sequents containing twice-annotated formulas, the following is an immediate consequence of (i), (ii) and (iv):

3.4.15. LEMMA. *For every finite or infinite branch $\tilde{\pi} = (u_n)_{n < N \leq \omega}$ of $\tilde{\mathcal{T}}$, there is an $N' \leq N$ and a strictly increasing sequence of natural numbers $(n_i)_{i < N'}$ such that $\pi := (f(u_{n_i}))_{i < N'}$ is a branch of \mathcal{T} . Moreover, π is infinite if, and only if, so is $\tilde{\pi}$.*

Abusing notation, for every infinite branch $\tilde{\pi}$ of $\tilde{\mathcal{T}}$ we denote by $f(\tilde{\pi})$ the corresponding infinite branch π of \mathcal{T} given by Lemma 3.4.15.

As a consequence of the definition of traces in the thinning rules, we additionally have:

3.4.16. LEMMA. *For every infinite branch $\tilde{\pi}$ of $\tilde{\mathcal{T}}$, every infinite sequent trace $\tau = (\mathbf{Q}_n \Phi_n)_{n < \omega}$ on $f(\tilde{\pi})$, and every infinite formula trace $\rho = (\varphi_n)_{n < \omega}$ on τ , there is a sequent trace $\tilde{\tau} = (\mathbf{Q}'_n \Psi_n)_{n < \omega}$ on $\tilde{\pi}$, a formula trace $\tilde{\rho} = (\psi_n^{a_n})_{n < \omega}$ on $\tilde{\tau}$, and a strictly increasing sequence of natural numbers $0 = n_0 < n_1 < \dots$ such that the following hold for all $i < \omega$ and $n_i \leq k < n_{i+1}$:*

$$(i) \quad \mathbf{b}(\mathbf{Q}'_k \Psi_k) = \mathbf{Q}_i \Phi_i;$$

$$(ii) \quad \psi_k = \varphi_k.$$

Abusing notation, for every infinite branch $\tilde{\pi}$ of $\tilde{\mathcal{T}}$, every infinite sequent trace τ on $f(\tilde{\pi})$ and every infinite formula trace ρ on τ , we denote by $f^{-1}(\tau)$ and $f^{-1}(\rho)$ the corresponding traces given by Lemma 3.4.16. Informally, $f^{-1}(\tau)$ and $f^{-1}(\rho)$ are the traces that result from τ and ρ after the addition of annotations (and instances of the thinning rules) to \mathcal{T} in the construction of $\tilde{\mathcal{T}}$.

3.4.17. OBSERVATION. Lemmas 3.4.15 and 3.4.16 provide a back-and-forth correspondence between \mathcal{T} and $\tilde{\mathcal{T}}$: to every infinite branch π of \mathcal{T} there corresponds an infinite branch $f(\pi)$ of $\tilde{\mathcal{T}}$, and then every trace on $f(\pi)$ yields a corresponding trace on π .

By inspection of the rules of CTL_\circ^* and the priority given to the thinning rules in the construction of $\tilde{\mathcal{T}}$, we have:

3.4.18. LEMMA. *Let $\Theta \dashv \Gamma$ be any hypersequent labelling a vertex in $\tilde{\mathcal{T}}$. The following hold:*

(i) *no formula is 3-annotated in Γ ;*

(ii) *there is no triple of pairwise different, pairwise similar sequents in Γ .*

3.4.19. PROPOSITION. *Only finitely many pairwise different names occur in $\tilde{\mathcal{T}}$.*

Proof. Let $n := |\varphi|$, and for every $u \in T$ let $\Theta_u \dashv \Gamma_u$ be the label of u . We define an equivalence relation \sim_u on the annotated sequents in Γ_u by setting $\mathbf{Q}\Phi \sim_u \mathbf{Q}'\Psi$ if, and only if, $\mathbf{Q}\Phi$ and $\mathbf{Q}'\Psi$ are similar. The \sim_u -equivalence class of a sequent is completely determined by its type and its underlying non-annotated multiset. Therefore, by Lemma 3.4.18(i) there are at most $2 \cdot 2^{2n} = 2^{2n+1}$ many different \sim_u -classes, and thus by (ii) of Lemma 3.4.18 the hypersequent Γ_u contains at most $2 \cdot 2^{2n+1} = 2^{2n+2}$ many pairwise different (annotated) sequents. Again by Lemma 3.4.18(i), each sequent in Γ_u contains at most $2n$ -many pairwise different annotated formulas. Therefore, any hypersequent labelling a vertex of $\tilde{\mathcal{T}}$ contains no more than $N := 2n \cdot 2^{2n+2} = n \cdot 2^{2n+3}$ many pairwise different names. By the choice of new names when applying \mathbf{AR}_1 in $\tilde{\mathcal{T}}$, it follows that the names occurring in any hypersequent in $\tilde{\mathcal{T}}$ are all among the first N -many names with respect to \leq_N . ■

3.4.20. COROLLARY. *Only finitely many pairwise different annotated hypersequents occur in $\tilde{\mathcal{T}}$.*

A *demotion* is a formula trace of the form $\alpha^x \triangleright \alpha^y$ with $x \neq y$. We say that x is *demoted* to y . Observe that demotions are only due to thinning.

The following ensures that the management of the annotations on $\tilde{\mathcal{T}}$ detects good \mathbf{A} -traces.

3.4.21. LEMMA. *Let π be an infinite branch of $\tilde{\mathcal{T}}$. If there is a good \mathbf{A} -trace τ on $f(\pi)$, then some name is eventually fixed on $f^{-1}(\tau)$.*

Proof. Since τ is a good \mathbf{A} -trace, there is some infinite \mathbf{R} -trace ρ on τ . Let $\pi = (u_i)_{i < \omega}$, $\tau' := f^{-1}(\tau) = (\mathbf{Q}_i \Psi_i)_{i < \omega}$, and $\rho' := f^{-1}(\rho) = (\psi_i^{a_i})_{i < \omega}$. Since ρ is of type \mathbf{R} , there is an \mathbf{R} -formula ψ and some $n < \omega$ such that $\psi_i \in \{\psi, \mathbf{X}\psi\}$ for every $i \geq n$.

Note that there are infinitely many $i \geq n$ such that rule \mathbf{AR}_0 or \mathbf{AR}_1 is applied at u_i with principal formula ψ^{a_i} , and that by construction of $\tilde{\mathcal{T}}$ each such application produces a formula of the form $\mathbf{X}\psi^y$, where $y \in \mathbf{N}$. The only way for y not to remain fixed on $\rho'_{>i}$ is to encounter an instance of \mathbf{iThin} or \mathbf{eThin} which demotes y to another name y' . So it suffices to show that ρ' only passes through finitely many demotions.

Suppose that $n < i < j$ are such that $\psi_i^{a_i} \triangleright \psi_{i+1}^{a_{i+1}}$ and $\psi_j^{a_j} \triangleright \psi_{j+1}^{a_{j+1}}$ are demotions and that ρ' encounters no demotion in between those two. Let \mathbf{R}_i and \mathbf{R}_j be the (thinning) rules applied at u_i and u_j , respectively. There are four possible cases. However, the effect that \mathbf{iThin} and \mathbf{eThin} have on non-empty annotations is

the same: either the name is unchanged (no demotion), or it is replaced by a name occurring strictly earlier in the control (demotion). We thus show only the case $R_i = R_j = \text{iThin}$; the other three have virtually identical proofs.

Assume, then, that $R_i = R_j = \text{iThin}$. Let $a_i = y$, $a_{i+1} = a_j = y'$, and $a_{j+1} = y''$. We then have:

$$\begin{array}{c} \text{iThin} \frac{\Theta''' \vdash \mathbf{A}\{\Phi', \psi_{j+1}^{y''}\}, \Delta'}{\Theta'' \vdash \mathbf{A}\{\Phi', \psi_j^{y'}, \psi_{j+1}^{y''}\}, \Delta'} \\ \vdots \\ \text{iThin} \frac{\Theta' \vdash \mathbf{A}\{\Phi, \psi_{i+1}^{y'}\}, \Delta}{\Theta \vdash \mathbf{A}\{\Phi, \psi_i^y, \psi_{i+1}^{y'}\}, \Delta} \end{array}$$

As $y' \prec_\Theta y$, we have $\Theta = ay'byc$. Analogously, $\Theta'' = a'y''b'y'c'$ because $y'' \prec_{\Theta''} y'$. By assumption, there is no demotion in between those two with principal formula in ρ' , so y' is fixed on $(\rho'(i+1), \dots, \rho'(j))$ and thus $|a'| < |a'y''b'| \leq |a|$ because new names are always appended to the right end of controls (reading the rules bottom-up).

Therefore, ρ' encounters only finitely many demotions. ■

Completeness of CTL_\circ^* follows now easily.

3.4.22. PROPOSITION (Completeness of CTL_\circ^*). *For any formula φ , if $\text{CTL}^* \models \varphi$, then $\text{CTL}_\circ^* \vdash \varphi$.*

Proof. Let \mathcal{T} be a CTL_∞^* proof of φ given by Corollaries 3.2.27 and 3.2.43, and let $\tilde{\mathcal{T}}$ and $f: \tilde{T} \rightarrow T$ be given as described above.

Note that \mathcal{T} is: discarding, by Definition 3.2.24; and regular and deterministic, by (the proof of) Proposition 3.3.2. Therefore, by construction $\tilde{\mathcal{T}}$ is also discarding, regular and deterministic.

Let π be any infinite branch of $\tilde{\mathcal{T}}$. By Proposition 3.2.25, Lemma 3.4.15, and Corollary 3.4.20, there is a vertex $v \in \pi$ such that for some $u <_{\tilde{\mathcal{T}}} v$ we have:

- (i) u and v have identical labels;
- (ii) there is a modal vertex in $[u, v)_{\tilde{\mathcal{T}}}$.

Let \mathcal{T}_r be the resulting tree with back-edges after pruning each infinite branch of \mathcal{T} at the encounter of the first such vertex v on the branch (keeping v as a repeat in \mathcal{T}_r). We claim that \mathcal{T}_r is a CTL_\circ^* proof of φ .

By the choice of the pruning vertices, for every repeat $l \in \text{Rep}_{\mathcal{T}_r}$ there is a modal vertex in $[c_l, l)_{\mathcal{T}_r}$. This establishes that \mathcal{T}_r is a CTL_\circ^* derivation (of φ).

By Lemma 3.4.15, every non-repeat leaf of \mathcal{T}_r is axiomatic.

Let now $\pi = (u_n)_{n < \omega}$ be an infinite path through \mathcal{T}_r° . Since $\tilde{\mathcal{T}}$ is deterministic, we can inductively build an infinite branch $\pi' = (u'_n)_{n < \omega}$ of $\tilde{\mathcal{T}}$ such that u_n and u'_n are labelled by the same instance of the same CTL_\circ^* rule, for every $n < \omega$. Let u'_0 be the root of $\tilde{\mathcal{T}}$, which has the same label as u_0 . For the inductive case, suppose that u'_n has been defined. Since any two vertices in $\tilde{\mathcal{T}}$ labelled by the same hypersequent are labelled by the same instance of the same CTL_\circ^* rule, the inductive hypothesis yields an immediate successor v of u'_n in $\tilde{\mathcal{T}}$ labelled by the same hypersequent as u_{n+1} and such that $\triangleright_{u_n}^{u_{n+1}} = \triangleright_{u'_n}^v$ and $\triangleright_{u_n}^{u_{n+1}} = \triangleright_{u'_n}^v$. We then let $u'_{n+1} := v$ and are done.

Consider $f(\pi')$, which by Lemma 3.4.15 is an infinite branch of \mathcal{T} . Since \mathcal{T} is a CTL_∞^* proof, there is a good sequent trace τ on $f(\pi')$. By Lemma 3.4.16, $f^{-1}(\tau)$ is an infinite sequent trace on π' , whence by construction of π' we conclude that $f^{-1}(\tau)$ is an infinite sequent trace on π as well. If τ is of type E, then $f^{-1}(\tau)$ is a good E-trace on π . And if τ is of type A, then by Lemma 3.4.21 there is a name eventually fixed on $f^{-1}(\tau)$.

This establishes that \mathcal{T}_r is indeed a CTL_\circ^* proof (of φ). ■

Combining Propositions 3.4.12 and 3.4.22, we have:

3.4.23. THEOREM. *A CTL* formula φ is valid if, and only if, there is a CTL_\circ^* proof of φ .*

3.5 Cycle elimination

In this section we show how to turn the cycles of (some) universal CTL_\circ^* proofs into acyclic, inductive arguments.¹¹ As explained in the introduction of the dissertation, cyclic proofs replace non-invertible inductive rules with invertible ones such as unfoldings of fixpoints. The reader may then wonder why anyone would be interested in the removal of cycles. One potential application of cycle elimination, and the one driving this section, is assisting with proving completeness of Hilbert-style calculi.

The problem of axiomatising CTL^* remained open for several years until Reynolds presented a sound and complete axiom system for it in [115]. The completeness proof is intricate and the system involves a rule with a considerably complicated side condition. We shall not be able to obtain a better axiomatisation here. Not only because we lack a ‘nice’, fully finitary cyclic calculus for the logic,

¹¹I am thankful to Sebastian Enqvist for suggesting this line of research as a continuation of the work in [6].

but also because finding adequate induction rules for E-sequents seems to be quite a difficult task.

We use annotations to find an inductive invariant within a cycle. A Park-style induction rule then ‘packs’ the invariant, dualised, inside an R-formula. It lays there until it is extracted at the repeat on the cycle, thereby turning the repeat into an instance of the law of the excluded middle. Essentially the same strategy is followed in [19] for LTL and CTL, and a similar one is found in [139], where the authors introduce two inductive systems for the first-order μ -calculus and then define translations between them. The idea, in fact, traces back at least to Kozen’s [81].

In the end, we arrive at a Hilbert-style system and compare it to a fragment of Reynolds’s axiomatisation.

3.5.1 Inductive acyclic proofs

Here we present the acyclic, ordinary sequent calculus $\text{CTL}_{\text{ind}}^*$ into which we shall later embed certain universal CTL_{\circ}^* proofs. Its rules are given in Figures 3.10 to 3.12. The system is essentially CTL_{∞}^* but with an inductive rule **ind** in place of the fixpoint unfolding rules **AR**, **EU** and **ER**. Note that we still need rule **AU** to characterise $\alpha\text{U}\beta$ as the least pre-fixpoint of (the function represented by) $\beta \vee [\alpha \wedge \text{X}(\cdot)]$. In addition, $\text{CTL}_{\text{ind}}^*$ contains an internal, formula-level cut rule for A-sequents. Cuts can also be performed at the external level by means of rule \perp^{-1} , which, when read bottom-up, introduces the sequent $\text{A}\emptyset$ ($\equiv \text{A}\vee\emptyset \equiv \perp$).

3.5.1. DEFINITION ($\text{CTL}_{\text{ind}}^*$ derivation). A $\text{CTL}_{\text{ind}}^*$ *derivation* of a formula φ is a finite tree \mathcal{T} built according to the rules in Figures 3.10 to 3.11 and such that the root of \mathcal{T} has label $\text{A}\varphi$. ⊣

By $\text{CTL}_{\text{ind}}^{*\,-}$ we denote the cut-free system $\text{CTL}_{\text{ind}}^* - \text{cut}$.

3.5.2. DEFINITION ($\text{CTL}_{\text{ind}}^*$ proof). A $\text{CTL}_{\text{ind}}^*$ *proof* of a formula φ is a $\text{CTL}_{\text{ind}}^*$ derivation of φ each of whose leaves is axiomatic. ⊣

Soundness of $\text{CTL}_{\text{ind}}^*$ follows almost immediately from inspection of the rules. We only need to make sure that rule **ind** is sound.¹²

3.5.3. LEMMA. *For every path σ through a serial model and all formulas α , β and γ , if $\sigma \not\models \alpha\text{R}\beta$ and $\sigma \not\models \gamma$, then there is a $k < \omega$ such that $(\sigma, k) \not\models \gamma$ and either $(\sigma, k) \not\models \beta$, or $(\sigma, k) \not\models \alpha \vee \text{X}\gamma^\partial$.*

¹²A rule is *sound* if its conclusion is valid whenever all premises are.

$$\begin{array}{l}
\text{ax} \frac{}{Q\rho, Q'\bar{\rho}, \Delta} \\
\text{ax}_{\text{AE}} \frac{}{Q\Phi, (Q\Phi)^\partial, \Delta} \\
\text{ALit} \frac{A\Phi, A\ell, \Delta}{A\{\Phi, \ell\}, \Delta} \\
A\vee \frac{A\{\Phi, \alpha, \beta\}, \Delta}{A\{\Phi, \alpha \vee \beta\}, \Delta} \\
A\wedge \frac{A\{\Phi, \alpha\}, \Delta \quad A\{\Phi, \beta\}, \Delta}{A\{\Phi, \alpha \wedge \beta\}, \Delta} \\
\text{ax}_\top \frac{}{E\emptyset, \Delta} \\
\text{ax}_{\text{RU}} \frac{}{A\{\Phi, \alpha R\beta, (\alpha R\beta)^\partial\}, \Delta} \\
\text{ELit} \frac{E\Phi, \Delta \quad E\lambda, \Delta}{E\{\Phi, \lambda\}, \Delta} \\
E\vee \frac{E\{\Phi, \alpha\}, E\{\Phi, \beta\}, \Delta}{E\{\Phi, \alpha \vee \beta\}, \Delta} \\
E\wedge \frac{E\{\Phi, \alpha, \beta\}, \Delta}{E\{\Phi, \alpha \wedge \beta\}, \Delta}
\end{array}$$

FIGURE 3.10: Some of the rules of the system $\text{CTL}_{\text{ind}}^*$.

Proof. By induction on $j < \omega$, we show that if j is the least natural number such that $(\sigma, j) \not\models \beta$ and $(\sigma, i) \not\models \alpha$ for every $i < j$, then there is a $k < \omega$ such that $(\sigma, k) \not\models \gamma$ and either $(\sigma, k) \not\models \beta$, or $(\sigma, k) \not\models \alpha \vee X\gamma^\partial$. Note that such a least j must exist because $\sigma \not\models \alpha R\beta$.

The base case $j = 0$ is clear by the assumptions on σ . For the inductive case, assume that $j > 0$ and that the claim holds for every $j' < j$. By the minimality of j , then, $\sigma \models \beta$, so $\sigma \not\models \alpha$ and $\sigma \not\models X(\alpha R\beta)$. If $\sigma \not\models X\gamma^\partial$ we are done. Otherwise, we have $(\sigma, 1) \not\models \alpha R\beta$ and $(\sigma, 1) \not\models \gamma$, so the claim follows by the inductive hypothesis for $j - 1$ applied to $(\sigma, 1)$. ■

3.5.4. COROLLARY. *The rule ind of system $\text{CTL}_{\text{ind}}^*$ is sound.*

Proof. Let σ such that $\sigma \not\models A\{\Phi, \alpha R\beta\}, \Delta$, and let $\sigma \sim \sigma'$ be such that $\sigma' \not\models \vee \Phi \vee \alpha R\beta$. If $\sigma' \not\models \beta$, then $\sigma \not\models A\{\Phi, \beta \wedge [\alpha \vee X(\vee \Phi \vee \Delta^\partial)^\partial]\}, \Delta$ and we are done. Otherwise, $\sigma' \not\models \alpha$ and $\sigma' \not\models X(\alpha R\beta)$. If $\sigma' \not\models X(\vee \Phi \vee \Delta^\partial)^\partial$ we are done, so assume otherwise. The claim then follows by applying Lemma 3.5.3 to $(\sigma', 1)$ with $\gamma := \vee \Phi \vee \Delta^\partial$. ■

We have thus established:

3.5.5. PROPOSITION (Soundness of $\text{CTL}_{\text{ind}}^*$). *For any formula α , if $\text{CTL}_{\text{ind}}^* \vdash \alpha$, then $\text{CTL}^* \models \alpha$.*

The following internal law of the excluded middle for $\text{CTL}_{\text{ind}}^*$ can be shown by a straightforward structural induction on α :

$$\begin{array}{c}
\text{AA} \frac{A\Phi, A\{\alpha\}, \Delta}{A\{\Phi, A\alpha\}, \Delta} \\
\text{AE} \frac{A\Phi, E\{\alpha\}, \Delta}{A\{\Phi, E\alpha\}, \Delta} \\
\text{AX} \frac{A\Phi, E\Psi_1, \dots, E\Psi_m}{A(X\Phi \cup \Xi), EX\Psi_1, \dots, EX\Psi_m, \Sigma} \\
\text{iW} \frac{A\Phi, \Delta}{A(\Phi \cup \Psi), \Delta}
\end{array}
\qquad
\begin{array}{c}
\text{EA} \frac{E\Phi, \Delta \quad A\{\alpha\}, \Delta}{E\{\Phi, A\alpha\}, \Delta} \\
\text{EE} \frac{E\Phi, \Delta \quad E\{\alpha\}, \Delta}{E\{\Phi, E\alpha\}, \Delta} \\
\text{EX} \frac{E\Psi_1, \dots, E\Psi_m}{EX\Psi_1, \dots, EX\Psi_m, \Sigma} \\
\text{eW} \frac{\Gamma}{\Gamma, \Delta}
\end{array}$$

FIGURE 3.11: More rules of the system $\text{CTL}_{\text{ind}}^*$.

3.5.6. LEMMA. *For every formula α , finite set of formulas Φ , and finite set of sequents Δ , we have $\text{CTL}_{\text{ind}}^* \vdash A\{\Phi, \alpha, \alpha^\partial\}, \Delta$.*

Observe that when read bottom-up, the induction rule ind , in contrast to the unfolding rule AR , does not preserve $\alpha R \beta$ in the premise under the scope of X . This makes it inconvenient for the proof transformations involved in the elimination of cycles. We shall thus utilise a stronger induction rule ind_s which turns out to be admissible in $\text{CTL}_{\text{ind}}^*$. To see this, though, we first need to establish a few properties of $\text{CTL}_{\text{ind}}^*$.

3.5.7. LEMMA. *For every finite set of formulas Φ and every finite set of sequents Δ , we have $\text{CTL}_{\text{ind}}^* \vdash A\{\Phi, (\bigvee \Phi \vee \Delta^\#)^\partial\}, \Delta$.*

Proof. We start by repeatedly applying rule $A\wedge$ in order to break down the conjunction $(\bigvee \Phi \vee \Delta^\#)^\partial$, obtaining in the end all the sequents in $\{A\{\Phi, \varphi^\partial\}, \Delta \mid \varphi \in \Phi\}$, $\{A\{\Phi, (A \vee \Psi)^\partial\}, \Delta \mid A\Psi \in \Delta\}$, and $\{A\{\Phi, (E \wedge \Psi)^\partial\}, \Delta \mid E\Psi \in \Delta\}$. The ones in the first collection are all provable in $\text{CTL}_{\text{ind}}^*$ by Lemma 3.5.6. The rest can be proved as follows:

$$\begin{array}{c}
\text{axAE} \frac{A\Phi, E\Psi^\partial, A\Psi, \Delta}{A\Phi, E\{\wedge \Psi^\partial\}, A\Psi, \Delta} \\
\text{E}\wedge \frac{\text{axAE} \frac{A\Phi, E\Psi^\partial, A\Psi, \Delta}{A\Phi, E\{\wedge \Psi^\partial\}, A\Psi, \Delta}}{A\{\Phi, (A \vee \Psi)^\partial\}, A\Psi, \Delta} \\
\text{AE} \frac{\text{axAE} \frac{A\Phi, E\Psi^\partial, A\Psi, \Delta}{A\Phi, E\{\wedge \Psi^\partial\}, A\Psi, \Delta}}{A\{\Phi, (A \vee \Psi)^\partial\}, A\Psi, \Delta}
\end{array}
\qquad
\begin{array}{c}
\text{axAE} \frac{A\Phi, A\Psi^\partial, E\Psi, \Delta}{A\Phi, A\{\vee \Psi^\partial\}, E\Psi, \Delta} \\
\text{A}\vee \frac{\text{axAE} \frac{A\Phi, A\Psi^\partial, E\Psi, \Delta}{A\Phi, A\{\vee \Psi^\partial\}, E\Psi, \Delta}}{A\{\Phi, (E \wedge \Psi)^\partial\}, E\Psi, \Delta} \\
\text{AA} \frac{\text{axAE} \frac{A\Phi, A\Psi^\partial, E\Psi, \Delta}{A\Phi, A\{\vee \Psi^\partial\}, E\Psi, \Delta}}{A\{\Phi, (E \wedge \Psi)^\partial\}, E\Psi, \Delta}
\end{array}$$

■

In what follows, for brevity we shall simply indicate invocations of Lemma 3.5.6 by writing ‘LEM’ as a rule name.

$$\begin{array}{c}
\text{AU} \frac{A\{\Phi, \beta \vee [\alpha \wedge X(\alpha U \beta)]\}, \Delta}{A\{\Phi, \alpha U \beta\}, \Delta} \quad \text{ind} \frac{A\{\Phi, \beta \wedge [\alpha \vee X(\vee \Phi \vee \Delta^\#)^\partial]\}, \Delta}{A\{\Phi, \alpha R \beta\}, \Delta} \\
\\
\perp^{-1} \frac{\Gamma, A\emptyset}{\Gamma} \quad \text{cut} \frac{A\{\Phi, \alpha\}, \Delta \quad A\{\Phi, \alpha^\partial\}, \Delta}{A\Phi, \Delta}
\end{array}$$

FIGURE 3.12: The rest of the rules of the system $\text{CTL}_{\text{ind}}^*$.

3.5.8. PROPOSITION. *The rule*

$$\text{VR} \frac{A\{\Phi, \gamma, \alpha R \beta\}, \Delta}{A\{\Phi, (\gamma \vee \alpha)R(\gamma \vee \beta)\}, \Delta}$$

is admissible in $\text{CTL}_{\text{ind}}^$.*

Proof. We simulate VR as follows. First, we cut against $\delta := \gamma \vee \alpha R \beta$:

$$\text{cut} \frac{A\{\Phi, (\gamma \vee \alpha)R(\gamma \vee \beta), \delta^\partial\}, \Delta \quad \text{iW} \frac{A\{\Phi, \gamma, \alpha R \beta\}, \Delta}{A\{\Phi, \gamma \vee \alpha R \beta\}, \Delta}}{A\{\Phi, (\gamma \vee \alpha)R(\gamma \vee \beta)\}, \Delta}$$

At the right branch there is nothing to do. At the left branch we proceed as follows:

$$\begin{array}{c}
\text{ind} \frac{\text{\#1} \frac{A\{(\gamma \vee \beta) \wedge [\gamma \vee \alpha \vee X(\gamma^\partial)^\partial], \gamma^\partial\}}{A\{(\gamma \vee \alpha)R(\gamma \vee \beta), \gamma^\partial\}} \quad \text{\#2} \frac{A\{(\gamma \vee \alpha)R(\gamma \vee \beta), \alpha^\partial U \beta^\partial\}}{A\{(\gamma \vee \alpha)R(\gamma \vee \beta), \gamma^\partial \wedge \alpha^\partial U \beta^\partial\}}}{A\{(\gamma \vee \alpha)R(\gamma \vee \beta), \gamma^\partial \wedge \alpha^\partial U \beta^\partial\}} \\
\text{iW} \frac{A\{(\gamma \vee \alpha)R(\gamma \vee \beta), \gamma^\partial \wedge \alpha^\partial U \beta^\partial\}}{A\{\Phi, (\gamma \vee \alpha)R(\gamma \vee \beta), \gamma^\partial \wedge \alpha^\partial U \beta^\partial\}} \\
\text{eW} \frac{A\{\Phi, (\gamma \vee \alpha)R(\gamma \vee \beta), \gamma^\partial \wedge \alpha^\partial U \beta^\partial\}}{A\{\Phi, (\gamma \vee \alpha)R(\gamma \vee \beta), \gamma^\partial \wedge \alpha^\partial U \beta^\partial\}, \Delta}
\end{array}$$

Continuation from #1:

$$\begin{array}{c}
\text{LEM} \frac{}{A\{\gamma, \beta, \gamma^\partial\}} \quad \text{LEM} \frac{}{A\{\gamma, \alpha, X(\gamma^\partial)^\partial, \gamma^\partial\}} \\
\text{AV} \frac{A\{\gamma, \beta, \gamma^\partial\}}{A\{(\gamma \vee \beta), \gamma^\partial\}} \quad \text{AV} \frac{A\{\gamma, \alpha, X(\gamma^\partial)^\partial, \gamma^\partial\}}{A\{\gamma \vee \alpha \vee X(\gamma^\partial)^\partial, \gamma^\partial\}} \\
\text{A}\wedge \frac{A\{(\gamma \vee \beta), \gamma^\partial\} \quad A\{\gamma \vee \alpha \vee X(\gamma^\partial)^\partial, \gamma^\partial\}}{A\{(\gamma \vee \beta) \wedge [\gamma \vee \alpha \vee X(\gamma^\partial)^\partial], \gamma^\partial\}}
\end{array}$$

Continuation from #2:

$$\text{AV} \frac{\text{AV} \frac{\text{A}\{\gamma, \beta, \alpha^\partial \text{U}\beta^\partial\}}{\text{A}\{\gamma \vee \beta, \alpha^\partial \text{U}\beta^\partial\}} \quad \text{AV} \frac{\text{A}\{\gamma, \alpha, \text{X}(\alpha^\partial \text{U}\beta^\partial)^\partial, \alpha^\partial \text{U}\beta^\partial\}}{\text{A}\{(\gamma \vee \alpha) \vee \text{X}(\alpha^\partial \text{U}\beta^\partial)^\partial, \alpha^\partial \text{U}\beta^\partial\}}}{\text{A}\{(\gamma \vee \beta) \wedge [(\gamma \vee \alpha) \vee \text{X}(\alpha^\partial \text{U}\beta^\partial)^\partial], \alpha^\partial \text{U}\beta^\partial\}} \quad \text{ind} \frac{\text{A}\{(\gamma \vee \alpha) \text{R}(\gamma \vee \beta), \alpha^\partial \text{U}\beta^\partial\}}{\text{A}\{(\gamma \vee \alpha) \text{R}(\gamma \vee \beta), \alpha^\partial \text{U}\beta^\partial\}}$$

Continuation from #3:

$$\text{AU} \frac{\text{LEM} \frac{\text{A}\{\gamma, \beta, \beta^\partial, \alpha^\partial\}}{\text{A}\{\gamma, \beta, \beta^\partial, \alpha^\partial\}} \quad \text{LEM} \frac{\text{A}\{\gamma, \beta, \beta^\partial, \text{X}(\alpha^\partial \text{U}\beta^\partial)^\partial\}}{\text{A}\{\gamma, \beta, \beta^\partial, \text{X}(\alpha^\partial \text{U}\beta^\partial)^\partial\}}}{\text{A}\{\gamma, \beta, \alpha^\partial \text{U}\beta^\partial\}}$$

Continuation from #4:

$$\text{AU} \frac{\text{LEM} \frac{\text{A}\{\gamma, \alpha, \beta^\partial, \alpha^\partial, \text{X}(\alpha^\partial \text{U}\beta^\partial)^\partial\}}{\text{A}\{\gamma, \alpha, \beta^\partial, \alpha^\partial, \text{X}(\alpha^\partial \text{U}\beta^\partial)^\partial\}} \quad \text{A}\{\gamma, \alpha, \beta^\partial, \text{X}(\alpha^\partial \text{U}\beta^\partial)^\partial, \text{X}(\alpha^\partial \text{U}\beta^\partial)^\partial\}}{\text{A}\{\gamma, \alpha, \text{X}(\alpha^\partial \text{U}\beta^\partial)^\partial, \alpha^\partial \text{U}\beta^\partial\}}$$

Continuation from #5:

$$\text{AX} \frac{\text{LEM} \frac{\text{A}\{\alpha \text{R}\beta, \alpha^\partial \text{U}\beta^\partial\}}{\text{A}\{\alpha \text{R}\beta, \alpha^\partial \text{U}\beta^\partial\}}}{\text{A}\{\gamma, \alpha, \beta^\partial, \text{X}(\alpha^\partial \text{U}\beta^\partial)^\partial, \text{X}(\alpha^\partial \text{U}\beta^\partial)^\partial\}}$$

■

3.5.9. PROPOSITION. *The strong induction rule*

$$\text{ind}_s \frac{\text{A}\{\Phi, \beta\}, \Delta \quad \text{A}\{\Phi, \alpha, \text{X}[(\xi \rightarrow \alpha) \text{R}(\xi \rightarrow \beta)]\}, \Delta}{\text{A}\{\Phi, \alpha \text{R}\beta\}, \Delta},$$

where $\xi := \vee \Phi \vee \Delta^\sharp$, is admissible in $\text{CTL}_{\text{ind}}^*$.

Proof. We simulate ind_s as follows. First, we cut against $(\xi \rightarrow \alpha) \text{R}(\xi \rightarrow \beta)$:

$$\text{iW} + \text{eW} \frac{\text{cut} \frac{\text{A}\{(\xi \wedge \alpha^\partial) \text{U}(\xi \wedge \beta^\partial), \alpha \text{R}\beta\}}{\text{A}\{\Phi, (\xi \wedge \alpha^\partial) \text{U}(\xi \wedge \beta^\partial), \alpha \text{R}\beta\}, \Delta} \quad \text{A}\{\Phi, (\xi \rightarrow \alpha) \text{R}(\xi \rightarrow \beta), \alpha \text{R}\beta\}, \Delta}{\text{A}\{\Phi, \alpha \text{R}\beta\}, \Delta} \quad \text{#1} \quad \text{#2}}$$

Continuation from #2:

$$\text{ind} \frac{\text{AV} \frac{\text{iW} \frac{A\{\Phi, \xi^\partial\}, \Delta}{A\{\Phi, \xi^\partial, \beta, \alpha R\beta\}, \Delta}}{A\{\Phi, \xi \rightarrow \beta, \alpha R\beta\}, \Delta}}{A\{\Phi, (\xi \rightarrow \alpha)R(\xi \rightarrow \beta), \alpha R\beta\}, \Delta}}{\text{AV} + \text{iW} \frac{A\{\Phi, \xi^\partial\}, \Delta}{A\{\Phi, \xi \rightarrow \alpha, X[\xi^\partial \wedge (\alpha R\beta)^\partial], \alpha R\beta\}, \Delta}}$$

By Lemma 3.5.7, $\text{CTL}_{\text{ind}}^* \vdash A\{\Phi, \xi^\partial\}, \Delta$.

Continuation from #1, where $\gamma := (\xi \wedge \alpha^\partial)U(\xi \wedge \beta^\partial)$:

$$\text{AU} \frac{\text{iW} \frac{A\{\xi \wedge \beta^\partial, \beta\}}{A\{\xi \wedge \alpha^\partial, \xi \wedge \beta^\partial, \beta\}}}{\text{AV} + \text{AQ} \frac{A\{\Phi, \beta\}, \Delta}{A\{\xi, \beta\}} \text{LEM} \frac{A\{\beta^\partial, \beta\}}{A\{\beta^\partial, \beta\}}} \frac{\text{iW} \frac{A\{\xi \wedge \beta^\partial, \beta\}}{A\{\xi \wedge \beta^\partial, X\gamma, \beta\}}}{A\{\gamma, \beta\}} \quad \#3}{\text{ind} + A\wedge \frac{A\{\gamma, \beta\}}{A\{\gamma, \alpha R\beta\}} \quad A\{\gamma, \alpha, X\gamma^\partial\}}$$

Note that we reach $A\{\xi \wedge \beta^\partial, \beta\}$ twice, so we omit one of the derivations.

Continuation from #3:

$$\text{AU} \frac{\text{AX} \frac{\text{LEM} \frac{A\{\gamma, \gamma^\partial\}}{A\{\gamma, \gamma^\partial\}}}{A\{\xi \wedge \beta^\partial, X\gamma, \alpha, X\gamma^\partial\}}}{A\{\xi \wedge \alpha^\partial, \xi \wedge \beta^\partial, \alpha, X\gamma^\partial\}} \quad \#4}{A\{\gamma, \alpha, X\gamma^\partial\}}$$

Continuation from #4:

$$A\wedge \frac{A\{\xi, \alpha, X\gamma^\partial\}}{A\{\xi, \xi \wedge \beta^\partial, \alpha, X\gamma^\partial\}} \frac{\text{AV} + \text{AQ} \frac{A\{\Phi, \alpha, X\gamma^\partial\}, \Delta}{A\{\xi, \alpha, X\gamma^\partial\}}}{\text{iW} \frac{A\{\xi, \beta^\partial, \alpha, X\gamma^\partial\}}{A\{\xi, \beta^\partial, \alpha, X\gamma^\partial\}}} \text{LEM} \frac{A\{\alpha^\partial, \xi \wedge \beta^\partial, \alpha, X\gamma^\partial\}}{A\{\xi \wedge \alpha^\partial, \xi \wedge \beta^\partial, \alpha, X\gamma^\partial\}}$$

Note that we reach $A\{\xi, \alpha, X\gamma^\partial\}$ twice, so we omit one of the derivations. Since $A\{\Phi, \alpha, X\gamma^\partial\}, \Delta = A\{\Phi, \alpha, X[(\xi \rightarrow \alpha)R(\xi \rightarrow \beta)]\}, \Delta$, we are done. \blacksquare

3.5.10. OBSERVATION. Read bottom-up, rule ind_s ‘packs’ the context formulas and sequents in Φ and Δ , respectively, together with the principal formula $\alpha R\beta$ into a single R-formula under an X. The context information is dualised and the result can be later extracted by rule $\forall R$ leaving the formula $\alpha R\beta$ intact.

We shall need a rule analogous to $\forall R$ but capable of reasoning under an X:

3.5.11. LEMMA. *The rule*

$$\text{X}\forall R \frac{A\{\Phi, X\gamma, X(\alpha R\beta)\}, \Delta}{A\{\Phi, X[(\gamma \vee \alpha)R(\gamma \vee \beta)]\}, \Delta}$$

is admissible in $\text{CTL}_{\text{ind}}^*$.

Proof. Let $\delta := (\gamma \vee \alpha)R(\gamma \vee \beta)$. We simulate $\text{X}\forall R$ as follows. We begin by cutting against $\text{X}\gamma \vee \text{X}(\alpha R\beta)$:

$$\text{cut} \frac{\begin{array}{c} \text{iW} \frac{A\{\Phi, X\gamma, X(\alpha R\beta)\}, \Delta}{A\{\Phi, X\gamma, X(\alpha R\beta), X\delta\}, \Delta} \\ \text{A}\vee \frac{A\{\Phi, X\gamma \vee \text{X}(\alpha R\beta), X\delta\}, \Delta}{A\{\Phi, X\gamma^\partial \wedge \text{X}(\alpha R\beta)^\partial, X\delta\}, \Delta} \end{array}}{A\{\Phi, X\delta\}, \Delta}$$

At the left branch we are done. At the right one we continue thus:

$$\text{A}\wedge \frac{\begin{array}{c} \text{AX} \frac{A\{\Phi, X\gamma^\partial, X\delta\}, \Delta}{A\{\Phi, X\gamma^\partial \wedge \text{X}(\alpha R\beta)^\partial, X\delta\}, \Delta} \\ \text{A}\wedge \frac{\begin{array}{c} \text{LEM} \frac{A\{\gamma^\partial, \gamma, \alpha R\beta\}}{A\{\gamma^\partial, \delta\}} \\ \text{VR} \frac{A\{\gamma^\partial, \delta\}}{A\{\gamma^\partial, \delta\}} \end{array}}{A\{\Phi, X\gamma^\partial \wedge \text{X}(\alpha R\beta)^\partial, X\delta\}, \Delta} \end{array}}{A\{\Phi, X\gamma^\partial \wedge \text{X}(\alpha R\beta)^\partial, X\delta\}, \Delta} \quad \begin{array}{c} \text{LEM} \frac{A\{(\alpha R\beta)^\partial, \gamma, \alpha R\beta\}}{A\{(\alpha R\beta)^\partial, \gamma, \alpha R\beta\}} \\ \text{VR} \frac{A\{(\alpha R\beta)^\partial, \gamma, \alpha R\beta\}}{A\{(\alpha R\beta)^\partial, \delta\}} \\ \text{AX} \frac{A\{(\alpha R\beta)^\partial, \delta\}}{A\{\Phi, \text{X}(\alpha R\beta)^\partial, X\delta\}, \Delta} \end{array}$$

■

Finally, we show the admissibility in $\text{CTL}_{\text{ind}}^*$ of several fixpoint unfolding rules.

3.5.12. PROPOSITION. *The rule*

$$\text{AR} \frac{A\{\Phi, \beta \wedge [\alpha \vee \text{X}(\alpha R\beta)]\}, \Delta}{A\{\Phi, \alpha R\beta\}, \Delta}$$

is admissible in $\text{CTL}_{\text{ind}}^*$.

Proof. Let $\gamma := \beta \wedge (\alpha \vee \text{X}(\alpha R\beta))$. We simulate AR as follows:

$$\text{iW} \frac{\text{LEM} \frac{\text{A}\{\Phi, \beta \wedge [\alpha \vee \text{X}(\alpha \text{R}\beta)]\}, \Delta}{\text{A}\{\Phi, \alpha \text{R}\beta, \gamma\}, \Delta}}{\text{cut} \frac{\text{A}\{\Phi, \alpha \text{R}\beta, \gamma^\partial\}, \Delta}{\text{A}\{\Phi, \alpha \text{R}\beta\}, \Delta}} \quad \#1$$

Continuation from #1, with $\xi := \vee \Phi \vee \gamma^\partial \vee \Delta^\sharp$:

$$\text{AV} \frac{\text{LEM} \frac{\text{A}\{\Phi, \beta, \beta^\partial, \alpha^\partial \wedge \text{X}(\alpha^\partial \text{U}\beta^\partial)\}, \Delta}{\text{A}\{\Phi, \beta, \gamma^\partial\}, \Delta}}{\text{inds}_s \frac{\text{A}\{\Phi, \alpha, \text{X}[(\xi \rightarrow \alpha)\text{R}(\xi \rightarrow \beta)], \gamma^\partial\}, \Delta}{\text{A}\{\Phi, \alpha \text{R}\beta, \gamma^\partial\}, \Delta}} \quad \#2$$

Continuation from #2:

$$\text{A}\wedge \frac{\text{LEM} \frac{\text{A}\{\alpha, \text{X}(\alpha \text{R}\beta), \beta^\partial, \alpha^\partial\}}{\text{A}\{\alpha, \text{X}(\alpha \text{R}\beta), \beta^\partial, \alpha^\partial \wedge \text{X}(\alpha \text{R}\beta)^\partial\}} \quad \text{LEM} \frac{\text{A}\{\alpha, \text{X}(\alpha \text{R}\beta), \beta^\partial, \text{X}(\alpha \text{R}\beta)^\partial\}}{\text{A}\{\alpha, \text{X}(\alpha \text{R}\beta), \gamma^\partial\}}}{\text{AV} \frac{\text{eW} \frac{\text{A}\{\alpha, \text{X}(\alpha \text{R}\beta), \gamma^\partial\}, \Delta}{\text{A}\{\Phi, \alpha, \text{X}\xi^\partial, \text{X}(\alpha \text{R}\beta), \gamma^\partial\}, \Delta}}{\text{X}\vee\text{R} \frac{\text{A}\{\Phi, \alpha, \text{X}[(\xi \rightarrow \alpha)\text{R}(\xi \rightarrow \beta)], \gamma^\partial\}, \Delta}{\text{A}\{\Phi, \alpha, \text{X}[(\xi \rightarrow \alpha)\text{R}(\xi \rightarrow \beta)], \gamma^\partial\}, \Delta}}}$$

■

3.5.13. PROPOSITION. *The rule*

$$\text{AU}^{-1} \frac{\text{A}\{\Phi, \alpha \text{U}\beta\}, \Delta}{\text{A}\{\Phi, \beta \vee [\alpha \wedge \text{X}(\alpha \text{U}\beta)]\}, \Delta}$$

is admissible in $\text{CTL}_{\text{ind}}^*$.

Proof. Let $\gamma := \beta \vee [\alpha \wedge \text{X}(\alpha \text{U}\beta)]$. We simulate AU^{-1} as follows:

$$\text{iW} \frac{\text{AR} \frac{\text{eW} \frac{\text{A}\{\gamma, \beta^\partial \wedge [\alpha^\partial \vee \text{X}(\alpha^\partial \text{R}\beta^\partial)]\}}{\text{A}\{\gamma, \alpha^\partial \text{R}\beta^\partial\}}}{\text{A}\{\Phi, \alpha \text{U}\beta\}, \Delta}}{\text{cut} \frac{\text{A}\{\Phi, \gamma, \alpha \text{U}\beta\}, \Delta}{\text{A}\{\Phi, \gamma\}, \Delta}} \quad \text{iW} \frac{\text{A}\{\Phi, \alpha \text{U}\beta\}, \Delta}{\text{A}\{\Phi, \gamma, \alpha \text{U}\beta\}, \Delta} \quad \text{iW} \frac{\text{A}\{\gamma, \alpha^\partial \text{R}\beta^\partial\}, \Delta}{\text{A}\{\Phi, \gamma, \alpha^\partial \text{R}\beta^\partial\}, \Delta} \quad \#1$$

Continuation from #1:

$$\frac{\text{LEM} \frac{\overline{A\{\beta, \alpha \wedge X(\alpha U\beta), \beta^\partial\}}}{A\vee} \quad \text{#2} \frac{A\{\gamma, \alpha^\partial, X(\alpha^\partial R\beta^\partial)\}}{A\vee \frac{A\{\gamma, \alpha^\partial \vee X(\alpha^\partial R\beta^\partial)\}}{A\wedge}}}{A\wedge \frac{A\{\gamma, \beta^\partial\}}{A\vee \frac{A\{\gamma, \beta^\partial \wedge [\alpha^\partial \vee X(\alpha^\partial R\beta^\partial)]\}}{A\wedge}}}$$

Continuation from #2:

$$\frac{\text{LEM} \frac{\overline{A\{\beta, \alpha, \alpha^\partial, X(\alpha^\partial R\beta^\partial)\}}}{A\wedge} \quad \text{LEM} \frac{\overline{A\{\beta, X(\alpha U\beta), \alpha^\partial, X(\alpha^\partial R\beta^\partial)\}}}{A\wedge}}{A\vee \frac{A\{\beta, \alpha \wedge X(\alpha U\beta), \alpha^\partial, X(\alpha^\partial R\beta^\partial)\}}{A\wedge} \frac{A\{\gamma, \alpha^\partial, X(\alpha^\partial R\beta^\partial)\}}{A\vee}}$$

■

3.5.14. PROPOSITION. *The rule*

$$\text{AR}^{-1} \frac{A\{\Phi, \alpha R\beta\}, \Delta}{A\{\Phi, \beta \wedge [\alpha \vee X(\alpha R\beta)]\}, \Delta}$$

is admissible in $\text{CTL}_{\text{ind}}^*$.

Proof. Let $\gamma := \beta \wedge [\alpha \vee X(\alpha R\beta)]$. We simulate AR^{-1} as follows:

$$\frac{\text{iW} \frac{A\{\Phi, \alpha R\beta\}, \Delta}{A\{\Phi, \gamma, \alpha R\beta\}, \Delta} \quad \text{cut} \quad \frac{\text{#1} \frac{A\{\gamma, \beta^\partial, \alpha^\partial \wedge X(\alpha^\partial U\beta^\partial)\}}{A\vee \frac{A\{\gamma, \beta^\partial \vee [\alpha^\partial \wedge X(\alpha^\partial U\beta^\partial)]\}}{A\wedge} \frac{A\{\gamma, \alpha^\partial U\beta^\partial\}}{eW} \frac{A\{\gamma, \alpha^\partial U\beta^\partial\}, \Delta}{A\wedge} \quad \text{iW} \frac{A\{\gamma, \alpha^\partial U\beta^\partial\}, \Delta}{A\{\Phi, \gamma, \alpha^\partial U\beta^\partial\}, \Delta}}{A\{\Phi, \gamma\}, \Delta}}$$

Continuation from #1:

$$\frac{\text{LEM} \frac{\overline{A\{\beta, \beta^\partial, \alpha^\partial\}}}{A\wedge} \quad \frac{\text{LEM} \frac{\overline{A\{\alpha, X(\alpha R\beta), \beta^\partial, \alpha^\partial\}}}{A\vee} \quad \frac{A\{\alpha \vee X(\alpha R\beta), \beta^\partial, \alpha^\partial\}}{A\wedge} \quad \text{#2} \frac{A\{\gamma, \beta^\partial, X(\alpha^\partial U\beta^\partial)\}}{A\wedge}}{A\wedge \frac{A\{\gamma, \beta^\partial, \alpha^\partial\}}{A\vee \frac{A\{\gamma, \beta^\partial, \alpha^\partial \wedge X(\alpha^\partial U\beta^\partial)\}}{A\wedge}}}$$

$$\begin{array}{c}
\text{E}\wedge \frac{\text{E}\{\Phi, \beta \vee (\alpha \wedge X(\alpha U\beta))\}, \Delta}{\text{E}\{\wedge \Phi \wedge [\beta \vee (\alpha \wedge X(\alpha U\beta))]\}, \Delta} \\
\text{eW} \frac{\text{E}\{\Phi, \alpha U\beta\}, \text{A}\emptyset, \text{E}\{\delta\}, \Delta}{\text{E}\{\Phi, \alpha U\beta\}, \text{A}\{\text{E}\delta\}, \Delta} \\
\text{AE} \frac{\text{E}\{\Phi, \alpha U\beta\}, \text{A}\emptyset, \text{E}\{\delta\}, \Delta}{\text{E}\{\Phi, \alpha U\beta\}, \text{A}\{\text{E}\delta\}, \Delta} \\
\text{cut} \frac{\text{E}\{\Phi, \alpha U\beta\}, \text{A}\emptyset, \Delta}{\perp^{-1} \frac{\text{E}\{\Phi, \alpha U\beta\}, \text{A}\emptyset, \Delta}{\text{E}\{\Phi, \alpha U\beta\}, \Delta}} \\
\text{AR}^{-1} \frac{\text{axAE} \frac{\text{E}\{\Phi, \alpha U\beta\}, \text{A}\{\Phi^\partial, \alpha^\partial \text{R}\beta^\partial\}}{\text{E}\{\Phi, \alpha U\beta\}, \text{A}\{\Phi^\partial, \gamma^\partial\}}}{\text{E}\{\Phi, \alpha U\beta\}, \text{A}\{\delta^\partial\}} \\
\text{AV} \frac{\text{E}\{\Phi, \alpha U\beta\}, \text{A}\{\delta^\partial\}}{\text{E}\{\Phi, \alpha U\beta\}, \text{A}\{\delta^\partial\}, \Delta} \\
\text{eW} \frac{\text{E}\{\Phi, \alpha U\beta\}, \text{A}\emptyset, \text{A}\{\delta^\partial\}, \Delta}{\text{E}\{\Phi, \alpha U\beta\}, \text{A}\{\text{A}\delta^\partial\}, \Delta} \\
\text{AA} \frac{\text{E}\{\Phi, \alpha U\beta\}, \text{A}\{\text{A}\delta^\partial\}, \Delta}{\text{E}\{\Phi, \alpha U\beta\}, \text{A}\{\text{A}\delta^\partial\}, \Delta}
\end{array}$$

■

3.5.2 Inductible cyclic proofs

We now isolate sufficient conditions for a CTL_\circ^* proof to be translatable into a $\text{CTL}_{\text{ind}}^*$ proof. In what follows, we let $X^0\alpha := \alpha$ and $X^{n+1}\alpha := X^n X\alpha$, for every $n < \omega$.

Let \mathcal{T} be a discarding CTL_\circ^* derivation. Let $u \in T$, and let $x \in \mathbf{N}$ be a name occurring in (the label of) u . By Lemma 3.4.2, the label of u is then of the form $\Theta \dashv \text{A}\{\Phi, X^j(\alpha \text{R}\beta)^x\}, \Delta$, where $j \leq 1$ and x does not occur in Φ or Δ . Let $v <_{\mathcal{T}} u$ be the $<_{\mathcal{T}}$ -greatest proper predecessor of u such that x is fixed on $(v, u]_{\mathcal{T}}$ and rule AR_0 or AR_1 is applied at v with active name x . Then, by Lemmas 3.4.2, 3.4.7 and 3.4.8 the label of v is of the form $\Omega \dashv \text{A}\{\Psi, (\alpha \text{R}\beta)^a\}, \Sigma$, where $\text{A}\{\Psi, (\alpha \text{R}\beta)^a\}$ and $(\alpha \text{R}\beta)^a$ are the principal sequent and formula, respectively, at v , $(\alpha \text{R}\beta)^a \notin \Psi$, and $\text{A}\{\Psi, (\alpha \text{R}\beta)^a\} \notin \Sigma$. The *inductive invariant* of x at u is the pair (Ψ, Σ) . We call Ψ the *internal (inductive) invariant* of x at u , and Σ the *external (inductive) invariant* of x at u . The *inductive base* of x at u is the vertex v .

A universal CTL_\circ^* proof \mathcal{T} has *unfolding companions* if for every repeat $l \in \text{Rep}_{\mathcal{T}}$ there is a name x fixed on $[c_l, l]_{\mathcal{T}}$ such that rule AR_0 is applied at c_l with principal name x . The name x is said to be the *prime witness* of l .

A universal CTL_\circ^* proof \mathcal{T} is *low-unfolding* if for every repeat $l \in \text{Rep}_{\mathcal{T}}$ rules AR_0 and AR_1 are not applied anywhere on $(c_l, l]_{\mathcal{T}}$.

A CTL_\circ^* proof \mathcal{T} is *weakly thinning* if in every instance of eThin in \mathcal{T} , the unique active sequent is one of the two principal sequents.

3.5.17. DEFINITION (Inductible CTL_\circ^* proof). A CTL_\circ^* proof \mathcal{T} is *inductible* if \mathcal{T} is universal, low-unfolding, and weakly thinning, and has unfolding companions. \dashv

3.5.18. EXAMPLE. The CTL_\circ^* proof depicted in Figure 3.8 is inductible. From our work in this section, then, it will follow that the unique cycle it contains can be replaced by an inductive argument inside $\text{CTL}_{\text{ind}}^*$ (see Figure 3.13 below).

If \mathcal{T} is inductible and $l \in \text{Rep}_{\mathcal{T}}$ has prime witness $x \in \mathbf{N}$, then by Lemma 3.4.2 the label of l is of the form $\Theta \dashv \mathbf{A}\{\Phi, (\alpha\mathbf{R}\beta)^x\}, \Delta$, where x does not occur in Φ or Δ . The set of formulas Φ is the *internal context* of l , and the set of sequents Δ is the *external context* of l . The *context* of l is the pair (Φ, Δ) .

The next proposition describes inductive invariants of repeats in inductible proofs.

3.5.19. PROPOSITION. *Let \mathcal{T} be an inductible CTL_\circ^* proof, and let $l \in \text{Rep}_{\mathcal{T}}$ be a repeat with prime witness $x \in \mathbf{N}$. The inductive base of x at l is c_l , and the inductive invariant of x at l is the context of l .*

Proof. Let (Φ, Δ) be the context of l . Then, the label of l and c_l is of the form $\Theta \dashv \mathbf{A}\{\Phi, (\alpha\mathbf{R}\beta)^x\}, \Delta$, where x does not occur in Φ or Δ .

Since x is the prime witness of l , rule AR_0 is applied at c_l with principal name x and x is fixed on $[c_l, l]_{\mathcal{T}}$. And, since \mathcal{T} is low-unfolding, rules AR_0 and AR_1 are not applied anywhere on $(c_l, l]_{\mathcal{T}}$. This shows that c_l is the inductive base of x at l .

By Lemma 3.4.2, $\mathbf{A}\{\Phi, (\alpha\mathbf{R}\beta)^x\}$ and $(\alpha\mathbf{R}\beta)^x$ are the principal sequent and formula, respectively, at c_l , so the inductive invariant of x at l is (Φ, Δ) . \blacksquare

Let \mathcal{T} be an inductible CTL_\circ^* proof. By induction on the height, for every vertex $u \in T$ we define a map t_u from well-annotated formulas to plain formulas. If u is the root of \mathcal{T} , we let $t_u(\varphi^a) := \varphi$ for every well-annotated formula φ^a . For the inductive case, assume that t_v has been defined for every vertex $v <_T u$. We extend each t_v to sequents and hypersequents by setting:

- $t_v(\mathbf{Q}\Phi) := \mathbf{Q}\{t_v(\varphi^a) \mid \varphi^a \in \Phi\}$, for every sequent $\mathbf{Q}\Phi$;
- $t_v(\Gamma) := \{t_v(\mathbf{Q}\Phi) \mid \mathbf{Q}\Phi \in \Gamma\}$, for every set of sequents Γ ;
- $t_v(\Theta \dashv \Gamma) := t_v(\Gamma)$, for every hypersequent $\Theta \dashv \Gamma$.

For the definition of t_u , we begin by setting $t_u(\varphi^a) := \varphi$ if $a = \varepsilon$. Assume now that $a = x$ for some $x \in \mathbf{N}$. So $\varphi = \mathbf{X}^j(\alpha\mathbf{R}\beta)$ for some $j \leq 1$. If neither φ^x nor $\mathbf{X}^{1-j}(\alpha\mathbf{R}\beta)^x$ occurs in u , we let $t_u(\mathbf{X}^k(\alpha\mathbf{R}\beta)^x) := \mathbf{X}^k(\alpha\mathbf{R}\beta)$ for $k = 0, 1$. Assume, finally, that one of φ^x or $\mathbf{X}^{1-j}(\alpha\mathbf{R}\beta)^x$ occurs in u . Let $v <_T u$ be the inductive base of x at u , and let (Φ, Δ) be the inductive invariant of x at u . Then, for every $k \leq 1$ we let

$$t_u(\mathbf{X}^k(\alpha\mathbf{R}\beta)^x) := \mathbf{X}^k[(\xi \rightarrow \alpha)\mathbf{R}(\xi \rightarrow \beta)],$$

where $\xi := \bigvee t_v(\Phi) \vee (t_v(\Delta))^\sharp$ and we have extended t_v to sets of formulas and

For every $u \in T$, we let $t_{\mathcal{T}}(u) := t_u(\Theta_u \dashv \Gamma_u)$, where $\Theta_u \dashv \Gamma_u$ is the hypersequent labelling u .

3.5.20. LEMMA. *Let \mathcal{T} be an inductible CTL_\circ^* proof. For every repeat $l \in \text{Rep}_{\mathcal{T}}$ and every well-annotated formula φ^a , if a is not the prime witness of l , then $t_l(\varphi^a) = t_{c_l}(\varphi^a)$.*

Proof. Assume first that $a = x$ for some $x \in \mathbf{N}$, whence $\varphi = \mathsf{X}^j\psi$ for some R-formula ψ and some $j \leq 1$, and that either φ^x or $\mathsf{X}^{1-j}\psi^x$ occurs in l and c_l . Then, since \mathcal{T} is low-unfolding and x is not the prime witness of l , the inductive invariants of x at l and c_l are identical and thus $t_l(\varphi^x) = t_{c_l}(\varphi^x)$. In all other cases we have $t_l(\varphi^a) = t_{c_l}(\varphi^a) = \varphi$ by definition and the fact that l and c_l have identical labels. \blacksquare

The next result shows that the translation of every repeat is provable in $\text{CTL}_{\text{ind}}^*$. This is the first major step towards cycle elimination.

3.5.21. LEMMA. *Let \mathcal{T} be an inductible CTL_\circ^* proof. For every repeat $l \in \text{Rep}_{\mathcal{T}}$, we have $\text{CTL}_{\text{ind}}^* \vdash t_{\mathcal{T}}(l)$.*

Proof. Let $\mathsf{A}\{\Phi, (\alpha\mathsf{R}\beta)^x\}, \Delta$ be the label of l , where x is the prime witness of l and (Φ, Δ) is the context of l . By Proposition 3.5.19, $t_l((\alpha\mathsf{R}\beta)^x) = (\xi \rightarrow \alpha)\mathsf{R}(\xi \rightarrow \beta)$, where $\xi := \bigvee t_{c_l}(\Phi) \vee t_{c_l}(\Delta)^\sharp$. It suffices to show that $\text{CTL}_{\text{ind}}^* \vdash \mathsf{A}\{t_l(\Phi), \xi^\partial\}, t_l(\Delta)$, for then $\text{CTL}_{\text{ind}}^* \vdash t_{\mathcal{T}}(l)$ follows by applying rules $\vee\mathsf{R}$ and $i\mathsf{W}$.

We first break down the conjunction $(\bigvee t_{c_l}(\Phi) \vee t_{c_l}(\Delta)^\partial)^\partial$ by repeated applications of rule $\mathsf{A}\wedge$, obtaining the following hypersequents:

- (i) $\Gamma_{\varphi^a} := \mathsf{A}\{t_l(\Phi), t_l(\varphi^a), t_{c_l}(\varphi^a)^\partial\}, t_l(\Delta)$, for every $\varphi^a \in \Phi$;
- (ii) $\Gamma_{\mathsf{Q}\Psi} := \mathsf{A}\{t_l(\Phi), (t_{c_l}(\mathsf{Q}\Psi)^\sharp)^\partial\}, t_l(\mathsf{Q}\Psi), t_l(\Delta)$, for every $\mathsf{Q}\Psi \in \Delta$.

By Lemmas 3.5.6 and 3.5.20, $\text{CTL}_{\text{ind}}^* \vdash \Gamma_{\varphi^a}$ for every $\varphi^a \in \Phi$.

Let $\mathsf{Q}\Psi \in \Delta$. We have:

$$\begin{aligned} (t_{c_l}(\mathsf{Q}\Psi)^\sharp)^\partial &= (\mathsf{Q}\{t_{c_l}(\psi^a) \mid \psi^a \in \Psi\}^\sharp)^\partial \\ &= (\mathsf{Q}\circ\{t_{c_l}(\psi^a) \mid \psi^a \in \Psi\})^\partial \\ &= \mathsf{Q}^\partial\circ^\partial\{t_{c_l}(\psi^a)^\partial \mid \psi^a \in \Psi\}, \end{aligned}$$

where $\circ := \vee$ and $\circ^\partial := \wedge$ if $\mathsf{Q} = \mathsf{A}$, and otherwise $\circ := \wedge$ and $\circ^\partial := \vee$. We may thus prove $\Gamma_{\mathsf{Q}\Psi}$ in $\text{CTL}_{\text{ind}}^*$ as follows:

$$\mathsf{AQ}^\partial \frac{\mathsf{eW} \frac{\mathsf{R} \frac{\mathsf{Q}^\partial\{t_{c_l}(\psi^a)^\partial \mid \psi^a \in \Psi\}, t_l(\mathsf{Q}\Psi)}{\mathsf{Q}^\partial\{\circ^\partial\{t_{c_l}(\psi^a)^\partial \mid \psi^a \in \Psi\}\}, t_l(\mathsf{Q}\Psi)}}{\mathsf{A}\{t_l(\Phi)\}, \mathsf{Q}^\partial\{\circ^\partial\{t_{c_l}(\psi^a)^\partial \mid \psi^a \in \Psi\}\}, t_l(\mathsf{Q}\Psi), t_l(\Delta)}}{\mathsf{A}\{t_l(\Phi), (t_{c_l}(\mathsf{Q}\Psi)^\sharp)^\partial\}, t_l(\mathsf{Q}\Psi), t_l(\Delta)}} ,$$

where $R := A\vee$ if $Q = E$, and otherwise $R := E\wedge$. By Lemma 3.5.20, $t_{c_l}(\psi^a) = t_l(\psi^a)$ for every $\psi^a \in \Psi$, so

$$\mathbf{Q}^\partial \{t_{c_l}(\psi^a)^\partial \mid \psi^a \in \Psi\} = \mathbf{Q}^\partial \{t_l(\psi^a)^\partial \mid \psi^a \in \Psi\} = (t_l(\mathbf{Q}\Psi))^\partial.$$

Therefore, $\mathbf{Q}^\partial \{t_{c_l}(\psi^a)^\partial \mid \psi^a \in \Psi\}$, $t_l(\mathbf{Q}\Psi)$ is an instance of \mathbf{ax}_{AE} and we are done. \blacksquare

We need two further results about the translation maps t_u before being ready to eliminate cycles.

3.5.22. LEMMA. *Let \mathcal{T} be an inductible CTL_\circ^* proof. For every vertex $u \in T$ and every well-annotated modal formula $X\varphi^a$, we have $t_u(X\varphi^a) = Xt_u(\varphi^a)$.*

Proof. If $a = \varepsilon$ then $t_u(X\varphi^a) = X\varphi = Xt_u(\varphi^a)$. Assume that $a = x$ for some $x \in \mathbf{N}$. Since $X\varphi^x$ is well-annotated, we have $\varphi = \alpha R\beta$ for some formulas α and β . If neither φ^x nor $X(\alpha R\beta)^x$ occurs in u , then again $t_u(X\varphi^a) = X\varphi = Xt_u(\varphi^a)$. Otherwise, for some formula η we have $Xt_u(\varphi^x) = Xt_u((\alpha R\beta)^x) = X[(\eta \rightarrow \alpha)R(\eta \rightarrow \beta)] = t_u(X(\alpha R\beta)^x) = t_u(X\varphi^a)$ and we are done. \blacksquare

3.5.23. LEMMA. *Let \mathcal{T} be an inductible CTL_\circ^* proof, $u \in T$ a non-final vertex, R the rule applied at u , and v any immediate successor of u . For every annotated formula φ^a occurring in v , if $R \notin \{\mathbf{AR}_0, \mathbf{AR}_1\}$ or φ^a is not active at v , then $t_u(\varphi^a) = t_v(\varphi^a)$.*

Proof. If $a = \varepsilon$ then $t_u(\varphi^a) = t_v(\varphi^a) = \varphi$. Assume that $a = x$ for some $x \in \mathbf{N}$. Then, $\varphi = X^j(\alpha R\beta)$ for some formulas α and β and some $j \leq 1$.

Suppose first that φ^x is not active at v . Then, φ^x occurs in u and is a side formula at u (by Lemma 3.4.2, rule \mathbf{del} is not applied at u with φ^x principal), whence the inductive invariants of x at v and u are identical and thus $t_u(\varphi^x) = t_v(\varphi^x)$.

Suppose now that φ^x is active at v . Then, $R \in \{\mathbf{AX}, \mathbf{iThin}, \mathbf{eThin}\}$. If $R \in \{\mathbf{iThin}, \mathbf{eThin}\}$, then $t_u(\varphi^x) = t_v(\varphi^x)$ because the inductive invariants of x at u and v are identical. And if $R = \mathbf{AX}$, then, given that the formula $X^{j+1}(\alpha R\beta)^x$ occurs in the label of u and is thus well-annotated, we have $j = 0$. Hence, $t_u(\varphi^x) = t_v(\varphi^x)$ because the inductive invariants of x at u and v are identical. \blacksquare

Fix an inductible CTL_\circ^* proof \mathcal{T} . Top-down, inductively, viewing \mathcal{T} as a tree without back-edges, we build a $\text{CTL}_{\text{ind}}^*$ proof of $t_{\mathcal{T}}(u)$ for each vertex $u \in T$. In particular, if the root of \mathcal{T} has label $A\varphi$, then $\text{CTL}_{\text{ind}}^* \vdash \varphi$.

Let l be a leaf of \mathcal{T} . If $l \in \text{Rep}_{\mathcal{T}}$, then $\text{CTL}_{\text{ind}}^* \vdash t_{\mathcal{T}}(l)$ by Lemma 3.5.21. Assume that l is axiomatic. If l has a label of the form $\Theta \dashv \text{Q}p, \text{Q}'\bar{p}, \Delta$, then $t_{\mathcal{T}}(l) = \text{Q}p, \text{Q}'\bar{p}, t_l(\Delta)$, which is axiomatic in $\text{CTL}_{\text{ind}}^*$. And if the label of l is of the form $\Theta \dashv \text{E}\emptyset, \Delta$, then $t_{\mathcal{T}}(l) = \text{E}\emptyset, t_l(\Delta)$ is again axiomatic in $\text{CTL}_{\text{ind}}^*$.

For the inductive step we distinguish cases according to the rule R applied at $u \in T$.¹³

Case R = ALit. Let v be the unique immediate successor of u . In \mathcal{T} we have:

$$\text{ALit} \frac{\Theta \dashv \text{A}\Phi, \text{A}\ell, \Delta}{\Theta \dashv \text{A}\{\Phi, \ell\}, \Delta}$$

We are to build a CTL_{\circ}^* proof of $\text{A}\{t_u(\Phi), \ell\}, t_u(\Delta)$. We have:

$$\text{ALit} \frac{t_u(\text{A}\Phi), \text{A}\ell, t_u(\Delta)}{\text{A}\{t_u(\Phi), \ell\}, t_u(\Delta)}$$

By Lemma 3.5.23, the premise is $t_{\mathcal{T}}(v)$, which by the inductive hypothesis is provable in CTL_{\circ}^* .

Case R = ELit. Let v_1, v_2 be the two immediate successors of u . In \mathcal{T} we have:

$$\text{ELit} \frac{\Theta \dashv \text{E}\Phi, \Delta \quad \Theta \dashv \text{E}\ell, \Delta}{\Theta \dashv \text{E}\{\Phi, \ell\}, \Delta}$$

We are to build a CTL_{\circ}^* proof of $t_u(\text{E}\{\Phi, \ell\}), t_u(\Delta) = \text{E}\{\Phi, \ell\}, t_u(\Delta)$. We have:

$$\text{ELit} \frac{\text{E}\Phi, t_u(\Delta) \quad \text{E}\ell, t_u(\Delta)}{\text{E}\{\Phi, \ell\}, t_u(\Delta)}$$

By Lemma 3.5.23, the premises are $t_{\mathcal{T}}(v_1)$ and $t_{\mathcal{T}}(v_2)$, which by the inductive hypothesis are both provable in CTL_{\circ}^* .

Case R = AV. Let v be the unique immediate successor of u . In \mathcal{T} we have:

$$\text{AV} \frac{\Theta \dashv \text{A}\{\Phi, \varphi, \psi\}, \Delta}{\Theta \dashv \text{A}\{\Phi, \varphi \vee \psi\}, \Delta}$$

We are to build a CTL_{\circ}^* proof of $\text{A}\{t_u(\Phi), t_u(\varphi \vee \psi)\}, t_u(\Delta) = \text{A}\{t_u(\Phi), \varphi \vee \psi\}, t_u(\Delta)$. We have:

$$\text{AV} \frac{\text{A}\{t_u(\Phi), \varphi, \psi\}, t_u(\Delta)}{\text{A}\{t_u(\Phi), \varphi \vee \psi\}, t_u(\Delta)}$$

By Lemma 3.5.23, the premise is $t_{\mathcal{T}}(v)$, which by the inductive hypothesis is provable in CTL_{\circ}^* .

Case R = EV. Let v be the unique immediate successor of u . In \mathcal{T} we have:

¹³We shall work at the same time in CTL_{\circ}^* and $\text{CTL}_{\text{ind}}^*$. The presence of controls or the absence thereof always distinguishes the two systems.

$$E\vee \frac{\Theta \dashv E\{\Phi, \varphi\}, E\{\Phi, \psi\}, \Delta}{\Theta \dashv E\{\Phi, \varphi \vee \psi\}, \Delta}$$

We are to build a CTL_o^* proof of $t_u(E\{\Phi, \varphi \vee \psi\}), t_u(\Delta) = E\{\Phi, \varphi \vee \psi\}, t_u(\Delta)$. We have:

$$E\vee \frac{E\{\Phi, \varphi\}, E\{\Phi, \psi\}, t_u(\Delta)}{E\{\Phi, \varphi \vee \psi\}, t_u(\Delta)}$$

By Lemma 3.5.23, the premise is $t_{\mathcal{T}}(v)$, which by the inductive hypothesis is provable in CTL_o^* .

Case R = $A\wedge$. Let v_1, v_2 be the two immediate successors of u . In \mathcal{T} we have:

$$A\wedge \frac{\Theta \dashv A\{\Phi, \varphi\}, \Delta \quad \Theta \dashv A\{\Phi, \psi\}, \Delta}{\Theta \dashv A\{\Phi, \varphi \wedge \psi\}, \Delta}$$

We are to build a CTL_o^* proof of $A\{t_u(\Phi), t_u(\varphi \wedge \psi)\}, t_u(\Delta) = A\{t_u(\Phi), \varphi \wedge \psi\}, t_u(\Delta)$. We have:

$$A\wedge \frac{A\{t_u(\Phi), \varphi\}, t_u(\Delta) \quad A\{t_u(\Phi), \psi\}, t_u(\Delta)}{A\{t_u(\Phi), \varphi \wedge \psi\}, t_u(\Delta)}$$

By Lemma 3.5.23, the premises are $t_{\mathcal{T}}(v_1)$ and $t_{\mathcal{T}}(v_2)$, which by the inductive hypothesis are both provable in CTL_o^* .

Case R = $E\wedge$. Let v be the unique immediate successor of u . In \mathcal{T} we have:

$$E\wedge \frac{\Theta \dashv E\{\Phi, \varphi, \psi\}, \Delta}{\Theta \dashv E\{\Phi, \varphi \wedge \psi\}, \Delta}$$

We are to build a CTL_o^* proof of $t_u(E\{\Phi, \varphi \wedge \psi\}), t_u(\Delta) = E\{\Phi, \varphi \wedge \psi\}, t_u(\Delta)$. We have:

$$E\wedge \frac{E\{\Phi, \varphi, \psi\}, t_u(\Delta)}{E\{\Phi, \varphi \wedge \psi\}, t_u(\Delta)}$$

By Lemma 3.5.23, the premise is $t_{\mathcal{T}}(v)$, which by the inductive hypothesis is provable in CTL_o^* .

Case R = AA . Let v be the unique immediate successor of u . In \mathcal{T} we have:

$$AA \frac{\Theta \dashv A\Phi, A\{\psi\}, \Delta}{\Theta \dashv A\{\Phi, A\psi\}, \Delta}$$

We are to build a CTL_o^* proof of $A\{t_u(\Phi), A\psi\}, t_u(\Delta)$. We have:

$$AA \frac{A(t_u(\Phi)), A\{\psi\}, t_u(\Delta)}{A\{t_u(\Phi), A\psi\}, t_u(\Delta)}$$

By Lemma 3.5.23, the premise is $t_{\mathcal{T}}(v)$, which by the inductive hypothesis is provable in CTL_o^* .

Case R = EA. Let v_1, v_2 be the two immediate successors of u . In \mathcal{T} we have:

$$\text{EA} \frac{\Theta \dashv \text{E}\Phi, \Delta \quad \Theta \dashv \text{A}\{\psi\}, \Delta}{\Theta \dashv \text{E}\{\Phi, \text{A}\psi\}, \Delta}$$

We are to build a CTL_o^* proof of $t_u(\text{E}\{\Phi, \text{A}\psi\}, t_u(\Delta)) = \text{E}\{\Phi, \text{A}\psi\}, t_u(\Delta)$. We have:

$$\text{EA} \frac{\text{E}\Phi, t_u(\Delta) \quad \text{A}\{\psi\}, t_u(\Delta)}{\text{E}\{\Phi, \text{A}\psi\}, t_u(\Delta)}$$

By Lemma 3.5.23, the premises are $t_{\mathcal{T}}(v_1)$ and $t_{\mathcal{T}}(v_2)$, which by the inductive hypothesis are both provable in CTL_o^* .

Case R = AE. Let v be the unique immediate successor of u . In \mathcal{T} we have:

$$\text{AE} \frac{\Theta \dashv \text{A}\Phi, \text{E}\{\psi\}, \Delta}{\Theta \dashv \text{A}\{\Phi, \text{E}\psi\}, \Delta}$$

We are to build a CTL_o^* proof of $\text{A}\{t_u(\Phi), \text{E}\psi\}, t_u(\Delta)$. We have:

$$\text{AE} \frac{\text{A}(t_u(\Phi)), \text{E}\{\psi\}, t_u(\Delta)}{\text{A}\{t_u(\Phi), \text{E}\psi\}, t_u(\Delta)}$$

By Lemma 3.5.23, the premise is $t_{\mathcal{T}}(v)$, which by the inductive hypothesis is provable in CTL_o^* .

Case R = EE. Let v_1, v_2 be the two immediate successors of u . In \mathcal{T} we have:

$$\text{EA} \frac{\Theta \dashv \text{E}\Phi, \Delta \quad \Theta \dashv \text{E}\{\psi\}, \Delta}{\Theta \dashv \text{E}\{\Phi, \text{E}\psi\}, \Delta}$$

We are to build a CTL_o^* proof of $t_u(\text{E}\{\Phi, \text{E}\psi\}, t_u(\Delta)) = \text{E}\{\Phi, \text{E}\psi\}, t_u(\Delta)$. We have:

$$\text{EA} \frac{\text{E}\Phi, t_u(\Delta) \quad \text{E}\{\psi\}, t_u(\Delta)}{\text{E}\{\Phi, \text{E}\psi\}, t_u(\Delta)}$$

By Lemma 3.5.23, the premises are $t_{\mathcal{T}}(v_1)$ and $t_{\mathcal{T}}(v_2)$, which by the inductive hypothesis are both provable in CTL_o^* .

Case R = AX. Let v be the unique immediate successor of u . In \mathcal{T} we have:

$$\text{AX} \frac{\Theta' \dashv \text{A}\Phi, \text{E}\Psi_1, \dots, \text{E}\Psi_m}{\Theta \dashv \text{A}(\text{X}\Phi \cup \Xi), \text{E}\text{X}\Psi_1, \dots, \text{E}\text{X}\Psi_m, \Sigma}$$

We are to build a CTL_o^* proof of $\text{A}\{t_u(\text{X}\Phi) \cup t_u(\Xi)\}, t_u(\text{E}\text{X}\Psi_1), \dots, t_u(\text{E}\text{X}\Psi_m), t_u(\Sigma)$. Note that $t_u(\text{E}\text{X}\Psi_i) = \text{E}\text{X}\Psi_i$ for every $1 \leq i \leq m$. And, by Lemma 3.5.22, $t_u(\text{X}\Phi) = \text{X}t_u(\Phi)$. We thus have:

$$\text{AX} \frac{A(t_u(\Phi)), E\Psi_1, \dots, E\Psi_m}{A\{Xt_u(\Phi) \cup t_u(\Xi)\}, EX\Psi_1, \dots, EX\Psi_m, t_u(\Sigma)}$$

By Lemma 3.5.23, the premise is $t_{\mathcal{T}}(v)$, which by the inductive hypothesis is provable in CTL_{\circ}^* .

Case R = EX. Let v be the unique immediate successor of u . In \mathcal{T} we have:

$$\text{EX} \frac{\Theta' \dashv E\Psi_1, \dots, E\Psi_m}{\Theta \dashv EX\Psi_1, \dots, EX\Psi_m, \Sigma}$$

We are to build a CTL_{\circ}^* proof of $t_u(EX\Psi_1), \dots, t_u(EX\Psi_m), t_u(\Sigma)$. Note that $t_u(EX\Psi_i) = EX\Psi_i$ for every $1 \leq i \leq m$. We thus have:

$$\text{EX} \frac{E\Psi_1, \dots, E\Psi_m}{EX\Psi_1, \dots, EX\Psi_m, \Sigma}$$

By Lemma 3.5.23, the premise is $t_{\mathcal{T}}(v)$, which by the inductive hypothesis is provable in CTL_{\circ}^* .

Case R = iW. Let v be the unique immediate successor of u . In \mathcal{T} we have:

$$\text{iW} \frac{\Theta' \dashv A\Phi, \Delta}{\Theta \dashv A(\Phi \cup \Psi), \Delta}$$

We are to build a CTL_{\circ}^* proof of $A\{t_u(\Phi) \cup t_u(\Psi)\}, t_u(\Delta)$. We have:

$$\text{iW} \frac{A(t_u(\Phi)), t_u(\Delta)}{A(t_u(\Phi) \cup t_u(\Psi)), t_u(\Delta)}$$

By Lemma 3.5.23, the premise is $t_{\mathcal{T}}(v)$, which by the inductive hypothesis is provable in CTL_{\circ}^* .

Case R = eW. Let v be the unique immediate successor of u . In \mathcal{T} we have:

$$\text{eW} \frac{\Theta' \dashv Q\Phi, \Delta}{\Theta \dashv Q\Phi, Q'\Psi, \Delta}$$

We are to build a CTL_{\circ}^* proof of $t_u(Q\Phi), t_u(Q'\Psi), t_u(\Delta)$. We have:

$$\text{eW} \frac{t_u(Q\Phi), t_u(\Delta)}{t_u(Q\Phi), t_u(Q'\Psi), t_u(\Delta)}$$

By Lemma 3.5.23, the premise is $t_{\mathcal{T}}(v)$, which by the inductive hypothesis is provable in CTL_{\circ}^* .

Case R = iThin. Let v be the unique immediate successor of u . In \mathcal{T} we have:

$$\text{iThin} \frac{\Theta' \dashv A\{\Phi, \varphi^x\}, \Delta}{\Theta \dashv A\{\Phi, \varphi^x, \varphi^a\}, \Delta} x \prec_{\Theta} a$$

We are to build a CTL_\circ^* proof of

$$t_u(\mathbf{A}\{\Phi, \varphi^x, \varphi^a\}), t_u(\Delta) = \mathbf{A}\{t_u(\Phi), t_u(\varphi^x), t_u(\varphi^a)\}, t_u(\Delta).$$

Since $t_{\mathcal{T}}(v) = \mathbf{A}\{t_v(\Phi), t_v(\varphi^x)\}, t_v(\Delta) = \mathbf{A}\{t_u(\Phi), t_u(\varphi^x)\}, t_u(\Delta)$, where the latter equality is given by Lemma 3.5.23, by iW and the inductive hypothesis we are done.

Case $\text{R} = \text{eThin}$. Let v be the unique immediate successor of u . Since \mathcal{T} is inductible, in \mathcal{T} we have:

$$\text{eThin} \frac{\Theta' \dashv \mathbf{A}\{\varphi_0^{a_0}, \dots, \varphi_n^{a_n}\}, \Delta}{\Theta \dashv \mathbf{A}\{\varphi_0^{a_0}, \dots, \varphi_n^{a_n}\}, \mathbf{A}\{\varphi_0^{b_0}, \dots, \varphi_n^{b_n}\}, \Delta}$$

We are to build a CTL_\circ^* proof of

$$t_{\mathcal{T}}(u) = \mathbf{A}\{t_u(\varphi_0^{a_0}), \dots, t_u(\varphi_n^{a_n})\}, \mathbf{A}\{t_u(\varphi_0^{b_0}), \dots, t_u(\varphi_n^{b_n})\}, t_u(\Delta).$$

Since $t_{\mathcal{T}}(v) = \mathbf{A}\{t_v(\varphi_0^{a_0}), \dots, t_v(\varphi_n^{a_n})\}, t_v(\Delta) = \mathbf{A}\{t_u(\varphi_0^{a_0}), \dots, t_u(\varphi_n^{a_n})\}, t_u(\Delta)$, where the latter equality is given by Lemma 3.5.23, by eW and the inductive hypothesis we are done.

Case $\text{R} = \text{del}$. Let v be the unique immediate successor of u . In \mathcal{T} we have:

$$\text{del} \frac{\Theta' \dashv \mathbf{A}\{\Phi, \varphi\}, \Delta}{\Theta \dashv \mathbf{A}\{\Phi, \varphi^x\}, \Delta}$$

We have $\varphi = \mathbf{X}^j(\alpha\mathbf{R}\beta)$ and $t_u(\varphi^x) = \mathbf{X}^j[(\eta \rightarrow \alpha)\mathbf{R}(\eta \rightarrow \beta)]$ for some formulas α , β and η , and some $j \leq 1$. We are to build a CTL_\circ^* proof of $\mathbf{A}\{t_u(\Phi), t_u(\varphi^x)\}, t_u(\Delta)$. By Lemma 3.5.23 and the inductive hypothesis, $t_{\mathcal{T}}(v) = \mathbf{A}\{t_u(\Phi), \varphi\}, t_u(\Delta)$ is provable in $\text{CTL}_{\text{ind}}^*$. We then have:

$$\text{iW} \frac{\mathbf{A}\{t_u(\Phi), \mathbf{X}^j(\alpha\mathbf{R}\beta)\}, t_u(\Delta)}{\mathbf{A}\{t_u(\Phi), \mathbf{X}^j\eta, \mathbf{X}^j(\alpha\mathbf{R}\beta)\}, t_u(\Delta)} \\ \text{X}^j\vee\mathbf{R} \frac{\mathbf{A}\{t_u(\Phi), \mathbf{X}^j[(\eta \rightarrow \alpha)\mathbf{R}(\eta \rightarrow \beta)]\}, t_u(\Delta)}{\mathbf{A}\{t_u(\Phi), \mathbf{X}^j\eta, \mathbf{X}^j(\alpha\mathbf{R}\beta)\}, t_u(\Delta)}$$

The premise is $t_{\mathcal{T}}(v)$, which by the inductive hypothesis is provable in CTL_\circ^* .

Case $\text{R} = \text{AU}$. Let v_1, v_2 be the two immediate successors of u . In \mathcal{T} we have:

$$\text{AU} \frac{\Theta \dashv \mathbf{A}\{\Phi, \varphi_1, \varphi_2\}, \Delta \quad \Theta \dashv \mathbf{A}\{\Phi, \varphi_2, \mathbf{X}(\varphi_1\mathbf{U}\varphi_2)\}, \Delta}{\Theta \dashv \mathbf{A}\{\Phi, \varphi_1\mathbf{U}\varphi_2\}, \Delta}$$

We are to build a CTL_\circ^* proof of $\mathbf{A}\{t_u(\Phi), \varphi_1\mathbf{U}\varphi_2\}, t_u(\Delta)$. We have:

$$\text{A}\wedge \frac{\mathbf{A}\{t_u(\Phi), \varphi_1, \varphi_2\}, t_u(\Delta) \quad \mathbf{A}\{t_u(\Phi), \varphi_2, \mathbf{X}(\varphi_1\mathbf{U}\varphi_2)\}, t_u(\Delta)}{\mathbf{A}\{t_u(\Phi), \varphi_2, \varphi_1 \wedge \mathbf{X}(\varphi_1\mathbf{U}\varphi_2)\}, t_u(\Delta)} \\ \text{A}\vee \frac{\mathbf{A}\{t_u(\Phi), \varphi_2, \varphi_1 \wedge \mathbf{X}(\varphi_1\mathbf{U}\varphi_2)\}, t_u(\Delta)}{\mathbf{A}\{t_u(\Phi), \varphi_2 \vee [\varphi_1 \wedge \mathbf{X}(\varphi_1\mathbf{U}\varphi_2)]\}, t_u(\Delta)} \\ \text{AU} \frac{\mathbf{A}\{t_u(\Phi), \varphi_2 \vee [\varphi_1 \wedge \mathbf{X}(\varphi_1\mathbf{U}\varphi_2)]\}, t_u(\Delta)}{\mathbf{A}\{t_u(\Phi), \varphi_1\mathbf{U}\varphi_2\}, t_u(\Delta)}$$

By Lemma 3.5.23, the premises are $t_{\mathcal{T}}(v_1)$ and $t_{\mathcal{T}}(v_2)$, which by the inductive hypothesis are both provable in CTL_o^* .

Case R = AR_0 with empty principal annotation. Let v_1, v_2 be the two immediate successors of u . In \mathcal{T} we have:

$$\text{AR}_0 \frac{\Theta \dashv \text{A}\{\Phi, \varphi_2\}, \Delta \quad \Theta \dashv \text{A}\{\Phi, \varphi_1, \text{X}(\varphi_1 \text{R}\varphi_2)\}, \Delta}{\Theta \dashv \text{A}\{\Phi, \varphi_1 \text{R}\varphi_2\}, \Delta}$$

We are to build a CTL_o^* proof of $\text{A}\{t_u(\Phi), \varphi_1 \text{R}\varphi_2\}, t_u(\Delta)$. We have:

$$\text{A}\wedge \frac{\text{A}\{t_u(\Phi), \varphi_2\}, t_u(\Delta) \quad \text{A}\vee \frac{\text{A}\{t_u(\Phi), \varphi_1, \text{X}(\varphi_1 \text{R}\varphi_2)\}, t_u(\Delta)}{\text{A}\{t_u(\Phi), \varphi_1 \vee \text{X}(\varphi_1 \text{R}\varphi_2)\}, t_u(\Delta)}}{\text{AR} \frac{\text{A}\{t_u(\Phi), \varphi_2 \wedge [\varphi_1 \vee \text{X}(\varphi_1 \text{R}\varphi_2)]\}, t_u(\Delta)}{\text{A}\{t_u(\Phi), \varphi_1 \text{R}\varphi_2\}, t_u(\Delta)}}$$

By Lemma 3.5.23, the premises are $t_{\mathcal{T}}(v_1)$ and $t_{\mathcal{T}}(v_2)$, which by the inductive hypothesis are both provable in CTL_o^* .

Case R $\in \{\text{EU}, \text{ER}\}$. Analogous to the previous one.

Case R = AR_0 with non-empty principal annotation x . Let v_1, v_2 be the two immediate successors of u . In \mathcal{T} we have:

$$\text{AR}_0 \frac{\Theta \dashv \text{A}\{\Phi, \varphi_2\}, \Delta \quad \Theta \dashv \text{A}\{\Phi, \varphi_1, \text{X}(\varphi_1 \text{R}\varphi_2)^x\}, \Delta}{\Theta \dashv \text{A}\{\Phi, (\varphi_1 \text{R}\varphi_2)^x\}, \Delta}$$

We are to build a CTL_o^* proof of $\text{A}\{t_u(\Phi), t_u((\varphi_1 \text{R}\varphi_2)^x)\}, t_u(\Delta)$. We have

$$t_u((\varphi_1 \text{R}\varphi_2)^x) = (\eta \rightarrow \varphi_1) \text{R}(\eta \rightarrow \varphi_2)$$

for some formula η . The inductive invariant of x at v_2 is (Φ, Δ) and

$$t_{v_2}(\text{X}(\varphi_1 \text{R}\varphi_2)^x) = \text{X}[(\xi \rightarrow \varphi_1) \text{R}(\xi \rightarrow \varphi_2)],$$

where $\xi := \vee t_u(\Phi) \vee t_u(\Delta)^\sharp$. We have:

$$\text{ind}_s \frac{\text{A}\{t_u(\Phi), \varphi_2\}, t_u(\Delta) \quad \text{A}\{t_u(\Phi), \varphi_1, \text{X}[(\xi \rightarrow \varphi_1) \text{R}(\xi \rightarrow \varphi_2)]\}, t_u(\Delta)}{\text{iW} \frac{\text{A}\{t_u(\Phi), \varphi_1 \text{R}\varphi_2\}, t_u(\Delta)}{\text{A}\{t_u(\Phi), \eta, \varphi_1 \text{R}\varphi_2\}, t_u(\Delta)}}}{\text{VR} \frac{\text{A}\{t_u(\Phi), (\eta \rightarrow \varphi_1) \text{R}(\eta \rightarrow \varphi_2)\}, t_u(\Delta)}$$

By Lemma 3.5.23, the premises are $t_{\mathcal{T}}(v_1)$ and $t_{\mathcal{T}}(v_2)$, which by the inductive hypothesis are both provable in CTL_o^* .

Case R = AR_1 . Let v_1, v_2 be the two immediate successors of u . In \mathcal{T} we have:

$$\text{AR}_1 \frac{\Theta \dashv \text{A}\{\Phi, \varphi_2\}, \Delta \quad \Theta x \dashv \text{A}\{\Phi, \varphi_1, \text{X}(\varphi_1 \text{R}\varphi_2)^x\}, \Delta}{\Theta \dashv \text{A}\{\Phi, \varphi_1 \text{R}\varphi_2\}, \Delta}$$

$$\begin{array}{c}
\text{aXRU} \frac{}{A\{\top U p, \top^\theta R \bar{p}\}} \\
\text{iW} \frac{}{A\{\top U p, \top^\theta R \bar{p}, \perp R \bar{p}\}} \\
\text{VR} \frac{}{A\{\top U p, ((\top U p) \rightarrow \perp) R((\top U p) \rightarrow \bar{p})\}} \\
\text{AX} \frac{}{A\{X(\top U p), p, X[((\top U p) \rightarrow \perp) R((\top U p) \rightarrow \bar{p})]\}} \\
\text{AU} \frac{}{A\{\top U p, p, X[((\top U p) \rightarrow \perp) R((\top U p) \rightarrow \bar{p})]\}} \\
\text{A}\wedge \frac{}{A\{\top U p, \perp, X[((\top U p) \rightarrow \perp) R((\top U p) \rightarrow \bar{p})]\}} \\
\text{ind}_s \frac{}{A\{\top U p, \perp R \bar{p}\}} \\
\text{iW} \frac{}{A\{\top U p, \top^\theta R \bar{p}, \perp R \bar{p}\}} \\
\text{VR} \frac{}{A\{\top U p, ((\top U p) \rightarrow \perp) R((\top U p) \rightarrow \bar{p})\}} \\
\text{AX} \frac{}{A\{X(\top U p), p, X[((\top U p) \rightarrow \perp) R((\top U p) \rightarrow \bar{p})]\}} \\
\text{AU} \frac{}{A\{\top U p, p, X[((\top U p) \rightarrow \perp) R((\top U p) \rightarrow \bar{p})]\}} \\
\text{A}\wedge \frac{}{A\{\top U p, \perp, X[((\top U p) \rightarrow \perp) R((\top U p) \rightarrow \bar{p})]\}} \\
\text{ind}_s \frac{}{A\{\top U p, \perp R \bar{p}\}} \\
\text{AV} \frac{}{A\{\top U p \vee \perp R \bar{p}\}}
\end{array}$$

FIGURE 3.13: The translation into $\text{CTL}_{\text{ind}}^*$ of the cycle in the CTL_o^* proof of $(\top U p) \vee (\perp R \bar{p})$ from Figure 3.8.

We are to build a $\text{CTL}_{\text{ind}}^*$ proof of $A\{t_u(\Phi), \varphi_1 R \varphi_2\}, t_u(\Delta)$. The inductive base of x at v_2 is u and the inductive invariant of x at v_2 is (Φ, Δ) . Thus, $t_{v_2}(X(\varphi_1 R \varphi_2)^x) = X[(\xi \rightarrow \varphi_1) R(\xi \rightarrow \varphi_2)]$, where $\xi := \vee t_u(\Phi) \vee t_u(\Delta)$ [‡]. We then have:

$$\text{ind}_s \frac{A\{t_u(\Phi), \varphi_2\}, t_u(\Delta) \quad A\{t_u(\Phi), \varphi_1, X[(\xi \rightarrow \varphi_1) R(\xi \rightarrow \varphi_2)]\}, t_u(\Delta)}{A\{t_u(\Phi), \varphi_1 R \varphi_2\}, t_u(\Delta)}$$

By Lemma 3.5.23, the premises are $t_{\mathcal{T}}(v_1)$ and $t_{\mathcal{T}}(v_2)$, which by the inductive hypothesis are both provable in CTL_o^* .

We have thus established:

3.5.24. THEOREM. *For any formula φ , if there is an inductible CTL_o^* proof of φ , then $\text{CTL}_{\text{ind}}^* \vdash \varphi$.*

3.5.25. EXAMPLE. Figure 3.13 depicts the translation into $\text{CTL}_{\text{ind}}^*$ of the cycle in Figure 3.8. As expected from such a general procedure, the resulting proof is much larger than it need be (especially considering the use of non-primitive rules such as VR and ind_s).

3.5.3 Hilbert-style proofs

We conclude the section on cycle elimination by presenting a Hilbert-style version of $\text{CTL}_{\text{ind}}^*$ and comparing its strength to a known axiomatisation of (a fragment of) CTL^* . We assume that the reader is familiar with proofs à la Hilbert: finite sequences of formulas such that every formula is either an instance of an axiom, or else is obtained by applying a rule to formulas occurring earlier in the sequence.

3.5.26. DEFINITION ($\text{CTL}_{\text{Hil}}^*$). The proof system $\text{CTL}_{\text{Hil}}^*$ is the Hilbert-style proof system whose rules are depicted in Figure 3.14 and with the following axiom schemata:

- Axioms for classical reasoning:
 - (i) $(\alpha \rightarrow \beta) \rightarrow ((\beta \rightarrow \gamma) \rightarrow (\alpha \rightarrow \gamma))$;
 - (ii) $(\alpha^\partial \rightarrow \alpha) \rightarrow \alpha$;
 - (iii) $\alpha \rightarrow (\alpha^\partial \rightarrow \beta)$;
- Axioms for quantifiers:
 - (iv) $A(\alpha \rightarrow \beta) \rightarrow (A\alpha \rightarrow A\beta)$.
 - (v) $A\alpha \rightarrow \alpha$.
 - (vi) $A\alpha \rightarrow AA\alpha$.
 - (vii) $E\alpha \rightarrow AE\alpha$.
- Axioms for literals:
 - (viii) $p \rightarrow Ap$.
 - (ix) $Ep \rightarrow p$.
- Axioms for the interaction between X and the quantifiers:
 - (x) $AX\alpha \rightarrow XA\alpha$.
- Temporal axioms:
 - (xi) $X(\alpha \rightarrow \beta) \rightarrow (X\alpha \rightarrow X\beta)$.
 - (xii) $(\beta \vee (\alpha \wedge X(\alpha U \beta))) \rightarrow \alpha U \beta$.

⊣

3.5.27. REMARK. We use axiom schemata rather than axioms plus a substitution rule because CTL^* is not closed under substitution: while $Ep \rightarrow p$ is valid, clearly $E\alpha \rightarrow \alpha$ is not always valid.

3.5.28. OBSERVATION. Formally, the universal quantifier A behaves like an S5 box (except on literals), where $\text{S5} := \text{S4} \oplus \diamond p \rightarrow \square \diamond p$.¹⁴

¹⁴Recall the definition of S4 from Chapter 2.

$$\begin{array}{ccc}
\text{MP} \frac{\alpha \quad \alpha \rightarrow \beta}{\beta} & & \text{ind} \frac{(\beta \vee (\alpha \wedge X\gamma)) \rightarrow \gamma}{\alpha \text{U} \beta \rightarrow \gamma} \\
\text{A} \frac{\alpha}{\text{A}\alpha} & & \text{X} \frac{\alpha}{\text{X}\alpha}
\end{array}$$

FIGURE 3.14: Non-axiomatic rules of the Hilbert-style system $\text{CTL}_{\text{Hil}}^*$.

The first three axiom schemata of $\text{CTL}_{\text{Hil}}^*$ correspond to the axioms for classical propositional logic (CPC) given in Chapter 2. We assume familiarity with CPC, and in particular with the Hilbert-style axiomatisation thereof based on said schemata and the rule of *modus ponens* (MP).

A CPC *tautology in the language of CTL^** is a CTL^* formula of the form $\sigma(\alpha)$, where α is a CPC tautology and $\sigma : \text{Prop} \rightarrow \text{Form}_{\text{CTL}^*}$ is a CTL^* substitution lifted to formulas in the usual manner. It is easy to see that the system $\text{CTL}_{\text{ind}}^*$ can reason classically in the language of CTL^* :

3.5.29. PROPOSITION. *Every CPC tautology in the language of CTL^* is provable in $\text{CTL}_{\text{Hil}}^*$.*

Proof. We show, by (strong) induction on $k < \omega$, that for every CPC tautology α having a (Hilbert-style) CPC proof of length k , and every substitution $\sigma : \text{Prop} \rightarrow \text{Form}_{\text{CTL}^*}$, we have $\text{CTL}_{\text{Hil}}^* \vdash \sigma(\alpha)$. If $k = 1$, then α is an instance of a CPC axiom and thus $\sigma(\alpha)$ is an instance of a $\text{CTL}_{\text{Hil}}^*$ axiom. For the inductive case, assume that the claim holds for any $k' < k$. If α is an instance of a CPC axiom, then we reason as in the base case $k = 1$. Otherwise, α results by MP from formulas β and $\beta \rightarrow \alpha$ having both CPC proofs of length at most k' for some $k' < k$. By the inductive hypothesis, then, $\text{CTL}_{\text{ind}}^* \vdash \sigma(\beta)$ and $\text{CTL}_{\text{ind}}^* \vdash \sigma(\beta \rightarrow \alpha)$, whence $\text{CTL}_{\text{ind}}^* \vdash \sigma(\alpha)$ by MP. ■

The axioms for $\text{CTL}_{\text{Hil}}^*$ correspond to the rules of $\text{CTL}_{\text{ind}}^*$ in a way that will become clear when we show how to embed the latter into the former. In order to do this, we first need to establish several properties of $\text{CTL}_{\text{Hil}}^*$.

3.5.30. PROPOSITION. *The rule*

$$\text{A}^{-1} \frac{\text{A}\alpha}{\alpha}$$

is admissible in $\text{CTL}_{\text{Hil}}^$.*

Proof. Clear. ■

3.5.31. PROPOSITION. *The rule*

$$A \rightarrow \frac{\alpha \rightarrow \beta}{A\alpha \rightarrow A\beta}$$

is admissible in CTL_{Hil}^* .

Proof. Assuming $CTL_{Hil}^* \vdash \alpha \rightarrow \beta$, we prove $A\alpha \rightarrow A\beta$ in CTL_{Hil}^* as follows:

1. $\alpha \rightarrow \beta$ (assumption)
2. $A(\alpha \rightarrow \beta)$ (A 1)
3. $A(\alpha \rightarrow \beta) \rightarrow (A\alpha \rightarrow A\beta)$ (axiom)
4. $A\alpha \rightarrow A\beta$ (MP 2,3)

■

3.5.32. PROPOSITION. *The rule*

$$E \rightarrow \frac{\alpha \rightarrow \beta}{E\alpha \rightarrow E\beta}$$

is admissible in CTL_{Hil}^* .

Proof. Assuming $CTL_{Hil}^* \vdash \alpha \rightarrow \beta$, we prove $E\alpha \rightarrow E\beta$ in CTL_{Hil}^* as follows:

1. $\alpha \rightarrow \beta$ (assumption)
2. $\beta^\partial \rightarrow \alpha^\partial$ (CPC 1)
3. $A\beta^\partial \rightarrow A\alpha^\partial$ (A \rightarrow 2)
4. $E\alpha \rightarrow E\beta$ (CPC 3)

■

3.5.33. PROPOSITION. *The rule*

$$\vee \rightarrow \frac{\alpha \vee \delta \quad \alpha \rightarrow \beta}{\beta \vee \delta}$$

is admissible in CTL_{Hil}^* .

Proof. Assuming $CTL_{Hil}^* \vdash \alpha \vee \delta$ and $CTL_{Hil}^* \vdash \alpha \rightarrow \beta$, we prove $\beta \vee \delta$ in CTL_{Hil}^* as follows:

1. $\alpha \vee \delta$ (assumption)

2. $\alpha \rightarrow \beta$ (assumption)
3. $\alpha \rightarrow (\beta \vee \delta)$ (CPC 2)
4. $\delta \rightarrow (\beta \vee \delta)$ (CPC)
5. $(\alpha \vee \delta) \rightarrow (\beta \vee \delta)$ (CPC 3,4)
6. $\beta \vee \delta$ (MP 1,5)

■

3.5.34. PROPOSITION. *For every $p \in \text{Prop}$, the following hold:*

- (i) $\text{CTL}_{\text{Hil}}^* \vdash Qp \leftrightarrow p$, for every $Q \in \{A, E\}$;
- (ii) $\text{CTL}_{\text{Hil}}^* \vdash Q\bar{p} \leftrightarrow \bar{p}$, for every $Q \in \{A, E\}$;
- (iii) $\text{CTL}_{\text{Hil}}^* \vdash Ap \leftrightarrow Ep$;
- (iv) $\text{CTL}_{\text{Hil}}^* \vdash A\bar{p} \leftrightarrow E\bar{p}$.

Proof. Clear. ■

3.5.35. PROPOSITION. *For all formulas α , β and γ , the following hold:*

- (i) if $\text{CTL}_{\text{Hil}}^* \vdash \alpha \vee \beta$, then $\text{CTL}_{\text{Hil}}^* \vdash \beta \vee \alpha$;
- (ii) if $\text{CTL}_{\text{Hil}}^* \vdash A(\alpha \vee \beta)$, then $\text{CTL}_{\text{Hil}}^* \vdash A(\beta \vee \alpha)$;
- (iii) if $\text{CTL}_{\text{Hil}}^* \vdash E(\alpha \wedge \beta)$, then $\text{CTL}_{\text{Hil}}^* \vdash E(\beta \wedge \alpha)$.

Proof. Claim (i) is clear by Proposition 3.5.29. For (ii), assume that $\text{CTL}_{\text{Hil}}^* \vdash A(\alpha \vee \beta)$. By Propositions 3.5.29 and 3.5.30, $\text{CTL}_{\text{Hil}}^* \vdash \beta \vee \alpha$, so by rule A we get $\text{CTL}_{\text{Hil}}^* \vdash A(\beta \vee \alpha)$. For (iii), assume that $\text{CTL}_{\text{Hil}}^* \vdash E(\alpha \wedge \beta)$. Then, $\text{CTL}_{\text{Hil}}^* \vdash (A(\alpha \rightarrow \beta^\partial))^\partial$ and thus:

1. $(A(\alpha \rightarrow \beta^\partial))^\partial$ (assumption)
2. $(\beta \rightarrow \alpha^\partial) \rightarrow (\alpha \rightarrow \beta^\partial)$ (CPC)
3. $A(\beta \rightarrow \alpha^\partial) \rightarrow A(\alpha \rightarrow \beta^\partial)$ (A \rightarrow 2)
4. $(A(\alpha \rightarrow \beta^\partial))^\partial \rightarrow (A(\beta \rightarrow \alpha^\partial))^\partial$ (CPC 3)
5. $(A(\beta \rightarrow \alpha^\partial))^\partial$ (MP 1,4)

Since $(A(\beta \rightarrow \alpha^\partial))^\partial = E(\beta \wedge \alpha)$, we are done. ■

3.5.36. PROPOSITION. *For all formulas α and β , we have:*

- (i) $\text{CTL}_{\text{Hil}}^* \vdash (\mathbf{A}\alpha \wedge \mathbf{A}\beta) \rightarrow \mathbf{A}(\alpha \wedge \beta)$;
- (ii) $\text{CTL}_{\text{Hil}}^* \vdash \mathbf{A}(\alpha \wedge \beta) \rightarrow (\mathbf{A}\alpha \wedge \mathbf{A}\beta)$;
- (iii) $\text{CTL}_{\text{Hil}}^* \vdash (\mathbf{E}\alpha \vee \mathbf{E}\beta) \rightarrow \mathbf{E}(\alpha \vee \beta)$;
- (iv) $\text{CTL}_{\text{Hil}}^* \vdash \mathbf{E}(\alpha \vee \beta) \rightarrow (\mathbf{E}\alpha \vee \mathbf{E}\beta)$.

Proof. Clearly, (ii) and (iv) follow by contraposition from (iii) and (i), respectively. To prove (i) we reason as follows:

1. $\alpha \rightarrow [\beta \rightarrow (\alpha \wedge \beta)]$ (CPC)
2. $\mathbf{A}\alpha \rightarrow \mathbf{A}[\beta \rightarrow (\alpha \wedge \beta)]$ ($\mathbf{A} \rightarrow 1$)
3. $\mathbf{A}[\beta \rightarrow (\alpha \wedge \beta)] \rightarrow [\mathbf{A}\beta \rightarrow \mathbf{A}(\alpha \wedge \beta)]$ (axiom)
4. $\mathbf{A}\alpha \rightarrow [\mathbf{A}\beta \rightarrow \mathbf{A}(\alpha \wedge \beta)]$ (CPC 2,3)
5. $(\mathbf{A}\alpha \wedge \mathbf{A}\beta) \rightarrow \mathbf{A}(\alpha \wedge \beta)$ (CPC 4)

And we can prove (iii) thus:

1. $\alpha \rightarrow (\alpha \vee \beta)$ (CPC)
2. $\mathbf{E}\alpha \rightarrow \mathbf{E}(\alpha \vee \beta)$ ($\mathbf{E} \rightarrow 1$)
3. $\beta \rightarrow (\alpha \vee \beta)$ (CPC)
4. $\mathbf{E}\beta \rightarrow \mathbf{E}(\alpha \vee \beta)$ ($\mathbf{E} \rightarrow 3$)
5. $(\mathbf{E}\alpha \vee \mathbf{E}\beta) \rightarrow \mathbf{E}(\alpha \vee \beta)$ (CPC 2,4)

■

3.5.37. PROPOSITION. *For all formulas α and β and every $\mathbf{Q} \in \{\mathbf{A}, \mathbf{E}\}$, we have:*

- (i) $\text{CTL}_{\text{Hil}}^* \vdash \mathbf{A}(\alpha \vee \mathbf{Q}\beta) \rightarrow (\mathbf{A}\alpha \vee \mathbf{Q}\beta)$;
- (ii) $\text{CTL}_{\text{Hil}}^* \vdash (\mathbf{A}\alpha \vee \mathbf{Q}\beta) \rightarrow \mathbf{A}(\alpha \vee \mathbf{Q}\beta)$;
- (iii) $\text{CTL}_{\text{Hil}}^* \vdash (\mathbf{E}\alpha \wedge \mathbf{Q}\beta) \rightarrow \mathbf{E}(\alpha \wedge \mathbf{Q}\beta)$;
- (iv) $\text{CTL}_{\text{Hil}}^* \vdash \mathbf{E}(\alpha \wedge \mathbf{Q}\beta) \rightarrow (\mathbf{E}\alpha \wedge \mathbf{Q}\beta)$.

Proof. Clearly, (i) and (iv) follow by contraposition from (iii) and (ii), respectively. To prove (iii) we reason as follows:

1. $\mathbf{A}(\mathbf{Q}\beta \rightarrow \alpha^\partial) \rightarrow (\mathbf{A}\mathbf{Q}\beta \rightarrow \mathbf{A}\alpha^\partial)$ (axiom)
2. $\mathbf{Q}\beta \rightarrow \mathbf{A}\mathbf{Q}\beta$ (axiom)

$$3. A(Q\beta \rightarrow \alpha^\theta) \rightarrow (Q\beta \rightarrow A\alpha^\theta) \quad (\text{CPC 1,2})$$

$$4. (Q\beta \wedge E\alpha) \rightarrow E(Q\beta \wedge \alpha) \quad (\text{CPC 3})$$

To prove (ii) we reason as follows:

$$1. A\alpha \rightarrow \alpha \quad (\text{axiom})$$

$$2. A\alpha \rightarrow (\alpha \vee Q\beta) \quad (\text{CPC 1})$$

$$3. AA\alpha \rightarrow A(\alpha \vee Q\beta) \quad (A \rightarrow 2)$$

$$4. A\alpha \rightarrow AA\alpha \quad (\text{axiom})$$

$$5. A\alpha \rightarrow A(\alpha \vee Q\beta) \quad (\text{CPC 3,4})$$

$$6. Q\beta \rightarrow (\alpha \vee Q\beta) \quad (\text{CPC})$$

$$7. AQ\beta \rightarrow A(\alpha \vee Q\beta) \quad (A \rightarrow 6)$$

$$8. Q\beta \rightarrow AQ\beta \quad (\text{axiom})$$

$$9. Q\beta \rightarrow A(\alpha \vee Q\beta) \quad (\text{CPC 7,8})$$

$$10. (A\alpha \vee Q\beta) \rightarrow A(\alpha \vee Q\beta) \quad (\text{CPC 5,9})$$

■

3.5.38. PROPOSITION. *The rule*

$$X \rightarrow \frac{\alpha \rightarrow \beta}{X\alpha \rightarrow X\beta}$$

is admissible in $\text{CTL}_{\text{Hil}}^*$.

Proof. Assuming $\text{CTL}_{\text{ind}}^* \vdash \alpha \rightarrow \beta$, we can prove $\text{CTL}_{\text{Hil}}^* \vdash X\alpha \rightarrow X\beta$ as follows:

$$1. \alpha \rightarrow \beta \quad (\text{assumption})$$

$$2. X(\alpha \rightarrow \beta) \quad (X 1)$$

$$3. X(\alpha \rightarrow \beta) \rightarrow (X\alpha \rightarrow X\beta) \quad (\text{axiom})$$

$$4. X\alpha \rightarrow X\beta \quad (\text{CPC 2,3})$$

■

3.5.39. PROPOSITION. *For all formulas* α *and* β , *we have:*

$$(i) \text{CTL}_{\text{Hil}}^* \vdash (X\alpha \wedge X\beta) \rightarrow X(\alpha \wedge \beta);$$

$$(ii) \text{CTL}_{\text{Hil}}^* \vdash X(\alpha \wedge \beta) \rightarrow (X\alpha \wedge X\beta);$$

(iii) $\text{CTL}_{\text{Hil}}^* \vdash (\mathbf{X}\alpha \vee \mathbf{X}\beta) \rightarrow \mathbf{X}(\alpha \vee \beta)$;

(iv) $\text{CTL}_{\text{Hil}}^* \vdash \mathbf{X}(\alpha \vee \beta) \rightarrow (\mathbf{X}\alpha \vee \mathbf{X}\beta)$.

Proof. Clearly, (ii) and (iv) follow by contraposition from (iii) and (i), respectively.

To prove (i) we reason as follows:

1. $\alpha \rightarrow [\beta \rightarrow (\alpha \wedge \beta)]$ (CPC)
2. $\mathbf{X}\alpha \rightarrow \mathbf{X}[\beta \rightarrow (\alpha \wedge \beta)]$ ($\mathbf{X} \rightarrow 1$)
3. $\mathbf{X}[\beta \rightarrow (\alpha \wedge \beta)] \rightarrow [\mathbf{X}\beta \rightarrow \mathbf{X}(\alpha \wedge \beta)]$ (axiom)
4. $\mathbf{X}\alpha \rightarrow [\mathbf{X}\beta \rightarrow \mathbf{X}(\alpha \wedge \beta)]$ (CPC 2,3)
5. $(\mathbf{X}\alpha \wedge \mathbf{X}\beta) \rightarrow \mathbf{X}(\alpha \wedge \beta)$ (CPC 4)

And we can prove (iii) thus:

1. $\alpha \rightarrow (\alpha \vee \beta)$ (CPC)
2. $\mathbf{X}\alpha \rightarrow \mathbf{X}(\alpha \vee \beta)$ ($\mathbf{X} \rightarrow 1$)
3. $\beta \rightarrow (\alpha \vee \beta)$ (CPC)
4. $\mathbf{X}\beta \rightarrow \mathbf{X}(\alpha \vee \beta)$ ($\mathbf{X} \rightarrow 3$)
5. $(\mathbf{X}\alpha \vee \mathbf{X}\beta) \rightarrow \mathbf{X}(\alpha \vee \beta)$ (CPC 2,4)

■

We are now ready to embed $\text{CTL}_{\text{ind}}^*$ into $\text{CTL}_{\text{Hil}}^*$. By (strong) induction on $n < \omega$, we show that, for every (unannotated) hypersequent Γ having a $\text{CTL}_{\text{ind}}^*$ proof of height n ,¹⁵ $\text{CTL}_{\text{Hil}}^* \vdash \Gamma^\sharp$. Note that, by Proposition 3.5.35, we may choose any ordering of the formulas and the sequents whenever we apply the translation map $(\cdot)^\sharp$.

If $n = 0$, then Γ is an instance of a $\text{CTL}_{\text{ind}}^*$ axiom.

Suppose that Γ is an instance of $\mathbf{ax}_{\mathbf{AE}}$. Then, Γ^\sharp is clearly provable in $\text{CTL}_{\text{Hil}}^*$ by classical reasoning alone.

Suppose that Γ is an instance of $\mathbf{ax}_{\mathbf{RU}}$, say $\Gamma = \mathbf{A}\{\Phi, \alpha\mathbf{R}\beta, (\alpha\mathbf{R}\beta)^\partial\}, \Delta$. Then, $\Gamma^\sharp = \mathbf{A}(\mathbf{V}\Phi \vee \alpha\mathbf{R}\beta \vee (\alpha\mathbf{R}\beta)^\partial) \vee \Delta^\sharp$ and thus Γ^\sharp is provable in $\text{CTL}_{\text{Hil}}^*$ by classical reasoning and rule \mathbf{A} .

Suppose that Γ is an instance of $\mathbf{ax}_{\mathbf{T}}$, say $\Gamma = \mathbf{E}\emptyset, \Delta$. Then, $\Gamma^\sharp = \mathbf{E}\mathbf{T} \vee \Delta^\sharp = \mathbf{E}(p \vee \bar{p}) \vee \Delta^\sharp$ and we can prove Γ^\sharp as follows:

1. $\mathbf{A}(p \wedge \bar{p}) \rightarrow (p \wedge \bar{p})$ (axiom)

¹⁵Strictly speaking, we only defined the notion of $\text{CTL}_{\text{ind}}^*$ proof for formulas. It extends to sequents and hypersequents in the obvious manner.

2. $(p \wedge \bar{p})^\partial \rightarrow E(p \wedge \bar{p})^\partial$ (CPC 1)
3. $(p \wedge \bar{p})^\partial$ (CPC)
4. $E(p \wedge \bar{p})^\partial$ (MP 2,3)
5. $E(p \wedge \bar{p})^\partial \vee \Delta^\sharp$ (CPC 4)

Since $E(p \wedge \bar{p})^\partial = E(\bar{p} \vee p)$, we are done.

Suppose that Γ is of the form $Qp, Q^\partial \bar{p}, \Delta$. Then, $\Gamma^\sharp = Qp \vee (Qp)^\partial \vee \Delta$, which is provable in $\text{CTL}_{\text{Hil}}^*$ by classical reasoning.

Suppose that Γ is of the form $A\bar{p}, Ap, \Delta$. Then, $\Gamma^\sharp = (Ep \rightarrow Ap) \vee \Delta^\sharp$ and we can prove Γ^\sharp in $\text{CTL}_{\text{Hil}}^*$ thus:

1. $Ep \rightarrow p$ (axiom)
2. $p \rightarrow Ap$ (axiom)
3. $Ep \rightarrow Ap$ (CPC 1,2)
4. $(Ep \rightarrow Ap) \vee \Delta^\sharp$ (CPC 3)

To finish the base case, suppose that Γ is of the form $E\bar{p}, Ep, \Delta$. Then, $\Gamma^\sharp = (Ap \rightarrow Ep) \vee \Delta^\sharp$ and we can prove Γ^\sharp in $\text{CTL}_{\text{Hil}}^*$ thus:

1. $A\bar{p} \rightarrow \bar{p}$ (axiom)
2. $p \rightarrow Ep$ (CPC 1)
3. $Ap \rightarrow p$ (axiom)
4. $Ap \rightarrow Ep$ (CPC 2,3)
5. $(Ap \rightarrow Ep) \vee \Delta^\sharp$ (CPC 4)

For the inductive case we distinguish cases according to the lowermost rule R applied in the $\text{CTL}_{\text{ind}}^*$ proof \mathcal{T} , say at vertex $u \in T$.

Case R = ALit. In \mathcal{T} we have:

$$\text{ALit} \frac{A\Phi, A\ell, \Delta}{A\{\Phi, \ell\}, \Delta}$$

We can prove $A(\vee \Phi \vee \ell) \vee \Delta^\sharp$ in $\text{CTL}_{\text{Hil}}^*$ thus:

1. $(A\vee \Phi \vee A\ell) \vee \Delta^\sharp$ (IH)
2. $\vee \Phi \rightarrow (A\vee \Phi \vee \ell)$ (CPC)
3. $A\vee \Phi \rightarrow A(A\vee \Phi \vee \ell)$ (A \rightarrow 2)
4. $\ell \rightarrow (A\vee \Phi \vee \ell)$ (CPC)
5. $A\ell \rightarrow A(A\vee \Phi \vee \ell)$ (A \rightarrow 4)

$$6. (A \vee \Phi \vee A\ell) \rightarrow A(\vee \Phi \vee \ell) \quad (\text{CPC 3,5})$$

$$7. A(\vee \Phi \vee \ell) \vee \Delta^\# \quad (\vee \rightarrow 6)$$

Case R = ELit. In \mathcal{T} we have:

$$\text{ELit} \frac{E\Phi, \Delta \quad E\ell, \Delta}{E\{\Phi, \ell\}, \Delta}$$

We can prove $E(\wedge \Phi \wedge \ell) \vee \Delta^\#$ in $\text{CTL}_{\text{Hil}}^*$ thus:

$$1. E\wedge \Phi \vee \Delta^\# \quad (\text{IH})$$

$$2. E\ell \vee \Delta^\# \quad (\text{IH})$$

$$3. (\Delta^\#)^\theta \rightarrow (E\wedge \Phi \wedge E\ell) \quad (\text{CPC 1,2})$$

$$4. (E\wedge \Phi \wedge E\ell) \rightarrow E(\wedge \Phi \wedge E\ell) \quad (\text{prop. 3.5.37})$$

$$5. (\Delta^\#)^\theta \rightarrow E(\wedge \Phi \wedge E\ell) \quad (\text{CPC 3,4})$$

$$6. E\ell \rightarrow \ell \quad (\text{prop. 3.5.34})$$

$$7. (\wedge \Phi \wedge E\ell) \rightarrow (\wedge \Phi \wedge \ell) \quad (\text{CPC 6})$$

$$8. E(\wedge \Phi \wedge E\ell) \rightarrow E(\wedge \Phi \wedge \ell) \quad (E \rightarrow 7)$$

$$9. (\Delta^\#)^\theta \rightarrow E(\wedge \Phi \wedge \ell) \quad (\text{CPC 5,8})$$

Case R = AV. In \mathcal{T} we have:

$$\text{AV} \frac{A\{\Phi, \alpha, \beta\}, \Delta}{A\{\Phi, \alpha \vee \beta\}, \Delta}$$

By the inductive hypothesis $\text{CTL}_{\text{Hil}}^* \vdash A(\vee \Phi \vee \alpha \vee \beta) \vee \Delta^\#$, so we are done.

Case R = EV. In \mathcal{T} we have:

$$\text{EV} \frac{E\{\Phi, \alpha\}, E\{\Phi, \beta\}, \Delta}{E\{\Phi, \alpha \vee \beta\}, \Delta}$$

We can prove $E(\wedge \Phi \wedge (\alpha \vee \beta)) \vee \Delta^\#$ in $\text{CTL}_{\text{Hil}}^*$ thus, where $\gamma := \wedge \Phi$:

$$1. E(\gamma \wedge \alpha) \vee E(\gamma \wedge \beta) \vee \Delta^\# \quad (\text{IH})$$

$$2. [E(\gamma \wedge \alpha) \vee E(\gamma \wedge \beta)] \rightarrow E[(\gamma \wedge \alpha) \vee (\gamma \wedge \beta)] \quad (\text{prop. 3.5.36})$$

$$3. [(\gamma \wedge \alpha) \vee (\gamma \wedge \beta)] \rightarrow [\gamma \wedge (\alpha \vee \beta)] \quad (\text{CPC})$$

$$4. E[(\gamma \wedge \alpha) \vee (\gamma \wedge \beta)] \rightarrow E[\gamma \wedge (\alpha \vee \beta)] \quad (E \rightarrow 3)$$

$$5. E(\gamma \wedge (\alpha \vee \beta)) \vee \Delta^\# \quad (\text{CPC 1,2,4})$$

Case R = A \wedge . In \mathcal{T} we have:

$$A\wedge \frac{A\{\Phi, \alpha\}, \Delta \quad A\{\Phi, \beta\}, \Delta}{A\{\Phi, \alpha \wedge \beta\}, \Delta}$$

We can prove $A(\bigvee \Phi \vee (\alpha \wedge \beta)) \vee \Delta^\sharp$ in $\text{CTL}_{\text{Hil}}^*$ thus, where $\gamma := \bigvee \Phi$:

- (i) $A(\gamma \vee \alpha) \vee \Delta^\sharp$ (IH)
- (ii) $A(\gamma \vee \beta) \vee \Delta^\sharp$ (IH)
- (iii) $[A(\gamma \vee \alpha) \wedge A(\gamma \vee \beta)] \rightarrow A[(\gamma \vee \alpha) \wedge (\gamma \vee \beta)]$ (prop. 3.5.36)
- (iv) $[(\gamma \vee \alpha) \wedge (\gamma \vee \beta)] \rightarrow [\gamma \vee (\alpha \wedge \beta)]$ (CPC)
- (v) $A[(\gamma \vee \alpha) \wedge (\gamma \vee \beta)] \rightarrow A[\gamma \vee (\alpha \wedge \beta)]$ ($A \rightarrow 4$)
- (vi) $A(\gamma \vee (\alpha \wedge \beta)) \vee \Delta^\sharp$ (CPC 1,2,3,5)

Case R = $E\wedge$. In \mathcal{T} we have:

$$E\wedge \frac{E\{\Phi, \alpha, \beta\}, \Delta}{E\{\Phi, \alpha \wedge \beta\}, \Delta}$$

By the inductive hypothesis $\text{CTL}_{\text{Hil}}^* \vdash E(\wedge \Phi \wedge \alpha \wedge \beta) \vee \Delta^\sharp$, so we are done.

Case R = AA . In \mathcal{T} we have:

$$AA \frac{A\Phi, A\{\alpha\}, \Delta}{A\{\Phi, A\alpha\}, \Delta}$$

We can prove $A(\bigvee \Phi \vee A\alpha) \vee \Delta^\sharp$ in $\text{CTL}_{\text{Hil}}^*$ thus:

- (i) $A\bigvee \Phi \vee A\alpha \vee \Delta^\sharp$ (IH)
- (ii) $\bigvee \Phi \rightarrow (\bigvee \Phi \vee A\alpha)$ (CPC)
- (iii) $A\bigvee \Phi \rightarrow A(\bigvee \Phi \vee A\alpha)$ ($A \rightarrow 2$)
- (iv) $A\alpha \rightarrow (\bigvee \Phi \vee A\alpha)$ (CPC)
- (v) $AA\alpha \rightarrow A(\bigvee \Phi \vee A\alpha)$ ($A \rightarrow 4$)
- (vi) $A\alpha \rightarrow AA\alpha$ (axiom)
- (vii) $A\alpha \rightarrow A(\bigvee \Phi \vee A\alpha)$ (CPC 5,6)
- (viii) $A(\bigvee \Phi \vee A\alpha) \vee \Delta^\sharp$ (CPC 1,3,7)

Case R = EA . In \mathcal{T} we have:

$$EA \frac{E\Phi, \Delta \quad A\{\alpha\}, \Delta}{E\{\Phi, A\alpha\}, \Delta}$$

We can prove $E(\wedge \Phi \wedge A\alpha) \vee \Delta^\sharp$ in $\text{CTL}_{\text{Hil}}^*$ thus:

- (i) $E\wedge \Phi \vee \Delta^\sharp$ (IH)

$$(ii) \ A\alpha \vee \Delta^\sharp \quad (IH)$$

$$(iii) \ (\Delta^\sharp)^\partial \rightarrow (E \wedge \Phi \wedge A\alpha) \quad (CPC \ 1,2)$$

$$(iv) \ (E \wedge \Phi \wedge A\alpha) \rightarrow E(\wedge \Phi \wedge A\alpha) \quad (\text{prop. 3.5.37})$$

$$(v) \ (\Delta^\sharp)^\partial \rightarrow E(\wedge \Phi \wedge A\alpha) \quad (CPC \ 3,4)$$

Case R = AE. In \mathcal{T} we have:

$$AE \frac{A\Phi, E\{\alpha\}, \Delta}{A\{\Phi, E\alpha\}, \Delta}$$

We can prove $A(\vee \Phi \vee E\alpha) \vee \Delta^\sharp$ in CTL_{Hil}^* thus:

$$(i) \ A \vee \Phi \vee E\alpha \vee \Delta^\sharp \quad (IH)$$

$$(ii) \ \vee \Phi \rightarrow (\vee \Phi \vee E\alpha) \quad (CPC)$$

$$(iii) \ A \vee \Phi \rightarrow A(\vee \Phi \vee E\alpha) \quad (A \rightarrow 2)$$

$$(iv) \ E\alpha \rightarrow (\vee \Phi \vee E\alpha) \quad (CPC)$$

$$(v) \ AE\alpha \rightarrow A(\vee \Phi \vee E\alpha) \quad (A \rightarrow 4)$$

$$(vi) \ E\alpha \rightarrow AE\alpha \quad (\text{axiom})$$

$$(vii) \ E\alpha \rightarrow A(\vee \Phi \vee E\alpha) \quad (CPC \ 5,6)$$

$$(viii) \ A(\vee \Phi \vee E\alpha) \vee \Delta^\sharp \quad (CPC \ 1,3,7)$$

Case R = EE. In \mathcal{T} we have:

$$EE \frac{E\Phi, \Delta \quad E\{\alpha\}, \Delta}{E\{\Phi, E\alpha\}, \Delta}$$

We can prove $E(\wedge \Phi \wedge E\alpha) \vee \Delta^\sharp$ in CTL_{Hil}^* thus:

$$(i) \ E \wedge \Phi \vee \Delta^\sharp \quad (IH)$$

$$(ii) \ E\alpha \vee \Delta^\sharp \quad (IH)$$

$$(iii) \ (\Delta^\sharp)^\partial \rightarrow (E \wedge \Phi \wedge E\alpha) \quad (CPC \ 1,2)$$

$$(iv) \ (E \wedge \Phi \wedge E\alpha) \rightarrow E(\wedge \Phi \wedge E\alpha) \quad (\text{prop. 3.5.37})$$

$$(v) \ (\Delta^\sharp)^\partial \rightarrow E(\wedge \Phi \wedge E\alpha) \quad (CPC \ 3,4)$$

Case R = AX. In \mathcal{T} we have:

$$AX \frac{A\Phi, E\Psi_1, \dots, E\Psi_m}{A(X\Phi \cup \Xi), EX\Psi_1, \dots, EX\Psi_m, \Sigma}$$

Let $\varphi := \bigvee \Phi$, $\xi := \bigvee \Xi$, $\psi_i := \bigwedge \Psi_i$ for $1 \leq i \leq m$, and $\psi := \psi_1 \vee \dots \vee \psi_m$. We can prove $A(X\varphi \vee \xi) \vee E \wedge X\Psi_1 \vee \dots \vee E \wedge X\Psi_m \vee \Sigma^\sharp$ in $\text{CTL}_{\text{Hil}}^*$ as follows:

1. $A\varphi \vee E\psi_1 \vee \dots \vee E\psi_m$ (IH)
2. $[E\psi_1 \vee \dots \vee E\psi_m] \rightarrow E\psi$ (prop. 3.5.36)
3. $A\psi^\partial \rightarrow A\varphi$ (CPC 1,2)
4. $X(A\psi^\partial \rightarrow A\varphi)$ (X 3)
5. $XA\psi^\partial \rightarrow XA\varphi$ (X \rightarrow 4)
6. $XE\varphi^\partial \rightarrow XE\psi$ (CPC 5)
7. $AX\psi^\partial \rightarrow XA\psi^\partial$ (axiom)
8. $XE\psi \rightarrow EX\psi$ (CPC 7)
9. $XE\varphi^\partial \rightarrow EX\psi$ (CPC 6,8)
10. $AX\psi^\partial \rightarrow XA\varphi$ (CPC 9)
11. $AAAX\psi^\partial \rightarrow AXA\varphi$ (A \rightarrow 10)
12. $AX\psi^\partial \rightarrow AAX\psi^\partial$ (axiom)
13. $AX\psi^\partial \rightarrow AXA\varphi$ (CPC 11,12)
14. $A\varphi \rightarrow \varphi$ (axiom)
15. $XA\varphi \rightarrow X\varphi$ (X \rightarrow 14)
16. $XA\varphi \rightarrow (X\varphi \vee \xi)$ (CPC 15)
17. $AXA\varphi \rightarrow A(X\varphi \vee \xi)$ (A \rightarrow 16)
18. $EX\psi \vee A(X\varphi \vee \xi)$ (CPC 13,17)
19. $X(\psi_1 \vee \dots \vee \psi_m) \rightarrow (\bigwedge X\Psi_1 \vee \dots \vee \bigwedge X\Psi_m)$ (prop. 3.5.39)
20. $EX\psi \rightarrow E(\bigwedge X\Psi_1 \vee \dots \vee \bigwedge X\Psi_m)$ (E \rightarrow 19)
21. $E(\bigwedge X\Psi_1 \vee \dots \vee \bigwedge X\Psi_m) \rightarrow (E \wedge X\Psi_1 \vee \dots \vee E \wedge X\Psi_m)$ (prop. 3.5.36)
22. $A(X\varphi \vee \xi) \vee E \wedge X\Psi_1 \vee \dots \vee E \wedge X\Psi_m$ (CPC 18,20,21)
23. $A(X\varphi \vee \xi) \vee E \wedge X\Psi_1 \vee \dots \vee E \wedge X\Psi_m \vee \Sigma^\sharp$ (CPC 22)

Case R = EX. In \mathcal{T} we have:

$$\text{EX} \frac{E\Psi_1, \dots, E\Psi_m}{E \wedge X\Psi_1, \dots, E \wedge X\Psi_m, \Sigma}$$

Let $\psi_i := \bigwedge \Psi_i$ for $1 \leq i \leq m$, and $\psi := \psi_1 \vee \dots \vee \psi_m$. We can prove $E \wedge X\Psi_1 \vee \dots \vee E \wedge X\Psi_m \vee \Sigma^\sharp$ in $\text{CTL}_{\text{Hil}}^*$ as follows:

1. $E\psi_1 \vee \dots \vee E\psi_m$ (IH)

2. $[E\psi_1 \vee \dots \vee E\psi_m] \rightarrow E\psi$ (prop. 3.5.36)
3. $E\psi$ (CPC 1,2)
4. $XE\psi$ (X 3)
5. $AX\psi^\partial \rightarrow XA\psi^\partial$ (axiom)
6. $XE\psi \rightarrow EX\psi$ (CPC 5)
7. $EX\psi$ (CPC 4,6)
8. $X(\psi_1 \vee \dots \vee \psi_m) \rightarrow (\wedge X\Psi_1 \vee \dots \vee \wedge X\Psi_m)$ (prop. 3.5.39)
9. $EX\psi \rightarrow E(\wedge X\Psi_1 \vee \dots \vee \wedge X\Psi_m)$ (E \rightarrow 8)
10. $E(\wedge X\Psi_1 \vee \dots \vee \wedge X\Psi_m)$ (CPC 7,9)
11. $E(\wedge X\Psi_1 \vee \dots \vee \wedge X\Psi_m) \rightarrow (E\wedge X\Psi_1 \vee \dots \vee E\wedge X\Psi_m)$ (prop. 3.5.36)
12. $E\wedge X\Psi_1 \vee \dots \vee E\wedge X\Psi_m$ (CPC 10,11)
13. $E\wedge X\Psi_1 \vee \dots \vee E\wedge X\Psi_m \vee \Sigma^\sharp$ (CPC 12)

Case R = iW. In \mathcal{T} we have:

$$iW \frac{A\Phi, \Delta}{A(\Phi \cup \Psi), \Delta}$$

We can prove $A(\vee \Phi \vee \vee \Psi) \vee \Delta^\sharp$ in CTL_{Hil}^* thus:

- (i) $A \vee \Phi \vee \Delta^\sharp$ (IH)
- (ii) $\vee \Phi \rightarrow (\vee \Phi \vee \vee \Psi)$ (CPC)
- (iii) $A \vee \Phi \rightarrow A(\vee \Phi \vee \vee \Psi)$ (A \rightarrow 2)
- (iv) $A(\vee \Phi \vee \vee \Psi) \vee \Delta^\sharp$ (CPC 1,3)

Case R = eW. In \mathcal{T} we have:

$$eW \frac{\Gamma}{\Gamma, \Delta}$$

By the inductive hypothesis, $CTL_{Hil}^* \vdash \Gamma^\sharp$, so $CTL_{Hil}^* \vdash \Gamma^\sharp \vee \Delta^\sharp$ by Proposition 3.5.29.

Case R = AU. In \mathcal{T} we have:

$$AU \frac{A\{\Phi, \beta \vee [\alpha \wedge X(\alpha U \beta)]\}, \Delta}{A\{\Phi, \alpha U \beta\}, \Delta}$$

We can prove $A(\vee \Phi \vee \alpha U \beta) \vee \Delta^\sharp$ in CTL_{Hil}^* thus:

- (i) $A[\vee \Phi \vee \beta \vee [\alpha \wedge X(\alpha U \beta)]] \vee \Delta^\sharp$ (IH)
- (ii) $[\beta \vee (\alpha \wedge X(\alpha U \beta))] \rightarrow \alpha U \beta$ (axiom)

- (iii) $[\bigvee \Phi \vee \beta \vee (\alpha \wedge \mathbf{X}(\alpha \mathbf{U} \beta))] \rightarrow (\bigvee \Phi \vee \alpha \mathbf{U} \beta)$ (CPC 2)
 (iv) $\mathbf{A}[\bigvee \Phi \vee \beta \vee (\alpha \wedge \mathbf{X}(\alpha \mathbf{U} \beta))] \rightarrow \mathbf{A}(\bigvee \Phi \vee \alpha \mathbf{U} \beta)$ ($\mathbf{A} \rightarrow 3$)
 (v) $\mathbf{A}(\bigvee \Phi \vee \alpha \mathbf{U} \beta) \vee \Delta^\sharp$ (CPC 1,4)

Case $\mathbf{R} = \text{ind}$. In \mathcal{T} we have:

$$\text{ind} \frac{\mathbf{A}\{\Phi, \beta \wedge [\alpha \vee \mathbf{X}(\bigvee \Phi \vee \Delta^\sharp)^\partial]\}, \Delta}{\mathbf{A}\{\Phi, \alpha \mathbf{R} \beta\}, \Delta}$$

We can prove $\mathbf{A}(\bigvee \Phi \vee \alpha \mathbf{R} \beta) \vee \Delta^\sharp$ in $\text{CTL}_{\text{Hil}}^*$ thus:

- (i) $\mathbf{A}[\bigvee \Phi \vee (\beta \wedge (\alpha \vee \mathbf{X}(\bigvee \Phi \vee \Delta^\sharp)^\partial))] \vee \Delta^\sharp$ (IH)
 (ii) $\mathbf{A}[\bigvee \Phi \vee (\beta \wedge (\alpha \vee \mathbf{X}(\bigvee \Phi \vee \Delta^\sharp)^\partial)) \vee \Delta^\sharp]$ (1, prop. 3.5.37)
 (iii) $[\beta^\partial \vee (\alpha^\partial \wedge \mathbf{X}(\bigvee \Phi \vee \Delta^\sharp))] \rightarrow [\bigvee \Phi \vee \Delta^\sharp]$ ($\mathbf{A}^{-1} 2$)
 (iv) $\alpha^\partial \mathbf{U} \beta^\partial \rightarrow [\bigvee \Phi \vee \Delta^\sharp]$ (ind 3)
 (v) $\mathbf{A}(\alpha \mathbf{R} \beta \vee \bigvee \Phi \vee \Delta^\sharp)$ ($\mathbf{A} 4$)
 (vi) $\mathbf{A}(\alpha \mathbf{R} \beta \vee \bigvee \Phi) \vee \Delta^\sharp$ (5, prop. 3.5.37)

Case $\mathbf{R} = \perp^{-1}$. In \mathcal{T} we have:

$$\perp^{-1} \frac{\Gamma, \mathbf{A}\emptyset}{\Gamma}$$

By the inductive hypothesis, $\text{CTL}_{\text{Hil}}^* \vdash \Gamma^\sharp \vee \mathbf{A}\emptyset$, i.e., $\text{CTL}_{\text{Hil}}^* \vdash \Gamma^\sharp \vee \mathbf{A}(p \wedge \bar{p})$. We can then prove Γ^\sharp in $\text{CTL}_{\text{Hil}}^*$ thus:

- (i) $\Gamma^\sharp \vee \mathbf{A}(p \wedge \bar{p})$ (IH)
 (ii) $\mathbf{A}(p \wedge \bar{p}) \rightarrow (p \wedge \bar{p})$ (axiom)
 (iii) $\Gamma^\sharp \vee (p \wedge \bar{p})$ (CPC 1,2)
 (iv) Γ^\sharp (CPC 3)

Case $\mathbf{R} = \text{cut}$. In \mathcal{T} we have:

$$\text{cut} \frac{\mathbf{A}\{\Phi, \alpha\}, \Delta \quad \mathbf{A}\{\Phi, \alpha^\partial\}, \Delta}{\mathbf{A}\Phi, \Delta}$$

We can prove $\mathbf{A}\bigvee \Phi \vee \Delta^\sharp$ in $\text{CTL}_{\text{Hil}}^*$ thus:

- (i) $\mathbf{A}(\bigvee \Phi \vee \alpha) \vee \Delta^\sharp$ (IH)
 (ii) $\mathbf{A}(\bigvee \Phi \vee \alpha^\partial) \vee \Delta^\sharp$ (IH)

- (iii) $(\bigvee \Phi \vee \alpha) \rightarrow [(\bigvee \Phi \vee \alpha^\partial) \rightarrow \bigvee \Phi]$ (CPC)
- (iv) $A(\bigvee \Phi \vee \alpha) \rightarrow A[(\bigvee \Phi \vee \alpha^\partial) \rightarrow \bigvee \Phi]$ (A \rightarrow 3)
- (v) $A[(\bigvee \Phi \vee \alpha^\partial) \rightarrow \bigvee \Phi] \rightarrow [A(\bigvee \Phi \vee \alpha^\partial) \rightarrow A\bigvee \Phi]$ (axiom)
- (vi) $A(\bigvee \Phi \vee \alpha) \rightarrow [A(\bigvee \Phi \vee \alpha^\partial) \rightarrow A\bigvee \Phi]$ (CPC 4,5)
- (vii) $[A(\bigvee \Phi \vee \alpha) \vee \Delta^\sharp] \rightarrow [[A(\bigvee \Phi \vee \alpha^\partial) \vee \Delta^\sharp] \rightarrow (A\bigvee \Phi \vee \Delta^\sharp)]$ (CPC 6)
- (viii) $[A(\bigvee \Phi \vee \alpha^\partial) \vee \Delta^\sharp] \rightarrow (A\bigvee \Phi \vee \Delta^\sharp)$ (MP 1,7)
- (ix) $A\bigvee \Phi \vee \Delta^\sharp$ (MP 2,8)

This finishes the inductive step of the proof. We have thus shown:

3.5.40. PROPOSITION. *For any formula φ , if $\text{CTL}_{\text{ind}}^* \vdash \varphi$, then $\text{CTL}_{\text{Hil}}^* \vdash A\varphi$.*

Recall that $\text{CTL}_{\text{ind}}^* \vdash \varphi$ means that there is a $\text{CTL}_{\text{ind}}^*$ proof with conclusion $A\varphi$. The axiom schema $A\alpha \rightarrow \alpha$ of $\text{CTL}_{\text{Hil}}^*$ bridges the gap from (universal) sequents to formulas, allowing us to obtain the following from Proposition 3.5.40:

3.5.41. COROLLARY. *For any formula φ , if $\text{CTL}_{\text{ind}}^* \vdash \varphi$, then $\text{CTL}_{\text{Hil}}^* \vdash \varphi$.*

We are now going to establish the converse.

3.5.42. PROPOSITION. *The rule*

$$A\vee^{-1} \frac{A\{\Phi, \alpha \vee \beta\}, \Delta}{A\{\Phi, \alpha, \beta\}, \Delta}$$

is admissible in $\text{CTL}_{\text{ind}}^$.*

Proof. We have:

$$\text{cut} \frac{\text{iW} \frac{A\{\Phi, \alpha \vee \beta\}, \Delta}{A\{\Phi, \alpha, \beta, \alpha \vee \beta\}, \Delta} \quad A\wedge \frac{\text{LEM} \frac{A\{\Phi, \alpha, \beta, \alpha^\partial\}, \Delta}{A\{\Phi, \alpha, \beta, \beta^\partial\}, \Delta} \quad \text{LEM} \frac{A\{\Phi, \alpha, \beta, \beta^\partial\}, \Delta}{A\{\Phi, \alpha, \beta, \alpha^\partial \wedge \beta^\partial\}, \Delta}}{A\{\Phi, \alpha, \beta\}, \Delta}}{A\{\Phi, \alpha, \beta\}, \Delta}$$

■

3.5.43. PROPOSITION. *The rule*

$$\text{MP} \frac{A\{\Phi, \alpha\}, \Delta \quad A\{\Phi, \alpha \rightarrow \beta\}, \Delta}{A\{\Phi, \beta\}, \Delta}$$

is admissible in $\text{CTL}_{\text{ind}}^*$.

Proof. We have:

$$\text{cut} \frac{\text{iW} \frac{\text{A}\{\Phi, \alpha\}, \Delta}{\text{A}\{\Phi, \alpha, \beta\}, \Delta} \quad \text{AV}^{-1} \frac{\text{A}\{\Phi, \alpha \rightarrow \beta\}, \Delta}{\text{A}\{\Phi, \alpha^\partial, \beta\}, \Delta}}{\text{A}\{\Phi, \beta\}, \Delta}$$

■

3.5.44. LEMMA. For every formula α , we have $\text{CTL}_{\text{ind}}^* \vdash \text{A}\alpha \rightarrow \text{E}\alpha$.

Proof. We have:

$$\begin{array}{c} \text{aXAE} \frac{\text{AE} \frac{\text{E}\{\alpha^\partial\}, \text{A}\{\alpha\}}{\text{A}\{\text{A}\alpha \rightarrow \alpha\}}}{\text{A}\{\text{A}\alpha \rightarrow \text{E}\alpha\}} \\ \text{AE} \frac{\text{A}\{\text{A}\alpha \rightarrow \alpha\}}{\text{A}\{\text{A}\alpha \rightarrow \text{E}\alpha\}} \\ \text{MP} \frac{\text{AE} \frac{\text{E}\{\alpha^\partial\}, \text{A}\{\alpha\}}{\text{A}\{\text{A}\alpha \rightarrow \alpha\}}}{\text{A}\{\text{A}\alpha \rightarrow \text{E}\alpha\}} \end{array} \quad \begin{array}{c} \text{aXAE} \frac{\text{AE} \frac{\text{E}\{\alpha^\partial\}, \text{E}\{\alpha\}}{\text{A}\{\text{A}\alpha, \text{E}\alpha^\partial, \text{E}\alpha\}}}{\text{A}\{\text{A}\alpha \wedge \alpha^\partial, \text{E}\alpha^\partial, \text{E}\alpha\}} \\ \text{AE} \frac{\text{E}\{\alpha^\partial\}, \text{E}\{\alpha\}}{\text{A}\{\alpha^\partial, \text{E}\alpha^\partial, \text{E}\alpha\}} \\ \text{iW} \frac{\text{AE} \frac{\text{E}\{\alpha^\partial\}, \text{E}\{\alpha\}}{\text{A}\{\alpha^\partial, \text{E}\alpha^\partial, \text{E}\alpha\}}}{\text{A}\{\alpha^\partial, \text{E}\alpha^\partial\}, \text{E}\{\alpha\}} \\ \text{aXAE} \frac{\text{AE} \frac{\text{E}\{\alpha^\partial\}, \text{E}\{\alpha\}}{\text{A}\{\alpha^\partial, \text{E}\alpha^\partial, \text{E}\alpha\}}}{\text{A}\{\text{A}\alpha \wedge \alpha^\partial, \text{E}\alpha^\partial, \text{E}\alpha\}} \\ \text{AV} \frac{\text{A}\{\text{A}\alpha \wedge \alpha^\partial, \text{E}\alpha^\partial, \text{E}\alpha\}}{\text{A}\{\text{A}\alpha \wedge \alpha^\partial, \text{E}\alpha^\partial \vee \text{E}\alpha\}} \\ \text{AV} \frac{\text{A}\{\text{A}\alpha \wedge \alpha^\partial, \text{E}\alpha^\partial \vee \text{E}\alpha\}}{\text{A}\{(\text{A}\alpha \rightarrow \alpha) \rightarrow (\text{A}\alpha \rightarrow \text{E}\alpha)\}} \\ \text{AV} \frac{\text{A}\{(\text{A}\alpha \rightarrow \alpha) \rightarrow (\text{A}\alpha \rightarrow \text{E}\alpha)\}}{\text{A}\{\text{A}\alpha \rightarrow \text{E}\alpha\}} \end{array}$$

■

3.5.45. LEMMA. We have $\text{CTL}_{\text{ind}}^* \vdash \text{A}(\alpha \rightarrow \beta) \rightarrow (\text{A}\alpha \rightarrow \text{A}\beta)$ for all formulas α and β .

Proof. We are to prove $\text{E}(\alpha \wedge \beta^\partial) \vee \text{E}\alpha^\partial \vee \text{A}\beta$ in $\text{CTL}_{\text{ind}}^*$. We begin as follows, where $\Phi := \{\text{E}(\alpha \wedge \beta^\partial), \text{E}\alpha^\partial, \text{A}\beta\}$:

$$\text{MP} \frac{\text{\#1} \quad \text{iW} \frac{\text{A}\{\text{A}(\alpha \wedge \beta^\partial) \rightarrow \text{E}(\alpha \wedge \beta^\partial)\}}{\text{A}\{\Phi, \text{A}(\alpha \wedge \beta^\partial) \rightarrow \text{E}(\alpha \wedge \beta^\partial)\}}}{\text{A}\{\text{E}(\alpha \wedge \beta^\partial), \text{E}\alpha^\partial, \text{A}\beta\}} \quad \text{AV} \frac{\text{A}\{\text{E}(\alpha \wedge \beta^\partial), \text{E}\alpha^\partial, \text{A}\beta\}}{\text{A}\{\text{E}(\alpha \wedge \beta^\partial) \vee \text{E}\alpha^\partial \vee \text{A}\beta\}}$$

The hypersequent labelling the right leaf is provable in $\text{CTL}_{\text{ind}}^*$ by Lemma 3.5.44. From #1 we proceed as follows:

$$\begin{array}{c}
\text{axAE} \frac{}{A\{\alpha\}, E\{\alpha \wedge \beta^\partial\}, E\{\alpha^\partial\}, A\{\beta\}} \quad \text{eW} \frac{E\{\alpha \wedge \beta^\partial\}, E\{\alpha^\partial\}, A\{\beta\}}{A\{\beta^\partial\}, E\{\alpha \wedge \beta^\partial\}, E\{\alpha^\partial\}, A\{\beta\}} \\
\text{A}\wedge \frac{}{A\{\alpha \wedge \beta^\partial\}, E\{\alpha \wedge \beta^\partial\}, E\{\alpha^\partial\}, A\{\beta\}} \\
\hline
A\{A(\alpha \wedge \beta^\partial), E(\alpha \wedge \beta^\partial), E\alpha^\partial, A\beta\}
\end{array}$$

And finally:

$$\begin{array}{c}
\text{axAE} \frac{}{E\{\alpha \wedge \beta^\partial\}, E\{\alpha^\partial\}, A\{\alpha\}} \quad \text{axAE} \frac{}{E\{\alpha, \beta^\partial\}, E\{\alpha^\partial\}, A\{\alpha^\partial, \beta\}} \\
\text{iW} \frac{}{E\{\alpha \wedge \beta^\partial\}, E\{\alpha^\partial\}, A\{\alpha, \beta\}} \quad \text{E}\wedge \frac{}{E\{\alpha \wedge \beta^\partial\}, E\{\alpha^\partial\}, A\{\alpha^\partial, \beta\}} \\
\text{cut} \frac{}{E\{\alpha \wedge \beta^\partial\}, E\{\alpha^\partial\}, A\{\beta\}}
\end{array}$$

■

It is straightforward to see that every $\text{CTL}_{\text{Hil}}^*$ rule is admissible in $\text{CTL}_{\text{ind}}^*$ (for ind we rely on the admissibility of rule AV^{-1}). Similarly, every $\text{CTL}_{\text{Hil}}^*$ axiom schema except for axiom (iv) is readily provable in $\text{CTL}_{\text{ind}}^*$. And Lemma 3.5.45 shows that axiom (iv) is also provable in $\text{CTL}_{\text{ind}}^*$. Therefore:

3.5.46. PROPOSITION. *For any formula φ , if $\text{CTL}_{\text{Hil}}^* \vdash \varphi$, then $\text{CTL}_{\text{ind}}^* \vdash \varphi$.*

Putting Corollary 3.5.41 and Proposition 3.5.46 together, we have:

3.5.47. THEOREM. *For any formula φ , $\text{CTL}_{\text{ind}}^* \vdash \varphi$ if, and only if, $\text{CTL}_{\text{Hil}}^* \vdash \varphi$.*

The system $\text{CTL}_{\text{ind}}^*$ (and thus, by Theorem 3.5.47, also $\text{CTL}_{\text{Hil}}^*$) is easily seen to be stronger than CTL_\circ^* restricted to inductible proofs. For example, there is clearly no inductible CTL_\circ^* proof of EGT , but we can prove EGT in $\text{CTL}_{\text{ind}}^*$ as follows:

$$\begin{array}{c}
\text{AX} \frac{}{A\{\top\}} \\
\text{AV} \frac{}{A\{\perp, \text{XT}\}} \\
\text{A}\wedge \frac{}{A\{\top\}} \frac{}{A\{\perp \vee \text{XT}\}} \\
\text{ind} \frac{}{A\{\top \wedge (\perp \vee \text{XT})\}} \\
\text{AA} \frac{}{A\{\perp \text{RT}\}} \\
\text{iW} \frac{}{A\{\text{AGT}\}} \\
\text{AV}^{-1} \frac{}{A\{\text{AGT} \rightarrow \text{EGT}\}} \\
\text{cut} \frac{}{A\{(\text{AGT})^\partial, \text{EGT}\}} \\
\hline
A\{\text{EGT}\}
\end{array}$$

The hypersequent labelling the left leaf is provable in $\text{CTL}_{\text{ind}}^*$ by Lemma 3.5.44, and clearly $\mathbf{A}\{\top\}$ is provable in $\text{CTL}_{\text{ind}}^*$ as well.

The use of *cut* is essential in the previous proof, for it is easy to see that there is no cut-free proof of $\text{EG}\top$ in $\text{CTL}_{\text{ind}}^*$. Due to our extensive recourse to *cut* in the translation from inductible $\text{CTL}_{\text{ind}}^*$ proofs into $\text{CTL}_{\text{ind}}^*$ proofs above, it is most likely the case that $\text{CTL}_{\text{ind}}^*$ is not equivalent to $\text{CTL}_{\text{ind}}^*$ restricted to inductible proofs. This raises the following question:

3.5.48. QUESTION. Does $\text{CTL}_{\text{ind}}^*$ embed into the inductible fragment of $\text{CTL}_{\text{ind}}^*$?

This is essentially asking: if there is an inductible cyclic proof of the premise of the induction rule *ind*, is there also an inductible cyclic proof of the conclusion?

As mentioned in Section 3.1, formulas of CTL^* admit other interpretations than the standard one. Following [115], we obtain a general semantics by evaluating formulas on *path frames*, that is, pairs of the form (S, Π) , where S is a non-empty set of states and $\Pi \subseteq S^\omega$ is a collection of (ω -)paths on S . Satisfaction of CTL^* formulas on paths of path frames (equipped with a valuation for propositional letters) is defined as for serial models (p. 67).

Different semantics for CTL^* formulas can be obtained from the path-frame one by restricting the collection of paths allowed in path frames. Of particular relevance are the following constraints, where (S, Π) is a path frame:

- *suffix closure* (SC): if $\pi \in \Pi$ and π' is a tail of π , then $\pi' \in \Pi$;
- *fusion closure* (FC): for every $\pi_0, \pi_1 \in S^{<\omega}$, every $s \in S$, and $\pi'_0, \pi'_1 \in S^\omega$, if $\pi_0 \hat{\ } (s) \hat{\ } \pi'_0 \in \Pi$ and $\pi_1 \hat{\ } (s) \hat{\ } \pi'_1 \in \Pi$, then $\pi_0 \hat{\ } (s) \hat{\ } \pi'_1 \in \Pi$;
- *limit closure* (LC): if $\pi \in S^\omega$ is such that for every $n < \omega$ there is a $\pi' \in \Pi$ satisfying $\pi_{\leq n} = \pi'_{\leq n}$, then $\pi \in \Pi$.

Clearly, the collection of all maximal paths on a serial model satisfies all three conditions. The converse is also true, i.e., the standard semantics based on serial models is exactly the one based on path frames satisfying SC, FC and LC [41].

Dropping the LC condition yields a strictly more generous notion of validity than the standard one. Let $\text{CTL}^* \text{- LC}$ be the logic obtained by evaluating CTL^* formulas on path frames satisfying SC and FC, and let us write $\text{CTL}^* \text{- LC} \models \varphi$ if φ is valid under this notion of validity. Then, $\text{CTL}^* \text{- LC} \models \varphi$ implies $\text{CTL}^* \models \varphi$, but the converse is not true in general: the formula $\alpha := \text{AG}(p \rightarrow \text{EX}p) \rightarrow (p \rightarrow \text{EG}p)$ is valid in the standard semantics, but $\text{CTL}^* \text{- LC} \not\models \alpha$ [115, § 4.8].

Stirling [142, 143] gave an axiomatisation of $\text{CTL}^* \text{- LC}$. The problem of axiomatising the standard, full CTL^* logic, however, remained open for several years until Reynolds presented a sound and complete axiom system for it in [115]. Reynolds's axiomatisation consists of a set of axioms and rules for $\text{CTL}^* \text{- LC}$ (slightly different than Stirling's) plus a limit closure axiom and an 'Auxiliary Atoms' rule.

We shall now see that our Hilbert-style system $\text{CTL}_{\text{Hil}}^*$ (hence also $\text{CTL}_{\text{ind}}^*$) is complete at least for $\text{CTL}^* - \text{LC}$.

Reynolds formulates his system for $\text{CTL}^* - \text{LC}$ in the language $\{\top, \neg, \wedge, X, U, E\}$, with the usual definitions for derived constants, connectives, operators and quantifiers. The system has rules

$$\text{MP} \frac{\alpha \quad \alpha \rightarrow \beta}{\beta} \quad \text{G} \frac{\alpha}{\text{G}\alpha} \quad \text{A} \frac{\alpha}{\text{A}\alpha}$$

and the following axioms:

- (C0) all CPC tautologies in the language of CTL^* ;
- (C1) $\text{F}\neg\alpha \leftrightarrow \text{F}\alpha$;
- (C2) $\text{G}(\alpha \rightarrow \beta) \rightarrow (\text{G}\alpha \rightarrow \text{G}\beta)$;
- (C3) $\text{G}\alpha \rightarrow (\alpha \wedge \text{X}\alpha \wedge \text{XG}\alpha)$;
- (C4) $\text{X}\neg\alpha \leftrightarrow \neg\text{X}\alpha$;
- (C5) $\text{X}(\alpha \rightarrow \beta) \rightarrow (\text{X}\alpha \rightarrow \text{X}\beta)$;
- (C6) $\text{G}(\alpha \rightarrow \text{X}\alpha) \rightarrow (\alpha \rightarrow \text{G}\alpha)$;
- (C7 \rightarrow) $(\alpha \text{U}\beta) \rightarrow (\beta \vee (\alpha \wedge \text{X}(\alpha \text{U}\beta)))$;
- (C7 \leftarrow) $(\alpha \text{U}\beta) \leftarrow (\beta \vee (\alpha \wedge \text{X}(\alpha \text{U}\beta)))$;
- (C8) $(\alpha \text{U}\beta) \rightarrow \text{F}\beta$;
- (C9) $\text{A}(\alpha \rightarrow \beta) \rightarrow (\text{A}\alpha \rightarrow \text{A}\beta)$;
- (C10) $\text{A}\alpha \rightarrow \text{AA}\alpha$;
- (C11) $\text{A}\alpha \rightarrow \alpha$;
- (C12) $\alpha \rightarrow \text{AE}\alpha$;
- (C13) $\text{A}\neg\alpha \leftrightarrow \neg\text{E}\alpha$;
- (C14) $p \rightarrow \text{A}p$;
- (C15) $\text{A}\text{X}\alpha \rightarrow \text{X}\text{A}\alpha$.

Rules MP and A of Reynolds's system are also rules of $\text{CTL}_{\text{Hil}}^*$. And rule G is admissible in $\text{CTL}_{\text{Hil}}^*$ because it is so in $\text{CTL}_{\text{ind}}^*$:

$$\text{A}\wedge \frac{\text{A}\{\alpha\} \quad \text{A}\wedge \frac{\text{A}\{\top\} \quad \text{A}\wedge \frac{\text{A}\{\perp, \text{X}\top\}}{\text{A}\{\perp \vee \text{X}\top\}}}{\text{A}\{\perp \vee \text{X}\top\}}}{\text{A}\{\alpha \wedge (\perp \vee \text{X}\perp^\theta)\}} \quad \text{ind} \frac{\text{A}\{\alpha \wedge (\perp \vee \text{X}\perp^\theta)\}}{\text{A}\{\perp \text{R}\alpha\}}$$

Let us then focus on the axioms. By Proposition 3.5.29, all instances of C0 are provable in $\text{CTL}_{\text{Hil}}^*$. So are C1, C4 and C13 because we work in positive normal form and $(\cdot)^\partial$ is involutive. C5, C7 \leftarrow , C9, C10, C11, C14 and C15 are axiom schemata of $\text{CTL}_{\text{Hil}}^*$, and C12 follows immediately. Hence, we just need to show that C2, C3, C6, C7 \rightarrow , and C8 are provable in $\text{CTL}_{\text{Hil}}^*$. By Theorem 3.5.47, we may work in $\text{CTL}_{\text{ind}}^*$ instead.

We shall use the following rules for F and G, easily seen to be admissible in $\text{CTL}_{\text{ind}}^*$:

$$\text{AF} \frac{A\{\Phi, \alpha, XFa\}, \Delta}{A\{\Phi, Fa\}, \Delta} \quad \text{AG} \frac{A\{\Phi, \alpha\}, \Delta \quad A\{\Phi, XGa\}, \Delta}{A\{\Phi, Ga\}, \Delta}$$

Proof in $\text{CTL}_{\text{ind}}^*$ of axiom schema C2, where dashed lines indicate the omission of instances of LEM:

$$\begin{array}{c} \text{LEM} \frac{}{A\{F(\alpha \wedge \beta^\partial), Fa^\partial, G(\alpha \rightarrow \beta)\}} \quad \text{LEM} \frac{}{A\{F(\alpha \wedge \beta^\partial), Fa^\partial, Ga\}} \\ \text{A}\wedge \frac{}{A\{F(\alpha \wedge \beta^\partial), Fa^\partial, G(\alpha \rightarrow \beta) \wedge Ga\}} \\ \text{AX} \frac{}{A\{\beta^\partial, XF(\alpha \wedge \beta^\partial), \alpha^\partial, XFa^\partial, \perp, X(G(\alpha \rightarrow \beta) \wedge Ga)\}} \\ \text{AV} \frac{}{A\{\beta^\partial, XF(\alpha \wedge \beta^\partial), \alpha^\partial, XFa^\partial, \perp \vee X(G(\alpha \rightarrow \beta) \wedge Ga)\}} \\ \text{A}\wedge \frac{}{A\{\beta^\partial, XF(\alpha \wedge \beta^\partial), \alpha^\partial, XFa^\partial, \beta \wedge [\perp \vee X(G(\alpha \rightarrow \beta) \wedge Ga)]\}} \\ \text{A}\wedge \frac{}{A\{\alpha \wedge \beta^\partial, XF(\alpha \wedge \beta^\partial), \alpha^\partial, XFa^\partial, \beta \wedge [\perp \vee X(G(\alpha \rightarrow \beta) \wedge Ga)]\}} \\ \text{AF} \frac{}{A\{\alpha \wedge \beta^\partial, XF(\alpha \wedge \beta^\partial), Fa^\partial, \beta \wedge [\perp \vee X(G(\alpha \rightarrow \beta) \wedge Ga)]\}} \\ \text{AF} \frac{}{A\{F(\alpha \wedge \beta^\partial), Fa^\partial, \beta \wedge [\perp \vee X(G(\alpha \rightarrow \beta) \wedge Ga)]\}} \\ \text{ind} \frac{}{A\{F(\alpha \wedge \beta^\partial), Fa^\partial, \perp R, \beta\}} \\ \text{AV} \frac{}{A\{G(\alpha \rightarrow \beta) \rightarrow (Ga \rightarrow G\beta)\}} \end{array}$$

Proof in $\text{CTL}_{\text{ind}}^*$ of axiom schema C3:

$$\begin{array}{c} \text{LEM} \frac{}{A\{\alpha^\partial, XFa^\partial, \alpha\}} \\ \text{AF} \frac{}{A\{Fa^\partial, \alpha\}} \\ \text{AX} \frac{}{A\{\alpha^\partial, XFa^\partial, X\alpha\}} \\ \text{A}\wedge \frac{}{A\{\alpha^\partial, XFa^\partial, X\alpha \wedge XGa\}} \\ \text{A}\wedge \frac{}{A\{\alpha^\partial, XFa^\partial, \alpha \wedge X\alpha \wedge XGa\}} \\ \text{AF} \frac{}{A\{Fa^\partial, \alpha \wedge X\alpha \wedge XGa\}} \\ \text{AV} \frac{}{A\{Ga \rightarrow (\alpha \wedge X\alpha \wedge XGa)\}} \end{array}$$

Proof in $\text{CTL}_{\text{ind}}^*$ of axiom schema C6:

$$\begin{array}{c}
\text{LEM} \frac{}{A\{\alpha^\partial, F(\alpha \wedge X\alpha^\partial), G(\alpha \rightarrow X\alpha)\}} \quad \text{LEM} \frac{}{A\{\alpha^\partial, F(\alpha \wedge X\alpha^\partial), \alpha\}} \\
\text{A}\wedge \frac{}{A\{\alpha^\partial, F(\alpha \wedge X\alpha^\partial), G(\alpha \rightarrow X\alpha) \wedge \alpha\}} \\
\text{AX} \frac{}{A\{X\alpha^\partial, XF(\alpha \wedge X\alpha^\partial), \alpha^\partial, \perp, X(G(\alpha \rightarrow X\alpha) \wedge \alpha)\}} \\
\text{A}\wedge \frac{}{A\{\alpha \wedge X\alpha^\partial, XF(\alpha \wedge X\alpha^\partial), \alpha^\partial, \perp, X(G(\alpha \rightarrow X\alpha) \wedge \alpha)\}} \\
\text{AF} \frac{}{A\{F(\alpha \wedge X\alpha^\partial), \alpha^\partial, \perp, X(G(\alpha \rightarrow X\alpha) \wedge \alpha)\}} \\
\text{A}\vee \frac{}{A\{F(\alpha \wedge X\alpha^\partial), \alpha^\partial, \perp \vee X(G(\alpha \rightarrow X\alpha) \wedge \alpha)\}} \\
\text{A}\wedge \frac{}{A\{F(\alpha \wedge X\alpha^\partial), \alpha^\partial, \alpha \wedge [\perp \vee X(G(\alpha \rightarrow X\alpha) \wedge \alpha)]\}} \\
\text{ind} \frac{}{A\{F(\alpha \wedge X\alpha^\partial), \alpha^\partial, \perp R\alpha\}} \\
\text{A}\vee \frac{}{A\{G(\alpha \rightarrow X\alpha) \rightarrow (\alpha \rightarrow G\alpha)\}}
\end{array}$$

Proof in $\text{CTL}_{\text{ind}}^*$ of axiom schema C7 \rightarrow :

$$\begin{array}{c}
\text{LEM} \frac{}{A\{\beta^\partial \wedge (\alpha^\partial \vee X(\alpha^\partial R\beta^\partial)), \beta \vee (\alpha \wedge X(\alpha U\beta))\}} \\
\text{AR} \frac{}{A\{\alpha^\partial R\beta^\partial, \beta \vee (\alpha \wedge X(\alpha U\beta))\}} \\
\text{A}\vee \frac{}{A\{(\alpha U\beta) \rightarrow (\beta \vee (\alpha \wedge X(\alpha U\beta)))\}}
\end{array}$$

Proof in $\text{CTL}_{\text{ind}}^*$ of axiom schema C8:

$$\begin{array}{c}
\text{LEM} \frac{}{A\{G\beta^\partial, F\beta\}} \\
\text{AX} \frac{}{A\{\alpha^\partial, X(G\beta^\partial), \beta, XF\beta\}} \\
\text{LEM} \frac{}{A\{\beta^\partial, \beta, XF\beta\}} \quad \text{A}\vee \frac{}{A\{\alpha^\partial \vee X(G\beta^\partial), \beta, XF\beta\}} \\
\text{A}\wedge \frac{}{A\{\beta^\partial \wedge (\alpha^\partial \vee X(G\beta^\partial)), \beta, XF\beta\}} \\
\text{AF} \frac{}{A\{\beta^\partial \wedge (\alpha^\partial \vee X(G\beta^\partial)), F\beta\}} \\
\text{ind} \frac{}{A\{\alpha^\partial R\beta^\partial, F\beta\}} \\
\text{A}\vee \frac{}{A\{(\alpha U\beta) \rightarrow F\beta\}}
\end{array}$$

We have thus established:

3.5.49. THEOREM. *For every formula φ , if $\text{CTL}^* - \text{LC} \models \varphi$, then $\text{CTL}_{\text{Hil}}^* \vdash \varphi$.*

As previously mentioned, Reynolds's axiomatisation for full CTL^* extends his system for $\text{CTL}^* - \text{LC}$ by adding a limit closure axiom and the Auxiliary Atoms rule.

We leave open the question of whether $\text{CTL}_{\text{Hil}}^*$ (equivalently, $\text{CTL}_{\text{ind}}^*$) can prove the axiom and derive the rule. The answer to the former is probably positive, but we conjecture the latter not to be the case: the (side condition of) the Auxiliary Atoms rule is, in Reynolds's words, 'slightly complicated' [115, § 7], and its power seems to be beyond our inductive system.

3.6 Conclusion

Full computation tree logic, denoted by CTL^* , is a branching-time logic which extends linear-time temporal logic (LTL) by means of the universal (A) and existential (E) path quantifiers from computation tree logic (CTL). In contrast to the latter, it imposes no restriction on the placement of quantifiers or temporal operators, resulting in a strictly more expressive logic than LTL and CTL.

We provided a cut-free, ill-founded hypersequent calculus CTL_{∞}^* for CTL^* . Proofs are trees of possibly infinite height and a correctness condition ensures that infinite branches yield valid conclusions. Hypersequents offer a natural framework for accommodating the existential and universal path quantifiers of the logic, as well as their interplay with the next operator X. Soundness of the system is shown by a signature-based argument. Completeness follows from a deterministic proof-search by means of game- and automata-theoretic tools.

Our proof-search mechanism for CTL_{∞}^* yields proofs which are both deterministic, in the sense that the rule applied at a vertex is uniquely determined by the hypersequent in its label, and regular. This allows a partially finitary presentation of CTL_{∞}^* proofs. We formalised this idea by introducing the infinitary cyclic system $\text{CTL}_{\text{reg}}^*$.

By annotating formulas to keep track of fixpoint unfoldings, as Jungteerapanich [76] and Stirling [144] did for the modal μ -calculus, we were able to isolate a finitary condition (universality) that suffices to guarantee that a derivation is a proof. Further work is required to see whether good existential traces admit a finitary description. Indeed, whereas infinite formula traces are easy to detect by finitary means, good E-traces require the *absence* of infinite U-traces. This seems to be beyond the capabilities of our annotations.

A natural step towards a fully finitary cyclic calculus would be encoding automata in the ill-founded system. Jungteerapanich and Stirling's annotations derive ultimately from Safra's determinisation construction for Büchi automata. As we have shown, there exists an NBA \mathcal{A} which recognises exactly the good branches of CTL_{∞}^* proof-search trees. It might then be possible to incorporate the states and transitions of a deterministic version of \mathcal{A} into the rules of CTL_{∞}^* by means of a more complex annotating mechanism than the one we used, thus obtaining a cyclic calculus with finitary correctness conditions. This idea is partially applied

in [47] to a tableau for satisfiability of CTL* formulas, and [37] investigates the correspondence between annotations and determinisation procedures other than Safra's construction.

In the second half of the chapter we isolated a class of 'inductible' cyclic proofs whose cycles can be transformed into inductive arguments. The resulting Hilbert-style axiom system is shown to be complete at least for CTL* - LC, a well-known variant of CTL* obtained by allowing the evaluation of formulas in a bigger class than that of serial models. This work was motivated as a means to examine the possibility of axiomatising CTL* by eliminating cycles. Our cyclic system is not fully finitary and it could not therefore yield an axiomatisation of CTL* in this manner. But obtaining a finitary system would only be the first step: in order to remove purely existential cycles, i.e., cycles whose correctness is not witnessed by any A-trace, we would have to find a suitable induction rule for E-sequents. While rules similar to *ind* can be shown to be sound for E-sequents, they seem to be too rigid for the proof transformations carried out in Section 3.5.

Lastly, and on a different note, we could try to extend our work to CTL* augmented with past-time operators. Several alternatives present themselves, for example: a finite or an infinite past; a branching or a linear past. We refer the reader to the introduction of [116] for a survey of these and other proposals that have been considered in the literature. In [116], Reynolds obtained an axiomatisation of CTL* with past-time operators for a linear, finite past. It is worth noting that the past operators allowed him to dispense with the Auxiliary Atoms rule of his axiom system for CTL*. They might thus yield simpler systems than the ones we have considered.

Chapter 4

A Cyclic Proof System for iLTL

This chapter is at the same time easier and more complicated than the first half of the previous one. Here we consider *intuitionistic* linear-time temporal logic iLTL, which combines Pnueli’s linear-time temporal logic and (propositional) intuitionistic logic. Models of iLTL are ‘two-dimensional’: every state has at least one temporal successor and zero or more intuitionistic successors. A simple confluence condition, expressing monotonicity of the temporal successor function with respect to the intuitionistic order, suffices to ensure that truth is preserved upwards in the intuitionistic dimension (see [10]). The temporal component of iLTL is thus much simpler than CTL* and will pose no difficulty. The intuitionistic component, however, requires considerable care.

We provide a cut-free, finitary cyclic system for iLTL. The calculus uses labelled formulas in order to accommodate the interplay between the temporal dimension, represented by the modal rule for the next operator X , and the intuitionistic dimension, corresponding to the right-implication rule.

Labelled calculi incorporate the semantics of the accessibility relation into the syntax of the sequents. The idea traces back to Kanger’s sequent calculus with ‘marks’ for S5 [79], and it has since proved to be adequate for the proof theory of modal logics, where traditional sequent calculi meet with considerable obstacles (see, e.g., [99, 100]). Recent applications of labelled calculi include the ill-founded labelled calculus for the intuitionistic version of Gödel–Löb provability logic presented in [35], and the proof of the decidability of intuitionistic S4 in [58].

Our approach, in a sense, is the opposite of the one followed in [5]. There, the authors introduce an ill-founded proof system for a proper fragment of iLTL based on nested sequents that reason under the temporal next operator X . Their proof-search mechanism is thus able to look arbitrarily far into the future, but in some cases it is unable to break down implications in the consequent due to the

non-invertibility of their right-implication rule. As a result, it outputs ill-founded proofs which need not be regular.

Labels, instead, yield an invertible right-implication rule. This makes it possible — though not without considerable care — to design a proof-search algorithm for an ill-founded calculus which explores all relevant regions of the proof-search space, spending only a finite amount of time on each ‘intuitionistic region’. The resulting ill-founded proofs can be regularised and made cyclic by a simple annotation procedure.

4.1 Intuitionistic linear-time temporal logic iLTL

The *language of iLTL*, denoted by $\mathcal{L}_{\text{iLTL}}$, consists of the following: countably many *propositional letters* drawn from a set \mathbf{Prop} ; the *constant* \perp (falsum); the *Boolean connectives* \wedge (conjunction), \vee (disjunction) and \rightarrow (implication); and the *temporal operators* \mathbf{X} (*next*), \mathbf{U} (*until*) and \mathbf{R} (*release*). The *formulas* of iLTL are given by the following grammar:

$$\varphi ::= \perp \mid p \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi) \mid (\varphi \rightarrow \varphi) \mid (\mathbf{X}\varphi) \mid (\varphi\mathbf{U}\varphi) \mid (\varphi\mathbf{R}\varphi),$$

where p ranges over \mathbf{Prop} . Formulas are denoted by small Greek letters $\alpha, \beta, \varphi, \dots$, and sets of formulas by capital Greek letters $\Gamma, \Delta, \Sigma, \dots$. We use \mathbf{O} to denote either \mathbf{U} or \mathbf{R} . The collection of all iLTL formulas is denoted by $\mathbf{Form}_{\text{iLTL}}$, and the set of all *subformulas* of a formula φ , defined as expected, is denoted by $\mathbf{Sub}(\varphi)$.

If no ambiguity arises, we drop the outer parenthesis and stipulate that \mathbf{X} bind more strongly than \wedge, \vee, \mathbf{U} and \mathbf{R} , and that these, in turn, bind more strongly than \rightarrow .

A *literal* is a formula of the form \perp or p . An *X-formula*, or *modal formula*, is a formula of the form $\mathbf{X}\varphi$. Similarly, an *O-formula*, for $\mathbf{O} \in \{\mathbf{U}, \mathbf{R}\}$, is a formula of the form $\varphi\mathbf{O}\psi$. The *negation* of a formula φ is defined as $\neg\varphi := \varphi \rightarrow \perp$. We let $\mathbf{T} := \neg\perp$. We also define the temporal operators \mathbf{F} (*eventually*) and \mathbf{G} (*henceforth*) as expected: $\mathbf{F}\varphi := \mathbf{T}\mathbf{U}\varphi$, and $\mathbf{G}\varphi := \perp\mathbf{R}\varphi$. The connective \neg and the operators \mathbf{F} and \mathbf{G} bind all as strongly as \mathbf{X} .

The *size*, or *complexity*, of a formula φ , denoted by $\langle\varphi\rangle$, is given by:

- $\langle\perp\rangle := \langle p\rangle := 1$, for every $p \in \mathbf{Prop}$;
- $\langle\varphi \star \psi\rangle := \langle\varphi\rangle + \langle\psi\rangle + 1$, for $\star \in \{\wedge, \vee, \rightarrow, \mathbf{U}, \mathbf{R}\}$;
- $\langle\mathbf{X}\varphi\rangle := \langle\varphi\rangle + 1$.

We extend $\langle\cdot\rangle$ to finite sets of formulas by setting $\langle\Phi\rangle := \sum\{\langle\varphi\rangle \mid \varphi \in \Phi\}$.

Formulas of iLTL are interpreted over intuitionistic propositional models equipped with a temporal successor function that respects the intuitionistic order.

An *intuitionistic (propositional) frame* is a Kripke frame $\mathcal{F} = (W, \leq)$ whose accessibility relation \leq is reflexive and transitive. We call \leq the *intuitionistic (pre-)order* of \mathcal{F} . For all $s, t \in W$, we say that t is an *intuitionistic successor* of s if $s \leq t$. An *intuitionistic (propositional) model* is a Kripke model $\mathcal{M} = (W, \leq, V)$ based on an intuitionistic propositional frame (W, R) and such that the valuation function $V : W \rightarrow 2^{\text{Prop}}$ is *persistent* with respect to the intuitionistic order: if $s \leq t$, then $V(s) \subseteq V(t)$.

An *iLTL frame* is a triple $\mathcal{F} = (W, \leq, S)$, where (W, \leq) is an intuitionistic frame and $S : W \rightarrow W$ is a function which is (*forward*) *confluent* with respect to \leq , that is: if $s \leq t$, then $S(s) \leq S(t)$. We call S is the (*temporal*) *successor* function. We inductively define the *n-th (temporal) successor function* $S^n : W \rightarrow W$ for every $n < \omega$ by setting $S^0 := S$ and $S^{n+1} := S \circ S^n$.

A *valuation* V on an iLTL frame $\mathcal{F} = (W, \leq, S)$ is a map $V : W \rightarrow 2^{\text{Prop}}$. An *iLTL model* is a tuple $\mathcal{M} = (W, \leq, S, V)$, where $\mathcal{F} = (W, \leq, S)$ is an iLTL frame and V is a valuation on \mathcal{F} .

Given a model $\mathcal{M} = (W, \leq, S, V)$, we inductively define a *satisfaction* or *forcing* relation \Vdash between states of \mathcal{M} and formulas in the expected manner:

- $\mathcal{M}, s \not\Vdash \perp$;
- $\mathcal{M}, s \Vdash p$ if, and only if, $p \in V(s)$, for every $p \in \text{Prop}$;
- $\mathcal{M}, s \Vdash \varphi \wedge \psi$ if, and only if, $\mathcal{M}, s \Vdash \varphi$ and $\mathcal{M}, s \Vdash \psi$;
- $\mathcal{M}, s \Vdash \varphi \vee \psi$ if, and only if, $\mathcal{M}, s \Vdash \varphi$ or $\mathcal{M}, s \Vdash \psi$;
- $\mathcal{M}, s \Vdash \varphi \rightarrow \psi$ if, and only if, for every $t \in W$ such that $s \leq t$, if $\mathcal{M}, t \Vdash \varphi$, then $\mathcal{M}, t \Vdash \psi$;
- $\mathcal{M}, s \Vdash X\varphi$ if, and only if, $\mathcal{M}, S(s) \Vdash \varphi$;
- $\mathcal{M}, s \Vdash \varphi U \psi$ if, and only if, there is a $j < \omega$ such that $\mathcal{M}, S^j(s) \Vdash \psi$ and $\mathcal{M}, S^i(s) \Vdash \varphi$ for every $i < j$;
- $\mathcal{M}, s \Vdash \varphi R \psi$ if, and only if, for every $j < \omega$, either $\mathcal{M}, S^j(s) \Vdash \psi$, or there is an $i < j$ such that $\mathcal{M}, S^i(s) \Vdash \varphi$.

If $\mathcal{M}, s \Vdash \varphi$, we say that φ is *true* at s , or that s *satisfies* φ . We write $\mathcal{M} \models \varphi$ if $\mathcal{M}, s \Vdash \varphi$ for every $s \in W$. Given a frame $\mathcal{F} = (W, \leq, S)$, we write $\mathcal{F} \models \varphi$, and say that φ is *valid* in \mathcal{F} , if $(W, \leq, S, V) \models \varphi$ for every valuation V on \mathcal{F} . Given a class of frames \mathcal{K} , we write $\mathcal{K} \models \varphi$, and say that φ is *\mathcal{K} -valid*, if $\mathcal{F} \models \varphi$ for every $\mathcal{F} \in \mathcal{K}$. We write $\text{iLTL} \models \varphi$, and say that φ is *valid*, if φ is valid in every frame. Finally, given any formulas φ and ψ , we write $\varphi \equiv \psi$, and say that φ and ψ are *equivalent*, if for every model $\mathcal{M} = (W, \leq, S, V)$ and every $s \in W$, we have $\mathcal{M}, s \Vdash \varphi$ if, and only if, $\mathcal{M}, s \Vdash \psi$.

It follows immediately from the definitions of F and G that we have:



FIGURE 4.1: An iLTL model \mathcal{M} refuting the classical dualities between U and R. The valuation is empty everywhere except at the states painted black, where q and only q holds. We then have $\mathcal{M}, s \Vdash pRq$ but $\mathcal{M}, s \not\Vdash \neg(\neg pU\neg q)$, and $\mathcal{M}, s \Vdash pUq$ but $\mathcal{M}, s \not\Vdash \neg(\neg pR\neg q)$. Analogously, we have $\mathcal{M}, t \Vdash \neg(\neg pU\neg q)$ but $\mathcal{M}, t \not\Vdash pRq$, and $\mathcal{M}, t \Vdash \neg(\neg pR\neg q)$ but $\mathcal{M}, t \not\Vdash pUq$.

- $\mathcal{M}, s \Vdash F\varphi$ if, and only if, there is some $n < \omega$ such that $\mathcal{M}, S^n(s) \Vdash \varphi$;
- $\mathcal{M}, s \Vdash G\varphi$ if, and only if, $\mathcal{M}, S^n(s) \Vdash \varphi$ for every $n < \omega$.

As in the classical setting, it is easy to see that U and R satisfy the fixpoint equivalences $\varphi U\psi \equiv \psi \vee [\varphi \wedge X(\varphi U\psi)]$ and $\varphi R\psi \equiv \psi \wedge [\varphi \vee X(\varphi R\psi)]$. However, in the intuitionistic setting U and R are not interdefinable [10] (see also Example 4.1.1 below).

Observe that iLTL frames are two-dimensional: every state has one temporal successor and, in addition, may have multiple intuitionistic successors other than itself. When depicting frames and models, we represent the temporal successor function with dashed, horizontal arrows, and the intuitionistic order with solid, vertical arrows, omitting reflexivity.

4.1.1. EXAMPLE. Figure 4.1 depicts an iLTL model showing that, in contrast to the classical case, we can neither express pRq as $\neg(\neg pU\neg q)$ nor pUq as $\neg(\neg pR\neg q)$. Moreover, in iLTL the operator U is not definable in terms of X and R, and neither is R definable in terms of X and U [10].

Intuitionistic truth should be propagated upwards in a frame with respect to the intuitionistic order. Forward confluence is exactly the condition that has to be imposed onto the temporal successor function in order for this to be the case [10, Prop. 2.1]. So, in particular, we have:

4.1.2. PROPOSITION. *For every model $\mathcal{M} = (W, \leq, S, V)$, all $s, t \in W$, and every formula φ , if $\mathcal{M}, s \Vdash \varphi$ and $s \leq t$, then $\mathcal{M}, t \Vdash \varphi$.*

4.2 The ill-founded system $iLTL_\infty$

Fix a countably infinite set **Worlds** of (*world*) labels. A labelled (iLTL) formula is a pair (w, φ) , henceforth written $w : \varphi$, where $w \in \mathbf{Worlds}$ and $\varphi \in \mathbf{Form}_{iLTL}$.

Let Γ be a set of labelled formulas. We define $X\Gamma := \{w : X\varphi \mid w : \varphi \in \Gamma\}$ and $\Gamma^- := \{\varphi \mid w : \varphi \in \Gamma \text{ for some } w \in \text{Worlds}\}$. And, for every set of labels W , we abuse notation and let $\Gamma \setminus W := \{w : \varphi \in \Gamma \mid w \notin W\}$. We say that a label w *occurs* in Γ if there is a formula φ such that $w : \varphi \in \Gamma$. When no ambiguity arises, we abuse notation and write $w \in \Gamma$ if w occurs in Γ . If Γ is a set of unlabelled formulas and $w \in \text{Worlds}$, we let $w : \Gamma := \{w : \varphi \mid \varphi \in \Gamma\}$.

A *world relation* is a pair (w, w') , henceforth written $w \preceq w'$, where $w, w' \in \text{Worlds}$. Given a world relation $w_0 \preceq w_1$ and labels w, w' , we define the *world substitution* $(w_0 \preceq w_1)[w/w']$ as follows:

- if $w_0 \neq w' \neq w_1$, then $(w_0 \preceq w_1)[w/w'] := w_0 \preceq w_1$;
- if $w_0 = w' \neq w_1$, then $(w_0 \preceq w_1)[w/w'] := w \preceq w_1$;
- if $w_0 \neq w' = w_1$, then $(w_0 \preceq w_1)[w/w'] := w_0 \preceq w$;
- if $w_0 = w' = w_1$, then $(w_0 \preceq w_1)[w/w'] := w \preceq w$.

A (*relational*) *control* is a finite set of world relations. When working with relational controls, we abbreviate $\{w \preceq w'\}$ to $w \preceq w'$, as well as $\Omega \cup \{w \preceq w'\}$ to $\Omega, w \preceq w'$, and $w \preceq w_0, \dots, w \preceq w_n$ to $w \preceq w_0, \dots, w_n$. A label w *occurs* in a control Ω if there is a label w' such that $\{w \preceq w', w' \preceq w\} \cap \Omega \neq \emptyset$. We abuse notation and write $w \in \Omega$ if w occurs in Ω . For any control Ω and labels w, w' , we let $\Omega[w/w'] := \{(w_0 \preceq w_1)[w/w'] \mid w_0 \preceq w_1 \in \Omega\}$.

Let $\mathcal{M} = (W, \leq, S, V)$ be an $i\text{LTL}$ model. A *label valuation* on \mathcal{M} is a map $\lambda : \text{Worlds} \rightarrow W$. We extend the notion of *satisfiability* to world relations, relational controls and labelled formulas as expected:

- (i) $\mathcal{M}, \lambda \Vdash w \preceq w'$ if, and only if, $\lambda(w) \leq \lambda(w')$;
- (ii) $\mathcal{M}, \lambda \Vdash \Omega$ if, and only if, $\mathcal{M}, \lambda \Vdash w \preceq w'$ for every $w \preceq w' \in \Omega$;
- (iii) $\mathcal{M}, \lambda \Vdash w : \varphi$ if, and only if, $\mathcal{M}, \lambda(w) \Vdash \varphi$.

Note that relational controls are interpreted conjunctively.

Given a model $\mathcal{M} = (W, \leq, S, V)$ and a label valuation λ on \mathcal{M} , for every $w \in \text{Worlds}$ and $s \in W$ we define the label valuation $\lambda[w \mapsto s]$ on \mathcal{M} by setting

$$\lambda[w \mapsto s](w') := \begin{cases} s & \text{if } w' = w \\ \lambda(w') & \text{otherwise} \end{cases}$$

for every $w' \in \text{Worlds}$.

Every relational control Ω determines a reflexive pre-order (W_Ω, \leq_Ω) , where W_Ω is the set of world labels occurring in Ω and \leq_Ω is the reflexive, transitive closure of the relation described by \preceq , i.e., $w \leq_\Omega w'$ holds if, and only if, there

are $w_0, \dots, w_n \in W_\Omega$ such that $w_0 = w$, $w_n = w'$, and, for every $i < n$, either $w_i = w_{i+1}$ or $w_i \preceq w_{i+1} \in \Omega$. By $w <_\Omega w'$ we mean $w \leq_\Omega w'$ and $w \neq w'$. And we write $w <_\Omega^0 w'$ if $w <_\Omega w'$ and there is no $w'' \in W_\Omega$ such that $w <_\Omega w''$ and $w'' <_\Omega w'$.

Let Ω be a relational control. For any label w and any set of labels W , we let:

- $\uparrow_{\leq_\Omega} w := \{w' \in W_\Omega \mid w \leq_\Omega w'\}$;
- $\Omega \setminus W := \{w_0 \preceq w_1 \in \Omega \mid w_0 \notin W \text{ and } w_1 \notin W\}$.

We abuse notation and write $\Omega \setminus w$ in place of $\Omega \setminus \{w\}$.

Relational controls that represent trees will be of particular importance to us. To this end, we say that a control Ω is:

- *reflexive*, if $w \preceq w \in \Omega$ for every $w \in \Omega$;
- *transitive*, if $w \preceq w' \in \Omega$ and $w' \preceq w'' \in \Omega$ implies $w \preceq w'' \in \Omega$;
- *rooted*, if there is a $w \in \Omega$ such that $w \preceq w' \in \Omega$ for every $w' \in \Omega$;
- and *tree-like*, if $(W_\Omega, <_\Omega)$ is a tree.

4.2.1. REMARK. Relational controls are finite by definition, so finitely many checks suffice to determine whether a control is tree-like or not.

A (*labelled iLTL*) *sequent* is a triple (Ω, Γ, Δ) , henceforth written $\Omega \dashv \Gamma \Rightarrow \Delta$, where Ω is a relational control and Γ and Δ are finite sets of labelled formulas. We call Γ the *antecedent* of the sequent, and Δ the *consequent* or *succedent*. The *size*, or *complexity*, of a sequent $S = \Omega \dashv \Gamma \Rightarrow \Delta$ is $\langle S \rangle := \langle \Gamma \rangle + \langle \Delta \rangle$.

A world label w *occurs* in a sequent $\Omega \dashv \Gamma \Rightarrow \Delta$ if w occurs in Ω or in $\Gamma \cup \Delta$.

Let \mathcal{M} be a model and λ a label valuation on \mathcal{M} . We extend *satisfiability* to labelled sequents by setting $\mathcal{M}, \lambda \Vdash \Omega \dashv \Gamma \Rightarrow \Delta$ if, and only if, the following holds: if $\mathcal{M}, \lambda \Vdash \Omega$ and $\mathcal{M}, \lambda \Vdash w : \varphi$ for every $w : \varphi \in \Gamma$, then $\mathcal{M}, \lambda \Vdash w' : \psi$ for some $w' : \psi \in \Delta$. A labelled sequent S is *valid* if we have $\mathcal{M}, \lambda \Vdash S$ for every model \mathcal{M} and every label valuation λ on \mathcal{M} .

A sequent $S = \Omega \dashv \Gamma \Rightarrow \Delta$ is *monotone* if the following conditions hold for all $w, w' \in \mathbf{Worlds}$ and $\varphi \in \mathbf{Form}_{\text{iLTL}}$:

- (i) if $w : \varphi \in \Gamma$ and $w \preceq w' \in \Omega$, then $w' : \varphi \in \Gamma$;
- (ii) if $w : \varphi \in \Delta$ and $w' \preceq w \in \Omega$, then $w' : \varphi \in \Delta$.

We say that S is *left-monotone* (*right-monotone*) if (i) holds (respectively, (ii)). So S is monotone if, and only if, S is both left- and right-monotone.

We are now ready to define the ill-founded system iLTL_∞ for iLTL. The system works on labelled sequents and allows infinite branches to deal with the fixpoint operators \mathbf{U} and \mathbf{R} .

$$\begin{array}{c}
\text{ax} \frac{}{\Omega \dashv \Gamma, w : p \Rightarrow w : p, \Delta} \quad \text{ax}_\perp \frac{}{\Omega \dashv \Gamma, w : \perp \Rightarrow \Delta} \\
\\
\text{L}\wedge \frac{\Omega \dashv \Gamma, w : \varphi, w : \psi \Rightarrow \Delta}{\Omega \dashv \Gamma, w : \varphi \wedge \psi \Rightarrow \Delta} \quad \text{R}\wedge \frac{\Omega \dashv \Gamma \Rightarrow w : \varphi, \Delta \quad \Omega \dashv \Gamma \Rightarrow w : \psi, \Delta}{\Omega \dashv \Gamma \Rightarrow w : \varphi \wedge \psi, \Delta} \\
\\
\text{L}\vee \frac{\Omega \dashv \Gamma, w : \varphi \Rightarrow \Delta \quad \Omega \dashv \Gamma, w : \psi \Rightarrow \Delta}{\Omega \dashv \Gamma, w : \varphi \vee \psi \Rightarrow \Delta} \quad \text{R}\vee \frac{\Omega \dashv \Gamma \Rightarrow w : \varphi, w : \psi, \Delta}{\Omega \dashv \Gamma \Rightarrow w : \varphi \vee \psi, \Delta} \\
\\
\text{L}\rightarrow \frac{\Omega \dashv \Gamma, w : \varphi \rightarrow \psi \Rightarrow w : \varphi, \Delta \quad \Omega \dashv \Gamma, w : \psi \Rightarrow \Delta}{\Omega \dashv \Gamma, w : \varphi \rightarrow \psi \Rightarrow \Delta} \\
\\
\text{R}\rightarrow \frac{\Omega, w \preceq w' \dashv \Gamma, w' : \varphi \Rightarrow w' : \psi, \Delta}{\Omega \dashv \Gamma \Rightarrow w : \varphi \rightarrow \psi, \Delta} \quad w' \notin \Omega \\
\\
\text{X} \frac{\Omega \dashv \Lambda \Rightarrow \Xi}{\Omega \dashv \Pi, \text{X}\Lambda \Rightarrow \text{X}\Xi, \Sigma}
\end{array}$$

FIGURE 4.2: Non-fixpoint, logical rules of the system $iLTL_\infty$.

4.2.2. DEFINITION ($iLTL_\infty$ derivation). An $iLTL_\infty$ *derivation* of a labelled sequent S is a finite or infinite tree \mathcal{T} built in accordance with the rules in Figures 4.2 to 4.4 and whose root has label S .

An $iLTL_\infty$ *derivation* of a formula φ is an $iLTL$ derivation of $\emptyset \dashv \emptyset \Rightarrow w : \varphi$ for some $w \in \text{Worlds}$. –

4.2.3. REMARK. The preservation of the principal formula $w : \varphi \rightarrow \psi$ in the left premise of rule $\text{L}\rightarrow$ is typical of ordinary sequent calculi for intuitionistic logic (see, e.g., [101, § 2.2]). But in contrast to those calculi, the use of labelled formulas allows the preservation of the set Δ in the left premise of rule $\text{L}\rightarrow$ and also in the premise of $\text{R}\rightarrow$. It thus follows that both $\text{L}\rightarrow$ and $\text{R}\rightarrow$ are semantically invertible.¹

Preserving the side formulas in the succedent when applying rule $\text{R}\rightarrow$ turns ordinary sequent calculi for intuitionistic logic into classical ones, for it allows to prove instances of the law of the excluded middle (see the introduction to [101, Ch. 3]). This is not the case when working with labelled formulas, however. For example, every attempt at proving $p \vee \neg p$ in $iLTL_\infty$ will look essentially like the following:

¹A rule is *semantically invertible* if its premises are valid when its conclusion is.

$$\begin{array}{c}
\text{LU} \frac{\Omega \dashv \Gamma, w : \psi \Rightarrow \Delta \quad \Omega \dashv \Gamma, w : \varphi, w : \mathbf{X}(\varphi \mathbf{U} \psi) \Rightarrow \Delta}{\Omega \dashv \Gamma, w : \varphi \mathbf{U} \psi \Rightarrow \Delta} \\
\text{RU} \frac{\Omega \dashv \Gamma \Rightarrow w : \varphi, w : \psi, \Delta \quad \Omega \dashv \Gamma \Rightarrow w : \psi, w : \mathbf{X}(\varphi \mathbf{U} \psi), \Delta}{\Omega \dashv \Gamma \Rightarrow w : \varphi \mathbf{U} \psi, \Delta} \\
\text{LR} \frac{\Omega \dashv \Gamma, w : \varphi, w : \psi \Rightarrow \Delta \quad \Omega \dashv \Gamma, w : \psi, w : \mathbf{X}(\varphi \mathbf{R} \psi) \Rightarrow \Delta}{\Omega \dashv \Gamma, w : \varphi \mathbf{R} \psi \Rightarrow \Delta} \\
\text{RR} \frac{\Omega \dashv \Gamma \Rightarrow w : \psi, \Delta \quad \Omega \dashv \Gamma \Rightarrow w : \varphi, w : \mathbf{X}(\varphi \mathbf{R} \psi), \Delta}{\Omega \dashv \Gamma \Rightarrow w : \varphi \mathbf{R} \psi, \Delta}
\end{array}$$

FIGURE 4.3: Fixpoint rules of the system iLTL_∞ .

$$\begin{array}{c}
\text{R}\rightarrow \frac{w \preceq w' \dashv w' : p \Rightarrow w : p, w' : \perp}{\emptyset \dashv \emptyset \Rightarrow w : p, w : p \rightarrow \perp} \\
\text{R}\vee \frac{\quad}{\emptyset \dashv \emptyset \Rightarrow w : p \vee (p \rightarrow \perp)}
\end{array}$$

Clearly, the sequent at the top is not provable.

Invertibility of rule $\text{R}\rightarrow$ appears to be necessary to obtaining regular ill-founded proofs, and thus in the end a cyclic calculus. If rule $\text{R}\rightarrow$ discarded the side formulas in Δ , then the proof-search procedure described below in Section 4.2.2 would not be able to apply rule $\text{R}\rightarrow$ without losing information that might be necessary to find a proof. Observe that rule $\text{R}\rightarrow$ is the only one that, when read bottom-up, creates intuitionistic successors. So, informally, if it were not invertible, the proof-search procedure would not be guaranteed to explore the intuitionistic regions above all worlds, because it could wait in vain for the right moment to break down implications in the succedent of a sequent. This problem seems to lie at the heart of the difficulties with turning the ill-founded calculus of [5], whose right-implication rule is not invertible, into a cyclic one (see, in particular, [5, § 5.2]). That our proof-search never runs into this problem is one of the main results of Section 4.2.2 (Proposition 4.2.41).

The distinguished formulas and world relations in the conclusions of the rules in Figures 4.2 to 4.4, such as $w : \varphi \rightarrow \psi$ in $\text{R}\rightarrow$ and $w \preceq w'$ and $w : \varphi$ in Lm , are said to be *principal*. In the special case of the *modal rule* \mathbf{X} , we consider all world relations and formulas in the conclusion to be principal. The distinguished formulas and world relations in the premises of the rules, for instance $w \preceq w'$, $w' : \varphi$ and $w' : \psi$ in $\text{R}\rightarrow$ and $w \preceq w'$, $w : \varphi$ and $w' : \varphi$ in Lm , are said to be *active*.

$$\begin{array}{c}
\text{refl} \frac{\Omega, w \preccurlyeq w \dashv\vdash \Gamma \Rightarrow \Delta}{\Omega \dashv\vdash \Gamma \Rightarrow \Delta} \quad \text{trans} \frac{\Omega, w \preccurlyeq w', w' \preccurlyeq w'', w \preccurlyeq w'' \dashv\vdash \Gamma \Rightarrow \Delta}{\Omega, w \preccurlyeq w', w' \preccurlyeq w'' \dashv\vdash \Gamma \Rightarrow \Delta} \\
\\
\text{Lm} \frac{\Omega, w \preccurlyeq w' \dashv\vdash \Gamma, w : \varphi, w' : \varphi \Rightarrow \Delta}{\Omega, w \preccurlyeq w' \dashv\vdash \Gamma, w : \varphi \Rightarrow \Delta} \quad \text{Rm} \frac{\Omega, w' \preccurlyeq w \dashv\vdash \Gamma \Rightarrow w : \varphi, w' : \varphi, \Delta}{\Omega, w' \preccurlyeq w \dashv\vdash \Gamma \Rightarrow w : \varphi, \Delta} \\
\\
\text{wk}_{\preccurlyeq} \frac{\Omega \dashv\vdash \Gamma \Rightarrow \Delta}{\Omega, w \preccurlyeq w' \dashv\vdash \Gamma \Rightarrow \Delta} \\
\\
\text{Lwk} \frac{\Omega \dashv\vdash \Gamma \Rightarrow \Delta}{\Omega \dashv\vdash \Gamma, w : \varphi \Rightarrow \Delta} \quad \text{Rwk} \frac{\Omega \dashv\vdash \Gamma \Rightarrow \Delta}{\Omega \dashv\vdash \Gamma \Rightarrow w : \varphi, \Delta}
\end{array}$$

FIGURE 4.4: Structural rules of the system $i\text{LTL}_\infty$.

For all of the rules other than X , the formulas occurring in Γ and Δ and the world relations occurring in Ω are, respectively, *side* formulas and *side* relations.

Since we work with set-based sequents, a side formula might also be principal. For example, the following is an instance of $\text{L}\wedge$ in which the principal formula $w : \varphi \wedge \psi$ is also a side formula:

$$\text{R}\wedge \frac{\Omega \dashv\vdash \Gamma, w : \varphi \wedge \psi, w : \varphi, w : \psi \Rightarrow \Delta}{\Omega \dashv\vdash \Gamma, w : \varphi \wedge \psi \Rightarrow \Delta}$$

An instance of a rule among $\text{L}\wedge$, $\text{L}\vee$, $\text{L}\rightarrow$, $\text{L}\cup$, LR , Lwk ($\text{R}\wedge$, $\text{R}\vee$, $\text{R}\rightarrow$, $\text{R}\cup$, RR , Rwk) is *preserving* if the unique principal formula is also a side formula contained in Γ (respectively, Δ). Otherwise, the instance is said to be *discarding*. Analogously, an instance of wk_{\preccurlyeq} is *preserving* if the unique principal relation is also a side relation, and *discarding* otherwise. Note that, by definition, instances of rules X , refl , trans , Lm and Rm are neither preserving nor discarding.

A vertex u of an $i\text{LTL}$ derivation is *modal* if rule X is applied at u .

Let \mathcal{T} be an $i\text{LTL}_\infty$ derivation, let $u, v \in T$ be such that $u <_T^0 v$, and let $\Omega \dashv\vdash \Gamma \Rightarrow \Delta$ and $\Omega' \dashv\vdash \Gamma' \Rightarrow \Delta'$ be the labels of u and v , respectively. We define the *formula trace* relation \triangleright_u^v between $(\Gamma \times \{0\}) \cup (\Delta \times \{1\})$ and $(\Gamma' \times \{0\}) \cup (\Delta' \times \{1\})$ by setting $(w : \varphi, b) \triangleright_u^v (w' : \varphi', b')$ if, and only if, one of the following holds:

- $w : \varphi$ is a side formula at u , and $(w' : \varphi', b') = (w : \varphi, b)$;
- a rule among $\text{L}\wedge$, $\text{L}\vee$, $\text{L}\cup$, LR , Lm ($\text{R}\wedge$, $\text{R}\vee$, $\text{R}\cup$, RR , Rm) is applied at u with $w : \varphi$ principal and $w' : \varphi'$ active, and $b = b' = 0$ (respectively, $b = b' = 1$);
- rule X is applied at u with $w' : \varphi'$ active, and $(w : \varphi, b) = (w' : \text{X}\varphi', b')$;

- rule $L\rightarrow$ is applied at u with $w : \varphi$ principal, say $\varphi = \alpha \rightarrow \beta$, and $b = 0$, and $(w' : \varphi', b') \in \{(w : \varphi, 0), (w : \alpha, 1), (w : \beta, 0)\}$;
- rule $R\rightarrow$ is applied at u with $w : \varphi$ principal and $w' : \varphi'$ active, say $\varphi = \alpha \rightarrow \beta$, and $b = 1$, and $(w' : \varphi', b') \in \{(w' : \alpha, 0), (w' : \beta, 1)\}$.

So, informally, $(w : \varphi, b) \triangleright_u^v (w' : \psi, b')$ holds if $w' : \psi$ arises from $w : \varphi$ in the rule applied at u . We may drop one or both indices from \triangleright_u^v if no ambiguity arises, as well as write $w : \varphi \triangleright_u^v w' : \psi$ in place of $(w : \varphi, b) \triangleright_u^v (w' : \psi, b')$ when b and b' are clear from context.

A (*formula*) *trace* on a finite or infinite path $\pi = (u_n)_{n < N}$ through \mathcal{T} , $N \leq \omega$, is a sequence $((w_n : \varphi_n, b_n))_{n < N}$ such that $w_n : \varphi_n$ occurs in the label of u_n and $(w_n : \varphi_n, b_n) \triangleright_{u_n}^{u_{n+1}} (w_{n+1} : \varphi_{n+1}, b_{n+1})$ for every $n < n+1 < N$.

Given vertices $u, v \in T$ such that $u \leq_T v$, say with $[u, v]_{\mathcal{T}} = (u_0, \dots, u_n)$, we write $w : \alpha \triangleright_u^v w' : \alpha'$ if there is a formula trace $(w_i : \alpha_i)_{i \leq n}$ on the path $[u, v]_{\mathcal{T}}$ such that $w_0 : \alpha_0 = w : \alpha$ and $w_n : \alpha_n = w' : \alpha'$.

Let τ be a finite trace, and τ' any trace. We say that τ' is *coherent* with τ if there is a path $\pi = \pi_0 \hat{\ } \pi_1$ through an iLTL derivation, with π_0 finite, such that τ is a trace on π_0 and τ' is a trace on π_1 . So τ' is coherent with τ if $\tau \hat{\ } \tau'$ is a trace.

A (*fixpoint*) *unfolding* is a trace of the form $(w : \varphi, b) \triangleright (w : X\varphi, b)$. Note that unfoldings can only be produced by rules LU, RU, LR and RR, and thus φ is of the form $\varphi_1 O \varphi_2$. Due to the presence of fixpoint unfoldings, the system iLTL_{∞} does not satisfy the *subformula property*: $(w : \varphi, b) \triangleright (w' : \psi, b')$ does not imply $\psi \in \text{Sub}(\varphi)$. However, ψ does belong to the *closure* of φ , which is the natural replacement of the notion of subformula in this context:

4.2.4. DEFINITION (iLTL closure). The (*iLTL*) *closure* of an iLTL formula φ is the smallest set of formulas $\text{Clos}(\varphi)$ satisfying:

- $\varphi \in \text{Clos}(\varphi)$;
- if $\psi_1 \star \psi_2 \in \text{Clos}(\varphi)$, for $\star \in \{\wedge, \vee, \rightarrow\}$, then $\psi_1, \psi_2 \in \text{Clos}(\varphi)$;
- if $\psi_1 O \psi_2 \in \text{Clos}(\varphi)$, for $O \in \{U, R\}$, then $\psi_1, \psi_2, X(\psi_1 O \psi_2) \in \text{Clos}(\varphi)$;
- if $X\psi \in \text{Clos}(\varphi)$, then $\psi \in \text{Clos}(\varphi)$.

⊥

By a straightforward structural induction, the closure of a formula can be characterised as follows:

4.2.5. LEMMA. *The following hold:*

- $\text{Clos}(\perp) = \{\perp\}$;
- $\text{Clos}(p) = \{p\}$, for every $p \in \text{Prop}$;

- (iii) $\text{Clos}(\psi_1 \star \psi_2) = \{\psi_1 \star \psi_2\} \cup \text{Clos}(\psi_1) \cup \text{Clos}(\psi_2)$, for $\star \in \{\wedge, \vee, \rightarrow\}$;
- (iv) $\text{Clos}(\mathbf{X}\psi) = \{\mathbf{X}\psi\} \cup \text{Clos}(\psi)$;
- (v) $\text{Clos}(\psi_1 \mathbf{O}\psi_2) = \{\psi_1 \mathbf{O}\psi_2, \mathbf{X}(\psi_1 \mathbf{O}\psi_2)\} \cup \text{Clos}(\psi_1) \cup \text{Clos}(\psi_2)$, for $\mathbf{O} \in \{\mathbf{U}, \mathbf{R}\}$.

4.2.6. COROLLARY. *For any $i\text{LTL}$ formula φ , the set $\text{Clos}(\varphi)$ is finite.*

For every formula φ , we let $|\varphi| := |\text{Clos}(\varphi)|$.

It is clear that $(w : \varphi, b) \triangleright (w' : \psi, b')$ implies $\psi \in \text{Clos}(\varphi)$. Moreover, since $(w : \mathbf{X}\varphi, b) \triangleright (w' : \psi, b')$ implies $\psi \in \{\mathbf{X}\varphi, \varphi\}$, arguing as in the proof of Lemma 3.2.8 yields:

4.2.7. PROPOSITION. *Let $\tau = ((w_n : \varphi_n, b_n))_{n < N \leq \omega}$ be a finite or infinite trace. For every $n < N$, either $\varphi_n \in \text{Sub}(\varphi_0)$, or $\varphi_n = \mathbf{X}\psi$ for some $\psi \in \text{Sub}(\varphi_0)$.*

We shall mainly concern ourselves with infinite traces. Let $\pi = (u_n)_{n < \omega}$ be an infinite path through an $i\text{LTL}_\infty$ derivation. For every $n < \omega$, let $\Omega_n \dashv \Gamma_n \Rightarrow \Delta_n$ be the label of u_n . Finally, let $\tau = ((w_n : \varphi_n, b_n))_{n < \omega}$ be a trace on π . We say that τ is *label-stable* (*side-stable*) if the sequence $(w_n)_{n < \omega}$ (respectively, $(b_n)_{n < \omega}$) is eventually constant. We say that τ is a *left* (*right*) trace if it is side-stable and $b_n = 0$ (respectively, $b_n = 1$) for infinitely many $n < \omega$. Finally, we say that τ is of *type* \mathbf{O} , or that τ is an *\mathbf{O} -trace*, for $\mathbf{O} \in \{\mathbf{U}, \mathbf{R}\}$, if there is an $N < \omega$ and an \mathbf{O} -formula α such that $\varphi_n \in \{\alpha, \mathbf{X}\alpha\}$ for every $n \geq N$ and the sequence $(\varphi_n)_{n < \omega}$ is not eventually constant. We call α the *dominating* formula in τ . Analogously, the *dominating* world label in a label-stable trace $\tau = ((w_n : \varphi_n, b_n))_{n < \omega}$ is the unique $w \in \text{Worlds}$ such that $w_n = w$ for infinitely many $n < \omega$.

4.2.8. PROPOSITION. *Any trace on an infinite path through an $i\text{LTL}_\infty$ derivation is side-stable.*

Proof. Let $\pi = (u_n)_{n < \omega}$ be an infinite path through an $i\text{LTL}_\infty$ derivation, and $\tau = ((w_n : \varphi_n, b_n))_{n < \omega}$ a trace on π . Let $n < \omega$ be such that $b_{n+1} \neq b_n$. By the definition of \triangleright , we then have $\langle \varphi_{n+1} \rangle \leq \langle \varphi_n \rangle - 2$. Moreover, by Proposition 4.2.7 for every $k > 0$ either $\varphi_{n+k} \in \text{Sub}(\varphi_{n+1})$, or $\varphi_{n+k} = \mathbf{X}\psi$ for some $\psi \in \text{Sub}(\varphi_{n+1})$. In the first case we have $\langle \varphi_{n+k} \rangle \leq \langle \varphi_{n+1} \rangle < \langle \varphi_n \rangle$; in the second, $\langle \varphi_{n+k} \rangle = 1 + \langle \psi \rangle \leq 1 + \langle \varphi_{n+1} \rangle < \langle \varphi_n \rangle$. Therefore, $b_{n+1} \neq b_n$ implies $\langle \varphi_{n+k} \rangle < \langle \varphi_n \rangle$ for every $k > 0$, so the sequence $(b_n)_{n < \omega}$ must be eventually constant. \blacksquare

The analogue of Proposition 4.2.8 for world labels, however, does not hold:

4.2.9. EXAMPLE. Figure 4.5 depicts an infinite branch of an $i\text{LTL}_\infty$ derivation containing an infinite formula trace which is not label-stable.

$$\begin{array}{c}
\vdots \\
\text{Lm} \frac{w_3 \preceq w_2 \dashv w_2 : (p \rightarrow q) \rightarrow r, w_3 : p, \mathbf{w}_3 : (\mathbf{p} \rightarrow \mathbf{q}) \rightarrow \mathbf{r} \Rightarrow w_3 : q}{\text{R} \rightarrow \frac{w_3 \preceq w_2 \dashv \mathbf{w}_2 : (\mathbf{p} \rightarrow \mathbf{q}) \rightarrow \mathbf{r}, w_3 : p \Rightarrow w_3 : q}{\text{L} \rightarrow \frac{\dashv \mathbf{w}_2 : (\mathbf{p} \rightarrow \mathbf{q}) \rightarrow \mathbf{r} \Rightarrow w_2 : p \rightarrow q}{\dashv \mathbf{w}_2 : (\mathbf{p} \rightarrow \mathbf{q}) \rightarrow \mathbf{r} \Rightarrow \emptyset}}} \\
\text{Lm} \frac{w_1 \preceq w_2 \dashv w_1 : (p \rightarrow q) \rightarrow r, w_2 : p, \mathbf{w}_2 : (\mathbf{p} \rightarrow \mathbf{q}) \rightarrow \mathbf{r} \Rightarrow w_2 : q}{\text{R} \rightarrow \frac{w_1 \preceq w_2 \dashv \mathbf{w}_1 : (\mathbf{p} \rightarrow \mathbf{q}) \rightarrow \mathbf{r}, w_2 : p \Rightarrow w_2 : q}{\text{wk}_{\preceq} \frac{\dashv \mathbf{w}_1 : (\mathbf{p} \rightarrow \mathbf{q}) \rightarrow \mathbf{r} \Rightarrow w_1 : p \rightarrow q}{\text{R} \rightarrow \frac{w_0 \preceq w_1 \dashv \mathbf{w}_1 : (\mathbf{p} \rightarrow \mathbf{q}) \rightarrow \mathbf{r} \Rightarrow w_1 : p \rightarrow q}{\dashv \emptyset \Rightarrow \mathbf{w}_0 : [(\mathbf{p} \rightarrow \mathbf{q}) \rightarrow \mathbf{r}] \rightarrow (\mathbf{p} \rightarrow \mathbf{q})}}} \\
\end{array}$$

FIGURE 4.5: An infinite branch of an iLTL_{∞} derivation containing an infinite trace (highlighted) which is not label-stable. The double line indicates the omission of some instances of weakening.

We have now all the ingredients necessary to characterise the iLTL_{∞} derivations that should be accepted as proofs.

4.2.10. DEFINITION (Good and bad trace). An infinite formula trace τ is *good* if τ is label-stable and either a left U-trace, or a right R-trace. Otherwise, τ is said to be *bad*. \dashv

4.2.11. OBSERVATION. Every good infinite formula trace passes through infinitely many modal vertices.

4.2.12. DEFINITION (iLTL_{∞} proof). An *iLTL_{∞} proof* of a labelled sequent S is an iLTL_{∞} derivation \mathcal{T} of S whose leaves are all axiomatic and such that on every infinite branch of \mathcal{T} there is a good trace.

An *iLTL_{∞} proof* of a formula φ is an iLTL_{∞} proof of $\emptyset \dashv \emptyset \Rightarrow w : \varphi$ for some $w \in \text{Worlds}$. \dashv

4.2.13. OBSERVATION. It follows from the definition of O-traces that every infinite branch of an iLTL_{∞} proof contains infinitely many modal vertices.

4.2.14. OBSERVATION. Let \mathcal{T} be an iLTL_{∞} derivation, and π an infinite branch of \mathcal{T} . Then, for every tail π' of π and every infinite trace τ' on π' , there is a trace τ on π such that τ' is a tail of τ . Therefore, requiring that there be a good trace on π is equivalent to requiring that there be a tail of π' containing a good trace.

4.2.15. EXAMPLE. Figure 4.6 depicts an infinite branch of an iLTL_{∞} proof: the branch contains a good left trace with dominating formula $p\mathbf{U}q$.

$$\begin{array}{c}
\vdots \\
\text{X} \frac{(\dagger) \quad w_0 \preceq w \dashv w : p\text{U}q \Rightarrow w : q, w : (\text{X}p)\text{U}(\text{X}q)}{w_0 \preceq w \dashv w : p, w : \text{X}(p\text{U}q) \Rightarrow w : q, w : \text{X}q, w : \text{X}((\text{X}p)\text{U}(\text{X}q))} \\
\text{RU} \frac{\text{LU} \frac{w_0 \preceq w \dashv w : p, w : \text{X}(p\text{U}q) \Rightarrow w : q, w : (\text{X}p)\text{U}(\text{X}q)}{(\dagger) \quad w_0 \preceq w \dashv w : p\text{U}q \Rightarrow w : q, w : (\text{X}p)\text{U}(\text{X}q)}}{w_0 \preceq w \dashv w : \text{X}(p\text{U}q) \Rightarrow w : \text{X}q, w : \text{X}((\text{X}p)\text{U}(\text{X}q))} \\
\text{RU} \frac{\text{X} \frac{w_0 \preceq w \dashv w : \text{X}(p\text{U}q) \Rightarrow w : \text{X}q, w : \text{X}((\text{X}p)\text{U}(\text{X}q))}{w_0 \preceq w \dashv w : \text{X}(p\text{U}q) \Rightarrow w : (\text{X}p)\text{U}(\text{X}q)}}{w_0 \preceq w \dashv w : \text{X}(p\text{U}q) \Rightarrow w : (\text{X}p)\text{U}(\text{X}q)} \\
\text{R}\rightarrow \frac{\text{RU} \frac{w_0 \preceq w \dashv w : \text{X}(p\text{U}q) \Rightarrow w : (\text{X}p)\text{U}(\text{X}q)}{\dashv \emptyset \Rightarrow w_0 : \text{X}(p\text{U}q) \rightarrow (\text{X}p)\text{U}(\text{X}q)}}{\dashv \emptyset \Rightarrow w_0 : \text{X}(p\text{U}q) \rightarrow (\text{X}p)\text{U}(\text{X}q)}
\end{array}$$

FIGURE 4.6: An infinite branch of an $i\text{LTL}_\infty$ proof of the formula $\text{X}(p\text{U}q) \rightarrow (\text{X}p)\text{U}(\text{X}q)$. The symbol ‘ \dagger ’ marks roots of identical subproofs.

The following follows immediately from the definition of F and G:

4.2.16. PROPOSITION. *The rules LF, RF, LG and RG, depicted in Figure 4.7, are all admissible in $i\text{LTL}_\infty$.*

4.2.1 Soundness of $i\text{LTL}_\infty$

In this section we establish the soundness of the system $i\text{LTL}_\infty$. We proceed as we did for CTL_∞^* in Section 3.2.1, that is: given a proof \mathcal{T} of a formula φ , by assuming that φ is not valid we find an infinite branch of \mathcal{T} whose existence contradicts the fact that \mathcal{T} is a proof.

For every $n < \omega$ we define the n -th approximation $\varphi\text{U}^n\psi$ of the formula $\varphi\text{U}\psi$ by setting $\varphi\text{U}^0\psi := \psi$ and $\varphi\text{U}^{n+1}\psi := \psi \vee (\varphi \wedge \text{X}(\varphi\text{U}^n\psi))$. Analogously, we define the n -th approximation $\varphi\text{R}^n\psi$ of the formula $\varphi\text{R}\psi$ by setting $\varphi\text{R}^0\psi := \psi$ and $\varphi\text{R}^{n+1}\psi := \psi \wedge (\varphi \vee \text{X}(\varphi\text{R}^n\psi))$.

Arguing as in the proof of Lemma 3.2.16, we get:

4.2.17. LEMMA. *For any model $\mathcal{M} = (W, \leq, S, V)$, any $s \in W$, any formulas φ, ψ and any $n < \omega$, the following hold:*

- (i) $\mathcal{M}, s \Vdash \varphi\text{U}^n\psi$ if, and only if, there is a $j \leq n$ such that $\mathcal{M}, S^j(s) \Vdash \psi$ and $\mathcal{M}, S^i(s) \Vdash \varphi$ for every $i < j$.
- (ii) $\mathcal{M}, s \Vdash \varphi\text{R}^n\psi$ if, and only if, for every $j \leq n$, either $\mathcal{M}, S^j(s) \Vdash \psi$, or there is an $i < j$ such that $\mathcal{M}, S^i(s) \Vdash \varphi$.

4.2.18. COROLLARY. *The following hold:*

- (i) $\mathcal{M}, s \Vdash \varphi\text{U}\psi$ if, and only if, $\mathcal{M}, s \Vdash \varphi\text{U}^n\psi$ for some $n < \omega$.

$$\begin{array}{c}
\text{LF} \frac{\Omega \dashv \Gamma, w : \varphi \Rightarrow \Delta \quad \Omega \dashv \Gamma, w : \text{XF}\varphi \Rightarrow \Delta}{\Omega \dashv \Gamma, w : \text{F}\varphi \Rightarrow \Delta} \quad \text{RF} \frac{\Omega \dashv \Gamma \Rightarrow w : \varphi, w : \text{XF}\varphi, \Delta}{\Omega \dashv \Gamma \Rightarrow w : \text{F}\varphi, \Delta} \\
\text{LG} \frac{\Omega \dashv \Gamma, w : \varphi, w : \text{XG}\varphi \Rightarrow \Delta}{\Omega \dashv \Gamma, w : \text{G}\varphi \Rightarrow \Delta} \quad \text{RG} \frac{\Omega \dashv \Gamma \Rightarrow w : \varphi, \Delta \quad \Omega \dashv \Gamma \Rightarrow w : \text{XG}\varphi, \Delta}{\Omega \dashv \Gamma \Rightarrow w : \text{G}\varphi, \Delta}
\end{array}$$

FIGURE 4.7: Admissible iLTL_∞ rules for F and G.

(ii) $\mathcal{M}, s \Vdash \varphi R\psi$ if, and only if, $\mathcal{M}, s \Vdash \varphi R^n\psi$ for all $n < \omega$.

4.2.19. PROPOSITION (Soundness of iLTL_∞). *For every formula φ , if iLTL_∞ $\vdash \varphi$, then iLTL $\models \varphi$.*

Proof. Let \mathcal{T} be an iLTL_∞ proof of φ , say with conclusion $\emptyset \dashv \emptyset \Rightarrow w_0 : \varphi$. Towards a contradiction, suppose that φ is not valid. Let $\mathcal{M} = (W, \leq, S, V)$ be a model and $\lambda : \text{Worlds} \rightarrow W$ such that $\mathcal{M}, \lambda \not\models w_0 : \varphi$. We inductively build an infinite path $(u_n)_{n < \omega}$ through \mathcal{T} and an infinite sequence $(\lambda_n)_{n < \omega}$ of label valuations on \mathcal{M} , respectively, such that:

- (i) $\mathcal{M}, \lambda_n \not\models \Omega_n \dashv \Gamma_n \Rightarrow \Delta_n$, where $\Omega_n \dashv \Gamma_n \Rightarrow \Delta_n$ is the label of u_n ;
- (ii) if a rule other than X and R \rightarrow is applied at u_n , then $\lambda_{n+1} = \lambda_n$;
- (iii) if rule R \rightarrow is applied at u_n , then $\lambda_{n+1} = \lambda_n[w \mapsto t]$ for some $w \in \text{Worlds}$ and $t \in W$;
- (iv) if rule X is applied at u_n , then $\lambda_{n+1} = S \circ \lambda_n$.

We begin by letting u_0 be the root of \mathcal{T} and $\lambda_0 := \lambda$. For the inductive case, assume that u_n and λ_n have been defined. Note that, by the induction hypothesis, u_n cannot be axiomatic because axioms are valid. We distinguish cases according to the rule R applied at u_n .

Case R $\in \{\text{L}\wedge, \text{RV}, \text{refl}, \text{trans}, \text{wk}_{\leq}, \text{Lwk}, \text{Rwk}\}$. We let u_{n+1} be the unique immediate successor of u_n and $\lambda_{n+1} := \lambda_n$. Clearly, $\mathcal{M}, \lambda_{n+1} \not\models \Omega_{n+1} \dashv \Gamma_{n+1} \Rightarrow \Delta_{n+1}$.

Case R = R \wedge , say with principal formula $w : \alpha_0 \wedge \alpha_1$. By the inductive hypothesis, there is some $i \leq 1$ such that $\mathcal{M}, \lambda_n \not\models w : \alpha_i$. We thus let u_{n+1} be the immediate successor of u_n with active formula $w : \alpha_i$, and $\lambda_{n+1} := \lambda_n$.

Case R = L \vee , say with principal formula $w : \alpha_0 \vee \alpha_1$. By the inductive hypothesis, there is some $i \leq 1$ such that $\mathcal{M}, \lambda_n \Vdash w : \alpha_i$. We thus let u_{n+1} be the immediate successor of u_n with active formula $w : \alpha_i$, and $\lambda_{n+1} := \lambda_n$.

Case R = L \rightarrow , say with principal formula $w : \alpha_0 \rightarrow \alpha_1$. By the inductive hypothesis, $\mathcal{M}, \lambda_n \Vdash w : \alpha_0 \rightarrow \alpha_1$, whence either $\mathcal{M}, \lambda_n \not\models w : \alpha_0$ or $\mathcal{M}, \lambda_n \Vdash$

$w : \alpha_1$. In the former case we let u_{n+1} be the immediate successor of u_n for the left premise; in the latter, the one for the right premise. In either case we let $\lambda_{n+1} := \lambda_n$.

Case $R = R \rightarrow$, say with principal formula $w : \alpha_0 \rightarrow \alpha_1$ and active world relation $w \preceq w'$. By the inductive hypothesis, $\mathcal{M}, \lambda_n \not\models w : \alpha_0 \rightarrow \alpha_1$, whence there is some $t \in W$ such that $\lambda_n(w) \leq t$, $\mathcal{M}, t \Vdash \alpha_0$, and $\mathcal{M}, t \not\models \alpha_1$. We let u_{n+1} be the unique immediate successor of u_n and $\lambda_{n+1} := \lambda_n[w' \mapsto t]$.

Case $R = X$, say with conclusion $\Omega \dashv \Gamma, X\Lambda \Rightarrow X\Xi, \Delta$ and premise $\Omega \dashv \Lambda \Rightarrow \Xi$. We let u_{n+1} be the unique immediate successor of u_n , and $\lambda_{n+1} := S \circ \lambda_n$. We show that $\mathcal{M}, \lambda_{n+1} \not\models \Omega \dashv \Lambda \Rightarrow \Xi$. Let $w \preceq w' \in \Omega$. By the inductive hypothesis, $\lambda_n(w) \leq \lambda_n(w')$, so by confluence $\lambda_{n+1}(w) = S(\lambda_n(w)) \leq S(\lambda_n(w')) = \lambda_{n+1}(w')$. Now let $w : \alpha \in \Lambda$. Then, $w : X\alpha \in X\Lambda$, so by the inductive hypothesis $\mathcal{M}, \lambda_n \Vdash w : X\alpha$ and thus $\mathcal{M}, S(\lambda_n(w)) \Vdash \alpha$, i.e., $\mathcal{M}, \lambda_{n+1} \Vdash w : \alpha$. Finally, let $w : \alpha \in \Xi$. Then, $w : X\alpha \in X\Xi$, so by the inductive hypothesis $\mathcal{M}, \lambda_n \not\models w : X\alpha$ and thus $\mathcal{M}, S(\lambda_n(w)) \not\models \alpha$, i.e., $\mathcal{M}, \lambda_{n+1} \not\models w : \alpha$.

Case $R = LU$, say with principal formula $w : \alpha_0 U \alpha_1$. By the inductive hypothesis, either $\mathcal{M}, \lambda_n \Vdash w : \alpha_1$, or $\mathcal{M}, \lambda_n \Vdash w : \alpha_0$ and $\mathcal{M}, \lambda_n \Vdash w : X(\alpha_0 U \alpha_1)$. In the former case we let u_{n+1} be the immediate successor of u_n for the left premise; in the latter, the one for the right premise. In either case we let $\lambda_{n+1} := \lambda_n$.

Case $R = RU$, say with principal formula $w : \alpha_0 U \alpha_1$. By the inductive hypothesis, either $\mathcal{M}, \lambda_n \not\models w : \alpha_0$ and $\mathcal{M}, \lambda_n \not\models w : \alpha_1$, or $\mathcal{M}, \lambda_n \not\models w : \alpha_1$ and $\mathcal{M}, \lambda_n \not\models w : X(\alpha_0 R \alpha_1)$. In the former case we let u_{n+1} be the immediate successor of u_n for the left premise; in the latter, the one for the right premise. In either case we let $\lambda_{n+1} := \lambda_n$.

Case $R = LR$, say with principal formula $w : \alpha_0 R \alpha_1$. By the inductive hypothesis, $\mathcal{M}, \lambda_n \Vdash w : \alpha_1$ and either $\mathcal{M}, \lambda_n \Vdash w : \alpha_0$, or $\mathcal{M}, \lambda_n \Vdash w : X(\alpha_0 R \alpha_1)$. In the former case we let u_{n+1} be the immediate successor of u_n for the left premise; in the latter, the one for the right premise. In either case we let $\lambda_{n+1} := \lambda_n$.

Case $R = RR$, say with principal formula $w : \alpha_0 R \alpha_1$. By the inductive hypothesis, either $\mathcal{M}, \lambda_n \not\models w : \alpha_1$, or $\mathcal{M}, \lambda_n \not\models w : \alpha_0$ and $\mathcal{M}, \lambda_n \not\models w : X(\alpha_0 R \alpha_1)$. In the former case we let u_{n+1} be the immediate successor of u_n for the left premise; in the latter, the one for the right premise. In either case we let $\lambda_{n+1} := \lambda_n$.

Case $R = Lm$, say with principal formula $w : \alpha$ and principal relation $w \preceq w'$. We let u_{n+1} be the unique immediate successor of u_n and $\lambda_{n+1} := \lambda_n$. By the inductive hypothesis, $\lambda_n(w) \leq \lambda_n(w')$ and $\mathcal{M}, \lambda_n \Vdash w : \alpha$, whence $\mathcal{M}, \lambda_{n+1} \Vdash w' : \alpha$ follows from Proposition 4.1.2.

Case $R = Rm$, say with principal formula $w : \alpha$ and principal relation $w' \preceq w$. We let u_{n+1} be the unique immediate successor of u_n and $\lambda_{n+1} := \lambda_n$. By the inductive hypothesis, $\lambda_n(w') \leq \lambda_n(w)$ and $\mathcal{M}, \lambda_n \not\models w : \alpha$, whence $\mathcal{M}, \lambda_{n+1} \not\models w' : \alpha$ follows from Proposition 4.1.2.

The path $\pi := (u_n)_{n < \omega}$ is an infinite branch of \mathcal{T} because axioms are valid. We show that the existence of such a branch contradicts the fact that \mathcal{T} is a proof. Note that, in the cases $R \in \{\text{LU}, \text{RU}, \text{LR}, \text{RR}\}$ in the construction of π , we prioritised the left, non-unfolding premise.

Since \mathcal{T} is a proof, π contains a good trace $\tau = ((w_n : \varphi_n, b_n))_{n < \omega}$, say with dominant formula $\gamma = \alpha\text{O}\beta$. Let then $N < \omega$ and $w \in \text{Worlds}$ be such that $\varphi_N = \gamma$, $\varphi_{N+1} = \text{X}\gamma$, and for every $n \geq N$ we have $w_n = w$, $b_n = b_{n+1}$ and $\varphi_n \in \{\gamma, \text{X}\gamma\}$. We distinguish cases according to the type of τ .

Case $\text{O} = \text{U}$. Then, τ is a left trace and thus $\mathcal{M}, \lambda_n \Vdash w : \varphi_n$ for every $n \geq N$. We inductively define a sequence of natural numbers $(k_n)_{n \geq N}$ such that:

- (i) if $\varphi_n = \gamma$, then $\mathcal{M}, \lambda_n \Vdash w : \alpha\text{U}^{k_n}\beta$;
- (ii) if $\varphi_n = \text{X}\gamma$, then $k_n > 0$ and $\mathcal{M}, \lambda_n \Vdash w : \text{X}(\alpha\text{U}^{k_n-1}\beta)$.

Let k_N be such that $\mathcal{M}, \lambda_N \Vdash w : \alpha\text{U}^{k_N}\beta$, which exists by Corollary 4.2.18.

If $\varphi_{n+1} = \varphi_n$, then $\lambda_{n+1}(w) = \lambda_n(w)$ and we let $k_{n+1} := k_n$.

Suppose that $\varphi_n = \gamma$ and $\varphi_{n+1} = \text{X}\gamma$. By the construction of π and the inductive hypothesis, we have $\lambda_{n+1}(w) = \lambda_n(w)$ and $\mathcal{M}, \lambda_n \not\Vdash w : \beta$, so $k_n > 0$ and $\mathcal{M}, \lambda_n \Vdash w : \text{X}(\alpha\text{U}^{k_n-1}\beta)$. Hence, $\mathcal{M}, \lambda_{n+1} \Vdash w : \text{X}(\alpha\text{U}^{k_n-1}\beta)$ and we let $k_{n+1} := k_n$.

Finally, suppose that $\varphi_n = \text{X}\gamma$ and $\varphi_{n+1} = \gamma$. So rule X is applied at u_n . Then, $\lambda_{n+1} = S \circ \lambda_n$. By the inductive hypothesis, $k_n > 0$ and $\mathcal{M}, \lambda_n \Vdash w : \text{X}(\alpha\text{U}^{k_n-1}\beta)$, whence $\mathcal{M}, \lambda_{n+1} \Vdash w : \alpha\text{U}^{k_n-1}\beta$. We let $k_{n+1} := k_n - 1$.

Observe that $(k_n)_{n < \omega}$ is non-increasing and that $k_{n+1} < k_n$ if u_n is a modal vertex. Therefore, since π passes through infinitely many applications of rule X , we have built an infinite descending chain of natural numbers.

Case $\text{O} = \text{R}$. Then, τ is a right trace and thus $\mathcal{M}, \lambda_n \not\Vdash w : \varphi_n$ for every $n \geq N$. We inductively define a sequence of natural numbers $(k_n)_{n \geq N}$ such that:

- (i) if $\varphi_n = \gamma$, then $\mathcal{M}, \lambda_n \not\Vdash w : \alpha\text{R}^{k_n}\beta$;
- (ii) if $\varphi_n = \text{X}\gamma$, then $k_n > 0$ and $\mathcal{M}, \lambda_n \not\Vdash w : \text{X}(\alpha\text{R}^{k_n-1}\beta)$.

Let k_N be such that $\mathcal{M}, \lambda_N \not\Vdash w : \alpha\text{R}^{k_N}\beta$, which exists by Corollary 4.2.18.

If $\varphi_{n+1} = \varphi_n$, then $\lambda_{n+1}(w) = \lambda_n(w)$ and we let $k_{n+1} := k_n$.

Suppose that $\varphi_n = \gamma$ and $\varphi_{n+1} = \text{X}\gamma$. By the construction of π and the inductive hypothesis, we have $\lambda_{n+1}(w) = \lambda_n(w)$ and $\mathcal{M}, \lambda_n \Vdash w : \beta$, so $k_n > 0$ and $\mathcal{M}, \lambda_n \not\Vdash w : \text{X}(\alpha\text{R}^{k_n-1}\beta)$. Hence, $\mathcal{M}, \lambda_{n+1} \not\Vdash w : \text{X}(\alpha\text{R}^{k_n-1}\beta)$ and we let $k_{n+1} := k_n$.

Finally, suppose that $\varphi_n = \text{X}\gamma$ and $\varphi_{n+1} = \gamma$. That is, rule X is applied at u_n . Then, $\lambda_{n+1} = S \circ \lambda_n$. By the inductive hypothesis, $k_n > 0$ and $\mathcal{M}, \lambda_n \not\Vdash w : \text{X}(\alpha\text{U}^{k_n-1}\beta)$, whence $\mathcal{M}, \lambda_{n+1} \not\Vdash w : \alpha\text{U}^{k_n-1}\beta$. We let $k_{n+1} := k_n - 1$.

As before, $(k_n)_{n < \omega}$ is non-increasing and $k_{n+1} < k_n$ if u_n is a modal vertex. Therefore, since π passes through infinitely many applications of rule X, we have built an infinite descending chain of natural numbers. ■

4.2.2 Completeness of $iLTL_\infty$

We now turn to proving the completeness of the system $iLTL_\infty$. The approach is the same as the one we followed for CTL_∞^* in Section 3.2.2, namely: given a formula φ , we build a proof-search tree for φ from which either a proof or a refutation of φ can be extracted, the latter implying the invalidity of φ . There are, however, two notable differences.

First, given that our ultimate purpose is to obtain a cyclic calculus, the proof-search for $iLTL_\infty$ is considerably more delicate than for CTL_∞^* . Care is needed to ensure that it explores all relevant possibilities without producing infinitely many different sequents, as otherwise it would fail to yield regularisable proofs in general. We must thus ensure that the trees represented by the relational controls in a proof-search do not grow indefinitely neither in height nor width, and also that only finitely many world labels are used in the whole process. As shown in Figure 4.5, this need not be the case if rules are applied in an arbitrary manner. Informally, we are to find the right balance between the intuitionistic and the temporal dimensions, the former represented syntactically by implications and the latter by modal formulas.

Second, whereas CTL_∞^* refutations are subtrees of proof-search trees, $iLTL_\infty$ refutations are simply branches thereof. On the one hand, this makes it possible to dispense altogether with game- and automata-theoretic arguments like the ones we used for CTL_∞^* , since either a proof-search tree is already a proof, or else it contains a ‘bad’ branch which suffices to build the countermodel. On the other hand, given that branches are one-dimensional but $iLTL$ models are two-dimensional, it follows that refutations are densely packed with information and thus several auxiliary functions are needed to process them. All such functions have simple definitions, but showing that they interact with each other appropriately is, to some degree, a long process.

Concordantly, this section is split into two parts. First, we define the proof-search procedure on which the completeness proof relies, and show that it never produces infinitely many sequents (Proposition 4.2.38) and that it explores all relevant regions of the proof-search space (Proposition 4.2.41). Second, we define $iLTL_\infty$ refutations and show how to untangle them into two-dimensional countermodels.

iLTL_∞ proof-search trees

If we are to eventually obtain a cyclic system for iLTL, we need a proof-search procedure which breaks down formulas in a controlled, sensible manner. In particular, it must know when a formula need not be broken down, as well as put a bound on the creation of world labels. The former is achieved by using implicit contraction to saturate sequents and following precise instructions to determine which formula to break down at any given step. To achieve the latter, we define three new, admissible rules: the *collapsing* rule **coll**, the *pruning* rule **prun**, and a restricted version of the **R**→ rule.

A formula $w : \alpha$ is *left-broken* in a sequent $\Omega \dashv \Gamma \Rightarrow \Delta$ if any of the following hold:

- (i) α is a literal or a **X**-formula;
- (ii) $\alpha = \varphi \wedge \psi$, and $w : \varphi, w : \psi \in \Gamma$;
- (iii) $\alpha = \varphi \vee \psi$, and either $w : \varphi \in \Gamma$ or $w : \psi \in \Gamma$;
- (iv) $\alpha = \varphi \rightarrow \psi$, and either $w : \psi \in \Gamma$ or $w : \varphi \in \Delta$;
- (v) $\alpha = \varphi \mathbf{U} \psi$, and either $w : \psi \in \Gamma$ or $w : \varphi, w : \mathbf{X}(\varphi \mathbf{U} \psi) \in \Gamma$;
- (vi) $\alpha = \varphi \mathbf{R} \psi$, and either $w : \varphi, w : \psi \in \Gamma$ or $w : \psi, w : \mathbf{X}(\varphi \mathbf{U} \psi) \in \Gamma$.

Similarly, $w : \alpha$ is *right-broken* in $\Omega \dashv \Gamma \Rightarrow \Delta$ if any of the following hold:

- (i) α is a literal or a **X**-formula;
- (ii) $\alpha = \varphi \wedge \psi$, and either $w : \varphi \in \Delta$ or $w : \psi \in \Delta$;
- (iii) $\alpha = \varphi \vee \psi$, and $w : \varphi, w : \psi \in \Delta$;
- (iv) $\alpha = \varphi \rightarrow \psi$, and there is some $w' \in \mathbf{Worlds}$ such that $w \preceq w' \in \Omega$, $w' : \varphi \in \Gamma$, and $w' : \psi \in \Delta$;
- (v) $\alpha = \varphi \mathbf{U} \psi \in \Delta$, and either $w : \varphi, w : \psi \in \Delta$ or $w : \psi, w : \mathbf{X}(\varphi \mathbf{U} \psi) \in \Delta$;
- (vi) $\alpha = \varphi \mathbf{R} \psi \in \Delta$, and either $w : \psi \in \Delta$ or $w : \varphi, w : \mathbf{X}(\varphi \mathbf{R} \psi) \in \Delta$.

A formula $w : \alpha$ is *broken* in a sequent $S = \Omega \dashv \Gamma \Rightarrow \Delta$ if $w : \alpha$ is either left- or right-broken in S . Note that we do not require $w : \alpha \in \Gamma \cup \Delta$ for $w : \alpha$ to be left- or right-broken in S .

We say that $w : \alpha$ is *left-unbroken* (*right-unbroken*) in $\Omega \dashv \Gamma \Rightarrow \Delta$ if $w : \alpha \in \Gamma$ (respectively, $w : \alpha \in \Delta$) but $w : \alpha$ is not left-broken (respectively, right-broken) in $\Omega \dashv \Gamma \Rightarrow \Delta$. We say that $w : \alpha$ is *unbroken* in a sequent S if $w : \alpha$ is either left- or right-unbroken in S .

A relational control Ω is *saturated* if Ω is reflexive and transitive. A sequent $S = \Omega \dashv \Gamma \Rightarrow \Delta$ is *saturated* if the following hold:

- (i) Ω is saturated;
- (ii) S is monotone;
- (iii) every formula in Γ is left-broken in S ;
- (iv) every formula in Δ is right-broken in S .

We now introduce the three new rules that will be used by the proof-search procedure.

4.2.20. PROPOSITION. *The collapsing rule*

$$\text{coll} \frac{\Omega[w/w'] \dashv \Gamma, w : \Lambda \Rightarrow w : \Xi, \Delta}{\Omega \dashv \Gamma, w : \Lambda, w' : \Lambda \Rightarrow w : \Xi, w' : \Xi, \Delta},$$

where Ω is tree-like, reflexive and transitive, $w, w' \notin \Gamma \cup \Delta$, and $w <_\Omega^0 w'$, is admissible in $iLTL_\infty$.

Proof. It suffices to show that under the aforementioned conditions we have $\Omega[w/w'] = \Omega \setminus w'$, for then rule **coll** can be simulated by **wk_≲**, **Lwk**, and **Rwk**.

If $w_0 \preceq w_1 \in \Omega \setminus w'$, then $w_0 \preceq w_1 \in \Omega$ and $w_0 \neq w' \neq w_1$, whence $w_0 \preceq w_1 = (w_0 \preceq w_1)[w/w'] \in \Omega[w/w']$.

For the other inclusion, let $w_0 \preceq w_1 \in \Omega[w/w']$. Then, there is some $w'_0 \preceq w'_1 \in \Omega$ such that $w_0 \preceq w_1 = (w'_0 \preceq w'_1)[w/w']$. If $w'_0 \neq w' \neq w'_1$, then $(w'_0 \preceq w'_1)[w/w'] = w'_0 \preceq w'_1 \in \Omega \setminus w'$. If $w'_0 = w' = w'_1$, then $(w'_0 \preceq w'_1)[w/w'] = w \preceq w \in \Omega \setminus w'$, where the inclusion follows from reflexivity of Ω and $w \neq w'$. If $w'_0 = w' \neq w'_1$, then we have $w \preceq w', w' \preceq w'_1 \in \Omega$, so the transitivity of Ω yields $(w'_0 \preceq w'_1)[w/w'] = w \preceq w'_1 \in \Omega \setminus w'$. Finally, if $w'_0 \neq w' = w'_1$, then we have $w \preceq w', w'_0 \preceq w' \in \Omega$. Since Ω is tree-like and $w <_\Omega^0 w'$, it follows that either $w = w'_0$ or $w'_0 \preceq w \in \Omega$. By reflexivity, in either case we have $(w'_0 \preceq w'_1)[w/w'] = w'_0 \preceq w \in \Omega \setminus w'$. ■

The intuition behind rule **coll** is the following: if w, w' are two states of a model that satisfy and refute the same formulas and w' is an immediate intuitionistic successor of w , then, given the transitivity of the intuitionistic order, we may ignore w' and focus solely on w . Consequently, rule **coll** will put a bound on the height of the trees represented by relational controls during proof-search.

A vertex u of an $iLTL_\infty + \text{coll}$ derivation is *collapsing* if rule **coll** is applied at u . Formula traces are defined for **coll** as expected from its simulation in terms of weakening rules. Looking at the instance of **coll** in the statement of Proposition 4.2.20 above, we say that w' *collapses* down to w , that w and w' are *principal* in the conclusion, and that w is *active* in the premise.

Observe that in the proof of Proposition 4.2.20 we have proved the following:

4.2.21. PROPOSITION. *In any instance of **coll**, say with control Ω in the conclusion and with w, w' principal and w active, we have $\Omega[w/w'] = \Omega \setminus w'$.*

4.2.22. PROPOSITION. *The pruning rule*

$$\text{prun} \frac{\Omega \setminus \uparrow w \dashv \Gamma \setminus \uparrow w \Rightarrow \Delta \setminus \uparrow w}{\Omega \dashv \Gamma \Rightarrow \Delta},$$

where $w \in \Omega$ and $\uparrow w := \uparrow_{\leq \Omega} w$, is admissible in iLTL_∞ .

Proof. Clearly, **prun** can be simulated by wk_{\preccurlyeq} , **Lwk**, and **Rwk**. ■

The intuition behind (our use below of) rule **prun** is the following: if w', w'' are two different *immediate* intuitionistic successors of a state w in a model, and the trees determined by $\uparrow w'$ and $\uparrow w''$ are ‘essentially the same’ (more formally, isomorphic with respect to the labelling induced by satisfaction and refutation of formulas), then we can ignore $\uparrow w''$ and keep only the information about $\uparrow w'$. Hence, whereas rule **coll** bounds the height of the trees represented by relational controls during proof-search, rule **prun** will bound their width.

A vertex u of an $\text{iLTL}_\infty + \text{prun}$ derivation is *pruning* if rule **prun** is applied at u . Formula traces are defined for **prun** as expected from its simulation in terms of weakening rules. Looking at the instance of **prun** in the statement of Proposition 4.2.22 above, we say that w is *active* in the premise, and that every label in $\uparrow_{\leq \Omega} w$ is *pruned*.

4.2.23. PROPOSITION. *The rule*

$$\text{R}\rightarrow^- \frac{\Omega \dashv \Gamma \Rightarrow w : \varphi \rightarrow \psi, w : \psi, \Delta}{\Omega \dashv \Gamma \Rightarrow w : \varphi \rightarrow \psi, \Delta}$$

is admissible in iLTL_∞ .

Proof. By relying on implicit contraction, $\text{R}\rightarrow^-$ can be simulated thus:

$$\begin{array}{c} \text{wk}_{\preccurlyeq} \frac{\Omega \dashv \Gamma \Rightarrow w : \varphi \rightarrow \psi, w : \psi, \Delta}{\Omega, w \preccurlyeq w' \dashv \Gamma \Rightarrow w : \varphi \rightarrow \psi, w : \psi, \Delta} \\ \text{Rwk} \frac{\Omega, w \preccurlyeq w' \dashv \Gamma \Rightarrow w : \varphi \rightarrow \psi, w : \psi, \Delta}{\Omega, w \preccurlyeq w' \dashv \Gamma \Rightarrow w : \varphi \rightarrow \psi, w' : \psi, w : \psi, \Delta} \\ \text{Lwk} \frac{\Omega, w \preccurlyeq w' \dashv \Gamma, w' : \varphi \Rightarrow w : \varphi \rightarrow \psi, w' : \psi, w : \psi, \Delta}{\Omega, w \preccurlyeq w' \dashv \Gamma, w' : \varphi \Rightarrow w : \varphi \rightarrow \psi, w' : \psi, w : \psi, \Delta} \\ \text{Rm} \frac{\Omega, w \preccurlyeq w' \dashv \Gamma, w' : \varphi \Rightarrow w : \varphi \rightarrow \psi, w' : \psi, \Delta}{\Omega \dashv \Gamma \Rightarrow w : \varphi \rightarrow \psi, \Delta} \\ \text{R}\rightarrow \end{array}$$

■

The intuition behind rule $\text{R}\rightarrow^-$ is the following: if a state w of a model refutes an implication $\varphi \rightarrow \psi$, then w must refute ψ because the intuitionistic order

preserves satisfaction of formulas. For the use of $R \rightarrow^-$ during proof-search, see Remark 4.2.25 below.

Formula traces are defined for $R \rightarrow^-$ as expected from its simulation in terms of $iLTL_\infty$ rules. Looking at the instance of *prun* in the statement of Proposition 4.2.23 above, we say that $w : \varphi \rightarrow \psi$ is *principal* in the conclusion and that $w : \varphi \rightarrow \psi$ and $w : \psi$ are *active* in the premise.

Before describing the proof-search procedure for $iLTL_\infty$ we still need to define a few notions.

Let \mathcal{T} be an $iLTL_\infty$ derivation and u a vertex of \mathcal{T} , say with label $\Omega \dashv \Gamma \Rightarrow \Delta$. A labelled U-formula $w : \alpha$ is *left-pending* in u if $w : \alpha \in \Gamma$ and, for every vertex $v <_{\mathcal{T}} u$ such that rule LU is applied at v with principal formula $w : \alpha$, there is a modal vertex in $(v, u)_{\mathcal{T}}$. Dually, a labelled R-formula $w : \alpha$ is *right-pending* in u if $w : \alpha \in \Delta$ and, for every vertex $v <_{\mathcal{T}} u$ such that rule RR is applied at v with principal formula $w : \alpha$, there is a modal vertex in $(v, u)_{\mathcal{T}}$.

A *proof-search guide* is a pair (\leq_W, \leq_{iLTL}) , where \leq_W is a well-ordering of Worlds and \leq_{iLTL} a well-ordering of Form_{iLTL} .

Every tree-like relational control determines a tree $(W_\Omega, <_\Omega)$. It is thus natural to associate, to every sequent with a tree-like relational control, a labelled tree defined as follows. Let $S = \Omega \dashv \Gamma \Rightarrow \Delta$ be a sequent with Ω tree-like, and let $\Lambda_S := (\Gamma^- \times \{0\}) \cup (\Delta^- \times \{1\})$. We define the Λ_S -labelled tree $W_S := (W_\Omega, <_\Omega, \lambda_S)$ by letting $\lambda_S : W_\Omega \rightarrow 2^{\Lambda_S}$ be given by:

$$\lambda_S(w) := \{(\gamma, 0) \mid w : \gamma \in \Gamma\} \cup \{(\delta, 1) \mid w : \delta \in \Delta\}.$$

We are finally ready to define the proof-search procedure for $iLTL_\infty$.

4.2.24. DEFINITION ($iLTL_\infty$ proof-search tree). Let φ be an $iLTL$ formula, and let $\mathcal{G} = (\leq_W, \leq_{iLTL})$ be a proof-search guide. An *$iLTL_\infty$ proof-search tree for φ guided by \mathcal{G}* is a finite or infinite tree \mathcal{T} built according to the rules of $iLTL_\infty$ plus *coll*, *prun* and $R \rightarrow^-$ such that the following hold:

- (i) The root of \mathcal{T} has label $w_{\mathcal{G}} \preceq w_{\mathcal{G}} \dashv \emptyset \Rightarrow w_{\mathcal{G}} : \varphi$, where $w_{\mathcal{G}}$, called the *root-label* of \mathcal{T} , is the \leq_W -least world label.
- (ii) A vertex of \mathcal{T} is a leaf if, and only if, it is axiomatic.
- (iii) The weakening rules wk_{\preceq} , Lwk and Rwk are not applied anywhere in \mathcal{T} .
- (iv) All instances in \mathcal{T} of the rules $L\wedge$, $R\wedge$, $L\vee$, $R\vee$, $L\rightarrow$, $R\rightarrow$, LU , RU , LR and RR are preserving.
- (v) For every non-final vertex $u \in T$, say with label $S_u = \Omega_u \dashv \Gamma_u \Rightarrow \Delta_u$, the following conditions hold:

- a) If Ω_u is not reflexive, then rule **refl** is applied at u with active relation $w \preceq w$, where w is the \leq_W -least label occurring in $\Omega_u \dashv \Gamma_u \Rightarrow \Delta_u$ such that $w \preceq w \notin \Omega$.
- b) Else, if Ω_u is not transitive, then rule **trans** is applied at u with principal relations $w \preceq w'$ and $w' \preceq w''$, where (w, w', w'') is the (\leq_W, \leq_W, \leq_W) -least triple of labels such that $w \preceq w', w' \preceq w'' \in \Omega_u$ but $w \preceq w'' \notin \Omega_u$.
- c) Else, if Ω_u is not left-monotone, then rule **Lm** is applied at u with principal relation $w \preceq w'$ and principal formula $w : \varphi$, where (w, w', φ) is the $(\leq_W, \leq_W, \leq_{\text{iLTL}})$ -least triple such that $w \preceq w' \in \Omega_u$ and $w : \varphi \in \Gamma_u$ but $w' : \varphi \notin \Gamma_u$.
- d) Else, if Ω_u is not right-monotone, then rule **Rm** is applied at u with principal relation $w' \preceq w$ and principal formula $w : \varphi$, where (w, w', φ) is the $(\leq_W, \leq_W, \leq_{\text{iLTL}})$ -least triple such that $w' \preceq w \in \Omega_u$ and $w : \varphi \in \Delta_u$ but $w' : \varphi \notin \Delta_u$.
- e) Else, if Ω_u is tree-like and W_{S_u} is not thin, then rule **prun** is applied at u with active label w , where w is the \leq_W -greatest label occurring in Ω_u such that there are $w_0, w' \in \Omega_u$ satisfying $w_0 <_{\Omega_u}^0 w$, $w_0 <_{\Omega_u}^0 w'$, $w \neq w'$, and $W_{S_u}, w \sim_{\Lambda_{S_u}} W_{S_u}, w'$.
- f) Else, if Ω_u is tree-like and W_{S_u} is not low, then rule **coll** is applied at u with principal labels w and w' and active label w , where (w, w') is the (\leq_W, \leq_W) -least pair of labels occurring in Ω_u such that $w <_{\Omega_u}^0 w'$ and $\lambda_u(w) = \lambda_u(w')$.
- g) Else, if there is a left-pending formula in u , then rule **LU** is applied at u with principal formula $w : \alpha$, where (w, α) is the $(\leq_W, \leq_{\text{iLTL}})$ -least pair such that $w : \alpha$ is left-pending in u .
- h) Else, if there is a right-pending formula in u , then rule **RR** is applied at u with principal formula $w : \alpha$, where (w, α) is the $(\leq_W, \leq_{\text{iLTL}})$ -least pair such that $w : \alpha$ is right-pending in u .
- i) Else, if there is a left-unbroken formula in S_u , then a rule among **L \wedge** , **L \vee** , **L \rightarrow** and **LR** is applied at u with principal formula $w : \alpha$, where (w, α) is the $(\leq_W, \leq_{\text{iLTL}})$ -least pair such that $w : \alpha$ is left-unbroken in S_u .
- j) Else, if there is a right-unbroken formula in S_u which is not an implication, then a rule among **R \wedge** , **R \vee** and **RU** is applied at u with principal formula $w : \alpha$, where (w, α) is the $(\leq_W, \leq_{\text{iLTL}})$ -least pair such that α is not an implication and $w : \alpha$ is right-unbroken in S_u .
- k) Else, if there is an implication $w : \alpha \rightarrow \beta \in \Delta_u$ such that $w : \beta \notin \Delta_u$, then rule **R \rightarrow $^-$** is applied at u with principal formula $w_0 : \alpha_0 \rightarrow \beta_0$,

where (w_0, α_0, β_0) is the $(\leq_W, \leq_{iLTL}, \leq_{iLTL})$ -least triple such that $w_0 : \alpha_0 \rightarrow \beta_0 \in \Delta_u$ but $w_0 : \beta_0 \notin \Delta_u$.

- l) Else, if there is an implication in Δ_u which is right-unbroken in S_u , then rule $R \rightarrow$ is applied at u with principal formula $w : \alpha \rightarrow \beta$ and active formulas $w' : \alpha$ and $w' : \beta$, where (w, α, β) is the $(\leq_W, \leq_{iLTL}, \leq_{iLTL})$ -least triple such that $w : \alpha \rightarrow \beta$ is right-unbroken in S_u and w' is the \leq_W -least label not occurring in Ω_u .
- m) In any other case, S_u is saturated and rule X is applied at u with premise $\Omega_u \dashv \{w : \gamma \mid w : X\gamma \in \Gamma_u\} \Rightarrow \{w : \delta \mid w : X\delta \in \Delta_u\}$.

For every pruning vertex $u \in T$, the corresponding distinguished label w_0 in Definition 4.2.24(e) is said to be the *root* of the pruning at u . ⊣

4.2.25. REMARK. The following remarks are meant to clarify requirements in Definition 4.2.24 that might be unclear.

- (i) Rules are applied preservingly during proof-search because doing so simplifies the construction of a countermodel from a ‘failed’ proof-search. It might be possible to replicate all of our arguments below but following a contraction-free approach.
- (ii) Rule **prun** has priority over **coll** in order to obtain Lemma 4.2.37 below, i.e., a global bound on the width of the trees represented by the controls in a proof-search tree. Applying **coll** might increase the width of the tree. If we prioritise pruning, then, we only collapse labels in a thin tree and can thus ensure that the width of the resulting tree does not exceed a fixed, global bound.
- (iii) Note that the active label in an instance of **prun** is the \leq_W -greatest possible one, and thus the remaining, isomorphic one is \leq_W -smaller. This allows us to derive Proposition 4.2.35 below, which, informally, states that every ‘label trace’ only passes through finitely many instances of **prun** and **coll** and thus eventually stabilises.
- (iv) Conditions (g) and (h) ensure that every U-formula on the left and R-formula on the right is unfolded in between any two instances of the modal rule, or the root and the first instance thereof (see Lemma 4.2.51). These formulas will become principal regardless of whether they are broken or not, hence ensuring that the proof-search mechanism never misses potential good traces. So if a good trace might be created, then it will be.
- (v) Rule $R \rightarrow^-$ has priority over $R \rightarrow$ in order to obtain Lemma 4.2.39 below: every infinite branch of a proof-search tree with infinitely many instances of $R \rightarrow$

passes through infinitely many modal vertices. Informally, for every intuitionistic step that the proof-search procedure takes, it eventually takes a temporal one. Consequently, it explores the entirety of the temporal dimension.

- (vi) It is not necessary to require that rule $R \rightarrow$ have less priority than the other non-modal rules: we only need $R \rightarrow^-$ to have higher priority than $R \rightarrow$. But deferring the applications of $R \rightarrow$ as much as possible ensures that we only create new world labels when we really have to.

Conditions (e) and (f) of Definition 4.2.24 are subject to the restriction that the relational control be tree-like. We now show that this is actually superfluous.

4.2.26. PROPOSITION. *Every relational control in a proof-search tree \mathcal{T} is tree-like and rooted. Moreover, the root of every control is the root-label of \mathcal{T} .*

Proof. Let w_g be the root-label of \mathcal{T} , and for every vertex $u \in T$ let $\Omega_u \dashv \Gamma_u \Rightarrow \Delta_u$ be the label of u . We abbreviate W_{Ω_u} and \leq_{Ω_u} to W_u and \leq_u , respectively.

We argue by induction on \mathcal{T} . The base case is clear. For the inductive case, let $u, v \in T$ be such that $u <_T^0 v$ and assume that Ω_u is tree-like and rooted, with root w^* . We distinguish cases according to the rule R applied at u . If $R \notin \{R \rightarrow, \text{prun}, \text{coll}\}$, then $(W_u, \leq_u) = (W_v, \leq_v)$ and we are done by the inductive hypothesis.

Suppose that $R = R \rightarrow$. Then, clearly (W_v, \leq_v) results from (W_u, \leq_u) by the addition of a new leaf, so Ω_v is tree-like and has root w_g .

Suppose now that $R = \text{prun}$. Let w be the active label, and w^* the root of the pruning at u . In particular, then, $\Omega_v = \Omega_u \setminus \uparrow_{\leq_u} w$.

4.2.26.1. CLAIM. $W_v = W_u \setminus \uparrow_{\leq_u} w$.

Proof of claim. Let $w_0 \in W_v$. Then, there is some $w_1 \in W_v$ such that $w_0 \preceq w_1 \in \Omega_v$ or $w_1 \preceq w_0 \in \Omega_v$. By definition, $\Omega_v = \Omega_u \setminus \uparrow_{\leq_u} w$, whence $w_0 \notin \uparrow_{\leq_u} w$ and $w_1 \notin \uparrow_{\leq_u} w$. So, since $\Omega_v \subseteq \Omega_u$, we have $w_0 \in W_u \setminus \uparrow_{\leq_u} w$.

Conversely, let $w_0 \in W_u \setminus \uparrow_{\leq_u} w$. Then, there is some $w_1 \in W_u$ such that $w_0 \preceq w_1 \in \Omega_u$ or $w_1 \preceq w_0 \in \Omega_u$. Assume the latter. Then, $w_1 \notin \uparrow_{\leq_u} w$ and thus $w_1 \preceq w_0 \in \Omega_v$, whence $w_0 \in W_v$. Now assume the former. If $w_1 \notin \uparrow_{\leq_u} w$, then $w_0 \preceq w_1 \in \Omega_v$ and thus $w_0 \in W_v$. Otherwise, we have $w_0 \leq_u w^*$ because (W_u, \leq_u) is a tree and $w^* <_u^0 w$. By the reflexivity and transitivity of Ω_u , then, $w_0 \preceq w^* \in \Omega_u$ and thus $w_0 \preceq w^* \in \Omega_v$, whence again $w_0 \in W_v$. \square

Moreover, since $\Omega_v \subseteq \Omega_u$, we have $\leq_v = \leq_u \upharpoonright W_v$. So (W_v, \leq_v) results from (W_u, \leq_u) by removing a principal upset $\uparrow_{\leq_u} w \neq W_u$, where the inequality follows from the

fact that, by the inductive hypothesis and Definition 4.2.24(e), $w_g \leq_u w^* <_u^0 w$. Therefore, Ω_v is clearly tree-like and has root w_g .

Finally, suppose that $R = \text{coll}$. Let w, w' be the principal labels, and w the active one. That Ω_v is tree-like follows immediately from Proposition 4.2.21 and the fact that $w <_u^0 w'$. And Ω_v has root w_g because, by the inductive hypothesis and the reflexivity and transitivity of Ω_u , we have $w_g \preceq w \in \Omega_v$ and thus $w_g \in W_v$. ■

In the proof of Proposition 4.2.26 we have also established the following:

4.2.27. PROPOSITION. *If $u \in T$ is a pruning vertex with active label w , then we have $W_{u^+} = W_u \setminus \uparrow_{\leq \Omega} w$, where Ω is the control of u and u^+ is the unique immediate successor of u .*

In the construction of a proof-search tree, rules coll and prun are never applied to non-monotone sequents or sequents with unsaturated controls. That is to say, rules refl , trans , Lm and Rm take precedence over coll and prun . Let us now see that, once a monotone sequent with saturated control has been obtained, both coll and prun preserve the monotonicity and saturation.

4.2.28. LEMMA. *If the control of the conclusion of an instance of coll or prun is saturated, then so is the control of the premise.*

Proof. Consider an instance of coll , say:

$$\text{coll} \frac{\Omega[w/w'] \dashv \Gamma, w : \Lambda \Rightarrow w : \Xi, \Delta}{\Omega \dashv \Gamma, w : \Lambda, w' : \Lambda \Rightarrow w : \Xi, w' : \Xi, \Delta}$$

By Proposition 4.2.21, $\Omega[w/w'] = \Omega \setminus w'$. Assume that $w_0 \in \Omega \setminus w'$. Then, $w_0 \neq w'$ and thus $w_0 \preceq w_0 = (w_0 \preceq w_0)[w/w'] \in \Omega[w/w']$ by the reflexivity of Ω . Assume now that $w_0 \preceq w_1, w_1 \preceq w_2 \in \Omega \setminus w'$. Then, $w' \notin \{w_0, w_1, w_2\}$ and thus $w_0 \preceq w_2 = (w_0 \preceq w_2)[w/w'] \in \Omega[w/w']$ by the transitivity of Ω .

Consider now an instance of prun , say:

$$\text{prun} \frac{\Omega \setminus \uparrow w \dashv \Gamma \setminus \uparrow w \Rightarrow \Delta \setminus \uparrow w}{\Omega \dashv \Gamma \Rightarrow \Delta},$$

where $\uparrow w := \uparrow_{\leq \Omega} w$. Assume that $w_0 \in \Omega \setminus \uparrow w$. Then, $w_0 \notin \uparrow w$ and thus $w_0 \preceq w_0 \in \Omega \setminus \uparrow w$ by the reflexivity of Ω . Assume now that $w_0 \preceq w_1, w_1 \preceq w_2 \in \Omega \setminus \uparrow w$. Then, $\uparrow w \cap \{w_0, w_1, w_2\} = \emptyset$ and thus $w_0 \preceq w_2 \in \Omega \setminus \uparrow w$ by the transitivity of Ω . ■

4.2.29. LEMMA. *If the conclusion of an instance of coll or prun is monotone, then so is the premise.*

Proof. Consider an instance of **coll**, say:

$$\text{coll} \frac{\Omega[w/w'] \dashv \Gamma, w : \Lambda \Rightarrow w : \Xi, \Delta}{\Omega \dashv \Gamma, w : \Lambda, w' : \Lambda \Rightarrow w : \Xi, w' : \Xi, \Delta}$$

By Proposition 4.2.21, $\Omega[w/w'] = \Omega \setminus w'$. Assume that $w_0 \preceq w_1 \in \Omega \setminus w'$ and $w_0 : \alpha \in \Gamma \cup w : \Lambda$. In particular, then, $w_1 \neq w'$. So, since the conclusion is monotone, $w_1 : \alpha \in \Gamma \cup w : \Lambda$. This establishes left-monotonicity. The proof of right-monotonicity is analogous.

Consider now an instance of **prun**, say:

$$\text{prun} \frac{\Omega \setminus \uparrow w \dashv \Gamma \setminus \uparrow w \Rightarrow \Delta \setminus \uparrow w}{\Omega \dashv \Gamma \Rightarrow \Delta},$$

where $\uparrow w := \uparrow_{\leq \Omega} w$. Assume that $w_0 \preceq w_1 \in \Omega \setminus \uparrow w$ and $w_0 : \alpha \in \Gamma \setminus \uparrow w$. In particular, then, $w_1 \notin \uparrow w$. So, since the conclusion is monotone, $w_1 : \alpha \in \Gamma \setminus \uparrow w$. This establishes left-monotonicity. The proof of right-monotonicity is analogous. ■

Observe that Definition 4.2.24 describes an algorithm for building a proof-search tree for φ guided by $(\leq_W, \leq_{\text{iLTL}})$. Clearly, the chain of conditions (a)–(m) in Definition 4.2.24(v) exhausts all possible cases, whence we get:

4.2.30. PROPOSITION. *For every formula φ and every proof-search guide $\mathcal{G} = (\leq_W, \leq_{\text{iLTL}})$, there is at least one iLTL_∞ proof-search tree for φ guided by \mathcal{G} .*

4.2.31. REMARK. In contrast to the case of CTL_∞^* , the rule applied at a vertex during iLTL_∞ proof-search is not uniquely determined by the sequent labelling the vertex. This is due to conditions (g) and (h) about pending formulas in Definition 4.2.24(v).

For the remaining of this section we fix an arbitrary proof-search tree \mathcal{T} , say for a formula φ and guided by the proof-search guide $\mathcal{G} = (\leq_W, \leq_{\text{iLTL}})$. For every vertex $u \in T$, we let $S_u = \Omega_u \dashv \Gamma_u \Rightarrow \Delta_u$ be the label of u . We abbreviate W_{Ω_u} and \leq_{Ω_u} to W_u and \leq_u , respectively. Finally, we let $w_{\mathcal{G}}$ be the \leq_W -least world label, which by Proposition 4.2.26 is the root of each Ω_u .

Observe that, for every $u \in T$ and every pair of immediate successors v_1, v_2 of u , we have $W_{v_1} = W_{v_2}$. This is because no branching rule alters the controls. Let then W_u^+ denote W_v for any immediate successor of a non-final vertex $u \in T$.

We are now going to define auxiliary functions operating on the trees W_u .

For every $u \in T$, we define the *height* function $h_u : W_u \rightarrow \omega$ by letting $h_u(w)$ be the height of w in the finite tree $(W_u, <_u)$. That is, $h_u(w)$ is the cardinality of the finite set $\{w' \in W_u \mid w' < w\}$. In particular, $h_u(w_{\mathcal{G}}) = 0$. The *height* of Ω_u , in symbols $h(\Omega_u)$, is the height of W_u . The *width* of Ω_u , in symbols $b(\Omega_u)$, is the width of W_u .

Let $u \in T$ be a pruning vertex, say with active label w_0 . Let w_r be the root of the pruning at u , and let $w_1 \in W_u \setminus \{w_0\}$ be such that $w_r <_u^0 w_0, w_1$ and $W_{S_u}, w_0 \sim_{\Lambda_{S_u}} W_{S_u}, w_1$. Fix any Λ_{S_u} -isomorphism f_u between $(w_0, \rightarrow)_{W_{S_u}}$ and $(w_1, \rightarrow)_{W_{S_u}}$. We define the *pruning function* $p_u : W_u \rightarrow W_u^+$ as follows, for every $w \in W_u$:

$$p_u(w) := \begin{cases} f_u(w) & \text{if } w_0 \leq_u w \\ w & \text{otherwise} \end{cases}$$

4.2.32. LEMMA. *For every pruning vertex $u \in T$ and all labels w, w' , the following hold, where u^+ is the unique immediate successor of u :*

- (i) if $w \preceq w' \in \Omega_u$, then $p_u(w) \preceq p_u(w') \in \Omega_{u^+}$;
- (ii) $h_{u^+}(p_u(w)) = h_u(w)$.

Proof. Item (ii) follows immediately from Proposition 4.2.27 and the fact that p_u is defined in terms of an order-isomorphism and the identity. We prove (i).

Let w_0 be the active label in u^+ , let w_r be the root of the pruning at u , and let $w_1 \in W_u$ be such that $w_r <_u^0 w_1$ and $f_u : W_{S_u}, w_0 \sim_{\Lambda_{S_u}} W_{S_u}, w_1$.

Assume that $w \preceq w' \in \Omega_u$. In particular, then, $w \leq_u w'$. By construction, we have $p_u(w^*) \notin \uparrow_{\leq_u} w_0$ for every label w^* . Therefore, by Proposition 4.2.27 and the reflexivity and transitivity of Ω_u it suffices to show that $p_u(w) \leq_u p_u(w')$.

Suppose first that $w, w' \notin \uparrow_{\leq_u} w_0$. Then $p_u(w) = w \leq w' = p_u(w')$.

Suppose now that $w, w' \in \uparrow_{\leq_u} w_0$. Then, since f is an order isomorphism, we have $p_u(w) = f(w) \leq_u f(w') = p_u(w')$.

Finally, suppose that $w \notin \uparrow_{\leq_u} w_0$ but $w' \in \uparrow_{\leq_u} w_0$. Then, $w \leq_u w_r <_u^0 w_0 \leq w'$, whence $p_u(w) = w \leq_u w_r <_u^0 w_1 \leq f(w') = p_u(w')$. \blacksquare

For every vertex $u \in T$, we now define the *successor map* $s_u : W_u \rightarrow W_u^+$ as follows. If u is a pruning vertex, we let $s_u := p_u$. Assume that u is not pruning, and let $w \in W_u$. If w collapses at u , say down to w' , let $s_u(w) := w'$, and otherwise let $s_u(w) := w$.

For every pair of vertices $u, v \in T$ such that $u \leq_T v$, we define the *iterated successor map* $s_u^v : W_u \rightarrow W_v$ as follows. If $u = v$, then s_u^v is the identity map on W_u . And if $u <_T v$, say with $[u, v]_T = (u_0, \dots, u_n)$, we let $s_u^v := s_{u_{n-1}} \circ \dots \circ s_{u_0}$. Clearly, if $u \leq_T v \leq_T v'$ then $s_u^{v'} = s_v^{v'} \circ s_u^v$.

The next result shows that the successor maps follow world labels as expected:

4.2.33. LEMMA. *For all vertices $u, v \in T$ such that $u <_T^0 v$, the following hold:*

- (i) if $w_0 \preceq w_1 \in \Omega_u$, then $s_u(w_0) \preceq s_u(w_1) \in \Omega_v$;
- (ii) if u is non-modal and $w : \alpha \in \Gamma_u$, then $s_u(w) : \alpha \in \Gamma_v$;

(iii) if u is non-modal and $w : \alpha \in \Delta_u$, then $s_u(w) : \alpha \in \Delta_v$.

Proof. Let us first prove (i). By Lemma 4.2.32, the claim holds if u is pruning. Assume, then, that u is not pruning and that we have $w_0 \preceq w_1 \in \Omega_u$. The only rules allowed in proof-search trees that, when read bottom-up, remove relations from the control are **prun** and **coll**, so we are done if either $w_0 = w_1$ or if neither w_0 nor w_1 collapses at u . Hence, assume otherwise and let $i \leq 1$ be such that w_i collapses at u . Then, $s_u(w_{1-i}) = w_{1-i}$ and w_i collapses down to $s_u(w_i)$, so $s_u(w_0) \preceq s_u(w_1) = (w_0 \preceq w_1)[s_u(w_i)/w_i] \in \Omega_u[s_u(w_i)/w_i] = \Omega_v$.

We now turn to the proof of (ii) and (iii). We only show (ii), for (iii) follows analogously. Let R be the rule applied at u . If $R \in \{\text{refl, trans, Lm, Rm, } R \rightarrow^-\}$, the claim clearly holds. If $R \in \{\text{L}\wedge, \text{R}\wedge, \text{L}\vee, \text{R}\vee, \text{L}\rightarrow, \text{R}\rightarrow, \text{LU, RU, LR, RR}\}$, then the claim follows from Definition 4.2.24(iv). The only remaining cases are $R = \text{prun}$ and $R = \text{coll}$.

Suppose first that $R = \text{coll}$. If w does not collapse at u we are done. And otherwise, by the definition of **coll** and s_u we have $w : \alpha \in \Gamma_u$ if, and only if, $s_u(w) : \alpha \in \Gamma_u$, whence the claim follows.

Suppose now that $R = \text{prun}$. As before, if w is not pruned at u we are done. Assume, then, that w is pruned at u . Let w_0 be the active label, w_r the root of the pruning, and w_1 such that $w_r <_u^0 w_1$ and $f_u : W_{S_u}, w_0 \sim_{\Lambda_{S_u}} W_{S_u}, w_1$. We then have $s_u(w) = p_u(w) = f(w)$. If $w : \alpha \in \Gamma_u$, then $(\alpha, 0) \in \lambda_{S_u}(w)$ and thus $(\alpha, 0) \in \lambda_{S_u}(f(w))$, whence $s_u(w) : \alpha \in \Gamma_v$. ■

Let us now see that the height of a successor never exceeds the height of the starting label.

4.2.34. LEMMA. *For all vertices $u, v \in T$ such that $u <_T^0 v$, and every $w \in W_u$, we have $h_v(s_u(w)) \leq h_u(w)$.*

Proof. Note that, by Lemma 4.2.33, $w \in W_u$ implies $s_u(w) \in W_v$. So h_v is indeed defined on $s_u(w)$. Let R be the rule applied at u . If $R \notin \{\text{prun, coll}\}$, then $s_u(w) = w$ and we clearly have $h_v(w) = h_u(w)$. And the case where $R = \text{prun}$ is given by Lemma 4.2.32.

Assume, then, that $R = \text{coll}$. Let w_0, w_1 be the principal labels, and w_0 the active one. So w_1 collapses down to w_0 . Clearly, if $w \notin \uparrow_{\leq_u} w_1$ we have $h_v(s_u(w)) = h_v(w) = h_u(w)$, and otherwise $h_v(s_u(w)) = h_u(w) - 1$. ■

Let $u \in T$ be any vertex. We define the *ancestor* function $a_u : W_u \rightarrow \text{Worlds}^{<\omega}$ by letting $a_u(w)$ be the unique path on the tree $(W_u, <_u)$ from the root up to and including w .

The well-order \leq_W induces a strict order $<_W^{<\omega}$ on $\text{Worlds}^{<\omega}$ defined as follows: $a <_W^{<\omega} b$ if, and only if, either $|a| < |b|$, or $|a| = |b|$ and there is an $i < |a|$

such that $a(i) <_W b(i)$ and $a(j) = b(j)$ for every $j < i$. It is straightforward to check that $\leq_W^{<\omega}$ is a well-order.

We arrive now at one of several results that will show that \mathcal{T} , despite possibly being infinite, contains only a finite amount of information.

4.2.35. PROPOSITION. *Let $(u_n)_{n<\omega}$ be an infinite path through \mathcal{T} , and $w^* \in W_{u_0}$. The infinite sequence $(s_{u_0}^{u_n}(w^*))_{n<\omega}$ is eventually constant.*

Proof. For every $n < \omega$, let $w_n := s_{u_0}^{u_n}(w)$ and $c_n := (h_{u_n}(w_n), a_{u_n}(w_n))$. Since $\leq^* := (\leq, \leq_W^{<\omega})$ is a well-order on $\omega \times \text{Worlds}^{<\omega}$, it suffices to show the following for every $n < \omega$:

- (i) if w_n collapses at u_n , then $h_{u_{n+1}}(w_{n+1}) < h_{u_n}(w_n)$ and thus $c_{n+1} <^* c_n$;
- (ii) if w_n is pruned at u_n , then $h_{u_{n+1}}(w_{n+1}) = h_{u_n}(w_n)$ and $a_{u_{n+1}}(w_{n+1}) <_W^{<\omega} a_{u_n}(w_n)$, so again $c_{n+1} <^* c_n$;
- (iii) in any other case, $c_{n+1} \leq^* c_n$.

Let then $n < \omega$ be arbitrary.

If any label in $\{w \in W_{u_n} \mid w \leq_{u_n} w_n\}$ collapses at u_n , then $h_{u_{n+1}}(w_{n+1}) < h_{u_n}(w_n)$. This establishes (i).

For (ii), suppose that w_n is pruned at $u := u_n$. Let w be the active label, w_r the root of the pruning, and $w' \neq w$ such that $w_r <_u^0 w'$ and $f_u : W_{S_u}, w \sim_{\Lambda S_u} W_{S_u}, w'$. So $w_n \in \uparrow_{\leq_u} w$. Then, $w_{n+1} = s_u(w_n) = p_u(w_n) = f_u(w_n)$ and thus $h_{n+1}(w_{n+1}) = h_n(w_n)$ by Lemma 4.2.32. Additionally, by Proposition 4.2.27 we have:

- a) $a_u(w_n) = a_u(w_r) \frown (w) \frown (w, w_n]_{W_u}$;
- b) $a_u(f_u(w_n)) = a_u(w_r) \frown (w') \frown (w', f_u(w_n)]_{W_u}$.

By Definition 4.2.24(e), $w' <_W w$, whence $a_{u_{n+1}}(w_{n+1}) = a_u(f_u(w_n)) <_W^{<\omega} a_u(w_n)$, where the equality follows again from Proposition 4.2.27. This establishes (ii).

For (iii), let R be the rule applied at u_n . If $R \notin \{R \rightarrow, \text{prun}, \text{coll}\}$, then $w_{n+1} = w_n$ and $(W_{u_{n+1}}, \leq_{u_{n+1}}) = (W_{u_n}, \leq_{u_n})$, so $c_{n+1} = c_n$.

If $R = R \rightarrow$, then $w_{n+1} = w_n$ and, clearly, $c_{n+1} = c_n$.

If $R = \text{coll}$ and no label in $\{w \in W_{u_n} \mid w \leq_{u_n} w_n\}$ collapses at u_n , then again $w_{n+1} = w_n$ and $c_{n+1} = c_n$. Otherwise, $c_{n+1} <^* c_n$ follows as in the proof of (i).

Finally, if $R = \text{prun}$ and w_n is not pruned at u_n , then $w_{n+1} = w_n$ and we have $c_{n+1} = c_n$ by Proposition 4.2.27. \blacksquare

Our aim now is to show that only finitely many pairwise different sequents occur in \mathcal{T} . We are thus to find global bounds for both the heights and the widths of the controls in \mathcal{T} .

4.2.36. LEMMA. *For every $u \in T$, we have $h(\Omega_u) \leq 2(|\varphi| + 1)$.*

Proof. Let $H := 2(|\varphi| + 1)$. We argue by induction on the height of u in \mathcal{T} . If u is the root of \mathcal{T} , then $h(\Omega_u) = 1 \leq H$. For the inductive case, assume that $h(\Omega_v) \leq H$, where v is the unique predecessor v of u .

The only rule that, when read bottom-up, might increase the height of the control is $\mathbf{R} \rightarrow$, so assume that it is applied at v . Hence, Ω_v is saturated and S_v is monotone. Note that $h(\Omega_u) \leq h(\Omega_v) + 1$. So if $h(\Omega_v) < H$ then $h(\Omega_u) \leq H$ and we are done. Assume, then, and towards a contradiction, that $h(\Omega_v) = H$.

Let w_0, \dots, w_{H-1} be such that $w_0 <_v^0 \dots <_v^0 w_{H-1}$. For every $i < H$, let $\Gamma_i := \{\gamma \mid w_i : \gamma \in \Gamma_v\}$ and $\Delta_i := \{\delta \mid w_i : \delta \in \Delta_v\}$. By the monotonicity of S_v , $\Gamma_i \subseteq \Gamma_{i+1}$ and $\Delta_i \supseteq \Delta_{i+1}$ for every $i < H$.

Let $m_i := |\text{Clos}(\varphi) \setminus \Gamma_i| + |\Delta_i|$. If $(\Gamma_i, \Delta_i) \neq (\Gamma_{i+1}, \Delta_{i+1})$, then either $\Gamma_i \subsetneq \Gamma_{i+1}$ or $\Delta_i \supsetneq \Delta_{i+1}$, so in either case $m_{i+1} < m_i$. Hence, if $(\Gamma_i, \Delta_i) \neq (\Gamma_{i+1}, \Delta_{i+1})$ for every $i < H$ we then have $m_i \leq 2|\varphi| - i$ for every $i < H$ and thus $m_{H-1} < 0$, which is impossible. So there is some $k < H$ such that $(\Gamma_k, \Delta_k) = (\Gamma_{k+1}, \Delta_{k+1})$.

Therefore, W_{S_v} is not low. By Definition 4.2.24, rule $\mathbf{R} \rightarrow$ is not applied at v and we have reached a contradiction. \blacksquare

Now we turn to bounding the width of the controls.

4.2.37. LEMMA. *There is a $B < \omega$ such that, for every $u \in T$, we have $b(\Omega_u) \leq B$.*

Proof. By Lemma 4.2.36, every (W_u, \leq_u) has height at most $H := 2(|\varphi| + 1)$.

Let $\Lambda := \text{Clos}(\varphi) \times \{0, 1\}$. Since $\Lambda_{S_u} \subseteq \Lambda$ for every $u \in T$, each W_{S_u} is a Λ -labelled rooted tree. For every $h \leq H$, by Lemma 1.1.4 there is an $N_h < \omega$ such that there are at most N_h -many thin Λ -labelled trees of height h up to Λ -isomorphism, say $\mathcal{C}_h = \{\mathcal{T}_1^h, \dots, \mathcal{T}_{N_h}^h\}$. Let then $B_h < \omega$ be the maximum of the widths of the trees in \mathcal{C}_h , which exists by Corollary 1.1.5. Let $B := 1 + \sum_{h \leq H} B_h$.

We show by induction that every vertex $u \in T$ has control of width at most B . This claim is true of the root. If u is final there is nothing to show, so assume that u has an immediate successor u^+ . The only rules that, when read bottom-up, might increase the width of the control are $\mathbf{R} \rightarrow$ and coll , so assume that one of them is applied at u . Then, W_{S_u} is a thin Λ -tree and thus $b(\Omega_u) \leq B_{k_u}$, where k_u is the height of W_{S_u} .

If rule $\mathbf{R} \rightarrow$ is applied at u , then $b(\Omega_{u^+}) \leq b(\Omega_u) + 1 \leq B_{k_u} + 1 \leq B$ and we are done.

Assume that rule coll is applied at u . Then, W_{S_u} is thin. Let $w, w' \in W_u$ be such that w' collapses onto w at u . Let c_w and $c_{w'}$ be the number of immediate successors of w and w' , respectively, in W_u . Let k_w and $k_{w'}$ be the heights of the

trees $[w, \rightarrow)_{W_u}$ and $[w', \rightarrow)_{W_u}$, respectively. Since $w <_{W_u}^0 w'$, we have $k_{w'} < k_w$. And, by thinness, $c_w \leq B_{k_w}$ and $c_{w'} \leq B_{k_{w'}}$. After the collapse, we have

$$b(\Omega_{u^+}) = \max\{b(\Omega_u), c_w - 1 + c_{w'}\}.$$

If $b(\Omega_{u^+}) = b(\Omega_u) \leq B$ we are done, so assume that $b(\Omega_{u^+}) = c_w - 1 + c_{w'}$. Then, $b(\Omega_{u^+}) \leq B_{k_w} + B_{k_{w'}} - 1 \leq B$, where the last inequality follows from the fact that $k_w \neq k_{w'}$. ■

4.2.38. PROPOSITION. \mathcal{T} contains only finitely many pairwise different sequents.

Proof. Since $\text{Clos}(\varphi)$ is finite, it suffices to see that only finitely many pairwise different world labels occur in \mathcal{T} .

By Proposition 4.2.26 and Lemmas 4.2.36 and 4.2.37, there are $H, B < \omega$ such that each W_u is a rooted tree of height at most H and width at most B . Clearly, there are only finitely many such (unlabelled) rooted trees up to isomorphism. And since $\Lambda := \text{Clos}(\varphi) \times \{0, 1\}$ is finite, there are also only finitely many Λ -labelled rooted trees of height at most H and width at most B up to Λ -isomorphism. There is thus an $N < \omega$ such that every such tree has at most N -many vertices.

We claim that the names occurring in \mathcal{T} are all among w_0, \dots, w_N . This is clear of the root. Assume that it holds for a non-final vertex u . If rule $R \rightarrow$ is not applied at u , then the claim is also true of any immediate successor of u . Suppose that rule $R \rightarrow$ is applied at u . By the canonical choice of names when applying $R \rightarrow$, it suffices to see that there is an $i \leq N$ such that $w_i \notin W_u$. Suppose not, so that $\{w_0, \dots, w_N\} \subseteq W_u$. Then, W_u has at least $(N + 1)$ -many vertices, contradiction. ■

Note that Proposition 4.2.38 also holds if we eliminate every instance of coll , prun and $R \rightarrow^-$ by writing them in terms of weakening and $R \rightarrow$. This is clear for coll and prun because they can be simulated solely by weakening rules. For $R \rightarrow^-$, we simply pick a fresh name w^* and use it in all simulations of $R \rightarrow^-$.

We are almost ready to show that the proof-search procedure explores all relevant regions of the proof-search space. We need Lemma 4.2.39 below, which guarantees that proof-search only spends a finite amount of time on each intuitionistic region. Its proof is a bit technical but the underlying idea is simple: on any infinite branch there emerges a ‘stable kernel’ of world labels (denoted K in the proof) which eventually gets ‘saturated’, thus putting an end to the intuitionistic exploration and forcing the proof-search mechanism to take a modal step (apply rule X) and start exploring the ‘next’ intuitionistic region.

For every $u \in T$, we define the *children* function $c_u : W_u \rightarrow \omega$ by letting $c_u(w)$ be the cardinality of the finite set $\{w' \in W_u \mid w <_u^0 w'\}$.

4.2.39. LEMMA. *Let π be an infinite path on \mathcal{T} . If rule $\mathbf{R}\rightarrow$ is applied infinitely often on π , then so is rule \mathbf{X} .*

Proof. Let $\pi = (u_n)_{n < \omega}$ and, for every $n < \omega$, let $\Omega_n \dashv \Gamma_n \Rightarrow \Delta_n$ be the label of u_n and $W_n := W_{u_n}$. Towards a contradiction, assume that on π rule $\mathbf{R}\rightarrow$ is applied infinitely often but rule \mathbf{X} only finitely many times. Let $W := \bigcup_{n < \omega} W_n$ be the collection of labels occurring on π . Since $\text{Clos}(\varphi)$ is finite and, by Proposition 4.2.38, so is W , there is an $N_0 < \omega$ such that:

- (i) rule \mathbf{X} is never applied on $\pi_0 := (u_n)_{n \geq N_0}$;
- (ii) for every $w \in W' := \bigcup_{n \geq N_0} W_n$, there are infinitely many $n \geq N_0$ such that $w \in W_n$;
- (iii) for every $w \in W$ and $\alpha \in \text{Clos}(\varphi)$, if $w : \alpha$ is the principal formula in an instance of $\mathbf{R}\rightarrow$ on π_0 , then there are infinitely many instances of $\mathbf{R}\rightarrow$ on π_0 with principal formula $w : \alpha$.

Let $K \subseteq W'$ be the collection of labels on π_0 which are not collapsed or pruned infinitely often on π_0 . Note that $K \neq \emptyset$ because $w_{\mathcal{G}} \in K$ by Proposition 4.2.26. Hence, by (ii) and Lemma 4.2.33(i), for every $w \in K$ there is an $n_w \geq N_0$ such that $w \in \bigcap_{k \geq n_w} W_k$. Let then $N_1 \geq N_0$ be such that:

- (iv) $K \subseteq W_k$ for every $k \geq N_1$;
- (v) no $w \in K$ collapses or is pruned on $\pi_1 := (u_n)_{n \geq N_1}$.

We may thus write $W_n = K \uplus F_n$ for every $n \geq N_1$, where $F_n := W_n \setminus K$.

We claim that no $w \in K$ is the principal label in an instance of $\mathbf{R}\rightarrow$ on π_1 . Indeed, suppose otherwise and let $n \geq N_1$ be such that u_n is labelled by the conclusion of an instance of $\mathbf{R}\rightarrow$ with principal formula $w : \alpha \rightarrow \beta$ for some $w \in K$. Let w' be the active name introduced at u_n , so that we have $w \preceq w' \in \Omega_{n+1}$, $w' : \alpha \in \Gamma_{n+1}$, and $w' : \beta \in \Delta_{n+1}$. By (iii), there is an $m > n$ such that rule $\mathbf{R}\rightarrow$ is also applied at u_m with principal formula $w : \alpha \rightarrow \beta$. Let $s := s_{u_{n+1}}^{u_m}$. By (i) and Lemma 4.2.33, $s(w) \preceq s(w') \in \Omega_m$, $s(w') : \alpha \in \Gamma_m$, and $s(w') : \beta \in \Delta_m$. But $s(w) = w$ because $w \in K$ and thus $w : \alpha \rightarrow \beta$ is right-broken in $\Omega_m \dashv \Gamma_m \Rightarrow \Delta_m$, whence by Definition 4.2.24(l) we have a contradiction.

It then follows that for every $w \in K$ there is some $m_w \geq N_1$ such that w is not the root of any pruning on $(u_n)_{n \geq m_w}$. Suppose otherwise, and let $w \in K$ be such that there are infinitely many $n \geq N_1$, say $n_0 < n_1 < \dots$, such that u_n is pruning with pruning root w . Since rule $\mathbf{R}\rightarrow$ is never applied on π_1 with principal name w , we have $c_{u_n}(w) \geq c_{u_{n+1}}(w)$ for every $n \geq N_1$. Moreover, $c_{u_{n_i}}(w) > c_{u_{n_{i+1}}}(w)$ for every $i < \omega$. We have thus built an infinite descending chain of natural numbers $(c_{u_{n_i}}(w))_{i < \omega}$.

Therefore, let $N_2 \geq N_1$ be such that:

(vi) no $w \in K$ is the root of a pruning on $\pi_2 := (u_n)_{n \geq N_2}$.

By assumption, $R \rightarrow$ is applied infinitely often on π_2 . Let then $n > N_2$ be such that u_n is labelled with the premise of an instance of $R \rightarrow$, say with active label w' . Then, $w' \notin K$, so $w' \in F_n$. By Proposition 4.2.26, Ω_n is tree-like and has root $w_\mathcal{G} \in K$, so there are $w_n^K \in K \cap W_n$ and $w_n^F \in F_n$ such that $w_n^K <_{u_n}^0 w_n^F$. Since W is finite, there are $w^K \in K$ and $w^F \in W' \setminus K$ such that for infinitely many $n \geq N_2$ we have $w^F \in F_n$ and $w^K <_{u_n}^0 w^F$.

To finish the proof, we inductively build pairs $(I_n, i_n) \in 2^{\varphi^\rightarrow} \times \omega$, where φ^\rightarrow denotes the collection of all implications in $\text{Clos}(\varphi)$, such that the following hold for every $n < \omega$:

- (1) $I_n \subsetneq I_{n+1}$;
- (2) $i_n < i_{n+1}$;
- (3) $i_0 \geq N_2$;
- (4) $w^F \in F_{i_n}$ and $w^K <_{u_{i_n}}^0 w^F$;
- (5) for every $\alpha \rightarrow \beta \in I_n$, we have $w^K : \alpha \in \Gamma_{i_n}$ and $w^K : \beta \in \Delta_{i_n}$.

Let $i_0 \geq N_2$ be such that $w^F \in F_{i_0}$ and $w^K <_{u_{i_0}}^0 w^F$, and let $I_0 := \emptyset$. Assume that (I_n, i_n) has been defined. Since $w^F \notin K$, there is some $j \geq i_n$ such that w^F is collapsed or pruned at u_j and such that for every $l = i_n, \dots, j$ we have $w^F \in W_l$ and $w^K <_{W_l}^0 w^F$. By (vi), w^F cannot be pruned at u_j and hence it is collapsed. In particular, $w^F \notin W_{j+1}$. By the definition of w^K and w^F , there is some $m > j + 1$ such that $w^F \in F_m$ and $w^K <_{u_m}^0 w^F$. Since w^K cannot be the principal name in an instance of $R \rightarrow$ on π_2 , there is some $m_1 > j + 1$ such that:

- (a) rule $R \rightarrow$ is applied at u_{m_1} with active name w^F and principal formula $w : \gamma \rightarrow \delta$ for some $w \neq w^K$ satisfying $w^K \preceq w \in \Omega_{m_1}$;
- (b) there is an $m_2 > m_1$ such that $w^F \in F_{m_2}$ and $w^K <_{u_{m_2}}^0 w^F$, and such that w^F is not collapsed or pruned at u_l for any $m_1 \leq l < m_2$.

Thus, $w^F : \gamma \in \Gamma_{m_2}$ and $w^F : \gamma \in \Delta_{m_2}$. As before, since $w^F \notin K$ there is some $m_3 \geq m_2$ such that w^F collapses at u_{m_3} and such that for every $l = m_2, \dots, m_3$ we have $w^F \in F_l$ and $w^K <_{u_l}^0 w^F$. We then let $i_{n+1} := m_3$ and $I_{n+1} := I_n \cup \{\gamma \rightarrow \delta\}$. Observe that, since w^F collapses down to w^K at u_{m_3} , we have $w^K : \gamma \in \Gamma_{m_3}$ and $w^K : \delta \in \Delta_{m_3}$. And, by the inductive hypothesis and Lemma 4.2.33, for every $\alpha \rightarrow \beta \in I_n$ we have $w^K : \alpha \in \Gamma_{m_3}$ and $w^K : \beta \in \Delta_{m_3}$. It only remains to see that $\gamma \rightarrow \delta \notin I_n$. Suppose otherwise. Then, by the inductive hypothesis we have $w^K : \gamma \in \Gamma_{i_n}$ and $w^K : \delta \in \Delta_{i_n}$, so by Lemma 4.2.33 we also have $w^K : \gamma \in \Gamma_{m_1}$ and $w^K : \delta \in \Delta_{m_1}$. By the priority given to Lm and $R \rightarrow^-$ in Definition 4.2.24,

then, $w : \gamma \in \Gamma_{m_1}$ and $w : \delta \in \Delta_{m_1}$, whence $w : \delta \rightarrow \gamma$ is right-broken in the label of u_{m_1} and thus rule $\mathbf{R}\rightarrow$ cannot be applied at u_{m_1} with principal formula $w : \gamma \rightarrow \delta$, contradiction. Hence, $I_n \subsetneq I_{n+1}$.

Item (1) contradicts the finitude of $\text{Clos}(\varphi)$. ■

4.2.40. LEMMA. *Every infinite branch of \mathcal{T} contains infinitely many vertices whose labels have saturated controls.*

Proof. By definition, either rule \mathbf{refl} or \mathbf{trans} is applied at any non-axiomatic vertex $u \in T$ with unsaturated control. And, clearly, finitely many applications of \mathbf{refl} and \mathbf{trans} suffice to saturate any control. ■

We get to the final result of this section, finally establishing that proof-search traverses the entire temporal dimension:

4.2.41. PROPOSITION. *Every infinite branch of \mathcal{T} contains infinitely many modal vertices.*

Proof. Let $\pi = (u_n)_{n < \omega}$ be an infinite branch of \mathcal{T} . For every $n < \omega$, let $S_n = \Omega_n \dashv \Gamma_n \Rightarrow \Delta_n$ be the label of u_n and $W_n := W_{\Omega_n}$.

Towards a contradiction, assume that π contains only finitely many modal vertices. Then, by Lemmas 4.2.39 and 4.2.40 there is an $N_0 < \omega$ such that rules \mathbf{X} and $\mathbf{R}\rightarrow$ are not applied anywhere on $\pi_0 := (u_n)_{n \geq N_0}$ and such that Ω_{N_0} is saturated. Note that rules \mathbf{refl} and \mathbf{trans} are not applied anywhere on π_0 because any rule other than $\mathbf{R}\rightarrow$ whose conclusion has saturated control is such that all of its premises have saturated controls. And, since the only rule which, when read bottom-up, creates new world labels is $\mathbf{R}\rightarrow$, it follows that $W_n \subseteq W_{N_0}$ for every $n \geq N_0$.

Given that the premise of an instance of \mathbf{coll} or \mathbf{prun} in a proof-search tree has strictly smaller size than the conclusion and that \mathbf{X} is the only rule that may be applied to the sequent $w_g \preceq w_g \dashv \emptyset \Rightarrow \emptyset$, there is an $N_1 \geq N_0$ such that rules \mathbf{coll} or \mathbf{prun} are never applied on $\pi_1 := (u_n)_{n \geq N_1}$.

We claim that rules \mathbf{Lm} , \mathbf{Rm} and $\mathbf{R}\rightarrow$ are only applied finitely often on π_1 .

Suppose that rule \mathbf{Lm} is applied infinitely often on π_1 . Since W_{N_0} and $\text{Clos}(\varphi)$ are both finite, there is a label $w \in W_{N_0}$ and a labelled formula $w' : \alpha$ such that infinitely many instances of \mathbf{Lm} on π_1 have $w : \alpha$ as principal formula and $w' : \alpha$ as an active formula. Let $N_1 \leq n < m$ be such that u_n and u_m are each labelled by conclusions of any such instances of \mathbf{Lm} , and let $s := s_{u_{n+1}}^{u_m}$. We then have $w \preceq w' \in \Omega_{n+1}$ and $w' : \alpha \in \Gamma_{n+1}$, so by Lemma 4.2.33 we get $s(w) \preceq s(w') \in \Omega_m$ and $s(w') : \alpha \in \Gamma_m$. Since w and w' are never collapsed or pruned on π_1 , $s(w) = w$ and $s(w') = w'$, whence by definition rule \mathbf{Lm} cannot be applied at u_m with principal formula $w : \alpha$ and active formula $w' : \alpha$.

The prove that Rm is not applied infinitely often on π_1 is analogous.

Finally, suppose that rule $R\rightarrow^-$ is applied infinitely often on π_1 . Since W_{N_0} and $Clos(\varphi)$ are both finite, there is a label $w \in W_{N_0}$ and a labelled formula $w : \alpha \rightarrow \beta$ such that infinitely many instances of $R\rightarrow^-$ on π_1 have $w : \alpha \rightarrow \beta$ as principal formula. Let $N_1 \leq n < m$ be such that u_n and u_m are each labelled by conclusions of any such instances of $R\rightarrow^-$, and let $s := s_{u_{n+1}}^{u_m}$. We then have $w : \beta \in \Delta_{n+1}$, so by Lemma 4.2.33 we get $s(w) : \beta \in \Delta_m$. Since w is never collapsed or pruned on π_1 , $s(w) = w$ and thus by definition rule $R\rightarrow^-$ cannot be applied at u_m with principal formula $w : \alpha \rightarrow \beta$.

Let then $N_2 \geq N_1$ be such that rules Lm , Rm and $R\rightarrow^-$ are never applied on $\pi_2 := (u_n)_{n \geq N_2}$. For every $n \geq N_2$ and every $w \in W_n$, let

$$g_n(w) := \sum \{ \langle \alpha \rangle \mid w : \alpha \text{ is unbroken in } S_n \}.$$

Let $g_n := \sum \{ g_n(w) \mid w \in W_n \}$.

No structural rule is ever applied on π_2 and, moreover, $g_{n+1} < g_n$ if any non-structural rule other than X and $R\rightarrow^-$ is applied at u_n . Since rules X and $R\rightarrow^-$ are not applied anywhere on π_2 , we have thus built an infinite descending chain of natural numbers $(g_n)_{n \geq N_2}$. ■

$iLTL_\infty$ refutations

In this section we show how to build a countermodel of φ from any proof-search tree for φ containing a ‘bad’ branch.

A branch π of a proof-search tree is called *good* if either π is finite, or π is infinite and contains a good infinite trace. Otherwise, π is said to be *bad*.

4.2.42. DEFINITION ($iLTL_\infty$ refutation). An *$iLTL_\infty$ refutation* of a formula φ is a bad branch of an $iLTL_\infty$ proof-search tree for φ . ⊣

4.2.43. OBSERVATION. Every $iLTL_\infty$ refutation is infinite because, by definition, finite branches of $iLTL_\infty$ proof-search trees end at axiomatic leaves.

Every proof-search tree which fails to be a proof contains a refutation:

4.2.44. PROPOSITION. *Let \mathcal{T} be an $iLTL_\infty$ proof-search tree for a formula φ . If \mathcal{T} is not an $iLTL_\infty$ proof of φ , then there is a branch of \mathcal{T} which is an $iLTL_\infty$ refutation of φ .*

Proof. By definition, all finite branches of \mathcal{T} end at axiomatic leaves. So, since \mathcal{T} is not a proof, there is some infinite branch π of \mathcal{T} such that there is no good trace on φ . Hence, π is an $iLTL_\infty$ refutation of φ . ■

The term *refutation*, which we borrow from [102], is justified by Proposition 4.2.59 and Corollary 4.2.62 below.

For the rest of this section we fix an arbitrary iLTL_∞ refutation $\mathcal{R} = (R, \leq_R)$. Let u^+ denote, for every vertex $u \in R$, the unique immediate $<_R$ -successor of u in R . For every $u \in R$, let $S_u = \Omega_u \dashv \Gamma_u \Rightarrow \Delta_u$ be the label of u , and let $W_u := W_{\Omega_u}$ and $\leq_u := \leq_{\Omega_u}$.

From \mathcal{R} we build a model $\mathcal{M}_\mathcal{R} = (W_\mathcal{R}, \leq_\mathcal{R}, S_\mathcal{R}, V_\mathcal{R})$ as follows. The construction is more complicated than in the classical case due to its holographic nature: we are to build a two-dimensional structure from a one-dimensional branch of a proof-search tree.

Define a binary relation $\sim_\mathcal{R}$ on R by setting $u \sim_\mathcal{R} v$ if, and only if, there is no modal vertex in $[u, v]_\mathcal{R} \cup [v, u]_\mathcal{R}$. It is straightforward to see that $\sim_\mathcal{R}$ is an equivalence relation. A (*temporal*) *stratum* is a $\sim_\mathcal{R}$ -equivalence class. This terminology, which we borrow from [10], is justified by the fact that each stratum will give rise to a ‘temporal stage’ in $\mathcal{M}_\mathcal{R}$. By Proposition 4.2.41, there are infinitely many instances of X in \mathcal{R} and thus every stratum is finite. For every $u \in R$, we let $C_u := [u]_{\sim_\mathcal{R}}$ be the unique stratum containing u . Since (R, \leq_R) is a linear order, each stratum C is a finite sequence of vertices $u_{C,0}, \dots, u_{C,n_C}$ such that $u_{C,i+1} = u_{C,i}^+$ for every $i < n_C$. The first and last vertices of a stratum C are denoted by $\text{fst}(C)$ and $\text{lst}(C)$, respectively. For every stratum C , we let $C^+ := C_{\text{lst}(C)^+}$ be the (unique) stratum whose first vertex is the immediate successor of the last vertex in C . The collection of all strata is denoted by $\mathcal{C}_\mathcal{R}$.

For every vertex $u \in R$, we define the *denotation* function $d_u : W_u \rightarrow W_{\text{lst}(C_u)}$ by setting $d_u := s_u^{\text{lst}(C_u)}$. Observe that, if u is a pruning vertex, then $d_u(w) = d_{u^+}(p_u(w))$ for every $w \in W_u$. World labels $w, w' \in W_u$ such that $d_u(w) = d_u(w')$ will have the same denotation in the intuitionistic propositional frame corresponding to C_u .

As an immediate consequence of Lemma 4.2.33, we have:

4.2.45. LEMMA. *Let C be a stratum and $u \in C$. The following hold:*

- (i) *If $w \preceq w' \in \Omega_u$, then $d_u(w) \preceq d_u(w') \in \Omega_{\text{lst}(C_u)}$.*
- (ii) *If $w : \alpha \in \Gamma_u$, then $d_u(w) : \alpha \in \Gamma_{\text{lst}(C)}$.*
- (iii) *If $w : \alpha \in \Delta_u$, then $d_u(w) : \alpha \in \Delta_{\text{lst}(C)}$.*

For every stratum C we define the intuitionistic propositional model $\mathcal{M}_C = (W_C, \leq_C, V_C)$ as follows. For every $w \in W_{\text{lst}(C)}$, let $t_{C,w}$ be a state in W_C . We let $t_{C,w} \leq_C t_{C,w'}$ if, and only if, $w \preceq w' \in \Omega_{\text{lst}(C)}$. Since $\Omega_{\text{lst}(C)}$ is reflexive and transitive, (W_C, \leq_C) is an intuitionistic frame. And we let $p \in V_C(t_{C,w})$ if, and only if, $w : p \in \Gamma_{\text{lst}(C)}$. Note that, if $t_{C,w} \leq_C t_{C,w'}$, then $V_C(t_{C,w}) \subseteq V_C(t_{C,w'})$ because $\Omega_{\text{lst}(C)}$ is monotone. Hence, (W_C, \leq_C, V_C) is indeed an intuitionistic model.

To simplify notation in the construction of $\mathcal{M}_\mathcal{R}$ below, we assume, without loss of generality, that $C \neq C'$ implies $W_C \cap W_{C'} = \emptyset$ for all strata C, C' .

We are now ready to define the iLTL model $\mathcal{M}_\mathcal{R} = (W_\mathcal{R}, \leq_\mathcal{R}, S_\mathcal{R}, V_\mathcal{R})$. We let $W_\mathcal{R} := \bigcup_{C \in \mathcal{C}_\mathcal{R}} W_C$ be the (disjoint) union of all the W_C for every stratum C . For every state $t \in W_\mathcal{R}$, say $t = t_{C,w}$, let $S_\mathcal{R}(t) := t_{C^+, d_{\text{fst}(C^+)}(w)}$. Given that $w \in \Omega_{\text{fst}(C^+)}$ if $w \in \Omega_{\text{lst}(C)}$, the function $S_\mathcal{R}$ is well-defined. Finally, let $\leq_\mathcal{R} := \bigcup_{C \in \mathcal{C}_\mathcal{R}} \leq_C$. Note that $t_{C,w} \leq_\mathcal{R} t_{C',w'}$ implies $C = C'$. It remains to define the valuation $V_\mathcal{R} : W_\mathcal{R} \rightarrow 2^{\text{Prop}}$. For every state $t_{C,w}$, let $V_\mathcal{R}(t_{C,w}) := V_C(t_{C,w})$. Monotonicity of each V_C implies that $V_\mathcal{R}$ is $\leq_\mathcal{R}$ -monotone.

4.2.46. PROPOSITION. *The function $S_\mathcal{R}$ is confluent with respect to $\leq_\mathcal{R}$: for all $t, t' \in W_\mathcal{R}$, if $t \leq_\mathcal{R} t'$, then $S_\mathcal{R}(t) \leq_\mathcal{R} S_\mathcal{R}(t')$.*

Proof. Let $t, t' \in W_\mathcal{R}$, say $t = t_{C,w}$ and $t' = t_{C',w'}$, be such that $t \leq_\mathcal{R} t'$. Then, $C = C'$ and, by the definition of $\leq_\mathcal{R}$, $t_{C,w} \leq_C t_{C,w'}$, i.e., $w \preceq w' \in \Omega_{\text{lst}(C)}$. By Lemma 4.2.33(i), we have $s_{\text{lst}(C)}(w) \preceq s_{\text{lst}(C)}(w') \in \Omega_{\text{lst}(C)^+}$. So, since $\text{lst}(C)$ is a modal vertex, $w \preceq w' \in \Omega_{\text{fst}(C^+)}$, whence by Lemma 4.2.45(i) we get $d_{\text{fst}(C^+)}(w) \preceq d_{\text{fst}(C^+)}(w') \in \Omega_{\text{lst}(C^+)}$. Therefore, $t_{C^+, d_{\text{fst}(C^+)}(w)} \leq_{C^+} t_{C^+, d_{\text{fst}(C^+)}(w')}$ and thus $S_\mathcal{R}(t_{C,w}) \leq_\mathcal{R} S_\mathcal{R}(t_{C,w'})$. \blacksquare

Therefore, $\mathcal{M}_\mathcal{R} = (W_\mathcal{R}, \leq_\mathcal{R}, S_\mathcal{R}, V_\mathcal{R})$ is indeed an iLTL model.

4.2.47. EXAMPLE. Figure 4.8 depicts an iLTL_∞ refutation of the invalid formula $(Xp \rightarrow Xq) \rightarrow X(p \rightarrow q)$. After finitely many steps, the sequent $w \preceq w \dashv \emptyset \Rightarrow \emptyset$ is reached and thus only rule X is applied thereafter. The corresponding countermodel is depicted in Figure 4.9.

4.2.48. EXAMPLE. Figure 4.10 depicts an iLTL_∞ refutation of the invalid formula $F(Xp \rightarrow q)$. Note the use of rule prun to bound the width of the trees represented by the controls. The corresponding countermodel is depicted in Figure 4.11.

For any vertices $u, v \in R$ such that $u \leq_R v$, the *modal distance from u to v* , denoted by $d_X(u, v)$, is the number of modal vertices in $[u, v]_\mathcal{R}$.

For every vertex $u \in R$, we define the label valuation $\lambda_u : W_u \rightarrow W_{C_u}$ by setting $\lambda_u(w) := t_{C_u, d_u(w)}$. So, as we mentioned, $d_u(w) = d_u(w')$ implies $\lambda_u(w) = \lambda_u(w')$. The next lemma relates the functions $S_\mathcal{R}$ and λ_u in the expected manner.

4.2.49. LEMMA. *Let $u \leq_R v$ and $n := d_X(u, v)$. Then, for every $w \in W_u$ we have $\lambda_v(S_u^v(w)) = S_\mathcal{R}^n(\lambda_u(w))$.*

$$\begin{array}{c}
\vdots \\
\text{X} \frac{w \preceq w \dashv \emptyset \Rightarrow \emptyset}{w \preceq w \dashv \emptyset \Rightarrow \emptyset} \\
\text{coll} \frac{w \preceq w' \dashv \emptyset \Rightarrow \emptyset}{w \preceq w', w' \preceq w'', w \preceq w'' \dashv \emptyset \Rightarrow \emptyset} \\
\text{X} \frac{\text{coll} \frac{w \preceq w', w' \preceq w'', w \preceq w'' \dashv \emptyset \Rightarrow \emptyset}{w \preceq w', w' \preceq w'', w \preceq w'' \dashv \emptyset \Rightarrow \emptyset}}{w \preceq w', w' \preceq w'', w \preceq w'' \dashv \emptyset \Rightarrow \emptyset} \\
\text{trans} \frac{\text{X} \frac{w \preceq w', w' \preceq w'', w \preceq w'' \dashv \emptyset \Rightarrow \emptyset}{w \preceq w', w' \preceq w'', w \preceq w'' \dashv \emptyset \Rightarrow \emptyset}}{w \preceq w', w' \preceq w'', w \preceq w'' \dashv \emptyset \Rightarrow \emptyset} \\
\text{R} \rightarrow \frac{w \preceq w' \dashv \emptyset \Rightarrow w' : p, w' : p \rightarrow q}{w \preceq w' \dashv \emptyset \Rightarrow w' : p, w' : p \rightarrow q} \\
\text{X} \frac{w \preceq w' \dashv \emptyset \Rightarrow w' : p, w' : p \rightarrow q}{w \preceq w' \dashv \emptyset \Rightarrow w' : p, w' : p \rightarrow q} \\
\text{L} \rightarrow \frac{\text{X} \frac{w \preceq w' \dashv \emptyset \Rightarrow w' : p, w' : p \rightarrow q}{w \preceq w' \dashv \emptyset \Rightarrow w' : p, w' : p \rightarrow q}}{w \preceq w' \dashv \emptyset \Rightarrow w' : p, w' : p \rightarrow q} \\
\text{R} \rightarrow \frac{w \preceq w' \dashv \emptyset \Rightarrow w' : p, w' : p \rightarrow q}{\dashv w \preceq w \Rightarrow w : (Xp \rightarrow Xq) \rightarrow X(p \rightarrow q)}
\end{array}$$

FIGURE 4.8: An iLTL_∞ refutation of the formula $\alpha := (Xp \rightarrow Xq) \rightarrow X(p \rightarrow q)$. To improve readability, instances of refl , as well as many relations of the form $w \preceq w$, are omitted. See Figure 4.9 for the corresponding model.

Proof. By induction on n . If $n = 0$, then $C_u = C_v$ and

$$d_u(w) = s_u^{\text{lst}(C_u)}(w) = (s_v^{\text{lst}(C_u)} \circ s_u^v)(w) = (s_v^{\text{lst}(C_v)} \circ s_u^v)(w) = d_v(s_u^v(w)),$$

so

$$\lambda_v(s_u^v(w)) = t_{C_v, d_v(s_u^v(w))} = t_{C_u, d_u(w)} = \lambda_u(w).$$

For the inductive case, assume that the claim holds for n and that $d_X(u, v) = n + 1$. Then, there is a $u \leq_R v' \leq_R v$ such that $d_X(u, v') = n$ and $d_X(v', v) = 1$. Therefore:

$$\begin{aligned}
S_{\mathcal{R}}^{n+1}(\lambda_u(w)) &= S_{\mathcal{R}}[S_{\mathcal{R}}^n(\lambda_u(w))] \\
&= S_{\mathcal{R}}(\lambda_{v'}(s_u^{v'}(w))) && \text{(IH)} \\
&= S_{\mathcal{R}}(t_{C_{v'}, d_{v'}(s_u^{v'}(w))}) && \text{(def. of } \lambda_{v'}) \\
&= t_{C_{v'}^+, d_{\text{fst}(C_{v'}^+)}(d_{v'}(s_u^{v'}(w)))} && \text{(def. of } S_{\mathcal{R}}) \\
&= t_{C_v, d_{\text{fst}(C_v)}(d_{v'}(s_u^{v'}(w)))} && (C_{v'}^+ = C_v) \\
&= t_{C_v, d_{\text{fst}(C_v)}(s_u^{\text{lst}(C_{v'})}(w))} && \text{(def. of } d_{v'}) \\
&= t_{C_v, d_{\text{fst}(C_v)}(s_u^{\text{fst}(C_{v'})}(w))} && (s_{\text{lst}(C_{v'})} \text{ is the identity)} \\
&= t_{C_v, s_u^{\text{lst}(C_v)}(w)} && \text{(def. of } d_{\text{fst}(C_v)}) \\
&= t_{C_v, d_v(s_u^v(w))} && \text{(def. of } d_v)
\end{aligned}$$

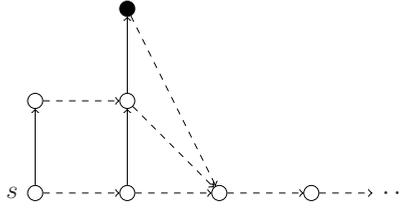


FIGURE 4.9: The model \mathcal{M} built from the refutation depicted in Figure 4.8. The valuation is empty everywhere except at the state painted black, where p and only p holds. We have $\mathcal{M}, s \not\models (Xp \rightarrow Xq) \rightarrow X(p \rightarrow q)$.

$$= \lambda_v(s_u^v(w)) \quad (\text{def. of } \lambda_v)$$

■

4.2.50. COROLLARY. $S_{\mathcal{R}}(\lambda_u(w)) = \lambda_{\text{fst}(C_u^+)}(d_u(w))$.

A world label w *persists* in a path $[u, v]_{\mathcal{R}}$ if there is no vertex $u' \in [u, v]_{\mathcal{R}}$ such that w collapses or is pruned at u' . A label w *persists* in a stratum C if w persists in $[\text{fst}(C), \text{lst}(C)]_{\mathcal{R}}$.

The following ensures that every pending formula at the first vertex in a stratum is eventually broken down in the stratum.

4.2.51. LEMMA. *Let C be a stratum. If a formula $w : \alpha$ is pending at $u := \text{fst}(C)$ and w persists in C , then there is a vertex $v \in C$ such that $w : \alpha \triangleright_u^v w : \alpha$ and in which rule LU or RR is applied with principal formula $w : \alpha$.*

Proof. By Definition 4.2.24, after finitely many applications of rules **refl**, **trans**, **Lm** and **Rm** we arrive at a vertex $u' \geq_R u$ whose label is monotone and has saturated control. The premise of an instance of **prun** or **coll** has strictly lower complexity than the conclusion, so after finitely many applications of rules **prun** and **coll** we arrive at a vertex $u'' \geq_R u'$ such that:

- (i) $W_{S_{u''}}$ is thin and low;
- (ii) $\Omega_{u''}$ is saturated, by Lemma 4.2.28;
- (iii) $S_{u''}$ is monotone, by Lemma 4.2.29.

Since w persists in $[u, u'']_{\mathcal{R}} \subseteq C$, the formula $w : \alpha$ is still pending at u'' . Hence, after finitely many instances of rules **LU** and **RR** with principal formulas some formulas other than $w : \alpha$ which are pending at u'' , we arrive at a vertex $v \geq_R u''$

$$\begin{array}{c}
\vdots \\
\text{prun} \frac{(\dagger) \quad w_0 \preceq w_1, w_2 \dashv w_2 : p \Rightarrow w_0 : \beta}{w_0 \preceq w_1, w_2, w_3 \dashv w_2 : p \Rightarrow w_0 : \beta} \\
\text{X} \frac{}{w_0 \preceq w_1, w_2, w_3 \dashv w_3 : p, w_2 : \text{X}p \Rightarrow w_2 : q, w_0 : \gamma, w_0 : \text{X}\beta, w_0 : \beta} \\
\text{R} \rightarrow \frac{\text{RF} \frac{w_0 \preceq w_1, w_3 \dashv w_3 : p \Rightarrow w_0 : \gamma, w_0 : \text{X}\beta, w_0 : \beta}{\text{prun} \frac{w_0 \preceq w_1, w_3 \dashv w_3 : p \Rightarrow w_0 : \beta}{w_0 \preceq w_1, w_2, w_3 \dashv w_3 : p \Rightarrow w_0 : \beta}}}{w_0 \preceq w_1, w_2, w_3 \dashv w_2 : p, w_3 : \text{X}p \Rightarrow w_3 : q, w_0 : \gamma, w_0 : \text{X}\beta, w_0 : \beta} \\
\text{R} \rightarrow \frac{\text{RF} \frac{w_0 \preceq w_1, w_2 \dashv w_2 : p \Rightarrow w_0 : \gamma, w_0 : \text{X}\beta, w_0 : \beta}{(\dagger) \quad w_0 \preceq w_1, w_2 \dashv w_2 : p \Rightarrow w_0 : \beta}}{\text{X} \frac{}{w_0 \preceq w_1, w_2 \dashv w_1 : p, w_2 : \text{X}p \Rightarrow w_2 : q, w_0 : \gamma, w_0 : \text{X}\beta, w_0 : \beta}} \\
\text{R} \rightarrow \frac{\text{RF} \frac{w_0 \preceq w_1 \dashv w_1 : p \Rightarrow w_0 : \gamma, w_0 : \text{X}\beta, w_0 : \beta}{w_0 \preceq w_1 \dashv w_1 : p \Rightarrow w_0 : \beta}}{\text{X} \frac{}{w_0 \preceq w_1 \dashv w_1 : \text{X}p \Rightarrow w_1 : q, w_0 : \gamma, w_0 : \text{X}\beta, w_0 : \beta}} \\
\text{R} \rightarrow \frac{\text{RF} \frac{w_0 \preceq w_0 \dashv \emptyset \Rightarrow w_0 : \gamma, w_0 : \text{X}\beta, w_0 : \beta}{w_0 \preceq w_0 \dashv \emptyset \Rightarrow w_0 : \beta}}{}
\end{array}$$

FIGURE 4.10: An iLTL_∞ refutation of the formula $\beta := \text{F}\gamma$, where $\gamma := \text{X}p \rightarrow q$. To improve readability, instances of refl , as well as many relations of the form $w \preceq w$, are omitted. The symbol ‘ \dagger ’ marks roots of identical subproofs. See Figure 4.11 for the corresponding model.

such that $\Omega_v = \Omega_{v'}$ and in which the \leq_{iLTL} -least pending formula is $w : \alpha$, whence rule LU or RR is applied at v with principal formula $w : \alpha$. Finally, since $w : \alpha$ is a side formula at every vertex in $[u, v]_{\mathcal{R}}$, we have $w : \alpha \triangleright_u^v w : \alpha$. ■

The following three lemmas about X -, U - and R -formulas with persistent labels on a stratum are oriented towards proving Lemmas 4.2.55 to 4.2.58 below.

4.2.52. LEMMA. *Let $u \in R$. If $w : \text{X}\alpha \in \Gamma_u \cup \Delta_u$ and w persists in $[u, \text{lst}(C_u)]_{\mathcal{R}}$, then $w : \text{X}\alpha \triangleright_u^{\text{lst}(C_u)} w : \text{X}\alpha$ and this trace is side-stable.*

Proof. By Definition 4.2.24, X -formulas may only be principal at instances of prun , coll and X . And, by assumption, w is not pruned or collapsed in $[u, \text{lst}(C_u)]_{\mathcal{R}}$, so the formula $w : \text{X}\alpha$ is a side formula at every vertex in $[u, \text{lst}(C_u)]_{\mathcal{R}}$ and thus $w : \text{X}\alpha \triangleright_u^{\text{lst}(C_u)} w : \text{X}\alpha$. This trace is clearly side-stable because no modal formula is ever principal when applying $\text{L} \rightarrow$ or $\text{R} \rightarrow$. ■

4.2.53. LEMMA. *Let C be a stratum. If a formula $w : \alpha \cup \beta$ is left-pending at $\text{fst}(C)$, w persists in C , and $w : \beta \notin \Gamma_u$ for any $u \in C$, then there is a trace τ on*

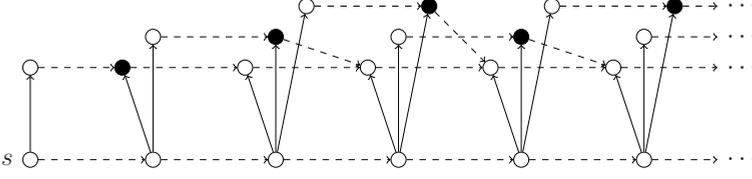


FIGURE 4.11: The model \mathcal{M} built from the refutation depicted in Figure 4.10. The valuation is empty everywhere except at the states painted black, where p and only p holds. We have $\mathcal{M}, s \not\models F(Xp \rightarrow q)$.

C of the form

$$(w : \alpha\text{U}\beta, 0) \triangleright \cdots \triangleright (w : \alpha\text{U}\beta, 0) \triangleright (w : X(\alpha\text{U}\beta), 0) \triangleright \cdots \triangleright (w : X(\alpha\text{U}\beta), 0).$$

Proof. By Lemma 4.2.51, there is a $v \in C$ such that $w : \alpha\text{U}\beta \triangleright_{\text{fst}(C)}^v w : \alpha\text{U}\beta$ and in which rule LU is applied with principal formula $w : \alpha\text{U}\beta$. Since by assumption $w : \alpha\text{U}\beta \not\triangleright_v w : \beta$, we have $w : \alpha\text{U}\beta \triangleright_v X(\alpha\text{U}\beta)$, whence by Lemma 4.2.52 $w : X(\alpha\text{U}\beta) \triangleright_{v^+}^{\text{lst}(C)} X(\alpha\text{U}\beta)$. ■

Analogously:

4.2.54. LEMMA. *Let C be a stratum. If a formula $w : \alpha\text{R}\beta$ is right-pending at $\text{fst}(C)$, w persists in C , and $w : \beta \notin \Delta_u$ for any $u \in C$, then there is a trace τ on C of the form*

$$(w : \alpha\text{R}\beta, 1) \triangleright \cdots \triangleright (w : \alpha\text{R}\beta, 1) \triangleright (w : X(\alpha\text{R}\beta), 1) \triangleright \cdots \triangleright (w : X(\alpha\text{R}\beta), 1).$$

The next lemma will be used to satisfy U-formulas occurring in antecedents of sequents when verifying the countermodel built from a refutation.

4.2.55. LEMMA. *For every vertex $u \in R$ and every U-formula $w : \alpha\text{U}\beta$, if $w : \alpha\text{U}\beta \in \Gamma_u$, then there is a $u \leq_R v$ such that:*

- (i) $d_v(s_u^v(w)) : \beta \in \Gamma_{\text{lst}(C_v)}$;
- (ii) for every $u \leq_R v'$, if $\text{lst}(C_{v'}) <_R \text{lst}(C_v)$, then $d_{v'}(s_u^{v'}(w)) : \alpha \in \Gamma_{\text{lst}(C_{v'})}$ and $d_{v'}(s_u^{v'}(w)) : X(\alpha\text{U}\beta) \in \Gamma_{\text{lst}(C_{v'})}$.

Proof. We begin by inductively constructing an infinite sequence $(u_n)_{n < \omega}$ of vertices in $[u, \rightarrow]_R$ such that the following hold for every $n < \omega$, where $w_n := s_u^{u_n}(w)$:

- (a) $u_0 = u$;

- (b) $w_n : \alpha \mathbf{U} \beta \in \Gamma_{u_n}$;
- (c) if u_n is modal and $w_n : \beta \in \Gamma_{u_n}$, then $u_{n+1} = u_n$;
- (d) otherwise, $u_{n+1} = u_n^+$.

Let $u_0 := u$. Assume that u_n has been defined. If u_n is not modal, we let $u_{n+1} := u_n^+$. If u_n is modal, then by (b) and saturation either $w_n : \beta \in \Gamma_{u_n}$ or $w_n : \alpha, w_n : \mathbf{X}(\alpha \mathbf{U} \beta) \in \Gamma_{u_n}$. If the former holds, we let $u_{n+1} := u_n$, and otherwise we let $u_{n+1} := u_n^+$.

Towards a contradiction, assume that the sequence $(u_n)_{n < \omega}$ is not eventually constant. Then, moreover, by construction we have $u_{n+1} \neq u_n$ for every $n < \omega$. If $n < \omega$ is such that u_n is modal, by (b), (c) and saturation, we have $w_n : \alpha, w_n : \mathbf{X}(\alpha \mathbf{U} \beta) \in \Gamma_{u_n}$. For every $n < \omega$, let $C_n := C_{u_n}$.

By Proposition 4.2.35, there is an $0 < N < \omega$ such that $w_i = w_j$ for all $i, j \geq N$ and $C_{N-1} \neq C_N$ and $u_N = \mathbf{fst}(C_N)$. Let $w^* := w_N$. By (b) and the fact that u_N is labelled by the premise of an instance of \mathbf{X} , the formula $w^* : \alpha \mathbf{U} \beta$ is left-pending at u_N .

We inductively build infinite sequences $(N_i)_{i < \omega}$ and $(\tau_i)_{i < \omega}$ of natural numbers and traces, respectively, such that the following hold for every $i < \omega$:

- (1) $N \leq N_i < N_{i+1}$;
- (2) $u_{N_i} = \mathbf{fst}(C_{N_i})$;
- (3) $C_{N_{i+1}} = C_{N_i}^+$;
- (4) τ_i is a trace on C_{N_i} of the form $(w^* : \alpha \mathbf{U} \beta, 0) \triangleright \cdots \triangleright (w^* : \alpha \mathbf{U} \beta, 0) \triangleright (w^* : \mathbf{X}(\alpha \mathbf{U} \beta), 0) \triangleright \cdots \triangleright (w^* : \mathbf{X}(\alpha \mathbf{U} \beta), 0)$.

Let $N_0 := N$, and let τ_0 be the trace on C_N starting from $w^* : \alpha \mathbf{U} \beta$ given by Lemma 4.2.53. For the inductive case, assume that N_i and τ_i have been defined. Since $w^* : \mathbf{X}(\alpha \mathbf{U} \beta) \triangleright_{\mathbf{lst}(C_{N_i})} w^* : \alpha \mathbf{U} \beta$, the formula $w^* : \alpha \mathbf{U} \beta$ is left-pending at $\mathbf{fst}(C_{N_i}^+)$. Let then N_{i+1} be such that $u_{N_{i+1}} = \mathbf{fst}(C_{N_i}^+)$, and let τ_{i+1} be the trace on $C_{N_{i+1}}$ starting from $w^* : \alpha \mathbf{U} \beta$ given by Lemma 4.2.53.

Since for every $i < \omega$ the trace τ_{i+1} is coherent with τ_i , the infinite trace $\tau := \tau_0 \widehat{\ } \tau_1 \widehat{\ } \tau_2 \widehat{\ } \cdots$ is a good left trace on $(u_n)_{n \geq N}$. This contradicts the fact that, by definition, refutations do not contain good infinite traces.

Therefore, the sequence $(u_n)_{n < \omega}$ must be eventually constant. By construction, then, there is a least $N < \omega$ such that u_N is modal and $w_N : \beta \in \Gamma_{u_N}$. We claim that we may take $v := u_N$ to satisfy (i) and (ii) above. That (i) holds is clear because $d_v(s_u^v(w)) = s_u^{\mathbf{lst}(C_v)}(w) = s_u^v(w) = w_N$. To see that (ii) is also the case, let v' be such that $u \leq_R v'$ and $\mathbf{lst}(C_{v'}) <_R \mathbf{lst}(C_v) = v$. Then, there is some $m < N$ such that $u_m = \mathbf{lst}(C_{v'})$. Since u_m is modal but $u_{m+1} \neq u_m$, by construction we

have $w_m : \beta \notin \Gamma_{u_m}$, so by saturation $w_m : \alpha, w_m : \text{X}(\alpha\text{U}\beta) \in \Gamma_{u_m}$. Given that $d_{v'}(s_u^{v'})(w) = s_u^{\text{lst}(C_{v'})}(w) = s_u^{u_m}(w) = w_m$, we are done. ■

Analogously, we can show:

4.2.56. LEMMA. *For every vertex $u \in R$ and every R-formula $w : \alpha\text{R}\beta$, if $w : \alpha\text{R}\beta \in \Delta_u$, then there is a $u \leq_R v$ such that:*

- (i) $d_v(s_u^v(w)) : \beta \in \Delta_{\text{lst}(C_v)}$;
- (ii) for every $u \leq_R v'$, if $\text{lst}(C_{v'}) <_R \text{lst}(C_v)$, then $d_{v'}(s_u^{v'})(w) : \alpha \in \Delta_{\text{lst}(C_{v'})}$ and $d_{v'}(s_u^{v'})(w) : \text{X}(\alpha\text{U}\beta) \in \Delta_{\text{lst}(C_{v'})}$.

The next lemma will be used to refute U-formulas occurring in succedents of sequents when verifying the countermodel built from a refutation.

4.2.57. LEMMA. *For every vertex $u \in R$ and every U-formula $w : \alpha\text{U}\beta$, if $w : \alpha\text{U}\beta \in \Delta_u$, then there is an infinite sequence $(u_n)_{n < \omega}$ of vertices in $[u, \rightarrow]_{\mathcal{R}}$ such that the following hold, where $w_n := s_u^{u_n}(w)$:*

- (i) $u_0 = u$;
- (ii) $w_n : \beta, w_n : \alpha\text{U}\beta \in \Delta_{u_n}$;
- (iii) if u_n is modal and $w_n : \alpha \in \Delta_{u_n}$, then $u_{n+1} = u_n$;
- (iv) otherwise, $u_{n+1} = u_n^+$.

Proof. We build the sequence $(u_n)_{n < \omega}$ inductively. Let $u_0 := \text{lst}(C_u)$. By Lemma 4.2.45 and saturation, we have $w_0 : \beta, w_0 : \alpha\text{U}\beta \in \Delta_{u_0}$. For the inductive case, assume that u_n has been defined. If u_n is modal, then by (ii) and saturation either $w_n : \alpha, w_n : \beta \in \Delta_{u_n}$ or $w_n : \beta, w_n : \text{X}(\alpha\text{U}\beta) \in \Delta_{u_n}$. If the former holds, we let $u_{n+1} := u_n$, and otherwise we let $u_{n+1} := u_n^+$. ■

Analogously:

4.2.58. LEMMA. *For every vertex $u \in R$ and every R-formula $w : \alpha\text{R}\beta$, if $w : \alpha\text{R}\beta \in \Gamma_u$, then there is an infinite sequence $(u_n)_{n < \omega}$ of vertices in $[u, \rightarrow]_{\mathcal{R}}$ such that the following hold, where $w_n := s_u^{u_n}(w)$:*

- (i) $u_0 = u$;
- (ii) $w_n : \beta, w_n : \alpha\text{R}\beta \in \Gamma_{u_n}$;
- (iii) if u_n is modal and $w_n : \alpha \in \Gamma_{u_n}$, then $u_{n+1} = u_n$;
- (iv) otherwise, $u_{n+1} = u_n^+$.

Since the formula φ and the refutation \mathcal{R} that we fixed were arbitrary, we can now obtain:

4.2.59. PROPOSITION. *For every formula φ , if there is an iLTL_∞ refutation of φ , then φ is not valid.*

Proof. Let $\mathcal{R} = (R, \leq_R)$ be an iLTL_∞ refutation of φ and, for every $u \in R$, let $\Omega_u \dashv \Gamma_u \Rightarrow \Delta_u$ be the label of u . By structural induction on χ we show that, for every formula χ , every vertex $u \in R$ and every world label $w \in \Omega_u$, if $w : \chi \in \Gamma_u$ then $\mathcal{M}_{\mathcal{R}}, \lambda_u(w) \Vdash \chi$, and if $w : \chi \in \Delta_u$ then $\mathcal{M}_{\mathcal{R}}, \lambda_u(w) \not\Vdash \chi$. In particular, therefore, $\mathcal{M}_{\mathcal{R}}, \lambda_r(w_{\mathcal{G}}) \not\Vdash \varphi$, where r is the root of \mathcal{R} and $w_{\mathcal{G}} \in \text{Worlds}$ is the unique world label occurring in the sequent labelling r .

Case $\chi = \perp$. We cannot have $w : \perp \in \Gamma_u$ because refutations contain no axiomatic vertices. And the case $w : \perp \in \Delta_u$ is clear.

Case $\chi = p$. Assume that $w : p \in \Gamma_u$. By Lemma 4.2.45, $d_u(w) : p \in \Gamma_{\text{lst}(C_u)}$, so $p \in V_{\mathcal{R}}(t_{C_u, d_u(w)})$ and thus $\mathcal{M}_{\mathcal{R}}, \lambda_u(w) \Vdash p$. Assume now that $w : p \in \Delta_u$. By Lemma 4.2.45, $d_u(w) : p \in \Delta_{\text{lst}(C_u)}$. Since refutations do not contain axiomatic vertices, $d_u(w) : p \notin \Gamma_{\text{lst}(C_u)}$ and thus $\mathcal{M}_{\mathcal{R}}, \lambda_u(w) \not\Vdash p$.

Case $\chi = \alpha \wedge \beta$. Assume that $w : \alpha \wedge \beta \in \Gamma_u$. By Lemma 4.2.45, $d_u(w) : \alpha \wedge \beta \in \Gamma_{\text{lst}(C_u)}$, so by saturation $d_u(w) : \alpha, d_u(w) : \beta \in \Gamma_{\text{lst}(C_u)}$. Hence, by the inductive hypothesis $\mathcal{M}_{\mathcal{R}}, \lambda_{\text{lst}(C_u)}(d_u(w)) \Vdash \chi$. Since $\lambda_u(w) = \lambda_{\text{lst}(C_u)}(d_u(w))$, we are done. Assume now that $w : \alpha \wedge \beta \in \Delta_u$. By Lemma 4.2.45, $d_u(w) : \alpha \wedge \beta \in \Delta_{\text{lst}(C_u)}$, so by saturation either $d_u(w) : \alpha \in \Delta_{\text{lst}(C_u)}$ or $d_u(w) : \beta \in \Delta_{\text{lst}(C_u)}$. Hence, by the inductive hypothesis $\mathcal{M}_{\mathcal{R}}, \lambda_{\text{lst}(C_u)}(d_u(w)) \not\Vdash \chi$ and we are done as before.

Case $\chi = \alpha \vee \beta$. Dually analogous to the previous case.

Case $\chi = X\alpha$. Assume that $w : X\alpha \in \Gamma_u$. By Lemma 4.2.45, we have $d_u(w) : X\alpha \in \Gamma_{\text{lst}(C_u)}$, whence $d_u(w) : \alpha \in \Gamma_{\text{fst}(C_u^+)}$. By the inductive hypothesis, then, $\mathcal{M}_{\mathcal{R}}, \lambda_{\text{fst}(C_u^+)}(d_u(w)) \Vdash \alpha$. Since $S_{\mathcal{R}}(\lambda_u(w)) = \lambda_{\text{fst}(C_u^+)}(d_u(w))$ by Corollary 4.2.50, we are done. Assume now that $w : X\alpha \in \Delta_u$. By Lemma 4.2.45, we have $d_u(w) : X\alpha \in \Delta_{\text{lst}(C_u)}$, whence $d_u(w) : \alpha \in \Delta_{\text{fst}(C_u^+)}$. By the inductive hypothesis, $\mathcal{M}_{\mathcal{R}}, \lambda_{\text{fst}(C_u^+)}(d_u(w)) \not\Vdash \alpha$ and we are done as before.

Case $\chi = \alpha \cup \beta$. Assume $w : \alpha \cup \beta \in \Gamma_u$. By Lemma 4.2.55, there is a $u \leq_R v$ such that:

- (i) $d_v(s_u^v(w)) : \beta \in \Gamma_{\text{lst}(C_v)}$;
- (ii) for every $u \leq_R v'$, if $\text{lst}(C_{v'}) <_R \text{lst}(C_v)$, then $d_{v'}(s_u^{v'}(w)) : \alpha \in \Gamma_{\text{lst}(C_{v'})}$ and $d_{v'}(s_u^{v'}(w)) : X(\alpha \cup \beta) \in \Gamma_{\text{lst}(C_{v'})}$.

By the inductive hypothesis, $\mathcal{M}_{\mathcal{R}}, \lambda_{\text{lst}(C_v)}(s_u^{\text{lst}(C_v)}(w)) \Vdash \beta$. Let $n := d_X(u, \text{lst}(C_v))$. By Lemma 4.2.49, $\mathcal{M}_{\mathcal{R}}, S_{\mathcal{R}}^n(\lambda_u(w)) \Vdash \beta$. Let $m < n$. Let $u \leq_R v'$ be such that

$d_X(u, v') = m$. Then, $\text{lst}(C_{v'}) <_R \text{lst}(C_v)$, so we have $s_u^{\text{lst}(C_{v'})}(w) : \alpha \in \Gamma_{\text{lst}(C_{v'})}$. By the inductive hypothesis, then, $\mathcal{M}_{\mathcal{R}}, \lambda_{\text{lst}(C_{v'})}(s_u^{\text{lst}(C_{v'})}(w)) \Vdash \alpha$. By Lemma 4.2.49, $\lambda_{\text{lst}(C_{v'})}(s_u^{\text{lst}(C_{v'})}(w)) = S_{\mathcal{R}}^m(\lambda_u(w))$, so $\mathcal{M}_{\mathcal{R}}, S_{\mathcal{R}}^m(\lambda_u(w)) \Vdash \alpha$. This establishes that $\mathcal{M}_{\mathcal{R}}, \lambda_u(w) \Vdash \alpha \cup \beta$.

Assume now that $w : \alpha \cup \beta \in \Delta_u$. By Lemma 4.2.57, there is an infinite sequence $(u_n)_{n < \omega}$ of vertices in $[u, \rightarrow)_{\mathcal{R}}$ such that the following hold, where $w_n := s_u^{u_n}(w)$:

- (i) $u_0 = u$;
- (ii) $w_n : \beta, w_n : \alpha \cup \beta \in \Delta_{u_n}$;
- (iii) if u_n is modal and $w_n : \alpha \in \Delta_{u_n}$, then $u_{n+1} = u_n$;
- (iv) otherwise, $u_{n+1} = u_n^+$.

Suppose first that the sequence $(u_n)_{n < \omega}$ is not eventually constant, and let $k < \omega$. Let $u \leq_R v$ be such that $d_X(u, v) = k$. Then, $s_u^v(w) : \beta \in \Delta_v$, so by the inductive hypothesis $\mathcal{M}_{\mathcal{R}}, \lambda_v(s_u^v(w)) \not\Vdash \beta$. By Lemma 4.2.49, $\lambda_v(s_u^v(w)) = S_{\mathcal{R}}^k(\lambda_u(w))$. So $\mathcal{M}_{\mathcal{R}}, S_{\mathcal{R}}^k(\lambda_u(w)) \not\Vdash \beta$ for every $k < \omega$, whence $\mathcal{M}_{\mathcal{R}}, \lambda_u(w) \not\Vdash \alpha \cup \beta$.

Suppose now that the sequence $(u_n)_{n < \omega}$ is eventually constant, and let $N < \omega$ be least such that $u_N = u_{N+1}$. So $w_N : \alpha, w_N : \beta \in \Delta_{u_N}$. By the inductive hypothesis, then, $\mathcal{M}_{\mathcal{R}}, \lambda_{u_N}(s_{u_N}^{u_N}(w)) \not\Vdash \alpha$ and $\mathcal{M}_{\mathcal{R}}, \lambda_{u_N}(s_{u_N}^{u_N}(w)) \not\Vdash \beta$ for every $n \leq N$. Let $k := d_X(u, u_N)$. By Lemma 4.2.49, $\lambda_{u_N}(s_{u_N}^{u_N}(w)) = S_{\mathcal{R}}^k(\lambda_u(w))$, so $\mathcal{M}_{\mathcal{R}}, S_{\mathcal{R}}^k(\lambda_u(w)) \not\Vdash \alpha$ and $\mathcal{M}_{\mathcal{R}}, S_{\mathcal{R}}^k(\lambda_u(w)) \not\Vdash \beta$. Let $j < k$, and let $n < N$ be such that $d_X(u, u_n) = j$. By Lemma 4.2.49, $\lambda_{u_n}(s_{u_n}^{u_n}(w)) = S_{\mathcal{R}}^j(\lambda_u(w))$, so $\mathcal{M}_{\mathcal{R}}, S_{\mathcal{R}}^j(\lambda_u(w)) \not\Vdash \beta$. This shows that $\mathcal{M}_{\mathcal{R}}, \lambda_u(w) \not\Vdash \alpha \cup \beta$.

Case $\chi = \alpha R \beta$. Dually analogous to the previous case, using Lemmas 4.2.56 and 4.2.58 in place of Lemmas 4.2.55 and 4.2.57, respectively.

Case $\chi = \alpha \rightarrow \beta$. Assume that $w : \alpha \rightarrow \beta \in \Gamma_u$. Let $t = t_{C', w'}$ be such that $\lambda_u(w) \leq_{\mathcal{R}} t$, i.e., $t_{C_u, d_u(w)} \leq_{\mathcal{R}} t_{C', w'}$.² Then, $C' = C_u$ and $d_u(w) \preceq w' \in \Omega_{\text{lst}(C_u)}$. By Lemma 4.2.45, $d_u(w) : \alpha \rightarrow \beta \in \Gamma_{\text{lst}(C_u)}$. By monotonicity, $w' : \alpha \rightarrow \beta \in \Gamma_{\text{lst}(C_u)}$. By saturation, either $w' : \alpha \in \Delta_{\text{lst}(C_u)}$ or $w' : \beta \in \Gamma_{\text{lst}(C_u)}$. By the inductive hypothesis, then, either $\mathcal{M}_{\mathcal{R}}, \lambda_{\text{lst}(C_u)}(w') \not\Vdash \alpha$ or $\mathcal{M}_{\mathcal{R}}, \lambda_{\text{lst}(C_u)}(w') \Vdash \beta$. Since $\lambda_{\text{lst}(C_u)}(w') = t_{C_{\text{lst}(C_u)}, d_{\text{lst}(C_u)}(w')} = t_{C_u, w'} = t_{C', w'} = t$ and t was an arbitrary intuitionistic successor of $\lambda_u(w)$, we have established $\mathcal{M}_{\mathcal{R}}, \lambda_u(w) \Vdash \alpha \rightarrow \beta$.

Assume now that $w : \alpha \rightarrow \beta \in \Delta_u$. By Lemma 4.2.45, $d_u(w) : \alpha \rightarrow \beta \in \Delta_{\text{lst}(C_u)}$, so by saturation there is a w' such that $d_u(w) \preceq w' \in \Omega_{\text{lst}(C_u)}$, $w' : \alpha \in \Gamma_{\text{lst}(C_u)}$ and $w' : \beta \in \Delta_{\text{lst}(C_u)}$. By the inductive hypothesis, $\mathcal{M}_{\mathcal{R}}, \lambda_{\text{lst}(C_u)}(w') \Vdash \alpha$ and $\mathcal{M}_{\mathcal{R}}, \lambda_{\text{lst}(C_u)}(w') \not\Vdash \beta$. Since

$$\lambda_u(w) = t_{C_u, d_u(w)} \leq_{\mathcal{R}} t_{C_u, w'} = t_{C_{\text{lst}(C_u)}, d_{\text{lst}(C_u)}(w')} = \lambda_{\text{lst}(C_u)}(w'),$$

²Recall that $\leq_{\mathcal{R}}$ is the intuitionistic order of the model $\mathcal{M}_{\mathcal{R}}$.

we have shown that $\mathcal{M}_{\mathcal{R}}, \lambda_u(w) \not\models \alpha \rightarrow \beta$. ■

As the formula φ , the proof-search \mathcal{G} and the proof-search tree \mathcal{T} were arbitrary, we have established the following:

4.2.60. COROLLARY. *Let φ be a valid formula. Then, for any proof-search guide \mathcal{G} , every proof-search tree for φ guided by \mathcal{G} is an iLTL_∞ proof of φ .*

Proof. Suppose that there is an iLTL_∞ proof-search tree \mathcal{T} for φ guided by \mathcal{G} which is not an iLTL_∞ proof of φ . By Proposition 4.2.44, there is a branch of \mathcal{T} which is an iLTL_∞ refutation of φ , contradicting, by Proposition 4.2.59, the validity of φ . ■

As an immediate consequence of Proposition 4.2.30 and Corollary 4.2.60, we get:

4.2.61. COROLLARY (Completeness of iLTL_∞). *For every formula φ , if $\text{iLTL} \models \varphi$, then $\text{iLTL}_\infty \vdash \varphi$.*

Moreover, we can now fully justify the use of the term ‘refutation’:

4.2.62. COROLLARY. *For every formula φ , there exists an iLTL refutation of φ if, and only if, φ is not valid.*

Proof. The left-to-right implication is Proposition 4.2.59. For the converse, assume that φ is not valid and let \mathcal{T} be any proof-search tree for φ (at least one exists by Proposition 4.2.30). By Proposition 4.2.19, \mathcal{T} is not an iLTL_∞ proof of φ , whence by Proposition 4.2.44 there is a branch of \mathcal{T} which is a refutation of φ . ■

4.3 Regular ill-founded proofs

As we did for CTL^* , before introducing the cyclic system for iLTL we discuss *regular* ill-founded proofs, which will serve as an intermediary between the cyclic system and the ill-founded one when proving the soundness of the former.

Even though iLTL_∞ proof-search trees are in general infinite, by Proposition 4.2.38 they only contain finitely many pairwise different sequents. It stands to reason, then, that proofs obtained via iLTL_∞ proof-search should admit at least a partially finitary presentation. We now formalise this idea. The question of whether the correctness condition imposed on infinite branches of iLTL_∞ proofs can also be finitised will be settled in the affirmative in Section 4.4, where we present the fully finitary, cyclic version of iLTL_∞ .

Infinite iLTL_∞ proofs obtained via proof-search may be folded down into finite trees with back-edges. To make this formal we introduce the system iLTL_{reg} as a partially finitary version of iLTL_∞ , in the sense that iLTL_{reg} derivations are finite objects but the correctness condition imposed on proofs remains infinitary.

4.3.1. DEFINITION (iLTL_{reg} derivation). An iLTL_{reg} derivation of a formula φ is a labelled tree with back-edges whose vertices are labelled according to the rules in Figures 4.2 to 4.4 and whose root has label $\emptyset \dashv \emptyset \Rightarrow w : \varphi$ for some $w \in \text{Worlds}$. \dashv

4.3.2. DEFINITION (iLTL_{reg} proof). An iLTL_{reg} proof of a formula φ is an iLTL_{reg} derivation \mathcal{T} of φ such that every non-repeat leaf of \mathcal{T} is axiomatic and any infinite path through \mathcal{T}° contains a good trace. \dashv

4.3.3. OBSERVATION. It follows from the definition of \mathbf{O} -traces that for every repeat l of an iLTL_{reg} proof \mathcal{T} there is a modal vertex in $[c_l, l)_{\mathcal{T}}$.

4.3.4. OBSERVATION. The second condition in Definition 4.3.2 is equivalent to requiring that any infinite rooted path through \mathcal{T}° contain a good trace (recall Observation 4.2.14).

Soundness of iLTL_{reg} is an immediate consequence of our work on iLTL_∞ :

4.3.5. PROPOSITION (Soundness of iLTL_{reg}). *For every formula φ , if $\text{iLTL}_{\text{reg}} \vdash \varphi$, then $\text{iLTL} \models \varphi$.*

Proof. Let \mathcal{T} be an iLTL_{reg} proof of φ . We claim that \mathcal{T}^ω is an iLTL_∞ proof of φ . Note that \mathcal{T}^ω is an iLTL_∞ derivation because, since every repeat $l \in \text{Rep}_{\mathcal{T}}$ is labelled by a non-empty sequent S_l , we may apply Lwk or Rwk preservingly to S_l to obtain S_l again, which also labels c_l .

For every leaf $l \in T^\omega$ there is a non-repeat leaf $l' \in T$ such that l and l' have identical labels, so l is axiomatic. And for every infinite branch $\pi = (u_n)_{n < \omega}$ of \mathcal{T}^ω there is an infinite path $\pi' = (u'_n)_{n < \omega}$ through \mathcal{T}° such that, for every $n < \omega$, u_n and u'_n are labelled by the same instance of the same iLTL_∞ rule, whence π contains a good trace. Therefore, \mathcal{T}^ω is an iLTL_∞ proof of φ and thus $\text{iLTL} \models \varphi$ by Proposition 4.2.19. \blacksquare

4.3.6. REMARK. In the proof of Proposition 4.3.5 we relied on the contraction implicitly built into the rules of iLTL_∞ in order to show that \mathcal{T}^ω itself is an iLTL_∞ derivation. This is clearly unnecessary because we could instead identify each repeat and its companion, but it simplifies the argument by allowing us to work directly with \mathcal{T}^ω .

As stated in Remark 4.2.31, $iLTL_\infty$ proof-search is not as deterministic as CTL_∞^* proof-search is, because the rule applied at a vertex u depends not only on the sequent labelling u but also on the U- and R-formulas pending at u . For this reason, we postpone the proof of the completeness of $iLTL_{reg}$ until Section 4.4, where it will follow from the soundness and completeness proofs of the cyclic version of $iLTL_\infty$.

Even though $iLTL_{reg}$ derivations are finite objects, the second condition in Definition 4.3.2 is infinitary, and thus the cyclic system $iLTL_{reg}$ is not entirely satisfactory. In the next section we obtain a fully finitary cyclic system by annotating sequents to keep track of fixpoint unfoldings.

4.4 The cyclic system $iLTL_\circ$

We introduce a finitary, cyclic version of the system $iLTL_\infty$. As we did for CTL_∞^* , we enrich formulas with annotations similar to the ones introduced by Jungteerapanich [76] and Stirling [144] for the modal μ -calculus. This allows us to keep track of fixpoint unfoldings and detect good traces on cycles. This time, however, the annotation mechanism is able to finitarily detect all good traces because in $iLTL$ there are no difficulties analogous to the ones posed by existential traces on CTL_∞^* proofs. In particular, we obtain a fully finitary calculus.

We fix a countably infinite set $\mathbf{N} = \{x, y, z, \dots\}$ of *names*, partitioned into two disjoint countably infinite sets \mathbf{N}_U and \mathbf{N}_R of, respectively, *U-names* and *R-names*. An *annotation* is either the empty string ε or a single name. Annotations are denoted by a, b, \dots . We identify a non-empty annotation with the unique name it contains.

An *annotated (labelled) formula* is a triple (w, φ, a) , henceforth written $w : \varphi^a$, where $w : \varphi$ is a labelled formula and a an annotation. We identify each unannotated formula $w : \varphi$ with the formula $w : \varphi^\varepsilon$ annotated by the empty annotation. A name x *occurs* in a set of annotated formulas Γ if there is a formula $w : \varphi$ such that $w : \varphi^x \in \Gamma$.

An *annotated (labelled iLTL) sequent* is a quadruple $(\Omega, \Theta, \Gamma, \Delta)$, henceforth written $\Omega ; \Theta \dashv \Gamma \Rightarrow \Delta$, where Γ and Δ are finite sets of annotated labelled formulas, Ω is a relational control, and Θ , called the *nominal control* of the sequent, is a linear ordering of the names occurring in $\Gamma \cup \Delta$.

The *base* of an annotated formula $w : \varphi^a$ is $\mathbf{b}(w : \varphi^a) := w : \varphi$. If Γ is a set of annotated formulas, we let $\mathbf{b}(\Gamma) := \{\mathbf{b}(w : \varphi^a) \mid w : \varphi^a \in \Gamma\}$. The *base* of an annotated sequent $\Omega ; \Theta \dashv \Gamma \Rightarrow \Delta$ is $\mathbf{b}(\Omega ; \Theta \dashv \Gamma \Rightarrow \Delta) := \Omega \dashv \mathbf{b}(\Gamma) \Rightarrow \mathbf{b}(\Delta)$. Finally, if Γ is a set of sequents we let $\mathbf{b}(\Gamma) := \{\mathbf{b}(S) \mid S \in \Gamma\}$.

Given a finite sequence of names Θ , we define the strict linear order \prec_Θ on $\{\varepsilon\} \cup \{x \in \mathbf{N} \mid x \in \Theta\}$ by letting $a \prec_\Theta b$ if, and only if, either $a \neq \varepsilon = b$, or both

$$\begin{array}{c}
\text{ax} \frac{}{\Omega ; \Theta \dashv \Gamma, w : p \Rightarrow w : p, \Delta} \quad \text{ax}_{\perp} \frac{}{\Omega ; \Theta \dashv \Gamma, w : \perp \Rightarrow \Delta} \\
\text{L}\wedge \frac{\Omega ; \Theta \dashv \Gamma, w : \varphi, w : \psi \Rightarrow \Delta}{\Omega ; \Theta \dashv \Gamma, w : \varphi \wedge \psi \Rightarrow \Delta} \quad \text{R}\vee \frac{\Omega ; \Theta \dashv \Gamma \Rightarrow w : \varphi, w : \psi, \Delta}{\Omega ; \Theta \dashv \Gamma \Rightarrow w : \varphi \vee \psi, \Delta} \\
\text{L}\vee \frac{\Omega ; \Theta \dashv \Gamma, w : \varphi \Rightarrow \Delta \quad \Omega ; \Theta \dashv \Gamma, w : \psi \Rightarrow \Delta}{\Omega ; \Theta \dashv \Gamma, w : \varphi \vee \psi \Rightarrow \Delta} \\
\text{R}\wedge \frac{\Omega ; \Theta \dashv \Gamma \Rightarrow w : \varphi, \Delta \quad \Omega ; \Theta \dashv \Gamma \Rightarrow w : \psi, \Delta}{\Omega ; \Theta \dashv \Gamma \Rightarrow w : \varphi \wedge \psi, \Delta} \\
\text{L}\rightarrow \frac{\Omega ; \Theta \dashv \Gamma, w : \varphi \rightarrow \psi \Rightarrow w : \varphi, \Delta \quad \Omega ; \Theta \dashv \Gamma, w : \psi \Rightarrow \Delta}{\Omega ; \Theta \dashv \Gamma, w : \varphi \rightarrow \psi \Rightarrow \Delta} \\
\text{R}\rightarrow \frac{\Omega, w \preceq w' ; \Theta \dashv \Gamma, w' : \varphi \Rightarrow w' : \psi, \Delta}{\Omega ; \Theta \dashv \Gamma \Rightarrow w : \varphi \rightarrow \psi, \Delta} \quad w' \notin \Omega \\
\text{X} \frac{\Omega ; \Theta' \dashv \Lambda \Rightarrow \Xi}{\Omega ; \Theta \dashv \Pi, \text{X}\Lambda \Rightarrow \text{X}\Xi, \Sigma}
\end{array}$$

FIGURE 4.12: Non-fixpoint, logical rules of the system $iLTL_{\circ}$.

a and b are non-empty strings and the name in a occurs in Θ strictly before the name in b .

The rules of the cyclic system $iLTL_{\circ}$, defined below, are given in Figures 4.12 to 4.15. In rules X , LU_0 , RR_0 , Lwk , Rwk , Lthin , Rthin , Ldel and Rdel , we denote by Θ' the result after removing from Θ all names not occurring in the premise with control Θ' . And in rules LU_1 and RR_1 , we define Θx as the concatenation $\Theta^{\frown}(x)$, if $x \notin \Theta$, and otherwise $\Theta x := \Theta$.

4.4.1. DEFINITION ($iLTL_{\circ}$ derivation). An $iLTL_{\circ}$ derivation of a formula φ is a finite labelled tree with back-edges \mathcal{T} built according to the rules in Figures 4.12 to 4.15 and whose root has label \emptyset ; $\emptyset \dashv \emptyset \Rightarrow w : \varphi^{\varepsilon}$ for some $w \in \text{Worlds}$. \dashv

Formula *traces*, as well as *principal* and *active* relations and formulas, follow the definition from the system $iLTL_{\infty}$ in the cases of the rules in Figures 4.12 to 4.14, and also of wk_{\preceq} , Lwk and Rwk .

$$\begin{array}{c}
\text{LU}_0 \frac{\Omega ; \Theta' \dashv \Gamma, w : \psi \Rightarrow \Delta \quad \Omega ; \Theta \dashv \Gamma, w : \varphi, w : X(\varphi \mathbf{U} \psi)^a \Rightarrow \Delta}{\Omega ; \Theta \dashv \Gamma, w : (\varphi \mathbf{U} \psi)^a \Rightarrow \Delta} \\
\text{LU}_1 \frac{\Omega ; \Theta \dashv \Gamma, w : \psi \Rightarrow \Delta \quad \Omega ; \Theta x \dashv \Gamma, w : \varphi, w : X(\varphi \mathbf{U} \psi)^x \Rightarrow \Delta}{\Omega ; \Theta \dashv \Gamma, w : \varphi \mathbf{U} \psi \Rightarrow \Delta} \quad x \in \mathbf{N}_U \setminus \Theta \\
\text{RU} \frac{\Omega ; \Theta \dashv \Gamma \Rightarrow w : \varphi, w : \psi, \Delta \quad \Omega ; \Theta \dashv \Gamma \Rightarrow w : \psi, w : X(\varphi \mathbf{U} \psi), \Delta}{\Omega ; \Theta \dashv \Gamma \Rightarrow w : \varphi \mathbf{U} \psi, \Delta} \\
\text{LR} \frac{\Omega ; \Theta \dashv \Gamma, w : \varphi, w : \psi \Rightarrow \Delta \quad \Omega ; \Theta \dashv \Gamma, w : \psi, w : X(\varphi \mathbf{R} \psi) \Rightarrow \Delta}{\Omega ; \Theta \dashv \Gamma, w : \varphi \mathbf{R} \psi \Rightarrow \Delta} \\
\text{RR}_0 \frac{\Omega ; \Theta' \dashv \Gamma \Rightarrow w : \psi, \Delta \quad \Omega ; \Theta \dashv \Gamma \Rightarrow w : \varphi, w : X(\varphi \mathbf{R} \psi)^a, \Delta}{\Omega ; \Theta \dashv \Gamma \Rightarrow w : (\varphi \mathbf{R} \psi)^a, \Delta} \\
\text{RR}_1 \frac{\Omega ; \Theta \dashv \Gamma \Rightarrow w : \psi, \Delta \quad \Omega ; \Theta x \dashv \Gamma \Rightarrow w : \varphi, w : X(\varphi \mathbf{R} \psi)^x, \Delta}{\Omega ; \Theta \dashv \Gamma \Rightarrow w : \varphi \mathbf{R} \psi, \Delta} \quad x \in \mathbf{N}_R \setminus \Theta
\end{array}$$

FIGURE 4.13: Fixpoint rules of the system iLTL_o.

In rules LU_0 and RR_0 with $a \neq \varepsilon$, the unique name in a is *principal* in the conclusion and *active* in the premise corresponding to the unfolding. Analogously, in rules LU_1 and RR_1 the name x is *active* in the unfolding premise.

In rules Lm and Rm with $a \neq \varepsilon$, the unique name in a is *principal* in the conclusion and *active* in the premise.

For Lthin , Rthin , Ldel and Rdel , principal and active relations and formulas are defined as expected: the distinguished relations and formulas in the conclusion are *principal*, the distinguished relations and formulas in the premise are *active*, every relation in Ω is a *side* relation, and every formula in $\Gamma \cup \Delta$ is a *side* formula.

In the *thinning* rules Lthin and Rthin , the name x is *principal* in the conclusion and *active* in the premise, and if $a \neq \varepsilon$ then the unique name in a is *principal* in the conclusion. In the *deletion* rules Ldel and Rdel , the name x is *principal* in the conclusion.

Traces are also defined as expected for Lthin , Rthin , Ldel and Rdel , with the (possible) exception that in Lthin and Rthin we let $w : \varphi^a \triangleright w : \varphi^x$. Observe, then, that even though the sequent $\Omega ; \Theta' \dashv \Gamma, w : \varphi^x \Rightarrow \Delta$ may be obtained from the sequent $\Omega ; \Theta \dashv \Gamma, w : \varphi^x, w : \varphi^a \Rightarrow \Delta$ via an application of either Lwk or Lthin , in

$$\begin{array}{c}
\text{refl} \frac{\Omega, w \preceq w ; \Theta \dashv \Gamma \Rightarrow \Delta}{\Omega ; \Theta \dashv \Gamma \Rightarrow \Delta} \quad \text{trans} \frac{\Omega, w \preceq w', w' \preceq w'', w \preceq w'' ; \Theta \dashv \Gamma \Rightarrow \Delta}{\Omega, w \preceq w', w' \preceq w'' ; \Theta \dashv \Gamma \Rightarrow \Delta} \\
\\
\text{Lm} \frac{\Omega, w \preceq w' ; \Theta \dashv \Gamma, w : \varphi^a, w' : \varphi^a \Rightarrow \Delta}{\Omega, w \preceq w' ; \Theta \dashv \Gamma, w : \varphi^a \Rightarrow \Delta} \\
\\
\text{Rm} \frac{\Omega, w' \preceq w ; \Theta \dashv \Gamma \Rightarrow w : \varphi^a, w' : \varphi^a, \Delta}{\Omega, w' \preceq w ; \Theta \dashv \Gamma \Rightarrow w : \varphi^a, \Delta}
\end{array}$$

FIGURE 4.14: Structural, non-weakening rules of the system $i\text{LTL}_\circ$.

the former case we have $w : \varphi^a \not\preceq w : \varphi^x$ but in the latter $w : \varphi^a \triangleright w : \varphi^x$; and similarly for **Rthin**.

The merging of traces in the thinning rules, on which our proofs of the soundness and completeness of $i\text{LTL}_\circ$ rely, corresponds to the fact that, if the annotations are dropped, then $w : \varphi^x$ and $w : \varphi^a$ become the same formula and their traces merge. We could dispense with the thinning rules by appropriately³ incorporating their trace merging into the weakening rules. Nevertheless, for the sake of clarity we keep thinning and weakening separate.

When unfolding an unannotated **U**-formula φ on the left, we may mark the unfolding by annotating the active formula $X\varphi$ (rule **LU**₁), or we may choose to leave $X\varphi$ unannotated (**LU**₀). If annotated, subsequent unfoldings may only be performed by rule **LU**₀ and thus the annotation is preserved unless we remove it by applying **Lwk**, **Lthin** or **Ldel**. Analogously for unfoldings of **R**-formulas on the right. This degree of freedom corresponds to the fact that, in order to identify the $i\text{LTL}_\circ$ derivations that should count as proofs, we shall look for infinite **U**-traces on the left and infinite **R**-traces on the right, and thus it does not matter when we begin to annotate unfoldings on such traces as long as we do so eventually (recall Observation 4.2.14).

A vertex u of an $i\text{LTL}_\circ$ derivation is *modal* if rule **X** is applied at u . And u is *thinning* if rule **Lthin** or **Rthin** is applied at u .

Let \mathcal{T} be an $i\text{LTL}_\circ$ derivation, $\pi = (u_n)_{n < N \leq \omega}$ a finite or infinite path through \mathcal{T} , and $\tau = ((w_n : \varphi_n^{a_n}, b_n))_{n < N}$ a trace on π . A name x is *fixed* on τ if $a_n = x$ for every $n < N$. Analogously, x is *fixed* on π if x occurs in the nominal control of every sequent labelling a vertex in π . If π and τ are infinite, we say that x is *eventually fixed* on π (τ) if there is a tail π' of π (respectively, a tail τ' of τ) such that x is fixed on π' (respectively, τ').

³Care would be needed to handle preserving instances of weakening.

$$\begin{array}{c}
\text{wk}_{\preceq} \frac{\Omega ; \Theta \vdash \Gamma \Rightarrow \Delta}{\Omega, w \preceq w' ; \Theta \vdash \Gamma \Rightarrow \Delta} \\
\\
\text{Lwk} \frac{\Omega ; \Theta' \vdash \Gamma \Rightarrow \Delta}{\Omega ; \Theta \vdash \Gamma, w : \varphi^a \Rightarrow \Delta} \quad \text{Rwk} \frac{\Omega ; \Theta' \vdash \Gamma \Rightarrow \Delta}{\Omega ; \Theta \vdash \Gamma \Rightarrow w : \varphi^a, \Delta} \\
\\
\text{Lthin} \frac{\Omega ; \Theta' \vdash \Gamma, w : \varphi^x \Rightarrow \Delta}{\Omega ; \Theta \vdash \Gamma, w : \varphi^x, w : \varphi^a \Rightarrow \Delta} x \prec_{\Theta} a \\
\\
\text{Rthin} \frac{\Omega ; \Theta' \vdash \Gamma \Rightarrow w : \varphi^x, \Delta}{\Omega ; \Theta \vdash \Gamma \Rightarrow w : \varphi^x, w : \varphi^a, \Delta} x \prec_{\Theta} a \\
\\
\text{Ldel} \frac{\Omega ; \Theta' \vdash \Gamma, w : \varphi \Rightarrow \Delta}{\Omega ; \Theta \vdash \Gamma, w : \varphi^x \Rightarrow \Delta} \quad \text{Rdel} \frac{\Omega ; \Theta' \vdash \Gamma \Rightarrow w : \varphi, \Delta}{\Omega ; \Theta \vdash \Gamma \Rightarrow w : \varphi^x, \Delta}
\end{array}$$

FIGURE 4.15: Structural, weakening rules of the system iLTL_0 .

4.4.2. DEFINITION (Successful repeat). Let \mathcal{T} be an iLTL_0 derivation. A repeat $l \in \text{Rep}_{\mathcal{T}}$ is *successful* if the following hold:

- (i) there is a modal vertex in $[c_l, l]_{\mathcal{T}}$;
- (ii) there is a name fixed on $[c_l, l]_{\mathcal{T}}$.

–

4.4.3. DEFINITION (iLTL_0 proof). An *iLTL₀ proof* of a formula φ is an iLTL_0 derivation \mathcal{T} of φ such that every non-axiomatic leaf of \mathcal{T} is a successful repeat. –

Note that the system iLTL_0 is finitary, for every iLTL_0 derivation is finite and finitely many checks suffice to determine whether a given iLTL_0 derivation is a proof.

Every successful repeat l in an iLTL_0 derivation \mathcal{T} has an associated *invariant*, denoted by $\text{inv}(l)$, defined as the shortest sequence of names wx such that wx is a prefix of every nominal control in $[c_l, l]_{\mathcal{T}}$. The existence of invariants follows immediately from Definition 4.4.2(ii) and the fact that, when reading iLTL_0 rules bottom-up, new names are always appended to the right end of nominal controls. The invariant map inv induces the following (reflexive) quasi-order \preceq on the successful repeats of \mathcal{T} : $l \preceq l'$ if, and only if, $\text{inv}(l) \sqsubseteq \text{inv}(l')$.

4.4.1 Soundness of iLTL_\circ

We prove soundness of iLTL_\circ by dropping the annotations from a given iLTL_\circ proof and showing that the resulting tree is an iLTL_{reg} proof.

As we did for CTL_\circ^* , we first prove some results relating traces and names in iLTL_\circ derivations.

4.4.4. LEMMA. *Let \mathcal{T} be an iLTL_\circ derivation, and π a finite path through \mathcal{T}° . If a name x is fixed on π , then there is a trace τ on π such that x is fixed on τ .*

Proof. We prove only the case $x \in \mathbf{N}_\mathbf{U}$, for the case $x \in \mathbf{N}_\mathbf{R}$ is analogous.

If π is empty or has only one vertex there is nothing to prove, so assume otherwise. Let then $\pi = (u_i)_{i \leq n}$, with $n > 0$, and let $S_i = \Omega_i ; \Theta_i \dashv \Gamma_i \Rightarrow \Delta_i$ be the label of u_i for every $i \leq n$. Since x is fixed on π and x is a \mathbf{U} -name, by the definition of nominal controls there is a formula $w : \alpha^x \in \Gamma_n$. It suffices to find a formula $w' : \beta^x \in \Gamma_{n-1}$ such that $(w' : \beta^x, 0) \triangleright (w : \alpha^x, 0)$, for then repeating the same argument finitely many times yields the desired trace τ .

Towards a contradiction, suppose that there is no such formula $w' : \beta^x \in \Gamma_{n-1}$. Let $w'' : \gamma^a \in \Gamma_{n-1}$ be such that $w'' : \gamma^a \triangleright w : \alpha^x$. Note that such a formula $w'' : \gamma^a$ exists by the definition of \triangleright . By assumption, $a \neq x$, so rule Lthin or LU_1 is applied at u_{n-1} with $w'' : \gamma^a$ principal and $w : \alpha^x$ active. The former is impossible because then we would have $\Gamma_{n-1} \ni w : \alpha^x \triangleright w : \alpha^x$. So LU_1 is applied at u_{n-1} and $a = \varepsilon$, whence $x \notin \Theta_{n-1}$. This contradicts the fact that x is fixed on π . ■

4.4.5. LEMMA. *Let \mathcal{T} be an iLTL_\circ derivation, π an infinite path through \mathcal{T}° , and $x \in \mathbf{N}$ any name. If there are infinitely many $n < \omega$ for which there exists a trace on $\pi_{\leq n}$ on which x is fixed, then there is an infinite trace on π on which x is fixed.*

Proof. We prove only the case $x \in \mathbf{N}_\mathbf{U}$, for the case $x \in \mathbf{N}_\mathbf{R}$ is analogous.

For every formula $w : \varphi^x$ in the label of $\pi(0)$ and every $n < \omega$, let $F(w : \varphi^x, n)$ be the collection of all traces on $\pi_{\leq n}$ starting from $(w : \varphi^x, 0)$ and on which x is fixed. As the label of $\pi(0)$ is finite, there is a formula $w : \varphi^x$ in the label of $\pi(0)$ such that $F(w : \varphi^x, n) \neq \emptyset$ for every $n < \omega$. Let $F := \bigcup_{n < \omega} F(w : \varphi^x, n)$. Then, (F, \sqsubset) is a tree of height ω whose levels are all finite. By König's lemma, (F, \sqsubset) has an infinite branch, whence there exists an infinite trace on π starting from $(w : \varphi^x, 0)$ and on which x is fixed. ■

4.4.6. PROPOSITION. *Let \mathcal{T} be an iLTL_\circ proof, and π an infinite path over \mathcal{T}° . There is a repeat $l \in \text{Rep}_\mathcal{T} \cap \text{Inf}(\pi)$ and an infinite trace on π on which the last name in $\text{inv}(l)$ is eventually fixed.*

Proof. Let $\pi' := (u_i)_{i < \omega}$ be a tail of π such that each u_i occurs infinitely often on π and $u_0 \in \text{Rep}_{\mathcal{T}}$. We may then write

$$\pi' = l_0 \widehat{\ } [c_0, l_1]_{\mathcal{T}} \widehat{\ } [c_1, l_2]_{\mathcal{T}} \widehat{\ } \cdots,$$

where each l_i is a (successful) repeat with companion c_i . Note that $l_i \triangleleft l_{i+1}$ for every $i < \omega$. By Proposition 1.2.11, there is some $k \geq 0$ such that $l_k \preceq l_i$ for every $i < \omega$. Since l_k is successful, $\text{inv}(l_k) = wx$ for some $w \in \mathbf{N}^{<\omega}$ and some name x . By Proposition 1.2.12, wx is a prefix of each nominal control on π' , and thus by Lemma 4.4.4 for every $j < \omega$ there is a trace on (u_0, \dots, u_j) on which x is fixed. Therefore, by Lemma 4.4.5 there exists an infinite trace on π' on which x is fixed. \blacksquare

4.4.7. PROPOSITION (Soundness of iLTL_o). *For every formula φ , if $\text{iLTL}_o \vdash \varphi$, then $\text{iLTL} \models \varphi$.*

Proof. Let $\mathcal{T} = (T, <_T, \lambda_T, l \mapsto c_l)$ be an iLTL_o proof of φ . Let $\mathcal{T}_r := (T, <_T, \lambda_r, l \mapsto c_l)$ be given by setting $\lambda_r(u) := \mathbf{b}(\lambda_T(u))$ for every $u \in T$. That is to say, \mathcal{T}_r results from replacing every sequent S labelling a vertex of \mathcal{T} by its base $\mathbf{b}(S)$. We claim that \mathcal{T}_r is an iLTL_{reg} proof of φ .

Let us first show that \mathcal{T}_r is an iLTL_{reg} derivation of φ . It suffices to see that instances of iLTL_o rules in \mathcal{T} turn into instances of iLTL_{reg} rules in \mathcal{T}_r . Let $u \in T_r$ be a non-final vertex of \mathcal{T}_r . Then, u is non-final in \mathcal{T} as well. Let R be the iLTL_o rule applied at u in \mathcal{T} . If R is among L \wedge , R \wedge , L \vee , R \vee , L \rightarrow , R \rightarrow , X, RU, LR, refl, trans, Lm, Rm, wk \preceq , then R is applied at u in \mathcal{T}_r as well. If $R \in \{\text{LU}_0, \text{LU}_1\}$ ($R \in \{\text{RR}_0, \text{RR}_1\}$), then rule LU (respectively, RR) is applied at u in \mathcal{T}_r . Suppose that $R = \text{Lwk}$. Let $w : \varphi^a$ be the unique principal formula at u in \mathcal{T} . If there is a side formula $w' : \psi^b$ in the antecedent of the sequent labelling u in \mathcal{T} such that $\mathbf{b}(w' : \psi^b) = \mathbf{b}(w : \varphi^a)$, then Lwk is applied preservingly at u in \mathcal{T}_r . Otherwise, Lwk is applied discardingly at u in \mathcal{T}_r . The case $R = \text{Rwk}$ is analogous. Finally, instances of Lthin and Ldel (Rthin and Rdel) in \mathcal{T} become preserving instances of Lwk (respectively, Rwk) in \mathcal{T}_r . This establishes that \mathcal{T}_r is an iLTL_{reg} derivation. And, since the roots of \mathcal{T} and \mathcal{T}_r have identical labels, \mathcal{T}_r is an iLTL_{reg} derivation of φ .

Every leaf of \mathcal{T}_r which is not a repeat is clearly axiomatic. It only remains to see that every infinite path through \mathcal{T}_r° contains a good trace. Let $\pi_r = (u_n)_{n < \omega}$ be an infinite path through \mathcal{T}_r° . Then, $\pi := (u_n)_{n < \omega} = \pi$ is an infinite path through \mathcal{T}° . By Proposition 4.4.6, there is a repeat $l \in \text{Rep}_{\mathcal{T}} \cap \text{Inf}(\pi)$ and an infinite trace $\tau = ((w_n : \varphi^{a_n}, b_n)_{n < \omega})$ on π on which the last name x in $\text{inv}(l)$ is eventually fixed. By the definition of traces in the thinning rules, then, $\tau_r := ((w_n : \varphi, b_n)_{n < \omega})$ is an

infinite trace on π_r . If $x \in \mathbf{N}_U$, then τ_r is a left U-trace. And if $x \in \mathbf{N}_R$, then τ_r is a right R-trace. We have thus found a good trace τ_r on π_r .

Therefore, \mathcal{T}_r is an iLTL_{reg} proof of φ , so φ is valid by Proposition 4.3.5. \blacksquare

Observe that in the proof of Proposition 4.4.7 we have furthermore established:

4.4.8. PROPOSITION. *For every formula φ , if $\text{iLTL}_\circ \vdash \varphi$, then $\text{iLTL}_{\text{reg}} \vdash \varphi$.*

4.4.9. COROLLARY. *If iLTL_\circ is complete, then so is iLTL_{reg} .*

4.4.10. REMARK. Like in the proof of the soundness of iLTL_{reg} , in the proof of Proposition 4.4.7 we relied on the contraction implicitly built into the rules of iLTL_{reg} (see Remark 4.3.6 above). This is not necessary, but it simplifies the argument because it allows us to keep all the vertices of \mathcal{T} in \mathcal{T}_r .

4.4.2 Completeness of iLTL_\circ

We establish the completeness of iLTL_\circ by annotating iLTL_∞ proofs obtained via proof-search and folding them down to trees with back-edges.

For the remainder of this section, fix an arbitrary well-order $\leq_{\mathbf{N}}$ on \mathbf{N} , a valid formula φ , and an iLTL_∞ proof \mathcal{T} of φ given by Proposition 4.2.30 and Corollary 4.2.60. In particular, \mathcal{T} is an iLTL_∞ proof-search tree for φ , say guided by $\mathcal{G} = (\leq_W, \leq_{\text{iLTL}})$.

Let Γ be a set of annotated formulas and $n < \omega$. A non-annotated formula $w : \alpha$ is *n-annotated* in Γ if there are pairwise different annotations a_1, \dots, a_n such that $w : \alpha^{a_i} \in \Gamma$ for every $1 \leq i \leq n$. A non-annotated formula is *twice-annotated* in Γ if it is 2-annotated in Γ .

We begin by inductively building a (possibly infinite) tree $\tilde{\mathcal{T}}$ according to the rules of iLTL_\circ , together with a function $f : \tilde{\mathcal{T}} \rightarrow \mathcal{T}$ such that the following hold for all $u, v \in \tilde{\mathcal{T}}$:

- (i) if $u <_{\tilde{\mathcal{T}}}^0 v$ and u is not thinning, then $f(u) <_{\mathcal{T}}^0 f(v)$;
- (ii) if $u <_{\tilde{\mathcal{T}}}^0 v$ and u is thinning, then $f(u) = f(v)$;
- (iii) if $f(u) <_{\mathcal{T}}^0 f(v)$, then $u <_{\tilde{\mathcal{T}}}^0 v$;
- (iv) if u has label $\Omega ; \Theta \dashv \Gamma \Rightarrow \Delta$, then $f(u)$ is labelled by $\Omega \dashv \mathbf{b}(\Gamma) \Rightarrow \mathbf{b}(\Delta)$.

For the base case, let the root of $\tilde{\mathcal{T}}$ be labelled by the same sequent as the one labelling the root of \mathcal{T} , and map it via f to the root of \mathcal{T} . For the inductive case, assume that a vertex $u \in \tilde{\mathcal{T}}$ has been defined, say with label $\Omega ; \Theta \dashv \Gamma \Rightarrow \Delta$, and that $f(u)$ is defined as well. We proceed as follows:

1. If a formula is twice-annotated in Γ , then we apply rule **Lthin** at u with principal formulas $w : \alpha^a$ and $w : \alpha^b$, where (w, α, a, b) is the $(\leq_W, \leq_{\text{iLTL}}, \prec_\Theta, \prec_\Theta)$ -least tuple such that $a \neq b$ and $w : \alpha^a, w : \alpha^b \in \Gamma$. We map the unique immediate successor of u via f to $f(u)$.
2. Else, if some formula is twice-annotated in Δ , then we apply rule **Rthin** at u with principal formulas $w : \alpha^a$ and $w : \alpha^b$, where (w, α, a, b) is the $(\leq_W, \leq_{\text{iLTL}}, \prec_\Theta, \prec_\Theta)$ -least tuple such that $a \neq b$ and $w : \alpha^a, w : \alpha^b \in \Delta$. We map the unique immediate successor of u via f to $f(u)$.
3. Otherwise, by the inductive hypothesis the base function \mathbf{b} induces bijections $\mathbf{b} : \Gamma \rightarrow \mathbf{b}(\Gamma)$ and $\mathbf{b} : \Delta \rightarrow \mathbf{b}(\Delta)$, and we distinguish cases according to the rule **R** applied at $f(u)$:
 - a) Case $\mathbf{R} = \mathbf{LU}$, say with principal formula $w : \alpha$. Let a be such that $b^{-1}(w : \alpha) = w : \alpha^a$. If $a = \varepsilon$, then we apply rule \mathbf{LU}_1 at u with principal formula $w : \alpha$ and active name the \leq_N -least **U**-name not occurring in Θ . Otherwise, we apply rule \mathbf{LU}_0 at u with principal formula $w : \alpha^a$.
 - b) Case $\mathbf{R} = \mathbf{RR}$, say with principal formula $w : \alpha$. Let a be such that $b^{-1}(w : \alpha) = w : \alpha^a$. If $a = \varepsilon$, then we apply rule \mathbf{RR}_1 at u with principal formula $w : \alpha$ and active name the \leq_N -least **R**-name not occurring in Θ . Otherwise, we apply rule \mathbf{RR}_0 at u with principal formula $w : \alpha^a$.
 - c) In any other case we apply rule **R** at u with principal and active formulas the b^{-1} -images of principal and active formulas, respectively, at $f(u)$, and principal and active relations the principal and active relations, respectively, at $f(u)$.

In all three cases (a), (b) and (c), we map each immediate successor of u via f to the corresponding immediate successor of $f(u)$.

Since the thinning rules are only applied in $\tilde{\mathcal{T}}$ to sequents containing twice-annotated formulas, the following is an immediate consequence of (i), (ii) and (iv):

4.4.11. LEMMA. *For every finite or infinite branch $\tilde{\pi} = (u_n)_{n < N \leq \omega}$ of $\tilde{\mathcal{T}}$, there is an $N' \leq N$ and a strictly increasing sequence of natural numbers $(n_i)_{i < N'}$ such that $\pi := (f(u_{n_i}))_{i < N'}$ is a branch of \mathcal{T} . Moreover, π is infinite if, and only if, so is $\tilde{\pi}$.*

Abusing notation, for every infinite branch $\tilde{\pi}$ of $\tilde{\mathcal{T}}$ we denote by $f(\tilde{\pi})$ the corresponding infinite branch π of \mathcal{T} given by Lemma 4.4.11.

As a consequence of the definition of traces in the thinning rules, we additionally have:

4.4.12. LEMMA. *For every infinite branch $\tilde{\pi}$ of $\tilde{\mathcal{T}}$ and every infinite trace $\tau = ((w_n : \varphi_n, b_n))_{n < \omega}$ on $f(\tilde{\pi})$, there is a trace $\tilde{\tau} = ((w'_n : \psi_n^{a_n}, b'_n))_{n < \omega}$ on $\tilde{\pi}$ and a strictly increasing sequence of natural numbers $0 = n_0 < n_1 < \dots$ such that the following hold for all $i < \omega$ and $n_i \leq k < n_{i+1}$:*

$$(i) \ w'_k : \psi_k = w_i : \varphi_i;$$

$$(ii) \ b'_k = b_i.$$

Abusing notation, for every infinite branch $\tilde{\pi}$ of $\tilde{\mathcal{T}}$ and every trace τ on $f(\tilde{\pi})$, we denote by $f^{-1}(\tau)$ the corresponding trace on $\tilde{\pi}$ given by Lemma 4.4.12. Informally, $f^{-1}(\tau)$ is the trace that results from τ after the addition of annotations (and instances of the thinning rules) to \mathcal{T} in the construction of $\tilde{\mathcal{T}}$.

4.4.13. OBSERVATION. Lemmas 4.4.11 and 4.4.12 provide a back-and-forth correspondence between \mathcal{T} and $\tilde{\mathcal{T}}$: to every infinite branch π of $\tilde{\mathcal{T}}$ there corresponds an infinite branch $f(\pi)$ of \mathcal{T} , and then every trace τ on $f(\pi)$ yields a corresponding trace $f^{-1}(\tau)$ on π .

By inspection of the rules of iLTL_\circ and the priority given to the thinning rules in the construction of $\tilde{\mathcal{T}}$, we have:

4.4.14. LEMMA. *Let $\Omega ; \Theta \dashv \Gamma \Rightarrow \Delta$ be any sequent labelling a vertex in $\tilde{\mathcal{T}}$. For every non-annotated formula $w : \alpha$, the following hold:*

$$(i) \ w : \alpha \text{ is not 3-annotated in } \Gamma;$$

$$(ii) \ w : \alpha \text{ is not 3-annotated in } \Delta.$$

4.4.15. COROLLARY. *Only finitely many pairwise different names occur in $\tilde{\mathcal{T}}$.*

Proof. We prove the claim for U-names only; the proof for R-names is analogous. Let $n := |\varphi|$. By Proposition 4.2.38, there is an $m < \omega$ such that only m -many pairwise different world labels occur in $\tilde{\mathcal{T}}$. By Lemma 4.4.14, then, any sequent labelling a vertex of $\tilde{\mathcal{T}}$ contains at most $2mn$ -many pairwise different U-names. By the choice of new names when applying LU_1 in $\tilde{\mathcal{T}}$, it follows that the U-names occurring in $\tilde{\mathcal{T}}$ are all among the first $2mn$ -many U-names with respect to $\leq_{\mathbb{N}}$. ■

As an immediate consequence of Proposition 4.2.38 and Corollary 4.4.15, we have:

4.4.16. PROPOSITION. *Only finitely many pairwise different annotated sequents occur in $\tilde{\mathcal{T}}$.*

A *demotion* is a formula trace of the form $w : \alpha^x \triangleright w : \alpha^y$ with $x \neq y$. We say that x is *demoted* to y . Observe that demotions are only due to thinning.

The following ensures that the management of the annotations on $\tilde{\mathcal{T}}$ detects all good traces.

4.4.17. LEMMA. *Let π be an infinite branch of $\tilde{\mathcal{T}}$. If there is a good trace τ on $f(\pi)$ of type \mathbf{O} , for $\mathbf{O} \in \{\mathbf{U}, \mathbf{R}\}$, then some \mathbf{O} -name is eventually fixed on $f^{-1}(\tau)$.*

Proof. We only prove the case $\mathbf{O} = \mathbf{U}$, for the case $\mathbf{O} = \mathbf{R}$ follows analogously. Let $\pi = (u_i)_{i < \omega}$ and $\tau' := f^{-1}(\tau) = ((w_i : \varphi_i^{a_i}, b_i))_{i < \omega}$. Since τ is a left, label-stable \mathbf{U} -trace, there is a world label w , a \mathbf{U} -formula ψ , and some $n < \omega$ such that $w_i : \varphi_i \in \{w : \psi, w : X\psi\}$ and $b_i = 0$ for every $i \geq n$.

Note that there are infinitely many $i \geq n$ such that rule LU_0 or LU_1 is applied at u_i with principal formula $w : \psi^{a_i}$, and that by construction of $\tilde{\mathcal{T}}$ each such application produces a formula of the form $w : X\psi^y$, where $y \in \mathbf{N}_U$. The only way for y not to remain fixed on $\tau'_{>i}$ is to encounter an instance of Lthin on $\tau'_{>i}$ which demotes y to another name y' . So it suffices to show that τ' only passes through finitely many demotions.

Suppose that $n < i < j$ are such that $w : \varphi_i^{a_i} \triangleright w : \varphi_{i+1}^{a_{i+1}}$ and $w : \varphi_j^{a_j} \triangleright w : \varphi_{j+1}^{a_{j+1}}$ are demotions and that τ' encounters no demotion in between those two. Let $a_i = y$, $a_{i+1} = a_j = y'$, and $a_{j+1} = y''$. We then have:

$$\begin{array}{c} \text{Lthin} \frac{\Omega' ; \Theta''' \dashv \Gamma', w : \varphi_{j+1}^{y''} \Rightarrow \Delta'}{\Omega' ; \Theta'' \dashv \Gamma', w : \varphi_j^{y'}, w : \varphi_{j+1}^{y''} \Rightarrow \Delta'} \\ \vdots \\ \text{Lthin} \frac{\Omega ; \Theta' \dashv \Gamma, w : \varphi_{i+1}^{y'} \Rightarrow \Delta}{\Omega ; \Theta \dashv \Gamma, w : \varphi_i^y, w : \varphi_{i+1}^{y'} \Rightarrow \Delta} \end{array}$$

As $y' \prec_{\Theta} y$, we have $\Theta = ay'byc$. Analogously, $\Theta'' = a'y''b'y'c'$ because $y'' \prec_{\Theta''} y'$. By assumption, there is no demotion in between those two with principal formula in τ' , so y' is fixed on $(\tau'(i+1), \dots, \tau'(j))$ and thus $|a'| < |a'y''b'| \leq |a|$ because new names are always appended to the right end of nominal controls (reading the rules bottom-up). Therefore, τ' encounters only finitely many demotions. \blacksquare

Completeness of iLTL_o now follows easily.

4.4.18. PROPOSITION (Completeness of iLTL_o). *For any formula φ , if $\text{iLTL} \models \varphi$, then $\text{iLTL}_o \vdash \varphi$.*

Proof. Let \mathcal{T} be an iLTL_∞ proof of φ given by Proposition 4.2.30 and Corollary 4.2.60, and let $\tilde{\mathcal{T}}$ and $f : \tilde{T} \rightarrow T$ be given as described above. By

Lemma 4.4.11, every leaf of $\tilde{\mathcal{T}}$ is axiomatic. It thus suffices to show that every infinite branch of $\tilde{\mathcal{T}}$ passes through a successful repeat, for then pruning each infinite branch of $\tilde{\mathcal{T}}$ at the encounter of the first successful repeat on the branch yields an iLTL_\circ proof of φ .

Let then π be an infinite branch of $\tilde{\mathcal{T}}$. By Lemma 4.4.11, $f(\pi)$ is an infinite branch of \mathcal{T} . So, since \mathcal{T} is a proof, there is a good infinite trace τ on $f(\pi)$, say of type \mathbf{O} . By Lemmas 4.4.12 and 4.4.17, then, $f^{-1}(\tau)$ is an infinite trace on π and there is an \mathbf{O} -name which is eventually fixed on $f^{-1}(\tau)$. Therefore, by Propositions 4.2.41 and 4.4.16 there are infinitely many successful repeats in π . ■

Completeness of iLTL_{reg} now follows from Corollary 4.4.9 and Proposition 4.4.18:

4.4.19. COROLLARY. *For any formula φ , if $\text{iLTL} \models \varphi$, then $\text{iLTL}_{\text{reg}} \vdash \varphi$.*

Putting Propositions 4.4.7 and 4.4.18 together, we get:

4.4.20. THEOREM. *An iLTL formula φ is valid if, and only if, there is an iLTL_\circ proof of φ .*

4.5 Conclusion

Intuitionistic linear-time temporal logic (iLTL) combines Pnueli's linear-time temporal logic and (propositional) intuitionistic logic. Models for iLTL are thus 'two-dimensional': every state has at least one temporal successor and zero or more intuitionistic successors. A simple confluence condition, expressing monotonicity of the temporal successor function with respect to the intuitionistic order, suffices to ensure that truth is preserved upwards in the intuitionistic dimension.

We have provided a cut-free, finitary cyclic system for iLTL . The calculus uses labelled formulas in order to accommodate the interplay between the temporal dimension, represented by the modal rule for the next operator \mathbf{X} , and the intuitionistic dimension, corresponding to the right-implication rule $\mathbf{R}\rightarrow$. In ordinary intuitionistic calculi, $\mathbf{R}\rightarrow$ is not invertible. Labels, in contrast, yield an invertible right-implication rule. This appears to be necessary to obtaining regular ill-founded proofs, and thus in the end a cyclic calculus.

Presently, no axiom system for iLTL is known [10]. A natural continuation of our work, then, would be to study cycle elimination in iLTL_\circ and see whether it yields an axiomatisation of the logic. In CTL_\circ^* , existential quantifiers made it difficult to remove all cycles. In the case of iLTL_\circ , the difficulty lies in finding adequate intuitionistic inductive rules.

We could also consider variations of iLTL . In the semantics that we gave for the logic, we only required the temporal successor function to be forward confluent

with respect to the intuitionistic order: if $s \leq t$, then $S(s) \leq S(t)$. This is exactly the condition required for monotonicity of intuitionistic truth. Another confluence property considered in the literature is *backward confluence*: if $S(s) \leq t'$, then there is a $t \geq s$ such that $S(t) = t'$. Imposing both conditions yields a logic, iLTL^P , which is different than iLTL [10].⁴ Consider the formula $(Xp \rightarrow Xq) \rightarrow X(p \rightarrow q)$, for example. We showed in Figure 4.9 that it is not valid in iLTL . But it is valid in iLTL^P [10, Prop. 4.2]. Note that the countermodel depicted in Figure 4.9 is not backward confluent.

We could therefore try to adapt our labelled calculus to accommodate backward confluence. A natural proposal is generalising world labels to expressions of the form $s^n w$, where $w \in \text{Worlds}$ and $n < \omega$. Rule $\text{R}\rightarrow$ would then become

$$\text{R}\rightarrow \frac{\Omega, w \preceq w' \dashv \Gamma, s^n w' : \varphi \Rightarrow s^n w' : \psi, \Delta}{\Omega \dashv \Gamma \Rightarrow s^n w : \varphi \rightarrow \psi, \Delta},$$

and we could replace X with the invertible rules:

$$\text{LX} \frac{\Omega \dashv \Gamma, s^{n+1} w : \varphi \Rightarrow \Delta}{\Omega \dashv \Gamma, s^n w : X\varphi \Rightarrow \Delta} \quad \text{and} \quad \text{RX} \frac{\Omega \dashv \Gamma \Rightarrow s^{n+1} w : \varphi, \Delta}{\Omega \dashv \Gamma \Rightarrow s^n w : X\varphi, \Delta}$$

The labels of this new system allow one to reason under X , just like the nested sequents of the system in [5]. This might make it difficult or even impossible to always obtain ‘exact’ repeats in a proof-search; if this is the case, we could relax the notion of repeat to accommodate uniform label substitutions of the form $s^n w \mapsto s^{n+c} w$. We leave the investigation of this calculus for future work.

Looking at our completeness proofs for iLTL_∞ and iLTL_\circ , it seems likely that a decision procedure could be extracted from them. The combinatorics on labelled trees involved in proof-search, however, would probably yield non-elementary complexity bounds, thus no better than the ones obtained in [10], where it is shown that iLTL enjoys the effective finite model property.

Lastly, it would be interesting to apply the label-based approach to intuitionistic versions of other temporal and fixpoint logics, such as CTL^* and the modal μ -calculus.

⁴We borrow the name iLTL^P from [10].

Cyclic Proof Systems and Interpolation

Chapter 5

Uniform Interpolants from Cyclic Proofs

This chapter differs from the preceding ones in that we shall not design here cyclic or ill-founded proof systems, but rather use an existing one to build uniform interpolants for the modal μ -calculus.

A logic has the (*Craig*) *interpolation* property if, informally, for every formulas α and β such that $\alpha \rightarrow \beta$ is valid, there exists a formula ι in the common vocabulary¹ of α and β such that both $\alpha \rightarrow \iota$ and $\iota \rightarrow \beta$ are valid. This property has its origins in Craig's [30], where it was shown that first-order logic enjoys interpolation. Later, Lyndon [91] strengthened this result showing that the vocabulary can take polarities into account (so, for example, if a propositional letter occurs positively in the interpolant, then it must also occur positively in both α and β). This property came to be known as *Lyndon interpolation*.

In [107], Pitts proved a stronger result than Craig interpolation, called *uniform interpolation*, for intuitionistic propositional logic: for every formula α and every $V \subseteq \text{Voc}(\alpha)$, where $\text{Voc}(\alpha)$ denotes the vocabulary of α , there exists a formula ι such that:

- (i) $\text{Voc}(\iota) \subseteq V$;
- (ii) $\alpha \rightarrow \iota$ is valid;
- (iii) for any β such that $\text{Voc}(\alpha) \cap \text{Voc}(\beta) \subseteq V$, if $\alpha \rightarrow \beta$ is valid, then so is $\iota \rightarrow \beta$.

To see that uniform interpolation entails Craig interpolation, let $V := \text{Voc}(\alpha) \cap \text{Voc}(\beta)$.

Looking again at Craig interpolation, it stands to reason that if a sufficiently 'nice' proof of a valid implication $\alpha \rightarrow \beta$ is available, one may succeed in defining an interpolant by induction on the proof-tree, starting from the leaves and proceeding

¹This notion differs slightly from logic to logic. In predicate logic, for example, it is the collection of all non-logical symbols in a formula.

to the implication at the root. This method has recently been applied even to fixpoint logics admitting cyclic proofs [134, 1].

In contrast, for uniform interpolation there is no single proof to work from but instead a collection of proofs to accommodate: a witness to each valid implication $\alpha \rightarrow \beta$, where the vocabulary of β is constrained as above. Working over a set of prospective proofs and relying on the structural properties of sequent calculus is the essence of Pitts’s aforementioned result on uniform interpolation for intuitionistic logic [107].

Proofs of uniform interpolation differ from one system to another. There have been efforts to find general frameworks to attack the problem. Notably, [70, 71, 148] identify sufficient conditions on the form of proof systems that entail uniform interpolation.

In this chapter we provide a proof of the uniform interpolation property for Kozen’s modal μ -calculus [81], a well-known extension of modal logic with explicit fixpoint quantifiers. The result was first established by D’Agostino and Hollenberg [32] by employing automata-theoretic tools to show the definability of bisimulation quantifiers in the logic. With this form of second-order quantification, uniform interpolants can readily be defined.

Our approach, in contrast, is purely proof-theoretic. We build uniform interpolants in an annotated, goal-oriented cyclic system due to Jungteerapanich [76] and Stirling [144]. A different cyclic system based on it was exploited in [1] to establish Lyndon interpolation for the modal μ -calculus.

The main ideas that we present are not specific to the μ -calculus but do rely on two of its features: the existence of (cyclic) analytic systems permitting ‘uniform’ proof-search; and the ability of the logic to express fixpoints. In the conclusion of the chapter we shall examine the applicability of our argument to other fixpoint logics.

Outline of the chapter. Section 5.1 introduces the modal μ -calculus. Section 5.2 examines the Jungteerapanich–Stirling cyclic system. Section 5.3 proves the uniform interpolation property for the μ -calculus by building (and verifying) uniform interpolants in the Jungteerapanich–Stirling system. Section 5.4 concludes the chapter and discusses some further lines of research based on the material therein.

5.1 The modal μ -calculus

The modal μ -calculus extends propositional (multi-)modal logic by adding two fixpoint quantifiers, μ and ν . Syntactically, they behave like quantifiers in predicate

logic, in the sense that they bound occurrences of variables in formulas. Semantically, μ and ν denote, respectively, least and greatest fixed points of functions, and are thus a kind of monadic second-order quantifiers.²

In its modern form, the modal μ -calculus was introduced by Kozen in [81]. It subsumes a wide range of temporal logics, such as Pnueli's linear-time temporal logic LTL [108]; Clarke and Emerson's computation tree logic CTL [26]; Emerson and Halpern's full computation tree logic CTL* [42];³ and Fischer and Ladner's propositional dynamic logic PDL [45, 46], among others. We assume that the reader is acquainted with the μ -calculus. An introduction may be found in [38, Ch. 8], and [16] surveys a range of results about it and related logics.

The *language of the (modal) μ -calculus*, denoted by \mathcal{L}_μ , consists of the following: countably many *propositional letters*, (*modal*) *actions* and *variables*, drawn respectively from sets **Prop**, **Act** and **Var**; the *constants* \perp (falsum) and \top (verum); the *Boolean connectives* \wedge (conjunction), \vee (disjunction) and $\bar{\cdot}$ (negation); the *modal operators* $[a]$ (box) and $\langle a \rangle$ (diamond), for each $a \in \mathbf{Act}$; and the *fixpoint quantifiers* μ and ν . The *formulas of the modal μ -calculus* are given by the following grammar:

$$\varphi ::= \top \mid \perp \mid p \mid \bar{p} \mid x \mid (\varphi \wedge \varphi) \mid (\varphi \vee \varphi) \mid ([a]\varphi) \mid (\langle a \rangle\varphi) \mid (\mu x \varphi) \mid (\nu x \varphi),$$

where p ranges over **Prop**, x over **Var**, and a over **Act**. Formulas are denoted by small Greek letters $\alpha, \beta, \varphi, \dots$, and sets of formulas by capital Greek letters $\Gamma, \Delta, \Sigma, \dots$. We use σ to denote either μ or ν . The collection of all μ -calculus formulas is denoted by \mathbf{Form}_μ , and the set of all *subformulas* of a formula φ , defined as usual (treating μx and νx like first-order quantifiers), is denoted by $\mathbf{Sub}(\varphi)$.

If no ambiguity arises, we drop the outer parenthesis and stipulate that $[a]$ and $\langle a \rangle$ bind more strongly than \wedge and \vee , and that these, in turn, bind more strongly than μ and ν .

A *literal* is either \top , \perp , or a formula of the form p or \bar{p} for some $p \in \mathbf{Prop}$. A *quantifier-free* formula is a formula built from literals and variables by means solely of modal operators and Boolean connectives. A σ -*formula*, for $\sigma \in \{\mu, \nu\}$, is a formula of the form $\sigma x \varphi$. A $[a]$ -*formula* ($\langle a \rangle$ -*formula*) is a formula of the form $[a]\varphi$ (respectively, $\langle a \rangle\varphi$). A $[\cdot]$ -*formula* ($\langle \cdot \rangle$ -*formula*) is a $[a]$ -formula (respectively, $\langle a \rangle$ -formula) for any $a \in \mathbf{Act}$. A *modal* formula is either a $[\cdot]$ - or a $\langle \cdot \rangle$ -formula. Given a set of formulas Γ , we define $[a]\Gamma := \{[a]\gamma \mid \gamma \in \Gamma\}$ and $\langle a \rangle\Gamma := \{\langle a \rangle\gamma \mid \gamma \in \Gamma\}$.

Observe that we allow negation to be applied to propositional letters only. This is not necessary and one could instead work with unrestricted negation and a small-

²In fact, the modal μ -calculus is the bisimulation-invariant fragment of monadic second-order logic [74].

³The translation from CTL* into the μ -calculus is far from straightforward. A direct translation is due to Dam [33].

ler language (for example, have μ as primitive and define $\nu x \varphi := \neg \mu x \neg \varphi(\neg x/x)$).⁴ Restricting negation to propositional letters is convenient when working with proof systems.

An occurrence of a variable x in a formula φ is *bound* if it is within the scope of a quantifier σx for some $\sigma \in \{\mu, \nu\}$, and it is *free* otherwise. A formula is *closed* if no variable occurs free in it. A set of formulas Γ is *closed* if every formula in Γ is closed. A formula φ is *well-named* if no variable occurs both free and bound in φ and, moreover, no variable is bound in φ more than once. Hence, each bound variable x in a well-named formula φ identifies a unique subformula $\sigma_x \varphi_x$ of φ . When φ is clear from context we express this correspondence as $x =_{\sigma_x} \varphi_x$. A finite set of formulas Γ is *well-named* if the formula $\bigwedge \Gamma$ is well-named.

Modal μ -calculus formulas are interpreted over labelled transition systems, essentially Kripke models with one or more accessibility relations. Informally, $\sigma\varphi$ denotes either the least (μ) or the greatest (ν) fixed point of the function represented by $\varphi(x)$.

A *labelled transition system (LTS)* is a triple $\mathcal{S} = (S, \{\overset{a}{\rightarrow} \mid a \in \text{Act}\}, \lambda)$, where S is a non-empty set of *states*, each $\overset{a}{\rightarrow} \subseteq S \times S$ is a *transition relation*, and $\lambda : \text{Prop} \rightarrow 2^S$ is the *labelling map* of \mathcal{S} . A *variable assignment* for \mathcal{S} is a function $V : \text{Var} \rightarrow 2^S$. If $T \subseteq S$, we denote by $V[x \mapsto T]$ the variable assignment for \mathcal{S} which sends x to T and every other variable $y \neq x$ to $V(y)$.

Given an LTS $\mathcal{S} = (S, \{\overset{a}{\rightarrow} \mid a \in \text{Act}\}, \lambda)$ and a variable assignment V for \mathcal{S} , we inductively define the sets $\llbracket \varphi \rrbracket_V^{\mathcal{S}}$ as follows:

- (i) $\llbracket \top \rrbracket_V^{\mathcal{S}} := S$ and $\llbracket \perp \rrbracket_V^{\mathcal{S}} := \emptyset$;
- (ii) $\llbracket p \rrbracket_V^{\mathcal{S}} := \lambda(p)$, for every $p \in \text{Prop}$;
- (iii) $\llbracket \bar{p} \rrbracket_V^{\mathcal{S}} := S \setminus \lambda(p)$, for every $p \in \text{Prop}$;
- (iv) $\llbracket x \rrbracket_V^{\mathcal{S}} := V(x)$, for every $x \in \text{Var}$;
- (v) $\llbracket \varphi \wedge \psi \rrbracket_V^{\mathcal{S}} := \llbracket \varphi \rrbracket_V^{\mathcal{S}} \cap \llbracket \psi \rrbracket_V^{\mathcal{S}}$;
- (vi) $\llbracket \varphi \vee \psi \rrbracket_V^{\mathcal{S}} := \llbracket \varphi \rrbracket_V^{\mathcal{S}} \cup \llbracket \psi \rrbracket_V^{\mathcal{S}}$;
- (vii) $\llbracket [a]\varphi \rrbracket_V^{\mathcal{S}} := \{s \in S \mid s \overset{a}{\rightarrow} t \text{ implies } t \in \llbracket \varphi \rrbracket_V^{\mathcal{S}} \text{ for all } t \in S\}$;
- (viii) $\llbracket \langle a \rangle \varphi \rrbracket_V^{\mathcal{S}} := \{s \in S \mid s \overset{a}{\rightarrow} t \text{ for some } t \in \llbracket \varphi \rrbracket_V^{\mathcal{S}}\}$;
- (ix) $\llbracket \mu x \varphi \rrbracket_V^{\mathcal{S}} := \bigcap \{T \subseteq S \mid \llbracket \varphi \rrbracket_{V[x \mapsto T]}^{\mathcal{S}} \subseteq T\}$;
- (x) $\llbracket \nu x \varphi \rrbracket_V^{\mathcal{S}} := \bigcup \{T \subseteq S \mid T \subseteq \llbracket \varphi \rrbracket_{V[x \mapsto T]}^{\mathcal{S}}\}$.

⁴A triple negation is necessary to ensure the monotonicity of the function represented by $\neg\varphi(\neg x/x)$, and thus the existence of least and greatest fixed points of the function.

We say that $s \in S$ *satisfies* φ with respect to V , in symbols $\mathcal{S}, s \Vdash_V \varphi$, if $s \in \llbracket \varphi \rrbracket_V^S$. We say that φ is *valid* if we have $\mathcal{S}, s \Vdash_V \varphi$ for every LTS \mathcal{S} , every variable assignment V for \mathcal{S} , and every state s of \mathcal{S} . Finally, given any formulas φ and ψ , we write $\varphi \equiv \psi$, and say that φ and ψ are *equivalent*, if for every LTS \mathcal{S} , every variable assignment V for \mathcal{S} and every state s of \mathcal{S} , we have $\mathcal{S}, s \Vdash_V \varphi$ if, and only if, $\mathcal{S}, s \Vdash_V \psi$.

By the well-known Knaster–Tarski theorem [150], $\llbracket \mu x \varphi \rrbracket_V^S$ and $\llbracket \nu x \varphi \rrbracket_V^S$ are, respectively, the least and the greatest fixed point of the function $T \mapsto \llbracket \varphi \rrbracket_{V[x \mapsto T]}^S$.

The *dual* of a formula φ , in symbols φ^∂ , is inductively defined as follows:

$$\begin{array}{ll}
 \perp^\partial := \top & \top^\partial := \perp \\
 p^\partial := \bar{p} & \bar{p}^\partial := p \\
 x^\partial := x & \\
 (\varphi \wedge \psi)^\partial := \varphi^\partial \vee \psi^\partial & (\varphi \vee \psi)^\partial := \varphi^\partial \wedge \psi^\partial \\
 ([\mathbf{a}]\varphi)^\partial := \langle \mathbf{a} \rangle \varphi^\partial & (\langle \mathbf{a} \rangle \varphi)^\partial := [\mathbf{a}]\varphi^\partial \\
 (\mu x \varphi)^\partial := \nu x \varphi^\partial & (\nu x \varphi)^\partial := \mu x \varphi^\partial
 \end{array}$$

We then define implication as $\varphi \rightarrow \psi := \varphi^\partial \vee \psi$.

The following is easy to see:

5.1.1. PROPOSITION. *Let \mathcal{S} be an LTS, V a variable assignment for \mathcal{S} , and s a state of \mathcal{S} . For every formula φ , we have $\mathcal{S}, s \Vdash_V \varphi$ if, and only if, $\mathcal{S}, s \not\Vdash_{V^\partial} \varphi^\partial$, where $V^\partial : \text{Var} \rightarrow 2^S$ is given by $V^\partial(x) := S \setminus V(x)$ for every $x \in \text{Var}$.*

A formula φ is *guarded* if, for any subformula $\sigma x \psi$ of φ , every occurrence of x in ψ is within a modal subformula of ψ . The following is well known (see, e.g., [102] or [38, § 8.3.5]):

5.1.2. PROPOSITION. *Every formula is equivalent to a guarded one.*

A finite set of formulas Γ is *guarded* if the formula $\bigwedge \Gamma$ is guarded.

We shall frequently restrict ourselves to (closed and well-named) guarded formulas. In particular, the Jungteerapanich–Stirling proof system for the μ -calculus, with which we shall build uniform interpolants below, assumes that formulas are guarded. Proposition 5.1.2 ensures that this carries no loss of generality. It is worth noting, nevertheless, that guarding a formula might yield an exponential blow-up in formula size (see, e.g., [38, Ex. 8.3.22]).

Bound variables in well-named formulas can be ordered in terms of their mutual dependencies. This is made precise in the following definition.

5.1.3. DEFINITION (Subsumption order). Let φ be a well-named formula, and V the set of all bound variables in φ . The *subsumption order* of φ , in symbols $<_{\varphi}$, is the irreflexive partial ordering of V defined by setting $x <_{\varphi} y$ if $\sigma_y \varphi_y$ is a proper subformula of $\sigma_x \varphi_x$.

If $x <_{\varphi} y$, we say that x *subsumes* y . ⊣

Intuitively, $x <_{\varphi} y$ if the evaluation of $\sigma_y \varphi_y$ potentially depends on the value of x , in the sense that x might occur free in $\sigma_y \varphi_y$. For example, if $\varphi = \mu x(p \vee \nu y(x \vee y))$, then $x <_{\varphi} y$. Note, however, that the same holds if $\varphi = \mu x(x \vee \nu y y)$, despite there not being any real dependency of $\nu y y$ on x .

Well-named formulas admit a natural presentation as systems of fixpoint equations. This alternative point of view will be convenient for our purposes. We provide a brief presentation of equational systems, referring the reader to [38, § 8.3.4] for more details.

5.1.4. DEFINITION (Modal equational system). Let $V_{\mathcal{E}}$ be a finite set of variables. A *modal equational system (MES)* over $V_{\mathcal{E}}$ is a triple $(\varphi, \mathcal{E}, V_{\mathcal{E}})$, where φ is a quantifier-free formula over the set of variables $V_{\mathcal{E}}$ and \mathcal{E} is a set of equations

$$\{x =_{p_x} \varphi_x \mid x \in V_{\mathcal{E}}\},$$

where each φ_x is a quantifier-free formula over $V_{\mathcal{E}}$ and $p_x < \omega$ for every $x \in V_{\mathcal{E}}$. We call p_x the *priority* of x , and say that x has a *higher priority* than y if $p_x < p_y$. ⊣

Every well-named formula φ determines a MES as follows. Let $V = \{x_1, \dots, x_n\}$ be the collection of bound variables in φ , where $x_i <_{\varphi} x_j$ implies $i < j$. For each $i = 1, \dots, n$, let $p_i := 2i$ if $\sigma_{x_i} = \nu$, and otherwise $p_i := 2i + 1$. Finally, let φ_i be the result after replacing each subformula $\sigma_{x_j} \varphi_{x_j}$ by x_j in φ_{x_i} . The MES corresponding to φ is $S_{\varphi} = (\varphi^*, \mathcal{E}, V)$, where $\mathcal{E} := \{x_i =_{p_i} \varphi_i \mid i = 1, \dots, n\}$ and φ^* results after replacing each subformula $\sigma_{x_j} \varphi_{x_j}$ by x_j in φ .

5.1.5. EXAMPLE. Consider the formula $\varphi = \nu x \mu y(x \vee y)$. The corresponding MES is $S_{\varphi} = (x, \{x =_2 y, y =_5 x \vee y\}, \{x, y\})$

Conversely, there is a recursive algorithm which flattens every MES S down to a μ -calculus formula φ_S and, moreover, in such a way that $\varphi_{S_{\varphi}} = \varphi$. We omit its definition and refer the reader to [38, § 8.3.4], as the internal details of this procedure are for the most part irrelevant for our purposes.⁵ It suffices to point out that, as expected, the flattening of a MES commutes with the Boolean connectives and the modal operators, that odd priorities correspond to μ and even ones to ν , and that the subsumption order of the resulting formula corresponds to the order of the priorities.

⁵Note that, in [38], our $x <_{\varphi} y$ is written $x >_{\varphi} y$ and higher priorities correspond to bigger natural numbers.

5.2 The Jungteerapanich–Stirling proof system

Already in his presentation of the μ -calculus, Kozen [81] proposed a natural axiomatisation of the logic in the form of an equational deductive system extending basic propositional multi-modal logic by the addition of the axiom

$$\varphi(\mu x \varphi/x) \rightarrow \mu x \varphi,$$

expressing that $\mu x \varphi$ is a pre-fixpoint of $\varphi(x)$, and the rule

$$\frac{\varphi(\psi/x) \rightarrow \psi}{\mu x \varphi \rightarrow \psi},$$

corresponding to Park’s well-known fixpoint induction principle [105]. Together, they characterise $\mu x \varphi$ as the least pre-fixpoint of $\varphi(x)$, and thus, by the Knaster–Tarski theorem [150], as the least fixpoint of $\varphi(x)$.

Kozen, however, was only able to establish in [81] the completeness of his axiomatisation for a proper fragment of the μ -calculus (which he termed ‘aconjunctive’). Later, he showed that the logic enjoys the finite model property and gave a sound and complete infinitary system with an ω -rule based on that result [82]. Walukiewicz [156] proposed an alternative, finitary system based on the small model theorem for the μ -calculus, which states that every satisfiable formula has a model of size exponential in the size of the formula [145].

The question of the completeness of Kozen’s original axiomatisation remained open for several years until Walukiewicz settled it in the affirmative [157]. As a step towards this result, Niwiński and Walukiewicz [102] presented a tableau system for checking the unsatisfiability of formulas. The corresponding dual sequent calculus is a natural sound and complete ill-founded proof system for the logic (see [37]).

Kozen’s infinitary system from [82] included a cut rule. A cut-free system with an ω -rule is given in [73], and [146] presents an embedding of said system into an ill-founded calculus dual to Niwiński and Walukiewicz’s tableau from [102].

More recently, Jungteerapanich [76] and Stirling [144] introduced a cut-free, finitary, sound and complete cyclic system for the μ -calculus by incorporating into the system Safra’s determinisation procedure for Büchi automata [123] in the form of ‘names’ annotating formulas (see, in particular, [77, § 4.3.5]). The annotations keep track of unfoldings of ν -formulas and allow detecting good traces. This idea is further explored in [37], where determinisation procedures other than Safra’s are considered. Jungteerapanich and Stirling’s annotations are the ultimate source of inspiration for the annotations that we used above for CTL* and iLTL. The μ -calculus, however, requires a considerably more involved annotation mechanism. This is in line with the fact that it corresponds to Safra’s construction, well known to be far from trivial.

We shall work with a two-sided version of the Jungteerapanich–Stirling system. As shown in (the completeness proofs in) [76, 144], the system admits terminating proof-search as long as a few requirements are met (see Definition 5.2.5 below). This will allow us to work with proofs having essentially the same ‘left-fragment’, i.e., proofs that treat the antecedents of the sequents in the same way. This uniformity lies at heart of our proof of uniform interpolation.

A (*plain μ -calculus*) *sequent* is a pair (Γ, Δ) , henceforth written $\Gamma \Rightarrow \Delta$, where Γ and Δ are finite sets of formulas. We call Γ the *antecedent* of the sequent, and Δ the *consequent* or *succedent*.

A sequent $\Gamma \Rightarrow \Delta$ is *closed* (*guarded*, *well-named*) if the set $\Gamma \cup \Delta$ is closed (respectively, guarded, well-named). The *interpretation* of a sequent $\Gamma \Rightarrow \Delta$ is the formula $(\Gamma \Rightarrow \Delta)^\# := \bigwedge \Gamma \rightarrow \bigvee \Delta$. A sequent S is *valid* if $S^\#$ is valid.

For every $x \in \text{Var}$, let $\mathbf{N}_x = \{x_0, x_1, \dots\}$ be a countably infinite set of *names* for x such that $\mathbf{N}_x \cap \mathbf{N}_y = \emptyset$ if $x \neq y$. We denote names for variables x, y, z, \dots by x, y, z, \dots , respectively (possibly with indices), and let $\mathbf{N} := \bigcup_{x \in \text{Var}} \mathbf{N}_x$. In addition, we assume that \mathbf{N} is well-ordered.

An *annotation* is a (possibly empty) finite sequence of pairwise different names in \mathbf{N} . An *annotated formula* is a pair (φ, a) , henceforth written φ^a , where φ is a formula and a is an annotation. The *base* of an annotated formula φ^a is $\mathbf{b}(\varphi^a) := \varphi$. Similarly, the *base* of a set of annotated formulas Γ is $\mathbf{b}(\Gamma) := \{\mathbf{b}(\varphi^a) \mid \varphi^a \in \Gamma\}$. A name x occurs in a set of annotated formulas Γ if x occurs in the annotation of some formula in Γ .

An (*annotated*) *sequent* is a triple (Θ, Γ, Δ) , henceforth written $\Theta \dashv \Gamma \Rightarrow \Delta$, where Θ is an annotation and Γ and Δ are finite sets of annotated formulas such that a name occurs in Θ if, and only if, it occurs $\Gamma \cup \Delta$. We call Θ the *control* of $\Theta \dashv \Gamma \Rightarrow \Delta$. Sequents with empty controls are identified with plain sequents. The *base* of an annotated sequent $\Theta \dashv \Gamma \Rightarrow \Delta$ is $\mathbf{b}(\Theta \dashv \Gamma \Rightarrow \Delta) := \mathbf{b}(\Gamma) \Rightarrow \mathbf{b}(\Delta)$.

Let φ be closed, guarded, and well-named. Fix an arbitrary linear ordering of the bound variables in φ , say $x_1 \prec \dots \prec x_n$, compatible with the subsumption order of φ , i.e., such that $i < j$ implies $x_i \not\prec_\varphi x_j$. If φ is given as an equational system, \prec can be chosen as any linear order such that $p_x < p_y$ implies $x \prec y$. Given an annotation a for bound variables in φ , we denote by $a \upharpoonright x_i$ the result of removing from a all names for x_{i+1}, \dots, x_n .

We now define the JS sequent calculus. Its rules are given in Figures 5.1 to 5.3.

In rules $\mathbf{L}\square$, $\mathbf{R}\square$, $\mathbf{L}\mu$, $\mathbf{R}\mu$, $\mathbf{L}\nu$, $\mathbf{R}\nu$, \mathbf{Lthin} , \mathbf{Rthin} , \mathbf{Lres} , \mathbf{Rres} and \mathbf{wk} , we denote by Θ' the subsequence of Θ given by removing any name which does not occur in the sequent whose control is Θ' .

In $\mathbf{L}\mu$ and $\mathbf{R}\nu$, x is a name for the variable x not occurring in Θ , and we define $\Theta'x := \Theta' \frown (x)$.

$$\begin{array}{ll}
\text{ax}_{\perp} \frac{}{\Theta \dashv \Gamma, \perp^a \Rightarrow \Delta} & \text{ax}_{\top} \frac{}{\Theta \dashv \Gamma \Rightarrow \top^a, \Delta} \\
\text{ax}_p \frac{}{\Theta \dashv \Gamma, p^a \Rightarrow p^b, \Delta} & \text{ax}_{\bar{p}} \frac{}{\Theta \dashv \Gamma, \bar{p}^a \Rightarrow \bar{p}^b, \Delta} \\
\text{L}\wedge \frac{\Theta \dashv \Gamma, \varphi^a, \psi^a \Rightarrow \Delta}{\Theta \dashv \Gamma, (\varphi \wedge \psi)^a \Rightarrow \Delta} & \text{R}\wedge \frac{\Theta \dashv \Gamma \Rightarrow \varphi^a, \Delta \quad \Theta \dashv \Gamma \Rightarrow \psi^a, \Delta}{\Theta \dashv \Gamma \Rightarrow (\varphi \wedge \psi)^a, \Delta} \\
\text{L}\vee \frac{\Theta \dashv \Gamma, \varphi^a \Rightarrow \Delta \quad \Theta \dashv \Gamma, \psi^a \Rightarrow \Delta}{\Theta \dashv \Gamma, (\varphi \vee \psi)^a \Rightarrow \Delta} & \text{R}\vee \frac{\Theta \dashv \Gamma \Rightarrow \varphi^a, \psi^a, \Delta}{\Theta \dashv \Gamma \Rightarrow (\varphi \vee \psi)^a, \Delta} \\
\text{L}\square \frac{\Theta' \dashv \Gamma, \varphi^b \Rightarrow \Delta}{\Theta \dashv \Pi, [a]\Gamma, \langle a \rangle \varphi^b \Rightarrow \langle a \rangle \Delta, \Sigma} & \text{R}\square \frac{\Theta' \dashv \Gamma \Rightarrow \varphi^b, \Delta}{\Theta \dashv \Pi, [a]\Gamma \Rightarrow [a]\varphi^b, \langle a \rangle \Delta, \Sigma}
\end{array}$$

FIGURE 5.1: Non-fixpoint, logical rules of the system JS.

In the *reset* rules **Lres** and **Rres**, the names x, x_1, \dots, x_k all name the same variable; the other annotations (a, a_1, \dots, a_k) are arbitrary. We say that x is *reset*. Furthermore, in **Lres** the name x may not occur in Γ , and in **Rres** the name x may not occur in Δ .

In the *thinning* rules **Lthin** and **Rthin**, the relation \prec_{Θ} is a total ordering on subsequences of Θ defined as follows. If Θ is an annotation and a, b are subsequences of Θ , set $a <_{\Theta} b$ if, and only if, a precedes b in the lexicographic ordering induced by Θ . Then, let $a \prec_{\Theta} b$ if, and only if, either $a <_{\Theta} b$, or there is some variable x such that $b \upharpoonright x \sqsubset a \upharpoonright x$. We refer the reader to [77, § 4.3] for the proof that \prec_{Θ} is indeed total.

In the unfolding rules **L μ** , **L ν** , **R μ** and **R ν** , we say that the variable x is *unfolded* to φ .

Our presentation of JS is essentially the two-sided version of the system in [144]. We have added an explicit weakening rule which is easily seen to be admissible by considering the ill-founded version of the system and the implicit weakening in the modal rules and the axioms. Additionally, our fixpoint unfolding rules use x and $\varphi(x)$ in place of $\varphi(\sigma x \varphi)$, where $x =_{\sigma} \varphi$. This can be considered to be a mere syntactic abbreviation, as done actually in [144].

5.2.1. DEFINITION (JS derivation). A *JS derivation* of a closed, well-named and guarded plain sequent $\alpha \Rightarrow \beta$ is a tree with back-edges built according to the rules of JS and whose root has label $\alpha \Rightarrow \beta$. ⊣

A sequence of names a is *preserved* on a path π through a JS derivation if a is a prefix of every control of a sequent labelling a vertex in π .

$$\begin{array}{ll}
L\sigma_0 \frac{\Theta \dashv \Gamma, x^a \Rightarrow \Delta}{\Theta \dashv \Gamma, \sigma x \varphi^a \Rightarrow \Delta} & R\sigma_0 \frac{\Theta \dashv \Gamma \Rightarrow x^a, \Delta}{\Theta \dashv \Gamma \Rightarrow \sigma x \varphi^a, \Delta} \\
L\mu \frac{\Theta' x \dashv \Gamma, \varphi^{(a|x)x} \Rightarrow \Delta}{\Theta \dashv \Gamma, x^a \Rightarrow \Delta} x =_{\mu} \varphi & R\mu \frac{\Theta' \dashv \Gamma \Rightarrow \varphi^{(a|x)x}, \Delta}{\Theta \dashv \Gamma \Rightarrow x^a, \Delta} x =_{\mu} \varphi \\
L\nu \frac{\Theta' \dashv \Gamma, \varphi^{a|x} \Rightarrow \Delta}{\Theta \dashv \Gamma, x^a \Rightarrow \Delta} x =_{\nu} \varphi & R\nu \frac{\Theta' x \dashv \Gamma \Rightarrow \varphi^{(a|x)x}, \Delta}{\Theta \dashv \Gamma \Rightarrow x^a, \Delta} x =_{\nu} \varphi
\end{array}$$

FIGURE 5.2: Fixpoint rules of the system JS.

A proof in the JS calculus is a derivation for which all repeat leaves fulfil a correctness condition that we now specify.

5.2.2. DEFINITION (Successful repeat). A repeat l of a JS derivation \mathcal{T} is *successful* if there is some name x such that:

- (i) x occurs in all controls on the path $[c_l, l]_{\mathcal{T}}$;
- (ii) x is reset in $[c_l, l]_{\mathcal{T}}$.

⊥

Every successful repeat l in a JS derivation \mathcal{T} has an associated *invariant*, denoted by $\text{inv}(l)$, defined as the shortest sequence of names wx such that wx is a prefix of every control in $[c_l, l]_{\mathcal{T}}$ and x is reset on $[c_l, l]_{\mathcal{T}}$. The existence of invariants follows immediately from Definition 5.2.2 and the fact that, when reading JS rules bottom-up, new names are always appended to the right end of controls. We extend the invariant map inv to unsuccessful repeats by letting $\text{inv}(l)$ be the longest (possibly empty) common prefix of all controls on $[c_l, l]_{\mathcal{T}}$, for every unsuccessful repeat $l \in \text{Rep}_{\mathcal{T}}$.

The invariant map inv induces the following (reflexive) quasi-order \preceq on the successful repeats of \mathcal{T} : $l \preceq l'$ if, and only if, $\text{inv}(l) \sqsubseteq \text{inv}(l')$.

5.2.3. DEFINITION (JS proof). A JS *proof* is a JS derivation whose non-axiomatic leaves are all successful repeats. ⊥

Note that every repeat l of a JS proof has an associated non-empty invariant $\text{inv}(l) = wx$ such that x is preserved along the path $[c_l, l]_{\mathcal{T}}$ and reset somewhere therein.

The following is straightforward to prove by reducing JS to its one-sided fragment and appealing to the soundness and completeness result in [144]:

5.2.4. THEOREM. *For every closed, well-named and guarded plain sequent $\Gamma \Rightarrow \Delta$, we have $\text{JS} \vdash \Gamma \Rightarrow \Delta$ if, and only if, $\Gamma \Rightarrow \Delta$ is valid.*

$$\begin{array}{c}
\text{Lthin} \frac{\Theta' \dashv \Gamma, \varphi^a \Rightarrow \Delta}{\Theta \dashv \Gamma, \varphi^a, \varphi^b \Rightarrow \Delta} a \prec_{\Theta} b \qquad \text{Rthin} \frac{\Theta' \dashv \Gamma \Rightarrow \varphi^a, \Delta}{\Theta \dashv \Gamma \Rightarrow \varphi^a, \varphi^b, \Delta} a \prec_{\Theta} b \\
\text{Lres} \frac{\Theta' \dashv \Gamma, \varphi_1^{ax}, \dots, \varphi_k^{ax} \Rightarrow \Delta}{\Theta \dashv \Gamma, \varphi_1^{axx_1a_1}, \dots, \varphi_k^{axx_ka_k} \Rightarrow \Delta} \qquad \text{Rres} \frac{\Theta' \dashv \Gamma \Rightarrow \varphi_1^{ax}, \dots, \varphi_k^{ax}, \Delta}{\Theta \dashv \Gamma \Rightarrow \varphi_1^{axx_1a_1}, \dots, \varphi_k^{axx_ka_k}, \Delta} \\
\text{wk} \frac{\Theta' \dashv \Gamma \Rightarrow \Delta}{\Theta \dashv \Gamma, \Pi \Rightarrow \Sigma, \Delta}
\end{array}$$

FIGURE 5.3: Structural rules of the system JS.

To build (and verify) uniform interpolants, it will be convenient to work with JS derivations satisfying several constraints specified below, where we refer to the form of the rules given in Figures 5.1 to 5.3:

5.2.5. DEFINITION (Normal derivation). A JS derivation \mathcal{T} is *normal* if the following hold:

- (i) There are no instances of **wk** in \mathcal{T} .
- (ii) Any vertex of \mathcal{T} labelled by an axiomatic sequent is a leaf.
- (iii) If a sequent $\Theta \dashv \Gamma \Rightarrow \Delta$ in \mathcal{T} can be realised as the conclusion of an instance of **Lthin**, **Rthin**, **Lres** or **Rres**, then the sequent is the conclusion of such a rule in \mathcal{T} , with the thinning rules having precedence over reset rules.
- (iv) In instances of **L μ** and **R ν** in \mathcal{T} , x is the first name in N_x not occurring in Θ .⁶
- (v) In instances of **L \square** or **R \square** in \mathcal{T} , the set Π consists only of \top , $\langle \cdot \rangle$ -formulas, and $[c]$ -formulas for $c \neq a$; and Σ of only \perp , $[\cdot]$ -formulas, and $\langle c \rangle$ -formulas for $c \neq a$.
- (vi) Any two non-repeat vertices of \mathcal{T} labelled by the same sequent are labelled by the same instance of the same rule.

–

The constraints in Definition 5.2.5 have their origin in the proof-search carried out in [144], as well as in the dual construction in [76]. Therefore, the following is a direct consequence of the completeness proofs in [76, 144]:

5.2.6. THEOREM. *A closed, well-named and guarded sequent $\Gamma \Rightarrow \Delta$ is valid if, and only if, there exists a normal JS proof of $\Gamma \Rightarrow \Delta$.*

⁶Recall that we fixed an arbitrary well-order on N .

The following can readily be established by considering the ω -unravelling of a JS proof, dropping the annotations therein, and showing that the resulting (possibly infinite) tree is a proof in the system dual to Niwiński and Walukiewicz's [102] (for an explicit presentation of the proof system, see, e.g., [37]). The proof is a routine imitation of the arguments in Sections 3.4 and 4.4 for soundness of CTL_\circ^* and iLTL_\circ , respectively.

5.2.7. LEMMA. *Let \mathcal{T} be a JS proof, and $\Theta \dashv \Gamma \Rightarrow \Delta$ an annotated sequent labelling a vertex of \mathcal{T} . Then, $\text{b}(\Gamma) \Rightarrow \text{b}(\Delta)$ is valid.*

5.3 Uniform interpolation

With the proof system JS now fixed, we present the statement of *uniform interpolation* that will be proved. The *vocabulary* of a formula φ , in symbols $\text{Voc}(\varphi)$, is the set of modal actions and literals, other than \top and \perp , occurring in φ . The vocabulary of a set of formulas Φ is $\text{Voc}(\Phi) := \bigcup_{\varphi \in \Phi} \text{Voc}(\varphi)$.

5.3.1. THEOREM (Uniform interpolation). *Let Γ be a finite well-named set of closed, guarded formulas, and $V \subseteq \text{Voc}(\Gamma)$. There exists a formula ι such that:*

- (i) $\text{Voc}(\iota) \subseteq V$;
- (ii) $\text{JS} \vdash \Gamma \Rightarrow \iota$;
- (iii) for every Δ such that the sequent $\Gamma \Rightarrow \Delta$ is closed, well-named and guarded, and $\text{Voc}(\Delta) \cap \text{Voc}(\Gamma) \subseteq V$, if $\text{JS} \vdash \Gamma \Rightarrow \Delta$, then $\text{JS} \vdash \iota \Rightarrow \Delta$.

We call the formula ι of Theorem 5.3.1 the (*uniform*) *interpolant* of Γ relative to V . Mention of V may be suppressed when clear from context.

The *Craig interpolation* property is a special case of uniform interpolation:

5.3.2. COROLLARY (Craig interpolation). *For every closed, well-named and guarded sequent $\Gamma \Rightarrow \Delta$, if $\text{JS} \vdash \Gamma \Rightarrow \Delta$, then there is a formula ι such that $\text{Voc}(\iota) \subseteq \text{Voc}(\Gamma) \cap \text{Voc}(\Delta)$, $\text{JS} \vdash \Gamma \Rightarrow \iota$, and $\text{JS} \vdash \iota \Rightarrow \Delta$.*

Proof. Apply Theorem 5.3.1 with $V := \text{Voc}(\Gamma) \cap \text{Voc}(\Delta)$. ■

The remainder of this section is concerned with the proof of Theorem 5.3.1.

The uniform interpolant for Γ will be designed to encode all information relevant to proofs of sequents $\Gamma \Rightarrow \Delta$ for Δ satisfying the aforementioned vocabulary restriction. This will be achieved by expressing as a formula of the μ -calculus an ‘interpolation template’ for the sequent $\Gamma \Rightarrow \emptyset$, essentially a proof-search for $\Gamma \Rightarrow \Delta$, where Δ is ‘generic’.

5.3.3. DEFINITION (Interpolation template). Let Γ be a finite well-named set of closed and guarded formulas, and let $V \subseteq \text{Voc}(\Gamma)$. An *interpolation template for Γ* (and V) is a normal derivation \mathcal{T}_Γ of $\Gamma \Rightarrow \emptyset$ such that:

- (i) for every $v \in T_\Gamma$, if there is a vertex $u <_{T_\Gamma} v$ such that u and v are labelled by the same sequent, then v is a leaf.
- (ii) the branching rule \square^* is applied in \mathcal{T} instead of $L\square$ and $R\square$.⁷

$$\frac{\{\Theta_{i,\alpha^b} \dashv \Gamma_i, \alpha^b \Rightarrow \mid i \leq n, \alpha^b \in \Pi_i\} \quad \{\Theta_i \dashv \Gamma_i \Rightarrow \mid i \leq n, \mathbf{a}_i \in V\} \quad \emptyset \Rightarrow}{\Theta \dashv [\mathbf{a}_0]\Gamma_0, \langle \mathbf{a}_0 \rangle \Pi_0, \dots, [\mathbf{a}_n]\Gamma_n, \langle \mathbf{a}_n \rangle \Pi_n, \Sigma \Rightarrow \emptyset},$$

where $\mathbf{a}_0, \dots, \mathbf{a}_n$ are distinct modal actions, each Θ_{i,α^b} (Θ_i) is the restriction of Θ to the names occurring in Γ_i, α^b (respectively, Γ_i), Σ is a (non-axiomatic) set of literals, and we have omitted the empty consequents of the premises.

As mentioned previously, an interpolation template is, informally, a proof-search for the sequent $\Gamma \Rightarrow \Delta$, where Δ is a ‘generic’, unspecified consequent. The uniform interpolant will be a formula representation of the template.

Unlike $L\square$ and $R\square$, rule \square^* is branching and involves three kinds of premise:

- sequents $\Theta_{i,\alpha^b} \dashv \Gamma_i, \alpha^b \Rightarrow \emptyset$, for $i \leq n$ and $\alpha^b \in \Pi_i$, called *active premises*;
- sequents $\Theta_i \dashv \Gamma_i \Rightarrow \emptyset$ for $i \leq n$ and $\mathbf{a}_i \in V$, called *passive premises*;
- the *trivial premise* $\emptyset \Rightarrow \emptyset$.

The active and passive premises encode instances of $L\square$ and $R\square$, respectively, assuming an appropriate (but unspecified) instantiation of the consequent. The trivial premise corresponds to an instance of $R\square$ for an action label in $\text{Act} \setminus V$, as any such application of $R\square$ yields a premise with empty antecedent. In practice, the trivial premise may be safely ignored because it takes no part in the construction of the uniform interpolant; its presence is merely a technical convenience for mapping paths through a proof of $\Gamma \Rightarrow \Delta$ onto the interpolation template.

A leaf of an interpolation template is *empty* if it has label $\emptyset \Rightarrow \emptyset$. By construction, empty leaves correspond exactly to trivial premises of \square^* .

The notion of successful repeat is defined for interpolation templates as expected from Definition 5.2.2.

The restriction to normal proofs has the effect that interpolation templates can be assumed to be finite. More precisely, conditions (iii) and (iv) of Definition 5.2.5 ensure that a maximal path through an interpolation template reaches either an axiom, the empty sequent, or a sequent which is repeated on the path. The argument is identical to the proof of termination of proof search in [76, 144].

⁷We assume that Definition 5.2.5 is generalised to derivations with \square^* . No additional restrictions are necessary to accommodate this rule.

5.3.4. PROPOSITION. *For every finite, well-named set of closed and guarded formulas Γ , and every $V \subseteq \text{Voc}(\Gamma)$, there is a finite interpolation template for Γ and V .*

Henceforth we assume a fixed finite well-named set Γ of closed guarded formulas, a fixed $V \subseteq \text{Voc}(\Gamma)$, and a fixed finite interpolation template \mathcal{T}_Γ for Γ and V given by Proposition 5.3.4.

We need another result which follows from the completeness arguments in [76, 144]. Informally, if $\Gamma \Rightarrow \Delta$ is valid, then we can build a normal proof of $\Gamma \Rightarrow \Delta$ whose left-fragment is essentially \mathcal{T}_Γ . More precisely:

5.3.5. LEMMA. *Let $\Gamma \Rightarrow \Delta$ be a valid closed, well-named and guarded sequent, with Δ such that $\text{Voc}(\Gamma) \cap \text{Voc}(\Delta) \subseteq V$. There exists a normal proof \mathcal{T} of $\Gamma \Rightarrow \Delta$ such that for every finite path $(v_i)_{i \leq N}$ through \mathcal{T} there is a sequence of vertices $(u_i)_{i \leq N' \leq N}$ of \mathcal{T}_Γ and a strictly increasing sequence of natural numbers $(k_i)_{i \leq N'}$ such that the following hold for each $i < N'$, where $\Theta_i \dashv \Gamma_i \Rightarrow \Delta_i$ is the label of v_i :*

- (i) *either u_{i+1} is an immediate successor of u_i , or u_i is a repeat and u_{i+1} is an immediate successor of the companion of u_i , or u_i is an empty leaf and $u_{i+1} = u_{i+2} = \dots = u_{N'} = u_i$;*
- (ii) *u_i has label $\Theta'_{k_i} \dashv \Gamma_{k_i} \Rightarrow \emptyset$, where Θ'_{k_i} is the restriction of Θ_{k_i} to names for variables in Γ_{k_i} ;*
- (iii) *for each $j < k_{i+1} - k_i$, we have $\Theta'_{k_i+j} = \Theta'_{k_i}$ and $\Gamma_{k_i+j} = \Gamma_{k_i}$.*

We are now ready to construct the uniform interpolant for Γ relative to V . The definition proceeds in two stages. First, we assign to each $u \in T_\Gamma$ a formula ι_u , called the *pre-interpolant for u* . These are defined by recursion through \mathcal{T}_Γ , from the leaves downwards to the root. Second, by considering the collection of all pre-interpolants we isolate a uniform interpolant for Γ .

For every set of formulas Π , let $\ell_V(\Pi)$ be the collection of literals, other than \top and \perp , occurring in $V \cap \Pi$, without their annotations.

Empty leaves are assigned pre-interpolant \top . Every instance of ax_\perp , say with label $\Theta \dashv \Pi, \perp^\alpha \Rightarrow \emptyset$, is associated pre-interpolant $\perp \wedge \ell_V(\Pi)$. Every repeat $l \in \text{Rep}_{\mathcal{T}_\Gamma}$ is pre-interpolated by a unique, fresh variable x_l . We refer to each x_l as an *interpolation variable*.

Once we have decided on pre-interpolants for the leaves of \mathcal{T}_Γ , it is convenient to deal with interpolants for some instances of \square^* as a special case, before the more general recursive construction. If $u \in T_\Gamma$ is the conclusion of an application of \square^* in \mathcal{T}_Γ such that one of the active premises is provable in JS as a plain sequent, let $\iota_u := \perp$.⁸ We call such instances of \square^* *trivial*.

⁸See Remark 5.3.6 below.

With the trivial instances of \Box^* pre-interpolated, we proceed with the recursive construction. Suppose that $u \in T_\Gamma$ has not yet been assigned a pre-interpolant. If there is a $v <_{T_\Gamma} u$ such that v is the conclusion of a trivial instance of \Box^* , let $\iota_u := \perp$.

Suppose now that u is the conclusion of a non-trivial application of \Box^* in T_Γ , say:

$$\frac{\{\Theta_{i,\alpha^b} \dashv \Gamma_i, \alpha^b \Rightarrow \mid i \leq n, \alpha^b \in \Pi_i\} \quad \{\Theta_i \dashv \Gamma_i \Rightarrow \mid i \leq n, \mathbf{a}_i \in V\} \quad \emptyset \Rightarrow}{\Theta \dashv [\mathbf{a}_0]\Gamma_0, \langle \mathbf{a}_0 \rangle \Pi_0, \dots, [\mathbf{a}_n]\Gamma_n, \langle \mathbf{a}_n \rangle \Pi_n, \Sigma \Rightarrow \emptyset}$$

Let the actions be ordered such that $V = \{\mathbf{a}_0, \dots, \mathbf{a}_{k-1}\}$ for some $k \leq n$. For each $i \leq n$ and $\alpha^b \in \Pi_i$, let v_{i,α^b} be the immediate successor of u for the active premise $\Theta_{i,\alpha^b} \dashv \Gamma_i, \alpha^b \Rightarrow \emptyset$. And for each $i \leq n$ with $\mathbf{a}_i \in V$, let v_i the immediate successor for the passive premise $\Theta_i \dashv \Gamma_i \Rightarrow \emptyset$. We let

$$\iota_u := \ell_V(\Sigma) \wedge \bigwedge_{i < k} \left([\mathbf{a}_i]_{\iota_{v_i}} \wedge \bigwedge_{\alpha^b \in \Pi_i} \langle \mathbf{a}_i \rangle_{\iota_{v_{i,\alpha^b}}} \right).$$

Restricting to $i < k$ ensures that modal actions in pre-interpolants are all in V .

Suppose now that u is the conclusion of a (non-axiomatic) rule R other than \Box^* . If $R \in \{\mathbf{L}\wedge, \mathbf{L}\sigma_0, \mathbf{L}\mu, \mathbf{L}\nu, \mathbf{L}\text{thin}, \mathbf{L}\text{res}\}$, we let $\iota_u := \iota_v$, where v is the unique immediate successor of u . Finally, suppose that $R = \mathbf{L}\vee$, and let $\Theta \dashv \Pi \Rightarrow \emptyset$ be the label of u . We let $\iota_u := \ell_V(\Pi) \wedge (\iota_{v_1} \vee \iota_{v_2})$, where v_1, v_2 are the two immediate successors of u .

We thus have a pre-interpolant ι_u for each $u \in T_\Gamma$. By construction, ι_u is a quantifier-free formula in the variables $\mathbf{x}_{l_1}, \dots, \mathbf{x}_{l_k}$, where l_1, \dots, l_k are the repeats in $[u, \rightarrow)_{T_\Gamma}$. Additionally, either $\iota_u \in \{\top, \perp\} \cup \mathbf{Var}$, or ι_u is of the form $\ell_V(\Pi) \wedge \iota$, where Π is the set of formulas in the label of u .

Fix an enumeration (l_1, \dots, l_n) of \mathbf{Rep}_{T_Γ} such that the following hold:

- (i) if $\mathbf{inv}(l_i) \preceq \mathbf{inv}(l_j)$, then $i \leq j$;
- (ii) if $\mathbf{inv}(l_i) = \mathbf{inv}(l_j)$ and $i < j$, then either l_i is successful, or l_j is unsuccessful.

Define $p_{l_i} := 2i - 1$ if l_i is successful, and $p_{l_i} := 2i$ otherwise.

Let r be the root of T_Γ . The uniform interpolant for Γ relative to V is defined as the formula represented by the modal equational system

$$\iota_\Gamma = (\iota_r, \mathcal{E}_\Gamma), \text{ where } \mathcal{E}_\Gamma := \{\mathbf{x}_l =_{p_l} \iota_{c_l} \mid l \in \mathbf{Rep}_{T_\Gamma}\}.$$

5.3.6. REMARK. In the construction of ι_Γ we distinguished between trivial and non-trivial instances of \Box^* by appealing to provability of the premises thereof (as plain sequents). This can be checked by carrying out a proof-search in accordance

with the constraints in Definition 5.2.5 (for the details, see the completeness proofs in [76, 144]).

Since \mathcal{T}_Γ is normal, it seems reasonable to suppose that provability of premises of \Box^* could be established by inspection of \mathcal{T}_Γ alone. This is indeed the case and it is not difficult to see; however, precisely because of the normality of \mathcal{T}_Γ , for all intents and purposes said inspection is a proof-search.

We now turn to the verification of the interpolant.

5.3.1 Verification

We verify that the formula ι_Γ built from the template \mathcal{T}_Γ is in fact a uniform interpolant for Γ (with respect to V). By construction, it is clear that $\mathbf{Voc}(\iota_\Gamma) \subseteq V$. It remains to see that $\mathbf{JS} \vdash \Gamma \Rightarrow \iota_\Gamma$ and that $\mathbf{JS} \vdash \iota_\Gamma \Rightarrow \Delta$, for an arbitrary Δ satisfying the conditions in Theorem 5.3.1.

5.3.7. REMARK. In what follows, we work in \mathbf{JS} with the equational presentation of ι_Γ given above. This is merely an abbreviation (which greatly improves readability, though), made possible by the unfolding rules of \mathbf{JS} and the fact that the flattening of an equational system commutes with Boolean connectives and modal operators.

In order to establish $\mathbf{JS} \vdash \Gamma \Rightarrow \iota_\Gamma$ we begin by inductively building a (possibly infinite) \mathbf{JS} derivation \mathcal{P}_Γ with conclusion $\Gamma \Rightarrow \iota_\Gamma$ from the interpolation template \mathcal{T}_Γ , together with a partial map $\tau : P_\Gamma \rightarrow T_\Gamma$ such that the following hold for every u in the domain of τ :

- (i) if $u <_{P_\Gamma}^0 v$ and v is in the domain of τ , then either $\tau(u) <_{T_\Gamma}^0 \tau(v)$, or $\tau(u) \in \mathbf{Rep}_{\mathcal{T}_\Gamma}$ and $\tau(v)$ is the companion of $\tau(u)$;
- (ii) $\tau(u)$ has a label of the form $\Theta \dashv \Lambda \Rightarrow \emptyset$, with $\Lambda \neq \emptyset$, and u has a label of the form $\Omega \dashv \Lambda \Rightarrow \iota_{\tau(u)}^a$, where Θ results from removing the names in Ω that occur in a .

Let the root of \mathcal{P}_Γ have label $\Gamma \Rightarrow \iota_\Gamma$ and map it via τ to the root of \mathcal{T}_Γ . For the inductive case, assume that $u \in P_\Gamma$ has been defined. We distinguish several cases:

- (i) $\tau(u)$ is an axiomatic leaf. Then, u is an instance of \mathbf{ax}_\perp and we stop.
- (ii) $\tau(u)$ is a repeat, say $\tau(u) = l \in \mathbf{Rep}_{\mathcal{T}_\Gamma}$ with companion c . By the inductive hypothesis, u has label $\Omega \dashv \Lambda \Rightarrow \mathbf{x}_l^a$. Apply $\mathbf{R}\mu$ or $\mathbf{R}\nu$ at u , obtaining $\Omega' \dashv \Lambda \Rightarrow \iota_c^{a'b}$, where $a' := a \setminus \mathbf{x}_l$, and $b = \emptyset$ if \mathbf{x}_l is of type μ and otherwise b is the first name for \mathbf{x}_l not occurring in Ω . Let v be the unique immediate successor of u . If \mathbf{Rres} is not applicable at t , let $\tau(t) := c$. Otherwise, apply \mathbf{Rres} at t and map both t and its unique immediate successor via τ to c .

- (iii) A rule R other than \square^* , LV is applied at $\tau(u)$. Apply R at u and map the unique immediate successor of u to the unique immediate successor of $\tau(u)$.
- (iv) Rule LV is applied at $\tau(u)$. The label of $\tau(u)$ is then of the form

$$\Theta \dashv \Lambda, (\varphi_0 \vee \varphi_1)^a \Rightarrow \emptyset,$$

where $(\varphi_0 \vee \varphi_1)^a$ is the principal formula. By the inductive hypothesis, u has label $\Omega \dashv \Lambda, (\varphi_0 \vee \varphi_1)^a \Rightarrow \ell_V(\Lambda) \wedge (\iota_0 \vee \iota_1)^b$, where ι_i is the pre-interpolant assigned to the immediate successor v_i of $\tau(u)$ with label $\Theta \dashv \Lambda, \varphi_i^a \Rightarrow \emptyset$. We apply the following rules at u :

$$\begin{array}{c} \text{wk} \frac{\Omega \dashv \Lambda, \varphi_0^a \Rightarrow \iota_0^b}{\Omega \dashv \Lambda, \varphi_0^a \Rightarrow \iota_0^b, \iota_1^b} \quad \text{wk} \frac{\Omega \dashv \Lambda, \varphi_1^a \Rightarrow \iota_1^b}{\Omega \dashv \Lambda, \varphi_1^a \Rightarrow \iota_0^b, \iota_1^b} \\ \text{LV} \frac{\text{wk} \frac{\Omega \dashv \Lambda, \varphi_0^a \Rightarrow \iota_0^b}{\Omega \dashv \Lambda, \varphi_0^a \Rightarrow \iota_0^b, \iota_1^b} \quad \text{wk} \frac{\Omega \dashv \Lambda, \varphi_1^a \Rightarrow \iota_1^b}{\Omega \dashv \Lambda, \varphi_1^a \Rightarrow \iota_0^b, \iota_1^b}}{\Omega \dashv \Lambda, (\varphi_0 \vee \varphi_1)^a \Rightarrow \iota_0^b, \iota_1^b} \\ \text{RV} \frac{\Omega \dashv \Lambda, (\varphi_0 \vee \varphi_1)^a \Rightarrow \iota_0^b, \iota_1^b}{\Omega \dashv \Lambda, (\varphi_0 \vee \varphi_1)^a \Rightarrow (\iota_0 \vee \iota_1)^b} \\ \text{R}\wedge \frac{\text{RV} \frac{\Omega \dashv \Lambda, (\varphi_0 \vee \varphi_1)^a \Rightarrow \iota_0^b, \iota_1^b}{\Omega \dashv \Lambda, (\varphi_0 \vee \varphi_1)^a \Rightarrow (\iota_0 \vee \iota_1)^b}}{\Omega \dashv \Lambda, (\varphi_0 \vee \varphi_1)^a \Rightarrow \ell_V(\Lambda) \wedge (\iota_0 \vee \iota_1)^b} \end{array}$$

The omitted vertices are all axiomatic. Let u_i be the vertex with label $\Omega \dashv \Lambda, \varphi_i^a \Rightarrow \iota_i^b$, for $i \leq 1$, and let $\tau(u_i) := v_i$.

- (v) Rule \square^* is applied at $\tau(u)$. In \mathcal{T}_τ we then have:

$$\square^* \frac{\{ \Theta_{i,\alpha^b} \dashv \Gamma_i, \alpha^b \Rightarrow \emptyset \mid i \leq n, \alpha^b \in \Pi_i \} \quad \{ \Theta_i \dashv \Gamma_i \Rightarrow \emptyset \mid i \leq n, \mathbf{a}_i \in V \}}{\Theta \dashv [\mathbf{a}_0]\Gamma_0, \langle \mathbf{a}_0 \rangle \Pi_0, \dots, [\mathbf{a}_n]\Gamma_n, \langle \mathbf{a}_n \rangle \Pi_n, \Sigma \Rightarrow \emptyset},$$

where we have omitted the empty leaf. Let v_{i,α^b} be the immediate successor of $\tau(u)$ for the active premise $\Theta_{i,\alpha^b} \dashv \Gamma_i, \alpha^b \Rightarrow \emptyset$, and v_i the immediate successor for the passive premise $\Theta_i \dashv \Gamma_i \Rightarrow \emptyset$. Also, let Φ be the antecedent of the label of $\tau(u)$. Let $\iota := \iota_{\tau(u)}$. By the inductive hypothesis, u has a label of the form $\Omega \dashv \Phi \Rightarrow \iota^a$. We distinguish two subcases.

Case 1: some active premise v_{i,α^b} of $\tau(u)$ is provable (as a plain sequent). Then, $\iota = \perp$. We proceed from u as follows:

$$\text{L}\square \frac{\Theta_{i,\alpha^b} \dashv \Gamma_i, \alpha^b \Rightarrow \emptyset}{\Omega \dashv [\mathbf{a}_0]\Gamma_0, \langle \mathbf{a}_0 \rangle \Pi_0, \dots, [\mathbf{a}_n]\Gamma_n, \langle \mathbf{a}_n \rangle \Pi_n, \Sigma \Rightarrow \perp^a}$$

Let u' be the unique immediate successor of u . Note that u' and v_{i,α^b} have the same label. Let \mathcal{T} be a JS proof of $\mathbf{b}(\Gamma_i), \alpha \Rightarrow \emptyset$. We insert a copy of \mathcal{T} on top of u' , propagating the annotations in u' upwards and ignoring them in the copy of \mathcal{T} . It is easy to see that every leaf in the resulting cone

$C := [u', \rightarrow]_{\mathcal{P}_\Gamma}$ is either axiomatic or a successful repeat with companion in C .

Case 2: no active premise of $\tau(u)$ is provable. Then, assuming the modal actions to be ordered such that $V = \{\mathbf{a}_0, \dots, \mathbf{a}_{k-1}\}$ for some $k \leq n$, we have:

$$\iota_u := \ell_V(\Sigma) \wedge \bigwedge_{i < k} \left([\mathbf{a}_i] \iota_{v_i} \wedge \bigwedge_{\alpha^b \in \Pi_i} \langle \mathbf{a}_i \rangle \iota_{v_i, \alpha^b} \right).$$

At u , we break down ι_u by successive applications of $R\wedge$, obtaining the following sequents: $\Omega \dashv \Phi \Rightarrow \ell$, for $\ell \in \ell_V(\Sigma)$; $\Omega \dashv \Phi \Rightarrow \langle \mathbf{a}_i \rangle \iota_{v_i, \alpha^b}^a$, for $i \leq n$ and $\alpha^b \in \Pi_i$; and $\Omega \dashv \Phi \Rightarrow [\mathbf{a}_i] \iota_{v_i}^a$, for $i < k$.

At each $\Omega \dashv \Phi \Rightarrow \langle \mathbf{a}_i \rangle \iota_{v_i, \alpha^b}^a$ we apply $L\Box$ with premise $\Omega' \dashv \Gamma_i, \alpha^b \Rightarrow \iota_{v_i, \alpha^b}^a$, and map the resulting vertex via τ to v_{i, α^b} .

Similarly, at each $\Omega \dashv \Phi \Rightarrow [\mathbf{a}_i] \iota_{v_i}^a$ we apply $R\Box$ with premise $\Omega'' \dashv \Gamma_i \Rightarrow \iota_{v_i}^a$, and map the resulting vertex via τ to v_i .

We are now ready to establish that $\text{JS} \vdash \Gamma \Rightarrow \iota_\Gamma$.

By construction, \mathcal{P}_Γ is a possibly infinite tree with conclusion $\Gamma \Rightarrow \iota_\Gamma$ and such that every leaf of \mathcal{P}_Γ is axiomatic or a successful repeat. We show that every infinite branch π in \mathcal{P}_Γ passes through a successful repeat, whence it follows that \mathcal{P}_Γ can be pruned into a JS proof of $\Gamma \Rightarrow \iota_\Gamma$.

Let (x_0, x_1, \dots) be the infinite sequence of interpolation variables unfolded, in that order, in π . For every $i < \omega$, let l_i be the repeat in \mathcal{T}_Γ pre-interpolated by x_i , and let c_i and p_i abbreviate c_{l_i} and p_{x_i} , respectively. Note that $l_i \triangleleft l_{i+1}$ for every $i < \omega$. Let $N \geq 0$ be such that all variables x_N, x_{N+1}, \dots are unfolded infinitely often in π , say in a tail π' of π such that all variables unfolded on π' are among x_N, x_{N+1}, \dots . By Proposition 1.2.11, there is some $k \geq N$ such that $l_k \preceq l_N, l_{N+1}, \dots$. Moreover, we may assume that $p_k < p_i$ for every $i \geq N$ with $x_k \neq x_i$.

Let $N \leq m < n$ be such that $x_m = x_n = x_k$, and let π'' be the finite subsequence of π' from the m -th unfolding of an interpolation variable up to, and including, the n -th one. We distinguish two cases.

Case 1: x_k is of type ν . The consequents at the first unfolding of x_k in π'' have the form $x_k^a, l_k^{a'x}$, where $\iota_k = \iota_{c_k}$, $a' := a|x_k$, and x is the first name for x_k not occurring in the control.

Suppose that Rres is not applied after the unfolding. Looking at the unfoldings of interpolation variables, the consequents in π'' then look like this:

$$x_k^a, l_k^{a'x}, \dots, x_{m+1}^{a'x}, l_{m+1}^{a'xb_{m+1}}, \dots, x_{n-1}^{a'xb_{n-2}}, l_{n-1}^{a'xb_{n-1}}, \dots, x_k^{a'xb_{n-1}}, l_k^{a'xx'}.$$

Note that $a'x$ is preserved due to the priority of x_k . By construction, the name x is reset at the last step, yielding a successful repeat with companion the second vertex in π'' .

Suppose now that **Rres** is applied after the first unfolding of x_k in π'' . In this case we have $a' = a''x'$ for some a'' , and the consequents in π'' look like this, again paying attention only at unfoldings of interpolation variables:

$$x_k^a, l_k^{a''x'}, l_k^{a''x'}, \dots, x_{n-1}^{a''x'b_{n-2}}, l_{n-1}^{a''x'b_{n-2}}, \dots, x_k^{a''x'b_{n-1}}, l_k^{a''x'b_{n-1}}.$$

As before, $a''x'$ is preserved due to the priority of x_k . By construction, x' is reset again at the last step, yielding a successful repeat with companion the third vertex in π'' .

Case 2: x_k is of type μ . The consequents at the first unfolding of x_k in π'' have the form $x_k^a, l_k^{a'}$, where $a' := a|x_k$. As **Rres** is not applied after the unfolding because no name is introduced, the consequents in π'' look like this, again paying attention only unfoldings of interpolation variables:

$$x_k^a, l_k^{a'}, \dots, x_{m+1}^{a'}, l_{m+1}^{a'b_{m+1}}, \dots, x_{n-1}^{a'b_{n-2}}, l_{n-1}^{a'b_{n-2}}, \dots, x_k^{a'b_{n-1}}, l_k^{a'}.$$

As before, a' is preserved due to the priority of x_k . We claim that this last repetition is successful with companion the second vertex in π'' .

Since x_k is a μ -variable, we know that $\text{inv}(l_k)$ is of the form wx , where x is preserved and reset in $[c_k, l_k]_{\mathcal{T}_\Gamma}$. Also, $\tau(\pi'')$ is a path through \mathcal{T}_Γ of the form

$$\tau(\pi'') = (l_k) \frown [c_k, l_{m+1}]_{\mathcal{T}_\Gamma} \frown \dots \frown [c_{n-2}, l_{n-1}]_{\mathcal{T}_\Gamma} \frown [c_{n-1}, l_k]_{\mathcal{T}_\Gamma}.$$

By Proposition 1.2.12, x is preserved throughout $\tau(\pi'')$ because $l_k \preceq l_{m+1}, \dots, l_n$, so x is also preserved throughout π'' . Finally, by Proposition 1.2.3 $[c_k, l_k]_{\mathcal{T}_\Gamma} \subseteq \tau(\pi'')$, whence x is reset in π'' .

We have thus established:

5.3.8. PROPOSITION. $\text{JS} \vdash \Gamma \Rightarrow \iota_\Gamma$.

It only remains to verify that $\text{JS} \vdash \iota_\Gamma \Rightarrow \Delta$ for a Δ given as in Theorem 5.3.1. In particular, $\text{Voc}(\Gamma) \cap \text{Voc}(\Delta) \subseteq V$.

Fix a normal JS proof \mathcal{P} of $\Gamma \Rightarrow \Delta$ given by Lemma 5.3.5, and let \mathcal{P}^\bullet be the result after identifying each repeat $l \in \text{Rep}_{\mathcal{P}}$ and its companion c_l .⁹ We define $\mathcal{T}_\Gamma^\bullet$ analogously.

If w is a prefix of the control of a sequent $\Theta \dashv \Pi \Rightarrow \Sigma$, we denote by w_L (w_R) the result after removing from w all names for variables in Σ (respectively, Π).

We inductively build a (possibly infinite) JS derivation \mathcal{P}_Δ with conclusion $\iota_\Gamma \Rightarrow \Delta$, together with partial maps $\gamma : \mathcal{P}_\Delta \rightarrow \mathcal{T}_\Gamma$ and $\delta : \mathcal{P}_\Delta \rightarrow \mathcal{P}^\bullet$ such that the following hold for every $u \in \mathcal{P}_\Delta$:

⁹Note the difference between \mathcal{P}^\bullet and \mathcal{P}° : the former identifies repeats and companions, whereas the latter does not.

- (i) The label of $\delta(u)$ is of the form $\Theta \dashv \Lambda \Rightarrow \Xi$, where Λ is such that $\gamma(u)$ has label $\Theta_L \dashv \Lambda \Rightarrow \emptyset$. If $\Lambda = \emptyset$, then u has label $\Theta \dashv \emptyset \Rightarrow \Xi$, and otherwise the label of u is of the form $\Omega \dashv \iota_{\gamma(u)}^a \Rightarrow \Xi$, where $\Omega_R = \Theta_R$.
- (ii) If π is a path through \mathcal{P}_Δ , then $\delta(\pi)$ is a path through \mathcal{P}^\bullet and $\gamma(\pi)$ is the path on \mathcal{T}_Γ° corresponding to $\delta(\pi)$ given by Lemma 5.3.5.¹⁰

Let the root of \mathcal{P}_Δ have label $\iota_\Gamma \Rightarrow \Delta$, and map it via γ to the root of \mathcal{T}_Γ and via δ to the root of \mathcal{P}^\bullet . For the inductive case, assume that $u \in \mathcal{P}_\Delta$ has been defined. We distinguish several cases, ordered below from greatest to lowest priority in the construction of \mathcal{P}_Δ :

- (i) $\delta(s)$ is a leaf. Note that $\delta(s)$ must be axiomatic by the definition of \mathcal{P}^\bullet and the fact that \mathcal{P} is a proof.

If $\delta(u)$ is an instance of \mathbf{ax}_ℓ for a literal $\ell \notin \{\perp, \top\}$, then $\ell \in \mathbf{Voc}(\Gamma) \cap \mathbf{Voc}(\Delta) \subseteq V$, whence by the construction of pre-interpolants we may apply $L\wedge$ finitely many times at u to reach an instance of \mathbf{ax}_ℓ .

If $\delta(u)$ is an instance of \mathbf{ax}_\top , then by the inductive hypothesis so is u . Otherwise, $\delta(u)$ is an instance of \mathbf{ax}_\perp and thus so is $\gamma(s)$, whence $\iota_{\gamma(s)} = \perp$ and u is also an instance of \mathbf{ax}_\perp .

- (ii) $\gamma(u)$ is an axiomatic leaf (so an instance of \mathbf{ax}_\perp). As before, u is an instance of \mathbf{ax}_\perp .
- (iii) $\gamma(u)$ is a repeat, say $\gamma(u) = l \in \mathbf{Rep}_{\mathcal{T}_\Gamma}$ with companion c . By the inductive hypothesis, u has label $\Omega \dashv x_i^a \Rightarrow \Xi$. We unfold x_i by applying $L\mu$ or $L\nu$ at u , obtaining $\Omega' \dashv \iota_c^{a'b} \Rightarrow \Xi$, where $a' := a \upharpoonright x_i$, and $b = \emptyset$ if x_i is of type ν and otherwise b is the first name for x_i not occurring in Ω . Let v be the unique immediate successor of u . If \mathbf{Lres} is not applicable at v , let $\delta(v) := \delta(u)$ and $\gamma(v) := c$. Otherwise, apply \mathbf{Lres} at v and map both v and its unique immediate successor to $\delta(u)$ via δ , and to c via γ .
- (iv) A rule R among $R\wedge$, $R\vee$, $R\sigma_0$, $R\mu$, $R\nu$, $R\text{thin}$, $R\text{res}$ and \mathbf{wk} is applied at $\delta(u)$. Note that \mathbf{wk} cannot affect the antecedent because it is not applied anywhere in \mathcal{T}_Γ . We apply R at u . For every immediate successor v of u , let $\delta(v)$ be the corresponding successor of $\delta(u)$ and set $\gamma(v) := \gamma(u)$.
- (v) A rule R among $L\wedge$, $L\sigma_0$, $L\mu$, $L\nu$, $L\text{thin}$, $L\text{res}$ is applied on the left at $\delta(u)$. By the inductive hypothesis, the label of u is of the form $\Omega \dashv \iota \Rightarrow \Xi$. We apply \mathbf{wk} at u with premise $\Omega \dashv \iota \Rightarrow \Xi$.¹¹ Let v be the unique immediate

¹⁰Strictly speaking, Lemma 5.3.5 yields paths on $\mathcal{T}_\Gamma^\bullet$. But, clearly, every such path determines a unique path through \mathcal{T}_Γ° .

¹¹This ‘degenerate’ application of \mathbf{wk} is needed for paths on \mathcal{P}_Δ to be mapped to paths on $\mathcal{T}_\Gamma^\bullet$ and \mathcal{P}^\bullet .

successor of u . We let $\gamma(v)$ ($\delta(v)$) be the unique immediate successor of $\gamma(u)$ (respectively, $\delta(u)$).

(vi) Rule LV is applied at $\delta(u)$. The label of $\delta(u)$ is of the form

$$\Theta \dashv \Lambda, (\varphi_0 \vee \varphi_1)^a \Rightarrow \Xi,$$

and by the inductive hypothesis u has label $\Omega \dashv \ell_V(\Lambda) \wedge (\iota_0 \vee \iota_1)^a \Rightarrow \Xi$ and LV is applied at $\gamma(u)$. We first apply L \wedge and wk at u to obtain the sequent $\Omega \dashv (\iota_0 \vee \iota_1)^a \Rightarrow \Xi$. Let u_0 and u_1 be the vertices thus created, with $u_0 <_{P_\Delta}^0 u_1$. We apply LV at u_1 and let v_i be the immediate successor of u_1 with label $\Omega \dashv \iota_i^a \Rightarrow \Xi$, for $i \leq 1$. Let $\gamma(u_0) := \gamma(u_1) := (u)$, $\delta(u_0) := \delta(u_1) := \delta(u)$. Finally, let $\delta(v_i)$ be the immediate successor of $\delta(u)$ with label $\Theta \dashv \Lambda, \varphi_i^a \Rightarrow \Xi$, and let $\gamma(v_i)$ be the immediate successor of $\gamma(u)$ with label $\Theta_L \dashv \Lambda, \varphi_i^a \Rightarrow \Xi$.

(vii) Rule R \square is applied at $\delta(u)$. Then, in \mathcal{P}^\bullet we have:

$$\text{R}\square \frac{\Theta' \dashv \Lambda \Rightarrow \varphi^e, \Xi}{\Theta \dashv [\mathbf{b}]\Lambda, \Phi \Rightarrow [\mathbf{b}]\varphi^e, \langle \mathbf{b} \rangle \Xi, \Psi}$$

Note that $\mathbf{b} \in \text{Voc}(\Delta)$. By the inductive hypothesis, \square^* is applied at $\gamma(u)$. If $\gamma(u)$ has pre-interpolant \perp , then by the inductive hypothesis u is an instance of ax_\perp and we stop, so assume that no active premise of $\gamma(u)$ is provable. In \mathcal{T}_Γ we have:

$$\square^* \frac{\{\Theta_{i,\alpha^c} \dashv \Gamma_i, \alpha^c \Rightarrow \emptyset \mid i \leq n, \alpha^c \in \Pi_i\} \quad \{\Theta_i \dashv \Gamma_i \Rightarrow \emptyset \mid i \leq n, \mathbf{a}_i \in V\}}{\Theta \dashv [\mathbf{a}_0]\Gamma_0, \langle \mathbf{a}_0 \rangle \Pi_0, \dots, [\mathbf{a}_n]\Gamma_n, \langle \mathbf{a}_n \rangle \Pi_n, \Sigma \Rightarrow \emptyset},$$

where we have omitted the empty leaf. Let v_{i,α^c} be the immediate successor of $\gamma(u)$ for the active premise $\Theta_{i,\alpha^c} \dashv \Gamma_i, \alpha^c \Rightarrow \emptyset$, and v_i the immediate successor for the passive premise $\Theta_i \dashv \Gamma_i \Rightarrow \emptyset$.

Assuming the modal actions to be ordered such that $V = \{\mathbf{a}_0, \dots, \mathbf{a}_{k-1}\}$ for some $k \leq n$, we have

$$\iota := \iota_{\gamma(s)} = \ell_V(\Sigma) \wedge \bigwedge_{i < k} \left([\mathbf{a}_i] \iota_{v_i} \wedge \bigwedge_{\alpha^c \in \Pi_i} \langle \mathbf{a}_i \rangle \iota_{v_{i,\alpha^c}} \right).$$

We distinguish two further cases:

a) $\mathbf{b} \in \text{Voc}(\Gamma)$. Since $\mathbf{b} \in \text{Voc}(\Delta)$, it follows that $\mathbf{b} = \mathbf{a}_i$ for some $i < k$. We apply the following rules at u , where the double line indicates finitely many applications of L \wedge and in the end an instance of wk to remove $\ell_V(\Sigma)$:

$$\text{R}\square \frac{\Omega' \dashv \iota_{v_i}^a \Rightarrow \varphi^e, \Xi}{\frac{\Omega \dashv \{\langle \mathbf{a}_i \rangle \iota_{v_i, \alpha^c}^a \mid i < k, \alpha^c \in \Pi_i\}, \{[\mathbf{a}_i] \iota_{v_i}^a \mid i < k\} \Rightarrow [\mathbf{b}] \varphi^e, \langle \mathbf{b} \rangle \Xi, \Psi}{\Omega \dashv \iota^a \Rightarrow [\mathbf{b}] \varphi^e, \langle \mathbf{b} \rangle \Xi, \Psi}}$$

We map the vertex at the top via δ to the unique immediate successor of $\delta(u)$, and via γ to v_i . All intermediate vertices are mapped to $\gamma(u)$ via γ and to $\delta(u)$ via δ .

- b) $b \notin \text{Voc}(\Gamma)$, i.e., $b \neq a_i$ for all $i \leq n$. We apply the following rules at u , where the double line is as in the previous case:

$$\text{R}\square \frac{\Omega' \dashv \emptyset \Rightarrow \varphi^e, \Xi}{\frac{\Omega \dashv \{\langle \mathbf{a}_i \rangle \iota_{v_i, \alpha^c}^a \mid i < k, \alpha^c \in \Pi_i\}, \{[\mathbf{a}_i] \iota_{v_i}^a \mid i < k\} \Rightarrow [\mathbf{b}] \varphi^e, \langle \mathbf{b} \rangle \Xi, \Psi}{\Omega \dashv \iota^a \Rightarrow [\mathbf{b}] \varphi^e, \langle \mathbf{b} \rangle \Xi, \Psi}}$$

We map the vertex at the top via δ to the unique immediate successor of $\delta(u)$, and via γ to the immediate successor of $\gamma(u)$ for the trivial premise $\emptyset \Rightarrow \emptyset$. All intermediate vertices are mapped to $\gamma(u)$ via γ and to $\delta(u)$ via δ .

- (viii) Rule $\text{L}\square$ is applied at $\delta(u)$. Then, in \mathcal{P}^\bullet we have:

$$\text{L}\square \frac{\Theta' \dashv \Lambda, \varphi^e \Rightarrow \Xi}{\Theta \dashv [\mathbf{b}] \Lambda, \Phi, \langle \mathbf{b} \rangle \varphi^e \Rightarrow \langle \mathbf{b} \rangle \Xi, \Psi}$$

By the inductive hypothesis, \square^* is applied at $\gamma(u)$. If $\gamma(u)$ has pre-interpolant \perp , then by the inductive hypothesis u is an instance of \mathbf{ax}_\perp and we stop. Assume, then, that no active premise of $\gamma(u)$ is provable.

In \mathcal{T}_Γ we have:

$$\square^* \frac{\{\Theta_{i, \alpha^c} \dashv \Gamma_i, \alpha^c \Rightarrow \emptyset \mid i \leq n, \alpha^c \in \Pi_i\} \{\Theta_i \dashv \Gamma_i \Rightarrow \emptyset \mid i \leq n, \mathbf{a}_i \in V\}}{\Theta \dashv [\mathbf{a}_0] \Gamma_0, \langle \mathbf{a}_0 \rangle \Pi_0, \dots, [\mathbf{a}_n] \Gamma_n, \langle \mathbf{a}_n \rangle \Pi_n, \Sigma \Rightarrow \emptyset},$$

where we have omitted the empty leaf. Let v_{i, α^c} be the immediate successor of $\gamma(u)$ for the active premise $\Theta_{i, \alpha^c} \dashv \Gamma_i, \alpha^c \Rightarrow \emptyset$, and v_i the immediate successor for the passive premise $\Theta_i \dashv \Gamma_i \Rightarrow \emptyset$.

Assuming the modal actions to be ordered such that $V = \{\mathbf{a}_0, \dots, \mathbf{a}_{k-1}\}$ for some $k \leq n$, we have

$$\iota := \iota_{\gamma(u)} = \ell_V(\Sigma) \wedge \bigwedge_{i < k} \left([\mathbf{a}_i] \iota_{v_i} \wedge \bigwedge_{\alpha^c \in \Pi_i} \langle \mathbf{a}_i \rangle \iota_{v_i, \alpha^c} \right).$$

Note that, by the inductive hypothesis, $\mathbf{b} = \mathbf{a}_i$ for some $i \leq n$ and $\Lambda = \Gamma_i$. We claim that, moreover, $\mathbf{a}_i \in V$. Suppose not. Then, $\mathbf{a}_i \notin \text{Voc}(\Delta)$

because $\text{Voc}(\Delta) \cap \text{Voc}(\Gamma) \subseteq V$, so the immediate successor of $\delta(u)$ has label $\Theta_{i,\varphi^e} \dashv \Gamma_i, \varphi^e \Rightarrow \emptyset$. By Lemma 5.2.7 and Theorem 5.2.4, $\mathbf{b}(\Gamma_i), \varphi \Rightarrow \emptyset$ is provable in JS, contradicting the assumption that no active premise of $\gamma(u)$ is provable (as a plain sequent). Hence, $\mathbf{b} = \mathbf{a}_i \in V$.

We apply the following rules at u , where the double line is as in the previous case:

$$\text{wk} \frac{\text{L}\square \frac{\Omega'' \dashv \iota_{v_i, \varphi^e}^a \Rightarrow \Xi}{\Omega' \dashv \langle \mathbf{b} \rangle \iota_{v_i, \varphi^e}^a \Rightarrow \langle \mathbf{b} \rangle \Xi, \Psi}}{\Omega \dashv \{ \langle \mathbf{a}_i \rangle \iota_{v_i, \alpha^c}^u \mid i < k, \alpha^c \in \Pi_i \}, \{ [\mathbf{a}_i] \iota_{v_i}^a \mid i < k \} \Rightarrow \langle \mathbf{b} \rangle \Xi, \Psi}}{\Omega \dashv \iota^a \Rightarrow \langle \mathbf{b} \rangle \Xi, \Psi}$$

We map the vertex at the top via δ to the unique immediate successor of $\delta(u)$, and via γ to v_{i,φ^e} . All intermediate vertices are mapped to $\gamma(u)$ via γ and to $\delta(u)$ via δ .

The following proposition captures the relation between the repeats in \mathcal{P} and the repeats in \mathcal{T}_Γ :

5.3.9. PROPOSITION. *Let $l \in P$ be a repeat with companion c . If $u, v \in P_\Delta$ are such that $u <_{P_\Delta} v$ and $\delta(u) = \delta(v) = l$, then there are repeats l_0, l_1, \dots, l_n in \mathcal{T}_Γ such that:*

- (i) $l_0 \triangleleft l_1 \triangleleft \dots \triangleleft l_n$;
- (ii) $\gamma([u, v]_{P_\Delta}) = [\gamma(u), l_0]_{\mathcal{T}_\Gamma} \frown (c_0, l_1)_{\mathcal{T}_\Gamma} \frown \dots \frown (c_{n-1}, l_n)_{\mathcal{T}_\Gamma} \frown (c_n, \gamma(v))_{\mathcal{T}_\Gamma}$,¹²

where c_i abbreviates c_{l_i} .

Proof. By construction, the labels of $\gamma(u)$ and $\gamma(v)$ are identical, so either $\gamma(v)$ is a repeat with companion $\gamma(u)$, or γ loops back from some repeat to its companion when going from $\gamma(u)$ to $\gamma(v)$, because every branch in \mathcal{T}_Γ ends at the first encounter of an axiom or a repeated label. In either case, there is some vertex in $[u, v]_{P_\Delta}$ whose γ -image is a repeat in \mathcal{T}_Γ .

Let (u_0, \dots, u_n) be the list of all such vertices in order of appearance in $[u, v]_{P_\Delta}$. Let $l_i := \gamma(u_i)$, and let c_i be the companion of l_i . Then, $\gamma([u, v]_{P_\Delta})$ is the path

$$[\gamma(u), \gamma(u_0)]_{\mathcal{T}_\Gamma} \frown (\gamma(u_0), \gamma(u_1))_{\mathcal{T}_\Gamma} \frown \dots \frown (\gamma(u_{n-1}), \gamma(u_n))_{\mathcal{T}_\Gamma} \frown (\gamma(u_n), \gamma(v))_{\mathcal{T}_\Gamma},$$

¹²For simplicity, we implicitly assume that multiple contiguous occurrences of a vertex in $\gamma([u, v]_{P_\Delta})$ are counted and treated as a single occurrence. This is clearly innocuous.

whence by construction of γ we have:

$$\gamma([u, v]_{\mathcal{P}_\Delta}) = [\gamma(u), l_0]_{\mathcal{T}_\Gamma} \hat{\wedge} (c_0, l_1]_{\mathcal{T}_\Gamma} \hat{\wedge} \cdots \hat{\wedge} (c_{n-1}, l_n]_{\mathcal{T}_\Gamma} \hat{\wedge} (c_n, \gamma(v)]_{\mathcal{T}_\Gamma}.$$

Clearly, $l_0 \triangleleft l_1 \triangleleft \cdots \triangleleft l_n$. ■

Finally, we are ready to show that $\text{JS} \vdash \iota_\Gamma \Rightarrow \Delta$.

By construction, \mathcal{P}_Δ is a possibly infinite JS derivation with conclusion $\iota_\Gamma \Rightarrow \Delta$. Moreover, all leaves of \mathcal{P}_Δ are axiomatic. It thus suffices to show that any infinite branch π of \mathcal{P}_Δ passes through a successful repeat.

We say that a repeat $l \in \text{Rep}_{\mathcal{P}}$ is encountered in π if π passes through some vertex u such that $\delta(u) = l$.

Case 1: there is a tail π' of π where all vertices have labels with empty antecedent. This is the case if, and only if, $\gamma(\pi)$ is finite (terminating at a trivial premise $\emptyset \Rightarrow \emptyset$ of some \square^* vertex). Let (l_0, l_1, \dots) be the infinite sequence of repeats in \mathcal{P} that are encountered, in that order, in π' . Since π' is infinite, there is some $N \geq 0$ such that l_N, l_{N+1}, \dots are all encountered in π' infinitely often and every repeat of \mathcal{P} encountered in π' is among l_N, l_{N+1}, \dots . Note that we have $l_N \triangleleft l_{N+1} \triangleleft \cdots$, so by Proposition 1.2.11 there is some $k \geq N$ such that $l_k \preceq l_N, l_{N+1}, \dots$, that is, $\text{inv}(l_k)$ is a prefix of $\text{inv}l_i$ for all $i \geq N$.

Let π'' be the finite subsequence of π' from the first encounter of l_k in π' up to, and including, the next encounter of l_k in π' . Since every label of a vertex in π has empty antecedent, the first and last vertices of π' are labelled with the same annotated sequent, so it suffices to show that some name is preserved and reset in π' . Let us write $\text{inv}l_k = wx$, where x is preserved and reset in $[c_k, l_k]_{\mathcal{P}}$. By Proposition 1.2.12, $\text{inv}l_k$ is preserved in $[c_i, l_i]_{\mathcal{P}}$ for every $i \geq N$, whence x is preserved in π' . And $[c_k, l_k]_{\mathcal{P}} \subseteq \pi'$ by Proposition 1.2.3, so x is also reset in π' and we are done.

Case 2: all the vertices in π have labels with non-empty antecedent. In this case, $\gamma(\pi)$ is an infinite path through \mathcal{T}_Γ° , so let (x_0, x_1, \dots) be the infinite sequence of interpolation variables unfolded, in that order, in π . Let l_i be the leaf in \mathcal{T}_Γ pre-interpolated by x_i , and let p_i and c_i abbreviate, respectively, p_{x_i} and c_{l_i} . As before, we have $l_i \triangleleft l_{i+1}$ for every $i < \omega$. Let $N \geq 0$ be such that all variables x_N, x_{N+1}, \dots are unfolded infinitely often in π , say on a tail π' of π such that every interpolation variable unfolded on π' is among x_N, x_{N+1}, \dots . By Proposition 1.2.11, there is some $k \geq N$ such that $l_k \preceq l_N, l_{N+1}, \dots$ and $p_k < p_i$ for every $i \geq N$ with $x_i \neq x_k$. We distinguish two further cases.

Case 2.1: x_k is of type μ . Since \mathcal{P} is finite, there are $k \leq m < n$ such that:

- (i) $x_k = x_m = x_n$;
- (ii) the m -th and n -th unfoldings of interpolation variables in π occur at vertices whose labels have identical consequent.

Let π'' be the finite subsequence of π from the m -th unfolding of an interpolation variable up to, and including, the n -th one. The antecedents at the first unfolding of \mathbf{x}_k in π'' have the form $\mathbf{x}_k^a, l_k^{a'x}$, where l_k is the pre-interpolant of c_k , $a' := a[\mathbf{x}_k]$, and x is a name for \mathbf{x}_k . Suppose that a' does not end in a name for \mathbf{x}_k , so that Lres is not applied immediately after the unfolding. Then, the antecedents in π'' look as follows, again paying attention only at unfoldings of interpolation variables:

$$\mathbf{x}_k^a, l_k^{a'x}, \dots, \mathbf{x}_{m+1}^{a'x}, l_{m+1}^{a'xb_1}, \dots, \mathbf{x}_{n-1}^{a'xb'_{n-m-2}}, l_{n-1}^{a'xb_{n-m-1}}, \dots, \mathbf{x}_k^{a'xb'_{n-m-1}}, l_k^{a'xx'}.$$

The sequence $a'x$ is preserved because $p_k < p_{m+1}, \dots, p_{n-1}$. By construction, x is reset at the last step, yielding a successful repeat with companion the second vertex in π'' . The case when a' ends in a name for \mathbf{x}_i is similar.

Case 2.2: \mathbf{x}_k is of type ν . Let (g_0, g_1, \dots) be the infinite sequence of repeats in \mathcal{P} that are encountered infinitely often, in that order, in a tail π'' of π' such that every repeat in \mathcal{P} encountered in π'' is among g_0, g_1, \dots . We have $g_i \triangleleft g_{i+1}$ for all $i < \omega$, so by Proposition 1.2.11 there is a $j \geq 0$ such that $g_j \preccurlyeq g_0, g_1, \dots$. Let us write $\text{inv}g_j = wx$ and denote by $(wx)_L$ the restriction of wx to names for variables in Γ .

We claim that x is a name for a variable in the consequent (i.e., in Δ). Suppose not. Let $\pi''' := \pi'' \cap \delta^{-1}(\bigcup_{i < \omega} [c_{g_i}, g_{i+1}]_{\mathcal{P}})$. Observe that π''' is a tail of π'' . By Proposition 1.2.12, wx is preserved in $[c_{g_i}, g_{i+1}]_{\mathcal{P}}$ for every $i \geq 0$, so wx is preserved in π''' .

Since all variables $\mathbf{x}_N, \mathbf{x}_{N+1}, \dots$ are unfolded infinitely often in π''' , by Proposition 1.2.3 we have $[c_i, l_i]_{\mathcal{T}_\Gamma} \subseteq \gamma(\pi''')$ for every $i \geq N$. Hence, by assumption $(wx)_L = w_Lx$ is preserved in $[c_i, l_i]_{\mathcal{T}_\Gamma}$ for all $i \geq N$. In particular, w_Lx is preserved in $[c_k, l_k]_{\mathcal{T}_\Gamma}$. So, since \mathbf{x}_k is of type ν , w_Lx is a prefix of $\text{inv}_{\mathcal{T}_\Gamma}(l_k)$, where $\text{inv}_{\mathcal{T}_\Gamma}$ denotes the invariant map for \mathcal{T}_Γ .

We shall now find a successful repeat $l_i \in \text{Rep}_{\mathcal{T}_\Gamma}$ such that $\text{inv}_{\mathcal{T}_\Gamma}(l_i)$ is a prefix of w_Lx , and thus of $\text{inv}_{\mathcal{T}_\Gamma}(l_k)$ as well. This will contradict the minimality of p_k and conclude the proof that x must name a variable in Δ . We find such a repeat l_i by tracing onto \mathcal{T}_Γ° a path in \mathcal{P}° starting and ending at the repeat g_j . We know that x is preserved and reset in $[c_{g_j}, g_j]_{\mathcal{P}}$, and we are assuming that x names a variable in Γ . Hence, x will also be preserved and reset on the path traced on the template.

Let $u, v \in \pi'''$ be such that $u <_{\mathcal{P}_\Delta} v$ and $\delta(u) = \delta(v) = g_j$. By Proposition 5.3.9, there are repeats l'_0, \dots, l'_n in \mathcal{T}_Γ such that:

- (i) $l'_0 \triangleleft l'_1 \triangleleft \dots \triangleleft l'_n$;
- (ii) $\gamma([u, v]_{\mathcal{P}_\Delta}) = [\gamma(u), l'_0]_{\mathcal{T}_\Gamma} \hat{\ } (c'_0, l'_1]_{\mathcal{T}_\Gamma} \hat{\ } \dots \hat{\ } (c'_{n-1}, l'_n]_{\mathcal{T}_\Gamma} \hat{\ } (c'_n, \gamma(v)]_{\mathcal{T}_\Gamma}$,

where c'_i abbreviates $c_{l'_i}$. For every $0 \leq i \leq n$, let \mathbf{y}_i be the interpolation variable assigned to l'_i . Note that every \mathbf{y}_i occurs in the sequence $(\mathbf{x}_N, \mathbf{x}_{N+1}, \dots)$

Since $\gamma([u, v]_{\mathcal{P}_\Delta}) \subseteq \gamma(\pi''')$, $w_L x$ is preserved in $\gamma([u, v]_{\mathcal{P}_\Delta})$. And x is reset in $[c_{g_j}, g_j]_{\mathcal{P}}$, which is included in $\delta([u, v]_{\mathcal{P}_\Delta})$ by Proposition 1.2.3. Hence, x is reset at some vertex $p \in \gamma([u, v]_{\mathcal{P}_\Delta})$. Let us see that there is some $i \geq N$ such that $p \in [c_i, l_i]_{\mathcal{T}_\Gamma}$.

Suppose that $p \in [\gamma(u), l'_0]_{\mathcal{T}_\Gamma}$. If $\gamma(v) \notin [p, \rightarrow)_{\mathcal{T}_\Gamma}$, then there is some $i \geq N$ such that $l_i \in [p, \rightarrow)_{\mathcal{T}_\Gamma}$ and $c_i \notin [p, \rightarrow)_{\mathcal{T}_\Gamma}$, so $p \in [c_i, l_i]_{\mathcal{T}_\Gamma}$. And if $\gamma(v) \in [p, \rightarrow)_{\mathcal{T}_\Gamma}$, then $\gamma(v)$ is a repeat with companion $\gamma(u)$ and, since $\gamma(v) \in \gamma([u, v]_{\mathcal{T}_\Gamma}) \subseteq \gamma(\pi''')$, there is some $i \geq N$ such that x_i is the interpolation variable assigned to $\gamma(v)$, whence again $p \in [c_i, l_i]_{\mathcal{T}_\Gamma}$. Now suppose that $p \in (c'_m, l'_{m+1}]_{\mathcal{T}_\Gamma}$ for some $0 \leq m \leq n-1$. If $p \in [c'_m, l'_m]_{\mathcal{T}_\Gamma}$ we are done because $y_m = x_i$ for some $i \geq N$. Otherwise $l'_m \notin [p, \rightarrow)_{\mathcal{T}_\Gamma}$ and thus, since we unfold y_m infinitely often in π''' , there must be some $i \geq N$ such that $l_i \in [p, \rightarrow)_{\mathcal{T}_\Gamma}$ and $c_i \notin [p, \rightarrow)_{\mathcal{T}_\Gamma}$, whence $p \in [c_i, l_i]_{\mathcal{T}_\Gamma}$. Finally, suppose that $p \in (c'_n, \gamma(v)]_{\mathcal{T}_\Gamma}$. We reason as in the previous case but with c'_n and l'_n in place of c'_m and l'_m , respectively.

We have thus found some $i \geq N$ such that x is preserved and reset in $[c_i, l_i]_{\mathcal{T}_\Gamma}$. Hence, l_i is a successful repeat in \mathcal{T}_Γ , so $\text{inv}_{\mathcal{T}_\Gamma}(l_i)$ is a prefix of $w_L x$ and $p_i < p_k$. This contradicts the minimality of p_k . Therefore, x cannot be a name for a variable in Γ .

Finally, we show that x witnesses a successful repetition on the infinite branch π . Since \mathcal{P} is finite, let $u, v \in \pi'''$ be such that:

- (i) $u <_{\mathcal{P}_\Delta} v$;
- (ii) x_k is unfolded at u and at v ;
- (iii) the labels of u and v have identical consequent.
- (iv) x is reset in $[u, v]_{\mathcal{P}_\Delta}$.

The unfolding of x_k at u has the form $x_k^a, l_k^{a'}$, where $a' := a|x_k$, and the unfolding of x_k at v looks like $x_k^{a'b}, l_k^{a'}$. Note that a' is preserved due to the priority of x_k . Therefore, the unique immediate successor of v is a successful repeat with companion the unique immediate successor of u .

Conclusion: every infinite branch in \mathcal{P}_Δ passes through a successful repeat.

We have thus established:

5.3.10. PROPOSITION. *For every Δ such that the sequent $\Gamma \Rightarrow \Delta$ is closed, well-named and guarded, and $\text{Voc}(\Delta) \cap \text{Voc}(\Gamma) \subseteq V$, if $\text{JS} \vdash \Gamma \Rightarrow \Delta$, then $\text{JS} \vdash \iota \Rightarrow \Delta$.*

Putting Propositions 5.3.8 and 5.3.10 together yields the uniform interpolation theorem stated above (Theorem 5.3.1).

5.4 Conclusion

Uniform interpolation for the μ -calculus was first established in [32] by automata-theoretic means. We have provided a purely proof-theoretic construction (and verification) of uniform interpolants in the cyclic system due to Jungteerapanich [76] and Stirling [144]. From a proof-search with generic consequent, a modal equational system is obtained the solution of which is the desired uniform interpolant.

Our strategy readily yields uniform interpolants for the Gödel–Löb provability logic **GL** by means of Shamkanov’s cyclic system [134] examined in Chapter 2.¹³ As in the case of the modal μ -calculus, uniform interpolants are represented as equational systems, solutions for which may be computed following a simple, purely syntactic algorithm relying on the well-known fixpoint theorem for **GL**,¹⁴ for instance the method due to Reidhaar–Olson [114].

One subtle difference between the cases of **GL** and the modal μ -calculus is that the verification of the interpolant in **GL** utilises an extension of the cyclic calculus with the rule:

$$\equiv \frac{\Gamma, \psi}{\Gamma, \varphi} \varphi \equiv \psi$$

This extension merely addresses the lack of explicit fixpoint quantifiers in **GL**.

The other cyclic systems examined in Chapter 2, however, do not seem adequate for this construction. Indeed, **GL** seems to be the exception rather than the norm.

Consider **S4Grz**, for instance, and the cyclic system **Grz_{circ}** for the logic due to Savateev and Shamkanov [130]. An appropriate notion of template can be defined, and the equational systems can be solved by appealing to the fixpoint theorem of **S4Grz** [93]. However, the fact that the (transitive) modal rule of **Grz_{circ}** has two premises, only one of which yielding ‘good’ cycles, coupled with the lack of explicit fixpoint quantifiers in **S4Grz**, does not allow the verification process to go through. We conjecture that similar problems arise for **K4Grz**.

Craig interpolation for **iGL** was recently shown in [57]. It is a syntactic proof relying on a cut-free sequent calculus for **iGL**. Uniform interpolation for **iGL**, on the other hand, remains an open problem (see [56, Ch. 3]). A cyclic calculus for **iGL** can be readily obtained from Shamkanov’s system for **GL** [69, 136]. Our method will probably not work in this setting either, for the same reasons that it fails for **S4Grz**.

We leave for future work investigating the adaptability of our strategy to these and similar logics.

¹³That **GL** enjoys uniform interpolation was first shown by Shavrukov [135].

¹⁴First proved, independently, by de Jongh and Sambin [124].

Another natural continuation of our work in this chapter is to study whether our method yields uniform *Lyndon* interpolants for the μ -calculus, i.e., interpolants that respect the polarities of the literals.

Conclusion

Cyclic and ill-founded proof-theory allow proofs with infinite branches or paths, as long as they satisfy some correctness conditions ensuring the validity of the conclusion. In this dissertation we have designed a few cyclic and ill-founded systems: a cyclic one for the modal logic $\mathbf{K4Grz}$; and ill-founded and cyclic ones for the temporal logics \mathbf{CTL}^* and \mathbf{iTL} . Lastly, we have used a cyclic system for the modal μ -calculus to obtain a proof of the uniform interpolation property for the logic which differs from the original, automata-based one [32].

Adding cycles to ordinary sequent calculi for $\mathbf{K4}$ and $\mathbf{S4}$ yields cyclic proof systems for the Gödel–Löb provability logic (\mathbf{GL}) and the Grzegorzcyk logic ($\mathbf{S4Grz}$), respectively [134, 130]. Rather than isolated contrivances, we argued in Chapter 2 that these systems arise from a natural correspondence between cycles in proofs and infinite chains in frames that enables the former to capture frame conditions involving the latter. We proposed to understand *cyclic companionship* by a combination of proof-theoretic and semantical considerations. According to our explanation, one should be able to obtain a cyclic system for the weak Grzegorzcyk logic $\mathbf{K4Grz}$ by adding cycles to a system for $\mathbf{K4}$. We showed that this is indeed the case, thus establishing that $\mathbf{K4Grz}$, like \mathbf{GL} , is a cyclic companion of $\mathbf{K4}$.

Our work in Chapter 2 likely applies to logics with frame conditions similar to the ones of \mathbf{GL} , $\mathbf{S4Grz}$ and $\mathbf{K4Grz}$, for example $\mathbf{S4.2Grz}$ and $\mathbf{S4.3Grz}$. But we suspect that not all flavours of converse well-foundedness are amenable to the cyclic approach. In this regard, an interesting line of research seems to us to be the study of the relation between cyclic companionship and modal definability.

In the first part of Chapter 3, we introduced a cut-free, cyclic hypersequent calculus for the full computation tree logic \mathbf{CTL}^* . Local soundness of inferences is immediate, and a global correctness condition ensures that cycles yield valid conclusions. Hypersequents offer a natural framework for accommodating the existential (\mathbf{E}) and universal (\mathbf{A}) path quantifiers of the logic, as well as their

interplay with the next operator X . Each ‘sequent’ in a hypersequent is a labelled set of formulas, either $A\Phi$ or $E\Phi$, interpreted as ‘along all paths, $\forall \Phi$ is the case’ and ‘along some path, $\wedge \Phi$ is the case’, 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.

A simple annotation mechanism on formulas allows us to isolate a finitary condition which suffices to guarantee that a derivation is a proof. Annotations, however, do not seem capable of handling all of the complexity of infinite branches in ill-founded derivations. Indeed, purely existential branches, i.e., those the success of which is not witnessed by any universal trace, seem beyond the scope of the simple annotations that we have used.

A promising alternative is provided by automata theory. Correctness conditions on branches are ω -regular, so it seems reasonable to suppose that one could encode a deterministic automaton recognising infinite branches in the syntax of the system, by means of a more complex annotating mechanism than the one we have used. This idea is partially applied in [47] to a tableau for satisfiability of CTL^* formulas, and [37] investigates the correspondence between annotations for the μ -calculus and automata determinisation procedures.

In the second half of Chapter 3, we isolated a class of ‘inductible’ cyclic proofs whose cycles can be transformed into inductive arguments based on the following Park-style characterisation of until:

$$\frac{}{(\beta \vee (\alpha \wedge X(\alpha U \beta))) \rightarrow \alpha U \beta} \qquad \frac{(\beta \vee (\alpha \wedge X\gamma)) \rightarrow \gamma}{\alpha U \beta \rightarrow \gamma}$$

In the end, we arrived at a Hilbert-style system and compared it to a fragment of a known axiomatisation for the full logic due to Reynolds [115]. Our axiom system is complete for a well-known variant of CTL^* obtained by allowing the evaluation of formulas in a bigger class than serial models.

Chapter 4 introduced a cut-free cyclic proof system for the intuitionistic linear-time temporal logic $iLTL$, with a fully finitary correctness condition. The calculus uses labelled formulas in order to accommodate the interplay between the ‘temporal dimension’, represented by the modal rule for the next operator X , and the ‘intuitionistic dimension’, corresponding to the right-implication rule $R \rightarrow$. Simple annotations on release and until formulas suffice to provide a finitary characterisation of good infinite branches. The most natural continuation of the work in Chapter 4 is to try to adapt our system to the version of $iLTL$ obtained by imposing not only forwards confluence, but backwards too.

Lastly, in Chapter 5 we gave a proof of the uniform interpolation theorem for the modal μ -calculus which differs from the original, automata-based one [32]. We

constructed uniform interpolants from cyclic derivations in the system for the μ -calculus due to Jungteerapanich [76] and Stirling [144]. The approach works also for **GL** by means of Shamkanov's cyclic system [134], but this seems to be the exception rather than the norm. In the absence of explicit fixpoint quantifiers, or so it seems, our method cannot be applied to systems more complex than Shamkanov's.

Bibliography

- [1] B. Afshari and G. E. Leigh. ‘Lyndon Interpolation for Modal μ -Calculus.’ In A. Özgün and Y. Zinova (eds.), *Language, Logic, and Computation*, volume 13206 of *Lecture Notes in Computer Science*, pages 197–213. Springer International Publishing, Cham, 2022. DOI: 10.1007/978-3-030-98479-3_10.
- [2] B. Afshari and D. Wehr. ‘Abstract Cyclic Proofs.’ In A. Ciabattoni, E. Pimentel, and R. J. G. B. de Queiroz (eds.), *Logic, Language, Information, and Computation*, volume 13468 of *Lecture Notes in Computer Science*, pages 309–325. Springer International Publishing, Cham, 2022. DOI: 10.1007/978-3-031-15298-6_20.
- [3] B. Afshari, G. E. Leigh, and G. Menéndez Turata. ‘Uniform Interpolation from Cyclic Proofs: The Case of Modal Mu-Calculus.’ In A. Das and S. Negri (eds.), *Automated Reasoning with Analytic Tableaux and Related Methods*, volume 12842 of *Lecture Notes in Computer Science*, pages 335–353. Springer International Publishing, Cham, 2021. DOI: 10.1007/978-3-030-86059-2_20.
- [4] B. Afshari, S. Enqvist, and G. E. Leigh. ‘Cyclic proofs for the first-order μ -calculus.’ In *Logic Journal of the IGPL*, page jzac053, 2022. DOI: 10.1093/jigpal/jzac053.
- [5] B. Afshari, L. Grotenhuis, G. E. Leigh, and L. Zenger. ‘Ill-Founded Proof Systems for Intuitionistic Linear-Time Temporal Logic.’ In R. Ramanayake and J. Urban (eds.), *Automated Reasoning with Analytic Tableaux and Related Methods*, volume 14278 of *Lecture Notes in Computer Science*, pages 223–241. Springer Nature Switzerland, Cham, 2023. DOI: 10.1007/978-3-031-43513-3_13.

- [6] B. Afshari, G. E. Leigh, and G. Menéndez Turata. ‘A Cyclic Proof System for Full Computation Tree Logic.’ In B. Klin and E. Pimentel (eds.), *31st EACSL Annual Conference on Computer Science Logic (CSL 2023)*, volume 252 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 5:1–5:19. Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, 2023. DOI: 10.4230/LIPIcs.CSL.2023.5.
- [7] M. Amerbauer. ‘Cut-free Tableau Calculi for some Propositional Normal Modal Logics.’ In *Studia Logica*, 57(2–3):359–372, 1996. DOI: 10.1007/BF00370840.
- [8] A. Avron. ‘On modal systems having arithmetical interpretations.’ In *The Journal of Symbolic Logic*, 49(3):935–942, 1984. DOI: 10.2307/2274147.
- [9] D. Baelde, A. Doumane, and A. Saurin. ‘Infinitary Proof Theory: the Multiplicative Additive Case.’ In J.-M. Talbot and L. Regnier (eds.), *25th EACSL Annual Conference on Computer Science Logic (CSL 2016)*, volume 62 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 42:1–42:17. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2016. DOI: 10.4230/LIPIcs.CSL.2016.42.
- [10] P. Balbiani, J. Boudou, M. Diéguez, and D. Fernández-Duque. ‘Intuitionistic Linear Temporal Logics.’ In *ACM Trans. Comput. Logic*, 21(2), 2019. DOI: 10.1145/3365833.
- [11] J. F. A. K. van Benthem. *Modal Logic and Classical Logic*. Number 3 in Monographs in Philosophical Logic and Formal Linguistics. Bibliopolis, Naples, 1983.
- [12] J. F. A. K. van Benthem and W. J. Blok. ‘Transitivity follows from Dummett’s axiom.’ In *Theoria*, 44(2):117–118, 1978. DOI: 10.1111/j.1755-2567.1978.tb00170.x.
- [13] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Number 53 in Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, Cambridge, 2001. DOI: 10.1017/CB09781107050884.
- [14] G. Boolos. *The Logic of Provability*. Cambridge University Press, Cambridge, 1993. DOI: 10.1017/CB09780511625183.
- [15] G. Boolos and G. Sambin. ‘Provability: the Emergence of a Mathematical Modality.’ In *Studia Logica*, 50(1):1–23, 1991. DOI: 10.1007/BF00370383.

- [16] J. Bradfield and C. Stirling. ‘Modal mu-calculi.’ In P. Blackburn, J. van Benthem, and F. Wolter (eds.), *Handbook of Modal Logic*, volume 3 of *Studies in Logic and Practical Reasoning*, pages 721–756. Elsevier, 2007. DOI: 10.1016/S1570-2464(07)80015-2.
- [17] J. Brotherston. *Sequent calculus proof systems for inductive definitions*. Ph.D. thesis, University of Edinburgh, 2006.
- [18] J. Brotherston and A. Simpson. ‘Sequent calculi for induction and infinite descent.’ In *Journal of Logic and Computation*, 21(6):1177–1216, 2010. DOI: 10.1093/logcom/exq052.
- [19] K. Brännler and M. Lange. ‘Cut-free sequent systems for temporal logic.’ In *The Journal of Logic and Algebraic Programming*, 76(2):216–225, 2008. DOI: 10.1016/j.jlap.2008.02.004.
- [20] J. R. Büchi. ‘On a decision method in restricted second order arithmetic.’ In E. Nagel, P. Suppes, and A. Tarski (eds.), *Logic, Methodology and Philosophy of Science: Proceedings of the 1960 International Congress*, pages 1–11. Stanford University Press, Stanford, California, 1962.
- [21] J. R. Büchi. ‘On a Decision Method in Restricted Second Order Arithmetic.’ In S. Mac Lane and D. Siefkes (eds.), *The Collected Works of J. Richard Büchi*, pages 425–435. Springer New York, New York, NY, 1990. DOI: 10.1007/978-1-4613-8928-6_23. Reprinted from [20].
- [22] R. A. Bull. ‘An algebraic study of Diodorean modal systems.’ In *The Journal of Symbolic Logic*, 30(1):58–64, 1965. DOI: 10.2307/2270582.
- [23] R. A. Bull. ‘That all normal extensions of S4.3 have the finite model property.’ In *Mathematical Logic Quarterly*, 12(1):341–344, 1966. DOI: 10.1002/malq.19660120129.
- [24] A. Chagrov and M. Zakharyashev. *Modal Logic*, volume 35 of *Oxford Logic Guides*. Clarendon Press, Midsomer Norton, 1997.
- [25] A. Chalki, C. D. Koutras, and Y. Zikos. ‘A quick guided tour to the modal logic S4.2.’ In *Logic Journal of the IGPL*, 26(4):429–451, 2018. DOI: 10.1093/jigpal/jzy008.
- [26] E. M. Clarke and E. A. Emerson. ‘Design and synthesis of synchronization skeletons using branching time temporal logic.’ In D. Kozen (ed.), *Logics of Programs*, volume 131 of *Lecture Notes in Computer Science*, pages 52–71. Springer Berlin Heidelberg, Berlin, Heidelberg, 1982. DOI: 10.1007/BFb0025774.

- [27] A. Colmerauer. ‘PROLOG and Infinite Trees.’ In *Logic Programming*, volume 16 of *APIC Studies in Data Processing*, pages 231–251. Academic Press, Orlando, 1982.
- [28] B. Courcelle. ‘Fundamental properties of infinite trees.’ In *Theoretical Computer Science*, 25(2):95–169, 1983. DOI: 10.1016/0304-3975(83)90059-2.
- [29] G. Cousineau. ‘An algebraic definition for control structures.’ In *Theoretical Computer Science*, 12(2):175–192, 1980. DOI: 10.1016/0304-3975(80)90028-6.
- [30] W. Craig. ‘Three uses of the Herbrand-Gentzen theorem in relating model theory and proof theory.’ In *The Journal of Symbolic Logic*, 22(3):269–285, 1957. DOI: 10.2307/2963594.
- [31] M. Cresswell, E. Mares, and A. Rini (eds.). *Logical Modalities from Aristotle to Carnap: The Story of Necessity*. Cambridge University Press, Cambridge, 2016. DOI: 10.1017/CB09781139939553.
- [32] G. D’Agostino and M. Hollenberg. ‘Logical Questions Concerning the μ -Calculus: Interpolation, Lyndon and Łoś-Tarski.’ In *The Journal of Symbolic Logic*, 65(1):310–332, 2000. DOI: 10.2307/2586539.
- [33] M. Dam. ‘Translating CTL* into the modal μ -calculus.’ Technical report ECS-LFCS-90-123, University of Edinburgh, Department of Computer Science, Laboratory for Foundations of Computer Science, 1990.
- [34] A. Das and M. Girlando. ‘Cyclic Hypersequent System for Transitive Closure Logic.’ In *Journal of Automated Reasoning*, 67(3):27, 2023. DOI: 10.1007/s10817-023-09675-1.
- [35] A. Das, I. van der Giessen, and S. Marin. ‘Intuitionistic Gödel–Löb logic, à la Simpson: labelled systems and birelational semantics.’ 2023. DOI: 10.48550/arXiv.2309.00532.
- [36] M. Davis. ‘Infinite Games of Perfect Information.’ In M. Dresher, Ll. S. Shapley, and A. W. Tucker (eds.), *Advances in Game Theory*, volume 52 of *Annals of Mathematics Studies*, pages 85–102. Princeton University Press, Princeton, 1964. DOI: 10.1515/9781400882014-008.
- [37] M. Dekker, J. Kloibhofer, J. Marti, and Y. Venema. ‘Proof Systems for the Modal μ -Calculus Obtained by Determinizing Automata.’ In R. Ramanayake and J. Urban (eds.), *Automated Reasoning with Analytic Tableaux and Related Methods*, pages 242–259. Springer Nature Switzerland, Cham, 2023. DOI: 10.1007/978-3-031-43513-3_14.

- [38] S. Demri, V. Goranko, and M. Lange. *Temporal Logics in Computer Science: Finite-State Systems*. Number 58 in Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, Cambridge, 2016. DOI: 10.1017/CB09781139236119.
- [39] M. A. E. Dummett and E. J. Lemmon. ‘Modal Logics Between S4 and S5.’ In *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 5 (14–24):250–264, 1959. DOI: 10.1002/ma1q.19590051405.
- [40] R. Dyckhoff and S. Negri. ‘A cut-free sequent system for Grzegorzcyk logic, with an application to the Gödel–McKinsey–Tarski embedding.’ In *Journal of Logic and Computation*, 26(1):169–187, 2013. DOI: 10.1093/logcom/ext036.
- [41] E. A. Emerson. ‘Alternative semantics for temporal logics.’ In *Theoretical Computer Science*, 26(1–2):121–130, 1983. DOI: 10.1016/0304-3975(83)90082-8.
- [42] E. A. Emerson and J. Y. Halpern. “‘Sometimes’ and ‘Not Never’ Revisited: On Branching versus Linear Time.” In *Proceedings of the 10th ACM SIGACT-SIGPLAN Symposium on Principles of Programming Languages*, pages 127–140. Association for Computing Machinery, New York, 1983. DOI: 10.1145/567067.567081.
- [43] S. Enqvist. ‘A circular proof system for the hybrid μ -calculus.’ In N. Olivetti, R. Verbrugge, S. Negri, and G. Sandu (eds.), *Advances in Modal Logic*, volume 13 of *Advances in Modal Logic*, pages 169–188. College Publications, 2020.
- [44] L. Esakia. ‘The modalized Heyting calculus: a conservative modal extension of the Intuitionistic Logic.’ In *Journal of Applied Non-Classical Logics*, 16 (3–4):349–366, 2006. DOI: 10.3166/jancl.16.349-366.
- [45] M. J. Fischer and R. E. Ladner. ‘Propositional Modal Logic of Programs.’ In *Proceedings of the Ninth Annual ACM Symposium on Theory of Computing*, pages 286–294. Association for Computing Machinery, New York, 1977. DOI: 10.1145/800105.803418.
- [46] M. J. Fischer and R. E. Ladner. ‘Propositional Dynamic Logic of Regular Programs.’ In *Journal of Computer and System Sciences*, 18(2):194–211, 1979. DOI: 10.1016/0022-0000(79)90046-1.
- [47] O. Friedmann, M. Latte, and M. Lange. ‘A Decision Procedure for CTL* Based on Tableaux and Automata.’ In J. Giesl and R. Hähnle (eds.),

- Automated Reasoning*, volume 6173 of *Lecture Notes in Computer Science*, pages 331–345. Springer Berlin Heidelberg, Berlin, Heidelberg, 2010. DOI: 10.1007/978-3-642-14203-1_28.
- [48] D. M. Gabbay and L. Maksimova. *Interpolation and Definability: Modal and Intuitionistic Logics*. Number 46 in Oxford Logic Guides. Oxford University Press, New York, 2005.
- [49] D. M. Gabbay and A. Pnueli. ‘A Sound and Complete Deductive System for CTL* Verification.’ In *Logic Journal of the IGPL*, 16(6):499–536, 2008. DOI: 10.1093/jigpal/jzn018.
- [50] D. M. Gabbay, A. Pnueli, S. Shelah, and J. Stavi. ‘On the Temporal Analysis of Fairness.’ In *Proceedings of the 7th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages*, pages 163–173. Association for Computing Machinery, New York, 1980. DOI: 10.1145/567446.567462.
- [51] D. Gabelaia. *Topological, Algebraic and Spatio-Temporal Semantics for Multi-Dimensional Modal Logics*. Ph.D. thesis, King’s College London, 2004.
- [52] D. Gale and F. M. Stewart. ‘Infinite Games with Perfect Information.’ In H. W. Kuhn and A. W. Tucker (eds.), *Contributions to the Theory of Games, Volume II*, volume 28 of *Annals of Mathematics Studies*, pages 245–266. Princeton University Press, Princeton, 1953. DOI: 10.1515/9781400881970-014.
- [53] R. Gaskin. *The Sea Battle and the Master Argument: Aristotle and Diodorus Cronus on the Metaphysics of the Future*. De Gruyter, Berlin, New York, 1995. DOI: doi:10.1515/9783110866346.
- [54] G. Gentzen. ‘Untersuchungen über das logische Schließen. I.’ In *Mathematische Zeitschrift*, 39(1):176–210, 1935. DOI: 10.1007/BF01201353. Reproduced in English translation in [147].
- [55] G. Gentzen. ‘Untersuchungen über das logische Schließen. II.’ In *Mathematische Zeitschrift*, 39(1):405–431, 1935. DOI: 10.1007/BF01201363. Reproduced in English translation in [147].
- [56] I. van der Giessen. *Uniform Interpolation and Admissible Rules: Proof-theoretic investigations into (intuitionistic) modal logics*. Ph.D. thesis, Universiteit Utrecht, 2022. DOI: 10.33540/1486.
- [57] I. van der Giessen and R. Iemhoff. ‘Sequent Calculi for Intuitionistic Gödel–Löb Logic.’ In *Notre Dame Journal of Formal Logic*, 62(2):221–246, 2021. DOI: 10.1215/00294527-2021-0011.

- [58] M. Girlando, R. Kuznets, S. Marin, M. Morales, and L. Straßburger. ‘Intuitionistic S4 is decidable.’ In *Proceedings of the 2023 38th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS)*. IEEE Computer Society Press, 2023. DOI: 10.1109/LICS56636.2023.10175684.
- [59] R. Goldblatt. ‘Diodorean Modality in Minkowski Spacetime.’ In *Studia Logica*, 39(2):219–236, 1980. DOI: 10.1007/BF00370321.
- [60] R. Goldblatt. ‘Mathematical modal logic: A view of its evolution.’ In *Journal of Applied Logic*, 1(5):309–392, 2003. DOI: 10.1016/S1570-8683(03)00008-9.
- [61] R. Goldblatt. ‘Modal Logics that Bound the Circumference of Transitive Frames.’ In J. Madarász and G. Székely (eds.), *Hajnal Andréka and István Németi on Unity of Science: From Computing to Relativity Theory Through Algebraic Logic*, pages 233–265. Springer International Publishing, Cham, 2021. DOI: 10.1007/978-3-030-64187-0_10.
- [62] R. Goré. ‘Semi-analytic tableaux for propositional normal modal logics with application to nonmonotonicity.’ In *Logique et Analyse*, 34(133–134):73–103, 1991.
- [63] R. Goré. ‘Cut-free Sequent and Tableau Systems for Propositional Normal Modal Logics.’ Technical report UCAM-CL-TR-257, University of Cambridge, Computer Laboratory, 1992.
- [64] R. Goré and R. Ramanayake. ‘Valentini’s cut-elimination for provability logic resolved.’ In C. Areces and R. Goldblatt (eds.), *Advances in Modal Logic*, volume 7 of *Advances in Modal Logic*, pages 67–86. College Publications, London, 2008.
- [65] A. Grzegorzcyk. ‘Some relational systems and the associated topological spaces.’ In *Fundamenta Mathematicae*, 60(3):223–231, 1967. DOI: 10.4064/fm-60-3-223-231.
- [66] J. D. Hamkins and B. Löwe. ‘The modal logic of forcing.’ In *Transactions of the American Mathematical Society*, 360(4):1793–1817, 2008. DOI: 10.1090/S0002-9947-07-04297-3.
- [67] H. J. Hoogeboom and G. Rozenberg. ‘Infinitary languages: Basic theory and applications to concurrent systems.’ In J. W. de Bakker, W.-P. de Roever, and G. Rozenberg (eds.), *Current Trends in Concurrency: Overviews and*

- Tutorials*”, volume 224 of *Lecture Notes in Computer Science*, pages 266–342. Springer Berlin Heidelberg, Berlin, Heidelberg, 1986. DOI: 10.1007/BFb0027043.
- [68] G. E. Hughes and M. J. Cresswell. *A New Introduction to Modal Logic*. Routledge, London and New York, 1996.
- [69] R. Iemhoff. ‘Reasoning in Circles.’ In J. van Eijck, R. Iemhoff, and J. J. Joosten (eds.), *Liber Amicorum Alberti. A Tribute to Albert Visser*, pages 165–176. College Publications, 2016.
- [70] R. Iemhoff. ‘Uniform interpolation and sequent calculi in modal logic.’ In *Archive for Mathematical Logic*, 58(1):155–181, 2019. DOI: 10.1007/s00153-018-0629-0.
- [71] R. Iemhoff. ‘Uniform interpolation and the existence of sequent calculi.’ In *Annals of Pure and Applied Logic*, 170(11), 2019. DOI: 10.1016/j.apal.2019.05.008.
- [72] Y. Imai and K. Iséki. ‘On Axiom Systems of Propositional Calculi, I.’ In *Proceedings of the Japan Academy*, 41(6):436–439, 1965. DOI: 10.3792/pja/1195522378.
- [73] G. Jäger, M. Kretz, and T. Studer. ‘Canonical completeness of infinitary μ .’ In *The Journal of Logic and Algebraic Programming*, 76(2):270–292, 2008. DOI: 10.1016/j.jlap.2008.02.005.
- [74] D. Janin and I. Walukiewicz. ‘On the expressive completeness of the propositional mu-calculus with respect to monadic second order logic.’ In U. Montanari and V. Sassone (eds.), *CONCUR ’96: Concurrency Theory*, volume 1119 of *Lecture Notes in Computer Science*, pages 263–277. Springer Berlin Heidelberg, Berlin, Heidelberg, 1996. DOI: 10.1007/3-540-61604-7_60.
- [75] T. Jech. *Set Theory. The Third Millennium Edition, Revised and Expanded*. Springer Monographs in Mathematics. Springer Berlin Heidelberg, Berlin, Heidelberg, 4th edition, 2006. DOI: 10.1007/3-540-44761-X.
- [76] N. Jungteerapanich. ‘A Tableau System for the Modal μ -Calculus.’ In M. Giese and A. Waaler (eds.), *Automated Reasoning with Analytic Tableaux and Related Methods*, volume 5607 of *Lecture Notes in Computer Science*, pages 220–234. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009. DOI: 10.1007/978-3-642-02716-1_17.

- [77] N. Jungteerapanich. *Tableau Systems for the Modal μ -Calculus*. Ph.D. thesis, University of Edinburgh, 2010.
- [78] J. A. W. Kamp. *Tense Logic and the Theory of Linear Order*. Ph.D. thesis, University of California, Los Angeles, 1968.
- [79] S. Kanger. *Provability in Logic*. Number 1 in Stockholm Studies in Philosophy. Almqvist & Wiksells, Uppsala, 1957.
- [80] A. S. Kechris. *Classical Descriptive Set Theory*. Number 156 in Graduate Texts in Mathematics. Springer New York, New York, 1995. DOI: 10.1007/978-1-4612-4190-4.
- [81] D. Kozen. ‘Results on the propositional μ -calculus.’ In *Theoretical Computer Science*, 27(3):333–354, 1983. DOI: 10.1016/0304-3975(82)90125-6.
- [82] D. Kozen. ‘A Finite Model Theorem for the Propositional μ -Calculus.’ In *Studia Logica*, 47(3):233–241, 1988. DOI: 10.1007/BF00370554.
- [83] K. Kunen. *Set Theory. An Introduction to Independence Proofs*. Number 102 in Studies in Logic and the Foundations of Mathematics. North Holland Publishing Company, Amsterdam, 1st edition, 1980.
- [84] G. E. Leigh and D. Wehr. ‘From GTC to Reset: Generating Reset Proof Systems from Cyclic Proof Systems.’ 2023. DOI: 10.48550/arXiv.2301.07544.
- [85] D. Leivant. ‘On the proof theory of the modal logic for arithmetic provability.’ In *The Journal of Symbolic Logic*, 46(3):531—538, 1981. DOI: 10.2307/2273755.
- [86] W. Lenzen. *Recent Work in Epistemic Logic*. North Holland Publishing Company, Amsterdam, 1978.
- [87] C. I. Lewis and C. H. Langford. *Symbolic Logic*. The Century Philosophy Series. The Century Company, New York, 1932.
- [88] T. Litak. ‘The non-reflexive counterpart of Grz.’ In *Bulletin of the Section of Logic*, 36(3–4):195–208, 2007.
- [89] J. Łukasiewicz. *Elementy logiki matematycznej*. Państwowe Wydawnictwo Naukowe, Warsaw, 2nd edition, 1958.
- [90] J. Łukasiewicz. *Elements of Mathematical Logic*, volume 31 of *International Series of Monographs in Pure and Applied Mathematics*. Pergamon Press, Oxford, 1963. English translation of [89].

- [91] R. C. Lyndon. ‘An interpolation theorem in the predicate calculus.’ In *Pacific Journal of Mathematics*, 9(1):129–142, 1959. DOI: 10.2140/pjm.1959.9.129.
- [92] L. Maksimova. ‘On Modal Grzegorzczk Logic.’ In *Fundamenta Informaticae*, 81(1–3):203–210, 2007.
- [93] S. Mardaev. ‘Least fixed points in Grzegorzczk’s logic and in the intuitionistic propositional logic.’ In *Algebra and Logic*, 32(5):279–288, 1993. DOI: 10.1007/BF02261708.
- [94] S. Mardaev. ‘Definable fixed points in modal and temporal logics — a survey.’ In *Journal of Applied Non-Classical Logics*, 17(3):317–346, 2007. DOI: 10.3166/janc1.17.317–346.
- [95] J. Marti and Y. Venema. ‘A Focus System for the Alternation-Free μ -Calculus.’ In A. Das and S. Negri (eds.), *Automated Reasoning with Analytic Tableaux and Related Methods*, volume 12842 of *Lecture Notes in Computer Science*, pages 371–388. Springer International Publishing, Cham, 2021. DOI: 10.1007/978-3-030-86059-2_22.
- [96] D. A. Martin. ‘Borel determinacy.’ In *Annals of Mathematics*, 102(2):363–371, 1975. DOI: 10.2307/1971035.
- [97] L. Mauborgne. ‘An Incremental Unique Representation for Regular Trees.’ In *Nordic Journal of Computing*, 7(4):290–311, 2000.
- [98] J. R. Munkres. *Topology*. Pearson Education Limited, Harlow, 2nd edition, 2014.
- [99] S. Negri. ‘Proof Analysis in Modal Logic.’ In *Journal of Philosophical Logic*, 34(5–6):507–544, 2005. DOI: 10.1007/s10992-005-2267-3.
- [100] S. Negri. ‘Proof Theory for Modal Logic.’ In *Philosophy Compass*, 6(8): 523–538, 2011. DOI: 10.1111/j.1747-9991.2011.00418.x.
- [101] S. Negri and J. von Plato. *Structural Proof Theory*. Cambridge University Press, Cambridge, 2001. DOI: 10.1017/CB09780511527340.
- [102] D. Niwiński and I. Walukiewicz. ‘Games for the μ -calculus.’ In *Theoretical Computer Science*, 163(1–2):99–116, 1996. DOI: 10.1016/0304-3975(95)00136-0.

- [103] R. Nollet, A. Saurin, and C. Tasson. ‘Local Validity for Circular Proofs in Linear Logic with Fixed Points.’ In D. Ghica and A. Jung (eds.), *27th EACSL Annual Conference on Computer Science Logic (CSL 2018)*, volume 119 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 35:1–35:23. Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, Dagstuhl, Germany, 2018. DOI: 10.4230/LIPIcs.CSL.2018.35.
- [104] P. Øhrstrøm and P. F. V. Hasle. *Temporal Logic*. Number 57 in Studies in Linguistics and Philosophy. Kluwer Academic Publishers, Dordrecht, 1995. DOI: 10.1007/978-0-585-37463-5.
- [105] D. Park. ‘Fixpoint Induction and Proofs of Program Properties.’ In B. Meltzer and D. Michie (eds.), *Machine Intelligence 5. Proceedings of the Fifth Annual Machine Intelligence Workshop*, number 5 in Machine Intelligence, pages 59–78. Edinburgh University Press, Edinburgh, 1969.
- [106] D. Perrin and J.-É. Pin. *Infinite Words: Automata, Semigroups, Logic and Games*. Number 141 in Pure and Applied Mathematics. Elsevier, Amsterdam, 2004.
- [107] A. M. Pitts. ‘On an interpretation of second order quantification in first order intuitionistic propositional logic.’ In *The Journal of Symbolic Logic*, 57(1):33–52, 1992. DOI: 10.2307/2275175.
- [108] A. Pnueli. ‘The temporal logic of programs.’ In *Proceedings of the 18th Annual Symposium on Foundations of Computer Science*, pages 46–57. 1977. DOI: 10.1109/SFCS.1977.32.
- [109] F. Poggiolesi. ‘A purely syntactic and cut-free sequent calculus for the modal logic of provability.’ In *The Review of Symbolic Logic*, 2(4):593–611, 2009. DOI: 10.1017/S1755020309990244.
- [110] A. N. Prior. ‘Diodoran Modalities.’ In *The Philosophical Quarterly*, 5(20): 205–213, 1955. DOI: 10.2307/2957434.
- [111] A. N. Prior. *Time and Modality*. Clarendon Press, Oxford, 1957.
- [112] C. Pulver. *Self-Referentiality in Contraction-free Fragments of Modal Logic S₄*. Master’s thesis, Universität Bern, 2010.
- [113] W. Rautenberg. ‘Modal tableau calculi and interpolation.’ In *Journal of Philosophical Logic*, 12(4):403–423, 1983. DOI: 10.1007/BF00249258.

- [114] L. Reidhaar-Olson. ‘A New Proof of the Fixed-Point Theorem of Provability Logic.’ In *Notre Dame Journal of Formal Logic*, 31(1):37–43, 1989. DOI: 10.1305/ndjf1/1093635331.
- [115] M. Reynolds. ‘An Axiomatization of Full Computation Tree Logic.’ In *The Journal of Symbolic Logic*, 66(3):1011–1057, 2001. DOI: 10.2307/2695091.
- [116] M. Reynolds. ‘An axiomatization of PCTL*.’ In *Information and Computation*, 201(1):72–119, 2005. DOI: 10.1016/j.ic.2005.03.005.
- [117] M. Reynolds. ‘A Tableau for CTL*.’ In A. Cavalcanti and D. R. Dams (eds.), *FM 2009: Formal Methods*, volume 5850 of *Lecture Notes in Computer Science*, pages 403–418. Springer Berlin Heidelberg, Berlin, Heidelberg, 2009. DOI: 10.1007/978-3-642-05089-3_26.
- [118] M. Reynolds. ‘A tableau-based decision procedure for CTL*.’ In *Formal Aspects of Computing*, 23(6):739–779, 2011. DOI: 10.1007/s00165-011-0193-4.
- [119] A. Rini. *Aristotle’s Modal Proofs: Prior Analytics A8–22 in Predicate Logic*. Number 68 in The New Synthese Historical Library. Springer Dordrecht, Dordrecht, 2011. DOI: 10.1007/978-94-007-0050-5.
- [120] J. Rooduijn. ‘Cyclic Hypersequent Calculi for Some Modal Logics with the Master Modality.’ In A. Das and S. Negri (eds.), *Automated Reasoning with Analytic Tableaux and Related Methods*, pages 354–370. Springer International Publishing, Cham, 2021. DOI: 10.1007/978-3-030-86059-2_21.
- [121] J. Rooduijn and Y. Venema. ‘Focus-Style Proofs for the Two-Way Alternation-Free μ -Calculus.’ In H. H. Hansen, A. Scedrov, and R. J. G. B. de Queiroz (eds.), *Logic, Language, Information, and Computation*, volume 13923 of *Lecture Notes in Computer Science*, pages 318–335. Springer Nature Switzerland, Cham, 2023. DOI: 10.1007/978-3-031-39784-4_20.
- [122] J. Rooduijn and L. Zenger. ‘An analytic proof system for common knowledge logic over S5.’ In D. Fernández-Duque, A. Palmigiano, and S. Pinchinat (eds.), *Advances in Modal Logic*, volume 14 of *Advances in Modal Logic*, pages 659–680. College Publications, London, 2022.
- [123] S. Safra. ‘On the complexity of ω -automata.’ In *Proceedings of the 29th Annual Symposium on Foundations of Computer Science*, pages 319–327. IEEE Computer Society Press, Los Alamitos, 1988. DOI: 10.1109/SFCS.1988.21948.

- [124] G. Sambin. ‘An effective fixed-point theorem in intuitionistic diagonalizable algebras.’ In *Studia Logica*, 35(4):345–361, 1976. DOI: 10.1007/BF02123402.
- [125] G. Sambin and S. Valentini. ‘A Modal Sequent Calculus for a Fragment of Arithmetic.’ In *Studia Logica*, 39(2–3):245–256, 1980. DOI: 10.1007/BF00370323.
- [126] G. Sambin and S. Valentini. ‘The Modal Logic of Provability. The Sequential Approach.’ In *Journal of Philosophical Logic*, 11(3):311–342, 1982. DOI: 10.1007/BF00293433.
- [127] A. Saurin. ‘A Linear Perspective on Cut-Elimination for Non-wellfounded Sequent Calculi with Least and Greatest Fixed-Points.’ In R. Ramanayake and J. Urban (eds.), *Automated Reasoning with Analytic Tableaux and Related Methods*, volume 14278 of *Lecture Notes in Computer Science*, pages 203–222. Springer Nature Switzerland, Cham, 2023. DOI: 10.1007/978-3-031-43513-3_12.
- [128] Y. Savateev and D. Shamkanov. ‘Cut-Elimination for the Modal Grzegorzczk Logic via Non-well-founded Proofs.’ In J. Kennedy and R. J. G. B. de Queiroz (eds.), *Logic, Language, Information, and Computation*, volume 10388 of *Lecture Notes in Computer Science*, pages 321–335. Springer Berlin Heidelberg, Berlin, Heidelberg, 2017. DOI: 10.1007/978-3-662-55386-2_23.
- [129] Y. Savateev and D. Shamkanov. ‘Cut Elimination for the Weak Modal Grzegorzczk Logic via Non-well-Founded Proofs.’ In R. Iemhoff, M. Moortgat, and R. de Queiroz (eds.), *Logic, Language, Information, and Computation*, volume 11541 of *Lecture Notes in Computer Science*, pages 569–583. Springer Berlin Heidelberg, Berlin, Heidelberg, 2019. DOI: 10.1007/978-3-662-59533-6_34.
- [130] Y. Savateev and D. Shamkanov. ‘Non-well-founded proofs for the Grzegorzczk modal logic.’ In *The Review of Symbolic Logic*, 14(1):22–50, 2021. DOI: 10.1017/S1755020319000510.
- [131] K. Segerberg. ‘Modal logics with linear alternative relations.’ In *Theoria*, 36(3):301–322, 1970. DOI: 10.1111/j.1755-2567.1970.tb00429.x.
- [132] K. Segerberg. *An Essay in Classical Modal Logic*. Number 13 in *Filosofiska studier*. Filosofiska föreningen och Filosofiska institutionen vid Uppsala universitet, Uppsala, 1971.

- [133] D. Shamkanov. ‘Interpolation properties for provability logics GL and GLP.’ In *Proceedings of the Steklov Institute of Mathematics*, 274(1):303–316, 2011. DOI: 10.1134/S0081543811060198.
- [134] D. Shamkanov. ‘Circular proofs for the Gödel–Löb provability logic.’ In *Mathematical Notes*, 96(3):575–585, 2014. DOI: 10.1134/S0001434614090326.
- [135] V. Y. Shavrukov. *Adventures in Diagonalizable Algebras*. Ph.D. thesis, Universiteit van Amsterdam, 1994.
- [136] B. Sierra Miranda. ‘Cyclic Proofs for iGL via Corecursion.’ 2023. DOI: 10.48550/arXiv.2310.10785.
- [137] A. Simpson. ‘Cyclic Arithmetic Is Equivalent to Peano Arithmetic.’ In J. Esparza and A. S. Murawski (eds.), *Foundations of Software Science and Computation Structures*, volume 10203 of *Lecture Notes in Computer Science*, pages 283–300, Berlin, Heidelberg, 2017. Springer Berlin Heidelberg. DOI: 10.1007/978-3-662-54458-7_17.
- [138] B. Sobociński. ‘Family K of the non-Lewis modal systems.’ In *Notre Dame Journal of Formal Logic*, 5(4):313–318, 1964. DOI: 10.1305/ndjfl/1093957981.
- [139] C. Sprenger and M. Dam. ‘On the Structure of Inductive Reasoning: Circular and Tree-Shaped Proofs in the μ -Calculus.’ In A. D. Gordon (ed.), *Foundations of Software Science and Computation Structures*, volume 2620 of *Lecture Notes in Computer Science*, pages 425–440. Springer Berlin Heidelberg, Berlin, Heidelberg, 2003. DOI: 10.1007/3-540-36576-1_27.
- [140] L. Staiger. ‘ ω -Languages.’ In G. Rozenberg and A. Salomaa (eds.), *Beyond Words*, volume 3 of *Handbook of Formal Languages*, pages 339–387. Springer Berlin Heidelberg, Berlin, Heidelberg, 1997. DOI: 10.1007/978-3-642-59126-6_6.
- [141] R. Stalnaker. ‘On Logics of Knowledge and Belief.’ In *Philosophical Studies*, 128(1):169–199, 2006. DOI: 10.1007/s11098-005-4062-y.
- [142] C. Stirling. ‘Comparing Linear and Branching Time Temporal Logics.’ In B. Banieqbal, H. Barringer, and A. Pnueli (eds.), *Temporal Logic in Specification*, volume 398 of *Lecture Notes in Computer Science*, pages 1–20. Springer Berlin Heidelberg, Berlin, Heidelberg, 1989. DOI: 10.1007/3-540-51803-7_19.

- [143] C. Stirling. ‘Modal and Temporal Logics.’ In S. Abramsky, D. M. Gabbay, and T. S. E. Maibaum (eds.), *Handbook of Logic in Computer Science. Volume 2. Background: Computational Structures*, volume 2 of *Handbook of Logic in Computer Science*, pages 477–563. Clarendon Press, New York, 1992.
- [144] C. Stirling. ‘A Tableau Proof System with Names for Modal Mu-calculus.’ In A. Voronkov and M. Korovina (eds.), *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. DOI: 10.29007/1wqm.
- [145] R. S. Streett and E. A. Emerson. ‘An Automata Theoretic Decision Procedure for the Propositional Mu-Calculus.’ In *Information and Computation*, 81(3):249–264, 1989. DOI: 10.1016/0890-5401(89)90031-X.
- [146] T. Studer. ‘On the Proof Theory of the Modal mu-Calculus.’ In *Studia Logica*, 89(3):343–363, 2008. DOI: 10.1007/s11225-008-9133-6.
- [147] M. E. Szabo (ed.). *The Collected Papers of Gerhard Gentzen*. North Holland Publishing Company, Amsterdam, 1969.
- [148] A. A. Tabatabai and R. Jalali. ‘Universal Proof Theory: Semi-analytic Rules and Uniform Interpolation.’ 2018. DOI: 10.48550/arXiv.1808.06258.
- [149] M. Takano. ‘A Modified Subformula Property for the Modal Logic S4.2.’ In *Bulletin of the Section of Logic*, 48(1):19–28, 2019. DOI: 10.18778/0138-0680.48.1.02.
- [150] A. Tarski. ‘A lattice-theoretical fixpoint theorem and its applications.’ In *Pacific Journal of Mathematics*, 5(2):285–309, 1955. DOI: 10.2140/pjm.1955.5.285.
- [151] W. Thomas. ‘Automata on Infinite Objects.’ In J. van Leeuwen (ed.), *Formal Models and Semantics*, volume B of *Handbook of Theoretical Computer Science*, pages 133–191. Elsevier, Amsterdam, 1990. DOI: 10.1016/B978-0-444-88074-1.50009-3.
- [152] A. S. Troelstra and H. Schwichtenberg. *Basic Proof Theory*. Number 43 in Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, Cambridge, 2nd edition, 2000. DOI: 10.1017/CB09781139168717.
- [153] F. Turbak and J. B. Wells. ‘Cycle Therapy: A Prescription for Fold and Unfold on Regular Trees.’ In *Proceedings of the 3rd ACM SIGPLAN International Conference on Principles and Practice of Declarative Programming*,

- PPDP '01, pages 137–149. Association for Computing Machinery, New York, 2001. DOI: 10.1145/773184.773200.
- [154] S. Valentini. ‘The Modal Logic of Provability: Cut-elimination.’ In *Journal of Philosophical Logic*, 12(4):471–476, 1983. DOI: 10.1007/BF00249262.
- [155] F. P. J. M. Voorbraak. *As Far as I Know: Epistemic Logic and Uncertainty*. Ph.D. thesis, Universiteit Utrecht, 1993.
- [156] I. Walukiewicz. ‘On Completeness of the μ -calculus.’ In *Proceedings of the Eighth Annual IEEE Symposium on Logic in Computer Science (LICS)*, pages 136–146. IEEE Computer Society Press, 1993. DOI: 10.1109/LICS.1993.287593.
- [157] I. Walukiewicz. ‘Completeness of Kozen’s Axiomatisation of the Propositional μ -Calculus.’ In *Information and Computation*, 157(1):142–182, 2000. DOI: 10.1006/inco.1999.2836.
- [158] H. Wansing. ‘Sequent Systems for Modal Logics.’ In D. M. Gabbay and F. Guenther (eds.), *Handbook of Philosophical Logic*, volume 8 of *Handbook of Philosophical Logic*, pages 61–145. Springer Netherlands, Dordrecht, 2002. DOI: 10.1007/978-94-010-0387-2_2.
- [159] A. Weil. *Number Theory: An approach through history from Hammurapi to Legendre*. Birkhäuser, Boston, 2007. Reprint of the 1984 edition.

A

accessibility relation, 33
action, 223
 α -sequence, *see* sequence
alphabet, 22
annotation, 99, 206, 228
antecedent, 36, 164, 228
ascending chain, *see* chain
automaton, *see* Büchi automaton
automaton skeleton, *see* skeleton

B

Büchi automaton, 23
 deterministic, 23
 final states of, 23
 initial state of, 23
 run of, *see* run
back-edge map, 11
bad branch, 75, 87, 193
base, 18, 99, 206, 228
Borel
 hierarchy, 19
 set, 18
branch, 9

C

calculus, 17
 cyclic, 18

 finitary, 18
Cantor topology, 19
chain, 33
chain condition
 strong, 34
 weak, 48
children function, 189
closure
 in CTL*, 73
 in iLTL, 168
coherent traces, 168
companion, 11
completeness, 33, 34
complexity, *see* size
concatenation, 8
cone, 9
confluence
 backward, 218
 forward, 161
consequent, 36, 164, 228
context, 126
 external, 126
 internal, 126
control, 99, 228
 relational, 163
control map, 14
Craig interpolation, 232
CTL*, 65
cul-de-sac, 17

cyclic companion, 37

D

DBA, 23

demotion, 112, 216

derivation, 17

CTL_o^{*} derivation, 100

CTL_{ind}^{*} derivation, 115

CTL_∞^{*} derivation, 69

CTL_{reg}^{*} derivation, 96

deterministic, 96

iLTL_o derivation, 207

iLTL_∞ derivation, 165

iLTL_{reg} derivation, 205

JS derivation, 229

K4^o derivation, 46

normal, 231

Diodorean logic, 57

Diodorean modality, 57

discarding derivation, 71, 102

dual, 32, 66, 68, 225

E

eventually, 67, 160

F

finite model property, 56

fixed name

eventually, 103

on a path, 103

on a trace, 103

fixpoint quantifiers, 223

fixpoint unfolding, 168

FMP, 56

forcing, 33

formula

active, 46, 71, 166

annotated, 228

[·]-, 223

[a]-, 223

closed, 224

of CTL^{*}, 66

⟨·⟩-, 223

⟨a⟩-, 223

dominating, 169

dual of, *see* dual

duplicated, 49

equivalence, 33, 68, 161, 225

guarded, 225

of iLTL, 160

\mathcal{K} -valid, 33

left-duplicated, 49

modal, 66, 160, 223

of modal logic, 31

of the μ -calculus, 223

μ -formula, 223

n -annotated, 109, 213

ν -formula, 223

principal, 46, 71, 166

quantifier-free, 223

R-, 66

R-formula, 160

right-duplicated, 49

side, 46, 71, 167

twice-annotated, 109, 213

U-, 66

U-formula, 160

valid, 33, 68, 161, 225

well-annotated, 99

well-named, 224

X-, 66

X-formula, 160

XR-, 66

XU-, 66

frame, 33

antisymmetric, 34

convergent, 57

conversely well-founded, 34

intuitionistic, 161

irreflexive, 34
 Noetherian, 34
 partially ordered, 34
 path frame, *see* path frame
 reflexive, 34
 strongly connected, 58
 transitive, 34
 weakly conversely well-founded, 34
 weakly Noetherian, 34
 fusion closure, 153

G

Gale–Stewart game, *see* game
 game, 20
 Δ_α^0 , 22
 Π_α^0 , 22
 Σ_α^0 , 22
 Borel, 22
 closed, 22
 determined, 21
 initial position of, 20
 open, 22
 position of, 20
 predecessor function of, 20
 space of, 20
 good branch, 75, 87, 193

H

height
 of a control, 184
 of a tree, 9
 of a vertex, 9
 henceforth, 67, 160
 hypersequent, 69

I

iLTL, 160
 iLTL frame, 161
 induction

 strong, 119
 inductive base, 125
 inductive invariant, 125
 external, 125
 internal, 125
 interpolation template, 233
 interpolation variable, 234
 intuitionistic order, 161
 intuitionistic pre-order, 161
 invariant, 230
 invariant map, 14

K

Kripke frame, *see* frame
 Kripke model, *see* model

L

label, *see* world label
 label valuation, *see* valuation, label
 valuation
 labelled transition system, 224
 labelling, 10
 labelling map, 10, 224
 labelling set, 10
 Λ -isomorphism, 10
 Λ -labelling, 10
 language
 see ω -language, 22
 leaf, 9
 bad, 50
 empty, 233
 good, 50
 letter, 22
 level, 9
 limit closure, 153
 literal, 31, 66, 160, 223
 LTS, 224

M

MES, *see* modal equational system

modal action, *see* action
 modal distance, 195
 modal equational system, 226
 modal logic, 32

- characterisation of, 34
- decidable, 56

 modal μ -calculus, 223
 modal pair, 50
 modal pruning, 51
 modal pruning map, 51
 model, 33

- iLTL, 161
- intuitionistic, 161
- serial, 67

modus ponens, 32
 μ -calculus, *see* modal μ -calculus

N

name, 99, 206, 228

- preserved, 229
- R-, 206
- reset, 229
- U-, 206

 NBA, 23
 necessitation, 32
 negation, 32
 next, 66, 160

O

ω -language, 22

- Büchi-recognisable, 23
- ω -regular, 23

 ω -unravelling, 12
 ω -word, 22

P

path, 9

- reachability path, *see* reachability path12

rooted, 9
 path frame, 153
 play, 20

- consistent, 21
- partial, 20
- total, 20
- win, 20

 pre-interpolant, 234
 predecessor, 9

- immediate, 9

 prefix, 8
 prefix topology, 19
 premise

- active, 233
- passive, 233
- trivial, 233

 prime witness, 125
 priority, 226

- higher, 226

 proof, 17

- CTL_o^{*} proof, 103
- CTL_{ind}^{*} proof, 115
- CTL_∞^{*} proof, 75
- CTL_{reg}^{*} proof, 96
- G3S4^o proof, 41
- GL_{circ} proof, 38
- Grz_{circ} proof, 42
- iLTL_o proof, 210
- iLTL_∞ proof, 170
- iLTL_{reg} proof, 205
- inductible, 125
- JS proof, 230
- K4^o proof, 46
- low-unfolding, 125
- normal, 231
- weakly thinning, 125

 proof system, *see* system17
 proof-search game, 87
 proof-search tree, 49, 84, 179

Q

- quantifier
 - existential, 66
 - universal, 66

R

- R-signature, *see* signature
- reachability, 12
- reachability path, 12
 - circular, 13
- refutation, 50, 86, 193
- relational control, *see* control, relational
 - reflexive, 164
 - rooted, 164
 - transitive, 164
 - tree-like, 164
- release, 66, 160
- repeat, 11
 - reachable, 12
 - successful, 46, 50, 230
- reset, 229
- root, 9
 - of a pruning, 181
- rule, 16
 - axiomatic, 17
 - conclusion of, 17
 - instance of, 17
 - name of, 17
 - premise of, 17
- rule schema
 - instance of, *see* rule, instance of
- run, 22, 23
 - accepting, 23

S

- satisfaction, 33, 67, 161, 163, 225
- SCC, 34
- sequence, 7

- finite, 7
- infinite, 7
- sequent
 - A-sequent, 68
 - abstract, 16
 - active, 71
 - annotated, 228
 - closed, 228
 - of CTL*, 68
 - E-sequent, 68
 - existential, 68
 - guarded, 228
 - of iLTL, 164
 - interpretation of, 36, 228
 - labelled, 164
 - modal, 36
 - in the μ -calculus, 228
 - principal, 71
 - provable, 17
 - side, 71
 - similar, 110
 - universal, 68
 - valid, 36, 228
- sequent calculus, *see* calculus
- sequent rule, *see* rule
- signature, 77
- size
 - of a CTL* formula, 66
 - of a CTL* hypersequent, 69
 - of a CTL* sequent, 68
 - of an iLTL formula, 160
 - of an iLTL sequent, 164
- skeleton, 22
 - run of, *see* run
 - states of, 22
 - transition relation of, 22
- sound rule, 115
- soundness, 33, 34
- state, 33, 224
- stationary premise, 45, 50

strategy, 21
 winning, 21
 strong chain condition, *see* chain condition, strong
 strong induction, *see* induction, strong
 substitution
 as a map, 32
 rule of, 32
 subsumption, 226
 subsumption order, 226
 subtree, 8
 succedent, 36, 164, 228
 successor, 9
 immediate, 9
 suffic closure, 153
 suffix, *see* tail
 symbol, 22
 system, 17
 cyclic, 18

T

tail, 8
 temporal successor function, 161
 thinning, 229
 trace
 bad, 75, 170
 of formulas, 168
 good, 75, 170
 label-stable, 169
 left, 169
 of type R, 169
 R-trace, 169
 right, 169
 side-stable, 169
 of type U, 169
 U-trace, 169
 transition relation, 224
 transition system, *see* labelled transition system
 transitional premise, 45, 50

tree, 8
 finitely branching, 9
 finitely-labelled, 10
 isomorphism, 10
 labelled, 10
 Λ -labelled, 10
 level of, *see* level
 low, 10
 regular, 11
 root of, *see* root
 rooted, 9
 sub-, *see* subtree
 thin, 10
 with back-edges, *see* tree with back-edges
 tree with back-edges, 11

U

U-signature, *see* signature
 unfolding companions, 125
 uniform interpolant, 232
 uniform interpolation, 232
 universal derivation, 108
 until, 66, 160

V

valuation, 33, 161
 label valuation, 163
 variable
 bound, 224
 free, 224
 variable assignment, 224
 vertex, 8
 axiomatic, 17
 branching, 9
 final, 9
 modal, 38, 46, 49, 69, 85, 103, 167, 209
 thinning, 103, 209
 vocabulary, 232

W

- WCC, 48
- weak chain condition, see chain condition, weak48
- well-named, 228
- width
 - of a control, 184
 - of a tree, 9
 - of a vertex, 9
- winning set, 20
- word, *see* ω -word
- world, 33
- world label, 162
 - dominating, 169
- world relation
 - active, 166
 - principal, 166
 - side, 167

List of Figures

1.1	The Borel hierarchy of a metrizable space.	19
2.1	Rules of the system $G3$	37
2.2	Rules of the system GL_{circ}	38
2.3	Additional rules for the system $G3S4$	40
2.4	A proof of the axiom gl in $G3S4^\circ$	41
2.5	Rules of the system Grz_{circ}	42
2.6	The modal rule of the system $G3K4$	45
2.7	A $K4^\circ$ proof of $\Box(\Box(p \rightarrow \Box p) \rightarrow p) \rightarrow \Box p$	46
3.1	Non-fixpoint, logical rules of the system CTL_∞^*	70
3.2	Fixpoint rules of the system CTL_∞^*	71
3.3	Structural rules of the system CTL_∞^*	72
3.4	A CTL_∞^* proof of $(\top U p) \vee (\perp R \bar{p})$	76
3.5	Non-fixpoint, logical rules of the system CTL_o^*	100
3.6	Fixpoint rules of the system CTL_o^*	101
3.7	Structural rules of the system CTL_o^*	102
3.8	A CTL_o^* proof of the formula $(\top U p) \vee (\perp R \bar{p})$	104
3.9	A non-discarding CTL_o^* derivation.	106
3.10	Some of the rules of the system CTL_{ind}^*	116
3.11	More rules of the system CTL_{ind}^*	117
3.12	The rest of the rules of the system CTL_{ind}^*	118
3.13	A CTL_{ind}^* proof of the formula $(\top U p) \vee (\perp R \bar{p})$	135
3.14	Non-axiomatic rules of the Hilbert-style system CTL_{Hil}^*	137
4.1	An $iLTL$ model refuting the classical dualities between U and R	162
4.2	Non-fixpoint, logical rules of the system $iLTL_\infty$	165
4.3	Fixpoint rules of the system $iLTL_\infty$	166

4.4	Structural rules of the system $iLTL_{\infty}$	167
4.5	An infinite, not label-stable trace.	170
4.6	An $iLTL_{\infty}$ proof of the formula $X(pUq) \rightarrow (Xp)U(Xq)$	171
4.7	Admissible $iLTL_{\infty}$ rules for F and G	172
4.8	An $iLTL_{\infty}$ refutation of the formula $(Xp \rightarrow Xq) \rightarrow X(p \rightarrow q)$	196
4.9	The model built from the refutation depicted in Figure 4.8.	197
4.10	An $iLTL_{\infty}$ refutation of the formula $F(Xp \rightarrow q)$	198
4.11	The model built from the refutation depicted in Figure 4.10.	199
4.12	Non-fixpoint, logical rules of the system $iLTL_{\circ}$	207
4.13	Fixpoint rules of the system $iLTL_{\circ}$	208
4.14	Structural, non-weakening rules of the system $iLTL_{\circ}$	209
4.15	Structural, weakening rules of the system $iLTL_{\circ}$	210
5.1	Non-fixpoint, logical rules of the system JS	229
5.2	Fixpoint rules of the system JS	230
5.3	Structural rules of the system JS	231

List of Tables

2.1	Some modal logic axioms.	35
2.2	Some modal logics and characteristic frames thereof.	36

List of Symbols

Axioms of modal logic

4, 35
gl, 35
grz, 35
grz₁, 35
k, 32, 35
t, 35

Logics

CPC, 31, 33
CTL*, 65
GL, 34, 36
Grz, 34
Grz.2, 57
Grz.3, 58
iTTL, 160
K, 33
K4, 36
K4Grz, 35, 36
S4, 36
S4.2, 57
S4.2Grz, 57
S4.3, 58
S4.3Grz, 58
S4Grz, 34, 36

Proof rules

\Box^* , 233

\Box_{K4}° , 45
coll, 177
ind_s, 119
prun, 178
 $R \rightarrow^-$, 178

Proof systems

CTL_o*, 99
CTL_{Hil}*, 136
CTL_{ind}*, 115
CTL_{ind}^{*}-, 115
CTL_∞*, 69
CTL_{reg}*, 96
G3, 37
G3K4, 45
G3S4, 40
G3S4^o, 41
GL_{circ}, 38
Grz_{circ}, 42
iTTL_o, 207
iTTL_∞, 164
iTTL_{reg}, 205
JS, 228
K4^o, 45
S4.2^o, 58

Other symbols

$(\cdot)^\omega$, 7
 $(\cdot)^{<\omega}$, 7

- (\leq_1, \dots, \leq_n) , 8
 $(u, v)_{\mathcal{T}}$, 9
 $(u, v]_{\mathcal{T}}$, 9
 $<^0_T$, 9
 $[\cdot]$, 49
 $[u, \rightarrow)_{\mathcal{T}}$, 9
 $[u, v)_{\mathcal{T}}$, 9
 $[u, v]_{\mathcal{T}}$, 9
 \Vdash , 33, 161
 $\langle \sigma, \tau \rangle$, 21
 \preceq , 14, 109, 230
 \leq_{Ω} , 163
 $\langle \cdot \rangle$, 66, 160
 \models , 33, 67
 \prec_{Θ} , 99, 206, 229
 \triangleleft , 12
 $\llbracket \varphi \rrbracket_V^S$, 224
 \sim_{Λ} , 10
 \sqsubset , 8
 \sqsubseteq , 8
 $\uparrow_{\leq_{\Omega}}$, 164
 \vdash , 17, 33
 $s_{<n}$, 8
 $s_{>n}$, 8
 $s_{\geq n}$, 8
 $s_{\leq n}$, 8
 $\widehat{\cdot}$, 8
 $\ell_V(\cdot)$, 234
 \mathbf{A} , 66
 Act , 223
 \mathbf{A}^{∂} , 68
 $\mathbf{b}(\cdot)$, 228
 $\mathcal{B}(\cdot)$, 18
 $\mathbf{b}(\cdot)$, 99
 $[\mathbf{a}]$, 223
 $\square\Gamma$, 31
 $B(X, T)$, 18
 c_i , 11
 $\text{Clos}(\cdot)$, 73, 168
 $c_u(\cdot)$, 189
 Δ_{α}^0 , 19
 $\Delta_{\alpha}^0(X)$, 18
 $\langle \mathbf{a} \rangle$, 223
 $\diamond\Gamma$, 31
 $d_X(\cdot, \cdot)$, 195
 $(\cdot)^{\partial}$, 32, 66–68, 225
 \mathbf{E} , 66
 \mathbf{E}^{∂} , 68
 \mathbf{F} , 67, 160
 FC , 153
 Form_{\square} , 31
 $\text{Form}_{\text{CTL}^*}$, 66
 Form_{LTL} , 160
 Form_{μ} , 223
 \triangleright , 72, 167
 \mathbf{G} , 67, 160
 $\text{Inf}(\cdot)$, 7
 inv. , 230
 $\text{inv}(\cdot)$, 108
 $\lambda[w \mapsto s]$, 163
 $\mathcal{L}(\cdot)$, 23
 \mathcal{L}_{\square} , 31
 LC , 153
 $\mathcal{L}_{\text{CTL}^*}$, 65
 $\text{Lev}_{\alpha}(\mathcal{T})$, 9
 \mathcal{L}_{LTL} , 160
 \mathcal{L}_{μ} , 223
 μ , 223
 \mathbf{N} , 99, 206
 $\mathbf{N}_{\mathbf{R}}$, 206
 $\mathbf{N}_{\mathbf{U}}$, 206
 ν , 223
 $\Omega \setminus W$, 164
 $\Omega \setminus w$, 164
 $\Omega[w/w']$, 163
 Π_{α}^0 , 19
 $\Pi_{\alpha}^0(X)$, 18
 \oplus , 33
 Prop , 31, 65, 160, 223
 Prov , 87

R, 66, 160
Ref, 87
Rep, 11
SC, 153
 \triangleright , 72
 Σ_α^0 , 19
 $\Sigma_\alpha^0(X)$, 18
 $\Sigma_I(\mathcal{G})$, 21
 $\Sigma_{II}(\mathcal{G})$, 21
Sub(\cdot), 31, 223
 \mathcal{T}° , 11
 \mathcal{T}^ω , 12
 $t_u(\cdot)$, 126
 $t_{\mathcal{T}}(\cdot)$, 126
U, 66, 160
Var, 223
Voc(\cdot), 232
 W_Ω , 163
X, 66, 160
 X^n , 125

Samenvatting

Dit proefschrift, waarvan de titel in het Nederlands luidt *Cyclische Bewijssystemen voor Modale Dekpuntlogica's*, gaat over cyclische en niet-welgefundeerde bewijssystemen voor modale dekpuntlogica's, met of zonder expliciete dekpuntkwantoren.

In de cyclische en niet-welgefundeerde bewijstheorie zijn bewijzen toegestaan met oneindige vertakkingen of paden, zolang ze maar voldoen aan bepaalde correctheidsvoorwaarden die de geldigheid van de conclusie garanderen. In dit proefschrift ontwerpen we een aantal cyclische en niet-welgefundeerde systemen: een cyclisch systeem voor de zwakke Grzegorzcyk modale logica $K4Grz$, gebaseerd op onze uitleg van het verschijnsel van *cyclische kompanen*; en niet-welgefundeerde en cyclische systemen voor de volledige berekeningsboom-logica CTL^* en de intuïtionistische lineaire temporele logica $iLTL$. Alle systemen zijn snedevrij, en de cyclische systemen voor $K4Grz$ en $iLTL$ hebben volstrekt eindige correctheidsvoorwaarden.

Ten slotte gebruiken we een cyclisch systeem voor de modale μ -calculus om een bewijs te verkrijgen van de uniforme interpolatie-eigenschap voor deze logica dat verschilt van het originele, op automaten gebaseerde bewijs.

Abstract

This thesis is about cyclic and ill-founded proof systems for modal fixpoint logics, with and without explicit fixpoint quantifiers.

Cyclic and ill-founded proof-theory allow proofs with infinite branches or paths, as long as they satisfy some correctness conditions ensuring the validity of the conclusion. In this dissertation we design a few cyclic and ill-founded systems: a cyclic one for the weak Grzegorzcyk modal logic $\mathbf{K4Grz}$, based on our explanation of the phenomenon of *cyclic companionship*; and ill-founded and cyclic ones for the full computation tree logic CTL^* and the intuitionistic linear-time temporal logic iLTL . All systems are cut-free, and the cyclic ones for $\mathbf{K4Grz}$ and iLTL have fully finitary correctness conditions.

Lastly, we use a cyclic system for the modal μ -calculus to obtain a proof of the uniform interpolation property for the logic which differs from the original, automata-based one.

Titles in the ILLC Dissertation Series:

- ILLC DS-2018-08: **Srinivasan Arunachalam**
Quantum Algorithms and Learning Theory
- ILLC DS-2018-09: **Hugo de Holanda Cunha Nobrega**
Games for functions: Baire classes, Weihrauch degrees, transfinite computations, and ranks
- ILLC DS-2018-10: **Chenwei Shi**
Reason to Believe
- ILLC DS-2018-11: **Malvin Gattinger**
New Directions in Model Checking Dynamic Epistemic Logic
- ILLC DS-2018-12: **Julia Ilin**
Filtration Revisited: Lattices of Stable Non-Classical Logics
- ILLC DS-2018-13: **Jeroen Zuiddam**
Algebraic complexity, asymptotic spectra and entanglement polytopes
- ILLC DS-2019-01: **Carlos Vaquero**
What Makes A Performer Unique? Idiosyncrasies and commonalities in expressive music performance
- ILLC DS-2019-02: **Jort Bergfeld**
Quantum logics for expressing and proving the correctness of quantum programs
- ILLC DS-2019-03: **Andras Gilyen**
Quantum Singular Value Transformation & Its Algorithmic Applications
- ILLC DS-2019-04: **Lorenzo Galeotti**
The theory of the generalised real numbers and other topics in logic
- ILLC DS-2019-05: **Nadine Theiler**
Taking a unified perspective: Resolutions and highlighting in the semantics of attitudes and particles
- ILLC DS-2019-06: **Peter T.S. van der Gulik**
Considerations in Evolutionary Biochemistry
- ILLC DS-2019-07: **Frederik Mollerstrom Lauridsen**
Cuts and Completions: Algebraic aspects of structural proof theory
- ILLC DS-2020-01: **Mostafa Dehghani**
Learning with Imperfect Supervision for Language Understanding

- ILLC DS-2020-02: **Koen Groenland**
Quantum protocols for few-qubit devices
- ILLC DS-2020-03: **Jouke Witteveen**
Parameterized Analysis of Complexity
- ILLC DS-2020-04: **Joran van Apeldoorn**
A Quantum View on Convex Optimization
- ILLC DS-2020-05: **Tom Bannink**
Quantum and stochastic processes
- ILLC DS-2020-06: **Dieuwke Hupkes**
Hierarchy and interpretability in neural models of language processing
- ILLC DS-2020-07: **Ana Lucia Vargas Sandoval**
On the Path to the Truth: Logical & Computational Aspects of Learning
- ILLC DS-2020-08: **Philip Schulz**
Latent Variable Models for Machine Translation and How to Learn Them
- ILLC DS-2020-09: **Jasmijn Bastings**
A Tale of Two Sequences: Interpretable and Linguistically-Informed Deep Learning for Natural Language Processing
- ILLC DS-2020-10: **Arnold Kochari**
Perceiving and communicating magnitudes: Behavioral and electrophysiological studies
- ILLC DS-2020-11: **Marco Del Tredici**
Linguistic Variation in Online Communities: A Computational Perspective
- ILLC DS-2020-12: **Bastiaan van der Weij**
Experienced listeners: Modeling the influence of long-term musical exposure on rhythm perception
- ILLC DS-2020-13: **Thom van Gessel**
Questions in Context
- ILLC DS-2020-14: **Gianluca Grilletti**
Questions & Quantification: A study of first order inquisitive logic
- ILLC DS-2020-15: **Tom Schoonen**
Tales of Similarity and Imagination. A modest epistemology of possibility

- ILLC DS-2020-16: **Ilaria Canavotto**
Where Responsibility Takes You: Logics of Agency, Counterfactuals and Norms
- ILLC DS-2020-17: **Francesca Zaffora Blando**
Patterns and Probabilities: A Study in Algorithmic Randomness and Computable Learning
- ILLC DS-2021-01: **Yfke Dulek**
Delegated and Distributed Quantum Computation
- ILLC DS-2021-02: **Elbert J. Booij**
The Things Before Us: On What it Is to Be an Object
- ILLC DS-2021-03: **Seyyed Hadi Hashemi**
Modeling Users Interacting with Smart Devices
- ILLC DS-2021-04: **Sophie Arnoult**
Adjunction in Hierarchical Phrase-Based Translation
- ILLC DS-2021-05: **Cian Guilfoyle Chartier**
A Pragmatic Defense of Logical Pluralism
- ILLC DS-2021-06: **Zoi Terzopoulou**
Collective Decisions with Incomplete Individual Opinions
- ILLC DS-2021-07: **Anthia Solaki**
Logical Models for Bounded Reasoners
- ILLC DS-2021-08: **Michael Sejr Schlichtkrull**
Incorporating Structure into Neural Models for Language Processing
- ILLC DS-2021-09: **Taichi Uemura**
Abstract and Concrete Type Theories
- ILLC DS-2021-10: **Levin Hornischer**
Dynamical Systems via Domains: Toward a Unified Foundation of Symbolic and Non-symbolic Computation
- ILLC DS-2021-11: **Sirin Botan**
Strategyproof Social Choice for Restricted Domains
- ILLC DS-2021-12: **Michael Cohen**
Dynamic Introspection

ILLC DS-2021-13: **Dazhu Li**

Formal Threads in the Social Fabric: Studies in the Logical Dynamics of Multi-Agent Interaction

ILLC DS-2022-01: **Anna Bellomo**

Sums, Numbers and Infinity: Collections in Bolzano's Mathematics and Philosophy

ILLC DS-2022-02: **Jan Czajkowski**

Post-Quantum Security of Hash Functions

ILLC DS-2022-03: **Sonia Ramotowska**

Quantifying quantifier representations: Experimental studies, computational modeling, and individual differences

ILLC DS-2022-04: **Ruben Brokkelkamp**

How Close Does It Get?: From Near-Optimal Network Algorithms to Suboptimal Equilibrium Outcomes

ILLC DS-2022-05: **Lwenn Bussière-Carae**

No means No! Speech Acts in Conflict

ILLC DS-2023-01: **Subhasree Patro**

Quantum Fine-Grained Complexity

ILLC DS-2023-02: **Arjan Cornelissen**

Quantum multivariate estimation and span program algorithms

ILLC DS-2023-03: **Robert Paßmann**

Logical Structure of Constructive Set Theories

ILLC DS-2023-04: **Samira Abnar**

Inductive Biases for Learning Natural Language

ILLC DS-2023-05: **Dean McHugh**

Causation and Modality: Models and Meanings

ILLC DS-2023-06: **Jialiang Yan**

Monotonicity in Intensional Contexts: Weakening and: Pragmatic Effects under Modals and Attitudes

ILLC DS-2023-07: **Yiyan Wang**

Collective Agency: From Philosophical and Logical Perspectives

ILLC DS-2023-08: **Lei Li**

Games, Boards and Play: A Logical Perspective

ILLC DS-2023-09: **Simon Rey**

Variations on Participatory Budgeting

ILLC DS-2023-10: **Mario Giulianelli**

Neural Models of Language Use: Studies of Language Comprehension and Production in Context



INSTITUTE FOR LOGIC, LANGUAGE AND COMPUTATION