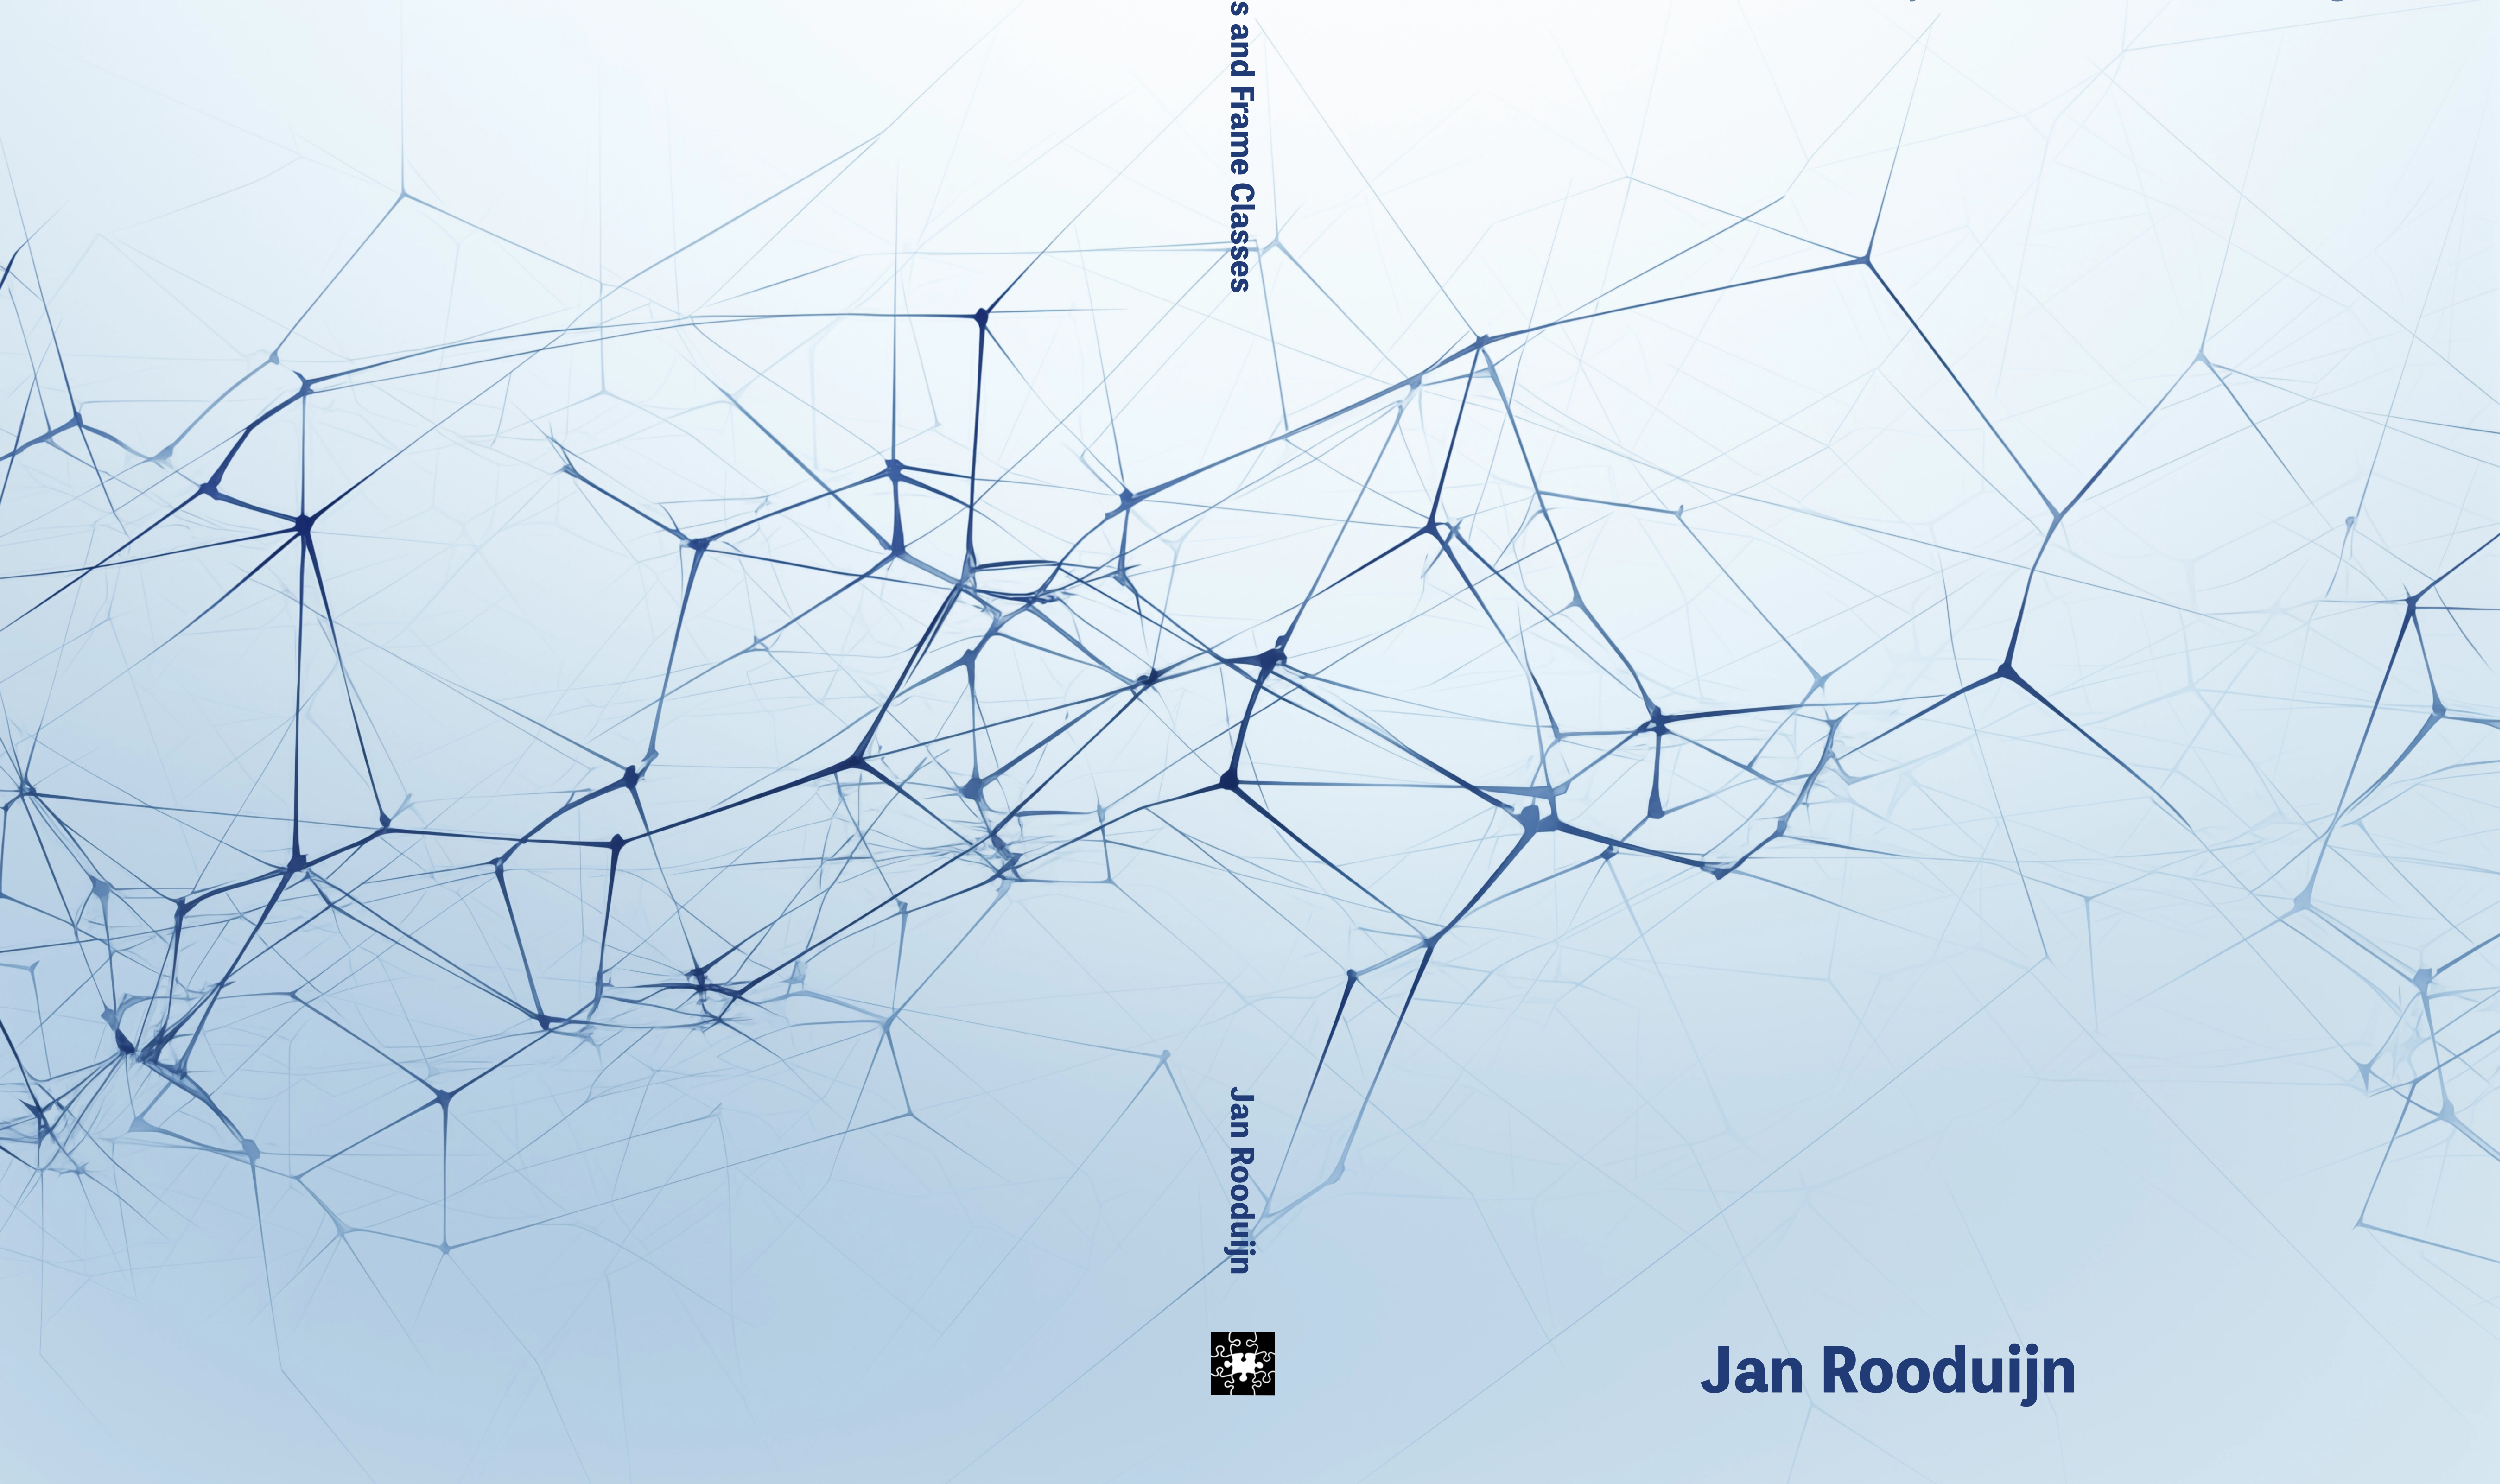

Fragments and Frame Classes

Towards a Uniform Proof Theory for Modal Fixed Point Logics

Jan Rooduijn

Fragments and Frame Classes
Jan Rooduijn

Fragments and Frame Classes

Towards a Uniform Proof Theory for Modal Fixed Point Logics

Jan Rooduijn

Fragments and Frame Classes

Towards a Uniform Proof Theory for Modal Fixed Point Logics

ILLC Dissertation Series DS-2024-01

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/

The research for this doctoral thesis has been made possible by a grant from the Dutch Research Council (NWO), project number 617.001.857.

Cover design by Midjourney V5.2.
Printed and bound by your Ipskamp Printing.

ISBN: 978-94-6473-362-4

Fragments and Frame Classes

Towards a Uniform Proof Theory for Modal Fixed Point 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 Aula der Universiteit
op woensdag 31 januari 2024, te 11.00 uur

door Johannes Maarten Willem Rooduijn
geboren te Amsterdam

Promotiecommissie

Promotor:	prof. dr. Y. Venema	Universiteit van Amsterdam
Copromotor:	dr. J.F. Marti	Universität Zürich
Overige leden:	prof. dr. J.F.A.K. van Benthem	Universiteit van Amsterdam
	dr. B. van den Berg	Universiteit van Amsterdam
	dr. A. Das	University of Birmingham
	dr. N.J.S. Enqvist	Stockholm University
	prof. dr. M. Lange	University of Kassel
	prof. dr. F. Roelofsen	Universiteit van Amsterdam
	prof. dr. S.J.L. Smets	Universiteit van Amsterdam

Faculteit der Natuurwetenschappen, Wiskunde en Informatica

Contents

Acknowledgments

Numerous people have contributed to this thesis in one way or another, whom I would like to thank.

I am most indebted to my promotor, Yde Venema. The privilege of meeting with you weekly has helped me become a better mathematician. Next to helping me choose interesting topics to work on, I have learnt a lot from your ability to combine deep intuition with strict rigor. You made me develop a habit of always trying to understand the essence of concepts, preferably by representing them pictorially, without compromising on precision. I also want to thank you for the unwavering moral support you have given me over the last years. Since the beginning of the project you have shown complete trust in me. From encouraging me to follow my interests and teaching my own course, to letting me spend a semester abroad. Because of your encouragement, I left each of our meetings with more confidence than I had when entering it.

Thanks, also, to my copromotor, Johannes Marti. You spent many days reading my - sometimes badly written - drafts, and your comments have greatly improved the presentation of this thesis. But your contribution comprises more than just that. I think we had already discussed each part of this thesis at the time I was first working on it. That is a testament both to your interest in our subject, and to your inclination to help other people. I hope that you can continue to inspire others with your academic enthusiasm (instead of becoming, in your words, a miserable JAVA-developer).

I have been lucky to be part of a relatively large research group at the ILLC. I thank Bahareh Afshari, Marianna Girlando, Lide Grotenhuis, Johannes Kloibhofer and Guillermo Menéndez Turata for the open and curious atmosphere, and the fun social activities.

Navigating a PhD can be a solitary journey, especially when your initial years coincide with a pandemic. Luckily, I had a solid network of support. Even though I found the research community intimidating at first, I soon understood it is comprised of many interesting, committed and welcoming researchers, which I would

also like to thank. In particular, I have fond memories of helpful discussions with Matteo Acclavio, Sergei Artemov, Lev Beklemishev, Gianluca Curzi, Anupam Das, Sebastian Enqvist, Malvin Gattinger, Iris van der Giessen, Helle Hansen, Tobias Kappé, Tadeusz Litak, Graham Leigh, Paul-André Melliès, Robin Martinot, Lukas Melgaard, Barteld Kooi, Revantha Ramanayake, Francisco Trucco Dalmas. A special mention to Lukas Zenger, who was an excellent co-author.

I am also grateful to other fellow mathematical logic PhD students at the ILLC. Robert Passmann and Ned Wontner, you have been there from the very beginning, and have provided a valuable source of companionship. Later on, Daniël Otten and Rodrigo Almeida joined our ranks, the latter of whom seems to have single-handedly orchestrated a doubling in both social and academic activities for our community.

One of the highlights of my PhD trajectory is the semester I spent in the United States. Even though only one month of this period comprised my stay at Cornell, it has led to the final chapter of this thesis. Thank you Alexandra Silva, Dexter Kozen, and Todd Schmid for the very warm welcome.

I would also like to thank the committee members for reading and evaluating this thesis: Johan van Benthem, Benno van den Berg, Anupam Das, Sebastian Enqvist, Martin Lange, Floris Roelofsen and Sonja Smets.

Lastly, I would like to thank my family and friends. My dear parents and sister, for always supporting me. And my friends, both in Utrecht and Amsterdam, for lending an ear or offering distraction whenever needed. Thank you Rik and Marie, for being my friends for life and now also my paranymphs. Thank you, dearest Eva, for always being there.

Chapter 1

Introduction

This thesis is about the proof theory of modal fixed point logics. In this introduction we shall give an informal introduction to this topic. Moreover, we will describe the particular goals of our research, existing related work, and our own contributions. Most of what is informally discussed here will be made formal in the next chapter.

1.1 Modal fixed point logics

Modal fixed point logics are a class of formalisms extending modal logic by so-called fixed point operators. We shall first introduce modal logic, and then a relatively simple modal fixed point logic called `PDL`. Thereafter we shall introduce the archetypal modal fixed point logic: the modal μ-calculus.

1.1.1 Modal logic

Modal logic was originally invented by philosophers to formalise the concepts of *possibility* and *necessitation*. It features two *modal operators* $\Diamond$ and $\Box$, where $\Diamond p$ means that some statement p is possibly true, and $\Box p$ means that p is necessarily true. An interesting example of a modal logical validity is given by the duality of $\Diamond$ and $\Box$:

$$\Diamond \neg p \text{ is equivalent to } \neg \Box p,$$

which means that it is possible that p is not true if and only if it is not necessary that p is true.

Since its inception, modal logic has been extended and reinterpreted in various ways. For instance, the modal operators have been interpreted as speaking about *belief*, *knowledge*, *provability* or *temporality*, rather than possibility and necessity. There is also a wealth of other modal operators, often related to each other in interesting ways. Mathematical tools have been developed to study the whole

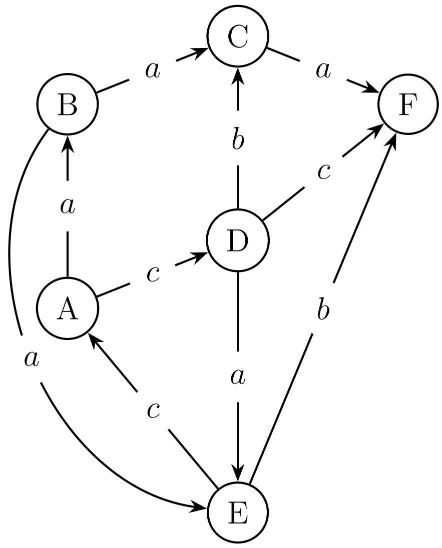

Figure 1.1: An example of a relational structure

landscape of modal logics. In this thesis we will only use a fraction of these tools. For an introduction to the field we refer the reader to [15]

Computer science is a particular field in which modal logic has found many applications. Modal logic is usually interpreted in *relational structures*, which are directed graphs, often labelled with some additional information.

Consider for instance the structure in Figure 1.1. A common view in computer science interprets the nodes of this graph as states of some machine, and the arrows as program executions. For instance, if the machine is in state A, then after executing program a it will be in state B, whereas executing program c will put it in state D. Under this interpretation the modal formula $\langle a \rangle x$ expresses that an execution of the program a *possibly* leads to a state where x is true, whereas the formula $[a]x$ expresses that it *necessarily* does. Formally, this respectively corresponds to *at least one* a-arrow pointing to the a state where x is true, or *every* arrow pointing to a state where x is true. If we say that A is true precisely at state A, and B is true precisely at state B, et cetera, one can thus verify that the formula $\langle a \rangle B$ is true at A and the formula $\langle a \rangle C$ is true at B. In contrast, while the formula $[a]B$ is true at A as well, the formula $[a]C$ is *not* true at B.

In the next section we will see how this perspective on modalities as programs has lead to extensions of the just-described basic modal logic.

1.1.2 Propositional Dynamic Logic

In [40], Fischer & Ladner introduced Propositional Dynamic Logic (or `PDL` for short). This is an extension of modal logic, based on the interpretation of the modalities as program executions. Characteristic of `PDL` is the fact that modalities (in the context of PDL also called *programs*) can be combined into new modalities, just like formulas can be combined into new formulas.

For instance, if a and b are programs, then there there is a program $a; b$, which first runs a and then runs b. One of the most interesting program constructors is the *Kleene star* $-^*$. Given a program a, the program a^* runs the program a a finite, but non-deterministically chosen, number of times. In other words, a^* might terminate immediately, or run the program a one time, or two times, et cetera.

Returning to Figure 1.1, this means that from the state A, every state except for D is reachable by an execution of the program a^*.

Let us again denote the modality corresponding to the program a^* by $\langle a^* \rangle$. This modality can be called an (implicit) *fixed point operator*, because the meaning of $\langle a^* \rangle p$ is the least fixed point of $x \mapsto p \vee \langle a \rangle x$. In other words, $\langle a^* \rangle p$ is the least solution for x satisfying

$$x \text{ is equivalent to } p \vee \langle a \rangle x.$$

This statement has two components. Firstly, $\langle a^* \rangle p$ is a *fixed point* of $x \mapsto p \vee \langle a \rangle x$, which means that applying it to $\langle a^* \rangle p$ returns $\langle a^* \rangle p$ itself. Spelling this out, we have

$$\langle a^* \rangle p \text{ is equivalent to } p \vee \langle a \rangle \langle a^* \rangle p.$$

Hence $\langle a^* \rangle p$ is true at some state s if and only if p is true at s, or $\langle a^* \rangle p$ is true at some state t reachable from s by an a-arrow, which means that in t either p is true or $\langle a^* \rangle p$ is true in some state u reachable by an a-arrow, which means that... and so on.

Secondly the fixed point $\langle a^* \rangle p$ is the *least* among all fixed points of the functions above. This means that for any φ such that φ is equivalent to $p \vee \langle a \rangle \varphi$, we have

$$\langle a^* \rangle p \text{ implies } \varphi.$$

The inclusion of this fixed point operator makes PDL much more expressive than basic modal logic, since it allows one to make statements about arbitrarily long paths. Nevertheless, many of the properties that make basic modal logic so well behaved are retained in the extension to PDL.

The modality $\langle a^* \rangle$ also plays an important role in the application of PDL in the field of *formal verification*. This branch of computer science aims to use mathematical tools to prove that certain programs behave correctly. Using the modality $\langle a^* \rangle$, one could for instance state that by repeatedly applying the program a, some desired state can be reached (this is called a *liveness* property). The dual modality, *i.e.* $[a^*]$ can be used to express that something bad never happens when repeatedly executing the program a (this is called a *safety* property).

1.1.3 The modal μ-calculus

The idea for the modal μ-calculus is to extend PDL by not only having the fixed point operator $-^*$, which is the least fixed point of the function $x \mapsto p \vee \langle a \rangle x$, but allowing one to take fixed points of *all* suitable functions.

To see how this works, consider for instance the function

$$x \mapsto [a]x \wedge [b]x \wedge [c]x$$

Recall that, for some arrow label r, the formula $[r]x$ is true at some state if for every state reachable by an r-arrow, the statement x is true. Hence, in Figure 1.1, the formula $[a]x \wedge [b]x \wedge [c]x$ is true in some state s, if x is true in every state reachable from s by *any* arrow. Now what might be a fixed point of this function? Certainly, if x is true in *every* state, then $[a]x \wedge [b]x \wedge [c]x$ is true in every state as well. This shows that the set of all states is a fixed point of the function above. In particular, it is the *greatest* fixed point, simply because it is the greatest possible set of states in this particular model. The modal μ-calculus can express this greatest fixed point using a quantifier-like operator ν, namely by $\nu x([a]x \wedge [b]x \wedge [c]x)$. Dually, there is a μ-operator which expresses the least fixed point. As we will see later in this thesis, it turns out that $\mu x([a]x \wedge [b]x \wedge [c]x)$ is true precisely in those states from which there is no infinite path. Hence, in the states C and F.

The modal μ-calculus was introduced in its current form by Dexter Kozen [59]. It can be seen as an extension of PDL, in the sense that every property expressible by PDL is also expressible by the modal μ-calculus. The converse does not hold. In fact, the least fixed point given as example above is a property which cannot be expressed by PDL.

The modal μ-calculus is a very interesting logic for several reasons. First, like PDL, it retains many of the desirable properties of basic modal logic, despite the even further gain in expressive power. Second, it has a very interesting theory, connecting it to combinatorial game theory, (co)algebra, automaton theory, and more standard techniques in basic modal logic. Finally, a seminal result by Janin & Walukiewicz characterises the modal μ-calculus as the bisimulation-invariant fragment of monadic second-order logic [53], where basic modal logic is the bisimulation-invariant fragment of first-order logic [10].

The model theory of the modal μ-calculus is relatively well understood. For instance, it has been known for a long time that the modal μ-calculus is bisimulation-invariant [53]. By exploiting the connection with automata, the small model property and decidability of the modal μ-calculus was first shown by Emerson & Streett [101]. Moreover, the modal μ-calculus enjoys *uniform interpolation*, as was shown by D'Agostina & Hollenberg [30, 31] (also using the connection with automata). More recently, Fontaine & Venema have used model-theoretic methods to obtain syntactic characterisations of semantic properties of formulas of the modal μ-calculus [42]. On the other hand, the proof theory of the modal

μ-calculus is a notoriously difficult and underdeveloped field. This will be the topic of the next section.

1.2 Proof theory

Once we have a logic, such as the modal μ-calculus, and a semantics, for instance in the form of relational structures, we can in principle check if a given formula is true in a given model. However, this does not directly give us a way to check if a formula is *valid*, *i.e.* true in every model.

To show that a formula is valid, one usually gives a *proof*. For instance, to show that the formula $\Box\top$ is valid, we might argue as follows. Suppose s is a state in some relational structure. By definition $\top$ is true in every state, so in particular it is true in every state reachable from s. Since s was chosen arbitrarily, the formula $\Box\top$ is true in every state of every model, and thus valid.

The goal of proof theory is to give a formal account of a proof such as the one informally given above. This involves specifying formal *axioms* and *rules* that can be used to form a proof, as well as specifying the shape that proofs should have. In the end, one obtains a formal notion of proof, which can then be the subject of mathematical study in its own right. Ideally, this formal notion of proof satisfies certain desirable properties, which makes it easier to prove results about proofs. For instance, in some nice proof systems every proof can be rewritten into a proof of a certain *normal form*. These proofs have a more predictable structure, making it easier to reason about them. In the end, the results one obtains about a proof system can sometimes even be used to obtain results about the logic it formalises, such as consistency, decidability, and interpolation.

In this thesis we shall mainly see two types of (formal) proof calculi, which together exemplify a common theme in proof theory. The first type is that of *Hilbert-style* calculi. Characteristically, these calculi have a *modus ponens* rule, allowing one to derive ψ from the hypotheses $\varphi \to \psi$ and φ. The other type of calculi are called *Gentzen-style*. These calculi are much better structured, owing in part to the fact that they manipulate lists or sets of formulas (usually called *sequents*) rather than a single formula at once. While Gentzen-style calculi sometimes also feature a rule akin to modus ponens, called the *cut rule*, ideally this rule is superfluous, in the sense that every valid formula has a proof that does not use the cut rule (this is an example of the aforementioned normal form). The well-structuredness of Gentzen-style proofs, especially those which are cut-free, makes them much more suitable for proof-theoretical analysis than Hilbert-style proofs. In particular they often satisfy a form of the *subformula property*: the formulas occurring in the premisses of any rule application are subformulas of the formulas occurring in the conclusion. As a consequence once usually obtains a bound on the number of formulas occurring in the whole proof, since they must, by transitivity, all be subformulas of formulas in the proof's conclusion.

1.2.1 ... of modal logic

Modal logics have traditionally been axiomatised using Hilbert-style proof systems. Every so-called *normal* modal logic contains the rule of $\frac{p}{\Box p}$ of *necessitation* and an axiom, called K, of the form $\Box(p \to q) \to \Box p \to \Box q$. Other modal logics are often formed by adding additional axioms. This will be discussed later in the section about frame conditions.

Although Hilbert-style proof systems of this kind have been very well studied, the tools used have been mostly model-theoretic. This is not surprising because, as mentioned above, Hilbert-style proof systems are not very suitable for proof-theoretic analysis. Completeness for these kind of proof systems is usually established through a Henkin-style canonical model construction. In this construction, one forms a model in which the states are maximally consistent sets of formulas, and one shows that every state satisfies exactly the formulas that it contains. Hence, every consistent set of formulas is satisfiable, implying that every valid formula is provable.

Amongst the earliest Gentzen-style proof systems for modal logic are the systems presented in [79]. They manipulate *sequents* of the from $\Gamma \Rightarrow \Delta$, where both Γ and Δ are finite sets of formulas. Nowadays, the modal logic K is most commonly axiomatised by adding the following rule to a sequent calculus for classical propositional logic:

$$\mathsf{K}\,\frac{\Gamma \Rightarrow \varphi}{\Box\Gamma \Rightarrow \Box\varphi}$$

Here $\Box\Gamma$ is shorthand for the set for formulas $\{\Box\varphi \mid \varphi \in \Gamma\}$.

The resulting sequent calculus is sound and complete for the modal logic K. In fact, it is even complete when one omits the cut rule, and this can be shown by systematically transforming proofs into cut-free proofs [80].

1.2.2 ... of program logics

Let us now turn to the proof theory of program logics, in particular of PDL. We first consider Hilbert-style proof systems. A crucial obstacle for using the canonical model method described above, is the fact that PDL is not *compact*. That is, there are unsatisfiable sets of PDL-formulas, of which every finite subset is satisfiable. Consider for instance the set

$$\{\langle a^* \rangle p, \neg p, [a]\neg p, [a][a]\neg p, [a][a][a]\neg p, \ldots\}.$$

Any state satisfying this set will satisfy $\langle a^* \rangle p$. So there is some number of executions of the program a, after which p is true. However, for any n, the state will also satisfy $[a]^n \neg p$, where $[a]^n$ is shorthand for n times $[a]$. This means that for each n, the proposition p is *not* true after n executions of a, a contradiction. We leave it to the reader to think about why every finite subset of the above set *is* satisfiable.

As a result, the canonical model method briefly sketched above is not applicable to PDL. Indeed, since derivations are finite objects, not every maximally consistent set is satisfiable. This problem turns out to have a relatively simple solution: one can use a finitary version of the canonical model construction [62]. This technique is closely related to the method of *filtration* and will play an important role in Chapter 4 of this thesis.

Constructing Gentzen-style proof systems for PDL is much more challenging. The primary obstacle is the inductive nature of the Kleene star. In a Hilbert-style system this operator is often axiomatised by the following *induction rule*

$$[a^*](p \to [a]p) \to (p \to [a^*]p)$$

It is difficult to translate this rule into a nice Gentzen-style rule. One often sees something that looks as follows (see *e.g.* [7, 84]).

$$\mathsf{ind}\,\frac{\Gamma \Rightarrow \psi, \Delta \qquad \psi \Rightarrow \varphi \qquad \psi \Rightarrow [a]\psi}{\Gamma \Rightarrow [a^*]\varphi, \Delta}$$

This rule is problematic because the ψ behaves as a cut formula, preventing a nice proof search procedure. This can also be explained as follows: where it is often said that the cut rule forces one to guess an appropriate *lemma*, possibly from some other mathematical field, the rule ind forces one to guess an *induction invariant*.

Alternatively, systems with a so-called *ω-rule* have been proposed. This rule is of the form

$$\omega\text{-}\mathsf{ind}\,\frac{\Gamma \Rightarrow [a]^n\varphi, \Delta}{\Gamma \Rightarrow [a^*]\varphi, \Delta}$$

The problem with this rule is that it does not support a finitary notion of proof, because it has infinitely many premisses. This can be somewhat salvaged by appealing to the finite model property of PDL, but this is generally agreed to be proof-theoretically unsatisfactory.

A third option is to consider proofs with infinitely long branches, rather than infinitely many premisses. This technique is more well-known in the context of the modal μ-calculus and will therefore be discussed in the next section.

We end this section by briefly mentioning another approach to axiomatising PDL, namely using *tableaux* [17, 47]. Although tableaux often axiomatise *satisfiability* rather than *validity*, they are in principle very similar to Gentzen-style proof systems. The difficulty described above for axiomatising PDL using a Gentzen-style proof system therefore also applies to tableaux for PDL. Moreover, solutions to this problem proposed in the tableaux community are sometimes similar to those of non-well-founded proof theory. A difference is that tableaux are often developed with computational efficiency as main objective, whereas proof theory cares more about the well-structuredness and readability of formal proofs.

1.2.3 ... of the modal μ-calculus

Already in the seminal paper [59], a very elegant Hilbert-style axiomatisation for the modal μ-calculus was proposed. It consists of an axiomatisation for the smallest normal modal logic K, together with an axiom and a rule characterising the μ-operator as a least fixed point (and dually for the ν-operator):

$$\varphi[\mu x\varphi/x] \to \mu x\varphi \qquad \frac{\varphi[\psi/x] \to \psi}{\mu x\varphi \to \psi} \qquad \nu x\varphi \to \varphi[\nu x\varphi/x] \qquad \frac{\psi \to \varphi[\psi/x]}{\psi \to \nu x\psi}$$

Since the modal μ-calculus can be seen as an extension of PDL, most of the proof-theoretical difficulties remain. In particular, compactness also fails for the modal μ-calculus, preventing the use of a standard canonical model construction.

Unlike PDL, however, the modal μ-calculus is not susceptible to the finitary canonical model method either. This was originally shown by Kozen in [59], and is caused by the fact that the method of filtration fails for the modal μ-calculus.

It turned out to be very difficult to prove the completeness of the Hilbert-style proof system presented above. Almost 20 years after its introduction by Kozen, completeness was finally obtained by Walukiewicz [107], building on joint work with Niwiński [77].

Central to Walukiewicz's proof is the use of automaton theory and certain combinatorial games called *parity games*. Interestingly, these are linked to the modal μ-calculus through certain tableaux. As already mentioned above, where Gentzen-style proof systems establish the validity of some set of formulas, tableau systems establish satisfiability. Apart from this dual perspective, tableau systems are in principle very similar to Gentzen-style proof systems.

The tableau system by Niwiński & Walukiewicz in [77] omits an induction rule akin to the rule ind above. Instead, it relies on infinite branches to axiomatise the recursive behaviour of the fixed-point operators. Proofs in this systems are therefore called *non-well-founded proofs*, and such proofs turn out to be a very powerful tool in the proof theory of modal fixed point logics in general. Importantly, despite the fact that non-well-founded proofs are a priori infinite objects, they often admit a finite representation in the form of finite trees with back edges (also called *cyclic proofs*).

The tableau system by Niwiński & Walukiewicz was later dualised into a proof system for validity by Dax, Hofmann & Lange [36], and Studer [102]. By further omitting bells and whistles such as *definition lists*, the resulting proof system acquires a form of the subformula property, called the *closure property*. One could therefore say that a crucial step for proving the completeness of a Hilbert-style proof system for the modal μ-calculus, was to first consider a well-behaved Gentzen-style calculus.

Following their successful application to the modal μ-calculus, non-well-founded proofs have become a popular topic of study the proof theory of modal fixed point logics and beyond. Such proof systems have been developed and studied for fragments and variants of the modal μ-calculus [39, 2, 73], linear logic [8, 78], Kleene algebra [34, 35], arithmetic [97, 13, 32] and more [1, 18, 63, 5]. One of the main goals of this thesis is to uniformly construct cyclic proof systems for (fragments of) the modal μ-calculus interpreted over various *frame classes*.

1.3 Frame conditions

A common theme in the study of modal logic, is to restrict attention to relational structures (specifically *frames*) satisfying certain properties, often called *frame conditions*. For instance, under the epistemic interpretation of the modal operator $\Box$, the formula $[a]p$ means "agent a knows that p" is true. One usually assumes that knowledge presupposes truth, *i.e.* that $[a]p \to p$ always holds. It turns out that this requirement corresponds to restricting attention to *reflexive* relational structures, that is, those where every state has an a-arrow to itself.

1.3.1 Hilbert-style systems

Hilbert-style systems are particularly well-suited for axiomatising frame conditions. While frame conditions are usually formulated using formulas of first-order logic, such a first-order formula in many cases has a modal *correspondent*, in the sense that the frames satisfying the first-order formula are exactly the frames in which the modal formula is valid. The field studying this connection is called *correspondence theory*, and for an overview we refer the reader to [15].

A central result in correspondence theory is Sahlqvist's Theorem [90]. This theorem gives a sufficient syntactic condition for a (basic) modal formula φ to have a first-order correspondent. In addition Sahlqvist's Theorem states that such a formula φ is *canonical*, meaning that it is valid in the *canonical frame* underlying the canonical model of any logic including φ. This can be used to show that the Hilbert-style proof system K, with φ as an additional axiom, is sound and complete with respect to the class of all frames that satisfy the first-order correspondent of φ. Below are three examples of axioms to which Sahlqvist's Theorem applies.

$$(\mathsf{T})\ \Box p \to p, \qquad (\mathsf{4})\ \Box p \to \Box\Box p, \qquad (\mathsf{B})\ p \to \Box\Diamond p.$$

These axioms correspond, respectively, to reflexivity, transitivity and symmetry.

Sahlqvist's Theorem is a sweeping result with applications to Hilbert-style proof systems. Unfortunately, it does not easily generalise to modal fixed point logics. As mentioned above, the canonical model method fails for most modal fixed point logics, because of the lack of compactness. Recently, Kikot, Shapirovsky & Zolin showed how to apply the *finitary* canonical model method to a class of `PDL`-like logics that admit the method of filtration [56]. In combination with Sahlqvist's Theorem this leads to a general completeness result for `PDL` interpreted over restricted classes of frames. In Chapter 3 of this thesis we show how to extend this result to a fragment of the modal μ-calculus, called the *continuous* modal μ-calculus.

For the modal μ-calculus itself, results about proof systems with frame conditions are rare. An important reason is that, as mentioned above, even the finitary canonical model method fails. Hemaspaandra in [51] obtained an important negative result. She showed that there is a frame class over which even a small fragment of the modal μ-calculus becomes highly undecidable. Remarkably, basic modal logic is very well behaved over the same frame class. In particular, it has a sound and complete Hilbert-style proof system. Since highly undecidable logics cannot have nice proof systems, this result puts a limit on the goals of this thesis: we cannot hope to obtain a nice proof system for the modal μ-calculus with respect to every frame class over which basic modal logic is well behaved.

Apart from this result, the question of proof systems for the modal μ-calculus over different frame classes has not gotten much attention. In part this is due to the complex nature of the field, but it may also be because the modal μ-calculus is mostly studied from a more practical computer science perspective, rather than a more mathematical perspective emphasising theory and general results. Only recently a general soundness and completeness result for Hilbert-style systems for the modal μ-calculus interpreted over certain *weakly transitive* frame classes was presented [9].

1.3.2 Gentzen-style systems

Already for basic modal logic, developing Gentzen-style proof systems for different frame conditions in a uniform and modular way has proven to be quite challenging. Below are three example of rules corresponding, respectively, to the Hilbert-style axioms T, $\mathsf{4}$, and B.

$$\mathsf{T}\,\frac{\varphi, \Gamma \Rightarrow \Delta}{\Box\varphi, \Gamma \Rightarrow \Delta} \qquad \mathsf{K4}\,\frac{\Gamma, \Box\Gamma \Rightarrow \varphi}{\Box\Gamma \Rightarrow \Box\varphi} \qquad \mathsf{KB}\,\frac{\Gamma \Rightarrow \varphi, \Box\Delta}{\Box\Gamma \Rightarrow \Box\varphi, \Delta}$$

Adding each of these axioms to a standard sequent calculus for K (with the cut rule!) yields a sound and complete calculus with respect to the respective frame condition [80]. However, these rules are not modular, in the sense that adding both $\mathsf{K4}$ and KB does not yield a complete calculus for the class of frames that are both transitive and symmetric. Rather, $\mathsf{K4}$ and KB must be combined in a non-trivial way into a *new* rule. Moreover, while the systems with T and $\mathsf{K4}$ remain complete when omitting the cut rule, the system with KB does not (by for instance the same counterexample as given for $\mathsf{S5}$ in [80]).

In the quest for a more satisfactory uniform proof theory for modal logics, much attention has been focused on extending the structure of sequents. A modest extension is the *hypersequent* framework. Hypersequents are simply sets of sequents. In [66], Ori Lahav uniformly constructs hypersequent calculi for a wide range of (basic) modal logics. Hypersequents are inspired by the semantics, as the multiple sequents in a hypersequent represent multiple states in a relational structure. This perspective has been taken further in the form of *nested sequents* [20], also appearing under name *tree hypersequents* [83]. Nested sequents can treat reflexivity, transitivity and symmetry in a modular way, without the need for a cut rule. Even closer to the semantics are the *labelled sequents* by Sara Negri [75]. while labelled sequents and nested sequents can both be used to treat many different frame conditions, a downside is that, even for proofs without the cut rule, there is no a priori bound on the number of sequents occurring in a given proof. As a result, each frame condition requires ad-hoc arguments to show that proof search is terminating.

So far, our discussion of Gentzen-style systems for modal logics satisfying various frame conditions has only been about basic modal logic. Results for modal fixed point logics are very scarce. There is a non-well-founded labelled Gentzen-calculus for `PDL` by Docherty & Rowe [37]. Although they do not pursue this direction explicitly, their choice for a labelled system is motivated by its ability to handle different frame conditions. Unfortunately, their system does not support a finite notion of proof, as the lack of a bound on the number of sequents occurring in a proof prevents them from turning their non-well-founded system into a cyclic system. Another example is a cyclic proof system for the two-way modal μ-calculus (which is similar to the modal μ-calculus interpreted over symmetric frames) by Afshari et. al [2]. Although this is a proper cyclic system for the modal μ-calculus interpreted over a certain frame class, it is specifically designed for this single frame class, and it does not provide a framework for uniformly treating multiple frame classes at once.

1.3.3 Our contributions

In Chapter 3, we extend Lahav's uniform hypersequents to a fragment of the modal μ-calculus called *modal logic with the master modality*. Our hypersequent calculi are non-well-founded and are made cyclic by using the method of *focus*, originally due to Lange & Stirling [68]. Like Lahav, we are only able to prove cut-free completeness for a subset of the frame conditions covered, and our subset is smaller than that of Lahav. To the best of our knowledge, this is the first uniform proof-theoretical treatment of modal fixed point logics characterised by frame conditions.

Before we move on to Chapter 4, there is an Intermezzo, in which we propose an abstract framework for so-called *annotated* non-well-founded proof systems. The main application of this abstract framework is to give a general proof of the *bounded proof property*. In a proof system with this property, every provable sequent has a proof whose size is bounded by a computable function of the size of the sequent. Consequently, any logic with the bounded proof property is decidable. The abstract tools developed in the Intermezzo apply to the hypersequent calculi of Chapter 3 and they were initially developed with this application in mind. Although it later turned out that the same results can also be obtained by known game-theoretical techniques, we still believe that the ideas presented in the Intermezzo are of independent value.

In Chapter 4, we consider Hilbert-style proof systems, in particular axiomatic extensions of the system by Kozen. As mentioned above, the only known completeness proof for this system is very complex and relies heavily on intricate automata-theoretic machinery. We show that for a certain known fragment, called the *continuous modal μ-calculus*, finitary canonical models can be used to show completeness. We moreover show that the continuous modal μ-calculus admits the method of filtration, and we extend both this result and the completeness result to a wide range of frame classes. As `PDL` properly embeds into the continuous modal μ-calculus, our result can be seen as a generalisation of the aforementioned result by Kikot et. al [56].

In the second-to-last chapter, Chapter 5, we consider the two-way modal μ-calculus. This is an extension of the modal μ-calculus, where each modality $\langle a \rangle$ is assigned a corresponding *backward* modality $\langle \breve{a} \rangle$. One can also see the resulting logic as the interpretation of the modal μ-calculus over a restricted frame class. Namely, the class of all frames where the relation interpreting the modality $\langle a \rangle$ is the converse of that interpreting $\langle \breve{a} \rangle$. We construct a sound and complete cyclic proof system for a fragment of the two-way modal μ-calculus, called the *alternation-free* fragment. For this we combine the multi-focus annotations originally due to Marti & Venema (see [73]) with the novel technique of *trace atoms*.

Chapter 6 is a somewhat of an outlier, because it is on Kleene Algebra, which is strictly not a *modal* logic. Before we go on to explain the contributions of this final chapter we will first give a brief informal introduction to Kleene Algebra.

1.4 Kleene Algebra

Although it was developed separately, at this point it is convenient to introduce Kleene Algebra as a certain reduction of the program logic `PDL`. Recall that in `PDL`, modalities are induced by programs. *Kleene Algebra* forgets about the modal logical aspect, and instead focuses on exclusively axiomatising the equivalence of programs. For instance, the program $a; a^*b \cup b$, which non-deterministically either runs a followed by a^* and then followed by b, or immediately runs b, is equivalent to the program a^*b.

Kleene Algebra originates as an axiomatisation of the equational theory of the algebra of regular languages [57]. After several non-algebraic axiomatisations, for instance by Salomaa [91] and Conway [28], the purely algebraic axiomatisation by Kozen is now most commonly used [60]. A *Kleene Algebra* is an algebraic structure satisfying the axioms and rules of this axiomatisation.

Kleene Algebra can be called a *fixed point logic*, because it uses the following rule to axiomatise the Kleene star.

$$\frac{a; x \cup b \leq x}{a^*b \leq x}$$

This rule states that the program a^*b is the least fixed point of the function $x \mapsto a; x \cup b$.

Recently Das & Pous constructed a cyclic proof system for Kleene Algebra [34]. In later joint work with Doumane, they gave an alternative proof of the completeness of Kozen's axiomatisation by translating their proofs into Kozen's system [33].

1.4.1 Guarded Kleene Algebra with Tests

Although Kleene Algebra axiomatises an abstract notion of programs, it does not capture the conventional programming constructs of if-then-else statements and while-loops. This feature can be obtained by augmenting Kleene Algebra with so-called *tests*. Formally, a test is a Boolean expression t. By adding these to the language of Kleene Algebra, we can construct the following expressions.

$$t; a \cup \neg t; b \qquad\qquad (t; a)^*; \neg t$$

The first expressions captures the statement `if` t `then` a `else` b, whereas the second captures the loop `wile` t `do` a. Remarkably, extending the language of Kleene Algebra by tests does not increase its computational complexity.

Guarded Kleene Algebra with Tests, or `GKAT` for short, is a reasonably expressive fragment of Kleene Algebra with Tests, with much lower computational complexity. This increase in efficiency is obtained by restricting the union operator $\cup$ and the Kleene star $-^*$ to their *guarded* counterparts. In other words, Guarded Kleene Algebra with Tests allows *only* if-then-else statements, instead of non-deterministic choice, and *only* while-loops, instead of the (non-deterministic) Kleene star. Since most practical programs are deterministic, `GKAT` retains much of the practical value of Kleene Algebra with Tests, while reducing the complexity of deciding program equivalence to *nearly linear*[1] time. This a great reduction from the PSPACE-completeness of deciding program equivalence in Kleene Algebra (whether with or without tests).

1.4.2 Our contributions

In Chapter 6 we propose a cyclic proof system for `GKAT`. Our system is inspired by the system for Kleene Algebra in [34], but we show that `GKAT` requires less complex sequents than Kleene Algebra. We show that the system is sound and complete. Moreover, we propose an inequational axiomatisation for `GKAT` and give

[1] This is a technical term which will be explained in Chapter 6.

a partial translation from the cyclic system into the ineqational system. This may be a first step towards solving the open problem of finding a purely algebraic proof system for GKAT.

1.5 Sources of the material

- Chapter 2 was written specifically for this thesis.
- Chapter 3 is based on the following two publications, the second of which is joint work with Lukas Zenger.

  [86] Jan Rooduijn. Cyclic hypersequent calculi for some modal logics with the master modality. In *30th International Conference on Automated Reasoning with Analytic Tableaux and Related Methods, TABLEAUX*, volume 12842 of *Lecture Notes in Computer Science*, pages 354–370. Springer, 2021

  [89] Jan Rooduijn and Lukas Zenger. An analytic proof system for common knowledge logic over S5. In *14th Conference on Advances in Modal Logic, AiML*, pages 659–680. College Publications, 2022

- The Intermezzo is based on unpublished work. It was presented at the 2021 Workshop on Proof Theory and its Applications in Funchal, Madeira.
- Chapter 4 is based on the following publication, which is joint work with Yde Venema.

  [87] Jan Rooduijn and Yde Venema. Filtration and canonical completeness for continuous modal mu-calculi. In *12th International Symposium on Games, Automata, Logics, and Formal Verification, GandALF*, volume 346 of *EPTCS*, pages 211–226, 2021

- Chapter 5 is based on the following publication, which is also joint work with Yde Venema.

  [88] Jan Rooduijn and Yde Venema. Focus-style proofs for the two-way alternation-free μ-calculus. In *29th International Workshop on Logic, Language, Information, and Computation, WoLLIC*, volume 13923 of *Lecture Notes in Computer Science*, pages 318–335. Springer, 2023

- Chapter 6 is as yet unpublished. It is based on joint work with Dexter Kozen and Alexandra Silva and was presented at the 2022 Workshop on Proof Theory and its Applications in Utrecht, The Netherlands.

Chapter 2

Introduction to the proof theory of modal fixed point logics

Most of this thesis is concerned with so-called *modal fixed point logics*. These are logics extending basic modal logic with operators capable of expressing recursive behaviour. Most modal fixed point logics, in particular all of those appearing in this thesis, can be interpreted in the archetypical modal fixed point logic: the modal μ-calculus.

In this chapter we shall introduce the modal μ-calculus, its proof theory, and some of the fragments that play a role in this thesis. Although our presentation will be largely self-contained, it is helpful if the reader has some familiarity with basic modal logic, for instance by having read the first four chapters of [15].

This chapter contains no original material, apart perhaps from how it is presented. The presentation is heavily inspired by Yde Venema's treatment in [106].

2.1 The modal μ-calculus

2.1.1 Syntax

For the rest of this thesis we fix a countably infinite set P of propositional variables. Recall that an occurrence of some propositional variable p in some formula φ is said to be *positive* if it is in the scope of an even number of negations.

2.1.1. Definition. Given a set D of *actions*, the syntax $\mu\mathtt{ML}(\mathsf{D})$ of the *modal μ-calculus* is generated by:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \langle a \rangle \varphi \mid [a]\varphi \mid \mu x \varphi \mid \nu x \varphi,$$

where $p \in \mathsf{P}$, $a \in \mathsf{D}$ and in the formation of $\eta x \varphi$, with η ranging over μ and ν, the variable x occurs only positively in φ.

2.1.2. Remark. If D is a singleton, say $\mathsf{D} = \{a\}$, we call $\mu\mathtt{ML}(\mathsf{D})$ the *monomodal μ-calculus* and write $\Diamond$ and $\Box$ rather than $\langle a \rangle$ and $[a]$. We denote this language simply by $\mu\mathtt{ML}$.

To avoid notational clutter, we will often work with $\mu\mathtt{ML}$ rather than with $\mu\mathtt{ML}(\mathsf{D})$ for some larger set D of actions. In almost every case the generalisation to $\mu\mathtt{ML}(\mathsf{D})$ is routine.

A formula of the form p or $\neg p$ is called a *literal*. A formula is called an *o-formula* if its main operator is o. The connectives $\{\neg, \vee, \wedge\}$ are called *propositional* and the connectives $\{\langle a \rangle, [a] \mid a \in \mathsf{D}\}$ are said to be *modal*. Finally, we use η to range over $\{\mu, \nu\}$, and denote by $\overline{\eta}$ the dual of η, *i.e.* $\overline{\mu} = \nu$ and $\overline{\nu} = \mu$.

We will use $\mu\mathtt{ML}(\mathsf{D})$*-formula*, or just *formula* when the specific language is clear from the context, to refer to a formula in $\mu\mathtt{ML}(\mathsf{D})$. Just like the quantifiers in first-order logic, the fixed point operators bind variables in a formula.

2.1.3. Definition. Given a $\mu\mathtt{ML}$-formula ξ, the sets $\mathrm{FV}(\xi)$ of *free variables* and $\mathrm{BV}(\xi)$ of *bound variables* of ξ are inductively defined by:

$$
\begin{array}{lll@{\qquad}lll}
\mathrm{FV}(p) & := & \{p\} & \mathrm{BV}(p) & := & \emptyset \\
\mathrm{FV}(\neg\varphi) & := & \mathrm{FV}(\varphi) & \mathrm{BV}(\neg\varphi) & := & \mathrm{BV}(\varphi) \\
\mathrm{FV}(\varphi \circ \psi) & := & \mathrm{FV}(\varphi) \cup \mathrm{FV}(\varphi) & \mathrm{BV}(\varphi \circ \psi) & := & \mathrm{BV}(\varphi) \cup \mathrm{BV}(\varphi) \\
\mathrm{FV}(\triangle\varphi) & := & \mathrm{FV}(\varphi) & \mathrm{BV}(\triangle\varphi) & := & \mathrm{BV}(\varphi) \\
\mathrm{FV}(\eta x\varphi) & := & \mathrm{FV}(\varphi) \setminus \{x\} & \mathrm{BV}(\eta x\varphi) & := & \mathrm{BV}(\varphi) \cup \{x\},
\end{array}
$$

where $\circ \in \{\vee, \wedge\}$, and $\triangle \in \{\Diamond, \Box\}$, and $\eta \in \{\mu, \nu\}$.

2.1.4. Example. Let $\varphi = \mu x(\Diamond x \vee p) \wedge x$. Then we have $\mathrm{FV}(\varphi) = \{p, x\}$ and $\mathrm{BV}(\varphi) = \{x\}$.

It will often be convenient to restrict attention to formulas with nice syntactic properties, such as those given in the following definitions.

2.1.5. Definition. A formula ξ is called *tidy* if $\mathrm{FV}(\xi) \cap \mathrm{BV}(\xi) = \emptyset$.

2.1.6. Remark. The modal μ-calculus is sometimes defined using variables of two sorts: *propositional variables* and *fixed point variables*. A *sentence* is then a formula where every fixed point variable is bound by some fixed point operator. Although our formulation only uses one sort of variables, tidy formulas can be seen as the analogue of sentences, where the free variables are the propositional variables and the bound variables are the fixed point variables.

Note that the formula φ in Example 2.1.4 is *not* tidy. The following proposition is immediate.

2.1.7. Proposition. *Any formula can be made tidy by uniformly renaming bound variables.*

2.1.8. Definition. A μML-formula is said to be in *negation normal form* if it belongs to the language generated by:

$$\varphi ::= p \mid \neg p \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \Diamond\varphi \mid \Box\varphi \mid \mu x\varphi \mid \nu x\varphi.$$

When working with formulas in negation normal form, we will often abbreviate $\neg p$ by $\overline{p}$.

In the next section we will define the semantics of μML-formulas. We will then see that every formula is equivalent to one in negation normal form. We will also see that every formula $\eta x\varphi$ is, in fact, a *fixed point* of the formula $\varphi(x)$. In other words, the formula $\eta x\varphi$ will be semantically equivalent to the formula $\varphi[\eta x\varphi/x]$, which is obtained by substituting $\eta x\varphi$ for x in φ. Substitution therefore plays a very important role in the syntax of the modal μ-calculus and deserves careful treatment.

2.1.9. Definition. We say that θ is *free for* x in ξ if no free occurrence of x in ξ occurs in the scope of a fixed point operator that binds a free variable of θ. More formally, we say that θ is *free for* x in ξ if one of the following holds:

- $\xi \in \mathsf{P}$;
- $\xi = \neg\varphi$ and θ is free for x in φ;
- $\xi = \varphi \circ \psi$ and θ is free for x in both φ and ψ;
- $\xi = \triangle\varphi$ and θ is free for x in φ;
- $\xi = \eta x\varphi$;
- $\xi = \eta y\varphi$ for some $y \neq x$ with $y \notin \mathrm{FV}(\theta)$, and θ is free for x in φ.

2.1.10. Definition. Suppose that θ is free for x in ξ. The *substitution* $\xi[\theta/x]$ of θ for x in ξ is obtained by replacing all free occurrences of x in ξ by θ. More formally, we define inductively:

- $x[\theta/x] := \theta$, and $y[\theta/x] := y$ for $y \neq x$.
- $(\neg\varphi)[\theta/x] := \neg\varphi[\theta/x]$.
- $(\varphi \circ \psi)[\theta/x] := \varphi[\theta/x] \circ \psi[\theta/x]$.
- $(\triangle\varphi)[\theta/x] := \triangle\varphi[\theta/x]$.
- $(\eta x\varphi)[\theta/x] := \eta x\varphi$.

- $(\eta y\varphi)[\theta/x] := \eta y(\varphi[\theta/x])$ for $y \neq x$.

The following, easily verifiable lemma, will be useful later on.

2.1.11. Lemma. *Let u be a variable not occurring in ξ. We have:*

- *if ξ is tidy and of the form $\eta x\varphi$, then $\varphi[u/x]$ is tidy;*
- *if θ is free for x in ξ, then θ is free for u in $\xi[u/x]$, and $\xi[\theta/x] = \xi[u/x][\theta/u]$.*

2.1.12. Remark. The above lemma will come in handy when inductively proving results about tidy formulas. If $\eta x\varphi$ is tidy, then φ need not be, whence we can in general not apply our induction hypothesis to φ. Instead, we take some u not occurring in φ and apply the induction hypothesis $\varphi[u/x]$, which is tidy.

The next lemma captures the main property making tidy formulas convenient to work with.

2.1.13. Lemma. *If $\eta x\varphi$ is tidy, then $\eta x\varphi$ is free for x in φ, and $\varphi[\eta x\varphi/x]$ is tidy as well. Moreover, if $\eta x\varphi$ is in negation normal form, then so is $\varphi[\eta x\varphi/x]$.*

Proof:
Suppose $y \in \mathrm{FV}(\eta x\varphi)$. Then by tidiness $y \notin \mathrm{BV}(\eta x\varphi)$, whence $y \notin \mathrm{BV}(\varphi)$. This shows that $\eta x\varphi$ is free for x in φ. For the tidiness of $\varphi[\eta x\varphi/x]$, suppose that $y \in \mathrm{FV}(\varphi[\eta x\varphi/x])$. Then either $y \in \mathrm{FV}(\varphi) \setminus x$, or $y \in \mathrm{FV}(\eta x\varphi)$. But these sets are equal by definition, whence by tidiness $y \notin \mathrm{BV}(\eta x\varphi)$. The result follows from the fact that $\mathrm{BV}(\varphi[\eta x\varphi/x]) \subseteq \mathrm{BV}(\eta x\varphi)$. Finally, the preservation of negation normal form follows directly from the positivity restriction on the bound variables of $\mu\mathsf{ML}$-formulas. $\square$

2.1.14. Definition. Let ξ be a formula. The set $\mathsf{Sfor}(\xi)$ of *subformulas* of ξ is the least set of formulas such that:

(i) $\xi \in \mathsf{Sfor}(\xi)$.

(ii) $\neg\varphi \in \mathsf{Sfor}(\xi)$ implies $\varphi \in \mathsf{Sfor}(\xi)$.

(iii) $\varphi \circ \psi \in \mathsf{Sfor}(\xi)$ implies $\varphi, \psi \in \mathsf{Sfor}(\xi)$ for $\circ \in \{\vee, \wedge\}$.

(iv) $\triangle\varphi \in \mathsf{Sfor}(\xi)$ implies $\varphi \in \mathsf{Sfor}(\xi)$ for $\triangle \in \{\Diamond, \Box\}$.

(v) $\eta x\varphi \in \mathsf{Sfor}(\xi)$ implies $\varphi \in \mathsf{Sfor}(\xi)$ for $\eta \in \{\mu, \nu\}$.

We write $\varphi \trianglelefteq \xi$ to indicate that φ is a subformula of ξ, and $\varphi \triangleleft \xi$ to indicate that it is a proper subformula.

The clause (v) of the definition of subformulas is an outlier from a semantic point of view. As we will see later, the meaning of all other operators are entirely truth functional: their truth-value depends on the truth-values of their subformulas. However, for a fixed point operator η, the truth of $\eta x\varphi$ depends on the truth of φ in *different models*. On the other hand, as mentioned above, $\eta x\varphi$ is always equivalent to $\varphi[\eta x\varphi/x]$. The following notion is therefore more appropriate for capturing formulas which are in some sense *semantically relevant* to ξ. It is known as the *(Fischer-Ladner) closure*.

2.1.15. DEFINITION. Let ξ be a tidy formula. The FL-*closure* $\mathsf{FL}(\xi)$ of ξ is the least set of formulas satisfying clause (i) - (iv) of Definition 2.1.14, as well as

(v') $\eta x\varphi \in \mathsf{FL}(\xi)$ implies $\varphi[\eta x\varphi/x] \in \mathsf{FL}(\xi)$ for $\eta \in \{\mu, \nu\}$.

If Ξ is a set of tidy formulas, we define $\mathsf{FL}(\Xi) := \bigcup_{\xi\in\Xi} \xi$. Using Lemma 2.1.13 it is immediate that if a formula is, respectively, tidy or in negation normal form, then so is every formula in its closure. It is also not hard to see that $\mathrm{FV}(\varphi) \subseteq \mathrm{FV}(\xi)$ and $\mathrm{BV}(\varphi) \subseteq \mathrm{BV}(\xi)$ for every $\varphi \in \mathsf{FL}(\xi)$.

We will now show that the FL-closure of a tidy formula ξ is always finite. In fact, the FL-closure has at most as many elements as the number of characters ξ has when viewed as a string. We first need the following auxiliary lemma.

2.1.16. LEMMA. *Let $\eta x\varphi$ be a tidy formula such that $x \in \mathrm{FV}(\varphi)$ and let u be some propositional variable not occurring in $\eta x\varphi$. Then:*

(i) $\eta x\varphi$ is free for u in every $\psi \in \mathsf{FL}(\varphi[u/x])$;

(ii) $\mathsf{FL}(\eta x\varphi) \subseteq \{\psi[\eta x\varphi/u] : \psi \in \mathsf{FL}(\varphi[u/x])\}$.

Proof:
For (i), suppose that $y \in \mathrm{BV}(\psi)$. Then $y \in \mathrm{BV}(\varphi[u/x])$ and thus $y \in \mathrm{BV}(\eta x\varphi)$. Hence by tidiness $y \notin \mathrm{FV}(\eta x\varphi)$, as required.

For (ii), it suffices to show that the set on the right-hand side, let us call it Σ, is closed under conditions (i) - (iv) of Definition 2.1.14 and (v') of Definition 2.1.15. For condition (i), note that, since $x \in \mathrm{FV}(\varphi)$, we have $u \in \mathsf{FL}(\varphi[u/x])$, whence $\eta x\varphi \in \Sigma$. Moreover, conditions (ii) - (iv) are satisfied by Σ, as the respective operators commute with substitution.

For condition (v'), suppose that $\lambda y\chi \in \Sigma$ with $\lambda \in \{\mu, \nu\}$. We first consider the case where $\lambda y\chi = \eta x\varphi$. In this case we indeed find $\varphi[\eta x\varphi/x] \in \Sigma$, since $\varphi[\eta x\varphi/x] = \varphi[u/x][\eta x\varphi/u]$. Now suppose that $\lambda y\chi \neq \eta x\varphi$. Because $\lambda y\chi \in \Sigma$, there must be some $\lambda y\theta \in \mathsf{FL}(\varphi[u/x])$ such that $\lambda y\chi = \lambda y\theta[\eta x\varphi/u]$. We have $\theta[\lambda y\theta/y] \in \mathsf{FL}(\varphi[u/x])$, and thus:

$$\Sigma \ni \theta[\lambda y\theta/y][\eta x\varphi/u] = \theta[\eta x\varphi/u][\lambda y\theta[\eta x\varphi/u]/y] = \chi[\lambda y\chi/y],$$

as required. □

Let us write $|\xi|_s$ for the length of ξ as a string.

2.1.17. Lemma. *For every tidy formula ξ it holds that $|\mathsf{FL}(\xi)| \leq |\xi|_s$.*

Proof:
We proceed by induction on the length of ξ (as a string). Suppose that the thesis has been proven for all φ of length smaller than ξ. We make a case distinction on the main connective of ξ.

- $\xi \in \mathsf{P}$. Then $\mathsf{FL}(\xi) = \{\xi\}$, whence $|\mathsf{FL}(\xi)| \leq |\xi|_s$.
- $\xi = \neg\varphi$. Then $\mathsf{FL}(\xi) \subseteq \{\xi\} \cup \mathsf{FL}(\varphi)$. It follows that:

$$|\mathsf{FL}(\xi)| \leq |\{\xi\} \cup \mathsf{FL}(\varphi)| \leq 1 + |\varphi|_s \leq |\xi|_s.$$

- $\xi = \varphi \circ \psi$. Then $|\mathsf{FL}(\xi)| \leq \{\xi\} \cup \mathsf{FL}(\varphi) \cup \mathsf{FL}(\psi) \leq 1 + |\varphi|_s + |\psi|_s \leq |\varphi \circ \psi|_s$.
- $\xi = \triangle\varphi$. Then $|\mathsf{FL}(\xi)| \leq |\{\xi\} \cup \mathsf{FL}(\varphi)| \leq 1 + |\varphi|_s \leq |\xi|_s$.
- $\xi = \eta x \varphi$. First consider the degenerate case where $x \notin \mathrm{FV}(\varphi)$. Then $|\mathsf{FL}(\xi)| = 1 + |\mathsf{FL}(\varphi)|$. By the induction hypothesis, it then follows that $|\mathsf{FL}(\xi)| \leq 1 + |\varphi|_s = |\xi|_s$. In case $x \in \mathrm{FV}(\varphi)$, we let u be a propositional variable not occurring in ξ. We have $\mathsf{FL}(\xi) \subseteq \{\psi[\xi/u] : \psi \in \mathsf{FL}(\varphi[u/x])\}$ by Lemma 2.1.16. Hence:

$$|\mathsf{FL}(\xi)| \leq |\mathsf{FL}(\varphi[u/x])| \leq |\varphi[u/x]|_s \leq 1 + |\varphi|_s \leq |\xi|_s,$$

  as required. □

We close this section by defining the notion of a *trace*. This notion gives another perspective on the closure. More importantly, it will be needed later to define the game semantics and a non-well-founded proof system for the modal μ-calculus.

2.1.18. Definition. A *trace* is a (possibly infinite) sequence (φ_n) of tidy formulas in negation normal form, such that for each two subsequent formulas φ_n, φ_{n+1} it holds that:

- φ_n is not of the form p or $\overline{p}$;
- $\varphi_n = \psi_1 \circ \psi_2$ implies $\varphi_{n+1} = \psi_i$ for some $i \in \{1, 2\}$;
- $\varphi_n = \triangle\psi$ implies $\varphi_{n+1} = \psi$;
- $\varphi_n = \eta x \varphi$ implies $\varphi_{n+1} = \varphi[\eta x \varphi/x]$.

We leave it to the reader to verify the following lemma.

2.1.19. Lemma. *For any tidy formula ξ in negation normal form,*

$$\mathsf{FL}(\xi) = \{\varphi \mid \text{ there is a trace } \xi = \varphi_0 \cdots \varphi_n = \varphi\}.$$

The following lemma will will be crucial in the next sections.

2.1.20. Lemma. *On every infinite trace $(\varphi_n)_{n\in\omega}$ there is a unique fixpoint formula $\eta x \chi$ which occurs infinitely often and is a subformula of φ_n for cofinitely many n.*

To prove it, we first need two technical lemmas.

2.1.21. Remark. The proof of Lemma 2.1.20 is quite tedious and not very relevant for the rest of this thesis. In fact, all modal fixed point logics we consider are *alternation free* (cf. Definition 2.1.31). These logics have a simpler trace structure for which Lemma 2.1.20 is an overkill. The reader only interested in the original results of this thesis is therefore encouraged to skip the following two lemmas and the proof of Lemma 2.1.20. The reason to nevertheless include them, is that they play a fundamental role in the (proof) theory of the modal μ-calculus.

2.1.22. Lemma. *Let $\eta x\varphi$ and ψ be tidy formulas in negation normal form such that $\psi \trianglelefteq \varphi[\eta x\varphi/x]$. Then we have either $\psi \trianglelefteq \eta x\varphi$, or $\eta x\varphi \trianglelefteq \psi$.*

Proof:
By induction on $\alpha \trianglelefteq \varphi$, we will show that $\psi \trianglelefteq \alpha[\eta x\varphi/x]$ implies $\psi \trianglelefteq \alpha$ or $\eta x\varphi \trianglelefteq \psi$. Note that this suffices, since $\psi \trianglelefteq \varphi$ implies $\psi \trianglelefteq \eta x\varphi$.

Let $\alpha \trianglelefteq \varphi$ be such that $\psi \trianglelefteq \alpha[\eta x\varphi/x]$ and suppose that the thesis holds for every proper subformula of α. We can assume that $x \in \mathrm{FV}(\alpha)$, for otherwise the lemma becomes trivial. We may further assume that $\psi \triangleleft \alpha[\eta x\varphi/x]$, *i.e.* that ψ is a *proper* subformula of $\alpha[\eta x\varphi/x]$, for otherwise the assumption that $x \in \mathrm{FV}(\alpha)$ implies that $\eta x\varphi \trianglelefteq \alpha[\eta x\varphi/x] = \psi$. We make a case distinction on the shape of α.

- Suppose that α is a literal. Since $x \in \mathrm{FV}(\alpha)$, we must have $\alpha = x$ or $\alpha = \overline{x}$. The latter is impossible, since that would mean that $\overline{x}$ is in the scope of ηx in $\eta x\varphi$. Therefore $\alpha = x$ and thus $\psi \trianglelefteq \alpha[\eta x\varphi/x] = \eta x\varphi$.

- Now suppose that $\alpha = \alpha_1 \circ \alpha_2$, for some $\circ \in \{\vee, \wedge\}$. Since $\psi \triangleleft \alpha[\eta x\varphi/x]$, we have $\psi \trianglelefteq \alpha_i[\eta x\varphi/x]$ for an $i \in \{1,2\}$. The result follows by the induction hypothesis.

- If $\alpha = \triangle\beta$ we argue similarly: since $\psi \triangleleft \alpha[\eta x\varphi/x]$ we have $\psi \trianglelefteq \beta[\eta x\varphi/x]$ and we can apply the induction hypothesis.

- Finally suppose that α is of the form $\lambda y\beta$. Since we assumed $x \in \mathrm{FV}(\alpha)$, it must be the case that $y \neq x$. We find that $\psi \triangleleft \alpha[\eta x\varphi/x] = \lambda y(\beta[\eta x\varphi/x])$. Hence $\psi \trianglelefteq \beta[\eta x\varphi/x]$ and we can again apply the induction hypothesis.

This finishes the proof. □

2.1.23. LEMMA. *Let $\varphi_0, \dots, \varphi_n$ be a finite trace. Then there is a single formula on the trace, which is a subformula of every formula on the trace.*

Proof:
We prove this by induction on n. The base case, where $n = 0$ is trivial. For the induction step, suppose the thesis holds for $n = k$. To prove it for $n = k + 1$, we apply the induction hypothesis to the trace $\varphi_1, \dots, \varphi_{k+1}$ to obtain an $i \in [1, k+1]$ such that $\varphi_i \trianglelefteq \varphi_j$ for every $j \in [1, k + 1]$.

We now make a case distinction on the shape of φ_0. If φ_0 is *not* a fixed point formula, then, by the fact that φ_0 has a successor in the trace, the main connective of φ_0 is either propositional or modal. In either case, we find $\varphi_i \trianglelefteq \varphi_1 \trianglelefteq \varphi_0$, which suffices.

Now suppose that φ_0 is a fixed point formula, say $\varphi_0 = \eta x \psi$. Then we must have $\varphi_1 = \psi[\eta x \psi / x]$ and thus $\varphi_i \trianglelefteq \psi[\eta x \psi / x]$. By the previous lemma, either $\varphi_i \trianglelefteq \varphi_0$ or $\varphi_0 \trianglelefteq \varphi_i$. In the former case we are done. In the latter case we have, by transitivity, that $\varphi_0 \trianglelefteq \varphi_j$ for each $j \in [0, k + 1]$, which also suffices. □

Proof of Lemma 2.1.20:
For every n, let $\psi_n \in \{\varphi_0, \dots \varphi_n\}$ be such that $\psi_n \trianglelefteq \varphi_i$ for every $i \in [0, n]$. The lemma follows from the fact that the sequence $(\psi_n)_{n \in \omega}$ must eventually be constant, for otherwise it would contain an infinite descending chain of proper subformulas. □

An infinite trace is called an *η-trace*, depending on the value of η in the previous lemma.

2.1.2 Semantics

Formulas of the modal μ-calculus are interpreted in the same structures as those of basic modal logic.

2.1.24. DEFINITION. A *Kripke model* $\mathbb{S}$ of type D consists of a set S of *states*, for each $a \in \mathsf{D}$ an *accessibility relation* $R_a \subseteq S \times S$, and a *valuation* $V : \mathsf{P} \to \mathcal{P}(S)$.

When $\mathsf{D} = \{a\}$ we often write R instead of R_a for the single accessibility relation of some given model.

Algebraic semantics

In this subsection we define a commonly used semantics of the modal μ-calculus. We must first introduce some notation.

2.1.25. Definition. Let $V : \mathsf{P} \to \mathcal{P}(S)$ be some valuation, and let $X \in \mathcal{P}(S)$. The valuation $V[x \mapsto X]$ is given by $V[x \mapsto X](x) = X$, and $V[x \mapsto X](p) = V(p)$ for every $p \neq x$. Given a Kripke model $\mathbb{S}$, we denote by $\mathbb{S}[x \mapsto X]$ the result of replacing its valuation function V by $V[x \mapsto X]$.

For $R \subseteq S \times S$, we write $R[s] := \{t \in S : sRt\}$ for the image of s under R.

2.1.26. Definition. The *meaning* $[\![\varphi]\!]^{\mathbb{S}} \subseteq S$ of a formula $\xi \in \mathcal{L}_\mu$ in $\mathbb{S}$ is inductively defined on the complexity of ξ:

$$
\begin{array}{llll}
[\![p]\!]^{\mathbb{S}} & := V(p) & [\![\neg\varphi]\!]^{\mathbb{S}} & := S \setminus [\![\varphi]\!]^{\mathbb{S}} \\
[\![\varphi \vee \psi]\!]^{\mathbb{S}} & := [\![\varphi]\!]^{\mathbb{S}} \cup [\![\psi]\!]^{\mathbb{S}} & [\![\varphi \wedge \psi]\!]^{\mathbb{S}} & := [\![\varphi]\!]^{\mathbb{S}} \cap [\![\psi]\!]^{\mathbb{S}} \\
[\![\langle a\rangle\varphi]\!]^{\mathbb{S}} & := \{s \in S \mid R_a[s] \cap [\![\varphi]\!]^{\mathbb{S}} \neq \emptyset\} & [\![[a]\varphi]\!]^{\mathbb{S}} & := \{s \in S \mid R_a[s] \subseteq [\![\varphi]\!]^{\mathbb{S}}\} \\
[\![\mu x\varphi]\!]^{\mathbb{S}} & := \bigcap\{X \subseteq S \mid [\![\varphi]\!]^{\mathbb{S}[x\mapsto X]} \subseteq X\} & [\![\nu x\varphi]\!]^{\mathbb{S}} & := \bigcup\{X \subseteq S \mid X \subseteq [\![\varphi]\!]^{\mathbb{S}[x\mapsto X]}\}
\end{array}
$$

We define the basic semantic notions of *truth*, *satisfiability*, *validity*, and *equivalence* in the usual way.

2.1.27. Remark. Recall that a prefixed point of an endofunction f on an ordered set L is an element x such that $f(x) \leq x$. Note that $\mu x\varphi$ is interpreted as the intersection of all prefixed points of the function

$$
\begin{aligned}
\varphi_x^{\mathbb{S}} &: \mathcal{P}(S) \to \mathcal{P}(S) \\
&: X \mapsto [\![\varphi]\!]^{\mathbb{S}[x\mapsto X]}.
\end{aligned}
$$

The positivity restriction on bounded variables guarantees that $\varphi_x^{\mathbb{S}}$ is a monotone function on the complete lattice $\mathcal{P}(S)$. By the Knaster-Tarski Theorem [104], the intersection of all prefixed point is the least fixed point of $\varphi_x^{\mathbb{S}}$. Dually, the interpretation of $\nu x\varphi$ is the greatest fixed point of $\varphi_x^{\mathbb{S}}$.

It is clear that the meaning of a formula does not change when uniformly renaming its bound variables. Hence, it follows from Proposition 2.1.7 that every formula is equivalent to a tidy formula. The following proposition shows that we may further assume that our formulas are in negation normal form.

2.1.28. Proposition. *There is a translation* $\mathsf{nnf} : \mu\mathtt{ML} \to \mu\mathtt{ML}$ *such that for every formula* φ*, the formula* $\mathsf{nnf}(\varphi)$ *is an equivalent formula in negation normal form.*

Proof:
We let nnf commute with every connective apart from negation. The translation of a formula of the form $\neg\psi$ depends on the main connective ψ:

$$
\begin{array}{llllll}
\mathsf{nnf}(\neg p) & := & \neg p & \mathsf{nnf}(\neg\neg\varphi) & := & \mathsf{nnf}(\varphi) \\
\mathsf{nnf}(\neg(\varphi \vee \psi)) & := & \mathsf{nnf}(\neg\varphi \wedge \neg\psi) & \mathsf{nnf}(\neg(\varphi \wedge \psi)) & := & \mathsf{nnf}(\neg\varphi \vee \neg\psi) \\
\mathsf{nnf}(\neg\Diamond\varphi) & := & \Box\mathsf{nnf}(\neg\varphi) & \mathsf{nnf}(\neg\Box\varphi) & := & \Diamond\mathsf{nnf}(\neg\varphi) \\
\mathsf{nnf}(\neg\mu x\varphi) & := & \nu x\mathsf{nnf}(\neg\varphi[\neg x/x]) & \mathsf{nnf}(\neg\nu x\varphi) & := & \mu x\mathsf{nnf}(\neg\varphi[\neg x/x])
\end{array}
$$

Clearly $\mathsf{nnf}(\varphi)$ is in negation normal form. We will demonstrate the equivalence of φ and $\mathsf{nnf}(\varphi)$ by showing that $\neg\mu x\varphi$ is equivalent to $\nu x\neg\varphi[\neg x/x]$, leaving the other cases to the reader. For any model $\mathbb{S} = (S, R, V)$, we have:

$$\begin{aligned}
[\![\neg\mu x\varphi]\!]^{\mathbb{S}} &= S \setminus [\![\mu x\varphi]\!] \\
&= S \setminus \bigcap\{X \subseteq S \mid [\![\varphi]\!]^{\mathbb{S}[x\mapsto X]} \subseteq X\} \\
&= \bigcup\{S \setminus X \mid [\![\varphi]\!]^{\mathbb{S}[x\mapsto X]} \subseteq X\} \\
&= \bigcup\{S \setminus X \mid (S \setminus X) \subseteq [\![\neg\varphi]\!]^{\mathbb{S}[x\mapsto X]}\} \\
&= \bigcup\{S \setminus X \mid (S \setminus X) \subseteq [\![\neg\varphi[\neg x/x]]\!]^{\mathbb{S}[x\mapsto S\setminus X]}\} \\
&= \bigcup\{Y \subseteq S \mid Y \subseteq [\![\neg\varphi[\neg x/x]]\!]^{\mathbb{S}[x\mapsto Y]}\} \\
&= [\![\nu x\varphi[\neg x/x]]\!]^{\mathbb{S}},
\end{aligned}$$

as required. □

As a consequence, we obtain a definable negation operator $\overline{}$ on the set of formulas in negation normal form, given by $\overline{\varphi} := \mathsf{nnf}(\neg\varphi)$. We leave it to the reader to verify that $\overline{\overline{\varphi}} = \varphi$.

2.1.29. Definition. Let Ξ be a set of tidy formulas in negation normal form. We write $\overline{\mathsf{FL}}(\Xi)$ for the least set of formulas such that for every $\xi \in \overline{\mathsf{FL}}(\Xi)$ it holds that $\overline{\xi} \in \overline{\mathsf{FL}}(\Xi)$ and $\mathsf{FL}(\xi) \subseteq \overline{\mathsf{FL}}(\Xi)$.

We sometimes write $\overline{\mathsf{FL}}(\xi)$ where we mean $\overline{\mathsf{FL}}(\{\xi\})$. We leave it to the reader to verify that $\overline{\mathsf{FL}}(\Xi) = \mathsf{FL}(\Xi) \cup \{\overline{\xi} \mid \xi \in \mathsf{FL}(\Xi)\}$. Finally, Ξ is said to be $\overline{\mathsf{FL}}$-*closed* whenever $\overline{\mathsf{FL}}(\Xi) = \Xi$.

Game semantics

In this thesis we shall mostly work with an alternative, equivalent, definition of the semantics of the modal μ-calculus, given by the *evaluation game*. This game-theoretic definition will only be given for tidy formulas in negation normal form. For now we will use game-theoretic notions in a rather informal way, but they will become formal in the next section.

Suppose we are given a formula φ, which is tidy and in negation normal form, and a model $\mathbb{S} = (S, R, V)$. The game $\mathcal{E}(\xi, \mathbb{S})$ is a board game played by two players called $\exists$ and $\forall$. The game is played on the board $\mathsf{FL}(\xi) \times S$, and at a position (φ, s) it is $\exists$'s objective to show that φ is true at s, and $\forall$'s objective to show the opposite. Whose turn it is (or who *owns*) at a particular position (φ, s) is determined by the main connective of φ, and given in the following table. The table also shows which moves are available to a given position's owner.

Position	Owner	Admissible moves
$(p, s), s \in V(p)$	$\forall$	$\emptyset$
$(p, s), s \notin V(p)$	$\exists$	$\emptyset$
$(\varphi \vee \psi, s)$	$\exists$	$\{(\varphi, s), (\psi, s)\}$
$(\varphi \wedge \psi, s)$	$\forall$	$\{(\varphi, s), (\psi, s)\}$
$(\Diamond\varphi, s)$	$\exists$	$\{\varphi\} \times R[s]$
$(\Box\varphi, s)$	$\forall$	$\{\varphi\} \times R[s]$
$(\eta x \varphi, s)$	$-$	$\{(\varphi[\eta x \varphi / x], s)\}$

Note that from a position of the form $(\eta x \varphi, s)$ there is only one available move, whence it does not matter who owns this position. If one of the players owns a position, but has no moves available to them, the match ends in a win for the other player. Due to the fixed point operators, the formulas occurring in a match are not necessarily strictly decreasing in length, and therefore matches can be of infinite length. Note that if $(\varphi_n, s_n)_{n\in\omega}$ is an infinite match, then its left projection $(\varphi_n)_{n\in\omega}$ is an infinite trace. By Lemma 2.1.20, this trace then is either a μ-trace or a ν-trace. We say that the infinite match $(\varphi_n, s_n)_{n\in\omega}$ is won by $\exists$ if and only if its left projection is a ν-trace.

The following theorem links the game semantics to the algebraic semantics. Its proof is out of the scope of this thesis, but can be found for instance in [106].

2.1.30. THEOREM. *For any model $\mathbb{S}$ and formula ξ, which is tidy and in negation normal form, it holds that:*

The player $\exists$ has a winning strategy at (ξ, s) in $\mathcal{E}(\xi, \mathbb{S})$ if and only if $s \in [\![\xi]\!]^{\mathbb{S}}$.

The main advantage of working with the game semantics, is that the evaluation game is a so-called *parity game*. Parity games have a well-developed theory and satisfy certain nice properties, which will be given in the next section. We close this section by discussing some fragments of the modal μ-calculus that play a role in the following chapters of this thesis.

2.1.3 Fragments and extensions

We consider, in decreasing order of expressiveness, one extension and three fragments of $\mu\mathtt{ML}$ that will feature in this thesis.

The two-way modal μ-calculus

The syntax $\mu_2\mathtt{ML}(\mathsf{D})$ of the *two-way modal μ-calculus* is precisely the same as that of $\mu\mathtt{ML}(\mathsf{D})$, with the additional assumption that there is an involution operator $\breve{}$ on D. That is, for every $a \in \mathsf{D}$ it holds that $\breve{a} \neq a$ and $\breve{\breve{a}} = a$.

The idea is that the modality $\breve{a}$ is the *converse* of the modality a. From a temporal perspective, this enables $\mu_2\mathtt{ML}$-formulas to express statements about the past. Formulas of $\mu_2\mathtt{ML}(\mathsf{D})$ are interpreted only over models which satisfy the following regularity property:

$$R_{\breve{a}} = \{(t,s) \mid (s,t) \in R_a\} \text{ for every } a \in \mathsf{D}.$$

In other words, the relation $R_{\breve{a}}$ must truly be the converse of the relation R_a.

Remarkably, $\mu_2\mathtt{ML}$ does not have the finite model property over the class of all models. This is exemplified by the formula

$$\nu x(\langle a\rangle x \wedge \mu y\langle \breve{a}\rangle y),$$

which expresses that there is an infinite forward path along which there is no infinite backward path.

The alternation-free modal μ-calculus

The alternation-free fragment of the modal μ-calculus is obtained by syntactically restricting the fixed point operators in such a way, that μ-variables do not depend on ν-variables and vice versa.

2.1.31. Definition. A formula ξ of $\mu\mathtt{ML}(\mathsf{D})$ is said to be *alternation free* if for every subformula $\eta x\varphi$ of ξ it holds that no free occurrence of x in φ is in the scope of an $\overline{\eta}$-operator in φ.

2.1.32. Example. The formula $\mu x\nu y(x \wedge \langle a\rangle y)$ is *not* alternation free, but the following formulas are:

$$\mu x\nu y(p \wedge \langle a\rangle y) \qquad \mu x(x \wedge \nu y\langle a\rangle y) \qquad \mu x\mu y(x \wedge \langle a\rangle y).$$

It will be useful to have an inductive definition of the alternation-free fragment of the $\mu\mathtt{ML}$. For this purpose, we first define the following class of formulas.

2.1.33. Definition. Let $\mathtt{F}$ be a set of formulas. The set $\mathsf{Noeth}_{\mathsf{X}}(\mathsf{D},\mathtt{F})$ of formulas *noetherian in* $\mathsf{X} \subseteq \mathsf{P}$ *built from* $\mathtt{F}$ contains all formulas generated by the following grammar:

$$\varphi ::= x \mid \alpha \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \langle a\rangle\varphi \mid [a]\varphi \mid \mu y\varphi',$$

where $x \in \mathsf{X}$, $a \in \mathsf{D}$, $y \in \mathsf{P}$, $\alpha \in \mathtt{F}$ is X-free, and $\varphi' \in \mathsf{Noeth}_{\mathsf{X}\cup\{y\}}(\mathsf{D},\mathtt{F})$.

Note that $\mathsf{Noeth}_{\mathsf{X}}(\mathsf{D},\mu\mathtt{ML}(\mathsf{D}))$ is a fragment of $\mu\mathtt{ML}(\mathsf{D})$. For the origin of the name 'noetherian' for this fragment, we refer the reader to [42]. Note that if φ is noetherian in X, then the variables in X do not occur within the scope of a ν-operator in φ. Dually, we define the following fragment.

2.1.34. DEFINITION. Let $\mathtt{F}$ be a set of formulas. The set $\mathsf{Conoeth}_{\mathsf{X}}(\mathsf{D},\mathtt{F})$ of formulas *conoetherian in* $\mathsf{X} \subseteq \mathsf{P}$ *built from* $\mathtt{F}$ contains all formulas generated by the following grammar:

$$\varphi ::= x \mid \alpha \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \langle a \rangle \varphi \mid [a]\varphi \mid \nu y \varphi',$$

where $x \in \mathsf{X}$, $y \in \mathsf{P}$, $a \in \mathsf{D}$, $\alpha \in \mathtt{F}$ is X-free, and $\varphi' \in \mathsf{Conoeth}_{\mathsf{X}\cup\{y\}}(\mathsf{D},\mathtt{F})$.

We are finally ready to given an inductive definition of the alternation-free modal μ-calculus.

2.1.35. DEFINITION. The syntax $\mu^{af}\mathtt{ML}(\mathsf{D})$ of the *alternation-free modal μ-calculus* is given by:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \langle a \rangle \varphi \mid [a]\varphi \mid \mu x.\varphi' \mid \nu x.\varphi'',$$

where $p, x \in \mathsf{P}$, $a \in \mathsf{D}$, and φ' belongs to $\mathsf{Noeth}_{\{x\}}(\mathsf{D}, \mu^{af}\mathtt{ML}(\mathsf{D}))$, and φ'' belongs to $\mathsf{Conoeth}_{\{x\}}(\mathsf{D}, \mu^{af}\mathtt{ML}(\mathsf{D}))$.

The (routine) proof of the following lemma is left to the reader.

2.1.36. LEMMA. *If φ is a $\mu^{af}\mathtt{ML}$-formula, then so are $\mathsf{nnf}(\varphi)$, $\overline{\varphi}$, and every formula in the closure of φ.*

A very important effect of restricting to the alternation-free fragment, is that the winning condition of the evaluation game becomes much simpler.

2.1.37. LEMMA. *Let $(\varphi_n)_{n\in\omega}$ be an infinite trace of alternation-free formulas. Then either infinitely many of the φ_i are μ-formulas, or infinitely many are ν-formulas, but not both.*

Before we go on to prove the above lemma, we first prove an auxiliary lemma.

2.1.38. LEMMA. *Let $\varphi_0, \ldots, \varphi_n$ be a finite trace of $\mu^{af}\mathtt{ML}$-formulas, such that only φ_n is an η-formula. Then $\varphi_n \trianglelefteq \varphi_i$ for every $i \in [0, n]$.*

Proof:
We proceed by induction on n. The case where $n = 0$ is trivial, so suppose that $n > 0$. By the induction hypothesis $\varphi_n \trianglelefteq \varphi_i$ for every $i \in [1, n]$. We make a case distinction on the shape of φ_0. If φ_0 is modal or propositional, it is clear that $\varphi_n \trianglelefteq \varphi_1 \trianglelefteq \varphi_0$. If φ_0 is a fixed point formula, then by assumption φ_0 is of the form $\overline{\eta} y \psi$.

Hence $\varphi_1 = \psi[\overline{\eta} y \psi / y]$. From the fact that $\varphi_n \trianglelefteq \varphi_1$, we obtain, by Lemma 2.1.22, that either $\varphi_n \trianglelefteq \varphi_0$, or $\varphi_0 \trianglelefteq \varphi_n$. In the former case we are done. We will now argue that the latter case is impossible. Recall that φ_n is an η-formula, say $\eta x \varphi$. Since φ is (co)noetherian, we know that x does not occur in φ in the scope of an $\overline{\eta}$-operator. Hence, if $\varphi_0 \trianglelefteq \varphi_n$, then $\varphi_0 \trianglelefteq \varphi$ and thus φ_0 is x-free. But then

φ_1 is x-free, contradicting the fact that $\varphi_n \trianglelefteq \varphi_1$. □

Proof of Lemma 2.1.37:
By Lemma 2.1.20 we know that the trace $(\varphi_n)_{n\in\omega}$ contains infinitely many fixed point formulas. Suppose, towards a contradiction, that it contains infinitely many μ-formulas as well as infinitely many ν-formulas. Then there are $k_0 > k_1 > k_2 \ldots$ such that each φ_{k_i} is a fixed point formula, and, in particular, if φ_{k_i} is an η-formula, then $\varphi_{k_{i+1}}$ is the first $\overline{\eta}$-formula after φ_{k_i}. Applying the previous lemma to each finite trace $\varphi_{k_i}, \ldots \varphi_{k_{i+1}}$, we obtain that $\varphi_{k_{i+1}} \trianglelefteq \varphi_{k_i}$. But if φ_{k_i} is a η-formula, then $\varphi_{k_{i+1}}$ is an $\overline{\eta}$-formula. Therefore $\varphi_{k_{i+1}} \neq \varphi_{k_i}$ and thus $\varphi_{k_{i+1}}$ is a *proper* subformula of φ_{k_i}. This is gives a descending chain of proper subformulas $\varphi_{k_0} \triangleright \varphi_{k_1} \triangleright \varphi_{k_2} \triangleright \cdots$, a contradiction. □

The *two-way alternation-free modal μ-calculus* $\mu_2^{af}\mathtt{ML}(\mathsf{D})$ is defined from $\mu^{af}\mathtt{ML}(\mathsf{D})$ in the same way as $\mu_2\mathtt{ML}(\mathsf{D})$ is defined from $\mu\mathtt{ML}(\mathsf{D})$. That is, by assuming that there is an involution operator $\breve{}$ on D. Clearly Lemma 2.1.37 also applies to $\mu_2^{af}\mathtt{ML}$.

The continuous modal μ-calculus

The continuous modal μ-calculus is another fragment of the modal μ-calculus, and, in particular, a fragment of the alternation-free modal μ-calculus. We again first define two dual classes parametrised in a set X of variables.

2.1.39. Definition. Let $\mathtt{F}$ be a set of formulas. The set $\mathsf{Con}_{\mathsf{X}}(\mathsf{D}, \mathtt{F})$ of formulas *continuous in* $\mathsf{X} \subseteq \mathsf{P}$ *built from* $\mathtt{F}$ contains all formulas generated by the following grammar:

$$\varphi ::= x \mid \alpha \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \langle a \rangle \varphi \mid \mu y \varphi',$$

where $x \in \mathsf{X}$, $a \in \mathsf{D}$, $y \in \mathsf{P}$, $\alpha \in \mu\mathtt{ML}(\mathsf{D})$ is X-free, and $\varphi' \in \mathsf{Con}_{\mathsf{X}}(\mathsf{D}, \mathtt{F})$

2.1.40. Definition. Let $\mathtt{F}$ be a set of formulas. The set $\mathsf{Cocon}_{\mathsf{X}}(\mathsf{D}, \mathtt{F})$ of formulas *cocontinuous in* $\mathsf{X} \subseteq \mathsf{P}$ *built from* $\mathtt{F}$ contains all formulas generated by the following grammar:

$$\varphi ::= x \mid \alpha \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid [a] \varphi \mid \nu y \varphi',$$

where $x \in \mathsf{X}$, $a \in \mathsf{D}$, $y \in \mathsf{P}$, $\alpha \in \mu\mathtt{ML}(\mathsf{D})$ is X-free, and $\varphi' \in \mathsf{Cocon}_{\mathsf{X}}(\mathsf{D}, \mathtt{F})$

2.1.41. Definition. The syntax $\mu^c\mathtt{ML}(\mathsf{D})$ of the *continuous modal μ-calculus* is given by:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \langle a \rangle \varphi \mid [a]\varphi \mid \mu x.\varphi' \mid \nu x.\varphi'',$$

where $p, x \in \mathsf{P}$, $a \in \mathsf{D}$, and φ' belongs to $\mathsf{Con}_{\{x\}}(\mathsf{D}, \mu^c\mathtt{ML}(\mathsf{D}))$, and φ'' belongs to $\mathsf{Cocon}_{\{x\}}(\mathsf{D}, \mu^c\mathtt{ML}(\mathsf{D}))$.

We write μ^c`ML` for the continuous monomodal μ-calculus. The proof of the following lemma is left to the reader.

2.1.42. LEMMA. *If φ is a μ^c`ML`-formula, then so is $\overline{\varphi}$ and every formula in the closure of φ.*

Since μ^c`ML` is a fragment of μ^{af}`ML`, Lemma 2.1.37 applies to it as well. In addition, the language μ`ML` satisfies the following even stronger property.

2.1.43. LEMMA. *Let $(\varphi_n)_{n\in\omega}$ be a trace of μ^c`ML`-formulas. Then:*

(i) if $(\varphi_n)_{n\in\omega}$ contains infinitely many μ-formulas, it contains at most finitely many $\Box$-formulas.

(ii) if $(\varphi_n)_{n\in\omega}$ contains infinitely many ν-formulas, it contains at most finitely many $\Diamond$-formulas.

Proof:
We only prove item (i), because item (ii) is dual. It suffices to prove the following claim, which is analogous to Lemma 2.1.38.

> Let $\varphi_0, \ldots, \varphi_n$ be a finite trace of μ^c`ML`-formulas, such that φ_n is a $\Box$-formula, and every other φ_i is neither a $\Box$-formula nor a ν-formula. Then $\varphi_n \trianglelefteq \varphi_i$ for each $i \in [0, n]$.

Indeed, once we have proven the claim above, the proof can be finished by using an argument analogous to the proof of Lemma 2.1.37. Like Lemma 2.1.38, we prove the above claim by induction on the length of the trace. The base case is again trivial, and the only interesting inductive step is the one where φ_0 is of the form $\mu x\psi$. Then φ_1 is of the form $\psi[\mu x\psi/x]$ and, by the induction hypothesis, we have $\varphi_n \trianglelefteq \psi[\mu x\psi/x]$. Hence, we obtain by Lemma 2.1.22 that $\varphi_n \trianglelefteq \mu x\psi$, in which case we are done, or $\mu x\psi \trianglelefteq \varphi_n$. Suppose the latter is the case. Then $\mu x\psi$ occurs within the scope of a $\Box$-operator in φ_n. Hence, since $\varphi_n \trianglelefteq \psi[\mu x\psi/x]$, $\mu x\psi$ occurs within the scope of an $\Box$-operator in $\psi[\mu x\psi/x]$. But this means that x occurs in the scope of an $\Box$-operator in ψ, contradicting the continuity of ψ in x. $\Box$

Modal logic with the master modality

We finish with the most simple fragment of this subsection.

2.1.44. DEFINITION. For D a finite set of actions, the syntax $\mathtt{ML}^*(\mathsf{D})$ of *modal logic with the master modality* is generated by:

$$\varphi ::= p \mid \bot \mid \varphi \to \varphi \mid [a]\varphi \mid [*]\varphi,$$

where $p \in \mathsf{P}$ and $a \in \mathsf{D}$.

The language $\mathtt{ML}^*$ is a fragment of $\mu\mathtt{ML}$ through the embedding $-^t$, inductively defined as follows, where $\mathsf{D} = \{a_1, \dots, a_n\}$.

$$\begin{aligned} p^t &:= p & \bot^t &:= \mu x(x) \\ (\varphi \to \psi)^t &:= \neg\varphi^t \vee \psi^t & ([a]\varphi)^t &:= [a]\varphi^t \\ ([*]\varphi)^t &:= \nu x([a_1]x \wedge \dots \wedge [a_n]x \wedge \varphi^t) \end{aligned}$$

The semantics of $\mathtt{ML}^*$ is inherited from $\mu\mathtt{ML}$ through this embedding. In particular, we have

$$\begin{aligned} \mathbb{S}, s \Vdash [*]\varphi &\Leftrightarrow \mathbb{S}, s \Vdash \nu x([a_1]x \wedge \dots \wedge [a_n]x \wedge \varphi^t) \\ &\Leftrightarrow \mathbb{S}, t \Vdash \varphi^t \text{ for all } sR^*t, \end{aligned}$$

where R^* is the reflexive-transitive closure of the union of the relations $R_{a_1}, \dots, R_{a_n}$.

2.2 Parity games

The goal of this section is threefold. First, we will put the game-theoretic notions used in the previous section, as well as in the rest of this thesis, on a formal footing. Second, we will define the notion of a parity game and state some of its most important properties. Third, we will use the fact that the evaluation game is a parity game to give a first example of how these properties of parity games can be used to prove facts about the modal μ-calculus.

2.2.1. Definition. A *(two-player) game* is a structure $\mathcal{G} = (B_0, B_1, E, W)$ where E is a binary relation on $B := B_0 \uplus B_1$, and W is a map $B^\omega \to \{0,1\}$.

The set B is called the *board* of $\mathcal{G}$, and its elements are called *positions*. Whether a position belongs to B_0 or B_1 determines which player *owns* that position. If a player $\Pi \in \{0,1\}$ owns a position q, it is their turn to play and the set of their *admissible moves* is given by the image $E[q]$ of q under E. .

2.2.2. Definition. A *match* in $\mathcal{G} = (B_0, B_1, E, W)$ (or simply a $\mathcal{G}$*-match*) is a path $\mathcal{M}$ through the graph (B, E). A match is said to be *full* if it is a maximal path.

Note that a full match $\mathcal{M}$ is either finite, in which case $E[\mathsf{last}(\mathcal{M})] = \emptyset$, or infinite. For a player $\Pi \in \{0,1\}$, we write $\overline{\Pi}$ for the other player, *i.e.* $\Pi = \Pi + 1 \mod 2$.

2.2.3. Definition. A full match $\mathcal{M}$ in $\mathcal{G} = (B_0, B_1, E, W)$ is *won* by player Π if either $\mathcal{M}$ is finite and $\mathsf{last}(\mathcal{M}) \in B_{\overline{\Pi}}$, or $\mathcal{M}$ is infinite and $W(\mathcal{M}) = \Pi$.

If a full match $\mathcal{M}$ is finite, and $\mathsf{last}(\mathcal{M})$ belongs to B_Π for $\Pi \in \{0,1\}$, we say that the player Π got *stuck*. A *partial match* is a match which is not full.

2.2.4. DEFINITION. In the context of a game $\mathcal{G}$, we denote by PM_Π the set of partial $\mathcal{G}$-matches $\mathcal{M}$ such that $\mathsf{last}(\mathcal{M})$ belongs to the player Π.

2.2.5. DEFINITION. A *strategy* for Π in a game $\mathcal{G}$ is a map $f : \mathrm{PM}_\Pi \to B$. Moreover, a $\mathcal{G}$-match $\mathcal{M}$ is said to be f*-guided* if for any $\mathcal{M}_0 \sqsubset \mathcal{M}$ with $\mathcal{M}_0 \in \mathrm{PM}_\Pi$ it holds that $\mathcal{M}_0 \cdot f(\mathcal{M}_0) \sqsubseteq \mathcal{M}$.

For a position q, the set $\mathrm{PM}_\Pi(q)$ contains all $\mathcal{M} \in \mathrm{PM}_\Pi$ such that $\mathsf{first}(\mathcal{M}) = q$.

2.2.6. DEFINITION. A strategy f for Π in $\mathcal{G}$ is *surviving* at a position q if $f(\mathcal{M})$ is admissible for every $\mathcal{M} \in \mathrm{PM}_\Pi(q)$, and *winning* at q if in addition all full f-guided matches starting at q are won by Π. A position q is said to be *winning* for Π if Π has a strategy winning at q. We denote the set of all positions in $\mathcal{G}$ that are winning for Π by $\mathrm{Win}_\Pi(\mathcal{G})$.

We write $\mathcal{G}@q$ for the game $\mathcal{G}$ *initialised* at the position q of $\mathcal{G}$. A strategy f for Π is *surviving* (*winning*) in $\mathcal{G}@q$ if it is surviving (winning) in $\mathcal{G}$ at q.

2.2.7. DEFINITION. A strategy f is *positional* if it only depends on the last move, *i.e.* if $f(\mathcal{M}) = f(\mathcal{M}')$ for all $\mathcal{M}, \mathcal{M}' \in \mathrm{PM}_\Pi$ with $\mathsf{last}(\mathcal{M}) = \mathsf{last}(\mathcal{M}')$.

We will often present a positional strategy for Π as a map $f : B_\Pi \to B$.

2.2.8. DEFINITION. A *priority map* on some board B is a map $\Omega : B \to \omega$ of finite range. A *parity game* is a game of which the winning condition is given by $W_\Omega(\mathcal{M}) = \max \mathit{Inf}(\Omega[\mathcal{M}]) \mod 2$, where $\mathit{Inf}(\Omega[\mathcal{M}])$ is the set of priorities occurring infinitely often in $\mathcal{M}$.

As mentioned before, the evaluation game $\mathcal{E}(\xi, \mathbb{S})$ is, in fact, a parity game. The priority map Ω can for instance be defined as follows. First note that $\trianglelefteq$ is a partial order on $\mathsf{FL}(\xi)$. Hence, it can be extended into a linear order $\preceq$. Let $\varphi_1, \ldots, \varphi_n$ be an enumeration of all fixed point formulas in $\mathsf{FL}(\xi)$ in the order of $\succeq$, *i.e.* such that $\varphi_1 \succeq \cdots \succeq \varphi_n$. We set:

$$\Omega(\varphi, s) := \begin{cases} 2i - 1 & \text{if } \varphi = \varphi_i \text{ for some } 1 \leq i \leq n \text{ and } \varphi_i \text{ is a } \mu\text{-formula,} \\ 2i & \text{if } \varphi = \varphi_i \text{ for some } 1 \leq i \leq n \text{ and } \varphi_i \text{ is a } \nu\text{-formula,} \\ 0 & \text{otherwise.} \end{cases}$$

We leave it to the reader to verify that, by Lemma 2.1.20, we have $W_\Omega(\mathcal{M}) = 0$ if and only if the infinite $\mathcal{E}(\xi, \mathbb{S})$-match $\mathcal{M}$ is won by $\exists$.

2.2.9. REMARK. The number of priorities, and therefore the complexity of the parity game, is sufficient for our purposes, but not optimal. For more details we refer the reader to [64].

The following theorem captures the key property of parity games: they are *positionally determined*. This means that at every position either of the two players has a winning strategy, which moreover is positional. In fact, each player Π has a positional strategy f_Π that is *optimal*, in the sense that f_Π is winning for Π in $\mathcal{G}@q$ for *every* $q \in \mathrm{Win}_\Pi(\mathcal{G})$.

2.2.10. Theorem ([74, 38]). *For any parity game $\mathcal{G}$, there are positional strategies f_Π for each player $\Pi \in \{0,1\}$, such that for every position q one of the f_Π is a winning strategy for Π in $\mathcal{G}@q$.*

The following proposition is an example application of positional determinacy. Recall that a model (S, R, V) is *image-finite* if the projection $R[s] = \{t \in S \mid sRt\}$ is finite for every $s \in S$.

2.2.11. Proposition. *Every satisfiable formula has an image-finite model.*

Proof:
Let $\mathbb{S} = (S, R, V)$ be a model and suppose that $s \in [\![\xi]\!]^{\mathbb{S}}$. Without loss of generality we assume that ξ is both tidy and in negation normal form. Hence, by Theorem 2.1.30, we know that $\exists$ has a winning strategy in the parity game $\mathcal{E}(\xi, \mathbb{S})@(\xi, s)$, and by Theorem 2.2.10 we may assume that this strategy is positional. Consider this strategy as a partial function $f : \mathsf{FL}(\xi) \times S \rightharpoonup \mathsf{FL}(\xi) \times S$. Since $\mathsf{FL}(\xi)$ is finite, there are for any state u of $\mathbb{S}$ at most finitely many states v such that $f(\Diamond\varphi, u) = (\varphi, v)$ for some $\Diamond\varphi \in \mathsf{FL}(\xi)$.

Let $\mathbb{S}' = (S, R', V)$ be the model obtained from $\mathbb{S}$ by defining

$$R'[u] := \{v \mid f(\Diamond\varphi, u) = (\varphi, v) \text{ for some } \Diamond\varphi \in \mathsf{FL}(\xi)\}.$$

By the reasoning above $\mathbb{S}'$ is image-finite. Moreover, the strategy f restricts to a winning strategy in $\mathcal{E}(\xi, \mathbb{S}')@(\xi, s)$, as required. □

2.3 Proof systems

In this section we introduce some proof systems for the modal μ-calculus. First, we will discuss a Hilbert-style proof system originally due to Dexter Kozen. Next, we will show how this Hilbert-style system can be translated into a finitary Gentzen-style counterpart. Finally, we describe how this Gentzen-style system can be adapted into a non-well-founded proof system, which is the kind of system that we will mostly work with in this thesis. We will prove the soundness and completeness of the non-well-founded system and discuss several related topics, such as cyclic proofs and the bounded proof property. We close this section with a short discussion on the goal of this thesis: adapting the proof systems in this section to accommodate various frame conditions.

2.3.1 Hilbert-style proof systems

In [59], Kozen presents a natural Hilbert-style axiomatisation for the modal μ-calculus.

2.3.1. Definition. The Hilbert-style proof system $\mu\mathsf{K}$ consists of a proof system for the least normal modal logic K, together with the following axioms and rules:

$$\varphi[\mu x\varphi/x] \to \mu x\varphi \qquad \frac{\varphi[\psi/x] \to \psi}{\mu x\varphi \to \psi} \qquad \nu x\varphi \to \varphi[\nu x\varphi/x] \qquad \frac{\psi \to \varphi[\psi/x]}{\psi \to \nu x\psi}$$

Like the algebraic semantics, these rules characterise $\mu x\varphi$ as the least prefixed point of $\varphi(x)$, and likewise for $\nu x\varphi$ as the greatest postfixed point.

Proving completeness for $\mu\mathsf{K}$ is notoriously difficult. The standard canonical model construction used in basic modal logic will not work, because compactness fails, as shown in the following proposition. The abbreviation $\Box^n$ is recursively defined by setting $\Box^0\varphi := \varphi$ and $\Box^{n+1}\varphi := \Box\Box^n\varphi$.

2.3.2. Proposition. *The set* $\{\mu x(\Diamond x \vee p)\} \cup \{\Box^n\overline{p} : n \in \omega\}$ *is finitely satisfiable, but not satisfiable.*

Proof:
Immediate from the fact that $\mu x(\Diamond x \vee p)$ expresses the reachability of a state where p is true. □

Hence, there will be (maximal) consistent sets which are not satisfiable. One way to remedy this is to consider *finitary* canonical models, which are closely related to the method of *filtration*. In Chapter 4, we will see that this method works for the continuous modal μ-calculus, but not for the full language. Another option is to consider non-well-founded proof systems, which are the topic of the next section.

2.3.2 Non-well-founded proof systems

To introduce non-well-founded proof systems, let us first consider a well-founded sequent-style reformulation of the Hilbert-style system $\mu\mathsf{K}$. Throughout this section we will assume that formulas are tidy and in negation normal form.

2.3.3. Definition. A *sequent* is a finite set of formulas.

Sequents should be read *disjunctively*. That is, a sequent $\Gamma = \{\gamma_1, \dots, \gamma_n\}$ represents the disjunction $\gamma_1 \vee \dots \vee \gamma_n$.

2.3.4. Remark. In this chapter we use so-called *one-sided* sequents. This is convenient when formulas are in negation normal form. Alternatively, one often encounters *two-sided* sequents, *i.e.* pairs of one-sided sequents. Two-sided sequents are better equipped to deal with negation and will be used both in Chapter 3 and in Chapter 6

$$\mathsf{id}\,\frac{}{p,\overline{p}} \qquad \vee\,\frac{\Gamma,\varphi,\psi}{\Gamma,\varphi\vee\psi} \qquad \wedge\,\frac{\Gamma,\varphi \qquad \Gamma,\psi}{\Gamma,\varphi\wedge\psi} \qquad \mu\,\frac{\Gamma,\varphi[\mu x\varphi/x]}{\Gamma,\mu x\varphi} \qquad \nu\,\frac{\Gamma,\varphi[\nu x\varphi/x]}{\Gamma,\nu x\varphi}$$

$$\mathsf{K}\,\frac{\Delta,\varphi}{\Gamma,\Diamond\Delta,\Box\varphi} \qquad \mathsf{ind}_\nu\,\frac{\Gamma,\varphi[\overline{\Gamma}/x]}{\Gamma,\nu x\varphi} \qquad \mathsf{cut}\,\frac{\Gamma,\varphi \qquad \Gamma,\overline{\varphi}}{\Gamma}$$

Figure 2.1: A sequent-style reformulation of $\mu\mathsf{K}$

The sequent-style reformulation of $\mu\mathsf{K}$ is given in Figure 2.1. The propositional rules id, $\vee$, and $\wedge$ take care of the propositional reasoning in $\mu\mathsf{K}$, assisted by cut, which simulates modus ponens. The modal rule K covers modal reasoning. To see that the rule K is sound, suppose that the conclusion is falsified by some state s. Then s must have a successor falsifying the premiss. Note that K has built-in weakening. The rules μ and ν correspond to their respective axioms in $\mu\mathsf{K}$, and it turns out we only need to have an additional *induction rule* ind_ν for the ν-operator. In this rule $\overline{\Gamma}$ stands for the negation of Γ, *i.e.* the conjunction of the negations of all formulas in Γ. Reading the premiss of ind_ν as $\overline{\Gamma} \to \varphi[\overline{\Gamma}/x]$, and similar for the conclusion, it is clear that ind_ν is a special case of the greatest postfixed point rule. We will later see why we do not need an analogous rule ind_μ corresponding to the least prefixed point rule.

Proof-theoretically, the sequent-style reformulation of $\mu\mathsf{K}$ lacks several desirable properties. First, it has a cut rule, which makes proof search infeasible, because it requires one to guess the cut formula. Second, both cut and ind_ν violate the *closure property*: their premisses may contain formulas outside of the Fischer-Ladner closure of their conclusions.

It is for the reasons above, combined with the difficulty of proving the completeness of $\mu\mathsf{K}$, that different kinds of proof systems have been developed. The following system, which simply drops the two problematic rules of the sequent-style reformulation of $\mu\mathsf{K}$, will play a central role in this section.

2.3.5. Definition. The system NW consists of id, $\vee$, $\wedge$, μ, ν, K from Figure 2.1.

Note that NW does have the closure property.

In the premiss or conclusion of some rule application, the formulas in Γ are called *inactive*, while the formulas outside of Γ are called *active*. Note that, since sequents are sets, a formula can simultaneously be both inactive and active.

The expressiveness lost by dropping cut and ind_ν, is made up for by allowing non-well-founded branches.

2.3.6. Definition. A NW-*derivation* is a (possibly infinite) tree generated by the rules of NW.

As usual, we say that a derivation is *closed* if every leaf is an axiom, which in the case of NW means that every leaf is an application of id. The root sequent of an NW-derivation is called its *conclusion*.

Unfortunately closed NW-derivations are not sound, in the sense that their conclusions need not be valid. Take for instance the formula $\mu x \Box x$. This formula expresses that there is no infinite path, which is certainly not true in every state of every model. However, it does have the following non-well-founded derivation:

$$\dfrac{\dfrac{\vdots}{\mu x \Box x}}{\dfrac{\Box \mu x \Box x}{\mu x \Box x}}$$

We will impose a certain sufficient condition on NW-derivations to ensure that their conclusions are valid. NW-derivations satisfying this condition will be called NW-*proofs*. In order to define this condition, we first require the proof-theoretical notion of *direct ancestry*.

It will convenient to have a slightly more formal definition of the notion of a rule instance.

2.3.7. Definition. A *rule instance* is a triple $i = (\Gamma, \mathsf{r}, \langle \Gamma_1, \ldots, \Gamma_n \rangle)$ such that

$$\mathsf{r}\,\dfrac{\Gamma_1 \cdots \Gamma_n}{\Gamma}$$

is a valid rule application in NW.

2.3.8. Remark. The above definition of a rule instance naturally generalises to different proof systems. We will therefore in the following chapters use this notion without redefining it for the proof system at hand. Moreover, we will sometimes use the *rule application* with rule instance.

2.3.9. Definition. Let $(\Gamma, \mathsf{r}, \langle \Gamma_1, \ldots, \Gamma_n \rangle)$ be a rule instance of NW. A formula ψ in Γ_i is said to be a *direct ancestor* of a formula φ in Γ if ψ 'comes from' φ. That is, if one of the following holds:

(i) φ and ψ are both inactive and $\varphi = \psi$;

(ii) φ and ψ are both active, and

  (a) $\mathsf{r} \neq \mathsf{K}$, or

  (b) $\mathsf{r} = \mathsf{K}$ and $\psi = \triangle\varphi$ for some $\triangle \in \{\Diamond, \Box\}$.

2.3.10. Remark. The above notion of direct ancestry is standard in proof theory (see *e.g.* [22, Definition 1.2.3.]). Note that direct ancestry does not necessarily determine a tree structure. Indeed, a single formula in the conclusion might have multiple direct ancestors in the same premiss, and a formula in some premiss might be a direct ancestor of multiple formulas in the conclusion.

2.3.11. Definition. Suppose some NW-derivation contains a (possibly infinite) path ρ:

$$\Gamma_0 \cdot \mathsf{r}_0 \cdot \Gamma_1 \cdot \mathsf{r}_1 \cdots (\Gamma_n)$$

A *trail* on ρ is a sequence of formulas (φ_i) such that for each φ_i belongs to Γ_i, and each for each two subsequent φ_i, φ_{i+1} it holds that φ_{i+1} is a direct ancestor of wφ_i.

Moreover, the *tightening* of a trail is the subsequence consisting of precisely those φ_i such that φ_i is active in the conclusion Γ_i, and φ_{i+1} is active in the premiss Γ_{i+1} of the application of r_i on the path ρ. In other words, such that φ_{i+1} is a direct ancestor of φ_i by virtue of item (ii) of Definition 2.3.9.

We leave it to the reader to verify that the tightening of a trail is always a trace. A trail whose tightening is an η-trace is called an η-*trail*. Note that being infinite is a necessary, but not a sufficient condition for some trail to be an η-trail. We are now ready to state the soundness condition for NW-derivations.

2.3.12. Definition. A closed NW-derivation is said to be an NW-*proof* if every infinite branch contains a ν-trail.

2.3.3 The proof search game

We will define a proof search game $\mathcal{G}(\Gamma)$ for the proof system NW. We write $\mathsf{conc}(i)$ for the conclusion, *i.e.* the first component of the rule instance i. Moreover, we use Seq_Γ and Inst_Γ respectively for the set of sequents and the set of valid rule instances, with the property that every formula belongs to $\mathsf{FL}(\Gamma)$.

The set of positions of the game $\mathcal{G}(\Gamma)$ is $\mathsf{Seq}_\Gamma \cup \mathsf{Inst}_\Gamma$. Since $\mathsf{FL}(\Gamma)$ is finite, the game $\mathcal{G}(\Gamma)$ has only finitely many positions. The ownership function and admissible moves of $\mathcal{G}(\Gamma)$ are as in the following table:

Position	Owner	Admissible moves
$\Delta \in \mathsf{Seq}_\Gamma$	Prover	$\{i \in \mathsf{Inst}_\Gamma \mid \mathsf{conc}(i) = \Delta\}$
$(\Delta, \mathsf{r}, \langle \Delta_1, \ldots, \Delta_n \rangle) \in \mathsf{Inst}_\Gamma$	Refuter	$\{\Delta_i \mid 1 \leq i \leq n\}$

The only thing we need to specify about $\mathcal{G}(\Gamma)$ is which infinite matches are won by whom. We say that an infinite match is won by Prover if and only if it contains a ν-trail. It is not hard to see that the strategy tree of a winning strategy for Prover is exactly the same as an NW-proof.

2.3.4 Soundness of NW

In this section we show that the non-well-founded proof system NW is sound. Our proof goes by showing that a winning strategy for $\forall$ in the evaluation game $\mathcal{E}(\bigvee\Gamma, \mathbb{S})@(\bigvee\Gamma, \mathbb{S})$ can be turned into a winning strategy for Refuter in $\mathcal{G}(\Gamma)@\Gamma$. It follows that, if Γ is invalid, it is not provable.

For $\mathcal{M}$ a finite match, we write $\mathcal{M}_<$ for the initial segment of $\mathcal{M}$ omitting only the last position $\mathsf{last}(\mathcal{M})$ of $\mathcal{M}$.

2.3.13. LEMMA. *Let f be positional and winning for $\forall$ in $\mathcal{E}(\bigvee\Gamma, \mathbb{S})@(\bigvee\Gamma, s)$. Then there is a strategy $\overline{f}$ for Refuter in $\mathcal{G}(\Gamma)@\Gamma$, and a function $s : \mathrm{PM}_P(\Gamma) \to \mathbb{S}$, such that the following hold for every finite $\overline{f}$-guided match $\mathcal{M}$:*

(i) If $\mathsf{last}(\mathcal{M})$ is an application of $\wedge$ with $\psi_1 \wedge \psi_2$ active in the conclusion, then $\overline{f}(\mathcal{M})$ corresponds to the conjunct selected by f at $(\psi_1 \wedge \psi_2, s(\mathcal{M}_<))$.

(ii) If $\mathsf{last}(\mathcal{M})$ is an application of K with $\Box\psi$ as active formula in the conclusion, then $s(\overline{f}(\mathcal{M}))$ is the state selected by f at $(\Box\psi, s(\mathcal{M}_<))$.

(iii) If $\mathsf{last}(\mathcal{M})$ is a sequent, then f is winning for $\forall$ in $\mathcal{E}(\bigvee\Gamma, \mathbb{S})@(\varphi, s(\mathcal{M}))$ for every $\varphi \in \mathsf{last}(\mathcal{M})$.

Proof:
For the match $\Gamma \in \mathrm{PM}_P$ consisting of only the initial position, we set $s_\Gamma := s$. Note that the required condition on s_Γ is met by assumption. By induction on the length $|\mathcal{M}|$, we will for every $\overline{f}$-guided match $\mathcal{M} \in \mathrm{PM}_R(\Gamma)$ simultaneously define the move $\overline{f}(\mathcal{M})$ and the state $s(\overline{f}(\mathcal{M}))$.

Let $\mathcal{M} \in \mathrm{PM}_R$ be $\overline{f}$-guided and let

$$\mathsf{r}\,\frac{\Delta_1 \cdots \Delta_n}{\mathsf{last}(\mathcal{M}_<)}$$

be the rule instance $\mathsf{last}(\mathcal{M})$. By the induction hypothesis (or by the above in case $|\mathcal{M}| = 2$), we know that f is winning for $\forall$ in $\mathcal{E}(\bigvee\Gamma, \mathbb{S})@(\varphi, s(\mathcal{M}_<))$ for every $\varphi \in \mathsf{last}(\mathcal{M}_<)$. We make a case distinction based on the rule r.

(i) $\mathsf{r} = \mathsf{id}$. This cannot happen, because f cannot be winning both at $(p, s(\mathcal{M}_<))$ and at $(\overline{p}, s(\mathcal{M}_<))$.

(ii) $\mathsf{r} \in \{\vee, \mu, \nu\}$. In these cases there is only a single choice for $\overline{f}(\mathcal{M})$, and it is not hard to see that it suffices to set $s(\overline{f}(\mathcal{M})) := s(\mathcal{M}_<)$.

(iii) $\mathsf{r} = \wedge$. Let $\psi_1 \wedge \psi_2$ be the active formula in the conclusion $\mathsf{last}(\mathcal{M}_<)$. We let $\overline{f}(\mathcal{M})$ be the premiss corresponding to the conjunct selected by f at $(\psi_1 \wedge \psi_2, s(\mathcal{M}_<))$. Again, it is not hard to see that it suffices to set $s(\overline{f}(\mathcal{M})) := s(\mathcal{M}_<)$.

(iv) $\mathsf{r} = \mathsf{K}$. There is only a single choice for $\overline{f}(\mathcal{M})$. Note that there is a unique active formula in the conclusion $\mathsf{last}(\mathcal{M}_<)$ of the form $\Box\psi$. We let $s(\overline{f}(\mathcal{M}))$ be the state selected by f at $(\Box\psi, s(\mathcal{M}_<))$ and leave it to the reader to verify that this suffices.

This finishes the proof. □

2.3.14. Lemma. *Let f be positional and winning for $\forall$ in $\mathcal{E}(\bigvee\Gamma, \mathbb{S})@(\bigvee\Gamma, s)$, and let $\overline{f}$ be the strategy for Refuter in $\mathcal{G}(\Gamma)$ given by Lemma 2.3.13. For each tightening $\varphi_{k_0}, \dots, \varphi_{k_n}$ of a finite trail on some $\overline{f}$-guided $\mathcal{G}(\Gamma)@\Gamma$ match, there is an f-guided $\mathcal{E}(\bigvee\Gamma, \mathbb{S})$-match*

$$(\bigvee\Gamma, s)\cdots(\varphi_{k_0}, s_0)\cdots(\varphi_{k_n}, s_n).$$

Proof:
Write $\mathcal{M}_n$ for the initial segment of $\mathcal{M}$ of length $2n+1$. That is,

$$\mathcal{M}_n = \Gamma_0 \cdot i_0 \cdot \Gamma_1 \cdot i_1 \cdots \Gamma_n.$$

We claim that we can prove the theorem by defining $s_i := s(\mathcal{M}_{k_i})$ for each $0 \le i \le n$. By induction on i, we will show that each $\mathcal{E}(\bigvee\Gamma, \mathbb{S})$-match

$$\mathcal{N}_i := (\bigvee\Gamma, s)\cdots(\varphi_{k_0}, s_0)\cdots(\varphi_{k_i}, s_i)$$

is f-guided.

For the base case note that φ_{k_0} belongs to $\Gamma_{k_0} = \Gamma$, and that $s_0 = s$, since φ_{k_0} is the first formula on the trail such that φ_{k_0} is active in the conclusion and its successor is active in the premiss. Hence there clearly is an f-guided match $(\bigvee\Gamma, s)\cdots(\varphi_{k_0}, s_0)$.

For the inductive step, suppose that we have proven that $\mathcal{N}_i$ is f-guided for $i < n$. We wish to show that $\mathcal{N}_i \cdot (\varphi_{k_{i+1}}, s_{i+1})$ is f-guided. By construction, the state s_{i+1} is accessible from the state s_i in $\mathbb{S}$. In case (φ_{k_i}, s_i) is a position that belongs to $\exists$, the result follows from the fact that the tightening of a trail is a trace. To finish the proof, suppose that (φ_{k_i}, s_i) belongs to $\forall$. We must show that $f(\varphi_{k_i}, s_i) = (\varphi_{k_{i+1}}, s_{i+1})$, but this is immediate from conditions (i), (ii), and (iii) of Lemma 2.3.13. □

2.3.15. Proposition. *If Prover has a winning strategy in $\mathcal{G}(\Gamma)@\Gamma$, then Γ is valid.*

Proof:
We will show this by contraposition. So suppose that Γ is invalid. Then there is a model $\mathbb{S}$ and state s of $\mathbb{S}$ such that $\forall$ has a winning strategy in $\mathcal{E}(\bigvee\Gamma, s)$.

By positional determinacy, we may assume without loss of generality that f is positional.

We claim that the strategy $\overline{f}$ given by Lemma 2.3.13 is winning for Refuter in $\mathcal{G}(\Gamma)@\Gamma$. Indeed, suppose that $\mathcal{M}$ is an $\overline{f}$-guided $\mathcal{G}(\Gamma)@\Gamma$-match. By condition (iii) of Lemma 2.3.13, the match cannot reach an axiom, whence $\mathcal{M}$ must be infinite. Suppose $(\varphi_{k_n})_{n\in\omega}$ is an infinite tightening of a trail on $\mathcal{M}$. By Lemma 2.3.14, there must be a f-guided $\mathcal{E}(\bigvee\Gamma, s)$-match

$$(\varphi_{k_0}, s_0) \cdot (\varphi_{k_1}, s_1) \cdot (\varphi_{k_2}, s_2) \cdots$$

Moreover, from the fact that f is winning for $\forall$, it follows that $(\varphi_{k_n})_{n\in\omega}$ must be a ν-trace. Hence $\overline{f}$ is winning for Refuter, which means that Prover cannot have a winning strategy in $\mathcal{G}(\Gamma)@\Gamma$. □

2.3.5 Completeness of NW

In this subsection we will prove that NW is complete. We will use the converse of the argument in the previous subsection. That is, we will show that a winning strategy f for Refuter in $\mathcal{G}@\Gamma$ induces a model $\mathbb{S}^f$ and strategy for $\forall$ in $\mathcal{E}(\bigvee\Gamma, \mathbb{S}^f)$, showing that $\mathbb{S}^f$ falsifies Γ.

2.3.16. Definition. A μML-formula ξ is said to be *guarded* if every subformula $\eta x\varphi$ of ξ only contains occurrences of x in the scope of a modality $\triangle \in \{\Diamond, \Box\}$.

2.3.17. Example. $\mu x(\Diamond x \wedge \Box x)$ is guarded, but $\nu y(\mu x(\Diamond x \wedge y))$ is not.

It turns out that every tidy μML-formula in negation normal form is equivalent to a tidy guarded formula in negation normal form. A proof of this result, originally by Kozen, can be found for instance in [107, Proposition 2].

2.3.18. Proposition. *Let φ be a tidy μML-formula in negation normal form. Then there is a guarded and tidy μML-formula φ' in negation normal form such that $\varphi \equiv \varphi'$.*

For the rest of this section we shall refer to formulas which are tidy, guarded and in negation normal form by the term *nice*.

The key property of nice formulas is that their traces have the following behaviour.

2.3.19. Lemma. *Suppose $(\varphi_n)_{n\in\omega}$ is an infinite trace starting at a nice formula φ. Then infinitely many φ_i have a modal operator as main operator.*

Proof:
Let the *exposure rank* $\mathsf{er}(\xi)$ of a nice formula ξ be the amount of occurrences of fixed point operators *outside* of the scope of a modal operator. Formally, er is defined by the following induction.

- If ξ is a literal or of the form $\triangle\varphi$, then $\mathsf{er}(\xi) = 0$.
- If $\xi = \varphi \circ \psi$, then $\mathsf{er}(\xi) = \mathsf{er}(\varphi) + \mathsf{er}(\psi)$.
- If $\xi = \eta x\varphi$, then $\mathsf{er}(\xi) = \mathsf{er}(\varphi) + 1$.

Let $\eta x\varphi$ be a nice formula. We claim that $\mathsf{er}(\varphi[\eta x\varphi/x]) = \mathsf{er}(\varphi)$. Indeed this follows from the fact that, by guardedness, every occurrence of $\eta x\varphi$ in $\varphi[\eta x\varphi/x]$ is in the scope of a modal operator. Hence $\mathsf{er}(\eta x\varphi) = \mathsf{er}(\varphi[\eta x\varphi/x]) + 1$.

Moreover, note that $\mathsf{er}(\psi_1 \circ \psi_2) \leq \mathsf{er}(\psi_i)$ for each $i \in \{1, 2\}$. Therefore the trace $(\varphi_n)_{n\in\omega}$ must contain infinitely many formulas whose main operator is modal, for otherwise there would be a final segment on which the exposure rank weakly decreases at every step, and strictly decreases at infinitely many steps. □

2.3.20. Definition. An application of K

$$\mathsf{K}\,\frac{\Delta, \varphi}{\Gamma, \Diamond\Delta, \Box\varphi}$$

is called *optimal* if Γ contains only literals and $\Box$-formulas.

2.3.21. Definition. An application of some rule $\mathsf{r} \in \{\vee, \wedge, \mu, \nu\}$ is called *reductive* if the active formula in the conclusion is not also inactive.

2.3.22. Example. Of the following two rules instances, the one left is reductive, whereas the one on the right is not.

$$\mu\,\frac{\Gamma, \Diamond\mu x(\Diamond x \vee p) \vee p}{\Gamma, \mu x(\Diamond x \vee p)} \qquad \mu\,\frac{\Gamma, \mu x(\Diamond x \vee p), \Diamond\mu x(\Diamond x \vee p) \vee p}{\Gamma, \mu x(\Diamond x \vee p)}$$

2.3.23. Definition. Let f be a strategy for Refuter in $\mathcal{G}(\Gamma)@\Gamma$. The *countermodel tree* T_f of f is the subtree of the strategy tree of f, where Prover only plays optimal applications of K, and only reductive applications of $\vee, \wedge, \mu, \nu$.

2.3.24. Definition. Let f be a strategy for Refuter in $\mathcal{G}(\Gamma)@\Gamma$. The *canonical model* $\mathbb{S}^f$ of f consists of:

- a set of states S^f containing precisely the maximal paths through the countermodel tree T^f which contain no application of K;
- a relation R^f given by $\rho_1 R^f \rho_2$ if and only if $\mathsf{last}(\rho_1)$ is connected to $\mathsf{first}(\rho_2)$ by an application of K in T^f;
- a valuation V^f given by $p \in V^f(\rho)$ if and only if p does *not* belong to any sequent Γ on the path ρ.

The following lemma captures the key reason for working with guarded formulas.

2.3.25. LEMMA. *Let f be a strategy for Refuter in some match $\mathcal{G}(\Gamma)@\Gamma$, where Γ is a set of nice formulas. Then every state ρ of $\mathbb{S}^f$ is finite.*

Proof:
Adding a single root, the tightenings of trails on ρ form a tree. If ρ were infinite, this tree would by Kőnig's Lemma have an infinite branch. But this is in contradiction with Lemma 2.3.19 and the fact that ρ contains no application of the modal rule. □

The next lemma follows from the restriction on the moves of Prover in T^f.

2.3.26. LEMMA. *If a formula of the form $\psi_1 \circ \psi_2$ or of the form $\eta x \psi$ occurs in some sequent Δ on a path $\rho \in S^f$, it will be the active formula in the conclusion of some rule application above Δ on ρ.*

2.3.27. PROPOSITION. *Let f be a winning strategy for Refuter in $\mathcal{G}(\Gamma)@\Gamma$ and let $\rho_0 \in S^f$ be any path containing the root of T^f. Then $\forall$ has a winning strategy $\underline{f}$ in $\mathcal{E}(\bigvee \Gamma, \mathbb{S}^f)@(\bigvee \Gamma, \rho_0)$.*

Proof:
It clearly suffices to show that $\underline{f}$ is winning in $\mathcal{E}(\bigvee \Gamma, \mathbb{S}^f)@(\varphi_0, \rho_0)$ for an arbitrary formula $\varphi_0 \in \Gamma$. By induction on n, we will simultaneously define $\underline{f}$ on every $\mathcal{E}(\bigvee \Gamma, \mathbb{S}^f)$-match

$$\mathcal{M} = (\varphi_0, \rho_0) \cdots (\varphi_n, \rho_n),$$

and show that for any $\underline{f}$-guided extension $\mathcal{M} \cdot (\varphi_{n+1}, \rho_{n+1})$ of $\mathcal{M}$ the sequence $\varphi_0 \cdots \varphi_{n+1}$ is the tightening of some trail on a $\mathcal{G}@\Gamma$-match of the form $\mathcal{N} \cdot \rho$, where ρ is an initial segment of ρ_{n+1}.

Suppose that the thesis holds for all $k < n$. If $n > 0$, we know, by the induction hypothesis, that $\varphi_0 \cdots \varphi_n$ is the tightening of some trail on a $\mathcal{G}@\Gamma$-match of the form $\mathcal{N} \cdot \rho$, where ρ is an initial segment of ρ_n. We let $\Delta = \mathsf{last}(\rho)$ be the premiss in which φ_n is active, which witnesses that φ_n is part of the tightening $\varphi_0 \cdots \varphi_n$. If $n = 0$ we let Δ be Γ. Note that in either case $\varphi_n \in \Delta$. We make a case distinction based on the shape of φ_n.

- If φ_n is a literal, there is nothing to do.

- If $\varphi_n = \psi_1 \vee \psi_2$, then the thesis holds, because, by Lemma 2.3.26, the formula φ_n must at some point in ρ_n above Δ be the active formula in the conclusion of an application of $\vee$, in which case each of the ψ_i is active in the premiss.

- If $\varphi_n = \psi_1 \wedge \psi_2$, there must again be some point in ρ_n above Δ such that $\psi_1 \wedge \psi_2$ is the active formula in the conclusion of an application of $\wedge$. We set $\underline{f}(\varphi_n, \rho_n) := (\psi_i, \rho_n)$, where ψ_i is the active formula in the *first* such application of $\wedge$ in ρ_n above Δ. The required property clearly holds.

- Suppose $\varphi_n = \Diamond\psi$. Let $\mathcal{M} \cdot (\psi, \rho_{n+1})$ be an extension of $\mathcal{M}$. It follows by definition that $\rho_n \cdot (\mathsf{last}(\rho_n), \mathsf{K}, \langle \mathsf{first}(\rho_{n+1})\rangle) \cdot \rho_{n+1}$ is a $\mathcal{G}@\Gamma$-match. The initial segment $\rho_n \cdot (\mathsf{last}(\rho_n), \mathsf{K}, \langle \mathsf{first}(\rho_{n+1})\rangle) \cdot \mathsf{first}(\rho_{n+1})$ satisfies the required condition, since ψ is a direct ancestor of $\Diamond\psi$.

- Suppose $\varphi_n = \Box\psi$. Since by the induction hypothesis $\Box\psi$ occurs in ρ_n, we have $\Box\psi \in \mathsf{last}(\rho_n)$ (which exists by Lemma 2.3.25). Hence, there is a ρ_{n+1} in T^f such that $\rho_n \cdot (\mathsf{last}(\rho_n), \mathsf{K}, \langle \mathsf{first}(\rho_{n+1})\rangle) \cdot \rho_{n+1}$ is a path in T^f. In particular, we have $\rho_n R^f \rho_{n+1}$. We set $\underline{f}(\varphi_n, \rho_n) := (\psi, \rho_{n+1})$. The initial segment $\rho_n \cdot (\mathsf{last}(\rho_n), \mathsf{K}, \langle \mathsf{first}(\rho_{n+1})\rangle) \cdot \mathsf{first}(\rho_{n+1})$ again satisfies the required condition.

- Suppose $\varphi = \eta x \psi$. By Lemma 2.3.26 there must be some point in ρ_n above Δ where φ is the active formula of an application of η, whence the result follows.

This defines $\underline{f}$. We now claim that $\underline{f}$ is, in fact, a winning strategy for $\forall$ in $\mathcal{E}(\bigvee \Gamma, \mathbb{S}^f)@(\varphi_0, \rho_0)$. Suppose first that some full $\underline{f}$-guided match $\mathcal{M}$ is finite, ending in, say, (φ_n, ρ_n). Then φ_n cannot be of the form $\Box\psi$, for then $\underline{f}(\varphi_n, \rho_n)$ would be defined and thus $\mathcal{M}$ would not be full. Suppose then, that φ_n is a literal. By the construction above, we know that φ_n is the final element of the tightening of some trace on some $\mathcal{G}@\Gamma$-match of the form $\mathcal{N} \cdot \rho$, where ρ is an initial segment of ρ_n. Hence, φ_n appears on the path ρ_n. Since T^f is built from a winning strategy f for Refuter, we know that $\overline{\varphi_n}$ does not appear on the path ρ_n. By the definition of V^f, we find that $\mathcal{M}$ is indeed won by Refuter.

Now suppose that $\mathcal{M}$ is an infinite $\underline{f}$-guided match

$$\mathcal{M} = (\varphi_0, \rho_0) \cdot (\varphi_1, \rho_1) \cdot (\varphi_2, \rho_2) \cdots$$

By the construction above, we know for every n that $\varphi_0 \cdots \varphi_n$ is the tightening of some trail on some $\mathcal{G}@\Gamma$-match of the form $\mathcal{N} \cdot \rho$, where ρ is an initial segment of ρ_n. Consider the infinite $\mathcal{G}@\Gamma$-match

$$\mathcal{N} = \rho_0 \cdot (\mathsf{last}(\rho_0), \mathsf{K}, \langle \mathsf{first}(\rho_1)\rangle) \cdot \rho_1 \cdot (\mathsf{last}(\rho_1), \mathsf{K}, \langle \mathsf{first}(\rho_2)\rangle) \cdot \rho_2 \cdots$$

By the above, the trails whose tightenings are of the form $\varphi_0 \cdots \varphi_n$ for some n, form an infinite, finitely branching tree, through $\mathcal{N}$. Hence, by Kőnig's Lemma, the match $\mathcal{N}$ contains an infinite trail whose tightening is the trace $(\varphi_n)_{n\in\omega}$. Since T^f is based on a winning strategy for Refuter, this trace must be a ν-trace, and therefore $\mathcal{M}$ is won by $\forall$. □

2.3.28. Corollary. *If Γ is a valid set of nice formulas, then Prover has a winning strategy in $\mathcal{G}(\Gamma)@\Gamma$.*

Proof:
As in the proof of soundness, we argue by contraposition. Suppose that Prover has no such winning strategy. By Theorem 2.2.10, it follows that Refuter has a winning strategy in $\mathcal{G}(\Gamma)@\Gamma$. But then the previous proposition implies that Γ is invalid, as required. □

2.3.29. Remark. The system NW is, in fact, also complete for unguarded formulas. However, to prove a general completeness result we would have to introduce technical machinery that is outside the scope of this thesis. In Chapter 5 we will introduce so-called *trace atoms*, whose purpose is to deal with the combinatorics of backward modalities in the two-way modal μ-calculus. We will see that, as a by-product, trace atoms facilitate a relatively easy completeness argument that covers unguarded formulas as well.

2.3.6 From trace-based to path-based proof systems

In this section we assume some familiarity with the theory of automata operating on infinite languages. For a good introduction we refer the reader to Chapter 1 of [49].

The non-well-founded proof system NW is of a kind that is often called *trace-based*. This is because the soundness condition on infinite branches is defined in terms of traces. A common theme in the literature is to turn trace-based systems into so-called *path-based* systems, where the soundness condition is instead defined on the branches themselves. Path-based proof systems are inspired by automaton theory, and the most direct way to construct them also uses automata. In this section we will sketch how to apply this direct construction to NW.

As usual, we call a language of infinite words *ω-regular* if it is recognised by some non-deterministic parity ω-automaton. A finite *game* $\mathcal{G} = (B_0, B_1, E, W)$ is called *ω-regular* if $W^{-1}(0)$ is ω-regular (or, equivalently, if $W^{-1}(1)$ is ω-regular).

We will use the following fact. while its proof is outside the scope of this thesis, the reader might not find it hard to imagine, seeing the correspondence between traces on branches of NW-proofs and matches in the evaluation game.

2.3.30. Fact. For every Γ, the proof search game $\mathcal{G}(\Gamma)$ is ω-regular, recognised by an automaton of size linear in $|\mathsf{FL}(\Gamma)|$.

It turns out that non-deterministic parity ω-automata have the same expressive power as their deterministic counterparts. This follows by combining several results about transformations of different types of ω-automata, the most important of which is a deterministation construction originally due to Shmuel Safra. For more information we refer the reader to Chapter 3 of [49].

2.3.31. THEOREM. *Every ω-regular language can be recognised by a deterministic parity automaton. In particular, for for every non-deterministic parity automaton $\mathbb{A}$, there is an equivalent parity automaton $\mathbb{B}$ such that the number of states of $\mathbb{B}$ is exponential in the number of states of $\mathbb{A}$.*

The following proposition game-theoretically captures the essence of the relationship between trace-based and path-based proof systems.

2.3.32. PROPOSITION. *Let $\mathcal{G} = (B_0, B_1, E, W)$ be an ω-regular game. Then there is a parity game $\mathcal{G}' = (B'_0, B'_1, E', W')$ and a surjection $\pi : B' \to B$ such that:*

(i) there is a k such that $|\pi^{-1}(a)| \leq k$ for every $a \in B$;

(ii) If aEb, then for every $a' \in \pi^{-1}(a)$ there is a unique $b' \in \pi^{-1}(b)$ with $a'E'b'$.

(iii) $a' \in B'_i$ if and only if $\pi(a') \in B_i$ for each $a' \in B'$ and $i \in \{0, 1\}$.

(iv) $W' = W \circ \pi^\omega$, where π^ω is defined pointwise.

Proof (sketch):
By Theorem 2.3.31 we may assume that $W^{-1}(0)$ is recognised by a deterministic parity automaton $(Q, \Sigma, \delta, q_I, \mathrm{Acc})$. The idea is to let $\mathcal{G}'$ be the following game.

- The board B' is $B \times Q$, where $B'_0 = B_0 \times Q$ and $B'_1 = B_1 \times Q$.
- The relation E' is given by $(a, p)E'(b, q)$ if and only if aEb and $q = \delta(p, a)$.
- The function W' is given by:

$$W'((b_n, q_n)_{n\in\mathbb{N}}) = 0 :\Leftrightarrow (q_n)_{n\in\mathbb{N}} \in \mathrm{Acc}.$$

Let $\pi : B' \to B$ be the left projection function. Clearly π satisfies conditions (i) - (iii). Condition (iv) follows from the fact that for each infinite $\mathcal{G}'$-match $(b_n, q_n)_{n\in\mathbb{N}}$, we have

$$W'((b_n, q_n)_{n\in\mathbb{N}}) = 0 \Leftrightarrow (q_n)_{n\in\mathbb{N}} \in \mathrm{Acc} \Leftrightarrow W((b_n)_{n\in\mathbb{N}}) = 0,$$

as required. □

Let us now return to the goal of obtaining a path-based counterpart to NW. By Fact 2.3.30, we know that the proof-search game $\mathcal{G}(\Gamma)$ for NW is ω-regular. Hence, by Proposition 2.3.32, there is a game $\mathcal{G}'(\Gamma)$ with a surjection π satisfying properties (i) - (iv). For p a position of $\mathcal{G}'(\Gamma)$, let us write $p \vdash \Delta$ whenever $\pi(p) = \Delta \in \mathsf{Seq}_\Gamma$, and $p \vdash i$ if $\pi(p) = i \in \mathsf{Inst}_\Gamma$.

Moreover, if $p \vdash \Delta$ and $i \in \mathsf{Inst}_\Gamma$ is such that $\mathsf{conc}(i) = \Delta$, we write $p * i$ for the unique $q \vdash i$ such that p sees q in $\mathcal{G}'(\Gamma)$ (cf. item (ii) of Proposition 2.3.32). Similarly, if $p \vdash (\Delta, \mathsf{r}, \langle \Delta_1, \ldots, \Delta_n \rangle)$, we write $p * \Delta_i$ for the unique $q \vdash \Delta_i$ such that p sees q in $\mathcal{G}'(\Gamma)$.

Position	Owner	Admissible moves
$p \vdash \Delta$	Prover	$\{p * i \vdash i \mid \mathsf{conc}(i) = \Delta\}$
$p \vdash (\Delta, \mathsf{r}, \langle \Delta_1, \dots, \Delta_n \rangle)$	Refuter	$\{p * \Delta_i \vdash \Delta_i \mid 1 \leq i \leq n\}$

We can now reconstruct a proof system, let us call it NW', from $\mathcal{G}'(\Gamma)$ by taking for each rule instance $i = (\Delta, \mathsf{r}, \langle \Delta_1, \dots, \Delta_n \rangle)$ and position p of $\mathcal{G}'$, the rule instances

$$\mathsf{r}\,\frac{p * i * \Delta_1 \vdash \Delta_1 \qquad \cdots \qquad p * i * \Delta_n \vdash \Delta_n}{p \vdash \Delta}$$

Note that one can view sequents in NW' as ordinary sequents, but *annotated* by a position of $\mathcal{G}'$. A closed NW'-*derivation* of $p \vdash \Gamma$ is said to be an NW'-*proof* if for every infinite branch $(p_n \vdash \Delta_n)_{n \in \omega}$ the induced $\mathcal{G}'(\Gamma)$-match

$$p_0 \cdot p_0 * i_0 \cdot p_1 \cdot p_1 * i_1 \cdot p_2 \cdot p_2 * i_2 \cdots$$

is winning for Prover, where for each k, we denote by i_k the rule instance of which p_k is a conclusion.

The system NW' has the nice property that its proof-search game $\mathcal{G}'(\Gamma)$ is a parity game and hence is positionally determined. In Section 2.3.8 we will see some consequences of this fact. On the other hand, the system NW' is not very practical. It is not directly clear how to construct proofs top-down, and even when constructing a proof bottom-up, we must, simultaneously, run an automaton to calculate the annotations. A considerable amount of research has been done on finding proof-theoretically more satisfactory ways of annotating sequents, in such a way that the natural corresponding proof-search game is still positionally determined. We give a few examples below.

- For some modal fixed point logics relatively simple annotations suffice. In Chapter 2 we will see that for $\mathtt{ML}^*$ it suffices to annotate (hyper)sequents by only a single focus annotation. This technique was originally developed by Lange & Stirling for the temporal logics $\mathtt{LTL}$ and $\mathtt{CTL}$ [68].

- It turns out that for the alternation-free modal μ-calculus slightly more complex, but still relatively simple annotations suffice. In [73], Marti & Venema present a system for $\mu^{af}\mathtt{ML}$ where sequents are of the form

  $$\varphi_1^{u_1}, \dots, \varphi_n^{u_1}$$

  and the $u_i \in \{\circ, \bullet\}$ indicate if a formula is *out of focus* ($\circ$) or *in focus* ($\bullet$).

- A relatively prominent path-based proof system for the modal μ-calculus is the system JS, by Jungteerapanich [54] & Stirling [100]. Its sequents have the shape

  $$\Theta \vdash \varphi^{\rho_1}, \dots, \varphi^{\rho_n}$$

  where Θ and the ρ_i are sequences of so-called *names*. The system JS moreover features some additional rules for annotation management. These additional rules are closely linked to the several stages which together compose Safra's construction for determinising ω-automata.

2.3.7 Cyclic proofs

So far we have only seen non-well-founded proofs of an infinitary nature. Indeed, both NW and NW' are only complete in case we allow certain infinite branches. In this section we will see how to obtain a so-called *cyclic* proof system for the modal μ-calculus.

2.3.33. Definition. A non-well-founded derivation is called *regular* if it has only finitely many subtrees.

The key property of regular proofs is that they can be finitely represented. To see this, we need the following definition.

2.3.34. Definition. A *finite tree with back edges* (T, f) consists of a finite tree T together with a partial function f from T to itself, such that (i) $\mathsf{dom}(f)$ consists of leaves of T, and (ii) $f(u)$ is an ancestor of u for every $u \in \mathsf{dom}(f)$.

An element $u \in \mathsf{dom}(f)$ is often called a *repeating leaf*, and $f(u)$ is then called its *companion*.

2.3.35. Definition. For $\mathsf{P} \in \{\mathsf{NW}, \mathsf{NW}'\}$, a *cyclic* P*-derivation* is a finite tree with back edges (π, f), where π is a finite P-derivation and for every leaf $u \in \mathsf{dom}(f)$ it holds that u and $f(u)$ are labelled by identical sequents.

Any cyclic derivation can be *unravelled* into a non-well-founded derivation. Informally, this works by recursively pasting the subtree generated by some companion at each of its leaves. In the Intermezzo following Chapter 2 we will give a formal definition of this construction. The proof of the following proposition is postponed to the Intermezzo as well.

2.3.36. Proposition. *A non-well-founded derivation π is regular if and only if it is the unravelling of some finite trees with back edges.*

Since, as we have seen, non-well-founded derivations are not necessarily sound, the same holds for cyclic derivations. We therefore need to impose a soundness condition similar to the ones on the branches of NW-proofs and NW'-proofs. The most straightforward way to do this is as follows.

2.3.37. Definition. A cyclic NW-derivation (π, f) is a *cyclic* NW*-proof* if its unravelling is an NW-proof.

In the next section we will see that cyclic NW-proofs, or equivalently, the subset of regular non-well-founded NW-proofs, are complete with respect to the modal μ-calculus. A drawback of this notion of cyclic proofs is that, in order to check if a cyclic derivation is a bona fide proof, one has to first unravel it and check the resulting non-well-founded proof.

For annotated proof systems, such as NW', there is often a soundness condition which can be defined in the terms of the paths between companions and their repeating leaves. As a result, one only has to check the finite cyclic derivation to see whether it is indeed a proof. All three examples of path-based proof systems given above feature such a simple soundness condition for cyclic proofs. The same holds for the systems we present in Chapter 3, Chapter 5, and Chapter 6.

2.3.8 The bounded model and proof properties

In this subsection we will exploit the fact that the proof-search game $\mathcal{G}'(\Gamma)$ for NW' is a parity game to infer some results about NW and the modal μ-calculus.

2.3.38. PROPOSITION. *Every valid sequent has a cyclic* NW*-proof.*

Proof:
Suppose Γ is a valid sequent. By Corollary 2.3.28, Prover has a winning strategy in $\mathcal{G}(\Gamma)@\Gamma$. Hence, by construction, Prover also has a winning strategy in the proof-search game $\mathcal{G}'(\Gamma)@(p \vdash \Gamma)$ of NW', where p is any position of $\mathcal{G}'(\Gamma)$ such that $\pi(p) = \Gamma$. Since $\mathcal{G}'(\Gamma)$ is a parity game, it follows by Theorem 2.2.10 that Prover has a positional strategy. Since $\mathcal{G}(\Gamma)$ has only finitely many positions, it follows from item (i) of Proposition 2.3.32 that $\mathcal{G}'(\Gamma)$ has only finitely many positions as well. Hence the positional winning strategy for Prover in $\mathcal{G}'(\Gamma)@(p \vdash \Gamma)$ corresponds to a regular NW'-proof τ'. Indeed, positional determinacy implies that equal positions generate the isomorphic subtrees. Dropping all automaton states annotating the sequents in τ', we obtain a regular NW-proof τ. Finally, by Proposition 2.3.36, the proof τ is the unravelling of some cyclic NW-proof of Γ. □

2.3.39. COROLLARY. *Every valid sequent* Γ *has a cyclic* NW*-proof whose size is doubly exponential in* $|\mathsf{FL}(\Gamma)|$*.*

Proof:
The key observation is that the derivation τ in the above proof is the unravelling of a certain small cyclic proof. Namely, we can take the cyclic proof obtained by drawing a back edge at every first repetition in τ'. The depth of this first repetition is bounded by the number of distinct annotated sequents. Since of the size of the deterministic automaton is exponential in $|\mathsf{FL}(\Gamma)|$, and the number of distinct (unannotated)sequents is as well, we find that the number of distinct annotated sequents is exponential in $|\mathsf{FL}(\Gamma)|$. Since the maximal branching is constant, the size of τ' is doubly exponential in $|\mathsf{FL}(\Gamma)|$. □

2.3.40. REMARK. The previous corollary establishes what is called the *bounded proof property*. In the Intermezzo following Chapter 2, we provide an alternative way of proving the bounded proof property for an abstract notion of an annotated non-well-founded proof system. Unlike the method of this chapter, our method does not rely on game-theoretic results.

By applying the same procedure to winning strategies for Refuter, rather than Prover, we obtain the *bounded model property*.

2.3.41. PROPOSITION. *Every invalid sequent* Γ *has a countermodel whose size is doubly exponential in* $|\mathsf{FL}(\Gamma)|$.

Proof (sketch):
By the same reasoning as above, we can obtain a representation of a winning strategy f for Refuter in $\mathcal{G}(\Gamma)@\Gamma$ as a finite tree with back edges, whose size is bounded by a computable function of Γ. This, in turn, can be used to construct a finite version of the canonical model $\mathbb{S}^f$, where the back edges in the representation of the refutation, become back edges in the model as well. Since both versions of the canonical model are bisimilar, we obtain a countermodel of Γ of the same size as the representation of the refutation. □

2.3.42. REMARK. By considering proofs and refutations as graphs rather than finite trees with back edges, the bounds of the previous corollary and proposition can be sharpened to become singly exponential. For more details we refer the reader to Section 6 of [77].

2.4 Frame conditions

The goal of this thesis is to extend the theory described in the chapter so far, to various fragments and variants of the modal μ-calculus. Moreover, we wish to do this in a uniform way. We will mostly generate variants of the modal μ-calculus by interpreting the language over restricted classes of frames. In this section we briefly introduce some relevant definitions, and discuss some known results in this field of study.

2.4.1 Preliminaries

2.4.1. DEFINITION. A *Kripke frame* of type D is a pair $(S, (R_a)_{a \in \mathsf{D}})$, where S is a set of *states* and for each $a \in \mathsf{D}$, $R_a \subseteq S \times S$ is an *accessibility relation*.

Note that a Kripke frame is simply a Kripke model without a valuation. We say that a formula φ is *valid* in some frame (S, R), and write $(S, R) \models \varphi$ if $(S, R, V), s \Vdash \varphi$ for every valuation $V : \mathsf{P} \to \mathcal{P}(S)$ and state $s \in S$.

2.4.2. Definition. A *(basic modal) logic* L is any set of formulas in the basic modal language $\mathtt{ML}$ closed under the following axioms and rules.

Axioms.

1. A sound and complete set of axioms for classical propositional logic.
2. Normality: $\neg\Diamond\bot$.
3. Additivity: $\Diamond(p \vee q) \leftrightarrow (\Diamond p \vee \Diamond q)$.
4. Dual for $\Box$: $\Box p \leftrightarrow \neg\Diamond\neg p$.

Rules.

1. Modus Ponens: from $\varphi \to \psi$ and φ, derive ψ.
2. Monotonicity: from $\varphi \to \psi$, derive $\Diamond\varphi \to \Diamond\psi$.
3. Uniform Substitution: from φ, derive $\varphi[\psi/x]$.

The smallest basic modal logic is denoted by K. Given a logic L, we say that (S,R) is an L-*frame* and write $(S,R) \models \mathsf{L}$ if $(S,R) \models \varphi$ for every $\varphi \in \mathsf{L}$.

2.4.3. Definition. A logic is *finitely axiomatisable* if is the smallest logic containing some finite set of axioms.

Let $\mathtt{L_1}(\mathsf{D})$ be the first-order language with equality and a a relation symbol R_a for every $a \in \mathsf{D}$. A *(first-order) frame condition* is simply an $\mathtt{L_1}(\mathsf{D})$-sentence. For Θ a set of frame conditions, a Kripke frame $(S,(R_a)_{a\in\mathsf{D}})$ is said to be a Θ-*frame* whenever, when regarded an $\mathtt{L_1}(\mathsf{D})$-structure, the frame $(S,(R_a)_{a\in\mathsf{D}})$ satisfies all sentences in Θ. A Kripke model will be called a Θ-*model* whenever its underlying frame is a Θ-frame.

2.4.2 A negative result

The following theorem is a reformulation of a result proven by Edith Hemaspaandra. Recall that a first-order formula is *universal* if it consists of a quantifier-free formula preceded by a string of universal quantifiers.

2.4.4. Theorem ([51, Theorem 3.5]). *There is a class* F *of frames such that:*

- $\mathsf{F} = \{(S,R) \mid (S,R) \models \psi\}$*, where* ψ *is a universal first-order formula;*
- $\mathsf{F} = \mathsf{Fr}(\mathsf{L})$*, for* L *a finitely axiomatisable and canonical*[1] *basic modal logic;*
- *the set* $\{\varphi \in \mathtt{ML} \mid \mathsf{F} \models \varphi\}$ *is decidable;*
- *the set* $\{\varphi \in \mathtt{ML}^* \mid \mathsf{F} \models \varphi\}$ *is not recursively enumerable.*

2.4.5. Remark. Hemaspaandra actually proves an even stronger result [51]. Namely, that the set $\{\varphi \in \mathtt{ML}^* \mid \mathsf{F} \not\models \varphi\}$ is Σ^1_1-complete.

[1] This is a standard notion in modal logic and will be defined in Chapter 4.

It follows that for certain frame conditions, such as the one witnessing Hemaspaandra's result, there is no hope of finding a nice proof system. Indeed, any reasonable (that is, computable) proof system, even without the bounded proof property, yields a method for recursively enumerating all validities by simply enumerating all proofs.

2.4.6. Remark. There are two notable positive results in the literature. First, Kikot, Shapirovsky & Zolin in [56] show the soundness and completeness of certain Hilbert-style proof systems for $\mathtt{ML}^*$ over several frame classes which admit the method of *filtration*. In Chapter 4 we extend their technique to $\mu^c\mathtt{ML}$.

Second, Baltag, Bezhanishvili & Fernández-Duque show the completeness of (also Hilbert-style) proof systems for the modal μ-calculus interpreted over several frame classes which are *weakly transitive* and define so-called *subframe logics* [9]. It was later shown that the modal μ-calculus collapses to its alternation-free fragment over all of frame classes to which their result applies [81].

It is also worth mentioning here the work by French [43, 44], and his joint work with D'Agostino, & Lenzi [29], on modal logics with bisimulation quantifiers. Their results entail that the property of *uniform interpolation*, originally proven for the modal μ-calculus over the class of all frames by D'Agostino and Hollenberg [30, 31], transfers to the so-called *idempotent transduction classes.* Although they do not apply their techniques to proof systems for the modal μ-calculus, their work is one of the few examples in the literature where the modal μ-calculus is considered in a general setting over different frame classes.

To the best of our knowledge, the literature contains no similar uniform treatment using non-well-founded proof systems, apart from the research that appears in the next chapters of this thesis.

Chapter 3

Modal logic with the master modality

3.1 Introduction

This chapter builds upon Ori Lahav's paper [66], where hypersequent calculi are constructed uniformly for the basic modal language over classes of frames satisfying *simple* first-order conditions. Lahav first presents hypersequent calculi for four basic frame classes: all frames, the transitive frames, the symmetric frames, and the transitive-symmetric frames. To each simple frame condition, a corresponding hypersequent rule is assigned. It is then shown that extending one of the basic calculi with all rules corresponding to some set of simple frame conditions yields a sound and complete calculus. Completeness is proven in a uniform way, through a canonical model construction, providing an analytic proof for any valid hypersequent. Furthermore, when the basic frame condition does not require symmetry, this proof is cut-free.

Many modal logics cannot be straightforwardly captured by a Gentzen-style sequent calculus. A notorious example is the modal logic S5, for which none of the proposed sequent calculi is entirely satisfactory (see [70] for an impossibility result). It is for this reason that Lahav uses hypersequents, which are finite sets of ordinary sequents. This minor increase in structure allows for a significant increase in expressive power, as illustrated by the existence of a natural cut-free hypersequent calculus for S5.[1] The literature also contains calculi with even more structured sequents, often based on the Kripke semantics (see Chapter 4 of [52] for an overview). Examples include nested sequents, labelled sequents, and display calculi. Unlike these other formalisms, hypersequents maintain the subformula property in its strongest form, ensuring that there are only finitely many hypersequents in any proof. As a consequence, decidability can be directly inferred from soundness and completeness.

[1]This calculus uses a rule corresponding to the simple frame condition of *universality*. Whilst universal frames characterise S5 in the unimodal language, this does not extend to the multi-modal language.

Our aim is to extend Lahav's calculi to accommodate fixed point operators. We focus on a relatively simple modal fixpoint language: *multimodal logic with the master modality.* For each set of frame conditions considered by Lahav, we uniformly construct both an infinitary and a cyclic hypersequent calculus. Sequents are annotated using a focus mechanism, originally due to Lange and Stirling (see *e.g.* [68]). All systems are proven to be sound and complete. Just like Lahav does for basic modal logic, we only obtain cut-free completeness if the basic frame condition does not require symmetry. However, we need as an additional requirement that the other frame conditions are all what we shall call *equable.* Although the equable frame conditions form a relatively small subset of the simple frame conditions, there are infinitely many of them, including seriality, reflexivity, directedness and universality. As will be explained later, all simple frame conditions admit filtration and therefore already PDL admits Hilbert-style proof systems over these frame classes (cf. Section 1.2.2). Hence, over these frame classes we do not have to fear for a negative result analogous to Theorem 2.4.4.

This chapter will not use game-theoretic methods. Instead, we will establish soundness using an argument by infinite descent, employing a measure similar to the *signatures* of [101]. More importantly, we will prove completeness by extending Lahav's canonical model construction. This approach allows us to use similar arguments to show that the canonical model satisfies the necessary frame conditions. We do not know if it is also possible to do this using game-theoretic methods.

Since the size of the canonical model of some given hypersequent H is bounded in the size of H, we obtain the small model property for each set of frame conditions. Decidability follows as a corollary. A natural question is whether there is also a bound on the size of proofs. This will be addressed in the intermezzo following the present chapter.

3.2 Simple and equable frame conditions

In this section we will introduce the frame conditions treated in this chapter. All of the material is from [66], unless specified otherwise.

Recall from Section 2.4 that a first-order frame condition is a sentence in the language $\mathsf{L}_1(\mathsf{D})$ of first-order logic with equality and relation symbols R_a for each $a \in \mathsf{D}$. In this chapter we will restrict attention to *unimodal* frame conditions, where there is only one relation symbol R. We will nevertheless impose these unimodal frame conditions on multimodal frames. That is, for Θ a set of such frame conditions, a Kripke frame $(S, (R_a)_{a \in \mathsf{D}})$ will be called a Θ-*frame* whenever, when regarded an L_1-structure, each frame (S, R_a) with $a \in \mathsf{D}$ satisfies all sentences in Θ. A Kripke model will be called a Θ-*model* whenever its underlying frame is a Θ-frame.

3.2.1. Remark. Note that we assume that each relation R_a in a Θ-model satisfies the same frame conditions. This is only for notational simplicity and our results can easily be extended to models in which different frame conditions are imposed on different accessibility relations. A further generalisation is to allow *mixed* frame conditions, where a single frame conditions may involve multiple different relation symbols. We leave it to future work to investigate whether our results extend to such frame conditions as well.

As mentioned in the introduction, we will consider sets of frame conditions comprised of a single *basic* frame condition, extended by a set of *simple* frame conditions. The four basic frame conditions are given in the following table.

name	L_1-sentence	frame class
K	$\top$	all frames
K4	$\forall x \forall y \forall z (xRy \wedge yRz \to xRz)$	transitive frames
B	$\forall x \forall y (xRy \to yRx)$	symmetric frames
B4	$\mathsf{B} \wedge \mathsf{K4}$	symmetric and transitive frames

The simple frame conditions are defined as follows.

3.2.2. Definition. A frame condition is called *n-simple* whenever it is of the form $\forall s_1 \cdots s_n \exists u \varphi$, where φ is built up using the connectives $\vee$ and $\wedge$ from atomic formulas of the form $s_i R u$ and of the form $s_i = u$, for any i with $1 \leq i \leq n$.

We will call a frame condition *simple* if it is n-simple for some $n \in \omega$. It turns out that the simple frame conditions have a convenient abstract representation.

3.2.3. Definition. Given $n \in \omega$, an *abstract n-simple frame condition* is a finite set C consisting of pairs $(C_R, C_=)$ of subsets $C_R, C_= \subseteq \{1, \ldots, n\}$.

3.2.4. Definition. The *interpretation* of some abstract n-simple frame condition C is the following simple first-order formula:

$$\forall s_1 \cdots s_n \exists u \bigvee_{(C_R, C_=) \in C} \Big(\bigwedge_{i \in C_R} s_i R u \wedge \bigwedge_{j \in C_=} s_j = u \Big).$$

Using disjunctive normal forms, the following proposition is immediate.

3.2.5. Proposition. *Any n-simple frame condition is equivalent to the interpretation of some abstract n-simple frame condition.*

3.2.6. Example. The following table shows some examples of simple frame conditions and their abstract representations. The table appears as Table I in [66].

Name	L_1-formula	Abstract representation
Seriality	$\forall s_1 \exists u(s_1 R u)$	$\{(\{1\}, \emptyset)\}$
Reflexivity	$\forall s_1 \exists u(s_1 R u \wedge s_1 = u)$	$\{(\{1\}, \{1\})\}$
Directedness	$\forall s_1 s_2 \exists u(s_1 R u \wedge s_2 R u)$	$\{(\{1, 2\}, \emptyset)\}$
Degenerateness	$\forall s_1 s_2 \exists u(s_1 = u \wedge s_2 = u)$	$\{(\emptyset, \{1, 2\})\}$
Universality	$\forall s_1 s_2 \exists u(s_1 R u \wedge s_2 = u)$	$\{(\{1\}, \{2\})\}$
Linearity	$\forall s_1 s_2 \exists u((s_1 R u \wedge s_2 = u) \vee (s_2 R u \wedge s_1 = u))$	$\{(\{1\}, \{2\}), (\{2\}, \{1\})\}$
Bounded Cardinality	$\forall s_1 \cdots s_n \exists u(\bigvee_{1 \leq i < j \leq n}(s_i = u \wedge s_j = u))$	$\{(\varnothing, \{i, j\}) : 1 \leq i < j \leq n\}$
Bounded Top Width	$\forall s_1 \cdots s_n \exists u(\bigvee_{1 \leq i < j \leq n}(s_i R u \wedge s_j R u))$	$\{(\{i, j\}, \emptyset) : 1 \leq i < j \leq n\}$
Bounded Acyclic Subgraph	$\forall s_1 \cdots s_n \exists u(\bigvee_{1 \leq i < j \leq n}(s_i R u \wedge s_j = u))$	$\{(\{i\}, \{j\}) : 1 \leq i < j \leq n\}$
Bounded Width	$\forall s_1 \cdots s_n \exists u(\bigvee_{1 \leq i, j \leq n, i \neq j}(s_i R u \wedge s_j = u))$	$\{(\{i\}, \{j\}) : 1 \leq i, j \leq n; i \neq j\}$

For the sake of readability we will blur the distinction between an abstract frame condition C and its interpretation. In particular, for $\mathcal{C}$ a set of abstract simple frame conditions and Θ the set of their interpretations, we will often use the terms $\mathcal{C}$-model and $\mathcal{C}$-frame instead of Θ-model and Θ-frame.

3.2.7. Remark. Frame classes definable by a simple (first-order) frame condition are not necessarily also *modally definable*. For instance, the class of frames satisfying the above condition of *linearity* is not closed under disjoint unions.

The following subsets of simple frame conditions will play an important role in this chapter. They are precisely the frame conditions for which we will be able to prove cut-free completeness (provided the basic frame conditions does not require symmetry). As such, the following definition is novel and does not already appear in [66].

3.2.8. Definition. An abstract n-simple frame condition C is called:

- *equality-free* if $C_= = \emptyset$ for all $(C_R, C_=) \in C$;
- *disjunction-free* if C is a singleton;
- *equable* if for some $U \subseteq \{1, \ldots, n\}$, we have $U = C_=$ for all $(C_R, C_=) \in C$.

Clearly if C is equality-free or disjunction-free, it is equable. It turns out that the converse is also true (up to logical equivalence). The verification of this fact is left to the reader. Some examples of equable frame conditions are reflexivity and *k-bounded top width*, which is given by $C = \{\langle \{i, j\}, \emptyset \rangle : 1 \leq i < j \leq k\}$ for any $k \geq 2$. A non-example is given by the simple frame condition of linearity.

3.2.9. Remark. Note that each simple frame condition is a positive first-order formula. Therefore, all simple frame conditions are preserved by surjective homomorphisms. It follows that basic modal logic admits filtration over all simple frame conditions (see *e.g.* [26, Theorem 5.28]). In Chapter 3 we will see that the same holds for $\mathtt{ML}^*$. As a consequence, for none of the simple frame conditions a negative result such as given in Section 2.4.2 holds.

Given a basic frame condition $\mathsf{X} \in \{\mathsf{K}, \mathsf{K4}, \mathsf{B}, \mathsf{B4}\}$ and a set $\mathcal{C}$ of simple frame conditions, we will use the term $\mathcal{C}\mathsf{X}$-*model* to refer to a model based on a frame satisfying both $\mathcal{C}$ and X.

3.2.10. EXAMPLE. Most of the logics for common knowledge discussed by Halpern and Moses in [50] correspond to $\mathtt{ML}^*$ interpreted over some class of $\mathcal{C}\mathsf{X}$-frames. More precisely, all of the logics K_n^C, T_n^C, $\mathsf{S4}_n^C$ and $\mathsf{S5}_n^C$ are captured by our framework, and only the logic $\mathsf{KD45}_n^C$ is not.

3.3 Infinitary and cyclic hypersequent calculi

In this section we introduce families of infinitary and cyclic hypersequent calculi for $\mathtt{ML}^*$ interpreted over classes of $\mathcal{C}\mathsf{X}$-models. As mentioned in the introduction, our calculi will be extensions of the hypersequent calculi from [66] for basic modal logic. The extension will be twofold. First, we extend the system to cover multimodal logic. Second, we include the master modality $[*]$. This involves adding left and right rules for $[*]$, as well as allowing infinite branches. We also annotate formulas using a simple focus mechanism and manage these annotations both within the rules and by adding two structural rules, fc and fm. The annotations will facilitate the use of a path-based soundness condition.

3.3.1 Hypersequents and derivations

An *annotated formula* is a formula φ, together with an annotation indicating whether φ is *in focus* or *out of focus*. If φ is in focus, it is denoted by $\varphi^\bullet$ and, if not, by φ°. We use u, v, w as variables ranging over $\{\circ, \bullet\}$.

3.3.1. DEFINITION. A *sequent* is an ordered pair (Γ, Δ) of finite sets of annotated formulas, written as $\Gamma \Rightarrow \Delta$. A *hypersequent* is a finite set $\{\sigma_0, \ldots, \sigma_n\}$ of sequents, written as $\sigma_0 \mid \cdots \mid \sigma_n$.

The idea for considering hypersequents for modal logic, is that one can think of the different sequents within a hypersequent as different states of a model. By doing so, it becomes possible to reason about multiple states simultaneously. In contrast, mere sequents only facilitate reasoning about a single state at a time.

Because the language considered in this chapter is of very restricted expressivity compared to the whole modal μ-calculus, it will suffice to only consider hypersequents where the focus annotations are distributed in a certain specific way. We first define the following syntactic abbreviations, which will also be of use later in this chapter.

$$[\mathsf{D}]\varphi := \bigwedge_{a \in \mathsf{D}} [a]\varphi, \qquad [x]^n\varphi := \underbrace{[x]\cdots[x]}_{n\text{-times}}\varphi \text{ (for } x \in \mathsf{D} \text{ or } x = \mathsf{D}\text{, and } n \geq 0)$$

Let us use the term ML*-*trace-formula* to refer to a formula of the form $[a]^i[*]\varphi$, where $a \in \mathsf{D}$ and $i \in \{0,1\}$. More explicitly, an ML*-trace-formula is any formula of the form $[*]\varphi$ or of the form $[a][*]\varphi$ for some $a \in \mathsf{D}$.

3.3.2. DEFINITION. An ML*-*hypersequent* is a hypersequent H in which at most one formula is in focus, in which case this formula is an ML*-trace-formula occurring at the right-hand side of some sequent $\Gamma \Rightarrow \Delta$ of H.

We adopt the convention of using shorthand notation for singleton formulas and sequents. For instance, we take $\Gamma, \varphi^u \Rightarrow \psi^v, \Delta$ to mean $\Gamma \cup \{\varphi^u\} \Rightarrow \{\psi^v\} \cup \Delta$, and the hypersequent $H \cup \{\sigma\}$ may be written as $H \mid \sigma$.

3.3.3. EXAMPLE. In the table below, the two hypersequents on the left are ML*-hypersequents, while those on the right are not.

$$[a]q^\circ \Rightarrow [a][*]q^\circ \mid p \wedge q^\circ \Rightarrow [*]p^\bullet, p^\circ, \qquad [a]q^\circ \Rightarrow [a][*]q^\bullet \mid p \wedge q^\circ \Rightarrow [*]p^\bullet, p^\circ,$$
$$[a]q^\circ \Rightarrow [a][*]q^\circ \mid p \wedge q^\circ \Rightarrow [*]p^\circ, p^\circ, \qquad [a]q^\circ \Rightarrow [a][*]q^\circ \mid p \wedge q^\circ \Rightarrow [*]p^\circ, p^\bullet,$$

For the rest of this chapter we will simply use the term hypersequents to refer to ML*-hypersequents.

We define the following operations on sets of annotated formulas, respectively taking all formulas out of focus, and stripping the focus of formulas entirely.

$$\Gamma^\circ := \{\varphi^\circ : \varphi^u \in \Gamma \text{ for some } u \in \{\circ, \bullet\}\},$$
$$\Gamma^- := \{\varphi \ : \varphi^u \in \Gamma \text{ for some } u \in \{\circ, \bullet\}\},$$

We extend these operations to sequents componentwise and to hypersequents sequent-wise. More precisely, we define:

$$(\Gamma \Rightarrow \Delta)^\circ := \Gamma^\circ \Rightarrow \Delta^\circ,$$
$$H^\circ := \{\sigma^\circ \mid \sigma \in H\},$$

and likewise for σ^- and H^-.

The interpretation of (hyper)sequents in Kripke models is defined as follows.

3.3.4. DEFINITION. Let $\mathbb{S}$ be a Kripke model. Then:

- A sequent $\Gamma \Rightarrow \Delta$ is said to be *satisfied* at a state s of $\mathbb{S}$ whenever:

  If $s \Vdash \varphi$ for all $\varphi \in \Gamma^-$, then $s \Vdash \psi$ for some $\psi \in \Delta^-$.

- A sequent is *valid* in $\mathbb{S}$ if it is satisfied at every state of $\mathbb{S}$.

- A hypersequent H is *valid* in $\mathbb{S}$ if there is a $\sigma \in H$ which is valid in $\mathbb{S}$.

A hypersequent valid in all $\mathcal{C}\mathsf{X}$-models will be called $\mathcal{C}\mathsf{X}$-*valid*.

3.3.5. REMARK. Note that the focus annotations play no role in Definition 3.3.4. They will become important later when defining the soundness conditions for non-well-founded derivations.

3.3.6. EXAMPLE. The hypersequent $\Rightarrow p, p \to \bot$ is valid in all models, but the hypersequent $\Rightarrow p \mid \Rightarrow p \to \bot$ is not.

For Γ a set of annotated formulas and a an action from D, we define the following two operations.

$$[a]\Gamma := \{[a]\varphi^u : \varphi^u \in \Gamma\} \qquad [a]^{-1}\Gamma := \{\varphi^u : [a]\varphi^u \in \Gamma\}.$$

Consequently, we have $[a][a]^{-1}\Gamma = \{[a]\varphi^u : [a]\varphi^u \in \Gamma\}$.

We are now ready to define our four basic hypersequent calculi. They are obtained from the four calculi in [66] by making the following adaptations:

- The modal rules are parametrised in an action $a \in \mathsf{D}$ in order to cover multimodal logic.
- The rules $[*]_L$ and $[*]_R$ are added.
- Annotation-management is added to all the rules.
- The structural rules fc of *focus change*, and fm of *focus merge*, are added.

Before we give the full definition, we will briefly give some intuition behind the modal rules of Figure 3.2. When read upside down, the idea is that they jump from a state to one of its a-successors. To see this, suppose that a state s of some model $\mathbb{S}$ falsifies $[a]\varphi$ (note that $[a]\varphi^u$ appears in the conclusion of every modal rule). Then s has an a-successor t that falsifies φ. But then for every set Γ such that s satisfies everything in $[a]\Gamma$, it holds that t satisfies everything in Γ. This explains the modal rule $[a]_{\mathsf{K}}$.

Now suppose that $\mathbb{S}$ is transitive. Then t will even satisfy everything in $[a]\Gamma$, explaining the premiss of the rule $[a]_{\mathsf{K4}}$. If $\mathbb{S}$ is symmetric, then we have that $tR_a s$. It follows that t falsifies everything in $[a]\Delta$ for every set of formulas Δ in which everything is falsified by s. Finally, if $\mathbb{S}$ is both transitive and symmetric, we can say all of the above, and even a little bit more. Indeed, if $[a]\psi$ is some formula falsified by s, we claim that t falsifies $[a]\psi$. To see this, suppose that w is some a-successor of s falsifying ψ. By symmetry we have $tR_a s$, whence by transitivity $tR_a w$, as required. This shows where the $[a][a]^{-1}\Delta^\circ$ comes from in the premiss of $[a]_{\mathsf{B4}}$.

Note that in the rules B and $\mathsf{B4}$ all formulas from Δ are taken out of focus. The reason is that keeping them in focus would make proving soundness more challenging, even though it is not necessary for completeness. We leave the following question for future work.

3.3.7. Question. Are the calculi HB^* and $\mathsf{HB4}^*$ still sound when one keeps in focus all the formulas in Δ in the conclusion of the modal rule?

We are now ready to define four basic hypersequent calculi, one for each basic frame condition.

3.3.8. Definition. The hypersequent calculus HX^* consists of all rules of Figure 3.1, together with the modal rule $[a]_\mathsf{X}$ from Figure 3.2.

Observe that all formulas on the left-hand side of a sequent in some rule of HX^* are out of focus. The reason for this is that all hypersequents are assumed to be $\mathtt{ML}^*$-hypersequents. Note, moreover, that in the systems HX^* there is no interaction between the different sequents within a hypersequent. In other words, the basic calculi HX^* for $\mathsf{X} \in \{\mathsf{K}, \mathsf{K4}, \mathsf{B}, \mathsf{B4}\}$ do not yet use the additional expressivity offered by the hypersequent framework and could be formulated as ordinary sequent calculi.

Following [66], we augment HX^* with rules corresponding to simple frame conditions. Figure 3.3 depicts these rules in their most general forms. For each simple frame condition $C = (C_R, C_=)$, and each basic frame condition X, we have a rule r^X_C. We discuss some examples, initially assuming that $\mathsf{X} = \mathsf{K}$.

Consider the condition $D = \{(\{1,2\}, \emptyset)\}$ of *directedness*. Its corresponding rule is:

$$\mathsf{r}^\mathsf{K}_D \frac{H \mid \Gamma'_1, \Gamma'_2 \Rightarrow}{H \mid [a]\Gamma'_1, \Gamma_1 \Rightarrow \Delta_1, \Delta'_1 \mid [a]\Gamma'_2, \Gamma_2 \Rightarrow \Delta_2, \Delta'_2}$$

Suppose that the conclusion is invalid in some directed model $\mathbb{S}$. Then every sequent in the conclusion is refuted by some state of $\mathbb{S}$. Hence in particular, for each $i \in \{1,2\}$, there is a state s_i in $\mathbb{S}$ such that $s_i \Vdash [a]\Gamma'_i$. By directedness, there is a state u such that $s_i R_a u$ for each i. Hence $u \Vdash \Gamma'_i$ for each i, showing that the premiss is invalid as well.

For another example, consider the condition $L = \{(\{1\},\{2\}), (\{2\},\{1\})\}$ of *linearity*. This gives the rule

$$\mathsf{r}^\mathsf{K}_L \frac{H \mid \Gamma'_1, \Gamma_2 \Rightarrow \Delta_2 \qquad H \mid \Gamma'_2, \Gamma_1 \Rightarrow \Delta_1}{H \mid [a]\Gamma'_1, \Gamma_1 \Rightarrow \Delta_1, \Delta'_1 \mid [a]\Gamma'_2, \Gamma_2 \Rightarrow \Delta_2, \Delta'_2}$$

Let us again suppose that the conclusion is invalid, but now in some linear model. Then there are states s_1 and s_2 such that $s_i \not\Vdash [a]\Gamma'_i, \Gamma_i \Rightarrow \Delta_i$, where $i \in \{1,2\}$. Suppose, without loss of generality, that $s_1 R_a s_2$. Then $s_2 \not\Vdash \Gamma'_1, \Gamma_2 \Rightarrow \Delta_2$, again showing that one of the premisses is invalid.

In the above two examples we have shown that the rules r^K_D and r^K_L are sound. We will see later that the same ideas generalise to show that all rules r^X_C are sound. Observe that the difference between the rules r^K_C and r^X_C closely resembles the difference between $[a]_\mathsf{K}$ and $[a]_\mathsf{X}$. The intuition for completeness will become more clear once we have introduced our canonical models in Section 3.5.

We are finally ready to define the calculi that will be the main topic of study in this chapter.

$$\mathsf{id}\ \frac{}{\varphi^\circ \Rightarrow \varphi^v} \qquad \bot\ \frac{}{\bot^\circ \Rightarrow}$$

$$\mathsf{iw}_L\ \frac{H \mid \Gamma \Rightarrow \Delta}{H \mid \Gamma, \varphi^\circ \Rightarrow \Delta} \qquad \mathsf{iw}_R\ \frac{H \mid \Gamma \Rightarrow \Delta}{H \mid \Gamma \Rightarrow \varphi^u, \Delta} \qquad \mathsf{ew}\ \frac{H}{H \mid \Gamma \Rightarrow \Delta}$$

$$\to_L\ \frac{H \mid \Gamma, \psi^\circ \Rightarrow \Delta \qquad H \mid \Gamma \Rightarrow \varphi^\circ, \Delta}{H \mid \Gamma, \varphi \to \psi^\circ \Rightarrow \Delta} \qquad \to_R\ \frac{H \mid \Gamma, \varphi^\circ \Rightarrow \psi^\circ, \Delta}{H \mid \Gamma \Rightarrow \varphi \to \psi^\circ, \Delta}$$

$$[*]_L\ \frac{H \mid \Gamma, \{\varphi^\circ, [a][*]\varphi^\circ : a \in \mathsf{D}\} \Rightarrow \Delta}{H \mid \Gamma, [*]\varphi^\circ \Rightarrow \Delta}$$

$$[*]_R\ \frac{H \mid \Gamma \Rightarrow \varphi^\circ, \Delta \qquad \{H \mid \Gamma \Rightarrow [a][*]\varphi^u, \Delta : a \in \mathsf{D}\}}{H \mid \Gamma \Rightarrow [*]\varphi^u, \Delta}$$

$$\mathsf{fc}\ \frac{H \mid \Gamma \Rightarrow [*]\varphi^v, \Delta}{H \mid \Gamma \Rightarrow [*]\varphi^u, \Delta} \qquad \mathsf{fm}\ \frac{H \mid \Gamma \Rightarrow \varphi^\bullet, \varphi^\circ, \Delta}{H \mid \Gamma \Rightarrow \varphi^\bullet, \Delta}$$

$$\mathsf{cut}\ \frac{H \mid \Gamma_1, \varphi^\circ \Rightarrow \Delta_1 \qquad H \mid \Gamma_2 \Rightarrow \varphi^\circ, \Delta_2}{H \mid \Gamma_1, \Gamma_2 \Rightarrow \Delta_1, \Delta_2}$$

Figure 3.1: The *local* rules.

$$[a]_{\mathsf{K}}\ \frac{H \mid \Gamma \Rightarrow \varphi^u}{H \mid [a]\Gamma \Rightarrow [a]\varphi^u} \qquad [a]_{\mathsf{K4}}\ \frac{H \mid \Gamma, [a]\Gamma \Rightarrow \varphi^u}{H \mid [a]\Gamma \Rightarrow [a]\varphi^u}$$

$$[a]_{\mathsf{B}}\ \frac{H \mid \Gamma \Rightarrow \varphi^u, [a]\Delta^\circ}{H \mid [a]\Gamma \Rightarrow [a]\varphi^u, \Delta} \qquad [a]_{\mathsf{B4}}\ \frac{H \mid \Gamma, [a]\Gamma \Rightarrow \varphi^u, [a]\Delta^\circ, [a][a]^{-1}\Delta^\circ}{H \mid [a]\Gamma \Rightarrow [a]\varphi^u, \Delta}$$

Figure 3.2: The *modal* rules

$$\mathsf{r}^{\mathsf{X}}_{C}\ \frac{\{H \mid \sigma^{\mathsf{X}}_{(C_R,C_=)} : (C_R, C_=) \in C\}}{H \mid [a]\Gamma'_1, \Gamma_1 \Rightarrow \Delta_1, \Delta'_1 \mid \cdots \mid [a]\Gamma'_n, \Gamma_n \Rightarrow \Delta_n, \Delta'_n}$$

where

$$\sigma^{\mathsf{K}}_{(I,J)} = \bigcup_{i\in I}\Gamma'_i, \bigcup_{j\in J}\Gamma_j \Rightarrow \bigcup_{j\in J}\Delta_j$$

$$\sigma^{\mathsf{K4}}_{(I,J)} = \bigcup_{i\in I}\Gamma'_i, [a]\bigcup_{i\in I}\Gamma'_i, \bigcup_{j\in J}\Gamma_j \Rightarrow \bigcup_{j\in J}\Delta_j$$

$$\sigma^{\mathsf{B}}_{(I,J)} = \bigcup_{i\in I}\Gamma'_i, \bigcup_{j\in J}\Gamma_j \Rightarrow \bigcup_{j\in J}\Delta_j, [a]\bigcup_{i\in I}(\Delta'_i)^\circ$$

$$\sigma^{\mathsf{B4}}_{(I,J)} = \bigcup_{i\in I}\Gamma'_i, [a]\bigcup_{i\in I}\Gamma'_i, \bigcup_{j\in J}\Gamma_j \Rightarrow \bigcup_{j\in J}\Delta_j, [a]\bigcup_{i\in I}(\Delta'_i)^\circ, [a][a]^{-1}\bigcup_{i\in I}(\Delta'_i)^\circ$$

Figure 3.3: The *frame condition* rules for some n-simple frame condition C.

3.3.9. Definition. Given a set $\mathcal{C}$ of simple frame conditions, we let $\mathsf{HX}^* + \mathsf{R}_{\mathcal{C}}$ be the system HX^*, augmented with the rules $\mathsf{r}^{\mathsf{X}}_{C}$ for each $C \in \mathcal{C}$.

It will be convenient to have a notion of active and inactive sequents and annotated formulas in some rule application of $\mathsf{HX}^* + \mathsf{R}_{\mathcal{C}}$. First, we will call the sequents outside of the context H *active*. In case of the local rules, the *active*active!formula annotated formulas of an active sequent are those that are mentioned individually, *i.e.* not as part of some set Γ or Δ (or Γ_i or Δ_i in the case of cut). In the case of the other rules, we call all annotated formulas of an active sequent *active*.

All other formulas and sequents are called *inactive*. Note that due to the fact that (hyper)sequents are based on sets rather than multisets, the context H might also contain active sequents. In the same way, the contexts Γ and Δ of an active sequent in a local rule might contain active annotated formulas. In the case of $\mathsf{r}^{\mathsf{X}}_{C}$, the i-th active sequent in the conclusion is said to have *index* i and the premiss corresponding to $(C_R, C_=) \in C$ is said to have *index* $(C_R, C_=)$. Here the fact that hypersequents are sets means that a single sequent might have multiple indices.

The following facts about rule applications of $\mathsf{HX}^* + \mathsf{R}_{\mathcal{C}}$ will be useful later on.

3.3.10. Fact. For any rule applications of $\mathsf{HX}^* + \mathsf{R}_{\mathcal{C}}$:

- If some annotated formula is active, it belongs to an active sequent.

- In applications of $\to_L$, $\to_R$, $[*]_R$, all conclusions and all premisses have precisely one active annotated formula.
- All premisses and conclusions of local and modal rules have precisely one active sequent, except for the premisses of id, $\bot$, and ew, which have none.

3.3.11. Definition. An $\mathsf{HX}^* + \mathsf{R}_{\mathcal{C}}$*-derivation* is a (possibly infinite) tree generated from the rules of $\mathsf{HX}^* + \mathsf{R}_{\mathcal{C}}$. Its root is also called its *conclusion*.

A derivation of which every leaf is an axiom is called *closed*. Other derivations are called *open*. For any $\mathsf{HX}^* + \mathsf{R}_{\mathcal{C}}$-derivation π with root H, we say that π is a $\mathsf{HX}^* + \mathsf{R}_{\mathcal{C}}$-derivation *of* H.

Just like in Chapter 2, we would like the calculi $\mathsf{HX}^* + \mathsf{R}_{\mathcal{C}}$ to satisfy the *closure property*. We explicitly define the notion of closure for $\mathtt{ML}^*$.

3.3.12. Definition. The *(Fischer-Ladner) closure* of a set Φ of formulas is the least $\Psi \supseteq \Phi$ such that:

(i) If $\varphi \to \psi \in \Psi$, then $\varphi, \psi \in \Psi$;

(ii) If $[a]\varphi \in \Psi$, then $\varphi \in \Psi$;

(iii) If $[*]\varphi \in \Psi$, then $\varphi \in \Psi$, and $[a][*]\varphi \in \Psi$ for every $a \in \mathsf{D}$.

We write $\mathsf{FL}(\Phi)$ for the closure of Φ. It is easy to see that FL is a closure operator and that the closure of any finite set of formulas is finite. A set Φ such that $\mathsf{FL}(\Phi) = \Phi$ will be called *closed*. The closure $\mathsf{FL}(H)$ of a hypersequent H is defined as the closure of the set all formulas occurring in H, *i.e.* all formulas φ such that φ^u occurs in some sequent $\sigma \in H$ for some $u \in \{\circ, \bullet\}$.

3.3.13. Definition. An $\mathsf{HX}^* + \mathsf{R}_{\mathcal{C}}$-derivation π is said to be *analytic* if every formula occurring in π belongs to the closure of its conclusion.

The following lemma can be verified by direct inspection of the rules.

3.3.14. Lemma. *For* $\mathsf{Y} \in \{\mathsf{K}, \mathsf{K4}\}$*, any cut-free* $\mathsf{HY}^* + \mathsf{R}_{\mathcal{C}}$*-derivation is analytic.*

3.3.2 Infinitary proofs

It is not hard to show that $\mathsf{HX}^* + \mathsf{R}_{\mathcal{C}}$-derivations need not be sound. In fact, already when $\mathsf{X} = \mathsf{K}$, the set $\mathcal{C}$ is empty, and $\mathsf{D} = \{a\}$ is a singleton, infinite derivations exist of invalid (singleton) hypersequents, as demonstrated by the following example.

$$\dfrac{\dfrac{\dfrac{\dfrac{}{\bot^\circ \Rightarrow}\,\mathsf{id}}{\bot^\circ, [a][*](p \to \bot)^\circ \Rightarrow}\,\mathsf{iw}_L \qquad \dfrac{\dfrac{\dfrac{\vdots}{[*]([a]p \to \bot)^\circ \Rightarrow}}{[*]([a]p \to \bot)^\circ \Rightarrow p^\circ}\,\mathsf{iw}_R}{[a][*]([a]p \to \bot)^\circ \Rightarrow [a]p^\circ}\,[a]_{\mathsf{K}}}{[a]p \to \bot^\circ, [a][*](p \to \bot)^\circ \Rightarrow}\,\to_L}{[*]([a]p \to \bot)^\circ \Rightarrow}\,[*]_L$$

We therefore need a way to distinguish valid from invalid derivations. To this end, we will introduce a path-based soundness condition (cf. Section 2.3.6). We first need to following auxiliary definition. Note that Fact 3.3.10 allows use to speak of *the* active formula.

3.3.15. Definition. Consider an instance

$$[*]_R \frac{H_1 \qquad \cdots \qquad H_n}{H}$$

of $[*]_R$. A conclusion-premiss pair (H, H_k) is said to be *a focused unfolding* if the active formula is in focus both in H and H_k.

3.3.16. Example. In the following rule application each premiss, except for the leftmost one, forms a focused unfolding with the conclusion.

$$[*]_R \frac{p^\circ \Rightarrow q^\circ \qquad \{p^\circ \Rightarrow q^\circ \mid p^\circ \Rightarrow [a][*]q^\bullet : a \in \mathsf{D}\}}{p^\circ \Rightarrow q^\circ \mid p^\circ \Rightarrow [*]q^\bullet}$$

Also note that the sequent of the leftmost premiss is an example of one that is at the same time both active and inactive.

We are now ready to define which derivations will be called proofs.

3.3.17. Definition. A closed $\mathsf{HX}^* + \mathsf{R}_\mathcal{C}$-derivation is an $\mathsf{HX}^* + \mathsf{R}_\mathcal{C}$-*proof* if every infinite branch β has a *good* final segment γ. That is, the rule fc is applied nowhere on γ and for infinitely many n it holds that $(\gamma(n), \gamma(n+1))$ is a focused unfolding.

The following proposition is often explicitly included in the path-based soundness condition for a proof system with focus annotations. This not needed in our case, as it is implied by our condition.

3.3.18. Proposition. *Suppose some branch β of an $\mathsf{HX}^* + \mathsf{R}_\mathcal{C}$-proof π has a good final segment γ. Then every hypersequent on γ has a formula in focus.*

Proof:
Suppose towards a contradiction that some $\gamma(n)$ does not have a formula in focus. Since the rule fc is not applied in γ, direct inspection of the rules yields that $\gamma(m)$ does not have a formula in focus for all $m \geq n$. But this contradicts the fact that a focused unfolding happens infinitely often on γ. □

3.3.19. Remark. It is not hard to see that every infinite branch of an $\mathsf{HX}^* + \mathsf{R}_\mathcal{C}$-proof in fact contains a good *trail*. Indeed, if in some conclusion-premiss pair both hypersequents have a formula in focus, direct inspection of the rules shows that the focused formula in the premiss is a *direct ancestor* of that in the conclusion (analogous to Definition 2.3.9). Hence, the formulas in focus on some final segment of a branch induce a trail, which is ensured to be a ν-trail by the requirement that a focused unfolding happens infinitely often. This fact will be exploited in the soundness proof of Section 3.4.

We close this section by introducing some notation that will improve the readability of the rest of this chapter.

3.3.20. DEFINITION. Let π be an $\mathsf{HX}^*+\mathsf{R}_\mathcal{C}$-derivation for some $\mathsf{X} \in \{\mathsf{K}, \mathsf{K4}, \mathsf{B}, \mathsf{B4}\}$ and set $\mathcal{C}$ of simple frame conditions. We say that:

- π is a $\mathcal{C}\mathsf{X}$*-proof* if π is an $\mathsf{HX}^* + \mathsf{R}_\mathcal{C}$-proof.
- π is a $\mathcal{C}\mathsf{X}^{cf}$*-proof* if π is an $\mathsf{HX}^* + \mathsf{R}_\mathcal{C}$-proof containing no applications of cut.
- π is a $\mathcal{C}\mathsf{X}^{an}$*-proof* if π is an analytic $\mathsf{HX}^* + \mathsf{R}_\mathcal{C}$-proof.

The notions of $\mathcal{C}\mathsf{X}$-provability, $\mathcal{C}\mathsf{X}^{cf}$-provability, and $\mathcal{C}\mathsf{X}^{an}$-provability of a hypersequent H are defined analogously.

We sometimes drop the $\mathcal{C}\mathsf{X}$-, $\mathcal{C}\mathsf{X}^{cf}$-, or $\mathcal{C}\mathsf{X}^{an}$- from notions concerning provability, whenever it is clear from the context which is meant.

3.3.21. EXAMPLE. Below we give a few examples of non-well-founded proofs. For readability we assume that $\mathsf{D} = \{a\}$, *i.e.* that there is only a single action. It is not hard to see how to generalise this to a larger set of actions. Also for the sake of readability, we drop the annotation $\circ$ from the formulas which are not in focus.

The following is a K^{cf}-proof of the *induction axiom*.

$$
\dfrac{
  \mathsf{iw}_L\,\dfrac{\mathsf{id}\,\dfrac{}{p \Rightarrow p}}{p, [*](p \to [a]p) \Rightarrow p}
  \qquad
  [*]_L\,\dfrac{
    \to_L\,\dfrac{
      \mathsf{iw}_L\,\dfrac{[a]_\mathsf{K}\,\dfrac{\dfrac{\vdots}{p, [*](p \to [a]p) \Rightarrow [*]p^\bullet}}{[a]p, [a][*](p \to [a]p) \Rightarrow [a][*]p^\bullet}}{p, [a]p, [a][*](p \to [a]p) \Rightarrow [a][*]p^\bullet}
      \qquad
      \mathsf{iw}_L\,\dfrac{\mathsf{iw}_R\,\dfrac{\mathsf{id}\,\dfrac{}{p \Rightarrow p}}{p \Rightarrow p, [a][*]p^\bullet}}{p, [a][*](p \to [a]p) \Rightarrow p, [a][*]p^\bullet}
    }{p, p \to [a]p, [a][*](p \to [a]p) \Rightarrow [a][*]p^\bullet}
  }{p, [*](p \to [a]p) \Rightarrow [a][*]p^\bullet}
}{p, [*](p \to [a]p) \Rightarrow [*]p^\bullet}\,[*]_R
$$

Let $U = \{(\{1\}, \{2\})\}$, *i.e.* U is the equable frame condition of *universality*. The following is an $U\mathsf{K}^{cf}$-proof of the fact that, on universal frames, the modalities $[*]$ and $[a]$ amount to the same thing.

$$
\mathsf{iw}_L\,\dfrac{\mathsf{fc}\,\dfrac{
  \dfrac{
    \mathsf{r}^\mathsf{K}_U\,\dfrac{\mathsf{id}\,\dfrac{}{p \Rightarrow p}}{[a]p \Rightarrow [*]p \mid \Rightarrow p}
    \qquad
    \dfrac{\dfrac{\vdots}{[a]p \Rightarrow [*]p \mid \Rightarrow [*]p^\bullet}}{[a]p \Rightarrow [*]p \mid \Rightarrow [a][*]p^\bullet}\,[a]_\mathsf{K}
  }{[a]p \Rightarrow [*]p \mid \Rightarrow [*]p^\bullet}\,[*]_R
}{[a]p \Rightarrow [*]p \mid \Rightarrow [*]p}}{[a]p \Rightarrow [*]p}
$$

The rule fc is often needed when there are multiple repeating leaves. For an example, let $D' = \{(\emptyset, \{1, 2\})\}$, *i.e.* D' is *degenerateness*. The following is a proof that in such models either $[*]p$ or $[*]\neg p$ is valid.

$$
{\scriptstyle [*]_R}\dfrac{{\scriptstyle \mathsf{fc}}\dfrac{{\scriptstyle [*]_{\mathsf{R}}}\dfrac{{\scriptstyle \mathsf{r}^{\mathsf{K}}_{D'}}\dfrac{{\scriptstyle \to_R}\dfrac{{\scriptstyle \mathsf{iw}_R}\dfrac{{\scriptstyle \mathsf{id}}\dfrac{}{p \Rightarrow p}}{p \Rightarrow p, \bot}}{\Rightarrow p, p \to \bot}}{\Rightarrow p \mid \Rightarrow p \to \bot} \qquad \dfrac{\dfrac{\vdots}{\Rightarrow p \mid \Rightarrow [*](p \to \bot)^\bullet}}{\Rightarrow p \mid \Rightarrow [a][*](p \to \bot)^\bullet}{\scriptstyle [a]_{\mathsf{K}}}}{\Rightarrow p \mid \Rightarrow [*](p \to \bot)^\bullet}}{\Rightarrow p \mid \Rightarrow [*](p \to \bot)} \qquad \dfrac{\dfrac{\vdots}{\Rightarrow [*]p^\bullet \mid \Rightarrow [*](p \to \bot)}}{\Rightarrow [a][*]p^\bullet \mid \Rightarrow [*](p \to \bot)}{\scriptstyle [\mathsf{a}]_{\mathsf{K}}}}{\Rightarrow [*]p^\bullet \mid \Rightarrow [*](p \to \bot)}
$$

In the examples above, the vertical dots indicate that the proof continues as it did when the same hypersequent appeared lower in the proof tree. In the following section we will make this presentation formal, by introducing cyclic proofs.

The following lemma shows that, as with validity, the focus annotations do not matter for the provability of some hypersequent H. Its proof is immediate by the presence of the rule fc.

3.3.22. Lemma. *For every hypersequent H it holds that H is provable iff H° is.*

3.3.3 Cyclic proofs

In this section we will define cyclic $\mathcal{C}\mathsf{X}$-proofs and show that they induce the same notion of provability as the (infinitary) $\mathcal{C}\mathsf{X}$-proofs of the previous section. Aside from facilitating a path-based soundness condition on infinite branches, annotations also make it easier to define cyclic proofs. Recall that finite trees with back edges were defined in Definition 2.3.34.

3.3.23. Definition. A *cyclic $\mathcal{C}\mathsf{X}$-derivation* is a finite tree with back edges (π, f) such that π is a $\mathcal{C}\mathsf{X}$-derivation, $\mathsf{dom}(f)$ consists only of leaves which are not axioms, and the hypersequent of each node $l \in \mathsf{dom}(f)$ is the same as that of the node $f(l)$.

Recall that a node $l \in \mathsf{dom}(f)$ is said to be a *repeating leaf*. The node $f(l)$ is its *companion*. Note that distinct repeating leaves may have the same companion. A cyclic derivation (π, f) is called *closed* if every non-axiom leaf belongs to $\mathsf{dom}(f)$ and *open* otherwise.

3.3.24. Definition. A *cyclic $\mathcal{C}\mathsf{X}$-proof* is a closed cyclic $\mathcal{C}\mathsf{X}$-derivation such that the path $[f(l), l]$ between each repeating leaf and its companion (inclusive) is *good*. That is, on this path the rule fc is not applied, and at least one focused unfolding happens.

3.3.25. Remark. A key feature of cyclic $\mathcal{C}\mathsf{X}$-proofs is that the soundness condition can be verified by only checking the paths between repeating leaves and their companions. This is possible by virtue of the focus annotations.

A naive approach, without focus annotations, would be to merely demand the existence of a good trace on each path $[f(l), l]$ where l is a repeating leaf.

This, however, is not sound. To get an idea for why this is the case, suppose, for instance, that there are two distinct repeating leaves l_1 and l_2 such that $f(l_1) = f(l_2)$. Then it might be possible that the paths $[f(l_1), l_1]$ and $[f(l_2), l_2]$ each contain a good trace, but that their concatenation does not.

Another approach would simply call a cyclic derivation a cyclic *proof* whenever its unravelling is a proof. This condition is clearly sound, but more laborious to verify.

3.3.26. Example. It is not hard to see that all proofs in Example 3.3.21 can be turned into cyclic proofs by assigning the appropriate back edges.

We obtain a proposition similar to Proposition 3.3.18.

3.3.27. Proposition. *Let l be a repeating leaf of some cyclic proof (π, f). Then every hypersequent between l and its companion $f(l)$ has a formula in focus.*

Proof:
Let v be a node on the path $[f(l), l]$ such that v is the conclusion of the focused unfolding between $f(l)$ and l. Since v has a formula in focus and fc is not applied between $f(l)$ and v, it follows that $f(l)$ must have a formula in focus. But then l has a formula in focus, and thus so must the whole path. □

The next two propositions show that exactly the same hypersequents are provable by analytic cyclic proofs and by analytic infinitary proofs. The proofs of these propositions are postponed to the Intermezzo following the present chapter, where they will be proven in a more general setting.

3.3.28. Proposition. *For every $\mathcal{C}\mathsf{X}^{an}$-proof π there is a cyclic $\mathcal{C}\mathsf{X}^{an}$-proof (π', f) with the same conclusion.*

Proof:
This is a special case of Proposition I.2.12. □

The converse also holds, *i.e.* each cyclic proof induces a proof. In fact, this proof can be obtained by unravelling.

3.3.29. Proposition. *The unravelling of a cyclic $\mathcal{C}\mathsf{X}$-proof is itself a $\mathcal{C}\mathsf{X}$-proof.*

Proof:
This is a special case of Proposition I.2.22. □

Throughout the rest of this chapter we shall work only with infinitary proofs, as they are more convenient for proving soundness and completeness. On the other hand, the advantage of cyclic proofs is that they are finite objects. This makes them more suitable for computational manipulations, such as the translation into Hilbert-style proofs, and extracting interpolants from proofs.

3.4 Soundness

This section is devoted to proving the following soundness theorem.

3.4.1. THEOREM. *Every $\mathcal{C}\mathsf{X}$-provable hypersequent is $\mathcal{C}\mathsf{X}$-valid.*

Our proof will go by contraposition and infinite descent. More precisely, assuming that some hypersequent is invalid, witnessed by a countermodel $\mathbb{S}$, we will suppose towards a contradiction that it nevertheless has a proof π. By the soundness of each individual rule (which we shall establish shortly), it then follows that π must have an infinite branch. We will then reach our contradiction by showing that this leads to the infinite decrease of some well-founded measure.

This measure, which can be seen as a very special case of the notion of *signature* in [101], is provided by the following definition. Recall the notion of $\mathtt{ML}^*$-trace-formula that was defined directly before Definition 3.3.2.

3.4.2. DEFINITION. Let $\varphi = [a]^i[*]\psi$ be an $\mathtt{ML}^*$-trace-formula formula and let $\mathbb{S}$ be a Kripke model. If s is a state of $\mathbb{S}$ such that $\mathbb{S}, s \nVdash \varphi$, we define the *signature* of φ at s as:

$$\mathsf{sig}_s(\varphi) := \min\{n \in \omega : \mathbb{S}, s \nVdash [a]^i[\mathsf{D}]^n\psi\}.$$

Note that, when it is defined, the signature of $[*]\psi$ at s is precisely the length of the shortest path from s to a state t refuting ψ.

3.4.3. DEFINITION. Let H be a hypersequent and let $\mathbb{S} = (S, R, V)$ be a Kripke model. A *countermodel state assignment (cmsa)* of H in $\mathbb{S}$ is a function $\alpha : H \to S$ assigning to each sequent σ of H a state $\alpha(\sigma)$ of $\mathbb{S}$ in which σ is not satisfied.

The notion of a cmsa allows us to express the soundness of a rule in the following manner: a rule r is sound, whenever the existence of a cmsa for the conclusion of an application of r implies the existence of a cmsa for one of its premisses.

We first show the soundness of the frame condition rules, proving a slightly stronger statement that we will later use to obtain an infinitely decreasing sequence of signatures. Note that in the rule application depicted below, each sequent σ_i refers to the i-th active sequent in the conclusion, and each sequent $\sigma_{(C_R,C_=)}$ is the single active sequent of the premiss with index $(C_R, C_=) \in C$.

3.4.4. LEMMA. *Let $\mathsf{X} \in \{\mathsf{K}, \mathsf{K4}, \mathsf{B}, \mathsf{B4}\}$ and let C be an n-simple frame condition. Then the rule $\mathsf{r}^{\mathsf{X}}_C$ is sound on all $C\mathsf{X}$-frames. In fact, given a rule application*

$$\mathsf{r}^{\mathsf{X}}_C \, \frac{\{H \mid \sigma_{(C_R,C_=)} : (C_R, C_=) \in C\}}{H \mid \sigma_1 \mid \cdots \mid \sigma_n}$$

if α is a cmsa for the conclusion in some $C\mathsf{X}$-model $\mathbb{S}$, then there is some $(C_R, C_=) \in C$, for which there is a cmsa α' of $H \mid \sigma_{(C_R,C_=)}$ in $\mathbb{S}$ such that α and α' agree on $H \setminus \sigma_{(C_R,C_=)}$ and moreover $\alpha'(\sigma_{(C_R,C_=)}) = \alpha(\sigma_j)$ for each $j \in C_=$.

Proof:
Suppose that α is a cmsa for the conclusion in some $C\mathsf{X}$-model $\mathbb{S}$. Since $\mathbb{S}$ is a C-model, it has a state s such that for some $(C_R, C_=) \in C$ it holds that $\alpha(\sigma_i)R_a s$ for each $i \in C_R$, and $\alpha(\sigma_j) = s$ for each $j \in C_=$. We define the following function α' on the premiss $H \mid \sigma_{(C_R,C_=)}$:

$$\alpha'(\sigma) := \begin{cases} s & \text{if } \sigma = \sigma_{(C_R,C_=)}, \\ \alpha(\sigma) & \text{otherwise.} \end{cases}$$

We claim that α' is a cmsa. It clearly suffices to show that $\mathbb{S}$ does not satisfy $\sigma_{(C_R,C_=)}$ at s. We treat the four options for X one-by-one. Note that the form of the rule $\mathsf{r}^{\mathsf{X}}_C$ dictates that each σ_i is of the form $[a]\Gamma'_i, \Gamma_i \Rightarrow \Delta_i, \Delta'_i$.

(K) In this case $\sigma_{(C_R,C=)}$ is of the form:

$$\bigcup_{i \in C_R} \Gamma'_i, \bigcup_{j \in C_=} \Gamma_j \Rightarrow \bigcup_{j \in C_=} \Delta_j.$$

Since for each $i \in C_R$, we have that $\alpha(\sigma_i) \Vdash [a]\Gamma'_i$, it follows $s \Vdash \Gamma'_i$. Moreover, for each $j \in C_=$, we have $s = \alpha(\sigma_j) \nVdash \sigma_j$. It follows that $s \nVdash \sigma_{(C_R,C=)}$, as required.

(K4) We have that $\sigma_{(C_R,C=)}$ is of the following form.

$$\bigcup_{i \in C_R} \Gamma'_i, [a] \bigcup_{i \in C_R} \Gamma'_i, \bigcup_{j \in C_=} \Gamma_j \Rightarrow \bigcup_{j \in C_=} \Delta_j.$$

Note that it suffices to show that $s \Vdash [a] \bigcup_{i \in C_R} \Gamma'_i$, as the rest of $\sigma_{(C_R,C=)}$ is already covered by the previous case. To that end, suppose that φ belongs to Γ'_i for some $i \in C_R$. For each t such that $sR_a t$, we have, by transitivity, that $\alpha(\sigma_i)R_a t$. Thus, since $\alpha(\sigma_i) \Vdash [a]\Gamma'_i$, it follows that $t \Vdash \varphi$, as required.

(B) When $\mathsf{X} = \mathsf{B}$, we have that $\sigma_{(C_R,C=)}$ is of the form:

$$\bigcup_{i \in C_R} \Gamma'_i, \bigcup_{j \in C_=} \Gamma_j \Rightarrow \bigcup_{j \in C_=} \Delta_j, [a] \bigcup_{i \in C_R} (\Delta'_i)^\circ.$$

Now, it suffices to show that s falsifies everything in $[a] \bigcup_{i \in C_R} (\Delta'_i)^\circ$, as the rest of $\sigma_{(C_R,C=)}$ is again covered by the case where $\mathsf{X} = \mathsf{K}$. So suppose that $\varphi^u \in \Delta'_i$ for some $i \in C_R$. Then $\alpha(\sigma_i) \nVdash \varphi$. But since $\mathbb{S}$ is symmetric and $\alpha(\sigma_i)R_a s$, we have $sR_a\alpha(\sigma_i)$, whence $s \nVdash [a]\varphi$.

(B4) For this final case, we have that $\sigma_{(C_R,C=)}$ is of the form

$$\bigcup_{i \in C_R} \Gamma'_i, [a] \bigcup_{i \in C_R} \Gamma'_i, \bigcup_{j \in C_=} \Gamma_j \Rightarrow \bigcup_{j \in C_=} \Delta_j, [a] \bigcup_{i \in C_R} (\Delta'_i)^\circ, [a][a]^{-1} \bigcup_{i \in C_R} (\Delta'_i)^\circ$$

Every part of $\sigma_{(C_R,C=)}$, except for $[a][a]^{-1}\bigcup_{i\in C_R}(\Delta_i')^\circ$, is already covered by the cases (K4) and (B). So suppose that $[a]\varphi^\circ \in \Delta_i'$ for some $i \in C_R$. Then $\alpha(\sigma_i) \nVdash [a]\varphi$, so there is some state t in $\mathbb{S}$ such that $\alpha(\sigma_i)R_a t$ and $t \nVdash \varphi$. Since by symmetry $sR_a\alpha(\sigma_i)$, we have by transitivity that $sR_a t$, whence $s \nVdash [a]\varphi$, as desired.

This finishes the proof. □

The following proposition shows that every individual rule is sound. We again prove something slightly stronger.

3.4.5. Proposition. *Every rule application*

$$\mathsf{r}\,\frac{H_1 \cdots H_k}{H}$$

of $\mathsf{HX}^* + \mathsf{R}_\mathcal{C}$ *is sound. In fact, if α is a cmsa in some $\mathcal{C}\mathsf{X}$-model $\mathbb{S}$ for H, then there is a premiss H_k and a cmsa α_k in $\mathbb{S}$, such that α and α_k agree on every sequent σ that occurs only inactively in H_k.*

Proof:
For each frame condition rule r_C^X, we can use the (stronger) statement of Proposition 3.4.4. For the other rules, we begin by choosing the premiss H_k and defining α_k only on the active sequent of H_k (if it exists). We make a case distinction on the rule r.

- For the irrefutable axioms id and $\bot$ the proposition holds vacuously.
- If $\mathsf{r} = \{\mathsf{ew}\}$ there is only a single premiss, which contains no active sequents.
- For any other local rule (*i.e.* from Figure 3.1), there is exactly one active sequent σ in H, and exactly on active sequent σ_k in each premiss H_k. We claim that $\mathbb{S}, \alpha(\sigma) \nVdash \sigma_k$ for each least one k. As example we show this for $\mathsf{r} = [*]_R$, leaving the other cases to the reader.

  Since $\alpha(\sigma) \nVdash [*]\varphi$, we have $\alpha(\sigma) \nVdash \varphi$ or $\alpha(\sigma) \nVdash [a][*]\varphi$ for some $a \in \mathsf{D}$. Hence $\alpha(\sigma) \nVdash \sigma_k$ for some appropriate k.

  We then pick H_k and set $\alpha_k(\sigma_k) := \alpha(\sigma)$.
- For any modal rule $[a]_\mathsf{X}$, there is only one single choice for H_1. Moreover, in each case the conclusion H has one active sequent, say σ, and the premiss H_1 has one active sequent, say σ_1. We claim that, as $\mathbb{S}, \alpha(\sigma) \nVdash \sigma$, there is a state s of $\mathbb{S}$ such that $\alpha(\sigma)R_a s$ and $\mathbb{S}, s \nVdash \sigma_1$. We will only show this for $\mathsf{X} = \mathsf{B4}$, leaving the other cases to the reader.

  First, since $\alpha(\sigma) \nVdash [a]\varphi$, there is some state $\alpha(\sigma)R_a t$ such that $t \nVdash \varphi$. Second, for each $[a]\psi \in [a]\Gamma$ we have $\alpha(\sigma) \Vdash [a]\psi$, and thus $t \Vdash \psi$. Moreover, by transitivity it holds that $t \Vdash [a]\psi$. Now, suppose that $\psi \in \Delta$. Then, by symmetry, we have that $t \nVdash [a]\psi$. Finally, if $[a]\psi \in \Delta$, then there is a state w such that $\alpha(\sigma)R_a w$ and $w \nVdash \psi$. It follows by symmetry and transitivity that $tR_a w$ and thus $t \nVdash [a]\psi$, as required.

We set $\alpha_1(\sigma_1) = s$.

Now let σ be a sequent that occurs in H_k, but only inactively. Then σ must also occur in H. We simply set $\alpha_k(\sigma) := \alpha(\sigma)$. □

Since every rule is sound, an easy inductive argument shows that every *well-founded* proof is sound. For extending this result to non-well-founded proofs, the following lemma is crucial.

3.4.6. Lemma. *Suppose the following is a rule application in* $\mathsf{HX}^* + \mathsf{R}_{\mathcal{C}}$

$$\mathsf{r}\,\frac{H_1 \cdots H_k}{H}$$

such that H has a cmsa α in some $\mathcal{C}\mathsf{X}$-model $\mathbb{S}$. Then there is a premiss H_k and cmsa α_k for H_k in $\mathbb{S}$, such that if $\varphi_k^\bullet \in \sigma_k \in H_k$ and $\varphi^\bullet \in \sigma \in H$, then

$$\mathsf{sig}_{\alpha_k(\sigma_k)}(\varphi_k) \leq \mathsf{sig}_{\alpha(\sigma)}(\varphi).$$

Moreover, if $\mathsf{r} = [*]_R$ and both $\varphi^\bullet$ and $\varphi_k^\bullet$ are active, this inequality is strict.

Proof:
Let H_k and α_k be as given by Proposition 3.4.5. If either H or H_k does not have a formula in focus, the result holds vacuously. So suppose there is $\varphi^\bullet \in \sigma \in H$ and $\varphi_k^\bullet \in \sigma_k \in H_k$.

Suppose first that $\varphi = \varphi_k$. We claim that $\alpha_k(\sigma_k) = \alpha(\sigma)$, from which it will follows that $\mathsf{sig}_{\alpha_k(\sigma_k)}(\varphi_k) = \mathsf{sig}_{\alpha(\sigma)}(\varphi)$.

If σ_k occurs only inactively in H_k, we must have $\sigma_k = \sigma$ and directly obtain from Proposition 3.4.5 that $\alpha_k(\sigma_k) = \alpha(\sigma)$. Let us therefore consider the case where σ_k is active. Then, as the only sequent with a formula in focus, σ must also be active in H. If r is local, it follows that $\alpha_k(\sigma_k) = \alpha(\sigma)$. Note that, as $\varphi^\bullet = \varphi_k^\bullet$ belongs to an active sequent in both H and H_k, the rule r cannot be modal. Finally, if r is a frame condition rule $\mathsf{r}_C^{\mathsf{X}}$, suppose that H_k is the premiss corresponding to $(C_R, C_=) \in C$. Then σ must be the sequent in the conclusion with index i for some $i \in C_=$. Hence again $\alpha_k(\sigma) = \alpha(\sigma)$.

Now suppose that $\varphi \neq \varphi_k$. Then both σ and σ_k must be active, and direct inspection of the rules yields that r is either $[*]_R$ or modal. Moreover, because the premiss can only contain a single formula in focus, we know that $\varphi^\bullet$ occurs neither in an inactive sequent of H, nor inactively in σ. We will now forget about the H_k and α_k given by Proposition 3.4.5 at the beginning of this proof, choosing our premiss with slightly more care.

Suppose that r is $[*]_R$. Since $\varphi^\bullet$ is active, it must be of the form $[*]\psi^\bullet$. Let $n := \mathsf{sig}_{\alpha(\sigma)}(\varphi)$. By definition $\alpha(\sigma) \not\Vdash [\mathsf{D}]^n\psi$. If $n = 0$, we let H_k be the leftmost premiss. If $n > 0$, then for some $a \in \mathsf{D}$ we have $\alpha(\sigma) \not\Vdash [a][\mathsf{D}]^{n-1}\psi$ and we let H_k be a premiss corresponding to this $a \in \mathsf{D}$. We set:

$$\alpha_k(\tau) := \begin{cases} \alpha(\sigma) & \text{if } \tau \text{ is active,} \\ \alpha(\tau) & \text{otherwise.} \end{cases}$$

We leave it to the reader to verify that α_k is a cmsa for H_k. If $n = 0$, this suffices to prove the theorem, because H_k has no formula in focus. If $n > 0$, then $[a][*]\psi^\bullet$ belongs to the active sequent σ_k of H_k, and we have

$$\mathsf{sig}_{\alpha_k(\sigma_k)}([a][*]\psi) = \mathsf{sig}_{\alpha(\sigma)}([a][*]\psi) = \mathsf{sig}_{\alpha(\sigma)}(\varphi) - 1 < \mathsf{sig}_{\alpha(\sigma)}(\varphi),$$

as required.

Finally, if r is $[a]_{\mathsf{X}}$, there is just a single premiss, say H_k. Since by assumption H_k has a formula in focus, unequal to $\varphi^\bullet$, we know that $\varphi^\bullet$ must be the principal formula in H. Hence φ is of the form $[a][*]\psi$. Again, we define $n := \mathsf{sig}_{\alpha(\sigma)}(\varphi)$. By definition we have $\alpha(\sigma) \not\Vdash [a][\mathsf{D}]^n\psi$. Let t be some state of $\mathbb{S}$ such that $\alpha(\sigma)R_a^{\mathsf{X}}t$ and $t \not\Vdash [\mathsf{D}]^n\psi$. We define:

$$\alpha_k(\tau) := \begin{cases} t & \text{if } \tau \text{ is active,} \\ \alpha(\tau) & \text{otherwise.} \end{cases}$$

We again leave it to the reader to verify that α_k is a cmsa for H_k. Note that H_k has $[*]\psi$ in focus, and

$$\mathsf{sig}_{\alpha_k(\sigma_k)}([*]\psi) = \mathsf{sig}_t([*]\psi) = n = \mathsf{sig}_{\alpha(\sigma)}(\varphi),$$

as required. □

With this is place, we are ready to prove the soundness theorem.

Proof of Theorem 3.4.1. Suppose, towards a contradiction, that some $\mathcal{C}\mathsf{X}$-provable hypersequent H is $\mathcal{C}\mathsf{X}$-invalid. Then there is a cmsa α of H in some model $\mathbb{S}$. Repeatedly applying Lemma 3.4.6, we obtain a branch $H = H_0 \cdot H_1 \cdots$ in the proof of H, with for each H_i a cmsa α_i of H_i in $\mathbb{S}$. Note that this branch must be infinite, for otherwise the final H_i is an axiom, contradicting the fact that it has a cmsa. Moreover, by the condition of infinite branches, there is some final segment, say $H_k \cdot H_{k+1} \cdots$, on which every hypersequent has a formula in focus (by Proposition 3.3.18) , and a focused unfolding happens infinitely often. But then, letting σ_i be the sequent in H_i containing a formula in focus, we have, by Lemma 3.4.6,

$$\mathsf{sig}_{\alpha_k(\sigma_k)}(\varphi_k) \geq \mathsf{sig}_{\alpha_{k+1}(\sigma_{k+1})}(\varphi_{k+1}) \geq \mathsf{sig}_{\alpha_{k+2}(\sigma_{k+2})}(\varphi_{k+2}) \geq \dots$$

where this inequality is strict infinitely often, contradicting the well-foundedness of ω.

3.5 Completeness

In this section we are concerned with the completeness of the systems $\mathsf{HX}^* + \mathsf{R}_\mathcal{C}$. The scope of our completeness result, *i.e.* whether or not we need analytic applications of the rule cut, depends on the frame conditions. The following is the precise statement that we will prove.

3.5.1. THEOREM. *Let H be $\mathcal{C}\mathsf{X}$-valid. Then H is $\mathcal{C}\mathsf{X}^{an}$-provable. If, in addition, $\mathsf{X} \in \{\mathsf{K}, \mathsf{K4}\}$ and $\mathcal{C}$ consists of equable frame conditions, then H is $\mathcal{C}\mathsf{X}^{cf}$-provable.*

Our proof strategy is essentially a Henkin construction, as is widely applied in first-order and modal logic. In Section 3.5.1, we will show how to construct a canonical X-model for a given hypersequent H. The states of this canonical model will be precisely the sequents in H. If the hypersequent H is sufficiently nice, its canonical model will satisfy the *Truth Lemma*: every sequent σ in H, regarded as a state of the canonical model of H, falsifies the sequent σ. Hence, the canonical model of H will be a countermodel for H itself.

In Section 3.5.2 we will define a certain maximality property for H which guarantees that H is sufficiently nice in above sense. Section 3.5.3 shows that any unprovable hypersequent can be extended to satisfy this maximality property. The rest of the chapter is concerned with showing the Truth Lemma, from which Theorem 3.5.1 will easily follow.

For technical reasons it will be convenient to prove Theorem 3.5.1 only for hypersequents that have no formula in focus. This clearly suffices, because H is valid iff H° is, and H is provable iff H° is. We will call such hypersequents *focus free.*

3.5.1 Canonical models

One interesting feature of hypersequents, as mentioned in the introduction of this chapter, is their ability to function as the carrier of a canonical model. The following definition is essentially taken from [66].

3.5.2. DEFINITION. Let H be a hypersequent. The *canonical* X-*model* $\mathbb{S}^{\mathsf{X}}_H$ of H is defined as follows:

- The set of states is H.
- For each X, the accessibility relation R^{X}_a is given by:
  - $(\Gamma_1 \Rightarrow \Delta_1) R^{\mathsf{K}}_a (\Gamma_2 \Rightarrow \Delta_2)$ iff $[a]^{-1}\Gamma_1 \subseteq \Gamma_2$;
  - $(\Gamma_1 \Rightarrow \Delta_1) R^{\mathsf{K4}}_a (\Gamma_2 \Rightarrow \Delta_2)$ iff $[a]^{-1}\Gamma_1 \subseteq \Gamma_2$ and $[a][a]^{-1}\Gamma_1 \subseteq \Gamma_2$;
  - $\sigma_1 R^{\mathsf{B}}_a \sigma_2$ iff $\sigma_1 R^{\mathsf{K}}_a \sigma_2$ and $\sigma_2 R^{\mathsf{K}}_a \sigma_1$.
  - $\sigma_1 R^{\mathsf{B4}}_a \sigma_2$ iff $\sigma_1 R^{\mathsf{K4}}_a \sigma_2$ and $\sigma_2 R^{\mathsf{K4}}_a \sigma_1$.

- The valuation function is given by $V(p) := \{\Gamma \Rightarrow \Delta \mid p^\circ \in \Gamma\}$.

3.5.3. REMARK. Note that the definition of R_a^{X} does not depend on the right-hand sides of the respective sequents. It might therefore be possible to instead define the states of the canonical model of H to consist only of the left-hand sides of sequents of H. As we will see later, in the presence of analytic applications of the cut rule, our notion of maximality implies that the left-hand side of some sequent in a maximal hypersequent determines the right-hand side (and vice versa). In that case this alternative definition of the canonical model would therefore be equivalent.

The following proposition almost follows by definition.

3.5.4. PROPOSITION. *The canonical* X*-model of* H *is indeed an* X*-model.*

Proof:
The case $\mathsf{X} = \mathsf{K}$ is clear. For $\mathsf{X} = \mathsf{K4}$, suppose that $\sigma_1 R_a^{\mathsf{K4}} \sigma_2 R_a^{\mathsf{K4}} \sigma_3$. We claim that $\sigma_1 R_a^{\mathsf{K4}} \sigma_3$. Indeed, writing $\Gamma_i \Rightarrow \Delta_i$ for each σ_i, suppose that $[a]\varphi^\circ$ belongs to Γ_1. Then $[a]\varphi^\circ$ belongs to Γ_2 and thus both φ° and $[a]\varphi^\circ$ belong to Γ_3. Both R_a^{B} and R_a^{B4} are by definition symmetric, and R_a^{B4} inherits its transitivity from R_a^{K4}. $\square$

3.5.2 $\mathcal{C}\mathsf{X}^i$-maximality

In this section we define our notion of maximality for hypersquents, namely $\mathcal{C}\mathsf{X}^i$-maximality, where $i \in \{\mathit{cf}, \mathit{an}\}$. The notion of $\mathcal{C}\mathsf{X}^{an}$-maximality, which is tailored to the availability of analytic applications of cut, is similar to that used in [66]. The notion of $\mathcal{C}\mathsf{X}^{cf}$-maximality, used for cut-free completeness, is substantially different and newly developed for handling the master modality in the absence of the cut rule. More details on the difference between the two notions of maximality will be given in Remark 3.5.18.

The following order on hypersequents also features in [66]. It is useful for comparing the (logical) strength of two hypersequents.

3.5.5. DEFINITION. Let $\Gamma_1 \Rightarrow \Delta_1$ and $\Gamma_2 \Rightarrow \Delta_2$ be sequents and let H_1, H_2 be hypersequents. We define:

- $(\Gamma_1 \Rightarrow \Delta_1) \sqsubseteq (\Gamma_2 \Rightarrow \Delta_2)$ if $\Gamma_1 \subseteq \Gamma_2$ and $\Delta_1 \subseteq \Delta_2$.

- $H_1 \sqsubseteq H_2$ if for every $\sigma_1 \in H_1$, there is some $\sigma_2 \in H_2$ such that $\sigma_1 \sqsubseteq \sigma_2$.

If $H \sqsubseteq K$, we say that H is *encompassed* by K.

Note that $\sqsubseteq$ is a preorder on the set of hypersequents. We will often use $\sigma \sqsubseteq H$ as a shorthand for $\{\sigma\} \sqsubseteq H$.

In the following it will be useful to restrict attention to hypersequents containing only formulas from a given finite set.

3.5.6. Definition. Let Σ be a finite closed set of formulas. A (hyper)sequent is said to be a Σ-*(hyper)sequent* if it contains only formulas from Σ.

We can now reformulate the notion of analyticity (Definition 3.3.13) as follows: a proof π of H is *analytic* if it only contains $\mathsf{FL}(H)$-hypersequents.

The following definition captures a first requirement that will feature in the definition of a $\mathcal{C}\mathsf{X}^i$-maximal hypersequent. Its main application will be in the inductive proof of the Truth Lemma.

3.5.7. Definition. A sequent $\Gamma \Rightarrow \Delta$ is said to be *propositionally saturated* if the following closure conditions hold:

(i) If $\varphi_1 \to \varphi_2 \in \Gamma^-$, then $\varphi_2 \in \Gamma^-$ or $\varphi_1 \in \Delta^-$.
(ii) If $\varphi_1 \to \varphi_2 \in \Delta^-$, then $\varphi_1 \in \Gamma^-$ and $\varphi_2 \in \Delta^-$.
(iii) If $[*]\varphi \in \Gamma^-$, then $\varphi \in \Gamma^-$ and $[a][*]\varphi \in \Gamma^-$ for every $a \in \mathsf{D}$.
(iv) If $[*]\varphi \in \Delta^-$, then $\varphi \in \Delta^-$ or $[a][*]\varphi \in \Delta^-$ for some $a \in \mathsf{D}$.

A hypersequent is *propositionally saturated* whenever each of its sequents is.

We also certainly want the canonical model of a $\mathcal{C}\mathsf{X}^i$-maximal hypersequent to be a $\mathcal{C}$-model. This is captured by the following definition.

3.5.8. Definition. Let $\mathcal{C}$ be a set of simple frame conditions and suppose that $\mathsf{X} \in \{\mathsf{K}, \mathsf{K4}, \mathsf{B}, \mathsf{B4}\}$. A hypersequent H is $\mathcal{C}\mathsf{X}$-*structured* if $\mathbb{S}^{\mathsf{X}}_H$ is a $\mathcal{C}$-model.

In the presence of cut, we can even require the following.

3.5.9. Definition. Let Σ be a finite and closed set of formulas. A sequent $\Gamma \Rightarrow \Delta$ is said to be Σ-*complete* if for every $\varphi \in \Sigma$ it holds that $\varphi \in \Gamma^-$ or $\varphi \in \Delta^-$. A hypersequent H is *complete* if every sequent in H is $\mathsf{FL}(H)$-complete.

Summing up, we now have the following saturation conditions that we want our $\mathcal{C}\mathsf{X}^i$-maximal hypersequents to satisfy.

3.5.10. Definition. A hypersequent H is called $\mathcal{C}\mathsf{X}^{cf}$-*saturated* if it is both $\mathcal{C}\mathsf{X}$-structured and propositionally saturated. If H, in addition, is complete, we say that H is $\mathcal{C}\mathsf{X}^{an}$-*saturated*.

The following property, also featuring in [66], will be very useful in the completeness proof. As mentioned in the introduction, the reason for working with focus-free sequents is technical convenience.

3.5.11. Definition. A hypersequent H is said to be $\mathcal{C}\mathsf{X}^i$-*full* with respect to a sequent σ if either $\sigma \sqsubseteq H$ or $H \mid \sigma$ is $\mathcal{C}\mathsf{X}^i$-provable. We say that H is $\mathcal{C}\mathsf{X}^i$-*full* if it is $\mathcal{C}\mathsf{X}^i$-full with respect to every focus-free $\mathsf{FL}(H)$-sequent σ.

It is not hard to show that fullness with respect to a given sequent is preserved by taking $\sqsubseteq$-extensions.

3.5.12. Lemma. *If H is $\mathcal{C}\mathsf{X}^i$-full with respect to σ, then so is every $H' \sqsupseteq H$.*

Proof:
If $\sigma \sqsubseteq H$, then by the transitivity of $\sqsubseteq$, also $\sigma \sqsubseteq H'$. If $H \mid \sigma$ is $\mathcal{C}\mathsf{X}^i$-provable, then, by the presence of iw_L, iw_R and ew, the same holds for $H' \mid \sigma$. □

Finally, we are ready to give our definition of $\mathcal{C}\mathsf{X}^i$-maximality. In addition to the preorder $\sqsubseteq$, this definition also uses the subset order $\subseteq$ on hypersequents.

3.5.13. Definition. A hypersequent H is called $\mathcal{C}\mathsf{X}^i$*-maximal* if H is:

(i) focus free, (ii) $\mathcal{C}\mathsf{X}^i$-unprovable, (iii) $\mathcal{C}\mathsf{X}^i$-full, (iv) $\mathcal{C}\mathsf{X}^i$-saturated,

and H is $\subseteq$-maximal as an $\mathsf{FL}(H)$-hypersequent satisfying conditions (i) - (iv).

In the proceeding, we will often drop the $\mathcal{C}\mathsf{X}^i$ from (un)provability, saturation, fullness or maximality, whenever it is clear from the context. In the next section, we will show that every unprovable and focus-free hypersequent has a maximal $\sqsubseteq$-extension. The following proposition shows that we will not have to worry about the $\subseteq$-maximality.

3.5.14. Lemma. *Suppose a hypersequent satisfies conditions (i) - (iv) of Definition 3.5.13. Then it can be $\subseteq$-extended to be $\mathcal{C}\mathsf{X}^i$-maximal.*

Proof:
Suppose that H satisfies conditions (i) - (iv) of Definition 3.5.13. Because $\mathsf{FL}(H)$ is finite, the set

$$\{H' : H \subseteq H' \text{ and } H' \text{ is an } \mathsf{FL}(H)\text{-sequent satisfying conditions (i) - (iv)}\}.$$

is finite, and we can simply let $\overline{H}$ be a $\subseteq$-maximal element of this set. □

3.5.3 The Extension Lemma

In this section we will show that every focus-free, unprovable hypersequent has a maximal $\sqsubseteq$-extension. Throughout the section we shall refer to the conditions (i), (ii), (iii) and (iv) of Definition 3.5.13 without explicitly mentioning this definition. The notion of *strong fullness* will be defined below and serves as an intermediate step in the extension into a maximal hypersequent,

Our extension procedure consists of a series of extensions, some of which depend on whether analytic cuts are available or not. Below is a diagrammatic representation of the entire argument. All arrows in this diagram preserve all of the earlier properties, except for the dotted arrow, which need not preserve strong fullness. For instance, after applying Lemma 3.5.17 to a hypersequent satisfying properties (i), (ii), (iii), we will obtain a strongly full $\sqsubseteq$-extension, which also still satisfies properties (i), (ii), (iii). Similarly, after applying Lemma 3.5.20 to a strongly full hypersequent, it remains strongly full, whence susceptible to Lemma 3.5.22. In contrast, the maximal hypersequent we end up with at the end of the procedure will in general *not* be strongly full. Note that the final step in the procedure is simply Lemma 3.5.14, which we have already established above.

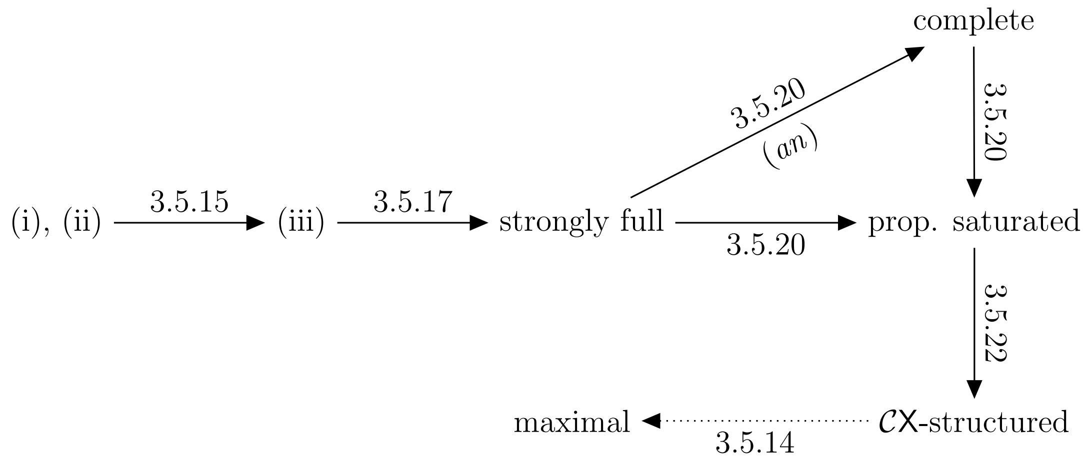

The first step is a standard Lindenbaum construction. Note that the following lemma is agnostic to whether $i = cf$ or $i = an$.

3.5.15. LEMMA. *For any hypersequent H satisfying conditions (i) and (ii), there exists an $\mathsf{FL}(H)$-hypersequent $\overline{H} \supseteq H$ satisfying conditions (i), (ii) and (iii).*

Proof:
Let $\sigma_1, \ldots, \sigma_n$ be an enumeration of all focus-free $\mathsf{FL}(H)$-sequents. Beginning with $H_0 := H$, we recursively define:

$$H_{k+1} := \begin{cases} H_k \mid \sigma_{k+1} & \text{if } H_k \mid \sigma_{k+1} \text{ is unprovable,} \\ H_k & \text{otherwise.} \end{cases}$$

Let $\overline{H} := H_n$. Clearly $\overline{H}$ satisfies condition (i) and (ii), since H_k is unprovable for any $0 \leq k \leq n$. Moreover, for each k we have $H_k \subseteq H_{k+1}$, whence $H \subseteq \overline{H}$. Finally, if σ_k is an $\mathsf{FL}(H)$-sequent such that $\sigma_k \not\sqsubseteq \overline{H}$, then certainly $\sigma_k \notin H_k$, so $H_{k-1} \mid \sigma_k$ is provable and, by the presence of ew, also $\overline{H} \mid \sigma$ is provable, as required. □

As mentioned above, we will go through hypersequents satisfying the following strong variant of fullness as an intermediate step.

3.5.16. DEFINITION. A hypersequent H is said to be *strongly $\mathcal{C}\mathsf{X}^i$-full*, if it is $\mathcal{C}\mathsf{X}^i$-full and for every $\Gamma \Rightarrow \Delta \in H$ and $\varphi \in \mathsf{FL}(H)$ it holds that:

(i) If $\varphi^\circ \notin \Gamma$, then $H \mid \Gamma, \varphi^\circ \Rightarrow \Delta$ is $\mathcal{C}\mathsf{X}^i$-provable.

(ii) If $\varphi^\circ \notin \Delta$, then $H \mid \Gamma \Rightarrow \varphi^\circ, \Delta$ is $\mathcal{C}\mathsf{X}^i$-provable.

The following lemma also appears in [66].

3.5.17. LEMMA. *Let H be a hypersequent satisfying conditions (i), (ii) and (iii). Then there is an $\mathsf{FL}(H)$-hypersequent $H' \sqsupseteq H$ that satisfies conditions (i), (ii) and (iii) and moreover is strongly full.*

Proof:
We proceed by induction on the number of sequents in H that do not satisfy the conditions of Definition 3.5.16. In the base case of our induction this number is zero and therefore H is indeed strongly $\mathcal{C}\mathsf{X}^i$-full.

For the induction step, let $\Gamma \Rightarrow \Delta$ be a sequent in H that does not satisfy the conditions of Definition 3.5.16. Write H_* for $H \setminus \{\Gamma \Rightarrow \Delta\}$ and fix an enumeration $\varphi_1, \ldots, \varphi_n$ of $\mathsf{FL}(H)$. We set $\Gamma_0 \Rightarrow \Delta_0 := \Gamma \Rightarrow \Delta$ and inductively define

$$\Gamma_{k+1} \Rightarrow \Delta_{k+1} := \begin{cases} \Gamma_k, \varphi^\circ_{k+1} \Rightarrow \Delta_k & \text{if } H_* \mid \Gamma_k, \varphi^\circ_{k+1} \Rightarrow \Delta_k \text{ is unprovable,} \\ \Gamma_k \Rightarrow \varphi^\circ_{k+1}, \Delta_k & \text{if } H_* \mid \Gamma_k \Rightarrow \varphi^\circ_{k+1}, \Delta_k \text{ is unprovable,} \\ \Gamma_k \Rightarrow \Delta_k & \text{otherwise,} \end{cases}$$

where we make an arbitrary choice if the first two cases both hold.

Let $\overline{H}$ be $H_* \mid \Gamma_n \Rightarrow \Delta_n$. By construction $\Gamma_n \Rightarrow \Delta_n$ satisfies the two conditions of Definition 3.5.16, whence by the induction hypothesis there is an unprovable and strongly full H' such that $H \sqsubseteq \overline{H} \sqsubseteq H'$, as required. □

3.5.18. REMARK. In [66], Lahav directly builds canonical models from strongly full hypersequents. Our new notion of maximality is needed for the inductive case of [$*$] in the proof of the Truth Lemma. Roughly, this proof requires us, for a given maximal hypersequent H and sequent $\Gamma \Rightarrow \Delta$ such that $H \mid \Gamma \Rightarrow \Delta$ is unprovable, to be able to extend $\Gamma \Rightarrow \Delta$ to some sequent $\Gamma' \Rightarrow \Delta' \in H$. This extension must not only preserve unprovability, but must in fact be obtained by directly applying the rules of $\mathsf{HX}^* + \mathsf{R}_{\mathcal{C}}$. In the presence of cut this is easy, because successive applications of this rule can always be used to obtain a sequent which is *complete*, *i.e.* where every formula occurs in either in the left-hand side or in the right-hand side. Without cut, however, we need to manually ensure that H has enough sequents to contain the required $\Gamma' \Rightarrow \Delta'$. For this, strong fullness does not suffice.

The key property of strongly full hypersequents, is that they are saturated, provided they also satisfy condition (i) and (ii). Before we prove this, we need the following lemma.

3.5.19. LEMMA. *Suppose H satisfies (i), (ii), and is strongly full. Then H forms a $\sqsubseteq$-antichain. That is, for each $\sigma_1, \sigma_2 \in H$: if $\sigma_1 \sqsubseteq \sigma_2$, then $\sigma_1 = \sigma_2$.*

Proof:
Write $\Gamma_i \Rightarrow \Delta_i$ for each σ_i and suppose that $\varphi^\circ \in \Gamma_2$. By the presence of the internal weakening rules, we have that $H \mid \Gamma_1, \varphi^\circ \Rightarrow \Delta_1$ is unprovable. Indeed, if we would have a proof, say π, we would be able to prove H in the following way.

$$\mathsf{iw}_L + \mathsf{iw}_R \frac{\dfrac{\overset{\pi}{H \mid \Gamma_1, \varphi^\circ \Rightarrow \Delta_1}}{\vdots}}{H \mid \Gamma_2 \Rightarrow \Delta_2}$$

From the unprovability of $H \mid \Gamma_1, \varphi^\circ \Rightarrow \Delta_1$ and the strong fullness of H, it follows that $\varphi^\circ \in \Gamma_1$. For $\varphi^u \in \Delta_2$ we can use an analogous argument, because, by condition (i), we know that $u = \circ$. □

3.5.20. LEMMA. *If H satisfies (i) and (ii) and is strongly $\mathcal{C}\mathsf{X}^i$-full, then H is propositionally saturated. Moreover, if $i =$ an then H is complete.*

Proof:
As all cases are similar, we will only treat the case where $[*]\varphi^u \in \Delta$ for some $\Gamma \Rightarrow \Delta \in H$. By condition (i) we have that $u = \circ$. We consider the following rule application

$$[*]_R \frac{H \mid \Gamma \Rightarrow \varphi^\circ, \Delta \qquad \{H \mid \Gamma \Rightarrow [a][*]\varphi^\circ, \Delta : a \in \mathsf{D}\}}{H \mid \Gamma \Rightarrow [*]\varphi^\circ, \Delta}$$

As the conclusion is equal to H, one of the premisses must be unprovable. Suppose first that the left premiss is unprovable. By the fullness of H, there is a $\sigma \in H$ such that $\Gamma \Rightarrow \varphi^\circ, \Delta \sqsubseteq \sigma$. But then it follows from Lemma 3.5.19 that $\sigma = \Gamma \Rightarrow \Delta$, whence $\varphi \in \Delta^-$. A similar argument can be used for the other premisses. □

Before we prove that the other part of $\mathcal{C}\mathsf{X}^i$-saturation also follows from strong $\mathcal{C}\mathsf{X}^i$-fullness, we first prove the following auxiliary lemma.

3.5.21. LEMMA. *Let H be a hypersequent. Given $\Gamma \Rightarrow \Delta$ and σ in H, we have*

(i) If $[a]^{-1}\Gamma \Rightarrow \varphi^\circ \sqsubseteq \sigma$, then $(\Gamma \Rightarrow \Delta) R_a^{\mathsf{K}} \sigma$.

(ii) If $[a]^{-1}\Gamma, [a][a]^{-1}\Gamma \Rightarrow \varphi^\circ \sqsubseteq \sigma$, then $(\Gamma \Rightarrow \Delta) R_a^{\mathsf{K4}} \sigma$.

If H is complete and $\mathcal{C}\mathsf{X}^i$-unprovable, then, moreover

(iii) If $[a]^{-1}\Gamma \Rightarrow \varphi^\circ, [a]\Delta_0^\circ \sqsubseteq \sigma$, then $(\Gamma \Rightarrow \Delta) R_a^{\mathsf{B}} \sigma$.

(iv) If $[a]^{-1}\Gamma, [a][a]^{-1}\Gamma \Rightarrow \varphi^\circ, [a]\Delta_0^\circ, [a][a]^{-1}\Delta_0^\circ \sqsubseteq \sigma$, then $(\Gamma \Rightarrow \Delta) R_a^{\mathsf{B4}} \sigma$.

where $\Delta_0 \subseteq \Delta$ consists of those $\psi^\circ \in \Delta$ such that $[a]\psi \in \mathsf{FL}(H)$.

Proof:
Write σ as $\Gamma_1 \Rightarrow \Delta_1$. Items (i) and (ii) follow directly from the definitions.

For item (iii), note that clearly $(\Gamma \Rightarrow \Delta)R_a^{\mathsf{K}}(\Gamma_1 \Rightarrow \Delta_1)$. For the converse, suppose that $[a]\psi^\circ \in \Gamma_1$. By unprovability, we have that $[a]\psi^\circ \notin \Delta_1$. Since $[a]\psi \in \mathsf{FL}(H)$, it follows that $\psi^\circ \notin \Delta$. The completeness of H gives $\psi^\circ \in \Gamma$, as required.

For item (iv) we reason similarly. It is clear that $(\Gamma \Rightarrow \Delta)R_a^{\mathsf{K4}}(\Gamma_1 \Rightarrow \Delta_1)$. For the converse, let $[a]\psi^\circ \in \Gamma_1$. By the same reasoning as before, we have $\psi^\circ \in \Gamma$. To see that also $[a]\psi^\circ \in \Gamma$, note that, since by unprovability $[a]\psi^\circ \notin \Delta_1$, we have $[a]\psi^\circ \notin \Delta_0$. Hence $[a]\psi^\circ \notin \Delta$ and therefore, by completeness, $[a]\psi^\circ \in \Gamma$. □

3.5.22. Lemma. *Suppose that H satisfies conditions (i) and (ii), and is strongly $\mathcal{C}\mathsf{X}^i$-full, where $\mathsf{X} \in \{\mathsf{B}, \mathsf{B4}\}$ only if $i = an$. Then H is also $\mathcal{C}\mathsf{X}$-structured.*

Proof:
Let $C \in \mathcal{C}$ be an n-simple frame condition. We must show that $\mathbb{S}_H^{\mathsf{X}}$ satisfies:

$$\forall s_1 \cdots s_n \exists u \bigvee_{(C_R, C_=) \in C} (\bigwedge_{i \in C_R} s_i R_a u \wedge \bigwedge_{j \in C_=} s_j = u).$$

To this end, let $\Gamma_1 \Rightarrow \Delta_1, \ldots, \Gamma_n \Rightarrow \Delta_n$ be states of $\mathbb{S}_H^{\mathsf{X}}$, or, in other words, elements of H. For $\mathsf{X} \in \{\mathsf{K}, \mathsf{B}\}$, consider, respectively, the following rule applications.

$$\mathsf{r}_C^{\mathsf{K}} \frac{\{H \mid \bigcup_{i \in C_R} [a]^{-1}\Gamma_i, \bigcup_{j \in C_=} \Gamma_j \Rightarrow \bigcup_{j \in C_=} \Delta_j : (C_R, C_=) \in C\}}{H \mid [a][a]^{-1}\Gamma_1, \Gamma_1 \Rightarrow \Delta_1 \mid \cdots \mid [a][a]^{-1}\Gamma_n, \Gamma_n \Rightarrow \Delta_n}$$

$$\mathsf{r}_C^{\mathsf{B}} \frac{\{H \mid \bigcup_{i \in C_R} [a]^{-1}\Gamma_1, \bigcup_{j \in C_=} \Gamma_j \Rightarrow \bigcup_{j \in C_=} \Delta_j, \bigcup_{i \in C_R} ([a]\Delta_i')^\circ : (C_R, C_=) \in C\}}{H \mid [a][a]^{-1}\Gamma_1, \Gamma_1 \Rightarrow \Delta_1, \Delta_1' \mid \cdots \mid [a][a]^{-1}\Gamma_n, \Gamma_n \Rightarrow \Delta_n, \Delta_n'}$$

where each Δ_i' consists of those $\psi^\circ \in \Delta_i$ such that $[a]\psi \in \mathsf{FL}(H)$.

Since the conclusion of each of the above two rules is H, they must both have at least one unprovable premiss. Let σ' be the active sequent of such a premiss. Because $H \mid \sigma'$ is unprovable, it follows from $\mathcal{C}\mathsf{X}^i$-fullness that $\sigma' \sqsubseteq \sigma$ for some $\sigma \in H$. Hence, by Lemma 3.5.21, we have $(\Gamma_i \Rightarrow \Delta_i)R_a^{\mathsf{X}}\sigma$ for each $i \in C_R$. Note that if $\mathsf{X} = \mathsf{B}$ we use the fact that H is complete by Lemma 3.5.20. Moreover, since for each $j \in C_=$ we have $\Gamma_j \Rightarrow \Delta_j \sqsubseteq \sigma$, Lemma 3.5.19 gives $\Gamma_j \Rightarrow \Delta_j = \sigma$, as required.

The cases where $\mathsf{X} \in \{\mathsf{K4}, \mathsf{B4}\}$ are similar. □

With this in place, we can now prove the main result of this section.

3.5.23. Lemma. *Any unprovable and focus-free hypersequent can be $\sqsubseteq$-extended to be maximal.*

Proof:
Suppose that H is $\mathcal{C}\mathsf{X}^i$-unprovable and focus free. First, we use Lemma 3.5.15 to extend H to $\overline{H} \supset H$ such that $\overline{H}$ is, in addition, $\mathcal{C}\mathsf{X}^i$-full. By Lemma 3.5.17 there is a $\mathcal{C}\mathsf{X}^i$-strongly full $H_1 \sqsupseteq \overline{H}$ satisfying conditions (i), (ii) and (iii), which by the lemmata 3.5.20 and 3.5.22 is $\mathcal{C}\mathsf{X}^i$-saturated. Finally, by Lemma 3.5.14, there is a $\mathcal{C}\mathsf{X}^i$-maximal $H_2 \sqsupseteq H_1 \sqsupseteq \overline{H} \sqsupseteq H$. □

3.5.4 The Existence Lemma for the basic modalities

In Section 3.5.6 we will prove the Truth Lemma by induction on formulas. The hardest clauses will be those where the main connective is a modality and the formula appears in the right-hand side some of sequent $\Gamma \Rightarrow \Delta \in H$. Say for instance, we have $[a]\psi \in \Delta$. We will have to show that the state $\Gamma \Rightarrow \Delta$ falsifies $[a]\psi$, and thus that there *exists* some a-successor falsifying ψ. It is for this reason that the following lemma is often called the Existence Lemma.

3.5.24. Lemma. *Let H be $\mathcal{C}\mathsf{X}^i$-maximal, with $\mathsf{X} \in \{\mathsf{B}, \mathsf{B4}\}$ only if $i = an$. Then for every $\Gamma \Rightarrow \Delta \in H$ with $[a]\varphi \in \Delta^-$ there is a sequent $\Gamma_1 \Rightarrow \Delta_1 \in H$ such that $(\Gamma \Rightarrow \Delta)R^{\mathsf{X}}_a(\Gamma_1 \Rightarrow \Delta_1)$ and $\varphi \in \Delta_1^-$.*

Proof:
Consider the following rule applications.

$$[\mathsf{a}]_{\mathsf{K}}\ \frac{H \mid [a]^{-1}\Gamma \Rightarrow \varphi^\circ}{H \mid [a][a]^{-1}\Gamma \Rightarrow [a]\varphi^\circ} \qquad [a]_{\mathsf{K4}}\ \frac{H \mid [a]^{-1}\Gamma, [a][a]^{-1}\Gamma \Rightarrow \varphi^\circ}{H \mid [a][a]^{-1}\Gamma \Rightarrow [a]\varphi^\circ}$$

$$[a]_{\mathsf{B}}\ \frac{H \mid [a]^{-1}\Gamma \Rightarrow \varphi^\circ, [a]\Delta_0^\circ}{H \mid [a][a]^{-1}\Gamma \Rightarrow [a]\varphi^\circ, \Delta_0} \quad [a]_{\mathsf{B4}}\ \frac{H \mid [a]^{-1}\Gamma, [a][a]^{-1}\Gamma \Rightarrow \varphi^\circ, [a]\Delta_0^\circ, [a][a]^{-1}\Delta_0^\circ}{H \mid [a][a]^{-1}\Gamma \Rightarrow [a]\varphi^\circ, \Delta_0}$$

where $\Delta_0 \subseteq \Delta$ consists of those $\psi^\circ \in \Delta$ such that $\psi \in \mathsf{FL}(H)$. By the presence of iw_L and iw_R, the conclusion of each $[a]_{\mathsf{X}}$ is $\mathcal{C}\mathsf{X}^i$-unprovable, whence so is the premiss. $\mathcal{C}\mathsf{X}^i$-fullnes gives a sequent $\sigma \in H$ which $\sqsubseteq$-extends the premiss. Finally, it follows by Lemma 3.5.21 that $(\Gamma \Rightarrow \Delta)R^{\mathsf{X}}_a\sigma$. □

3.5.5 The Existence Lemma for the master modality

In this section we prove another Existence Lemma, but this time for the master modality. This clause of the inductive proof of the Truth Lemma is shown by contradiction. We show that if some formula $[*]\varphi$ in Δ^- does not have a witness $\varphi \in \Delta_1^-$ in some $\Gamma_1 \Rightarrow \Delta_1$ reachable from $\Gamma \Rightarrow \Delta$, then there exists a proof of H. As we will see, the key to constructing this proof of H is that we can saturate sequents without losing focus. In the presence of cut this is easy, as captured by the following proposition.

3.5.25. Proposition. *Let H be a focus-free hypersequent and let σ a sequent such that $\sigma \notin H$. If σ is not $\mathsf{FL}(H \mid \sigma)$-complete, there is a rule application*

$$\mathsf{r}\,\frac{H \mid \sigma_1 \qquad \cdots \qquad H \mid \sigma_n}{H \mid \sigma}$$

such that for each σ_i it holds that $\sigma \sqsubset \sigma_i$.

Proving a similar saturation result without the presence of cut is considerably harder. We first show how to saturate sequents propositionally.

3.5.26. Proposition. *Let H be a focus-free hypersequent and let σ be a sequent such that $\sigma \notin H$. If σ is not propositionally saturated, there is a rule application*

$$\mathsf{r}\,\frac{H \mid \sigma_1 \qquad \cdots \qquad H \mid \sigma_n}{H \mid \sigma}$$

such that r is not cut and for each σ_i it holds that $\sigma \sqsubset \sigma_i$.

Proof:
Write σ as $\Gamma \Rightarrow \Delta$. Let r be a rule corresponding to a clause of Definition 3.5.7 of propositional saturation, which is not satisfied by $\Gamma \Rightarrow \Delta$. The idea is to apply r with $\Gamma \Rightarrow \Delta$ as active sequent, in such a way that all formulas in $\Gamma \Rightarrow \Delta$ are preserved (this is sometimes called a *cumulative* rule application).

Since each clause is very similar, we will only treat clause (iv). Suppose that $[*]\varphi \in \Delta^-$, but neither $\varphi \in \Delta^-$ not $[a][*]\varphi \in \Delta^-$. We make a case distinction on whether $[*]\varphi^\circ \in \Delta$. If not, then $[*]\varphi^\bullet \in \Delta$, and we take as our rule application

$$\mathsf{fm}\,\frac{H \mid \Gamma \Rightarrow [*]\varphi^\bullet, [*]\varphi^\circ, \Delta}{H \mid \Gamma \Rightarrow [*]\varphi^\bullet, \Delta}$$

This suffices, since $(\Gamma \Rightarrow [*]\varphi^\bullet, [*]\varphi^\circ) \sqsubset (\Gamma \Rightarrow [*]\varphi^\bullet, \Delta)$. If, on the other hand, it holds that $[*]\varphi^\circ \in \Delta$, we take the following rule application

$$[*]_R\,\frac{H \mid \Gamma \Rightarrow \varphi^\circ, \Delta \qquad \{H \mid \Gamma \Rightarrow [a][*]\varphi^\circ, \Delta : a \in \mathsf{D}\}}{H \mid \Gamma \Rightarrow [*]\varphi^\circ, \Delta}$$

Since neither φ nor $[a][*]\varphi$ belongs to Δ^- for any $a \in \mathsf{D}$, each premiss satisfies the required condition. □

For $\mathcal{C}\mathsf{X}$-structuredness, we require that $\mathcal{C}$ consists of only equable frame conditions.

3.5.27. Proposition. *Let C be an equable frame condition let $\mathsf{X} \in \{\mathsf{K}, \mathsf{K4}\}$. Suppose that H is $C\mathsf{X}$-structured and focus free. If σ is a sequent such that $\sigma^\circ \sqsubseteq H$ and $H \mid \sigma^\circ$ is not $C\mathsf{X}$-structured, there is a rule application of $\mathsf{HX}^* + \mathsf{R}_{\mathcal{C}}$*

$$\mathsf{r}\,\frac{H \mid \sigma_1 \qquad \cdots \qquad H \mid \sigma_m}{H \mid \sigma}$$

such that r is not cut and for each σ_i it holds that $\sigma \sqsubset \sigma_i$.

Proof:
We assume that $\mathsf{X} = \mathsf{K}$. The case where $\mathsf{X} = \mathsf{K4}$ is similar. Let n be such that C is n-simple. Recall that the equability of C means that there is some fixed $U \subseteq \{1, \dots n\}$ such that for every $(C_R, C_=) \in C$ it holds that $C_= = U$. We will therefore simply write $(C_R, U) \in C$ for an arbitrary element in C.

Since $H \mid \sigma^\circ$ is not $C\mathsf{K}$-structured, there is a list $(\Gamma_k \Rightarrow \Delta_k)_{1 \leq k \leq n}$ of sequents in $H \cup \{\sigma^\circ\}$ such that for every $\Gamma \Rightarrow \Delta \in H \cup \{\sigma^\circ\}$ and $(C_R, U) \in C$ there is an $i \in C_R$ such that $[a]^{-1}\Gamma_i \not\subseteq \Gamma$ or a $j \in U$ such that $\Gamma_j \Rightarrow \Delta_j \neq \Gamma \Rightarrow \Delta$. For the rest of this proof we fix this list $(\Gamma_k \Rightarrow \Delta_k)_{1 \leq k \leq n}$.

Because $\sigma^\circ \sqsubseteq H$, there is a sequent $\overline{\sigma} \in H$ such that $\sigma^\circ \sqsubseteq \overline{\sigma}$. Let $(\overline{\Gamma_k} \Rightarrow \overline{\Delta_k})_{1 \leq k \leq n}$ be the list obtained by replacing in $(\Gamma_k \Rightarrow \Delta_k)_{1 \leq k \leq n}$ each occurrence of σ° by $\overline{\sigma}$. By the $C\mathsf{K}$-structuredness of H, there must be some $(C_R^0, U) \in C$ and $\overline{\Gamma} \Rightarrow \overline{\Delta} \in H$ such that for each $i \in C_R^0$ we have $[a]^{-1}\overline{\Gamma_i} \subseteq \overline{\Gamma}$, and for each $j \in U$ we have $\overline{\Gamma_j} \Rightarrow \overline{\Delta_j} = \overline{\Gamma} \Rightarrow \overline{\Delta}$.

It follows for every $i \in C_R^0$ that $[a]^{-1}\Gamma_i \subseteq [a]^{-1}\overline{\Gamma_i} \subseteq \overline{\Gamma}$. Thus, by the fact that $H \mid \sigma^\circ$ is not $C\mathsf{K}$-structured, there is a $k \in U$ such that $\Gamma_k \Rightarrow \Delta_k \neq \overline{\Gamma_k} \Rightarrow \overline{\Delta_k}$. By construction this can only be the case if $\Gamma_k \Rightarrow \Delta_k = \sigma^\circ$.

Now consider the following rule instance.

$$\mathsf{r}_C^{\mathsf{HK}^*}\ \frac{\{H \mid \bigcup_{i \in C_R}[a]^{-1}\Gamma_i, \bigcup_{j \in U}\Gamma_j \Rightarrow \bigcup_{j \in U}\Delta_j : (C_R, U) \in C\}}{H \mid \Gamma_k \Rightarrow \Delta_k}$$

We claim that for any $(C_R, U) \in C$, the sequent

$$\sigma_R := \bigcup_{i \in C_R}[a]^{-1}\Gamma_i \cup \bigcup_{j \in U}\Gamma_j \Rightarrow \bigcup_{j \in U}\Delta_j$$

is such that $\sigma^\circ \sqsubset \sigma_R$. As $k \in U$, we already have $\sigma^\circ \sqsubseteq \sigma_R$. Now suppose, towards a contradiction, that $\sigma^\circ = \sigma_R$. Then by the fact that $H \mid \sigma^\circ$ is not $C\mathsf{K}$-structured, there must be some $j \in U$ such that $\Gamma_j \Rightarrow \Delta_j \neq \sigma^\circ$. It follows that

$$\begin{aligned}
\Gamma_j \Rightarrow \Delta_j &= \overline{\Gamma_j} \Rightarrow \overline{\Delta_j} && (\text{Definition of } \overline{\cdot},\ \Gamma_j \Rightarrow \Delta_j \neq \sigma^\circ)\\
&= \overline{\Gamma_k} \Rightarrow \overline{\Gamma_k} && (j, k \in U)\\
&= \overline{\sigma}.
\end{aligned}$$

But then $\overline{\sigma} = \Gamma_j \Rightarrow \Delta_j \sqsubseteq \sigma_R = \sigma^\circ$. Since, by construction, also $\sigma^\circ \sqsubseteq \overline{\sigma}$, we have $\sigma^\circ = \overline{\sigma}$ and thus $H \mid \sigma^\circ = H$, contradicting the assumption that $H \mid \sigma^\circ$ is not $C\mathsf{K}$-structured.

To finish the proof, let $(\widehat{\Gamma_k} \Rightarrow \widehat{\Delta_k})_{1 \leq k \leq n}$ be the result of replacing in the list $(\Gamma_k \Rightarrow \Delta_k)_{1 \leq k \leq n}$ each occurrence of σ° by σ. Consider the following rule instance:

$$\mathsf{r}_C^{\mathsf{HK}^*}\ \frac{\{H \mid \bigcup_{i \in C_R}[a]^{-1}\widehat{\Gamma}_i, \bigcup_{j \in U}\widehat{\Gamma}_j \Rightarrow \bigcup_{j \in U}\widehat{\Delta_j} : (C_R, U) \in C\}}{H \mid \widehat{\Gamma_k} \Rightarrow \widehat{\Delta_k}}$$

Let $(C_R, U) \in C$ be arbitrary and define:

$$\widehat{\sigma_R} := \bigcup_{i \in C_R} [a]^{-1}\widehat{\Gamma}_i \cup \bigcup_{j \in U} \widehat{\Gamma}_j \Rightarrow \bigcup_{j \in U} \widehat{\Delta_j}.$$

Clearly $\sigma = \widehat{\Gamma_k} \Rightarrow \widehat{\Delta_k} \sqsubseteq \widehat{\sigma_R}$. We have seen above that $\sigma^\circ \sqsubset \sigma_R$. It follows that σ_R has no formula in focus, and there is a some formula φ° in either the right-hand side of the left-hand side of σ_R, which does not belong the same side of σ°. Without loss of generality, suppose that φ° belongs to the right-hand side of σ_R and $\varphi^\circ \notin \Delta_k$. Then clearly $\varphi^\circ \notin \widehat{\Delta_k}$, and we claim that φ° belongs to the right-hand side of $\widehat{\sigma_R}$. Indeed, since $\varphi^\circ \notin \Delta_k$, we must have $\varphi^\circ \in \Delta_j$ for some $j \in U$ such that $\Gamma_j \Rightarrow \Delta_j \neq \sigma^\circ$. But then $\varphi^\circ \in \widehat{\Delta_j}$, as required. This shows that $\sigma \sqsubset \widehat{\sigma_R}$ for each $(C_R, U) \in C$. □

We are now ready to prove the Existence Lemma for the master modality.

3.5.28. LEMMA. *Let H be a $\mathcal{C}\mathsf{X}^i$-maximal hypersequent, such that $i = cf$ implies both that $\mathsf{X} \in \{\mathsf{K}, \mathsf{K4}\}$ and that $\mathcal{C}$ consists of only equable frame conditions. Then for every sequent $\Gamma \Rightarrow \Delta \in H$ with $[*]\varphi \in \Delta^-$, there is a sequent $\Gamma_1 \Rightarrow \Delta_1 \in H$ such that $(\Gamma \Rightarrow \Delta)R_*^{\mathsf{X}}(\Gamma_1 \Rightarrow \Delta_1)$ and $\varphi \in \Delta_1^-$.*

Proof:
Let $\mathcal{S}$ be the subset of H consisting of those sequents $\Gamma_1 \Rightarrow \Delta_1$ for which it holds that $\Gamma \Rightarrow \Delta R_*^{\mathsf{X}} \Gamma_1 \Rightarrow \Delta_1$, and either $\varphi \in \Delta_1^-$ or $[a][*]\varphi \in \Delta_1^-$ for some $a \in \mathsf{D}$. We must show that $\mathcal{S}$ contains a sequent $\Gamma_1 \Rightarrow \Delta_1$ with $\varphi \in \Delta_1^-$. Assume that this is not the case. We will reach a contradiction by constructing a $\mathcal{C}\mathsf{X}^i$-proof π of H.

Since $\Gamma \Rightarrow \Delta \in \mathcal{S}$, our assumption gives $[a][*]\varphi \in \Delta^-$ for some $a \in \mathsf{D}$. We begin the construction of π as follows:

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\pi_1}{H \mid \Gamma'_{\mathsf{X}} \Rightarrow \varphi^\circ, \Delta'_{\mathsf{X}}}}{H \mid \Gamma'_{\mathsf{X}} \Rightarrow \varphi^\circ, \Delta'_{\mathsf{X}}}\,\mathsf{fc} \qquad \dfrac{\{\pi_b : b \in \mathsf{D}\}}{\{H \mid \Gamma'_{\mathsf{X}} \Rightarrow [b][*]\varphi^\bullet, \Delta'_{\mathsf{X}} \mid b \in \mathsf{D}\}}}{H \mid \Gamma'_{\mathsf{X}} \Rightarrow [*]\varphi^\bullet, \Delta'_{\mathsf{X}}}\,[*]_R}{H \mid \Gamma_{\mathsf{X}} \Rightarrow [a][*]\varphi^\bullet, \Delta_{\mathsf{X}}}\,[a]_{\mathsf{X}}}{\dfrac{\vdots}{H \mid \Gamma \Rightarrow [a][*]\varphi^\bullet, \Delta}\,\mathsf{iw}_L + \mathsf{iw}_R}}{H \mid \Gamma \Rightarrow [a][*]\varphi^\circ, \Delta}\,\mathsf{fc}$$

Here the rule application $[a]_\mathsf{X}$ is similar to the one in Lemma 3.5.24. That is:

$$\Gamma_\mathsf{X} := [a][a]^{-1}\Gamma \qquad \Gamma'_\mathsf{X} := \begin{cases} [a]^{-1}\Gamma & \text{if } \mathsf{X} \in \{\mathsf{K}, \mathsf{K4}\} \\ [a]^{-1}\Gamma, [a][a]^{-1}\Gamma & \text{if } \mathsf{X} \in \{\mathsf{B}, \mathsf{B4}\} \end{cases}$$

$$\Delta_\mathsf{X} := \begin{cases} \emptyset & \text{if } \mathsf{X} \in \{\mathsf{K}, \mathsf{K4}\} \\ \Delta_0 & \text{if } \mathsf{X} \in \{\mathsf{B}, \mathsf{B4}\} \end{cases} \qquad \Delta'_\mathsf{X} := \begin{cases} \emptyset & \text{if } \mathsf{X} \in \{\mathsf{K}, \mathsf{K4}\} \\ [a]\Delta_0^\circ & \text{if } \mathsf{X} = \mathsf{B} \\ [a]\Delta_0^\circ, [a][a]^{-1}\Delta_0^\circ & \text{if } \mathsf{X} = \mathsf{B4} \end{cases}$$

where $\Delta_0 \subseteq \Delta^\circ$ consists of those $\psi^\circ \in \Delta^\circ$ such that $[a]\psi \in \mathsf{FL}(H)$. Note that, by Lemma 3.5.21, if for some $\sigma \in H$ it holds that $(\Gamma'_\mathsf{X} \Rightarrow \Delta'_\mathsf{X}) \sqsubseteq \sigma$, then $(\Gamma \Rightarrow \Delta)R_a^\mathsf{X}\sigma$. This will be useful later in our proof.

The proof π_1 is obtained by the $\mathcal{C}\mathsf{X}^i$-fullness of H together with the fact that $\Gamma'_\mathsf{X} \Rightarrow \varphi^\circ, \Delta'_\mathsf{X} \not\sqsubseteq H$. The latter must be the case, for otherwise there would be a $\Gamma_1 \Rightarrow \Delta_1 \in \mathcal{S}$ with $\varphi \in \Delta_1^-$, which we assumed not to be the case.

For the construction of the derivations π_b, we make a case distinction on whether $i = cf$ or $i = an$.

Suppose first that $i = an$. Then each derivation π_b is constructed by repeatedly applying Proposition 3.5.25 to the active sequent, until we have reached a sequent $\Gamma_2 \Rightarrow [b][*]\varphi^\bullet, \Delta_2$ such that one of the following holds:

- $H \mid \Gamma_2 \Rightarrow [b][*]\varphi^\circ, \Delta_2^\circ$ is $\mathcal{C}\mathsf{X}^{an}$-provable. In this case we append its proof to to our derivation, with an application of

  $$\mathsf{fc}\,\frac{H \mid \Gamma_2 \Rightarrow [b][*]\varphi^\circ, \Delta_2^\circ}{H \mid \Gamma_2 \Rightarrow [b][*]\varphi^\bullet, \Delta_2}$$

  in between.

- $\Gamma_2 \Rightarrow [b][*]\varphi^\circ, \Delta_2^\circ \in H$. Since $\Gamma \Rightarrow \Delta R_a^\mathsf{X}\Gamma_2 \Rightarrow [b][*]\varphi^\circ, \Delta_2^\circ$, we have in this case that $\Gamma_2 \Rightarrow [b][*]\varphi^\circ, \Delta_2^\circ \in \mathcal{S}$. We then repeat the same process but now applied to $\Gamma_2 \Rightarrow [b][*]\varphi^\bullet, \Delta_2$, minus the first step of applying fc.

- $H \mid \Gamma_2 \Rightarrow [b][*]\varphi^\circ, \Delta_2^\circ$ is complete. We claim that in this case we must already be in either one of the two previous cases. Indeed, by Lemma 3.5.12, the hypersequent $H \mid \Gamma_2 \Rightarrow [b][*]\varphi^\circ, \Delta_2^\circ$ is $\mathcal{C}\mathsf{X}^{an}$-full. In fact, by completeness, it is strongly $\mathcal{C}\mathsf{X}^{an}$-full. Thus, if it is $\mathcal{C}\mathsf{X}^{an}$-unprovable, it must by Lemma 3.5.20 and Lemma 3.5.22 be $\mathcal{C}\mathsf{X}^{an}$-saturated. But then it is equal to H by $\mathcal{C}\mathsf{X}^{an}$-maximality.

Note that the above process must terminate, since Proposition 3.5.25 properly extends our sequent, which, by analyticity, can only happen finitely often.

Now suppose that $i = cf$. We construct each π_b in a similar way as in the previous case, this time repeatedly applying Proposition 3.5.26 and Proposition 3.5.27 until a sequent $\Gamma_2 \Rightarrow [b][*]\varphi^\bullet, \Delta_2^\circ$ is reached such that one of the following holds:

- $H \mid \Gamma_2 \Rightarrow [b][*]\varphi^\circ, \Delta_2^\circ$ is $\mathcal{C}\mathsf{X}^{cf}$-provable. We append its proof exactly as in the case of $i = 1$.

- $\Gamma_2 \Rightarrow [b][*]\varphi^\circ, \Delta_2^\circ \in H$. Here, similar to the case where $i = an$, we loop back and repeat the process with $\Gamma_2 \Rightarrow [b][*]\varphi^\bullet, \Delta_2^\circ$.

- $H \mid \Gamma_2 \Rightarrow [b][*]\varphi^\circ, \Delta_2^\circ$ is $\mathcal{C}\mathsf{X}^{cf}$-saturated. As with $i = an$, in this case we must already be in either one of the two previous cases. By Lemma 3.5.12, the hypersequent $H \mid \Gamma_2 \Rightarrow [b][*]\varphi^\circ, \Delta_2^\circ$ is $\mathcal{C}\mathsf{X}^{cf}$-full. Thus, if it $\mathcal{C}\mathsf{X}^{cf}$-unprovable and $\mathcal{C}\mathsf{X}^{cf}$-saturated, it must be equal to H by $\mathcal{C}\mathsf{X}^{cf}$-maximality.

Again, this process terminates for similar reasons.

In either case we end up with a (possibly infinite) $\mathcal{C}\mathsf{X}^i$-derivation π of H. We claim that π is in fact a $\mathcal{C}\mathsf{X}^i$-proof. Indeed, the rule fc is only applied at the root of π, and every infinite branch must have a final segment on which there is always a formula in focus and a focused unfolding happens infinitely often. Hence, we have obtained our desired contradiction. □

3.5.6 The Truth Lemma

With the existence lemmata in place, the Truth Lemma can now be proved by a routine induction.

3.5.29. Lemma. *Let H be $\mathcal{C}\mathsf{X}^i$-maximal. Then $\mathbb{S}^{\mathsf{X}}_H, \sigma \not\Vdash \sigma$, for every $\sigma \in H$.*

Proof:
We will show by induction on φ that for every $\Gamma \Rightarrow \Delta \in H$ we have $\Gamma \Rightarrow \Delta \Vdash \varphi$ if $\varphi \in \Gamma^-$, and $\Gamma \Rightarrow \Delta \not\Vdash \varphi$ if $\varphi \in \Delta^-$. We make a case distinction on the main connective of φ.

- $\varphi = \bot$.

  In this case $\varphi \notin \Gamma^-$, by the unprovability of H.

  If $\varphi \in \Delta^-$, then $\Gamma \Rightarrow \Delta \not\Vdash \varphi$, as required.

- $\varphi = \psi_1 \to \psi_2$.

  If $\varphi \in \Gamma^-$, we have by propositional saturation that $\psi_2 \in \Gamma^-$ or $\psi_1 \in \Delta^-$. By the induction hypothesis, this given $\Gamma \Rightarrow \Delta \Vdash \psi_2$ or $\Gamma \Rightarrow \Delta \not\Vdash \psi_1$, *i.e.* $\Gamma \Rightarrow \Delta \Vdash \psi_1 \to \psi_2$.

  If $\varphi \in \Delta^-$, propositional saturation gives $\psi_1 \in \Gamma^-$ and $\psi_2 \in \Delta^-$. Hence, by the induction hypothesis, we have $\Gamma \Rightarrow \Delta \Vdash \psi_1$ and $\Gamma \Rightarrow \Delta \not\Vdash \psi_2$, whence $\Gamma \Rightarrow \Delta \not\Vdash \psi_1 \to \psi_2$.

- $\varphi = [a]\psi$.

  Suppose that $\varphi \in \Gamma^-$ and $\Gamma \Rightarrow \Delta R^{\mathsf{X}}_a \Gamma_1 \Rightarrow \Delta_1$. Then $\psi \in \Gamma_1^-$, whence by the induction hypothesis $\Gamma_1 \Rightarrow \Delta_1 \Vdash \psi$. Since $\Gamma_1 \Rightarrow \Delta_1$ was taken arbitrarily, we find $\Gamma \Rightarrow \Delta \Vdash \varphi$.

  If $\varphi \in \Delta^-$, then by Lemma 3.5.24, there is some $\Gamma \Rightarrow \Delta R^{\mathsf{X}}_a \Gamma_1 \Rightarrow \Delta_1$ with $\psi \in \Delta_1^-$. The induction hypothesis gives $\Gamma_1 \Rightarrow \Delta_1 \nVdash \psi$, whence $\Gamma \Rightarrow \Delta \nVdash \varphi$.

- $\varphi = [*]\psi$.

  If $\varphi \in \Gamma^-$, let

  $$\Gamma \Rightarrow \Delta =: \Gamma_0 \Rightarrow \Delta_0 R^{\mathsf{X}}_{a_1} \Gamma_1 \Rightarrow \Delta_1 R^{\mathsf{X}}_{a_2} \Gamma_2 \Rightarrow \Delta_2 R^{\mathsf{X}}_{a_3} \ldots R^{\mathsf{X}}_{a_n} \Gamma_n \Rightarrow \Delta_n$$

  be a path in $\mathbb{S}^{\mathsf{X}}_H$ starting at $\Gamma \Rightarrow \Delta$. We will prove by induction on n that $\varphi \in \Gamma_n^-$, whence by propositional saturation $\psi \in \Gamma_n^-$, and thus by the induction hypothesis $\Gamma_n \Rightarrow \Delta_n \Vdash \psi$. Since this then holds for arbitrary paths, we find $\Gamma \Rightarrow \Delta \Vdash [*]\psi$. The induction base holds by assumption. For the induction step, suppose that $\varphi \in \Gamma_k^-$. Then by propositional saturation $[a_{k+1}]\varphi \in \Gamma_k^-$. Hence by the same reasoning as in the previous case, we have $\varphi \in \Gamma_{k+1}^-$, as required.

  Finally, if $\varphi \in \Delta^-$, we use Lemma 3.5.28 to obtain some $\Gamma \Rightarrow \Delta R^{\mathsf{X}}_* \Gamma_1 \Rightarrow \Delta_1$ with $\psi \in \Delta_1^-$. Since by the induction hypothesis $\Gamma_1 \Rightarrow \Delta_1 \nVdash \psi$, we find $\Gamma \Rightarrow \Delta \nVdash [*]\psi$.

This finishes the proof. □

3.5.7 Wrapping up

We now have everything needed to prove Theorem 3.5.1.

Proof of Theorem 3.5.1:
Suppose that H is $\mathcal{C}\mathsf{X}^i$-unprovable. By Lemma 3.5.23, there is a $\mathcal{C}\mathsf{X}^i$-maximal $\overline{H} \sqsupseteq H$. But then it holds by Lemma 3.5.29 that the assignment $\alpha(\overline{\sigma}) = \overline{\sigma}$ is a cmsa of $\overline{H}$ in $\mathbb{S}^{\mathsf{X}}_{\overline{H}}$. Hence we find by Proposition 3.5.4 that $\overline{H}$, and thus also $H \sqsubseteq \overline{H}$, is not $\mathcal{C}\mathsf{X}$-valid. □

As a corollary we obtain the small model property for all frame conditions under consideration.

3.5.30. Corollary (Small Model Property). *If φ is not $\mathcal{C}\mathsf{X}$-valid, then it is falsified in a $\mathcal{C}\mathsf{X}$-model of size exponential in $\mathsf{FL}(\varphi)$.*

Proof:
Suppose φ is not valid. Then the hypersequent $\Rightarrow \varphi$ not valid. In the same way as in the proof of Theorem 3.5.1, we obtain a $\mathsf{FL}(\varphi)$-hypersequent H such that $\mathbb{S}^{\mathsf{X}}_{H}$ falsifies φ. We claim that, by its maximality, the hypersequent H contains at most 3^n sequents. Indeed, for each sequent σ in H, and each formula ψ in $\mathsf{FL}(\varphi)$, precisely one of the following holds: (i) ψ° belongs to the left-hand side of σ, (ii) ψ° belongs to the right-hand side of σ, or (iii) ψ° belongs to neither side of σ. Moreover, each sequent σ in H is precisely determined by the which of the previous three items it satisfies for each formula $\psi \in \mathsf{FL}(\varphi)$. Therefore H contains at most 3^n sequents, where $n := \mathsf{FL}(\varphi)$. □

3.6 Conclusion

We have constructed sound and complete non-well-founded sequent calculi for modal logic with the master modality interpreted over classes of $\mathcal{C}\mathsf{X}$-frames. This is an extension of the method and a generalisation of the results by Lahav in [66]. The following gives an overview of our contributions.

- We extended the calculi from unimodal to multimodal logic.
- We extended the calculi from multimodal logic to multimodal logic with the master modality (*aka* Common Knowledge Logic), by (i) adding left and right rules for the master modality $[*]$, and (ii) imposing a soundness condition on infinite branches using a focus mechanism. The resulting calculi are denoted by $\mathsf{HX}^* + \mathsf{R}_{\mathcal{C}}$.
- We established soundness for each of the $\mathsf{HX}^* + \mathsf{R}_{\mathcal{C}}$.
- We established analytic completeness for each of the $\mathsf{HX}^* + \mathsf{R}_{\mathcal{C}}$. This required a novel argument for the case of the modality $[*]$ in the Truth Lemma.
- We established cut-free completeness for those calculi $\mathsf{HX}^* + \mathsf{R}_{\mathcal{C}}$ for modal logic with the master modality, where $\mathsf{X} \in \{\mathsf{K}, \mathsf{K4}\}$ and $\mathcal{C}$ consists only of what we called *equable* frame conditions. For this we needed to introduce the notion of equability, as well as a new notion of maximality, as explained in Remark 3.5.18. Although we do not obtain cut-free completeness for all frame conditions that Lahav obtains cut-free completeness for with respect to the basic modal language, our result still covers infinitely many different frame conditions.

The following table sums up what we now know about the completeness of the hypersequent calculi at hand, where $\mathsf{HX} + \mathsf{R}_{\mathcal{C}}$ denotes the calculus for the basic modal language presented by Lahav in [66].

	$\mathsf{X} \in \{\mathsf{K}, \mathsf{K4}\}$, $\mathcal{C}$ equable	$\mathsf{X} \in \{\mathsf{K}, \mathsf{K4}\}$, $\mathcal{C}$ *not* equable	$\mathsf{X} \in \{\mathsf{B}, \mathsf{B4}\}$
$\mathsf{HX} + \mathsf{R}_{\mathcal{C}}$	*cf*	*cf*	*an*
$\mathsf{HX}^* + \mathsf{R}_{\mathcal{C}}$	*cf*	*an*	*an*

In other words, the only frame conditions for which Lahav obtains a cut-free calculus and we do not, are those where $\mathsf{X} \in \{\mathsf{K}, \mathsf{K4}\}$ and $\mathcal{C}$ contains conditions which are not equable. This begs the following question.

3.6.1. QUESTION. For which other sets $\mathcal{C}$ of simple (not necessarily equable) frame conditions is $\mathsf{HK}^* + \mathsf{R}_{\mathcal{C}}$ cut-free complete?

Along the same lines, it would be interesting to see whether we could cover frame conditions *not* covered by Lahav's method. For instance:

3.6.2. QUESTION. Is it possible to construct an analytic cyclic proof system for $\mathtt{ML}^*$ interpreted over the class of $\mathsf{KD45}$-frames?

Another pressing question is whether the same techniques can be applied to more expressive fragments of the modal μ-calculus. We conjecture this is the case for PDL (possibly with converse). In [67] a system for PDL with converse is presented, which shares many similarities to our hypersequent calculi. Unfortunately, an error was found in this paper, which is explained in [89, Section 7.3]. If our conjecture holds, it might provide a way to repair the error in [67] and restore its results.

Along similar lines, it might be interesting to consider a fragment with specific syntactic properties, such as the *weakly aconjunctive* or *disjunctive* formulas from [107]. In the same paper these formulas are related to so-called *thin refutations*. Those, in turn, are related to proof systems with a single focus annotation. This suggests that our techniques might be applicable to these fragments as well.

Related to the above, it would be interesting to see if our cyclic proofs could be translated into Hilbert-style proofs with an explicit induction rule. Such translations would provide alternative proofs of completeness for well-known Hilbert-style calculi for $\mathtt{ML}^*$. This will be challenging, as the language $\mathtt{ML}^*$ lacks the necessary expressive power for the *strengthened induction rule* of Afshari & Leigh in [4]. For more details about this issue, we refer the reader to Section 7.1 of [89]. This further points in the direction of extending the language, where an ultimate goal would be to prove completeness for Kozen's axiomatisation of the modal μ-calculus interpreted over various frame classes.

On a different note, we would like to investigate whether our proof systems can be used to establish the interpolation property for $\mathtt{ML}^*$. Thomas Studer, based on earlier work by Maksimova [72], shows in [103] that $\mathtt{ML}^*$ does not enjoy this property over several frame classes, including the class of all frames. For many other frame classes, such as the class $\mathsf{S5}_n$ which plays an important role in epistemic logic, this question remains open. Although the negative results by

Maksimova and Studer are not encouraging, if ML* does turn out to have the interpolation property over some simple frame class, then perhaps our calculus can be used to prove it.

Another interesting avenue for further research would be to connect the hypersequent calculi of this chapter to algebraic approaches to proof theory. For instance, the paper [14] constructs analytic proof calculi for basic modal logic using algebraic techniques. In combination with our ideas this could perhaps be generalised to modal fixed point logics.

Lastly, whilst we showed the small model property for all of the logics considered this chapter, we have yet to provide a bound on the size of *proofs*. This topic will be addressed in the Intermezzo following this chapter.

Intermezzo

In this intermezzo we introduce an abstract framework for reasoning about the cyclic proof systems in this thesis. As main application, we will prove the *bounded proof property* for (abstract) proof systems satisfying certain sufficient conditions. Even though in all known concrete cases this can already be proven using the positional determinacy of parity games, often with a sharper bound, we still think our framework is interesting for the following reasons:

Special purpose. Our framework is tailored to the types of conditions on infinite branches one encounters in the cyclic proof theory literature. Even though in practice these very often turn out to be parity conditions, this is not always the case. Moreover, as we will see later, our argument for the bounded proof property has a proof-theoretic rather than an automata-theoretic flavour. Like the standard arguments for cut-elimination, it works by pushing some unwanted structure in a proof upwards, until it eventually disappears.

Generality. It seems that the generality of our framework cannot be captured by that of parity games, in the sense that the bounded proof property result given below cannot be obtained using the positional determinacy of some parity game. We make no claim in the other direction, *i.e.* that our framework says something about the theory of parity games.

Unification. The bounded proof property is often proven for a specific path-based non-well-founded proof system, *e.g.* in [54] and in [73]. Our framework allows one to unify these arguments and to prove results about multiple systems at once. Although much more sophisticated frameworks already exist for trace-based non-well-founded proof systems [6, 19], we believe our framework is the first to axiomatise path-based non-well-founded proof systems.

We begin this intermezzo by defining (non-well-founded) trees and some of their constructions as sets of sequences. In the section thereafter, we will give our abstract notion of a path-based non-well-founded proof system, define infinitary and cyclic proofs, and prove some basic properties. In the third section we will prove our main result: any infinitary proof with only finitely many distinct sequents, can be transformed into a cyclic proof of bounded size. We conclude in the final section.

I.1 Trees

Let $\mathbb{N}^*$ be the set of finite strings of natural numbers. For $a \in \mathbb{N}^*$ we use $|a|$ to denote the length of a. Moreover, we use $\leq$ for the prefix relation on $\mathbb{N}^*$.

A *tree* is a non-empty subset T of $\mathbb{N}^*$ such that (i) T is closed under taking prefixes, and (ii) $u \cdot m \in T$ entails $u \cdot n \in T$ for all $n, m \in \mathbb{N}$ with $n < m$. Elements of a tree are called *nodes*, and the empty string ϵ is called the *root*. If u, v are nodes of a tree such that $u \leq v$ we say that u is an *ancestor* of v, and that v is a *descendant* of u. If, moreover, $u \neq v$, then u and v, respectively, are said to be a *proper* ancestor and descendant. Finally, if there is no w such that $u < w < v$, then u and v, respectively are called *direct ancestors* and *direct descendants*. The *depth* $|u|$ of a node u is defined to be its length as a string.

A *path* in a tree T is a chain $u_0 < u_1 < \cdots (< u_n)$ such that u_{i+1} is a direct descendant of u_i for each i. If $u < v$, we write $[u, v]$ for the unique path $u = u_0 < u_1 < \cdots < u_n = v$ and $[u, v)$ for the same path minus u_n. A *branch* β of a tree T is a maximal $<$-chain in T. Note that every branch is also a path. We write $\beta(m)$ for the unique element in β of length m (if it exists).

We will now introduce a way of labelling the nodes of some tree. A *ranked alphabet* is a set Σ of *characters* together with a function $\mathsf{ar} : \Sigma \to \mathbb{N}$ assigning to each character an arity. Given a ranked alphabet Σ, a Σ*-labelled tree* is a tree T together with a labelling function $l : T \to \Sigma$ such that for every $u \in T$ and $n \in \mathbb{N}$:

$$u \cdot n \in T \Leftrightarrow n < \mathsf{ar}(l(u)).$$

A *tree language* over a ranked alphabet Σ is any set of Σ-labelled trees. Note that, by definition, every Σ-labelled tree is finitely branching.

Given a Σ-labelled tree T with labelling function l, the subtree *generated* by some node u of T is the Σ-labelled tree $\langle u \rangle$ with nodes

$$\langle u \rangle := \{v \in \mathbb{N}^* : u \cdot v \in T\},$$

and the labelling function l_u given by $l_u(v) := l(u \cdot v)$. Now suppose that T_1 and T_2 are both Σ-labelled trees, with labelling functions l_1 and l_2, respectively. We define the *substitution* $T_1[T_2/u]$ of the node $u \in T_1$ by the tree T_2 as follows:

$$T_1[T_2/u] := \{u \cdot v \mid v \in T_2\} \cup \{w \in T_1 \mid u \not\leq w\},$$

with the labelling function l given by $l(u \cdot v) = l_2(v)$ for every $v \in T_2$ and $l(w) = l_1(w)$ for every $w \in T_1$ with $u \not\leq w$. Note that, in particular, we have $l(u) = l(u \cdot \epsilon) = l_2(\epsilon)$.

A *finite tree with back edges* (T, f) consists of a finite tree T together with a partial function f from T to itself, such that (i) $\mathsf{dom}(f)$ consists of $\leq$-maximal elements in T, and (ii) $f(u) < u$ for every $u \in \mathsf{dom}(f)$. A *Σ-labelled finite tree with back edges* is a finite tree with back edges together with a labelling function $l : T \to \Sigma$ such that for every $u \in T \setminus \mathsf{dom}(f)$ it holds that

$$u \cdot n \in T \Leftrightarrow n < \mathsf{ar}(l(u)).$$

Note that the $\leq$-maximal elements in a tree T are its *leaves*. It follows from the definitions that in a Σ-labelled finite tree with back edges (T, f) with labelling function l, every leaf u either satisfies $\mathsf{ar}(l(u)) = 0$ or $u \in \mathsf{dom}(f)$. An element $u \in \mathsf{dom}(f)$ is often called a *repeating leaf*, and $f(u)$ is then called its *companion*. It will become clear later why we do not require that $l(u) = l(f(u))$.

A *path* in a finite tree with back edges (T, f) is a chain $(u_i)_{i \in I}$ with $I \in \omega$ or $I = \omega$ such that for each i, either u_{i+1} is a direct descendant of u_i, or $u_{i+1} = f(u)$, for some repeating leaf $u \in \mathsf{dom}(f)$ which is a direct descendant of u_i. Note that if $I = \omega$ then no u_i on the path belongs to $\mathsf{dom}(f)$.

We wish to define the tree obtained by unravelling a Σ-labelled finite tree with back edges (T, f) with labelling function l. For this we need the following definition. Let $\epsilon = u_0, \ldots, u_m$ be a finite path in (T, f) starting at the root. By definition, for each $1 \leq i \leq m$ there is an $n_i \in \mathbb{N}$ such that u_{i+1} is either of the form $u_i \cdot n_i$ or of the form $f(u_i \cdot n_i)$. The *projection* of this path is the finite string $n_1 \cdots n_m \in \mathbb{N}^*$. We define

$$\mathsf{un}(T, f) := \{u : u \text{ is a projection of a finite path in } (T, f) \text{ starting at the root}\}$$

The labelling function $l_{\mathsf{un}(T,f)}$ is defined as follows: if u is a projection of a non-empty path $u_0, \ldots, u_m$, we set $l_{\mathsf{un}(T,f)}(u) := l(u_m)$. If $u = \epsilon$, *i.e.* u is a projection of the empty path, we set $l_{\mathsf{un}(T,f)}(u) := l(\epsilon)$.

We close this section by defining words induced by paths in Σ-labelled trees. If $(u_i)_{i \in I}$ for $I \in \omega \cup \{\omega\}$ is a path in a tree T with labelling function $l : T \to \Sigma$, the word *induced* by this path is the word $l(u_0) \cdot l(u_1) \cdots l(u_n)$ if $I = n \in \omega$, and $l(u_0) \cdot l(u_1) \cdot l(u_2) \cdots$ if $I = \omega$.

I.2 Path-based non-well-founded proof systems

The following gives an abstract definition of a path-based non-well-founded proof system. Our definition is intentionally broad, offering substantial flexibility in the specification of both finite and infinite *good* paths, whereas usually in a cyclic proof system those two are more restricted and interrelated. On the one hand this is a disadvantage, because it means that our definition is too broad to precisely capture the concept of a path-based non-well-founded proof system, as is it often informally used in the literature. On the other hand, the advantage of proving results about these more general objects, of course, is that those results are more general as well.

We write Σ^* for the set of finite sequences of characters in Σ, and Σ^∞ for the set of infinite sequences of characters in Σ.

I.2.1. DEFINITION. A *(path-based) proof system* P is a ranked alphabet Σ together with:

(i) An equivalence relation $\equiv$ on Σ.

(ii) A relation $R \subseteq \Sigma \times \Sigma^*$ such that for all $a \in \Sigma$ and $w \in \Sigma^*$:

  (a) If aRw, then $\mathsf{length}(w) = \mathsf{ar}(a)$.

  (b) If aRw and $w' \in \Sigma^*$ is such that $\mathsf{length}(w') = \mathsf{length}(w)$ and moreover $w'(n) \equiv w(n)$ for all $n < \mathsf{ar}(a)$, then aRw'.

(iii) A set $G \subseteq \Sigma^*$ such that if $w_1 \cdot w_2 \cdot w_3 \in G$ and $w_2 \notin G$, then $w_1 \cdot w_3 \in G$.

(iv) A set $I \subseteq \Sigma^\infty$ such that if $w_0 \cdot w_1 \cdot w_2 \cdots \in I$, then $w_i \in G$ for infinitely many i.

In the context of a proof system P, elements of G are called *good words*, and elements of I are called *good infinite words*.

The intuition behind this definition is that Σ consists of rules instances of the form:

$$\mathsf{r}\,\frac{\Gamma_1 \quad \cdots \quad \Gamma_n}{\Gamma}$$

Note that the *axioms* are precisely the rule instances in Σ of arity 0.

The equivalence relation $\equiv$ identifies rule instances with the same conclusion. The relation R determines whether the premisses of some rule instance $i \in \Sigma$ match the conclusions the rule instances $i_1 \cdots i_n$. Accordingly, for $iRi_1 \cdots i_n$ to hold, the arity of i should be n. Moreover, if $i'_1, \ldots, i'_n$ are such that $i_k \equiv i'_k$ for every $1 \leq k \leq n$, then we have

$$iRi_1 \cdots i_n \Leftrightarrow iRi'_1 \cdots i'_n.$$

As will become clear later, the purpose of the relation $\equiv$ is to allow us to express that some branch in a proof tree contains multiple occurrences of the same sequent (even though it need not contain multiple occurrences of the same rule instance).

The third and fourth conditions are needed to determine which cyclic, respectively infinitary, derivations will count as proofs. The third condition is inspired by the fact that path-based cyclic proof systems often deem some path good if something good persists on the path (*e.g.* there is always a formula in focus) and something good happens at least once (*e.g.* a focussed unfolding happens). Hence, if $w_1 \cdot w_2 \cdot w_3$ is good and w_2 is not good, then the good thing must happen either in w_1 or in w_3, whence $w_1 \cdot w_3$ is good. Note that the condition on G forces $\epsilon \in G$. This will have no impact on our definition of proofs, because we will require that the path between a repeating leaf and its companion is of non-zero length.

Finally, the fourth condition links the infinite good paths to finite good paths. It is a weakened version of the observation that infinite good paths often have final segments that are concatenations of finite good paths. For instance, in the focus systems of the previous chapter, a finite good path is one where the focus always persists and a focussed unfolding happens at least once, and an infinite good path is one with a final segment where the focus always persists and a focussed unfolding happens infinitely often.

Before we go on to define what a *proof* is in some proof system P, we first establish a consequence of the definition that will be useful later on.

I.2.2. Lemma. *Let G be the set of good words of some proof system P. If $w_1 \cdot w_2$ belongs to G, then so does either w_1 or w_2.*

Proof:
By the hypothesis we have $\epsilon \cdot w_1 \cdot w_2 \in G$. Now suppose that $w_1 \notin G$. Then by condition (iii) of Definition I.2.1, we have $\epsilon \cdot w_2 \in G$, whence $w_2 \in G$. □

The following definition specifies what it means to be an infinitary proof in some proof system P.

I.2.3. Definition. Let P be a proof system with alphabet Σ. A Σ-labelled tree T with labelling function l is said to be a P-*derivation* if for every node u of T it holds that $l(u)Rl(u \cdot 0) \cdots l(u \cdot (\mathsf{ar}(u) - 1))$. A P-derivation is called a P-*proof* if every word induced by an infinite branch belongs to the set I of infinite good words.

Cyclic proofs are defined similarly, using the set G of finite good words instead of the set I of infinite good words. Recall that if u, v are nodes in a tree such that $u < v$, we write $[u, v)$ for the finite upward path from u to v that includes u but not v

I.2.4. Definition. Let P be a proof system with alphabet Σ. A Σ-labelled finite tree with back edges (T, f) is a *cyclic* P-*derivation* if for every node u of $T \backslash \mathsf{dom}(f)$ it holds that $l(u)Rl(u \cdot 0) \cdots l(u \cdot (\mathsf{ar}(u) - 1))$, and for every $u \in \mathsf{dom}(f)$ it holds that $l(f(u)) \equiv l(u)$. A cyclic P-derivation is called a *cyclic* P-*proof* if for every $u \in \mathsf{dom}(f)$ the word induced by the path $[f(u), u)$ belongs to G.

Note that we do not require $l(f(u)) = l(u)$, but merely $l(f(u)) \equiv l(u)$. The idea is that a repeating leaf only has to repeat a sequent, not the entire rule instance.

We will often speak about a path $[u, v)$, or about some infinite branch β, in a P-derivation being *good* or *bad*. This is a slight abuse of language, because we are actually talking about whether the *words* induced by these paths belong to G or, respectively, to I.

I.2.5. Example. Many of the path-based non-well-founded proof systems appearing in the literature fall within the scope of our definition. For each of the examples below, the appropriate ranked alphabet Σ consists of the respective system's rule instances (including its axioms), the equivalence relation $\equiv$ identifies rule instances with the same conclusion, and the relation R connects a rule instance i with n premisses, to n rule instances i_k such that each k-th premiss of the instance i is the same as the conclusion of the instance i_k.

- Each hypersequent calculus $\mathsf{HX}^* + \mathsf{R}_\mathcal{C}$ from the previous chapter.
- The multi-focus system Focus for the alternation-free modal μ-calculus given by Marti & Venema in [73]. The good finite words are those in which there is always a formula in focus, the focus rule is never applied, and the modal rule is applied at least once. The good infinite words are those where there is always a formula in focus, the focus rule is never applied, and the modal rule is applied infinitely often.
- The Jungteerapanich-Stirling system for the modal μ-calculus [54, 100]. The good finite words are those where some name z persists in the control throughout and is *reset* at least once. The good infinite words are those where some name z persists throughout and is reset infinitely often.
- A recent preprint by Leigh & Wehr gives a general method for constructing path-based proof systems from trace-based proof systems [69]. All of the resulting path-based proof systems are captured by our definition. The good words are those which contain an application of the reset rule to some part of the control that persists throughout. The good infinite words are those in which some part of the control is preserved throughout and is reset infinitely often.
- The cyclic proof system for Gödel-Löb logic by Shamkanov [96]. All finite and infinite words are good.
- The cyclic proof system for Grzegorczyk logic by Savateev & Shamkanov [92]. The good finite words are those on which there is a right premiss of the modal rule. The good infinite words are those where this is the case infinitely often.

Note, on the other hand, that a *trace-based* non-well-founded proof system generally does not fall within the scope of our definition. The reason is that our definition does not have a notion of *formula*, let alone of *trace*.

In many of the above examples the set of good infinite words can be generated from the set of good finite words in a canonical way.

I.2.6. Definition. Let $G \subseteq \Sigma^*$ be a set of good words. The set $I(G) \subseteq \Sigma^\infty$ of infinite words *generated* by G is given by:

$$I(G) := \{w_0 \cdot w_1 \cdot w_2 \cdots \mid w_i \in G \text{ for each } i \in \mathbb{N}\}.$$

I.2.7. Definition. A proof system P with G as set of good finite words and I as set of good infinite words, is called *simple* if the following two conditions hold:

(i) $I = I(G)$;

(ii) if $w_1 \cdot w_3 \in G$ and $w_2 \in G$, then $w_1 \cdot w_2 \cdot w_3 \in G$.

I.2.8. Example. All of the proof systems of example I.2.5 are simple, except for those featuring a *reset* rule.

I.2.9. Lemma. *Let* P *be a simple proof system. If* w_1 *and* w_2 *are good finite words, then so is* $w_1 \cdot w_2$*.*

Proof:
If w_1 is good, then clearly $w_1 \cdot \epsilon$ is as well. Applying item (ii) of Definition I.2.7, we obtain that $w_1 \cdot w_2 = w_1 \cdot \epsilon \cdot w_2$ is good. □

The following definition captures those proofs that have only finitely many distinct sequents.

I.2.10. Definition. A P-proof T is *frugal* whenever $\{l(u) \mid u \in T\}/\equiv$ is finite.

I.2.11. Example. Any $\mathcal{C}\mathsf{X}^{an}$-proof from the previous chapter is frugal. On the other hand, it is not hard to see that not every $\mathcal{C}\mathsf{X}$-proof is frugal. Consider, for instance, any infinitary $\mathcal{C}\mathsf{X}$-proof π and interleave some branch with applications of cut that together add infinitely many new formulas to that branch. For every proof system of Example I.2.5 it holds that every provable sequent has a frugal proof.

A node v of a P-proof T is called a *repeat* if $l(v) \equiv l(u)$ for some $u < v$. The repeat v is called *good* whenever there is a $u < v$ such that $l(v) \equiv l(u)$ and the word induced by $[u, v)$ is good. Note that a good repeat v might nonetheless have another ancestor $u' < v$ such that $l(v) \equiv l(u')$ and the word induced by $[u', v)$ is bad. Finally, if v is a repeat, we write $\widehat{v}$ for the minimal $u < v$ such that $l(v) \equiv l(u)$ and call $\widehat{v}$ the *companion* of v. Note that v might be a good repeat, even though the word induced by $[\widehat{v}, v)$ is not good. Namely, when the word induced by $[u, v)$ is good for some $\widehat{v} < u < v$ with $l(v) \equiv l(u)$.

For any frugal proof there is a cyclic proof with the same conclusion.

I.2.12. Proposition. *Let T be a frugal* P*-proof. Then there is a cyclic* P*-proof* (T_0, f) *such that* $l^T(\epsilon) \equiv l^{T'}(\epsilon)$.

Proof:
Let T_0 be the subtree of T obtained removing all proper descendants of good repeats. We claim that T_0 is finite. Indeed, suppose not. Then, since T is finitely branching, it follows by Kőnig's Lemma that T_0 has an infinite branch β. Moreover, since T is frugal, it follows by the pigeonhole principle that there are $n_0 < n_1 < n_2 < \cdots$ such that $l(\beta(n_i)) \equiv l(\beta(n_j))$ for every i, j. Consider the infinite word induced by the following infinite path (where concatenation of paths is defined in the obvious way):

$$[0, \beta(n_0)) \cdot [\beta(n_0), \beta(n_1)) \cdot [\beta(n_1), \beta(n_2)) \cdots$$

Since β is a branch of T, there must by Definition I.2.1.(iv) be at least one (in fact, infinitely many) i such that the word induced by $[\beta(n_i), \beta(n_{i+1}))$ is good, contradicting the assumption that β does not contain a good repeat.

The cyclic proof (T_0, f) is then defined by, for each repeating leaf v of T_0, letting $f(v)$ be the $u < v$ such that $l(v) \equiv l(u)$ and the word induced by $[u, v)$ is good. Note that such u exists, since, by construction, v is a good repeat. □

Although the previous proposition tells us that there is a cyclic proof for each frugal proof, it does not say anything about the size of this cyclic proof. The frugality ensures that a repeat happens at some point on each branch. However, it might take many repeats before a *good* repeat is reached. One way to ensure that the resulting cyclic proof is small, is by requiring that the *first* repeat is good.

A repeat v is said to be *minimal* if no $u < v$ is a repeat.

I.2.13. Definition. A P-proof is *concise* if all of its minimal repeats are good.

In the next section we will show how to transform frugal proofs into concise proofs. First, we will show that (over frugal proofs) conciseness is a weak form of the more well-known property of *uniformity*.

I.2.14. Definition. A P-proof T is *uniform* if $l(u) \equiv l(v)$ implies $l(u) = l(v)$ for every $u, v \in T$.

It follows that a proof is uniform if and only if the subtrees generated by two occurrences of the same sequent are isomorphic. In game-theoretic terms, this corresponds to the proof being a positional winning strategy for Prover.

I.2.15. Proposition. *Let T be a frugal* P*-proof. Then T is concise if uniform.*

Proof:
We must show that each minimal repeat of T is good. In fact, we will show that for every $u < v$ such that $l(u) \equiv l(v)$, the path $[u, v)$ is good. Suppose, towards a contradiction, that this is not the case. Write v as $v = u \cdot x$. By assumption the word induced by x is bad. Consider the final segment of an infinite branch

$$[u, u \cdot x) \cdot [v, v \cdot x) \cdot [v \cdot x, v \cdot x \cdot x) \cdots$$

By Definition I.2.1.(iv) at least one (in fact, infinitely many) of the words induced by these segments must be good, contradicting the fact that the word induced by x is bad. □

Another important property of infinitary proofs is *regularity*. Regular proofs are precisely the proofs that can be obtained by unravelling finite trees with back edges.

I.2.16. DEFINITION. A P-proof is *regular* if it has at most finitely many non-isomorphic subtrees.

The following proposition relates the notion of regularity to the notions we have discussed so far.

I.2.17. PROPOSITION. *Let T be a P-proof. Then:*

(i) If T is regular, then T is frugal.

(ii) If T is frugal and uniform, then T is regular.

Proof:
Item (i) is immediate by the definitions. For item (ii), note that, by uniformity, equivalent nodes generate isomorphic subtrees. Since, by frugality, T only contains finitely many nodes up to equivalence, it follows that T is regular. □

In the non-well-founded proof theory literature, it is often shown that the unravelling of a cyclic proof is a (regular) infinitary proof. Generally this does not hold for our abstract P-proofs, because Definition I.2.1 does not enforce a sufficiently strict relation between the set G of good words and the set I of good infinite words. It does, however, hold for simple proof systems. The rest of this section is devoted to proving this result.

I.2.18. DEFINITION. Let (T, f) be a finite tree with back edges. The *one-step dependency order* $\preceq_1$ on $\mathsf{ran}(f)$ is given by:

$$u \preceq_1 v :\Leftrightarrow v \leq u < v' \text{ for some } v' \in f^{-1}(v).$$

The *dependency order* $\preceq$ on $\mathsf{ran}(f)$ is defined as the transitive closure of $\preceq_1$. If $u \preceq v$, we say that u *depends* on v.

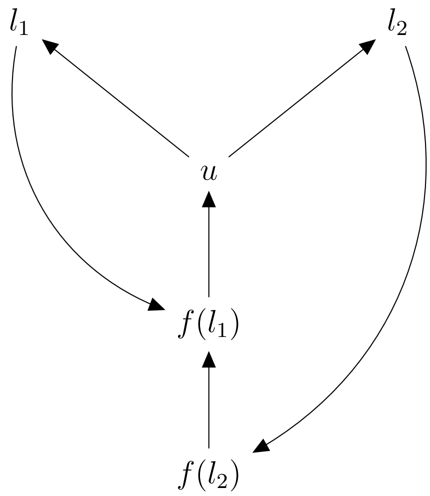

Figure I.4: A finite tree with back edges where $f(l_1) \preceq_1 f(l_2)$.

I.2.19. EXAMPLE. Figure I.4 shows an example of a finite tree with back edges where $f(l_1) \preceq_1 f(l_2)$.

Note that $u \preceq_1 v$ implies that there is a path from v to u, and therefore so does $u \preceq v$. It immediately follows that $\preceq$ is antisymmetric. Since $\preceq$ is also reflexive, it defines a partial order on $\mathsf{ran}(f)$ for any finite tree with back edges (T, f).

We write $\mathrm{Inf}(\alpha)$ for the elements occurring infinitely often in some given infinite sequence α.

I.2.20. LEMMA. *For any infinite path α through some finite tree with back edges (T, f), the set $\mathrm{Inf}(\alpha) \cap \mathsf{ran}(f)$ has a $\preceq$-greatest element.*

Proof:
An infinite path through (T, f) must pass at least one back edge, whence the set $\mathrm{Inf}(\alpha) \cap \mathsf{ran}(f)$ is both non-empty and finite. It therefore suffices to prove that all $\preceq$-maximal nodes in this set are equal.

Let u be a $\preceq$-maximal node in $\mathrm{Inf}(\alpha) \cap \mathsf{ran}(f)$. We claim that for all nodes v in $\mathrm{Inf}(\alpha)$ it holds that $u \leq v$. Note that it suffices to show this for each

$$f(l) \in \mathrm{Inf}(\alpha) \cap \{f(l) : l \in \mathsf{dom}(f) \text{ and } u \leq l\}$$

To this end, let l be an arbitrary such repeating leaf. Then both u and $f(l)$ lie somewhere on the path p_l from the root of T to l, so that either $u \leq f(l)$ or $f(l) \leq u$. Moreover, we must have $f(l) \in \mathrm{Inf}(\alpha) \cap \mathsf{ran}(f)$ and thus, by the maximality of u, it follows that $u \not\prec f(l)$. But this means that $u \leq f(l)$, as required.

Since u was chose arbitrarily, we find that $\mathrm{Inf}(\alpha) \cap \mathsf{ran}(f)$ indeed has a unique $\preceq$-maximal element. $\square$

The following lemma makes essential use of the assumption that P is a simple proof system.

I.2.21. LEMMA. *Let* P *be a simple proof system and let* (T, f) *be a cyclic* P*-proof. Suppose that* $u_0, \ldots, u_n$ *is a finite path in* (T, f) *such that the following hold:*

(i) $u_0 \in \mathsf{ran}(f)$*;*

(ii) $u_0 = u_n$*;*

(iii) $u_0 \leq u_i$ *for every* $0 \leq i \leq n$*.*

Then the word induced by the path $u_0, \ldots, u_{n-1}$ *is good.*

Proof:
Let K be the set of indices of nodes on $u_0, \ldots, u_n$ which belong to $\mathsf{ran}(f)$ and have already occurred earlier on the path. That is,

$$K := \{k \mid u_k \in \mathsf{ran}(f) \text{ and } u_k = u_j \text{ for some } j \in [0, k)\}$$

We proceed by induction on $|K|$. So suppose that the thesis has been proven for all K' such that $|K'| < |K|$. Let k be the minimum of K and let $j \in [0, k)$ be such that $u_j = u_k$ (note that these exist by assumptions (i) and (ii)).

We claim that the word induced by $u_j, \ldots, u_{k-1}$ is good. Note first that for every $q \in [j, k-1)$ it holds that u_{q+1} is a direct descendant of u_q. Indeed, if not, then by the definition of a path through a cyclic proof, we have $u_{q+1} = f(u)$ for some $u \in \mathsf{dom}(f)$ such that u is a direct descendant of u_q. But then, by assumption (iii) above and the definition of a cyclic proof, we have $u_0 \leq u_{q+1} \leq u_q$. Hence, there must be some $p \leq q$ such that $u_p = u_{q+1}$. But $q+1 < k$, contradicting the minimality of k.

It follows by transitivity that $u_j \leq u_{k-1}$. Hence $u_k = u_j$ is not a direct descendant of u_{k-1}, which means that $u_k = f(u)$, where u is a direct descendant of u_{k-1}. Therefore, the word induced by $u_j, \ldots, u_{k-1}$ is the word induced by $[u_k, u)$, which is indeed good by the fact that (T, f) is a cyclic proof.

Next, we claim that the path $u_0, \ldots, u_{j-1}, u_k, \ldots u_{n-1}$ is good. Note that this is indeed a path through (T, f), by the fact that $u_j = u_k$. Hence our claim follows directly from the induction hypothesis. This application of the induction hypothesis is justified by the fact that, by the minimality of k in K, the index k does not appear in the analogous set, say K', of the path $u_0, \ldots, u_{j-1}, u_k, \ldots, u_{n-1}$.

Finally, by item (ii) of Definition I.2.7, we find that the word induced by $u_0, \ldots, u_{n-1}$ is good. □

I.2.22. PROPOSITION. *Let* P *be a simple proof system and let* (T, f) *be a cyclic* P*-proof. Then the unravelling* $\mathsf{un}(T, f)$ *is a regular* P*-proof.*

Proof:
We must show that every infinite branch β of $\mathsf{un}(T, f)$ induces a good infinite word. Note that any such β can be seen as an infinite path through (T, f). By Lemma I.2.20, the set $\mathrm{Inf}(\beta) \cap \mathsf{ran}(f)$ must contain a $\preceq$-greatest element u. For every $n \in \mathbb{N}$, let u_n be the n-th occurrence on β of u,. Then the path β can be written as the infinite concatenation of all finite paths $u_i, \ldots, u_{i-1}$. By Lemma I.2.21 each of these finite paths induced a good word. Hence it holds by item (i) of Definition I.2.7 that β induces a good infinite word. □

I.2.23. Corollary. *Let* P *be a simple proof system. Then for every frugal* P*-proof T there is a regular* P*-proof T' with equivalent root.*

Proof:
Take $T' = \mathsf{un}(T_0, f)$, where (T_0, f) is as given by Proposition I.2.12. By the previous proposition this is a P-proof, which is regular because it is the unravelling of a finite tree with back edges. □

I.3 From frugal to concise proofs

In this section we will show that for every frugal proof in some proof system P, there is a concise proof with the same conclusion. Roughly, our strategy will be to push the bad repeats upwards, until every *first* repeat must be good. Here we use the fact that, by frugality, there is a bound on the depth of the first repeat.

More precisely, we will measure non-conciseness by the depth of $\widehat{v}$, where v is a repeat such that the word induced by $[\widehat{v}, v)$ is bad and the minimal repeat below v is also bad. Recall that $\widehat{v}$ is the minimal $u < v$ such that $l(v) \equiv l(u)$, and that $\widehat{v}$ is called the *companion* of v.

I.3.1. Definition. Let v be a repeat of T. The *conciseness-rank* (or c-rank) of v is given by:

$$\mathsf{c}(v) := \begin{cases} \infty & \text{if the word induced by } [\widehat{v}, v) \text{ is good,} \\ \infty & \text{if the minimal repeat } u \leq v \text{ is a good repeat,} \\ |\widehat{v}| & \text{otherwise.} \end{cases}$$

The c-*rank* $\mathsf{c}(T)$ of T is defined to be the minimum of all the c-ranks of repeats of T (which we assume to be ∞ if T has no repeats).

Note that $\mathsf{c}(T) = \infty$ if and only if there is no bad minimal repeat, *i.e.* if and only if T is concise.

We define a preorder $\preceq$ on proofs such that, intuitively, $T_1 \preceq T_2$ means that T_1 has the same conclusion as T_2 and every sequent occurring in T_2 also occurs in T_1. The reason that we write $T_1 \preceq T_2$, instead of the other way around, is that we care about making proofs more concise. The higher in the $\preceq$-order, the sharper the bound on the depth of the first repeat.

Formally, we define $\preceq$ as follows.

I.3.2. DEFINITION. The preorder $\preceq$ on P-proofs is given by: $T_1 \preceq T_2$ if and only if both of the following hold:

(i) $l^{T_1}(\epsilon) \equiv l^{T_2}(\epsilon)$;

(ii) for every $v \in T_2$ there is a $u \in T_1$ such that $l^{T_1}(u) \equiv l^{T_2}(v)$.

Clearly, if T_1 is frugal and $T_1 \preceq T_2$, then T_2 is frugal as well. The following lemma makes formal the idea that any frugal proof of sufficiently high c-rank is concise.

I.3.3. LEMMA. *For any frugal* P*-proof* T_1 *there is a number* k *such that for every* T_2 *such that* $T_1 \preceq T_2$ *it holds that* $\mathsf{c}(T_2) \geq k$ *implies* $\mathsf{c}(T_2) = \infty$.

Proof:
Let k be the (finite) cardinality of $\{l(u) \mid u \in T_1\}/\equiv$ and let T_2 be a P-proof with $T_1 \preceq T_2$ and $\mathsf{c}(T_2) \geq k$. Let u be some minimal repeat of T_2. By the pigeonhole principle, we must have $k \geq |u| > |\widehat{u}|$. Hence, since $\mathsf{c}(T_2) \geq k$, the repeat u must be good. □

The previous lemma tells us that, for making a frugal proof T concise, it suffices to find an infinite chain

$$T := T_0 \preceq T_1 \preceq T_2 \preceq \cdots \tag{I.1}$$

such that $i < j$ implies $\mathsf{c}(T_i) < \mathsf{c}(T_j)$.

The following definition captures which nodes contribute to the c-rank of T.

I.3.4. DEFINITION. A repeat u of T is said to be *critical* if $\mathsf{c}(u) = \mathsf{c}(T)$.

We write $\widehat{C_T} \subseteq T$ for the set of companions of critical repeats of T:

$$\widehat{C_T} := \{\widehat{u} \mid \ u \text{ is a critical repeat of } T\}.$$

I.3.5. LEMMA. *If* $\mathsf{c}(T) < \infty$*, then* $\widehat{C_T}$ *is finite.*

Proof:
Since every node in $\widehat{C_T}$ has depth $\mathsf{c}(T)$, this follows from the fact that T is finitely branching. □

The following lemma is the first ingredient for building the chain (I.1). Its statement might be a bit opaque right now, but in the next lemma we will see that it is exactly what we need.

I.3.6. Lemma. *Suppose* $\mathsf{c}(T) < \infty$*. For every critical repeat* u *of* T *there is a node* $v \geq u$ *such that* $l(v) \equiv l(u)$*, the path* $[\widehat{u}, v)$ *is bad, and* $\mathsf{c}(\langle v \rangle) > 0$*.*

Proof:
Suppose, towards a contradiction, that there is no such $v \geq u$. As u is critical, we have $\mathsf{c}(u) = \mathsf{c}(T)$. But by hypothesis $\mathsf{c}(T) < \infty$, and thus we know that $[\widehat{u}, u)$ is bad. It follows that $\mathsf{c}(\langle u \rangle) = 0$, for otherwise we could just have chosen $v = u$.

Hence, in $\langle u \rangle$ there must be a witness of $\mathsf{c}(\langle u \rangle) = 0$. That is, in T there is a $u_1 > u$ such that $l(u) \equiv l(u_1)$ and $[u, u_1)$ is bad. Since $[\widehat{u}, u)$ is also bad, we find by Lemma I.2.2 that $[\hat{u}, u_1)$ is bad. Again, it follows that $\mathsf{c}(\langle u_1 \rangle) = 0$, for otherwise we could have set $v = u_1$.

Continuing in this way, we find an infinite upward path

$$[u_0, u_1) \cdot [u_1, u_2) \cdot [u_2, u_3) \cdots$$

where $u = u_0$. Since T is a proof, this path must be a final segment of some good branch. But then $[u_i, u_{i+1})$ must be good for some i (in fact, for infinitely many i), a contradiction. □

The next lemma builds upon the previous lemma. It either directly shows how to increase the c-rank of T, or, if not, at least tells us how to decrease some well-founded measure while leaving $\mathsf{c}(T)$ untouched.

I.3.7. Lemma. *Suppose that* $\mathsf{c}(T) < \infty$ *and that* u *is a critical repeat of* T*. Let* v *be as given by the previous lemma and consider the tree* $T' := T[\langle v \rangle / \widehat{u}]$*. Then* T' *is a* P*-proof such that* $T \preceq T'$ *and, moreover, either* $\mathsf{c}(T') > \mathsf{c}(T)$*, or* $\mathsf{c}(T') = \mathsf{c}(T)$ *and* $|\widehat{C_{T'}}| < |\widehat{C_T}|$*.*

Proof:
It is clear that T' is a P-proof such that $T \preceq T'$. Suppose that $\mathsf{c}(T') \not> \mathsf{c}(T)$, *i.e.* $\mathsf{c}(T') \leq \mathsf{c}(T)$. We will show that $\mathsf{c}(T') = \mathsf{c}(T)$ and $|\widehat{C_{T'}}| < |\widehat{C_T}|$.

For the former, note that it suffices to show that $\mathsf{c}(T) \leq \mathsf{c}(T')$. To that end, take an arbitrary critical repeat w in T'. We will show that $\mathsf{c}(T) \leq \mathsf{c}(w)$. Recall that by the definition of T' as a substitution, we either have $w = \widehat{u} \cdot x$ for some $x \in \langle v \rangle$, or $w = y$ for some $y \in T$ with $\widehat{u} \not\leq y$.

In the latter case, we have $\widehat{u} \not\leq w$ and, since $\widehat{w} < w$ in T', also $\widehat{u} \not\leq \widehat{w}$. Hence w is a repeat in T as well. Since $\mathsf{c}(w)$ only depends on w and its ancestors, the value of $\mathsf{c}(w)$ is the same regardless whether it is calculated in T or in T'. Hence the desired inequality $\mathsf{c}(T) \leq \mathsf{c}(w)$ simply follows from the definition of $\mathsf{c}(T)$.

For the rest of the proof of the claim that $\mathsf{c}(T) \leq \mathsf{c}(T')$, we can thus assume that w is of the form $w = \widehat{u} \cdot x$ for some $x \in \langle v \rangle$. We make another case distinction, namely on whether $\widehat{u} \leq \widehat{w}$. If so, we find, by the fact that u is critical in T and w is critical in T':

$$\mathsf{c}(T) = |\widehat{u}| \leq |\widehat{w}| = \mathsf{c}(T'),$$

and we are done.

Hence we may further assume that $\widehat{u} \nleq \widehat{w}$, *i.e.* $\widehat{w} < \widehat{u}$. The path under consideration in T' now looks like this:

$$\widehat{w} \text{ ——————— } \widehat{u} \text{ ——————— } w$$

We will consider the node $w^+ := v \cdot x \in T$. The corresponding path in T then looks like this:

$$\widehat{w} \text{ ——————— } \widehat{u} \text{ ——————— } u \text{ ——————— } v \text{ ——————— } w^+$$

where u and v are possibly equal. We prove several properties about w^+.

(i) w^+ is a repeat in T. Indeed, since $\widehat{w} < \widehat{u}$, we have $\widehat{w} \in T$. The result then follows from the fact that

$$l_T(\widehat{w}) = l_{T'}(\widehat{w}) \equiv l_{T'}(w) = l_{T'}(\widehat{u} \cdot x) = l_{\langle v \rangle}(x) = l_T(v \cdot x) = l_T(w^+),$$

and $\widehat{w} < w^+$.

(ii) $\widehat{w^+} = \widehat{w}$. Like the previous item, this follows from $l_T(w^+) \equiv l_T(w)$ together with $\widehat{w} < w^+$.

(iii) The minimal repeat below w^+ is bad. Indeed, we have that u is a repeat such that $u \leq v \leq w^+$. Thus, the minimal repeat below w^+ is the minimal repeat below u, which is bad because $\mathsf{c}(u) < \infty$.

(iv) $[\widehat{w}, w^+)$ is bad in T. To see this, note that $[\widehat{w}, w^+)$ can be written as

$$[\widehat{w}, w^+) = [\widehat{w}, \widehat{u}) \cdot [\widehat{u}, v) \cdot [v, w^+)$$

By the construction of Lemma I.3.6 the path $[\widehat{u}, v)$ in T is bad. Moreover, since w is critical in T', we have that $[\widehat{w}, w) = [\widehat{w}, \widehat{u}) \cdot [\widehat{u}, w)$ in T' is bad. Since the word induced by $[\widehat{w}, \widehat{u}) \cdot [\widehat{u}, w)$ in T' is the same as the word induced by $[\widehat{w}, \widehat{u}) \cdot [v, w^+)$ in T, the latter is bad as well. Hence it follows from Definition I.2.1.(iii) that $[\widehat{w}, w^+)$ must be bad.

By fact (i) above, we find that $\mathsf{c}(T) \leq \mathsf{c}(w^+)$. Moreover, by facts (iii) and (iv), we have $\mathsf{c}(w^+) = |\widehat{w^+}|$. Combining this with fact (ii), we obtain $\mathsf{c}(T) \leq |\widehat{w}|$. As w is critical by assumption, we have $\mathsf{c}(T') = |\widehat{w}|$ and thus $\mathsf{c}(T) \leq \mathsf{c}(T')$, as required.

It remains to show that $|\widehat{C_T}| > |\widehat{C_{T'}}|$. In fact, we will show that $\widehat{C_{T'}} \subseteq \widehat{C_T} \setminus \{\widehat{u}\}$. To this end, let $\widehat{y} \in \widehat{C_{T'}}$, where y is a critical repeat.

Suppose first, towards a contradiction, that $\widehat{y} = \widehat{u}$. Since $y > \widehat{y}$, we can write y as $y = \widehat{u} \cdot z$. By the definition of T', we have that $z \in \langle v \rangle$. Observe that z is a repeat in $\langle v \rangle$ with $\widehat{z} = \epsilon$ such that $[\epsilon, z)$ is bad. We claim that the minimal repeat below z in $\langle v \rangle$ is bad as well. To see this, consider the minimal repeat s below y in T'. As y is critical, we know that s must be a bad repeat. But since we have established that $\mathsf{c}(T') = \mathsf{c}(T) = |\widehat{u}|$, it follows that $s > \widehat{s} \geq \widehat{u}$. The situation so far can be depicted as follows. In T', we have

$$\epsilon \longrightarrow \widehat{u} \longrightarrow \widehat{s} \longrightarrow s \longrightarrow y = \widehat{u} \cdot z$$

where $\widehat{s}$ may be equal to $\widehat{u}$, and s may be equal to y. Hence we can write $\widehat{s}$ as $\widehat{s} = \widehat{u} \cdot a$, and s as $s = \widehat{u} \cdot a \cdot b$. In $\langle v \rangle$, this looks as follows:

$$\epsilon \longrightarrow a \longrightarrow a \cdot b \longrightarrow z$$

where a may be equal to ϵ, and $a \cdot b$ may be equal to z. It follows that $a \cdot b$ is the minimal repeat below z in $\langle v \rangle$. Moreover, this path is bad, since the path induced by $[a, a \cdot b)$ in $\langle v \rangle$ is the same as the path induced by $[\widehat{s}, s)$ in T'. But then $\mathsf{c}(z) = 0$ in $\langle v \rangle$, whence $\mathsf{c}(\langle v \rangle) = 0$, contradicting Lemma I.3.6.

Thus we must have $\widehat{y} \neq \widehat{u}$. As $|\widehat{y}| = |\widehat{u}|$, it follows that $\widehat{u} \not\leq \widehat{y}$ and thus $y, \widehat{y} \in T$, whence $\widehat{y} \in \widehat{C_T}$. □

We are finally ready to prove the main theorem of this intermezzo.

I.3.8. THEOREM. *For any frugal* P*-proof, there is a concise* P*-proof with equivalent root.*

Proof:
Let T be a frugal P-proof which is not concise. Let k be the number given by Lemma I.3.3. Since $\mathsf{c}(T) < \infty$, we can repeatedly apply I.3.7 to obtain a chain

$$T =: T_0 \preceq T_1 \preceq T_2 \preceq \cdots$$

We claim that for some n it must hold that $\mathsf{c}(T_n) \geq k$. Indeed, this follows from the fact that the well founded measure $|\widehat{C_{T_i}}|$ can only decrease finitely often. Hence, by Lemma I.3.3, there is some n such that $\mathsf{c}(T_n) > \infty$ and thus such that T_n is concise. □

I.4 Conclusion

We mention some questions for future work:

- Positional strategies in (parity) proof search games correspond to uniform, rather than concise proofs. Can we also obtain uniform proofs in our axiomatic setting?
- Is there a way of improving our axiomatisation such that also for the systems featuring a *reset* rule the good infinite words can be generated from the good finite words (like with our *simple* proof systems)?
- In the same spirit: which axioms do we need to add in order to be able to prove that the unravelling of a cyclic P-proof is always a P-proof? That is, which axioms should we add in order to prove a variant of Proposition I.2.22 without the requirement that the proof system P is simple?

Chapter 4

Filtration and canonical completeness for continuous modal μ-calculi

4.1 Introduction

In this chapter we take a more traditional approach to the proof theory of modal fixed point logics. Rather than working with Gentzen-style proof systems, we will directly prove the completeness of Hilbert-style proof systems. As in Chapter 3, the key tool will be a canonical model construction.

As already mentioned in the introduction of this thesis, the fact that compactness fails for modal fixed point logics, prevents the use of the (infinitary) canonical model construction, which is often used to prove completeness of basic modal logics. A common solution is to work with *finitary* canonical models, as was first done by Kozen & Parikh for PDL in [62] and by Segerberg for GL in [95]. In the book [45], Goldblatt applies the same procedure to several modal fixed point logics, including CTL.

These finitary canonical models are closely linked to the most common method of proving the finite model property for basic modal logic, *i.e.* the technique of *filtration*. This roughly works as follows. One begins by taking the canonical model $\mathbb{S}^{\mathsf{L}}$ of some non-compact logic L. Due to the compactness failure, this model is *non-standard*, meaning that the frame underlying $\mathbb{S}^{\mathsf{L}}$ fails to satisfy some desired properties. However, by applying filtration to $\mathbb{S}^{\mathsf{L}}$ one obtains a finite model whose underlying frame often does satisfy these desired properties

In the recent paper [56], Kikot, Shapirovsky, and Zolin prove a result of this kind that is relatively wide in scope. They show that if a basic modal logic L allows the method of filtration, then so does its expansion with the transitive closure modality. By iterating this procedure they show the same for the expansion of L by all modalities of PDL. Subsequently, if the original basic modal logic L moreover is canonical, the completeness of this PDL-expansion of L can be obtained by applying filtration to its canonical model.

As PDL can be seen a fragment of the modal μ-calculus [25], a natural question is whether similar techniques can be applied to a more expressive fragment. It is well known that the formula $\mu x \Box x$ provides a counterexample against the applicability of filtration, so any candidate fragment will have to omit this formula.

In this paper we consider the methods of filtration and canonical models for the continuous modal μ-calculus $\mu^c\mathtt{ML}$ (cf. Definition 2.1.41). Fontaine, in [41], shows that there are two equivalent ways to define $\mu^c\mathtt{ML}$. First semantically, as the fragment of the modal μ-calculus where the application of fixed point operators is restricted to formulas whose functional interpretation is Scott-continuous, rather than merely monotone. And second syntactically as, roughly, the fragment where the modal operator $\Box$ and the greatest fixed point operator ν are not allowed to occur in the scope of a μ-operator (and dually, $\Diamond$ and μ are not allowed in the scope of a ν-operator). To the best of our knowledge, the logic $\mu^c\mathtt{ML}$ was mentioned first in van Benthem [11] under the name 'ω-μ-calculus'. It is related, and perhaps equivalent in expressive power, to the logic of concurrent propositional dynamic logic, cf. Carreiro [23, section 3.2] for more information.

It is also worth mentioning here that imposing syntactic continuity restrictions on fixed point logics dates back at least 50 years. In [82], a syntactic continuity restriction was imposed on a first-order fixed point logic in order to identify a fragment that embeds into the infinitary logic $L_{\omega_1,\omega}$. More similar to our work is the least root calculus by Pratt in [85]. Like us, Pratt shows that his calculus admits filtration. The fact that Pratt's calculus is formulated as a least root calculus rather than a least fixed point calculus makes no substantial difference for our purposes. An important difference, however, is that Pratt's syntactic continuity restriction features an *aconjunctivity* restriction (see [59, Definition 3.3.4]). Because of this it aligns more closely to what in modern language is called the *completely additive* fragment of modal μ-calculus, as opposed to the continuous fragment. We refer the reader to, for instance, the paper [41] for a definition of the completely additive fragment. Also in [41], an example from [11] is given, showing that the completely additive fragment is strictly less expressive than the continuous fragment.

There are at least two reasons why the continuous μ-calculus is an interesting logic; first, the continuity condition that is imposed on the formation of fixed point formulas ensures that the construction of a definable fixed point using its ordinal approximations will always be finished after ω many steps. And second, in the same manner that the modal μ-calculus is the bisimulation-invariant fragment of monadic second-order logic [53], $\mu^c\mathtt{ML}$ has the same expressive power as *weak* monadic second-order logic, when it comes to bisimulation-invariant properties [24].

In the present chapter we show that we can add two more desirable properties to this list: (i) the Filtration Theorem holds for $\mu^c\mathtt{ML}$ and (ii) completeness for Kozen's axiomatisation adapted to sufficiently nice logics in the language of $\mu^c\mathtt{ML}$ can be proven using finitary canonical models.

4.2 Filtration

Filtration is a well-known method in the theory of basic modal logic. In this section we define filtration and related notions for the continuous modal μ-calculus and show that some of their most important properties transfer to this more expressive language. Introduced by Lemmon & Scott in [71], filtration is a technique for shrinking a Kripke model into a finite one, by identifying states that agree on the truth of some given finite set of formulas. The *Filtration Theorem* then states that the equivalence classes in the finite model satisfy the same formulas from this finite set as their members do in the original model. Filtration is an important tool for proving the finite model property and the decidability of modal logics. Not only does it entail the finite model property, but also the *small* model property: every satisfiable formula φ is satisfied in a model of size exponential in the size of φ. Satisfiability and validity can then be decided by simply checking all models up to this size. For an overview of recent developments in the theory of filtration, see [12].

Throughout this chapter we will without loss of generality assume that all formulas are tidy and in negation normal form (cf. Proposition 2.1.28).

Filtration As mentioned above, the main idea of the technique of filtration is to identify states of a model that agree on some finite set of formulas. This is captured by the following definition.

4.2.1. Definition. Let Σ be a set of formulas and let $\mathbb{S}$ be a Kripke model. The equivalence relation $\sim_{\Sigma}^{\mathbb{S}}$ is given by:

$$s \sim_{\Sigma}^{\mathbb{S}} s' \text{ if and only if } \mathit{Th}_{\mathbb{S}}(s) \cap \Sigma = \mathit{Th}_{\mathbb{S}}(s') \cap \Sigma.$$

A filtration of $\mathbb{S}$ through Σ is then obtained by taking, as set of states, the quotient set of the equivalence relation $\sim_{\Sigma}^{\mathbb{S}}$. Note that, although the following definition imposes restrictions on the accessibility relation and valuation of a filtration, it does leave room for multiple distinct filtrations through the same set of formulas. Recall that $\overline{\mathsf{FL}}$ was defined in Definition 2.1.29.

4.2.2. Definition. Let $\mathbb{S} = (S, R, V)$ be a Kripke model and let Σ be a finite and $\overline{\mathsf{FL}}$-closed set of formulas. A *filtration* of $\mathbb{S}$ *through* Σ is any model $\overline{\mathbb{S}} = (\overline{S}, \overline{R}, \overline{V})$ such that:

(i) $\overline{S} = S/{\sim_{\Sigma}^{\mathbb{S}}}$;

(ii) $R_{\Sigma}^{\mathsf{min}} \subseteq \overline{R} \subseteq R_{\Sigma}^{\mathsf{max}}$;

(iii) $\overline{V}(p) = \{\overline{s} : s \in V(p)\}$ for every $p \in \Sigma$.

where:

$$R^{\mathsf{min}}_{\Sigma} := \{(\overline{s}, \overline{t}) : \text{there are } s' \sim^{\mathbb{S}}_{\Sigma} s \text{ and } t' \sim^{\mathbb{S}}_{\Sigma} t \text{ such that } Rs't'\},$$
$$R^{\mathsf{max}}_{\Sigma} := \{(\overline{s}, \overline{t}) : \text{for all } \Box\varphi \in \Sigma; \text{ if } s \Vdash \Box\varphi, \text{ then } t \Vdash \varphi\}.$$

where $\overline{s}$ denotes the equivalence class with representative s.

4.2.3. Remark. The above definition of filtration is the standard definition for basic modal logic, as for instance found in [15, Definition 2.36]. Applying it to the continuous modal μ-calculus requires no adaptation.

Filtration Theorem for the continuous modal μ-calculus We will now prove the Filtration Theorem. It states that, restricted to the set Σ, the state s of a model and the state $\overline{s}$ of its filtration satisfy the same formulas.

4.2.4. Theorem (Filtration Theorem). *Let Σ be a finite and $\overline{\mathsf{FL}}$-closed set of formulas and let $\mathbb{S} = (S, R, V)$ be a Kripke model. For any filtration $\overline{\mathbb{S}} = (\overline{S}, \overline{R}, \overline{V})$ of $\mathbb{S}$ through Σ it holds that $Th_{\mathbb{S}}(s) \cap \Sigma = Th_{\overline{\mathbb{S}}}(\overline{s}) \cap \Sigma$ for every $s \in S$.*

Proof:
We must show that for every formula $\xi \in \Sigma$ and for every state $s \in S$ it holds that:

$$\mathbb{S}, s \Vdash \xi \Leftrightarrow \overline{\mathbb{S}}, \overline{s} \Vdash \xi.$$

Because Σ is negation closed, it suffices to prove just one direction of the bi-implication, which in our case will be the direction $\Rightarrow$. Throughout this proof we will write $\mathcal{G}$ for the game $\mathcal{E}(\xi, \mathbb{S})$ and $\overline{\mathcal{G}}$ for the game $\mathcal{E}(\xi, \overline{\mathbb{S}})$. As hypothesis we assume that $\exists$ has a winning strategy f in the game $\mathcal{G}$ initialised at position (ξ, s); we wish to show that $(\xi, \overline{s}) \in \text{Win}_{\exists}(\overline{\mathcal{G}})$.

The main idea of the proof is to obtain a winning strategy for $\exists$ in $\overline{\mathcal{G}}$ by playing a 'shadow match' in $\mathcal{G}$. That is, we will simulate in $\mathcal{G}$ every move played by $\forall$ in our $\overline{\mathcal{G}}$-match, and, to determine a move for $\exists$ in $\overline{\mathcal{G}}$, we copy the move dictated in $\mathcal{G}$ by the strategy f. If we manage to do this, then whenever the match in $\overline{\mathcal{G}}$ is at some position $(\varphi, \overline{s})$, the shadow match in $\mathcal{G}$ will be at a position (φ, s) (note that this is indeed the case for the initial positions). It turns out that this works well for all positions, except those of the form $(\Box\varphi, \overline{s})$. At those positions, a problem arises when $\forall$ chooses a position $(\varphi, \overline{t})$ such that $\overline{s}\overline{R}\overline{t}$, but not sRt. This move by $\forall$ in $\overline{\mathcal{G}}$ can then not be simulated in the shadow match, because (φ, t) is not an admissible move for $\forall$ in $\mathcal{G}$. However, using the fact that $\overline{R} \subseteq R^{\mathsf{max}}_{\Sigma}$, it will nonetheless hold that $\overline{s}\overline{R}\overline{t}$ and $(\Box\varphi, s) \in \text{Win}_{\exists}(\mathcal{G})$ together imply $(\varphi, t) \in \text{Win}_{\exists}(\mathcal{G})$. We will use this to initiate a *new* shadow match in $\mathcal{G}$ whenever we encounter a position of the form $(\Box\varphi, \overline{s})$. The key observation will be that we only need to initiate a new shadow match at most finitely often, because formulas of the form $\Box\varphi$ do not occur within the scope of least fixed point operators in the language $\mu^c\mathtt{ML}$.

More formally, we say for $I \in \omega \cup \{\omega\}$ that a $\overline{\mathcal{G}}$-match $\overline{\mathcal{M}} = (\varphi_i, \overline{t_i})_{i\in I}$ is *linked* to some $\mathcal{G}$-match $\mathcal{M} = (\psi_i, s_i)_{i\in I}$ whenever for every $i \in I$ it holds that $\varphi_i = \psi_i$ and $\overline{s_i} = \overline{t_i}$. Moreover, we say that $\overline{\mathcal{M}}$ *follows* $\mathcal{M}$ whenever some final segment of $\overline{\mathcal{M}}$ is linked to $\mathcal{M}$.

Claim. Let $\overline{\mathcal{M}}$ be a finite $\overline{\mathcal{G}}$-match that follows some f-guided $\mathcal{G}$-match $\mathcal{M}$, where f is a winning strategy for $\exists$ in $\mathcal{G}$ initialised at $\mathsf{first}(\mathcal{M})$. Then precisely one of the following holds:

- Both $\mathsf{last}(\mathcal{M})$ and $\mathsf{last}(\overline{\mathcal{M}})$ belong to $\exists$, and there is an admissible move $(\varphi_{i+1}, \overline{t_{i+1}})$ in $\overline{\mathcal{G}}$, such that $\overline{\mathcal{M}} \cdot (\psi_{i+1}, \overline{t_{i+1}})$ follows $\mathcal{M} \cdot (\varphi_{i+1}, s_{i+1})$, where (φ_{i+1}, s_{i+1}) is the move instructed by f in $\mathcal{G}$.
- Both $\mathsf{last}(\mathcal{M})$ and $\mathsf{last}(\overline{\mathcal{M}})$ belong to $\forall$, the formula in $\mathsf{last}(\mathcal{M})$ is *not* of the form $\Box\chi$, and for every admissible move $(\psi_{i+1}, \overline{t_{i+1}})$ for $\forall$ in $\overline{\mathcal{G}}$, there is an admissible move (φ_{i+1}, s_{i+1}) for $\forall$ in $\mathcal{G}$ such that $\overline{\mathcal{M}} \cdot (\psi_{i+1}, \overline{t_{i+1}})$ follows $\mathcal{M} \cdot (\varphi_{i+1}, s_{i+1})$.
- The formula in $\mathsf{last}(\mathcal{M})$ is of the form $\Box\chi$ and for every admissible move $(\psi_{i+1}, \overline{t_{i+1}})$ for $\forall$ in $\overline{\mathcal{G}}$, there is a position (φ_{i+1}, s_{i+1}) in $\mathcal{G}$ such that f is winning for $\exists$ in $\mathcal{G}@(\varphi_{i+1}, s_{i+1})$, and $\overline{\mathcal{M}} \cdot (\psi_{i+1}, \overline{t_{i+1}})$ follows (φ_{i+1}, s_{i+1}). Note that only in this case we lose the link with $\mathcal{M}$, and instead start following a new $\mathcal{G}$-match.

The above claim is proven by a case distinction on the main connective of the formula ψ in $\mathsf{last}(\mathcal{M})$. Since by assumption $\overline{\mathcal{M}}$ follows $\mathcal{M}$, the formula in $\mathsf{last}(\overline{\mathcal{M}})$ is also ψ.

Suppose first that ψ is a literal. Since f is assumed to be winning for $\exists$, the position $\mathsf{last}(\mathcal{M})$ must belong to $\forall$. Hence $\mathsf{last}(\mathcal{M}) = (\psi, s)$ for some $s \not\Vdash \psi$. Because $\psi \in \mathsf{FL}(\xi)$, we have $\psi \in \Sigma$. It follows by the restriction on $\overline{V}$ that $[s] \not\Vdash \psi$. Hence $\forall$ has no admissible move in $\overline{\mathcal{G}}$ and the claim holds vacuously.

Now suppose that ψ is of the form $\psi_1 \vee \psi_2$. Let (ψ_i, s) be the position instructed by f in $\mathcal{G}$. Then $\mathsf{last}(\overline{\mathcal{M}}) = (\psi, \overline{s})$ and thus $\exists$ can simply choose $(\psi_i, \overline{s})$ in $\overline{\mathcal{G}}$. The case where ψ is of the form $\psi_1 \wedge \psi_2$ is similar. Indeed, if $\forall$ chooses (ψ_i, s) in $\mathcal{G}$, he can also choose $(\psi_i, \overline{s})$ in $\overline{\mathcal{G}}$.

Now suppose that $\mathsf{last}(\mathcal{M})$ is of the form $(\Diamond\theta, s_n)$. Let (θ, s_{n+1}) be the next move instructed by the assumed winning strategy f. Then $s_n R s_{n+1}$ and thus, because $R^{\mathsf{min}}_{\Sigma} \subseteq \overline{R}$ and $s_n \sim t_n$, we have $\overline{t_n} \overline{R} \overline{s_{n+1}}$. Therefore in $\overline{\mathcal{G}}$ we have that $\exists$ can simply choose the position $(\theta, \overline{s_{n+1}})$.

If $\mathsf{last}(\mathcal{M})$ is of the form $(\Box\theta, t_n)$, consider the move $(\theta, \overline{t_{n+1}})$ chosen by $\forall$ in $\overline{\mathcal{G}}$. We have,

$$
\begin{aligned}
(\Box\theta, t_n) \in \mathrm{Win}_{\exists}(\mathcal{G}) &\Rightarrow \mathbb{S}, t_n \Vdash \Box\theta && \text{(definition)} \\
&\Rightarrow \mathbb{S}, t_{n+1} \Vdash \theta && (\overline{t_n} R^{\mathsf{max}}_{\Sigma} \overline{t_{n+1}},\ \Box\theta \in \Sigma) \\
&\Rightarrow (\theta, t_{n+1}) \in \mathrm{Win}_{\exists}(\mathcal{G}). && \text{(definition)}
\end{aligned}
$$

Thus we may choose (θ, t_{n+1}) as the new match that is followed by $\overline{\mathcal{M}} \cdot (\theta, \overline{t_{n+1}})$.

Using the fact that (ξ, s) is linked to $(\xi, \overline{s})$ as induction base, and the above claim as induction step, we obtain a strategy g for $\exists$ in $\overline{\mathcal{G}}$ initialised at $(\xi, \overline{s})$. We claim that g is a winning strategy. Indeed, if a g-guided match $\overline{\mathcal{M}}$ ends in finitely many steps, then $\forall$ must have gotten stuck.

If a g-guided match $\overline{\mathcal{M}}$ lasts infinitely long, then by item (1) of Lemma 2.1.43, there must be some point after which either only μ-variables, or only ν-variables, are unfolded. In the latter case the match is indeed winning for $\exists$. We will now argue that the former case cannot occur. The reason is that, if from some point on in $\overline{\mathcal{M}}$ only μ-variables are unfolded, then, by item (2) of Lemma 2.1.43, from some point no formula of the form $\Box\theta$ will occur. By construction, this means that the infinite $\overline{\mathcal{G}}$-match $\overline{\mathcal{M}}$ follows an infinite $\mathcal{G}$-match $\mathcal{M}$ which is guided by a strategy f for $\exists$, such that f is winning at $\mathsf{first}(\mathcal{M})$. But this is a contradiction, because the match $\mathcal{M}$, by the fact that it is linked to an infinite final segment of $\overline{\mathcal{M}}$, contains infinitely many μ-unfoldings. □

4.2.5. Remark. Note that the above argument would not go through for the alternation-free modal μ-calculus, since we would no longer be able to guarantee that we create at most finitely many shadow matches in the case of infinitely many μ-unfoldings. As mentioned above, a well-known counterexample to the Filtration Theorem for the alternation free modal μ-calculus is the formula $\mu x \Box x$.

Admissibility of filtration Having established that filtrations preserve satisfaction of $\mu^c\mathtt{ML}$-formulas, we will now investigate to which classes of models filtration can be applied.

4.2.6. Definition. A class of models M is said to *admit filtration* with respect to a language D if for every model $\mathbb{S}$ in M and every finite set of D-formulas Σ, the class M contains a filtration of $\mathbb{S}$ through some finite $\overline{\mathsf{FL}}$-closed set $\Theta \supseteq \Sigma$. A class of frames F is said to *admit filtration* if the class of models $\{(S, R, V) : (S, R) \in \mathsf{F}\}$ does.

One might expect that admitting filtration with respect to the basic modal language is a weaker property than admitting filtration with respect to a proper extension of the language. However, for the language $\mu^c\mathtt{ML}$ it turns out that this is not the case, at least for classes of models that are *closed under substitution*.

Recall that a *substitution* is a function $\sigma : \mathsf{P} \to \mu^c\mathtt{ML}$. For $\varphi \in \mathtt{ML}$, we write $\sigma(\varphi)$ for the result of applying the substitution σ to φ. That is, we take the unique extension $\sigma : \mathtt{ML} \to \mu^c\mathtt{ML}$ which commutes with the propositional and modal operators. Note that, because we are working in negation normal form, we define $\sigma(\overline{p})$ as $\overline{\sigma(p)}$, where the $\overline{}$ is the definable negation operator from Section 2.1.2, and not the explicit negation symbol $\neg$. We will only apply these substitutions to basic modal formulas, so we do not have to worry about free variables being captured. Let us verify that substitution commutes with negation in our setting.

4.2.7. Lemma. *Let $\sigma : \mathsf{P} \to \mu^c\mathtt{ML}$ be a substitution. Then for every* ML*-formula φ it holds that $\sigma(\overline{\varphi}) = \overline{\sigma(\varphi)}$.*

Proof:
We proceed by induction on φ. If φ is a literal, the statement holds by definition. We write $\circ$ for a connective in $\{\vee, \wedge\}$ and $\overline{\circ}$ for its dual. Then we have

$$\begin{aligned}\sigma(\overline{\psi_1 \circ \psi_2}) = \sigma(\overline{\psi_1} \mathbin{\overline{\circ}} \overline{\psi_2}) &= \sigma(\overline{\psi_1}) \mathbin{\overline{\circ}} \sigma(\overline{\psi_2}) \\ &= \overline{\sigma(\psi_1)} \mathbin{\overline{\circ}} \overline{\sigma(\psi_2)} = \overline{\sigma(\psi_1) \circ \sigma(\psi_2)} = \overline{\sigma(\psi_1 \circ \psi_2)},\end{aligned}$$

as required. Similarly, writing $\triangle$ for a modality in $\{\Diamond, \Box\}$ and $\overline{\triangle}$ for its dual, we find

$$\sigma(\overline{\triangle\psi}) = \sigma(\overline{\triangle}\overline{\varphi}) = \overline{\triangle}\sigma(\overline{\varphi}) = \overline{\triangle}\,\overline{\sigma(\varphi)} = \overline{\triangle\sigma(\varphi)},$$

as desired. □

Given a model $\mathbb{S} = (S, R, V)$, we let a substitution σ act on $\mathbb{S}$ by setting $\mathbb{S}[\sigma] := (S, R, V[\sigma])$, where $V[\sigma]$ is given by $V[\sigma](p) := [\![\sigma(p)]\!]^{\mathbb{S}}$. A class M of models is said to be *closed under substitution* if $\mathbb{S}[\sigma] \in \mathsf{M}$ for every substitution σ and model $\mathbb{S} \in \mathsf{M}$. Note that for any class of frames F, the class of models based on a frame in F is closed under substitution.

4.2.8. Remark. The following proposition resembles Theorem 3.8 of [56]. One important difference in its statement is that we work with the standard notion of filtration, rather than their notion of *definable filtration*. An important difference in its proof is that our translation acts only on fixed point formulas, and commutes with all other operators. In contrast, the translation in [56] assigns a propositional variable q_φ to each formula φ.

4.2.9. Proposition. *Let M be a class of models which is closed under substitution. If M admits filtration with respect to* ML*, then also with respect to $\mu^c\mathtt{ML}$.*

Proof:
Let Σ be a finite set of $\mu^c\mathtt{ML}$-formulas. Without loss of generality we may assume that Σ is $\overline{\mathsf{FL}}$-closed. Since the assumption only tells us that M admits filtration with respect to ML, we want to represent Σ by a set of ML-formulas. To this end, we let $\varphi_1, \dots, \varphi_n$ be an enumeration of the μ-formulas in Σ. Note that, by negation closure, it follows that $\overline{\varphi_1}, \dots, \overline{\varphi_n}$ is an enumeration of the ν-formulas in Σ. For every formula φ_i, we pick a unique propositional variable p_i not occurring in Σ.

Now let $\tau : \Sigma \to \mathtt{ML}$ be the translation that commutes with all propositional and modal operators, and acts on fixed point operators in the following way:

$$\tau(\varphi_i) := p_i, \qquad \tau(\overline{\varphi_i}) = \overline{p_i}.$$

We let $\sigma : \mathsf{P} \to \mu^c\mathtt{ML}$ be the substitution given by $\sigma(p) := \varphi_i$ if $p = p_i$, and $\sigma(p) := p$ otherwise. We will use the following two facts:

(i) $\sigma(\tau(\xi)) = \xi$, for every $\xi \in \Sigma$.

(ii) $\mathbb{S}, s \Vdash \sigma(\varphi) \Leftrightarrow \mathbb{S}[\sigma], s \Vdash \varphi$, for every $\varphi \in \mu^c\mathtt{ML}$. (Substitution Lemma)

Fact (ii) is a standard lemma in modal logic, and fact (i) is proven by induction on ξ. If ξ is a literal, then $\sigma(\tau(\xi)) = \sigma(\xi) = \xi$. The modal and propositional cases follow from the fact that both τ and σ commute with those operators. The fixed point cases are given by:

$$\sigma(\tau(\varphi_i)) = \sigma(p_i) = \varphi_i, \qquad \sigma(\tau(\overline{\varphi_i})) = \sigma(\overline{p_i}) = \overline{\varphi_i}.$$

By hypothesis, there is a filtration $\overline{\mathbb{S}[\sigma]}$ of $\mathbb{S}[\sigma]$ through some finite $\overline{\mathsf{FL}}$-closed set Θ such that $\tau[\Sigma] \subseteq \Theta \subseteq \mathtt{ML}$. We claim that $\overline{\mathbb{S}[\sigma]}$ is simultaneously a filtration of $\mathbb{S}$ through $\sigma[\Theta]$. This finishes the proof as, by (i) above, we have $\Sigma \subseteq \sigma[\Theta] \subseteq \mu^c\mathtt{ML}$.

Before we show that $\overline{\mathbb{S}[\sigma]}$ is indeed a filtration of $\mathbb{S}$ through $\sigma[\Theta]$, we will first show that $\sigma[\Theta]$ is $\overline{\mathsf{FL}}$-closed. For negation closure, note that $\sigma(\varphi) \in \sigma[\Theta]$ implies $\sigma(\overline{\varphi}) \in \sigma[\Theta]$ and thus, by Lemma 4.2.7, also $\overline{\sigma(\varphi)} \in \sigma[\Theta]$. Likewise, the propositional and modal clauses follow from the FL-closure of Θ and the commuting properties of σ. For the fixed point formulas, we argue as follows. Suppose $\eta x\psi \in \sigma[\Theta]$. Then either $\eta x\psi = \sigma(p_i)$ or $\eta x\psi = \sigma(\overline{p_i})$. In both cases we find $\eta x\psi \in \Sigma$, whence $\psi[\eta x\psi/x] \in \Sigma$. So $\sigma(\tau(\psi[\eta x\psi/x])) = \psi[\eta x\psi/x] \in \sigma[\Theta]$, as required.

Now let us write $(\overline{S}, \overline{R}, \overline{V})$ for $\overline{\mathbb{S}[\sigma]}$. We first show that $\overline{S} = S/{\sim^{\mathbb{S}}_{\sigma[\Theta]}}$. Since we know that $\overline{S} = S/{\sim^{\mathbb{S}[\sigma]}_{\Theta}}$, it suffices to show that $\sim^{\mathbb{S}}_{\sigma[\Theta]} = \sim^{\mathbb{S}[\sigma]}_{\Theta}$. But this follows directly from the fact that, by the substitution lemma,

$$Th_{\mathbb{S}}(s) \cap \sigma[\Theta] = Th_{\mathbb{S}[\sigma]}(s) \cap \Theta$$

for every $s \in S$. From $\sim^{\mathbb{S}}_{\sigma[\Theta]} = \sim^{\mathbb{S}[\sigma]}_{\Theta}$ we moreover obtain that $R^{\mathsf{min}}_{\sigma[\Theta]} \subseteq \overline{R}$.

We claim that $\overline{R} \subseteq R^{\mathsf{max}}_{\sigma[\Theta]}$. To this end, suppose that $\overline{s}\overline{R}\overline{t}$ and $\Box\varphi \in \sigma[\Theta]$ is such that $\mathbb{S}, s \Vdash \Box\varphi$. Let $\psi \in \Theta$ be such that $\sigma(\psi) = \Box\varphi$. By definition ψ must be of the form $\Box\delta$ with $\sigma(\delta) = \varphi$. The substitution lemma gives $\mathbb{S}[\sigma], s \Vdash \Box\delta$, whence, since $\overline{\mathbb{S}[\sigma]}$ is a filtration, it follows that $\mathbb{S}[\sigma], t \Vdash \delta$. Thus $\mathbb{S}, t \Vdash \sigma(\delta)$, as required.

Finally, we must show that for every $p \in \sigma[\Theta]$ it holds that

$$\overline{V}(p) = \{\overline{s} : s \in V(p)\}.$$

This follows from the fact that none of the p_i belongs to $\sigma[\Theta]$. Hence, we have $V(p) = V[\sigma](p)$ and $p \in \Theta$. Thus, we find

$$\overline{V}(p) = \{\overline{s} : s \in V[\sigma](p)\} = \{\overline{s} : s \in V(p)\},$$

as required. □

Note that the above proof does not rely on any specific properties of the language $\mu^c\mathtt{ML}$. In fact, it could also have been carried out for the full language $\mu\mathtt{ML}$ of the modal μ-calculus.

In the presence of the Filtration Theorem, we obtain the finite model property as a corollary.

4.2.10. Corollary (Finite Model Property). *Let* M *be a class of models that is closed under substitution and admits filtration with respect to* $\mathtt{ML}$*, and let* φ *be a formula of the continuous* μ*-calculus. Then* φ *is valid in every model in* M *if and only if* φ *is valid in every finite model in* M*.*

Proof:
Let φ be a formula such that $\mathbb{S} \not\models \varphi$ for some $\mathbb{S} \in \mathsf{M}$. By Proposition 4.2.9 and the assumption that M admits filtration with respect to $\mathtt{ML}$, there is a filtration $\overline{\mathbb{S}}$ of $\mathbb{S}$ through some finite $\Sigma \supseteq \{\varphi\}$ such that $\overline{\mathbb{S}} \in \mathsf{M}$. Observe that the number of states of $\overline{\mathbb{S}}$ is at most 2^Σ and thus finite. By Theorem 4.2.4, it holds that $\overline{\mathbb{S}} \not\models \varphi$, as required. □

For instance, since the class of symmetric frames admits filtration with respect to the basic modal language, we find that the continuous modal μ-calculus has the finite model property over this class.

4.3 Canonical completeness

In this section we prove our completeness result. In the first paragraph we will define the logics that our completeness proof applies to, which we shall call *μ_c-logics.* The paragraph thereafter defines the finitary canonical models of an arbitrary μ_c-logic L and proves the Truth Lemma. In the third paragraph we will show that a finitary canonical model can be obtained for the logic $\mu_c\mathsf{L}$, where L is any canonical basic modal logic such that the class of L-frames admits filtration. As a direct consequence we obtain that $\mu_c\mathsf{L}$ is sound and complete with respect to the class of L-frames.

Axiomatisation We will now tailor Kozen's axiomatisation of the full modal μ-calculus (cf. Section 2.3.1) to the continuous modal μ-calculus. Recall that $\bot$, $\rightarrow$, and $\leftrightarrow$ are definable in our language.

4.3.1. Definition. The logic $\mu_c\mathsf{K}$ is the least logic containing the following axioms and closed under the following rules.

Axioms.

1. A complete set of axioms for classical propositional logic.
2. Normality: $\neg\Diamond\top$.
3. Additivity: $\Diamond(p \vee q) \leftrightarrow (\Diamond p \vee \Diamond q)$.
4. Dual for $\Box$: $\Box p \leftrightarrow \neg\Diamond\neg p$.
5. Dual for ν: $\nu x\varphi \leftrightarrow \neg\mu x\neg\varphi[\neg x/x]$.
6. For every $\varphi \in \mathsf{Con}_x(\mu\mathtt{ML}) \cap \mu^c\mathtt{ML}$, the prefixed point axiom:

$$\varphi[\mu x\varphi/x] \rightarrow \mu x\varphi.$$

Rules.

1. Modus Ponens: from $\varphi \rightarrow \psi$ and φ, derive ψ.
2. Monotonicity: from $\varphi \rightarrow \psi$, derive $\Diamond\varphi \rightarrow \Diamond\psi$.
3. Uniform Substitution: from φ, derive $\varphi[\psi/x]$.
4. The least prefixed point rule: from $\varphi[\gamma/x] \rightarrow \mathcal{M}$ with $\varphi \in \mathsf{Con}_x(\mu\mathtt{ML}) \cap \mu^c\mathtt{ML}$, derive $\mu x\varphi \rightarrow \gamma$.

We will consider axiomatic extensions of $\mu_c\mathsf{K}$ that are closed under the rules above. We will use *μ_c-logic* to refer to such an extension. The term *logic* will be used to refer to any normal modal logic. If L is a logic in the basic modal language, we use $\mu_c\mathsf{L}$ to denote the least μ_c-logic containing L. Moreover, we will use $\mathsf{Mod}(\mathsf{L})$ ($\mathsf{Fr}(\mathsf{L})$) to denote the class of models (frames) on which every formula in L is valid. If (S, R, V) belongs to $\mathsf{Mod}(\mathsf{L})$ ((S, R) belongs to $\mathsf{Fr}(\mathsf{L})$) we say that (S, R, V) is an L-model ((S, R) is an L-frame) and write $(S, R, V) \models \mathsf{L}$ ($(S, R) \models \mathsf{L}$).

Recall that in Proposition 2.1.28, we showed that for every formula φ, there is an equivalent formula $\mathsf{nnf}(\varphi)$ which is in negation normal form. The followig lemma, which can be proven by an easy induction on formulas, states that if φ belongs to $\mu^c\mathtt{ML}$, then so does $\mathsf{nnf}(\varphi)$. Moreover, the two formulas φ and $\mathsf{nnf}(\varphi)$ are not only semantically equivalent, but also provably in $\mu_c\mathsf{K}$.

4.3.2. Lemma. *Let $\varphi \in \mu^c\mathtt{ML}$. Then $\mathsf{nnf}(\varphi) \in \mu^c\mathtt{ML}$ and $\mu_c\mathsf{K} \vdash \varphi \leftrightarrow \mathsf{nnf}(\varphi)$.*

The above lemma allows us for the rest of this section to restrict attention to formulas in negation normal form.

Finitary canonical models For the entirety of this paragraph we fix an arbitrary μ_c-logic L. A set Γ of formulas is said to be L-*inconsistent* whenever $\mathsf{L} \vdash (\gamma_1 \wedge \ldots \wedge \gamma_n) \rightarrow \bot$ for some $\gamma_1, \ldots, \gamma_n \in \Gamma$. We say of a formula φ that it is L-inconsistent if the set $\{\varphi\}$ is.

4.3.3. Definition. A set of formulas Γ is called *maximally* L*-consistent*indexL-inconsistent!maximally if it is consistent and maximal in that respect, *i.e.* for every other set of formulas Γ':

$$\text{If } \Gamma \subset \Gamma', \text{ then } \Gamma' \text{ is } \mathsf{L}\text{-inconsistent.}$$

Note that, because we work with a countably infinite set P of propositional variables, an L-maximally consistent set of formulas is necessarily countably infinite. The proofs of the following two lemmas are standard and therefore left to the reader. The first lemma is often called Lindenbaum's Lemma.

4.3.4. Lemma (Lindenbaum). *Every* L*-consistent set* Γ *of* $\mu^c\mathtt{ML}$*-formulas has a maximally* L*-consistent extension* $\overline{\Gamma} \supseteq \Gamma$ *of* $\mu^c\mathtt{ML}$*-formulas.*

4.3.5. Lemma. *Let* Γ *be a maximally* L*-consistent set. Then:*

(i) *If* $\mathsf{L} \vdash \varphi$*, then* $\varphi \in \Gamma$*;*

(ii) $\overline{\varphi} \in \Gamma$ *if and only* $\varphi \notin \Gamma$*;*

(iii) $\varphi \vee \psi \in \Gamma$ *if and only* $\varphi \in \Gamma$ *or* $\psi \in \Gamma$*;*

(iv) $\mu x\varphi \in \Gamma$ *if and only if* $\varphi[\mu x\varphi/x] \in \Gamma$.

4.3.6. Definition. Let Σ be a finite $\overline{\mathsf{FL}}$-closed set of formulas. A *model over* Σ *with respect to* L is any model (S, R, V) such that:

- $S = \{\Gamma \cap \Sigma : \Gamma \text{ is maximally } \mathsf{L}\text{-consistent}\}$.
- $R^{\mathsf{min}}_{\mathsf{L}} \subseteq R \subseteq R^{\mathsf{max}}_{\mathsf{L}}$, where:

$$AR^{\mathsf{min}}_{\mathsf{L}}B :\Leftrightarrow \bigwedge A \wedge \Diamond \bigwedge B \text{ is } \mathsf{L}\text{-consistent}$$
$$AR^{\mathsf{max}}_{\mathsf{L}}B :\Leftrightarrow \text{ for all } \Box\varphi \in \Sigma : \ \Box\varphi \in A \Rightarrow \varphi \in B.$$

- $V(p) = \{A \in S : p \in A\}$ for all $p \in \Sigma$.

For A some finite set of formulas, we will usually write ψ_A for the conjunction $\bigwedge A$. In the following we will assume a fixed model over some finite and $\overline{\mathsf{FL}}$-closed set Σ with respect to L, which will be denoted by $\mathbb{S} = (S, R, V)$. If we refer to provability or consistency, this will be tacitly assumed to be with respect to the logic L.

The idea of our completeness proof is to show a *Truth Lemma* for $\mathbb{S}$. More precisely, we will show that $\varphi \in A$, implies $\mathbb{S}, A \Vdash \varphi$ for every $A \in S$. The first step is the following lemma, often called the Existence Lemma. Since it is a standard lemma in the context of (finitary) canonical models for modal logics, the proof is left to the reader.

4.3.7. Lemma. *For any formula* $\varphi \in \mu_c\mathtt{ML}$ *and state* $A \in S$*:*

$\psi_A \wedge \Diamond\varphi$ *is consistent if and only if* $\psi_B \wedge \varphi$ *is consistent for some* ARB.

In particular, it follows that for all $\Diamond\varphi \in \Sigma$ we have $\Diamond\varphi \in A$ if and only if $\varphi \in B$ for some ARB.

The following lemma is a consequence of the fact that Σ is negation closed.

4.3.8. Lemma. *For every* $A, B \in S$ *it holds that* $\psi_A \wedge \psi_B$ *is consistent iff* $A = B$.

Given a finite collection U of finite sets of formulas, we write ψ_U for the disjunction of all ψ_X for $X \in U$, *i.e.*

$$\psi_U = \bigvee_{X \in U} \psi_X.$$

Note that by the previous lemma, for any $U \subseteq S$ and $A \in S$, the formula $\psi_U \wedge \psi_A$ is consistent if and only if $A \in U$.

The following lemma, often called the Context Lemma, will be very useful. It was originally proven in [59], where it appears as Proposition 5.(vi).

4.3.9. Lemma. *If* $\gamma \wedge \mu x\varphi$ *is consistent, then so is* $\gamma \wedge \varphi[\mu x.\overline{\gamma} \wedge \varphi/x]$.

4.3.10. Remark. We provide some intuition for the lemma by sketching why it holds semantically. We argue by contraposition. So suppose that $\gamma \wedge \varphi[\mu x.\overline{\gamma} \wedge \varphi/x]$ is invalid. Given an arbitrary model $\mathbb{S}$ and state s of $\mathbb{S}$ such that $\mathbb{S}, s \Vdash \gamma$, we will give a winning strategy for $\forall$ in the game $\mathcal{E}@(\mu x\varphi, \mathbb{S})@(\mu x\varphi, s)$. Note that the match immediately proceeds to the position $(\varphi[\mu x\varphi/x], s)$. Moreover, by assumption, we have a winning strategy for $\forall$ in the game $\mathcal{E}@(\varphi[\mu x.\overline{\gamma} \wedge \varphi/x], \mathbb{S})@(\varphi[\mu x.\overline{\gamma} \wedge \varphi/x], s)$. The idea is to exploit the similarity of the two games by applying the winning strategy for $\forall$ in the second game in the first game, and playing a shadow match in the second game. If the first match winds up at a position of the form $(\mu x\varphi, t)$, then the second match is at a position of the form $(\mu x.\overline{\gamma} \wedge \varphi, t)$, Hence, since the second match is guided by a winning strategy for $\forall$, it holds that $\mathbb{S}, t \Vdash \gamma$, and we can repeat the same argument. We leave it to the reader to convince themselves that this indeed describes a winning strategy for $\forall$ in $\mathcal{E}(\mu x\varphi, \mathbb{S})@(\mu x\varphi, s)$.

In our completeness proof we will apply the Context Lemma with ψ_U in place of γ, where U is some set of states of $\mathbb{S}$. This will help us to show the Truth Lemma for μ-formulas. For instance, given some $\mu x\varphi \in A$, the idea will be to show that $\mathbb{S}, A \Vdash \mu x\varphi$ by describing a winning strategy for $\exists$ in the game $\mathcal{E}(\mu x\varphi, \mathbb{S})@(\mu x\varphi, A)$. Since $\mu x\varphi \in A$, we have that $\psi_A \wedge \mu x\varphi$ is consistent, whence by the Context Lemma $\psi_A \wedge \varphi[\mu x.\overline{\psi_A} \wedge \varphi/x]$ is consistent. This allows us to construct the strategy for $\exists$ in such a way that the position $(\mu x\varphi, A)$ will not be reached again. By iterating this argument, and the fact that $\mathbb{S}$ only has finitely many states, it follows that $\mu x\varphi$ will be unfolded only finitely often, which is an important step in showing that the strategy is indeed winning for $\exists$.

To keep track of the extra information provided by the Context Lemma, we will use the following syntax extension.

4.3.11. Definition. We define the syntax $\mu_c^a\mathtt{ML}$ of *adorned* formulas in exactly the same way as $\mu^c\mathtt{ML}$, but now letting the least fixed point operators to be adorned by a set U of states of $\mathbb{S}$.

4.3.12. Example. $\mu^{\emptyset}x.\Diamond x \vee p$ and $\mu^{\{A,B\}}x.\Diamond x \vee p$ are adorned formulas, where A and B are elements of S.

4.3.13. Definition. The *interpretation function* $\iota : \mu_c^a\mathtt{ML} \to \mu^c\mathtt{ML}$ acts on least fixed point operators by

$$\iota(\mu^U x\varphi) := \mu x.\overline{\psi_U} \wedge \iota(\varphi),$$

and commutes with all other operators.

The *forgetful translation* $\cdot^- : \mu_c^a\mathtt{ML} \to \mu^c\mathtt{ML}$ acts on least fixed point operators by

$$(\mu^U x\varphi)^- := \mu x\varphi^-,$$

and commutes with all other operators.

The following lemma collects some useful facts concerning the provability of adorned formulas.

4.3.14. Lemma. *For every* $\varphi \in \mu_c^a\mathtt{ML}$*:*

(i) If $\psi_A \wedge \iota(\mu^U x\varphi)$ *is consistent, then so is* $\psi_A \wedge \iota(\varphi)[\iota(\mu^{U\cup\{A\}}x\varphi)/x]$*.*

(ii) $\vdash \iota(\varphi) \to \varphi^-$*.*

Proof:
For (i), suppose that $\psi_A \wedge \iota(\mu^U x\varphi)$ is consistent. Writing out the definition of ι, this means that $\psi_A \wedge \mu x.\overline{\psi_U} \wedge \iota(\varphi)$ is consistent. Hence, we can apply the Context Lemma, to obtain that

$$\psi_A \wedge \overline{\psi_U} \wedge \iota(\varphi)[\mu x.\overline{\psi_A} \wedge \overline{\psi_U} \wedge \iota(\varphi)/x]$$

is consistent as well. In particular $\psi_A \wedge \overline{\psi_U}$ is consistent and therefore $A \notin U$. But this means that $\psi_A \wedge \psi_U$ is propositionally equivalent to ψ_A. Moreover, $\overline{\psi_A} \wedge \overline{\psi_U}$ is propositionally equivalent to $\overline{\psi_{U\cup\{A\}}}$. We thus find that

$$\psi_A \wedge \iota(\varphi)[\mu x.\overline{\psi_{U\cup\{A\}}} \wedge \iota(\varphi)/x]$$

is consistent. By the definition of ι this means that $\psi_A \wedge \iota(\varphi)[\iota(\mu^{U\cup\{A\}}x\varphi)/x]$ is consistent, as required.

We prove (ii) by induction on φ. The only non-trivial case is where φ is of the form $\mu^U x\delta$, as in all other cases the two translations behave equally. So suppose that $\vdash \iota(\delta) \to \delta^-$. Then by propositional reasoning, it also holds that $\vdash \overline{\psi_U} \wedge \iota(\delta) \to \delta^-$. Using uniform substitution, we have

$$\vdash \overline{\psi_U} \wedge \iota(\delta)[\mu x \delta^-/x] \to \delta^-[\mu x \delta^-/x].$$

Moreover, the prefixed point axiom gives $\vdash \delta^-[\mu x \delta^-/x] \to \mu x \delta^-$ and thus, by propositional reasoning,

$$\vdash \overline{\psi_U} \wedge \iota(\delta)[\mu x \delta^-/x] \to \mu x \delta^-.$$

Applying the least prefixed point rule, we obtain $\vdash \mu x(\overline{\psi_U} \wedge \iota(\delta)) \to \mu x \delta^-$, or in other words,

$$\vdash \iota(\mu^U x \delta) \to (\mu^U x \delta)^-,$$

as required. □

We are now ready to prove the main result of this section. Note that its proof very much resembles the proof of the Filtration Theorem (Theorem 4.2.4 above). We will further comment on this resemblance in the conclusion of the present chapter.

4.3.15. LEMMA (Truth Lemma). *For every $A \in S$: $\xi \in A$ implies $A \in [\![\xi]\!]$.*

Proof:
We will prove this by directly defining a winning strategy f for $\exists$ in the game $\mathcal{E}(\xi, \mathbb{S})$ initialised at (ξ, A). By induction on $|\mathcal{M}|$, we will, for every f-guided partial $\mathcal{E}$-match $\mathcal{M}$, simultaneously define the following:

- a formula $\varphi_{\mathcal{M}} \in \mu_c^a\mathtt{ML}$;
- a move $f(\mathcal{M})$ whenever $\mathsf{last}(\mathcal{M})$ belongs to $\exists$.

The adorned formula $\varphi_{\mathcal{M}}$ should be thought of as providing auxiliary information guiding the definition of the strategy f. At every stage of the induction, we will show that for $(\varphi, B) := \mathsf{last}(\mathcal{M})$ it holds that:

(i) $\varphi^-_{\mathcal{M}} = \varphi$;

(ii) $\psi_B \wedge \iota(\varphi_{\mathcal{M}})$ is consistent.

Note that, by Lemma 4.3.14.(ii), it then follows that $\varphi \in B$.

For the induction base, we set $\varphi_{(\xi,A)} := \xi^{\emptyset}$, where $\xi^{\emptyset}$ is the $\mu_c^a\mathtt{ML}$-formula obtained by adorning every least fixed point operator in ξ by the empty set. It is not hard to see that $\varphi^-_{(\xi,A)} = \xi$, and that $\iota(\varphi_{(\xi,A)})$ is provably equivalent to ξ, whence consistent with ψ_A.

For the induction step, suppose that $\mathcal{M}$ is an f-guided match and that $\varphi_{\mathcal{M}}$ has been defined. We treat matches extending $\mathcal{M}$ by a single position, making a case distinction on the shape of $\varphi_{\mathcal{M}}$. In the following we denote by (φ, B) the position $\mathsf{last}(\mathcal{M})$.

- If $\varphi_{\mathcal{M}}$ is a literal there is nothing to do, because then φ is a literal as well and so the match $\mathcal{M}$ will not be extended.

- Suppose $\varphi_{\mathcal{M}}$ is of the form $\varphi_1 \vee \varphi_2$. Because $\psi_B \wedge \iota(\varphi_1 \vee \varphi_2)$ is consistent, there must be a disjunct φ_i such that $\psi_B \wedge \iota(\varphi_i)$ is consistent. We set $f(\mathcal{M}) := (\varphi_i^-, B)$. Note that the only f-guided match extending $\mathcal{M}$ by a single position is $\mathcal{N} := \mathcal{M} \cdot f(\mathcal{M})$. We define $\varphi_{\mathcal{N}} := \varphi_i$.

- Now suppose that $\varphi_{\mathcal{M}}$ is of the form $\varphi_1 \wedge \varphi_2$. As (φ, B) is owned by $\forall$, we do not have to define $f(\mathcal{M})$. Suppose that $\mathcal{N} = \mathcal{M} \cdot (\varphi_i^-, B)$. We then define $\varphi_{\mathcal{N}} := \varphi_i$.

- Now suppose that $\varphi_{\mathcal{M}}$ is of the form $\Diamond\chi$. Since $\psi_B \wedge \Diamond\iota(\chi)$ is consistent, there is, by the Existence Lemma, some state C such that BRC and moreover $\psi_C \wedge \iota(\chi)$ is consistent. We let $f(\mathcal{M}) := (\chi^-, C)$ and, for $\mathcal{N} := \mathcal{M} \cdot f(\mathcal{M})$, define $\varphi_{\mathcal{N}} = \chi$.

- If $\varphi_{\mathcal{M}}$ is of the form $\Box\chi$, things are less nice than in the other cases. The reason is that the consistency of $\psi_B \wedge \Box\iota(\chi)$ does not guarantee the consistency of $\psi_C \wedge \iota(\chi)$ for every C such that BRC. However, we do know that $\Box\chi^- \in B$, whence $\chi^- \in C$ for every such C. Hence, for every $\mathcal{N} = \mathcal{M} \cdot (\chi^-, C)$, we define $\varphi_{\mathcal{N}} := \chi^{\emptyset}$, where $\chi^{\emptyset}$ is defined as in the base case of this induction.

- Now suppose that $\varphi_{\mathcal{M}}$ is of the form $\mu^U x\delta$. By Lemma 4.3.14.(i), we find that $\psi_B \wedge \iota(\delta)[\iota(\mu^{U\cup\{B\}}x\delta)/x]$ is consistent. Therefore, under the assumption that $\mathcal{N} := \mathcal{M} \cdot (\delta^-[\mu x\delta^-/x], B)$, we may define $\varphi_{\mathcal{N}} := \delta[\mu^{U\cup\{B\}}x\delta/x]$.

- Finally, suppose that $\varphi_{\mathcal{M}}$ is of the form $\nu x\delta$. Since $\psi_B \wedge \varphi_{\mathcal{M}}$ is consistent, so is $\psi_B \wedge \delta[\nu x\delta/x]$. For $\mathcal{N} := \mathcal{M} \cdot (\delta^-[\nu x\delta^-/x])$, we set $\varphi_{\mathcal{N}} := \delta[\nu x\delta/x]$.

We claim that the strategy f is winning for $\exists$. To this end, let $\mathcal{M}$ be a full f-guided $\mathcal{E}(\xi, \mathbb{S})@(\xi, A)$-match. Suppose first that $\mathcal{M}$ is finite. We claim that $(\varphi, B) := \mathsf{last}(\mathcal{M})$ must belong to $\forall$, whence $\mathcal{M}$ is won by $\exists$. Indeed, note that φ cannot be of the form $\Diamond\chi$, for otherwise $f(\mathcal{M})$ would be defined and the match $\mathcal{M}$ would not be full. Moreover, if φ is a literal, it follows from the fact that $\varphi \in B$ that $\mathcal{M}$ is won by $\exists$.

Now suppose that $\mathcal{M}$ is infinite. Suppose, towards a contradiction, that $\mathcal{M}$ is won by $\forall$. By Lemma 2.1.43, we find that $\mathcal{M}$ has a final segment

$$(\varphi_0, B_0) \cdot (\varphi_1, B_1) \cdot (\varphi_2, B_2) \cdots$$

such that the main connective of φ_i is *not* amongst $\{\Box, \nu\}$ for any i. By the pigeonhole principle there must be n, m with $n < m$ such that $(B_n, \varphi_n) = (B_m, \varphi_m)$, and $\varphi_n = \varphi_m$ is of the form $\mu x\delta$.

For each i, let us write $\mathcal{M}_i$ for the match $\mathcal{M}$ up to (and including) the position (φ_i, B_i). By construction $\varphi_{\mathcal{M}_n}$ is of the form $\mu^U x\delta$, and $\varphi_{\mathcal{M}_{n+1}} = \delta[\mu^{U\cup\{B_n\}}x\delta)/x]$. Since $\Box$ does not occur on the segment from (φ_n, B_n) to (φ_m, B_m) it follows that $\varphi_{\mathcal{M}_m} = \mu^V x\delta$ for some $V \subseteq S$ with $B_m \in V$. But the consistency of $\psi_{B_m} \wedge \iota(\varphi_{M_m})$ then entails that $\psi_{B_m} \wedge \overline{\psi_V}$ is consistent, a contradiction. □

Completeness The goal of this paragraph is to prove completeness for certain well-behaved μ_c-logics.

Given a logic L, we define its canonical model as usual.

4.3.16. Definition. The *canonical model* $(S^{\mathsf{L}}, R^{\mathsf{L}}, V^{\mathsf{L}})$ of a logic L is given by:

- $S^{\mathsf{L}} := \{\Gamma : \Gamma \text{ is maximally } \mathsf{L}\text{-consistent}\}$.
- $\Gamma R^{\mathsf{L}} \Delta :\Leftrightarrow (\Box\varphi \in \Gamma \Rightarrow \varphi \in \Delta)$.
- $V^{\mathsf{L}}(p) := \{\Gamma : p \in \Gamma\}$.

We denote this canonical model by $\mathbb{S}^{\mathsf{L}}$. The *canonical frame* of L is the frame $(S^{\mathsf{L}}, R^{\mathsf{L}})$ underlying $\mathbb{S}^{\mathsf{L}}$.

For (infinitary) canonical models there is also a standard Existence Lemma:

4.3.17. Lemma. *For any state Γ of a canonical model $\mathbb{S}^{\mathsf{L}}$:*

If $\Diamond\varphi \in \Gamma$, then there is a state Δ such that $\Gamma R^{\mathsf{L}} \Delta$ and $\varphi \in \Delta$.

Generally, a μ_c-logic L will lack the compactness property. It is well-known that this prevents one to prove a Truth Lemma for the (standard) canonical model of L. Indeed, if there are unsatisfiable sets of formulas which are finitely satisfiable, then, because derivations are finite objects, there will be unsatisfiable maximally consistent sets. Recall that a concrete example of such a set was given by Proposition 2.3.2.

4.3.18. Remark. It is interesting to compare our situation to that of PDL. Although PDL is a modal fixed point logic, its fixed point behaviour is *implicit*. That is, PDL can be described purely as a modal logic, without the need for explicit fixed point operators.

In this view, a model for PDL is a multimodal Kripke model $(S, R_{\pi\in\mathsf{Prog}}, V)$, where Prog is a set of programs. Formulas of PDL are interpreted over the subclass of *standard models*, where the R_π are required to stand in a certain relation to each other. For instance, the relation R_{a^*} should be the reflexive-transitive closure of the relation R_a.

Using the ordinary methods of basic modal logic, one can obtain a canonical model $\mathbb{S}^{\mathsf{PDL}}$ which, despite the compactness failure, will satisfy the Truth Lemma.

However, this model will *not* be standard. To obtain completeness of PDL with respect to standard models, one then applies filtration to this non-standard canonical model.

In a sense, our canonical model $\mathbb{S}^{\mathsf{L}}$ is similar to the non-standard canonical model for PDL. The only difference is that we lack a non-standard interpretation for the fixed point operators, under which a Truth Lemma holds. We will circumvent this by applying a form of filtration which does not identify states that satisfy the same formulas, but instead identifies states (of the canonical model) which *contain* the same formulas. Note that the Truth Lemma precisely says that these two identifications are the same.

The following lemma is analogous to Proposition 4.2.9.

4.3.19. LEMMA. *Let* L *be a logic and let* F *be a class of frames that admits filtration and contains the canonical frame* $(S^{\mathsf{L}}, R^{\mathsf{L}})$. *For any finite set* Σ *of* $\mu^c\mathtt{ML}$*-formulas, the class* F *contains a frame underlying some model over* Θ *with respect to* L*, where* Θ *is a finite* $\overline{\mathsf{FL}}$*-closed extension of* Σ.

Proof:
Without loss of generality, we assume that Σ is $\overline{\mathsf{FL}}$-closed. The idea is to apply filtration to some particular model $\mathbb{S}'$ based on the canonical frame $(S^{\mathsf{L}}, R^{\mathsf{L}})$. As in the proof of Proposition 4.2.9, we let $\varphi_1, \ldots, \varphi_n$ be an enumeration of the formulas of the form $\mu x\delta$ in Σ. For each such formula φ_i, we pick a unique propositional variable p_i not occurring in Σ.

We define the valuation $V' : \mathsf{P} \to \mathcal{P}(S^{\mathsf{L}})$ of $\mathbb{S}'$ as follows:

$$V'(p) := \begin{cases} \{\Gamma : \varphi_i \in \Gamma\} & \text{if } p = p_i \text{ for some } \varphi_i \in \Sigma; \\ V(p) & \text{otherwise.} \end{cases}$$

Note that this is similar to the model $\mathbb{S}[\sigma]$ in the proof of Proposition 4.2.9, but now we do not let the valuation of p_i be the *meaning* of φ_i in $\mathbb{S}^{\mathsf{L}}$, but rather we let p_i be true at precisely those Γ where $\varphi_i \in \Gamma$. Were a Truth Lemma to hold for $\mathbb{S}^{\mathsf{L}}$, these two options would be equivalent.

We define the translation $\tau : \Sigma \to \mathtt{ML}$ in exactly the same way as in the proof of Proposition 4.2.9.

Since the frame underlying $\mathbb{S}'$ belongs to F, we can apply the assumed admissibility of filtration to obtain a filtration $\overline{\mathbb{S}} = (\overline{S}, \overline{R}, \overline{V})$ of $\mathbb{S}'$ through some finite $\overline{\mathsf{FL}}$-closed set $\Theta \supseteq \tau[\Sigma]$ of ML-formulas such that the frame $(\overline{S}, \overline{R})$ belongs to F.

Let $\sigma : \mathsf{P} \to \mu^c\mathtt{ML}$ be exactly as in the proof of Proposition 4.2.9. A straightforward induction shows that for every $\varphi \in \Theta$:

$$\mathbb{S}', \Gamma \Vdash \varphi \Leftrightarrow \sigma(\varphi) \in \Gamma. \tag{4.1}$$

We will finish the proof by showing that $\overline{\mathbb{S}}$ is isomorphic to a model over $\sigma[\Theta]$ (note that $\sigma[\Theta]$ is $\overline{\mathsf{FL}}$-closed by the same argument as in the proof of Proposition 4.2.9).

We define the set of states $S^{\sigma[\Theta]} := \{\Gamma \cap \sigma[\Theta] : \Gamma \text{ is maximally } \mathsf{L}\text{-consistent}\}$ and claim that the map

$$h : [\Gamma] \mapsto \Gamma \cap \Sigma$$

is a well-defined bijection from $S^{\mathsf{L}}/{\sim_\Theta^{\mathbb{S}'}}$ to $S^{\sigma[\Theta]}$. For well-definedness, suppose $\Gamma \sim_\Theta^{\mathbb{S}'} \Gamma'$ and let $\varphi \in \Theta$. By using the equivalence (4.1), we find for every $\sigma(\varphi) \in \sigma[\Theta]$:

$$\sigma(\varphi) \in \Gamma \Leftrightarrow \mathbb{S}', \Gamma \Vdash \varphi \Leftrightarrow \mathbb{S}', \Gamma' \Vdash \varphi \Leftrightarrow \sigma(\varphi) \in \Gamma',$$

as required. Injectivity is similar: if $\Gamma \cap \sigma[\Theta] = \Gamma' \cap \sigma[\Theta]$, then for all $\varphi \in \Theta$, we have:

$$\mathbb{S}', \Gamma \Vdash \varphi \Leftrightarrow \sigma(\varphi) \in \Gamma \Leftrightarrow \sigma(\varphi) \in \Gamma' \Leftrightarrow \mathbb{S}', \Gamma' \Vdash \varphi.$$

For surjectivity, take $\Gamma \cap \sigma[\Theta]$ for some any $\Gamma \in S^{\mathsf{L}}$. Then $h([\Gamma]) = \Gamma \cap \sigma[\Theta]$, as required.

Now we let $R^{\sigma[\Theta]} \subseteq S^{\sigma[\Theta]} \times S^{\sigma[\Theta]}$ and $V^{\sigma[\Theta]} : \mathsf{P} \to \mathcal{P}(S^{\sigma[\Theta]})$ be given by transporting the structure of $\overline{\mathbb{S}}$ along h. More precisely, we let

$$AR^{\sigma[\Theta]}B :\Leftrightarrow h^{-1}(A)\overline{R}h^{-1}(B).$$

We claim that $R_{\mathsf{L}}^{\mathsf{min}} \subseteq R^{\sigma[\Theta]} \subseteq R_{\mathsf{L}}^{\mathsf{max}}$.

First, suppose that $AR_{\mathsf{L}}^{\mathsf{min}}B$. Then $\psi_A \wedge \Diamond\psi_B$ is L-consistent. Pick some $\Gamma \in S^{\mathsf{L}}$ containing both ψ_A and $\Diamond\psi_B$. By Lemma 4.3.17, there is a $\Delta \in S^{\mathsf{L}}$ such that $\Gamma R^{\mathsf{L}} \Delta$ and $\psi_B \in \Delta$. Since $\overline{\mathbb{S}}$ is a filtration, we have $R_\Theta^{\mathsf{min}} \subseteq \overline{R}$. Hence $[\Gamma]\overline{R}[\Delta]$ and thus $h([\Gamma])R^{\sigma[\Theta]}h([\Delta])$. The required result follows from the fact that $h([\Gamma]) = A$ and $h([\Delta]) = B$.

Now suppose that $AR^{\sigma[\Theta]}B$. We will show that $AR_{\mathsf{L}}^{\mathsf{max}}B$. To that end, let $\Box\sigma(\varphi) \in \sigma[\Theta]$ such that $\Box\sigma(\varphi) \in A$. Pick $\Gamma \supset A$ and $\Delta \supset B$ from S^{L}. Since $[\Gamma] = h^{-1}(A)$ and $[\Delta] = h^{-1}(B)$, we have $[\Gamma]\overline{R}[\Delta]$. We now use the fact that $\overline{R} \subseteq R_\Theta^{\mathsf{max}}$. This means that for all $\Box\psi \in \Theta$ such that $\mathbb{S}', \Gamma \Vdash \Box\psi$, we have $\mathbb{S}', \Delta \Vdash \psi$. By assumption we have $\sigma(\Box\varphi) \in \Gamma$, whence the equivalence (4.1) gives $\mathbb{S}', \Gamma \Vdash \Box\varphi$. It follows that $\mathbb{S}', \Delta \Vdash \varphi$. Finally, another application of the equivalence (4.1) yields $\sigma(\varphi) \in \Delta$, hence $\sigma(\varphi) \in B$, as required. Since every $\Box$-formula in $\sigma[\Theta]$ is of the form $\Box\sigma(\varphi)$ for some $\varphi \in \Theta$, this suffices to show that $AR_{\mathsf{L}}^{\mathsf{max}}B$.

Lastly, for any $p \in \sigma[\Theta]$, we have $p \in \Theta$ and $p \neq p_i$ for every $1 \leq i \leq n$. We define:

$$V^{\sigma[\Theta]}(p) := \{A \in S^{\sigma[\Theta]} : h^{-1}(A) \in \overline{V}(p)\} = \{A \in S^{\sigma[\Theta]} : p \in A\},$$

which suffices. □

4.3.20. THEOREM. *Let* L *be a canonical logic in the basic modal language such* $\mathsf{Fr}(\mathsf{L})$ *admits filtration. Then* μ_c-L *is sound and complete with respect to* $\mathsf{Fr}(\mathsf{L})$.

Proof:
Soundness follows from the fact the fixed point axioms and rules are sound on the class of all frames. For completeness, let $\varphi \in \mu^c\mathtt{ML}$ be L-consistent; we will show that φ is satisfiable in a model based on an L-frame. Without loss of generality we may assume that φ is tidy and in negation normal form. Note that by canonicity the canonical frame $(S^{\mathsf{L}}, R^{\mathsf{L}})$ is contained in $\mathsf{Fr}(\mathsf{L})$. Therefore, we can use Lemma 4.3.19 to obtain a model $\mathbb{S}^{\Sigma}$ over some $\Sigma \supseteq \{\varphi\}$ with respect to L whose frame belongs to $\mathsf{Fr}(\mathsf{L})$. By the L-consistency of φ, there is a state $A \in S^{\Sigma}$ such that $\varphi \in A$. Finally, Lemma 4.3.15 gives $\mathbb{S}^{\Sigma}, A \Vdash \varphi$, as required. □

For instance, the logic μ_c-KB is sound and complete with respect to the class of symmetric frames. Some other examples of basic modal logics that satisfy the hypotheses of the above theorem are: K, T, $\mathsf{K4}$, $\mathsf{S4}$ and $\mathsf{S5}$.

4.4 Conclusion

We have shown that the methods of filtration and finitary canonical models generalise from $\mathtt{PDL}$ to the continuous modal μ-calculus. To the best of our knowledge, this is the first completeness proof for Kozen's axiomatisation restricted to the continuous modal μ-calculus, even over the class of all frames.

Since $\mu^c\mathtt{ML}$ is strictly more expressive than $\mathtt{PDL}$ [41, 23], this is a proper generalisation of the results in [56]. On the other hand, because the failure of filtration for $\mu\mathtt{ML}$ is witnessed by the formula $\mu x \Box x$, the syntactic restrictions characterising $\mu^c\mathtt{ML}$ seem to be not only sufficient, but also necessary for filtration. This indicates that $\mu^c\mathtt{ML}$ might be positioned as a maximal filtration-allowing language between the basic modal language and the full language of the modal μ-calculus. We leave it for future work to make this statement mathematically precise and to investigate its correctness.

Another important remaining open question is that of the unification of the two techniques. As mentioned in the introduction and Remark 4.3.18, this is clear in the case of $\mathtt{PDL}$: finitary canonical models arise as filtrations of infinitary canonical models. In contrast, the proofs in this chapter, although very similar, are carried out independently. It would be interesting to see if the finitary canonical model of some μ_c-logic L could explicitly be obtained as the filtration of some (non-standard, perhaps topological) infinitary canonical model for L.

Chapter 5

Focus-style proofs for the two-way alternation-free μ-calculus

5.1 Introduction

In this chapter, we introduce a non-well-founded proof system for the two-way alternation-free modal μ-calculus $\mu_2^{af}\mathtt{ML}$. As mentioned in the Chapter 2, this logic extends the alternation-free modal μ-calculus with backward modalities. Already without fixed point operators, backward modalities are known to require more expressivity than offered by a cut-free Gentzen system [79]. A common solution is to add more structure to sequents, as *e.g.* the nested sequents of Kashima [55]. This approach, however, does not combine well with cyclic proofs, as the number of possible sequents in a given proof becomes unbounded. We therefore opt for the alternative approach of still using ordinary sequents, but allowing analytic applications of the cut rule (see [46] for more on the history of this approach). We have already seen in Chapter 3 that this can be fruitfully combined with cyclic proofs. Choosing analytic cuts over sequents with extended structure has recently also been gaining interest in the proof theory of logics without fixed point operators [27].

Although allowing analytic cuts handles the backward modalities on a local level, further issues arise on a global level in the combination with non-well-founded branches. The main challenge is that the progress condition should not just hold on infinite branches, but also on paths that can be constructed by moving both up and down a proof tree. Our solution takes inspiration from Vardi's reduction of alternating two-way automata to deterministic one-way automata [105]. Roughly, the idea is to view these paths simply as upward paths, only interrupted by several detours, each returning to the same state as where it departed. One of the main insights of the present research is that such detours have a natural interpretation in terms of the game semantics of the modal μ-calculus. We exploit this by extending the syntax with so-called *trace atoms*, whose semantics corresponds with this interpretation. Our sequents will then be one-sided Gentzen sequents containing annotated formulas, trace atoms, and negations of trace atoms.

During the development of this work, the preprint [2] by Enqvist et al. appeared, in which a proof system is presented for the two-way modal μ-calculus (with alternation). Like our system, their system is cyclic. Moreover, they also extend the syntax in order to apply the techniques from Vardi in a proof-theoretical setting. However, their extension, which uses so-called *ordinal variables*, is substantially different from ours, which uses trace atoms. It would be interesting to see whether the two approaches are intertranslatable.

Section 5.2 is devoted to introducing the proof system, after which in Section 5.3 we define the proof search game. In Section 5.4 we prove soundness and completeness. The concluding Section 5.5 contains a short summary and some ideas for further research.

5.2 The proof system

Recall that a set Σ of tidy formulas in negation-normal form is called *negation-closed* if for every $\xi \in \Sigma$ it holds that $\overline{\xi} \in \Sigma$ and $\mathsf{Clos}(\xi) \subseteq \Sigma$. For the remainder of this chapter we fix a finite and negation-closed set Σ of $\mu_2^{af}\mathtt{ML}$-formulas. For reasons of technical convenience, we will assume that every formula is drawn from Σ. This does not restrict the scope of our results, as any formula is equivalent to one contained in some finite negation-closed set of tidy formulas in negation-normal form.

5.2.1 Sequents

Syntax

In [73], Marti & Venema show that for the alternation-free modal μ-calculus, a path-based soundness condition can already be obtained by only annotating formulas by a single bit of information, namely a *focus annotation* in $\{\circ, \bullet\}$. We follow their approach.

5.2.1. Definition. An *annotated formula* is a formula plus a focus annotation.

The letters $b, c, d, \ldots$ are used as variables ranging over the annotations $\circ$ and $\bullet$. An annotated formula φ^b is said to be *out of focus* if $b = \circ$, and *in focus* if $b = \bullet$.

5.2.2. Remark. Note that the annotations are the same as those in Chapter 3. However, in that chapter the limited expressivity of the language allowed us to restrict attention to hypersequents of a very specific shape. In particular, those hypersequents had at most one formula in focus and any formula in focus had a specific shape. For the more expressive language $\mu_2^{af}\mathtt{ML}$ we cannot allow ourselves the same restriction, and therefore also consider sequents which have multiple formulas in focus, where the formulas can be of any shape.

Where traces usually only move upward in a proof, the backward modalities of our language will be enable them to go downward as well. We will handle this in our proof system by further enriching our sequents with the following additional information.

5.2.3. Definition. For any two formulas φ, ψ, there is a *trace atom* $\varphi \rightsquigarrow \psi$ and a *negated trace atom* $\varphi \not\rightsquigarrow \psi$.

The concept of trace atoms will become clearer later, but for now one can think of $\varphi \rightsquigarrow \psi$ as expressing that there is some kind of trace going from φ to ψ, and of $\varphi \not\rightsquigarrow \psi$ as its negation. Finally, our sequents are built from the above three entities.

5.2.4. Definition. A *sequent* is a finite set consisting of annotated formulas, trace atoms, and negated trace atoms.

Whenever we want to refer to general elements of a sequent Γ, without specifying whether we mean annotated formulas or (negated) trace atoms, we will use the capital letters $A, B, C, \ldots$.

Semantics

Unlike annotations, which do not affect the semantics but only serve as bookkeeping devices, the trace atoms have a well-defined interpretation. We will work with a refinement of the usual satisfaction relation that is defined with respect to a strategy for $\forall$ in the evaluation game. Most of the time, this strategy will be both *optimal* and *positional* (recall that the precise definition of these terms was given in Section 2.2). Because we will frequently need to mention such optimal positional strategies, we will refer to them by the abbreviation *ops*. We first define the interpretation of annotated formulas. Note that the focus annotations play no role in this definition.

5.2.5. Definition. Let $\mathbb{S}$ be a model, f an ops for $\forall$ in $\mathcal{E}@(\bigwedge \Sigma, \mathbb{S})$, and φ^b an annotated formula. We write $\mathbb{S}, s \Vdash_f \varphi^b$ if f is *not* winning for $\forall$ at (φ, s).

The following proposition, which is an immediate consequence of Theorem 2.2.10, relates $\Vdash_f$ to the usual satisfaction relation $\Vdash$.

5.2.6. Proposition. $\mathbb{S}, s \Vdash \varphi$ *iff for every ops f for $\forall$ in $\mathcal{E}(\bigwedge \Sigma, \mathbb{S})$:* $\mathbb{S}, s \Vdash_f \varphi^b$.

The semantics of trace atoms is also given relative to an ops for $\forall$ in the game $\mathcal{E}(\bigwedge \Sigma, \mathbb{S})$ (in the following often abbreviated to $\mathcal{E}$).

5.2.7. Definition. Given an ops f for $\forall$ in $\mathcal{E}$, we say that $\varphi \rightsquigarrow \psi$ is *satisfied* in $\mathbb{S}$ at s with respect to f (and write $\mathbb{S}, s \Vdash_f \varphi \rightsquigarrow \psi$) if there is an f-guided match

$$(\varphi, s) = (\varphi_0, s_0) \cdot (\varphi_1, s_1) \cdots (\varphi_n, s_n) = (\psi, s) \quad (n \geq 0)$$

such that for no $0 \leq i < n$ the formula φ_i is a μ-formula. We say that $\mathbb{S}$ *satisfies* $\varphi \not\rightsquigarrow \psi$ at s with respect to f (and write $\mathbb{S}, s \Vdash_f \varphi \not\rightsquigarrow \psi$) iff $\mathbb{S}, s \not\Vdash_f \varphi \rightsquigarrow \psi$.

Note that in the witnessing match $(\varphi_0, s_0) \cdots (\varphi_n, s_n)$, the formula φ_0 is only allowed to be a μ-formula in case $n = 0$, *i.e.* in case the match has length 1.

The idea behind the satisfaction of a trace atom $\varphi \rightsquigarrow \psi$ at a state s is that $\exists$ can take the match from (φ, s) to (ψ, s) without passing through a μ-formula. This is good for the player $\exists$. For instance, if $\varphi \rightsquigarrow \psi$ and $\psi \rightsquigarrow \varphi$ are satisfied at s with respect to f for some $\varphi \neq \psi$, then f is necessarily losing for $\forall$ at the position (φ, s). We will later relate trace atoms to traces in infinitary proofs.

Note that, in a match witnessing that $\mathbb{S}, s \Vdash_f \varphi \rightsquigarrow \psi$, only the final formula allowed be a μ-formula. For example, for every φ it vacuously holds that $\mathbb{S}, s \Vdash \varphi \rightsquigarrow \varphi$, including formulas φ of the form $\mu x \psi$. However, as shown in the first example below, a trace atom of the form $\chi \rightsquigarrow \mu x \psi$ can only be satisfied in this vacuous way.

5.2.8. Example. We illustrate the workings of trace formulas by highlighting some facts.

(i) $\mu x \varphi \rightsquigarrow \chi$ is satisfiable if and only if $\chi = \mu x \varphi$.

This follows directly from the definition: if $\chi \neq \mu x \varphi$, then $\mu x \varphi$ is not the final formula in the witnessing match, contradicting the requirement that the non-final formulas are not μ-formulas. Conversely, the satisfiability of the trace atom $\mu x \varphi \rightsquigarrow \mu x \varphi$ is witnessed by any match consisting of only one position $(\mu x \varphi, s)$.

(ii) $\nu x \varphi \rightsquigarrow \varphi[\nu x \varphi / x]$ is always true.

If $(\nu x \varphi, s)$ is some position in $\mathcal{E}$, then $(\varphi[\nu x \varphi/x], s)$ is always the next position. Hence the two-position match $(\nu x \varphi, s) \cdot (\varphi[\nu x \varphi/x], s)$ can be used to witness $\mathbb{S}, s \Vdash_f \nu x \varphi \rightsquigarrow \varphi[\nu x \varphi/x]$ for any strategy f.

(iii) $\mathbb{S}, s \Vdash_f \varphi \rightsquigarrow \langle a \rangle \psi$ implies $\mathbb{S}, t \Vdash_f \langle \breve{a} \rangle \varphi \rightsquigarrow \psi$ for every a-successor f of s.

Suppose $\mathbb{S}, s \Vdash_f \varphi \rightsquigarrow \langle a \rangle \psi$. This must be witnessed by an f-guided match

$$(\varphi, s) = (\varphi_0, s_0) \cdot (\varphi_1, s_1) \cdots (\varphi_n, s_n) = (\langle a \rangle \psi, s) \quad (n \geq 0)$$

Note that no formula in this match is a μ-formula, as the final formula is not a μ-formula. Since t is an a-successor of s, we can extend the match on both sides in the following way

$$(\langle \breve{a} \rangle \varphi, t) \cdot (\varphi, s) = (\varphi_0, s_0) \cdot (\varphi_1, s_1) \cdots (\varphi_n, s_n) = (\langle a \rangle \psi, s) \cdot (\psi, t)$$

which gives us $\mathbb{S}, s \Vdash \langle \breve{a} \rangle \varphi \rightsquigarrow \psi$, as $\langle \breve{a} \rangle \varphi$ is not a μ-formula.

We interpret sequents disjunctively, that is: $\mathbb{S}, s \Vdash_f \Gamma$ whenever $\mathbb{S}, s \Vdash_f A$ for some $A \in \Gamma$. The sequent Γ is said to be *valid* whenever $\mathbb{S}, s \Vdash_f \Gamma$ for every model $\mathbb{S}$, state s of $\mathbb{S}$, and ops f for $\forall$ in $\mathcal{E}$. Recall from Chapter 1 that the alternation-free two-way modal μ-calculus is interpreted over models which are *standard* in the following sense

$$R_{\breve{a}} = \{(t,s) \mid (s,t) \in R_a\} \text{ for every } a \in \mathsf{D}. \qquad (*)$$

In other words, this means that $R_{\breve{a}}$ must always be the converse of the relation R_a. Throughout this chapter we assume that all models satisfy the property (*).

5.2.9. Remark. There is another way in which one could interpret sequents, which corresponds to what one might call *strong validity*, and which the reader should note is different from our notion of validity. Spelling it out, we say that Γ is *strongly valid* if for every model $\mathbb{S}$ and state s there is an A in Γ that such that for every ops f for $\forall$ in $\mathcal{E}$ it holds that $\mathbb{S}, s \Vdash_f A$. While these two notions coincide for sequents containing only annotated formulas, an example of a valid, but not strongly valid sequent is given by $\{\varphi \wedge \psi \rightsquigarrow \varphi, \varphi \wedge \psi \rightsquigarrow \psi\}$.

It is not hard to see that the sequent above is valid. Indeed, any f-guided match must from $(\varphi \wedge \psi, s)$ either proceed to (φ, s) or to (ψ, s). To see that it is not strongly valid, suppose for instance that $\varphi = p$ and $\psi = q$ with $p \neq q$. If both p and q are true at some state s of some model $\mathbb{S}$, then there are ops's f_1 and f_2 for $\forall$ such that $f_1(\varphi \wedge \psi, s) = (p, s)$ and $f_2(\varphi \wedge \psi, s) = (q, s)$. But then $\mathbb{S}, s \nVdash_{f_1} \varphi \wedge \psi \rightsquigarrow \psi$, and $\mathbb{S}, s \nVdash_{f_2} \varphi \wedge \psi \rightsquigarrow \varphi$, showing that the sequent is not strongly valid.

As will become clear in Section 5.4, our soundness proof requires the notion of validity, rather than strong validity. The proof works by contraposition, showing that every invalid sequent is unprovable. Crucially, we use the fact that for every invalid sequent Γ there exists a model $\mathbb{S}$, a state s and a *particular* ops f for $\forall$ such that for every $A \in \Gamma$ it holds that $\mathbb{S}, s \nVdash_f A$. For a sequent which is merely not strongly valid, we would not get such a particular ops.

We finish this subsection by defining three operations on sequents that, respectively, extract the formulas contained annotated in some sequent, take all annotated formulas out of focus, and put all formulas into focus.

$$\begin{aligned}
\Gamma^- &:= \{\chi \mid \chi^b \in \Gamma \text{ for some } b \in \{\circ, \bullet\}\},\\
\Gamma^\circ &:= \{\varphi \rightsquigarrow \psi \mid \varphi \rightsquigarrow \psi \in \Gamma\} \cup \{\varphi \not\rightsquigarrow \psi \mid \varphi \not\rightsquigarrow \psi \in \Gamma\} \cup \{\chi^\circ \mid \chi \in \Gamma^-\},\\
\Gamma^\bullet &:= \{\varphi \rightsquigarrow \psi \mid \varphi \rightsquigarrow \psi \in \Gamma\} \cup \{\varphi \not\rightsquigarrow \psi \mid \varphi \not\rightsquigarrow \psi \in \Gamma\} \cup \{\chi^\bullet \mid \chi \in \Gamma^-\}.
\end{aligned}$$

5.2.2 Proofs

In this subsection we give the rules of our proof system. Because the rule for modalities is quite involved, its details are given in a separate definition. Recall that Σ is the finite negation-closed set of formulas fixed at the beginning of Section 5.2.

5.2.10. DEFINITION. Let Γ be a sequent and let $[a]\varphi^b$ be an annotated formula. The *jump* $\Gamma^{[a]\varphi^b}$ of Γ with respect to $[a]\varphi^b$ consists of:

1. (a) $\varphi^{s([a]\varphi,\Gamma)}$;
   (b) $\psi^{s(\langle a\rangle\psi,\Gamma)}$ for every $\langle a\rangle\psi^c \in \Gamma$;
   (c) $[\breve{a}]\chi^\circ$ for every $\chi^d \in \Gamma$ such that $[\breve{a}]\chi \in \Sigma$;

2. (a) $\varphi \rightsquigarrow \langle\breve{a}\rangle\chi$ for every $[a]\varphi \rightsquigarrow \chi \in \Gamma$ such that $\langle\breve{a}\rangle\chi \in \Sigma$;
   (b) $\langle\breve{a}\rangle\chi \not\rightsquigarrow \varphi$ for every $\chi \not\rightsquigarrow [a]\varphi \in \Gamma$ such that $\langle\breve{a}\rangle\chi \in \Sigma$;
   (c) $\psi \rightsquigarrow \langle\breve{a}\rangle\chi$ for every $\langle a\rangle\psi \rightsquigarrow \chi \in \Gamma$ such that $\langle\breve{a}\rangle\chi \in \Sigma$;
   (d) $\langle\breve{a}\rangle\chi \not\rightsquigarrow \psi$ for every $\chi \not\rightsquigarrow \langle a\rangle\psi \in \Gamma$ such that $\langle\breve{a}\rangle\chi \in \Sigma$,

where $s(\xi,\Gamma)$ is defined by:

$$s(\xi,\Gamma) = \begin{cases} \bullet & \text{if } \xi^\bullet \in \Gamma, \\ \bullet & \text{if } \theta \not\rightsquigarrow \xi \in \Gamma \text{ for some } \theta^\bullet \in \Gamma, \\ \circ & \text{otherwise.} \end{cases}$$

Before we go on to provide the rest of the proof system, we will give some intuition for the modal rule, by proving the lemma below. This lemma essentially expresses that the modal rule is sound. Since the annotations play no role in the soundness of an individual rule, we suppress the annotations in the proof below for the sake of readability. Intuition for the annotations in the modal rule, and in particular for the function s, will be given later.

5.2.11. LEMMA. *Given a model $\mathbb{S}$, a state s of $\mathbb{S}$, and an ops f for $\forall$ in $\mathcal{E}$ such that $\mathbb{S}, s \not\Vdash_f [a]\varphi^b, \Gamma$, let t be such that $f(([a]\varphi, s) = (\varphi, t)$. Then $\mathbb{S}, t \not\Vdash_f \Gamma^{[a]\varphi^b}$.*

Proof:
First note that t is well-defined. Indeed, by assumption $\mathbb{S}, s \not\Vdash_f [a]\varphi^b, \Gamma$. Hence in particular $\mathbb{S}, s \not\Vdash_f [a]\varphi^b$. This means that the strategy f is winning in $\mathcal{E}$ at the position $([a]\varphi, s)$. In particular $\forall$ does not get stuck at this position, and thus f must select a position (φ, t), where t is an a-successor of s.

Now, we claim that $\mathbb{S}, t \not\Vdash_f \Gamma^{[a]\varphi^b}$. To start with, since f is winning, we have $\mathbb{S}, t \not\Vdash_f \varphi$. Moreover, if $\langle a\rangle\psi$ belongs to Γ, then $\mathbb{S}, s \not\Vdash_f \langle a\rangle\psi$ and thus $\mathbb{S}, t \not\Vdash_f \psi$. Thirdly, if χ belongs to Γ and $[\breve{a}]\chi \in \Sigma$, then, by optimality, it holds that $\mathbb{S}, t \not\Vdash_f [\breve{a}]\chi$.

With this we have shown all conditions under item 1 of Definition 5.2.10. For the conditions under item 2, suppose that $\langle\breve{a}\rangle\chi \in \Sigma$. We only show 2(a), because the others are similar (note that 2(d) is essentially the third item of Example 5.2.8). For 2(a), we reason by contraposition. So suppose that $\mathbb{S}, t \Vdash_f \varphi \rightsquigarrow \langle\breve{a}\rangle\chi$. This is witnessed by an f-guided $\mathcal{E}$-match

$$(\varphi, t) = (\varphi_0, s_0) \cdot (\varphi_1, s_1) \cdots (\varphi_n, s_n) = (\langle\breve{a}\rangle\chi, t).$$

$$\dfrac{}{\varphi^b, \overline{\varphi}^c, \Gamma}\ \mathsf{Ax1} \qquad \dfrac{}{\varphi \rightsquigarrow \psi, \varphi \not\rightsquigarrow \psi, \Gamma}\ \mathsf{Ax2} \qquad \dfrac{}{\varphi \rightsquigarrow \varphi, \Gamma}\ \mathsf{Ax3}$$

$$\dfrac{(\varphi \vee \psi) \not\rightsquigarrow \varphi, (\varphi \vee \psi) \not\rightsquigarrow \psi, \varphi^b, \psi^b, \Gamma}{\varphi \vee \psi^b, \Gamma}\ \mathsf{R}_\vee \qquad \dfrac{\varphi^\circ, \Gamma \qquad \overline{\varphi}^\circ, \Gamma}{\Gamma}\ \mathsf{cut}$$

$$\dfrac{(\varphi \wedge \psi) \not\rightsquigarrow \varphi, \varphi^b, \Gamma \qquad (\varphi \wedge \psi) \not\rightsquigarrow \psi, \psi^b, \Gamma}{\varphi \wedge \psi^b, \Gamma}\ \mathsf{R}_\wedge \qquad \dfrac{\varphi[\mu x\varphi/x]^\circ, \Gamma}{\mu x\varphi^b, \Gamma}\ \mathsf{R}_\mu$$

$$\dfrac{\nu x\varphi \not\rightsquigarrow \varphi[\nu x\varphi/x], \varphi[\nu x\varphi/x] \rightsquigarrow \nu x\varphi, \varphi[\nu x\varphi/x]^b, \Gamma}{\nu x\varphi^b, \Gamma}\ \mathsf{R}_\nu \qquad \dfrac{\Gamma^{[a]\varphi^b}}{[a]\varphi^b, \Gamma}\ \mathsf{R}_{[a]}$$

$$\dfrac{\Gamma^\bullet}{\Gamma^\circ}\ \mathsf{F} \qquad \dfrac{\varphi \not\rightsquigarrow \psi, \psi \not\rightsquigarrow \chi, \varphi \not\rightsquigarrow \chi, \Gamma}{\varphi \not\rightsquigarrow \psi, \psi \not\rightsquigarrow \chi, \Gamma}\ \mathsf{trans} \qquad \dfrac{\varphi \rightsquigarrow \psi, \Gamma \qquad \varphi \not\rightsquigarrow \psi, \Gamma}{\Gamma}\ \mathsf{tc}$$

Figure 5.1: The proof rules of the system Focus^2.

But then the f-guided $\mathcal{E}$-match

$$([a]\varphi, s) \cdot (\varphi_0, s_0) \cdots (\varphi_n, s_n) \cdot (\langle a \rangle \chi, s),$$

witnesses that $\mathbb{S}, s \Vdash_f [a]\varphi \rightsquigarrow \chi$, as required. □

The rules of the system Focus^2 are given in Figure 5.1. In each rule except of the modal rule, in the conclusion and each premiss, the annotated formulas occurring in the set Γ are called *inactive*. Moreover, the conclusions and premisses of the rules in $\{\mathsf{R}_\vee, \mathsf{R}_\wedge, \mathsf{R}_\mu, \mathsf{R}_\nu\}$ have precisely one *active formula*, which by definition is the annotated formula appearing to the left of Γ. The single active formula in the conclusion is often called *principal*. Note that, due to the fact that sequents are taken to be sets, an annotated formula may at the same time be both active and inactive. For the modal rule, the *active* formulas in the conclusion are those that are referred to by Definition 5.2.10. That is, the formula $[a]\varphi^b$, all formulas of the form $\langle a \rangle \psi^c$, and all formulas χ^d such that $[\breve{a}] \in \Sigma$. All other formulas are *inactive* in the conclusion of the modal rule are *inactive*. Finally, the premiss of the modal rule contains only active formulas.

5.2.12. Remark. Note that being active is only defined for annotated formulas, and not for (negated) trace atoms. The same holds for the notion of direct ancestry, which we will define below.

We will now define the relation of *direct ancestry* between formulas in the conclusion and formulas in the premisses of some arbitrary rule application. For any inactive formula in the conclusion of some rule, we let its *direct ancestors* be the corresponding inactive formulas in the premisses. For every rule except $\mathsf{R}_{[a]}$, if some formula in the conclusion is an active formula, its *direct ancestors* are the active formulas in the premisses. Finally, for the *modal rule* $\mathsf{R}_{[a]}$, we stipulate that $\varphi^{s([a]\varphi,\Gamma)}$ is an *direct ancestor* of the principal formula $[a]\varphi^b$, and that each $\psi^{s(\langle a\rangle\psi,\Gamma)}$ contained in $\Gamma^{[a]\varphi^b}$ due to clause 1(b) of Definition 5.2.10 is an *direct ancestor* of $\langle a\rangle\psi^b \in \Gamma$.

As mentioned before, the purpose of the focus annotations is to keep track of trails of formulas on branches (in the sense of Definition 2.3.11. Usually, a trail is a sequence of formulas $(\varphi_n)_{n<\omega}$ such that each φ_k is an direct ancestor of φ_{k+1}. The idea is then that whenever an infinite branch has cofinitely many sequents with a formula in focus, this branch contains a trail on which infinitely many formulas are ν-formulas. Disregarding the backward modalities for now, this can be seen as follows. As long as the focus rule is not applied, any focussed formula is an direct ancestor of some earlier focussed formula. Since the principal formula of R_μ loses focus, while that of R_ν preserves focus, a straightforward application of Kőnig's Lemma shows that every infinite branch contains a trail with infinitely many ν-formulas. We refer the reader to [73] for more details on this argument.

Our setting is slightly more complicated, because the function s in Definition 5.2.10 additionally allows the focus to transfer along negated trace atoms, rather than just from a formula to one of its direct ancestors. This is inspired by [105], as are the conditions in the second part of Definition 5.2.10. The main idea is that, because of the backward modalities, traces may move not only up, but also down a proof tree. To get a grip on these more complex traces, we cut them up in segments consisting of upward paths, which are the same as ordinary traces, and loops, which are captured by the negated trace atoms. This intuitive idea will become explicit in the proof of completeness in Section 5.4.

5.2.13. Remark. The reader might be surprised by Clause 1(c) of Definition 5.2.10. Since $[\breve{a}]\chi^\circ$ in the premiss is closely related to χ^d in the conclusion, one would expect there to be some transfer of focus between the two. The crucial points is that this focus transfer would have to be *backwards*, in the sense that if $[\breve{a}]\chi$ is in focus, then χ must be in focus as well. The role of the trace atoms is precisely to capture these dynamics.

5.2.14. Remark. We will often reason about the proof system Focus^2 contrapositively. As a result, the negative trace atoms are often considered as positive tace atoms and vice versa. This for instance explains the formulation of the transitivity rule.

5.2.15. Remark. Note that, except for the rule tc, or *trace cut*, the only rule with a positive trace atom in the premiss is the rule R_ν. The idea is that if $\nu x\varphi$ is false at some state in some model with respect to some ops f for $\forall$, then

$\varphi[\nu x\varphi/x] \rightsquigarrow \nu x\varphi$ must be false as well. Indeed, if it were true, then, since by Example 5.2.8 $\varphi[\nu x\varphi/x] \rightsquigarrow \nu x\varphi$ is also true, there would be a winning match for $\exists$ going back and forth between $\nu x\varphi$ and $\varphi[\nu x\varphi/x]$. Hence $\nu x\varphi$ would be winning for $\exists$. Similar positive trace atoms could be added to the premisses of the other rules. For instance, it would be sound to weaken the rule $\mathsf{R}_\vee$ by adding the positive trace atoms $\varphi \rightsquigarrow (\varphi \vee \psi)$ and $\psi \rightsquigarrow (\varphi \vee \psi)$ to the premiss. It turns out, however, that only having a positive trace atom in the conclusion of R_ν is sufficient for completeness. Why it is necessary for completeness will become clear in the completeness proof of Section 5.4.2.

We are now ready to define a notion of infinitary proofs in Focus^2. Recall that a derivation is *closed* if every leaf is an axiom.

5.2.16. Definition. A Focus^2_∞*-proof* is a closed Focus^2-derivation such that:

1. Every infinite branch has infinitely many applications of $\mathsf{R}_{[a]}$.
2. On every infinite branch cofinitely many sequents have a formula in focus.

Condition 1 ensures that the tightening of every infinite trail is infinite. Conditon 2 guarantees that this infinite tightening is a ν-trace. These properties will be used in Section 5.4 to show that infinitary proofs are sound. The key idea is to relate the traces in a proof to matches in the evaluation game on a purported countermodel of the proof's conclusion.

We close this section with two examples of Focus^2_∞-proofs. The first example demonstrates cut and item 1(c) of Definition 5.2.10. The second example demonstrates trace atoms.

5.2.17. Example. Define the following two formulas:

$$\varphi := \mu x(\langle \breve{a}\rangle x \vee p), \qquad\qquad \psi := \nu y([a]x \wedge \varphi).$$

The formula φ expresses 'there is a backward a-path to some state where p holds'. The formula ψ expresses 'φ holds at every state reachable by a forwards a-path'. As our context Σ we take least negation-closed set containing φ and ψ:

$$\{\varphi, \langle\breve{a}\rangle\varphi \vee p, \langle\breve{a}\rangle\varphi, p, \psi, [a]\psi \wedge \varphi, [a]\psi, \overline{\varphi}, [\breve{a}]\overline{\varphi} \wedge \overline{p}, \overline{p}, [\breve{a}]\overline{\varphi}, \overline{\psi}, \langle a\rangle\overline{\psi} \vee \overline{\varphi}, \langle a\rangle\overline{\psi}\}.$$

The implication $p \to \psi$ is valid, and below we give a Focus^2_∞-proof. As this particular proof does not rely on trace atoms, we omit them for readability.

$$\dfrac{\dfrac{\dfrac{\dfrac{}{\overline{p}^\bullet, \psi^\bullet, \langle\breve{a}\rangle\varphi^\circ, p^\circ}\,\mathsf{Ax1}}{\overline{p}^\bullet, \psi^\bullet, \langle\breve{a}\rangle\varphi \vee p^\circ}\,\mathsf{R}_\vee}{\overline{p}^\bullet, \psi^\bullet, \varphi^\circ}\,\mathsf{R}_\mu \qquad \dfrac{\dfrac{\dfrac{\dfrac{\pi}{\psi^\bullet, [\breve{a}]\overline{\varphi}^\circ}}{\overline{p}^\bullet, [a]\psi^\bullet, \overline{\varphi}^\circ}\,\mathsf{R}_{[a]} \qquad \dfrac{}{\overline{p}^\bullet, \varphi^\bullet, \overline{\varphi}^\circ}\,\mathsf{Ax1}}{\overline{p}^\bullet, [a]\psi \wedge \varphi^\bullet, \overline{\varphi}^\circ}\,\mathsf{R}_\wedge}{\overline{p}^\bullet, \psi^\bullet, \overline{\varphi}^\circ}\,\mathsf{R}_\nu}{\overline{p}^\bullet, \psi^\bullet}\,\mathsf{cut}$$

In the above proof, the proof π is given by

$$
\dfrac{
\dfrac{
\dfrac{
\dfrac{
\dfrac{}{\overline{\varphi}^\circ, \varphi^\circ}\,\mathsf{Ax1}
}{[a]\psi^\bullet, [\breve{a}]\overline{\varphi}^\circ, \langle\breve{a}\rangle\varphi^\circ, p^\circ}\,\mathsf{R}_{[\breve{a}]}
}{[a]\psi^\bullet, [\breve{a}]\overline{\varphi}^\circ, \langle\breve{a}\rangle\varphi \vee p^\circ}\,\mathsf{R}_{\vee}
}{[a]\psi^\bullet, [\breve{a}]\overline{\varphi}^\circ, \varphi^\circ}\,\mathsf{R}_{\mu}
\qquad
\dfrac{
\dfrac{\vdots}{\psi^\bullet, [\breve{a}]\overline{\varphi}^\circ}
}{[a]\psi^\bullet, [\breve{a}]\overline{\varphi}^\circ, \overline{\varphi}^\circ}\,\mathsf{R}_{[a]}
}{[a]\psi^\bullet, [\breve{a}]\overline{\varphi}^\circ}\,\mathsf{cut}
\qquad
\dfrac{
\dfrac{
\dfrac{
\dfrac{}{\varphi^\circ, \overline{\varphi}^\circ}\,\mathsf{Ax1}
}{\langle\breve{a}\rangle\varphi^\circ, p^\circ, [\breve{a}]\overline{\varphi}^\circ}\,\mathsf{R}_{[\breve{a}]}
}{\langle\breve{a}\rangle\varphi \vee p^\circ, [\breve{a}]\overline{\varphi}^\circ}\,\mathsf{R}_{\vee}
}{\varphi^\bullet, [\breve{a}]\overline{\varphi}^\circ}\,\mathsf{R}_{\mu}
}{
\dfrac{[a]\psi \wedge \varphi^\bullet, [\breve{a}]\overline{\varphi}^\circ}{\psi^\bullet, [\breve{a}]\overline{\varphi}^\circ}\,\mathsf{R}_{\nu}
}\,\mathsf{R}_{\wedge}
$$

where the vertical dots indicate that the proof continues by repeating what happens at the root of π. The resulting proof of $\overline{p}^\bullet, \psi^\bullet$ has a single infinite branch, which can easily be seen to satisfy the conditions of Definition 5.2.16.

5.2.18. Example. Define $\varphi := \nu x\langle a\rangle\langle\breve{a}\rangle x$, *i.e.* φ expresses that there is an infinite path of alternating a and $\breve{a}$ transitions. Clearly this holds at every state with an a-successor. Hence the implication $\langle a\rangle p \to \varphi$ is valid. As context Σ we consider the least negation-closed set containing both $\langle a\rangle p$ and φ, *i.e.*,

$$\{\langle a\rangle p, p, \varphi, \langle a\rangle\langle\breve{a}\rangle\varphi, \langle\breve{a}\rangle\varphi, [a]\overline{p}, \overline{p}, \overline{\varphi}, [a][\breve{a}]\overline{\varphi}, [\breve{a}]\overline{\varphi}\}.$$

The following is a Focus^2_∞-proof of $\langle a\rangle p \to \varphi$.

$$
\dfrac{
\dfrac{
\dfrac{}{\overline{p}^\bullet, \langle\breve{a}\rangle\varphi^\bullet, \langle\breve{a}\rangle\varphi \not\rightsquigarrow \langle\breve{a}\rangle\varphi, \langle\breve{a}\rangle\varphi \rightsquigarrow \langle\breve{a}\rangle\varphi}\,\mathsf{Ax2}
}{[a]\overline{p}^\bullet, \langle a\rangle\langle\breve{a}\rangle\varphi^\bullet, \varphi \not\rightsquigarrow \langle a\rangle\langle\breve{a}\rangle\varphi, \langle a\rangle\langle\breve{a}\rangle\varphi \rightsquigarrow \varphi}\,\mathsf{R}_{[a]}
}{[a]\overline{p}^\bullet, \varphi^\bullet}\,\mathsf{R}_{\nu}
$$

Note that it is also possible to use $\mathsf{Ax3}$ instead of $\mathsf{Ax2}$ in the above proof.

5.3 The proof search game

We will define a proof search game $\mathcal{G}(\Gamma)$ for the proof system Focus^2_∞ analogous to the game of Section 2.3.3. First, we require a slightly more formal definition of the notion of a rule instance.

For Γ a sequent, the set of positions of $\mathcal{G}(\Gamma)$ is $\mathsf{Seq}_\Gamma \cup \mathsf{Inst}_\Gamma$, where Seq_Γ is the set of sequents and Inst_Γ is the set of valid rule instances, containing only formulas in the negation-closure of the formula occurring in Γ (either as an annotated formula or as part of a, possibly negated, trace atom).

Since Γ is finite, the game $\mathcal{G}(\Gamma)$ has only finitely many positions. In particular, since by assumption every formula in Γ belongs to the set Σ we fixed at the beginning of Section 5.2, every formula occurring in some position of $\mathcal{G}(\Gamma)$ belongs to Σ as well. Note that this implies that cut and tc only introduce annotated formulas and trace atoms build from formulas in Σ.

The ownership function and admissible moves of $\mathcal{G}(\Gamma)$ are as in the following table:

Position	Owner	Admissible moves
$\Gamma \in \mathsf{Seq}_\Gamma$	Prover	$\{i \in \mathsf{Inst}_\Gamma \mid \mathsf{conc}(i) = \Gamma\}$
$(\Gamma, \mathsf{r}, \langle \Delta_1, \ldots, \Delta_n \rangle) \in \mathsf{Inst}_\Gamma$	Refuter	$\{\Delta_i \mid 1 \leq i \leq n\}$

In the above table, the expression $\mathsf{conc}(i)$ stands for the conclusion (*i.e.* the first element of the triple) of the rule instance i. As usual, a finite match is lost by the player who got stuck. An infinite $\mathcal{G}(\Gamma)$-match is won by Prover if and only it has a final segment

$$\Gamma_0 \cdot i_0 \cdot \Gamma_1 \cdot i_1 \cdots$$

on which each Γ_k has at least one formula in focus and the instance i_k is an application of $\mathsf{R}_{[a]}$ for infinitely many k. The two main observations about $\mathcal{G}(\Gamma)$ that we will use are the following:

1. A Focus^2_∞-proof of Γ is the same as a winning strategy for Prover in $\mathcal{G}(\Gamma)@\Gamma$.
2. $\mathcal{G}(\Gamma)$ is a parity game, whence positionally determined.

The first observation is immediate when viewing a winning strategy as a subtree of the full game tree. To make the second observation more explicit, we give the parity function Ω for $\mathcal{G}(\Gamma)$. On Seq_Γ, we simply set $\Omega(\Gamma) := 0$ for every $\Gamma \in \mathsf{Seq}_\Gamma$. On Inst_Γ, we define:

$$\Omega(\Gamma, \mathsf{r}, \langle \Delta_1, \ldots, \Delta_n \rangle) := \begin{cases} 3 & \text{if } \Gamma \text{ has no formula in focus,} \\ 2 & \text{if } \Gamma \text{ has a formula in focus and } \mathsf{r} = \mathsf{R}_{[a]}, \\ 1 & \text{if } \Gamma \text{ has a formula in focus and } \mathsf{r} \neq \mathsf{R}_{[a]}. \end{cases}$$

As a result we immediately obtain a method to reduce general non-well-founded proofs to a kind of cyclic proofs. Indeed, if Prover has a winning strategy, she also has positional winning strategy, which clearly corresponds to a regular Focus^2_∞-proof.

5.3.1. Remark. Using the terminology of the Intermezzo, let us call *good* those finite paths in Focus^2-derivations which always have a formula in focus, and on which the modal rule is applied at least once. It is not hard to see that the *simple* infinitary proof system (cf. Definition I.2.7) generated by this notion of good finite paths is precisely Focus^2_∞. Hence, we obtain a notion of cyclic Focus^2-proofs with a soundness condition that can be checked by looking only at the paths between each repeating leaf and its companion. This cyclic system proves exactly the same sequents as Focus^2_∞. Hence the soundness and completeness theorems of the next section transfer to the cyclic system.

5.4 Soundness and completeness

In this section we will prove the soundness and completeness of the system Focus^2_∞. More specifically, for soundness we will show that if Γ is invalid, then Refuter has a winning strategy in $\mathcal{G}(\Gamma)@\Gamma$. Our completeness result is slightly less wide in scope, showing only that if Refuter has a winning strategy in $\mathcal{G}(\Gamma)@\Gamma$, then Γ^- is invalid.

5.4.1 Soundness

For soundness, we assume an ops f for $\forall$ in $\mathcal{E} := \mathcal{E}(\bigwedge \Sigma, \mathbb{S})$ for some $\mathbb{S}$ and s such that $\mathbb{S}, s \not\Vdash_f \Gamma$. The goal is to construct from f a strategy $\overline{f}$ for Refuter in $\mathcal{G} := \mathcal{G}(\Gamma)$. The key idea is to assign to each position p reached in $\mathcal{G}$ a state s such that whenever $p = \Delta \in \mathsf{Seq}_\Gamma$ it holds that $\mathbb{S}, s \not\Vdash_f \Delta$. For $i \in \mathsf{Inst}_\Gamma$, the choice of $\overline{f}$ is then based on $f(\varphi, s)$ where φ is a formula determined by the rule instance i. The existence of such an s implies that i cannot be an axiom and thus that Refuter never gets stuck. For infinite matches, the proof works by showing that an $\overline{f}$-guided $\mathcal{G}@\Gamma$-match lost by Refuter induces an f-guided $\mathcal{E}@\varphi$-match lost by $\forall$. As mentioned above, the key idea here is to relate an f-guided $\mathcal{E}@\varphi$-match to a trail through the $\overline{f}$-guided $\mathcal{G}@\Gamma$-match. If the $\mathcal{G}@\Gamma$-match is losing for Refuter, it must contain a ν-trail, which gives us an $\mathcal{E}@\varphi$-match lost by $\forall$. A novel challenge here is that not all steps in a trail necessarily go from a formula to one of its direct ancestors, but may instead transfer along a negated trace atom. When this happens, say from φ_n to φ_{n+1}, it holds for Δ as above that both $\varphi_n^\bullet$ and $\varphi_n \not\rightsquigarrow \varphi_{n+1}$ belong to Δ. Since, by the above, it holds that $\mathbb{S}, s \not\Vdash_f \Delta$, we use the fact that $\mathbb{S}, s \Vdash_f \varphi_n \rightsquigarrow \varphi_{n+1}$ to take the $\mathcal{E}@\varphi$-match from (φ_n, s) to (φ_{n+1}, s).

Recall that Lemma 5.2.11 showed that the modal rule is sound. The next proposition shows that every other rule of Focus^2 is sound as well. In fact, if the conclusion is falsified in some state, one of the premisses is falsified in the same state.

5.4.1. Lemma. *Let*

$$\mathsf{R}\,\frac{\Delta_1 \qquad \cdots \qquad \Delta_n}{\Gamma}$$

be an instance of any rule apart from $\mathsf{R}_{[a]}$. *Given a model* $\mathbb{S}$*, a state* s *of* $\mathbb{S}$*, and an ops* f *for* $\forall$ *in* $\mathcal{E}$ *such that* $\mathbb{S}, s \not\Vdash_f \Gamma$*, there is an* $1 \leq i \leq n$ *such that* $\mathbb{S}, s \not\Vdash_f \Delta_i$.

In particular, if $\mathsf{R} = \mathsf{R}_\wedge$ *and* $\varphi_1 \wedge \varphi_2^b$ *is the principal formula, then it holds that* $\mathbb{S}, s \not\Vdash_f (\varphi_1 \wedge \varphi_2) \rightsquigarrow \varphi_i, \varphi_i^b, \Gamma$*, where* φ_i *is such that* $f(\varphi_1 \wedge \varphi_2, s) = (\varphi_i, s)$.

Proof:
Suppose that $\mathbb{S}, s \not\Vdash_f \Gamma$. We make a case distinction on the rule R.

- ($\mathsf{Ax1}$) In this case there is a formula φ such that $\varphi^b, \overline{\varphi}^c \in \Gamma$. Since φ and $\overline{\varphi}$ are Boolean duals, it either holds that $\mathbb{S}, s \Vdash_f \varphi$, or $\mathbb{S}, s \Vdash_f \overline{\varphi}$. Hence it is not possible that $\mathbb{S}, s \not\Vdash_f \Gamma$ and thus the implication is vacuous in this case.

- ($\mathsf{Ax2}$) This case is similar to the previous case: either $\varphi \rightsquigarrow \psi$ or $\varphi \not\rightsquigarrow \psi$ belongs to Γ. Hence $\mathbb{S}, s \Vdash_f \Gamma$.

(Ax3) We have $\mathbb{S}, s \Vdash_f \varphi \rightsquigarrow \varphi$, witnessed by the one-position match (φ, s). Hence again $\mathbb{S}, s \Vdash_f \Gamma$.

($\mathsf{R}_\vee$) By assumption $\mathbb{S}, s \nVdash_f \varphi \vee \psi$. This means that $\mathbb{S}, s \nVdash_f \varphi$ and $\mathbb{S}, s \nVdash_f \psi$. Moreover, both $\mathbb{S}, s \Vdash_f (\varphi \vee \psi) \rightsquigarrow \varphi$ and $\mathbb{S}, s \Vdash_f (\varphi \vee \psi) \rightsquigarrow \varphi$ hold. Indeed, from the position $(\varphi \vee \psi, s)$, it is possible for $\exists$ to proceed to either (φ, s) or (ψ, s). Hence $\mathbb{S}, s \nVdash_f (\varphi \vee \psi) \not\rightsquigarrow \varphi, (\varphi \vee \psi) \not\rightsquigarrow \psi, \varphi^b, \psi^b, \Gamma$, as required.

($\mathsf{R}_\wedge$) Since $\mathbb{S}, s \nVdash_f \varphi \wedge \psi$, it holds that f is a winning strategy in $\mathcal{E}$ at $(\varphi \wedge \psi, s)$. Suppose without loss of generality that $f(\varphi \wedge \psi, s) = (\varphi, s)$. Then f is winning in $\mathcal{E}$ at (φ, s) as well. Moreover, the f-guided match $(\varphi \wedge \psi, s) \cdot (\varphi, s)$ witnesses that $\mathbb{S}, s \Vdash_f (\varphi \wedge \psi) \rightsquigarrow \varphi$. So $\mathbb{S}, s \nVdash_f (\varphi \wedge \psi) \not\rightsquigarrow \varphi, \varphi^b, \Gamma$.

(R_μ) Since f is winning for $\forall$ in $\mathcal{E}$ at $(\mu x \varphi, s)$, and the next position necessarily is $(\varphi[\mu x \varphi / x], s)$, it follows that f is winning for $\forall$ at $(\varphi[\mu x \varphi / x], s)$ as well. Hence $\mathbb{S}, s \nVdash_f \varphi[\mu x \varphi / x]^b, \Gamma$.

(R_ν) By the same argument as in the previous case, we find $\mathbb{S}, s \nVdash_f \varphi[\nu x \varphi / x]$. Moreover, item (ii) of Example 5.2.8 shows that $\mathbb{S}, s \Vdash_f \nu x \varphi \rightsquigarrow \varphi[\nu x \varphi / x]$. Finally, suppose that we would also have $\mathbb{S}, s \Vdash_f \varphi[\nu x \varphi / x] \rightsquigarrow \nu x \varphi$. Then there would be an infinite f-guided $\mathcal{E}$-match

$$(\nu x \varphi, s) \cdots (\varphi[\nu x \varphi / x], s) \cdots (\nu x \varphi, s) \cdots (\varphi[\nu x \varphi / x], s) \cdots$$

which does not go through a μ-formula. But then $\mathbb{S}, s \Vdash_f \nu x \varphi$, contradicting the assumption. Hence $\mathbb{S}, s \nVdash_f \varphi[\nu x \varphi / x] \rightsquigarrow \nu x \varphi$. We have thus found that $\mathbb{S}, s \nVdash_f \nu x \varphi \not\rightsquigarrow \varphi, \varphi[\nu x \varphi / x] \rightsquigarrow \nu x \varphi, \varphi[\nu x \varphi / x]^b, \Gamma$, as required.

(F) This is trivial, as the focus annotation do no impact satisfiability.

(trans) Note that it suffices to show that $\mathbb{S}, s \nVdash_f \varphi \not\rightsquigarrow \chi$, *i.e.* $\mathbb{S}, s \Vdash_f \varphi \rightsquigarrow \chi$. By assumption, we have $\mathbb{S}, s \Vdash_f \varphi \rightsquigarrow \psi$ and $\mathbb{S}, s \Vdash_f \psi \rightsquigarrow \chi$. Hence there indeed exists an f-guided $\mathcal{E}$-match

$$(\varphi, s) \cdots (\psi, s) \cdots (\chi, s),$$

which does not go through a μ-formula.

(cut) By the optimality of f, we know that f must either be winning for $\forall$ in $\mathcal{E}$ at the position (φ, s) or at the position $(\overline{\varphi}, s)$, Hence, we either have $\mathbb{S}, s \nVdash_f \varphi^\circ, \Gamma$, or $\mathbb{S}, s \nVdash_f \overline{\varphi}^\circ, \Gamma$.

(tc) This final case is similar to the previous case: we must either have $\mathbb{S}, s \nVdash_f \varphi \rightsquigarrow \psi$, or $\mathbb{S}, s \nVdash_f \varphi \not\rightsquigarrow \psi$.

□

Together with Lemma 5.2.11, the previous lemma entails that well-founded Focus_∞^2-proofs are sound.

5.4.2. PROPOSITION. *Well-founded* Focus_∞^2*-proofs are sound.*

The rest of this section is devoted to generalising the previous proposition to also include non-well-founded Focus_∞^2-proofs. We first establish an auxiliary lemma.

5.4.3. LEMMA. *Let $\mathcal{M}$ be some infinite $\mathcal{G}(\Gamma)$-match won by Prover. Then $\mathcal{M}$ has a final segment*

$$\mathcal{N} = \Gamma_0 \cdot i_0 \cdot \Gamma_1 \cdot i_1 \cdot \Gamma_2 \cdot i_2 \cdots$$

for which there is a sequence of formulas $\varphi_0, \varphi_1, \varphi_2, \cdots$ such that for every $n \geq 0$ it holds that $\varphi_n^\bullet \in \Gamma_n$, and, in addition, at least one of the following holds:

- *$\varphi_{n+1}^\bullet \in \Gamma_{n+1}$ is an direct ancestor of $\varphi_n^\bullet \in \Gamma_n$;*
- *$i_n = \mathsf{R}_{[a]}$ and Γ_n contains some $\varphi_n \not\rightsquigarrow \xi$ such that $\varphi_{n+1}^\bullet \in \Gamma_{n+1}$ is an direct ancestor of some $\xi^b \in \Gamma_n$.*

Proof:
First note that, by winning condition on infinite matches of $\mathcal{G}(\Gamma)$, there is a final segment $\mathcal{N} = \Gamma_0 \cdot i_0 \cdot \Gamma_1 \cdot i_1 \cdots$ of $\mathcal{M}$ on which every sequent Γ_n has a formula in focus, and the rule instance i_n is modal for infinitely many n. Since every annotated formula in the conclusion of F is out of focus, we know that the rule instance i_n is not an application of F for any $n \geq 0$. By direct inspection of the rules, one can then see that a formula $\varphi^\bullet$ in some Γ_{n+1} can only be in focus for one of the following two reasons:

(i) $\varphi^\bullet$ is an direct ancestor of some formula $\psi^\bullet$ in Γ_n.

(ii) $i_n = \mathsf{R}_{[a]}$ and $\varphi^\bullet$ is an direct ancestor of some ξ^b in Γ_n such that Γ_n contains $\psi \not\rightsquigarrow \chi$ for some $\psi^\bullet$ in Γ_n.

Hence, we can build a tree where the root r has as children all formulas which are in focus in Γ_0, and each formula ψ in focus in Γ_n has as children each formula φ such that $\varphi^\bullet \in \Gamma_{n+1}$, whose focus can be justified from ψ either through item (i) or item (ii) above. Since each formula in focus in Γ_n can be traced back to the root r, it is contained in this tree. Hence, because there are infinitely many Γ_n, each of which as a formula in focus, the tree must be infinite. By Kőnig's Lemma, the tree has an infinite branch $\varphi_0, \varphi_1, \varphi_2, \ldots$ such that $\varphi_n^\bullet \in \Gamma_n$ for each n. Moreover, each φ_n satisfies at least one of the required conditions by construction. □

We are now ready to prove the full soundness theorem.

5.4.4. PROPOSITION. *If Γ is the conclusion of a* Focus_∞^2*-proof, then Γ is valid.*

Proof:
Our proof will go by contraposition, so suppose that some sequent Γ is invalid. This means that there is a model $\mathbb{S}$ with a state s and ops f for $\forall$ in the game $\mathcal{E} := \mathcal{E}(\bigwedge \Sigma, \mathbb{S})$, such that $\mathbb{S}, s \not\Vdash_f \Gamma$. We will construct a (positional) winning strategy $\overline{f}$ for Refuter in the game $\mathcal{G} := \mathcal{G}(\Gamma)$ initialised at Γ.

Formally, this strategy is a function $\overline{f} : \mathrm{PM}_R(\Gamma) \to \mathsf{Seq}_\Gamma$. In addition, we will define a function $s_f : \mathrm{PM}(\Gamma) \to \mathbb{S}$, from partial $\mathcal{G}$-matches starting at Γ to states of $\mathbb{S}$, such that $\mathbb{S}, s_f(\mathcal{M}) \not\Vdash_f \mathsf{last}(\mathcal{M})$ for every $\overline{f}$-guided $\mathcal{M} \in \mathrm{PM}_P(\Gamma)$, and $\mathbb{S}, s_f(\mathcal{M}) \not\Vdash_f \overline{f}(\mathcal{M})$ for every $\overline{f}$-guided $\mathcal{M} \in \mathrm{PM}_R(\Gamma)$.

We define $\overline{f}$ and s_f by induction on the length $|\mathcal{M}|$ of a match $\mathcal{M} \in \mathrm{PM}(\Gamma)$. For the base case, *i.e.* where $|\mathcal{M}| = 1$, we have $\mathcal{M} = \Gamma$. Since in this case $\mathcal{M} \in \mathrm{PM}_P(\Gamma)$, we only have to define $s_f(\mathcal{M})$ and not $\overline{f}(\mathcal{M})$. We set $s_f(\mathcal{M}) := s$. Note that this suffices, because $\mathsf{last}(\mathcal{M}) = \Gamma$ and by assumption $\mathbb{S}, s_f \not\Vdash_f \Gamma$.

Now suppose that $\overline{f}$ and s_f have been defined for all matches up to length n, and that $|\mathcal{M}| = n + 1$. We assume that $\mathcal{M}$ is $\overline{f}$-guided, for otherwise we may just assign $\overline{f}(\mathcal{M})$ and $s_f(\mathcal{M})$ some arbitrary value.

Suppose first that $\mathcal{M}$ belongs to $\mathrm{PM}_P(\Gamma)$. Writing $\mathcal{M}_{\leq n} \in \mathrm{PM}_R(\Gamma)$ for the initial segment of $\mathcal{M}$ consisting of the first n moves, we set $s_f(\mathcal{M}) := s_f(\mathcal{M}_{\leq n})$. Since $\mathcal{M}$ is $\overline{f}$-guided, we have $\mathsf{last}(\mathcal{M}) = \overline{f}(\mathcal{M}_{\leq n})$. Hence it holds by the induction hypothesis that $\mathbb{S}, s_f(\mathcal{M}) \not\Vdash_f \mathsf{last}(\mathcal{M})$, as required.

If $\mathcal{M}$ belongs to $\mathrm{PM}_R(\Gamma)$, then $\mathsf{last}(\mathcal{M})$ is a rule instance and we distinct cases based on the rule R of $\mathsf{last}(\mathcal{M}) \in \mathsf{Inst}_\Gamma$. If R is the modal rule $\mathsf{R}_{[a]}$, we let $s_f(\mathcal{M})$ be the state in $f([a]\varphi, s_f(\mathcal{M}_{\leq n}))$, where $[a]\varphi$ is principal in $\mathsf{last}(\mathcal{M})$. For $\overline{f}$ there is only a single choice, say Δ. We set $\overline{f}(\mathcal{M}) := \Delta$. Note that by Lemma 5.2.11, it indeed follows that $\mathbb{S}, s_f(\mathcal{M}) \not\Vdash_f \overline{f}(\mathcal{M})$. If R is any rule but the modal rule, we set $s_f(\mathcal{M}) := s_f(\mathcal{M}_{\leq n})$ and invoke Lemma 5.4.1 to obtain a premiss Δ_i such that $\mathbb{S}, s_f(\mathcal{M}) \not\Vdash_f \Delta_i$. We set $\overline{f}(\mathcal{M}) := \Delta_i$. In particular, if $\mathsf{R} = \mathsf{R}_\wedge$ and $\varphi_1 \wedge \varphi_2^b$ is the principal formula, then we set $\overline{f}(\mathcal{M}) := (\varphi_1 \wedge \varphi_2) \rightsquigarrow \varphi_i, \varphi_i^b, \Gamma$, where φ_i is such that $f(\varphi_1 \wedge \varphi_2, s_f(\mathcal{M}_{\leq n})) = (\varphi_i, s_f(\mathcal{M}_{\leq n}))$.

We will now show that $\overline{f}$ is indeed a winning strategy for Refuter in $\mathcal{G}@\Gamma$. To that end, suppose towards a contradiction that Refuter loses a $\overline{f}$-guided $\mathcal{G}@\Gamma$-match $\mathcal{M}$. We already know that Refuter does not get stuck, as an axiom is never reached and all other rule instances have a non-zero number of premisses. Hence, the match $\mathcal{M}$ must be infinite. Let $\mathcal{N} = \Gamma_0 \cdot i_0 \cdot \Gamma_1 \cdot i_1 \cdots$ be a final segment of $\mathcal{M}$ as given by Lemma 5.4.3. We use $\mathcal{K}$ to denote the initial segment of $\mathcal{M}$ occurring before $\mathcal{N}$, *i.e.* such that $\mathcal{M} = \mathcal{K} \cdot \mathcal{N}$. Without loss of generality we assume that $|\mathcal{K}| > 0$.

As before, we write $\mathcal{N}_{\leq n}$ for the initial segment of $\mathcal{N}$ up to the first n moves. Note that $\overline{f}(\mathcal{K} \cdot \mathcal{N}_{\leq 2n}) = \Gamma_n$ for every $n \geq 0$. For convenience we will denote $\mathcal{K} \cdot \mathcal{N}_{\leq 2n}$ by $\mathcal{M}_n$. We will reach a contradiction by showing that $\mathbb{S}, s_f(\mathcal{M}_0) \Vdash_f \varphi_0$, which contradicts the fact that $\mathbb{S}, s_f(\mathcal{M}_0) \not\Vdash_f \overline{f}(\mathcal{M}_0) = \Gamma_0$.

The crucial claim is that for every n there is an f-guided $\mathcal{E}$-match starting at $(\varphi_n, s_f(\mathcal{M}_n))$ and ending at $(\varphi_{n+1}, s_f(\mathcal{M}_{n+1}))$, without passing through a μ-unfolding. More precisely, we will show that there is an f-guided $\mathcal{E}$-match

$$(\varphi_n, s_f(\mathcal{M}_n)) = (\psi_0, s_0) \cdots (\psi_m, s_m) = (\varphi_{n+1}, s_f(\mathcal{M}_{n+1})) \quad (m \geq 0)$$

such that for no $i < m$ the formula ψ_i is a μ-formula. By pasting together these finite segments, it will then follow that the strategy f is not winning for $\forall$ in $\mathcal{E}@(\varphi_0, s_f(\mathcal{M}_0))$, reaching the desired contradiction.

We will first show the above claim under the assumption that $\varphi^\bullet_{n+1}$ is an direct ancestor of $\varphi^\bullet_n$, and $\varphi_n = \varphi_{n+1}$. Note that in this case i_n is not the modal rule. Hence $s_f(\mathcal{M}_n) = s_f(\mathcal{M}_{n+1})$ and thus $(\varphi_n, s_f(\mathcal{M}_n)) = (\varphi_{n+1}, s_f(\mathcal{M}_{n+1}))$, by which the result holds vacuously.

Now suppose that $\varphi^\bullet_{n+1}$ is an direct ancestor of $\varphi^\bullet_n$ and $\varphi_n \neq \varphi_{n+1}$. Note that in this case $\varphi^\bullet_n$ must be active in the conclusion Γ_n of the rule instance i_n, and $\varphi^\bullet_{n+1}$ must be active in the premiss Γ_{n+1} of same rule instance i_n. We will show, by a case distinction on the main connective of φ_n, that the match proceeds to the desired position $(\varphi_{n+1}, s_f(\mathcal{M}_{n+1}))$ after a single round.

- First note that φ_n cannot be atomic, for atomic formulas can only have direct ancestors when they are inactive.

- Suppose φ_n is of the form $\psi_1 \vee \psi_2$. Then $\varphi^\bullet_n$ must be principal and we have $\varphi_{n+1} = \psi_i$ for some $i \in \{1, 2\}$. We let $\exists$ simply choose the position $(\psi_i, s_f(\mathcal{M}_n)$. Since the rule of i_n must be $\mathsf{R}_\vee$, we have $s_f(\mathcal{M}_n) = s_f(\mathcal{M}_{n+1})$ and thus reach the desired position in $\mathcal{E}$.

- Suppose φ_n is of the form $\psi_1 \wedge \psi_2$. Again we find that $\varphi^\bullet_n$ must be principal, the rule of i_n now being $\mathsf{R}_\wedge$. Recall that, when invoking Lemma 5.4.1, we chose the premiss Γ_{n+1} in such a way that the active formula in Γ_{n+1} is $\psi^\bullet_i$, where $f(\psi_1 \wedge \psi_2, s_f(\mathcal{M}_n)) = (\psi_i, s_f(\mathcal{M}_n))$. Hence $\varphi_{n+1} = f(\varphi_n, s_f(\mathcal{M}_n))$, and the next position in $\mathcal{E}$ again suffices.

- Suppose $\varphi_n = \langle a \rangle \psi$. Then the rule of i_n must be $\mathsf{R}_{[a]}$ and $\varphi_{n+1} = \psi$. Recall that $s_f(\mathcal{M}_{n+1})$ was obtained from Lemma 5.2.11, and therefore $f([a]\chi, s_f(\mathcal{M}_n)) = (\chi, s_f(\mathcal{M}_{n+1})$, where $[a]\chi^b$ is the principal formula of the rule instance i_n. In particular it follows that $s_f(\mathcal{M}_{n+1})$ is an a-successor of $s_f(\mathcal{M}_n)$ in $\mathbb{S}$ and thus we can let $\exists$ choose $(\varphi_{n+1}, s_f(\mathcal{M}_{n+1}))$, as required.

- If $\varphi_n = [a]\chi$, then the rule of i_n must be $\mathsf{R}_{[a]}$ and $\varphi^\bullet_n$ must be the principal formula of this rule instance. As explained in the previous case, we have $f([a]\chi, s_f(\mathcal{M}_n)) = (\chi, s_f(\mathcal{M}_{n+1})$. Therefore the next position in $\mathcal{E}$ will be $(\chi, s_f(\mathcal{M}_{n+1})$, as required.

- $\varphi_n = \mu x \psi$ is not possible, because any direct ancestor of $\mu x \psi^\bullet$ that is not a side formula, will be out of focus.

- Finally, suppose that $\varphi_n = \nu x\psi$. We have that $\varphi_{n+1} = \psi[\nu x\psi/x]$ and the rule of i_n is R_ν. Because $s_f(\mathcal{M}_{n+1}) = s_f(\mathcal{M}_n)$, the required position is reached immediately.

Finally, suppose that $\varphi^\bullet_{n+1}$ is not an direct ancestor of $\varphi^\bullet$. Then it must be the case that $i_n = \mathsf{R}_{[a]}$ and Γ_n contains some $\varphi_n \not\rightsquigarrow \xi$ such that $\varphi^\bullet_{n+1}$ is an direct ancestor of some $\xi^b \in \Gamma_n$. By assumption $\mathbb{S}, s_f(\mathcal{M}_n) \nVdash_f \Gamma_n$, and thus in particular $\mathbb{S}, s_f(\mathcal{M}_n) \Vdash_f \varphi_n \rightsquigarrow \xi$. Hence $\exists$ can take the f-guided match from $(\varphi_n, s_f(\mathcal{M}_n))$ to $(\xi, s_f(\mathcal{M}_n))$ without passing through a μ-unfolding. Since ξ^b has an direct ancestor (namely $\varphi^\bullet_{n+1}$), we find that ξ must be of the form $\langle a\rangle\psi$ or of the form $[a]\chi$, where $[a]\chi$ is the principal formula of i_n. In either case $\exists$ can ensure that the next position after $(\xi, s_f(\mathcal{M}_n))$ is $(\varphi_{n+1}, s_f(\mathcal{M}_{n+1}))$ by using the same strategy as above for the $\langle a\rangle$ and $[a]$ cases, respectively.

Since the modal rule is applied infinitely often in $\mathcal{M}$, the segments constructed above must be of non-zero length infinitely often. Hence, we obtain an infinite f-guided $\mathcal{E}@(\varphi_0, s_f(\mathcal{M}_0))$-match won by $\exists$, a contradiction. □

5.4.2 Completeness

For completeness we conversely show that from a winning strategy f for Refuter in $\mathcal{G}@\Gamma$, we can construct a model $\mathbb{S}^f$ and a positional strategy $\underline{f}$ for $\forall$ in $\mathcal{E}(\bigwedge\Sigma, \mathbb{S}^f)$ such that $\mathbb{S}^f$ falsifies Γ^- with respect to $\underline{f}$. The strategy $\underline{f}$ we construct will not necessarily be optimal but, by Theorem 2.2.10, it follows that there must also be an ops g such that $\mathbb{S}^f \nVdash_g \Gamma^-$. We will view f as a tree, and restrict attention a certain subtree. We first need to define two relevant properties of rule applications.

5.4.5. Definition. A rule application is *cumulative* if all of the premisses are supersets of the conclusion. A rule application is *productive* if all of the premisses are distinct from the conclusion.

Without renaming f, we restrict f to its subtree where the strategy of Prover is to go through the following stages in succession:

1. Exhaustively apply productive instances of cut and tc.
2. If applicable, apply the focus rule.
3. Exhaustively take applications of $\mathsf{R}_\vee$, $\mathsf{R}_\wedge$, R_μ, R_ν, trans that are both cumulative and productive.
4. If applicable, apply an axiom.
5. If applicable, apply a modal rule and loop back to stage (1).

Note that this strategy for Prover is non-deterministic. Most importantly, to a single sequent it might be possible to apply different instances of the modal rule. Hence the resulting strategy tree f does not only branch at positions owned by Refuter, but may also branch at positions owned by Refuter.

It is not hard to see that each of the above phases terminates. More precisely, phases (2), (4) and (5) either terminate immediately or after applying a single rule. By the productivity requirement and the finiteness of Σ, phases (1) and (3) must terminate after a finite number of rule applications as well. Note also that non-cumulative rule applications can only happen in phases (2) or (5).

We will now define the model $\mathbb{S}^f$. The set S^f of states consists of maximal paths in f not containing a modal rule. Here we mean that different paths in f correspond to different states in $\mathbb{S}^f$, even if they happen to be labelled by the same sequents. We write $\Gamma(\rho)$ for $\bigcup\{\Gamma : \Gamma \text{ occurs in } \rho\}$. Note that, since the only possibly non-cumulative rule application in ρ is the focus rule, $\Gamma(\rho)^\bullet = \mathsf{last}(\rho)^\bullet$ for every state ρ of $\mathbb{S}^f$. Moreover, we write $\rho_1 \xrightarrow{a} \rho_2$ if ρ_2 is directly above ρ_1 in f, separated only by an application of $\mathsf{R}_{[a]}$ (we assume that trees grow upward). We write $\to$ for the union $\bigcup\{\xrightarrow{a} : a \in \mathsf{D}\}$. Clearly, under the relation $\to$ the states of $\mathbb{S}^f$ form a forest (not necessarily a tree!). We write $\rho \leq \tau$ if τ is a descendant of ρ in this forest, *i.e.* $\leq$ is the reflexive-transitive closure of $\to$. The accessibility relations R^f_a of $\mathbb{S}^f$ are defined as follows:

$$\rho_1 R^f_a \rho_2 \text{ if and only if } \rho_1 \xrightarrow{a} \rho_2 \text{ or } \rho_2 \xrightarrow{\breve{a}} \rho_1.$$

Note under these accessibility relations $\mathbb{S}^f$ indeed satisfies the regularity property (*) above. We define the valuation $V^f : S^f \to \mathcal{P}(\mathsf{P})$ by

$$V^f(\rho) := \{p : \overline{p} \in \Gamma(\rho)^-\}.$$

The model $\mathbb{S}^f$ inherits much of the tree structure of f. There are two main differences. First, a path in f between two modal rules is collapsed into a single state in $\mathbb{S}^f$. Second, for every path ρ_2 directly above some path ρ_1, with an application of $\mathsf{R}_{[a]}$ in between, in the model $\mathbb{S}^f$ it does not only hold that $\rho_1 R^f_a \rho_2$, but also that $\rho_2 R^f_{\breve{a}} \rho_1$.

By the restriction on f, and in particular the fact that each of the stages (1) - (5) terminates, each state ρ of $\mathbb{S}^f$ is based on a *finite* path. Together with the fact that f is winning for Refuter, the restriction on f further guarantees that each $\Gamma(\rho)$ satisfies certain saturation properties, which are spelled out in the following lemma. We will later use these saturation conditions to construct our positional strategy $\underline{f}$ for $\forall$ in $\mathcal{E}(\bigwedge \Sigma, \mathbb{S}^f)$ and to show that $\mathbb{S}^f$ falsifies Γ^- with respect to $\underline{f}$.

5.4.6. Lemma. *For every state ρ of $\mathbb{S}^f$, the set $\Gamma(\rho)$ is* saturated. *That is, it satisfies all of the following conditions:*

- *For no φ it holds that $\varphi, \overline{\varphi} \in \Gamma(\rho)^-$.*

- *For all φ it holds that $\varphi^\circ \in \Gamma(\rho)$ if and only if $\overline{\varphi}^\circ \notin \Gamma(\rho)$*

- *For all φ it holds that $\varphi \rightsquigarrow \psi \in \Gamma(\rho)$ if and only if $\varphi \not\rightsquigarrow \psi \notin \Gamma(\rho)$.*

- *For no φ it holds that $\varphi \rightsquigarrow \varphi \in \Gamma(\rho)$.*

- *If $\psi_1 \vee \psi_2 \in \Gamma(\rho)^-$, then for both i: $\psi_1 \vee \psi_2 \not\rightsquigarrow \psi_i \in \Gamma(\rho)$ and $\psi_i \in \Gamma(\rho)^-$.*

- *If $\psi_1 \wedge \psi_2 \in \Gamma(\rho)^-$, then for some i: $\psi_1 \wedge \psi_2 \not\rightsquigarrow \psi_i \in \Gamma(\rho)$ and $\psi_i \in \Gamma(\rho)^-$.*

- *If $\mu x\varphi \in \Gamma(\rho)^-$, then $\varphi[\mu x\varphi/x] \in \Gamma(\rho)^-$.*

- *If $\nu x\varphi \in \Gamma(\rho)^-$, then $\nu x\varphi \not\rightsquigarrow \varphi[\nu x\varphi/x] \in \Gamma(\rho)$ and $\varphi[\nu x\varphi/x] \in \Gamma(\rho)^-$.*

- *If $\nu x\varphi \in \Gamma(\rho)^-$, then $\varphi[\nu x\varphi/x] \rightsquigarrow \nu x\varphi \in \Gamma(\rho)$.*

- *If $\varphi \not\rightsquigarrow \psi, \psi \not\rightsquigarrow \chi \in \Gamma(\rho)$, then $\varphi \not\rightsquigarrow \chi \in \Gamma(\rho)$.*

Proof:
We will prove two illustrative items, leaving the other items to the reader. For instance, the fact that $\varphi \rightsquigarrow \varphi \notin \Gamma(\rho)$ follows from the fact that f is winning for Refuter and the presence of the axiom Ax3. Now suppose that $\nu x\varphi \in \Gamma(\rho)^-$. Then $\nu x\varphi^b$ occurs in some Γ on ρ. It follows that $\nu x\varphi^c$ must occur in some Δ on ρ which is in stage 3 of Prover's strategy. If the saturation conditions for $\nu x\varphi$ do not already hold, Prover will in stage 3 cumulatively and productively apply R_ν with $\nu x\varphi$ as prinicipal formula. As a result, we will have $\nu x\varphi \not\rightsquigarrow \varphi[\nu x\varphi/x] \in \Gamma(\rho)$ and $\varphi[\nu x\varphi/x] \in \Gamma(\rho)^-$, and $\varphi[\nu x\varphi/x] \rightsquigarrow \nu x\varphi \in \Gamma(\rho)$, as required. □

Now let ρ_0 be any state of $\mathbb{S}^f$ containing the root Γ of f. We wish to show that Γ^- is *not* satisfied at ρ_0 in $\mathbb{S}^f$. To this end, we will construct a strategy $\underline{f}$ for $\forall$ in the game $\mathcal{E} := \mathcal{E}(\bigwedge \Sigma, \mathbb{S}^f)$ which is winning (φ_0, ρ_0) for every $\varphi_0 \in \Gamma^-$. The strategy $\underline{f}$ is defined as follows:

- At $(\psi_1 \wedge \psi_2, \rho)$, pick a conjunct $\psi_i \in \Gamma(\rho)^-$ such that $\psi_1 \wedge \psi_2 \not\rightsquigarrow \psi_i \in \Gamma(\rho)$.

- At $([a]\varphi, \rho)$, choose (φ, τ) for some τ such that $\rho \xrightarrow{a} \tau$ by virtue of some application of $\mathsf{R}_{[a]}$ with $[a]\varphi^b$ principal.

Before we show that $\underline{f}$ is winning for $\forall$, we must first argue that it is well defined. By saturation, for every formula $\psi_1 \wedge \psi_2$ contained in $\Gamma(\rho)^-$, there is a $\psi_i \in \Gamma(\rho)^-$ with $\psi_1 \wedge \psi_2 \not\rightsquigarrow \psi_i \in \Gamma(\rho)$. Likewise, for every formula $[a]\varphi^b \in \Gamma(\rho)$, there is a τ directly above ρ in f, separated only by an application of $\mathsf{R}_{[a]}$ with $[a]\varphi^b$ principal. The following lemma therefore shows that $\underline{f}$ is well-defined, at least for $\mathcal{E}$ initialised at a position (ψ_0, τ_0) such that $\psi_0 \in \tau_0^-$.

5.4.7. Lemma. *Let $\mathcal{M}$ be an $\underline{f}$-guided $\mathcal{E}$-match initialised at some position (ψ_0, τ_0) such that $\psi_0 \in \Gamma(\tau_0)^-$. Then for any position (ψ, τ) occurring in $\mathcal{M}$ it holds that $\psi \in \Gamma(\tau)^-$. Moreover, if (ψ, τ) comes directly after a modal step and the focus rule is applied on τ, then $\psi^\bullet \in \Gamma(\tau)$.*

Proof:
Denote the n-th position of $\mathcal{M}$ by (ψ_n, τ_n). We proceed by induction on n. The base case is simply the assumption that $\psi_0 \in \Gamma(\tau_0)^-$. For the induction step, suppose (ψ_n, τ_n) is such that $\psi_n \in \Gamma(\tau_n)^-$, and the next position is (ψ_{n+1}, τ_{n+1}). We make a case distinction based on the shape of ψ_n. Note that $\psi_n \notin \{p, \overline{p}\}$, for otherwise there would not be a next position (ψ_{n+1}, τ_{n+1}).

If the main connective of ψ_n is among $\{\vee, \mu, \nu\}$, it follows directly from saturation that ψ_{n+1} belongs to $\Gamma(\tau_{n+1})^-$. If ψ_n is a conjunction, then ψ_{n+1} is the conjunct of $\underline{f}(\psi_n, \tau)$, which by the definition of $\underline{f}$ belongs to $\Gamma(\tau_{n+1})^-$.

Now suppose ψ_n is of the form $\langle a \rangle \chi$. Then $\tau_n R^f_a \tau_{n+1}$, so either $\tau_n \xrightarrow{a} \tau_{n+1}$ or $\tau_{n+1} \xrightarrow{\breve{a}} \tau_n$. If $\tau_n R^f_a \tau_{n+1}$ we clearly have $\psi_{n+1} = \chi \in \Gamma(\tau_{n+1})^-$, by case 1(b) of Definition 5.2.10. Moreover, since in particular $\psi^b_{n+1} \in \mathsf{first}(\tau_{n+1})$, it follows from the restriction on f that in case the focus rule is applied in τ_{n+1}, we have $\psi^\bullet_{n+1} \in \Gamma(\tau_{n+1})$. If $\tau_{n+1} \xrightarrow{\breve{a}} \tau_n$, we argue by contradiction:

$$\begin{aligned}
\chi \notin \Gamma(\tau_{n+1})^- &\Rightarrow \overline{\chi} \in \Gamma(\tau_{n+1})^- && \text{(Saturation)}\\
&\Rightarrow [a]\overline{\chi} \in \Gamma(\tau_n)^- && \text{(1(c) of Definition 5.2.10)}\\
&\Rightarrow \langle a \rangle \chi \notin \Gamma(\tau_n)^-, && ([a]\overline{\chi} = \overline{\langle a \rangle \chi} \in \Sigma, \text{ Saturation})
\end{aligned}$$

which indeed contradicts the inductive hypothesis that $\langle a \rangle \chi \in \Gamma(\tau_n)^-$. Moreover, if the focus rule is applied in τ_{n+1}, we again argue by contradiction. Suppose $\chi^\bullet \notin \Gamma(\tau_{n+1})$. Then τ^-_{n+1} does not contain χ° after phase (1), whence because of the exhaustively applying productive instance of cut the formula $\overline{\chi}^\circ$ must have been added in stage (1). Hence $\overline{\chi} \in \Gamma(\tau_{n+1})^-$. But then saturation gives $\chi \notin \Gamma(\tau_{n+1})^-$, and we can use the same argument as before. Finally, the case where ψ_n is of the form $[a]\psi$ is similar to the easy part of the previous case and therefore left to the reader. □

The following lemma is key to the completeness proof. It shows that if an $\underline{f}$-guided $\mathcal{E}$-match can loop from some state ρ to itself, without passing through a μ-formula, then this information is already contained in ρ in the form of a negated trace atom.

5.4.8. Lemma. *Let $\rho \in S^f$ and let φ be such that $\varphi \in \Gamma(\rho)^-$. Then for every ψ such that $\mathbb{S}^f, \rho \Vdash_f \varphi \rightsquigarrow \psi$ it holds that $\varphi \not\rightsquigarrow \psi \in \Gamma(\rho)$.*

Proof:
Let $\mathcal{N}$ be the match witnessing that $\mathbb{S}^f, \rho \Vdash_f \varphi \rightsquigarrow \psi$. Recall that, by definition,

this means that the match $\mathcal{N}$ is an f-guided match starting at (φ, ρ) and ending at (ψ, ρ), such that only the formula in the final position, namely ψ, may be a μ-formula. We proceed by induction on the number of distinct states occurring in $\mathcal{N}$.

For the base case, we assume that ρ is the only state visited in $\mathcal{N}$. We proceed by induction on the length $n+1$ of $\mathcal{N}$. For the (inner) base case, where $|\mathcal{N} = 1|$, we have $\mathsf{first}(\mathcal{N}) = (\varphi, \rho) = \mathsf{last}(\mathcal{N})$. By saturation $\varphi \rightsquigarrow \varphi \notin \Gamma(\rho)$ and thus $\varphi \not\rightsquigarrow \varphi \in \Gamma(\rho)$, as required. For the inductive step, suppose the claim holds for every match up to size $n+1$. Suppose $|\mathcal{N}| = n+2$ and consider the final transition $(\chi, \rho) \cdot (\psi, \rho)$ of $\mathcal{N}$. Since the match proceeds after the position (χ, ρ), but does not move to a new state of $\mathbb{S}^f$, it follows from the irreflexivity of $\mathbb{S}^f$ that the main connective of χ must be among $\{\vee, \wedge, \mu, \nu\}$. Moreover, by Lemma 5.4.7, we have $\chi \in \Gamma(\rho)^-$. We claim that $\chi \not\rightsquigarrow \psi \in \Gamma(\rho)^-$. When the main connective of χ is in $\{\vee, \mu, \nu\}$, this follows directly from saturation. If χ is a conjunction, we have, since $\mathcal{M}$ is $\underline{f}$-guided, that $(\psi, \rho) = \underline{f}(\chi, \rho)$. By the definition of $\underline{f}$, it follows that $\chi \not\rightsquigarrow \psi \in \Gamma(\rho)$, as required. We finish the proof of this special case of the lemma by applying the induction hypothesis to the initial segment of $\mathcal{N}$ obtained by removing the last position (ψ, ρ). This gives $\varphi \not\rightsquigarrow \chi \in \Gamma(\rho)$, hence by saturation $\varphi \not\rightsquigarrow \psi \in \Gamma(\rho)$.

For the (outer) inductive step, suppose that $n > 1$ states are visited in $\mathcal{N}$. We write $\mathcal{N}$ as $\mathcal{A}_1 \cdot \mathcal{B}_1 \cdot \mathcal{A}_2 \cdot \mathcal{B}_2 \cdots \mathcal{A}_m$, where for every (χ, τ) in $\mathcal{A}_i$ it holds that $\tau = \rho$ and for every (χ, τ) in $\mathcal{B}_i$ it holds that $\tau \neq \rho$. As $\mathbb{S}^f$ is a forest, there must for each $\mathcal{B}_i$ be some γ_i, δ_i, and τ_i such that $\mathsf{first}(\mathcal{B}_i) = (\gamma_i, \tau_i)$ and $\mathsf{last}(\mathcal{B}_i) = (\delta_i, \tau_i)$. Denote $\mathsf{first}(\mathcal{A}_i) = (\alpha_i, \rho)$ and $\mathsf{last}(\mathcal{A}_i) = (\beta_i, \rho)$. Summing up, we will we use the following notation for each $i \in [1, m)$:

$$\mathsf{first}(\mathcal{A}_i) = (\alpha_i, \rho), \quad \mathsf{last}(\mathcal{A}_i) = (\beta_i, \rho), \quad \mathsf{first}(\mathcal{B}_i) = (\gamma_i, \tau_i), \quad \mathsf{last}(\mathcal{B}_i) = (\delta_i, \tau_i).$$

Let $i \in [1, m)$ be arbitrary. Since $\mathcal{B}_i$ does not visit ρ, it must visit strictly less states than $\mathcal{N}$. By the induction hypothesis we find that $\gamma_i \not\rightsquigarrow \delta_i \in \Gamma(\tau_i)$. We claim that $\beta_i \not\rightsquigarrow \alpha_{i+1} \in \Gamma(\rho)$. Since the match $\mathcal{N}$ transitions from the state ρ to the state τ_i, there must be some $a \in \mathsf{D}$ such that $\rho R^f_a \tau_i$.

We first assume that $\rho \xrightarrow{a} \tau_i$. Then by the nature of the game, β_i must be of the form $\beta_i = \langle a \rangle \gamma_i$ or of the form $\beta_i = [a]\gamma_i$, and, since by definition $\underline{f}$ only moves upward in $\mathbb{S}^f$, we must have $\delta_i = \langle \breve{a} \rangle \alpha_{i+1}$. We only cover the case where $\beta_i = [a]\gamma_i$ (the case where $\beta_i = \langle a \rangle \gamma_i$ is almost the same, but uses 2(c) instead of 2(a) of Definition 5.2.10). We indeed find:

$$\begin{array}{ll}
\gamma_i \not\rightsquigarrow \langle \breve{a} \rangle \alpha_{i+1} \in \Gamma(\tau_i) & \text{(Induction hypothesis, } \delta_i = \langle \breve{a} \rangle \alpha_{i+1}\text{)} \\
\Rightarrow \gamma_i \rightsquigarrow \langle \breve{a} \rangle \alpha_{i+1} \notin \Gamma(\tau_i) & \text{(Saturation)} \\
\Rightarrow [a]\gamma_i \rightsquigarrow \alpha_{i+1} \notin \Gamma(\rho) & \text{(Case 2(a) of Definition 5.2.10)} \\
\Rightarrow \beta_i \not\rightsquigarrow \alpha_{i+1} \in \Gamma(\rho), & \text{(Saturation, } \beta_i = [a]\gamma_i\text{)}
\end{array}$$

Now suppose that $\tau_i \xrightarrow{\breve{a}} \rho$. Then β_i must be of the form $\beta_i = \langle a \rangle \gamma_i$, because the strategy f moves only upward in $\mathbb{S}^f$. Moreover, we have $\delta_i = [\breve{a}]\alpha_{i+1}$ or $\delta_i = \langle \breve{a} \rangle \alpha_{i+1}$. An argument similar to the one above, respectively using cases 2(b) and 2(d) of Definition 5.2.10, shows that $\langle a \rangle \gamma_i \not\rightsquigarrow \alpha_{i+1} \in \Gamma(\rho)$.

Since ρ is the only state visited in $\mathcal{A}_i$, we can apply the (outer) base case of the induction hypothesis to the $\mathcal{A}_i$, to obtain $\alpha_i \not\rightsquigarrow \beta_i \in \Gamma(\rho)$ for every $1 \leq i \leq m$. Hence, by saturation, we find $\gamma_1 \not\rightsquigarrow \delta_m \in \Gamma(\rho)$, as required. □

Before we proceed to the completeness theorem, we first prove a final lemma. It shows that any infinite match in $\mathcal{E}(\bigwedge \Sigma, \mathbb{S}^f)$ either visits some state τ of $\mathbb{S}^f$ infinitely often, or can be split up into an upward path, interspersed with several detours, each of which returns to the same state as it departed from. Recall that, for two states τ_1 and τ_2 of $\mathbb{S}^f$, the order $\tau_1 \leq \tau_2$ means that the path τ_1 occurs below the path τ_2 in the strategy tree f (or they are exactly the same paths). More formally, we defined the relation $\tau_1 \leq \tau_2$ as the reflexive-transtive closure of the relation $\tau_1 \to \tau_2$, which in turn was the union of all relations $\tau_1 \xrightarrow{a} \tau_2$ where $a \in \mathsf{D}$.

5.4.9. Lemma. *Let $(\tau_0, \psi_0), (\tau_1, \psi_1), \ldots$ be an infinite $\mathcal{E}(\bigwedge \Sigma, \mathbb{S}^f)$-match such that for every $n \geq 0$ the following hold:*

(i) $\tau_n \geq \tau_0$;

(ii) there are only finitely many $m \geq 0$ such that $\tau_n = \tau_m$.

Then there is a subsequence

$$(\tau_{\alpha(0)}, \psi_{\alpha(0)}), (\tau_{\alpha(1)}, \psi_{\alpha(1)}), (\tau_{\alpha(2)}, \psi_{\alpha(2)}), \ldots$$

such that $\alpha(0) > 0$, and for every $i \geq 0$ the formula $\psi_{\alpha(i)-1}$ is modal, the state $\tau_{\alpha(i+1)-1}$ is equal to the state $\tau_{\alpha(i)}$, and there is an $a_i \in \mathsf{D}$ such that $\tau_{\alpha(i)} \xrightarrow{a_i} \tau_{\alpha(i+1)}$.

Proof:
We define the sequence $\alpha(0), \alpha(1), \ldots$ by recursion. In addition to the properties required by the lemma, we will show that the following holds for each $\alpha(i)$:

$$\text{For every } n \geq \alpha(i) \text{ it holds that } \tau_n \geq \tau_{\alpha(i)}. \tag{\dagger}$$

We define $\alpha(-1) := 0$, so that we can cover the recursion base and recursion step in one go. Note that $\alpha(-1)$ satisfies (†) by assumption (i).

Now suppose that $\alpha(i)$ has been defined. We let $\alpha(i+1)$ be least such that $\tau_n > \tau_{\alpha(i)}$ for every $n \geq \alpha(i+1)$. Such an $\alpha(i+1)$ must exist, because by (†) we have $\tau_n \geq \tau_{\alpha(i)}$ for all $n \geq \alpha(i)$ and by (ii) there are only finitely many $n \geq \alpha(i)$ such that $\tau_n = \tau_{\alpha(i)}$.

We will now show that $\alpha(i+1)$ satisfies the required conditions. First, note that $\alpha(i+1) > \alpha(-1) = 0$. Next, the formula $\psi_{\alpha(i+1)-1}$ must be modal, for otherwise $\tau_{\alpha(i+1)-1} = \tau_{\alpha(i+1)}$, contradicting the minimality of $\alpha(i+1)$. Moreover, we claim that $\tau_{\alpha(i+1)-1} = \tau_{\alpha(i)}$. Indeed, if not, then it will by (†) hold that $\tau_{\alpha(i+1)-1} > \tau_{\alpha(i)}$, again contradicting the minimality of $\alpha(i+1)$.

Hence there exists some $a_i \in \mathsf{D}$ such that either $\psi_{\alpha(i+1)-1} = \langle a_i \rangle \psi_{\alpha_{i+1}}$ or $\psi_{\alpha(i+1)-1} = [a_i]\psi_{\alpha_{i+1}}$. As a result, the rules of the game dictate that $\tau_{\alpha(i)} R^f_{a_i} \tau_{\alpha(i+1)}$. By the definition of $R^f_{a_i}$, it follows that

$$\tau_{\alpha(i)} \xrightarrow{a_i} \tau_{\alpha(i+1)}, \text{ or } \tau_{\alpha(i+1)} \xrightarrow{a_i} \tau_{\alpha(i)}.$$

If $\tau_{\alpha(i+1)} \xrightarrow{a_i} \tau_{\alpha(i)}$, then $\tau_{\alpha(i+1)} < \tau_{\alpha(i)}$, contradicting the definition of $\alpha(i+1)$. It thus follows that $\tau_{\alpha(i)} \xrightarrow{a_i} \tau_{\alpha(i+1)}$, as required. □

With the above lemmata in place, we are ready to prove that ∀ wins every full $\underline{f}$-guided $\mathcal{E}@(\varphi_0, \rho_0)$-match $\mathcal{M}$. If $\mathcal{M}$ is finite, it is not hard to show that it must be ∃ who got stuck. If $\mathcal{M}$ is infinite, the proof depends on whether $\mathcal{M}$ visits some single state infinitely often. If it does, one can show that if ∃ would win the match $\mathcal{M}$, then $\mathcal{M}$ would visit some state ρ with $\nu x\varphi, \varphi[\nu x\varphi/x] \not\rightsquigarrow \varphi \in \Gamma(\rho)^-$, contradicting saturation. If, on the other hand, $\mathcal{M}$ visits each state at most finitely often, the proof works by showing that a win for ∃ in $\mathcal{M}$ would imply that f contains an infinite branch won by Prover, which is also a contradiction.

5.4.10. Proposition. *Let ρ_0 be a state of $\mathbb{S}^f$ containing the root Γ of f and let $\varphi_0 \in \Gamma^-$. Then the strategy $\underline{f}$ is winning for ∀ in $\mathcal{E}@(\varphi_0, \rho_0)$.*

Proof:
Let $\mathcal{M}$ be an arbitrary $\underline{f}$-guided and full $\mathcal{E}@(\varphi_0, \rho_0)$-match. Since $\varphi_0 \in \Gamma(\rho_0)^-$, it follows from Lemma 5.4.7 that $\overline{f}$ is well-defined on $\mathcal{E}@(\varphi_0, \rho_0)$. By positional determinacy, we may without loss of generality assume that ∃ adheres to some positional strategy in $\mathcal{M}$. First suppose that $\mathcal{M}$ is finite, ending in, say (φ, ρ). We make a case distinction on the shape of φ.

If φ is a propositional letter p, we find:

$$\varphi = p \Rightarrow p \in \Gamma(\rho)^- \Rightarrow \overline{p} \notin \Gamma(\rho)^- \Rightarrow \mathbb{S}^f, \rho \not\Vdash p,$$

where the first implication holds due to Lemma 5.4.7, the second due to saturation, and the third by the definition of the valuation function of $\mathbb{S}^f$. It follows that in this case ∃ gets stuck.

Similarly, if φ is a negated propositional letter $\overline{p}$, we find:

$$\varphi = \overline{p} \Rightarrow \overline{p} \in \Gamma(\rho)^- \Rightarrow \mathbb{S}^f, \rho \Vdash p \Rightarrow \mathbb{S}^f, \rho \not\Vdash \overline{p},$$

hence again ∃ gets stuck.

Finally, we claim that φ is not of the form $[a]\psi$. Indeed, in that case the fact that $[a]\psi \in \Gamma(\rho)^-$ would entail that the modal rule is applicable. Hence $\underline{f}(\varphi, \rho)$ would be defined, contradicting the assumed fullness of $\mathcal{M}$.

Now suppose that $\mathcal{M}$ is infinite, say $\mathcal{M} = (\varphi_n, \rho_n)_{n\in\omega}$. Suppose first that some state ρ is visited infinitely often in $\mathcal{M}$. By the pigeonhole principle, there must be a formula φ and segment $\mathcal{N}$ of $\mathcal{M}$ such that $\mathsf{first}(\mathcal{N}) = \mathsf{last}(\mathcal{N}) = (\varphi, \rho)$. Since both players follow a positional strategy, we can write the match $\mathcal{M}$ as $\mathcal{K}\mathcal{N}^*$, where $\mathcal{K}$ is some initial segment of $\mathcal{M}$. But this means that only finitely many states of $\mathbb{S}^f$ occur in $\mathcal{M}$. As $\mathcal{M}$ is winning for $\exists$, there must, by Lemma 2.1.37, be some formula $\nu x\psi$ occurring infinitely often in $\mathcal{M}$. Therefore, there must be a position $(\nu x\psi, \tau)$ occurring infinitely often in $\mathcal{M}$. But then Lemma 5.4.8 gives $\varphi[\nu x\varphi/x] \not\rightsquigarrow \nu x\psi \in \Gamma(\tau)$. But by saturation we also have $\varphi[\nu x\varphi/x] \rightsquigarrow \nu\varphi \in \Gamma(\tau)$, which is in contradiction with the third item of the Lemma 5.4.6.

Hence we may assume that $\mathcal{M}$ visits each state ρ at most finitely often. Suppose, towards a contradiction, that $\mathcal{M}$ is won by $\exists$. Then, by Lemma 2.1.37, there is some $k \geq 0$ such that no formula φ_n with $n \geq k$ is a μ-formula. Moreover, since $\mathcal{M}$ visits each state at most finitely often, there must be some $l \geq k$ such that for every ρ_n with $n \geq l$ it holds that $\rho_n \geq \rho_l$. Let

$$\mathcal{N} = (\varphi_l, \rho_l) \cdot (\varphi_{l+1}, \rho_{l+1}) \cdot (\varphi_{l+2}, \rho_{l+2}) \cdots$$

be the final segment of $\mathcal{M}$ generated by (φ_l, ρ_l). Since for every ρ_n with $n \geq l$ it holds that $\rho_n \geq \rho_l$, we can apply Lemma 5.4.9. Let

$$(\varphi_{\alpha(0)}, \rho_{\alpha(0)}), (\varphi_{\alpha(1)}, \rho_{\alpha(1)}), (\varphi_{\alpha(2)}, \rho_{\alpha(2)}), \ldots$$

be a subsequence of $\mathcal{N}$ as given by Lemma 5.4.9. Then for every i there is an $a_i \in \mathsf{D}$ with $\rho_{\alpha(i)} \xrightarrow{a_i} \rho_{\alpha(i+1)}$. Hence, we have an f-guided $\mathcal{G}(\Gamma)$-match

$$\mathcal{K} = \rho_{\alpha(0)} \cdot \mathsf{R}_{[a_0]} \cdot \rho_{\alpha(1)} \cdot \mathsf{R}_{[a_1]} \cdot \rho_{\alpha(2)} \cdot \mathsf{R}_{[a_2]} \cdots$$

Note that $\mathcal{K}$ is infinite, as $\mathcal{N}$ visits infinitely many states. Because f is by assumption winning for Refuter, the focus rule must be applied infinitely often on $\mathcal{K}$.

Let $\rho_{\alpha(i)}$ with $i > 0$ be a segment on which the focus rule is applied. By Lemma 5.4.9, we have that $\varphi_{\alpha(i)-1}$ is modal, hence we obtain by Lemma 5.4.7 that $\varphi^\bullet_{\alpha(i)} \in \Gamma(\rho_{\alpha(i)})$. We claim that for every $j > i$ it holds that every sequent in $\rho_{\alpha(j)}$ has a formula in focus. With this we reach the desired contradiction, because it means that the focus rule cannot be applied on this final segment of $\mathcal{K}$ after all.

In particular, we will show that $\varphi^\bullet_{\alpha(j)} \in \mathsf{first}(\rho_{\alpha(j)})$ for every $j > i$, which suffices by the restriction of f to cumulative rule applications. For this, in turn, it is enough to show that for every $j \geq i$: if we have $\varphi^\bullet_{\alpha(j)} \in \Gamma(\rho_{\alpha(j)})$, then we have $\varphi^\bullet_{\alpha(j+1)} \in \mathsf{first}(\rho_{\alpha(j+1)})$.

To that end, consider the following submatch of $\mathcal{N}$.

$$\mathcal{J} = (\varphi_{\alpha(j)}, \rho_{\alpha(j)}) \cdots (\varphi_{\alpha(j+1)-1}, \rho_{\alpha(j+1)-1}).$$

Recall that $\mathcal{N}$ was constructed in such a way that it contains no μ-formulas. Hence $\mathcal{J}$ contains no μ-formulas. Moreover, by Lemma 5.4.9, we have $\rho_{\alpha(j+1)-1} = \rho_{\alpha(j)}$. Therefore may apply Lemma 5.4.8 to obtain $\varphi_{\alpha(j)} \not\rightsquigarrow \varphi_{\alpha(j+1)-1} \in \Gamma(\rho_{\alpha(i)})$. By Lemma 5.4.9, the formula $\varphi_{\alpha(j+1)-1}$ must be of the form $\langle a_j \rangle \varphi_{\alpha(j+1)}$ or of the form $[a_j]\varphi_{\alpha(j+1)}$. In either case, part 1 of Definition 5.2.10 gives $\varphi^{\bullet}_{\alpha(j+1)} \in \mathsf{first}(\rho_{\alpha(j+1)})$, as required. □

Since 5.4.10 holds for an arbitrary φ_0 in Γ^-, we find that $\mathbb{S}^f \not\Vdash_{\underline{f}} \Gamma^-$. Hence, by Theorem 2.2.10, we obtain completeness for the formula part of sequents.

5.4.11. PROPOSITION. *If* Γ^- *is valid, then* Γ *has a* Focus^2_∞*-proof.*

5.5 Conclusion

We have constructed a non-well-founded proof system Focus^2_∞ for the two-way alternation-free modal μ-calculus $\mathcal{L}^{af}_{2\mu}$. This system naturally reduces to a cyclic system when restricting to positional strategies in the proof search game.

Using the proof search game and the game semantics for the modal μ-calculus, we have shown that the system is sound for all sequents, and complete for those sequents not containing trace atoms. A natural first question for future research is to see if a full completeness result can be obtained. For this, a logic of trace atoms would have to be developed. One could for instance think of a rule like

$$\dfrac{\varphi \rightsquigarrow \chi, \Gamma \qquad \psi \rightsquigarrow \chi, \Gamma}{\varphi \wedge \psi \rightsquigarrow \chi, \Gamma}\ \mathsf{R}_\wedge$$

Following on this, we think it would be interesting to properly include trace atoms in the syntax by allowing the Boolean, modal and perhaps even the fixed point operators to apply to trace atoms. An example of a valid formula in this syntax is given by $((\varphi \rightsquigarrow \langle a \rangle \psi) \wedge [a](\psi \rightsquigarrow \langle \breve{a} \rangle \varphi)) \rightarrow \varphi$.

Another pressing question is whether our system could be used to prove interpolation, as has been done for language without backward modalities in [73]. To the best of our knowledge it is currently an open question whether $\mathcal{L}^{af}_{2\mu}$ has interpolation. At the same time, it is known that analytic applications of the cut rule do not necessarily interfere with the process of extracting interpolants from proofs [58, 76].

Finally, it would be interesting to see if our system can be extended to the full language $\mathcal{L}_{2\mu}$. The main challenge would be to keep track of the most important fixed point formula being unfolded on a trace. Perhaps this could be done by employing an annotation system such as the one by Jungteerapanich and Stirling [100, 54], together with trace atoms that record the most important fixed point variable unfolded on a loop.

Chapter 6

A cyclic proof system for Guarded Kleene Algebra with Tests

In this chapter we introduce a cyclic proof system for Guarded Kleene Algebra with Tests, or GKAT for short. This is a formal system used for reasoning about imperative programs. It draws from a long tradition, which we will briefly sketch here.

The first origin for GKAT is *Kleene Algebra*. Recall that, given a finite alphabet Σ, a *regular expression over* Σ is a string generated by the grammar:

$$e, f ::= 0 \mid 1 \mid a \in \Sigma \mid e + f \mid e \cdot f \mid e^*$$

Under the standard interpretation, regular expressions denote *languages over* Σ, *i.e.* sets of words over Σ. The constant 0 is interpreted as the empty language $\emptyset$, and the constant 1 as the language $\{\epsilon\}$ containing just the empty string. The interpretation of a is the language $\{a\}$ containing only a. The operators $+, \cdot, ^*$ are, respectively, interpreted as union, pairwise concatenation, and *Kleene closure*, which is the smallest language extension closed under concatenation. A language denoted by some regular expression is said to be a *regular language*.

Although Kleene Algebra appears under different interpretations in various areas of logic and computer science, it is most commonly understood as a generalisation of the algebra of regular languages under the operations $(0, 1, +, \cdot, ^*)$. There has been much interest in finding a nice axiomatisation of the equational validities of this algebra. The fragment without the Kleene star *, it turns out, is finitely axiomatised by the equational axioms of a certain algebraic structure called an *idempotent semiring*.

Extending this axiomatisation to incorporate the Kleene star has posed a formidable challenge. Numerous proposed axiomatisations exist, with one of the most prominent coming from Salomaa [91]. To capture the behaviour of the Kleene star, this system features the following rule:

$$\frac{eg + f \equiv g \qquad e \text{ does not have the empty word property}}{e^* f \equiv g}$$

where e is said to have the *empty word property* if the language denoted by e contains the empty word ϵ. Although Salomaa's system is sound and complete with respect to the algebra of regular languages, it is not entirely satisfactory, because it is not *algebraic*. More precisely, the empty word property is not closed under substitution and therefore itself not equationally axiomatisable. As a consequence, this system does not give rise a notion of "a Kleene algebra", in which equations may be true or false.

A solution to this problem was provided by Kozen in [60]. Using $e \leq f$ as a shorthand for $e + f \equiv f$, Kozen axiomatised e^*f as a least fixed point by adding the following axiom and rule:[1]

$$1 + ee^* \leq e^* \qquad\qquad \frac{eg + f \leq g}{e^*f \leq g}$$

Kozen showed that this system is complete with respect to the algebra of regular languages. With this, a *Kleene algebra* is then defined to be any algebra $(K, 0, 1, +, \cdot, ^*)$ such that $(K, 0, 1, +, \cdot)$ is an idempotent semiring, and the axiom and rule above are satisfied. It turns out that the algebra of regular language over Σ is then the free Kleene algebra generated by Σ.

Another way to interpret Kleene Algebra is as a semantics of programs. Under this interpretation, characters in Σ are seen as primitive programs, 0 is a program without any valid behaviour, and 1 is *skip*, which simply does nothing. The concatenation $e \cdot f$ is thought of as first running program e, and then running program f. The union $e + f$ non-deterministically runs e or f. Finally, e^* repeats the program e a finite number of times, possibly zero.

A shortcoming of using Kleene Algebra for this purpose is that is unable to express common programming constructs such as if-then-else statements and while loops. This inspired the development of an extension of Kleene Algebra, called Kleene Algebra with Tests [61], or KAT for short. KAT is a finite quasi-equational theory with two sorts, namely *programs* and a subset thereof consisting of *tests*, such that the programs form a Kleene algebra under the operations $(+, \cdot, *, 0, 1)$ and the tests form a Boolean algebra under the operations $(+, \cdot, ^-, 0, 1)$. The inclusion of tests allows one to express if-then-else statements and while loops. Despite the gain in expressive power, the complexity of deciding KAT-equations remains the same as for Kleene Algebra, namely PSPACE-complete.

Finally, the system GKAT, introduced in [98], is a fragment of KAT, obtained by replacing the operations $+$ and $*$ by their *guarded* counterparts $+_{(b)}$ and $-^{(b)}$. Roughly, this restricts the language to *only* if-then-else statements, rather than general non-deterministic choice, and to *only* while loops, rather than the general (non-deterministic) Kleene star. As a result, we will see that the languages denoted by GKAT-expressions satisfy a certain determinacy property. The main

[1] Kozen's original axiomatisation also has a dual axiom and rule characterising fe^* as a least fixed point, but it turns out those can be derived from the other rules.

advantage of GKAT over KAT lies in the efficiency of deciding program equivalence. For GKAT this can be done in *nearly linear* time, *i.e.* in $\mathcal{O}(n \cdot \alpha(n))$, where α is the very slow-growing inverse Ackermann function.

We will introduce a cyclic system SGKAT for GKAT. This system is inspired by a cyclic system for Kleene Algebra by Das & Pous in [34]. Our proofs of soundness and completeness are inspired by the same paper. Throughout the chapter we will remark on the differences and similarities between the two systems and the proofs of metalogical results. An important difference is that the determinacy property of GKAT allows us to use sequents with a simpler structure. More precisely, the succedents of our sequents will be lists of expressions, whereas the system in [34] needs multisets of lists to capture Kleene Algebra with a cyclic proof system.

In Section 6.1 we formally define the syntax and semantics of GKAT, and give a brief overview the foundational results about GKAT. In Section 6.2 we will present our non-well-founded proof system SGKAT^∞. Section 6.3 proves that SGKAT^∞ is sound with respect to the language model. In Section 6.4 we show that every proof is frugal, in the sense that it contains only finitely many distinct sequents, which we use to prove completeness in Section 6.5. By the results in the Intermezzo regular completeness follows: every sequent valid in the language model, has a regular proof in SGKAT^∞. By the theory in the Intermezzo, this immediately gives rise to a notion of cyclic proofs.

At the time of writing the only known axiomatisation for GKAT suffers from a similar defect as Salomaa's system for Kleene Algebra [98]. Even worse, a single axiom is not sufficient but an axiom schema is needed. More details about this will be given in Section 6.1.3. In Section 6.6 we propose an inequational axiomatisation PoGKAT for GKAT and document an attempt to prove its completeness by translating cyclic SGKAT-proofs into finite PoGKAT-proofs. This is inspired by a recent alternative proof of Kozen's completeness result for Kleene Algebra, by Das, Doumane & Pous [33], which likewise works by translating proofs from the aforementioned cyclic proof system for Kleene algebra into Kozen's system.

6.1 Preliminaries

6.1.1 Syntax

The language of GKAT has two sorts, namely *programs* and a subset thereof consisting of *tests*. It is built from a finite and non-empty set T of *primitive tests* and a non-empty set Σ of *primitive programs*, where T and Σ are disjoint. For the rest of this chapter we fix such sets T and Σ. We reserve the letters t and p to refer, respectively, to arbitrary primitive tests and programs. The first of the following grammars defines the *tests* of the language of GKAT and the second its *expressions.*

$$b, c ::= 0 \mid 1 \mid t \mid \overline{b} \mid b \vee c \mid b \cdot c \qquad e, f ::= b \mid p \mid e \cdot f \mid e +_b f \mid e^{(b)},$$

where $t \in T$ and $p \in \Sigma$.

Note that the tests are simply propositional formulas. It is convention to use $\cdot$ instead of $\wedge$ for conjunction. As we will later see, the interpretation of $b \cdot c$ where b and c are regarded as tests, is the same as the interpretation of $b \cdot c$, where b and c are regarded as expressions.

6.1.1. Example. The idea of GKAT is to model imperative programs. For instance, the expression $(p +_b q)^{(a)}$ represents the following imperative program:

```
while a do
  if b then
    p
  else
    q
  end
end
```

6.1.2. Remark. As mentioned in the introduction, GKAT is a fragment of Kleene Algebra with Tests, or KAT [61]. The syntax of KAT is the same as that of GKAT, but with unrestricted union $+$ instead of guarded union $+_b$, and unrestricted iteration $*$ instead of the while loop operator (b).

The embedding φ of GKAT into KAT acts on guarded union and guarded iteration as follows, and commutes with all other operators.

$$\varphi(e +_b f) = b \cdot \varphi(e) + \overline{b} \cdot \varphi(f) \qquad\qquad \varphi(e^{(b)}) = (b \cdot \varphi(e))^* \cdot \overline{b}$$

The restriction to guarded union and guarded iteration can be seen as restricting to *deterministic* programs. This point will be made precise after defining the semantics.

6.1.2 Semantics

There are different kinds of semantics for GKAT. In [98], a *language* semantics, a *relational* semantics, and a *probabilistic* semantics are given. In this chapter we will only be concerned with the language semantics.

We denote by At the set of *atoms* of the free Boolean algebra generated by $T = \{t_1, \dots t_n\}$. That is, At consists of all tests of the form $c_1 \cdot \ldots \cdot c_n$, where $c_i \in \{t_i, \overline{t_i}\}$ for each $1 \leq i \leq n$. Lowercase Greek letters $(\alpha, \beta, \gamma, \ldots)$ will be used to denote elements of At. A *guarded string* is an element of the regular set $\mathsf{At} \cdot (\Sigma \cdot \mathsf{At})^*$, that is, a string of the form

$$\alpha_1 p_1 \alpha_2 p_2 \cdot \ldots \cdot \ldots \cdot \alpha_n p_n \alpha_{n+1}.$$

We will interpret expressions as languages (formally just sets) of guarded strings. The interpretation of the sequential composition operator $\cdot$ is in terms of the *fusion product* $\diamond$ on languages of guarded strings, given by:

$$L \diamond K := \{x\alpha y \mid x\alpha \in L \text{ and } \alpha y \in K\}.$$

For the interpretation of $+_b$, we define the following operation of *guarded union* on languages for every set of atoms $B \subseteq \mathsf{At}$:

$$L +_B K := (B \diamond L) \cup (\overline{B} \diamond K),$$

where $\overline{B}$ is $\mathsf{At} \setminus B$. Finally, for the interpretation of (b), we stipulate:

$$L^0 := \mathsf{At} \qquad L^{n+1} := L^n \diamond L \qquad L^B := \bigcup_{n \geq 0} (B \diamond L)^n \diamond \overline{B}$$

The semantics of `GKAT` is now defined as follows:

$$[\![b]\!] := \{\alpha \in \mathsf{At} : \alpha \leq b\} \qquad [\![p]\!] := \{\alpha p \beta : \alpha, \beta \in \mathsf{At}\} \qquad [\![e \cdot f]\!] := [\![e]\!] \diamond [\![f]\!]$$
$$[\![e +_b f]\!] := [\![e]\!] +_{[\![b]\!]} [\![f]\!] \qquad [\![e^{(b)}]\!] := [\![e]\!]^{[\![b]\!]}$$

Note that the interpretation of $\cdot$ between tests, regarded as tests, is the same as the interpretation of $\cdot$ between tests, regarded as programs.

$$[\![b \cdot c]\!] = [\![b]\!] \cap [\![c]\!] = [\![b]\!] \diamond [\![c]\!].$$

6.1.3. Remark. While the semantics of expressions is explicitly defined, the semantics of tests is derived implicitly through the free Boolean algebra generated by T. It is conventional in the `GKAT` literature to address the Boolean content in this manner.

6.1.4. Example. In a guarded string, the atoms can be thought of as states of a machine, and the programs as executions. For instance, the guarded string $\alpha p \beta$ can be read as: the machine starts in state α, then executes program p, and ends in state β.

Let us briefly check which guarded strings of, say, the form $\alpha p \beta q \gamma$ belong to the interpretation $[\![(p +_b q)^{(a)}]\!]$ of the program of Example 6.1.1. First, we must have $\alpha \leq a$, for otherwise we would not enter the loop at all. Moreover, we must have $\alpha \leq b$, for otherwise q rather than p would be executed. Similarly, we find that $\beta \leq a, \overline{b}$. Since the loop is exited after two iterations, we must have $\gamma \leq \overline{a}$. Hence, we find

$$\alpha p \beta q \gamma \in [\![(p +_b q)^{(a)}]\!] \Leftrightarrow \alpha \leq a, b \text{ and } \beta \leq a, \overline{b} \text{ and } \gamma \leq \overline{a}.$$

The following lemmas will be useful later on. Note that it does not hold by definition, as $\diamond$ is not commutative.

6.1.5. Lemma. *$L^{n+1} = L \diamond L^n$ for every language L of guarded strings.*

Proof:
Since At is the identity element for the fusion operator, we have

$$L^{n+1} = \mathsf{At} \diamond \underbrace{L \diamond \cdot \ldots \cdot \diamond L}_{n+1 \text{ times}} = L \diamond \mathsf{At} \diamond \underbrace{L \diamond \cdot \ldots \cdot \diamond L}_{n \text{ times}} = L \diamond L^n,$$

as required. □

6.1.6. LEMMA. *Let p be a primitive program and let L and K be languages of guarded strings. Then $[\![p]\!] \diamond L = [\![p]\!] \diamond K$ implies $L = K$.*

Proof:
Since $[\![p]\!] = \{\alpha p \beta : \alpha, \beta \in \mathsf{At}\}$, we have

$$\gamma y \in L \Leftrightarrow \gamma p \gamma y \in [\![p]\!] \diamond L \Leftrightarrow \gamma p \gamma y \in [\![p]\!] \diamond K \Leftrightarrow \gamma y \in K,$$

as required. □

The fact that GKAT models deterministic programs is reflected in the fact that interpretations of GKAT-expressions satisfy the following *determinacy property*.

6.1.7. DEFINITION. A language L of guarded strings is said to be *deterministic* if for every $x\alpha y$ and $x\alpha z$ in L, either y and z are both empty, or both begin with the same primitive program.

6.1.8. EXAMPLE. For $\alpha \neq \beta$ and $p \neq q$, the languages on the left of the following table satisfy the determinacy property, while those on the right do not.

$\{\alpha p \beta, \beta\}$,	$\{\alpha p \beta, \alpha\}$.
$\{\alpha p \alpha, \alpha p \beta\}$	$\{\alpha p \alpha, \alpha q \beta\}$

The following proposition can be directly shown by a somewhat tedious induction on expressions. We omit this proof because nothing in the rest of this chapter formally depends on it. For a more conceptual proof we refer the reader to [98], where it follows as a corollary from their Theorem 5.8 and their automaton model.

6.1.9. PROPOSITION. *The denotation $[\![e]\!]$ of a GKAT-expression e is deterministic.*

6.1.10. REMARK. The language semantics of GKAT is the same as that of KAT (see [61]), in the sense that

$$[\![e]\!] = [\![\varphi(e)]\!],$$

where φ is the embedding from Remark 6.1.2 and e is any GKAT-expression.

6.1.3 Foundational results

In this subsection we briefly summarise some of the foundational results presented in [98]. Reading this subsection is not strictly necessary for understanding the rest of this chapter, but it may provide some helpful context and intuition about GKAT.

Automaton model

Automata for GKAT are given as coalgebras for the functor $G : X \mapsto (2+\Sigma\times X)^{\mathsf{At}}$. That is, a state $s \in X$ of a G-coalgebra, when given an atom $\alpha \in \mathsf{At}$, does one of three things: halt and accept, halt and reject, or execute a program $p \in \Sigma$ and move to a new state in X. A *G-automaton* is simply a G-coalgebra with a designated initial state.

In [98], it is shown that the languages of guarded strings accepted by some G-automaton, possibly with infinitely many states, are precisely the languages which satisfy the determinacy property of Definition 6.1.7. However, there are G-automata, even with finitely many states, whose language is not denoted by any GKAT-program.

To remedy this situation, the authors of [98] introduce the notion of *well-nestedness* of G-automata, the definition of which falls outside the scope of this thesis. They show for any GKAT-expression e how to construct a (finite) well-nested G-automaton $\mathbb{A}_e$ such that the language accepted by $\mathbb{A}_e$ is precisely $[\![e]\!]$. Conversely, they describe for any given well-nested G-automaton $\mathbb{A}$ a GKAT-expression $e_{\mathbb{A}}$ such that the language of $\mathbb{A}$ is precisely $[\![e_{\mathbb{A}}]\!]$.

Decision procedure

One of the main advantages of GKAT over KAT lies in the efficiency of deciding program equivalence, *i.e.* whether $[\![e]\!] = [\![f]\!]$ holds for two given expressions e and f. Roughly, the decision procedure for GKAT-expressions presented in [98] works by first converting e and f into G-automata $\mathbb{A}_e$ and $\mathbb{A}_f$. By construction the number of states of $\mathbb{A}_e$ and $\mathbb{A}_f$ will be linear in, respectively, the sizes of the expressions e and f. After applying a certain normalisation procedure on $\mathbb{A}_e$ and $\mathbb{A}_f$, a general algorithm for checking bisimilarity of coalgebras can be used to check whether their initial states are bisimilar.

Since G is a so-called *polynomial functor*, and the set of deterministic languages carries a G-coalgebra structure under which it is the final coalgebra for normal coalgebras, general coalgebraic theory entails that bisimilarity and language equivalence coincide on normal coalgebras. Hence, the given decision procedure is correct.

By virtue of the relatively small size of automata for GKAT-expressions, the decision procedure runs in time $\mathcal{O}(n \cdot \alpha(n))$, when $|\mathsf{At}|$ is constant, n is the sum of the sizes of the expressions e and f, and α is the inverse of the Ackermann function. Recall that the Ackermann function is a computable function which grows so fast that it is not definable by primitive recursion. Hence, the inverse of the Ackermann function is an extremely slow-growing function. One therefore also says that this procedure runs in *nearly linear* time. This is much more efficient than deciding KAT-equivalence, which is PSPACE-complete, even when the number of atoms is constant.

Axiomatisation

In [98] an axiomatisation for GKAT-equivalence was put forward. While there it is presented from a more algebraic perspective, we will present it explicitly as a proof system. For this will use the following definition of substitution.

6.1.11. Definition. A substitution is a function $\sigma : \Sigma \to \texttt{GKAT}$, assigning a GKAT-expression to each primitive program.

Given a substitution σ, we let $\widehat{\sigma} : \texttt{GKAT} \to \texttt{GKAT}$ be the unique map which extends σ such that $\widehat{\sigma}$ commutes with the guarded union, concatenation and while-loop operators, and such that $\sigma(b) = b$ for every test b.

The system is based on *equational logic*, of which the axioms and rules are given in Figure 6.1. For background we refer the reader to [21].

$$\frac{}{e \equiv e} \qquad \frac{e \equiv f}{f \equiv e} \qquad \frac{e \equiv f \qquad f \equiv g}{e \equiv g} \qquad \frac{e \equiv f}{\widehat{\sigma}(e) \equiv \widehat{\sigma}(g)}$$

$$\frac{e_1 \equiv f_1 \qquad e_2 \equiv f_2}{e_1 +_b e_2 \equiv f_1 +_b f_2} \qquad \frac{e_1 \equiv f_1 \qquad e_2 \equiv f_2}{e_1 \cdot e_2 \equiv f_1 \cdot f_2} \qquad \frac{e \equiv f}{e^{(b)} \equiv f^{(b)}}$$

Figure 6.1: The axioms and rules of equational logic in the signature of GKAT.

It moreover contains all of the following axioms (cf. [98, Figure 1]).

U1.	$e +_b e \equiv e$	S1.	$(e \cdot f) \cdot g \equiv e \cdot (f \cdot g)$
U2.	$e +_b f \equiv f +_{\overline{b}} e$	S2.	$0 \cdot e \equiv 0$
U3.	$(e +_b f) +_c g \equiv e +_{bc} (f +_c g)$	S3.	$e \cdot 0 \equiv 0$
U4.	$e +_b f \equiv be +_b f$	S4.	$1 \cdot e \equiv e$
U5.	$eg +_b fg \equiv (e +_b f) \cdot g$	S5.	$e \cdot 1 \equiv e$
W1.	$e^{(b)} \equiv ee^{(b)} +_b 1$	W2.	$(e +_c 1)^{(b)} \equiv (ce)^{(b)}$

Figure 6.2: The GKAT axioms from [98, Fig. 1].

6.1.12. Definition. The system EGKAT consists of the following axioms and rules.

1. All axioms and rules of equational logic, as given in Figure 6.1.

2. The axiom $b \equiv c$ for all tests b, c such that b is equivalent to c in classical logic.

3. All axioms from [98, Fig. 1], *i.e.* all axioms in Figure 6.2 above.

4. A fixed point rule of the form

$$\frac{g \equiv eg +_b f}{g \equiv e^{(b)} f}\ (\dagger)$$

   with a side condition (†).

We will not go into the technical details of the side condition (†) in the above definition. Roughly, it is a syntactic restriction on the loop body e, guaranteeing that e is *strictly productive*, *i.e.*, always executes at least one primitive program. It is shown in [98] that for every expression e there is a strictly productive expression f such that $e^{(b)}$ and $f^{(b)}$ are provably equivalent.

The soundness of the above axiomatisation, *i.e.* that $\mathsf{EGKAT} \vdash e \equiv f$ implies $[\![e]\!] = [\![f]\!]$ is not hard to show by induction on the length of derivations. The completeness is an open question, although completeness has been shown for an extension by a stronger fixed point axiom. We refer the reader to [98, Section 6] for more details. Note, however, that even if the above system were complete, it would still suffer from the same drawback as Salomaa's system discussed in the introduction: it is not algebraic, because the strict productivity condition in the fixed point rule is not closed under substitution.

More explicitly, if we instantiate the fixed point rule with $b = 1$ and with $e = f = g = p$, then the side condition is met and the rule is sound, as p is strictly productive. However, if we apply the substitution $p \mapsto 1$, the side condition is no longer met. In fact, the resulting instance of the rule is unsound, because the premiss $1 \equiv 1 \cdot 1 +_1 1$ is true in the language model, but the conclusion $1 \equiv 1^{(1)} \cdot 1$ is false. Indeed, it holds that $[\![1^{(1)}]\!] = [\![0]\!]$, because $1^{(1)}$ represents a never-ending loop.

6.2 The non-well-founded proof system SGKAT^∞

In this section we commence our proof-theoretical study of GKAT. We will present a cyclic sequent system for GKAT, which is inspired by the cyclic sequent system for Kleene Algebra presented in [34]. In passing, we will comment on the similarities and differences between our system and the earlier system.

6.2.1. Definition. A *sequent* is a triple (Γ, A, Δ), usually written $\Gamma \Rightarrow_A \Delta$, where $A \subseteq \mathsf{At}$ and Γ and Δ are (possibly empty) lists of GKAT-expressions.

The list on the left-hand side of a sequent is called its *antecedent*, and the list on the right-hand side its *succedent*. The symbol ϵ is used to refer to the empty list.

6.2.2. Definition. We say that a sequent $e_1, \ldots, e_n \Rightarrow_A f_1, \ldots, f_m$ is *valid* whenever $A \diamond [\![e_1 \cdot \ldots \cdot e_n]\!] \subseteq [\![f_1 \cdot \ldots \cdot f_n]\!]$.

We will often abuse notation by writing $[\![\Gamma]\!]$ instead of $[\![e_1 \cdot \ldots \cdot e_n]\!]$, where Γ is some list of expressions $e_1, \ldots, e_n$.

6.2.3. Example. An example of a valid sequent is given by

$$(cp)^{(b)} \Rightarrow_{\mathsf{At}} (p(cp +_b 1))^{(b)}.$$

The left-hand side denotes guarded strings of the form $\alpha_1 p \alpha_2 p \cdot \ldots \cdot \alpha_n p \alpha_{n+1}$ for which $\alpha_i \leq b, c$ for each $1 \leq i \leq n$, and $\alpha_{n+1} \leq \bar{b}$. Similarly, the right-hand side denotes guarded strings of the form $\alpha_1 p \alpha_2 p \cdot \ldots \cdot \alpha_n p \alpha_{n+1}$ such that for each $1 \leq i \leq n$ it holds that $\alpha_i \leq b$ and, in addition, $\alpha_i \leq c$ if i is even, and $\alpha_{n+1} \leq \bar{b}$. Clearly the antecedent is contained in the succedent.

6.2.4. Remark. Like the sequents for Kleene Algebra in [34], our sequents express language *inclusion*, rather than language equivalence. For Kleene Algebra this difference is insignificant, as the two notions are interdefinable using unrestricted union:

$$[\![e]\!] \subseteq [\![f]\!] \Leftrightarrow [\![e + f]\!] = [\![f]\!].$$

For GKAT, however, it is not clear how to define language inclusion in terms of language equivalence. As a result, an advantage of axiomatising language inclusion rather than language equivalence, is that the while-operator can be axiomatised as a *least* fixed point, eliminating the need for a *strict productivity* requirement as is present in the axiomatisation in [98]. This will become more clear in Section 6.6, where we propose an algebraic axiomatisation of GKAT.

Given a set of atoms A and a test b, we write $A \restriction b$ for the set $\{\alpha \in A : \alpha \leq b\}$. Note that this is the same as $A \diamond [\![b]\!]$.

The rules of the sequent system SGKAT are given in Figure 6.3. Importantly, the rules are always applied to the leftmost expression in a list (whether in the antecedent or in the succedent). Also note that the system has no propositional rules for tests, since the propositional reasoning is tucked away in the set of atoms labelling a sequent. This makes the sequent system much simpler, and is in line with the ordinary way of treating the (finitely many) tests in the GKAT literature.

6.2.5. REMARK. Following [34], we call k a 'modal' rule. The reason is simply that it looks like the rule k (sometimes called K or $\Box$) in the standard sequent calculus for basic modal logic. Our system also features a second modal rule, called k_0. Like k, this rule adds a primitive program p to the antecedent of the sequent. Since its premiss entails that $[\![\Gamma]\!] = [\![0]\!]$, the antecedent of its conclusion will denote the language $\emptyset$, and is therefore included in any antecedent Δ.

As usual, an SGKAT$^\infty$-*derivation* is a possibly infinite tree generated by the rules of SGKAT. Such a derivation is said to be *closed* if every leaf is an axiom.

6.2.6. DEFINITION. A closed SGKAT$^\infty$-derivation is said to be an SGKAT$^\infty$-*proof* if every infinite branch is *fair* for (b)-l, *i.e.* contains infinitely many applications of the rule (b)-l.

Left logical rules

$$\frac{\Gamma \Rightarrow_{A\restriction b} \Delta}{b, \Gamma \Rightarrow_A \Delta}\ b\text{-}l \qquad \frac{e, g, \Gamma \Rightarrow_A \Delta}{e \cdot g, \Gamma \Rightarrow_A \Delta}\ \cdot\text{-}l$$

$$\frac{e, \Gamma \Rightarrow_{A\restriction b} \Delta \qquad f, \Gamma \Rightarrow_{A\restriction \bar{b}} \Delta}{e +_b f, \Gamma \Rightarrow_A \Delta}\ +_b\text{-}l \qquad \frac{e, e^{(b)}, \Gamma \Rightarrow_{A\restriction b} \Delta \qquad \Gamma \Rightarrow_{A\restriction \bar{b}} \Delta}{e^{(b)}, \Gamma \Rightarrow_A \Delta}\ (b)\text{-}l$$

Right logical rules

$$(\dagger)\ \frac{\Gamma \Rightarrow_A \Delta}{\Gamma \Rightarrow_A b, \Delta}\ b\text{-}r \qquad \frac{\Gamma \Rightarrow_A e, f, \Delta}{\Gamma \Rightarrow_A e \cdot f, \Delta}\ \cdot\text{-}r$$

$$\frac{\Gamma \Rightarrow_{A\restriction b} e, \Delta \qquad \Gamma \Rightarrow_{A\restriction \bar{b}} f, \Delta}{\Gamma \Rightarrow_A e +_b f, \Delta}\ +_b\text{-}r \qquad \frac{\Gamma \Rightarrow_{A\restriction b} e, e^{(b)}, \Delta \qquad \Gamma \Rightarrow_{A\restriction \bar{b}} \Delta}{\Gamma \Rightarrow_A e^{(b)}, \Delta}\ (b)\text{-}r$$

Axioms and modal rules

$$\frac{}{\epsilon \Rightarrow_A \epsilon}\ \mathsf{id} \qquad \frac{}{\Gamma \Rightarrow_\emptyset \Delta}\ \bot \qquad \frac{\Gamma \Rightarrow_{\mathsf{At}} \Delta}{p, \Gamma \Rightarrow_A p, \Delta}\ \mathsf{k} \qquad \frac{\Gamma \Rightarrow_{\mathsf{At}} 0}{p, \Gamma \Rightarrow_A \Delta}\ \mathsf{k}_0$$

Figure 6.3: The rules of SGKAT. The side condition (†) requires that $A \restriction b = A$.

6.2.7. REMARK. In [93] a variant of GKAT is studied which omits the axiom called (S3) in Figure 6.2. This axiom, also called the *early termination axiom*, equates all programs which eventually fail. A denotational model of this variant of GKAT is given in the form of certain kinds of trees. We conjecture that SGKAT$^\infty$ without the rule k_0 is sound and complete with respect to this denotational model.

6.2.8. Example. Let $\Delta_1 := (p(cp +_b 1))^{(b)}$ and $\Delta_2 := cp +_b 1, \Delta_1$. The following proof Π_1 is an example of an SGKAT^∞-proof of the sequent of Example 6.2.3. We use $(\bullet)$ to indicate that the proof repeats itself at these leaves and, for the sake of readability, omit branches that can be immediately closed by an application of $\bot$. Note that Π_1 is, in fact, regular.

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{(cp)^{(b)} \Rightarrow_{\mathsf{At}} \Delta_1 \quad (\bullet)}{p, (cp)^{(b)} \Rightarrow_{\mathsf{At}\restriction bc} p, \Delta_1}\,\mathsf{k}}{p, (cp)^{(b)} \Rightarrow_{\mathsf{At}\restriction bc} c, p, \Delta_1}\,c\text{-}r}{p, (cp)^{(b)} \Rightarrow_{\mathsf{At}\restriction bc} cp, \Delta_1}\,\cdot\text{-}r}{p, (cp)^{(b)} \Rightarrow_{\mathsf{At}\restriction bc} \Delta_2}\,+_b\text{-}r}{c, p, (cp)^{(b)} \Rightarrow_{\mathsf{At}\restriction b} \Delta_2}\,c\text{-}l}{cp, (cp)^{(b)} \Rightarrow_{\mathsf{At}\restriction b} \Delta_2}\,\cdot\text{-}l \qquad \dfrac{\dfrac{\dfrac{\dfrac{}{\epsilon \Rightarrow_{\mathsf{At}\restriction \bar{b}} \epsilon}\,\mathsf{id}}{\epsilon \Rightarrow_{\mathsf{At}\restriction \bar{b}} \Delta_1}\,(b)\text{-}r}{\epsilon \Rightarrow_{\mathsf{At}\restriction \bar{b}} 1, \Delta_1}\,1\text{-}r}{\epsilon \Rightarrow_{\mathsf{At}\restriction \bar{b}} \Delta_2}\,+_b\text{-}r}{(cp)^{(b)} \Rightarrow_{\mathsf{At}} \Delta_2}\,(b)\text{-}l$$

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{(cp)^{(b)} \Rightarrow_{\mathsf{At}} \Delta_2}{p, (cp)^{(b)} \Rightarrow_{\mathsf{At}\restriction bc} p, (cp +_b 1), \Delta_1}\,\mathsf{k}}{p, (cp)^{(b)} \Rightarrow_{\mathsf{At}\restriction bc} p \cdot (cp +_b 1), \Delta_1}\,\cdot\text{-}r}{p, (cp)^{(b)} \Rightarrow_{\mathsf{At}\restriction bc} \Delta_1}\,(b)\text{-}r}{c, p, (cp)^{(b)} \Rightarrow_{\mathsf{At}\restriction b} \Delta_1}\,c\text{-}l}{cp, (cp)^{(b)} \Rightarrow_{\mathsf{At}\restriction b} \Delta_1}\,\cdot\text{-}l \qquad \dfrac{\dfrac{}{\epsilon \Rightarrow_{\mathsf{At}\restriction \bar{b}} \epsilon}\,\mathsf{id}}{\epsilon \Rightarrow_{\mathsf{At}\restriction \bar{b}} \Delta_1}\,(b)\text{-}r}{(cp)^{(b)} \Rightarrow_{\mathsf{At}} \Delta_1 \quad (\bullet)}\,(b)\text{-}l$$

To illustrate how the omission of branches that can be immediately closed by an application of $\bot$ works, let us write out the two applications of $+_b$-r in Π_1.

$$\dfrac{\epsilon \Rightarrow_{\mathsf{At}\restriction bc} cp, \Delta_1 \qquad \dfrac{}{\epsilon \Rightarrow_{\emptyset} 1, \Delta_1}\,\bot}{\epsilon \Rightarrow_{\mathsf{At}\restriction bc} \Delta_2}\,+_b\text{-}r \qquad \dfrac{\bot\,\dfrac{}{\epsilon \Rightarrow_{\emptyset} cp, \Delta_1} \qquad \epsilon \Rightarrow_{\mathsf{At}\restriction \bar{b}} 1, \Delta_1}{\epsilon \Rightarrow_{\mathsf{At}\restriction \bar{b}} \Delta_2}\,+_b\text{-}r$$

It can also be helpful to think of the set of atoms as *selecting* one of the premisses.

6.2.9. Remark. SGKAT^∞ is a path-based non-well-founded proof system, as defined by Definition I.2.1 of the Intermezzo. In fact, if one calls those finite paths *good* which contain an application of (b)-l, it is a *simple* proof system (cf. Definition I.2.7). This immediately gives rise to a notion of *cyclic* SGKAT-proofs based on finite trees with back edges.

For the sake of simplicity we shall in this chapter mostly be concerned with infinitary non-well-founded proofs. However, in Section 6.4, we shall show that every SGKAT^∞-proof is frugal. It will then follow by I.2.23 that completeness implies regular completeness.

6.2.10. Remark. Note that the rules of SGKAT are highly symmetric. Indeed, the only rules that behave differently on the left than on the right, are the b-rules and k_0. For the b-rules, note that b-l changes the set of atoms, while b-r uses a side condition. The asymmetry of k_0 is clear: the succedent of the premiss has a 0, whereas the antecedent does not. A third and final asymmetry will be introduced in Definition 6.2.6, where a soundness condition is imposed on infinite branches which is sensitive to (b)-l but not to (b)-r.

6.2.11. Remark. Recall that our system is inspired by the system in [34] for Kleene Algebra (without tests). In the first part of [34], sequents similar to ours are considered *i.e.* pairs of lists of expressions (without set of atoms, because there are no tests in ordinary Kleene Algebra). It turns out, however, that the resulting system is complete, but not *regularly complete*, in the sense that not every valid sequent has a *regular* proof.

Consider for instance the valid sequent $p^* \Rightarrow (pp)^* + (pp)^*p$. In words, this sequent expresses that any finite string of p's is either of even or of odd length. To the antecedent, one can only apply the following rule, which corresponds to our rule (b)-l.

$$\frac{p, p^* \Rightarrow (pp)^* + (pp)^*p \qquad \epsilon \Rightarrow (pp)^* + (pp)^*p}{p^* \Rightarrow (pp)^* + (pp)^*p} *\text{-}l$$

The premiss on the right is not hard to prove. Since ϵ is of even length, it can be shown to be included in $(pp)^*$. For the other premiss, however, we cannot choose one of $(pp)^*$ and $(pp)^*p$ because we do not yet know the parity of the antecedent's length (only that it is non-zero). This situation can be slightly improved by allowing rules to apply to expressions *inside* lists, rather than only to the leftmost expression. One could then keep applying $*$-l to the p^* on the antecedent, obtaining a derivation which looks like this.

$$\dfrac{\dfrac{\stackrel{\vdots}{p, p, p^* \Rightarrow (pp)^* + (pp)^*p} \qquad \epsilon \Rightarrow (pp)^* + (pp)^*p}{p, p^* \Rightarrow (pp)^* + (pp)^*p} *\text{-}l \qquad \epsilon \Rightarrow (pp)^* + (pp)^*p}{p^* \Rightarrow (pp)^* + (pp)^*p} *\text{-}l$$

This derivation will be a proof, because it is fair for $*$-l. However, it is not regular, because it does contain infinitely many distinct sequents. In general, as mentioned before, the resulting system will be complete, but not regularly complete.

The authors of [34] remedy this situation by moving to *hypersequents*. The right-hand side of a hypersequent is a multiset of lists, rather than just a list. This allows reasoning 'underneath' the union operator $+$, enabling one to break down the $(pp)^*$ and $(pp)^*p$ on the right-hand side, without having to choose either one.

Fortunately for us, this problem does not arise with GKAT. The reason is that we do not have the unrestricted union operator $+$, but only the guarded union operator $+_b$. We will therefore stick with ordinary sequents, rather than hypersequents. We will show in Section 6.5 that, even though it uses sequents rather than hypersequents, SGKAT^∞ is regularly complete.

We end this section with a definition and a lemma which will be useful in the proofs of both soundness and completeness.

6.2.12. Definition. A list Γ of expressions is said to be *exposed* if it is either empty or begins with a primitive program.

Recall that the sets of primitive tests and primitive programs are disjoint. Hence an exposed list Γ cannot start with a test.

6.2.13. Remark. Coming from the modal μ-calculus, the reader might be surprised by the notion of exposure, as an antonym of *guardedness*. Because the *guards* in the modal μ-calculus are modalities, one might be tempted to think the primitive programs in GKAT as guards as well. Rather, in the context of GKAT it is typical to refer to the tests as guards.

6.2.14. Lemma. *Let Γ and Δ be exposed lists of expressions. Then:*

(i) $\alpha x \in [\![\Gamma]\!] \Leftrightarrow \beta x \in [\![\Gamma]\!]$ *for all* $\alpha, \beta \in \mathsf{At}$

(ii) $\Gamma \Rightarrow_{\mathsf{At}} \Delta$ *is valid if and only if* $\Gamma \Rightarrow_A \Delta$ *is valid for some* $A \neq \emptyset$.

Proof:
For item (i), we make a case distinction on whether $\Gamma = \epsilon$ or $\Gamma = p, \Theta$ for some list Θ. If $\Gamma = \epsilon$, the result follows immediately from the fact that $[\![\epsilon]\!] = \mathsf{At}$. If $\Gamma = p, \Theta$, we have

$$[\![\Gamma]\!] = [\![p]\!] \diamond [\![\Theta]\!] = \{\gamma p \delta y : \gamma \in \mathsf{At}, \delta y \in [\![\Theta]\!]\},$$

which also suffices.

For item (ii), the only non-trivial implication is the one from right to left. So suppose $\Gamma \Rightarrow_A \Delta$ for some $A \neq \emptyset$. Let $\alpha \in \mathsf{At}$ and let $\beta \in A$ be arbitrary. We find:

$$\begin{aligned} \alpha x \in [\![\Gamma]\!] &\Rightarrow \beta x \in [\![\Gamma]\!] && \text{(item (i))} \\ &\Rightarrow \beta x \in [\![\Delta]\!] && (\beta \in A, \text{hypothesis}) \\ &\Rightarrow \alpha x \in [\![\Delta]\!], && \text{(item (i))} \end{aligned}$$

as required. □

6.3 Soundness

In this section we prove that SGKAT^∞ is sound. We will first prove that *well-founded* SGKAT^∞-proofs are sound. For the sake of readability we will write $[\![\Gamma]\!]$ to abbreviate $[\![\gamma_1 \cdot \ldots \cdot \gamma_n]\!]$ for some list Γ of expressions $\gamma_1, \ldots, \gamma_n$.

6.3.1. Lemma. *Let A be a set of atoms, let b be a test, and let Θ be a list of expressions. We have:*

1. $A \restriction b = A \diamond [\![b]\!]$;
2. $[\![e +_b f, \Theta]\!] = ([\![b]\!] \diamond [\![e, \Theta]\!]) \cup ([\![\bar{b}]\!] \diamond [\![f, \Theta]\!])$;
3. $[\![e^{(b)}, \Theta]\!] = ([\![b]\!] \diamond [\![e, e^{(b)}, \Theta]\!]) \cup ([\![\bar{b}]\!] \diamond [\![\Theta]\!])$.

Proof:
Each item is shown by simply unfolding the definitions. We will use the fact $\diamond$ distributes over $\cup$. Note that $\cup$ is not the same as *guarded* union, over which $\diamond$ is merely right-distributive. First, we have $A \restriction b = \{\alpha \in A : \alpha \leq b\} = A \diamond [\![b]\!]$.

For the second item, we calculate

$$\begin{aligned}
[\![e +_b f, \Theta]\!] &= [\![e +_b f]\!] \diamond [\![\Theta]\!] && \text{(sequent interpretation)}\\
&= (([\![b]\!] \diamond [\![e]\!]) \cup (\overline{[\![b]\!]} \diamond [\![f]\!])) \diamond [\![\Theta]\!] && \text{(interpretation of } +_b)\\
&= ([\![b]\!] \diamond [\![e]\!] \diamond [\![\Theta]\!]) \cup (\overline{[\![b]\!]} \diamond [\![f]\!] \diamond [\![\Theta]\!]) && (\diamond \text{ distributes over } \cup)\\
&= ([\![b]\!] \diamond [\![e]\!] \diamond [\![\Theta]\!]) \cup ([\![\bar{b}]\!] \diamond [\![f]\!] \diamond [\![\Theta]\!]) && (\overline{[\![b]\!]} = [\![\bar{b}]\!])\\
&= ([\![b]\!] \diamond [\![e, \Theta]\!]) \cup ([\![\bar{b}]\!] \diamond [\![f, \Theta]\!]). && \text{(sequent interpretation)}
\end{aligned}$$

Finally, for the third item, we have

$$\begin{aligned}
[\![e^{(b)}, \Theta]\!] &= [\![e^{(b)}]\!] \diamond [\![\Theta]\!] && \text{(sequent int.)}\\
&= \bigcup_{n \geq 0} ([\![b]\!] \diamond [\![e]\!])^n \diamond \overline{[\![b]\!]} \diamond [\![\Theta]\!] && (\text{int. } -^{(b)})\\
&= (\bigcup_{n \geq 1} ([\![b]\!] \diamond [\![e]\!])^n \cup \mathsf{At}) \diamond \overline{[\![b]\!]} \diamond [\![\Theta]\!] && (\text{split } \textstyle\bigcup)\\
&= (\bigcup_{n \geq 1} ([\![b]\!] \diamond [\![e]\!])^n \diamond \overline{[\![b]\!]} \diamond [\![\Theta]\!]) \cup (\mathsf{At} \diamond \overline{[\![b]\!]} \diamond [\![\Theta]\!]) && (\diamond \text{ dist. } \cup)\\
&= (\bigcup_{n \geq 1} ([\![b]\!] \diamond [\![e]\!])^n \diamond [\![\bar{b}]\!] \diamond [\![\Theta]\!]) \cup ([\![\bar{b}]\!] \diamond [\![\Theta]\!]) && (\overline{[\![b]\!]} = [\![\bar{b}]\!])\\
&= (\bigcup_{n \geq 0} [\![b]\!] \diamond [\![e]\!] \diamond ([\![b]\!] \diamond [\![e]\!])^n \diamond [\![\bar{b}]\!] \diamond [\![\Theta]\!]) \cup ([\![\bar{b}]\!] \diamond [\![\Theta]\!]) && \text{(Lem. 6.1.5)}\\
&= ([\![b]\!] \diamond [\![e]\!] \diamond \bigcup_{n \geq 0} ([\![b]\!] \diamond [\![e]\!])^n \diamond [\![\bar{b}]\!] \diamond [\![\Theta]\!]) \cup ([\![\bar{b}]\!] \diamond [\![\Theta]\!]) && (\diamond \text{ dist. } \textstyle\bigcup)\\
&= ([\![b]\!] \diamond [\![e]\!] \diamond [\![e^{(b)}]\!] \diamond [\![\Theta]\!]) \cup ([\![\bar{b}]\!] \diamond [\![\Theta]\!]) && (\text{int. } -^{(b)})\\
&= ([\![b]\!] \diamond [\![e, e^{(b)}, \Theta]\!]) \cup ([\![\bar{b}]\!] \diamond [\![\Theta]\!]), && \text{(sequent int.)}
\end{aligned}$$

as required. □

We prioritise the rules of SGKAT in order of occurrence in Figure 6.3, reading left-to-right, top-to-bottom, *i.e.* in normal English reading order. For instance, each left logical rule is of higher priority than each right logical rule, which in turn is of higher priority than each axiom or modal rule.

We will use the following property of the system SGKAT, which follows from direct inspection of the rules and the fact that sequents are lists.

6.3.2. LEMMA. *Let* $\Gamma \Rightarrow_A \Delta$ *be a sequent, and let* r *be any rule of* SGKAT. *Then there is at most one rule instance of* r *with conclusion* $\Gamma \Rightarrow_A \Delta$.

Therefore, the following is well-defined.

6.3.3. DEFINITION. A rule instance of r with conclusion $\Gamma \Rightarrow_A \Delta$ is said to *have priority* if any other rule instance, say of r', with conclusion $\Gamma \Rightarrow_A \Delta$ is of lower priority (that is, the rule r' appears after r in Figure 6.3).

Recall that a rule is *sound* if the validity of all its premisses implies the validity of its conclusion. Conversely, a rule is *invertible* if the validity of its conclusion implies the validity of all of its premisses.

The above notion of priority will be used in the completeness proof of Section 6.5 to guide a proof-search procedure. Conveniently, the following proposition entails that every rule instance which has priority is invertible, allowing this proof-search procedure to be deterministic.

6.3.4. PROPOSITION. *Every rule of* SGKAT *is sound. Moreover, every rule is invertible except for* k *and* k_0*, which are invertible whenever they have priority.*

Proof:
We will cover the rules of SGKAT one-by-one.

(b-l) This is immediate by Lemma 6.3.1.1.

(b-r) We have:

$$\begin{aligned} A \diamond [\![\Gamma]\!] \subseteq [\![b, \Delta]\!] &\Leftrightarrow A \diamond [\![\Gamma]\!] \subseteq [\![b]\!] \diamond [\![\Delta]\!] && \text{(sequent int.)} \\ &\Leftrightarrow A \restriction b \diamond [\![\Gamma]\!] \subseteq [\![b]\!] \diamond [\![\Delta]\!] && \text{(by (†))} \\ &\Leftrightarrow A \restriction b \diamond [\![\Gamma]\!] \subseteq [\![\Delta]\!] && (A \restriction b \subseteq [\![b]\!]) \\ &\Leftrightarrow A \diamond [\![\Gamma]\!] \subseteq [\![\Delta]\!] && \text{(by (†))} \end{aligned}$$

($\cdot$-l) Immediate, since $A \diamond [\![e \cdot f, \Gamma]\!] = A \diamond [\![e, f, \Gamma]\!]$.

($\cdot$-r) Likewise, but by $[\![e \cdot f, \Delta]\!] = [\![e, f, \Delta]\!]$.

($+_b$-l) This follows directly from the fact that

$$\begin{aligned} A \diamond [\![e +_b f, \Gamma]\!] &= A \diamond [\![e +_b f]\!] \diamond [\![\Gamma]\!] && \text{(sequent int.)} \\ &= A \diamond (([\![b]\!] \diamond [\![e, \Gamma]\!]) \cup ([\![\bar{b}]\!] \diamond [\![f, \Gamma]\!])) && \text{(Lem. 6.3.1.2)} \\ &= (A \diamond [\![b]\!] \diamond [\![e, \Gamma]\!]) \cup (A \diamond [\![\bar{b}]\!] \diamond [\![f, \Gamma]\!]) && \text{(distrib.)} \\ &= (A \upharpoonright b \diamond [\![e, \Gamma]\!]) \cup (A \upharpoonright \bar{b} \diamond [\![f, \Gamma]\!]) && \text{(Lem. 6.3.1.1)} \end{aligned}$$

($+_b$-r) We find

$$\begin{aligned} & A \diamond [\![\Gamma]\!] \subseteq [\![e +_b f]\!] \diamond [\![\Delta]\!] \\ \Leftrightarrow\ & A \diamond [\![\Gamma]\!] \subseteq ([\![b]\!] \diamond [\![e, \Delta]\!]) \cup ([\![\bar{b}]\!] \diamond [\![f, \Delta]\!]) \\ \Leftrightarrow\ & A \upharpoonright b \diamond [\![\Gamma]\!] \subseteq [\![e, \Delta]\!] \text{ or } A \upharpoonright \bar{b} \subseteq [\![f, \Delta]\!], \end{aligned}$$

where the first equivalence holds due to Lemma 6.3.1.2, and the second due to $A \diamond [\![\Gamma]\!] = ([\![b]\!] \diamond A \diamond [\![\Gamma]\!]) \cup ([\![\bar{b}]\!] \diamond A \diamond [\![\Gamma]\!])$ and Lemma 6.3.1.1.

((b)-l) This follows directly from the fact that

$$\begin{aligned} A \diamond [\![e^{(b)}, \Gamma]\!] &= A \diamond [\![e^{(b)}]\!] \diamond [\![\Gamma]\!] && \text{(sequent int.)} \\ &= A \diamond (([\![b]\!] \diamond [\![e, e^{(b)}, \Gamma]\!]) \cup ([\![\bar{b}]\!] \diamond [\![f, \Gamma]\!])) && \text{(Lem. 6.3.1.3)} \\ &= (A \diamond [\![b]\!] \diamond [\![e, e^{(b)}, \Gamma]\!]) \cup (A \diamond [\![\bar{b}]\!] \diamond [\![f, \Gamma]\!]) && \text{(distrib.)} \\ &= (A \upharpoonright b \diamond [\![e, e^{(b)}, \Gamma]\!]) \cup (A \upharpoonright \bar{b} \diamond [\![f, \Gamma]\!]) && \text{(Lem. 6.3.1.1)} \end{aligned}$$

((b)-r) We find

$$\begin{aligned} & A \diamond [\![\Gamma]\!] \subseteq [\![e^{(b)}, \Delta]\!] \\ \Leftrightarrow\ & A \diamond [\![\Gamma]\!] \subseteq ([\![b]\!] \diamond [\![e, e^{(b)}, \Delta]\!]) \cup ([\![\bar{b}]\!] \diamond [\![\Delta]\!]) \\ \Leftrightarrow\ & A \upharpoonright b \diamond [\![\Gamma]\!] \subseteq [\![b]\!] \diamond [\![e, e^{(b)}, \Delta]\!] \text{ and } A \upharpoonright \bar{b} \subseteq [\![\bar{b}]\!] \diamond [\![\Delta]\!], \end{aligned}$$

where the first equivalence holds due to Lemma 6.3.1.3, and the second due to $A \diamond [\![\Gamma]\!] = ([\![b]\!] \diamond A \diamond [\![\Gamma]\!]) \cup ([\![\bar{b}]\!] \diamond A \diamond [\![\Gamma]\!])$ and Lemma 6.3.1.1.

(id) This follows from $A \diamond [\![1]\!] = A \diamond \mathsf{At} = A \subseteq \mathsf{At} = [\![1]\!]$.

($\bot$) We have $\emptyset \diamond [\![\Gamma]\!] = \emptyset \subseteq [\![\Delta]\!]$.

(k) Suppose first that some application of k does *not* have priority. The only rule of higher priority than k which can have a conclusion of the form $p, \Gamma \Rightarrow_A p, \Delta$ is $\bot$, whence we must have $A = \emptyset$. As shown in the previous case, this conclusion must be valid. Hence under this restriction the rule application is vacuously sound. It is, however, not invertible, as the following rule instance demonstrates

$$\mathsf{k}\,\frac{1 \Rightarrow_{\mathsf{At}} 0}{p, 1 \Rightarrow_{\emptyset} p, 0}$$

Next, suppose that some application of k does have priority. This means that the set A of atoms in the conclusion $p, \Gamma \Rightarrow_A p, \Delta$ is *not* empty. We will show that under this restriction the rule is both sound and invertible. Let $\alpha \in A$. We have

$$
\begin{aligned}
A \diamond [\![p, \Gamma]\!] \subseteq [\![p, \Delta]\!] &\Leftrightarrow A \diamond [\![p]\!] \diamond [\![\Gamma]\!] \subseteq [\![p]\!] \diamond [\![\Delta]\!] && \text{(seq. int.)} \\
&\Leftrightarrow \alpha \diamond [\![p]\!] \diamond [\![\Gamma]\!] \subseteq [\![p]\!] \diamond [\![\Delta]\!] && (\alpha \in A, \text{Lem. } 6.2.14) \\
&\Leftrightarrow [\![p]\!] \diamond [\![\Gamma]\!] \subseteq [\![p]\!] \diamond [\![\Delta]\!] && \text{(Lem. } 6.2.14) \\
&\Leftrightarrow [\![\Gamma]\!] \subseteq [\![\Delta]\!], && \text{(Lem. } 6.1.6)
\end{aligned}
$$

as required.

(k_0) For the final rule k_0, we will first show the soundness of all instances, and then the invertibility of those instances which have priority. For soundness, suppose that the premiss is valid. Since

$$[\![\Gamma]\!] = \mathsf{At} \diamond [\![\Gamma]\!] \subseteq [\![0]\!] = \emptyset,$$

it follows that $[\![\Gamma]\!] = \emptyset$. Hence

$$A \diamond [\![p, \Gamma]\!] = A \diamond [\![p]\!] \diamond [\![\Gamma]\!] = A \diamond [\![p]\!] \diamond \emptyset = \emptyset \subseteq [\![\Delta]\!],$$

as required.

For invertibility, suppose that some instance of k_0 has priority. Then the conclusion $p, \Gamma \Rightarrow_A \Delta$ cannot be the conclusion of any other rule application.

Suppose that $p, \Gamma \Rightarrow_A \Delta$ is valid. We wish to show that $\Gamma \Rightarrow_{\mathsf{At}} 0$ is valid, or, in other words, that $[\![\Gamma]\!] = \emptyset$.

First note that, as in the previous case, from the assumption that our instance of k_0 has priority, it follows that $A \neq \emptyset$.

We now make a case distinction on the shape of Δ. Suppose first that $\Delta = \epsilon$. Then

$$A \diamond [\![p, \Gamma]\!] \subseteq [\![\Delta]\!] = [\![\epsilon]\!] = \mathsf{At}.$$

As $A \diamond [\![p, \Gamma]\!] = \{\alpha p \beta x : \alpha \in A \text{ and } \beta x \in [\![\Gamma]\!]\}$, we must have $[\![\Gamma]\!] = \emptyset$.

Next, suppose that Δ has a leftmost expression e. By the assumption that the rule instance has priority, we know that e is not of the form $e_0 \cdot e_1$, $e_0 +_b e_1$, or $e^{(b)}$, for otherwise a right logical rule could be applied. Hence, the expression e must either be a test or a primitive program.

If e is a test, say b, we know that $A \restriction b \neq A$, for otherwise b-r could be applied. Recall that it suffices to show that $[\![\Gamma]\!] = \emptyset$. So suppose, towards a contradiction, that there is some $\beta x \in [\![\Gamma]\!]$. Let $\alpha \in A$ such that $\alpha \not\leq b$. Then $\alpha p \beta x \in [\![p, \Gamma]\!] \subseteq [\![\Delta]\!]$. But this contradicts the fact that $[\![\Delta]\!] \subseteq \{\alpha y : \alpha \leq b\}$.

Finally, suppose that e is a primitive program, say q. Write $\Delta = q, \Theta$. First note that assumption that the rule instance has priority implies $p \neq q$, for otherwise the rule k could be applied. We have:

$$A \diamond [\![p, \Gamma]\!] \subseteq [\![\Delta]\!] = \{\alpha q \beta x : \beta x \in [\![\Theta]\!]\},$$

As $A \diamond [\![p, \Gamma]\!] = \{\alpha p \beta x : \alpha \in A \text{ and } \beta x \in [\![\Gamma]\!]\}$ and $p \neq q$, we again find that $[\![\Gamma]\!] = \emptyset$.

This finishes the proof. □

Our extension of the soundness result to also include non-well-founded proofs closely follows the treatment in [34]. We first recursively define the following syntactic abbreviations:

$$e^{(b)^0} := \bar{b}, \qquad\qquad e^{(b)^{n+1}} := bee^{(b)^n}.$$

6.3.5. Lemma. *For every $n \in \mathbb{N}$: if we have $\mathsf{SGKAT}^\infty \vdash e^{(b)}, \Gamma \Rightarrow_A \Delta$, then we also have $\mathsf{SGKAT}^\infty \vdash e^{(b)^n}, \Gamma \Rightarrow_A \Delta$.*

Proof:
We assume that $A \neq \emptyset$, for otherwise the lemma is trivial. Let π be the assumed SGKAT^∞-proof of $e^{(b)}, \Gamma \Rightarrow_A \Delta$. Note that, since all succedents referred to in the lemma are equal to Δ, it suffices to prove the lemma under the assumption that the last rule applied in π is *not* a right logical rule. Hence, we may assume that the last rule applied in π is (b)-l, for that is the only remaining rule with a sequent of this shape as conclusion. This means that π is of the form:

$$\dfrac{\overset{\pi_1}{e, e^{(b)}, \Gamma \Rightarrow_{A \restriction b} \Delta} \qquad \overset{\pi_2}{\Gamma \Rightarrow_{A \restriction \bar{b}} \Delta}}{e^{(b)}, \Gamma \Rightarrow_A \Delta}\ (b)\text{-}l$$

We show the lemma by induction on n. For the induction base, we take the following proof:

$$\dfrac{\overset{\pi_2}{\Gamma \Rightarrow_{A \restriction \bar{b}} \Delta}}{e^{(b)^0}, \Gamma \Rightarrow_A \Delta}\ \bar{b}\text{-}l$$

For the inductive step $n+1$, we construct from π_1 a proof τ of $e, e^{(b)^n}, \Gamma \Rightarrow_{A \restriction b} \Delta$. To that end, we first replace in π_1 every occurrence of $e^{(b)}, \Gamma$ as a final segment of the antecedent by $e^{(b)^n}, \Gamma$ and cut off all branches at sequents of the form

$e^{(b)^n}, \Gamma \Rightarrow_B \Theta$. This may be depicted as follows, where to the left of the arrow $\rightsquigarrow$ we have a branch of π_1, and to right the resulting branch of τ.

$$\dfrac{\dfrac{\vdots}{e^{(b)}, \Gamma \Rightarrow_B \Theta}}{\dfrac{\vdots}{e, e^{(b)}, \Gamma \Rightarrow_{A \restriction b} \Delta}} \quad \rightsquigarrow \quad \dfrac{\overline{e^{(b)^n}, \Gamma \Rightarrow_B \Theta}}{\dfrac{\vdots}{e, e^{(b)^n}, \Gamma \Rightarrow_{A \restriction b} \Delta}}$$

Note that every remaining infinite branch in the resulting derivation τ satisfies the fairness condition. Therefore, to turn τ into a proper SGKAT^∞-proof, we only need to close each open leaf, which by construction is of the form $e^{(b)^n}, \Gamma \Rightarrow_B \Delta$. Note that π_1 must contain a proof of $e^{(b)}, \Gamma \Rightarrow_B \Delta$, whence by the induction hypothesis the sequent $e^{(b)^n}, \Gamma \Rightarrow_B \Delta$ is provable. We can thus close the leaf by simply appending the witnessing proof.

Letting τ be the resulting proof, we finish the induction step by taking:

$$\dfrac{\begin{array}{c}\tau \\ e, e^{(b)^n}, \Gamma \Rightarrow_{A \restriction b} \Delta\end{array}}{e^{(b)^{n+1}}, \Gamma \Rightarrow_A \Delta} \ \textit{b-l}$$

which gives us the required SGKAT^∞-proof. □

We let the *while-height* $\mathsf{wh}(e)$ be the maximal nesting of while loops in a given expression e. Formally,

- $\mathsf{wh}(b) = \mathsf{wh}(p) = 0$;
- $\mathsf{wh}(e \cdot f) = \mathsf{wh}(e +_b f) = \max\{\mathsf{wh}(e), \mathsf{wh}(f)\}$;
- $\mathsf{wh}(e^{(b)}) = \mathsf{wh}(e) + 1$.

Given a list Γ, the *weighted while-height* $\mathsf{wwh}(\Gamma)$ of Γ is defined to be the multiset $[\mathsf{wh}(e) : e \in \Gamma]$. We order such multisets using the Dershowitz–Manna ordering:

$$\begin{aligned} N < M \text{ iff } N \neq M &\text{ and for any } n \text{ with } N(n) > M(n), \\ &\text{there is an } n' > n \text{ such that } N(n') < M(n'). \end{aligned}$$

Given any partial order $(S, <_S)$, the Dershowitz-Manna ordering can be used to give a well-founded partial order on the set of finite multisets of S. Since expressions are, in fact, *linearly* ordered by while-height, the Dershowitz-Manna ordering admits a more simple description in our case. We say that $N < M$ if and only if $N \neq M$ and for the greatest n such that $N(n) \neq M(n)$, it holds that $N(n) < M(n)$.

Note that in any SGKAT-derivation the weighted while-height of the antecedent does not increase when reading bottom-up. Moreover, we have:

6.3.6. Lemma. $\mathsf{wwh}(e^{(b)^n}, \Gamma) < \mathsf{wwh}(e^{(b)}, \Gamma)$ *for every* $n \in \mathbb{N}$.

Proof:
Let $k := \mathsf{wh}(e^{(b)})$. Note that the maximum while-height in $e^{(b)^n}$ is that of e. Hence, we have $\mathsf{wwh}(e^{(b)^n})(k) = 0 < 1 = \mathsf{wwh}(e^{(b)})(k)$. Therefore:

$$\begin{aligned}\mathsf{wwh}(e^{(b)^n}, \Gamma)(k) &= \mathsf{wwh}(e^{(b)^n})(k) + \mathsf{wwh}(\Gamma)(k) \\ &< \mathsf{wwh}(e^{(b)})(k) + \mathsf{wwh}(\Gamma)(k) = \mathsf{wwh}(e^{(b)}, \Gamma)(k).\end{aligned}$$

Hence $\mathsf{wwh}(e^{(b)^n}, \Gamma) \neq \mathsf{wwh}(e^{(b)}, \Gamma)$. Now suppose that for some $l \in \mathbb{N}$ we have $\mathsf{wwh}(e^{(b)^n}, \Gamma)(l) > \mathsf{wwh}(e^{(b)}, \Gamma)(l)$. We leave it to the reader to verify that in this case we must have $l < k$. As $\mathsf{wwh}(e^{(b)^n}, \Gamma)(k) < \mathsf{wwh}(e^{(b)}, \Gamma)(k)$, we find $\mathsf{wwh}(e^{(b)^n}, \Gamma) < \mathsf{wwh}(e^{(b)}, \Gamma)$. □

We are now ready to prove the soundness theorem.

6.3.7. Theorem. *If* $\mathsf{SGKAT}^\infty \vdash \Gamma \Rightarrow_A \Delta$, *then* $A \diamond [\![\Gamma]\!] \subseteq [\![\Delta]\!]$.

Proof:
We prove this by induction on $\mathsf{wwh}(\Gamma)$. Given a proof π of $\Gamma \Rightarrow_A \Delta$, let $\mathcal{B}$ contain for each infinite branch of π the node of least depth to which a rule (b)-l is applied. Note that $\mathcal{B}$ must be finite, for otherwise, by Kőnig's Lemma, the proof π cut off along $\mathcal{B}$ would have an infinite branch that does not satisfy the fairness condition.

Note that Proposition 6.3.4 entails that of every finite derivation with valid leaves the conclusion is valid. Hence, it suffices to show that each of the nodes in $\mathcal{B}$ is valid. To that end, consider an arbitrary such node labelled $e^{(b)}, \Gamma' \Rightarrow_{A'} \Delta'$ and the subproof π' it generates. By Lemma 6.3.5, we have that $e^{(b)^n}, \Gamma' \Rightarrow_{A'} \Delta'$ is provable for every n. Lemma 6.3.6 gives $\mathsf{wwh}(e^{(b)^n}, \Gamma') < \mathsf{wwh}(e^{(b)}, \Gamma') \leq \mathsf{wwh}(\Gamma)$, and thus we may apply the induction hypothesis to obtain

$$A' \diamond [\![e^{(b)^n}]\!] \diamond [\![\Gamma]\!] \subseteq [\![\Delta]\!]$$

for every $n \in \mathbb{N}$. Then by

$$\bigcup_n (A' \diamond [\![e^{(b)^n}]\!] \diamond [\![\Gamma]\!]) = A' \diamond \bigcup_n ([\![e^{(b)^n}]\!]) \diamond [\![\Gamma]\!] = A' \diamond [\![e]\!]^{[\![b]\!]} \diamond [\![\Gamma]\!],$$

we obtain that $e^{(b)}, \Gamma' \Rightarrow_{A'} \Delta'$ is valid, as required. □

6.4 Frugality

Before we show that SGKAT^∞ is not only sound, but also complete, we will first show that every SGKAT^∞-proof is *frugal*. Recall that this notion was defined in the Intermezzo, and that a frugal proof is one in which only finitely many distinct sequents appear.

Our treatment is again similar to that in [34] for Kleene Algebra, but presented in a slightly different way, namely using the standard notion of a *syntax tree.*

6.4.1. DEFINITION. The *syntax tree* (T_e, l_e) of an expression e is a well-founded, labelled and ordered tree, defined by the following inducton on e.

- If e is a test or primitive program, its syntax tree only has a root node ϵ, with label $l_e(\epsilon) := e$.
- If $e = f_1 \circ f_2$ where $\circ = \cdot$ or $\circ = +_b$, its syntax tree again has a root node ϵ with label $l_e(\epsilon) = e$, and with two outgoing edges. The first edge connects ϵ to (T_{f_1}, l_{f_1}), the second edge connects it to (T_{f_2}, l_{f_2}).
- If $e = f^{(b)}$, its syntax tree again has a root node ϵ with label $l_e(\epsilon) = e$, but now with just one outgoing edge. This edge connects ϵ to (T_f, l_f).

Formally there is a sibling ordering implicit in the syntax tree (T_e, l_e), but for the sake of readability we omit this in our notation. There are several ways to extend this sibling ordering to a total order on the set of all nodes of T_e. The correct ordering for our purposes is the following.

6.4.2. DEFINITION. Let (T_e, l_e) be the syntax tree of e, and let u, v be nodes in of T_e. We let $u \prec_e v$ if u comes before v in the depth-first traversal of T_e.

6.4.3. EXAMPLE. The syntax tree of $a \cdot (p +_b q)^{(a)}$ is given by

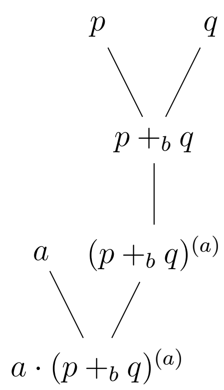

and $a \cdot (p +_b q)^{(a)} \prec a \prec (p +_b q)^{(a)} \prec p +_b q \prec p \prec q$ is the order of depth-first traversal.

Observe that $\prec_e$ is indeed a total order on the nodes of T_e. It will be convenient to have the following abstract notion of sequents.

6.4.4. DEFINITION. Let e and f be GKAT-expressions. An (e, f)-sequent is a triple (Γ, A, Δ), where Γ is a list of nodes in T_e, and Δ is a list of nodes in T_f, and A is a set of atoms.

Let $(u_1, \ldots, u_n, A, v_1, \ldots, v_m)$ be an (e, f)-sequent. Its *realisation* is the sequent $l_e(u_1), \ldots, l_e(u_n) \Rightarrow_A l_f(v_1), \ldots, l_f(v_m)$. It is *strictly increasing* if $u_1, \ldots u_n$ is strictly increasing under $\prec_e$, and $v_1, \ldots, v_m$ is strictly increasing under $\prec_f$. .

6.4.5. REMARK. In the above definition n and m are allowed to be 0, in which case they realise the empty list ϵ. We regard the empty list as strictly increasing.

The following lemma embodies the key idea for establishing the frugality of SGKAT^∞-proofs. By ϵ_e we will denote the root of the syntax tree (T_e, l_e). Moreover, recall that 0 is an expression, whence the notion of $(e, 0)$-sequent is well-defined.

6.4.6. LEMMA. *Let π be an SGKAT^∞-derivation of a sequent of the form $e \Rightarrow_A f$. Then every node of π is either the realisation of some strictly increasing (e, f)-sequent, or of some strictly increasing $(e, 0)$-sequent.*

Proof:
We will prove this by bottom-up induction on π. For the base case, note that the root of π is the realisation of the strictly increasing (e, f)-sequent $(\epsilon_e, A, \epsilon_f)$.

For the inductive step, suppose that π contains a rule instance of the rule r such that the thesis holds for its conclusion $\Gamma_0 \Rightarrow_{A_0} \Delta_0$. We make a case distinction on r. We will only treat two illustrative cases, leaving the other cases to the reader.

Suppose that $\mathsf{r} = (b)$-l. Then the rule instance is of the form

$$\frac{g, g^{(b)}, \Gamma \Rightarrow_{A_0 \restriction b} \Delta \qquad \Gamma \Rightarrow_{A_0 \restriction \bar{b}} \Delta}{g^{(b)}, \Gamma \Rightarrow_{A_0} \Delta}\ (b)\text{-}l$$

Let $(u_1, \ldots, u_n, A_0, v_1, \ldots, v_m)$ be the (e, f)-sequent or $(e, 0)$-sequent realising $g^{(b)}, \Gamma \Rightarrow_{A_0} \Delta$. Then $l_e(u_1) = g^{(b)}$ and thus the node u_1 of T_e has a child u_0 such that $l(u_0) = g$. Hence, the premisses of this rule instance are realised by, respectively,

$$(u_0, u_1, \ldots, u_n, A_0 \restriction b, v_1, \ldots, v_m), \text{ and } (u_2, \ldots, u_n, A_0 \restriction b, v_1, \ldots, v_m),$$

which are clearly both strictly increasing (e, f)-sequents or $(e, 0)$-sequents.

Now suppose that $\mathsf{r} = \mathsf{k}_0$. Then the rule instance is of the form

$$\frac{\Gamma \Rightarrow_{\mathsf{At}} 0}{p, \Gamma \Rightarrow_{A_0} \Delta}\ \mathsf{k}_0$$

Let $(u_1, \ldots, u_n, A_0, v_1, \ldots, v_m)$ be an (e, f)-sequent or $(e, 0)$-sequent which realises $p, \Gamma \Rightarrow_{A_0} \Delta$. Then $(u_2, \ldots, u_n, \mathsf{At}, \epsilon_0)$ is a strictly increasing $(e, 0)$-sequent which realises $\Gamma \Rightarrow_{\mathsf{At}} 0,$. □

As a corollary we obtain a bound on the number of sequents occurring in a proof.

6.4.7. Corollary. *Any* SGKAT^∞*-derivation is frugal.*

Proof:
Let π be an SGKAT^∞-derivation. Without loss of generality we may assume that the conclusion of π is of the form $e \Rightarrow_A f$ (for otherwise we can simply append a series of applications of $\cdot$-l and $\cdot$-r to the root of π). Hence, by the previous proposition every sequent occurring in π is the realisation of an (e, f)-sequent or an $(e, 0)$ sequent. This means that there are at most n distinct antecedents in π, where n is the number of nodes in the syntax tree of e. Moreover, there are at most $m + 1$ distinct succedents, where m number of nodes in the syntax tree of f. Since there are $2^{|(\mathsf{At})|}$ different sets of atoms, it follows that there are at most $2^{|(\mathsf{At})|} \cdot n(m + 1)$ distinct sequents in π. □

Finally, we find by Corollary I.2.23 that SGKAT^∞ is complete, then it is regularly complete.

6.4.8. Corollary. *If* $\Gamma \Rightarrow_A \Delta$ *has an* SGKAT^∞*-proof, then it also has a regular* SGKAT^∞*-proof.*

6.5 Completeness

In this section we prove the completeness of SGKAT^∞. Our treatment is again inspired by that in [34] for ordinary Kleene Algebra, but requires several modifications to treat the tests present in GKAT. We first prove some auxiliary lemmas.

The following lemma is, of course, a minimal requirement for completeness.

6.5.1. Lemma. *Any valid sequent is the conclusion of some rule application.*

Proof:
We prove this lemma by contraposition. So suppose $\Gamma \Rightarrow_A \Delta$ is *not* the conclusion of *any* rule application. We make a few observations:

- Both Γ and Δ are exposed, for otherwise $\Gamma \Rightarrow_A \Delta$ would be the conclusion of an application of a left, respectively right, logical rule.
- A is non-empty, for otherwise $\Gamma \Rightarrow_A \Delta$ would be the conclusion of an application of $\perp$.
- The leftmost expression of Γ is not a primitive program, for otherwise our sequent $\Gamma \Rightarrow_A \Delta$ would be the conclusion of an application of k_0.

- The leftmost expression of Δ *is* a primitive program, for otherwise, by the previous items, the sequent $\Gamma \Rightarrow_A \Delta$ would be the conclusion of an application of id.

Hence $\Gamma \Rightarrow_A \Delta$ is of the form $\epsilon \Rightarrow_A p, \Theta$. Let $\alpha \in A$. Then $\alpha \in A \diamond [\![\epsilon]\!]$. However, since α is not of the form $\beta p \gamma y$, we have $\alpha \notin [\![p, \Theta]\!]$. This shows that $\Gamma \Rightarrow_A \Delta$ is not valid, as required. □

Note the occurrence of two, possibly distinct, sets A and B of atoms in the formulation of the following lemma.

6.5.2. Lemma. *Let π be a derivation using only right logical rules and containing a branch of the form:*

$$\frac{\begin{array}{c}\Gamma \Rightarrow_B e^{(b)}, \Delta \\ \vdots\end{array}}{\Gamma \Rightarrow_A e^{(b)}, \Delta}\ (b)\text{-}r \qquad (*)$$

such that:

1. $\Gamma \Rightarrow_A e^{(b)}, \Delta$ *is valid, and*
2. *Every succedent on the branch has* $e^{(b)}, \Delta$ *as a final segment.*

Then $\Gamma \Rightarrow_B 0$ *is valid.*

Proof:
We claim that $e^{(b)} \Rightarrow_B 0$ is provable. We will show this by exploiting the symmetry of the left and right logical rules of SGKAT (cf. Remark 6.2.10). Since on the branch (*) every rule is a right logical rule, and $e^{(b)}, \Delta$ is preserved throughout, we can construct a derivation π' of $e^{(b)} \Rightarrow_B 0$ from π by applying the analogous left logical rules to $e^{(b)}$. Note that the set of atoms B precisely determines the branch (*), in the sense that for every leaf $\Gamma \Rightarrow_C \Theta$ of π it holds that $C \cap B = \emptyset$. Hence, as the root of π' is $e^{(b)} \Rightarrow_B 0$, every branch of π' except for the one corresponding to (*) can be closed directly by an application of $\bot$. The branch corresponding to (*) is of the form

$$\frac{\begin{array}{c}e^{(b)} \Rightarrow_B 0 \\ \vdots\end{array}}{e^{(b)} \Rightarrow_B 0}\ (b)\text{-}l \qquad (*)$$

and can thus be closed by a back edge. The resulting finite tree with back edges clearly represents an SGKAT^∞-proof.

Now by soundness, we have $B \diamond [\![e^{(b)}]\!] = \emptyset$. Moreover, by the invertibility of the right logical rules and hypothesis (1), we get

$$B \diamond [\![\Gamma]\!] \subseteq B \diamond [\![e^{(b)}]\!] \diamond [\![\Delta]\!] = \emptyset,$$

as required. □

6.5.3. LEMMA. *Let* $(\Gamma_n \Rightarrow_{A_n} \Delta_n)_{n\in\omega}$ *be an infinite branch of some* SGKAT^∞*-derivation on which the rule* (b)*-*r *is applied infinitely often. Then there are* n, m *with* $n < m$ *such that the following hold:*

(i) the sequents $\Gamma_n \Rightarrow_{A_n} \Delta_n$ *and* $\Gamma_m \Rightarrow_{A_m} \Delta_m$ *are equal;*

(ii) the sequent $\Gamma_n \Rightarrow_{A_n} \Delta_n$ *is the conclusion of an application of* (b)*-*r *in* π*;*

(iii) for every $i \in [n, m)$ *it holds that* Δ_n *is a final segment of* Δ_i*.*

Proof:
First note that k_0 is not applied on this branch, because if it were then there could not be infinitely many applications of (b)-r.

By frugality (cf. Corollary 6.4.7), there must be a $k \geq 0$ be such that every Δ_i with $i \geq k$ occurs infinitely often on the branch above. Denote by $|\Delta|$ the length of a given list Δ and let l be minimum of $\{|\Delta_i| : i \geq k\}$. In other words, l is the minimal length of the Δ_i with $i \geq k$.

To prove the lemma, we first claim that there is an $n \geq k$ such that $|\Delta_n| = l$ and the leftmost expression in Δ_n is of the form $e^{(b)}$ for some e. Suppose, towards a contradiction, that this is not the case. Then there must be a $u \geq k$ such that $|\Delta_u| = l$ and the leftmost expression in Δ_u is *not* of the form $e^{(b)}$ for any e. Note that (b)-r is the only rule apart from k_0 that can increase the length of the succedent (when read bottom-up). It follows that for no $w \geq u$ the leftmost expression in Δ_w is of the form $e^{(b)}$, contradicting the fact that (b)-r is applied infinitely often.

Now let $n \geq k$ be such that $|\Delta_n| = l$ and the leftmost expression of Δ_n is $e^{(b)}$. Since the rule (b)-r must at some point after Δ_n be applied to $e^{(b)}$, we may assume without loss of generality that $\Gamma_n \Rightarrow_{A_n} \Delta_n$ is the conclusion of an application of (b)-r. By the pigeonhole principle, there must be an $m > n$ such that $\Gamma_n \Rightarrow_{A_n} \Delta_n$ and $\Gamma_m \Rightarrow_{A_m} \Delta_m$ are the same sequents. We claim that these sequents satisfy the three properties above. Properties (i) and (ii) directly hold by construction. Property (iii) follows from the fact that Δ_n is of minimal length and has $e^{(b)}$ as leftmost expression. □

We are now ready for the completeness proof.

6.5.4. THEOREM. *Every valid sequent is provable in* SGKAT^∞*.*

Proof:
Given a valid sequent, we do a bottom-up proof search with the following strategy. Throughout the procedure all leaves remain valid, in most cases by an appeal to invertibility.

1. Apply left logical rules as long as possible. If this stage terminates, it will be at a leaf of the form $\Gamma \Rightarrow_A \Delta$, where Γ is exposed. We then go to stage (2). If left logical rules remain applicable, we stay in this stage (1) forever and create an infinite branch.

2. Apply right logical rules until one of the following happens:

   (a) We reach a leaf at which no right logical rule can be applied. This means that the leaf must be a valid sequent of the form $\Gamma \Rightarrow_A \Delta$ such that Γ is exposed, and Δ is either exposed or begins with a test b such $A \upharpoonright b \neq A$. We go to stage (4).

   (b) If (a) does not happen, then at some point we must reach a valid sequent of the $\Gamma \Rightarrow_A e^{(b)}, \Delta$ which together with an ancestor satisfies properties (i) - (iii) of Lemma 6.5.3. In this case Lemma 6.5.2 is applicable. Hence we must be at a leaf of the form $\Gamma \Rightarrow_A e^{(b)}, \Delta$ such that $e^{(b)} \Rightarrow_A 0$ is valid. We then go to stage (3).

   Since at some point either (a) or (b) must be the case, stage (2) always terminates.

3. We are at a valid leaf of the form $\Gamma \Rightarrow_A e^{(b)}, \Delta$, where Γ is exposed. If $A = \emptyset$, we apply $\bot$. Otherwise, if $A \neq \emptyset$, we use the validity of $\Gamma \Rightarrow_A e^{(b)}, \Delta$ and $e^{(b)} \Rightarrow_A 0$ to find:

   $$A \diamond [\![\Gamma]\!] \subseteq A \diamond [\![e^{(b)}]\!] \diamond [\![\Delta]\!] = \emptyset.$$

   We claim that $[\![\Gamma]\!] = \emptyset$. Indeed, suppose towards a contradiction that $\alpha x \in [\![\Gamma]\!]$. By the exposedness of Γ and item (i) of Lemma 6.2.14, we would have $\beta x \in [\![\Gamma]\!]$ for some $\beta \in A$, contradicting the statement above. Therefore, the sequent $\Gamma \Rightarrow_{\mathsf{At}} 0$ is valid. We apply the rule k_0 and loop back to stage (1).

   Stage (3) only comprises a single step and thus always terminates.

4. Let $\Gamma \Rightarrow_A \Delta$ be the current leaf. By construction $\Gamma \Rightarrow_A \Delta$ is valid, Γ is exposed, and Δ is either exposed or begins with a test b such that $A \upharpoonright b \neq A$. Note that only rules id, $\bot$, k, and k_0 can be applicable. By Lemma 6.5.1, at least one of them must be applicable. If id is applicable, apply id. If $\bot$ is applicable, apply $\bot$. If k is applicable, apply k and loop back to stage (1). Note that this application of k will have priority (cf. Definition 6.3.3), and is therefore invertible.

   Finally, suppose that only k_0 is applicable. We claim that, by validity, the list Γ is not ϵ. Indeed, since A is non-empty, and Δ either begins with a primitive program p or a test b such that $A \upharpoonright b \neq A$, the sequent

   $$\epsilon \Rightarrow_A \Delta$$

   must be invalid. Hence Γ must be of the form p, Θ. We apply k_0, which has priority and thus is invertible, and loop back to stage (1).

   Like stage (3), stage (4) only comprises a single step and thus always terminates.

We claim that the constructed derivation is fair. Indeed, every stage except stage (1) terminates. Therefore, every infinite branch must either eventually remain in stage (1), or pass through stages (3) or (4) infinitely often. Since k and k_0 shorten the antecedent, and no left logical rule other than (b)-l lengthens it, such branches must be fair. □

By Corollary 6.4.8 we obtain that SGKAT^∞ is regularly complete.

6.5.5. Corollary. *Every valid sequent has a regular* SGKAT^∞*-proof.*

6.6 An inequational axiomatisation

In this section we propose an inequational axiomatisation PoGKAT for GKAT and sketch partial translations from SGKAT^∞ into PoGKAT. The question of whether PoGKAT is complete with respect to the language model is left open. In future work we wish to investigate whether this question can be settled by completing our partial translation into a full translation.

The base of our axiomatisation is provided by inequational logic, whose axioms and rules are given in Figure 6.4. For background on inequational logic we refer the reader to [16] and [65].

$$e \leq e \qquad \frac{e \leq f \quad f \leq g}{e \leq g} \qquad \frac{e \leq f}{\widehat{\sigma}(e) \leq \widehat{\sigma}(g)}$$

$$\frac{e_1 \leq f_1 \quad e_2 \leq f_2}{e_1 +_b e_2 \leq f_1 +_b f_2} \qquad \frac{e_1 \leq f_1 \quad e_2 \leq f_2}{e_1 \cdot e_2 \leq f_1 \cdot f_2} \qquad \frac{e \leq f}{e^{(b)} \leq f^{(b)}}$$

Figure 6.4: The axioms and rules of inequational logic in the signature of GKAT.

In the following every equation $e \equiv f$ should be read as a shorthand for the pair of inequations $e \leq f$ and $f \leq e$.

6.6.1. Definition. The system PoGKAT consists of the following axioms and rules.

1. All axioms and rules of inequational logic, as given in Figure 6.4.
2. The axiom $b \leq c$ for all tests b, c such that c is a Boolean consequence of b.
3. All axioms from [98, Sect. 3], *i.e.* all axioms in Figure 6.2 above.
4. The least fixed point rule: if $eg +_b f \leq g$, then $e^{(b)} f \leq g$.

The following lemma collects some useful PoGKAT-derivable (in)equations.

6.6.2. Lemma. *The following equations are provable in* PoGKAT.

$$b(e +_b f) \equiv be \qquad \overline{b}(e +_b f) \equiv \overline{b}f \qquad be^{(b)} \equiv bee^{(b)} \qquad \overline{b}e^{(b)} \equiv \overline{b}$$

Proof:
For the first equation, we refer the reader to the proof of Fact (U8) in [98], which directly transfers to our system. The other equations readily follow using axioms (U2) and (W1) from Figure 6.2. □

The following proposition will be useful later, and at the same time serves as an example of a more involved PoGKAT-proof.

6.6.3. Proposition. PoGKAT $\vdash e^{(b)} \leq (e(e +_b 1))^{(b)}$.

Proof:
Let us abbreviate $(e(e +_b 1))^{(b)}$ by f. We first deduce

$$\begin{aligned}
\mathsf{PoGKAT} \vdash\ & e(e +_b 1)(e +_b 1)f +_b 1 \\
&\leq e(e(e +_b 1)f +_b 1(e +_b 1)f) +_b 1 && \text{(U5)} \\
&\leq e(e(e +_b 1)f +_b \overline{b} \cdot 1(e +_b 1)f) +_b 1 && \text{(U4, U2)} \\
&\leq e(e(e +_b 1)f +_b \overline{b}(e +_b 1)f) +_b 1 && \text{(S4)} \\
&\leq e(e(e +_b 1)f +_b \overline{b} \cdot 1 \cdot f) +_b 1 && \text{(Lem. 6.6.2)} \\
&\leq e(e(e +_b 1)f +_b \overline{b} \cdot f) +_b 1 && \text{(S4)} \\
&\leq e(e(e +_b 1)f +_b \overline{b}) +_b 1 && \text{(Lem. 6.6.2, def. of } f\text{)} \\
&\leq e(e(e +_b 1)f +_b \overline{b} \cdot 1) +_b 1 && \text{(S5)} \\
&\leq e(e(e +_b 1)f +_b 1) +_b 1 && \text{(U4, U2)} \\
&\leq ef +_b 1 && \text{(W1)} \\
&\leq ef +_b f && \text{(U4, U2, S1, Lem 6.6.2)} \\
&\leq (e +_b 1)f. && \text{(U5)}
\end{aligned}$$

Hence, by the least fixed point rule, we have PoGKAT $\vdash f \cdot 1 \leq (e +_b 1)f$ and therefore:

$$\mathsf{PoGKAT} \vdash f \leq (e +_b 1)f.$$

By monotonicity, it then follows that

$$\mathsf{PoGKAT} \vdash ef +_b 1 \leq e(e +_b 1)f +_b 1,$$

which by (W1) means that

$$\mathsf{PoGKAT} \vdash ef +_b 1 \leq f.$$

A final application of the least fixed point rule yields $\mathsf{PoGKAT} \vdash e^{(b)} \cdot 1 \leq f$, from which it follows that

$$\mathsf{PoGKAT} \vdash e^{(b)} \leq f,$$

as required. □

We wish to show that PoGKAT is complete with respect to the language model, by translating proofs from SGKAT^∞ into PoGKAT. For technical reasons, it will be useful to augment SGKAT by the following admissible axiom

$$\frac{}{\Gamma \Rightarrow_A \Gamma}\ \mathsf{id_s}$$

We will call the resulting system $\mathsf{SGKAT_s}$. The system $\mathsf{SGKAT_s^\infty}$, supporting infinitary proofs, is defined using the same fairness condition on infinite branches as we used to define SGKAT^∞ from SGKAT.

The following proposition entails that all *well-founded* $\mathsf{SGKAT_s}$-proofs admit such a translation. We will slightly abuse notation by, for instance, writing

$$\mathsf{PoGKAT} \vdash A \upharpoonright b \cdot \Gamma \leq e \cdot \Delta,$$

when strictly we mean

$$\mathsf{PoGKAT} \vdash \alpha_1 \cdot \ldots \cdot \alpha_n \cdot e_1 \cdot \ldots \cdot e_m \leq e \cdot f_1 \cdot \ldots \cdot f_k,$$

where $A \upharpoonright b$ is $\{\alpha_1, \ldots, \alpha_n\}$, and Γ and Δ resp. are $\{e_1, \ldots, e_m\}$ and $\{f_1, \ldots, f_k\}$.

6.6.4. PROPOSITION. *If all of the premisses of some* $\mathsf{SGKAT_s}$*-rule are derivable in* PoGKAT*, then so is the conclusion.*

Proof:
We will treat $+_b$-r as illustrative case. Suppose that the following hold:

$$\mathsf{PoGKAT} \vdash A \upharpoonright b \cdot \Gamma \leq e \cdot \Delta \tag{6.1}$$
$$\mathsf{PoGKAT} \vdash A \upharpoonright \overline{b} \cdot \Gamma \leq f \cdot \Delta \tag{6.2}$$

Then we have:

$$\begin{aligned}
\mathsf{PoGKAT} \vdash A \cdot \Gamma &\leq A \cdot \Gamma +_b A \cdot \Gamma && \text{(U1)} \\
&\leq b \cdot A \cdot \Gamma +_b \overline{b} \cdot A \cdot \Gamma && \text{(U2, U4)} \\
&\leq A \upharpoonright b \cdot \Gamma +_b A \upharpoonright \overline{b} \cdot \Gamma && \text{(Boolean, monotonicity)} \\
&\leq e \cdot \Delta +_b f \cdot \Delta && \text{(6.1, 6.2, monotonicity)} \\
&\leq (e +_b f) \cdot \Delta, && \text{(U5)}
\end{aligned}$$

as required. □

The previous proposition implies that well-founded $\mathsf{SGKAT_s}$-proofs can be translated into PoGKAT-proofs. Following [33], our goal is to extend this translation to also include the non-well-founded, but regular $\mathsf{SGKAT}_\mathsf{s}^\infty$-proofs. It will be convenient to assume that our $\mathsf{SGKAT_s}$-proofs are explicitly given as cyclic proofs, *i.e.* as finite trees with back edges.

We can directly obtained a notion of cyclic $\mathsf{SGKAT_s}$-proofs by the observation that $\mathsf{SGKAT_s}^\infty$ is a simple path-based proof system (cf. Definition I.2.7). For the sake of completeness, we spell it out here.

6.6.5. Definition. A *cyclic* $\mathsf{SGKAT_s}$*-derivation* (π, f) is a pair consisting of a finite $\mathsf{SGKAT_s}$-derivation, together with a partial function $f : \pi \rightharpoonup \pi$ from the set of nodes of π to itself, such that for every $u \in \mathsf{dom}(f)$: (i) u is a leaf of π, and (ii) $f(u)$ is a proper ancestor of u labelled by exactly the same sequent.

A *cyclic* $\mathsf{SGKAT_s}$*-proof* is a cyclic $\mathsf{SGKAT_s}$-derivation such that every leaf l either belongs to $\mathsf{dom}(f)$ or is an axiom, and for every $u \in \mathsf{dom}(f)$ the path $[f(u), u)$ contains an application of (b)-l.

As usual, a node $u \in \mathsf{dom}(f)$ is called a *repeating leaf*, and $f(u)$ its *companion.*

It will be convenient to restrict attention to cyclic proofs of a certain nice shape. For this we need the following definition.

6.6.6. Definition. A cyclic proof (π, f) is *oriented* if for every $u \in \mathsf{dom}(f)$ it holds that $f(u)$ is the conclusion of an application of (b)-l.

A cyclic proof (π, f) is *monotone* if for every $u \in \mathsf{dom}(f)$ the antecedent of u is a final segment of every antecedent on the path $[f(u), u)$.

6.6.7. Remark. The concept of monotonicity is not new in the cyclic proof theory literature. An analogous condition appears for instance in [3] under the same name. A related condition was earlier considered by Sprenger & Dam under the name *tree-dischargeability* [99]. In both cases monotone (resp. tree-dischargeable) proofs are used as an intermediate step in the translation of cyclic proofs into a system with an explicit induction rule.

6.6.8. Example. Let $e = p^{(ab)}$, and $f = qe$, and $g = p +_b q$, and consider the following cyclic $\mathsf{SGKAT_s}$-proof Π_2. As usual we omit branches that can be immediately closed by an application of $\bot$. The $(\bullet)$ indicates repeating leaves and their companions. Note that $\mathsf{At} \restriction \overline{ab} \restriction a = \mathsf{At} \restriction a\bar{b}$, and that $\mathsf{At} \restriction \overline{ab} \restriction \overline{a} = \mathsf{At} \restriction \overline{a}$.

$$
\begin{prooftree}
\AxiomC{$e, f^{(a)} \Rightarrow_{\mathsf{At}} g^{(a)} \quad (\bullet)$}
\LeftLabel{k}
\UnaryInfC{$p, e, f^{(a)} \Rightarrow_{\mathsf{At}\restriction ab} p, g^{(a)}$}
\LeftLabel{$+_b$-r}
\UnaryInfC{$p, e, f^{(a)} \Rightarrow_{\mathsf{At}\restriction ab} g, g^{(a)}$}
\LeftLabel{(a)-r}
\UnaryInfC{$p, e, f^{(a)} \Rightarrow_{\mathsf{At}\restriction ab} g^{(a)}$}
\AxiomC{$e, f^{(a)} \Rightarrow_{\mathsf{At}} g^{(a)} \quad (\bullet)$}
\RightLabel{k}
\UnaryInfC{$q, e, f^{(a)} \Rightarrow_{\mathsf{At}\restriction a\bar{b}} q, g^{(a)}$}
\RightLabel{$+_b$-r}
\UnaryInfC{$q, e, f^{(a)} \Rightarrow_{\mathsf{At}\restriction a\bar{b}} g, g^{(a)}$}
\RightLabel{(a)-r}
\UnaryInfC{$q, e, f^{(a)} \Rightarrow_{\mathsf{At}\restriction a\bar{b}} g^{(a)}$}
\RightLabel{$\cdot$-l}
\UnaryInfC{$qe, f^{(a)} \Rightarrow_{\mathsf{At}\restriction a\bar{b}} g^{(a)}$}
\AxiomC{$\epsilon \Rightarrow_{\mathsf{At}\restriction\overline{a}} \epsilon$}
\RightLabel{(a)-r}
\UnaryInfC{$\epsilon \Rightarrow_{\mathsf{At}\restriction\overline{a}} g^{(a)}$}
\RightLabel{(a)-l}
\BinaryInfC{$f^{(a)} \Rightarrow_{\mathsf{At}\restriction\overline{ab}} g^{(a)}$}
\LeftLabel{(ab)-l}
\BinaryInfC{$e, f^{(a)} \Rightarrow_{\mathsf{At}} g^{(a)} \quad (\bullet)$}
\end{prooftree}
$$

Although Π_2 is oriented, it is not monotone. Indeed, the antecedent $e, f^{(a)}$ is not a final segment of the antecedent $f^{(a)}$.

The following proof, which we shall call Π_3 is an example of a proof with the same conclusion as Π_2 which *is* monotone. In fact, Π_2 and Π_3 represent exactly the same regular infinitary proof, in the sense that their infinitary unravelings are equal.

$$\dfrac{\dfrac{\dfrac{\dfrac{e, f^{(a)} \Rightarrow_{\mathsf{At}} g^{(a)} \quad (\bullet)}{p, e, f^{(a)} \Rightarrow_{\mathsf{At}\restriction ab} p, g^{(a)}}\,\mathsf{k}}{p, e, f^{(a)} \Rightarrow_{\mathsf{At}\restriction ab} g, g^{(a)}}\,+_b\text{-}r}{p, e, f^{(a)} \Rightarrow_{\mathsf{At}\restriction ab} g^{(a)}}\,(a)\text{-}r \qquad \Pi_4}{e, f^{(a)} \Rightarrow_{\mathsf{At}} g^{(a)} \quad (\bullet)}\,(ab)\text{-}l$$

where Π_4 is the proof

$$\dfrac{\dfrac{\dfrac{\dfrac{e, f^{(a)} \Rightarrow_{\mathsf{At}} g^{(a)} \quad (\bullet)}{p, e, f^{(a)} \Rightarrow_{\mathsf{At}\restriction ab} p, g^{(a)}}\,\mathsf{k}}{p, e, f^{(a)} \Rightarrow_{\mathsf{At}\restriction ab} g, g^{(a)}}\,+_b\text{-}r}{p, e, f^{(a)} \Rightarrow_{\mathsf{At}\restriction ab} g^{(a)}}\,(a)\text{-}r \qquad f^{(a)} \Rightarrow_{\mathsf{At}\restriction \overline{ab}} g^{(a)} \quad (\bullet\bullet)}{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{e, f^{(a)} \Rightarrow_{\mathsf{At}} g^{(a)} \quad (\bullet)}{q, e, f^{(a)} \Rightarrow_{\mathsf{At}\restriction a\bar{b}} q, g^{(a)}}\,\mathsf{k}}{q, e, f^{(a)} \Rightarrow_{\mathsf{At}\restriction a\bar{b}} g, g^{(a)}}\,+_b\text{-}r}{q, e, f^{(a)} \Rightarrow_{\mathsf{At}\restriction a\bar{b}} g^{(a)}}\,(a)\text{-}r}{qe, f^{(a)} \Rightarrow_{\mathsf{At}\restriction a\bar{b}} g^{(a)}}\,\cdot\text{-}l \qquad \dfrac{\epsilon \Rightarrow_{\mathsf{At}\restriction \bar{a}} \epsilon}{\epsilon \Rightarrow_{\mathsf{At}\restriction \bar{a}} g^{(a)}}\,(a)\text{-}r}{f^{(a)} \Rightarrow_{\mathsf{At}\restriction \overline{ab}} g^{(a)} \quad (\bullet\bullet)}\,(a)\text{-}l}\,(ab)\text{-}l$$

The following proposition shows that for every cyclic proof we can find a monotone and oriented cyclic proof which represents the same regular infinitary proof, generalising Example 6.6.8.

6.6.9. PROPOSITION. *Every regular* $\mathsf{SGKAT}_{\mathsf{s}}^{\infty}$*-proof* π *is the unravelling of an oriented and monotone cyclic* $\mathsf{SGKAT}_{\mathsf{s}}$*-proof.*

Proof:
Note that it suffices to show that for every infinite branch

$$\Gamma_0 \Rightarrow_{A_0} \Delta_0; \Gamma_1 \Rightarrow_{A_1} \Delta_1; \Gamma_2 \Rightarrow_{A_2} \Delta_2; \cdot \ldots \cdot$$

through π, there are $n < m$ such that (i) $\Gamma_n \Rightarrow_{A_n} \Delta_n$ and $\Gamma_m \Rightarrow_{A_m} \Delta_m$ are the same sequents, (ii) $\Gamma_n \Rightarrow_{A_n} \Delta_n$ is the conclusion of an application of (b)-l in π, and (iii) for every $i \in [n, m)$ it holds that Γ_n is a final segment of Γ_i. Indeed, the proposition then follows from an easy appeal to Kőnig's lemma, such as in the proof of Proposition I.2.12.

Now note that the required statement is exactly the same as Lemma 6.5.3, but about the antecedent rather than the succedent. Its proof is also entirely analogous. □

Suppose that (π, f) is an oriented and monotone cyclic $\mathsf{SGKAT_s}$-proof. It might of course happen that there is some f', distinct from f, for which we have $\mathsf{dom}(f) = \mathsf{dom}(f')$ and such that (π, f') is also an oriented and monotone cyclic $\mathsf{SGKAT_s}$-proof. Later in this section it will be convenient to work with proofs where f is *minimal*, in the sense that for every f' as above, it holds that $f(l)$ is an ancestor of $f'(l)$ for every repeating leaf $l \in \mathsf{dom}(f)$. It is not hard to see how to obtain the following corollary from Proposition 6.6.9.

6.6.10. Corollary. *Every regular* $\mathsf{SGKAT_s^\infty}$*-proof is the unravelling of an oriented and monotone cyclic* $\mathsf{SGKAT_s}$*-proof* (π, f)*, which moreover is* minimal*. That is, for every* $f' : \mathsf{dom}(f) \to \pi$ *such that* (π, f') *is an oriented and monotone cyclic* $\mathsf{SGKAT_s}$*-proof, it holds that* $f(l)$ *is an ancestor of* $f'(l)$ *for every* $l \in \mathsf{dom}(f)$*.*

We want to inductively show how to translate each oriented and monotone cyclic $\mathsf{SGKAT_s}$-proof π with endsequent $\Gamma \Rightarrow_A \Delta$ into a PoGKAT-proof. For this we will use a measure given by the following definition.

6.6.11. Definition. Let (π, f) be a cyclic $\mathsf{SGKAT_s}$-proof which is both oriented and monotone. The *size* $|(\pi, f)|$ of (π, f) is defined to be its number of nodes. The *while-height* $\mathsf{wh}(\pi, f)$ of (π, f) is the maximal value of $\mathsf{wh}(e^{(b)})$ of all sequents of the form $e^{(b)}, \Gamma \Rightarrow_A \Delta$ in $\mathsf{ran}(f)$, which is 0 if $\mathsf{ran}(f) = \emptyset$.

In the following we will for notational convenience often suppress the back edge function f when referring to a cyclic proof (π, f). Our argument will go by induction on the measure $\langle \mathsf{wh}(\pi), |\pi| \rangle$, ordered lexicographically. Given an oriented and monotone cyclic proof π, we make a case distinction as to whether the root of π is a companion or not.

If the root is *not* a companion, let $(\pi_i)_{1 \leq i \leq n}$ be the subproofs generated by the premisses of the final rule application of π. Then we have $|\pi_i| < |\pi|$ for each i. As moreover $\mathsf{wh}(\pi_i) \leq \mathsf{wh}(\pi)$, we can invoke the induction hypothesis to obtain PoGKAT-proofs of the premisses of the last rule application of π. Finally, Proposition 6.6.4 gives the desired PoGKAT-proof of $\Gamma \Rightarrow_A \Delta$.

The difficult case is where the root of π is a companion. By the fact that π is oriented, the last rule applied in π must then be (b)-l. Hence π looks as follows:

$$\dfrac{\overset{\pi_1}{e, e^{(b)}, \Gamma \Rightarrow_{A \restriction b} \Delta} \qquad \overset{\pi_2}{\Gamma \Rightarrow_{A \restriction \bar{b}} \Delta}}{e^{(b)}, \Gamma \Rightarrow_A \Delta}\ (b)\text{-}l$$

Since π is monotone, there is no back edge from π_2 to the endsequent of π. We can therefore apply the induction hypothesis to π_2 to obtain $\mathsf{PoGKAT} \vdash A \cdot \bar{b} \cdot \Gamma \leq \Delta$.

However, we cannot apply the induction hypothesis to π_1, since π_1 is not a subproof. To proceed, we intend to use an idea from [33], namely to compute an *invariant* for e. Unfortunately, at the time of writing we do not have a technique that works for all cyclic $\mathsf{SGKAT_s}$-proofs. We do, however, know how to proceed under certain additional assumptions, which are discussed below. Extending these techniques to the general case is left for future work.

The following lemma, in [33] called the Invariant Lemma, is not difficult to prove but conceptually important.

6.6.12. LEMMA. *For all expressions e, I, Γ, Δ:*

$$\textit{if} \begin{cases} \mathsf{PoGKAT} \vdash \bar{b} \cdot \Gamma \leq I \\ \mathsf{PoGKAT} \vdash b \cdot e \cdot I \leq I \\ \mathsf{PoGKAT} \vdash I \leq \Delta \end{cases} \quad \textit{then} \;\; \mathsf{PoGKAT} \vdash e^{(b)} \cdot \Gamma \leq \Delta.$$

Proof:
From the first two assumptions it is easy to derive that $\mathsf{PoGKAT} \vdash eI +_b \Gamma \leq \Delta$. Using the least fixed point rule, we then find $\mathsf{PoGKAT} \vdash e^{(b)} \cdot \Gamma \leq I$. Finally, by the third assumption and transitivity, it follows that $\mathsf{PoGKAT} \vdash e^{(b)} \cdot \Gamma \leq \Delta$, as required. □

Because of its second property, the expression I in the above lemma is called an *invariant for $b \cdot e$ (on the left).*

Our first application of the invariant lemma requires the assumptions (a) and (b) in the proposition below. Assumption (b) is based on a very similar example in [33].

6.6.13. PROPOSITION. *Let (π, f) be an oriented and monotone cyclic $\mathsf{SGKAT_s}$-proof such that for every companion $u \in \mathsf{ran}(f)$, labelled by, say, $\Theta \Rightarrow_A \Sigma$, the following hold:*

(a) $A = \mathsf{At}$.

(b) For each sequent of the form $\Theta \Rightarrow_{A'} \Sigma'$ in the subtree generated by u, it holds both that $A' = A$ and that $\Sigma' = \Sigma$.

Then there is a PoGKAT-proof of the root of π.

Proof (sketch):
We proceed by induction on $\langle \mathsf{wh}(\pi), |\pi| \rangle$. As we have seen above, the only interesting case is where the root of π is a companion. By the fact that π is oriented and monotone, and assumption (a), we know that π looks as follows:

$$\dfrac{\begin{array}{c}\pi_1 \\ e, e^{(b)}, \Gamma \Rightarrow_{\mathsf{At}\restriction b} \Delta\end{array} \qquad \begin{array}{c}\pi_2 \\ \Gamma \Rightarrow_{\mathsf{At}\restriction \bar{b}} \Delta\end{array}}{e^{(b)}, \Gamma \Rightarrow_{\mathsf{At}} \Delta}\;(b)\text{-}l$$

As argued above, we know by monotonicity that π_2 cannot contain a back edge to the root of π. This means that π_2 is a subproof such that $|\pi_2| < |\pi|$. Moreover, π_2 clearly inherits conditions (a) and (b) from π. Hence, we can apply the induction hypothesis to π_2, which gives $\mathsf{PoGKAT} \vdash \mathsf{At} \upharpoonright b \cdot \Gamma \leq \Delta$.

To show that also $\mathsf{PoGKAT} \vdash \mathsf{At} \upharpoonright b \cdot e \cdot e^{(b)} \cdot \Gamma \leq \Delta$, we consider the proof $\pi_1[\Delta/e^{(b)}, \Gamma]$ obtained by replacing in π_1 every occurrence of $e^{(b)}, \Gamma$ as a final segment in an antecedent, by Δ. If in π_1 a sequent of the form $e^{(b)}, \Gamma \Rightarrow_A \Delta'$ is reached, then the corresponding node in $\pi_1[\Delta/e^{(b)}, \Gamma]$ is of the form $\Delta \Rightarrow_A \Delta'$. By assumption (b), we then know that this node must be of the form $\Delta \Rightarrow_A \Delta$, and we close it with an application of $\mathsf{id_s}$. This transformation can be depicted as follows, where to the left of the arrow $\rightsquigarrow$ a branch of π is shown, and the to right the corresponding branch of $\pi_1[\Delta/e^{(b)}, \Gamma]$.

$$
\dfrac{\dfrac{\vdots}{e^{(b)}, \Gamma \Rightarrow_A \Delta}}{\dfrac{\vdots}{e, e^{(b)}, \Gamma \Rightarrow_{\mathsf{At}\upharpoonright b} \Delta}}
\quad \rightsquigarrow \quad
\dfrac{\dfrac{}{\Delta \Rightarrow_A \Delta}\,\mathsf{id_s}}{\dfrac{\vdots}{e, \Delta \Rightarrow_{\mathsf{At}\upharpoonright b} \Delta}}
$$

Clearly $\pi_1[\Delta/e^{(b)}, \Gamma]$ is indeed a cyclic $\mathsf{SGKAT_s}$-proof of $e, \Delta \Rightarrow_{\mathsf{At}\upharpoonright b} \Delta$. Moreover, since by assumption the endsequent of π is a companion, and by construction the final segment Δ of each antecedent in $\pi_1[\Delta/e^{(b)}, \Gamma]$ is never explored, we have $\mathsf{wh}(\pi_1[\Delta/e^{(b)}, \Gamma]) < \mathsf{wh}(\pi)$. As before, $\pi_1[\Delta/e^{(b)}, \Gamma]$ inherits conditions (a) and (b) from π. Hence, we can apply the induction hypothesis, from which we obtain $\mathsf{PoGKAT} \vdash b \cdot e \cdot \Delta \leq \Delta$. Since trivially $\mathsf{PoGKAT} \vdash \Delta \leq \Delta$, we can now apply Lemma 6.6.12 to obtain $\mathsf{PoGKAT} \vdash e^{(b)} \cdot \Gamma \leq \Delta$, as required. □

6.6.14. Remark. The use of the invariant lemma in the above proof is an overkill, because the invariant I is simply Δ itself. The sufficiency of this simple invariant is enabled by assumption (b). In [33], in the context of Kleene Algebra, a proposition similar to our Proposition 6.6.13 is refined to eliminate the need of assumption (b). If we were to apply their technique to our proof system, roughly the idea would be to extract from a cyclic SGKAT-proof π of $e^{(b)}, \Gamma \Rightarrow_{\mathsf{At}} \Delta$ an expression that, provably in PoGKAT, corresponds to the intersection of all Δ' such that $e^{(b)}, \Gamma \Rightarrow_A \Delta'$ occurs in π (note that this is indeed Δ itself in the presence of assumption (b)). This expression can then be used as the invariant I in an application of Lemma 6.6.12.

It is unclear whether their method for computing this expression I, by viewing proofs as automata and calculating the so-called *minimal solution* to this automaton, can be applied to SGKAT. In fact, it is to the best of our knowledge an open question whether languages recognised by `GKAT`-expressions are closed under taking intersections at all.

The proof Π_1 of Example 6.2.8 satisfies assumption (a) of Proposition 6.6.13, but not assumption (b). Indeed, the companion labelled by $(cp)^{(b)} \Rightarrow_{\mathsf{At}} \Delta_1$ in Π_1 has a descendant labelled $(cp)^{(b)} \Rightarrow_{\mathsf{At}} \Delta_2$, where $\Delta_1 \neq \Delta_2$.

However, after applying (b)-l to $(cp)^{(b)}$ *twice*, we reach in Π_1 a repetition of $(cp)^{(b)} \Rightarrow_{\mathsf{At}} \Delta_1$. We will provide a strengthening of Proposition 6.6.13, featuring a weakening of assumption (b), which *is* satisfied by Π_1. The key idea is to require a certain uniformity in the amount of applications of (b)-l needed to reach a repetition. To make this formal, we will use the following definition.

6.6.15. Definition. Let π be an $\mathsf{SGKAT_s}$-derivation of $e^{(b)}, \Gamma \Rightarrow_{\mathsf{At}} \Delta$, and let u be a node of π, such that $e^{(b)}, \Gamma$ is a final segment of every antecedent from the root of π to u, and equal to the antecedent of u itself. Then the *unfolding depth* of u is the number of those rule instances of (b)-l on the path from the root to u, for which the antecedent of the conclusion of this rule instance is $e^{(b)}, \Gamma$.

For any node u that does not meet this condition, the *unfolding depth* is undefined.

6.6.16. Example. In the proof Π_1 of Example 6.2.8, viewed as a cyclic proof, the node labelled $(cp)^{(b)} \Rightarrow_{\mathsf{At}} \Delta_2$ has unfolding depth 1, and the repeating leaf labelled $(cp)^{(b)} \Rightarrow_{\mathsf{At}} \Delta_1$ has unfolding depth 2.

The assumption (b') in the following proposition has two components. First, it requires for every companion u that there is some number $n > 0$ such that after unfolding the leftmost expression of the antecedent of u (which is of the form $e^{(b)}$ by orientedness) n times, one ends up at a node labelled by the same sequent as u. Second, it requires that all repeating leaves with u as companion occur sufficiently deep in the proof. Note that when assumption (b) is satisfied, we can simply choose $n = 1$. This shows that assumption (b') is indeed a weakening of assumption (b).

Note also that the following proposition is formulated for *minimal* oriented and monotone cyclic SGKAT_s-proofs (cf. Corollary 6.6.10).

6.6.17. Proposition. *Let (π, f) be a minimal oriented and monotone cyclic $\mathsf{SGKAT_s}$-proof such that for every companion $u \in \mathsf{ran}(f)$, labelled by say $\Theta \Rightarrow_A \Sigma$, the following hold:*

(a) $A = \mathsf{At}$.

(b') In the subderivation π' of π generated by u, it holds for some $n > 0$ that:

(i) every node in π' of unfolding depth n is of the form $\Theta \Rightarrow_A \Sigma$;

(ii) for every node $l \in f^{-1}(u)$ contained in π', the unfolding depth of l is at least n.

Proof (sketch):
We proceed by induction on $\langle \mathsf{wln}(\pi), |\pi| \rangle$, where $\mathsf{wln}(\pi)$ is the number of applications of the left while rule (c)-l in π (for any test c). Again, the only interesting case is when the root of π is a companion. As we have seen before, we know that π then looks as follows:

$$\dfrac{\begin{matrix}\pi_1 \\ e, e^{(b)}, \Gamma \Rightarrow_{\mathsf{At}\restriction b} \Delta\end{matrix} \qquad \begin{matrix}\pi_2 \\ \Gamma \Rightarrow_{\mathsf{At}\restriction \bar{b}} \Delta\end{matrix}}{e^{(b)}, \Gamma \Rightarrow_{\mathsf{At}} \Delta}\ (b)\text{-}l$$

We will now sketch how to show that $\mathsf{PoGKAT} \vdash e^{(b)} \cdot \Gamma \leq \Delta$ under assumptions (a) and (b'). Let n be as given by assumption (b') applied to the root of π (which we assumed to be a companion), and define the expression $I := (e +_b 1)^n \cdot (\Delta +_b \Gamma)$. Here $(e +_b 1)^0 = 1$ and $(e +_b 1)^{k+1} := (e +_b 1) \cdot (e +_b 1)^k$. We claim that

$$\mathsf{PoGKAT} \vdash (e +_b 1)^n \cdot (\Delta +_b \Gamma) \leq \Delta +_b \Gamma. \qquad (*)$$

Note that trivially $\mathsf{PoGKAT} \vdash \bar{b} \cdot \Gamma \leq \Delta +_b \Gamma$ and $\mathsf{PoGKAT} \vdash \Delta +_b \Gamma \leq \Delta +_b \Gamma$. Hence, if (*) holds, then, of course, also

$$\mathsf{PoGKAT} \vdash b \cdot (e +_b 1)^n \cdot (\Delta +_b \Gamma) \leq \Delta +_b \Gamma$$

and thus we can apply Lemma 6.6.12 with the arguments

$$(e +_b 1)^n, \qquad \Delta +_b \Gamma, \qquad \Gamma, \qquad \Delta +_b \Gamma,$$

in place of e, I, Γ, Δ. We then obtain:

$$\mathsf{PoGKAT} \vdash ((e +_b 1)^n)^{(b)} \cdot \Gamma \leq \Delta +_b \Gamma.$$

By a straightforward generalisation of the proof of Lemma 6.6.3, it can be shown that $\mathsf{PoGKAT} \vdash e^{(b)} \leq ((e +_b 1)^n)^{(b)}$. It then follows by monotonicity that $\mathsf{PoGKAT} \vdash e^{(b)} \cdot \Gamma \leq \Delta +_b \Gamma$. Applying the induction hypothesis to π_2 (which we can do by monotonicity and the fact that π_2 inherits both conditions (a) and (b')), we know that $\mathsf{PoGKAT} \vdash \bar{b} \cdot \Gamma \leq \Delta$. Hence we obtain $\mathsf{PoGKAT} \vdash e^{(b)} \cdot \Gamma \leq \Delta$, as required.

To finish the proof, it thus suffices to show that (*) holds. The idea is to build a cyclic $\mathsf{SGKAT_s}$-proof (π', f') of

$$(e +_b 1)^n, \Delta +_b \Gamma \Rightarrow_{\mathsf{At}} \Delta +_b \Gamma$$

to which we can apply the induction hypothesis. This cyclic $\mathsf{SGKAT_s}$-proof π' can be depicted as follows, suppressing branches which can be immediately closed by an application of $\bot$.

$$\dfrac{\dfrac{\begin{matrix}\pi_1[(e +_b 1)^{n-1}, (\Delta +_b \Gamma)/e^{(b)}, \Gamma] \\ \vdots\end{matrix}}{e, (e +_b 1)^{n-1}, \Delta +_b \Gamma \Rightarrow_{\mathsf{At}\restriction b} \Delta} \qquad \dfrac{\dfrac{\begin{matrix}\dfrac{}{\Delta +_b \Gamma \Rightarrow_{\mathsf{At}\restriction \bar{b}} \Delta +_b \Gamma}\ \mathsf{id_s} \\ \vdots\end{matrix}}{(e +_b 1)^{n-1}, \Delta +_b \Gamma \Rightarrow_{\mathsf{At}\restriction \bar{b}} \Delta +_b \Gamma}\ +_b\text{-}l}{1, (e +_b 1)^{n-1}, \Delta +_b \Gamma \Rightarrow_{\mathsf{At}\restriction \bar{b}} \Delta +_b \Gamma}\ 1\text{-}l}{(e +_b 1)^n, \Delta +_b \Gamma \Rightarrow_{\mathsf{At}} \Delta +_b \Gamma}\ +_b\text{-}l$$

where the proof $\pi_1[(e +_b 1)^{n-1}, (\Delta +_b \Gamma)/e^{(b)}, \Gamma]$ is constructed as follows. We first replace in π_1 every final segment $e^{(b)}, \Gamma$ of an antecedent by $(e +_b 1)^{n-1}, (\Delta +_b \Gamma)$. If $e^{(b)}, \Gamma \Rightarrow_A \Delta'$ is the conclusion of an application of (b)-l in π of unfolding depth $m < n$, the proof continues as follows:

$$\dfrac{\begin{array}{c}\pi_1[(e +_b 1)^{n-m-1}, \Delta +_b \Gamma)/e^{(b)}, \Gamma] \\ \vdots \\ \hline e, (e +_b 1)^{n-m-1}, (\Delta +_b \Gamma) \Rightarrow_{A \restriction b} \Delta \end{array} \quad \begin{array}{c} \pi_2' \\ \dfrac{\Gamma \Rightarrow_{A \restriction \bar{b}} \Delta'}{\Delta +_b \Gamma \Rightarrow_{A \restriction \bar{b}} \Delta'}\ {+_b\text{-}l} \\ \vdots \\ \hline (e +_b 1)^{n-m-1}, \Delta +_b \Gamma \Rightarrow_{A \restriction \bar{b}} \Delta' \quad {+_b\text{-}l} \\ \hline 1, (e +_b 1)^{n-m-1}, \Delta +_b \Gamma \Rightarrow_{A \restriction \bar{b}} \Delta' \quad 1\text{-}l \end{array}}{(e +_b 1)^{n-m}, \Delta +_b \Gamma \Rightarrow_A \Delta'}\ {+_b\text{-}l}$$

where π_2' is the subproof of $\Gamma \Rightarrow_{A \restriction \bar{b}} \Delta'$ contained in π_1. To see that such a π_2' exists, note that by construction π_1 contains a subtree of the form

$$\dfrac{\begin{array}{c}\pi_1' \\ e, e^{(b)}, \Gamma \Rightarrow_{A \restriction b} \Delta'\end{array} \qquad \begin{array}{c}\pi_2' \\ \Gamma \Rightarrow_{A \restriction \bar{b}} \Delta'\end{array}}{e^{(b)}, \Gamma \Rightarrow_A \Delta'}\ (b)\text{-}l$$

such that every ancestor of the root of this subtree in π_1 has $e^{(b)}, \Gamma$ as a final segment. The fact that π_2' is a *subproof* and not a mere *subderivation* follows from the monotonicity of π.

If $e^{(b)}, \Gamma \Rightarrow_A \Delta'$ is the conclusion of an application of (b)-l in π of unfolding depth n, then by assumption (b) we have $A = \mathsf{At}$ and $\Delta' = \Delta$. We then finish the proof as follows:

$$\dfrac{\mathsf{id_s}\,\dfrac{}{\Delta \Rightarrow_{\mathsf{At} \restriction b} \Delta} \qquad \begin{array}{c}\pi_2 \\ \Gamma \Rightarrow_{\mathsf{At} \restriction \bar{b}} \Delta\end{array}}{\Delta +_b \Gamma \Rightarrow_{\mathsf{At}} \Delta}\ {+_b\text{-}l}$$

Having constructed the cyclic $\mathsf{SGKAT_s}$-proof (π', f'), it remains to show that it is susceptible to the induction hypothesis. We claim that $\mathsf{wln}(\pi') < \mathsf{wln}(\pi)$, *i.e.* that π' contains strictly less applications of the rule (c)-l (where c can be any test). This can be seen by noting that every application of (c)-l in π' has a unique corresponding application of (c)-l in π, but the first application of (b)-l in π does not appear in π'. $\square$

6.6.18. Remark. Just like the proof of Proposition 6.6.13, the above proof only uses a trivial instance of the invariant lemma. In particular, the invariant is $\Delta +_b \Gamma$. It is a topic of future work to see whether a generalisation of Proposition 6.6.17 can be obtained by harnessing the full power of the invariant lemma (cf. Remark 6.6.14).

6.6.19. Example. The proof Π_1 of Example 6.2.8 satisfies both assumptions of Proposition 6.6.17. Indeed, for the single companion it suffices to take $n = 2$. Hence we obtain $\mathsf{PoGKAT} \vdash (cp)^{(b)} \Rightarrow_{\mathsf{At}} (p(cp +_b 1))^{(b)}$.

Unfortunately, it is not hard to find examples of cyclic $\mathsf{SGKAT_s}$-proofs failing to satisfy assumption (a) or assumption (b'). For instance, the proof Π_3 (as well as Π_4) of Example 6.6.8 fails to satisfy assumption (a). One strategy for eliminating assumption (a) would be to prove a more sophisticated version of the Invariant Lemma (Lemma 6.6.12) that works for sets of atoms other than At.

6.7 Conclusion

In this chapter we have presented a non-well-founded proof system SGKAT^∞ for GKAT. Our system is similar to the system for Kleene Algebra in [34], but the deterministic nature of GKAT allows us to use ordinary sequents rather than hypersequents. To deal with the tests of GKAT every sequent is annotated by a set of atoms. We proved soundness and regular completeness with respect to the language model.

We proposed an algebraic inequational counterpart to our system, called PoGKAT, based on the equational system in [98]. We have presented a partial translation of cyclic SGKAT proofs into PoGKAT-proofs. This may be a first step towards proving that PoGKAT is complete with respect to the language model, which to the best of our knowledge is an open question.

There are many interesting questions left to explore. Perhaps the most pressing question is whether our partial translation into PoGKAT can be completed. A first step would be to try to further adapt the method in [33], which, as mentioned in Remark 6.6.14 comes with some difficulties. Even if the method from [33] can be adapted, a challenge remains in eliminating assumption (a) from the propositions 6.6.13 and 6.6.17 above, perhaps by employing a more sophisticated version of Lemma 6.6.12.

Very recently a completeness result for a certain fragment of GKAT was obtained by Kappé, Schmid & Silva [94]. This fragment, called *skip-free* GKAT, omits programs which can accept immediately without performing any action. In particular, every skip-free GKAT-expression is strictly productive (cf. Section 6.1.3). The completeness proof in [94] works by reducing skip-free GKAT to another formal system, of *1-free star expressions modulo bisimulation*, which was recently shown to be complete by Grabmayer and Fokkink in [48]. As an intermediate step, the authors of [94] show completeness for an axiomatisation of skip-free GKAT with respect to its so-called *bisimulation semantics*. Characteristic of the bisimulation semantics is that it does not satisfy an *early termination* axiom of the form $x \cdot 0 \equiv 0$. As mentioned in Remark 6.2.7, we conjecture that removing the axiom k_0 from SGKAT^∞ will make it sound and complete with respect to the bisimulation semantics.

Another interesting question is to determine the optimal complexity for proof-search in SGKAT^∞. Proof-search for the cyclic system for Kleene Algebra in [34] is in PSPACE, which is optimal since Kleene Algebra is PSPACE-complete. Because the decision problem for GKAT-equations is of very low complexity, the optimal decision procedure for GKAT-inequations is expected to be very efficient as well. Hence, there might not be proof-search procedure for SGKAT^∞ which is optimal for deciding GKAT-inequations. Nevertheless, we wonder whether a proof-search procedure exists that is at least more efficient than PSPACE.

Bibliography

[1] Bahareh Afshari, Sebastian Enqvist, and Graham E. Leigh. Cyclic proofs for the first-order μ-calculus. *Logic Journal of the IGPL*, 2022.

[2] Bahareh Afshari, Sebastian Enqvist, Graham E. Leigh, Johannes Marti, and Yde Venema. Proof systems for two-way modal mu-calculus, 2023. ILLC Prepublication Series. PP-2023-03. To appear in *Journal of Symbolic Logic*.

[3] Bahareh Afshari and Graham E. Leigh. Finitary proof systems for Kozen's μ. 2016. Technical Report, Mathematisches Forschungsinstitut Oberwolfach.

[4] Bahareh Afshari and Graham E. Leigh. Cut-free completeness for modal mu-calculus. In *32nd Annual Symposium on Logic in Computer Science, LICS*, pages 1–12. IEEE, 2017.

[5] Bahareh Afshari, Graham E. Leigh, and Guillermo Menéndez Turata. A cyclic proof system for full computation tree logic. In *31st Annual Conference on Computer Science Logic, CSL*, volume 252 of *LIPIcs*, pages 5:1–5:19. Schloss Dagstuhl, 2023.

[6] Bahareh Afshari and Dominik Wehr. Abstract cyclic proofs. In *28th International Workshop on Logic, Language, Information, and Computation, WoLLIC*, volume 13468 of *Lecture Notes in Computer Science*, pages 309–325. Springer, 2022.

[7] Luca Alberucci and Gerhard Jäger. About cut elimination for logics of common knowledge. *Annals of Pure and Applied Logic*, 133(1-3):73–99, 2005.

[8] David Baelde, Amina Doumane, Denis Kuperberg, and Alexis Saurin. Bouncing threads for circular and non-wellfounded proofs: Towards compositionality with circular proofs. In *37th Annual Symposium on Logic in Computer Science, LICS*, pages 63:1–63:13. IEEE, 2022.

[9] Alexandru Baltag, Nick Bezhanishvili, and David Fernández-Duque. The topological mu-calculus: completeness and decidability. In *36th Annual Symposium on Logic in Computer Science, LICS*, pages 1–13. IEEE, 2021.

[10] Johan van Benthem. *Modal correspondence theory.* PhD thesis, Universiteit van Amsterdam, Instituut voor Logica en Grondslagenonderzoek van de Exacte Wetenschappen, 1976.

[11] Johan van Benthem. Modal frame correspondences and fixed-points. *Studia Logica*, 83(1-3):133–155, 2006.

[12] Johan van Benthem and Nick Bezhanishvili. Modern faces of filtration, 2019. ILLC Prepublication PP-2019-13.

[13] Stefano Berardi and Makoto Tatsuta. Equivalence of inductive definitions and cyclic proofs under arithmetic. In *32nd Annual Symposium on Logic in Computer Science, LICS*, pages 1–12. IEEE, 2017.

[14] Nick Bezhanishvili and Silvio Ghilardi. The bounded proof property via step algebras and step frames. *Annals of Pure and Applied Logic*, 165(12):1832–1863, 2014.

[15] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press, 2001.

[16] Stephen L. Bloom. Varieties of ordered algebras. *Journal of Computer and System Sciences*, 13(2):200–212, 1976.

[17] Manfred Borzechowski. Tableau-kalkül für PDL und interpolation. Diplomarbeit, Freie Universität Berlin, 1988.

[18] James Brotherston. Cyclic proofs for first-order logic with inductive definitions. In *14th International Conference on Automated Reasoning with Analytic Tableaux and Related Methods, TABLEAUX*, volume 3702 of *Lecture Notes in Computer Science*, pages 78–92. Springer, 2005.

[19] James Brotherston. *Sequent calculus proof systems for inductive definitions.* PhD thesis, University of Edinburgh, 2006.

[20] Kai Brünnler and Lutz Straßburger. Modular sequent systems for modal logic. In *18th International Conference on Automated Reasoning with Analytic Tableaux and Related Methods, TABLEAUX*, volume 5607 of *Lecture Notes in Computer Science*, pages 152–166. Springer, 2009.

[21] Stanley Burris and Hanamantagouda P. Sankappanavar. *A course in universal algebra*, volume 78 of *Graduate Texts in Mathematics*. Springer, 1981.

[22] Samuel R. Buss. An introduction to proof theory. In *Handbook of proof theory*, volume 137 of *Studies in Logic and The Foundations of Mathematics*, pages 1–78. Elsevier, 1998.

[23] Facundo Carreiro. *Fragments of fixpoint logics*. PhD thesis, University of Amsterdam, 2015.

[24] Facundo Carreiro, Alessandro Facchini, Yde Venema, and Fabio Zanasi. The power of the weak. *ACM Transactions on Computational Logic*, 21(2):15:1–15:47, 2020.

[25] Facundo Carreiro and Yde Venema. PDL inside the μ-calculus: A syntactic and an automata-theoretic characterization. In *10th Conference on Advances in Modal Logic, AiML*, pages 74–93. College Publications, 2014.

[26] Alexander V. Chagrov and Michael Zakharyaschev. *Modal Logic*, volume 35 of *Oxford logic guides*. Oxford University Press, 1997.

[27] Agata Ciabattoni, Timo Lang, and Revantha Ramanayake. A theory of cut-restriction: first steps, 2022. arXiv preprint 2203.01600.

[28] John Horton Conway. *Regular algebra and finite machines*. Chapman and Hall, 1971.

[29] Giovanna D'Agostino. Uniform interpolation, bisimulation quantifiers, and fixed points. In *6th International Tbilisi Symposium on Logic, Language, and Computation*, volume 4363 of *Lecture Notes in Computer Science*, pages 96–116. Springer, 2005.

[30] Giovanna D'Agostino and Marco Hollenberg. Uniform interpolation, automata and the modal μ-calculus. In *1st Conference on Advances in Modal Logic, AiML*, pages 73–84.

[31] Giovanna D'Agostino and Marco Hollenberg. Logical questions concerning the μ-calculus: Interpolation, Lyndon and Łoś-Tarski. *Journal of Symbolic Logic*, 65(1):310–332, 2000.

[32] Anupam Das. On the logical complexity of cyclic arithmetic. *Logical Methods in Computer Science*, 16(1), 2020.

[33] Anupam Das, Amina Doumane, and Damien Pous. Left-handed completeness for Kleene algebra, via cyclic proofs. In *22nd International Conference on Logic for Programming, Artificial Intelligence and Reasoning, LPAR*, volume 57 of *EPiC Series in Computing*, pages 271–289, 2018.

[34] Anupam Das and Damien Pous. A cut-free cyclic proof system for Kleene algebra. In *26th International Conference on Automated Reasoning with Analytic Tableaux and Related Methods, TABLEAUX*, volume 10501 of *Lecture Notes in Computer Science*, pages 261–277. Springer, 2017.

[35] Anupam Das and Damien Pous. Non-wellfounded proof theory for (Kleene+action)(algebras+lattices). In *27th Annual Conference on Computer Science Logic, CSL*, volume 119 of *LIPIcs*, pages 19:1–19:18. Schloss Dagstuhl, 2018.

[36] Christian Dax, Martin Hofmann, and Martin Lange. A proof system for the linear time μ-calculus. In *26th International Conference on Foundations of Software Technology and Theoretical Computer Science FSTTCS*, volume 4337 of *Lecture Notes in Computer Science*, pages 273–284. Springer, 2006.

[37] Simon Docherty and Reuben N. S. Rowe. A non-wellfounded, labelled proof system for propositional dynamic logic. In *28th International Conference on Automated Reasoning with Analytic Tableaux and Related Methods, TABLEAUX*, volume 11714 of *Lecture Notes in Computer Science*, pages 335–352. Springer, 2019.

[38] E. Allen Emerson and Charanjit S. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *32nd Annual Symposium on Foundations of Computer Science*, pages 368–377. IEEE, 1991.

[39] Sebastian Enqvist. A circular proof system for the hybrid μ-calculus. In *13th Conference on Advances in Modal Logic, AiML*, pages 169–188. College Publications, 2020.

[40] Michael J. Fischer and Richard E. Ladner. Propositional dynamic logic of regular programs. *Journal of Computer and System Sciences*, 18(2):194–211, 1979.

[41] Gaëlle Fontaine. Continuous fragment of the mu-calculus. In *22nd Annual Conference on Computer Science Logic, CSL*, volume 5213 of *Lecture Notes in Computer Science*, pages 139–153. Springer, 2008.

[42] Gaëlle Fontaine and Yde Venema. Some model theory for the modal μ-calculus: syntactic characterisations of semantic properties. *Logical Methods in Computer Science*, 14(1), 2018.

[43] Tim French. Bisimulation quantified modal logics: Decidability. In *6th Conference on Advances in Modal Logic, AiML*, pages 147–166. College Publications, 2006.

[44] Tim French. Idempotent transductions for modal logics. In *6th International Symposium on Frontiers of Combining Systems, FroCoS*, volume 4720 of *Lecture Notes in Computer Science*, pages 178–192. Springer, 2007.

[45] Robert Goldblatt. *Logics of time and computation*. CSLI Lecture Notes. University of Chicago Press, 1987.

[46] Rajeev Goré. Tableau methods for modal and temporal logics. In *Handbook of Tableau Methods*, pages 297–396. Springer, 1999.

[47] Rajeev Goré and Florian Widmann. An optimal on-the-fly tableau-based decision procedure for PDL-satisfiability. In *22nd International Conference on Automated Deduction, CADE*, volume 5663 of *Lecture Notes in Computer Science*, pages 437–452. Springer, 2009.

[48] Clemens Grabmayer and Wan Fokkink. A complete proof system for 1-free regular expressions modulo bisimilarity. In *35th Annual Symposium on Logic in Computer Science, LICS*, pages 465–478. IEEE, 2020.

[49] Erich Grädel, Wolfgang Thomas, and Thomas Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, volume 2500 of *Lecture Notes in Computer Science*. Springer, 2002.

[50] Joseph Y. Halpern and Yoram Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54(2):319–379, 1992.

[51] Edith Hemaspaandra. The price of universality. *Notre Dame Journal of Formal Logic*, 37(2):174–203, 1996.

[52] Andrzej Indrzejczak. *Sequents and Trees: An Introduction to the Theory and Applications of Propositional Sequent Calculi*. Studies in Universal Logic. Springer Nature, 2020.

[53] David Janin and Igor Walukiewicz. On the expressive completeness of the propositional mu-calculus with respect to monadic second order logic. In *7th International Conference on Concurrency Theory, CONCUR*, volume 1119 of *Lecture Notes in Computer Science*, pages 263–277. Springer, 1996.

[54] Natthapong Jungteerapanich. A tableau system for the modal μ-calculus. In *18th International Conference on Automated Reasoning with Analytic Tableaux and Related Methods, TABLEAUX*, volume 5607 of *Lecture Notes in Computer Science*, pages 220–234. Springer, 2009.

[55] Ryo Kashima. Cut-free sequent calculi for some tense logics. *Studia Logica*, 53(1):119–136, 1994.

[56] Stanislav Kikot, Ilya Shapirovsky, and Evgeny Zolin. Modal logics with transitive closure: Completeness, decidability, filtration. In *13th Conference on Advances in Modal Logic, AiML*, pages 369–388. College Publications, 2020.

[57] Stephen Cole Kleene. Representation of events in nerve nets and finite automata. In *Automata Studies*, number 34 in Annals of Mathematics Studies, pages 3–41. Princeton University Press, 1956.

[58] Tomasz Kowalski and Hiroakira Ono. Analytic cut and interpolation for bi-intuitionistic logic. *Review of Symbolic Logic*, 10(2):259–283, 2017.

[59] Dexter Kozen. Results on the propositional mu-calculus. *Theoretical Computer Science*, 27:333–354, 1983.

[60] Dexter Kozen. A completeness theorem for Kleene algebras and the algebra of regular events. *Information and Computation*, 110(2):366–390, 1994.

[61] Dexter Kozen. Kleene algebra with tests. *ACM Transactions on Programming Languages and Systems*, 19(3):427–443, 1997.

[62] Dexter Kozen and Rohit Parikh. An elementary proof of the completeness of PDL. *Theoretical Computer Science*, 14:113–118, 1981.

[63] Denis Kuperberg, Laureline Pinault, and Damien Pous. Cyclic proofs, system T, and the power of contraction. *48th Annual Symposium on Principles of Programming Languages, POPL*, pages 1–28, 2021.

[64] Clemens Kupke, Johannes Marti, and Yde Venema. Size matters in the modal μ-calculus, 2020. arXiv preprint 2010.14430.

[65] Alexander Kurz and Jiří Velebil. Quasivarieties and varieties of ordered algebras: regularity and exactness. *Mathematical Structures in Computer Science*, 27(7):1153–1194, 2017.

[66] Ori Lahav. From frame properties to hypersequent rules in modal logics. In *28th Annual Symposium on Logic in Computer Science, LICS*, pages 408–417. IEEE, 2013.

[67] M. Lange. Satisfiability and completeness of Converse-PDL replayed. In *26th Annual German Conference on AI, KI*, volume 2821 of *Lecture Notes in Computer Science*, pages 79–92. Springer, 2003.

[68] Martin Lange and Colin Stirling. Focus games for satisfiability and completeness of temporal logic. In *16th Annual Symposium on Logic in Computer Science, LICS*, pages 357–365. IEEE Computer Society, 2001.

[69] Graham E. Leigh and Dominik Wehr. From GTC to reset: Generating reset proof systems from cyclic proof systems, 2023. arXiv preprint 2301.07544.

[70] Björn Lellmann and Dirk Pattinson. Correspondence between modal Hilbert axioms and sequent rules with an application to S5. In *26th International Conference on Automated Reasoning with Analytic Tableaux and Related Methods, TABLEAUX*, volume 8123 of *Lecture Notes in Computer Science*, pages 219–233. Springer, 2013.

[71] John Lemmon and Dana Scott. *An introduction to modal logic*. American Philosophical Quarterly Monograph Series. Basil Blackwell, 1977.

[72] Larisa Maksimova. Temporal logics of "the next" do not have the Beth property. *Journal of Applied Non-Classical Logics*, 1(1):73–76, 1999.

[73] Johannes Marti and Yde Venema. A focus system for the alternation-free μ-calculus. In *30th International Conference on Automated Reasoning with Analytic Tableaux and Related Methods, TABLEAUX*, volume 12842 of *Lecture Notes in Computer Science*, pages 371–388. Springer, 2021.

[74] A.M. Mostowski. Games with forbidden positions. Technical report, Instytut Matematyki, Uniwersytet Gdański, Poland, 1991.

[75] Sara Negri. Proof analysis in modal logic. *Journal of Philosophical Logic*, 34:507–544, 2005.

[76] Linh Anh Nguyen. Analytic tableau systems and interpolation for the modal logics KB, KDB, K5, KD5. *Studia Logica*, 69(1):41–57, 2001.

[77] Damian Niwinski and Igor Walukiewicz. Games for the mu-calculus. *Theoretical Computer Science*, 163:99–116, 1996.

[78] Rémi Nollet, Alexis Saurin, and Christine Tasson. PSPACE-completeness of a thread criterion for circular proofs in linear logic with least and greatest fixed points. In *28th International Conference on Automated Reasoning with Analytic Tableaux and Related Methods, TABLEAUX*, volume 11714 of *Lecture Notes in Computer Science*, pages 317–334. Springer, 2019.

[79] Masao Ohnishi and Kazuo Matsumoto. Gentzen method in modal calculi. *Osaka Mathematical Journal*, 9(2):113 – 130, 1957.

[80] Hiroakira Ono. Proof-theoretic methods in nonclassical logic–an introduction. In *Theories of Types and Proofs*, volume 2, pages 207–254. Mathematical Society of Japan, 1998.

[81] Leonardo Pacheco and Kazuyuki Tanaka. The alternation hierarchy of the μ-calculus over weakly transitive frames. In *28th International Workshop on Logic, Language, Information, and Computation, WoLLIC*, volume 13468 of *Lecture Notes in Computer Science*, pages 207–220. Springer, 2022.

[82] David Park. Finiteness is mu-ineffable. *Theoretical Computer Science*, 3:173–181, 1976.

[83] Francesca Poggiolesi. The method of tree-hypersequents for modal propositional logic. In *Towards Mathematical Philosophy*, volume 28 of *Trends in Logic*, pages 31–51. Springer, 2009.

[84] Francesca Poggiolesi and Brian Hill. Common knowledge: a finitary calculus with a syntactic cut-elimination procedure. *Logique et Analyse*, pages 279–306, 2015.

[85] Vaughan R. Pratt. A decidable mu-calculus: Preliminary report. In *22nd Annual Symposium on Foundations of Computer Science*, pages 421–427. IEEE, 1981.

[86] Jan Rooduijn. Cyclic hypersequent calculi for some modal logics with the master modality. In *30th International Conference on Automated Reasoning with Analytic Tableaux and Related Methods, TABLEAUX*, volume 12842 of *Lecture Notes in Computer Science*, pages 354–370. Springer, 2021.

[87] Jan Rooduijn and Yde Venema. Filtration and canonical completeness for continuous modal mu-calculi. In *12th International Symposium on Games, Automata, Logics, and Formal Verification, GandALF*, volume 346 of *EPTCS*, pages 211–226, 2021.

[88] Jan Rooduijn and Yde Venema. Focus-style proofs for the two-way alternation-free μ-calculus. In *29th International Workshop on Logic, Language, Information, and Computation, WoLLIC*, volume 13923 of *Lecture Notes in Computer Science*, pages 318–335. Springer, 2023.

[89] Jan Rooduijn and Lukas Zenger. An analytic proof system for common knowledge logic over S5. In *14th Conference on Advances in Modal Logic, AiML*, pages 659–680. College Publications, 2022.

[90] Henrik Sahlqvist. Completeness and correspondence in the first and second order semantics for modal logic. In *3rd Scandinavian Logic Symposium*, volume 82 of *Studies in Logic and the Foundations of Mathematics*, pages 110–143. Elsevier, 1975.

[91] Arto Salomaa. Two complete axiom systems for the algebra of regular events. *Journal of the ACM*, 13(1):158–169, 1966.

[92] Yury Savateev and Daniyar S. Shamkanov. Non-well-founded proofs for the Grzegorczyk modal logic. *Review of Symbolic Logic*, 14(1):22–50, 2021.

[93] Todd Schmid, Tobias Kappé, Dexter Kozen, and Alexandra Silva. Guarded Kleene algebra with tests: Coequations, coinduction, and completeness. In *48th International Colloquium on Automata, Languages, and Programming, ICALP*, volume 198 of *LIPIcs*, pages 142:1–142:14. Schloss Dagstuhl, 2021.

[94] Todd Schmid, Tobias Kappé, and Alexandra Silva. A complete inference system for skip-free guarded Kleene algebra with tests. In *32nd European Symposium on Programming, ESOP*, volume 13990 of *Lecture Notes in Computer Science*, pages 309–336. Springer, 2023.

[95] Karl Krister Segerberg. *An essay in classical modal logic.* PhD thesis, Stanford University, 1971.

[96] Daniyar S. Shamkanov. Circular proofs for the Gödel-Löb provability logic. *Mathematical Notes*, 96:575–585, 2014.

[97] Alex Simpson. Cyclic arithmetic is equivalent to Peano arithmetic. In *20th International Conference on Foundations of Software Science and Computation Structures, FOSSACS*, volume 10203 of *Lecture Notes in Computer Science*, pages 283–300. Springer, 2017.

[98] Steffen Smolka, Nate Foster, Justin Hsu, Tobias Kappé, Dexter Kozen, and Alexandra Silva. Guarded Kleene algebra with tests: verification of uninterpreted programs in nearly linear time. *47th Annual Symposium on Principles of Programming Languages, POPL*, pages 61:1–61:28, 2020.

[99] Christoph Sprenger and Mads Dam. On the structure of inductive reasoning: Circular and tree-shaped proofs in the μ-calculus. In *6th International Conference on Foundations of Software Science and Computation Structures, FOSSACS*, volume 2620 of *Lecture Notes in Computer Science*, pages 425–440. Springer, 2003.

[100] Colin Stirling. A tableau proof system with names for modal mu-calculus. In *HOWARD-60: A Festschrift on the Occasion of Howard Barringer's 60th Birthday*, volume 42 of *EPiC Series in Computing*, pages 306–318. 2014.

[101] Robert S. Streett and E. Allen Emerson. An automata theoretic decision procedure for the propositional mu-calculus. *Information and Computation*, 81(3):249–264, 1989.

[102] Thomas Studer. On the proof theory of the modal mu-calculus. *Studia Logica*, 89(3):343–363, 2008.

[103] Thomas Studer. Common knowledge does not have the Beth property. *Information Processing Letters*, 109(12):611–614, 2009.

[104] Alfred Tarski. A lattice-theoretical fixpoint theorem and its applications. *Pacific Journal of Mathematics*, 5(2):285 – 309, 1955.

[105] Moshe Y. Vardi. Reasoning about the past with two-way automata. In *25th International Colloquium on Automata, Languages and Programming, ICALP*, volume 1443 of *Lecture Notes in Computer Science*, pages 628–641. Springer, 1998.

[106] Yde Venema. Lectures on the modal μ-calculus. Lecture Notes, ILLC, University of Amsterdam. 2020.

[107] Igor Walukiewicz. Completeness of Kozen's axiomatisation of the propositional μ-calculus. *Information and Computation*, 157(1-2):142–182, 2000.

Index

Samenvatting

Dit proefschrift gaat over de bewijstheorie van modale dekpuntlogica's. In het bijzonder bevat het constructies van bewijssystemen voor verschillende fragmenten van de modale mu-calculus, geïnterpreteerd over verschillende klassen van frames. Dit proefschrift beoogt de relatief onderontwikkelde bewijstheorie van de modale mu-calculus dichter bij de gevestigde bewijstheorie van basismodale logica te brengen, met een nadruk op uniforme constructies en algemene resultaten. Twee benaderingen staan centraal. Ten eerste, het veralgemeniseren van bestaande methoden voor basismodale logica naar fragmenten van de modale mu-calculus. Deze methode wordt gebruikt om Hilbert-stijl bewijssystemen te ontwikkelen. Ten tweede, het aanpassen van bestaande methoden op het gebied van de modale mu-calculus zodat ze werken voor verschillende klassen van frames. Deze methode geeft bewijssystemen die *niet-welgefundeerd* of *cyclisch* zijn.

Hoofdstuk 1 bevat een informele introductie, waarin de doelen en resultaten van dit proefschrift worden besproken. In Hoofdstuk 2 wordt de noodzakelijke voorkennis geïntroduceerd, inclusief een fundamenteel hulpmiddel: pariteitsspelen.

In Hoofdstuk 3 ontwikkelen we cyclische hypersequentencalculi voor een relatief simpel fragment van de modale mu-calculus: modale logica met de mastermodaliteit. Voortbouwend op eerder werk van Ori Lahav voor basismodale logica [66], construeren we op uniforme wijze hypersequentencalcui voor verschillende, zogenaamd *simpele*, frameklassen. Snedevrije volledigheid wordt alleen bewezen voor bepaalde simpele frameklassen, die we *gelijkmatig* noemen.

In het Intermezzo dat volgt op Hoofdstuk 3 introduceren we een algemeen raamwerk om zogenaamde *pad-gebaseerde* niet-welgefundeerde bewijssystemen te bestuderen. In dit raamwerk geven we een voldoende voorwaarde om te laten zien dat elke bewijsbare sequent een *beknopt* bewijs heeft. In de meeste gevallen leidt dit tot de *begrensdebewijseigenschap*: elke bewijsbare sequent heeft een cyclisch bewijs waarvan de grootte begrensd wordt door een waarde berekenbaar vanuit de grootte van de sequent.

In Hoofdstuk 4 veralgemeniseren we bestaand onderzoek op het gebied van Hilbert-stijl bewijssystemen voor PDL. Recentelijk lieten Kikot, Shapirovsky & Zolin in [56] zien hoe het oorspronkelijke volledigheidsbewijs van Kozen & Parikh [62] voor PDL uitgebreid kan worden naar verschillende frameklassen die de methode van filtratie toelaten. Wij laten zien dat de continue modale μ-calculus, die strikt meer expressief is dan PDL, filtratie toelaat. Daarnaast veralgemeniseren we de resultaten van Kikot, Shapirovsky & Zolin, zodat ze toepasbaar zijn op de continue modale μ-calculus

In Hoofdstuk 5 beschouwen we de tweezijdige modale μ-calculus. Voortbouwend op eerder onderzoek van Vardi naar tweezijdige automaten [105], construeren we een cyclisch bewijsssyteem voor de alternatievrije tweezijdige modale μ-calculus. De kern van onze methode is het nieuwe concept van een *spooratoom*; een additioneel element binnen een sequent dat de mogelijke sporen door een bewijs bijhoudt.

In het laatste hoofdstuk, Hoofdstuk 6, wijken we lichtelijk af van het hoofdonderwerp door een dekpuntlogica te beschouwen die strikt genomen niet modaal is, namelijk Guarded Kleene Algebra met Tests. Dit is een computationeel efficiënte variant van de beter bekende Kleene Algbera met Tests. We construeren een cyclisch bewijssysteem voor Guarded Kleene Algebra met Tests, geïnspireerd op een eerder systeem ontwikkeld door Das & Pous voor Kleene Algebra [34]. Ook nemen we een eerste stap richting het vertalen van onze cyclische bewijzen naar een algebraïsch systeem voor Guarded Kleene Algebra met Tests.

Abstract

This thesis studies the proof theory of modal fixed point logics. In particular, we construct proof systems for various fragments of the modal μ-calculus, interpreted over various classes of frames. With an emphasis on uniform constructions and general results, we aim to bring the relatively underdeveloped proof theory of modal fixed point logics closer to the well-established proof theory of basic modal logic. We employ two main approaches. First, we seek to generalise existing methods for basic modal logic to accommodate fragments of the modal μ-calculus. We use this approach for obtaining Hilbert-style proof systems. Secondly, we adapt existing proof systems for the modal μ-calculus to various classes of frames. This approach yields proof systems which are *non-well-founded* or *cyclic*.

In Chapter 1 we give an informal introduction to the goals and results of this thesis. Chapter 2 introduces the necessary preliminaries, including the fundamental tool of infinite (parity) games.

In Chapter 3 we construct cyclic hypersequent calculi for a relatively simple fragment of the modal μ-calculus: modal logic with the master modality. Building upon prior work by Ori Lahav for basic modal logic [66], we uniformly construct hypersequent calculi for various, so-called *simple*, frame classes. We are only able to prove cut-free completeness for certain specific simple frame classes, which we call *equable*.

In the Intermezzo following Chapter 3, we introduce a general framework for studying so-called *path-based* non-well-founded proof systems. In this framework we establish a sufficient condition for showing that every provable sequent admits a *concise* proof. In most cases this leads to the *bounded proof property*: every provable sequent has a cyclic proof for which a size bound can be calculated from the size of the sequent.

In Chapter 4 we generalise existing work on Hilbert-style proof systems for `PDL`. Recently it was shown by Kikot, Shapirovsky & Zolin in [56] how to extend the original completeness proof by Kozen & Parikh [62] of `PDL` to several frame classes which admit the method of filtration. We show that the continuous modal

μ-calculus, which is strictly more expressive than PDL, admits filtration. Moreover, we generalise the results by Kikot, Shapirovsky & Zolin to the continuous modal μ-calculus.

In Chapter 5 we consider the two-way modal μ-calculus. Building on previous work by Vardi on two-way automata [105], we construct a cyclic proof system for the alternation-free two-way modal μ-calculus. At the heart of our method is the novel concept of a *trace atom*, an additional element within a sequent that records possible traces through a proof.

In the final chapter, Chapter 6, we slightly diverge from our main topic by considering a fixed point logic which is not strictly modal, known as Guarded Kleene Algebra with Tests. This logic is a computationally efficient fragment of the more well known Kleene Algebra with Tests. We construct a cyclic proof system for Guarded Kleene Algebra with Tests, inspired by an earlier system developed by Das & Pous for Kleene Algebra [34]. Furthermore, we take a first step towards translating our cyclic proofs into an algebraic system for Guarded Kleene Algebra with Tests.

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: **András Gilyén**
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 Möllerström 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