

Completeness for the modal μ -calculus:
separating the combinatorics from the dynamics

Sebastian Enqvist Fatemeh Seifan*
sebastian.enqvist@fil.lu.se fateme.sayfan@gmail.com

Yde Venema
y.venema@uva.nl

Institute for Logic, Language and Computation
University of Amsterdam
PO Box 94242, 1090 GE Amsterdam, The Netherlands

August 2016

*The research of this author has been made possible by *Vrije Competitie* grant 612.001.115 of the Netherlands Organisation for Scientific Research (NWO).

Abstract

The *modal mu-calculus* is a very expressive formalism extending basic modal logic with least and greatest fixpoint operators. In the seminal paper introducing the formalism in the shape known today, Dexter Kozen also proposed an elegant axiom system, and he proved a partial completeness result with respect to the Kripke-style semantics of the logic. The problem of proving Kozen’s axiom system complete for the full language remained open for about a decade, until it was finally resolved by Igor Walukiewicz. Walukiewicz’ proof is notoriously difficult however, and the result has remained somewhat isolated from the standard theory of completeness for modal (fixpoint) logics. Our aim in this paper is to develop a framework that will let us clarify and simplify parts of Walukiewicz’s proof. We hope that this will also help to facilitate future research into completeness of modal fixpoint logics, including fragments, variants and extensions of the modal mu-calculus.

Our main contribution is to take the automata-theoretic viewpoint, already implicit in Walukiewicz’s proof, much more seriously by bringing automata explicitly into the proof theory. Thus we further develop the theory of *modal parity automata* as a mathematical framework for proving results about the modal mu-calculus. Once the connection between automata and derivations is in place, large parts of the completeness proof can be reformulated as purely automata-theoretic theorems. From a conceptual viewpoint, our automata-theoretic approach lets us distinguish two key aspects of the mu-calculus: the *one-step dynamics* encoded by the modal operators, and the *combinatorics* involved in dealing with nested fixpoints. This “deconstruction” allows us to work with these two features in a largely independent manner.

More in detail, prominent roles in our proof are played by two classes of modal automata: next to the *disjunctive automata* that are known from the work of Janin & Walukiewicz, we introduce here the class of *semi-disjunctive* automata that roughly correspond to the fragment of the mu-calculus for which Kozen proved completeness. We will establish a connection between the proof theory of Kozen’s system, and two kinds of games involving modal automata: a *satisfiability game* involving a single modal automaton, and a *consequence game* relating two such automata. In the key observations on these games we bring the dynamics and combinatorics of parity automata together again, by proving some results that witness the nice behaviour of disjunctive and semi-disjunctive automata in these games. As our main result we prove that every formula of the modal mu-calculus *provably* implies the translation of a disjunctive automaton; from this the completeness of Kozen’s axiomatization is immediate.

Mathematics Subject Classification (MSC2010): 03B45; 03B70; 68Q60; 91A43.

Keywords: *modal mu-calculus; modal fixpoint logic; completeness; parity automata; infinite games*

1 Introduction

The modal μ -calculus The modal μ -calculus μ ML is an extension of basic modal logic with least and greatest fixpoint operators. In the shape known today, it was introduced by Dexter Kozen [36], building on earlier work by others, including de Bakker & Scott [3], Park [49], and Pratt [54]. Since then it has been the subject of extensive research, see [1, 27, 8] for some surveys.

The μ -calculus has its roots in computer science, where it serves to unify and provide a foundation for the various modal logics that have featured as specification languages for programs and reactive systems. For example, linear temporal logic LTL [52, 23], computation tree logic CTL [13] and its extension CTL* [19] can all be embedded into the μ -calculus (see [15] for the non-trivial case of CTL*), as can propositional dynamic logic PDL [53] and Parikh’s dynamic logic of games GL [48]. In this sense, the μ -calculus serves as a “universal” modal specification language.

Besides this important role in formal verification, it has become increasingly clear over the years that the modal μ -calculus also has a rich and beautiful meta-theory, and deserves a place in “pure” (mathematical) logic as well as in computer science. A paper that highlights this is the work by d’Agostino & Hollenberg [14], which shows that the μ -calculus enjoys several nice model-theoretic properties such as a Lyndon theorem and a Łos-Tarski preservation theorem. Among the results obtained, the most striking is perhaps the interpolation theorem, which shows that the μ -calculus has the very strong *uniform* interpolation property. This property fails for first-order logic, but has previously been shown for a select number of logics, most famously intuitionistic logic [51], but also basic modal logic [24, 70].

The modal μ -calculus is a natural extension of basic modal logic, and retains many of its good properties. For example, Kozen [37] showed that the finite model property is preserved and, improving on a series of earlier results, Emerson & Jutla [20] could pin down the complexity of the satisfiability problem for the μ -calculus as being EXPTIME-complete. While basic modal logic has a PSPACE-complete satisfiability problem, already minor extensions of it are EXPTIME-hard, including the logic obtained by adding the global modality [66]. Moreover, the μ -calculus retains what is often viewed as *the* defining property of modal logic, namely, its tight link with the notion of *bisimilarity*. For modal logic, this is highlighted by van Benthem’s celebrated characterization theorem [4], exhibiting modal logic as the bisimulation-invariant fragment of first-order logic. It turns out that the μ -calculus is also the bisimulation-invariant fragment of an important system, namely monadic second-order logic (MSO). In the context of applications in process theory, this result, due to Janin & Walukiewicz [31], can be seen as an expressive completeness theorem, stating that all “relevant” MSO-formulas can be expressed in the modal μ -calculus.

Completeness Besides the topics we have mentioned so far – expressive power, complexity, preservation theorems etc. – one of the first problems normally considered for any logical formalism is to provide a sound and complete deductive system of axioms and rules. Indeed this was among the very first problems raised about the modal μ -calculus. In his seminal paper, Kozen already suggested an axiomatization which is as simple and elegant as the μ -calculus itself: on top of the usual rules and axioms of the least normal modal logic **K**, Kozen’s

axiom system adds a single axiom schema and a rule schema to handle the fixpoint operators. Together, the new axiom and rule schemas express in a straightforward way the equivalent characterization of the least fixpoint, well known from the Knaster-Tarski theorem, as the *least pre-fixpoint*. The axiom captures the pre-fixpoint property:

$$\varphi[\mu x.\varphi/x] \rightarrow \mu x.\varphi.$$

The rule schema, which is sometimes referred to as Park’s *induction rule*, expresses in an equally simple way that $\mu x.\varphi$ is indeed the *least* pre-fixpoint:

$$\frac{\varphi[\gamma/x] \rightarrow \gamma}{\mu x.\varphi \rightarrow \gamma}$$

The problem of proving the completeness of this axiom system turned out to be rather hard: Kozen presented a proof only for a fragment of the μ -calculus, which he called the *aconjunctive* fragment, and the completeness problem for the full language remained open for more than a decade until it was finally solved by Walukiewicz in [72] — in the sequel we shall refer to the extended journal publication [73].

Theorem 1 (Kozen-Walukiewicz) *Kozen’s deductive system provides a sound and complete axiomatization for the set of valid formulas of the modal μ -calculus.*

Walukiewicz’ proof is widely considered to be very hard to understand, and while the Kozen-Walukiewicz completeness theorem is often cited and generally recognized as a landmark in the theory of the modal μ -calculus, it has remained something of an isolated point in the completeness theory of modal (fixpoint) logic. Modal logic has a well established theoretical framework for completeness theory [6], with many standardized techniques such as canonical models and filtration, and sweeping general results like Sahlqvist’s theorem. The completeness theorem for the modal μ -calculus, on the other hand, does not seem to have given rise to any comparable follow-up literature (some exceptions to this rule will be discussed in a moment).

If we try to diagnose this situation, two obvious differences between the μ -calculus and the modal logics that are covered by the general completeness theory come to mind. First, the μ -calculus lacks the *compactness property*, which is tightly related to the Stone-like duality between modal algebras and (descriptive) frames that underlies much of modal completeness theory [32, 26]. This obviously prevents the possibility of having a strongly complete and finitary system of axioms, but besides this the lack of compactness does not appear to be the heart of the matter: for example, the non-compact logic PDL has an elegant weak completeness proof using essentially a finitary version of the canonical model construction [38]. Likewise, Emerson & Halpern’s completeness proof for CTL is relatively straightforward [18], and Lange & Stirling [40] present a game-theoretic framework for establishing completeness results for simple modal fixpoint logics such as LTL and CTL. For certain fragments of the μ -calculus (Santocanale & Venema’s “flat fixpoint logics” [62]) a generic completeness proof is known that stays much closer to familiar algebraic techniques in modal logic than Walukiewicz’ proof for the full μ -calculus.

Rather, what is special about the modal μ -calculus are the *combinatorial* issues that arise with simultaneous fixpoints, especially in the presence of *alternation* between least and greatest fixpoints (i.e., mutual dependencies between least and greatest fixpoint operators). There is some evidence that this is indeed the crucial threshold: PDL, CTL, and flat fixpoint logics all sit inside the alternation free fragment. Also quite telling is that for the dual-free fragment of Parikh’s game logic, completeness is proved by a gentle modification of the proof for PDL [48], while for the full logic GL (which spans the whole alternation hierarchy, see [5]), finding a complete system of axioms remains an intriguing open problem. In addition, the “exceptions” that we just mentioned, that is, the few papers that do address completeness questions for modal logics with *explicit* fixpoint operators and do make a link with the Kozen-Walukiewicz result, seem to support our claim. For instance, all papers that report positive results based on (relatively) easy proofs, seem to focus on settings where the alternation hierarchy collapses: ten Cate & Fontaine [2] prove a completeness result for the set of μ ML-validities on the class of finite trees (where nontrivial least and greatest fixpoints coincide), while Kaivola [34] proves completeness for the *linear time* modal μ -calculus, that is, the modal μ -calculus interpreted on the structure of the natural numbers with the successor relation (where the alternation hierarchy collapses at the alternation-free level). Doumane & coauthors [17] obtain stronger positive results for this linear-time interpretation of μ ML, but only for fragments of the language. Similarly, Santocanale’s analysis of Kozen’s axiomatization in [61] indicates that familiar algebraic methods can only prove completeness for the alternation-free fragment of the language.

It is also worth mentioning that the only known completeness proof for CTL^{*}, due to Reynolds [57], is quite complex. The logic CTL^{*} can be seen as a fragment of the μ -calculus, and as such it is not alternation-free. However, it is less clear how the case of CTL^{*} fits into our description. Rather than a fragment of the μ -calculus, it might be more natural to think of CTL^{*} as a sort of combination of LTL with **S5** modal logic, and furthermore the difficulties that Reynolds deals with in his proof do not seem to be primarily concerned with the combinatorics of simultaneous fixpoints. Similarly, the completeness proof for ECTL^{*}, due to Kaivola [35], seems to be based on regarding this extension of CTL^{*} as a combination of **S5** with the linear μ -calculus, rather than as a fragment of μ ML.

Finally, in passing we note that in this paper we focus on finitary proof systems. If one is happy to work with infinitary proof systems (or finitary systems derived from these by an appeal to the small model property of the modal μ -calculus), then completeness for the full language of μ ML can be obtained in more direct ways than by the Kozen-Walukiewicz proof, see for instance the work of Kozen [37], Walukiewicz [71], or Jäger, Kretz & Studer [29].

Logic and Automata A mathematical framework for the modal μ -calculus that is tailor suited precisely to deal with the combinatorics of fixpoint alternation is the theory of finite automata. This places the μ -calculus in a long tradition connecting logic and automata theory, going back to the seminal work of Büchi, Rabin and others. As two landmark results in this tradition we mention Büchi’s result showing that finite automata and monadic second-order logic have the same expressive power over infinite words [9], and Rabins’ decidability theorem [55] for the monadic second-order theory S_nS .

More specifically, in the context of fixpoint logics, the natural and most frequently used

type of automata are the *parity automata*, independently introduced by Mostowski [45] and Emerson & Jutla [20]. And indeed, most of the deep results on the modal μ -calculus have used parity automata in one way or another; in particular, Walukiewicz’s completeness proof heavily uses automata-theoretic ideas and insights. Specifically, the proof proceeds from the observation that the satisfaction problem is easy for formulas of a certain normal form, called *disjunctive* formulas, that correspond to the μ -automata (here called disjunctive automata) introduced by Janin & Walukiewicz in [30]. The strategy of Walukiewicz’s proof is thus to prove that every formula φ of the μ -calculus can be rewritten as a semantically equivalent disjunctive formula $\hat{\varphi}$, such that the implication $\varphi \rightarrow \hat{\varphi}$ is *provable* in Kozen’s system. In a sense, this can be seen as re-establishing the equivalence of μ -calculus formulas and μ -automata as a *proof-theoretic* result.

Our aim Our main goal in this paper is to streamline, clarify and, where possible, simplify the proof of the Kozen-Walukiewicz completeness theorem for the modal μ -calculus, by exhibiting and further developing the key mathematical concepts underlying the proof. In particular, we set up a *framework* for dealing with the completeness problem, in the hope that this will help to facilitate future research into completeness of modal fixpoint logics, including fragments, variants and extensions of the modal μ -calculus. In addition, our approach leads to a number of new automata-theoretic concepts and results that we believe to be of independent interest. In the remainder of this introduction we outline some of our main ideas, novel concepts and technical results.

Automata, coalgebra and proof theory The main conceptual difference with the approach taken by both Kozen and Walukiewicz is that automata feature far more prominently in our proof. That is, one of the main novelties in our approach is that we make an explicit and mathematically precise connection between the proof theory for the μ -calculus and the theory of modal parity automata. Following this idea, we may rework large parts of Kozen’s and Walukiewicz’s arguments in an entirely automata-theoretic framework, where we may clearly distinguish what we take to be the two main parallel aspects of the completeness proof: the *combinatorics* involved in reasoning with fixpoints, and the *dynamics* encoded in the semantics of the modal operators. The combinatorics is dealt with by a purely combinatorial framework that we will call “trace theory”. The dynamics is understood by studying the so-called “one-step logic”, a concept originating from the theory of coalgebraic logic.

In fact, much of our approach here has been inspired by coalgebraic ideas. Universal coalgebra [58] is a categorical theory that has been developed to provide a uniform mathematical framework for point-based evolving systems, such as, indeed, traces (streams) and Kripke models. In particular, much of the theory of automata operating on infinite structures is essentially coalgebraic in nature, in the sense that many key results have a natural generalization at a coalgebraic level of abstraction [39]. And with Kripke structures providing key examples of coalgebras, the same applies to modal logic, which has been recognized as the natural branch of logic for coalgebras [11]. The power of coalgebra lies in its combination of mathematical simplicity with broad applicability, and the approach we take here can be generalized to a much wider setting (we will come back to this below). Since we want to focus on the classical case of the μ -calculus for standard Kripke structures here, and in order to

suppress the use of categorical methods, we have chosen to keep the coalgebraic perspective implicit, with three exceptions.

First of all, it will often be convenient for us to work with a coalgebraic presentation of a Kripke structure \mathbb{S} , viz., as a function $\sigma_{\mathbb{S}}$ mapping each state s to a small window on the Kripke structure, consisting of the pair, formed by the set of proposition letters that s satisfies, together with the collection of its (immediate) successors. Such “one-step unfoldings” can then be examined by means of formulas of the one-step language that forms the domain of the transition map of our automata — this is what the game-theoretic semantics of our automata will be based on. Starting from this, we will investigate the *one-step logic* of Kripke structures, involving notions like one-step equivalence and one-step completeness. One-step logic, stemming from the work on coalgebraic logic by Cirstea, Pattinson, Schröder and others [12, 50, 64, 65], is then the second important tool that we will take from coalgebra. And third, we will make extensive use of the coalgebraic *cover* modality ∇ ; this modality, which was introduced by Janin & Walukiewicz [30] as the natural modality of their μ -automata, (and which features also prominently in Walukiewicz’s completeness proof), was independently introduced in the context of coalgebraic logic by Moss [43].

Returning to the global picture of our proof, on the one hand its “deconstruction” allows us to deal with the combinatorial and the dynamic concepts in largely separate frameworks. On the other hand, the use of modal parity automata will allow us to combine these two features, to understand where and how the two perspectives interact, and how they connect to each other. In particular, we will see that the trace theory of an automaton is largely determined by the shape of the formulas of the one-step language.

Technically, the way we achieve this is to work with the wider class of *modal automata* (introduced under the name of “alternating automata” by Wilke [74]), rather than passing directly to the μ -automata used by Walukiewicz. One might say that we introduce automata into the picture at an *earlier stage*: as mentioned, the main goal in Walukiewicz’s proof strategy is to show that every formula of the μ -calculus proves some formula that corresponds to a syntactic representation of a disjunctive automaton. By working with arbitrary modal automata, we can prove an analogous result by much more elementary techniques: every formula is provably equivalent to a formula in a normal form, that is the syntactic representation of some *modal* automaton. Formally we provide a modal automaton \mathbb{A}_{φ} for each formula φ and a formula $\text{tr}(\mathbb{A})$ for each modal automaton \mathbb{A} , and prove the following proposition (with \equiv_K denoting provable equivalence with respect to Kozen’s axiomatization):

Theorem 2 *For every formula $\varphi \in \mu\text{ML}$, we have $\varphi \equiv_K \text{tr}(\mathbb{A}_{\varphi})$.*

Technically, Theorem 2 is not a very deep result, but we see it as an important *conceptual* contribution of our approach. Bringing automata into the proof theory, so to say, it enables us to apply proof-theoretic notions such as derivability and consistency to automata, and thus it takes us “half-way” towards Walukiewicz’s result, where the remainder of the distance can now be addressed by wholly automata-theoretic methods. In some sense then, Theorem 2 answers a question raised by Lange & Stirling [40], who wrote “When proving completeness, one needs to establish that a finite consistent set of formulas is satisfiable. It is not known, in general, how to plug into such a proof automata theoretic constructions (such as product and determinisation) for satisfiability”.

Games and special automata The main tools that we employ in our automata-theoretic approach towards Kozen’s deductive system are two kinds of games for modal automata: the satisfiability game and the consequence game, and two special kinds of modal automata: next to the disjunctive automata, the class of semi-disjunctive automata.

The *satisfiability game* $\mathcal{S}(\mathbb{A})$ related to a modal automaton \mathbb{A} , was introduced in [22] in the more general setting of the coalgebraic μ -calculus. It is an infinite two-player game, that can be seen as a streamlined, game-theoretic analog for automata to what tableaux are for formulas. In this game, the *dynamics* of the semantics appears in the moves of the player \exists (Éloise) who has the role of “model builder”, and attempts to construct a satisfying model one layer at a time, while constrained by the one-step transition structure of the automaton.

The *combinatorics* of the trace theory enters the picture through the winning condition for infinite matches. As we shall see, each infinite match naturally induces a *trace graph*, an intricate graph structure of which the finite and infinite paths correspond to \mathbb{A} -*traces*: finite and infinite sequences of states of the automaton \mathbb{A} . The winning condition of $\mathcal{S}(\mathbb{A})$ states that, for \exists to win the infinite match, all infinite traces, corresponding to full branches through this graph, need to satisfy the acceptance condition of \mathbb{A} . Intuitively then, the smaller and simpler the trace graph, the easier it is for her to win. In particular, it will be to her advantage if we restrict the use of conjunctions in the one-step language, since these correspond to branching in the trace graph.

As we shall prove, the satisfiability game is *adequate* in the sense that \exists has a winning strategy in $\mathcal{S}(\mathbb{A})$ iff the automaton \mathbb{A} is satisfiable, that is, has a non-empty language. Hence our overall approach towards the completeness proof will be to prove that for any *consistent* automaton \mathbb{A} , \exists has a winning strategy in the satisfiability game associated with \mathbb{A} . This makes that we will generally “take sides” with \exists in the satisfiability game.

Before moving on to the other game featuring in our proof, we mention that the game $\mathcal{S}(\mathbb{A})$ comes in two flavours: the standard and the *thin* satisfiability game $\mathcal{S}_{thin}(\mathbb{A})$. The two versions of the game have identical sets of positions for \exists and her opponent \forall (Abélard), the only difference being that in $\mathcal{S}_{thin}(\mathbb{A})$ we curtail the power of \forall by restricting his moves.

The *consequence game* $\mathcal{C}(\mathbb{A}, \mathbb{A}')$, an original contribution of this paper, is an infinite two-player game that can be seen as a kind of implication game between the two automata \mathbb{A} and \mathbb{A}' . Its moves revolve around one of the players, prosaically named “player II”, trying to establish some structural connection between the two automata to support the claim that \mathbb{A} implies \mathbb{A}' . We write $\mathbb{A} \models_{\mathcal{G}} \mathbb{A}'$ in case he succeeds, in the sense of having a winning strategy in the game $\mathcal{C}(\mathbb{A}, \mathbb{A}')$. This relation $\models_{\mathcal{G}}$ is a strong consequence relation between automata, indicating a close structural relation between the automata; for instance, we shall see that $\mathbb{A} \models_{\mathcal{G}} \mathbb{A}'$ implies that \mathbb{A}' is a semantic consequence of \mathbb{A} , but not vice versa.

The consequence game $\mathcal{C}(\mathbb{A}, \mathbb{A}')$ is tightly connected to the satisfiability games $\mathcal{S}(\mathbb{A})$ and $\mathcal{S}(\mathbb{A}')$; for instance, we will see that an infinite match Σ of the consequence game naturally induces infinite matches of the two respective satisfiability games, and that the winning condition on infinite matches of the consequence game is formulated accordingly.

As a third novelty of this paper, related to these games, we identify a new class of automata, which we call *semi-disjunctive*. These structures can be viewed as an automata-theoretic counterpart to the “weakly aconjunctive formulas” introduced by Walukiewicz. They are much less constrained than disjunctive automata, but their one-step formulas are

still of a shape that guarantees the trace theory of an infinite match of the satisfiability game to be well behaved. In fact, by restricting the use of conjunctions in the one-step formulas of a semi-disjunctive automaton, we can guarantee that the collection of *bad* traces associated with a match of the satisfiability game is finite (modulo a natural equivalence relation of cofinal equality). Another important observation is that the standard and the thin satisfiability game for a semi-disjunctive automaton \mathbb{A} are equivalent in the sense that \exists has a winning strategy in the one game iff this applies to the other game as well.

Main results Concerning the concepts that we just discussed, we prove a number of results that we consider to be of independent interest. Here we outline three contributions that stand out as lemmas in the main proof.

The first result, which involves the Boolean operations of conjunction and negation that we shall define on modal automata, provides an essential link between the consequence game and the thin satisfiability game. We say that an automaton has a (thin) refutation if player \forall has a winning strategy in the (thin) satisfiability game for the automaton.

Theorem 3 *Let \mathbb{A} and \mathbb{D} be respectively be a semi-disjunctive and an arbitrary modal automaton such that $\mathbb{A} \models_{\mathcal{G}} \mathbb{D}$. Then the automaton $\mathbb{A} \wedge \neg \mathbb{D}$ has a thin refutation.*

This result can be seen both as an automata-theoretic counterpart and as a significant strengthening of a key result in Walukiewicz’ proof, viz., his Lemma 36.

Our second main auxiliary result, which also concerns the consequence game, can be seen as a strengthening of a classic result, viz., the *simulation theorem* for modal (or alternating) automata. Simulation theorems are among the pillarstones of automata theory. They generally show that an automaton of certain given type can be transformed into, or “simulated” by, an automaton that recognizes the same language, but in which the transition structure is of a conceptually simpler kind. Typically this either means that the transition structure is deterministic rather than non-deterministic or non-deterministic rather than alternating. The simplest result in this category is the standard powerset construction for (finite) word automata [56]; other examples include the simulation of non-deterministic Büchi automata for streams (infinite words) by deterministic Rabin automata due to Safra [59], or the simulation of alternating parity tree automata by non-deterministic ones due to Emerson & Jutla [20]. In our terminology, the simulation theorem says that any modal automaton \mathbb{A} can be simulated by a disjunctive automaton $\text{sim}(\mathbb{A})$. Our contribution is to show that this can be strengthened as follows.

Theorem 4 *The map $\text{sim}(\cdot)$ assigns to each modal automaton \mathbb{A} a disjunctive modal automaton $\text{sim}(\mathbb{A})$ such that*

- (1) $\mathbb{A} \models_{\mathcal{G}} \text{sim}(\mathbb{A})$ and $\text{sim}(\mathbb{A}) \models_{\mathcal{G}} \mathbb{A}$;
- (2) $\mathbb{B}[\text{sim}(\mathbb{A})/p] \models_{\mathcal{G}} \mathbb{B}[\mathbb{A}/p]$, for any modal automaton \mathbb{B} which is positive in p .

Here the substitution of an automaton for the variable p in \mathbb{B} is a well-defined operation on automata that will be introduced in Section 4. This result also has a counterpart from Walukiewicz’ proof: it roughly corresponds to his Lemma 39 in [73], but again ours is a

stronger result since we place no restrictions at all on the automaton \mathbb{B} that appears as a parameter (apart from requiring that all occurrences of the proposition letter p are positive).

Note that Theorem 3 and Theorem 4 can be formulated and proved completely independently of any deductive system for the modal μ -calculus.

The third result that we want to mention here crucially involves Kozen’s axiomatization. Recall that as the main goal of our completeness proof, we stated our intention to prove that \exists has a winning strategy in the satisfiability game for any consistent automaton \mathbb{A} . Theorem 5 below states a slightly weaker version of this, phrased in terms of the *thin* satisfiability game.

Theorem 5 *\exists has a winning strategy in the thin satisfiability game for any consistent modal automaton \mathbb{A} .*

We will informally refer to this observation as “Kozen’s Lemma”, since it is an automata-theoretic version of Kozen’s partial completeness result for the aconjunctive fragment of the modal μ -calculus [36]. As a consequence of the fact that for semi-disjunctive automata the standard and the thin satisfiability game are equivalent, our Theorem 5 also yields a partial completeness result, stating that any consistent (formula corresponding to a) semi-disjunctive automaton is satisfiable.

Proof of completeness Bringing all these ideas and results together, as the principal lemma in our proof we obtain the following version of Walukiewicz’ main technical result, with \leq_K denoting provable implication with respect to Kozen’s axiomatization.

Theorem 6 *For every formula $\varphi \in \mu\text{ML}$ there is a semantically equivalent disjunctive automaton \mathbb{D} such that $\varphi \leq_K \mathbb{D}$.*

We prove this theorem by a formula induction, and it should not come as a surprise that the key inductive cases are those concerning the fixpoint operators. In particular, the case where φ is of the form $\varphi = \mu x.a$ requires all of the machinery developed earlier on.

Finally, from Theorem 6, the completeness theorem is almost immediate. If φ is an arbitrary consistent formula φ , then by Theorem 6 it is semantically equivalent to a consistent disjunctive automaton \mathbb{D} . But for disjunctive automata it is easy to prove that consistency implies satisfiability, and so we are done.


Future work Our hope is that our “deconstruction” of the results and methods involved in the proof of the Kozen-Walukiewicz completeness theorem will lead to a better widespread understanding of a difficult result, but also that it will serve as a stepping stone for future research. In particular, we believe that an important direction for future work is to provide complete axiomatizations of several extensions, variations of and systems related to the modal μ -calculus. An example that we already mentioned is Parikh’s game logic, more directly related examples are the μ -calculus with converse [68] and the closely related guarded fixpoint logic [28], hybrid μ -calculus [63], probabilistic variants as in e.g. [42], or even inflationary modal fixpoint logic [16]. We should mention here that a complete axiom system for the hybrid μ -calculus was recently presented by Tamura [67], but this system uses a rule that directly

involves a bound on model sizes for satisfiable formulas, rather than the Park induction rule used in Kozen’s system. We are currently studying the problem of proving completeness for coalgebraic generalizations of the μ -calculus, and a result that covers all weak pullback- and finite set preserving functors is forthcoming in [21]. We hope to be able in the near future to extend this result further, covering systems like the monotone μ -calculus [41].

Two additional possible directions of research deserve to be mentioned. First, an interesting task would be to aim for generic completeness results for *fragments* of the μ -calculus. This would be a continuation of the work of Santocanale & Venema in [62] which provides generic completeness results for flat fixpoint logics, and the main question would be whether it is possible to push these results beyond the flat and even alternation-free fragments. (Again, this naturally ties in with the problem of proving completeness for game logic.) Second, we would like to study the problem of completeness for *axiomatic* extensions of the μ -calculus (as opposed to the *expressive* extensions mentioned earlier). This would be another step towards bridging the gap between the study of modal fixpoint logics and general research in modal logic, where the study of axiomatic extensions of the minimal normal or classical modal logic is usually at the centre of attention. Although such extensions of the μ -calculus are relatively rarely mentioned in the literature, some research does exist that suggests they are worthy of investigation, see for example [25] for a study of the least fixpoint extension of the logic **S4** in a topological context.

Overview of paper In section 2 we fix some notation and terminology on infinite games, and on elementary mathematics. Section 3 introduces the syntax and semantics of the modal μ -calculus, and we define Kozen’s deductive system $\mathbf{K}\mu$. Section 4 sees the appearance of the main characters of our work, viz., the modal automata; we also discuss their one-step logic, and define the translation from μML -formulas to modal automata which is based on automata-theoretic operations that correspond to syntactic operators of μML . In section 5 we define the two satisfiability games and the consequence game. Section 6 is pivotal to our paper: here we introduce the disjunctive and semi-disjunctive automata, and we prove Theorem 3. Section 7 is devoted to the proof of our strong simulation result, Theorem 4. In section 8, which can be read independently of the sections 5 – 7, we provide the translation back from automata to formulas, and we prove Theorem 2. In section 9 we focus on the proof of Kozen’s Lemma, Theorem 5. We wrap things up in the final section 10, where we prove our main lemma, Theorem 6, and we show how to derive the Kozen-Walukiewicz result, Theorem 1, from this. Finally, while we have made an effort to provide all main results with detailed proofs, we have put some of the more tedious arguments and derivations to an appendix.

The structure of the paper is shown in the following dependency graph, where the arrows represent the order in which individual sections may be read:



2 Preliminaries

We assume familiarity with the basic notions concerning infinite games [27]. Here we fix some notation and terminology, also regarding elementary mathematical concepts.

2.1 Basic mathematical concepts and notation

Definition 2.1 The power set of a set A is denoted as PA , and its size as $|A|$. \triangleleft

Since binary relations play an important role in our work, we will frequently use the following notation.

Definition 2.2 The collection of binary relations over a set A is denoted as A^\sharp .

Given a relation $R \subseteq A \times A'$, we let $\mathsf{Dom}R$ and $\mathsf{Ran}R$ denote its domain and range, respectively; for a subset $B' \subseteq A'$, we define $\mathsf{Ran}_{B'}R := \mathsf{Ran}R \cap B'$. Furthermore, we denote the converse relation of R as $R^{-1} := \{(a', a) \in A' \times A \mid (a, a') \in R\}$, and we set $R[a] := \{a' \in A' \mid Raa'\}$. Given a relation $R \subseteq A \times A$ and a subset $B \subseteq A$, we let $\mathsf{Res}_B R := R \cap (B \times B)$ denote the *restriction* of R to B . \triangleleft

Definition 2.3 Given a relation $R \subseteq A \times A'$, we define the following relations between PA and PA' :

$$\begin{aligned} \vec{\mathsf{P}}R &:= \{(B, B') \in \mathsf{PA} \times \mathsf{PA}' \mid \text{for all } b \in B \text{ there is a } b' \in B' \text{ with } Rbb'\} \\ \overleftarrow{\mathsf{P}}R &:= \{(B, B') \in \mathsf{PA} \times \mathsf{PA}' \mid \text{for all } b' \in B' \text{ there is a } b \in B \text{ with } Rbb'\} \\ \overline{\mathsf{P}}R &:= \vec{\mathsf{P}}R \cap \overleftarrow{\mathsf{P}}R. \end{aligned}$$

The relation $\overline{\mathsf{P}}R$ is called the *Egli-Milner* lifting of R . \triangleleft

Definition 2.4 We write $f : A \overset{\circ}{\rightarrow} B$ to denote that f is a partial map from A to B , and we denote the *graph* of f as $\mathsf{Gr}f := \{(a, fa) \mid a \in \mathsf{Dom}f\}$; here $\mathsf{Dom}f$ denotes the domain of f . The composition of two (partial) functions $f : A \overset{\circ}{\rightarrow} B$ and $g : B \overset{\circ}{\rightarrow} C$ is denoted as $g \circ f : A \overset{\circ}{\rightarrow} C$. \triangleleft

Definition 2.5 Given a set A , we let A^* and A^ω denote, respectively, the set of *words* (finite sequences) and *streams* (infinite sequences) over A . We will write both ww' and $w \cdot w'$ to denote the concatenation of the words w and w' , and similar for the concatenation of a word and a stream. The last symbol of a word w is denoted as $\mathsf{last}(w)$.

Two A -streams σ and τ are *eventually equal*, denoted as $\sigma =_\infty \tau$, if there is a $k \in \omega$ such that $\sigma(j) = \tau(j)$ for all $j \geq k$. \triangleleft

2.2 Graph games

Definition 2.6 A *board game* is a tuple $\mathbb{G} = (G_\exists, G_\forall, E, W)$ where G_\exists and G_\forall are disjoint sets, and, with $G := G_\exists \cup G_\forall$ denoting the *board* of the game, the binary relation $E \subseteq G^2$ encodes the moves that are admissible to the respective players, and $W \subseteq G^\omega$ denotes the

winning condition of the game. In a *parity game*, the winning condition is determined by a parity map $\Omega : G \rightarrow \omega$ with finite range, in the sense that the set W_Ω is given as the set of G -streams $\rho \in G^\omega$ such that the maximum value occurring infinitely often in the stream $(\Omega\rho_i)_{i \in \omega}$ is even.

Elements of G_\exists and G_\forall are called *positions* for the players \exists and \forall , respectively; given a position p for player $\Pi \in \{\exists, \forall\}$, the set $E[p]$ denotes the set of *moves* that are *legitimate* or *admissible* to Π at p . In case $E[p] = \emptyset$ we say that player Π *gets stuck* at p .

An *initialized board game* is a pair consisting of a board game \mathbb{G} and a *initial* position p , usually denoted as $\mathbb{G}@p$. \triangleleft

Definition 2.7 A *match* of a graph game $\mathbb{G} = (G_\exists, G_\forall, E, W)$ is nothing but a (finite or infinite) path through the graph (G, E) . Such a match ρ is called *partial* if it is finite and $E[\text{last}\rho] \neq \emptyset$, and *full* otherwise. We let PM_Π denote the collection of partial matches ρ ending in a position $\text{last}(\rho) \in G_\Pi$, and define $\text{PM}_\Pi@p$ as the set of partial matches in PM_Π starting at position p .

The *winner* of a full match ρ is determined as follows. If ρ is finite, then by definition one of the two players got stuck at the position $\text{last}(\rho)$, and so this player loses ρ , while the opponent wins. If ρ is infinite, we declare its winner to be \exists if $\rho \in W$, and \forall otherwise. \triangleleft

Definition 2.8 A *strategy* for a player $\Pi \in \{\exists, \forall\}$ is a map $\chi : \text{PM}_\Pi \rightarrow G$. A strategy is *positional* if it only depends on the last position of a partial match, i.e., if $\chi(\rho) = \chi(\rho')$ whenever $\text{last}(\rho) = \text{last}(\rho')$; such a strategy can and will be presented as a map $\chi : G_\Pi \rightarrow G$.

A match $\rho = (p_i)_{i < \kappa}$ is *guided* by a Π -strategy χ if $\chi(p_0 p_1 \dots p_{n-1}) = p_n$ for all $n < \kappa$ such that $p_0 \dots p_{n-1} \in \text{PM}_\Pi$ (that is, $p_{n-1} \in G_\Pi$). A Π -strategy χ is *legitimate* in $\mathbb{G}@p$ if the moves that it prescribes to χ -guided partial matches in $\text{PM}_\Pi@p$ are always admissible to Π , and *winning* for Π in $\mathbb{G}@p$ if in addition all χ -guided full matches starting at p are won by Π .

A position p is a *winning position* for player $\Pi \in \{\exists, \forall\}$ if Π has a winning strategy in the game $\mathbb{G}@p$; the set of these positions is denoted as Win_Π . The game $\mathbb{G} = (G_\exists, G_\forall, E, W)$ is *determined* if every position is winning for either \exists or \forall . \triangleleft

When defining a strategy χ for one of the players in a board game, we can and in practice will confine ourselves to defining χ for partial matches that are themselves guided by χ .

The following fact, independently due to Emerson & Jutla [20] and Mostowski [44], will be quite useful to us.

Fact 2.9 (Positional Determinacy) *Let $\mathbb{G} = (G_\exists, G_\forall, E, W)$ be a graph game. If W is given by a parity condition, then \mathbb{G} is determined, and both players have positional winning strategies.*

3 The modal μ -calculus

Although we assume the reader is familiar with the syntax and semantics of the modal μ -calculus, here we provide a quick recapitulation of the main notions that play a role in this paper. More detail on the formalism can be found in [27, 69].

3.1 Syntax

Throughout this paper we fix an (unnamed) infinite set of propositional variables.

Definition 3.1 The language μML of the modal μ -calculus is given by the following grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \diamond\varphi \mid \mu x.\varphi,$$

where p and x are propositional variables, and the formation of the formula $\mu x.\varphi$ is subject to the constraint that the variable x is *positive* in φ , i.e., all occurrences of x in φ are in the scope of an even number of negations. Elements of μML will be called *modal fixpoint formulas*, *μ -formulas*, or simply *formulas*.

The collection of *subformulas* of a formula is defined as usual, as are the sets of, respectively, its *free* and *bound* variables. We let $\mu\text{ML}(\mathbf{X})$ denote the set of μ -formulas of which all free variables belong to the set \mathbf{X} . \triangleleft

Remark 3.2 In order to focus completely on the hard and intricate parts of the completeness proof, we restrict attention to *monomodal* logic here, that is, we consider the version of modal logic with one single primitive modality \diamond . The completeness proof for the polymodal μ -calculus, where one has a family $\{\langle d \rangle \mid d \in D\}$ of modal diamonds, can be obtained by a straightforward adaptation of the monomodal case, and follows from our more general result in [21]. \triangleleft

As a convention, the free variables of a formula φ are denoted by the symbols p, q, r, \dots , and referred to as *proposition letters*, while we use the symbols x, y, z, \dots for the bound variables of a formula. Throughout the paper we will use standard abbreviations, including the symbols $\top, \perp, \wedge, \rightarrow, \square, \bigwedge$, and \bigvee (where the latter two symbols are used to denote arbitrary but finite conjunctions and disjunctions, respectively).

Definition 3.3 Let φ and $\{\psi_z \mid z \in Z\}$ be modal fixpoint formulas, where Z is a set of variables that are of free in φ . Then we let

$$\varphi[\psi_z/z \mid z \in Z]$$

denote the formula obtained from φ by simultaneously substituting each formula ψ_z for z in φ (with the usual understanding that no free variable in any of the ψ_z will get bound by doing so). In case Z is a singleton z , we will simply write $\varphi[\psi_z/z]$, or $\varphi[\psi]$ if z is clear from context. If $Z = Y_1 \uplus Y_2$, it will occasionally be convenient to write $\varphi[\psi_z/z \mid z \in Y_1, \psi_z/z \mid z \in Y_2]$ instead of $\varphi[\psi_z/z \mid z \in Z]$. \triangleleft

Fact 3.4 Let $\{\psi_y \mid y \in \mathbf{Y}\}$ and $\{\chi_z \mid z \in \mathbf{Z}\}$ be sets of formulas that are indexed by two disjoint sets of variables \mathbf{Y} and \mathbf{Z} . Then for every formula φ we have

- (1) $\varphi[\psi_y/y \mid y \in \mathbf{Y}][\chi_z/z \mid z \in \mathbf{Z}] = \varphi[\psi_y[\chi_z/z \mid z \in \mathbf{Z}]/y \mid y \in \mathbf{Y}, \chi_z/z \mid z \in \mathbf{Z}]$
- (2) $\varphi[\psi_y/y \mid y \in \mathbf{Y}][\chi_z/z \mid z \in \mathbf{Z}] = \varphi[\psi_y/y \mid y \in \mathbf{Y}, \chi_z/z \mid z \in \mathbf{Z}]$, provided no $z \in \mathbf{Z}$ occurs freely in any ψ_y .

We will sometimes make the assumption (but always explicitly) that our formulas are in negation normal form.

Definition 3.5 A formula of the modal μ -calculus is in *negation normal form* if it belongs to the language given by the following grammar:

$$\varphi ::= p \mid \neg p \mid \varphi \vee \varphi \mid \varphi \wedge \varphi \mid \diamond \varphi \mid \square \varphi \mid \mu x.\varphi \mid \nu x.\varphi,$$

where p and x are propositional variables, and the formation of the formulas $\mu x.\varphi$ and $\nu x.\varphi$ is subject to the constraint that the variable x is *positive* in φ , i.e., no occurrence of x in φ is in the scope of a negation.

We use the symbol η to range over μ and ν . \triangleleft

3.2 Structures

Definition 3.6 Given a set S , an *A-marking* on S is a map $m : S \rightarrow \mathbf{PA}$; an *A-valuation* on S is a map $V : A \rightarrow \mathbf{PS}$. Any valuation $V : A \rightarrow \mathbf{PS}$ gives rise to its *transpose marking* $V^\dagger : S \rightarrow \mathbf{PA}$ defined by $V^\dagger(s) := \{a \in A \mid s \in V(a)\}$, and dually each marking gives rise to a valuation in the same manner. \triangleleft

Since markings and valuations are interchangeable notions, we will often switch from one perspective to the other, based on what is more convenient in context.

Definition 3.7 A *Kripke structure* over a set \mathbf{X} of proposition letters is a triple $\mathbb{S} = (S, R, V)$ such that S is a set of objects called *points*, $R \subseteq S \times S$ is a binary relation called the *accessibility* relation, and V is an \mathbf{X} -valuation on S .

Given a Kripke structure $\mathbb{S} = (S, R, V)$, a propositional variable x and a subset U of S , we define $V[x \mapsto U]$ as the $\mathbf{X} \cup \{x\}$ -valuation given by

$$V[x \mapsto U](p) := \begin{cases} V(p) & \text{if } p \neq x \\ U & \text{otherwise,} \end{cases}$$

and we let $\mathbb{S}[x \mapsto U]$ denote the structure $(S, R, V[x \mapsto U])$. \triangleleft

Remark 3.8 Occasionally it will be convenient to take a coalgebraic perspective on Kripke structures. With \mathbf{X} denoting a set of proposition letters, for a given set S we define

$$\mathbf{K}_{\mathbf{X}}S := \mathbf{PX} \times \mathbf{PS},$$

that is, $\mathbf{K}_{\mathbf{X}}S$ denotes the set of pairs (\mathbf{Y}, U) with $\mathbf{Y} \subseteq \mathbf{X}$ and $U \subseteq S$. In practice we will usually write \mathbf{K} rather than $\mathbf{K}_{\mathbf{X}}$, assuming that the set \mathbf{X} of proposition letters is clear from context.

A Kripke structure $\mathbb{S} = (S, R, V)$ over the set \mathbf{X} can then be represented as a map $\sigma_{\mathbb{S}} : S \rightarrow \mathbf{KS}$ given by

$$\sigma_{\mathbb{S}}(s) := (V^\dagger(s), R[s]).$$

This map $\sigma_{\mathbb{S}}$ will be called the *(coalgebraic) unfolding map* of \mathbb{S} .

The operation \mathbf{K} is in fact a *functor* on the category of sets with functions, and while we do not focus on this, in order to have compact notation it will be useful to borrow the following bit of category theory. Note that any map $f : S \rightarrow S'$ gives rise to a map $\mathbf{K}f$ from \mathbf{KS} to \mathbf{KS}' , defined by

$$\mathbf{K}f : (\mathbf{Y}, U) \mapsto (\mathbf{Y}, f[U]).$$

The only fact about this map that we shall need is that it satisfies the composition law, stating that

$$\mathbf{K}(g \circ f) = \mathbf{K}g \circ \mathbf{K}f.$$

for any pair of composeable maps g, f . ◁

3.3 Semantics

Definition 3.9 By induction on the complexity of modal fixpoint formulas, we define a meaning function $\llbracket \cdot \rrbracket$, which assigns to a formula $\varphi \in \mu\text{ML}$ its *meaning* $\llbracket \varphi \rrbracket^{\mathbb{S}} \subseteq S$ in any Kripke structure $\mathbb{S} = (S, R, V)$. The clauses of this definition are standard:

$$\begin{aligned} \llbracket p \rrbracket^{\mathbb{S}} &:= V(p) \\ \llbracket \neg\varphi \rrbracket^{\mathbb{S}} &:= S \setminus \llbracket \varphi \rrbracket^{\mathbb{S}} \\ \llbracket \varphi \vee \psi \rrbracket^{\mathbb{S}} &:= \llbracket \varphi \rrbracket^{\mathbb{S}} \cup \llbracket \psi \rrbracket^{\mathbb{S}} \\ \llbracket \diamond\varphi \rrbracket^{\mathbb{S}} &:= \{s \in S \mid R[s] \cap \llbracket \varphi \rrbracket^{\mathbb{S}} \neq \emptyset\} \\ \llbracket \mu x.\varphi \rrbracket^{\mathbb{S}} &:= \bigcap \{U \in \text{PS} \mid \llbracket \varphi \rrbracket^{\mathbb{S}[x \rightarrow U]} \subseteq U\}. \end{aligned}$$

If a point $s \in S$ belongs to the set $\llbracket \varphi \rrbracket^{\mathbb{S}}$, we write $\mathbb{S}, s \Vdash \varphi$, and say that φ is *true at s* or *holds at s* , or that s *satisfies* φ . ◁

Definition 3.10 A modal fixpoint formula φ is *valid*, notation: $\models \varphi$, if $\llbracket \varphi \rrbracket^{\mathbb{S}} = S$ for any structure $\mathbb{S} = (S, R, V)$, and *satisfiable* if $\llbracket \varphi \rrbracket^{\mathbb{S}} \neq \emptyset$ for some structure \mathbb{S} . Two formulas φ and ψ are *equivalent*, notation: $\varphi \equiv \psi$, if $\llbracket \varphi \rrbracket^{\mathbb{S}} = \llbracket \psi \rrbracket^{\mathbb{S}}$ for any structure \mathbb{S} . ◁

3.4 The cover modality

As mentioned in the introduction, the *cover modality* ∇ , which was independently introduced in coalgebraic logic [43] and in automata theory [30], plays a prominent role in our proof, just as in Walukiewicz'. It is a slightly non-standard connective that takes a finite set Φ of formulas as its argument.

Definition 3.11 Given a finite set Φ , we let $\nabla\Phi$ abbreviate the formula

$$\nabla\Phi := \bigwedge \diamond\Phi \wedge \square \bigvee \Phi,$$

where $\diamond\Phi$ denotes the set $\{\diamond\varphi \mid \varphi \in \Phi\}$. ◁

Remark 3.12 Observe that the semantics of the cover modality can be expressed in terms of the Egli-Milner lifting of the satisfaction relation \Vdash :

$$\mathbb{S}, s \Vdash \nabla\Phi \text{ iff } R[s] \bar{\Vdash} \Phi.$$

In words, $\nabla\Phi$ holds at s iff every successor of s satisfies some formula in Φ and every formula in Φ holds in some successor of s . From this observation it is easy to derive that, conversely, the standard modal operators can be expressed in terms of the cover modality:

$$\begin{aligned} \diamond\varphi &\equiv \nabla\{\varphi, \top\} \\ \square\varphi &\equiv \nabla\{\varphi\} \vee \nabla\emptyset, \end{aligned}$$

where we note that $\nabla\emptyset$ holds at a point s iff s is a ‘blind’ world, that is, $R[s] = \emptyset$. \triangleleft

3.5 Axiomatics

As mentioned in the introduction, Kozen’s axiomatization for the modal μ -calculus is obtained by adding the (pre-)fixpoint axiom and rule to the basic modal logic \mathbf{K} . For completeness’ sake we give the definition of Kozen’s system here, taking a standard \diamond -based axiomatization for \mathbf{K} [6, Remark 4.7].

Definition 3.13 The *axioms* of the basic modal logic \mathbf{K} are the following:

- (C) a complete set of axioms for classical propositional logic;
- (NA) axioms stating that \diamond is *normal* ($\neg\diamond\perp$) and *additive* ($\diamond(p \vee q) \leftrightarrow \diamond p \vee \diamond q$),

while its *derivation rules* are

- (MP) modus ponens: from φ and $\varphi \rightarrow \psi$, derive ψ ;
- (Mon) a monotonicity rule: from $\varphi \rightarrow \psi$ derive $\diamond\varphi \rightarrow \diamond\psi$;
- (US) uniform substitution: from φ derive $\varphi[\sigma]$, for any substitution σ .

Kozen’s *deductive system* $\mathbf{K}\mu$ is obtained by this by adding the following axiom schema and rule to those of \mathbf{K} :

- (A $_{\mu}$) all prefixpoint axioms of the form $\varphi[\mu x.\varphi/x] \rightarrow \mu x.\varphi$;
- (R $_{\mu}$) Park’s prefixpoint rule: from $\varphi[\gamma/x] \rightarrow \gamma$ derive $\mu x.\varphi \rightarrow \gamma$.

A derivation in $\mathbf{K}\mu$ is a finite list of μ ML-formulas, such that each formula on the list is either an axiom of $\mathbf{K}\mu$ or obtained from earlier formulas by applying one of the derivation rules of $\mathbf{K}\mu$. \triangleleft

Definition 3.14 A μ -formula φ is *derivable* or *provable*, notation: $\vdash_K \varphi$, if there is a $\mathbf{K}\mu$ -derivation leading up to φ . Given two formulas φ and ψ , we say that φ *provably implies* ψ , notation: $\varphi \leq_K \psi$, if the formula $\varphi \rightarrow \psi$ is derivable. The formulas φ and ψ are *provably equivalent*, notation: $\varphi \equiv_K \psi$, if $\varphi \leq_K \psi$ and $\psi \leq_K \varphi$. A formula is *consistent* if its negation is not provable. \triangleleft

Without proof we mention the following facts on our proof system.

Fact 3.15 *Let φ be a modal μ -formula, and let Φ, Ψ be sets of modal μ -formulas. Then*

- (1) $\nabla\Phi \wedge \nabla\Psi \equiv_K \bigvee\{\nabla\{\varphi \wedge \psi \mid \varphi R\psi\} \mid R \subseteq \Phi \times \Psi \text{ and } (\Phi, \Psi) \in \overline{PR}\}$;
- (2) $\nabla\{\varphi_0 \vee \varphi_1 \mid \varphi \in \Phi\} \equiv_K \bigvee\{\nabla\{\varphi_i \mid (\varphi, i) \in Z\} \mid Z \subseteq \Phi \times \{0, 1\}, \text{Dom}Z = \Phi\}$;
- (3) $\mu x.\mu y.\varphi \equiv_K \mu y.\mu x.\varphi$;
- (4) $\varphi \equiv_K \varphi'$ for some effectively obtainable formula φ' in negation normal form;
- (5) $\varphi \leq_K \psi$ only if $\eta x.\varphi \leq_K \eta x.\psi$;
- (6) $\neg\eta x.\varphi \equiv_K \lambda x.\neg\varphi[\neg x/x]$, where $\{\eta, \lambda\} = \{\mu, \nu\}$.

4 Modal automata and their one-step logic

As mentioned in the introduction, one of the main goals of the present paper is to further strengthen the role of automata theory in the completeness proof for the modal μ -calculus. While Walukiewicz’s proof works directly with what we will call disjunctive automata, we will work with the wider class of *modal automata* [74] that we will introduce in this section.

A further goal in this section is to define a translation transforming a formula of the μ -calculus into an equivalent modal parity automata. Of course, there are already a few different methods available for this transformation. Janin & Walukiewicz [30] first construct a tableau for the formula, which is then transformed into an automaton in which the states are certain distinguished nodes of the tableau. This method already produces a *non-deterministic* automaton (what we will call a “disjunctive” automaton), and so is not suited for our purposes since we want to work with the wider class of alternating modal automata introduced by Wilke [74]. It seems that the standard approach to this (see for instance [27], following Wilke), is to transform a μ ML-formula φ into an automaton in one go, by taking the states of the automaton to be syntactic items related to φ (such as its subformulas or bound variables), and then (possibly) perform some postprocessing in order to get the device into the right shape. Our preferred method here will be to define the translation by induction on the complexity of formulas, making use of certain effective closure conditions on the class of modal automata. In fact, most (but not all) of the operations on automata that we will use to take care of the inductive step of the translation, are the ones used by Wilke in order to prove the correctness of his translation.

Before turning to the introduction of the modal automata themselves, we first define and discuss the one-step logic that determines the shape of their transition function. As mentioned in the introduction, the notion of a “one-step logic” stems from the literature on coalgebra [12], where it is used to obtain a modular approach to defining and studying logics for specifying the behaviour of a wide variety of coalgebras, or state-based evolving systems. The idea behind this logic is that it provides the syntax and semantics to extract information about the one-step behaviour of such a system, that is, the properties of one single unfolding of a state in the system. One-step logic thus comes with the notion of one-step syntax (a language consisting of one-step formulas), one-step semantics, and (possibly) one-step derivation systems and one-step model theory. This perspective is very compatible with the theory of automata operating on infinite objects, and Fontaine, Leal & Venema [22] introduced a notion of coalgebra automata of which the transition function maps states of the automaton to one-step formulas.

4.1 One-step logic

Modal automata are based on the *modal one-step language*. This language consists of modal formulas of rank 1, built up from *proposition letters* (which must appear unguarded) and *variables* (which must appear guarded). In practice, the variables of a one-step formula will always be states of some automaton.

Definition 4.1 Given a set P , we define the set $\text{Latt}(P)$ of *lattice terms* over P through the

following grammar:

$$\pi ::= \perp \mid \top \mid p \mid \pi \wedge \pi \mid \pi \vee \pi,$$

where $p \in P$. Given two sets \mathbf{X} and A , we define the set $\mathbf{1ML}(\mathbf{X}, A)$ of *modal one-step formulas* over A with respect to \mathbf{X} inductively by

$$\alpha ::= \perp \mid \top \mid p \mid \neg p \mid \diamond \pi \mid \square \pi \mid \alpha \wedge \alpha \mid \alpha \vee \alpha,$$

with $p \in \mathbf{X}$ and $\pi \in \mathbf{Latt}(A)$. ◁

Observe that the set of modal one-step formulas over A with respect to \mathbf{X} corresponds to the set of lattice terms over the set $\{p, \neg p \mid p \in \mathbf{X}\} \cup \{\diamond \pi, \square \pi \mid \pi \in \mathbf{Latt}(P)\}$. Note too that elements from the two parameter sets, \mathbf{X} and A , are treated quite differently in the syntax of one-step formulas: all occurrences of elements of \mathbf{X} , corresponding to the proposition letters, must be unguarded, whereas the elements of A , corresponding to bound variables of a formula and to states of our modal automata, may only occur in the scope of exactly one modality.

Modal one-step formulas will serve to provide the type of the transition map of a modal automaton, which will map states of the automaton to modal one-step formulas over the set of states. Intuitively, in a successful “run” of a modal automaton on a Kripke structure, given that some state s in a Kripke structure is visited by some state a of the automaton, the run should provide a “local valuation” of the states of the automaton over the set of successors of s , in such a way that the formula assigned to the state a becomes true. So a run of a modal automaton proceeds step-by-step, at each moment considering a local window into the Kripke structure as seen from the particular state s that is currently being visited. In order to make these intuitions precise, we need a semantics for modal one-step formulas. This is given by *one-step models*.

Definition 4.2 Fix sets \mathbf{X} and A . A *one-step frame* is a pair (Y, S) where S is any set, and $Y \subseteq \mathbf{X}$. A *one-step model* is a triple (Y, S, m) such that (Y, S) is a one-step frame and m is an A -marking on S . ◁

Observe that with this definition, the coalgebraic representation of a Kripke structure (S, R, V) can now be seen as a function σ_S mapping any state $s \in S$ to a one-step frame of which the carrier is a subset of S .

We now turn to the semantics based on one-step models:

Definition 4.3 The *one-step satisfaction relation* \Vdash^1 between one-step models and one-step formulas is defined as follows. Fix a one-step model (Y, S, m) . First, we define the value $\llbracket \pi \rrbracket$ of a lattice formula π over A by induction, setting $\llbracket a \rrbracket = \{s \in S \mid a \in m(s)\}$ for $a \in A$, and treating conjunctions and disjunctions in the obvious manner.

Now we define the one-step satisfaction relation by giving the usual clauses for conjunction and disjunction, and the following clauses for the literals and modal operators:

- $(Y, S, m) \Vdash^1 \square \pi$ iff $\llbracket \pi \rrbracket = S$,
- $(Y, S, m) \Vdash^1 \diamond \pi$ iff $\llbracket \pi \rrbracket \neq \emptyset$,
- $(Y, S, m) \Vdash^1 p$ iff $p \in Y$,

- $(Y, S, m) \Vdash^1 \neg p$ iff $p \notin Y$.

Two one-step formulas α and α' are *(one-step) equivalent*, notation: $\alpha \equiv_1 \alpha'$, if they are satisfied by the same one-step models. \triangleleft

Examples of one-step equivalent pairs of formulas include the familiar axioms of modal logic, such as $\Box(a \wedge b) \equiv_1 \Box a \wedge \Box b$, but also formulas involving the nabla modality, such as $\nabla B \wedge \nabla B' \equiv_1 \bigvee \{ \nabla \{ b \wedge b' \mid bRb' \} \mid R \subseteq B \times B' \text{ and } (B, B') \in \overline{PR} \}$ (cf. Fact 3.15(1)).

One particular kind of one-step models will be of special interest to us.

Definition 4.4 Given a set A , we define the *canonical A -marking on PA* as the map $I_A : B \mapsto B$ (that is, the identity map on PA). More generally, for any subset \mathcal{B} of PA , we consider the marking $I_A \upharpoonright_{\mathcal{B}} : \mathcal{B} \rightarrow \mathsf{PX}$.

For any one-step formula $\alpha \in \mathbf{1ML}(\mathbf{X}, A)$ and any element $\Gamma \in \mathsf{KPA}$, say $\Gamma = (Y, \mathcal{B})$, we abbreviate $(\Gamma, I_A \upharpoonright_{\mathcal{B}}) \Vdash^1 \alpha$ as $\Gamma \Vdash^1 \alpha$, and we denote $\llbracket \alpha \rrbracket^1 := \{ \Gamma \in \mathsf{KPA} \mid \Gamma \Vdash^1 \alpha \}$. \triangleleft

The main result about the modal one-step language that we shall need later is the following one-step version of the usual bisimulation invariance result for modal logic, i.e. all one-step formulas are invariant for bisimulations between one-step models in a precise sense. Observe that the definition below makes use of the (Egli-Milner) relation lifting of Definition 2.3.

Definition 4.5 Let (Y, S, m) and (Y', S', m') be one-step models with respect to A and \mathbf{X} . We say that these models are *one-step bisimilar* if they satisfy the following conditions:

(atomic) $Y = Y'$;

(forth) for all $s \in S$, there is $s' \in S'$ with $m(s) = m'(s')$;

(back) for all $s' \in S'$, there is $s \in S$ with $m(s) = m'(s')$.

We write $(Y, S, m) \Leftrightarrow^1 (Y', S', m')$ to say that (Y, S, m) and (Y', S', m') are one-step bisimilar. \triangleleft

We can now state the one-step bisimulation invariance theorem:

Proposition 4.6 (One-step Bisimulation Invariance) *Let (Y, S, m) and (Y', S', m') be any two one-step models with respect A and \mathbf{X} . If $(Y, S, m) \Leftrightarrow^1 (Y', S', m')$, then both one-step models satisfy the same formulas in $\mathbf{1ML}(\mathbf{X}, A)$.*

We consider a useful instance of the one-step bisimulation invariance theorem: pick any set W , let $\Gamma = (Y, S)$ be any one-step frame in KW , and let $f : W \rightarrow W'$. Then any A -marking m on $f[S]$ gives rise to the marking $m \circ f$ on S , and clearly the one-step models $(\Gamma, m \circ f)$ and $(\mathsf{K}f(\Gamma), m)$ are one-step bisimilar:

$$(\Gamma, m \circ f) \Leftrightarrow^1 (\mathsf{K}f(\Gamma), m).$$

So by one-step bisimulation invariance, these two one-step models satisfy precisely the same one-step formulas: $(\Gamma, m \circ f) \Vdash^1 \alpha$ iff $(\mathsf{K}f(\Gamma), m) \Vdash^1 \alpha$ for all $\alpha \in \mathbf{1ML}(\mathbf{X}, A)$. This particular instance of the principle of one-step bisimulation invariance will play an important role in some of our main proofs.

We also remark on a variant of the one-step bisimulation invariance theorem that we will make use of later. It can be thought of as a combination of one-step bisimulation invariance and a monotonicity property of one-step formulas, which we obtain since all variables in A appear positively in any one-step formula in $\mathbf{1ML}(\mathbf{X}, A)$.

Definition 4.7 Let (Y, S, m) and (Y', S', m') be one-step models with respect to A and \mathbf{X} . We say that (Y', S', m') *one-step simulates* (Y, S, m) , notation: $(Y, S, m) \rightleftarrows^1 (Y', S', m')$, if

- (atomic) $Y = Y'$;
- (forth) for all $s \in S$, there is $s' \in S'$ with $m(s) \subseteq m'(s')$;
- (back) for all $s' \in S'$, there is $s \in S$ with $m(s) \subseteq m'(s')$. ◁

Proposition 4.8 (Preservation) *Let (Y, S, m) and (Y', S', m') be any two one-step models with respect to A and \mathbf{X} . If $(Y, S, m) \rightleftarrows^1 (Y', S', m')$, then any formula in $\mathbf{1ML}(\mathbf{X}, A)$ that is satisfied by (Y, S, m) is also satisfied by (Y', S', m') .*

Finally, we introduce the key property of the one-step logic for the purposes of our completeness proof: the one-step logic already enjoys a sort of completeness property with respect to the Kozen proof system, which we will lift to a completeness result for the full mu-calculus. This is the *one-step completeness theorem*, stated below.

Theorem 4.9 (One-step completeness) *Let $\alpha \in \mathbf{1ML}(\mathbf{X}, A)$ be a one-step formula, and let $\sigma : A \rightarrow \mu\mathbf{ML}(\mathbf{X})$ be a substitution such that the formula $\alpha[\sigma]$ is consistent. Then there is a one-step model (Y, S, m) such that $(Y, S, m) \Vdash^1 \alpha$ and for all $s \in S$, the set of formulas $\{\sigma(a) \mid a \in m(s)\}$ is consistent.*

Proof. We begin by rewriting the one-step formula α as a disjunction of conjunctions of the shape:

$$\gamma \wedge \heartsuit_1 \pi_1 \wedge \dots \wedge \heartsuit_n \pi_n$$

where γ is a conjunction of literals over \mathbf{X} , each \heartsuit -operator is a modality $\heartsuit_i \in \{\diamond, \square\}$, and each π_i is a lattice formula over A . Using the equivalences in Remark 3.12, we can rewrite α into a disjunction of formulas of the form

$$\gamma \wedge \nabla \Pi_1 \wedge \dots \wedge \nabla \Pi_n,$$

where each Π_i is a finite set of lattice formulas over A . Now apply Fact 3.15(1) repeatedly, and then distribute the conjunct γ over disjunctions, to obtain a disjunction of formulas of the shape $\gamma \wedge \nabla \Pi$ where Π is a finite set of lattice formulas over A .

Focusing on a disjunct of the shape $\gamma \wedge \nabla \Pi$, we may assume that each member of Π is in disjunctive normal form. Now we can apply Fact 3.15(2) repeatedly to each member of Π (and again, distribute the conjunct γ over disjunctions) to pull the disjunctions outside the scope of the modalities. In the end we obtain a formula $\beta \equiv_K \alpha$ in a certain normal form, viz., β is a disjunction of formulas of the form

$$\gamma \wedge \nabla \{ \bigwedge B_1, \dots, \bigwedge B_k \},$$

where γ is a conjunction of literals over \mathbf{X} and B_1, \dots, B_k are subsets of A . We thus find that $\alpha[\sigma] \equiv_K \beta[\sigma]$, so $\beta[\sigma]$ is consistent, which by propositional logic means that at least one disjunct

$$\gamma \wedge \nabla\{\bigwedge \sigma[B_1], \dots, \bigwedge \sigma[B_k]\}$$

is consistent. We construct a one-step model \mathbf{Y}, S, m by setting $S = \{B_1, \dots, B_k\}$, m is the identity map on S , and \mathbf{Y} is the set of proposition letters p in \mathbf{X} that appear as a conjunct in γ . It is easy to see that

$$\mathbf{Y}, S, m \Vdash \gamma \wedge \nabla\{\bigwedge B_1, \dots, \bigwedge B_k\}$$

and it follows that $\mathbf{Y}, S, m \Vdash \alpha$. It remains only to check that each conjunction $\bigwedge \sigma[B_i]$ is consistent - but were it not, the formula $\diamond \bigwedge \sigma[B_i]$ would be inconsistent by the normality axiom for the diamond, and we reach a contradiction since $\diamond \bigwedge \sigma[B_i]$ is a conjunct of the consistent formula $\gamma \wedge \nabla\{\bigwedge \sigma[B_1], \dots, \bigwedge \sigma[B_k]\}$. QED

4.2 Modal automata

We now formally introduce modal automata.

Definition 4.10 Fix a set of proposition letters \mathbf{X} . A *modal \mathbf{X} -automaton* \mathbb{A} is a quadruple (A, Θ, Ω, a_I) where A is a finite set of *states*, a_I is the *start state*, $\Omega : A \rightarrow \omega$ is the *priority map*, while the *transition map*

$$\Theta : A \rightarrow \mathbf{1ML}(\mathbf{X}, A)$$

maps states to one-step formulas. ◁

Some basic concepts concerning modal automata are introduced in the following definitions:

Definition 4.11 The *(directed) graph* of \mathbb{A} is the structure $(G, E_{\mathbb{A}})$, where $aE_{\mathbb{A}}b$ if a occurs in $\Theta(b)$, and we let $\triangleleft_{\mathbb{A}}$ denote the transitive closure of $E_{\mathbb{A}}$. If $a \triangleleft_{\mathbb{A}} b$ we say that a is *active* in b . We write $a \bowtie_{\mathbb{A}} b$ if $a \triangleleft_{\mathbb{A}} b$ and $b \triangleleft_{\mathbb{A}} a$.

A *cluster* of \mathbb{A} is a cell of the equivalence relation generated by $\bowtie_{\mathbb{A}}$ (i.e., the smallest equivalence relation on A containing $\bowtie_{\mathbb{A}}$); a cluster C is *degenerate* if it is of the form $C = \{a\}$ with $a \not\bowtie_{\mathbb{A}} a$. The unique cluster to which a state $a \in A$ belongs is denoted as C_a . ◁

Definition 4.12 Fix a modal \mathbf{X} -automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$. The *size* of \mathbb{A} is defined as the cardinality of its carrier A .

With $b \in A$, let $\mathbb{A}\langle b \rangle$ denote the variant of \mathbb{A} that takes b as its starting state, i.e., $\mathbb{A}\langle b \rangle = (A, \Theta, \Omega, b)$.

We write $a \sqsubset_{\mathbb{A}} b$ if $\Omega(a) < \Omega(b)$, and $a \sqsubseteq_{\mathbb{A}} b$ if $\Omega(a) \leq \Omega(b)$. When clear from context we sometimes write \sqsubset and \sqsubseteq instead, dropping the explicit reference to \mathbb{A} .

Given a state a of \mathbb{A} , we write $\eta_a = \mu$ if $\Omega(a)$ is odd, and $\eta_a = \nu$ if $\Omega(a)$ is even, and we call a an η_a -state. The sets of μ - and ν -states are denoted with A^μ and A^ν , respectively.

We say that \mathbb{A} is *positive* in a proposition letter $p \in \mathbf{X}$ if each occurrence of p in each formula $\Theta(a)$ is positive. ◁

Modal automata run on Kripke structures, and acceptance is defined in terms of a two-player game, the *acceptance game*. The two players have opposing goals: one player \exists (or “Eloise”) wants to defend the claim that \mathbb{A} accepts \mathbb{S} , and the opposing Player \forall (“Abelard”) wants to establish the opposite. The game connects the one-step logic to Kripke structures, and the main thing to notice is that for any Kripke structure $\mathbb{S} = (S, R, V)$ over \mathbf{X} , any given state $s \in S$ gives rise to a one-step frame $(V^\dagger(s), R[s])$ consisting of the variables in \mathbf{X} that are true at s , together with the set of R -successors of s .

Definition 4.13 Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be any modal \mathbf{X} -automaton, and let $\mathbb{S} = (S, R, V)$ be any Kripke structure. The *acceptance game* $\mathcal{A}(\mathbb{A}, \mathbb{S})$ for \mathbb{A} with respect to \mathbb{S} is defined as in the following table:

Position	Player	Admissible moves
$(a, s) \in A \times S$	\exists	$\{m : R[s] \rightarrow \text{PA} \mid (V^\dagger(s), R[s], m) \Vdash^1 \Theta(a)\}$
m	\forall	$\{(b, t) \mid b \in m(t)\}$

Winning conditions are the usual ones for parity games. That is, the loser of a finite match is the player who got stuck. An infinite match $(a_1, s_1)m_1(a_2, s_2)m_2(a_3, s_3)m_3 \dots$ induces a stream $a_1 a_2 a_3 \dots$ over the alphabet A , and we declare the winner of this match to be \exists if the highest priority state that appears infinitely often in the word $a_1 a_2 a_3 \dots$ has an even priority, and \forall is the winner otherwise.

We say that \mathbb{A} *accepts* the pointed structure (\mathbb{S}, s) if (a_I, s) is a winning position in the acceptance game $\mathcal{A}(\mathbb{A}, \mathbb{S})$, and write $\mathbb{S}, s \Vdash \mathbb{A}$ to denote that \mathbb{A} accepts (\mathbb{S}, s) . We define $[\mathbb{A}]^{\mathbb{S}} := \{s \in S \mid \mathbb{S}, s \Vdash \mathbb{A}\}$, and we define $L(\mathbb{A})$ (the “language recognized by \mathbb{A} ”) to be the class of pointed Kripke structures accepted by \mathbb{A} . \triangleleft

Definition 4.14 Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ and $\mathbb{A}' = (A', \Theta', \Omega', a'_I)$ be two modal automata. We say that \mathbb{A} (*semantically*) *implies* \mathbb{A}' , notation: $\mathbb{A} \leq \mathbb{A}'$, if $L(\mathbb{A}) \subseteq L(\mathbb{A}')$, and that \mathbb{A} and \mathbb{A}' are *equivalent*, notation: $\mathbb{A} \equiv \mathbb{A}'$, if they recognize the same language, i.e., if $L(\mathbb{A}) = L(\mathbb{A}')$. \triangleleft

In the sequel we will need the following strong version of equivalence between automata.

Definition 4.15 Two modal automata $\mathbb{A} = (A, \Theta, \Omega, a_I)$ and $\mathbb{A}' = (A', \Theta', \Omega', a'_I)$ are *one-step equivalent*, notation: $\mathbb{A} \equiv_1 \mathbb{A}'$, if $A = A'$, $\Omega = \Omega'$, $a_I = a'_I$, and $\Theta(a) \equiv_1 \Theta'(a)$ for all $a \in A$. \triangleleft

It is obvious that one-step equivalence implies equivalence.

4.3 Operations on modal automata

We now introduce the logical operations on modal automata that will enable us to translate formulas to modal automata, and later to connect proof theoretic concepts with automata theory. Most of these operations, like the Boolean and modal ones, and substitution, are standard [46, 74]. Our definitions of least and greatest fixpoints of modal automata, are new as far as we know.

Conjunction and disjunction

Suppose we are given modal automata $\mathbb{A} = (A, \Theta_A, \Omega_A, a_I)$ and $\mathbb{B} = (B, \Theta_B, \Omega_B, b_I)$. We define the automaton $\mathbb{A} \wedge \mathbb{B} = (C, \Theta_C, \Omega_C, a_C)$ as follows:

- a_C is some arbitrarily chosen object, and C is defined to be $A \uplus B \uplus \{a_C\}$.
- $\Theta_C(a_C) := \Theta_A(a_I) \wedge \Theta_B(b_I)$ and $\Omega_C(a_C) := k + 1$ where k is the maximum priority of \mathbb{A}, \mathbb{B} .

- For $a \in A$, $\Theta_C(a) := \Theta_A(a)$ and $\Omega_C(a) := \Omega_A(a)$.
- For $b \in B$, $\Theta_C(a) := \Theta_B(b)$ and $\Omega_C(b) := \Omega_B(b)$.

Disjunction is handled in precisely the same manner, setting $\Theta_C(a_C) := \Theta_A(a_I) \vee \Theta_B(b_I)$ instead.

Negation

Negation corresponds to *complementation* on the side of automata, and for this we need the concept of the *boolean dual* α^∂ of a one-step formula α :

Definition 4.16 First, we define the (*boolean*) *dual* of a lattice term over A , by setting:

$$\begin{aligned} a^\partial &:= a \\ (\pi \wedge \pi')^\partial &:= \pi^\partial \vee \pi'^\partial \\ (\pi \vee \pi')^\partial &:= \pi^\partial \wedge \pi'^\partial \end{aligned}$$

With this definition in place, by putting

$$\begin{array}{llll} p^\partial &:= \neg p & (\Box \pi)^\partial &:= \Diamond \pi^\partial & (\alpha \wedge \beta)^\partial &:= \alpha^\partial \vee \beta^\partial \\ (\neg p)^\partial &:= p & (\Diamond \pi)^\partial &:= \Box \pi^\partial & (\alpha \vee \beta)^\partial &:= \alpha^\partial \wedge \beta^\partial \end{array}$$

we inductively define the (*boolean*) *dual* of one-step formulas. ◁

Observe that in this definition we see another clear example of the different role of the proposition letters and the automaton states in one-step formulas.

Given a modal automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$ we define the automaton $\neg \mathbb{A} := (A, \Theta', \Omega', a_I)$ by setting, for each $a \in A$:

- $\Theta'(a) := \Theta(a)^\partial$
- $\Omega'(a) := \Omega(a) + 1$.

Modal operators

Given a modal automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$, pick an arbitrary object c , and define $\Diamond \mathbb{A} = (A', \Theta', \Omega', a'_I)$ by setting:

- $A' := A \uplus \{c\}$,
- $a'_I := c$,
- $\Theta'(a) := \Theta(a)$ for $a \in A$, and $\Theta'(c) := \Diamond a_I$,
- $\Omega'(a) := \Omega(a)$ for $a \in A$, and $\Omega'(c) := k + 1$ where k is the maximum priority of \mathbb{A} .

The definition of $\Box \mathbb{A}$ is similar, the only difference being that now we set $\Theta'(c) := \Box a_I$.

Substitution

Let $\mathbb{A} = (A, \Theta, a_I, \Omega)$ and $\mathbb{B} = (B, \Lambda, \Psi, b_I)$ be modal automata, and assume that \mathbb{A} is positive in x . We define the modal automaton $\mathbb{A}[\mathbb{B}/x]$ as the structure $(D, \Theta', \Omega', d_I)$, where $D := A \uplus B$, Θ' is given by

$$\Theta'(d) := \begin{cases} \Theta(d)[\Lambda(b_I)/x] & \text{if } d \in A \\ \Lambda(d) & \text{if } d \in B. \end{cases}$$

The priority map Ω' could have been defined as $\Omega \uplus \Psi$, but we find it convenient for later proofs to define Ω' instead so that all states in A get higher priority than all states in B . So let n be the least even number greater than any priority in \mathbb{B} . Then we set $\Omega'(b) = \Psi(b)$ for $b \in B$, and $\Omega'(a) = \Omega(a) + n$ for $a \in A$. (Clearly this will preserve the priority order among states in A and will not change the parities.) Finally, set $d_I = a_I$.

Fixpoint operators

We now turn to the definition of fixpoint operators on automata. For at least two reasons this is the most difficult case to handle. First, recall that in the one-step language associated with a modal automaton, the proposition letters (corresponding to the free variables of a formula) are treated rather differently from the states of the automaton (which correspond to the bound variables of a formula). We have good reasons to do so, but when constructing the automaton $\eta x.\mathbb{A}$ from an automaton \mathbb{A} there is a price to pay for this, related to the different status of the variable x in the two automata: while x is a free proposition letter in \mathbb{A} , and so appears only in unguarded positions in the one-step formulas, it is treated as a *state* of $\mu x.\mathbb{A}$ and must therefore appear only guarded in $\mu x.\mathbb{A}$. For this reason it will be necessary to *pre-process* the automaton \mathbb{A} putting it in a shape \mathbb{A}^x in which x is, in some sense, guarded.

Second, we have to be careful about how we go about this “pre-processing” of \mathbb{A} . The reason for this will become clearer once we consider the satisfiability game for modal automata in Section 5. The game is played between Eloise, who wants to show that the automaton accepts some model, and Abelard, who wants to show that the automaton does not accept any model. It is important to realize that the roles of Eloise and Abelard are not treated symmetrically here, for the following reason: a match of the satisfiability game can be viewed as a collection of “virtual matches” of the acceptance game played at once. We shall see that the combinatorial difficulties involved in the completeness proof all stem from a common source: choices made by Abelard will generally cause the number of virtual matches we need to consider to *multiply*, making the combinatorics of the game harder. For this reason, we want to take Eloise’s side as much as possible, and restrict the power of Abelard. In particular, since Abelard is in charge of conjunctions, we need to carefully control the shape of conjunctions that we introduce when we pre-process the automaton \mathbb{A} into \mathbb{A}^x .

Let us now turn to the construction of the auxiliary structure \mathbb{A}^x , for which we shall require the following observation.

Proposition 4.17 *For every modal \mathbb{X} -automaton \mathbb{A} positive in $x \in \mathbb{X}$, and any state $a \in A$, there are formulas θ_0^a and θ_1^a in which x does not appear, such that*

$$\Theta(a) \equiv_K (x \wedge \theta_0^a) \vee \theta_1^a$$

Proof. First rewrite $\Theta(a)$ as a disjunction

$$(x \wedge \psi_0) \vee \dots \vee (x \wedge \psi_n) \vee \psi'_0 \vee \dots \vee \psi'_m$$

where each ψ_i and each ψ'_j is a conjunction consisting of literals distinct from x and formulas of the form $\Box\pi, \Diamond\pi$. This is then equivalent to

$$(x \wedge (\psi_0 \vee \dots \vee \psi_n)) \vee (\psi'_0 \vee \dots \vee \psi'_m)$$

and so we are done. QED

Convention 4.18 Relying on the previous observation, we fix from now on for every automaton \mathbb{A} and $a \in A$, one-step formulas θ_0^a, θ_1^a such that $\Theta(a) \equiv_K (x \wedge \theta_0^a) \vee \theta_1^a$.

The construction of \mathbb{A}^x is based on the following four ideas. First, since we do not formally allow proposition letters to appear guarded in the one-step formulas in the image of the transition map of an automaton, we introduce a new state \underline{x} that we use to represent the variable x , in the sense that we put $\Theta^x(\underline{x}) := x$. Second, we will “split” each state a into two states a_0 and a_1 , taking care of the θ_0^a - and the θ_1^a -part of $\Theta(a)$, respectively. Thus we define $A^x := (A \times \{0, 1\}) \cup \{\underline{x}\}$. Third, after this “change of base” of the automaton, we need to ensure that the transition map Θ^x has the right co-domain (A^x). We can take care of this by *substituting*, in every one-step formula $\alpha \in \mathbf{1ML}(\mathbf{X}, A)$, each occurrence of a state a by the formula $(\underline{x} \wedge a_0) \vee a_1$. We shall denote the resulting substitution as $\kappa : A \rightarrow A^x$. Fourth, while we are mostly interested in the underlying automaton structure $(A^x, \Theta^x, \Omega^x)$ of \mathbb{A}^x , we do need to assign it an initial state. Our choice of $(a_I)_1$ is guided by the role of \mathbb{A}^x in the proof of our main Lemma, Theorem 6.

Definition 4.19 Let \mathbb{A} be any modal \mathbf{X} -automaton which is positive in $x \in \mathbf{X}$, and assume without loss of generality that the smallest priority in the image of Ω is greater than 0 (otherwise just start by raising all priorities in \mathbb{A} by 2). Pick a new state $\underline{x} \notin A$. Then we define the \mathbf{X} -automaton $\mathbb{A}^x = (A^x, \Theta^x, \Omega^x, a_I^x)$ as follows:

- $A^x := (A \times \{0, 1\}) \cup \{\underline{x}\}$. We write (a, i) as a_i , for $i \in \{0, 1\}$.
- $\Theta^x(a_0) := \theta_0^a[\kappa]$ and $\Theta^x(a_1) := \theta_1^a[\kappa]$,
- $\Theta^x(\underline{x}) := x$,
- $a_I^x := (a_I)_1$,
- $\Omega^x(a_i) := \Omega(a)$ and $\Omega^x(\underline{x}) := 0$.

Here, κ is defined to be the substitution $a \mapsto (\underline{x} \wedge a_0) \vee a_1$. ◁

Note that the substitution κ involved in this construction does introduce new conjunctions, but in a very controlled manner: the only new conjunctions are of the form $\underline{x} \wedge a_0$ for $a \in A$, i.e., we don't introduce any conjunctions between states a_i , for $a \in A$. This would not be the case if we worked for example with the dual substitution $\kappa^\partial : a \mapsto (\underline{x} \vee a_0) \wedge a_1$. So the pre-processing of \mathbb{A} into \mathbb{A}^x has indeed been set up in such a way that conjunctions are of a restricted shape, and this is crucial.

Remark 4.20 The automaton \mathbb{A}^x is *not equivalent* to \mathbb{A} , in the sense that it does not accept the same pointed Kripke structures as \mathbb{A} does. On the other hand, it does contain all information that \mathbb{A} does, and vice versa. The precise connection between \mathbb{A} and \mathbb{A}^x can best be expressed using the *translation map* that we will define in section 8. Running ahead of this, assume that we have defined, for each modal automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$ a map $\text{tr}_{\mathbb{A}} : A \rightarrow \mu\text{ML}$ assigning to each state $a \in A$ an *equivalent* μ -calculus formula $\text{tr}_{\mathbb{A}}(a)$ in the sense that $\mathbb{A}\langle a \rangle \equiv \text{tr}_{\mathbb{A}}(a)$ for each state a .

Phrased in terms of this translation map, the relation between \mathbb{A} and \mathbb{A}^x is given by the equivalences

$$\text{tr}_{\mathbb{A}}(a) \equiv (x \wedge \text{tr}_{\mathbb{A}^x}(a_0)) \vee \text{tr}_{\mathbb{A}^x}(a_1)$$

and

$$\text{tr}_{\mathbb{A}^x}(a_i) \equiv \theta_i^a[\text{tr}_{\mathbb{A}}(b)/b \mid b \in A]$$

which hold for all $a \in A$ and $i \in \{0, 1\}$. ◁

We now turn to the definition of the automata $\mu x.\mathbb{A}$ and $\nu x.\mathbb{A}$; both constructs are variations of the auxiliary structure \mathbb{A}^x . The key to understanding the definitions, and to proving correctness of the construction is the following proposition. We shall make use of it later on, when we consider the converse translation from automata to formulas. Since there we will be concerned with *provable* equivalence, we formulate the next two propositions using the relation \equiv_K rather than the semantic equivalence relation \equiv . Note that the semantic versions of the statements follow by the soundness of the axiom system.

Proposition 4.21 *Let φ_0, φ_1 be any formulas in which the variable x appears positively. Then:*

$$\mu x.(x \wedge \varphi_0) \vee \varphi_1 \equiv_K \mu x.\varphi_1$$

and

$$\nu x.(x \wedge \varphi_0) \vee \varphi_1 \equiv_K \nu x.\varphi_0 \vee \varphi_1$$

Proof. We consider the case for μ first. One direction of the equivalence is immediate, since we have $\varphi_1 \leq_K (x \wedge \varphi_0) \vee \varphi_1$. For the converse, we show that $\mu x.\varphi_1$ is a pre-fixpoint for the formula $(x \wedge \varphi_0) \vee \varphi_1$. To see this, we have:

$$\begin{aligned} ((x \wedge \varphi_0) \vee \varphi_1)[\mu x.\varphi_1/x] &\equiv_K ((\mu x.\varphi_1) \wedge \varphi_0[\mu x.\varphi_1/x]) \vee \varphi_1[\mu x.\varphi_1/x] \\ &\equiv_K ((\mu x.\varphi_1) \wedge \varphi_0[\mu x.\varphi_1/x]) \vee \mu x.\varphi_1 \\ &\equiv_K \mu x.\varphi_1 \end{aligned}$$

For the ν -case, again one direction is immediate since we have $(x \wedge \varphi_0) \vee \varphi_1 \leq_K \varphi_0 \vee \varphi_1$. For the other direction we need to show that $\nu x.\varphi_0 \vee \varphi_1$ is a post-fixpoint for $(x \wedge \varphi_0) \vee \varphi_1$. We reason as follows:

$$\begin{aligned} \nu x.\varphi_0 \vee \varphi_1 &\equiv_K (\nu x.\varphi_0 \vee \varphi_1) \wedge (\nu x.\varphi_0 \vee \varphi_1) \\ &\equiv_K (\nu x.\varphi_0 \vee \varphi_1) \wedge (\varphi_0[\nu x.\varphi_0 \vee \varphi_1/x] \vee \varphi_1[\nu x.\varphi_0 \vee \varphi_1/x]) \\ &\leq_K ((\nu x.\varphi_0 \vee \varphi_1) \wedge \varphi_0[\nu x.\varphi_0 \vee \varphi_1/x]) \vee \varphi_1[\nu x.\varphi_0 \vee \varphi_1/x] \\ &= ((x \wedge \varphi_0) \vee \varphi_1)[\nu x.\varphi_0 \vee \varphi_1/x] \end{aligned}$$

and the proof is finished. QED

We are now ready to define fixpoint operations on automata in which the proposition letter x appears positively:

Definition 4.22 Let \mathbb{A} be any modal \mathbf{X} -automaton which is positive in $x \in \mathbf{X}$. The $\mathbf{X} \setminus \{x\}$ -automaton $\mu x.\mathbb{A} = (A', \Theta', \Omega', a'_I)$ is defined by setting:

- $A' := A^x$
- $\Theta'(a_i) := \Theta^x(a_i)$ for $a \in A$
- $\Theta'(\underline{x}) := \theta_1^{a_I}[\kappa]$
- $a'_I := \underline{x}$
- $\Omega'(a_i) := \Omega^x(a_i)$ and $\Omega^x(\underline{x}) := 2 \cdot \max(\Omega^x[A^x]) + 1$.

Similarly, the $\mathbf{X} \setminus \{x\}$ -automaton $\nu x.\mathbb{A} = (A', \Theta', \Omega', a'_I)$ is defined as follows:

- $A' := A^x$
- $\Theta'(a_i) := \Theta^x(a_i)$ for $a \in A$
- $\Theta'(\underline{x}) := \theta_0^{a_I}[\kappa] \vee \theta_1^{a_I}[\kappa]$
- $a'_I := \underline{x}$
- $\Omega'(a_i) := \Omega^x(a_i)$ and $\Omega^x(\underline{x}) := 2 \cdot \max(\Omega^x[A^x]) + 2$. \triangleleft

Remark 4.23 We finish this subsection with noting that all the constructions defined above are semantically correct, in the sense that $L(\mathbb{A} \wedge \mathbb{B}) = L(\mathbb{A}) \cap L(\mathbb{B})$, etc. We leave the formal proofs of these statements, which follow from our later results, as exercises to the reader. \triangleleft

4.4 Translating formulas to automata

We finish this section on modal automata by providing a translation associating an equivalent modal parity automaton with every μ -calculus formula. As mentioned in the introduction to this section, our definition will proceed by induction on the complexity of formulas, applying the operations that we just defined to handle the inductive cases of this definition.

Definition 4.24 By induction on the complexity of a modal μ -formula φ we define a modal automaton \mathbb{A}_φ .

First of all, we need to consider atomic formulas: given any propositional variable p , we take some arbitrary object a distinct from p to be the one and only state of \mathbb{A}_p , and define $\Theta_p(a) = p$, and $\Omega_p(a) = 0$.

With this in place, we can complete the translation as follows:

$$\begin{aligned}
 \mathbb{A}_{\neg\varphi} &:= \neg\mathbb{A}_\varphi \\
 \mathbb{A}_{\varphi\vee\psi} &:= \mathbb{A}_\varphi \vee \mathbb{A}_\psi \\
 \mathbb{A}_{\diamond\varphi} &:= \diamond\mathbb{A}_\varphi \\
 \mathbb{A}_{\mu x.\varphi} &:= \mu x.\mathbb{A}_\varphi,
 \end{aligned}$$

i.e., by applying the operations we have defined above to handle the various connectives of the μ -calculus. \triangleleft

We finish by stating the semantic correctness of this definition. Since this proposition is not needed in the sequel, we leave the details of its proof, which proceeds by a routine induction on the complexity of formulas, as an exercise to the reader.

Proposition 4.25 *Let φ be a formula of the modal μ -calculus. Then*

$$\varphi \equiv \mathbb{A}\varphi. \tag{1}$$

5 Games for automata

5.1 Introduction

In this section we introduce two of our main tools, viz., the satisfiability game $\mathcal{S}(\mathbb{A})$ related to a modal automaton \mathbb{A} , and the consequence game $\mathcal{C}(\mathbb{A}, \mathbb{B})$ related to two automata \mathbb{A} and \mathbb{B} . Before we turn to the technicalities of the definitions we start with an intuitive explanation of the satisfiability game.

The satisfiability game $\mathcal{S}(\mathbb{A})$, introduced in [22] in the more general setting of the coalgebraic μ -calculus, can be seen as a streamlined, game-theoretic analog for automata to what tableaux are for formulas. To understand the game, which is played by two players, \forall and \exists , it helps to think of \exists as aiming to construct a (tree) structure for the automaton \mathbb{A} — in fact, we will provide close connections between \exists 's winning strategies in $\mathcal{S}(\mathbb{A})$ and models of \mathbb{A} in the proof of Proposition 5.10. The role of \forall in $\mathcal{S}(\mathbb{A})$ is rather different: he acts as a *path finder* in the (partial) structure constructed by \exists , his task being to challenge \exists to construct ever more detail of the structure. What distinguishes the satisfiability game from tableaux is that, because of the uniform internal structure of modal automata as compared to formulas, the interaction between the two players can be shaped in a highly regulated pattern. The satisfiability game does not have separate rules dealing with specific connectives; in particular, all rules/moves dealing with Boolean connectives have been encapsulated in the streamlined interaction between \exists and \forall .

For two reasons, it is also useful to relate the satisfiability game $\mathcal{S}(\mathbb{A})$ to the acceptance games associated with \mathbb{A} . First, similar to the acceptance games for \mathbb{A} , the satisfiability game proceeds in *rounds*: one round of $\mathcal{S}(\mathbb{A})$ consists of first \exists constructing (or aiming to construct) one more level of the tree structure for \mathbb{A} , and then \forall picking one of the newly created nodes for further inspection. Second, and more in particular, every match of $\mathcal{S}(\mathbb{A})$ can be seen as a *bundle* of matches of the acceptance game played on exactly the structure that \exists is constructing.

In somewhat more detail, positions of the satisfiability game $\mathcal{S}(\mathbb{A})$ will represent *macro-states* of \mathbb{A} , that is, subsets of the state space A . Intuitively, a partial $\mathcal{S}(\mathbb{A})$ -match Σ ending in a position p corresponds to a node t_Σ in the Kripke structure \mathbb{T} that \exists is constructing, and if p represents the macro-state B , this indicates that the subtree starting at t_Σ needs to be accepted by the automaton $\mathbb{A}\langle b \rangle$ for *each* $b \in B$.

Concretely, what \exists has to come up with in such a partial match Σ is (1) a one-step frame (Y, W) that will be the local continuation of the tree structure \mathbb{T} (in the sense that the coalgebraic unfolding at t_Σ can be defined as $\sigma_{\mathbb{T}}(t_\Sigma) := (Y, W)$), together with (2) a local move m_a at position (t_Σ, a) in the acceptance game $\mathcal{A}(\mathbb{A}, \mathbb{T})$ for *each* of the states $a \in B$, comprising a *family* of markings $m_a : W \rightarrow \mathcal{P}A$ such that $(Y, W, m_a) \Vdash^1 \Theta(a)$ for each $a \in B$.¹ It is then up to \forall to pick an element w of W , after which the satisfiability game is continued at a position representing the macro-state $\bigcup \{m_a(w) \mid a \in B\}$.

Crucial in our definition of the satisfiability game is the particular nature of the carrier sets W of the one-step frames that \exists can pick: We will restrict \exists 's moves to one-step frames

¹Note that she does not need one single marking m taking care of all the states $a \in B$ at once, in the sense that $(Y, W, m) \Vdash^1 \bigwedge \{\Theta(a) \mid a \in B\}$.

of the form (Y, \mathcal{R}) where $\mathcal{R} \subseteq A^\sharp$ is a collection of *binary relations* on A . One advantage of this approach is that the family of markings on \mathcal{R} (needed as the second ingredient of \exists 's move) is naturally induced by the first ingredient: the local strategy associated with a state a can simply be defined as the *natural a -marking*

$$n_a : R \mapsto R[a],$$

mapping an arbitrary element $R \in \mathcal{R}$ to the set of its a -successors. In other words, pairs of the form $\Gamma = (Y, \mathcal{R})$ provide both a one-step frame *and* a family of markings on this very same frame. For any $R \in \mathcal{R}$, the intuitive understanding of $(a, b) \in R$ is that b holds at R ‘because of’ a .

Thus our set-up will be as follows. Binary relations over A , i.e., elements of the set A^\sharp , will provide the *basic* positions of $\mathcal{S}(\mathbb{A})$. The macro-state represented by a basic position $R \in A^\sharp$ is simply given as the range $\text{Ran}R \subseteq A$ of R . The basic positions are the ones where \exists has to move, and her set of admissible moves at position R will consist of those elements $\Gamma = (Y, \mathcal{R}) \in \text{KA}^\sharp = \text{PX} \times \text{PA}^\sharp$ that provide legitimate moves in all associated acceptance games, in the sense that $(Y, \mathcal{R}, n_a) \Vdash^1 \Theta(a)$ for all $a \in \text{Ran}R$. Positions of the form (Y, \mathcal{R}) are for \forall , and his set of admissible move is simply given by \mathcal{R} itself, that is, the moves available to him are provided by the binary relations in \mathcal{R} .² A *round* of the satisfiability game thus starts at a basic position $R \in A^\sharp$, and consists of \exists choosing a suitable one-step model $\Gamma = (Y, \mathcal{R}) \in \text{KA}^\sharp$, followed by \forall picking a next basic position $Q \in \mathcal{R}$.

A second advantage of our approach using binary relations shows up when, finally, we consider the winning condition of the satisfiability game. Clearly, the sequence of basic positions of an infinite $\mathcal{S}(\mathbb{A})$ -match Σ provides an A^\sharp -stream, that is, a stream of binary relations over A . As we mentioned before, the idea is that Σ corresponds to a *bundle* of matches of the acceptance game relating \mathbb{A} to the tree structure \exists has been constructing. The requirement that \exists needs to win *all* of these matches can be nicely expressed in terms of the collection of *traces* over the A^\sharp -stream associated with Σ .

5.2 Traces

We first need some notation and terminology concerning streams of binary relations and the traces they carry. Coming back to the title of our paper, this is where the combinatorics of our proof will be located.

Definition 5.1 Fix a set A . We let A^\sharp denote the set of binary relations over A , that is, $A^\sharp := \text{P}(A \times A)$.

Given a finite word $\Sigma = R_1 R_2 R_3 \dots R_k$ over the set A^\sharp , a *trace* through Σ is a finite A -word $\alpha = a_0 a_1 a_2 \dots a_k$ such that $a_i R_{i+1} a_{i+1}$ for all $i < k$. A *trace* through a A^\sharp -stream $\Sigma = R_1 R_2 R_3 \dots$ is an A -stream $\alpha = a_0 a_1 a_2 \dots$, such that $a_i R_{i+1} a_{i+1}$ for all $i < \omega$). In both cases we denote the set of traces through Σ as Tr_Σ .

Given a stream $\Sigma = R_1 R_2 R_3 \dots$ over A^\sharp we denote by $\Sigma|_k$ the word $R_1 \dots R_k$, and for a trace $\tau = a_0 a_1 a_2 \dots$ on Σ we denote by $\tau|_k$ the restricted trace $a_0 \dots a_k$ on $\Sigma|_k$. We use similar notation for restrictions of words over A^\sharp of length $\geq k$. \triangleleft

²For technical reasons, the actual definition of \forall 's moves in $\mathcal{S}(\mathbb{A})$ is a slight modification of this, see Remark 5.7.

It is often convenient to think of the set of finite traces providing a graph structure. Formally we define the trace graph of an A^\sharp -stream as follows. Observe that the infinite Σ -traces are in 1-1 correspondences with the maximal infinite paths through this graph.

Definition 5.2 Given an A^\sharp -stream $\Sigma = (R_n)_{n \geq 1}$, we define the *trace graph* \mathbb{G}_Σ as the directed graph with vertices $\omega \times A$ and edges $E_{\mathbb{G}} := \{((i, a), (j, b)) \mid j = i + 1 \text{ and } R_i ab\}$. \triangleleft

Definition 5.3 Fix a finite set A and a priority map $\Omega : A \rightarrow \omega$. We let NBT_Ω denote the set of A^\sharp -streams that contain no bad trace, that is, no trace $\tau = a_0 a_1 \dots$ such that $\max(\Omega[\text{Inf}(\tau)])$, the highest priority occurring infinitely often on τ , is odd. \triangleleft

It is not difficult to show that NBT_Ω is an ω -regular subset of $(A^\sharp)^\omega$.

Proposition 5.4 *Given a finite set A and a priority map $\Omega : A \rightarrow \omega$, there is a parity stream automaton recognizing the set NBT_Ω , seen as a stream language over A^\sharp .*

Proof. It is easy to construct a nondeterministic parity stream automaton \mathbb{A} recognizing the complement of NBT_Ω , that is, the set of A^\sharp -streams that do contain a bad trace. The Proposition is then immediate by the fact that the collection of ω -regular language is closed under taking complementation. QED

5.3 The satisfiability game

We are now ready for the formal definition of the satisfiability game. First of all we consider the one-step models based on the set A^\sharp of binary relations over A .

Definition 5.5 Given a set A , the *natural a -marking* on the set A^\sharp is defined as the map $n_a : A^\sharp \rightarrow \text{PA}$ given by

$$n_a : R \mapsto R[a].$$

Its transpose, i.e., the corresponding *natural a -valuation* $U_a : A \rightarrow \text{PA}^\sharp$ is given by

$$U_a : b \mapsto \{R \in A^\sharp \mid (a, b) \in R\}.$$

Any object $\Gamma = (Y, \mathcal{R}) \in \text{KA}^\sharp$ can be seen as a one-step model, by restricting the valuation n_a to the domain $\mathcal{R} \subseteq A^\sharp$. For a one-step formula $\alpha \in \mathbf{1ML}(X, A)$, we write $\Gamma \Vdash_a^{-1} \alpha$ to denote that $A^\sharp, n_a, \Gamma \Vdash^{-1} \alpha$, and we define $\llbracket \alpha \rrbracket_a^{-1} := \{\Gamma \in \text{KA}^\sharp \mid \Gamma \Vdash_a^{-1} \alpha\}$.

Given an object $\Gamma \in \text{KA}^\sharp$, we let Y_Γ and \mathcal{R}_Γ denote the unique objects such that $\Gamma = (Y_\Gamma, \mathcal{R}_\Gamma)$. \triangleleft

Clearly then, we may indeed think of $\Gamma \in \text{KA}^\sharp$ as a *family* of one-step models *on the same one-step frame*.

Definition 5.6 Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be a modal automaton. Then the *satisfiability game* $\mathcal{S}(\mathbb{A})$ is the graph game of which the moves are given by Table 1. Positions of the form $R \in A^\sharp$ are called *basic*.

The winner of an infinite match of the satisfiability game is given by the induced stream $\Sigma = R_0 R_1 \dots \in (A^\sharp)^\omega$ of basic positions. This winner is \exists if Σ belongs to the set NBT_Ω , that is, if Σ contains no bad traces, and it is \forall otherwise. A winning strategy of \forall in $\mathcal{S}(\mathbb{A})$ may be called a *refutation* of \mathbb{A} . \triangleleft

Position	Player	Admissible moves
$R \in A^\sharp$	\exists	$\bigcap_{a \in \text{Ran} R} \llbracket \Theta(a) \rrbracket_a^1$
$(Y, \mathcal{R}) \in \text{KA}^\sharp$	\forall	$\{Q \in A^\sharp \mid Q \subseteq R \text{ for some } R \in \mathcal{R}\}$

Table 1: Admissible moves in the satisfiability game $\mathcal{S}(\mathbb{A})$

Remark 5.7 An alternative and perhaps more natural version of $\mathcal{S}(\mathbb{A})$ would restrict the moves available to \forall at position $(Y, \mathcal{R}) \in \text{KA}^\sharp$ to the actual *elements* of \mathcal{R} , instead of allowing subsets of elements of \mathcal{R} . It is not so difficult to prove, however, that this version of the game is in fact equivalent to $\mathcal{S}(\mathbb{A})$ itself. Roughly, the reason for this is that in $\mathcal{S}(\mathbb{A})$ it never will be to \forall 's advantage at a position $(Y, \mathcal{R}) \in \text{KA}^\sharp$ to pick a *strict* subset Q of some relation $Q' \in \mathcal{R}$: the bigger the relations that he picks, the more opportunities he has to obtain a bad trace.

Our motivation for taking $\mathcal{S}(\mathbb{A})$ as the standard version of our satisfiability game is simply that in some cases $\mathcal{S}(\mathbb{A})$ is technically more convenient to work with than its apparently simpler variant. \triangleleft

Remark 5.8 It may be useful to cover three special cases of \exists 's move at a position $R \in A^\sharp$.

First, suppose that $R = \emptyset$, that is, R is the empty relation. In this case we have that $\bigcap_{a \in \text{Ran} R} \llbracket \Theta(a) \rrbracket_a^1 = \text{KA}^\sharp$, so that \exists could pick any pair of the form (Y, \mathcal{R}) where $\mathcal{R} = \emptyset$ and win (as we will see right now).

Second, consider the situation where \exists picks a pair (Y, \mathcal{R}) where $\mathcal{R} = \emptyset$. This is the one-step version of a ‘blind world’ (a state in a Kripke that has no successors) and thus such a move is required in case one of the one-step formulas contains or implies the formula $\Box \perp$. Observe that any position of the form (Y, \emptyset) is *winning* for \exists since in the next move it forces \forall to pick an element from the empty set.

Finally, the situation where $\emptyset \in \mathcal{R}$, that is, where \mathcal{R} *contains* the empty set, is different again. In this case, the empty relation would be available as a move to \forall , and as a next basic position after R , no trace would be continued. But as we have just seen, should \forall indeed pick the empty relation, then he would loose at the very next step of the match. \triangleleft

Convention 5.9 Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be a modal automaton. Since we will only consider matches of the satisfiability game $\mathcal{S}(\mathbb{A})$ that take the singleton $\{(a_I, a_I)\}$ as their starting position, we will often be sloppy and blur the difference between $\mathcal{S}(\mathbb{A})$ and the initialized game $\mathcal{S}(\mathbb{A})@_{\{(a_I, a_I)\}}$.

The following proposition expresses the adequacy of the satisfiability game. Although this proposition is not needed for proving the main result of this paper, we sketch its proof since this may be useful to obtain further intuitions on the satisfiability game. It is here that we see the tight connection between \exists 's winning strategies in the satisfiability game, and models for the automaton.

Proposition 5.10 (Adequacy) *Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be a modal automaton \mathbb{A} . Then \exists has a winning strategy in $\mathcal{S}(\mathbb{A})$ iff the language recognized by \mathbb{A} is non-empty.*

Proof. For the direction from left to right, assume that \exists has a winning strategy χ in the satisfiability game $\mathcal{S}(\mathbb{A})$ starting at position $R_I := \{(a_I, a_I)\}$. Given the nature of the winning condition of $\mathcal{S}(\mathbb{A})$, it is obvious that without loss of generality we may take $\chi(\Sigma)$ to only depend on the subsequence of *basic* positions of the partial match $\Sigma \in \text{PM}_{\exists}^X @ R_I$.

We will now show that basically, χ itself can be seen as a Kripke structure. Define $\text{PM}_{\exists}^X @ R_I$ as the set of partial χ -guided $\mathcal{S}(\mathbb{A})$ -matches starting at R_I and ending in a position for \exists , and let S be the set of basic reducts of these matches, i.e.,

$$S := \{R_0 R_1 \dots R_k \mid R_0 \Gamma_1 R_1 \dots \Gamma_k R_k \in \text{PM}_{\exists}^X @ R_I\}.$$

\exists 's strategy can then be seen as a map $\chi : S \rightarrow \text{KA}^{\sharp}$ and so it naturally induces a map $\sigma_{\chi} : S \rightarrow \text{KS}$ given by

$$\sigma_{\chi}(\Sigma) := \left(\mathbf{Y}_{\chi(\Sigma)}, \{\Sigma \cdot Q \mid Q \in \mathcal{R}_{\chi(\Sigma)}\} \right),$$

where the notations $\mathbf{Y}_{\chi(\Sigma)}$ and $\mathcal{R}_{\chi(\Sigma)}$ are as introduced at the end of Definition 5.5. But then the pair (S, σ_{χ}) is (the coalgebraic representation of) a Kripke structure \mathbb{S}_{χ} . We leave it as an exercise for the reader to check, finally, that \mathbb{A} accepts the pointed Kripke structure (\mathbb{S}, R_I) .

For the opposite direction, from right to left, assume that \mathbb{A} accepts the pointed Kripke structure (\mathbb{S}, s_1) , where $\mathbb{S} = (S, R, V)$. Since the acceptance game is a parity game, by positional determinacy (Fact 2.9) we may assume that \exists has a positional winning strategy m starting at (a_I, s_1) (or at any winning position, for that matter). This strategy assigns to each pair $(a, s) \in \text{Win}_{\exists}$ an A -marking $m_{a,s}$ on $R[s]$ such that the induced one-step model satisfies $(V^{\dagger}(s), R[s], m_{a,s}) \Vdash^1 \Theta(a)$ (*).

We will now use this positional strategy m to define a strategy χ for \exists in $\mathcal{S}(\mathbb{A}) @ R_I$, where $R_I := \{(a_I, a_I)\}$. We will define χ by induction on the length of a partial χ -guided match $\Sigma = R_0 \dots R_k$, where $R_0 = R_I$. By a simultaneous induction, with any such match we will associate a path $s_0 \dots s_k$ through \mathbb{S} such that every trace $a_I a_I a_1 \dots a_k$ on $R_0 \dots R_k$ corresponds to an m -guided match $(a_I, s_0)(a_1, s_1) \dots (a_k, s_k)$ of the acceptance game. Clearly, if we can maintain this condition indefinitely, \exists will be the winner of the resulting infinite match, by our assumption that her strategy m is winning in $\mathcal{A}(\mathbb{A}, \mathbb{S}) @ (a_I, s_0)$. Hence, all we need to show is that \exists can maintain the inductive condition one round.

To see how to do this, consider a partial χ -guided match $\Sigma = R_0 \dots R_k$, where $R_0 = R_I$. By the inductive hypothesis it follows that all states $a \in \text{Ran} R_k$ are such that $(a, s_k) \in \text{Win}_{\exists}(\mathcal{A}(\mathbb{A}, \mathbb{S}))$. In particular, \exists 's winning strategy m provides a marking $m_{a, s_k} : R[s_k] \rightarrow \text{PA}$ (recall that R denotes the accessibility relation of the Kripke structure \mathbb{S}). Given any successor $t \in R[s_k]$, set $Q_t := \{(a, b) \mid b \in m_{a, s_k}(t)\}$, and let

$$\chi(\Sigma) := ((V^{\dagger}(s_k), \{Q_t \mid t \in R[s_k]\}))$$

be the definition of \exists 's strategy χ . Observe that each move $Q \in A^{\sharp}$ of \forall in response to the position $\chi(\Sigma)$ is by definition of the form Q_t for some successor t of s_k . We then set s_{k+1} to be any such t , and leave the routine verification that the partial match $\Sigma \cdot \chi(\Sigma) \cdot Q$ and the \mathbb{S} -path $s_0 \dots s_k s_{k+1}$ satisfy the condition (*) as an exercise to the reader. QED

Remark 5.11 In general $\mathcal{S}(\mathbb{A})$ is not a parity game, but we saw in Proposition 5.4 that the winning condition NBT_Ω is an ω -regular subset of $(A^\sharp)^\omega$. It follows from a result by Büchi & Landweber [10] that we may assume that winning strategies in $\mathcal{S}(\mathbb{A})$ only use finite memory. This observation can be used to prove the finite model property of modal automata, and hence, of the modal μ -calculus, cf. [22] for a proof in the more general setting of the coalgebraic μ -calculus. \triangleleft

In the sequel it will often be convenient to make some simplifying assumptions on the moves picked by \exists . Most of these assumptions can be justified by the observation that it is in \exists 's interest to keep the set of traces in an $\mathcal{S}(\mathbb{A})$ -match as small as possible. One way to formulate this more precisely is the following.

Proposition 5.12 *Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be a modal automaton, and let $\mathcal{N} \subseteq A^\sharp$ be some set of relations. Assume that for every basic position $R \in A^\sharp$ of the satisfiability game, and every legitimate move (Y, \mathcal{R}) of \exists there is a legitimate move (Y, \mathcal{R}') such that $\mathcal{R}' \subseteq \mathcal{N}$ and $\mathcal{R}' \xrightarrow{\vec{P}} \subseteq \mathcal{R}$. Then for any winning position in $\mathcal{S}(\mathbb{A})$ \exists has a winning strategy that restricts her moves to pairs (Y, \mathcal{R}) with $\mathcal{R} \subseteq \mathcal{N}$.*

Proof. Assume that \exists has a winning strategy Γ in the game $\mathcal{S}(\mathbb{A})$ initialized at position R_0 . We need to provide her with a winning \mathcal{N} -strategy, that is, a strategy Γ' that always selects moves (Y, \mathcal{R}) with $\mathcal{R} \subseteq \mathcal{N}$.

We will define this strategy Γ' by induction on the length of partial $\mathcal{S}(\mathbb{A})$ -matches. Simultaneously, for any such match

$$\Sigma = R_0 \Gamma_0 R_1 \Gamma_1 \dots R_k$$

which is Γ' -guided, we will define a parallel match

$$\Sigma^* = R_0 \Gamma_0^* R_1 \Gamma_1^* \dots R_k$$

which is guided by \exists 's winning strategy Γ . If we can maintain such a shadow match infinitely long, is routine to prove that Γ' is winning for \exists .

For the case where $k = 0$ there is nothing to prove, so assume inductively that there are matches Σ and Σ^* as above. Observe that since the last positions of Σ and Σ^* are identical, the set of \exists 's legitimate moves in Σ and Σ^* are the same. Let (Y, \mathcal{R}) be the move prescribed by \exists 's winning strategy Γ in the partial match Σ^* , then by assumption there is a legitimate move (Y, \mathcal{R}') such that $\mathcal{R}' \subseteq \mathcal{N}$ and $\mathcal{R}' \xrightarrow{\vec{P}} \subseteq \mathcal{R}$. Then we let

$$\Gamma'_\Sigma := (Y, \mathcal{R}')$$

be \exists 's move in Σ . This defines the strategy Γ' .

To finish the inductive step, consider an arbitrary continuation of the match $\Sigma \cdot (Y, \mathcal{R}')$, say, where \forall plays some relation Q . By definition, Q is a subset of some $Q' \in \mathcal{R}'$, while by $\mathcal{R}' \xrightarrow{\vec{P}} \subseteq \mathcal{R}$ we find some $Q'' \in \mathcal{R}$ such that $Q' \subseteq Q''$. But then it follows from $Q \subseteq Q''$ that Q is also a legitimate move for \forall in $\Sigma^* \cdot (Y, \mathcal{R})$. In other words, the two $k + 1$ -length matches $\Sigma \cdot (Y, \mathcal{R}') \cdot Q$ and $\Sigma^* \cdot (Y, \mathcal{R}) \cdot Q$ satisfy the required conditions. QED

Remark 5.13 As a consequence of Proposition 5.12, we can always make some minimality assumptions on \exists 's strategy in the satisfiability game. In particular, suppose that \exists , at some position $R \in A^\sharp$ in a match of $\mathcal{S}(\mathbb{A})$, picks a move $\Gamma = (\mathcal{Y}, \mathcal{R})$. Then without loss of generality we can assume that:

- (1) $\text{Dom}(Q) \subseteq \text{Ran}(R)$, for all $Q \in \mathcal{R}$.
- (2) b occurs in $\Theta(a)$, for all $Q \in \mathcal{R}$ and $(a, b) \in Q$.

We leave it for the reader to verify this claim. \triangleleft

Of course, we will need to connect the satisfiability game to Kozen's proof system for the μ -calculus to be able to make use of it in the completeness proof. Recall from our discussion surrounding Theorem 2 in the introduction to this paper that in our completeness proof we will apply proof-theoretic notions, and in particular that of consistency, to automata (for the details, see section 8). Ideally then, we would want that whenever an automaton is consistent, we can find a winning strategy for \exists in the associated satisfiability game. This will indeed follow from the completeness theorem, together with the adequacy of the satisfiability game, but it is very hard to verify directly. In Section 9 we will establish the result for a special class of automata, called semi-disjunctive automata, that will be introduced in Section 6.

But we can already say something about the connection between the satisfiability game and provability: the following corollary of Theorem 4.9 connects the one-step completeness theorem to the satisfiability game.

Corollary 5.14 *Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be a modal automaton, let $R \in A^\sharp$ be a position in $\mathcal{S}(\mathbb{A})$ and let $\{\sigma_b \mid b \in \text{Ran}R\}$ be a family of substitutions such that the formula*

$$\bigwedge_{b \in \text{Ran}R} \Theta(b)[\sigma_b]$$

is consistent. Then \exists has an admissible move $(\mathcal{Y}, \mathcal{R})$ such that for all $Q \in \mathcal{R}$ the formula

$$\bigwedge_{(b,d) \in Q} \sigma_b(d)$$

is consistent.

Proof. We will derive this from Theorem 4.9. The idea is to consider pairs $(b, d) \in A \times A$ as one-step variables, so that we may combine the family of substitutions $\{\sigma_b \mid b \in \text{Ran}R\}$ into one single substitution over $A \times A$.

In detail, let $\lambda_b : A \rightarrow A \times A$ denote the substitution $\lambda_b : d \mapsto (b, d)$ that "tags" an arbitrary variable d with the state b , and let $\sigma : A \times A \rightarrow \mu\text{ML}(\mathbf{X})$ be the substitution given by putting

$$\sigma(b, d) := \sigma_b(d),$$

then clearly we have that

$$\sigma_b = \sigma \circ \lambda_b$$

for all b . Hence we may read the assumption as making a statement about the formula $\bigwedge\{\Theta(b)[\lambda_b] \mid b \in \text{Ran}R\} \in \mathbf{1ML}(X, A \times A)$ and the substitution σ , namely, that the formula

$$\left(\bigwedge\{\Theta(b)[\lambda_b] \mid b \in \text{Ran}R\}\right)[\sigma]$$

is consistent. It then follows from Theorem 4.9 that there is a one-step model (Y, S, m) such that $(Y, S, m) \Vdash^1 \bigwedge\{\Theta(b)[\lambda_b] \mid b \in \text{Ran}R\}$, and for all $s \in S$ the formula $\bigwedge\{\sigma(b, d) \mid (b, d) \in m(s)\}$ is consistent.

Note that $m : S \rightarrow A^\sharp$. It is then straightforward to verify that the move $(Y, m[S])$ satisfies the required properties. QED

As a special case of this result, consider the situation where the substitutions σ_b are all the same, given by $\sigma_b(a) = \text{tr}_\mathbb{A}(a)$ for each $b, a \in A$, where $\text{tr}_\mathbb{A} : A \rightarrow \mu\text{ML}$ denotes the map translating states in \mathbb{A} to their corresponding equivalent μML -formulas. Then Corollary 5.14 can be used to build a *surviving* strategy for \exists in the satisfiability game for a consistent automaton — in fact, all she has to do is to maintain the consistency of the formula $\bigwedge_{b \in \text{Ran}R} \text{tr}_\mathbb{A}(b)$.

But more importantly, Corollary 5.14 will be a key ingredient in our later of a winning strategy in the satisfiability game for a consistent semi-disjunctive automaton, and will allow us to maintain consistency of formulas recording some information about the traces of partial matches of the satisfiability game.

5.4 The consequence game

The consequence game for \mathbb{A} and \mathbb{A}' , $\mathcal{C}(\mathbb{A}, \mathbb{A}')$, is a graph game for two players, that we will simply call I and II. For convenience we will think of player I as being female, and player II as being male. One may think of the game being about player II trying to show that automaton \mathbb{A} *implies* \mathbb{A}' by establishing some kind of a *simulation relation* between the automata \mathbb{A} and \mathbb{A}' . Taking a more proof-theoretic perspective: with respect to a *basic* position in $\mathcal{C}(\mathbb{A}, \mathbb{A}')$ of the form $(R, R') \in A^\sharp \times A'^\sharp$, player II tries to show that ‘ R implies R' ’, in the sense that the conjunction of $\text{Ran}R$ implies the conjunction³ of $\text{Ran}R'$. Here we take the ‘conjunction of R ’ to be the conjunction of the set of automata $\{\mathbb{A}(b) \mid b \in \text{Ran}R\}$.

The consequence game is tightly linked to the satisfiability games of the two associated automata: $\mathcal{C}(\mathbb{A}, \mathbb{A}')$ also proceeds in rounds and these can be associated with rounds of the satisfiability games $\mathcal{S}(\mathbb{A})$ and $\mathcal{S}(\mathbb{A}')$. In words, one round of $\mathcal{C}(\mathbb{A}, \mathbb{A}')$ consists of three moves. At the start of the round, at basic position (R, R') , player I picks a local model $\Gamma \in A^\sharp$ for R , as if she was player \exists in $\mathcal{S}(\mathbb{A})$. Second, player II then maps this one-step model for R into a one-step model for R' , inducing a move in the game $\mathcal{S}(\mathbb{A}')$. Concretely, for a one-step model $\Gamma = (Y, \mathcal{R})$, player II provides a map $F : \mathcal{R} \rightarrow A'^\sharp$, turning Γ into the one-step model $(KF)\Gamma = (Y, \{FQ \mid Q \in \mathcal{R}\})$. Player I then finishes the round by picking a pair (Q, FQ) from the graph of F as the next basic position (note that in Table 2 such a map F is represented as a partial map from A^\sharp to A'^\sharp).

³This interpretation should be explicitly contrasted to the standard approach in proof theory regarding sequents, where $\Gamma \Rightarrow \Delta$ is interpreted as stating that the conjunction of Γ implies the *disjunction* of Δ .

Position	Player	Admissible moves
$(R, R') \in A^\sharp \times A'^\sharp$	I	$\bigcap_{a \in \text{Ran} R} \llbracket \Theta(a) \rrbracket_a^1 \times \{R'\}$
$(\Gamma, R') \in \text{KA}^\sharp \times A'^\sharp$	II	$\{F : \mathcal{R}_\Gamma \rightarrow A'^\sharp \mid (\text{KF})\Gamma \in \bigcap_{a' \in \text{Ran} R'} \llbracket \Theta(a') \rrbracket_{a'}^1\}$
$F : A^\sharp \xrightarrow{\circ} A'^\sharp$	I	$\text{Gr}(F)$

Table 2: Admissible moves in the consequence game $\mathcal{C}(\mathbb{A}, \mathbb{A}')$

The tight link with the satisfiability games extends to the winning conditions of $\mathcal{C}(\mathbb{A}, \mathbb{A}')$, which can be defined in terms of the winning conditions of $\mathcal{S}(\mathbb{A})$ and $\mathcal{S}(\mathbb{A}')$, since any infinite match of $\mathcal{C}(\mathbb{A}, \mathbb{A}')$ naturally induces infinite matches of the latter two games.

Remark 5.15 In fact, the consequence game can be seen as a kind of communication or implication game between the satisfiability games of the two automata involved. As such, the construction of $\mathcal{C}(\mathbb{A}, \mathbb{A}')$ from $\mathcal{S}(\mathbb{A})$ and $\mathcal{S}(\mathbb{A}')$ is vaguely reminiscent of the operation $\langle -, - \rangle$ on games, defined by Santocanale [60], where Santocanale’s construction in its turn is the result of enriching fixpoint theory with ideas from the game semantics of linear logic (see, e.g., Blass [7] or Joyal [33]). Note however that the actual moves of our game crucially involve modal one-step logic, in a way that makes $\mathcal{C}(\mathbb{A}, \mathbb{A}')$ rather different from the game $\langle \mathcal{S}(\mathbb{A}), \mathcal{S}(\mathbb{A}') \rangle$ one would obtain by applying Santocanale’s construction. \triangleleft

Definition 5.16 Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ and $\mathbb{A}' = (A', \Theta', \Omega', a'_I)$ be modal automata. Then the rules of the *consequence game* $\mathcal{C}(\mathbb{A}, \mathbb{A}')$ are given by Table 2. Positions of the form $(R, R') \in A^\sharp \times A'^\sharp$ are called *basic*. For the *winning conditions* of this game, consider an infinite match Σ of $\mathcal{C}(\mathbb{A}, \mathbb{A}')$, and let

$$(R_0, R'_0)(R_1, R'_1)(R_2, R'_2) \dots$$

be the induced stream of basic positions in Σ . Then player I is the winner of Σ if $R_0 R_1 \dots \in \text{NBT}_\Omega$ but $R'_0 R'_1 \dots \notin \text{NBT}_{\Omega'}$; that is, if there is a bad trace on the \mathbb{A}' -side but not on the \mathbb{A} -side.

If the position $(\{(a_I, a_I)\}, \{(a'_I, a'_I)\})$ is a winning position for player II in $\mathcal{C}(\mathbb{A}, \mathbb{A}')$, we say that \mathbb{A}' is a *game consequence* of \mathbb{A} , notation: $\mathbb{A} \models_{\mathcal{G}} \mathbb{A}'$. \triangleleft

The following proposition can be seen as stating a soundness result for the consequence game.

Proposition 5.17 *For any two modal automata \mathbb{A} and \mathbb{A}' it holds that*

$$\mathbb{A} \models_{\mathcal{G}} \mathbb{A}' \text{ implies } \mathbb{A} \models \mathbb{A}'. \quad (2)$$

Proof. Fix a pointed Kripke model (\mathbb{S}, s_I) with $\mathbb{S} = (S, R, V)$, a winning strategy χ for \exists in the acceptance game for \mathbb{A} with respect to \mathbb{S} at the start position (a_I, s_I) , and a winning strategy f for Player II in $\mathcal{C}(\mathbb{A}, \mathbb{A}')$. For simplicity we assume without loss of generality that the strategy χ is positional (recalling that $\mathcal{A}(\mathbb{A}, \mathbb{S})$ is a parity game). Our goal is to provide a

winning strategy χ' for \exists in the acceptance game for game \mathbb{A}' at the start position (a'_I, s_I) . By induction on the length of a χ' -guided match with basic positions $(a'_I, s_I), (a'_1, s_1) \dots (a'_n, s_n)$, we shall define an f -guided shadow match $(R_I, R'_I)(R_1, R'_1) \dots (R_n, R'_n)$ such that the following conditions hold:

- (1) $a_I a_1 \dots a_n$ is a trace through $R_I R_1 \dots R_n$ iff $(a_I, s_I)(a_1, s_1) \dots (a_n, s_n)$ is a χ -guided match; furthermore, each $b \in \text{Ran}(R_n)$ is the last element of some trace on $R_I R_1 \dots R_n$.
- (2) $a'_I a'_1 \dots a'_n$ is a trace through $R'_I R'_1 \dots R'_n$.

Furthermore, we shall associate these shadow matches in a uniform manner, so that the shadow match of an initial segment of a partial match Σ is an initial segment of the shadow match associated with Σ . First, note that this means that \exists wins all χ' -guided infinite matches: if $(a'_I, s_I)(a'_1, s_1)(a'_2, s_2) \dots$ is a loss for \exists then $a'_I a'_1 a'_2 \dots$ is a bad trace through $R'_I R'_1 R'_2 \dots$ in the shadow match $(R_I, R'_I)(R_1, R'_1)(R_2, R'_2) \dots$ by condition (1). Since this match was f -guided and f is a winning strategy, this means there must be some bad trace $a_I a_1 a_2 \dots$ through $R_I R_1 R_2 \dots$, and by condition (2) we get that $(a_I, s_I)(a_1, s_1)(a_2, s_2) \dots$ is an infinite χ -guided match, which furthermore is a loss for \exists . This is a contradiction since χ was a winning strategy by assumption.

We now show how \exists can respond to any move by \forall while maintaining the induction hypothesis. Suppose we are given a χ' -guided partial match Σ consisting of positions $(a'_I, s_I), (a'_1, s_1), \dots, (a'_n, s_n)$ with a shadow match $(R_I, R'_I)(R_1, R'_1) \dots (R_n, R'_n)$ satisfying the conditions (1) and (2). For each $a \in \text{Ran} R_n$, by (1) there is a χ -guided partial match with last position (a, s_n) . So we see that the move $\chi(a, s_n) : S \rightarrow \mathcal{P}(A)$ prescribed by the winning strategy χ is legitimate for each such a . Define the map $h : S \rightarrow A^\sharp$ by

$$h(v) = \{(b, b') \mid b \in \text{Ran} R_n \ \& \ b' \in \chi(b, s_n)(v)\}$$

Then it follows by the one-step bisimulation invariance theorem (Proposition 4.6) that the pair $\text{Kh}(V^\dagger(s_n), R[s_n]) = (V^\dagger(s_n), h[R[s_n]])$ is a legitimate move for Player I in the consequence game at the position (R_n, R'_n) . So the winning strategy f provides a map $g : A^\sharp \rightarrow A^\sharp$ which, since it is a legitimate move, satisfies the condition $\text{K}(g \circ h)(V^\dagger(s_n), R[s_n]), m_{a'_n} \Vdash_1 \Theta'(a'_n)$, where $m_{a'_n}$ is the natural marking at a'_n . It follows by Proposition 4.6 that $V^\dagger(s_n), R[s_n], m_{a'_n} \circ g \circ h \Vdash_1 \Theta(a'_n)$. So we set $\chi'(\Sigma) = m_{a'_n} \circ g \circ h$, and this is a legitimate move. Furthermore, if $b' \in m_{a'_n}(g(h(v)))$, then $(a'_n, b') \in g(h(v))$, and we can continue the shadow match with the pair $(h(v), g(h(v)))$. Then the extended shadow match

$$(R_I, R'_I)(R_1, R'_1) \dots (R_n, R'_n)(h(v), g(h(v)))$$

satisfies both conditions (1) and (2), so the proof is finished. QED

Similar to the satisfiability game, we will often want to make some simplifying assumptions on the moves picked by player I. These are justified by the following analog of Proposition 5.18. We omit the equally analogous proof.

Proposition 5.18 *Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ and $\mathbb{A}' = (A', \Theta', \Omega', a'_I)$ be modal automata, and let $\mathcal{N} \subseteq A^\sharp$ be some set of relations. Assume that for every basic position $(R, R') \in A^\sharp \times A^\sharp$ of the consequence game, and every legitimate move (Y, \mathcal{R}) of player I there is a legitimate move (Y, \mathcal{R}') such that $\mathcal{R}' \subseteq \mathcal{N}$ and $\mathcal{R}' \overline{\text{P}} \subseteq \mathcal{R}$. Then for any winning position in $\mathcal{S}(\mathbb{A})$ player I has a winning strategy that restricts her moves to pairs (Y, \mathcal{R}) with $\mathcal{R} \subseteq \mathcal{N}$.*

It should be stressed that the converse direction of Proposition 5.17 does *not* hold in general. If \mathbb{A}' is a game consequence of \mathbb{A} , the existence of a winning strategy for player II in the consequence game indicates a close *structural* relation between \mathbb{A} and \mathbb{A}' , far tighter than what is required for \mathbb{A} being merely a semantic consequence of \mathbb{A}' . Below we will see an example of two automata such that $\mathbb{A} \models \mathbb{A}'$ but $\mathbb{A} \not\models_{\mathbf{G}} \mathbb{A}'$, but first we give an example of two automata that do satisfy the game consequence relation. Note that this example is closely linked to the fixpoint rule of Kozen's axiom system.

Proposition 5.19 *For all modal automata \mathbb{A} that are positive in x , we have $\mathbb{A}_1^x[\mu x.\mathbb{A}/x] \models_{\mathbf{G}} \mu x.\mathbb{A}$.*

Proof. Recall that the automaton $\mu x.\mathbb{A}$ has the same carrier as the automaton \mathbb{A}_1^x , and that the automaton $\mathbb{A}_1^x[\mu x.\mathbb{A}/x]$ is built from $\mu x.\mathbb{A}$ together with one disjoint copy of \mathbb{A}_1^x , so $\mathbb{A}_1^x[\mu x.\mathbb{A}/x]$ will contain for each state a in $\mu x.\mathbb{A}$ an extra state a' corresponding to a belonging to the disjoint copy of \mathbb{A}_1^x . With this in mind, we define a map f from states of $\mathbb{A}_1^x[\mu x.\mathbb{A}/x]$ to states of $\mu x.\mathbb{A}$ by mapping a' to a , and a to itself, for each state a in $\mu x.\mathbb{A}$. This map induces a map F from relations over the states of $\mathbb{A}_1^x[\mu x.\mathbb{A}/x]$ to relations over the states of $\mu x.\mathbb{A}$ by the assignment:

$$F : R \mapsto \{(f(a), f(b)) \mid (a, b) \in R\}$$

We also have a map F_0 defined by:

$$F_0 : R \mapsto \{(\underline{x}, f(b)) \mid ((a_I)_1', b) \in R\}$$

Thus we get a strategy for Player II in the game $\mathcal{C}(\mathbb{A}_1^x[\mu x.\mathbb{A}/x], \mu x.\mathbb{A})$ defined by choosing the map $F_0 \upharpoonright_{\mathcal{R}}$ in response to the first move (Y, \mathcal{R}) made by Player I, and choosing the map $F \upharpoonright_{\mathcal{R}}$ as a response to every other move (Y, \mathcal{R}) made by Player I. It can be checked that this is a winning strategy for Player II. QED

Note that this proposition is stated in terms of the guarded automaton \mathbb{A}_1^x rather than \mathbb{A} . It is not too hard to convince oneself that the automata $\mathbb{A}[\mu x.\mathbb{A}]$ and $\mathbb{A}_1^x[\mu x.\mathbb{A}]$ are *semantically* equivalent, but the consequence game is a stronger concept and in general it is somewhat surprisingly not true that $\mathbb{A}[\mu x.\mathbb{A}/x] \models_{\mathbf{G}} \mu x.\mathbb{A}$. For a simple counterexample, consider the following:

Example 5.20 Let p, q, x be any three propositional variables and let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ where $A = \{a, b\}$, $a_I = a$ and $\Theta(a) = \diamond b$, $\Theta(b) = (x \wedge p) \vee q$. Hence, we will have $\theta_0^b = p$ and $\theta_1^b = q$, and we can take $\theta_0^a = \perp$ and $\theta_1^a = \Theta(a)$. (The priority map Ω is irrelevant to the example and hence not specified.) We shall show that $\mathbb{A}[\mu x.\mathbb{A}] \not\models_{\mathbf{G}} \mu x.\mathbb{A}$.

Note that $\mu x.\mathbb{A}$ will have five states $\underline{x}, a_0, a_1, b_0, b_1$ and $\mathbb{A}[\mu x.\mathbb{A}/x]$ will have the two additional states a and b . For convenience we list the transition maps of the automata \mathbb{A} , \mathbb{A}^x , $\mu x.\mathbb{A}$ and $\mathbb{A}[\mu x.\mathbb{A}]$ in Table 3; the last row of the table provides the starting states of the respective automata. We let Θ' denote the transition map of $\mu x.\mathbb{A}$ and Θ'' the transition map of $\mathbb{A}[\mu x.\mathbb{A}/x]$.

To prove that that $\mathbb{A}[\mu x.\mathbb{A}] \not\models_{\mathbf{G}} \mu x.\mathbb{A}$ we supply Player I with a winning strategy in the consequence game $\mathcal{C}(\mathbb{A}[\mu x.\mathbb{A}], \mu x.\mathbb{A})$.

State	$\mathbb{A}(\Theta)$	$\mathbb{A}^x(\Theta^x)$	$\mu x.\mathbb{A}(\Theta')$	$\mathbb{A}[\mu x.\mathbb{A}](\Theta'')$
a	$\diamond b$	—	—	$\diamond b$
b	$(x \wedge p) \vee q$	—	—	$(\diamond((\underline{x} \wedge b_0) \vee b_1) \wedge p) \vee q$
a_0	—	\perp	\perp	\perp
a_1	—	$\diamond((x \wedge b_0) \vee b_1)$	$\diamond((x \wedge b_0) \vee b_1)$	$\diamond((x \wedge b_0) \vee b_1)$
b_0	—	p	p	p
b_1	—	q	q	q
\underline{x}	—	x	$\diamond((\underline{x} \wedge b_0) \vee b_1)$	$\diamond((\underline{x} \wedge b_0) \vee b_1)$
St.st.	a	a_1	\underline{x}	a

Table 3: The transition map and starting states of the automata \mathbb{A} , \mathbb{A}^x , $\mu x.\mathbb{A}$ and $\mathbb{A}[\mu x.\mathbb{A}]$

In fact we shall show that Player I can win this game in just a few moves. Since the free variables of the two automata $\mathbb{A}[\mu x.\mathbb{A}/x]$ and $\mu x.\mathbb{A}$ are p and q , she first needs to pick an element of the set

$$\mathcal{P}(\{p, q\}) \times \mathcal{P}(\{\underline{x}, a_0, a_1, b_0, b_1, a, b\}^\#)$$

We let the first move of Player I be the pair $(\emptyset, \{\{(a, b)\}\})$, which satisfies the formula $\Theta''(a) = \diamond b$ as required (recall that a is the starting state of $\mathbb{A}[\mu x.\mathbb{A}]$).

Now Player II has to come up with a map F such that $(\emptyset, F[\{\{(a, b)\}\}]) = (\emptyset, \{F(\{(a, b)\})\})$ satisfies the formula $\diamond((\underline{x} \wedge b_0) \vee b_1)$, since \underline{x} is the start state of $\mu x.\mathbb{A}$. Player II's move is thus completely determined by a single relation $Q = F(\{(a, b)\})$. To satisfy the mentioned formula, Player II is forced to define the relation Q such that either (i) $(\underline{x}, b_0), (\underline{x}, \underline{x}) \in Q$ or (ii) $(\underline{x}, b_1) \in Q$. Suppose now that player II's move is given by the relation $Q = F(\{(a, b)\})$.

In either case, Player I first has to make a trivial move, picking the single element of the graph of F , that is, the pair $(\{(a, b)\}, Q)$. To see that from this position she can win the game, we make a case distinction, as to the nature of the earlier move by \forall .

If (i) we have $(\underline{x}, b_0), (\underline{x}, \underline{x}) \in Q$ then we let Player I choose the pair $(\{q\}, \emptyset)$ which is a legal move since it satisfies the disjunct q of $\Theta''(b)$. Since $p \notin \{q\}$ there is no map G that Player II can respond with such that $(\{q\}, G[\emptyset])$ satisfies the formula $\Theta'(b_0)$, which is just p . Hence Player II is stuck and Player I wins the game.

On the other hand, if (ii) $(\underline{x}, b_1) \in Q$ then we can let Player I choose the pair $(\{p\}, \{(b, b_1)\})$ since it satisfies the disjunct $\Theta'(\underline{x}) \wedge p$ of $\Theta''(b)$, where we recall that $\Theta'(\underline{x})$ was the formula $\diamond((\underline{x} \wedge b_0) \vee b_1)$. But now, since $q \notin \{p\}$, there is no map G that Player II can choose such that $(\{p\}, G[\{(b, b_1)\}])$ satisfies the formula $\Theta'(b_1)$, which is just q . So again Player II is stuck, and Player I wins. \triangleleft

Intuitively, what is driving the previous example is that the construction \mathbb{A}^x splits states of \mathbb{A} into disjunctions, which gives Player I some extra power in the game $\mathcal{C}(\mathbb{A}[\mu x.\mathbb{A}], \mu x.\mathbb{A})$ since she can choose which disjunct of a one-step formula to make true on the left-hand side of the game, while the choice may be already made for Player II on the right-hand side. This illustrates our earlier point that $\mathbb{A} \models_{\mathcal{G}} \mathbb{A}'$ indicates a rather tight structural relation between the two automata.

To finish off this section, we mention two basic facts about the consequence game, stating that the game consequence relation is reflexive and transitive.

Proposition 5.21 *Let \mathbb{A} , \mathbb{A}' and \mathbb{A}'' be modal automata.*

- (1) $\mathbb{A} \models_{\mathbf{G}} \mathbb{A}$;
- (2) if $\mathbb{A} \models_{\mathbf{G}} \mathbb{A}'$ and $\mathbb{A}' \models_{\mathbf{G}} \mathbb{A}''$ then $\mathbb{A} \models_{\mathbf{G}} \mathbb{A}''$.

Proof. Clearly, the proof of the first item is trivial. Concerning the transitivity of $\models_{\mathbf{G}}$, it is a routine exercise to verify that player II can compose any two winning strategies in the games $\mathcal{C}(\mathbb{A}, \mathbb{A}')$ and $\mathcal{C}(\mathbb{A}', \mathbb{A}'')$, respectively, to obtain a winning strategy in the game $\mathcal{C}(\mathbb{A}, \mathbb{A}'')$. QED

6 Taming the traces - one step at a time

We have seen that the winner of an infinite match of the satisfiability game for a modal automaton \mathbb{A} is determined by checking whether the induced A^\sharp -stream of relations over A contains any bad traces. Hence, the problem of determining membership of the no-bad-trace language $NBT_{\mathbb{A}} \subseteq (A^\sharp)^\omega$ will be of central importance to us. In Proposition 5.4 we already saw that $NBT_{\mathbb{A}}$ is an ω -regular language, but to find an actual stream automaton recognizing this language is a nontrivial task, involving nontrivial operations such as the Safra construction. Put in different terms: the combinatorics of the trace graph of a stream over the alphabet A^\sharp , as defined in Section 5, is rather intricate. For this reason, we shall be interested in special classes of modal automata for which the trace graphs of matches in the satisfiability game have significantly simpler structure. As we said in the introduction, this can be achieved by imposing certain restrictions on the one-step formulas of the automata. In this section we make this precise, thus bringing together the two different aspects of Walukiewicz's completeness proof that we aim to distinguish in our analysis, the *combinatorics of traces* on the one hand and the *one-step dynamics* of automata on the other.

We shall start by isolating, for a fixed automaton \mathbb{A} , certain subsets of the alphabet A^\sharp (that is, certain kinds of binary relations over A), such that the trace graphs of streams over these restricted alphabets are in a precise sense simpler than in the general case. We then proceed to introduce the corresponding restrictions on the one-step formulas that produce modal automata for which infinite matches of the satisfiability game can indeed be assumed without loss of generality to produce streams over those restricted classes of relations in A^\sharp .

6.1 Functional, clusterwise functional and thin relations

The simplest class of relations that we shall consider are the *functional* ones:

Definition 6.1 Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be a modal automaton, and let $R \in A^\sharp$ be a relation. We call R *functional*, if each $a \in A$ has at most one R -successor. This element, if it exists, is denoted as a_R^+ , or as a^+ in case R is clear from context. The set of functional relations in A^\sharp will be denoted by A^\sharp_f . \triangleleft

The trace combinatorics of streams of functional relations is *trivial*, as described by the following proposition. Although the result is obvious, we state it explicitly for emphasis:

Proposition 6.2 *Fix a modal automaton \mathbb{A} and let $R_1 R_2 R_3 \dots$ be any stream over A^\sharp_f . Then for each $a \in \text{Dom}(R_1)$ there is at most one infinite trace on this stream beginning with a , and if each R_i is total then the correspondence is one-to-one.*

In particular, there are at most $|A|$ many infinite traces on any stream on A^\sharp_f . This is the best possible scenario one could hope for: it is easy to check whether there is any bad trace on such a stream, since there is in total only a bounded number of traces to consider.

A wider class of relations that maintains some of this simplicity is the following:

Definition 6.3 Given a fixed modal automaton \mathbb{A} , a relation $R \in A^\sharp$ is said to be *clusterwise functional* if:

- (1) for all $a, b \in A$ with aRb , we have $b \triangleleft a$;
- (2) for all $a \in A$, there is at most one $b \in C_a$ such that aRb .

The set of clusterwise functional relations in A^\sharp will be denoted by A^\sharp_{cf} . \triangleleft

Generally, a stream over the alphabet A^\sharp_{cf} will have infinitely many traces. However, the trace combinatorics of streams over A^\sharp_{cf} is still much simpler than the general case, in a sense made precise by the following proposition. We recall that two streams σ, τ over any alphabet are said to be eventually equal if there is a $k \in \omega$ such that $\sigma(j) = \tau(j)$ for all $j \geq k$.

Proposition 6.4 *Given a modal automaton \mathbb{A} , let $\Sigma = R_1R_2R_3\dots$ be any stream over A^\sharp_{cf} . Then there exists a collection F of at most $|A|$ many infinite traces over $\Sigma = R_1R_2R_3\dots$, such that every infinite trace on this stream is eventually equal to some trace in F .*

Proof. By the first condition on clusterwise functionality, every infinite trace τ through Σ eventually ends up in a cluster C , in the sense that $\exists n. \forall k \geq n. \tau(k) \in C$. It thus suffices to prove that the relation of eventual equality, taken over the set of traces that eventually end up in an arbitrary but fixed cluster C , is an equivalence relation of index at most $|C|$.

Suppose for contradiction that this is not the case, i.e., there are traces $\{\tau_i \mid 0 \leq i \leq |C|\}$, all ending up in C , and such that τ_i and τ_j are eventually equal only if $i = j$. Then we can find an $n \in \omega$ such that for all $k \geq n$ each $\tau_i(k)$ belongs to C . By the pigeon hole principle then there must be distinct indices i and j such that $\tau_i(n) = \tau_j(n)$. But by clusterwise functionality this implies that $\tau_i(k) = \tau_j(k)$ for all $k \geq n$, so that τ_i and τ_j are eventually equal after all. QED

In particular this means that we only have to examine the $|A|$ many traces in F to decide whether there is a bad trace on $R_0R_1R_2\dots$, since two eventually equal traces are clearly either both bad or both not bad.

Cluster-wise functional relations are *almost* the key concept that we need, but it turns out that we are going to require a little bit of extra generality. While the infinite traces of a stream over A^\sharp_{cf} are essentially finite in the sense of Proposition 6.4, we shall finally consider a wider class of relations for which the corresponding streams have the property that there are essentially only finitely many *bad traces*.

Definition 6.5 Fix a modal automaton \mathbb{A} . A state c belonging to some cluster C of \mathbb{A} is called a *safe state* of C if $\Omega(c)$ is even, and no μ -state in C has a higher priority than c . A subset $B \subseteq A$ is *C-safe* if there is at most one state in $B \cap C$ that is not safe in C .

Given a state $a \in A$, we call a relation $R \in A^\sharp$ *thin* with respect to \mathbb{A} and a , or \mathbb{A} -*thin with respect to a* , if:

- (1) for all $b \in A$ with aRb , we have $b \triangleleft a$;
- (2) $R[a] \subseteq A$ is C_a -safe.

We call R \mathbb{A} -thin if it is \mathbb{A} -thin with respect to all $a \in A$. We denote the collection of thin relations in A^\sharp by A^\sharp_{thin} . \triangleleft

For streams over the set of \mathbb{A} -thin relations, we have the following result:

Proposition 6.6 *Given a modal automaton \mathbb{A} , let $(R_i)_{i \geq 1}$ be a stream over A^\sharp_{thin} . Then there exists a collection F of at most $|A|$ many infinite traces over $(R_i)_{i \geq 1}$, such that any given bad trace over $(R_i)_{i \geq 1}$ is eventually equal to some trace in F .*

Proof. We prove this proposition by a similar argument as used for Proposition 6.4, the differences being that we restrict attention to bad traces, and, in the reductio argument, let $n \in \omega$ satisfy the additional requirement that for all $k \geq n$, and all i , $\tau_i(k)$ is *not* a safe state in C . QED

Again, this combinatorial property greatly simplifies the problem of checking whether there is some bad trace on $R_0 R_1 R_2 \dots$, since we only have to check whether the bounded collection F contains a bad trace. In order to exploit this nice property of thin relations, we will introduce a second version of the satisfiability game:

Definition 6.7 Given a modal automaton \mathbb{A} , the *thin satisfiability game* for \mathbb{A} , denoted $\mathcal{S}_{thin}(\mathbb{A})$, is defined as the satisfiability game $\mathcal{S}(\mathbb{A})$ except that the moves of \forall are constrained so that \forall may only choose \mathbb{A} -thin relations. That is, R is a legitimate move for \forall at some position in $\mathcal{S}_{thin}(\mathbb{A})$ iff R is a legitimate move at the same position in $\mathcal{S}(\mathbb{A})$, and $R \in A^\sharp_{thin}$. A winning strategy for \forall in $\mathcal{S}_{thin}(\mathbb{A})$ will be called a *thin refutation* of \mathbb{A} . \triangleleft

In general, the game $\mathcal{S}_{thin}(\mathbb{A})$ is not equivalent to $\mathcal{S}(\mathbb{A})$ in the sense that there is always a winning strategy for the same player in both games: since the moves of \forall are restricted in $\mathcal{S}_{thin}(\mathbb{A})$, it may be that \exists has a winning strategy in $\mathcal{S}_{thin}(\mathbb{A})$ but not in $\mathcal{S}(\mathbb{A})$. In the following subsection, we shall arrive at a class of modal automata for which the equivalence does hold.

6.2 Disjunctive and semi-disjunctive automata

The first class of automata that we introduce is well-known from the literature: disjunctive automata were introduced under the name of μ -automata in [31]. The definition of disjunctive automata is based on the cover modality introduced in Definition 3.11. Recall that for a set of formulas Ψ , the formula $\nabla\Psi$ is given as $\nabla\Psi := \Box \bigvee \Psi \wedge \bigwedge \Diamond\Psi$.

Definition 6.8 Let X be a given set of proposition letters and A any finite set. We first define the language $\text{LitC}(X)$ to be generated by π in the grammar:

$$\pi ::= \top \mid p \mid \neg p \mid \pi \wedge \pi$$

where $p \in X$. We now define the set of disjunctive formulas in $1\text{ML}(X, A)$, which we denote by $1\text{ML}_d(X, A)$, as follows:

$$\alpha ::= \perp \mid \alpha \vee \alpha \mid \pi \wedge \nabla B$$

where $\pi \in \text{LitC}(X)$ and $B \subseteq A$. \triangleleft

Definition 6.9 A modal X -automaton \mathbb{A} is said to be *disjunctive* if the range of the transition map Θ is contained in $1\text{ML}_d(X, A)$. \triangleleft

Remark 6.10 Since the cover modality ∇ plays a key role in the semantics of disjunctive formulas, we recall its meaning in the specific setting of the satisfiability game. Given a subset $B \subseteq A$ and a pair $\Gamma = (\mathbf{Y}, \mathcal{R}) \in \mathbf{KA}^\sharp$, observe that

$$(\mathbf{Y}, \mathcal{R}) \Vdash_a^1 \nabla B \text{ iff } B \subseteq \bigcup \{R[a] \mid R \in \mathcal{R}\} \text{ and } R[a] \cap B \neq \emptyset \text{ for all } R \in \mathcal{R}. \quad (3)$$

This boils down to:

$$(\mathbf{Y}, \mathcal{R}) \Vdash_a^1 \nabla \emptyset \text{ iff } \mathcal{R} = \emptyset, \quad (4)$$

in the specific case where $B = \emptyset$. \triangleleft

We note that our definition does not admit the formula \top as a disjunctive one-step formula. This is in no way a restriction on the expressive power of disjunctive automata, for the following reason: given a disjunctive automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$, we call a state $a \in A$ a *true state* if $\Theta(a) = \nabla \emptyset \vee \nabla \{a\}$ and $\Omega(a) = 0$. It is easy to verify that in this case, the automaton $\mathbb{A}\langle a \rangle$ accepts all pointed Kripke structures, so that we may think of the one-step formula $\Theta(a)$ as internally representing the formula \top .

The basic observation about disjunctive one-step formulas is given by Proposition 6.11 below.

Proposition 6.11 *Let $\alpha \in \mathbf{1ML}_d(\mathbf{X}, A)$ and $(\mathbf{Y}, \mathcal{R}) \in \mathbf{KA}^\sharp$ be such that $(\mathbf{Y}, \mathcal{R}) \Vdash_a^1 \alpha$. Then there is some $\mathcal{R}' \in \mathbf{PA}^\sharp$ such that $\mathcal{R}' \xrightarrow{\vec{P}} \subseteq \mathcal{R}$, $(\mathbf{Y}, \mathcal{R}') \Vdash_a^1 \alpha$ and $R[a] = 1$ for all $R \in \mathcal{R}$.*

Proof. Since every $\alpha \in \mathbf{1ML}_d(\mathbf{X}, A)$ is a finite disjunction of formulas of the form $\pi \wedge \nabla B$, it suffices to prove the claim for α a one-step formula of the form ∇B , with $B \subseteq A$. Roughly, the idea of the proof is to ‘split’ the elements of \mathcal{R} if needed.

For more details, assume that $(\mathbf{Y}, \mathcal{R}) \Vdash_a^1 \nabla B$, and distinguish cases. If $B = \emptyset$ the claim holds trivially since by (4) we find $\mathcal{R} = \emptyset$ and so we may take $\mathcal{R}' := \mathcal{R}$.

Hence we may focus on the case where $B \neq \emptyset$. First of all it follows from (3) that $\mathcal{R} \neq \emptyset$ and that $R[a] \cap B \neq \emptyset$ for all $R \in \mathcal{R}$. Second, we ensure that $R[a] \subseteq B$ for all $R \in \mathcal{R}$. If this would not be the case, by (3) we may replace every $R \in \mathcal{R}$ with the relation $R_{a,B} := R \setminus \{(a, b) \mid b \in R[a] \setminus B\}$. The resulting set $\mathcal{R}'' := \{R_{a,B} \mid R \in \mathcal{R}\}$ satisfies $\mathcal{R}'' \xrightarrow{\vec{P}} \subseteq \mathcal{R}$ and $(\mathbf{Y}, \mathcal{R}'') \Vdash_a^1 \nabla B$.

Finally, define $\mathcal{Q} := \{R \in \mathcal{R}'' \mid R[a] > 1\}$ as the set of relations for which $R[a]$ is too big, and for each $R \in \mathcal{Q}$, and each $b \in R[a]$, define

$$R_b := R \setminus \{(a, b') \mid b' \neq b\},$$

so that $R = \bigcup \{R_b \mid b \in R[a]\}$, while $R_b[a]$ is a singleton for each $b \in R[a]$. Then put

$$\mathcal{R}' := (\mathcal{R} \setminus \mathcal{Q}) \cup \bigcup \{R_b \mid R \in \mathcal{Q}, b \in R[a]\},$$

that is, we ‘split’ every $R \in \mathcal{Q}$. Using (3) it is then a matter of routine to verify that \mathcal{R}' meets the required conditions. QED

We will state a number of results along the lines of Proposition 6.11, and the proofs of these results will be omitted since they all follow essentially the same line of reasoning as the previous one. The first of these generalizes Proposition 6.11 from looking at only single disjunctive formulas to considering the entire range of the transition map of a disjunctive automaton:

Proposition 6.12 *Let $\mathbb{A} = (A, a_I, \Theta, \Omega)$ be a disjunctive automaton, and let $R \in A^\sharp$ be a basic position of the satisfiability game. Then for every legitimate move (Y, \mathcal{R}) of \exists at R there is a legitimate move (Y, \mathcal{R}') such that $\mathcal{R}' \xrightarrow{\vec{P}} \subseteq \mathcal{R}$ and $\mathcal{R}' \subseteq A^\sharp_f$.*

From this it follows by Proposition 5.12 that we may always assume without loss of generality that \exists restricts her moves to pairs (Y, \mathcal{R}) where $\mathcal{R} \subseteq A^\sharp_f$, so that the stream over A^\sharp consisting of the basic positions of an infinite match in the satisfiability game consists only of functional relations.

It is possible to define a more general class of modal automata that would aptly be called “clusterwise disjunctive automata”, for which a similar result could be proved with “clusterwise functional relations” in place of “functional relations”. These automata will not play any essential role in the completeness proof however, so we proceed directly to introduce the natural class of automata for which matches in the satisfiability game produce streams over A^\sharp_{thin} . Unlike disjunctive automata, these automata do not already feature in the literature. We call them *semi-disjunctive automata*, and their transition function is clearly linked to the notion of thinness.

Definition 6.13 Let \mathbb{A} be a modal automaton and let C be a cluster of \mathbb{A} . A *C-safe conjunction* is a formula of the form $\bigwedge B$, where B is C -safe. The grammar

$$\alpha ::= \top \mid \pi \wedge \nabla\{\bigwedge B \mid B \in \mathcal{B}\} \mid \alpha \vee \alpha,$$

where $\pi \in \text{LitC}(X)$ and \mathcal{B} is a collection of C -safe sets, defines the set $\text{1ML}_C(X, A)$ of *C-safe one-step formulas*.

A modal automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$ is said to be *semi-disjunctive* if $\Theta(a)$ is a C_a -safe formula for all $a \in A$ (where we recall that C_a denotes the cluster of a). \triangleleft

Put informally: a C -safe one-step formula is a disjunctive formula over the set of C -safe conjunctions, and a semi-disjunctive automaton is one in which the one-step formula assigned to each state is safe with respect to the cluster of that state. Semi-disjunctive automata are to modal automata what the *weakly aconjunctive formulas* introduced by Walukiewicz in [73] are to formulas of the μ -calculus. These were in turn introduced as a generalized variant of the *aconjunctive formulas* for which Kozen proved his partial completeness result in [36].

The definition of semi-disjunctive automata is tailored towards the following proposition, the proof of which is similar to that of Proposition 6.12.

Proposition 6.14 *Let $\mathbb{A} = (A, a_I, \Theta, \Omega)$ be a semi-disjunctive automaton, and let $R \in A^\sharp$ be a basic position of the satisfiability game. Then for every legitimate move (Y, \mathcal{R}) of \exists at R , there is a legitimate move (Y, \mathcal{R}') such that $\mathcal{R}' \xrightarrow{\vec{P}} \subseteq \mathcal{R}$ and $\mathcal{R}' \subseteq A^\sharp_{thin}$.*

This means that, by appealing to Proposition 5.12, we can assume without loss of generality that in the satisfiability game of a semi-disjunctive automaton, \exists always plays a strategy such that all the infinite matches guided by this strategy induce streams over A^\sharp consisting entirely of thin relations.

We can make this point a bit more formal by considering the thin satisfiability game of Definition 6.7. We now obtain the following result, as an immediate consequence of Propositions 6.14 and 5.12:

Corollary 6.15 *Let $\mathbb{A} = (A, a_I, \Theta, \Omega)$ be a semi-disjunctive automaton. Then each player $\Pi \in \{\exists, \forall\}$ has a winning strategy in $\mathcal{S}_{thin}(\mathbb{A})$ iff she/he has one in $\mathcal{S}(\mathbb{A})$. Hence, \mathbb{A} is either satisfiable or admits a thin refutation.*

Proof. It is clear that any winning strategy for \exists in $\mathcal{S}(\mathbb{A})$ is still a winning strategy in $\mathcal{S}_{thin}(\mathbb{A})$. Conversely, suppose \exists has a winning strategy χ in $\mathcal{S}_{thin}(\mathbb{A})$. By Proposition 6.14 and Proposition 5.12 (or, to be more precise, the version of the latter Proposition formulated for the thin satisfiability game) we may without loss of generality assume that χ only picks moves (Y, \mathcal{R}) such that \mathcal{R} is a collection of thin relations. But then it is easy to see that this strategy χ is winning for \exists in $\mathcal{S}(\mathbb{A})$ as well. First, χ is obviously legitimate; to prove that it is winning the key observation is that at any position of the form (Y, \mathcal{R}) with $\mathcal{R} \subseteq A^\sharp_{thin}$, \forall 's moves in $\mathcal{S}(\mathbb{A})$ and $\mathcal{S}_{thin}(\mathbb{A})$ are exactly the same.

This shows the equivalence of the two games for \exists . The equivalence for \forall now follows by determinacy. QED

We note the following closure properties for disjunctive and semi-disjunctive automata. Here we say that an automaton is (semi-)disjunctive *modulo one-step equivalence* if it is one-step equivalent to a (semi-)disjunctive automaton.

Proposition 6.16 *Let \mathbb{A} and \mathbb{B} be two modal automata.*

- (1) *If \mathbb{A} is disjunctive, then it is also semi-disjunctive.*
- (2) *If \mathbb{A} and \mathbb{B} are disjunctive, then so is $\mathbb{A} \vee \mathbb{B}$,*
- (3) *If \mathbb{A} and \mathbb{B} are semi-disjunctive, then so is $\mathbb{A} \vee \mathbb{B}$.*
- (4) *If \mathbb{A} and \mathbb{B} are semi-disjunctive, then so is $\mathbb{A} \wedge \mathbb{B}$, modulo one-step equivalence.*
- (5) *If \mathbb{A} and \mathbb{B} are semi-disjunctive, then so is $\mathbb{A}[\mathbb{B}/x]$, modulo one-step equivalence.*
- (6) *If \mathbb{A} is disjunctive and positive in x , then \mathbb{A}^x and $\nu x.\mathbb{A}$ are semi-disjunctive, modulo one-step equivalence.*

Proof. The first three statements are immediate consequences of the definitions. We skip the proof of the fourth statement: it is similar to but simpler than that of (5), since in the case of $\mathbb{A} \wedge \mathbb{B}$ there is only one state, where we have to replace a conjunction of ∇ -formulas by a disjunction of appropriate ∇ -formulas, viz., the initial one.

For the fifth item, first observe that the states from \mathbb{B} cause no problem whatsoever: for $b \in B$ we have $\Theta_{\mathbb{A}[\mathbb{B}/x]}(b) = \Theta_{\mathbb{B}}(b)$, so the one-step formulas assigned to any $b \in B$ are b -safe since \mathbb{B} is semi-disjunctive. For a state a from \mathbb{A} , using the (Boolean) distributive law we can rewrite every formula $\Theta_{\mathbb{A}[\mathbb{B}/x]}(a)$ as a one-step equivalent finite disjunction of formulas of the form

$$\pi \wedge \nabla\{\wedge A \mid A \in \mathcal{A}\} \wedge \nabla\{\wedge B \mid B \in \mathcal{B}\}, \quad (5)$$

where $\pi \in \text{LitC}(\mathbf{X})$, \mathcal{A} is a collection of a -safe subsets of A , and \mathcal{B} is a family of subsets B .

We use the distributive law for the cover modality (Fact 3.15(1)) to observe that any conjunction of the shape $\nabla\{\bigwedge A \mid A \in \mathcal{A}\} \wedge \nabla\{\bigwedge B \mid B \in \mathcal{B}\}$ can be rewritten as an equivalent disjunction of formulas of the shape

$$\nabla\{\bigwedge(A \cup B) \mid A \in \mathcal{A}, B \in \mathcal{B}\}.$$

Since each set $A \cup B$ with $A \in \mathcal{A}, B \in \mathcal{B}$ is still a -safe in $\mathbb{A}[\mathbb{B}/x]$ (observe that no cluster in $\mathbb{A}[\mathbb{B}/x]$ contains states from both A and B), it is now a simple exercise to show that each formula $\Theta_{\mathbb{A}[\mathbb{B}/x]}(a)$ can be rewritten into a one-step equivalent a -safe formula.

Concerning item (6), assume that the automaton \mathbb{A} is disjunctive. We need to consider the precise shape of the automata \mathbb{A}^x and $\eta x.\mathbb{A}$. It is not hard to see that the formulas θ_i^a from Convention 4.18 themselves are disjunctive. The problem is that when we define the transition maps for the automata \mathbb{A}^x and $\nu x.\mathbb{A}$, the substitution $\kappa : a \mapsto (\underline{x} \wedge a_0) \vee a_1$ introduces conjunctions in the scope of modalities. We claim, however, that

$$\text{if } \alpha \in \text{ML}_d(\mathbf{X}, A), \text{ then } \alpha[\kappa] \text{ is equivalent to a } C_a\text{-safe formula, for any } a \in A. \quad (6)$$

Clearly we may restrict our attention to formulas α of the form ∇B with $B \subseteq A$. The key to proving (6) is the following one-step equivalence:

$$\nabla\{\gamma_0 \vee \gamma_1 \mid \gamma \in \Gamma\} \equiv_1 \bigvee \left\{ \nabla\{\gamma_i \mid (\gamma, i) \in Z\} \mid Z \subseteq \Gamma \times \{0, 1\}, \text{Dom}Z = \Gamma \right\}. \quad (7)$$

It follows from (7) that any formula of the form $\nabla B[\kappa]$ is one-step equivalent to a disjunction of formulas $\nabla \Pi$, where $\Pi \subseteq \{\underline{x} \wedge b_0, b_1 \mid b \in B\}$. Thus it remains to prove that formulas of the form $\underline{x} \wedge b_0$ are C_a -safe, for any $a \in A$. In the case of \mathbb{A}^x this follows from the fact that \underline{x} forms a degenerate cluster on its own, so that $\underline{x} \neq b_0$ implies that \underline{x} and b_0 belong to different clusters. In the case of $\nu x.\mathbb{A}$ the state \underline{x} is by construction the maximal even element of its cluster, so that again the set $\{\underline{x}, b_0\}$ is C_a -safe for any $a \in A$. QED

Note that $\mu x.\mathbb{A}$ is *not* generally semi-disjunctive, even if \mathbb{A} is disjunctive.

6.3 A key lemma

We now come to one of the key lemmas of the paper, phrased as Theorem 3 in the introduction. The role of this lemma in the overall completeness proof is to establish a *link* between the two games we have introduced for modal automata, the satisfiability game and the consequence game.

Theorem 3 Let \mathbb{A} and \mathbb{D} be respectively be a semi-disjunctive and an arbitrary modal automaton such that $\mathbb{A} \models_{\mathbf{G}} \mathbb{D}$. Then the automaton $\mathbb{A} \wedge \neg \mathbb{D}$ has a thin refutation.

This result is analogous to the result labelled as Lemma 36 in Walukiewicz' paper, but differs in two ways. First, it is formulated in automata-theoretic terms. But also, it is more general: the result in Walukiewicz' paper is stated for a weakly aconjunctive formula and a disjunctive formula, and so one should expect our result to be stated analogously for a semi-disjunctive automaton \mathbb{A} and a disjunctive automaton \mathbb{D} . It turns out that disjunctiveness

of \mathbb{D} is not needed; in fact this assumption doesn't even simplify the proof, so although the more restricted version of the result would have sufficed for the completeness proof we have preferred to state it for an arbitrary automaton \mathbb{D} .

We recall that the transition map of the automaton $\neg\mathbb{D}$ is defined by taking boolean duals of the formulas assigned by the transition map of \mathbb{D} , and the priority map is defined by simply raising all priorities by 1. We shall need the following fact on boolean duals, which is a straightforward consequence of the definitions:

Proposition 6.17 *Let S be any set, α any one-step formula in $1\text{ML}(\mathbf{X}, A)$ and let $m, m' : S \rightarrow \text{P}(A)$ be markings such that $(\mathbf{Y}, S, m) \Vdash^1 \alpha$ and $(\mathbf{Y}, S, m') \Vdash^1 \alpha^\partial$. Then for some $a \in A$ and some $s \in S$ we have $a \in m(s) \cap m'(s)$.*

Proof of Theorem 3. To fix notation, let $\mathbb{A} = (A, \Theta_{\mathbb{A}}, \Omega_{\mathbb{A}}, a_I)$, $\mathbb{D} = (D, \Theta_{\mathbb{D}}, \Omega_{\mathbb{D}}, d_I)$ and let \mathbb{B} denote the automaton $\mathbb{A} \wedge \neg\mathbb{D}$. We write $\mathbb{B} = (B, \Theta_{\mathbb{B}}, \Omega_{\mathbb{B}}, b_I)$ and recall that $B = A \uplus D \uplus \{b_I\}$.

Assume that player II has a winning strategy χ in the consequence game $\mathcal{C}(\mathbb{A}, \mathbb{D})$ starting at position $(\{(a_I, a_I)\}, \{(d_I, d_I)\})$. Our aim is to provide a thin refutation for the automaton \mathbb{B} , that is, a winning strategy for player \forall in the thin satisfiability game for the automaton $\mathbb{A} \wedge \neg\mathbb{D}$.

We shall make use of the following claim, which is based on another variation of Proposition 6.11. Call a relation $R \subseteq B^\sharp$ *\mathbb{A} -thin* if the relation $\text{Res}_A(R)$ is thin with respect to \mathbb{A} .

CLAIM 1 Without loss of generality we may assume that in any match of $\mathcal{S}_{\text{thin}}(\mathbb{A} \wedge \neg\mathbb{D})$, \exists only picks moves $(\mathbf{Y}, \mathcal{R})$ such that each $R \in \mathcal{R}$ is \mathbb{A} -thin.

We will now define a strategy σ for \forall in $\mathcal{S}(\mathbb{B})$, inductively making sure that the following two conditions are maintained, for any σ -guided partial match $\Sigma = R_0 \dots R_n$:

- (\dagger) R_n is thin, and for $n \geq 1$ satisfies $|\text{Ran}(R_n) \cap D| = 1$;
- (\ddagger) There is a χ -guided shadow $\mathcal{C}(\mathbb{A}, \mathbb{D})$ -match of the form $(S_0, S'_0)(S_1, S'_1) \dots (S_n, S'_n)$, where
 - (a) $S_0 = \{(a_I, a_I)\}$ and $S'_0 = \{(d_I, d_I)\}$;
 - (b) $S_1 = \{(a_I, a) \in A \times A \mid (b_I, a) \in R_1\}$ and $\{(d_I, d) \in D \times D \mid (b_I, d) \in R_1\} \subseteq S'_1$;
 - (c) for each $i > 1$ we have $R_i \cap (A \times A) = S_i$ and $R_i \cap (D \times D)$ is a singleton $\{(d, d')\}$ with $d \in \text{Ran}(R_{i-1}) \cap D$ and $(d, d') \in S'_i$.

For $n = 0$ by definition we have $R_0 = \{b_I, b_I\}$, $S_0 = \{(a_I, a_I)\}$ and $S'_0 = \{(d_I, d_I)\}$, so that the conditions (\dagger) and (\ddagger) hold. We leave it for the reader to verify that the case where $n = 1$ can be seen as a notational variant of the general case, and focus on showing how \forall can extend the match $R_0 \dots R_n$ to $R_0 \dots R_n R_{n+1}$ and maintain the above two conditions in the case that $n > 1$.

Suppose that the inductive hypothesis has been maintained for the partial match Σ consisting of the positions $R_0 R_1 \dots R_n$ where $n > 1$, with shadow match $(S_0, S'_0)(S_1, S'_1) \dots (S_n, S'_n)$, and let $\Gamma = (\mathbf{Y}, \mathcal{R}) \in \text{KB}^\sharp$ be an arbitrary move chosen by \exists at Σ . Recall that by Claim 1 we may assume that each member of the family \mathcal{R} is thin relative to \mathbb{A} . We have:

$$(\Gamma, n_b) \Vdash^1 \Theta_B(b) \text{ for all } b \in \text{Ran} R_n \tag{8}$$

where we recall that $n_b : \mathcal{R} \rightarrow PB$ denotes the natural b -marking on \mathcal{R} , mapping R to $R[b]$. In particular, we obtain that

$$(\Gamma, n_d) \Vdash^1 \Theta_D(d)^\partial \quad (9)$$

where d is the unique element of $\text{Ran}(R_n) \cap D$.

Recall that $\text{Res}_A : B^\sharp \rightarrow A^\sharp$ is the map sending a relation R to $R \cap (A \times A)$, so that $(\text{KRes}_A)\Gamma$ is the pair $(Y, \{R \cap (A \times A) \mid R \in \mathcal{R}\})$. By the one-step bisimulation invariance theorem, we may infer from (8) that

$$((\text{KRes}_A)\Gamma, n_a) \Vdash^1 \Theta_A(a), \text{ for all } a \in A \cap \text{Ran}R_n, \quad (10)$$

so that $(\text{KRes}_A)\Gamma$ is an admissible move for player I in the consequence game at position (S_n, S'_n) . Thus we find a map $F : \{S \cap (A \times A) \mid S \in \mathcal{R}\} \rightarrow D^\sharp$ dictated by Player II's winning strategy χ in the consequence game, such that

$$((\text{KF})(\text{KRes}_A)(\Gamma), n_d) \Vdash^1 \Theta_D(d). \quad (11)$$

We shall prove the following claim:

CLAIM 2 There is some $S \in \mathcal{R}$ and some $c \in D$ with $(d, c) \in F(S \cap A \times A) \cap (S \cap D \times D)$.

PROOF OF CLAIM By the one-step bisimulation invariance theorem and (11) we get that:

$$(\Gamma, n_d \circ F \circ \text{Res}_A) \Vdash^1 \Theta_D(d), \quad (12)$$

But then by Proposition 6.17 it follows from (9) and (12) that there is some $c \in D$ and some $S \in \mathcal{R}$ such that $c \in n_d(S) \cap n_d(F(\text{Res}_A(S)))$. From this the claim is immediate. \blacktriangleleft

With this claim in place, we define the next move for \forall prescribed by the strategy σ to be the relation $R_{n+1} := \text{Res}_A S \cup \{(d, c)\}$, where $S \in \mathcal{R}$ and $c \in D$ are as described in the claim, so that $(d, c) \in F(\text{Res}_A S) \cap \text{Res}_D S$. Note that this is a legitimate move in response to (Y, \mathcal{R}) since $R_{n+1} \subseteq S \in \mathcal{R}$. The shadow match is then extended by the pair $(S_{n+1}, S'_{n+1}) := (\text{Res}_A S, F(\text{Res}_A S))$ so that condition $(\ddagger c)$ of the induction hypothesis holds as an immediate consequence of the claim. For condition (\ddagger) , it is obvious that $|\text{Ran}(R_n) \cap D| = 1$; thinness of the relation R_{n+1} follows from the assumption that $S \in \mathcal{R}$ was thin relative to \mathbb{A} .

To show that the thus defined strategy σ is winning for \forall , first observe that he never gets stuck, so that we may focus on infinite matches. It suffices to prove that every infinite σ -guided match contains a bad trace, so consider an arbitrary such match $\Sigma = (R_i)_{i \geq 0}$.

Clearly we may assume that all initial parts of Σ , corresponding to the partial matches $(R_i)_{0 \leq i \leq n}$, satisfy the conditions (\ddagger) and (\ddagger) . From this it follows that Σ itself has an infinite χ -guided shadow match $(S_i, S'_i)_{i \geq 0}$ satisfying the condition $(\ddagger a-c)$. In addition, it follows from (\ddagger) that Σ will contain a *unique* trace in D , which by (\ddagger) will also be a trace on the right side of the shadow match in the consequence game. That is, the match $R_0 R_1 R_n \dots$ contains a unique trace of the form $b_I d_1 d_2 d_3 \dots$ with each d_i in D , and this is a trace through the stream $S'_0 S'_1 S'_2 \dots$ as well. If this trace is *not* bad given the priorities assigned to states in $\neg \mathbb{D}$, then it must be bad as a trace in \mathbb{D} since parities are swapped in $\neg \mathbb{D}$. Hence there must be a bad trace on the left side $S_0 S_1 S_2 \dots$ of the shadow match in the consequence game, since

this shadow match was guided by the winning strategy χ of Player II. But since every trace $b_I a_1 a_2 a_3 \dots$ through $S_0 S_1 S_2 \dots$ is also a trace in the match $R_0 R_1 R_2 \dots$ in the satisfiability game, with $a_1, a_2, a_3 \dots \in A$ and thus with the same priorities as in the consequence game, we see that either the unique trace through D in Σ is bad or there is some bad trace through A in Σ . In either case, Σ must contain at least one bad trace, and so Σ is a loss for \exists as required. QED

7 A strong simulation theorem

In this section we present a construction that turns an arbitrary, i.e., alternating, modal automaton \mathbb{A} into an equivalent disjunctive, i.e., nondeterministic, modal automaton $\text{sim}(\mathbb{A})$. In other words, we are concerned with the *simulation theorem* for modal automata here. In the setting of finite tree automata, the concept of alternation was introduced by Muller & Schupp [46]; they also mention the simulation problem and hint at a solution, but do not provide details. Emerson & Jutla [20] showed that the simulation problem for alternating tree automata becomes somewhat easier if acceptance is given by a *parity* condition (in fact, a concept introduced in this paper, independently from [44]). A fairly general construction, for tree automata with various acceptance conditions, was given by Muller & Schupp [47]. All of the work mentioned above dealt with automata operating on trees of a fixed, finite branching degree, which is slightly different from our setting of Kripke structures, where in particular the successors of a state form a (completely unstructured) *set*. As mentioned, the μ -automata introduced in this setting by Janin & Walukiewicz [30] are *nondeterministic*, and although the authors do not define alternating automata explicitly, their construction can be seen as a simulation theorem.

Our definition of the simulation of a modal automaton more or less follows the approach of Arnold & Niwiński [1]. However, we shall present a strengthened version of this simulation theorem: rather than merely showing \mathbb{A} and $\text{sim}(\mathbb{A})$ to be semantically equivalent, we shall prove stronger claims involving the consequence game. In one direction, we have $\mathbb{A} \models_G \text{sim}(\mathbb{A})$, from which it follows that $\text{sim}(\mathbb{A})$ accepts every model accepted by \mathbb{A} . In the other direction, we have not only that $\text{sim}(\mathbb{A}) \models_G \mathbb{A}$, but in fact given any modal automaton \mathbb{B} which is positive in the proposition letter x , we have:

$$\mathbb{B}[\text{sim}(\mathbb{A})/x] \models_G \mathbb{B}[\mathbb{A}/x].$$

This is tightly connected to one of the key lemmas in Walukiewicz' completeness proof (labelled "Lemma 39" in [73]), and can be seen as an automata-theoretic counterpart of this result.

We shall begin by providing an explicit definition of the transformation $\text{sim}(\cdot)$ from modal automata to disjunctive modal automata.

Definition 7.1 Fix a modal X -automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$. Given a subset Y of X , let \widehat{Y} denote the formula:

$$\bigwedge \{p \mid p \in Y\} \wedge \bigwedge \{\neg p \mid p \notin Y\}.$$

The *pre-simulation* of \mathbb{A} is defined to be the structure $\text{pre}(\mathbb{A}) = (A^\sharp, \Theta', NBT_\Omega, a'_I)$ where $A^\sharp := P(A \times A)$ as always, $a'_I := \{(a_I, a_I)\}$,

$$\Theta'(R) = \bigvee \{ \widehat{Y} \wedge \nabla \mathcal{R} \mid (Y, \mathcal{R}) \in \bigcap_{b \in \text{Ran} R} \llbracket \Theta(b) \rrbracket_b^1 \},$$

and NBT_Ω is the set of streams over A^\sharp that do not contain any bad traces.

Since the acceptance condition NBT_Ω is an ω -regular language with alphabet A^\sharp as we noted in Section 5, we may pick some deterministic parity automaton $\mathbb{Z} = (Z, \zeta, \Omega', z_I)$ that recognizes NBT_Ω . Finally we define $\text{sim}(\mathbb{A})$ to be the structure $(D, \Theta'', \Omega'', d_I)$ where:

- $D = A^\sharp \times Z$,
- $d_I = (a'_I, z_I)$,
- $\Theta''(R, z) = \Theta'(R)[(Q, \zeta(R, z)/Q \mid Q \in A^\sharp]$ and
- $\Omega''(R, z) = \Omega'(z)$.

We shall let $G_{\mathbb{A}} : D \rightarrow A^\sharp$ denote the map that is defined to be the projection map sending a pair (R, z) in the product $A^\sharp \times Z$ to its left component R . \triangleleft

The main result of this section is the following result, which we already mentioned in the introduction as one of the key lemmas in our completeness proof.

Theorem 4 *The map $\text{sim}(\cdot)$ assigns to each modal automaton \mathbb{A} a disjunctive modal automaton $\text{sim}(\mathbb{A})$ such that*

- (1) $\mathbb{A} \models_{\mathcal{G}} \text{sim}(\mathbb{A})$ and $\text{sim}(\mathbb{A}) \models_{\mathcal{G}} \mathbb{A}$;
- (2) $\mathbb{B}[\text{sim}(\mathbb{A})/p] \models_{\mathcal{G}} \mathbb{B}[\mathbb{A}/p]$, for any modal automaton \mathbb{B} which is positive in p .

Observe that the semantic equivalence of an automaton and its simulation follows from the first part of this theorem, together with Proposition 5.17.

To get a hint of the proof of Theorem 4(1), one should observe that the game trees of $\mathcal{S}(\mathbb{A})$ and $\mathcal{S}(\text{sim}(\mathbb{A}))$ are structurally very similar. In fact, it is clear that the simulation construction is very tightly related to the satisfiability game: the states of the pre-simulation of \mathbb{A} just *are* the basic positions of the satisfiability game for \mathbb{A} , and the acceptance condition for the pre-simulation of \mathbb{A} is exactly the winning condition in $\mathcal{S}(\mathbb{A})$.

Proof of Theorem 4(1). To show that $\mathbb{A} \models_{\mathcal{G}} \text{sim}(\mathbb{A})$ is easy: fix the stream automaton \mathbb{Z} that recognizes NBT_{Ω} . Then every finite word $R_0 \dots R_k$ over A^\sharp determines an associated state of \mathbb{Z} by simply running \mathbb{Z} on the word $R_0 \dots R_k$; so for R_0 the associated state is z_I , for $R_0 R_1$ the associated state is $\zeta(R_0, z_I)$ etc. Since every k -length partial match Σ of the consequence game $\mathcal{C}(\mathbb{A}, \text{sim}(\mathbb{A}))$ determines a word $R_0 \dots R_k$ over A^\sharp in the obvious way, we can associate a state z_{Σ} of \mathbb{Z} with each such partial match. If Player I continues the match Σ consisting of basic positions $(R_0, R'_0) \dots (R_k, R'_k)$ by choosing the move $(Y, \mathcal{R}) \in \text{KA}^\sharp$, then we let Player II respond with the map $F : \mathcal{R} \rightarrow (A^\sharp \times Z)^\sharp$ that is defined by mapping $R \in \mathcal{R}$ to the singleton $\{((R_k, z_{\Sigma}), (R, \zeta(R_k, z_{\Sigma})))\}$. It can be checked that this defines a winning strategy for Player II, and we leave the details to the reader.

The direction $\text{sim}(\mathbb{A}) \models_{\mathcal{G}} \mathbb{A}$ of clause (1), which can be seen as a simple special case of clause (2), will follow from the Propositions 7.3 and 7.5. QED

The difficult part of Theorem 4 is to prove clause (2), and this will be the focus of the rest of this section. It will be convenient to state more abstractly what the crucial properties are of the automaton $\text{sim}(\mathbb{A})$ that we have associated with an arbitrary automaton \mathbb{A} :

Definition 7.2 Let \mathbb{A} and \mathbb{D} be an arbitrary and a disjunctive modal automaton, respectively. We say that \mathbb{D} is a *disjunctive companion* of \mathbb{A} if there is a map $G : D \rightarrow A^\sharp$ satisfying the following conditions:

- (DC1) $G(d_I) = \{(a_I, a_I)\}$
- (DC2) Let $\delta = (Y, E) \in \text{KD}$ be such that $(Y, E, \text{sing}) \Vdash^1 \Theta_{\mathbb{D}}(d)$, where sing is the singleton map given by $e \mapsto \{e\}$. Then $(\text{KG})\delta \in \text{KA}^\sharp$ satisfies $(\text{KG})\delta \Vdash_a^1 \Theta_{\mathbb{A}}(a)$ for all $a \in \text{Ran}(Gd)$.
- (DC3) If $G(d_i)_{i \in \omega} \in (A^\sharp)^\omega$ contains a bad \mathbb{A} -trace, then $(d_i)_{i \in \omega}$ is itself a bad \mathbb{D} -trace. \triangleleft

The map G in this definition is intended as a witness of a tight structural relationship between the automaton \mathbb{D} and the satisfiability game for \mathbb{A} . In particular the map G captures the intuition that every state of a disjunctive companion represents a macro-state of \mathbb{A} (i.e. a position of the satisfiability game), plus possibly some extra information. In the concrete case of the automaton $\text{sim}(\mathbb{A})$, this “extra information” is a state of the stream automaton that detects bad traces. Informally one can think of a state $d \in D$ as a conjunction of the states in $\text{Ran}G(d)$, or differently put: for each for $a \in \text{Ran}G(d)$, think of a as being “implied by” d .

Each of the clauses of this definition can thus be given an informal explanation that is consistent with this idea. The first clause (DC1) simply expresses that the start state of \mathbb{D} is a representation of the start position of the satisfiability game $\mathcal{S}(\mathbb{A})$. The second clause (DC2) captures the idea that any state $a \in \text{Ran}G(d)$ is “entailed” by d in the following sense. Given an object $\delta = (Y, \{d_1, \dots, d_k\}) \in \mathbf{KD}$, we can see δ as a one-step model over D by taking the singleton map $\text{sing} : D \rightarrow PD$ (restricted to the set $E = \{d_1, \dots, d_k\}$) as a D -marking. Similarly, applying the map KG to δ , we obtain the object $(KG)\delta = (Y, \{Gd_1, \dots, Gd_k\}) \in \mathbf{KA}^\sharp$ which, as we have seen, may be taken as an A -indexed family of one-step models $(Y, \{Gd_1, \dots, Gd_k\}, n_a)$ over A . Now the condition (DC3) requires that if $(Y, \{d_1, \dots, d_k\}, \text{sing})$ satisfies the one-step formula $\Theta_{\mathbb{D}}(d)$, then the one-step model $(Y, \{Gd_1, \dots, Gd_k\}, n_a)$ satisfies $\Theta_{\mathbb{A}}(a)$, for each $a \in \text{Ran}(Gd)$. Finally, the clause (DC3) makes sure that if $(d_i)_{i \in \omega}$ is a “good” \mathbb{D} -stream, in the sense that it satisfies the acceptance condition of \mathbb{D} , then the A^\sharp -stream $(Gd_i)_{i \in \omega}$ satisfies the $NBT_{\mathbb{A}}$ -condition, that is, each of its traces satisfies the acceptance condition of \mathbb{A} , and thus provides a win for \exists in the satisfiability game.

Proposition 7.3 *The simulation map $\text{sim}(\cdot)$ assigns a disjunctive companion to any modal automaton.*

Proof. It is fairly straightforward to check that the projection map $G_{\mathbb{A}} : D \rightarrow A^\sharp$ specified in Definition 7.1, which simply forgets the states of the stream automaton used in the product construction, has all the properties required to witness that $\text{sim}(\mathbb{A})$ is a disjunctive companion of \mathbb{A} . QED

The main technical result of this section is the following:

Proposition 7.4 *Let \mathbb{A} and \mathbb{B} be arbitrary modal automata, let \mathbb{D} be a disjunctive companion of \mathbb{A} , and assume that \mathbb{B} is positive in p . Then*

$$\mathbb{B}[\mathbb{D}/p] \models_{\mathbf{G}} \mathbb{B}[\mathbb{A}/p].$$

To warm up, we first prove a simplified version of the Proposition (which immediately yields the statement $\text{sim}(\mathbb{A}) \models_{\mathbf{G}} \mathbb{A}$ in Theorem 4(1)).

Proposition 7.5 *Let \mathbb{D} be a disjunctive companion of the modal automaton \mathbb{A} . Then*

$$\mathbb{D} \models_{\mathbf{G}} \mathbb{A}.$$

Proof. Since \mathbb{D} is disjunctive, by the Propositions 6.12 and 5.12 we may assume without loss of generality that player I only picks moves $\Gamma = (\mathbf{Y}, \mathcal{R})$ such that all relations in \mathcal{R} are functional. As a result, all basic positions that we encounter will be of the form (R, R') where R consists of a single pair (d, d_R^+) .

Player II's strategy will be completely determined by the map G . More specifically, suppose that at a position $(\{(e, d)\}, R)$, player I plays $(\mathbf{Y}, \mathcal{R})$; by our assumption, every relation $Q \in \mathcal{R}$ is of the form $Q = \{(d, d_Q^+)\}$. Now Player II's strategy χ is given by the assignment

$$\chi : \{(d, d_Q^+)\} \mapsto G(d_Q^+).$$

It is a routine verification that by condition (DC2) this strategy is legitimate.

From the legitimacy of χ it follows that II cannot loose any χ -guided finite match. To show that χ is winning for II, consider an infinite χ -guided match $\Sigma = (R_n, R'_n)_{n \in \omega}$. It follows from our assumptions on player I's strategy that $R_0 R_1 \dots$ carries only a single trace, say, $d_0 d_1 \dots$, and by our definition of player II's strategy we then have that $R_i = G d_i$ for all i . From this it is immediate by the trace reflection clause (DC3) that Σ is a win for player II, as required. QED

Remark 7.6 In the light of the proof of Proposition 7.5, we now see that the map G witnessing that \mathbb{D} is a disjunctive companion of \mathbb{A} can be seen as encoding a particularly simple winning strategy for Player II in the consequence game $\mathcal{C}(\mathbb{D}, \mathbb{A})$. The trick of proving Proposition 7.4 is to turn this winning strategy encoded by G into a new winning strategy for Player II in $\mathcal{C}(\mathbb{B}[\mathbb{D}/p], \mathbb{B}[\mathbb{A}/p])$.

Before turning to Proposition 7.4 itself, let us first see why its proof is not so straightforward as one might expect on the basis of that of Proposition 7.5. To see where the difficulties lie, consider an arbitrary infinite match $\Sigma = (R_n, R'_n)_{n \in \omega}$ of the consequence game for $\mathbb{B}[\mathbb{D}/p]$ and $\mathbb{B}[\mathbb{A}/p]$. Given the shape of these two automata, we may assume that traces on $\Sigma_l := R_0 R_1 \dots$ consist of either a \mathbb{B} -trace or a finite \mathbb{B} -trace followed by an infinite \mathbb{D} -trace, and that, similarly, traces on $\Sigma_r := R'_0 R'_1 \dots$ consist of either a \mathbb{B} -trace or a finite \mathbb{B} -trace followed by an \mathbb{A} -trace. Our purpose will be to associate with each Σ_r -trace

$$\tau = b_0 b_1 \dots b_n a_{n+1} a_{n+2} a_{n+3} \dots,$$

a Σ_l -trace

$$\tau_l = b_0 b_1 \dots b_n d_{n+1} d_{n+2} d_{n+3} \dots,$$

such that we can use the trace reflection clause of Definition 7.2 on the \mathbb{D} - and \mathbb{A} -tail of τ and τ_l , respectively. For this purpose we will define, for each partial match leading to final position (R_n, R'_n) , a map $g_n : \text{Ran}_A R'_n \rightarrow \text{Ran}_D R_n$. Intuitively, for $a \in A$, $g_n(a)$ represents a state $d \in D$ that 'implies' a . Ideally, we would like to show that the τ -tail $(a_i)_{i > n}$ is in fact a trace on the A^\sharp -stream $(G(g_i a_i))_{i > n}$, while $(g_i a_i)_{i > n}$ is a tail of a Σ_l -trace, so that (DC3) applies indeed.

Unfortunately, this is too good to be true, due to complications that are caused by \mathbb{A} -traces *merging*: the point is that *trace jumps* may occur, that is, situations where for some pair $(a, a') \in R'_{j+1}$ it does not hold that $(g_j a, g_{j+1} a') \in R_{j+1}$. Our solution to this problem will be to ensure that every Σ_r -trace can suffer only *finitely many* trace jumps. Thus,

what we *can* show is that any \mathbb{A} -trace $(a_i)_{i>n}$ has a *tail* $a_k a_{k+1} a_{k+2} \cdots$ which is a trace on $G(g_k a_k)G(g_{k+1} a_{k+1})G(g_{k+2} a_{k+2}) \cdots$. This suffices to prove that if there is a bad trace on Σ_r , then there is also a bad trace on Σ_l , so that player II indeed wins the match Σ .

The tool that we employ to guarantee this consists of a total order on the collection of those Σ_l -traces that arrive to the \mathbb{D} -part of the automaton $\mathbb{B}[\mathbb{D}/p]$. The definition of this order crucially involves the disjunctivity of \mathbb{D} . \triangleleft

In the proof of Proposition 7.4, the following technical definition and proposition will be useful.

Definition 7.7 With $\Gamma, \Gamma' \in \mathbf{KA}^\sharp$, $a, a' \in A$, and $B \subseteq A$, we write $\Gamma \xrightarrow{B}_{a, a'} \Gamma'$ to abbreviate that $(\Gamma, m) \xrightarrow{1} (\Gamma', m')$ where $m : R \mapsto R[a] \cap B$ and $m' : R \mapsto R[a'] \cap B$. Similarly, we write $\Gamma \xleftrightarrow{B}_{a, a'} \Gamma'$ as an abbreviation for $(\Gamma, m) \xleftrightarrow{1} (\Gamma', m')$. \triangleleft

Proposition 7.8 Let $\Gamma, \Gamma' \in \mathbf{KA}^\sharp$. (1) If $\Gamma \xrightarrow{B}_{a, a'} \Gamma'$ then for all $\beta \in \mathbf{1ML}(\mathbf{X}, B)$ we have

$$\Gamma \Vdash_a^1 \beta \text{ implies } \Gamma' \Vdash_{a'}^1 \beta.$$

(2) If $\Gamma \xleftrightarrow{B}_{a, a'} \Gamma'$ then for all $\beta \in \mathbf{1ML}(\mathbf{X}, B)$ we have

$$\Gamma \Vdash_a^1 \beta \text{ iff } \Gamma' \Vdash_{a'}^1 \beta.$$

Proof. Immediate from the one-step preservation and bisimulation invariance theorems (Propositions 4.8 and 4.6, respectively) — or by a direct inductive proof. QED

We are now ready to prove Proposition 7.4.

Proof of Proposition 7.4. Starting with notation, let $\mathbb{A} = (A, \Theta_A, \Omega_A, a_I)$, $\mathbb{B} = (B, \Theta_B, \Omega_B, b_I)$ and $\mathbb{D} = (D, \Theta_D, \Omega_D, d_I)$, and let $G : D \rightarrow A^\sharp$ be the map witnessing that \mathbb{D} is a disjunctive companion of \mathbb{A} . We recall our notation Ran_A , Res_A and d_Q^+ from Definition 2.2.

Our goal is to provide player II with a winning strategy χ in the consequence game \mathcal{C} between $\mathbb{B}[\mathbb{D}/p]$ and $\mathbb{B}[\mathbb{A}/p]$. It will be convenient to make some simplifying assumptions on player I's moves in this game.

CLAIM 1 Without loss of generality we may assume that in any partial match Σ ending with (R, R') , player I always picks an element $Q = (Y, \mathcal{R})$ such that

(Ass1) $\text{Dom}(Q) \subseteq \text{Ran}(R)$ for all $Q \in \mathcal{R}$;

(Ass2) $Q \cap (D \times B) = \emptyset$, for all $Q \in \mathcal{R}$;

(Ass3) $|Q[d] \cap D| = 1$ for all $d \in D$ and all $Q \in \mathcal{R}$;

(Ass4) for all $b \in \text{Ran}R \cap B$, either $Q[b] \cap D = \emptyset$ for all $Q \in \mathcal{R}$, or $|Q[b] \cap D| = 1$ for all $Q \in \mathcal{R}$.

PROOF OF CLAIM We reason similarly as in Remark 5.13, which *mutatis mutandis* applies to the moves of player I in the consequence game as well, cf. Proposition 5.18. In particular: (Ass1) was explicitly listed as an assumption in Remark 5.13, and (Ass2) is based on the

second assumption in Remark 5.13, together with the fact that in the transition relation Θ_{BD} of $\mathbb{B}[\mathbb{D}/p]$, no state in B occurs in the one-step formula $\Theta_{BD}(d)$ of any $d \in D$. (Ass3) follows from the disjointivity of \mathbb{D} (cf. Proposition 6.11, we leave the details of this clause to the reader).

For (Ass4), consider $b \in \text{Ran}R \cap B$ and suppose that $\Gamma = (Y, \mathcal{R}) \in \mathcal{K}(B \cup D)$ is a legitimate move for Player I, which we may assume already satisfies (Ass2). We recall that $\Theta_{BD}(b) = \Theta_B(b)[\Theta_D(d_I)/p]$. If $\Gamma \notin \llbracket \Theta(d_I) \rrbracket_b^1$ then we set $\mathcal{R}' = \{Q \setminus (\{b\} \times D) \mid Q \in \mathcal{R}\}$. If $\Gamma \in \llbracket \Theta(d_I) \rrbracket_b^1$, then pick some disjoint $\hat{Y} \wedge \nabla C_b$ of $\Theta_D(d_I)$ such that $(B \cup D)^\sharp, \Gamma \Vdash_b^1 \hat{Y} \wedge \nabla C_b$. If $C_b = \emptyset$ then we must have $\mathcal{R} = \emptyset$ and set $\mathcal{R}_b = \mathcal{R} = \emptyset$. Otherwise, if $C_b \neq \emptyset$, set

$$\mathcal{R}_b = \{(Q \setminus (\{b\} \times D)) \cup \{(b, d)\} \mid Q \in \mathcal{R} \text{ and } (b, d) \in Q \cap (B \times C_b)\}$$

In each of these cases, we can verify that either $Q[b] \cap D$ is empty for all $Q \in \mathcal{R}_b$ or $Q[b] \cap D$ is a singleton for all $Q \in \mathcal{R}_b$, $\mathcal{R}_b \xrightarrow{\vec{P}} \subseteq \mathcal{R}$, and $(B \cup D)^\sharp, (Y, \mathcal{R}_b) \Vdash_b^1 \Theta_{BD}(b)$. The key observation is that if (Y, \mathcal{R}) satisfies $\Theta_D(d_I)$, then so does (Y, \mathcal{R}_b) , and this is proved by simply verifying that the required back-and-forth properties hold for C_b and \mathcal{R}_b . By repeating the procedure for all $b \in \text{Ran}R \cap B$, we will finally find some $\mathcal{R}' \xrightarrow{\vec{P}} \subseteq \mathcal{R}$ with the required properties. \blacktriangleleft

To appreciate the above claim, consider an arbitrary partial match

$$\Sigma = (R_0, R'_0), \dots, (R_k, R'_k),$$

with $R_0 = R'_0 = \{(b_I, b_I)\}$. It follows by Claim 1 that we may assume each element $c \in \text{Ran}R_k$ to lie on some trace through R_0, \dots, R_k , and that every trace through R_0, \dots, R_k is either a \mathbb{B} -trace, or else it consists of an initial, non-empty \mathbb{B} -trace, followed by a non-empty \mathbb{D} -trace. By the second and third assumption of the claim, traces are *D-functional*, that is, if $d \in \text{Ran}R_n$ for some $n < k$, then d has exactly one R_{n+1} -successor, that we will denote as d^+ . As a consequence, every trace τ on R_0, \dots, R_n ending at d has *exactly* one continuation through R_{n+1}, \dots, R_k . (This does not imply that all matches and traces are infinite.) The use of (Ass4) will be rather technical, uniformizing the transition of traces from the \mathbb{B} -part to the \mathbb{D} -part of the automaton $\mathbb{B}[\mathbb{D}/p]$.

A key role in our proof is played by a Σ -induced total order on $\text{Ran}_D R_k$ that we will introduce now. Intuitively, we say, for $d, d' \in \text{Ran}_D R_k$, that d is Σ -older than d' if d lies on a trace τ that entered D at an earlier stage than any trace arriving at d' .

For a formal definition of this ordering, we need to assume some arbitrary but fixed total order on D , given by an injective map $\text{mb} : D \rightarrow \omega$; we call $\text{mb}(d)$ the *birth minute* of d . The reason is that there may be “ties”, i.e situations where the longest D -trace leading to two different states in D are of the same length. Following the analogy: we can have cases where two states have the same “birth date”, and we then refer to the birth minute to decide which is the oldest.

Given a state $d \in \text{Ran}_D R_k$, by Claim 1(1) there is a trace τ through R_0, \dots, R_k such that $\tau(k) = d$. By Claim 1(2), all such traces start in B and at some moment j move to the \mathbb{D} -part of the automaton. We let $\text{tb}_\Sigma(d)$ be the smallest pair of natural numbers (j, l) in the lexicographic order on $\omega \times \omega$ such that there is some $e \in \text{Ran}_D R_j$ with $\text{mb}(e) = l$ and such

that the unique trace on $R_j \dots R_k$ beginning with e ends with d (this trace is unique because of trace functionality in D). The pair $\text{tb}_\Sigma(d) = (j, l)$ is called the *time of birth* of d relative to the match Σ ; we simply write $\text{tb}(d)$ if Σ is clear from context.

Note that tb_Σ is always an injective map. To see this, suppose that $\text{tb}_\Sigma(d) = \text{tb}_\Sigma(d') = (j, l)$. Then there are $e, e' \in \text{Ran}_D R_j$ such that the unique trace on R_j, \dots, R_k beginning with e ends with d , and the unique trace beginning with e' ends with d' , and such that $\text{mb}(e) = \text{mb}(e') = l$. By injectivity of mb , we get $e = e'$, and so we get $d = d'$ by uniqueness of traces in the \mathbb{D} -part of R_0, \dots, R_k .

Finally, we define a strict total ordering on $\text{Ran}_D R_k$ relative to Σ by saying that d is Σ -older than d' if $\text{tb}(d)$ is smaller than $\text{tb}(d')$ (in the lexicographic order). We leave it for the reader to verify that, for $d \in \text{Ran} R_n$ with $n < k$, it holds that $\text{tb}(d^+) \leq \text{tb}(d)$.

We now turn to the definition of player II's winning strategy χ . By a simultaneous induction on the length of a partial χ -match

$$\Sigma = (R_0, R'_0), \dots, (R_n, R'_n),$$

with $R_0 = R'_0 = \{(b_I, b_I)\}$, we will define maps

$$F_n : (B \cup D)^\# \rightarrow (B \cup A)^\#$$

and

$$g_n : \text{Ran}_A R'_n \rightarrow \text{Ran}_D R_n.$$

We let the F -maps determine player II's strategy in the following sense. Suppose that in the mentioned partial match Σ , player I legitimately picks an element $\Gamma = (\mathbb{Y}, \mathcal{R}) \in \mathbb{K}(B \cup D)^\#$. Then player II's response will be the map $F_{n+1} \upharpoonright_{\mathcal{R}}$, that is, the map F_{n+1} , restricted to the set $\mathcal{R} \subseteq (B \cup D)^\#$.

Inductively we will ensure that the following conditions are maintained:

- (*) $F_n R_n = R'_n$,
- (†0) $R'_n = \text{Res}_B R'_n \cup (R'_n \cap (B \times A)) \cup \text{Res}_A R'_n$,
- (†1) $\text{Res}_B R'_n = \text{Res}_B R_n$,
- (†2) $R'_n \cap (B \times A) \subseteq \bigcup_{d \in D} \{(b, a) \mid (b, d) \in R_n \cap (B \times D) \ \& \ (a_I, a) \in G(d)\}$
- (†3) $\text{Res}_A R'_n \subseteq \bigcup \{G(d) \mid d \in \text{Ran}_D R_n\}$,
- (‡) $a \in \text{Ran} G(g_n a)$, for all $a \in \text{Ran}_A R_n$.

For some explanation and motivation of these conditions, observe that (*) indicates that Σ itself is indeed χ -guided. For condition (†), first observe that while by Claim 1, all $\mathbb{B}[\mathbb{D}/p]$ -traces consist of an initial \mathbb{B} -part followed by an \mathbb{D} -tail, condition (†0) states that similarly, all $\mathbb{B}[\mathbb{A}/p]$ -traces consist of an initial \mathbb{B} -part followed by an \mathbb{A} -tail. Condition (†1) then states that the \mathbb{B} -part on the left and right side of a $\mathcal{C}(\mathbb{B}[\mathbb{D}/p], \mathbb{B}[\mathbb{A}/p])$ -match is the same, and condition (†3) states that every pair $(a, b) \in \text{Res}_A \text{Ran} R'_n$ is 'covered' or 'implied' by some $d \in \text{Ran}_D R_n$. Finally, (‡) states that, for every $a \in \text{Ran} R'_n$, the map g_n picks a specific element $d \in \text{Ran}_D R_n$ such that $a \in \text{Ran}(Gd)$. As we will see in Claim 4 below, it will be this condition, together with the condition on the reflection of traces in Definition 7.2 and the actual definition of the maps g_n , that is pivotal in proving that player II wins all infinite matches.

Setting up the induction, observe that $R_0 = R'_0 = \{(b_I, b_I)\}$. Defining F_0 as the map $R \mapsto \text{Res}_B R$ and g_0 as the empty map, we can easily check that $(*)$, (\dagger) and (\ddagger) hold.

In the inductive case we will define the maps F_{n+1} and g_{n+1} for a partial match Σ as above. For the definition of $F_{n+1} : (B \cup D)^\# \rightarrow (B \cup A)^\#$, first observe that that by $(\dagger 0)$ we are only interested in relations $R \in (B \cup D)^\#$ that are of the form $R = \text{Res}_B R \cup (R \cap (B \times D)) \cup \text{Res}_D R$. We will define F_{n+1} by treating these three parts of R separately, using, respectively, the identity map on $B^\#$ and two auxiliary maps that we define now.

For the D -part of R , we define an auxiliary map $H_{n+1} : D \times D \rightarrow A^\#$:

$$H_{n+1}(d, d') := G(d') \cap (g_n^{-1}(d) \times A),$$

that is, $H_{n+1}(d, d')$ consists of those pairs $(a, a') \in G(d')$ for which $g_n(a) = d$. For the $B \times D$ -part of R , we need a second auxiliary map $L : B \times D \rightarrow \mathcal{P}(B \times A)$, given by

$$L(b, d) := \{(b, a) \in B \times A \mid (a_I, a) \in G(d)\}.$$

Now we define $F_{n+1} : (B \cup D)^\# \rightarrow (B \cup A)^\#$ as follows:

$$\begin{aligned} F_{n+1}(R) &:= \text{Res}_B R \\ &\cup \bigcup \{L(b, d) \mid (b, d) \in R \cap (B \times D)\} \\ &\cup \bigcup \{H_{n+1}(d, d^+) \mid (d, d^+) \in \text{Res}_D R\}. \end{aligned}$$

That is, we define $F_{n+1}(R)$ as the union of three disjoint parts: a $B \times B$ -part, a $B \times A$ -part and an $A \times A$ -part.

For the definition of g_{n+1} , let (R_{n+1}, R'_{n+1}) be an arbitrary next basic position following the partial match Σ . Note that we may assume that R_{n+1} satisfies the assumptions formulated in Claim 1, and that we have $R'_{n+1} = F_{n+1}(R_{n+1})$ by the fact that player II's strategy is given by the map F_{n+1} . Given $a \in \text{Ran}_A R'_{n+1}$, distinguish cases:

Case 1 If a has no R'_{n+1} -predecessor in A , then by definition of F_{n+1} and L , the set of states $d \in D$ for which there is a $b \in B$ with $(b, d) \in R_{n+1}$ and $(a_I, a) \in G(d)$ is non-empty. We define $g_{n+1}a$ to be the *oldest* element of this set, that is, in this case, the element with the earliest birth minute.

Case 2 If a does have an R'_{n+1} -predecessor in A , that is, the set $\{b \in A \mid (b, a) \in R'_{n+1}\}$ is non-empty, then we can define $g_{n+1}a$ to be the *oldest* element (with respect to the match $\Sigma \cdot (R_{n+1}, R'_{n+1})$) of the set $\{(g_n b)^+ \mid (b, a) \in R'_{n+1}\} \subseteq D$.

To gain some intuitions concerning this definition, observe that in the first case, we cannot define $g_{n+1}a$ inductively on the basis of the map g_n applied to an R'_{n+1} -predecessor of a : we have to start from scratch. This case only applies, however, in a situation where a does have an R'_{n+1} -successor $b \in B$ such that in R_{n+1} , this same b has a R_{n+1} -successor $d \in D$ such that $(a_I, a) \in Gd$. In this case we simply define $g_{n+1}a := d$, and if there are more such pairs (b, d) , then for $g_{n+1}a$ we may pick any of these d 's, for instance the one with the earliest birth minute.

We now turn to the second clause of the definition of g_{n+1} — here lies, in fact, the heart of the proof of Proposition 7.4. Consider a situation where a_0 and a_1 , both in A , are the two R_{n+1} -predecessors of $a \in A$. Both $g_n a_0$ and $g_n a_1$ are states in D , and therefore they have unique R_{n+1} -successors in D , denoted by $(g_n a_0)^+$ and $(g_n a_1)^+$, respectively. We want to define $g_{n+1} a$ as either $(g_n a_0)^+$ or $(g_n a_1)^+$, but then we are facing a *choice* between these two states of D in case they are *distinct*. It is here that our match-dependent ordering of states in D comes in: we will define $g_{n+1} a$ as the *oldest* element of the two, relative to the (extended) match $\Sigma \cdot (R_{n+1}, R'_{n+1})$. Suppose (without loss of generality) it holds that $(g_n a_0)^+$ is older than $(g_n a_1)^+$, so that we put $g_{n+1} a := (g_n a_0)^+$. In this case we say that the trace through $g_n a_0$ is *continued*, while there is also a *trace jump* witnessed by the fact that $(a_1, a) \in R_{n+1}$ but $(g_n a_1, g_{n+1} a) \notin R'_{n+1}$ (see Figure 1, where the dashed lines represent the g -maps, and the partial trace of white points on the right is not mapped to a partial trace on the left, due to a trace jump).




Figure 1: A trace merge results in a trace jump.

CLAIM 2 By playing according to the strategy χ , player II indeed maintains the conditions (*), (\dagger) and (\ddagger).

PROOF OF CLAIM Let Σ be a partial χ -match satisfying the conditions (*), (\dagger) and (\ddagger), and let $(R_{n+1}, R'_{n+1}) \in \text{Gr}(F_{n+1})$ be any possible next position. It suffices to show that (R_{n+1}, R'_{n+1}) also satisfies (*), (\dagger) and (\ddagger).

The conditions (*), ($\dagger 0$), ($\dagger 1$) and ($\dagger 2$) are direct consequences of the definition of F_{n+1} , while ($\dagger 3$) is immediate by the fact that

$$(b, a) \in F_{n+1} R_{n+1} \iff (b, a) \in G((g_n b)^+). \quad (13)$$

for all $b, a \in A$. To prove (13), consider the following chain of equivalences, which hold for all

$b, a \in A$:

$$\begin{aligned}
(b, a) \in F_{n+1}R_{n+1} &\iff (b, a) \in H_{n+1}(d, d^+), \text{ some } (d, d^+) \in \text{Res}_D R_n && \text{(Def. } F_{n+1}) \\
&\iff (b, a) \in G(d^+), \text{ some } (d, d^+) \in \text{Res}_D R_n \text{ with } d = g_n b && \text{(Def. } H_{n+1}) \\
&\iff (b, a) \in G((g_n b)^+). && \text{(obvious)}
\end{aligned}$$

Finally, for condition (\ddagger) , let $a \in \text{Ran}_A R'_{n+1}$ be arbitrary. If a has an R'_{n+1} -predecessor in A , then we are in case 2 of the definition of $g_{n+1}a$, where $g_{n+1}a$ is of the form $(g_n b)^+$ for some b with $(b, a) \in \text{Res}_A R'_{n+1}$. But then $(b, a) \in G((g_n b)^+)$ by (13), so that indeed we find $a \in \text{Ran}G(g_{n+1}a)$. If, on the other hand, a has no R_{n+1} -predecessor in A , then we are in case 1 of the definition of $g_{n+1}a$. In this case, $g_{n+1}a$ is an element of a set, each of whose elements d satisfies $a \in \text{Ran}G(d)$; so we certainly have $a \in \text{Ran}G(g_{n+1}a)$. \blacktriangleleft

CLAIM 3 The moves for player II prescribed by the strategy χ are legitimate.

PROOF OF CLAIM Let Θ_{BD} and Θ_{BA} denote the transition maps of the automata $\mathbb{B}[\mathbb{D}/p]$ and $\mathbb{B}[\mathbb{A}/p]$, respectively. Consider a partial match Σ ending with the position (R_n, R'_n) and a subsequent move $\Gamma = (Y, \mathcal{R}) \in \mathcal{K}(B \cup D)^\sharp$ by player I such that $\Gamma \Vdash_e^1 \Theta_{BD}(e)$ for all $e \in \text{Ran}R_n$. We need to show that

$$(\mathcal{K}F_{n+1})\Gamma \Vdash_c^1 \Theta_{BA}(c) \quad (14)$$

for an arbitrary element $c \in \text{Ran}R'_n$. Since $c \in B \cup A$ by definition of $\mathbb{B}[\mathbb{A}/p]$, one of the following two cases applies.

Case 1 $c \in A$. Then by (\ddagger) we find $c \in \text{Ran}(Gd)$, where $d := g_n c$ belongs to $\text{Ran}_D R_n$. Let the map $\text{succ}_d : \mathcal{R} \rightarrow D$ be given by $\text{succ}_d(Q) := d_Q^+$ — this is well-defined by (Ass4) in Claim 1. We will apply the second clause of Definition 7.2 with $\delta = (\mathcal{K}\text{succ}_d)\Gamma$.

As an immediate consequence of the assumption that Γ is a legitimate move of player I and the fact that $\Theta_{BD}(d) = \Theta_D(d)$, we find

$$\Gamma \Vdash_d^1 \Theta_D(d). \quad (15)$$

From this and the fact that $\Theta_D(d)$ is a one-step formula in D , it easily follows that

$$(\mathcal{K}\text{sing})(\mathcal{K}\text{succ}_d)\Gamma \Vdash_J^1 \Theta_D(d). \quad (16)$$

Now we can use the assumption that (\mathbb{D}, d) is a disjunctive companion of (\mathbb{A}, a) , obtaining from clause (DC2) that

$$(\mathcal{K}G)(\mathcal{K}\text{succ}_d)\Gamma \Vdash_c^1 \Theta_A(c). \quad (17)$$

By functoriality of \mathcal{K} and the fact that $\Theta_A(c) = \Theta_{BA}(c)$, this is equivalent to

$$(\mathcal{K}(G \circ \text{succ}_d))\Gamma \Vdash_c^1 \Theta_{BA}(c). \quad (18)$$

We now claim that

$$\text{for all } Q \in \mathcal{R}, a \in A : (c, a) \in (G \circ \text{succ}_d)(Q) \implies (c, a) \in F_{n+1}Q. \quad (19)$$

For a proof of (19), assume that $(c, a) \in (G \circ \text{succ}_d)(Q) = G(d_Q^+)$. Then (c, a) belongs to $H_{n+1}(d, d_Q^+)$ by definition of H_{n+1} , and to $F_{n+1}Q$ by definition of F_{n+1} .

It easily follows from (19) and the observations that (with $\Gamma = \mathbf{Y}, \mathcal{R}$) we have $\mathbf{K}(G \circ \text{succ}_d)\Gamma = (\mathbf{Y}, \{(G \circ \text{succ}_d)(Q) \mid Q \in \mathcal{R}\})$ and $(\mathbf{K}F_{n+1})\Gamma = (\mathbf{Y}, \{F_{n+1}(Q) \mid Q \in \mathcal{R}\})$, that

$$\mathbf{K}(G \circ \text{succ}_d)\Gamma \xrightarrow{A}_{c,c} (\mathbf{K}F_{n+1})\Gamma. \quad (20)$$

But from this Proposition 7.8 yields (14), as required.

Case 2 $c \in B$. Note that in this case we have $\Theta_{BA}(c) = \Theta_B(c)[\Theta_A(a_I)/p]$ and $\Theta_{BD}(c) = \Theta_B(c)[\Theta_D(d_I)/p]$. Thus by assumption, we know that $\Gamma \Vdash_c^1 \Theta_B(c)[\Theta_D(d_I)/p]$, while we need to establish that $(\mathbf{K}F_{n+1})\Gamma \Vdash_c^1 \Theta_B(c)[\Theta_A(a_I)/p]$. To achieve this it clearly suffices to show that

$$\Gamma \Vdash_c^1 \alpha[\Theta_D(d_I)/p] \text{ implies } (\mathbf{K}F_{n+1})\Gamma \Vdash_c^1 \alpha[\Theta_A(a_I)/p] \quad (21)$$

for all $\alpha \in \mathbf{1ML}(\mathbf{X}, B)$. We will prove (21) by induction on the one-step formula α , taken as a lattice term over the set $\{p\} \cup \mathbf{1ML}(\mathbf{X} \setminus \{p\}, B)$. This perspective allows us to distinguish the following two cases in the induction base.

Base Case a: $\alpha = p$. Here we find $\alpha[\Theta_D(d_I)/p] = \Theta_D(d_I)$ and $\alpha[\Theta_A(a_I)/p] = \Theta_A(a_I)$.

We first prove that

$$|Q[c] \cap D| = 1, \text{ for all } Q \in \mathcal{R}. \quad (22)$$

To see this, note that from $\Gamma \Vdash_c^1 \Theta_D(d_I)$ it follows by the shape of disjunctive one-step formulas that either $\Gamma \Vdash_c^1 \nabla \emptyset$ or $\Gamma \Vdash_c^1 \nabla E$ for some non-empty $E \subseteq D$. In the first case we find by (4) in Remark 6.10 that $\mathcal{R} = \emptyset$, which clearly satisfies (22). In the second case we obtain by (3) in the same Remark that $Q[c] \cap D \neq \emptyset$ for some $Q \in \mathcal{R}$, so that (22) follows by (Ass4) in Claim 1.

But by (22) we may assume the existence of a map $\text{succ}_c : \mathcal{R} \rightarrow D$ such that $Q[c] = \{\text{succ}_c(Q)\}$ for all $Q \in \mathcal{R}$. (This covers the case where $\mathcal{R} = \emptyset$.) It easily follows from $\Gamma \Vdash_c^1 \Theta_D(d_I)$ that $(\mathbf{K}\text{sing})(\mathbf{K}\text{succ}_c)\Gamma \Vdash_I^1 \Theta_D(d_I)$. Hence, by Definition 7.2 and functoriality of \mathbf{K} we obtain

$$\mathbf{K}(G \circ \text{succ}_c)\Gamma \Vdash_{a_I}^1 \Theta_A(a_I). \quad (23)$$

We leave it for the reader to verify that the definition of the maps L and F_{n+1} implies

$$(a_I, a) \in G(d) \ \& \ (c, d) \in Q \implies (c, a) \in L(c, d) \subseteq F_{n+1}(Q),$$

and hence, similar to the situation in case 1 above, is tailored towards establishing

$$\mathbf{K}(G \circ \text{succ}_c)\Gamma \xrightarrow{A}_{a_I,c} (\mathbf{K}F_{n+1})\Gamma. \quad (24)$$

But from (23) and (24) it is immediate by Proposition 7.8 that

$$(\mathsf{KF}_{n+1})\Gamma \Vdash_c^1 \Theta_A(a_I), \quad (25)$$

as required.

Base Case b: $\alpha \in \mathsf{1ML}(\mathsf{X} \setminus \{p\}, B)$, that is, α is a p -free one-step formula over B . This case is in fact easy, first of all because $\alpha[\Theta_D(d_I)/p] = \alpha[\Theta_A(a_I)/p] = \alpha$. Furthermore, by ($\dagger 1$) Γ and $(\mathsf{KF}_{n+1})\Gamma$ coincide as one-step models over B . Formally:

$$\begin{aligned} \Gamma \Vdash_c^1 \alpha &\iff (\mathsf{KRes}_B)\Gamma \Vdash_c^1 \alpha && \text{(Proposition 7.8)} \\ &\iff (\mathsf{K}(\mathsf{Res}_B \circ F_{n+1}))\Gamma \Vdash_c^1 \alpha && (\dagger 1) \\ &\iff (\mathsf{KRes}_B)(\mathsf{KF}_{n+1})\Gamma \Vdash_c^1 \alpha && \text{(functoriality of } \mathsf{K}) \\ &\iff (\mathsf{KF}_{n+1})\Gamma \Vdash_c^1 \alpha && \text{(Proposition 7.8)} \end{aligned}$$

Here we leave it for the reader to verify that $\Gamma \xleftrightarrow{cc}^B (\mathsf{KRes}_B)\Gamma$ and $(\mathsf{KRes}_B)(\mathsf{KF}_{n+1})\Gamma \xleftrightarrow{cc}^B (\mathsf{KF}_{n+1})\Gamma$, in the two respective steps where we apply Proposition 7.8.

Inductive case The inductive case in the proof of (21) is trivial.

This finishes the proof of Claim 3. ◀

CLAIM 4 Suppose Σ is an infinite χ -guided match with basic positions

$$(R_0, R'_0)(R_1, R'_1)(R_2, R'_2) \dots$$

such that the stream $R'_0 R'_1 R'_2 \dots$ contains a bad trace. Then there is a bad trace on $R_0 R_1 R_2 \dots$ as well.

PROOF OF CLAIM Fix a χ -guided match $\Sigma = (R_i, R'_i)_{i \geq 0}$ and a bad trace τ on $(R'_i)_{i \geq 0}$, as above. We will show that there is a bad trace on the stream $(R_i)_{i \geq 0}$ as well.

There are two possibilities for τ . In case τ stays entirely in B , then by ($\dagger 1$), τ is also a trace on $R_0 R_1 R_2 \dots$, and so we are done. Hence we may focus on the second case, where from some finite stage onwards, τ stays entirely in A . So suppose τ is an infinite trace of the form

$$\tau = b_0 b_1 \dots b_n a_{n+1} a_{n+2} a_{n+3} \dots,$$

where the b_j are all in B , and the a_i are all in A . Our key claim is the following:

$$\text{there exists an index } k > n \text{ such that } g_{j+1} a_{j+1} = (g_j a_j)^+ \text{ for all } j \geq k. \quad (26)$$

In order to prove (26), recall that a *trace jump* occurs at the index $j > n$ if we have $g_{j+1} a_{j+1} \neq (g_j a_j)^+$. We want to show that there can only be finitely many j at which a trace jump occurs. If no trace jump occurs at j , then we have

$$\text{tb}(g_j a_j) \geq \text{tb}((g_j a_j)^+) = \text{tb}(g_{j+1} a_{j+1}).$$

Hence, it suffices to prove that if a trace jump occurs at j then $\text{tb}(g_{j+1}a_{j+1})$ is strictly smaller than $\text{tb}(g_j a_j)$ in the lexicographic order. It then follows that the stream

$$\text{tb}(g_k a_k), \text{tb}(g_{k+1} a_{k+1}), \text{tb}(g_{k+2} a_{k+2}), \dots$$

is a stream of pairs of natural numbers that never increases, and strictly decreases at each j at which a trace jump occurs. By well-foundedness of the lexicographic order on $\omega \times \omega$ this can therefore only happen finitely many times, as required.

So we are left with the task of proving that tb is strictly decreasing at each index j for which a trace jump occurs. To see that this is indeed so, suppose that $g_{j+1}a_{j+1} \neq (g_j a_j)^+$. Recall that we defined $g_{j+1}a_{j+1}$ to be the oldest element of the set

$$\{(g_j c)^+ \mid (c, a_{j+1}) \in R'_{j+1}\}.$$

But since $(a_j, a_{j+1}) \in R'_{j+1}$, it follows that $g_{j+1}a_{j+1}$ must be older than $(g_j a_j)^+$, with respect to the age relation induced by the match $(R_0, R'_0), \dots, (R_{j+1}, R'_{j+1})$, and so $\text{tb}(g_{j+1}a_{j+1})$ must be strictly smaller than $\text{tb}((g_j a_j)^+) \leq \text{tb}(g_j a_j)$, as required. This completes the proof of (26).

Let us finally see how (26) entails Claim 4. Suppose there exists an index k as in (26), and consider $g_k a_k \in \text{Ran}_D R_k$. Pick an arbitrary initial trace $b_0 \dots b_n d_{n+1} \dots d_k$ of $R_0 \dots R_k$ leading up to $g_k a_k = d_k$ (as mentioned already after Claim 1, the existence of such a trace follows from our assumptions on player I's strategy). Then the stream

$$b_0, \dots, d_{k-1}, g_k a_k, g_{k+1} a_{k+1}, g_{k+2} a_{k+2}, \dots$$

is a trace of $R_0 R_1 R_2 \dots$ by the property of the index k described in (26). Furthermore, it follows that $a_k a_{k+1} a_{k+2} \dots$ is a trace of the stream

$$G(g_k a_k), G(g_{k+1} a_{k+1}), G(g_{k+2} a_{k+2}), \dots$$

To see why, consider the pair (a_j, a_{j+1}) where $j \geq k$. Then $(a_j, a_{j+1}) \in R'_{j+1} = F_k(R_{j+1})$, so there is some $(d, d') \in R_{j+1}$ with $(a_j, a_{j+1}) \in H_{j+1}(d, d')$. Hence $d = g_j a_j$ and $(a_j, a_{j+1}) \in G(d')$.

But $d' = d^+$ by functionality of traces on D (which follows from the third assumption in Claim 1), and so we find $d' = d^+ = (g_j a_j)^+ = g_{j+1} a_{j+1}$. From this we get $(a_j, a_{j+1}) \in G(g_{j+1} a_{j+1})$ as required. Note too that $a_k a_{k+1} a_{k+2} \dots$ has the same tail as τ , and hence it is a *bad* trace too. It now follows from the trace reflection clause of Definition 7.2 that $g_k a_k, g_{k+1} a_{k+1}, g_{k+2} a_{k+2}, \dots$ is itself a bad trace, and so we have found a bad trace on $R_0 R_1 R_2 \dots$ as required. \blacktriangleleft

Finally, the proof of the Proposition is immediate by the last two claims: it follows from Claim 3 that player II never gets stuck, so that we need not worry about finite matches. But Claim 4 states that II wins all infinite matches of $\mathcal{C}(\mathbb{B}[\mathbb{D}/p], \mathbb{B}[\mathbb{A}/p])$ as well. QED

8 From automata to formulas

In section 4.3 we defined an inductive translation from formulas to modal automata, based on operations on automata corresponding to Boolean connectives, modalities and fixpoint operators. In this section we provide a translation tr in the opposite direction, that is, from automata to formulas, and we establish some properties of this translation. Our definition of the translation map is based on a more or less standard [27] induction on the complexity of the automaton.

The main purpose of this translation is the following statement that we already mentioned in the introduction to the paper as one of our main lemmas:

Theorem 2 *For every formula $\varphi \in \mu\text{ML}$, we have $\varphi \equiv_K \text{tr}(\mathbb{A}_\varphi)$.*

The proof of this proposition will proceed by induction on the complexity of formulas. As a central auxiliary result (Proposition 8.15 below) we will show that the translation commutes with the logical operations on automata and formulas, and with the operation of substitution.

The point is that, allowing us to apply proof-theoretic notions such as derivability or consistency to automata, it is Theorem 2 that opens the door to proof theory for automata.

Definition 8.1 A modal automaton \mathbb{A} will be called *consistent* if the formula $\text{tr}(\mathbb{A})$ is consistent. Given two modal automata \mathbb{A} and \mathbb{B} , we say that \mathbb{A} provably implies \mathbb{B} , notation: $\mathbb{A} \leq_K \mathbb{B}$, if $\text{tr}(\mathbb{A}) \leq_K \text{tr}(\mathbb{B})$, and that \mathbb{A} and \mathbb{B} are provably equivalent if $\text{tr}(\mathbb{A}) \equiv_K \text{tr}(\mathbb{B})$. We will use similar notation and terminology relating formulas and automata, for instance we will say that φ *provably implies* \mathbb{A} and write $\varphi \leq_K \mathbb{A}$ if $\varphi \leq_K \text{tr}(\mathbb{A})$, etc. \triangleleft

In order to provide the translation $\text{tr}(\mathbb{A})$ of an automaton \mathbb{A} , we first define a map $\text{tr}_\mathbb{A}$ assigning a formula to *each* state of \mathbb{A} . The formula $\text{tr}(\mathbb{A})$ is then obtained by applying the map $\text{tr}_\mathbb{A}$ to the initial state of \mathbb{A} . Three minor modifications of our earlier definitions will turn out to be convenient for a smooth inductive proof.

First, it will be convenient to generalize the definition of a modal automaton to the extent that we allow guarded occurrences of proposition letters in the range of the transition map.

Definition 8.2 A *generalized modal automaton* is a structure $\mathbb{A} = (A, \Theta, \Omega, a_I)$ where A , Ω and a_I are as in the definition of standard modal automata, and the transition map Θ is of type $\Theta : A \rightarrow 1\text{ML}(\mathbf{X}, A \cup \mathbf{X})$. \triangleleft

The notion of acceptance for generalized automata is a straightforward generalization of the one for standard modal automata. For completeness we provide a definition here — one that stays close to our approach in terms of one-step models is the following.

Definition 8.3 A *generalized one-step model* is a structure (\mathbf{Y}, S, m) such that S is some set, $\mathbf{Y} : S \uplus \{\star\} \rightarrow \text{PX}$ is a \mathbf{X} -marking on the set $S \uplus \{\star\}$ and m is an A -marking on the set S . The one-step satisfaction relation \Vdash^1 for generalized one-step formulas in $1\text{ML}(\mathbf{X}, A \cup \mathbf{X})$ is defined in the most obvious way: we treat a generalized one-step model (\mathbf{Y}, S, m) as if it were the standard one-step model $(\mathbf{Y}(\star), S, \mathbf{Y}|_S \cup m)$ over $(\mathbf{X}, \mathbf{X} \cup A)$.

Then given a generalized modal automaton \mathbb{A} and Kripke model $\mathbb{S} = (S, R, V)$, the rules of the *acceptance game* $\mathcal{A}(\mathbb{A}, \mathbb{S})$ for \mathbb{A} with respect to \mathbb{S} can be defined using the following table:

Position	Player	Admissible moves
$(a, s) \in A \times S$	\exists	$\{m : R[s] \rightarrow \text{PA} \mid (V^\dagger \upharpoonright_{\{s\} \cup R[s]}, R[s], m) \Vdash^{-1} \Theta(a)\}$
m	\forall	$\{(b, t) \mid b \in m(t)\}$

The *winning conditions* and the notion of *acceptance* are as in the acceptance game for standard modal automata. \triangleleft

Remark 8.4 This generalization of modal automata is for technical convenience only. Similar to the approach taken in Definition 4.19, given a generalized automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$ we may define the structure $\mathbb{A}^s = (A^s, \Theta^s, \Omega^s, a_I)$, by putting $A^s := A \cup \{\underline{a} \mid a \in A\}$, $\Theta^s(a) := \Theta(a)[b/b \mid b \in A]$, $\Theta^s(\underline{a}) := a$, $\Omega^s(a) := \Omega(a)$, and $\Omega^s(\underline{a}) := 0$. It is easy to see that \mathbb{A}^s is always a standard modal automaton, equivalent to \mathbb{A} . We do not pursue this approach here, since it would lead to some technical complications that obscure the important issues. \triangleleft

Second, it will make sense to define the mentioned translation map $\text{tr}_{\mathbb{A}}$ for ‘uninitialized’ automata, i.e., structures (A, Θ, Ω) that could be called (generalised) automata if they did not lack an initial state.

Definition 8.5 An *automaton structure* is a triple $\mathbb{A} = (A, \Theta, \Omega)$ such that A is a finite, non-empty set endowed with a transition map $\Theta : A \rightarrow \mathbf{1ML}(\mathbf{X}, A \cup \mathbf{X})$ and a priority function $\Omega : A \rightarrow \omega$.

The *underlying automaton structure* of a (generalized) modal automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$ is given as the triple $\underline{\mathbb{A}} := (A, \Theta, \Omega)$. Conversely, given an automaton structure $\mathbb{A} = (A, \Theta, \Omega)$ and a state a in A , we let $\mathbb{A}(a)$ denote the initialized automaton (A, Θ, Ω, a) . \triangleleft

Many concepts that we defined for automata in fact apply to automaton structures in the most obvious way, and we will use this observation without further notice.

Finally, the restriction that we announced is that for our definition of the translation map $\text{tr}_{\mathbb{A}}$ we will first confine our attention to so-called *linear* automaton structures.

Definition 8.6 An automaton structure $\mathbb{A} = (A, \Theta, \Omega)$ will be called *linear* if the relation $\sqsubset_{\mathbb{A}}$ is a strict linear order satisfying $(\triangleleft_{\mathbb{A}} \setminus \triangleright_{\mathbb{A}}) \subseteq \sqsubset_{\mathbb{A}}$.

Given two automaton structures $\mathbb{A} = (A, \Theta, \Omega)$ and $\mathbb{A}' = (A, \Theta, \Omega')$, we say that \mathbb{A}' is a *refinement* of \mathbb{A} if

- (1) the partial order $\sqsubseteq_{\mathbb{A}}$ is clusterwise contained in $\sqsubseteq_{\mathbb{A}'}$, i.e., $a \bowtie b$ and $a \sqsubseteq_{\mathbb{A}} b$ imply $a \sqsubseteq_{\mathbb{A}'} b$; and
- (2) $\Omega'(a')$ has the same parity as $\Omega(a)$, for all $a \in A$.

A linear refinement is called a *linearization*. \triangleleft

In words: linear automata structures have an injective priority map Ω , and satisfy the condition that if one state a is active in another state b , but not vice versa, then $a \sqsubset b$. In other words, the priority of states goes down if a match of the acceptance game passes from one cluster to the next. Our focus on linear automaton structures is justified by the following proposition.

Proposition 8.7 *Every automaton structure \mathbb{A} has a linearization \mathbb{A}^l such that, for all $a \in A$*

(1) $\mathbb{A}\langle a \rangle \models_{\mathbb{G}} \mathbb{A}^l\langle a \rangle$ and $\mathbb{A}^l\langle a \rangle \models_{\mathbb{G}} \mathbb{A}\langle a \rangle$;

(2) each player $\Pi \in \{\exists, \forall\}$ has a winning strategy in $\mathcal{S}(\mathbb{A}\langle a \rangle)$ (resp. $\mathcal{S}_{thin}(\mathbb{A}\langle a \rangle)$) iff she/he has a winning strategy in $\mathcal{S}(\mathbb{A}^l\langle a \rangle)$ (resp. $\mathcal{S}_{thin}(\mathbb{A}^l\langle a \rangle)$).

Proof. One may easily obtain a linearization \mathbb{A}^l of \mathbb{A} , so it suffices to prove that the statements in (1) and (2) hold for an arbitrary refinement \mathbb{A}' of \mathbb{A} and an arbitrary state a in \mathbb{A} . To prove (1), it is straightforward to verify that the identity map on A^\sharp provides a winning strategy for player I in both $\mathcal{C}(\mathbb{A}\langle a \rangle, \mathbb{A}'\langle a \rangle)$ and $\mathcal{C}(\mathbb{A}'\langle a \rangle, \mathbb{A}\langle a \rangle)$. And to prove (2), it is equally straightforward to verify that a winning strategy for \exists in the (thin) satisfiability game for $\mathbb{A}\langle a \rangle$ is also a winning strategy for her in the (thin) satisfiability game for $\mathbb{A}'\langle a \rangle$, and vice versa. Part (2) then easily follows by the determinacy of the (thin) satisfiability game. QED

The advantage of working with linear automaton structures is that we may define the translation map by a simple induction on the *size* of the structure.

Definition 8.8 By induction on the size of a linear modal \mathbf{X} -automaton structure \mathbb{A} we define a map $\mathbf{tr}_{\mathbb{A}} : A \rightarrow \mu\mathbf{ML}(\mathbf{X})$. Recall that our notation for formula substitution has been given in Definition 3.3.

In the base case of the induction we are dealing with an automaton structure \mathbb{A} based on a single state a . Then we define

$$\mathbf{tr}_{\mathbb{A}}(a) := \eta_a a. \Theta(a),$$

where $\eta_a \in \{\mu, \nu\}$ denotes the type of a .

In the inductive case, where $|\mathbb{A}| > 1$, by injectivity of Ω there is a unique state $m \in A$ that reaches the maximal priority, that is, with $\Omega(m) = \max \Omega[A]$. Let $\eta = \eta_m$ be the fixpoint type of m . Define \mathbb{A}^- to be the $\mathbf{X} \cup \{m\}$ -automaton structure $(A^-, \Theta^-, \Omega^-)$ with

$$A^- := A \setminus \{m\}$$

$$\Theta^- := \Theta \upharpoonright_{A^-}$$

$$\Omega^- := \Omega \upharpoonright_{A^-} .$$

Clearly we have $|\mathbb{A}^-| < |\mathbb{A}|$, so that inductively we may assume a map $\mathbf{tr}_{\mathbb{A}^-} : A^- \rightarrow \mu\mathbf{ML}(\mathbf{X} \cup \{m\})$. (Our motivation for introducing generalized modal automata stems from the observation that $\Theta^-(a)$ generally will have guarded occurrences of m , which in \mathbb{A}^- is no longer a state of the automaton but a proposition letter.)

The map $\mathbf{tr}_{\mathbb{A}}$ is now defined in two steps. First we define $\mathbf{tr}_{\mathbb{A}}(m)$ as follows:

$$\mathbf{tr}_{\mathbb{A}}(m) := \eta m. \Theta(m)[\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-].$$

Second, by putting

$$\mathbf{tr}_{\mathbb{A}}(a) := \mathbf{tr}_{\mathbb{A}^-}(a)[\mathbf{tr}_{\mathbb{A}}(m)/m]$$

we define $\mathbf{tr}_{\mathbb{A}}(a)$ for each $a \neq m$. ◁

Remark 8.9 An alternative approach would be to define the translation by induction on the *index* of an automaton, i.e., the size of the range of the priority map. In this approach, one would not have a unique maximal state, but a *set* of maximal states $\{m_1, \dots, m_n\}$, and the automaton structure \mathbb{A}^- would remove all the maximal states. We would then get a set of “equations” $m_i := \Theta(m_i)[\text{tr}_{\mathbb{A}^-}(b) \mid b \sqsubseteq_{\mathbb{A}} m_i]$, which is solved by a formula of the *vectorial μ -calculus* [1], and this formula can then be translated into the one-dimensional μ -calculus using the Bekič principle for simultaneous fixpoints. \triangleleft

We now turn to the translation map for arbitrary automaton structures. By standard order theory every automaton structure has at least one linearization. Furthermore, by the following result the translation maps of different linearizations of the same structure are provably equivalent.

Proposition 8.10 *Let $\mathbb{A}' = (A, \Theta, \Omega')$ and $\mathbb{A}'' = (A, \Theta, \Omega'')$ be two linearizations of the automaton structure $\mathbb{A} = (A, \Theta, \Omega)$. Then*

$$\text{tr}_{\mathbb{A}'}(a) \equiv_K \text{tr}_{\mathbb{A}''}(a)$$

for all $a \in A$.

Proof. The proof of this proposition is conceptually straightforward, boiling down to the observation in Fact 3.15 that $\mu x \mu y. \varphi(x, y) \equiv_K \mu y \mu x. \varphi(x, y)$, for any formula $\varphi(x, y)$. We leave the technical details to the reader. QED

Proposition 8.10 ensures that modulo provable equivalence the following definition of $\text{tr}(\mathbb{A})$ for an arbitrary automaton \mathbb{A} does not depend on the particular choice of a linearization for the underlying automaton structure of \mathbb{A} .

Definition 8.11 With each automaton structure $\mathbb{A} = (A, \Theta, \Omega)$ we associate an arbitrary but fixed linearization \mathbb{A}^l of \mathbb{A} (with the understanding that $\mathbb{A}^l = \mathbb{A}$ in case \mathbb{A} itself is linear). We then define $\text{tr}_{\mathbb{A}} := \text{tr}_{\mathbb{A}^l}$.

Finally, given an arbitrary modal automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$, we let

$$\text{tr}(\mathbb{A}) := \text{tr}_{\mathbb{A}}(a_I)$$

define the translation of the automaton \mathbb{A} itself. \triangleleft

The following lemma gives two useful representations of the translation map $\text{tr}_{\mathbb{A}}$ associated with an automaton structure \mathbb{A} . The point of the second result is that it displays each formula $\text{tr}_{\mathbb{A}}(a)$ as a fixpoint formula; this characterization will be of crucial importance in the next section. For its formulation we need to consider *restrictions* of linear automaton structures, and it is for this definition that we needed to introduce the notion of an automaton structure: initialized automata will not necessarily be closed under this operation, but automata structures are.

Definition 8.12 Let $\mathbb{A} = (A, \Theta, \Omega)$ be a linear automaton structure, and let $a \in \mathbb{A}$. The *a-restriction* of \mathbb{A} is the automaton structure $\mathbb{A} \downarrow a := (B, \Theta \upharpoonright_B, \Omega \upharpoonright_B)$ of which the carrier is given as $B := \{b \in A \mid b \sqsubseteq a\}$. \triangleleft

Proposition 8.13 *Let \mathbb{A} be any automaton structure and let $a \in A$. Then:*

$$\mathbf{tr}_{\mathbb{A}}(a) \equiv_K \Theta(a)[\mathbf{tr}_{\mathbb{A}}(b)/b \mid b \in A]. \quad (27)$$

If \mathbb{A} is linear, we have in addition

$$\mathbf{tr}_{\mathbb{A}}(a) \equiv_K \eta_a a. \Theta(a)[\mathbf{tr}_{(\mathbb{A} \downarrow a)^-}(b)/b \mid b \sqsubset a][\mathbf{tr}_{\mathbb{A}}(b)/b \mid a \sqsubset b] \quad (28)$$

Before moving on to prove this proposition, we quickly note that for a linear automaton structure \mathbb{A} , a is the maximal priority state of $\mathbb{A} \downarrow a$, so that we find

$$\mathbf{tr}_{\mathbb{A} \downarrow a}(a) = \eta_a a. \Theta(a)[\mathbf{tr}_{(\mathbb{A} \downarrow a)^-}(b)/b \mid b \sqsubset a]$$

by definition of $\mathbf{tr}_{(\mathbb{A} \downarrow a)^-}$. Hence, we may read (28) as stating that

$$\mathbf{tr}_{\mathbb{A}}(a) \equiv_K \mathbf{tr}_{\mathbb{A} \downarrow a}(a)[\mathbf{tr}_{\mathbb{A}}(b)/b \mid a \sqsubset b],$$

which may be of help to understand this characterization.

Proof. For the first part of the proposition, we reason by induction on the size of \mathbb{A} . By Proposition 8.10 we may without loss of generality assume that \mathbb{A} is linear. The case for automaton structures of size 1 is simple, so we focus on the case of a structure \mathbb{A} with $|\mathbb{A}| > 1$. Let m be the (by linearity unique) state that reaches the maximal priority of \mathbb{A} , that is, $\Omega(m) = \max \Omega[A]$. For this state m we obtain:

$$\begin{aligned} \mathbf{tr}_{\mathbb{A}}(m) &= \eta_m m. \Theta(m)[\mathbf{tr}_{\mathbb{A}^-}(b)/b \mid b \sqsubset m] && \text{(Definition } \mathbf{tr}_{\mathbb{A}}) \\ &\equiv_K \Theta(m)[\mathbf{tr}_{\mathbb{A}^-}(b)/b \mid b \sqsubset m][\mathbf{tr}_{\mathbb{A}}(m)/m] && \text{(fixpoint unfolding)} \\ &= \Theta(m)[\mathbf{tr}_{\mathbb{A}^-}(b)[\mathbf{tr}_{\mathbb{A}}(m)/m]/b \mid b \sqsubset m, \mathbf{tr}_{\mathbb{A}}(m)/m] && \text{(Fact 3.4)} \\ &= \Theta(m)[\mathbf{tr}_{\mathbb{A}}(b)/b \mid b \sqsubset m, \mathbf{tr}_{\mathbb{A}}(m)/m] && \text{(Definition } \mathbf{tr}_{\mathbb{A}}) \\ &= \Theta(m)[\mathbf{tr}_{\mathbb{A}}(a)/a \mid a \in A] && \text{(obvious)} \end{aligned}$$

For $a \neq m$, we have:

$$\begin{aligned} \mathbf{tr}_{\mathbb{A}}(a) &= \mathbf{tr}_{\mathbb{A}^-}(a)[\mathbf{tr}_{\mathbb{A}}(m)/m] && \text{(Definition } \mathbf{tr}_{\mathbb{A}}) \\ &\equiv_K \Theta(a)[\mathbf{tr}_{\mathbb{A}^-}(b)/b \mid b \sqsubset m][\mathbf{tr}_{\mathbb{A}}(m)/m] && \text{(inductive hypothesis)} \\ &= \Theta(a)[\mathbf{tr}_{\mathbb{A}^-}(b)[\mathbf{tr}_{\mathbb{A}}(m)/m]/b \mid b \sqsubset m, \mathbf{tr}_{\mathbb{A}}(m)/m] && \text{(Fact 3.4)} \\ &= \Theta(a)[\mathbf{tr}_{\mathbb{A}}(b)/b \mid b \sqsubset m, \mathbf{tr}_{\mathbb{A}}(m)/m] && \text{(Definition } \mathbf{tr}_{\mathbb{A}}) \\ &= \Theta(a)[\mathbf{tr}_{\mathbb{A}}(b)/b \mid b \in A] && \text{(obvious)} \end{aligned}$$

The second part of the proposition is also proved by induction on the size of the automaton structure, and again we only consider the inductive case of the argument. Supposing that the result holds for automaton structures smaller than \mathbb{A} , we prove the result for \mathbb{A} .

For the unique state m of maximal priority, the result is immediate from the definition since in this case $\mathbb{A} \downarrow m = \mathbb{A}$.

For a non-maximal state a , assuming that the induction hypothesis holds for states b with $b \sqsubset a$, we get:

$$\begin{aligned}
& \mathbf{tr}_{\mathbb{A}}(a) \\
& \equiv_K \mathbf{tr}_{\mathbb{A}^-}(a)[\mathbf{tr}_{\mathbb{A}}(m)/m] && \text{(Definition } \mathbf{tr}_{\mathbb{A}}) \\
& = \eta_a a. \Theta^-(a)[\mathbf{tr}_{(\mathbb{A} \downarrow a)}(b)/b \mid b \sqsubset a][\mathbf{tr}_{\mathbb{A}^-}(b)/b \mid a \sqsubset b \sqsubset m][\mathbf{tr}_{\mathbb{A}}(m)/m] && \text{(inductive hyp.)} \\
& = \eta_a a. \Theta^-(a)[\mathbf{tr}_{(\mathbb{A} \downarrow a)^-}(b)/b \mid b \sqsubset a][\mathbf{tr}_{\mathbb{A}^-}(b)/b \mid a \sqsubset b \sqsubset m][\mathbf{tr}_{\mathbb{A}}(m)/m] && ((\mathbb{A} \downarrow a)^- = \mathbb{A}^- \downarrow a) \\
& = \eta_a a. \Theta(a)[\mathbf{tr}_{(\mathbb{A} \downarrow a)^-}(b)/b \mid b \sqsubset a][\mathbf{tr}_{\mathbb{A}^-}(b)/b \mid a \sqsubset b \sqsubset m][\mathbf{tr}_{\mathbb{A}}(m)/m] && (\Theta(a) = \Theta^-(a)) \\
& = \eta_a a. \Theta(a)[\mathbf{tr}_{(\mathbb{A} \downarrow a)^-}(b)/b \mid b \sqsubset a][\mathbf{tr}_{\mathbb{A}}(b)/b \mid a \sqsubset b] && \text{(Fact 3.4, Def. } \mathbf{tr}_{\mathbb{A}})
\end{aligned}$$

as required. QED

The translation map interacts well with the operation on automata that we defined in section 4.3. As an auxiliary result we need the following observation, the proof of which we defer to the appendix.

Proposition 8.14 *Let \mathbb{A} be a modal automaton with x free and positive. Then we have:*

$$\mathbf{tr}(\mathbb{A}) \equiv_K (x \wedge \mathbf{tr}_{\mathbb{A}^x}((a_I)_0) \vee \mathbf{tr}_{\mathbb{A}^x}((a_I)_1)) \quad (29)$$

$$\mathbf{tr}(\mu x. \mathbb{A}) \equiv_K \mu x. \mathbf{tr}_{\mathbb{A}^x}((a_I)_1) \quad (30)$$

$$\mathbf{tr}(\nu x. \mathbb{A}) \equiv_K \nu x. (\mathbf{tr}_{\mathbb{A}^x}((a_I)_0) \vee \mathbf{tr}_{\mathbb{A}^x}((a_I)_1)). \quad (31)$$

Note that we can alternatively write Proposition 8.14((30)) as:

$$\mathbf{tr}(\mu x. \mathbb{A}) \equiv_K \mu x. \mathbf{tr}(\mathbb{A}^x)$$

since we chose $(a_I)_1$ as the start state of \mathbb{A}^x . As mentioned, the central result of this section is the following.

Proposition 8.15 *The following claims hold, for all modal automata \mathbb{A}, \mathbb{B} :*

- (1) $\mathbf{tr}(\mathbb{A} \wedge \mathbb{B}) \equiv_K \mathbf{tr}(\mathbb{A}) \wedge \mathbf{tr}(\mathbb{B})$ and $\mathbf{tr}(\mathbb{A} \vee \mathbb{B}) \equiv_K \mathbf{tr}(\mathbb{A}) \vee \mathbf{tr}(\mathbb{B})$;
- (2) $\mathbf{tr}(\neg \mathbb{A}) \equiv_K \neg \mathbf{tr}(\mathbb{A})$;
- (3) $\mathbf{tr}(\diamond \mathbb{A}) \equiv_K \diamond \mathbf{tr}(\mathbb{A})$ and $\mathbf{tr}(\square \mathbb{A}) \equiv_K \square \mathbf{tr}(\mathbb{A})$;
- (4) if \mathbb{A} is positive in p then $\mathbf{tr}(\eta p. \mathbb{A}) \equiv_K \eta p. \mathbf{tr}(\mathbb{A})$ for $\eta \in \{\mu, \nu\}$;
- (5) if \mathbb{A} is positive in p then $\mathbf{tr}(\mathbb{A}[\mathbb{B}/p]) \equiv_K \mathbf{tr}(\mathbb{A})[\mathbf{tr}(\mathbb{B})/p]$.

Proof. A full proof can be found in the appendix. We include only the proof for Clause (4) here, for which we will use Proposition 8.14. We first consider the case where $\eta = \mu$. We have:

$$\mathbf{tr}(\mu x. \mathbb{A}) \equiv_K \mu x. \mathbf{tr}_{\mathbb{A}^x}((a_I)_1) \quad \text{(Proposition 8.14(30))}$$

$$= \mu x. (x \wedge \mathbf{tr}_{\mathbb{A}^x}((a_I)_0)) \vee \mathbf{tr}_{\mathbb{A}^x}((a_I)_1) \quad \text{(Proposition 4.21)}$$

$$\equiv_K \mu x. \mathbf{tr}(\mathbb{A}) \quad \text{(Proposition 8.14(29))}$$

Next, for the case of $\eta = \nu$, we have:

$$\begin{aligned}
\mathbf{tr}(\nu x.\mathbb{A}) &\equiv_K \nu x.(\mathbf{tr}_{\mathbb{A}^x}((a_I)_0) \vee \mathbf{tr}_{\mathbb{A}^x}((a_I)_1)) && \text{(Proposition 8.14(31))} \\
&= \nu x.((x \wedge \mathbf{tr}_{\mathbb{A}^x}((a_I)_0)) \vee \mathbf{tr}_{\mathbb{A}^x}((a_I)_1)) && \text{(Proposition 4.21)} \\
&\equiv_K \nu x.\mathbf{tr}(\mathbb{A}) && \text{(Proposition 8.14(29))}
\end{aligned}$$

and the proof is done. QED

From this result, Theorem 2 follows easily.

Proof of Theorem 2. By induction on the complexity of a formula. For atomic formulas the result is easily checked, and for the inductive clauses we use the properties established in Proposition 8.15. For example, for a fixpoint formula $\mu x.\varphi(x)$, we have $\mathbb{A}_{\mu x.\varphi(x)} = \mu x.\mathbb{A}_{\varphi(x)}$ by definition, and we get

$$\mathbf{tr}(\mathbb{A}_{\mu x.\varphi(x)}) = \mathbf{tr}(\mu x.\mathbb{A}_{\varphi(x)}) \equiv_K \mu x.\mathbf{tr}(\mathbb{A}_{\varphi(x)}) \equiv_K \mu x.\varphi(x).$$

The other cases are similar. QED

We finish this section with a proposition, stating that one-step equivalent automata are in fact provably equivalent. The proof of this result, which we leave as an exercise to the reader, is conceptually simple, based on the facts that Kozen's axiomatization of the modal μ -calculus is an extension of the basic modal logic \mathbf{K} , from which it follows that equivalent one-step formulas, seen as formulas of basic modal logic, are in fact provably equivalent.

Proposition 8.16 *Let \mathbb{A} and \mathbb{B} be two modal automata. If $\mathbb{A} \equiv_1 \mathbb{B}$ then $\mathbb{A} \equiv_K \mathbb{B}$.*

This proposition will be used in the completeness proof, when we need to show that the closure properties mentioned in Proposition 6.16 in fact hold modulo *provable* equivalence. For instance, it follows from clause (4) of the mentioned proposition that the conjunction of two semi-disjunctive automata is *provably* equivalent to a semi-disjunctive automaton.

9 Kozen’s Lemma

The aim of this section is to show that if the formula associated with a modal automaton is consistent, then \exists has a winning strategy in the *thin* satisfiability game associated with the automaton. This result was formulated in the introduction of this paper as one of the main lemmas underlying our completeness proof.

Theorem 5 *\exists has a winning strategy in the thin satisfiability game for any consistent modal automaton \mathbb{A} .*

We will informally refer to this observation as “Kozen’s Lemma”, since it is an automata-theoretic version of Kozen’s partial completeness result for the aconjunctive fragment of the modal μ -calculus [36]. An analogous lemma also features prominently in Walukiewicz’ completeness proof [73, Theorem 31]. Observe that, in line with Kozen’s approach, as an immediate consequence of our Proposition and Corollary 6.15, we also obtain a partial completeness result, stating that every consistent semi-disjunctive automaton is satisfiable.

For the proof of Theorem 5, throughout this section we fix a modal automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$.

9.1 Intuitions

Before we turn to the technical details, we first provide some intuitions underlying our proof of Theorem 5. Assume that our automaton \mathbb{A} is consistent. Our goal will be to define a winning strategy for \exists in the thin satisfiability game for \mathbb{A} . We may and will assume that \exists ensures that at every position $R \in A^\sharp$, every next position $Q \in A^\sharp$ satisfies $\text{Dom}Q \subseteq \text{Ran}R$.

First of all, it is immediate by the definitions (Definition 8.1 and 8.11) and the Propositions 8.7 and 8.10 that without loss of generality we may assume \mathbb{A} to be linear. That means that the relation $\sqsubset = \sqsubset_{\mathbb{A}}$ linearly orders the states of \mathbb{A} , and in addition satisfies that $a \sqsubset b$ if $a \triangleleft_{\mathbb{A}} b$ but not $b \triangleleft_{\mathbb{A}} a$ (that is, if a is active in b but not vice versa).

\exists ’s winning strategy will be based on ensuring that a certain formula remains consistent throughout the match of the thin satisfiability game. This formula, which she will dynamically associate with the position under scrutiny, will encode certain information on the history of the match played so far. More in detail, given a partial match Σ , with current position $R = \text{last}(\Sigma) \in A^\sharp$, \exists associates with every state $a \in \text{Ran}R$ a ‘private’ formula $\text{tr}_\Sigma(a)$ that *tightens* the ‘public’ formula $\text{tr}_{\mathbb{A}}(a)$ in the sense that $\text{tr}_\Sigma(a) \leq_K \text{tr}_{\mathbb{A}}(a)$. \exists ’s strategy will then be geared towards keeping the formula

$$\psi_\Sigma := \bigwedge \{ \text{tr}_\Sigma(a) \mid a \in \text{Ran}(\text{last}\Sigma) \}$$

consistent throughout the match. As in Kozen’s approach, the key tool guaranteeing this strategy to be winning, is the *context rule* that we formulate as the following proposition. The proof of the proposition, which appears as Proposition 5.7(vi) in [36], will be given in the Appendix.

Proposition 9.1 *Suppose that $\gamma \wedge \mu x. \neg \gamma \wedge \varphi/x$ is consistent. Then so is $\gamma \wedge \varphi[\mu x. \neg \gamma \wedge \varphi/x]$.*

To see how this context rule can be used in \exists 's strategy in $\mathcal{S}_{thin}(\mathbb{A})$, consider a partial match Σ with $R = \text{last}(\Sigma)$ and inductively assume that the mentioned formula ψ_Σ is consistent. Suppose that τ is some trace on Σ , leading up to some μ -state $a \in \text{Ran}R$, so that it is one of \exists 's tasks to avoid a being unfolded infinitely often on any continuation of τ . The idea is now to think of the states in $\text{Ran}R \setminus \{a\}$ as providing the current *context* of a , and to ensure that there is no trace continuation from a leading to a future occurrence of a in the *same* context. If we can subsequently do this for all possible contexts of a (of which there are only finitely many because \mathbb{A} itself is finite), it follows that on any trace continuation from a , the state a will appear only finitely often.

To implement this idea, we encode the context of a as the formula

$$\gamma = \gamma_\Sigma := \bigwedge \{\text{tr}_\Sigma(b) \mid b \in \text{Ran}(\text{last}\Sigma), b \neq a\},$$

and at the same time ensure that the formula $\text{tr}_\Sigma(a)$ is a least fixpoint formula, that is, of the form $\mu a.\varphi$ (cf. (28) in Proposition 8.13). It then follows by the context rule of Proposition 9.1 that not only the formula $\gamma \wedge \mu a.\varphi$ is consistent, but also its tightening, $\gamma \wedge \varphi[\mu a.\neg\gamma \wedge \varphi/a]$. Suppose now that \exists tags the pair (τ, a) with the formula $\neg\gamma$, in such a way that, should a be visited again, in a partial match Σ' extending Σ with $\text{last}(\Sigma') = \text{last}(\Sigma) = R$, by a Σ' -trace τ' that is a continuation of τ , then \exists can guarantee that $\text{tr}_{\Sigma'}(a) \leq_K \neg\gamma$. Hence, if in such a situation we would have that $\text{tr}_{\Sigma'}(b) = \text{tr}_\Sigma(b)$ for all $b \in \text{Ran}R \setminus \{a\}$, we would find that

$$\begin{aligned} \psi_{\Sigma'} &= \bigwedge \{\text{tr}_{\Sigma'}(a) \mid a \in \text{Ran}(\text{last}\Sigma')\} \\ &= \text{tr}_{\Sigma'}(a) \wedge \bigwedge \{\text{tr}_\Sigma(b) \mid b \in \text{Ran}(\text{last}\Sigma), b \neq a\} \\ &\leq_K \neg\gamma \wedge \gamma \\ &\leq_K \perp. \end{aligned}$$

In this way \exists can guarantee that, provided she maintains the consistency of the formula ψ_Σ , in fact such a situation *cannot* occur; in other words, if on any trace from a in the current position she would encounter the state a again, it will be in a different context indeed.

In our automata-theoretic approach, the idea of *tightening* a formula can be realized neatly and simply by *decorating* the automaton \mathbb{A} .

Definition 9.2 A *decoration* of a linear modal \mathbf{X} -automaton $\mathbb{A} = (A, \Theta, \Omega, a_I)$ is a map $\delta : A \rightarrow \mu\text{ML}(\mathbf{X})$. Given such a decoration, by putting

$$\text{tr}_\mathbb{A}^\delta(a) := \eta_a a.\delta(a) \wedge \Theta(a)[\text{tr}_{(\mathbb{A}\downarrow a)^-}(b)/b \mid b \sqsubset a][\text{tr}_\mathbb{A}^\delta(b)/b \mid a \sqsubset b] \quad (32)$$

we define the *tightening* map $\text{tr}_\mathbb{A}^\delta : A \rightarrow \mu\text{ML}(\mathbf{X})$ associated with δ . \triangleleft

To obtain an understanding of this definition, it makes sense to compare it to the characterization of the translation map $\text{tr}_\mathbb{A}$ in (28):

$$\text{tr}_\mathbb{A}(a) \equiv_K \eta_a a.\Theta(a)[\text{tr}_{(\mathbb{A}\downarrow a)^-}(b)/b \mid b \sqsubset a][\text{tr}_\mathbb{A}(b)/b \mid a \sqsubset b],$$

and to note that the map $\text{tr}_\mathbb{A}^\delta$ is defined by a downward induction on the priority of states in \mathbb{A} . The definition of $\text{tr}_\mathbb{A}^\delta$ and its comparison with (28) also reveal that in case $\delta(a) = \top$

we are dealing with a vacuous tightening. As we will see, ν -states will always be vacuously tightened, in the sense that we have $\delta(a) = \top$ whenever $a \in A^\nu$.

For future reference we gather some simple facts on decorations in the following Proposition; the routine proofs are omitted. Here we extend the order \leq_K on μML to μML -valued maps (such as decorations), in the obvious, pointwise, manner.

Proposition 9.3 *Let δ, δ' be decorations of the linear modal automaton \mathbb{A} . Then we have*

- (1) $\text{tr}_{\mathbb{A}}^\delta \leq_K \text{tr}_{\mathbb{A}}$;
- (2) $\text{tr}_{\mathbb{A}}^\delta \leq_K \delta$;
- (3) if $\delta \leq_K \delta'$ then $\text{tr}_{\mathbb{A}}^\delta \leq_K \text{tr}_{\mathbb{A}}^{\delta'}$.

A key feature of our approach is that in a partial match Σ of the thin satisfiability game, we associate a decoration δ_ρ with each *trace* ρ on Σ . Since traces encode part of the history of the match Σ , this enables us to dynamically update the decoration by tightening it with relevant information about contexts. The ‘private’ formula $\text{tr}_\Sigma(a)$ that, as mentioned earlier on, we want to associate with a state a in the range of the relation $\text{last}(\Sigma)$, can now be defined by means of the decoration that we associate with a selected trace on Σ , the so-called *most significant trace* for a , notation: $\text{mst}_\Sigma(a)$. The map mst_Σ , associating a trace on Σ with each state in $\text{Ran}(\text{last}(\Sigma))$, is another dynamically defined entity maintained by \exists .

What complicates the proof is that \exists has to make sure that *all* traces associated with an infinite match are good (in the sense that the highest parity occurring infinitely often is even); thus, she needs to maintain a (separate) decoration for *each* trace. But since the context of one state a is encoded by the formulas associated with other states, it is nontrivial to guarantee that the formula $\bigwedge\{\text{tr}_\Sigma(b) \mid b \in \text{Ran}(\text{last}\Sigma), b \neq a\}$, representing the context of a state $a \in \text{Ran}(\text{last}\Sigma)$ in one partial match Σ , still represents the same context in an extension Σ' of Σ with $\text{Ran}(\text{last}\Sigma') = \text{Ran}(\text{last}(\Sigma))$. In order to let the number of context formulas not grow too large, we will not only update decorations by *tightening* them with negated context formulas, another operation that we will perform on decorations is a (partial) *reset*, i.e., we may set some of the values of the updated decoration to \top .

The exact definitions of the decoration associated with a trace on a partial match Σ will depend on a dynamically maintained linear order $<_\Sigma$ of all pairs consisting of a Σ -trace ρ and a state a of \mathbb{A} , the so-called *priority list*. (In fact the only traces that are relevant to us are the selected ones of the form $\text{mst}_\Sigma(a)$ for some state a belonging to the range of the last relational position of Σ , but our definitions are somewhat simpler if we take all traces into account.) Basically then, the updating of decorations proceeds as follows. Given an continuation $\Sigma \cdot Q$ of a partial match Σ of $\mathcal{S}_{\text{thin}}(\mathbb{A})$ of length k , and a trace $\rho \cdot a$ on $\Sigma \cdot Q$ continuing the Σ -trace ρ , inductively we assume that the decoration δ_ρ has been given; the decoration $\delta_{\rho \cdot a}(b)$ is defined as an update of δ_ρ . The value $\delta_{\rho \cdot a}(b)$ of a state b of the automaton under the updated decoration is determined by the value of $\delta_\rho(b)$, but also by the priority of the pair $(\rho \cdot a, b)$ relative to the ordering $<_{\Sigma \cdot Q}$. In particular, $\delta_{\rho \cdot a}(b)$ could be the tightening of δ_ρ with the negation of the current context formula, in case (ρ, b) was the *most significant item* of the priority list $<_\Sigma$. We also make the following adjustment: the decoration $\delta_{\rho \cdot a}(b)$ is reset to \top if the pair $(\rho \cdot a, b)$ has a low priority, while $\delta_{\rho \cdot a}(b)$ will keep the value of $\delta_\rho(b)$

if (ρ, b) has a high priority. As a rule of thumb, it is the most significant item that separates the items of high and low priority, respectively.

Before we turn to the technical details of the data structure that \exists will maintain throughout the play of any match of the thin satisfiability game, let us briefly comment on the role of *thinness* in the definition of \exists 's strategy. Similar to the proof of Proposition 6.6, the relative simplicity of the trace graph of any infinite $\mathcal{S}_{thin}(\mathbb{A})$ -match enables \exists to exercise a fairly tight control over the family of priority lists that we may associate with the initial, finite parts of Σ . Furthermore, we will also keep track of a *most significant trace* associated with each state in the range of the last relation of a partial match, and the thinness constraint will allow us to focus on a small number of infinite traces on any infinite match, with the property that each initial segment of each trace in this small collection is the most significant trace of its last element. The thinness constraint will be crucial here, since it guarantees that we can always find a bad trace in this small set of “continuous” traces, given that we can find a bad trace at all. In other words, we can safely focus on continuous traces without running the risk of not detecting the existence of a bad trace even if there is one.

Convention 9.4 Throughout this section we will restrict our attention to matches of the thin satisfiability game of the form $\Sigma = (R_i)_{i < \kappa}$ where $\text{Dom}R_{i+1} \subseteq \text{Ran}R_i$ for all i . Recall that by Remark 5.13 \exists always has a strategy that guarantees this.

9.2 Trace combinatorics

9.2.1 Finite traces

In order to assign the right decoration to each relevant trace on a partial match of $\mathcal{S}_{thin}(\mathbb{A})$, player \exists dynamically maintains an intricate data structure, associating with each partial match Σ the following entities:

- a set $V_\Sigma \subseteq \text{Tr}_\Sigma$ of *selected* traces on Σ ;
- a *most significant trace* map $\text{mst}_\Sigma : \text{Ran}(\text{last}\Sigma) \rightarrow V_\Sigma$;
- a total *relevance order* \prec_Σ on the set Tr_Σ ;
- a total *priority order* $<_\Sigma$ on the set $\text{Tr}_\Sigma \times A$.

Doing so, we will ensure that the following conditions are met throughout:

(*Shuffle-merge Condition:*) for every pair of states $a, b \in A$ and any trace τ on Σ , we have

$$(\tau, a) <_\Sigma (\tau, b) \text{ iff } a \sqsubset b.$$

(*Compatibility Condition:*) $\text{mst}_\Sigma : \text{Ran}(\text{last}\Sigma) \rightarrow V_\Sigma$ is a bijection, with $\text{mst}^{-1} = \text{last}$ and, for all $\tau \in \text{Tr}_\Sigma$:

$$\tau \preceq_\Sigma \text{mst}_\Sigma(\text{last}(\tau)).$$

Here are some first intuitions concerning this structure. To start with, the set V_Σ is a collection of selected Σ -traces. This set is in 1-1 correspondence with the collections of states in the range of the relation $\text{last}(\Sigma)$, with the map $\text{mst}_\Sigma : \text{Ran}(\text{last}\Sigma) \rightarrow V_\Sigma$ selecting a *most*

significant trace $\text{mst}_\Sigma(a) \in \text{Tr}_\Sigma$ for each state a . The compatibility condition on the map mst_Σ ensures that for each $a \in \text{Ran}(\text{last}\Sigma)$ there is a unique trace $\tau \in V_\Sigma$ such that $a = \text{last}(\tau)$, viz., the trace $\tau = \text{mst}_\Sigma(a)$.

The family of relevance relations \prec_Σ is used to guide the definition of the map mst : in cases where there are several candidates to pick the most significant trace associated with some state a , we may always choose the *most relevant* one, that is, the highest one according to the relevance order \prec_Σ . This explains the second part of the compatibility condition, stating that among all traces ending in some state a , the most significant trace of a is indeed the most relevant.

Finally, the main tool in the dynamic assignment of decorations to traces is the order $<_\Sigma$ on trace-state pairs, which can be thought of as arranging these pairs in a priority list for each partial match Σ . We make sure that this ranking is compatible with the priority $\sqsubset_{\mathbb{A}}$ induced by Ω , as expressed by the shuffle merge condition. We think of this list as a vertical ordering, with higher ranking items in the top and lower ranking items in the bottom.

Remark 9.5 Before moving on we note that we could have restricted the definition of \prec_Σ to the set V_Σ of *selected* Σ -traces, and the definition of $<_\Sigma$ to the set $V_\Sigma \times A$ of so-called Σ -*items*, respectively, since these are the only objects that we are interested in. We chose to consider the full set of traces instead because this makes the definitions somewhat simpler. \triangleleft

Before we can give the definitions of the actual structures $(V_\Sigma, \text{mst}_\Sigma, \prec_\Sigma, <_\Sigma)$, we need a few auxiliary notions.

Definition 9.6 Given a trace τ on a partial $\mathcal{S}_{\text{thin}}(\mathbb{A})$ -match Σ , let $C_\tau := C_{\text{last}(\tau)}$ denote the (*final*) *cluster* of τ .

Now let $\Sigma \cdot Q$ be a continuation of Σ with the thin relation Q , and suppose that Σ is of length k . Recall that if ρ is a trace on $\Sigma \cdot Q$ then $\rho|_k$ denotes the initial Σ -part of ρ , so that $\rho = \rho|_k \cdot \text{last}(\rho) = \rho|_k \cdot \rho(k+1)$. We say that a trace ρ on $\Sigma \cdot Q$ *stays in the same cluster* if $C_\rho = C_{\rho|_k}$, that ρ *enters a new cluster* if, on the contrary, its last state $\rho(k+1)$ belongs to a different cluster than the last state $\rho(k)$ of $\rho|_k$, and that ρ is *refreshed* if it either enters a new cluster or $\text{last}(\rho)$ is a safe state in its cluster.

Given a trace ρ , we define the *last refreshment date* of ρ , denoted $\text{lrd}(\rho)$, to be the smallest natural number k such that either $k = 0$, or else $k > 0$ and $\rho|_k$ is refreshed while $\rho|_j$ is refreshed for no later $j > k$. \triangleleft

Note that if $\text{lrd}(\rho) > 0$ then the trace $\rho|_{\text{lrd}(\rho)}$ is indeed refreshed. On the other hand, if $\text{lrd}(\rho) = 0$ then ρ is the unique trace that has never been refreshed, that is, $\rho|_k$ is not refreshed for any k . From this it is easily seen that if $\text{lrd}(\sigma) = \text{lrd}(\sigma')$ for *distinct* traces σ, σ' , then this common last refreshment date must be *bigger* than 0.

We can now define the data structure associated with a partial match Σ , starting with the relevance order \prec_Σ .

Definition 9.7 Let Σ be a partial match of the thin satisfiability game for \mathbb{A} . We define the relation $\prec_\Sigma \subseteq \text{Tr}_\Sigma \times \text{Tr}_\Sigma$ by the following case distinction.

$$\sigma \prec_\Sigma \sigma' \text{ iff } \begin{cases} \text{either} & \text{lrd}(\sigma) > \text{lrd}(\sigma') \\ \text{or} & \text{lrd}(\sigma) = \text{lrd}(\sigma') \text{ and } \sigma|_k \prec_{\Sigma|_k} \sigma'|_k \\ \text{or} & \text{lrd}(\sigma) = \text{lrd}(\sigma') \text{ and } \sigma|_k = \sigma'|_k \text{ and } \sigma(k+1) \sqsubset \sigma'(k+1), \end{cases}$$

where in the last two cases we let k be such that $k+1 = \text{lrd}(\sigma) = \text{lrd}(\sigma')$. \triangleleft

The choice to rely on the priority order \sqsubset to arbitrate between σ and σ' in the last case in this definition is not essential, any linear ordering of the states in A would have done just as well. But since we have assumed that \mathbb{A} is linear, it seems natural to use the one linear order of A that we already have in place.

We are now ready to define, for each partial match Σ , the collection V_Σ of selected traces and the map mst associated with Σ . In the latter definition we use the fact that the relation \prec_Σ is a strict total order, see Proposition 9.10(1).

Definition 9.8 For any given partial match Σ and $a \in \text{Ran}(\text{last}\Sigma)$, set $\text{mst}_\Sigma(a)$ to be the greatest trace σ with $\text{last}(\sigma) = a$ according to the relevance order \prec_Σ , which clearly exists since the collection Tr_Σ is always finite and \prec_Σ is a strict total order. With this definition in place, we simply set

$$V_\Sigma := \{\text{mst}_\Sigma(b) \mid b \in \text{Ran}(\text{last}\Sigma)\}$$

to be the range of the map mst_Σ . \triangleleft

Finishing the definition of the data structure associated with a partial match Σ , we have now arrived at the most fundamental relation, viz., the priority order $<_\Sigma$; its definition is given by induction on the length of the partial match Σ .

Definition 9.9 For Σ being the unique initial match consisting of the single position $R = \{(a_I, a_I)\}$, there is only a single trace τ to consider, so the order $<_\Sigma$ is simply set to agree with the order \sqsubset over A . In other words, set $(\tau, a) <_\Sigma (\tau, b)$ iff $a \sqsubset b$.

Now suppose that $<_\Sigma$ has been defined for some match Σ of length k , and let $Q \subseteq A \times A$ be a thin relation. We define $<_{\Sigma.Q}$ in a series of two steps.

(Step 1) First we define a new order $<_\Sigma^0$ on the set $\text{Tr}_{\Sigma.Q} \times A$. Basically, $<_{\Sigma.Q}^0$ is a natural “continuation” of the order $<_\Sigma$, with the proviso that refreshed traces will be moved to the bottom of the list. Formally, we put:

- (a) $(\sigma, b) <_{\Sigma.Q}^0 (\sigma', b')$,
in case σ is refreshed and σ' is not;
- (b) $(\sigma, b) <_{\Sigma.Q}^0 (\sigma', b')$ iff either $(\sigma|_k, b) <_\Sigma (\sigma'|_k, b')$ or $\sigma|_k = \sigma'|_k$ and $\text{last}(\sigma) \sqsubset \text{last}(\sigma')$,
in case $\sigma \neq \sigma'$ and σ and σ' are either both refreshed or both not refreshed.
- (c) $(\sigma, b) <_{\Sigma.Q}^0 (\sigma, b')$ iff $b \sqsubset b'$.

(Step 2) Second, we move all erasable $(\sigma, b) \in \text{Tr}_{\Sigma, Q} \times A$ to the bottom of the list, where a pair (σ, b) is *erasable* if either $b \sqsubset \text{last}(\sigma)$, or $b = \text{last}(\sigma) \in A^\nu$. More precisely, we define the order $<_{\Sigma, Q}$ by setting:

- (a) $(\sigma, b) <_{\Sigma, Q} (\sigma', b')$,
in case (σ, b) is erasable and (σ', b') is not;
- (b) $(\sigma, b) <_{\Sigma, Q} (\sigma', b')$ iff $(\sigma, b) <_{\Sigma, Q}^0 (\sigma', b')$,
in case (σ, b) and (σ', b') are either both erasable or both not erasable.

◁

It is not very hard to verify that with these conditions, for each partial match Σ , the data structure $(V_\Sigma, \text{mst}_\Sigma, \prec_\Sigma, <_\Sigma)$ is well defined and satisfies the required conditions.

Proposition 9.10 *The following hold for any partial match $\Sigma = (R_i)_{i \leq k}$ of $\mathcal{S}_{\text{thin}}(\mathbb{A})$:*

- (1) *the relation \prec_Σ is a strict total order on Tr_Σ ;*
- (2) *the map mst_Σ is a bijection from $\text{Ran}R_k$ to V_Σ satisfying the compatibility condition;*
- (3) *the relation $<_\Sigma$ is a strict total order on $\text{Tr}_\Sigma \times A$ satisfying the shuffle merge condition.*

Proof. For part (1), the only non-trivial clause to prove is totality, and for this one proceeds by induction on the length of the match Σ . The case for the unique match consisting of the relation $\{(a_I, a_I)\}$ is trivial, so we focus on the induction step: let Σ be a match of a given length for which the induction hypothesis holds for all shorter matches, and let σ, τ be traces on Σ . Suppose that we have neither $\sigma \prec_\Sigma \tau$ nor $\tau \prec_\Sigma \sigma$. If $\text{lrd}(\sigma) = \text{lrd}(\tau) = 0$ then we must have $\tau = \sigma$, and the only other possibility is that $\text{lrd}(\sigma) = \text{lrd}(\tau) = k + 1$ for some $k \in \omega$. So we must have neither $\tau|_k \prec_{\Sigma|_k} \sigma|_k$ nor $\sigma|_k \prec_{\Sigma|_k} \tau|_k$. By the induction hypothesis we have $\tau|_k = \sigma|_k$, and we now see that we can have neither $\sigma(k+1) \sqsubset \tau(k+1)$ nor $\tau(k+1) \sqsubset \sigma(k+1)$. By linearity of \mathbb{A} we get $\sigma(k+1) = \tau(k+1)$, and since neither $\tau|_j$ nor $\sigma|_j$ are refreshed for any $j > k + 1$ it now follows by thinness of all relations in Σ that $\sigma = \tau$.

For part (2), the compatibility condition holds since by definition the most significant trace of $\text{last}(\tau)$ is maximal among the traces ending with $\text{last}(\tau)$, hence $\tau \preceq_\Sigma \text{mst}_\Sigma(\text{last}(\tau))$. Also note that the map mst_Σ is surjective by definition of V_Σ , and injective for the trivial reason that two distinct states cannot both be the last element of the same trace.

Finally, part (3) can be proved by a straightforward induction on the length of Σ . QED

Next we establish some basic properties of the relevance order and most significant trace.

Proposition 9.11 *Let Σ be any partial match of the thin satisfiability game, Q any thin relation and σ, τ any two traces on Σ with the same last element a . If $\sigma \prec_\Sigma \tau$ and b is any state with $(a, b) \in Q$, then $\sigma \cdot b \prec_{\Sigma, Q} \tau \cdot b$.*

Proof. Since the last two elements ab of the traces $\sigma \cdot b$ and $\tau \cdot b$ are the same, these traces are either both refreshed or both not refreshed. In the former case, the last refreshment date of both $\sigma \cdot b$ and $\tau \cdot b$ is $k + 1$, where k is the length of the match Σ . Furthermore, we have

$(\sigma \cdot b)|_k = \sigma$ and $(\tau \cdot b)|_k = \tau$, so since $\sigma \prec_\Sigma \tau$ we get $\sigma \cdot b \prec_{\Sigma.Q} \tau \cdot b$ from the definition of the relevance order.

In the latter case, we distinguish cases:

(Case 1:) $\text{lrd}(\tau) < \text{lrd}(\sigma)$. Then since the traces $\sigma \cdot b$ and $\tau \cdot b$ are both not refreshed we have $\text{lrd}(\tau \cdot b) = \text{lrd}(\tau) < \text{lrd}(\sigma) = \text{lrd}(\sigma \cdot b)$, and thus we get $\sigma \cdot b \prec_{\Sigma.Q} \tau \cdot b$ as required.

(Case 2:) $\text{lrd}(\tau) = \text{lrd}(\sigma) = k + 1$ for some $k \in \omega$, and either $\sigma|_k \prec_{\Sigma|_k} \tau|_k$ or $\sigma|_k = \tau|_k$ and $\sigma(k+1) \sqsubset \tau(k+1)$. (We cannot have $\text{lrd}(\tau) = \text{lrd}(\sigma) = 0$ since it would follow that $\tau = \sigma$.) In both these cases it follows that $\sigma \cdot b \prec_{\Sigma.Q} \tau \cdot b$, since $\text{lrd}(\tau \cdot b) = \text{lrd}(\tau) = \text{lrd}(\sigma) = \text{lrd}(\sigma \cdot b)$, $\tau \cdot b|_k = \tau|_k$ and $\sigma \cdot b|_k = \sigma|_k$, $\tau \cdot b(k+1) = \tau(k+1)$ and $\sigma \cdot b(k+1) = \sigma(k+1)$. This concludes the proof. QED

The following almost immediate consequence of the previous proposition expresses a downward coherence property of selected traces and the map mst .

Proposition 9.12 *Let Σ be any partial match of the thin satisfiability game. Then for any selected trace $\tau \in V_\Sigma$ and any k smaller than the length of Σ , the trace $\tau|_k$ is a selected trace of $\Sigma|_k$, and in particular, it holds that $\text{mst}(\text{last}(\tau|_k)) = \tau|_k$.*

Proof. Suppose for a contradiction that there is some trace σ on $\Sigma|_k$ with $\text{last}(\sigma) = \text{last}(\tau|_k)$ and $\tau|_k \prec_{\Sigma|_k} \sigma$. It follows from Proposition 9.11 that $\tau|_k \cdot \tau(k+1) \prec_{\Sigma|_{k+1}} \sigma \cdot \tau(k+1)$, and since these traces also have the same last elements we can repeat the same argument for these two traces to find that $\tau|_k \cdot \tau(k+1) \cdot \tau(k+2) \prec_{\Sigma|_{k+2}} \sigma \cdot \tau(k+1) \cdot \tau(k+2)$... Continuing in this way we find that $\tau \prec_\Sigma \sigma \cdot \rho$, where ρ is the segment of τ starting with $\tau(k+1)$. But this contradicts our assumption that τ is the most significant trace of its last element, so we are done. QED

The corresponding *upward* coherence condition does not hold: due to the occurrence of *trace merging* it is not always the case that $\sigma \cdot a \in V_{\Sigma.Q}$ whenever $\sigma \in V_\Sigma$ and $(\text{last}(\sigma), a) \in Q$. In case we have $\sigma \cdot a \neq \text{mst}_{\Sigma.Q}(a)$ (and thus $\sigma \cdot a \notin V_{\Sigma.Q}$), we say that a *trace jump* occurs.

9.2.2 Infinite traces

The data structure $(V_\Sigma, \text{mst}_\Sigma, \prec_\Sigma, <_\Sigma)$ and the procedure for updating it provides a combinatorial device that allows us to exercise some control over the collection of bad traces on an infinite match in the thin satisfiability game. To see how this works, we turn to infinite matches, but first we adopt a notational convention.

Convention 9.13 Let Σ be a (full or partial) match of the thin satisfiability game for \mathbb{A} . When no confusion is likely to arise, we will frequently use simplified notation, writing V_k rather than $V_{\Sigma|_k}$, mst_k rather than $\text{mst}_{\Sigma|_k}$, etc.

Definition 9.14 Let Σ be an infinite match of $\mathcal{S}_{\text{thin}}(\mathbb{A})$. We let V_Σ denote the set of *continuous* traces on Σ , that is, infinite traces τ such that $\tau|_k \in V_k$ for all $k \in \omega$. ◁

Continuous traces have various nice properties; in particular, any continuous τ satisfies $\tau|_k = \text{mst}_k(\tau(k))$ and thus $\tau|_k \cdot \tau(k+1) = \text{mst}_{k+1}(\tau|_{k+1})$, for all $k \in \omega$. That is, continuous traces have *no trace jumps* — the property explaining the name. What make the continuous traces also nice to work with is the fact that there are only finitely many of them.

Proposition 9.15 *For any infinite match Σ of the thin satisfiability game for \mathbb{A} the set V_Σ of continuous traces over Σ satisfies $|V_\Sigma| \leq |A|$.*

Proof. Suppose for contradiction that the set V_Σ would contain $n+1$ distinct traces $\sigma_0, \dots, \sigma_n$, where $n := |A|$ is the number of states of \mathbb{A} . Let $k \in \omega$ be such that all $\sigma_i|_k$ for $0 \leq i \leq n$ are distinct.

Observe that the set $\{\text{last}(\sigma_i|_k) \mid 0 \leq i \leq n\}$ is a subset of $\text{Ran}(\text{last}_{\Sigma|_k}) \subseteq A$, and so there must be indices $i \neq j$ such that $\text{last}(\sigma_i|_k) = \text{last}(\sigma_j|_k)$. But then it follows by the compatibility condition that $\sigma_i|_k = \text{mst}_k(\text{last}(\sigma_i|_k)) = \text{mst}_k(\text{last}(\sigma_j|_k)) = \sigma_j|_k$, which provides the desired contradiction. QED

For the above-mentioned reasons it will be convenient for us to restrict attention to continuous traces as much as possible, and here the following observation (Proposition 9.17), stating that every (infinite) bad trace is eventually equal to a continuous trace, will be immensely useful. The key property of bad traces that allow us to prove this is the following.

Proposition 9.16 *Let Σ be an infinite match of the thin satisfiability game and let τ be a bad trace on Σ . Then there are at most finitely many $k \in \omega$ such that $\tau|_k$ is refreshed.*

Proof. There are two ways that a trace may be refreshed: by entering into a new cluster, or by visiting a safe state of some cluster. The first of these cases can generally only occur finitely many times on any infinite trace, and on a bad trace it is clear that the second case can also only occur finitely many times. QED

Recall that two A -streams σ and τ are *eventually equal*, notation: $\sigma =_\infty \tau$, if there is a $k \in \omega$ such that $\sigma(j) = \tau(j)$ for all $j \geq k$. In particular, if two traces are eventually equal, then they will be either both good or both bad. The next proposition ensures that if an infinite $\mathcal{S}_{\text{thin}}(\mathbb{A})$ -match carries a bad trace, then it also carries a bad trace that is in addition continuous.

Proposition 9.17 *Let Σ be an infinite match of the thin satisfiability game for \mathbb{A} . Then for every bad trace $\tau \in \text{Tr}_\Sigma$ there is a continuous trace $\hat{\tau}$ that is eventually equal to τ and, hence, bad as well.*

Proof. Fix a bad trace τ . Say that $k+1 \in \omega$ is a *discontinuous point* of τ if the stages $k, k+1$ constitute a trace jump with respect to the most significant traces associated with the corresponding entries of τ , that is, if $\text{mst}(\tau(k)) \cdot \tau(k+1) \neq \text{mst}(\tau(k+1))$.

It suffices to prove that the bad trace τ has only finitely many discontinuous points.

CLAIM 1 If τ has only finitely many discontinuous points, then there is a continuous trace $\hat{\tau} =_\infty \tau$.

PROOF OF CLAIM Let $k \in \omega$ be such that no $j \geq k$ is a discontinuous point of τ , let σ be the part of τ following $\tau(k)$, and let $\rho = \text{mst}_k(\tau(k))$. Then

$$\hat{\tau} := \rho \cdot \sigma$$

is a continuous trace that is eventually equal to τ . Since it is obvious that $\hat{\tau}$ is eventually equal to τ , it suffices to show that

$$\hat{\tau}|_j = \text{mst}(\hat{\tau}(j)) \tag{33}$$

for all $j \in \omega$. First, by a simple induction we prove (33) for $j \geq k$. For $j = k$ we have $\text{mst}_k(\tau(k)) = \hat{\tau}|_k$ by definition, and assuming that the induction hypothesis holds for j we have:

$$\begin{aligned} \hat{\tau}|_{j+1} &= \hat{\tau}|_j \cdot \hat{\tau}(j+1) && \text{(obvious)} \\ &= \text{mst}_j(\hat{\tau}(j)) \cdot \hat{\tau}(j+1) && \text{(induction hypothesis)} \\ &= \text{mst}_j(\tau(j)) \cdot \tau(j+1) && \text{(definition } \hat{\tau} \text{)} \\ &= \text{mst}_{j+1}(\tau(j+1)) && (j+1 \text{ not discontinuous)} \\ &= \text{mst}_{j+1}(\hat{\tau}(j+1)) && \text{(definition } \hat{\tau} \text{)} \end{aligned}$$

Second, for $j < k$ we have $\hat{\tau}|_j = (\hat{\tau}|_k)|_j = \rho|_j$, where $\rho = \text{mst}_j(\tau(k))$, and so by Proposition 9.12 we obtain from $\rho \in V_k$ that $\hat{\tau}|_j \in V_j$, meaning that $\hat{\tau}|_j = \text{mst}_j(\hat{\tau}(j))$ as required. \blacktriangleleft

We now turn to prove the main claim, that there are only finitely many discontinuous points for τ . Since τ is a bad trace it is only refreshed finitely many times by Proposition 9.16, so pick $k_0 \in \omega$ for which $\tau|_j$ is not refreshed for any $j \geq k_0$. For our first step of the proof, we define a function $f : \omega \rightarrow \omega$ by setting

$$f(i) := \text{lrd}(\text{mst}_i(\tau(i))).$$

CLAIM 2 The function f is antitone above k_0 , that is: $n \leq m$ implies $f(m) \leq f(n)$ whenever $n, m \geq k_0$.

PROOF OF CLAIM It suffices to prove that $f(n+1) \leq f(n)$ for all $n \geq k_0$. So pick $n \geq k_0$. We know that $\tau|_{n+1}$ is not refreshed, and it clearly follows that the trace $\text{mst}_n(\tau(n)) \cdot \tau(n+1)$ is not refreshed either since it has the same last and next-to-last entries as $\tau|_{n+1}$. This means that:

$$\text{lrd}(\text{mst}_n(\tau(n)) \cdot \tau(n+1)) = \text{lrd}(\text{mst}_n(\tau(n))) = f(n).$$

Hence, if $f(n) < f(n+1)$, then we get

$$\text{lrd}(\text{mst}_n(\tau(n)) \cdot \tau(n+1)) < \text{lrd}(\text{mst}_{n+1}(\tau(n+1))),$$

and it immediately follows that

$$\text{mst}_{n+1}(\tau(n+1)) \prec_{\Sigma|_{n+1}} \text{mst}_n(\tau(n)) \cdot \tau(n+1).$$

But this directly contradicts the compatibility condition for the most significant trace. \blacktriangleleft

From Claim 2 it follows that there is a $k_1 \geq k_0$ such that $f(k_1) = f(j)$ for all $j \geq k_1$. We assume that $f(k_1) > 0$ since the other case is easier, and we let q denote the predecessor of $f(k_1)$ so that $f(k_1) = q + 1$. We define a new function $g : \omega \rightarrow \text{Tr}_q \times A$, setting:

$$g(n) := (\text{mst}_n(\tau(n))|_q, \text{mst}_n(\tau(n))(q + 1))$$

for $n > k_1$ – we can set $g(n)$ to be any fixed arbitrary pair for $n \leq k_1$. Note that $k_1 \geq q + 1$, since the function f clearly satisfies $f(m) \leq m$ for all $m \in \omega$. Intuitively, the value $g(n)$ of the map g at n records two pieces of information about the trace $\text{mst}_n(\tau(n))$ that will determine its place in the relevance order among traces in Tr_n : the place of the restricted trace $\text{mst}_n(\tau(n))|_q$ in the relevance order at that stage, which is the last stage before the last refreshment date of $\text{mst}_n(\tau(n))$, and the state visited by $\text{mst}_n(\tau(n))$ at its last refreshment date. Since we already know what the last refreshment date of $\text{mst}_n(\tau(n))$ is (namely $q + 1$), these two pieces of information indeed suffice to determine the place of $\text{mst}_n(\tau(n))$ in the relevance order.

We order the elements of $\text{Tr}_q \times A$ lexicographically with respect to the relevance order and the priority order. More precisely put, we define the strict total order $\prec_q| \sqsubset$ on $\text{Tr}_q \times A$ by setting $(\sigma, b) \prec_q| \sqsubset (\sigma', b')$ iff $\sigma \prec_q \sigma'$ or $\sigma = \sigma'$ and $b \sqsubset b'$.

CLAIM 3 For all $n > k_1$, we have $g(n) \prec_q| \sqsubset g(n + 1)$ or $g(n) = g(n + 1)$. Furthermore, if $n + 1$ is a discontinuous point, then in fact $g(n) \prec_q| \sqsubset g(n + 1)$.

PROOF OF CLAIM It suffices to prove the second part of the claim, since if $n + 1$ is not a discontinuous point then $\text{mst}_{n+1}(\tau(n + 1)) = \text{mst}_n(\tau(n)) \cdot \tau(n + 1)$, and if $n > k_1 \geq q$ it follows that

$$(\text{mst}_{n+1}(\tau(n + 1)))|_q = (\text{mst}_n(\tau(n)))|_q$$

and

$$(\text{mst}_{n+1}(\tau(n + 1)))(q + 1) = (\text{mst}_n(\tau(n)))(q + 1)$$

and hence $g(n) = g(n + 1)$.

So let n be a discontinuous point. Then by the compatibility condition for most significant traces we have:

$$\text{mst}_n(\tau(n)) \cdot \tau(n + 1) \prec_{n+1} \text{mst}_{n+1}(\tau(n + 1)). \quad (34)$$

But the trace $\text{mst}_n(\tau(n)) \cdot \tau(n + 1)$ is not refreshed since $\tau|_{n+1}$ is not refreshed, so

$$\text{lrd}(\text{mst}_n(\tau(n)) \cdot \tau(n + 1)) = f(n) = q + 1,$$

and since $f(n + 1) = q + 1$ we have $\text{lrd}(\text{mst}_{n+1}(\tau(n + 1))) = q + 1$ by definition of the map f . Hence from the definition of the relevance order there are two possibilities for (34). Either we have

$$(\text{mst}_n(\tau(n)) \cdot \tau(n + 1))|_q \prec_q (\text{mst}_{n+1}(\tau(n + 1)))|_q,$$

or else $(\text{mst}_n(\tau(n)) \cdot \tau(n + 1))|_q = (\text{mst}_{n+1}(\tau(n + 1)))|_q$ and:

$$(\text{mst}_n(\tau(n)) \cdot \tau(n + 1))(q + 1) \sqsubset (\text{mst}_{n+1}(\tau(n + 1)))(q + 1).$$

But since $n > q$ we have:

$$(\mathbf{mst}_n(\tau(n)) \cdot \tau(n+1))|_q = (\mathbf{mst}_n(\tau(n)))|_q$$

and

$$(\mathbf{mst}_n(\tau(n)) \cdot \tau(n+1))(q+1) = (\mathbf{mst}_n(\tau(n)))(q+1)$$

and so in each of the two cases we get $g(n) \prec_q | \sqsubset g(n+1)$ as required. \blacktriangleleft

It now follows that there can be only finitely many discontinuous points of τ above k_1 , simply because the set $\text{Tr}_q \times A$ is finite and hence the map g can only strictly increase finitely many times with respect to the strict total order $\prec_q | \sqsubset$, as it is never decreasing with respect to this order.

From this and Claim 1 the Proposition is immediate. \square

Motivated by Proposition 9.17 we will from now on focus on continuous traces and on selected items.

Definition 9.18 Let Σ be a partial match; a Σ -item is defined as a pair $(\tau, a) \in \text{Tr}_\Sigma \times A$ that is *selected* in the sense that $\tau \in V_\Sigma$. A Σ -item is called *in focus* if $\tau = \mathbf{mst}_\Sigma(a)$, a μ -item if $\Omega(a)$ is odd, and a ν -item otherwise. The *most significant item* (MSI) of Σ , denoted \mathbf{msi}_Σ , is defined to be the highest ranking Σ -item (τ, a) in the priority order $<_\Sigma$ which is in focus at Σ . \triangleleft

Note that the MSI of Σ must be of the form $(\mathbf{mst}_\Sigma(a), a)$ for some $a \in \text{Ran}(\text{last}(\Sigma))$.

Definition 9.19 Let Σ be an infinite match of the thin satisfiability game for \mathbb{A} . A Σ -item is nothing but a pair consisting of an infinite continuous trace on Σ , together with a state of \mathbb{A} ; the notions of μ -item and ν -item apply as before.

We say that such a Σ -item (τ, a) *stabilizes* at the index $k \in \omega$ if there is a finite set $S_\Sigma = \{(\sigma_1, d_1), \dots, (\sigma_m, d_m)\}$ of Σ -items such that, for all $j \geq k$:

- (1) the MSI of $\Sigma|_j$ is equal to or smaller than $(\tau|_j, a)$ in the priority order $<_{\Sigma|_j}$;
- (2) the $\Sigma|_j$ -items above $(\tau|_j, a)$ are precisely $(\sigma_1|_j, d_1), \dots, (\sigma_m|_j, d_m)$ (in some fixed order).

\triangleleft

The framework of orderings $<_\Sigma$ associated with partial matches Σ is designed to make the following proposition true.

Proposition 9.20 *Let Σ be an infinite match of the thin satisfiability game for \mathbb{A} . If Σ has a bad trace, then there is a continuous bad trace τ on Σ , such that, with $a \in A$ denoting the highest priority state appearing infinitely often on τ , the following hold:*

- (1) the pair (τ, a) stabilizes at some $k < \omega$;
- (2) $(\tau|_k, a)$ is the MSI of $\Sigma|_k$, for infinitely many $k < \omega$.

Proof. Fix a \mathcal{S}_{thin} -match Σ , and recall that for $j \in \omega$ we will abbreviate $V_{\Sigma|_j}$ as V_j , etc. Let F denote the set of all pairs $(\rho, b) \in V_{\Sigma} \times A^{\mu}$ such that $\rho \in V_{\Sigma}$ is a continuous bad trace and $b \in A^{\mu}$ is the highest priority state visited infinitely many times on ρ . If Σ has a bad trace then F is non-empty by Proposition 9.17, and by Proposition 9.15 F is finite.

We first show that the priority ordering among members of F will eventually stabilize, in the following sense.

CLAIM 1 For any pair $(\rho, b), (\rho', b')$ of items in F there is a $k < \omega$ such that for all $j \geq k$:

$$(\rho|_j, b) <_j (\rho'|_j, b') \text{ iff } (\rho|_k, b) <_k (\rho'|_k, b'). \quad (35)$$

PROOF OF CLAIM Fix two Σ -items $(\rho, b), (\rho', b')$ in F . If $\rho = \rho'$, then by the shuffle merge condition (35) holds in fact for all $j \in \omega$, so we may focus on the case where ρ and ρ' are *distinct* traces. Let $k_0 \in \omega$ be such that $\rho|_{k_0} \neq \rho'|_{k_0}$.

Now suppose that for some stage $k \geq k_0$ we have

$$(\rho|_{k+1}, b) <_{k+1} (\rho'|_{k+1}, b') \text{ but } (\rho|_k, b) >_k (\rho'|_k, b'). \quad (36)$$

It follows from the construction of $<_{k+1}$ out of $<_k$ that there are only two possibilities for this swap to happen: (i) if the items were swapped in step 1, then $\rho|_{k+1}$ must be refreshed, (ii) if the swap took place in step 2, then we must have $b \sqsubset \rho(k)$, since (ρ_{k+1}, b) is a μ -item.

It is not hard to see, however, that each of the two mentioned possibilities can only occur for finitely many $k \geq k_0$. In the case of (i), note that ρ , just like any trace, can enter a new cluster only finitely often, and that $\rho(k)$ can be a safe state of its cluster for finitely many k only – otherwise ρ would not be a bad trace. Similarly, situation (ii) can apply to finitely many k only, since b is by assumption the greatest priority state that ρ visits infinitely often.

Hence there are only finitely many indices k satisfying (36). From this the claim is immediate. \blacktriangleleft

On the basis of Claim 1 we can and will speak unambiguously of “the priority ordering over F relative to Σ ”, and this order is a strict total order just like the priority orders associated with partial matches. Hence, by finiteness of $F \subseteq V_{\Sigma} \times A$ it is immediate that F must have a greatest element (τ, a) under the priority order $<_{\Sigma}$.

We now consider the elements above the item (τ, a) in the priority orders.

CLAIM 2 There is an index k and a finite set S_{Σ} of Σ -items such that for every $j \geq k$, every element above $(\tau|_j, a)$ in the priority list $<_j$ is of the form $(\sigma|_j, d)$ for some pair $(\sigma, d) \in S_{\Sigma}$.

PROOF OF CLAIM Let k_0 be a point at which the trace τ has stopped visiting states with higher priority than a , and has arrived in its final cluster at least one stage ago. Then from this moment on τ is not refreshed and (τ, a) is not erasable. That is, there is no $j \geq k_0$ such that the trace $\tau|_j$ is refreshed or the item $(\tau|_j, a)$ is erasable; the latter holds because $\rho(j) \sqsubseteq a$ and a is a μ -state.

For each $j \geq k_0$, let $S_j \subseteq V_j \times A$ be the set of $\Sigma|_j$ -items (σ, d) of higher priority than $(\tau|_j, a)$ in the order $<_j$. We first show that every element of S_{j+1} is of the form $(\sigma \cdot b, d)$ for some

$(\sigma, d) \in S_j$. To see this, let (σ, d) be an arbitrary element of S_{j+1} . Since the item $(\tau|_{j+1}, a)$ is not erasable, we can only have $(\tau|_{j+1}, a) <_{j+1} (\sigma, d)$ if already $(\tau|_{j+1}, a) <_{j+1}^0 (\sigma, d)$. And since $\tau|_{j+1}$ is not refreshed, this can only be the case if already $(\tau|_j, a) <_j (\sigma|_j, d)$. So the item (σ, d) is of the form $(\sigma|_j \cdot \text{last}(\sigma), d)$ with $(\tau|_j, a) <_{\Sigma|_j} (\sigma|_j, d)$. Furthermore by Proposition 9.12 we have that $\sigma|_j = \text{mst}(\sigma(j+1))|_j = \text{mst}(\sigma(j)) \in V_j$ and we get $(\sigma|_j, d) \in S_j$ as required.

At the same time, each $\Sigma|_j$ -item $(\sigma, d) \in S_j$ has at most one Σ_{j+1} -continuation of the form $(\sigma \cdot b, d) \in S_{j+1}$. To see this, suppose that σ as a trace has two Σ_{j+1} -continuations $\sigma \cdot b_0$ and $\sigma \cdot b_1$; it then suffices to show that at least one item $(\sigma \cdot b_i, d)$ does not belong to S_{j+1} . But it follows by thinness of the relation Q (defined by $\Sigma|_{j+1} = \Sigma|_j \cdot Q$), that one of the two states, say, b_i , must be a safe state of its cluster. Then the $\Sigma|_{j+1}$ -trace $\sigma \cdot b_i$ is refreshed, so that in step 1 of the update procedure we make sure that $(\sigma \cdot b_i, d) <_{j+1}^0 (\tau, a)$. Subsequently, step 2 will not swap these two items since (τ, a) is not erasable. This means that we obtain $(\sigma \cdot b_i, d) <_{j+1} (\tau|_{j+1}, a)$ as well. In other words, we find $(\sigma \cdot b_i, d) \notin S_{j+1}$, as required.

From these two observations it follows by some basic combinatorics that there is a $k_1 \in \omega$, and a finite set S_Σ of Σ -items, such that for every $j \geq k_1$, the set of $\Sigma|_j$ -items in $V_j \times A$ of higher priority than $(\tau|_j, a)$ is given as $S_j = \{(\sigma|_j, d) \mid (\sigma, d) \in S_\Sigma\}$ indeed. \blacktriangleleft

Our next claim states that the priority ordering on the set S_Σ eventually stabilizes.

CLAIM 3 For any pair $(\sigma, d), (\sigma', d')$ of items in S_Σ there is a $k \in \omega$ such that for all $j \geq k$:

$$(\sigma|_j, d) <_j (\sigma'|_j, d') \text{ iff } (\sigma|_k, d) <_k (\sigma'|_k, d'). \quad (37)$$

PROOF OF CLAIM This claim can be proved by an argument similar to the proof of Claim 1, using the observation that no two items among $\{(\sigma|_{j+1}, d) \mid (\sigma, d) \in S_\Sigma\}$ for $j \geq k$ can have been swapped at stage $j+1$. To see why, it suffices to observe that any such swap would place one of the mentioned items not only below the other one, but also below the greatest element (τ, a) of F , since the trace τ is not refreshed, and the item (τ, a) not erasable. \blacktriangleleft

On the basis of Claim 3 we can extend the order $<_\Sigma$ to include the members of S_Σ as well. Thus $<_\Sigma$ is now an order defined over the set $F \cup S_\Sigma$. Our final claim about the set S_Σ is the following.

CLAIM 4 There is some $k \in \omega$ such that no item $(\sigma, b) \in S_\Sigma$ is in focus for any $j \geq k$.

PROOF OF CLAIM Suppose that, on the contrary, some item in S_Σ is in focus for infinitely many j . To derive a contradiction from this, we let (σ, d) be the highest priority item with this property among S_Σ in the ordering $<_\Sigma$ - such an item clearly exists since S_Σ is finite. We make a case distinction as to the nature of the state d .

In case d is a μ -state, it follows by the shuffle-merge condition that d must be the highest priority state such that $(\sigma|_j, d)$ is in focus for infinitely many j . But this means that (σ, d) is a member of F , and thus contradicts our choice of (τ, a) as the highest priority member of F .

On the other hand, (σ, d) cannot be a ν -item either, since then each time it is in focus, Step 2 of the update procedure of the priority order would apply to it, placing (σ, d) below the item (τ, a) . \blacktriangleleft

It follows that (τ, a) is the MSI each of the infinitely many times that it is in focus after the point k given by Claim 2, and so the proof is done. QED

9.3 Decorations

Our aim is now to define, by induction on the length of a partial match Σ , for every trace ρ on Σ two decorations $\delta_\rho, \delta_\rho^+ : A \rightarrow \mu\text{ML}$. For the definition of the tightened decoration δ_ρ^+ we need to introduce some notation.

Definition 9.21 Assume that decorations $\delta_\rho, \delta_\rho^+$ have been defined. We define the formulas $\text{tr}_\mathbb{A}^\rho(a)$ and $\text{tr}_\mathbb{A}^{\rho^+}(a)$ for $a \in A$ simply by:

$$\begin{aligned} \text{tr}_\Sigma^\rho(a) &:= \text{tr}_\mathbb{A}^{\delta_\rho}(a) \\ \text{tr}_\Sigma^{\rho^+}(a) &:= \text{tr}_\mathbb{A}^{\delta_\rho^+}(a) \end{aligned}$$

In addition, the following abbreviations will help to avoid notational clutter:

$$\begin{aligned} \text{tr}_\Sigma^b(a) &:= \text{tr}_\Sigma^{\text{mst}_\Sigma(b)}(a) & \text{tr}_\Sigma(a) &:= \text{tr}_\Sigma^a(a) \\ \text{tr}_\Sigma^b(a)^+ &:= \text{tr}_\Sigma^{\text{mst}_\Sigma(b)^+}(a) & \text{tr}_\Sigma^+(a) &:= \text{tr}_\Sigma^a(a)^+ \end{aligned}$$

The *context formula* for a finite partial match Σ is defined to be the formula

$$\gamma(\Sigma) := \bigwedge \{ \text{tr}_\Sigma(b) \mid b \in \text{Ran}(\text{last}(\Sigma)) \text{ and } b \neq a \},$$

where a is the unique state such that for some trace τ , the pair (τ, a) is the MSI of Σ . \triangleleft

We are now ready for the inductive definition of the decorations δ_ρ and δ_ρ^+ , where ρ is a trace on a partial match Σ . Recall that msi_Σ is the highest priority Σ -item in focus.

Definition 9.22 For the unique trace ρ on the initial match consisting only of the relation $\{(a_I, a_I)\}$ we define $\delta_\rho(b) = \top$ for all $b \in A$. This ensures that $\text{tr}_\Sigma^\rho(a) \equiv_K \text{tr}_\mathbb{A}(a)$ for all $a \in A$.

Inductively, suppose that decorations δ_ρ and δ_ρ^+ have been defined for all traces on the partial match Σ' . Let $\Sigma = \Sigma' \cdot Q$ be a continuation of Σ' , and let $\rho = \sigma \cdot a$ be an arbitrary trace on Σ .

$$\delta_\rho(b) := \begin{cases} \top & \text{if } \Omega(b) \text{ is even} \\ \top & \text{if } b \sqsubset \text{last}(\rho) \\ \top & \text{if } (\rho, b) <_\Sigma \text{msi}_\Sigma \\ \delta_\sigma^+(b) & \text{if } \text{msi}_\Sigma \leq_\Sigma (\rho, b). \end{cases}$$

In all cases, setting

$$\delta_\rho^+(b) := \begin{cases} \delta_\rho(b) \wedge \neg\gamma(\Sigma) & \text{if } (\rho, b) = \text{msi}_\Sigma \text{ and } b \in A^\mu \\ \delta_\rho(b) & \text{otherwise.} \end{cases}$$

defines the decoration δ_ρ^+ in terms of δ_ρ . \triangleleft

The decoration δ_ρ^+ is defined as the tightening of δ_ρ , according to a simple principle: we just tighten the formula associated with the MSI at Σ by the negation of the context formula (only if the MSI is a μ -item), and leave everything else the same. When we then define the

updated decoration with respect to any Σ -continuation $\sigma \cdot a$ of a Σ' -trace σ , the decoration δ_σ^+ is our starting point. A naive approach would be to set $\delta_{\sigma \cdot a}(b) := \delta_\sigma^+(b)$ for each b and each trace $\sigma \cdot a$. Instead, we need to reset the formulas associated with certain items to \top . In particular we do this for those $\Sigma' \cdot Q$ -items that end up below the new MSI of the extended match $\Sigma' \cdot Q$, and those items $(\sigma \cdot a, b)$ such that $\Omega(b)$ has a lower priority than $\text{last}(\sigma \cdot a) = a$.

The following proposition is in some sense the heart of our proof of Kozen's Lemma, since it is here that the context rule of Proposition 9.1 is actually used.

Proposition 9.23 *Let Σ be a partial match in the (thin) satisfiability game for \mathbb{A} such that the formula*

$$\bigwedge_{b \in \text{Ran}(\text{last}(\Sigma))} \text{tr}_\Sigma(b)$$

is consistent. Then so is

$$\bigwedge_{b \in \text{Ran}(\text{last}(\Sigma))} \Theta(b)[\text{tr}_\Sigma^b(d)^+ / d \mid d \in A].$$

Proof. We only treat the case where msi_Σ is of the form (τ, a) for a a μ -state, since the other case is easier. We can write the first conjunction as:

$$\gamma(\Sigma) \wedge \mu a. \theta,$$

where θ is an abbreviation for the formula:

$$\delta_\tau(a) \wedge \Theta(a)[\text{tr}_{(\mathbb{A} \downarrow a)^-}(b)/b \mid b \sqsubset a][\text{tr}_\Sigma^\tau(b)/b \mid a \sqsubset b].$$

By Proposition 9.1 we get that the following conjunction is consistent:

$$\gamma(\Sigma) \wedge \theta[\mu a. \neg \gamma(\Sigma) \wedge \theta / a].$$

It remains to prove the following two claims.

CLAIM 1 All $b \neq a$ in the range of $\text{last}(\Sigma)$ satisfy

$$\text{tr}_\Sigma(b) \leq \Theta(b)[\text{tr}_\Sigma^b(d)^+ / d \mid d \in A].$$

CLAIM 2 For $a \in \text{Ran}(\text{last}(\Sigma))$ we have

$$\theta[\mu a. \neg \gamma(\Sigma) \wedge \theta / a] \leq_K \Theta(a)[\text{tr}_\Sigma^\tau(d)^+ / d \mid d \in A].$$

The proof of both claims will be given in the appendix.

QED

9.4 The proof of Kozen's Lemma

We are now ready for the proof of Theorem 5. As a first step, from Proposition 9.23 we shall derive the following proposition, which states that \exists has a surviving strategy which maintains the consistency of the formula $\bigwedge_{b \in \text{Ran}(\text{last}(\Sigma))} \text{tr}_\Sigma(b)$.

Proposition 9.24 *Let Σ be a partial match of length k in the (thin) satisfiability game for \mathbb{A} such that the formula*

$$\bigwedge_{b \in \text{Ran}(\text{last}(\Sigma))} \text{tr}_\Sigma(b)$$

is consistent. Then \exists has a legitimate move (Y, \mathcal{R}) such that, for all $Q \in \mathcal{R}$, the formula

$$\bigwedge_{b \in \text{Ran}(Q)} \text{tr}_{\Sigma \cdot Q}(b)$$

is consistent.

Proof. It follows by Proposition 9.23 that the formula

$$\bigwedge_{b \in \text{Ran}(\text{last}(\Sigma))} \Theta(b)[\text{tr}_\Sigma^b(d)^+ / d \mid d \in A]$$

is consistent. By Corollary 5.14 we now find an admissible move (Y, \mathcal{R}) for \exists such that, for all $Q \in \mathcal{R}$, the formula

$$(*) \quad \bigwedge_{(b,d) \in Q} \text{tr}_\Sigma^b(d)^+$$

is consistent, so it suffices to show that this implies the required conjunction $\bigwedge_{b \in \text{Ran}(Q)} \text{tr}_{\Sigma \cdot Q}(b)$. For this, it suffices to show for all $d \in \text{Ran}(Q)$ that the formula $\text{tr}_{\Sigma \cdot Q}(d)$ is implied by some conjunct of $(*)$.

So let $d \in \text{Ran}(Q)$. Then by Proposition 9.12 the trace $\text{mst}_{\Sigma \cdot Q}(d)$ is of the form $\text{mst}_\Sigma(d') \cdot d$ for some d' with $(d', d) \in Q$. So we either have $\delta_{\text{mst}_{\Sigma \cdot Q}(d)} = \delta_{\text{mst}_\Sigma(d')}^+$ or $\delta_{\text{mst}_{\Sigma \cdot Q}(d)} = \top$. In both cases we have $\text{tr}_\Sigma^{d'}(d)^+ \leq_K \text{tr}_{\Sigma \cdot Q}(d)$, and so we are done. QED

We are now ready for the proof of Theorem 5 itself.

Proof of Theorem 5. We shall define the winning strategy χ for \exists , and simultaneously maintain the induction hypothesis that for every partial χ -guided match Σ , the formula

$$\bigwedge \{ \text{tr}_\Sigma(a) \mid a \in \text{Ran}(\text{last}(\Sigma)) \}$$

is consistent.

It is clear from Proposition 9.24 that \exists has a strategy χ guaranteeing that she never gets stuck and that the induction hypothesis remains true for all χ -guided partial matches. It remains to show that this strategy is actually winning, i.e., that \exists wins every infinite χ -guided match. So suppose for a contradiction that Σ is an infinite χ -guided match containing a bad

trace. By Proposition 9.20, there is a continuous bad trace τ on Σ such that, with a denoting the the highest priority state appearing infinitely often on τ , (τ, a) stabilizes at some $k < \omega$, and $(\tau|_j, a)$ is the MSI of $\Sigma|_j$ for infinitely many $j \in \omega$. Let $S_\Sigma \subseteq V_\Sigma \times A$ be the finite set of Σ -items such that for all $j \geq k$, the set of Σ_j -items in $V_{\Sigma|_j} \times A$ that are above $(\tau|_j, a)$ is of the form $\{(\sigma|_j, d) \mid (\sigma, d) \in S_\Sigma\}$.

Let $j < j'$ be the first two indices above k for which the following hold:

- (1) $\text{Ran}(\text{last}(\Sigma|_j)) = \text{Ran}(\text{last}(\Sigma|_{j'}))$;
- (2) $(\tau|_j, a) = \text{msi}_{\Sigma|_j}$ and $(\tau|_{j'}, a) = \text{msi}_{\Sigma|_{j'}}$;
- (3) $\sigma(j) = \sigma(j')$, for all σ such that $(\sigma, d) \in S_\Sigma$ for some $d \in A$.

Clearly such indices must exist by the pigeon-hole principle, since $(\tau|_l, a)$ is the MSI for infinitely many $l < \omega$, while for each $m < \omega$ we have that $\text{Ran}(\text{last}(\Sigma|_m))$ is an element of the finite set PA , and for $m \geq k$ each of the finitely many objects $\sigma(k)$ for $(\sigma, d) \in S_\Sigma$ belongs to the finite set A .

Since (τ, a) stabilizes at k , we obtain the following claim which intuitively states that the context formulas, as expressed by the decoration values for items of higher priority than (τ, a) , get *frozen*.

CLAIM 1 For any Σ -item $(\sigma, d) \in S_\Sigma$ and any $l \geq j$ we have

$$\delta_{\sigma|_l}(d) = \delta_{\sigma|_j}(d).$$

PROOF OF CLAIM This claim can be established by a straightforward inductive proof, where the inductive step is taken care of by showing that

$$\delta_{\sigma|_{n+1}}(d) = \delta_{\sigma|_n}(d) \tag{38}$$

for all $n \geq j$. But since we both have $\text{msi}_{\Sigma|_n} \leq_{\Sigma|_n} (\tau|_n, a) <_{\Sigma|_n} (\sigma|_n, d)$ and $\text{msi}_{n+1} \leq_{n+1} (\tau|_{n+1}, a) <_{n+1} (\sigma|_{n+1}, d)$, this is an immediate consequence of the definition of $\delta_{\tau|_{n+1}}$ from $\delta_{\tau|_n}$. \blacktriangleleft

On the basis of this we can prove the following key claim.

CLAIM 2 For any $b \in \text{Ran}(\text{last}(\Sigma|_j))$, we have

$$\text{tr}_{\Sigma|_{j'}}(b) \leq_K \text{tr}_{\Sigma|_j}(b).$$

The key observation underlying the proof of Claim 2 is that Σ -items (σ, d) *above* (τ, a) in the respective priority orderings of Σ_l , $j \leq l \leq j'$, *stabilize*, implying that $\delta_{\sigma|_l}(d) = \delta_{\sigma|_j}(d)$, while items (ρ, b) *below* (τ, a) are *reset* at stages where (τ, a) provides the MSI, so that in particular at stage j we find $\delta_{\sigma|_j}(d) = \top$. Technical details are given below.

PROOF OF CLAIM Abbreviate $B := \text{Ran}(\text{last}(\Sigma|_j)) = \text{Ran}(\text{last}(\Sigma|_{j'}))$. Our proof of Claim 2 is based on the following observation.

$$\delta_{j'}^b(d) \leq_K \delta_j^b(d) \text{ for all } (b, d) \in B \times A, \tag{39}$$

where in order to avoid notational clutter, we abbreviate $\delta_j^b := \delta_{\text{mst}_j(b)}$, $\delta_{j'}^b := \delta_{\text{mst}_{j'}(b)}$, $\text{mst}_j := \text{mst}_{\Sigma|_j}$, and $\text{mst}_{j'} := \text{mst}_{\Sigma|_{j'}}$. In order to prove (39), make a case distinction.

(**case 1**) If $d \sqsubset b$ then (39) is immediate since $\delta_j^b(d) = \top$.

(**case 2**) If $b \sqsubseteq d$ and $(\text{mst}_j(b), d) <_{\Sigma_j} (\tau, a)$ then (39) is immediate as well since $\delta_j^b(d) = \top$.

(**case 3**) If $b \sqsubseteq d$ and $(\tau, a) <_{\Sigma_j} (\text{mst}_j(b), d)$ then $\text{mst}_j(b)$ is of the form $\sigma|_j$ for some $\sigma \in V_{\Sigma}$ with $(\sigma, d) \in S_{\Sigma}$. We claim that this very same Σ -trace σ also provides the most significant trace for b at stage j' , that is:

$$\text{mst}_{j'}(b) = \sigma|_{j'}. \quad (40)$$

To see this, first observe that $\text{last}(\sigma|_{j'}) = b$, since $\text{last}(\sigma|_{j'}) = \sigma(j')$ by immediate unravelling of the definitions, $\sigma(j') = \sigma(j)$ by our earlier assumption on k , and $\sigma(j) = b$ by the fact that $\sigma(j) = \text{last}(\sigma|_j) = \text{last}(\text{mst}_j(b)) = b$. But from $\text{last}(\sigma|_{j'}) = b$ and the fact that σ is a continuous trace, we immediately obtain (40).

Finally we derive (39) as follows:

$$\begin{aligned} \delta_{j'}^b(d) &= \delta_{\sigma|_{j'}}(d) && \text{(immediate by (40))} \\ &= \delta_{\sigma|_j}(d) && \text{(Claim 1)} \\ &= \delta_j^b(d) && (\sigma|_j = \text{mst}_j(b)) \end{aligned}$$

(**case 4**) If $b \sqsubseteq d$ and $(\text{mst}_j(b), d) = (\tau, a)$ then what we need to show is that $\delta_{\tau|_{j'}}(a) \leq_K \delta_{\tau|_j}(a)$. But in fact, we can prove that

$$\delta_{\tau|_l}(a) \leq_K \delta_{\tau|_j}(a) \text{ for all } l \geq j, \quad (41)$$

by a straightforward inductive proof, where the inductive step is taken care of by showing that

$$\delta_{\tau|_{n+1}}(a) \leq_K \delta_{\tau|_n}(a) \text{ for all } n \geq j. \quad (42)$$

But this is not difficult to see, since $\text{msi}_{\Sigma_n} \leq (\tau, a)$, for all $n \geq j$.

Finally, Claim 2 follows immediately from (39) by Proposition 9.3(3). \blacktriangleleft

The remaining part of the proof is the argument that we already sketched when we gave the intuitions underlying the proof of Kozen's Lemma. Recall that $\gamma(\Sigma|_j)$ is the context formula at stage j , that is

$$\gamma(\Sigma|_j) = \bigwedge \{ \text{tr}_{\Sigma|_j}(b) \mid b \in \text{Ran}(\text{last}(\Sigma|_j)), b \neq a \}.$$

Since $\text{Ran}(\text{last}(\Sigma|_j)) = \text{Ran}(\text{last}(\Sigma|_{j'}))$, Claim 2 gives:

$$\bigwedge_{b \in \text{Ran}(\text{last}(\Sigma|_{j'})) \setminus \{a\}} \text{tr}_{\Sigma|_{j'}}(b) \leq_K \gamma(\Sigma|_j), \quad (43)$$

On the other hand we claim that

$$\text{tr}_{\Sigma|_{j'}}(a) \leq_K \neg \gamma(\Sigma|_j). \quad (44)$$

To see this, observe that since $(\tau|_j, a) = \text{msi}_{\Sigma|_j}$ and since $\text{msi}_{\Sigma|_{j''}}$ is never of higher priority than $(\tau|_{j''}, a)$ for $j \leq j''$, it follows that

$$\delta_{\tau|_{j'}}(a) \leq_K \delta_{\tau|_j}^+(a) \leq_K \neg\gamma(\Sigma|_j).$$

But $\text{tr}_{\Sigma|_{j'}}$ is the tightened translation induced by $\delta_{\tau|_{j'}}$ and so by Proposition 9.3(2) we obtain

$$\text{tr}_{\Sigma|_{j'}}(a) \leq_K \delta_{\tau|_{j'}}(a).$$

From this (44) is immediate.

Finally then, since $a \in \text{Ran}(\text{last}(\Sigma|_{j'}))$ we get from (43) and (44) that

$$\bigwedge_{b \in \text{Ran}(\text{last}(\Sigma|_{j'}))} \text{tr}_{\Sigma|_{j'}}(b) \leq_K \gamma(\Sigma|_j) \wedge \neg\gamma(\Sigma|_j) \leq_K \perp,$$

directly contradicting the fact that the formula $\bigwedge_{b \in \text{Ran}(\text{last}(\Sigma|_{j'}))} \text{tr}_{\Sigma|_{j'}}(b)$ is consistent. This finishes the proof of Theorem 5. QED

10 Completeness for the modal μ -calculus

With all the pieces in place, we are ready for the main result of the paper: we shall show that every formula of the modal μ -calculus provably implies the translation of some semantically equivalent disjunctive automaton. From this result the completeness of Kozen's proof system for the modal μ -calculus follows almost immediately.

We have already come half way towards this result in Theorem 2: using arbitrary modal automata rather than disjunctive automata, we were able to prove, using comparatively elementary techniques, that every formula of the μ -calculus is provably equivalent to (the translation of) a modal automaton, i.e., $\varphi \equiv_K \mathbb{A}_\varphi$ for each formula φ . (Observe that here, and in the sequel, we will use the notation of Definition 8.1.) We now want to apply the automata-theoretic machinery that we developed in previous sections, to strengthen this result, showing that for any formula φ there is an equivalent *disjunctive* automaton \mathbb{D}_φ such that $\varphi \leq_K \mathbb{D}_\varphi$. The following proposition shows that whenever φ is the translation of a semi-disjunctive automaton this result can be proved.

Proposition 10.1 *Let \mathbb{A} be any semi-disjunctive modal automaton. Then $\mathbb{A} \leq_K \text{sim}(\mathbb{A})$.*

Proof. It is clear from Theorem 4 that there is a winning strategy for Player II in the consequence game $\mathcal{C}(\mathbb{A}, \text{sim}(\mathbb{A}))$. Since \mathbb{A} is semi-disjunctive it follows by Theorem 3 that \forall has a winning strategy in the thin satisfiability game for $\mathbb{A} \wedge \neg \text{sim}(\mathbb{A})$. By Kozen's Lemma (Theorem 5) it follows that the automaton $\mathbb{A} \wedge \neg \text{sim}(\mathbb{A})$ is inconsistent. From this and the clauses 1 and 2 of Proposition 8.15, it is immediate that $\mathbb{A} \leq_K \text{sim}(\mathbb{A})$. QED

We are now ready for the statement and proof of our main result.

Theorem 6 *For every formula $\varphi \in \mu\text{ML}$ there is a semantically equivalent disjunctive automaton \mathbb{D} such that $\varphi \leq_K \mathbb{D}$.*

Proof. By Fact 3.15 any modal fixpoint formula is provably equivalent to a formula in negation normal form. Hence without loss of generality we may prove the theorem for formulas in this shape, and proceed by an induction on the complexity of such formulas. That is, the base cases of the induction are the literals, and we need to consider induction steps for conjunctions, disjunctions, both modal operators and both fixpoint operators.

The base case for literals follows immediately since it is easy to see that the modal automaton \mathbb{A}_φ corresponding to a literal φ is already disjunctive. Disjunctions are easy since the operation \vee on automata preserves the property of being disjunctive. For conjunctions: given formulas φ, φ' we have semantically equivalent disjunctive automata \mathbb{D}, \mathbb{D}' such that $\varphi \leq_K \mathbb{D}$ and $\varphi' \leq_K \mathbb{D}'$. By the first clause of Proposition 8.15 we get $\varphi \wedge \varphi' \leq_K \mathbb{D} \wedge \mathbb{D}'$. But by the Propositions 6.16(4) and 8.16 the automaton $\mathbb{D} \wedge \mathbb{D}'$ is semi-disjunctive modulo provable equivalence, and we can apply Proposition 10.1 to obtain the desired conclusion. The cases for the modalities are easy since boxes and diamonds as operations on automata preserve the property of being disjunctive.

For the greatest fixpoint operator, consider the formula $\varphi = \nu x. \alpha(x)$, and assume inductively that there is a disjunctive automaton \mathbb{A} for α such that $\alpha \equiv \mathbb{A}$ and $\alpha \leq_K \mathbb{A}$. It follows

by Proposition 8.15(4) that $\varphi = \nu x.\alpha \leq_K \nu x.\mathbb{A}$, and since $\nu x.\mathbb{A}$ is semidisjunctive modulo provable equivalence by the Propositions 6.16(6) and 8.16, by Proposition 10.1 we are done.

Finally, we cover the crucial case for $\varphi = \mu x.\alpha(x)$. By the induction hypothesis there is a semantically equivalent disjunctive automaton \mathbb{A} for α such that $\alpha \leq_K \mathbb{A}$. Let $\mathbb{D} := \text{sim}(\mu x.\mathbb{A})$. This automaton is clearly semantically equivalent to φ . We want to show that

$$\mu x.\mathbb{A} \leq_K \mathbb{D}, \tag{45}$$

from which the result follows since $\varphi = \mu x.\alpha \leq_K \mu x.\mathbb{A}$ by Proposition 8.15(4) and the induction hypothesis.

In order to prove (45) we will work with the automaton \mathbb{A}^x . First observe that

$$\mathbb{A}^x[\mathbb{D}/x] \models_{\mathcal{G}} \mathbb{A}^x[\mu x.\mathbb{A}/x],$$

by Theorem 4, and that

$$\mathbb{A}^x[\mu x.\mathbb{A}/x] \models_{\mathcal{G}} \mu x.\mathbb{A}$$

by Proposition 5.19. But since

$$\mu x.\mathbb{A} \models_{\mathcal{G}} \text{sim}(\mu x.\mathbb{A}) = \mathbb{D}$$

by Theorem 4 again, we find by transitivity of the game consequence relation (Proposition 5.21) that

$$\mathbb{A}^x[\mathbb{D}/x] \models_{\mathcal{G}} \mathbb{D}.$$

By the Propositions 6.16(5) and 8.16 the automaton $\mathbb{A}^x[\mathbb{D}/x]$ is semi-disjunctive modulo provable equivalence, and so by Theorem 3 the automaton $\mathbb{A}^x[\mathbb{D}/x] \wedge \neg \mathbb{D}$ has a thin refutation, whence by Kozen's Lemma (Theorem 5) and Proposition 8.15 this automaton is inconsistent. In other words, we have

$$\mathbb{A}^x[\mathbb{D}/x] \leq_K \mathbb{D}.$$

Then by Proposition 8.15(5) we obtain that

$$\text{tr}(\mathbb{A}^x[\text{tr}(\mathbb{D})/x]) \leq_K \text{tr}(\mathbb{D}),$$

so that one application of the fixpoint rule yields

$$\mu x.\text{tr}(\mathbb{A}^x) \leq_K \mathbb{D}.$$

By (30) in Proposition 8.14 this suffices to prove (45). QED

Finally we see how Theorem 6 implies completeness.

Theorem 10.2 (Completeness) *Every consistent formula $\varphi \in \mu\text{ML}$ is satisfiable.*

Proof. Given a consistent formula φ , by Theorem 6 there exists a semantically equivalent disjunctive automaton \mathbb{D} such that $\varphi \leq_K \mathbb{D}$. Clearly then, \mathbb{D} is consistent too, whence by Theorem 5, \exists has a winning strategy in the thin satisfiability game for \mathbb{D} . But \mathbb{D} is disjunctive and hence semi-disjunctive, and so by Proposition 6.15 \exists also has a winning strategy in $\mathcal{S}(\mathbb{D})$. It then follows by the adequacy of the satisfiability game (Proposition 5.10) that \mathbb{D} is satisfiable, and so φ , being semantically equivalent to \mathbb{D} , is satisfiable as well. QED

A Proof details of section 8

A.1 Vectorial notation

In this appendix we will have use of the vectorial notation for fixpoints. For simplicity we shall only use it for binary vectors, which is all we need. Given variables x_0, x_1 and formulas φ_0 and φ_1 in which these variables may occur, we let the expression:

$$\eta \begin{pmatrix} x_0 \\ x_1 \end{pmatrix} \cdot \begin{pmatrix} \varphi_0 \\ \varphi_1 \end{pmatrix}$$

abbreviate the pair of formulas $\langle \eta x_0.\varphi_0[\eta x_1.\varphi_1/x_1], \eta x_1.\varphi_1[\eta x_0.\varphi_0/x_0] \rangle$. It is a standard fact [1], known as the Bekič principle, that this pair of formulas defines a multi-variate fixpoint of the formulas φ_0, φ_1 , but the details of this will not concern us here. We shall think of pairs of formulas as column vectors, so that we may write:

$$\eta \begin{pmatrix} x_0 \\ x_1 \end{pmatrix} \cdot \begin{pmatrix} \varphi_0 \\ \varphi_1 \end{pmatrix} = \begin{pmatrix} \eta x_0.\varphi_0[\eta x_1.\varphi_1/x_1] \\ \eta x_1.\varphi_1[\eta x_0.\varphi_0/x_0] \end{pmatrix}.$$

Generally, given vectors $\langle \varphi_0, \varphi_1 \rangle$ and $\langle \psi_0, \psi_1 \rangle$ of formulas, we write

$$\begin{pmatrix} \varphi_0 \\ \varphi_1 \end{pmatrix} \leq_K \begin{pmatrix} \psi_0 \\ \psi_1 \end{pmatrix}$$

to say that $\varphi_0 \leq_K \psi_0$ and $\varphi_1 \leq_K \psi_1$. Similarly we write

$$\begin{pmatrix} \varphi_0 \\ \varphi_1 \end{pmatrix} \equiv_K \begin{pmatrix} \psi_0 \\ \psi_1 \end{pmatrix}$$

to say that $\varphi_0 \equiv_K \psi_0$ and $\varphi_1 \equiv_K \psi_1$.

In what follows, we need to extend the definition of \mathbb{A}^- to automaton structures that are not necessarily linear. Given such a structure $\mathbb{A} = (A, \Theta, \Omega)$ we let \mathbb{A}^- denote the automaton structure $(A^-, \Theta^-, \Omega^-)$ where $A^- = A \setminus M$ is obtained from A by removing the set M of states in A with maximal priority, and Θ^-, Ω^- are the restrictions of the maps Θ, Ω to the set A^- .

We shall have use of the observation below.

Proposition A.1 *Let \mathbb{A} be any modal automaton with exactly two maximal states m_0, m_1 , of the same priority. Then:*

$$\begin{pmatrix} \text{tr}_{\mathbb{A}}(m_0) \\ \text{tr}_{\mathbb{A}}(m_1) \end{pmatrix} \equiv_K \eta \begin{pmatrix} m_0 \\ m_1 \end{pmatrix} \cdot \begin{pmatrix} \Theta(m_0)[\text{tr}_{\mathbb{A}^-}(b)/b \mid b \in A^-] \\ \Theta(m_1)[\text{tr}_{\mathbb{A}^-}(b)/b \mid b \in A^-] \end{pmatrix}$$

where η is μ or ν depending on the parity of m_0, m_1 .

Proof. Just recall that $\text{tr}_{\mathbb{A}}$ is defined with reference to some arbitrarily chosen linearization of \mathbb{A} , and the result is unique up to provable equivalence. So by picking a linearization \mathbb{A}^l

such that $m_1 \sqsubset_{\mathbb{A}^l} m_0$ (and noting that $(\mathbb{A}^l)^{-}$ is a linearization of \mathbb{A}^-) we get:

$$\begin{aligned}
& \mathbf{tr}_{\mathbb{A}}(m_0) \\
\equiv_K & \eta m_0. \Theta(m_0) [\mathbf{tr}_{(\mathbb{A}^l)^-}(b) \mid b \sqsubset_{\mathbb{A}^l} m_0] && \text{(Def. } \mathbf{tr}_{\mathbb{A}}) \\
= & \eta m_0. \Theta(m_0) [\mathbf{tr}_{(\mathbb{A}^l)^-}(m_1)/m_1, \mathbf{tr}_{(\mathbb{A}^l)^-}(b)/b \mid b \sqsubset_{\mathbb{A}} m_1] && \text{(obvious)} \\
= & \eta m_0. \Theta(m_0) [\mathbf{tr}_{(\mathbb{A}^l)^-}(m_1)/m_1, \mathbf{tr}_{(\mathbb{A}^l)^-}(b) [\mathbf{tr}_{(\mathbb{A}^l)^-}(m_1)/m_1] / b \mid b \sqsubset_{\mathbb{A}} m_1] \\
& \hspace{15em} \text{(Def. } \mathbf{tr}_{(\mathbb{A}^l)^-}(b) \text{ for } b \sqsubset m_1) \\
= & \eta m_0. \Theta(m_0) [\mathbf{tr}_{(\mathbb{A}^l)^-}(b)/b \mid b \sqsubset_{\mathbb{A}} m_1] [\mathbf{tr}_{(\mathbb{A}^l)^-}(m_1)/m_1] && \text{(Fact 3.4)} \\
\equiv_K & \eta m_0. \Theta(m_0) [\mathbf{tr}_{\mathbb{A}^-}(b)/b \mid b \in A^-] [\mathbf{tr}_{(\mathbb{A}^l)^-}(m_1)/m_1] && \text{(above remark)} \\
= & \eta m_0. \Theta(m_0) [\mathbf{tr}_{\mathbb{A}^-}(b)/b \mid b \in A^-] [\eta m_1. \Theta(m_1) [\mathbf{tr}_{(\mathbb{A}^l)^-}(d)/d \mid d \sqsubset_{\mathbb{A}} m_1] / m_1] \\
& \hspace{15em} \text{(Def. } \mathbf{tr}_{(\mathbb{A}^l)^-}(m_1)) \\
= & \eta m_0. \Theta(m_0) [\mathbf{tr}_{\mathbb{A}^-}(b)/b \mid b \in A^-] [\eta m_1. \Theta(m_1) [\mathbf{tr}_{\mathbb{A}^-}(d)/d \mid d \in A^-] / m_1]
\end{aligned}$$

as required. The case for $\mathbf{tr}_{\mathbb{A}}(m_1)$ is similar, picking a linearization with $m_0 \sqsubset_{\mathbb{A}^l} m_1$. QED

The next observation we shall require for the vectorial notation is the following vectorial formulation of the induction rule. Here, and in what follows, we take a substitution applied to a vector to produce as a result the vector obtained by applying the substitution to each entry.

Proposition A.2 *For all formulas φ_0, φ_1 and all formulas ψ_0, ψ_1 with no free occurrences of variables x_0, x_1 , we have:*

$$\left(\begin{array}{c} \varphi_0 \\ \varphi_1 \end{array} \right) [\psi_i/x_i] \leq_K \left(\begin{array}{c} \psi_0 \\ \psi_1 \end{array} \right) \text{ implies } \mu \left(\begin{array}{c} x_0 \\ x_1 \end{array} \right) \cdot \left(\begin{array}{c} \varphi_0 \\ \varphi_1 \end{array} \right) \leq_K \left(\begin{array}{c} \psi_0 \\ \psi_1 \end{array} \right).$$

Dually, we have:

$$\left(\begin{array}{c} \psi_0 \\ \psi_1 \end{array} \right) \leq_K \left(\begin{array}{c} \varphi_0 \\ \varphi_1 \end{array} \right) [\psi_i/x_i] \text{ implies } \left(\begin{array}{c} \psi_0 \\ \psi_1 \end{array} \right) \leq_K \nu \left(\begin{array}{c} x_0 \\ x_1 \end{array} \right) \cdot \left(\begin{array}{c} \varphi_0 \\ \varphi_1 \end{array} \right).$$

Proof. We only show the case for the least fixpoint, since the case for greatest fixpoints is dual. Suppose that $\varphi_0[\psi_i/x_i] \leq_K \psi_0$ and $\varphi_1[\psi_i/x_i] \leq_K \psi_1$. We need to show that $\mu x_0. \varphi_0[\mu x_1. \varphi_1/x_1] \leq_K \psi_0$ and $\mu x_1. \varphi_1[\mu x_0. \varphi_0/x_0] \leq_K \psi_1$. We only prove the first item since the second is similar.

From $\varphi_1[\psi_i/x_i] \leq_K \psi_1$ we get $\mu x_1. \varphi_1[\psi_0/x_0] \leq_K \psi_1$ by the induction rule, since $\varphi_1[\psi_i/x_i] = \varphi_1[\psi_0/x_0][\psi_1/x_1]$ (which holds because x_1 is not free in ψ_0). So we get:

$$\begin{aligned}
\varphi_0[\mu x_1. \varphi_1/x_1][\psi_0/x_0] &= \varphi_0[\psi_0/x_0, \mu x_1. \varphi_1[\psi_0/x_0]/x_1] && \text{(Fact 3.4)} \\
&\leq_K \varphi_0[\psi_0/x_0, \psi_1/x_1] && \text{(monotonicity)} \\
&\leq_K \psi_0 && \text{(assumption)}
\end{aligned}$$

Hence by the induction rule we get:

$$\mu x_0. \varphi_0[\mu x_1. \varphi_1/x_1] \leq_K \psi_0$$

as required. QED

We will also need the following:

Proposition A.3 *Let \mathbb{A} be any modal automaton with exactly two maximal states m_0, m_1 , of the same priority. Then for all $a \in A$, $a \notin \{m_1, m_2\}$, we have:*

$$\mathbf{tr}_{\mathbb{A}}(a) \equiv_K \mathbf{tr}_{\mathbb{A}^-}(a)[\mathbf{tr}_{\mathbb{A}}(m_0)/m_0, \mathbf{tr}_{\mathbb{A}}(m_1)/m_1]$$

Proof. Let \mathbb{A}^l be an arbitrarily chosen linearization of \mathbb{A} with $m_0 \sqsubset_{\mathbb{A}^l} m_1$. First note that we have:

$$\mathbf{tr}_{\mathbb{A}^l}(m_0) \equiv_K \mathbf{tr}_{\mathbb{A}^{l-}}(m_0)[\mathbf{tr}_{\mathbb{A}^l}(m_1)/m_1] \quad (46)$$

by definition of the translation map $\mathbf{tr}_{\mathbb{A}^l}$, since m_1 is the maximal state. Now pick a $a \notin \{m_0, m_1\}$. Since \mathbb{A}^{l--} is a linearization of \mathbb{A}^- , by Proposition 8.10 we get that:

$$\mathbf{tr}_{\mathbb{A}^-}(a) \equiv_K \mathbf{tr}_{\mathbb{A}^{l--}}(a) \quad (47)$$

Hence, we get:

$$\begin{aligned} & \mathbf{tr}_{\mathbb{A}^-}(a)[\mathbf{tr}_{\mathbb{A}}(m_0)/m_0, \mathbf{tr}_{\mathbb{A}}(m_1)/m_1] \\ \equiv_K & \mathbf{tr}_{\mathbb{A}^{l--}}(a)[\mathbf{tr}_{\mathbb{A}}(m_0)/m_0, \mathbf{tr}_{\mathbb{A}}(m_1)/m_1] && (47) \\ \equiv_K & \mathbf{tr}_{\mathbb{A}^{l--}}(a)[\mathbf{tr}_{\mathbb{A}^l}(m_0)/m_0, \mathbf{tr}_{\mathbb{A}^l}(m_1)/m_1] && (\text{Proposition 8.10}) \\ \equiv_K & \mathbf{tr}_{\mathbb{A}^{l--}}(a)[\mathbf{tr}_{\mathbb{A}^{l-}}(m_0)[\mathbf{tr}_{\mathbb{A}^l}(m_1)/m_1]/m_0, \mathbf{tr}_{\mathbb{A}^l}(m_1)/m_1] && (46) \\ \equiv_K & \mathbf{tr}_{\mathbb{A}^{l--}}(a)[\mathbf{tr}_{\mathbb{A}^{l-}}(m_0)][\mathbf{tr}_{\mathbb{A}^l}(m_1)/m_1] && (\text{Fact 3.4}) \\ \equiv_K & \mathbf{tr}_{\mathbb{A}^{l-}}(a)[\mathbf{tr}_{\mathbb{A}^l}(m_1)/m_1] && (\text{definition } \mathbf{tr}_{\mathbb{A}^{l-}}) \\ \equiv_K & \mathbf{tr}_{\mathbb{A}^l}(a) && (\text{definition } \mathbf{tr}_{\mathbb{A}^l}) \\ \equiv_K & \mathbf{tr}_{\mathbb{A}}(a) && (\text{Proposition 8.10}) \end{aligned}$$

QED

A.2 Proofs of results in section 8

By Proposition 8.10 we may assume without loss of generality that \mathbb{A} is a linear automaton.

We shall establish the following proposition, from which Proposition 8.14 (29) immediately follows. Note that the equivalences in this proposition, characterizing the relation between the automata \mathbb{A} and \mathbb{A}^x , were already given in Remark 4.20. Recall that κ denotes the substitution defined by $b \mapsto (\underline{x} \wedge b_0) \vee b_1$.

Proposition A.4 *For a linear modal automaton \mathbb{A} it holds that:*

$$\mathbf{tr}_{\mathbb{A}}(a) \equiv_K (x \wedge \mathbf{tr}_{\mathbb{A}^x}(a_0)) \vee \mathbf{tr}_{\mathbb{A}^x}(a_1) \quad (48)$$

and

$$\mathbf{tr}_{\mathbb{A}^x}(a_i) \equiv_K \theta_i^a[\mathbf{tr}_{\mathbb{A}}(b)/b \mid b \in A] \quad (49)$$

for all $a \in A$ and $i \in \{0, 1\}$.

Proof. We shall prove the proposition by induction on the size of \mathbb{A} . Let m denote the unique maximal state of \mathbb{A} in the order $\sqsubset_{\mathbb{A}}$, which exists because we assumed \mathbb{A} to be linear. We first compare the structures \mathbb{A}^{x-} and \mathbb{A}^{-x} .

CLAIM 1

$$\mathbb{A}^{x-} = \mathbb{A}^{-x}[\kappa(m)/m]$$

PROOF OF CLAIM The claim can be proved by a simple direct verification. Note that when constructing \mathbb{A}^{x-} we first duplicate the maximal state m to get two maximal states m_0, m_1 in \mathbb{A}^x , which are then both removed in \mathbb{A}^{x-} . When we form $\mathbb{A}^{-x}[\kappa(m)/m]$ we first remove the maximal state m to get \mathbb{A}^- in which m now appears as a guarded propositional variable in some one-step formulas, hence the new states m_0, m_1 are not created when we construct \mathbb{A}^{-x} , and this is why we have to perform the substitution $m \mapsto \kappa(m)$. ◀

We now get the following technical claim:

CLAIM 2 For all $a \in A^-$, $i = 0, 1$ we have

$$\mathbf{tr}_{\mathbb{A}^-}(a)[\kappa(m)/m] \equiv_K (x \wedge \mathbf{tr}_{\mathbb{A}^{x-}}(a_0)) \vee \mathbf{tr}_{\mathbb{A}^{x-}}(a_1) \quad (50)$$

and

$$\mathbf{tr}_{\mathbb{A}^{x-}}(a_i) \equiv_K \theta_i^a[\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-][\kappa(m)/m] \quad (51)$$

PROOF OF CLAIM For (50), note that

$$\mathbf{tr}_{\mathbb{A}^-}(a)[\kappa(m)/m] \equiv_K ((x \wedge \mathbf{tr}_{\mathbb{A}^{-x}}(a_0)) \vee \mathbf{tr}_{\mathbb{A}^{-x}}(a_1))[\kappa(m)/m]$$

by the induction hypothesis on (48). Now observe that $\mathbf{tr}_{\mathbb{A}^{-x}}(a_i)[\kappa(m)/m] \equiv_K \mathbf{tr}_{\mathbb{A}^{x-}}(a_i)$ by Claim 1, and (50) follows.

Furthermore, observe that

$$\begin{aligned} \mathbf{tr}_{\mathbb{A}^{x-}}(a_i) &= \mathbf{tr}_{\mathbb{A}^{-x}[\kappa(m)/m]}(a_i) && \text{(Claim 1)} \\ &\equiv_K \mathbf{tr}_{\mathbb{A}^{-x}}(a_i)[\kappa(m)/m] && \text{(obvious)} \\ &\equiv_K \theta_i^a[\mathbf{tr}_{\mathbb{A}^-}(b)/b \mid b \in A^-][\kappa(m)/m] && \text{(IH on (49))} \end{aligned}$$

which proves (51). ◀

From this claim we get:

CLAIM 3

$$\begin{pmatrix} \mathbf{tr}_{\mathbb{A}^x}(m_0) \\ \mathbf{tr}_{\mathbb{A}^x}(m_1) \end{pmatrix} \equiv_K \eta \begin{pmatrix} m_0 \\ m_1 \end{pmatrix} \cdot \begin{pmatrix} \theta_0^m \\ \theta_1^m \end{pmatrix} [\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-][\kappa(m)/m] \quad (52)$$

PROOF OF CLAIM Note that the maximal states of \mathbb{A}^x are precisely the states m_0 and m_1 .

Hence we obtain:

$$\begin{aligned}
& \begin{pmatrix} \mathbf{tr}_{\mathbb{A}^x}(m_0) \\ \mathbf{tr}_{\mathbb{A}^x}(m_1) \end{pmatrix} \\
\equiv_K & \eta \begin{pmatrix} m_0 \\ m_1 \end{pmatrix} \cdot \begin{pmatrix} \Theta^x(m_0) \\ \Theta^x(m_1) \end{pmatrix} [\mathbf{tr}_{\mathbb{A}^{x-}}(a_i)/a_i \mid (a, i) \in A^- \times \{0, 1\}] & \text{(Proposition A.1)} \\
= & \eta \begin{pmatrix} m_0 \\ m_1 \end{pmatrix} \cdot \begin{pmatrix} \theta_0^m[\kappa] \\ \theta_1^m[\kappa] \end{pmatrix} [\mathbf{tr}_{\mathbb{A}^{x-}}(a_i)/a_i \mid (a, i) \in A^- \times \{0, 1\}] & \text{(definition } \Theta^x) \\
= & \eta \begin{pmatrix} m_0 \\ m_1 \end{pmatrix} \cdot \begin{pmatrix} \theta_0^m[((x \wedge \mathbf{tr}_{\mathbb{A}^{x-}}(a_0)) \vee \mathbf{tr}_{\mathbb{A}^{x-}}(a_1))/a \mid a \in A^-, \kappa(m)/m] \\ \theta_1^m[((x \wedge \mathbf{tr}_{\mathbb{A}^{x-}}(a_0)) \vee \mathbf{tr}_{\mathbb{A}^{x-}}(a_1))/a \mid a \in A^-, \kappa(m)/m] \end{pmatrix} & \text{(Fact 3.4)} \\
\equiv_K & \eta \begin{pmatrix} m_0 \\ m_1 \end{pmatrix} \cdot \begin{pmatrix} \theta_0^m[\mathbf{tr}_{\mathbb{A}^-}(a)[\kappa(m)/m]/a \mid a \in A^-, \kappa(m)/m] \\ \theta_1^m[\mathbf{tr}_{\mathbb{A}^-}(a)[\kappa(m)/m]/a \mid a \in A^-, \kappa(m)/m] \end{pmatrix} & \text{(Claim 2(50))} \\
\equiv_K & \eta \begin{pmatrix} m_0 \\ m_1 \end{pmatrix} \cdot \begin{pmatrix} \theta_0^m[\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-][\kappa(m)/m] \\ \theta_1^m[\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-][\kappa(m)/m] \end{pmatrix} & \text{(Fact 3.4)}
\end{aligned}$$

which suffices to prove the claim. \blacktriangleleft

We now prove (48) and (49) by a case distinction, depending on whether the priority of m in \mathbb{A} is odd or even. Suppose first that it is odd. We start by establishing the left-to-right inequality of (48):

CLAIM 4 If $m \in A^\mu$ then $\mathbf{tr}_{\mathbb{A}}(a) \leq_K (x \wedge \mathbf{tr}_{\mathbb{A}^x}(a_0)) \vee \mathbf{tr}_{\mathbb{A}^x}(a_1)$ for all $a \in A$.

PROOF OF CLAIM We first consider the case where $a = m$. Given the definition of $\mathbf{tr}_{\mathbb{A}}(m)$ as

$$\mathbf{tr}_{\mathbb{A}}(m) = \mu m. \Theta(m) [\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-],$$

it suffices to show that the formula $\rho := (x \wedge \mathbf{tr}_{\mathbb{A}^x}(m_0)) \vee \mathbf{tr}_{\mathbb{A}^x}(m_1)$ is a prefixpoint of the formula $\theta' := \Theta(m) [\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-]$. This we can prove directly, as we will see now. For succinctness we abbreviate the substitution $[\mathbf{tr}_{\mathbb{A}^x}(m_i)/m_i \mid i = 0, 1]$ as τ .

$$\begin{aligned}
& \theta'[\rho/m] \\
= & \Theta(m) [\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-] [(x \wedge \mathbf{tr}_{\mathbb{A}^x}(m_0)) \vee \mathbf{tr}_{\mathbb{A}^x}(m_1)/m] & \text{(definitions } \theta', \rho) \\
= & \Theta(m) [\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-] [\kappa(m)/m][\tau] & \text{(direct verification)} \\
= & \Theta(m) [\mathbf{tr}_{\mathbb{A}^-}(a)[\kappa(m)/m]/a \mid a \in A^-, \kappa(m)/m][\tau] & \text{(Fact 3.4)} \\
\equiv_K & \Theta(m) [(x \wedge \mathbf{tr}_{\mathbb{A}^{x-}}(a_0)) \vee \mathbf{tr}_{\mathbb{A}^{x-}}(a_1)/a \mid a \in A^-, \kappa(m)/m][\tau] & \text{(Claim 2(50))} \\
= & \Theta(m) [(x \wedge \mathbf{tr}_{\mathbb{A}^{x-}}(a_0)[\tau] \vee \mathbf{tr}_{\mathbb{A}^{x-}}(a_1)[\tau])/a \mid a \in A^-, \\ & \hspace{15em} \kappa(m)[\tau]/m] & \text{(Fact 3.4)} \\
= & \Theta(m) [(x \wedge \mathbf{tr}_{\mathbb{A}^{x-}}(a_0)[\tau]) \vee \mathbf{tr}_{\mathbb{A}^{x-}}(a_1)[\tau]/a \mid a \in A^-, \\ & \hspace{15em} (x \wedge \mathbf{tr}_{\mathbb{A}^x}(m_0)) \vee \mathbf{tr}_{\mathbb{A}^x}(m_1)/m] & \text{(obvious)} \\
\equiv_K & \Theta(m) [(x \wedge \mathbf{tr}_{\mathbb{A}^x}(a_0)) \vee \mathbf{tr}_{\mathbb{A}^x}(a_1)/a \mid a \in A] & \text{(Proposition A.3)} \\
= & \Theta(m) [\kappa][\mathbf{tr}_{\mathbb{A}^x}(a_i)/a_i \mid a_i \in A^x] & \text{(obvious)} \\
\equiv_K & (x \wedge \Theta^x(m_0) [\mathbf{tr}_{\mathbb{A}^x}(a_i)/a_i \mid a_i \in A^x]) \\ & \hspace{15em} \vee \Theta^x(m_1) [\mathbf{tr}_{\mathbb{A}^x}(a_i)/a_i \mid a_i \in A^x] & \text{(definitions } \Theta^x, \Theta(m)) \\
\equiv_K & (x \wedge \mathbf{tr}_{\mathbb{A}^x}(m_0)) \vee \mathbf{tr}_{\mathbb{A}^x}(m_1) & \text{(Proposition 8.13(27))} \\
= & \rho.
\end{aligned}$$

In what follows we shall abbreviate substitutions like $[\mathbf{tr}_{\mathbb{A}^x}(m_i)/m_i \mid i = 0, 1]$ by $[\mathbf{tr}_{\mathbb{A}^x}(m_i)/m_i]$. For the case where $a \neq m$, we reason as follows:

$$\begin{aligned}
\mathbf{tr}_{\mathbb{A}}(a) &= \mathbf{tr}_{\mathbb{A}^-}(a)[\mathbf{tr}_{\mathbb{A}}(m)/m] && \text{(definition } \mathbf{tr}_{\mathbb{A}}) \\
&\leq_K \mathbf{tr}_{\mathbb{A}^-}(a)[((x \wedge \mathbf{tr}_{\mathbb{A}^x}(m_0)) \vee \mathbf{tr}_{\mathbb{A}^x}(m_1))/m] && \text{(previous case)} \\
&= \mathbf{tr}_{\mathbb{A}^-}(a)[\kappa(m)/m][\mathbf{tr}_{\mathbb{A}^x}(m_i)/m_i] && \text{(obvious)} \\
&\equiv_K ((x \wedge \mathbf{tr}_{\mathbb{A}^x}(a_0)) \vee \mathbf{tr}_{\mathbb{A}^x}(a_1))[\mathbf{tr}_{\mathbb{A}^x}(m_i)/m_i] && \text{(Claim 2(50))} \\
&\equiv_K (x \wedge \mathbf{tr}_{\mathbb{A}^x}(a_0))[\mathbf{tr}_{\mathbb{A}^x}(m_i)/m_i] \vee \mathbf{tr}_{\mathbb{A}^x}(a_1)[\mathbf{tr}_{\mathbb{A}^x}(m_i)/m_i] && \text{(obvious)} \\
&\equiv_K (x \wedge \mathbf{tr}_{\mathbb{A}^x}(a_0)) \vee \mathbf{tr}_{\mathbb{A}^x}(a_1) && \text{(Proposition A.3)}
\end{aligned}$$

◀

Before turning to the opposite inequality of (48) we prove the left-to-right inequality of (49).

CLAIM 5 If $m \in A^\mu$ then $\mathbf{tr}_{\mathbb{A}^x}(a_i) \leq_K \theta_i^a[\mathbf{tr}_{\mathbb{A}}(b)/b \mid b \in A]$ for all $a \in A$.

PROOF OF CLAIM By Claim 3, it suffices to show that $\begin{pmatrix} \theta_0[\mathbf{tr}_{\mathbb{A}}(b)/b \mid b \in A] \\ \theta_1[\mathbf{tr}_{\mathbb{A}}(b)/b \mid b \in A] \end{pmatrix}$ is a two-dimensional prefixpoint for $\begin{pmatrix} m_0 \\ m_1 \end{pmatrix}$ in $\begin{pmatrix} \theta_0^m \\ \theta_1^m \end{pmatrix} [\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-][\kappa(m)/m]$.

This we prove as follows, abbreviating $[\mathbf{tr}_{\mathbb{A}}(b)/b \mid b \in A]$ as σ :

$$\begin{aligned}
&\begin{pmatrix} \theta_0^m \\ \theta_1^m \end{pmatrix} [\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-][\kappa(m)/m][\theta_i[\sigma]/m_i \mid i = 0, 1] \\
&= \begin{pmatrix} \theta_0^m[\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-][((x \wedge \theta_0^m[\sigma]) \vee \theta_1^m[\sigma])/m] \\ \theta_1^m[\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-][((x \wedge \theta_0^m[\sigma]) \vee \theta_1^m[\sigma])/m] \end{pmatrix} && \text{(obvious)} \\
&\equiv_K \begin{pmatrix} \theta_0^m[\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-][\Theta(m)[\sigma]/m] \\ \theta_1^m[\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-][\Theta(m)[\sigma]/m] \end{pmatrix} && \text{(Convention 4.18)} \\
&\equiv_K \begin{pmatrix} \theta_0^m[\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-][\mathbf{tr}_{\mathbb{A}}(m)/m] \\ \theta_1^m[\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-][\mathbf{tr}_{\mathbb{A}}(m)/m] \end{pmatrix} && \text{(Proposition 8.13(27))} \\
&\equiv_K \begin{pmatrix} \theta_0^m[\mathbf{tr}_{\mathbb{A}}(b)/b \mid b \in A] \\ \theta_1^m[\mathbf{tr}_{\mathbb{A}}(b)/b \mid b \in A] \end{pmatrix}
\end{aligned}$$

The last equivalence is just by definition of $\mathbf{tr}_{\mathbb{A}}$ together with Fact 3.4. ◀

On the basis of this we easily prove the right-to-left inequality of (48).

CLAIM 6 If $m \in A^\mu$ then $(x \wedge \mathbf{tr}_{\mathbb{A}^x}(a_0)) \vee \mathbf{tr}_{\mathbb{A}^x}(a_1) \leq_K \mathbf{tr}_{\mathbb{A}}(a)$ for all $a \in A$.

PROOF OF CLAIM

$$\begin{aligned}
&(x \wedge \mathbf{tr}_{\mathbb{A}^x}(a_0)) \vee \mathbf{tr}_{\mathbb{A}^x}(a_1) \\
&\leq_K (x \wedge \theta_0^a[\mathbf{tr}_{\mathbb{A}}(b)/b \mid b \in A]) \vee \theta_1^a[\mathbf{tr}_{\mathbb{A}}(b)/b \mid b \in A] && \text{(Claim 5)} \\
&\equiv_K \Theta(a)[\mathbf{tr}_{\mathbb{A}}(b)/b \mid b \in A] && \text{(Convention 4.18)} \\
&\equiv_K \mathbf{tr}_{\mathbb{A}}(a) && \text{(Proposition 8.13(27))}
\end{aligned}$$

◀

Our final claim takes care of the right-to-left inequality of (49) (in fact by proving the full equivalence):

CLAIM 7 If $m \in A^\mu$ then $\theta_i^a[\mathbf{tr}_\mathbb{A}(b)/b \mid b \in A] \leq_K \mathbf{tr}_{\mathbb{A}^x}(a_i)$ for all $a \in A$.

PROOF OF CLAIM

$$\begin{aligned}
\theta_i^a[\mathbf{tr}_\mathbb{A}(b)/b \mid b \in A] &\equiv_K \theta_i^a[(x \wedge \mathbf{tr}_{\mathbb{A}^x}(b_0)) \vee \mathbf{tr}_{\mathbb{A}^x}(b_1)/b \mid b \in A] && (48) \\
&\equiv_K \theta_i^a[\kappa][\mathbf{tr}_{\mathbb{A}^x}(b_i)/b_i \mid b_i \in A^x] && (\text{obvious}) \\
&\equiv_K \Theta^x(a_i)[\mathbf{tr}_{\mathbb{A}^x}(b_i)/b_i \mid b_i \in A^x] && (\text{definition } \Theta^x) \\
&\equiv_K \mathbf{tr}_{\mathbb{A}^x}(a_i) && (\text{Proposition 8.13(27)})
\end{aligned}$$

◀

Now, for the second case, suppose the priority of m in \mathbb{A} is even. We start by establishing the right-to-left inequality of (48):

CLAIM 8 If $m \in A^\nu$ then $(x \wedge \mathbf{tr}_{\mathbb{A}^x}(a_0)) \vee \mathbf{tr}_{\mathbb{A}^x}(a_1) \leq_K \mathbf{tr}_\mathbb{A}(a)$ for all $a \in A$.

PROOF OF CLAIM We first consider the case where $a = m$. Given the definition of $\mathbf{tr}_\mathbb{A}(m)$ as

$$\mathbf{tr}_\mathbb{A}(m) = \nu m. \Theta(m)[\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-],$$

it suffices to show that the formula $\rho := (x \wedge \mathbf{tr}_{\mathbb{A}^x}(m_0)) \vee \mathbf{tr}_{\mathbb{A}^x}(m_1)$ is a postfixpoint of the formula $\theta' := \Theta(m)[\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-]$. But, using the same argument as before in Claim 4, we can prove that $\theta'[\rho/m] \equiv_K \rho$, so we are done.

For the case where $a \neq m$, we reason as follows:

$$\begin{aligned}
\mathbf{tr}_\mathbb{A}(a) &= \mathbf{tr}_{\mathbb{A}^-}(a)[\mathbf{tr}_\mathbb{A}(m)/m] && (\text{definition } \mathbf{tr}_\mathbb{A}) \\
&\geq_K \mathbf{tr}_{\mathbb{A}^-}(a)[((x \wedge \mathbf{tr}_{\mathbb{A}^x}(m_0)) \vee \mathbf{tr}_{\mathbb{A}^x}(m_1))/m] && (\text{previous case}) \\
&= \mathbf{tr}_{\mathbb{A}^-}(a)[\kappa(m)/m][\mathbf{tr}_{\mathbb{A}^x}(m_i)/m_i] && (\text{obvious}) \\
&\equiv_K ((x \wedge \mathbf{tr}_{\mathbb{A}^x}(a_0)) \vee \mathbf{tr}_{\mathbb{A}^x}(a_1))[\mathbf{tr}_{\mathbb{A}^x}(m_i)/m_i] && (\text{Claim 2(50)}) \\
&\equiv_K (x \wedge \mathbf{tr}_{\mathbb{A}^x}(a_0))[\mathbf{tr}_{\mathbb{A}^x}(m_i)/m_i] \vee \mathbf{tr}_{\mathbb{A}^x}(a_1)[\mathbf{tr}_{\mathbb{A}^x}(m_i)/m_i] && (\text{obvious}) \\
&\equiv_K (x \wedge \mathbf{tr}_{\mathbb{A}^x}(a_0)) \vee \mathbf{tr}_{\mathbb{A}^x}(a_1) && (\text{Proposition A.3})
\end{aligned}$$

◀

Before turning to the opposite inequality of (48) we prove the left-to-right inequality of (49).

CLAIM 9 If $m \in A^\nu$ then $\theta_i^a[\mathbf{tr}_\mathbb{A}(b)/b \mid b \in A] \leq_K \mathbf{tr}_{\mathbb{A}^x}(a_i)$ for all $a \in A$.

PROOF OF CLAIM By Claim 3, it suffices to show that $\begin{pmatrix} \theta_0[\mathbf{tr}_\mathbb{A}(b)/b \mid b \in A] \\ \theta_1[\mathbf{tr}_\mathbb{A}(b)/b \mid b \in A] \end{pmatrix}$ is a two-dimensional postfixpoint for $\begin{pmatrix} m_0 \\ m_1 \end{pmatrix}$ in $\begin{pmatrix} \theta_0^m \\ \theta_1^m \end{pmatrix}[\mathbf{tr}_{\mathbb{A}^-}(a)/a \mid a \in A^-][\kappa(m)/m]$. But exactly the same argument as before establishes that it is actually a fixpoint for this vector of formulas, and so a postfixpoint.

◀

On the basis of this we can prove the left-to-right inequality of (48).

CLAIM 10 If $m \in A^\nu$ then $\text{tr}_{\mathbb{A}}(a) \leq_K (x \wedge \text{tr}_{\mathbb{A}^x}(a_0)) \vee \text{tr}_{\mathbb{A}^x}(a_1)$ for all $a \in A$.

PROOF OF CLAIM

$$\begin{aligned}
& (x \wedge \text{tr}_{\mathbb{A}^x}(a_0)) \vee \text{tr}_{\mathbb{A}^x}(a_1) \\
\geq_K & (x \wedge \theta_0^a[\text{tr}_{\mathbb{A}}(b)/b \mid b \in A]) \vee \theta_1^a[\text{tr}_{\mathbb{A}}(b)/b \mid b \in A] && \text{(Claim 9)} \\
\equiv_K & \Theta(a)[\text{tr}_{\mathbb{A}}(b)/b \mid b \in A] && \text{(Convention 4.18)} \\
\equiv_K & \text{tr}_{\mathbb{A}}(a) && \text{(Proposition 8.13(27))}
\end{aligned}$$

◀

Finally, we can now prove (49) just as before, and so the proof of Proposition A.4 is done. This also concludes the proof of Proposition 8.14 (29). QED

Proof of Proposition 8.14 (30). Let $\eta \in \{\mu, \nu\}$ and let \mathbb{A}_η^x be defined as \mathbb{A}^x except we give the state \underline{x} the same priority as in $\eta x.\mathbb{A}$. One realizes by comparing $\eta x.\mathbb{A}$ and \mathbb{A}^x that

$$(\eta x.\mathbb{A})^- = (\mathbb{A}_\eta^x)^-. \quad (53)$$

Furthermore, since $\text{tr}_{\mathbb{A}^x}(\underline{x}) \equiv_K x$ it is easy to check that for each state $a \neq \underline{x}$ in \mathbb{A}^x , we have

$$\text{tr}_{(\mathbb{A}_\eta^x)^-}(a)[x/\underline{x}] \equiv_K \text{tr}_{\mathbb{A}^x}(a). \quad (54)$$

Using this we get:

$$\begin{aligned}
& \text{tr}(\mu x.\mathbb{A}) \\
= & \text{tr}_{\mu x.\mathbb{A}}(\underline{x}) && \text{(definition } \text{tr}(\mu x.\mathbb{A})) \\
= & \mu \underline{x}.\Theta_{\mu x.\mathbb{A}}(\underline{x})[\text{tr}_{(\mu x.\mathbb{A})^-}(a)/a \mid a \sqsubset \underline{x}] && \text{(definition } \text{tr}_{\mu x.\mathbb{A}}) \\
= & \mu \underline{x}.\theta_1^{a_I}[\kappa][\text{tr}_{(\mu x.\mathbb{A})^-}(a)/a \mid a \sqsubset \underline{x}] && \text{(definition } \Theta_{\mu x.\mathbb{A}}) \\
= & \mu \underline{x}.\theta_1^{a_I}[\kappa][\text{tr}_{(\mathbb{A}_\eta^x)^-}(a)/a \mid a \sqsubset \underline{x}] && (53) \\
\equiv_K & \mu \underline{x}.\theta_1^{a_I}[\kappa][\text{tr}_{(\mathbb{A}_\eta^x)^-}(a)/a \mid a \sqsubset \underline{x}][x/\underline{x}] && \text{(renaming bound variables)} \\
= & \mu \underline{x}.\theta_1^{a_I}[\kappa][\text{tr}_{(\mathbb{A}_\eta^x)^-}(a)[x/\underline{x}]/a \mid a \sqsubset \underline{x}, x/\underline{x}] && \text{(Fact 3.4)} \\
\equiv_K & \mu \underline{x}.\theta_1^{a_I}[\kappa][\text{tr}_{\mathbb{A}^x}(a)/a \mid a \sqsubset \underline{x}, x/\underline{x}] && (54) \\
\equiv_K & \mu \underline{x}.\theta_1^{a_I}[\kappa][\text{tr}_{\mathbb{A}^x}(a)/a \mid a \sqsubset \underline{x}, \text{tr}_{\mathbb{A}^x}(\underline{x})/x] && (\text{tr}_{\mathbb{A}^x} \equiv_K x) \\
\equiv_K & \mu \underline{x}.\Theta^x((a_I)_1)[\text{tr}_{\mathbb{A}^x}(a)/a \mid a \sqsubset \underline{x}, \text{tr}_{\mathbb{A}^x}(\underline{x})/x] && \text{(definition } \Theta^x) \\
= & \mu \underline{x}.\text{tr}_{\mathbb{A}^x}((a_I)_1) && \text{(Proposition 8.13(27))}
\end{aligned}$$

as required. QED

Proof of Proposition 8.14 (31). This is similar to the previous:

$$\begin{aligned}
& \mathbf{tr}(\nu x.\mathbb{A}) \\
= & \mathbf{tr}_{\nu x.\mathbb{A}}(\underline{x}) && \text{(definition } \mathbf{tr}(\nu x.\mathbb{A})) \\
= & \nu \underline{x}.\Theta_{\nu x.\mathbb{A}}(\underline{x})[\mathbf{tr}_{(\nu x.\mathbb{A})^-}(a)/a \mid a \sqsubset \underline{x}] && \text{(definition } \mathbf{tr}_{\nu x.\mathbb{A}}) \\
= & \nu \underline{x}.\theta_0^{a_I} \vee \theta_1^{a_I}[\kappa][\mathbf{tr}_{(\nu x.\mathbb{A})^-}(a)/a \mid a \sqsubset \underline{x}] && \text{(definition } \Theta_{\nu x.\mathbb{A}}) \\
= & \nu \underline{x}.\theta_0^{a_I} \vee \theta_1^{a_I}[\kappa][\mathbf{tr}_{(\mathbb{A}_0^x)^-}(a)/a \mid a \sqsubset \underline{x}] && (53) \\
\equiv_K & \nu x.\theta_0^{a_I} \vee \theta_1^{a_I}[\kappa][\mathbf{tr}_{(\mathbb{A}_0^x)^-}(a)/a \mid a \sqsubset \underline{x}][x/\underline{x}] && \text{(renaming)} \\
= & \nu x.\theta_0^{a_I} \vee \theta_1^{a_I}[\kappa][\mathbf{tr}_{(\mathbb{A}_0^x)^-}(a)[x/\underline{x}]/a \mid a \sqsubset \underline{x}, x/\underline{x}] && \text{(Fact 3.4)} \\
\equiv_K & \nu x.\theta_0^{a_I} \vee \theta_1^{a_I}[\kappa][\mathbf{tr}_{\mathbb{A}^x}(a)/a \mid a \sqsubset \underline{x}, x/\underline{x}] && (54) \\
\equiv_K & \nu x.\theta_0^{a_I} \vee \theta_1^{a_I}[\kappa][\mathbf{tr}_{\mathbb{A}^x}(a)/a \mid a \sqsubset \underline{x}, \mathbf{tr}_{\mathbb{A}^x}(\underline{x})/\underline{x}] && (\mathbf{tr}_{\mathbb{A}^x}(\underline{x}) \equiv_K x) \\
= & \nu x.\theta_0^{a_I} \vee \theta_1^{a_I}[\kappa][\mathbf{tr}_{\mathbb{A}^x}(a)/a \mid a \in A^x] && \text{(obvious)} \\
= & \nu x.\theta_0^{a_I}[\kappa] \vee \theta_1^{a_I}[\kappa][\mathbf{tr}_{\mathbb{A}^x}(a)/a \mid a \in A^x] && \text{(obvious)} \\
= & \nu x.(\Theta^x((a_I)_0) \vee \Theta^x((a_I)_1))[\mathbf{tr}_{\mathbb{A}^x}(a)/a \mid a \in A^x] && \text{(definition } \Theta^x) \\
= & \nu x.(\Theta^x((a_I)_0)[\mathbf{tr}_{\mathbb{A}^x}(a)/a \mid a \in A^x]) \vee (\Theta^x((a_I)_1)[\mathbf{tr}_{\mathbb{A}^x}(a)/a \mid a \in A^x]) && \\
& && \text{(obvious)} \\
= & \nu x.\mathbf{tr}_{\mathbb{A}^x}((a_I)_0) \vee \mathbf{tr}_{\mathbb{A}^x}((a_I)_1) && \text{(Proposition 8.13(27))}
\end{aligned}$$

QED

Proof of Proposition 8.15(1). We only consider the case of the conjunction of two automata, the case of disjunction being completely similar.

Let c be the start state of $\mathbb{A} \wedge \mathbb{B}$. Note that c will not appear anywhere in $\Theta_{\mathbb{A} \wedge \mathbb{B}}(c)$, nor in any of the formulas $\mathbf{tr}_{\mathbb{A}}(d)$ for $d \in A$ or $\mathbf{tr}_{\mathbb{B}}(d)$ for $d \in B$. Also, recall that c is the highest ranking state in $\mathbb{A} \wedge \mathbb{B}$ in the order \sqsubset . So we can calculate:

$$\begin{aligned}
& \mathbf{tr}_{\mathbb{A} \wedge \mathbb{B}}(c) \\
= & \eta_c c.\Theta_{\mathbb{A} \wedge \mathbb{B}}(c)[\mathbf{tr}_{(\mathbb{A} \wedge \mathbb{B})^-}(d)/d \mid d \sqsubset c] && \text{(definition } \mathbf{tr}_{\mathbb{A} \wedge \mathbb{B}}(c)) \\
\equiv_K & \Theta_{\mathbb{A} \wedge \mathbb{B}}(c)[\mathbf{tr}_{(\mathbb{A} \wedge \mathbb{B})^-}(d)/d \mid d \sqsubset c] && \text{(vacuous fixpoint)} \\
= & \Theta_{\mathbb{A}}(a_I)[\mathbf{tr}_{(\mathbb{A} \wedge \mathbb{B})^-}(d)/d \mid d \in A] \wedge \Theta_{\mathbb{B}}(b_I)[\mathbf{tr}_{(\mathbb{A} \wedge \mathbb{B})^-}(d)/d \mid d \in B] && \text{(definition } \Theta_{\mathbb{A} \wedge \mathbb{B}}) \\
= & \Theta_{\mathbb{A}}(a_I)[\mathbf{tr}_{\mathbb{A}}(d)/d \mid d \in A] \wedge \Theta_{\mathbb{B}}(b_I)[\mathbf{tr}_{\mathbb{B}}(d)/d \mid d \in B] && \text{(obvious)} \\
\equiv_K & \mathbf{tr}_{\mathbb{A}}(a_I) \wedge \mathbf{tr}_{\mathbb{B}}(b_I) && \text{(Proposition 8.13(27))}
\end{aligned}$$

as required.

QED

Proof of Proposition 8.15(2). For this clause we extend the negation operation to generalized modal \mathbf{X} -automata and automaton structures. First we consider the dualization operator on the one-step formulas of these structures, where proposition letters may have *guarded* occurrences. Note that when we form boolean duals of these one-step formulas, such guarded occurrences of propositional variables will be treated like proposition letters rather than variables/states.

In more detail, recall that an automaton structure over a set \mathbf{X} of proposition letters is a triple $\mathbb{A} = (A, \Theta, \Omega)$ with $\Theta : A \rightarrow \mathbf{1ML}(\mathbf{X}, A \cup \mathbf{X})$, where we may assume that \mathbf{X} and A are disjoint. The (boolean) dual α^∂ of a formula $\alpha \in \mathbf{1ML}(\mathbf{X}, A \cup \mathbf{X})$ is obtained from α by replacing all occurrences of \wedge with \vee , of \diamond with \square , of $p \in \mathbf{X}$ with $\neg p$, and vice versa. As an example,

with $p, q \in \mathbf{X}$ and $a \in A$ we get $(\neg p \vee \diamond(a \wedge q))^\partial = p \wedge \square(a \vee \neg q)$. Given a (generalized) modal \mathbf{X} -automaton structure \mathbb{A} , we define $\neg\mathbb{A}$ by dualizing all one-step formulas and raising all priorities by 1.

With this in mind, our purpose is to prove that for every automaton structure \mathbb{A} and every $a \in A$, we have

$$\mathbf{tr}_{\neg\mathbb{A}}(a) \equiv_K \neg\mathbf{tr}_{\mathbb{A}}(a). \quad (55)$$

Clause (2) then follows for the special case where \mathbb{A} is a non-generalized automaton, and a is the initial state of \mathbb{A} .

Our proof of (55) proceeds by induction on the size of \mathbb{A} , and we only consider the induction step here. So let \mathbb{A} be an automaton structure, and inductively assume that (55) holds for all automaton structures smaller than \mathbb{A} . Let m be the maximum priority state of \mathbb{A} , and note that \mathbb{A}^- is an automaton structure over $\mathbf{X} \cup \{m\}$ if \mathbb{A} is an \mathbf{X} -automaton structure. We leave it for the reader to verify that

$$\neg(\mathbb{A}^-) = (\neg\mathbb{A})^-[\neg m/m], \quad (56)$$

where the substitution $[\neg m/m]$ is needed because in $(\neg\mathbb{A})^-$ we dualize while m is still a state/variable (so $m^\partial = m$), while in $\neg(\mathbb{A}^-)$ we dualize when m is already a proposition letter (so all occurrences of m will get negated).

Furthermore, to simplify notation we assume that the parity of $\Omega_{\mathbb{A}}(m)$ is odd - the other case is handled in a completely symmetric fashion. Finally, since $\Theta_{\neg\mathbb{A}}(a)$ is obtained by dualizing $\Theta_{\mathbb{A}}(a)$, it is easy to see that for all $a \in A$ we get

$$\Theta_{\neg\mathbb{A}}(a) \equiv_K \neg\Theta_{\mathbb{A}}(a)[\neg b/b \mid b \in A]. \quad (57)$$

We can now calculate:

$$\begin{aligned} & \mathbf{tr}_{\neg\mathbb{A}}(m) \\ = & \nu m. \Theta_{\neg\mathbb{A}}(m) [\mathbf{tr}_{(\neg\mathbb{A})^-}(b)/b \mid b \sqsubset m] && \text{(definition } \mathbf{tr}_{\neg\mathbb{A}}) \\ \equiv_K & \nu m. \Theta_{\neg\mathbb{A}}(m) [\mathbf{tr}_{\neg(\mathbb{A}^-)}(b)[\neg m/m]/b \mid b \sqsubset m] && (56) \\ \equiv_K & \nu m. \Theta_{\neg\mathbb{A}}(m) [\neg\mathbf{tr}_{\mathbb{A}^-}(b)[\neg m/m]/b \mid b \sqsubset m] && \text{(IH)} \\ \equiv_K & \nu m. \neg\Theta_{\mathbb{A}}(m) [\neg b/b \mid b \in A] [\neg\mathbf{tr}_{\mathbb{A}^-}(b)[\neg m/m]/b \mid b \sqsubset m] && (57) \\ = & \nu m. \neg\Theta_{\mathbb{A}}(m) [\neg\neg\mathbf{tr}_{\mathbb{A}^-}(b)[\neg m/m]/b \mid b \sqsubset m, \neg m/m] && \text{(Fact 3.4)} \\ \equiv_K & \nu m. \neg\Theta_{\mathbb{A}}(m) [\mathbf{tr}_{\mathbb{A}^-}(b)[\neg m/m]/b \mid b \sqsubset m, \neg m/m] && \text{(propositional logic)} \\ \equiv_K & \neg\mu m. \Theta_{\mathbb{A}}(m) [\mathbf{tr}_{\mathbb{A}^-}(b)[\neg m/m]/b \mid b \sqsubset m, \neg m/m] [\neg m/m] && \text{(Fact 3.15(6))} \\ = & \neg\mu m. \Theta_{\mathbb{A}}(m) [\mathbf{tr}_{\mathbb{A}^-}(b)[\neg\neg m/m]/b \mid b \sqsubset m, \neg\neg m/m] && \text{(Fact 3.4)} \\ \equiv_K & \neg\mu m. \Theta_{\mathbb{A}}(m) [\mathbf{tr}_{\mathbb{A}^-}(b)/b \mid b \sqsubset m] && \text{(propositional logic)} \\ = & \neg\mathbf{tr}_{\mathbb{A}}(m) && \text{(definition } \mathbf{tr}_{\mathbb{A}}) \end{aligned}$$

Now that the statement has been proved for the maximal element m , we extend it to

arbitrary $a \in A$ as follows:

$$\begin{aligned}
& \mathbf{tr}_{\neg\mathbb{A}}(a) \\
= & \mathbf{tr}_{(\neg\mathbb{A})^-}(a)[\mathbf{tr}_{\neg\mathbb{A}}(m)/m] && \text{(definition } \mathbf{tr}_{\neg\mathbb{A}}) \\
\equiv_K & \mathbf{tr}_{\neg(\mathbb{A}^-)}(a)[\neg m/m][\mathbf{tr}_{\neg\mathbb{A}}(m)/m] && (56) \\
\equiv_K & \neg\mathbf{tr}_{\mathbb{A}^-}(a)[\neg m/m][\mathbf{tr}_{\neg\mathbb{A}}(m)/m] && \text{(IH)} \\
\equiv_K & \neg\mathbf{tr}_{\mathbb{A}^-}(a)[\neg m/m][\neg\mathbf{tr}_{\mathbb{A}}(m)/m] && \text{(previous case)} \\
\equiv_K & \neg\mathbf{tr}_{\mathbb{A}^-}(a)[\neg\neg\mathbf{tr}_{\mathbb{A}}(m)/m] && \text{(Fact 3.4)} \\
\equiv_K & \neg\mathbf{tr}_{\mathbb{A}^-}(a)[\mathbf{tr}_{\mathbb{A}}(m)/m] && \text{(propositional logic)} \\
= & \neg\mathbf{tr}_{\mathbb{A}}(a) && \text{(definition } \mathbf{tr}_{\mathbb{A}})
\end{aligned}$$

This concludes the proof of Clause (2). QED

Proof of Proposition 8.15(3). This case is very similar to the one for the conjunction of two automata (Proposition 8.15(1)). We only cover the case of the diamond modality.

Let c be the start state of $\diamond\mathbb{A}$, and observe that $(\diamond\mathbb{A})^- = \mathbb{A}$, that c is the maximum priority state of $\diamond\mathbb{A}$, that c will not appear anywhere in any of the formulas $\mathbf{tr}_{\mathbb{A}}(a)$ for $a \in A$. So we can calculate:

$$\begin{aligned}
& \mathbf{tr}_{\diamond\mathbb{A}}(c) \\
= & \eta_c c. \Theta_{\diamond\mathbb{A}}(c)[\mathbf{tr}_{(\diamond\mathbb{A})^-}(a)/a \mid a \sqsubset c] && \text{(definition } \mathbf{tr}_{\diamond\mathbb{A}}(c)) \\
= & \eta_c c. \diamond a_I[\mathbf{tr}_{(\diamond\mathbb{A})^-}(a)/a \mid a \sqsubset c] && \text{(definition } \Theta_{\diamond\mathbb{A}}(c)) \\
= & \diamond a_I[\mathbf{tr}_{(\diamond\mathbb{A})^-}(a)/a \mid a \sqsubset c] && \text{(vacuous fixpoint)} \\
= & \diamond a_I[\mathbf{tr}_{\mathbb{A}}(a)/a \mid a \sqsubset c] && ((\diamond\mathbb{A})^- = \mathbb{A}) \\
= & \diamond\mathbf{tr}_{\mathbb{A}}(a_I) && \text{(obvious)}
\end{aligned}$$

as required. QED

Clause (4) was proved in the main text.

Proof of Proposition 8.15(5). Finally, for clause (5), suppose that p is free and positive in \mathbb{A} . We shall first prove that for all such \mathbb{A} and all automata \mathbb{B} we have

$$\mathbf{tr}_{\mathbb{A}[\mathbb{B}/p]}(b) = \mathbf{tr}_{\mathbb{B}}(b) \tag{58}$$

for all $b \in B$. Recall that no state in A appears in $\Theta_B(b)$ for any $b \in B$, and furthermore the states in B have lower priority than the states in A . Also, note that

$$(\mathbb{A}[\mathbb{B}/p] \downarrow b)^- = (\mathbb{B} \downarrow b)^- \tag{59}$$

since all states in \mathbb{A} have higher priority in $\mathbb{A}[\mathbb{B}/p]$ than all states in \mathbb{B} . Using these facts, we can reason by an inner (downward) induction on the priorities of states in B . Supposing the inner induction hypothesis holds for all $b' \in B$ with $b \sqsubset_{\mathbb{B}} b'$ (in particular this holds vacuously for b equal to the maximum priority state m), we get:

$$\begin{aligned}
& \mathbf{tr}_{\mathbb{A}[\mathbb{B}/p]}(b) \\
\equiv_K & \eta_b b. \Theta_B(b) [\mathbf{tr}_{(\mathbb{A}[\mathbb{B}/p]\downarrow b)^-}(d)/d \mid d \sqsubset_{\mathbb{A}[\mathbb{B}/p]} b] [\mathbf{tr}_{\mathbb{A}[\mathbb{B}/p]}(b') \mid b \sqsubset_{\mathbb{A}[\mathbb{B}/p]} b'] \quad (\text{Proposition 8.13(27)}) \\
\equiv_K & \eta_b b. \Theta_B(b) [\mathbf{tr}_{(\mathbb{B}\downarrow b)^-}(d)/d \mid d \sqsubset_{\mathbb{B}} b] [\mathbf{tr}_{\mathbb{A}[\mathbb{B}/p]}(b') \mid b \sqsubset_{\mathbb{A}[\mathbb{B}/p]} b'] \quad (59) \\
\equiv_K & \eta_b b. \Theta_B(b) [\mathbf{tr}_{(\mathbb{B}\downarrow b)^-}(d)/d \mid d \sqsubset_{\mathbb{B}} b] [\mathbf{tr}_{\mathbb{A}[\mathbb{B}/p]}(b') \mid b' \in B \ \& \ b \sqsubset_{\mathbb{B}} b'] \quad (a \in A \text{ not in } \Theta_B(b)) \\
\equiv_K & \eta_b b. \Theta_B(b) [\mathbf{tr}_{(\mathbb{B}\downarrow b)^-}(d)/d \mid d \sqsubset_{\mathbb{B}} b] [\mathbf{tr}_{\mathbb{B}}(b') \mid b \sqsubset_{\mathbb{B}} b'] \quad (\text{IH}) \\
\equiv_K & \mathbf{tr}_{\mathbb{B}}(b) \quad (\text{Proposition 8.13(28)})
\end{aligned}$$

This proves (58), and in particular it shows that $\mathbf{tr}_{\mathbb{A}[\mathbb{B}/p]}(b_I) \equiv_K \mathbf{tr}(\mathbb{B})$.

We shall now prove, by induction on the size of an automaton \mathbb{A} , that for all $a \in A$ we have

$$\mathbf{tr}_{\mathbb{A}[\mathbb{B}/p]}(a) \equiv_K \mathbf{tr}_{\mathbb{A}}(a) [\mathbf{tr}(\mathbb{B})/p]. \quad (60)$$

Leaving the case where \mathbb{A} has size 1 as an exercise, we focus on the induction step. Given that the result holds for automata that are smaller than \mathbb{A} , we reason as follows. Clearly, we may assume without loss of generality that the variable p does not occur in the automaton \mathbb{B} . Let Θ' be the transition map for $\mathbb{A}[\mathbb{B}/p]$, and let b_I be the start state of \mathbb{B} . Again, we reason by an inner (downwards) induction on the priority of the states in \mathbb{A} . Supposing that the result holds for all states a' with $a \sqsubset_{\mathbb{A}} a'$, let A_a denote the set of states a' in A with $a' \sqsubset a$ and let A^a denote the set of states a' in A with $a \sqsubset a'$. Note that we have:

$$(\mathbb{A}[\mathbb{B}/p]\downarrow a)^- = (\mathbb{A}\downarrow a)^- [\mathbb{B}/p] \quad (61)$$

We get:

$$\begin{aligned}
& \text{tr}_{\mathbb{A}[\mathbb{B}/p]}(a) \\
\equiv_K & \eta_a a. \Theta'(a) [\text{tr}_{(\mathbb{A}[\mathbb{B}/p] \downarrow a)^-} (b)/b \mid b \sqsubset_{\mathbb{A}[\mathbb{B}/p]} a] \\
& \quad [\text{tr}_{\mathbb{A}[\mathbb{B}/p]}(b)/b \mid a \sqsubset_{\mathbb{A}[\mathbb{B}/p]} b] \tag{Proposition 8.13(27)} \\
\equiv_K & \eta_a a. \Theta'(a) [\text{tr}_{(\mathbb{A}[\mathbb{B}/p] \downarrow a)^-} (b)/b \mid b \sqsubset_{\mathbb{A}[\mathbb{B}/p]} a] \\
& \quad [\text{tr}_{\mathbb{A}}(b) [\text{tr}(\mathbb{B})/p] / b \mid a \sqsubset_{\mathbb{A}[\mathbb{B}/p]} b] \tag{inner IH} \\
= & \eta_a a. \Theta'(a) [\text{tr}_{(\mathbb{A} \downarrow a)^- [\mathbb{B}/p]} (b)/b \mid b \sqsubset_{\mathbb{A}[\mathbb{B}/p]} a] \\
& \quad [\text{tr}_{\mathbb{A}}(b) [\text{tr}(\mathbb{B})/p] / b \mid a \sqsubset_{\mathbb{A}[\mathbb{B}/p]} b] \tag{61} \\
= & \eta_a a. \Theta'(a) [\text{tr}_{(\mathbb{A} \downarrow a)^- [\mathbb{B}/p]} (b)/b \mid b \in A_a, \text{tr}_{(\mathbb{A} \downarrow a)^- [\mathbb{B}/p]} (b)/b \mid b \in B] \\
& \quad [\text{tr}_{\mathbb{A}}(b) [\text{tr}(\mathbb{B})/p] / b \in A^a] \tag{obvious} \\
\equiv_K & \eta_a a. \Theta'(a) [\text{tr}_{(\mathbb{A} \downarrow a)^- [\mathbb{B}/p]} (b)/b \mid b \in A_a, \text{tr}_{\mathbb{B}}/b \mid b \in B] \\
& \quad [\text{tr}_{\mathbb{A}}(b) [\text{tr}(\mathbb{B})/p] / b \in A^a] \tag{58} \\
\equiv_K & \eta_a a. \Theta'(a) [\text{tr}_{(\mathbb{A} \downarrow a)^-} (b) [\text{tr}(\mathbb{B})/p] / b \mid b \in A_a, \text{tr}_{\mathbb{B}}(b)/b \mid b \in B] \\
& \quad [\text{tr}_{\mathbb{A}}(b) [\text{tr}(\mathbb{B})/p] / b \in A^a] \tag{outer IH} \\
= & \eta_a a. \Theta_{\mathbb{A}}(a) [\Theta_{\mathbb{B}}(b_I)/p] \\
& \quad [\text{tr}_{(\mathbb{A} \downarrow a)^-} (b) [\text{tr}(\mathbb{B})/p] / b \mid b \in A_a, \text{tr}_{\mathbb{B}}(b)/b \mid b \in B] \\
& \quad [\text{tr}_{\mathbb{A}}(b) [\text{tr}(\mathbb{B})/p] / b \mid b \in A^a] \tag{definition } \Theta' \\
= & \eta_a a. \Theta_{\mathbb{A}}(a) [\Theta_{\mathbb{B}}(b_I) [\text{tr}_{(\mathbb{A} \downarrow a)^-} (b) [\text{tr}(\mathbb{B})/p] / b \mid b \in A_a, \text{tr}_{\mathbb{B}}(b)/b \mid b \in B] / p, \\
& \quad \text{tr}_{(\mathbb{A} \downarrow a)^-} (b) [\text{tr}(\mathbb{B})/p] / b \mid b \in A_a, \text{tr}_{\mathbb{B}}(b)/b \mid b \in B] \\
& \quad [\text{tr}_{\mathbb{A}}(b) [\text{tr}(\mathbb{B})/p] / b \mid b \in A^a] \tag{Fact 3.4} \\
= & \eta_a a. \Theta_{\mathbb{A}}(a) [\Theta_{\mathbb{B}}(b_I) [\text{tr}_{\mathbb{B}}(b)/b \mid b \in B] / p, \\
& \quad \text{tr}_{(\mathbb{A} \downarrow a)^-} (b) [\text{tr}(\mathbb{B})/p] / b \mid b \in A_a, \text{tr}_{\mathbb{B}}(b)/b \mid b \in B] \\
& \quad [\text{tr}_{\mathbb{A}}(b) [\text{tr}(\mathbb{B})/p] / b \mid b \in A^a] \tag{b \in A_a not in } \Theta_{\mathbb{B}}(b_I)) \\
\equiv_K & \eta_a a. \Theta_{\mathbb{A}}(a) [\text{tr}(\mathbb{B})/p, \\
& \quad \text{tr}_{(\mathbb{A} \downarrow a)^-} (b) [\text{tr}(\mathbb{B})/p] / b \mid b \in A_a, \text{tr}_{\mathbb{B}}(b)/b \mid b \in B] \\
& \quad [\text{tr}_{\mathbb{A}}(b) [\text{tr}(\mathbb{B})/p] / b \mid b \in A^a] \tag{Proposition 8.13(27)} \\
= & \eta_a a. \Theta_{\mathbb{A}}(a) [\text{tr}(\mathbb{B})/p, \\
& \quad \text{tr}_{(\mathbb{A} \downarrow a)^-} (b) [\text{tr}(\mathbb{B})/p] / b \mid b \in A_a] \\
& \quad [\text{tr}_{\mathbb{A}}(b) [\text{tr}(\mathbb{B})/p] / b \mid b \in A^a] \tag{b \in B not in } \Theta_{\mathbb{A}}(a)) \\
= & \eta_a a. \Theta_{\mathbb{A}}(a) [\text{tr}_{(\mathbb{A} \downarrow a)^-} (b)/b \mid b \in A_a] [\text{tr}(\mathbb{B})/p] \\
& \quad [\text{tr}_{\mathbb{A}}(b) [\text{tr}(\mathbb{B})/p] / b \mid b \in A^a] \tag{Fact 3.4} \\
= & \eta_a a. \Theta_{\mathbb{A}}(a) [\text{tr}_{(\mathbb{A} \downarrow a)^-} (b)/b \mid b \in A_a] \\
& \quad [\text{tr}(\mathbb{B}) [\text{tr}_{\mathbb{A}}(b) [\text{tr}(\mathbb{B})/p] / b \mid b \in A^a] / p, \\
& \quad \text{tr}_{\mathbb{A}}(b) [\text{tr}(\mathbb{B})/p] / b \mid b \in A^a] \tag{Fact 3.4} \\
= & \eta_a a. \Theta_{\mathbb{A}}(a) [\text{tr}_{(\mathbb{A} \downarrow a)^-} (b)/b \mid b \in A_a] \\
& \quad [\text{tr}(\mathbb{B})/p, \text{tr}_{\mathbb{A}}(b) [\text{tr}(\mathbb{B})/p] / b \mid b \in A^a] \tag{b \in A^a not in } \text{tr}(\mathbb{B})) \\
= & \eta_a a. \Theta_{\mathbb{A}}(a) [\text{tr}_{(\mathbb{A} \downarrow a)^-} (b)/b \mid b \in A_a] \\
& \quad [\text{tr}_{\mathbb{A}}(b)/b \mid b \in A^a] [\text{tr}(\mathbb{B})/p] \tag{Fact 3.4} \\
\equiv_K & \text{tr}_{\mathbb{A}}(a) [\text{tr}(\mathbb{B})/p] \tag{Proposition 8.13 (28)}
\end{aligned}$$

This proves (60). The desired result now follows by putting $a = a_I$, from which we get $\text{tr}(\mathbb{A}[\mathbb{B}/p]) = \text{tr}(\mathbb{A})[\text{tr}(\mathbb{B})/p]$. QED

B Proof details of section 9

Proof of Proposition 9.1. Contrapositively we will show that

$$\varphi[\mu x.\delta \wedge \varphi] \leq_K \delta \text{ implies } \mu x.\varphi \leq_K \delta, \quad (62)$$

for arbitrary μ ML-formulas $\varphi(x)$ and δ . The proposition follows from this by taking $\neg\gamma$ for δ .

In order to prove (62), assume that

$$\varphi[\mu x.(\delta \wedge \varphi)] \leq_K \delta. \quad (63)$$

We claim that the formula $\delta \wedge \mu x.(\delta \wedge \varphi[\delta \wedge x])$ is a prefixpoint of $\varphi(x)$:

$$\varphi[\delta \wedge \mu x.(\delta \wedge \varphi[\delta \wedge x])] \leq_K \delta \wedge \mu x.(\delta \wedge \varphi[\delta \wedge x]). \quad (64)$$

To arrive at this, we reason as follows. From the assumption (63) it follows by monotonicity and classical logic that

$$\varphi[\delta \wedge \mu x.(\delta \wedge \varphi[\delta \wedge x])] \leq_K \delta, \quad (65)$$

from which we obtain by classical logic that

$$\varphi[\delta \wedge \mu x.(\delta \wedge \varphi[\delta \wedge x])] \leq_K \delta \wedge \varphi[\delta \wedge \mu x.(\delta \wedge \varphi[\delta \wedge x])]. \quad (66)$$

From this, an application of the pre-fixpoint axiom to the formula $\mu x.(\delta \wedge \varphi[\delta \wedge x])$ yields

$$\varphi[\delta \wedge \mu x.(\delta \wedge \varphi[\delta \wedge x])] \leq_K \mu x.(\delta \wedge \varphi[\delta \wedge x]), \quad (67)$$

so that we obtain (64) by (65) and (67).

But from (64), an application of the pre-fixpoint rule shows that

$$\mu x.\varphi \leq_K \delta \wedge \mu x.(\delta \wedge \varphi[\delta \wedge x]), \quad (68)$$

and from this, (62) is immediate. QED

Proof of Proposition 9.23. We supply the proofs of the two claims used in the proposition.

First, note that for all $b \neq a$ in $\text{Ran}(\text{last}(\Sigma))$, and all $d \in A$, we have $\delta_{\text{mst}_\Sigma(b)}^+(d) = \delta_{\text{mst}_\Sigma(b)}(d)$ simply because $(\text{mst}_\Sigma(b), d) \neq \text{msi}_\Sigma$. It follows that for all d we have

$$\text{tr}_\Sigma^b(d) = \text{tr}_\Sigma^{\text{mst}_\Sigma(b)}(d) = \text{tr}_\mathbb{A}^{\text{mst}_\Sigma(b)}(d)^+ = \text{tr}_\Sigma^b(d)^+.$$

So we now want to prove: for all $b \neq a$ in the range of $\text{last}(\Sigma)$, we have

$$\text{tr}_\Sigma(b) \leq \Theta(b)[\text{tr}_\Sigma^b(d)/d \mid d \in A].$$

We can prove this as follows. First, by recalling that

$$\text{tr}_\Sigma(b) = \eta_b b.\delta_{\text{mst}_\Sigma(b)}(b) \wedge \Theta(b)[\text{tr}_{(\mathbb{A}\downarrow b)^-}(d)/d \mid d \sqsubset b][\text{tr}_\Sigma^b(d)/d \mid b \sqsubset d]$$

according to Definition 9.2 (and the notational conventions we have introduced subsequently), and by unfolding the fixpoint in the right-hand formula we may prove that

$$\begin{aligned} \mathbf{tr}_\Sigma(b) &\equiv_K \delta_{\mathbf{mst}_\Sigma(b)}(b) \wedge \Theta(b)[\mathbf{tr}_{(\mathbb{A}\downarrow b)^-}(d)/d \mid d \sqsubset b][\mathbf{tr}_\Sigma^b(d)/d \mid b \sqsubset d][\mathbf{tr}_\Sigma(b)/b] \\ &= \delta_{\mathbf{mst}_\Sigma(b)}(b) \wedge \Theta(b)[\mathbf{tr}_{(\mathbb{A}\downarrow b)^-}(d)/d \mid d \sqsubset b][\mathbf{tr}_\Sigma^b(d)/d \mid b \sqsubseteq d] \\ &\leq_K \Theta(b)[\mathbf{tr}_{(\mathbb{A}\downarrow b)^-}(d)/d \mid d \sqsubset b][\mathbf{tr}_\Sigma^b(d)/d \mid b \sqsubseteq d] \end{aligned}$$

Here, the equality on the second line uses that $\mathbf{tr}_\Sigma^b(b) = \mathbf{tr}_\Sigma(b)$, and that no element of A appears as a free variable in any of the formulas $\mathbf{tr}_\Sigma^b(d)$, so that the two nested substitutions can be viewed as a simultaneous substitution.

The first claim now follows if we can prove that, for every $d \sqsubset b$, we have

$$\mathbf{tr}_{(\mathbb{A}\downarrow b)^-}(d)[\mathbf{tr}_\Sigma^b(d')/d' \mid b \sqsubseteq d'] \equiv_K \mathbf{tr}_\Sigma^b(d).$$

Abbreviating $\sigma = [\mathbf{tr}_\Sigma^b(d')/d' \mid b \sqsubseteq d']$, we prove this by downwards induction on the priority of $d \sqsubset b$:

$$\begin{aligned} &\mathbf{tr}_{(\mathbb{A}\downarrow b)^-}(d)[\mathbf{tr}_\Sigma^b(d')/d' \mid b \sqsubseteq d'] \\ &= \mathbf{tr}_{(\mathbb{A}\downarrow b)^-}(d)[\sigma] \\ &\equiv_K \eta_d d. \Theta(d)[\mathbf{tr}_{(\mathbb{A}\downarrow d)^-}(e)/e \mid e \sqsubset d][\mathbf{tr}_{(\mathbb{A}\downarrow b)^-}(e)/e \mid d \sqsubset e \sqsubset b][\sigma] && \text{(Proposition 8.13)} \\ &= \eta_d d. \Theta(d)[\mathbf{tr}_{(\mathbb{A}\downarrow d)^-}(e)/e \mid e \sqsubset d][\mathbf{tr}_{(\mathbb{A}\downarrow b)^-}(e)[\sigma]/e \mid d \sqsubset e \sqsubset b, \sigma] && \text{(Fact 3.4(1))} \\ &\equiv_K \eta_d d. \Theta(d)[\mathbf{tr}_{(\mathbb{A}\downarrow d)^-}(e)/e \mid e \sqsubset d][\mathbf{tr}_\Sigma^b(e)/e \mid d \sqsubset e \sqsubset b, \sigma] && \text{(induction hyp.)} \\ &= \eta_d d. \Theta(d)[\mathbf{tr}_{(\mathbb{A}\downarrow d)^-}(e)/e \mid e \sqsubset d][\mathbf{tr}_\Sigma^b(e)/e \mid d \sqsubset e] && \text{(obvious)} \\ &\equiv_K \eta_d d. \top \wedge \Theta(d)[\mathbf{tr}_{(\mathbb{A}\downarrow d)^-}(e)/e \mid e \sqsubset d][\mathbf{tr}_\Sigma^b(e)/e \mid d \sqsubset e] && \text{(obvious)} \\ &= \eta_d d. \delta_{\mathbf{mst}_\Sigma(b)}(d) \wedge \Theta(d)[\mathbf{tr}_{(\mathbb{A}\downarrow d)^-}(e)/e \mid e \sqsubset d][\mathbf{tr}_\Sigma^b(e)/e \mid d \sqsubset e] && \text{(def. } \delta_{\mathbf{mst}_\Sigma(b)}(d)\text{)} \\ &= \mathbf{tr}_\Sigma^b(d) \end{aligned}$$

For the second claim, we first note that for each b with $a \sqsubset b$, we have $\mathbf{tr}_\Sigma^a(b) = \mathbf{tr}_\Sigma^\tau(b) \equiv_K \mathbf{tr}_\Sigma^\tau(b)^+ = \mathbf{tr}_\Sigma^a(b)^+$. This follows by an easy downwards induction on the priority of b since, for each b with $a \sqsubset b$, we have $\delta_\tau^+(b) = \delta_\tau(b)$. We now have:

$$\theta \equiv_K \delta_\tau(a) \wedge \Theta(a)[\mathbf{tr}_{(\mathbb{A}\downarrow a)^-}(b)/b \mid b \sqsubset a][\mathbf{tr}_\Sigma^a(b)^+/b \mid a \sqsubset b], \quad (69)$$

so that:

$$\begin{aligned} &\theta[\mu a. \neg \gamma(\Sigma) \wedge \theta/a] \\ &\leq_K \Theta(a)[\mathbf{tr}_{(\mathbb{A}\downarrow a)^-}(b)/b \mid b \sqsubset a][\mathbf{tr}_\Sigma^a(b)^+/b \mid a \sqsubset b][\mu a. \neg \gamma(\Sigma) \wedge \theta/a] && \text{(Obvious)} \\ &\equiv_K \Theta(a)[\mathbf{tr}_{(\mathbb{A}\downarrow a)^-}(b)/b \mid b \sqsubset a][\mathbf{tr}_\Sigma^a(b)^+/b \mid a \sqsubset b] \\ &\quad \left[\mu a. \neg \gamma(\Sigma) \wedge \delta_\tau(a) \wedge \Theta(a)[\mathbf{tr}_{(\mathbb{A}\downarrow a)^-}(b)/b \mid b \sqsubset a][\mathbf{tr}_\Sigma^a(b)^+/b \mid a \sqsubset b]/a \right] && (69) \\ &= \Theta(a)[\mathbf{tr}_{(\mathbb{A}\downarrow a)^-}(b)/b \mid b \sqsubset a][\mathbf{tr}_\Sigma^a(b)^+/b \mid a \sqsubset b] \\ &\quad \left[\mu a. \delta_\tau^+(a) \wedge \Theta(a)[\mathbf{tr}_{(\mathbb{A}\downarrow a)^-}(b)/b \mid b \sqsubset a][\mathbf{tr}_\Sigma^a(b)^+/b \mid a \sqsubset b]/a \right] && \text{(Def. } \delta_\tau^+(a)\text{)} \\ &= \Theta(a)[\mathbf{tr}_{(\mathbb{A}\downarrow a)^-}(b)/b \mid b \sqsubset a][\mathbf{tr}_\Sigma^a(b)^+/b \mid a \sqsubset b][\mathbf{tr}_\Sigma^a(a)^+/a] && \text{(Def. } \mathbf{tr}_\Sigma^a(a)^+\text{)} \\ &= \Theta(a)[\mathbf{tr}_{(\mathbb{A}\downarrow a)^-}(b)/b \mid b \sqsubset a][\mathbf{tr}_\Sigma^a(b)^+/b \mid a \sqsubseteq b] && \text{(Fact 3.4(2))} \end{aligned}$$

where the last step uses the fact that none of the formulas $\mathbf{tr}_\Sigma^a(b)^+$ contain any free variables among A . It now suffices to prove, for all $d \sqsubset a$, that

$$\mathbf{tr}_{(\mathbb{A}\downarrow a)^-}(d)[\mathbf{tr}_\Sigma^a(b)^+/b \mid a \sqsubseteq b] \equiv_K \mathbf{tr}_\Sigma^a(d)^+.$$

This is proved by downwards induction on the priority of $d \sqsubset a$, using the fact that $\delta_\tau(d) = \top$ for all $d \sqsubset a$, using precisely the same argument as in the proof of the first claim. QED

References

- [1] A. Arnold and D. Niwiński. *Rudiments of μ -calculus*, volume 146 of *Studies in Logic and the Foundations of Mathematics*. North-Holland Publishing Co., Amsterdam, 2001.
- [2] B. ten Cate and G. Fontaine. An easy completeness proof for the modal μ -calculus on finite trees. In C.-H. Luke Ong, editor, *Proceedings of the 13th International Conference on Foundations of Software Science and Computational Structures, FOSSACS 2010*, volume 6014 of *LNCS*, pages 161–175. Springer, 2010.
- [3] J.W. de Bakker and D. Scott. A theory of programs (unpublished manuscript). IBM seminar, Vienna, 1969.
- [4] J. van Benthem. *Modal Correspondence Theory*. PhD thesis, Mathematisch Instituut & Instituut voor Grondslagenonderzoek, University of Amsterdam, 1976.
- [5] D. Berwanger. Game logic is strong enough for parity games. *Studia Logica*, 75(2):205–219, 2003.
- [6] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Number 53 in Cambridge Tracts in Theoretical Computer Science. Cambridge University Press, 2001.
- [7] A. Blass. A game semantics for linear logic. *Annals of Pure and Applied Logic*, 56(1-3):183–220, 1992.
- [8] J. Bradfield and C. Stirling. Modal μ -calculi. In J. van Benthem, P. Blackburn, and F. Wolter, editors, *Handbook of Modal Logic*, pages 721–756. Elsevier, 2006.
- [9] J.R. Büchi. Weak second-order arithmetic and finite automata. *Mathematical Logic Quarterly*, 6(1-6):66–92, 1960.
- [10] J.R. Büchi and L.H. Landweber. Solving sequential conditions by finite state strategies. *Transactions of the American Mathematical Society*, 138:295–311, 1969.
- [11] C. Cirstea, A. Kurz, D. Pattinson, L. Schröder, and Y. Venema. Modal logics are coalgebraic. *The Computer Journal*, 54:524–538, 2011.
- [12] C. Cirstea and D. Pattinson. Modular construction of modal logics. In Ph. Gardner and N. Yoshida, editors, *Proceedings of the 15th International Conference on Concurrency Theory (CONCUR 2004)*, volume 3170 of *LNCS*, pages 258–275. Springer, 2004.
- [13] E.M. Clarke and E.A. Emerson. Design and synthesis of synchronisation skeletons using branching time temporal logic. In D. Kozen, editor, *Logics of Programs*, pages 52–71. Springer, 1981.
- [14] G. D’Agostino and M. Hollenberg. Logical questions concerning the μ -calculus. *Journal of Symbolic Logic*, 65:310–332, 2000.
- [15] M. Dam. CTL* and ECTL* as fragments of the modal μ -calculus. *Theoretical Computer Science*, 126(1):77–96, 1994.
- [16] A. Dawar, E. Grädel, and S. Kreutzer. Inflationary fixed points in modal logic. *ACM Transactions on Computational Logic (TOCL)*, 5(2):282–315, 2004.
- [17] A. Doumane, D. Baelde, L. Hirschi, and A. Saurin. Towards completeness via proof search in the linear time μ -calculus. In *Proceedings of the 31st Annual ACM/IEEE Symposium on Logic In Computer Science (LICS’16)*. ACM Press, 2016. To appear.
- [18] E. Emerson and J.Y. Halpern. Decision procedures and expressiveness in the temporal logic of branching time. In *Proceedings of 14th ACM Symposium on Theory of Computing*, pages 169–180. 1982.

- [19] E.A. Emerson and J. Halpern. sometimes and not never revisited: on branching versus linear time temporal logic. *Journal of the ACM*, 33(1):151–178, 1986.
- [20] E.A. Emerson and C.S. Jutla. Tree automata, mu-calculus and determinacy (extended abstract). In *Proceedings of the 32nd Symposium on the Foundations of Computer Science*, pages 368–377. IEEE Computer Society Press, 1991.
- [21] S. Enqvist, F. Seifan, and Y. Venema. Completeness for coalgebraic fixpoint logics. Unpublished manuscript, 2016.
- [22] G. Fontaine, R. Leal, and Y. Venema. Automata for coalgebras: An approach using predicate liftings. In *Automata, Languages and Programming: 37th International Colloquium ICALP’10*, pages 381–392, 2010.
- [23] D.M. Gabbay, A. Pnueli, S. Shelah, and J. Stavi. On the temporal analysis of fairness. In *Proc. 7th ACM Symposium on Principles of Programming Languages*, pages 163–173, 1980.
- [24] S. Ghilardi. An algebraic theory of normal forms. *Annals of Pure and Applied Logic*, 71(3):189–245, 1995.
- [25] R. Goldblatt and I. Hodkinson. Spatial logic of modal mu-calculus and tangled closure operators. *arXiv preprint arXiv:1603.01766*, 2016.
- [26] R.I. Goldblatt. Metamathematics of modal logic I. *Reports on Mathematical Logic*, 6:41–78, 1976.
- [27] E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logic, and Infinite Games*, volume 2500 of *LNCS*. Springer, 2002.
- [28] E. Grädel and I. Walukiewicz. Guarded fixed point logic. In *Logic in Computer Science, 1999. Proceedings. 14th Symposium on*, pages 45–54. IEEE, 1999.
- [29] G. Jäger, M. Kretz, and T. Studer. Canonical completeness of infinitary mu. *Journal of Logic and Algebraic Programming*, 76(2):270–292, 2008.
- [30] D. Janin and I. Walukiewicz. Automata for the modal μ -calculus and related results. In J. Wiedermann and P. Hájek, editors, *Mathematical Foundations of Computer Science 1995, 20th International Symposium (MFCS’95)*, volume 969 of *LNCS*, pages 552–562. Springer, 1995.
- [31] D. Janin and I. Walukiewicz. On the expressive completeness of the propositional μ -calculus w.r.t. monadic second-order logic. In *Proceedings CONCUR ’96*, 1996.
- [32] B. Jónsson and A. Tarski. Boolean algebras with operators, Part I. *American Journal of Mathematics*, 73:891–939, 1952.
- [33] A. Joyal. Free lattices, communication and money games. In *Proceedings of the 10th International Congress of Logic, Methodology, and Philosophy of Science*, 1995.
- [34] R. Kaivola. Axiomatizing linear time mu-calculus. In I. Lee and S.A. Smolka, editors, *Proceedings of the 6th International Conference on Concurrency Theory (CONCUR ’95)*, volume 962 of *LNCS*, pages 423–437. Springer, 1995.
- [35] R. Kaivola. Axiomatizing extended computation tree logic. *Theoretical Computer Science*, 190(1):41–60, 1998.
- [36] D. Kozen. Results on the propositional μ -calculus. *Theoretical Computer Science*, 27:333–354, 1983.
- [37] D. Kozen. A finite model theorem for the propositional μ -calculus. *Studia Logica*, 47:233–241, 1988.

- [38] D. Kozen and R. Parikh. An elementary proof of the completeness of PDL. *Theoretical Computer Science*, 14:113–118, 1981.
- [39] C. Kupke and Y. Venema. Coalgebraic automata theory: basic results. *Logical Methods in Computer Science*, 4, 2008. (An earlier version of this paper appeared in LICS 2005.).
- [40] M. Lange and C. Stirling. Focus games for satisfiability and completeness of temporal logic. In *Proceedings of the 16th International Conference on Logic in Computer Science (LICS 2001)*, pages 357–365. IEEE Computer Society, 2001.
- [41] J. Marti, F. Seifan, and Y. Venema. Uniform interpolation for coalgebraic fixpoint logic. *arXiv preprint arXiv:1503.02319*, 2015.
- [42] M. Mio. Probabilistic modal μ -calculus with independent product. In *International Conference on Foundations of Software Science and Computational Structures*, pages 290–304. Springer, 2011.
- [43] L. Moss. Coalgebraic logic. *Annals of Pure and Applied Logic*, 96:277–317, 1999. (Erratum published *APAL* 99:241–259, 1999).
- [44] A.M. Mostowski. Games with forbidden positions. Technical Report 78, Instytut Matematyki, Uniwersytet Gdański, Poland, 1991.
- [45] A.W. Mostowski. Regular expressions for infinite trees and a standard form of automata. In A. Skowron, editor, *Computation Theory*, LNCS, pages 157–168. Springer-Verlag, 1984.
- [46] D.E. Muller and P.E. Schupp. Alternating automata on infinite trees. *tcs*, 54:267–276, 1987.
- [47] D.E. Muller and P.E. Schupp. Simulating alternating tree automata by nondeterministic automata. *Theoretical Computer Science*, 141:69–107, 1995.
- [48] R. Parikh. The logic of games and its applications. *Annals of Discrete Mathematics*, 24:111–139, 1985.
- [49] D.P.R. Park. Fixpoint induction and proof of program semantics. *Machine Intelligence*, 5:59–78, 1970.
- [50] D. Pattinson. Coalgebraic modal logic: soundness, completeness and decidability of local consequence. *Theoretical Computer Science*, 309(1–3):177–193, 2003.
- [51] A.M. Pitts. On an interpretation of second-order quantification in first-order intuitionistic propositional logic. *Journal of Symbolic Logic*, 57(1):33–52, 1992.
- [52] A. Pnueli. The temporal logic of programs. In *Proc. 18th Symp. Foundations of Computer Science*, pages 46–57, 1977.
- [53] V.R. Pratt. Semantical considerations on Floyd-Hoare logic. In *Proc. 17th IEEE Symposium on Computer Science*, pages 109–121, 1976.
- [54] V.R. Pratt. A decidable μ -calculus: preliminary report. In *Proc. 22th IEEE Symposium on Foundations of Computer Science*, pages 421–427, 1981.
- [55] M.O. Rabin. Decidability of second-order theories and automata on infinite trees. *Transactions of the American Mathematical Society*, 141:1–35, 1969.
- [56] M.O. Rabin and D. Scott. Finite automata and their decision problems. *IBM Journal of Research and Development*, 3:115–125, 1959.
- [57] M. Reynolds. An axiomatization of full computation tree logic. *Journal of Symbolic Logic*, pages 1011–1057, 2001.

- [58] J. Rutten. Universal coalgebra: a theory of systems. *Theoretical Computer Science*, 249:3–80, 2000.
- [59] S. Safra. On the complexity of ω -automata. In *Proceedings of the 29th Symposium on the Foundations of Computer Science*, pages 319–327. IEEE Computer Society Press, 1988.
- [60] L. Santocanale. Free μ -lattices. *Journal of Pure and Applied Algebra*, 168(2-3):227–264, 2002.
- [61] L. Santocanale. Completions of μ -algebras. *Annals of Pure and Applied Logic*, 154(1):27–50, 2008.
- [62] L. Santocanale and Y. Venema. Completeness for flat modal fixpoint logics. *Annals of Pure and Applied Logic*, 162:55–82, 2010.
- [63] U. Sattler and M.Y. Vardi. The hybrid μ -calculus. In *International Joint Conference on Automated Reasoning*, pages 76–91. Springer, 2001.
- [64] L. Schröder and D. Pattinson. Pspace bounds for rank-1 modal logics. *ACM Transactions on Computational Logic (TOCL)*, 10(2):13, 2009.
- [65] L. Schröder and D. Pattinson. Rank-1 modal logics are coalgebraic. *Journal of Logic and Computation*, 20(5):1113–1147, 2010.
- [66] E. Spaan. *Complexity of Modal Logics*. PhD thesis, ILLC, University of Amsterdam, 1993.
- [67] K. Tamura. A small model theorem for the hybrid μ -calculus. *Journal of Logic and Computation*, page exs052, 2013.
- [68] M.Y. Vardi. Reasoning about the past with two-way automata. In *International Colloquium on Automata, Languages, and Programming*, pages 628–641. Springer, 1998.
- [69] Yde Venema. Lectures on the modal μ -calculus. Lecture Notes, ILLC, University of Amsterdam, 2012.
- [70] A. Visser. Uniform interpolation and layered bisimulation. In *Gödel '96, (Brno)*, volume 6 of *Lecture Notes in Logic*, pages 139–164, 1996.
- [71] I. Walukiewicz. On completeness of the mu-calculus. In *Proceedings of the Eighth Annual Symposium on Logic in Computer Science (LICS '93)*, pages 136–146. IEEE Computer Society, 1993.
- [72] I. Walukiewicz. Completeness of Kozen’s axiomatisation of the propositional μ -calculus. In *Proceedings of the 10th Annual IEEE Symposium on Logic in Computer Science (LICS'95)*, pages 14–24. IEEE Computer Society Press, 1995.
- [73] I. Walukiewicz. Completeness of Kozen’s axiomatisation of the propositional μ -calculus. *Information and Computation*, 157:142–182, 2000.
- [74] T. Wilke. Alternating tree automata, parity games, and modal μ -calculus. *Bulletin of the Belgian Mathematical Society*, 8:359–391, 2001.