Theoretical Computer Science ••• (••••) •••–•••

Contents lists available at ScienceDirect

Theoretical Computer Science

www.elsevier.com/locate/tcs

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

Sebastian Enqvist, Fatemeh Seifan [1], Yde Venema *

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

A R T I C L E   I N F O

Article history:
Received 27 October 2016
Received in revised form 6 February 2018
Accepted 1 March 2018
Available online xxxx
Communicated by D. Sannella

Keywords:
Modal mu-calculus
Modal fixpoint logic
Completeness
Parity automata
Infinite games

A B S T R A C T

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, 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 Walukiewicz. In this paper we develop a framework that will let us clarify and simplify parts of Walukiewicz' proof.

Our main contribution is to take the automata-theoretic viewpoint, already implicit in Walukiewicz' 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.

* Corresponding author.
E-mail addresses: sebastian.enqvist@fil.lu.se (S. Enqvist), fateme.sayfan@gmail.com (F. Seifan), y.venema@uva.nl (Y. Venema).

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

https://doi.org/10.1016/j.tcs.2018.03.001
0304-3975/© 2018 Published by Elsevier B.V.

1. Introduction

The modal μ-calculus The modal μ-calculus $\mu\mathtt{ML}$ is an extension of basic modal logic with fixpoint operators. In its simplest form, we can define its formulas via the following grammar[2]

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

where p is a proposition letter, and μ is a *least fixpoint operator*, indicating that $\mu p\, \varphi$ is to be interpreted as the least 'solution' of the 'equation' $p \equiv \varphi(p)$. In the shape known today, the modal μ-calculus was introduced by Kozen [34], building on earlier work by others, including de Bakker & Scott [3], Park [46], and Pratt [51]. Since then it has been the subject of extensive research, see [2,27,7] 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 [49,24], computation tree logic CTL [11] and its extension CTL* [16] can all be embedded into the μ-calculus (see [13] for the non-trivial case of CTL*), as can propositional dynamic logic PDL [50] and Parikh's dynamic logic of games GL [45]. 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 [12], which shows that the μ-calculus enjoys several nice model-theoretic properties such as a Lyndon theorem and a Łos–Tarski theorem; perhaps most strikingly, they prove that, just like intuitionistic logic [48] and basic modal logic [25,66], the modal μ-calculus has the *uniform* interpolation property.

The modal μ-calculus is a natural extension of basic modal logic, and retains many of its good properties. For example, Kozen [35] showed that the finite model property is preserved and, improving on a series of earlier results, Emerson & Jutla [17] 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 [60]. 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 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, model theory, 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 [34], 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:

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

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:

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

The problem of proving the completeness of this axiom system turned out to be rather hard: Kozen presented a completeness proof, but only for a fragment of the μ-calculus, which he called the *aconjunctive* fragment. The completeness problem for the full language remained open for more than a decade. After proving the completeness of a different axiomatization, where the induction rule was replaced with a somewhat less elegant derivation rule [67], Walukiewicz finally provided a positive solution to the completeness question of Kozen's system for the full language [68] — in the sequel we shall refer to the extended journal publication [69].

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.*

[2] There is a syntactic proviso that the formation of the formula $\mu x.\varphi$ is subject to the constraint that the variable x is *positive* in φ.

Please cite this article in press as: S. Enqvist et al., Completeness for the modal μ-calculus: Separating the combinatorics from the dynamics, Theoret. Comput. Sci. (2018), https://doi.org/10.1016/j.tcs.2018.03.001

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.

This is not for lack of positive results. Indeed, complete derivation systems have been found for many modal fixpoints logics; important examples include the results of Kozen & Parikh on $\mathtt{PDL}$ [36], of Emerson & Halpern for $\mathtt{CTL}$ [15], of Reynolds for $\mathtt{CTL}^{\star}$ [53], and of Kaivola for the linear time μ-calculus [33], that is, the modal μ-calculus interpreted on the structure of the natural numbers with the successor relation.

However, the mentioned results are all proved on a case-by-case basis, and almost no uniform results are known. Exceptions are perhaps Lange & Stirling's game-theoretic framework of *focus games* [37], which applies to $\mathtt{LTL}$ and $\mathtt{CTL}$ alike. Also, some general methods for proving completeness are available, such as the technique of *filtration* [5], or the algebraic techniques used by Santocanale & Venema to prove a general result for "flat" fixpoint logics [56]. But then, these approaches only seem to apply to relatively simple fragments of the modal μ-calculus; for the algebraic approach, this is the message of Santocanale's [55]. And for filtration: it was shown by Kozen [35] that this technique does not apply to the μ-calculus.

The picture that seems to emerge from the literature — see, in particular, the work of Niwiński & Walukiewicz [44] on tableau games for the modal μ-calculus — is that the source of the difficulties in proving completeness for the modal μ-calculus and other fixpoint logics, lies in the combinatorial issues involved in dealing with infinite *traces*, i.e., possible histories of formulas, recording unfoldings of fixpoint variables. In the presence of simultaneous and possibly alternating fixpoints such combinatorics can get quite intricate.

In passing we note that we focus on finitary proof systems here. Allowing infinitary proof systems (or finitary systems derived from these by an appeal to the small model property of the modal μ-calculus), one may obtain completeness for the full language of $\mu\mathtt{ML}$ in more direct ways than by the Kozen–Walukiewicz proof, see for instance Kozen [35] or Jäger, Kretz & Studer [29].

Our goal 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, where we put traces and their combinatorics in the foreground, 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 (we will say more about this in our concluding section 11).. In the remainder of this introduction we outline some of the main themes of the paper.

Logic and Automata A mathematical framework for the modal μ-calculus that is tailor-suited to deal with the combinatorics of traces is the theory of finite automata and infinite games. This places the μ-calculus in a long tradition connecting logic and automata theory, going back to the seminal work of Büchi [8], Rabin [52] and others. 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 [41] and Emerson & Jutla [17]. And indeed, most of the deep results on the modal μ-calculus have used parity automata in one way or another; in particular, Walukiewicz' completeness proof, which builds on Kozen's result for aconjunctive formulas, heavily uses automata-theoretic ideas and insights. His proof starts from the observation that the satisfaction problem is easy for so-called *disjunctive* formulas, which correspond to the disjunctive automata introduced by Janin & Walukiewicz [30] (under the name "μ-automata"). Walukiewicz' strategy then is to prove that every μ-calculus formula φ can be rewritten as a semantically equivalent disjunctive formula $\widehat{\varphi}$, such that the implication $\varphi \to \widehat{\varphi}$ is *provable* in Kozen's system.

Roughly speaking, we follow the approach taken by Kozen and Walukiewicz, but with some differences, of which we mention three. First of all, we work with a wider class of automata than the disjunctive ones, and formulate a precise connection between the proof theory for the modal μ-calculus and the theory of these so-called modal automata. This enables us to rework large parts of Kozens's and Walukiewicz' arguments in an entirely automata-theoretic framework. As a second difference, we set up a framework where we may clearly distinguish what we take to be the two main orthogonal aspects of the completeness proof, viz., the *combinatorics* involved in reasoning with fixpoints, and the *dynamics* encoded in the semantics of the modal operators. And third, our approach is thoroughly game-theoretic in nature.

Modal automata The disjunctive automata that Walukiewicz is working with are rather special, viz., nondeterministic, in nature. It is much easier to establish the correspondence with μ-formulas for Wilke's more general class of *alternating* automata [70] that we shall simply call *modal automata*. Seen from this perspective, Walukiewicz' proof does three things at once: not only does it connect formulas to equivalent automata, it implicitly includes a simulation theorem stating that every modal/alternating automaton can be simulated by an equivalent disjunctive/nondeterministic one, and on top of that it establishes both equivalences as *proof-theoretic* result.

Working with arbitrary modal automata, we can link formulas and automata 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. We shall provide a recursive construction providing a modal automaton $\mathbb{A}_{\varphi}$ for each formula φ, and a translation in the converse direction providing a formula $\mathtt{tr}(\mathbb{A})$ for each modal automaton $\mathbb{A}$. We then prove the following proposition (with $\equiv_K$ denoting provable equivalence with respect to Kozen's axiomatization):

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

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 takes us "half-way" towards Walukiewicz' result, where the remainder of the distance can now be addressed by wholly automata-theoretic methods.

Dynamics versus combinatorics Our automata-theoretic approach allows for an explicit study of the interaction between the combinatorial and dynamical aspects of the formalism. We will deal with the combinatorics in a graph-theoretical framework, well-known from the theory of automata operating on infinite words [27], where a trace can simply be taken as a path through a sequence of binary relations. The dynamics is studied via the concept of a *one-step logic*, which stems from the work on coalgebraic logic by Cîrstea, Pattinson, Schröder and others [10,47,58,59]. One-step logics are extremely simple logical formalisms that feature as the codomain of the transition function of automata. Their importance lies in the observation that many results on modal fixpoint logics, often involving nontrivial automata-theoretic phenomena, can in fact be understood at the basic level of one-step logic [23,21].

On the one hand, the "deconstruction" into dynamics versus combinatorics allows us to deal with the combinatorial and the dynamic concepts in largely separate frameworks. On the other hand, the use of modal 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. For example one of our key concepts, that of a *semi-disjunctive* automaton, is defined in terms of a syntactic form of the one-step formulas, but is *motivated* by certain results about the structure on traces for such automata.

Games Most of our results will be formulated in terms of two infinite two-player games that may be associated with modal automata. The *satisfiability game*, which has previously been studied in a wider coalgebraic context [64,22], is our automata-theoretic version of Niwiński & Walukiewicz' tableau game [44]. It is *adequate* with respect to satisfiability in the sense that one of the players has a winning strategy in the game $\mathcal{S}(\mathbb{A})$ (associated with the modal automaton $\mathbb{A}$) iff the automaton $\mathbb{A}$ is satisfiable, that is, has a non-empty language. The *consequence game* is inspired by Walukiewicz' notion of tableau consequence [69]. Given two automata $\mathbb{A}$ and $\mathbb{A}'$, the game $\mathcal{C}(\mathbb{A}, \mathbb{A}')$ is about establishing a structural connection, denoted as $\mathbb{A} \models_{\mathsf{G}} \mathbb{A}'$, that is much stronger than $\mathbb{A}'$ just being a semantic consequence of $\mathbb{A}$.

We make a number of observations about these games in this paper. Perhaps the most interesting new result is a strengthened version of the *simulation theorem* for modal automata. We show that the semantic equivalence between a modal automaton $\mathbb{A}$ and its disjunctive simulation $\mathsf{sim}(\mathbb{A})$ can be upgraded to the level of the game consequence relation:

$$\mathbb{A} \models_{\mathsf{G}} \mathsf{sim}(\mathbb{A}) \quad \text{and} \quad \mathsf{sim}(\mathbb{A}) \models_{\mathsf{G}} \mathbb{A}.$$

In fact, this result still holds if we only simulate a "part" of a modal automaton; formulated precisely, for any modal automaton of the form $\mathbb{B}[\mathbb{A}/x]$ we show that:

$$\mathbb{B}[\mathbb{A}/x] \models_{\mathsf{G}} \mathbb{B}[\mathsf{sim}(\mathbb{A})/x], \text{ and } \mathbb{B}[\mathsf{sim}(\mathbb{A})/x] \models_{\mathsf{G}} \mathbb{B}[\mathbb{A}/x]$$

where the notion of substitution for automata ($[\mathbb{A}/x]$) will be formally introduced below.

Related work During the refereeing process of this paper, two publications presenting related work saw the light. Afshari & Leigh [1] presented a breakthrough result, proving completeness of a new cut free and finitary proof system for the modal μ-calculus. Another interesting contribution was made by Doumane [14], who gave a new completeness proof for the *linear-time* μ-calculus. In the concluding section of this paper we will say a bit more about the connection between our work and these results.

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 a translation from $\mu\mathtt{ML}$-formulas to modal automata which is based on automata-theoretic operations that correspond to syntactic operators of $\mu\mathtt{ML}$. In section 5 we define the satisfiability game and the consequence game. Section 6 is pivotal to our paper: here we introduce *disjunctive* and *semi-disjunctive* automata, and we prove a key result relating the consequence game and the satisfiability game (Theorem 3). Section 7 is devoted to the proof of our previously mentioned "strong simulation theorem" for modal automata, Theorem 4. In section 8, which can be read almost completely independently of the sections 5–7, we provide the translation back from automata to formulas, and we prove that our translations from formulas to automata and back interact well with the proof theory of Kozen's proof system (Theorem 2). In section 9 we focus on the proof of what we call *Kozen's Lemma*, Theorem 5. We wrap things up in section 10, where we prove our main lemma, Theorem 6, and we show how to derive the Kozen–Walukiewicz result, Theorem 1, from this. We end the paper with a short concluding section, where we mention some directions for future research. 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 in an appendix to an earlier report version of this article [19].

To help the reader navigate through this paper, we show the dependency of the sections in the following dependency graph, where the arrows represent the order in which individual sections may be read:

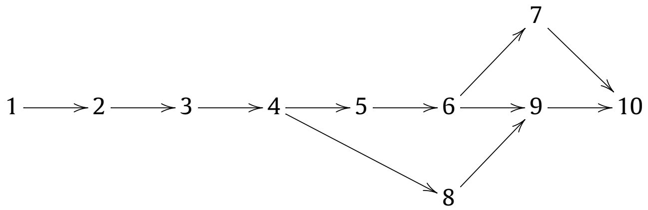

In addition, at the end of the paper we provide a table with some of the most important notations used in this paper.

2. Preliminaries

In this section we fix some notation and terminology.

2.1. Basic mathematical concepts and notation

Definition 2.1. Let A be some set. We denote its size as $|A|$, and its power set as $\mathsf{P}A$. $\lhd$

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. $\lhd$

Definition 2.3. Given a relation $R \subseteq A \times A'$, we define the following relations between $\mathsf{P}A$ and $\mathsf{P}A'$:

$$\overrightarrow{\mathsf{P}}\, R := \{(B, B') \in \mathsf{P}A \times \mathsf{P}A' \mid \text{for all } b \in B \text{ there is a } b' \in B' \text{ with } Rbb'\}$$
$$\overleftarrow{\mathsf{P}}\, R := \{(B, B') \in \mathsf{P}A \times \mathsf{P}A' \mid \text{for all } b' \in B' \text{ there is a } b \in B \text{ with } Rbb'\}$$
$$\overline{\mathsf{P}}R := \overrightarrow{\mathsf{P}}\, R \cap \overleftarrow{\mathsf{P}}\, R.$$

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

Definition 2.4. Given a function $f : A \to B$, we denote the *graph* of f as $\mathsf{Gr}f := \{(a, fa) \mid a \in A\}$. The composition of functions $f : A \to B$ and $g : B \to C$ is denoted as $g \circ f : A \to C$. $\lhd$

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 $\mathsf{last}(w)$.

The j-th word in a stream σ is denoted by $\sigma(j)$. 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 $j \geq k$. $\lhd$

2.2. Infinite games

For an introduction to infinite (parity) games and their use in the theory of automata operating on infinite objects, we refer the reader to [27]. Here we list the notions that are used in the paper.

Definition 2.6. A *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* or *arena* 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 priority map $\Omega : G \to \omega$ 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 elements of $E[p]$ are called the *legitimate* or *admissible moves* of player Π at p. In case $E[p] = \varnothing$ we say that player Π *gets stuck* at p.

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

Definition 2.7. A *play* 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 play ρ is called *partial* if it is finite and $E[\mathsf{last}\rho] \neq \varnothing$, and *full* otherwise. We let PM_Π denote the collection of partial plays ρ ending in a position $\mathsf{last}(\rho) \in G_\Pi$, and define $\mathrm{PM}_\Pi@p$ as the set of partial plays in PM_Π starting at position p.

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

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

A play $\rho = (p_i)_{i<\kappa}$ 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 $p_0 \dots p_{n-1} \in \mathrm{PM}_\Pi$ (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 plays in $\mathrm{PM}_\Pi@p$ are always admissible to Π, and *winning for* Π in $\mathbb{G}@p$ if in addition all χ-guided full plays 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 $\forall$. ◁

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

The following fact, independently due to Emerson & Jutla [17] and Mostowski [40], 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 that the reader is familiar with the syntax and semantics of the modal μ-calculus, here we provide a brief recapitulation of the main notions that play a role in this paper. More detail on the formalism can be found in [27, 65].

3.1. Syntax

Throughout this paper we fix an (unnamed) infinite set of propositional variables. In the introduction we formulated a "minimal" grammar for the language of the modal μ-calculus. As usual in logic, it is quite standard to use various abbreviations on top of this basic language, including the symbols $\top$, $\bot$, $\wedge$, $\to$, $\leftrightarrow$, $\Box$, $\bigwedge$, and $\bigvee$ (where the latter two symbols are used to denote arbitrary but finite conjunctions and disjunctions, respectively), and the greatest fixpoint operator ν. For our purposes it will be convenient to introduce $\wedge$ and $\Box$ as *primitive connectives*, and similarly, ν as a *primitive fixpoint operator*. That is, throughout the paper we will work with the following extended language.

Definition 3.1. The language $\mu\mathtt{ML}$ of the modal μ-calculus is given by the following grammar:

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

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. Elements of this language will be called *modal fixpoint formulas*, *μ-formulas*, or simply *formulas*. We use the symbol η to range over μ and ν.

The collection of *subformulas* of a given formula is defined as usual, as are the sets of, respectively, its *free* and *bound* variables — with the understanding that in the formula $\eta x \varphi$, ηx binds all free occurrences of x in the formula φ. We let $\mu\mathtt{ML}(\mathtt{Prop})$ denote the set of μ-formulas of which all free variables belong to the set $\mathtt{Prop}$. ◁

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 pair of primitive modalities $\Diamond$ and $\Box$. The completeness proof for the polymodal μ-calculus, where one has a family $\{\Diamond_d, \Box_d \mid d \in D\}$ of modal diamonds and boxes, can be obtained by a straightforward adaptation of the monomodal case, and follows from our more general result in [18]. ◁

Please cite this article in press as: S. Enqvist et al., Completeness for the modal μ-calculus: Separating the combinatorics from the dynamics, Theoret. Comput. Sci. (2018), https://doi.org/10.1016/j.tcs.2018.03.001

As a convention, the free variables of a formula φ are denoted by the symbols $p, q, r, \ldots$, and referred to as *proposition letters*, while we use the symbols $x, y, z, \ldots$ for the bound variables of a formula. (Note however, that the bound variables of a formula will generally have free occurrences in some of its subformulas, so this convention cannot be a strict rule.)

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 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 the context. If $Z = Y_1 \uplus Y_2$, with $\uplus$ denoting the disjoint union, 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 $\varphi[\psi_z/z \mid z \in Z]$. $\lhd$

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

(1) $\varphi[\psi_y/y \mid y \in \mathtt{Y}][\chi_z/z \mid z \in \mathtt{Z}] = \varphi\big[\psi_y[\chi_z/z \mid z \in \mathtt{Z}]/y \mid y \in \mathtt{Y}, \chi_z/z \mid z \in \mathtt{Z}\big]$
(2) $\varphi[\psi_y/y \mid y \in \mathtt{Y}][\chi_z/z \mid z \in \mathtt{Z}] = \varphi\big[\psi_y/y \mid y \in \mathtt{Y}, \chi_z/z \mid z \in \mathtt{Z}\big]$, *provided no $z \in \mathtt{Z}$ 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:

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

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. $\lhd$

3.2. Structures

Definition 3.6. Given a set S, an A-*marking* on S is a map $m : S \to \mathsf{P}A$; an A-*valuation* on S is a map $V : A \to \mathsf{P}S$. Any valuation $V : A \to \mathsf{P}S$ gives rise to its *transpose marking* $V^\dagger : S \to \mathsf{P}A$ 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. $\lhd$

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 $\mathtt{Prop}$ 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 $\mathtt{Prop}$-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 $\mathtt{Prop} \cup \{x\}$-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 $(S, R, V[x \mapsto U])$. $\lhd$

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

$$\mathsf{K}_{\mathtt{Prop}}S := \mathsf{P}\mathtt{Prop} \times \mathsf{P}S,$$

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

A Kripke structure $\mathbb{S} = (S, R, V)$ over the set $\mathtt{Prop}$ can then be represented as a map $\sigma_{\mathbb{S}} : S \to \mathsf{K}S$ given by

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

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

Please cite this article in press as: S. Enqvist et al., Completeness for the modal μ-calculus: Separating the combinatorics from the dynamics, Theoret. Comput. Sci. (2018), https://doi.org/10.1016/j.tcs.2018.03.001

The operation 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 $f : S \to S'$ gives rise to a map $\mathsf{K}f$ from $\mathsf{K}S$ to $\mathsf{K}S'$, defined by

$$\mathsf{K}f : (\Upsilon, U) \mapsto (\Upsilon, f[U]).$$

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

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

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 $[\![\cdot]\!]$, which assigns to a formula φ its *meaning* $[\![\varphi]\!]^{\mathbb{S}} \subseteq S$ in any Kripke structure $\mathbb{S} = (S, R, V)$. The clauses of this definition are standard:

$$\begin{aligned}
[\![p]\!]^{\mathbb{S}} &:= V(p) \\
[\![\neg\varphi]\!]^{\mathbb{S}} &:= S \setminus [\![\varphi]\!]^{\mathbb{S}} \\
[\![\varphi \vee \psi]\!]^{\mathbb{S}} &:= [\![\varphi]\!]^{\mathbb{S}} \cup [\![\psi]\!]^{\mathbb{S}} \\
[\![\varphi \wedge \psi]\!]^{\mathbb{S}} &:= [\![\varphi]\!]^{\mathbb{S}} \cap [\![\psi]\!]^{\mathbb{S}} \\
[\![\Diamond\varphi]\!]^{\mathbb{S}} &:= \{s \in S \mid R[s] \cap [\![\varphi]\!]^{\mathbb{S}} \neq \varnothing\} \\
[\![\Box\varphi]\!]^{\mathbb{S}} &:= \{s \in S \mid R[s] \subseteq [\![\varphi]\!]^{\mathbb{S}}\} \\
[\![\mu x \varphi]\!]^{\mathbb{S}} &:= \bigcap\{U \in \mathsf{P}S \mid [\![\varphi]\!]^{\mathbb{S}[x \mapsto U]} \subseteq U\}. \\
[\![\nu x \varphi]\!]^{\mathbb{S}} &:= \bigcup\{U \in \mathsf{P}S \mid U \subseteq [\![\varphi]\!]^{\mathbb{S}[x \mapsto U]}\}.
\end{aligned}$$

If a point $s \in S$ belongs to the set $[\![\varphi]\!]^{\mathbb{S}}$, 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*, denoted by: $\models \varphi$, if $[\![\varphi]\!]^{\mathbb{S}} = S$ for any structure $\mathbb{S} = (S, R, V)$, and *satisfiable* if $[\![\varphi]\!]^{\mathbb{S}} \neq \varnothing$ for some structure $\mathbb{S}$. Two formulas φ and ψ are *equivalent*, notation: $\varphi \equiv \psi$, if $[\![\varphi]\!]^{\mathbb{S}} = [\![\psi]\!]^{\mathbb{S}}$ for any structure $\mathbb{S}$. ◁

3.4. Axiomatics

We saw in the introduction that Kozen's axiomatization for the modal μ-calculus is obtained by adding the (pre-)fixpoint axiom and rule to the basic modal logic **K**. For completeness' sake we give the definition of Kozen's system here, taking a standard $\Diamond$-based axiomatization for **K** [5, Remark 4.7]. The price to pay for taking $\wedge$, $\Box$ and ν as primitive symbols (rather than as abbreviations) is that we need to add their definitions as axioms to the derivation system, in the form of the (Dual) axioms below. In the formulation below we use the connectives $\to$ and $\leftrightarrow$ as abbreviations (with their standard definitions).

Definition 3.11. The *axioms* of the basic modal logic **K** are the following:
(C) a complete set of axioms for classical propositional logic;
(NA) axioms stating that $\Diamond$ is *normal* ($\neg\Diamond\bot$) and *additive* ($\Diamond(p \vee q) \leftrightarrow \Diamond p \vee \Diamond q$),
(Dual$_\wedge$) $p \wedge q \leftrightarrow \neg(\neg p \vee \neg q)$
(Dual$_\Box$) $\Box p \leftrightarrow \neg\Diamond\neg p$
while its *derivation rules* are
(MP) modus ponens: from φ and $\varphi \to \psi$, derive ψ;
(Mon) monotonicity: from $\varphi \to \psi$ derive $\Diamond\varphi \to \Diamond\psi$;
(US) uniform substitution: from φ derive $\varphi[\sigma]$, for any substitution σ.
Kozen's deductive system is obtained from this by adding the following axiom schemes and rule to those of **K**:
(Dual$_\nu$) $\nu x \varphi \leftrightarrow \neg\mu x(\neg\varphi[\neg x/x])$,
(A$_\mu$) all prefixpoint axioms of the form $\varphi[\mu x \varphi/x] \to \mu x \varphi$;
(R$_\mu$) Park's prefixpoint rule: from $\varphi[\gamma/x] \to \gamma$ derive $\mu x \varphi \to \gamma$. ◁

Definition 3.12. A *derivation* in $\mathbf{K}\mu$ is a finite list of 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 $\mathbf{K}\mu$. 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 $\varphi \to \psi$ 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. ◁

The main goal of this paper will be to prove the completeness of the derivation system $\mathbf{K}\mu$ for the collection of valid $\mu\mathtt{ML}$-formulas.

Remark 3.13. Our formulation of the completeness theorem in the introduction (Theorem 1) was in terms of the basic language, based on $\{\vee, \Diamond, \mu\}$. Readers who worry about a possible mismatch with the goal just described, may rest assured.

Let $\mathbf{K}\mu^-$ be the derivation system for the basic language that we obtain from $\mathbf{K}\mu$ by dropping the three (Dual) axioms, i.e., $\mathbf{K}\mu^-$ is the derivation system we implicitly defined in the introduction. Then one may show that $\mathbf{K}\mu$ is a *conservative extension* of $\mathbf{K}\mu^-$, that is:

$$\varphi \text{ is derivable in } \mathbf{K}\mu^- \text{ iff } \varphi \text{ is derivable in } \mathbf{K}\mu. \tag{1}$$

for any formula φ in the basic language.

For a proof of this, the direction from left to right is immediate since $\mathbf{K}\mu$ is an extension of $\mathbf{K}\mu^-$. For the opposite direction we confine ourselves to a brief sketch. Define an inductive translation $(\cdot)^-$ mapping $\mu\mathtt{ML}$ to the basic language, based on the clauses $(\varphi \vee \psi)^- := \varphi^- \vee \psi^-$, $(\varphi \wedge \psi)^- := \neg(\neg\varphi^- \vee \neg\psi^-)$, etc. It is then straightforward to show, by induction on the length of a derivation, that an arbitrary formula $\varphi \in \mu\mathtt{ML}$ is derivable in $\mathbf{K}\mu$ only if its translation φ^- is derivable in the system $\mathbf{K}\mu^-$. The direction from right to left of (1) follows by the observation that $\varphi = \varphi^-$ for formulas in the basic language.

The point of (1) is that the completeness of $\mathbf{K}\mu$ for the extended language $\mu\mathtt{ML}$ implies the completeness of $\mathbf{K}\mu^-$ for the basic language. That is, the goal formulated above indeed suffices to prove Theorem 1 as formulated in the introduction. ◁

Finally, in the sequel we shall need the following facts on our proof system.

Fact 3.14. *Let φ and ψ be modal μ-formulas. Then*

(1) $\nu x \varphi \leq_K \varphi[\nu x \varphi/x]$;
(2) if $\psi \leq_K \varphi[\psi/x]$ then $\psi \leq_K \nu x \varphi$;
(3) $\eta x \eta y \varphi \equiv_K \eta y \eta x \varphi$;
(4) $\varphi \leq_K \psi$ implies $\eta x \varphi \leq_K \eta x \psi$;
(5) $\varphi \equiv_K \varphi'$ for some effectively obtainable formula φ' in negation normal form.

Proof. The first two items can be proved via a routine argument.

The easiest way to prove the statements (3) and (4) is to proceed algebraically — here we assume some basic familiarity with the theory of modal algebras (as in Chapter 5 of [5]). Let a *modal μ-algebra* be a modal algebra $\mathbb{A} = (A, \vee, \wedge, \neg, \bot, \Diamond, \Box)$ such that, relative to a given assignment $h : \mathtt{Prop} \to A$, every formula (or term) $\varphi \in \mu\mathtt{ML}$ finds an interpretation $[\![\varphi]\!]^h$, in such a way that for every formula $\mu x \psi$ the interpretation $[\![\mu x \psi]\!]^h$ is the least fixpoint of the map $a \mapsto [\![\psi]\!]^{h[x \mapsto a]}$ (and similarly for greatest fixpoint formulas $\nu x \psi$). It is then a routine exercise in algebraic logic to show that Kozen's axiomatization is sound and complete for the class of modal μ-algebras. (Note that this algebraic completeness does not imply the completeness of $\mathbf{K}\mu$ with respect to its Kripke semantics, since the latter corresponds to a rather special *subclass* of the modal μ-algebras.) We may then prove, for instance, item (3) by showing that the equation $\mu x \mu y \varphi = \mu y \mu x \varphi$ holds in every modal μ-algebra; for this one may use a standard argument, see for instance [2, Proposition 1.3.2]. Similarly, item (4) can be proved as in Proposition 1.2.18 of *ibid*.

The last item can be shown by a direct induction on the complexity of φ. □

3.5. The cover modality

The *cover modality* ∇ (also known as *nabla modality*), which was independently introduced in coalgebraic logic [39] 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 3.15. Given a finite set of formulas Φ, we let $\nabla\Phi$ abbreviate the formula

$$\nabla\Phi := \bigwedge \Diamond\Phi \wedge \Box \bigvee \Phi,$$

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

Remark 3.16. Observe that the semantics of the cover modality can be expressed in terms of the Egli-Milner lifting ($\overline{\mathsf{P}}\Vdash$) of the satisfaction relation $\Vdash$:

$$\mathbb{S}, s \Vdash \nabla\Phi \text{ iff } (R[s], \Phi) \in \overline{\mathsf{P}}\Vdash.$$

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:

$$\begin{aligned} \Diamond\varphi &\equiv \nabla\{\varphi, \top\} \\ \Box\varphi &\equiv \nabla\{\varphi\} \vee \nabla\varnothing, \end{aligned}$$

where we note that $\nabla\varnothing$ holds at a point s iff s is a 'blind' world, that is, $R[s] = \varnothing$. ◁

Proposition 3.17. *Let* Φ,Ψ *and* $\{\varphi_{i,0}, \varphi_{i,1} \mid i \in I\}$ *be finite sets of modal* μ*-formulas. Then*

(1) $\nabla\Phi \wedge \nabla\Psi \equiv_K \bigvee_{R \in \Phi\otimes\Psi} \nabla\{\varphi \wedge \psi \mid (\varphi, \psi) \in R\}$, *where* $\Phi \otimes \Psi = \{R \subseteq \Phi \times \Psi \mid (\Phi, \Psi) \in \overline{\mathsf{P}}R\}$*;*

(2) $\nabla\{\varphi_{i,0} \vee \varphi_{i,1} \mid i \in I\} \equiv_K \bigvee_{Z \in \mathcal{C}} \nabla\{\varphi_{i,j} \mid (i, j) \in Z\}$, *where* $\mathcal{C} = \{Z \subseteq I \times \{0, 1\} \mid \mathsf{Dom} Z = I\}$.

Proof. Again we confine ourselves to a sketch. It suffices to observe that the mentioned equivalences are semantically valid and do not involve the fixpoint operators. From this it follows that they are already provable in the basic modal logic **K**, of which Kozen's system is an extension. □

4. Modal automata and their one-step logic

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' proof works directly with what we will call disjunctive automata, we will work with the wider class of *modal automata* [70] that we will introduce in this section.

A goal in this section is to define a translation transforming a formula of the μ-calculus into an equivalent modal parity automaton. 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 [70]. It seems that the standard approach to this (see for instance [27], following Wilke), is to transform a $\mu\mathtt{ML}$-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 [10], 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 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, i.e. formulas 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 A, we define the set $\mathtt{Latt}(A)$ of *lattice terms*, or *positive propositional formulas*, over A through the following grammar:

Please cite this article in press as: S. Enqvist et al., Completeness for the modal μ-calculus: Separating the combinatorics from the dynamics, Theoret. Comput. Sci. (2018), https://doi.org/10.1016/j.tcs.2018.03.001

$$\pi ::= \bot \mid \top \mid a \mid \pi \wedge \pi \mid \pi \vee \pi,$$

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

$$\alpha ::= \bot \mid \top \mid p \mid \neg p \mid \Diamond\pi \mid \Box\pi \mid \alpha \wedge \alpha \mid \alpha \vee \alpha,$$

with $p \in \mathtt{Prop}$ and $\pi \in \mathtt{Latt}(A)$. $\lhd$

Note that elements from the two parameter sets, $\mathtt{Prop}$ and A, are treated quite differently in the syntax of one-step formulas: all occurrences of elements of $\mathtt{Prop}$, 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.

One-step formulas will be interpreted in *one-step models*.

Definition 4.2. Fix sets $\mathtt{Prop}$ and A. A *one-step frame* is a pair (Υ, S) where S is any set called the carrier of (Υ, S), and $\Upsilon \subseteq \mathtt{Prop}$. A *one-step model* is a triple (Υ, S, m) such that (Υ, S) is a one-step frame and m is an A-marking on S. $\lhd$

Observe that with this definition, the coalgebraic representation of a Kripke structure (S, R, V) can now be seen as a function $\sigma_{\mathbb{S}}$ 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 (Υ, S, m) and a one-step modal language $\mathtt{1ML}(\mathtt{Prop}, A)$. First, we define the value $[\![\pi]\!]$ of a lattice formula π over A by induction, setting $[\![a]\!] = \{s \in S \mid a \in m(s)\}$ 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:

- $(\Upsilon, S, m) \Vdash^1 \Box\pi$ iff $[\![\pi]\!] = S$,
- $(\Upsilon, S, m) \Vdash^1 \Diamond\pi$ iff $[\![\pi]\!] \neq \emptyset$,
- $(\Upsilon, S, m) \Vdash^1 p$ iff $p \in \Upsilon$,
- $(\Upsilon, S, m) \Vdash^1 \neg p$ iff $p \notin \Upsilon$.

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. $\lhd$

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 $\nabla B \wedge \nabla B' \equiv_1 \bigvee\{\nabla\{b \wedge b' \mid Rbb'\} \mid R \subseteq B \times B' \text{ and } (B, B') \in \overline{\mathsf{P}}R\}$ (cf. Fact 3.14(1)).

Readers who want to move on to the definition of modal automata can now proceed to the next subsection. In the remainder of this subsection we mention some further concepts and results related to one-step logic. 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* $\mathsf{P}A$ as the map $I_A : B \mapsto B$ (that is, the identity map on $\mathsf{P}A$). More generally, for any subset $\mathcal{B}$ of $\mathsf{P}A$, we consider the marking $I_A{\upharpoonright_{\mathcal{B}}} : \mathcal{B} \to \mathsf{P}A$.

For any one-step formula $\alpha \in \mathtt{1ML}(\mathtt{Prop}, A)$ and any element $\Gamma \in \mathsf{KP}A$, say $\Gamma = (\Upsilon, \mathcal{B})$, we abbreviate $(\Gamma, I_A{\upharpoonright_{\mathcal{B}}}) \Vdash^1 \alpha$ as $\Gamma \Vdash^1_I \alpha$, and we denote $[\![\alpha]\!]^1 := \{\Gamma \in \mathsf{KP}A \mid \Gamma \Vdash^1_I \alpha\}$. $\lhd$

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 (Υ, S, m) and (Υ', S', m') be one-step models with respect to A and $\mathtt{Prop}$. We say that these models are *one-step bisimilar* if they satisfy the following conditions:

(atomic) $\Upsilon = \Upsilon'$;
(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 $(\Upsilon, S, m) \leftrightarroweq^1 (\Upsilon', S', m')$ to say that (Υ, S, m) and (Υ', S', m') are one-step bisimilar. $\lhd$

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

Please cite this article in press as: S. Enqvist et al., Completeness for the modal μ-calculus: Separating the combinatorics from the dynamics, Theoret. Comput. Sci. (2018), https://doi.org/10.1016/j.tcs.2018.03.001

Proposition 4.6 *(One-step bisimulation invariance). Let $(\mathtt{Y}, S, m)$ and $(\mathtt{Y}', S', m')$ be any two one-step models with respect to A and $\mathtt{Prop}$. If $(\mathtt{Y}, S, m) \leftrightarrow^1 (\mathtt{Y}', S', m')$, then both one-step models satisfy the same formulas in $\mathtt{1ML}(\mathtt{Prop}, A)$.*

We consider a useful instance of the one-step bisimulation invariance theorem: pick any set W, let $\Gamma = (\mathtt{Y}, S)$ be any one-step frame in the set $\mathsf{K}W$, and let $f : W \to W'$. 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 $\alpha \in \mathtt{1ML}(\mathtt{Prop}, A)$. 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 $\mathtt{1ML}(\mathtt{Prop}, A)$.

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

(atomic) $\mathtt{Y} = \mathtt{Y}'$;
(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 $(\mathtt{Y}, S, m)$ and $(\mathtt{Y}', S', m')$ be any two one-step models with respect to A and $\mathtt{Prop}$. If $(\mathtt{Y}, S, m) \rightrightarrows^1 (\mathtt{Y}', S', m')$, then any formula in $\mathtt{1ML}(\mathtt{Prop}, A)$ that is satisfied by $(\mathtt{Y}, S, m)$ is also satisfied by $(\mathtt{Y}', S', m')$.*

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 μ-calculus. This is the *one-step completeness theorem*, stated below.

Theorem 4.9 *(One-step completeness). Let $\alpha \in \mathtt{1ML}(\mathtt{Prop}, A)$ be a one-step formula, and let $\sigma : A \to \mu\mathtt{ML}(\mathtt{Prop})$ be a substitution such that the formula $\alpha[\sigma]$ is consistent. Then there is a one-step model $(\mathtt{Y}, S, m)$ such that $(\mathtt{Y}, S, m) \Vdash^1 \alpha$ 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:

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

where γ is a conjunction of literals over $\mathtt{Prop}$, each $\heartsuit$-operator is a modality $\heartsuit_i \in \{\Diamond, \Box\}$, and each π_i is a lattice formula over A. Using the equivalences in Remark 3.16, we can rewrite α into a disjunction of formulas of the form

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

where each Π_i is a finite set of lattice formulas over A. Now apply Fact 3.14(1) repeatedly, and then distribute the conjunct γ over disjunctions, to obtain a disjunction of formulas of 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 3.14(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

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

where γ is a conjunction of literals over $\mathtt{Prop}$ and $B_1, \ldots, B_k$ 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

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

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

$$(\mathtt{Y}, S, m) \Vdash^1 \gamma \wedge \nabla\{\bigwedge B_1, \ldots, \bigwedge B_k\}$$

and it follows that $(\mathtt{Y}, S, m \Vdash^1 \alpha)$. 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 $\gamma \wedge \nabla\{\bigwedge \sigma[B_1], \ldots, \bigwedge \sigma[B_k]\}$. □

4.2. Modal automata

We now formally introduce modal automata.

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

$$\Theta : A \to 1\mathtt{ML}(\mathtt{Prop}, A)$$

maps states to one-step formulas. $\lhd$

Modal automata run on Kripke structures, and acceptance is defined in terms of a two-player game, the *acceptance game* $\mathcal{A}(\mathbb{A}, \mathbb{S})$ associated with an automaton $\mathbb{A}$ and a Kripke structure $\mathbb{S}$. The two players have opposing goals: one player $\exists$ (or "Eloise") defends the claim that $\mathbb{A}$ accepts $\mathbb{S}$, and the opposing player $\forall$ ("Abelard") wants to establish the opposite. A play of the game proceeds in *rounds*, moving from one basic position (i.e., of the form $(a, s) \in A \times S$) to another. Each round connects the one-step logic to the Kripke structure; the key observation here is that for any Kripke structure $\mathbb{S} = (S, R, V)$ over $\mathtt{Prop}$, any given state $s \in S$ gives rise to a one-step frame $(V^\dagger(s), R[s])$ consisting of the proposition letters in $\mathtt{Prop}$ that are true at s, together with the set of R-successors of s. At the basic position (a, s), it is $\exists$'s challenge to come up with a marking m for this one-step frame $(V^\dagger(s), R[s])$ such that, in the resulting one-step model, the formula $\Theta(a) \in 1\mathtt{ML}(\mathtt{Prop}, A)$ holds. The round of the play then closes with $\forall$ picking a next basic position from the set $\{(b, t) \mid b \in m(t)\}$. All of this is summarized in the table below.

Definition 4.11. Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be any modal $\mathtt{Prop}$-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
$(a, s) \in A \times S$	$\exists$	$\{m : R[s] \to \mathsf{P}A \mid (V^\dagger(s), R[s], m) \Vdash^1 \Theta(a)\}$
$m : R[s] \to \mathsf{P}A$	$\forall$	$\{(b, t) \mid b \in m(t)\}$

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

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

Definition 4.12. 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 $\lhd_{\mathbb{A}}$ denote the transitive closure of $E_{\mathbb{A}}$. If $a \lhd_{\mathbb{A}} b$ we say that a is *active* in b. We write $a \bowtie_{\mathbb{A}} b$ if $a \lhd_{\mathbb{A}} b$ and $b \lhd_{\mathbb{A}} a$.

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. $\lhd$

Definition 4.13. Fix a modal $\mathtt{Prop}$-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 call a a μ*-state*, writing $\eta_a = \mu$, if $\Omega(a)$ is odd, and a ν*-state*, writing $\eta_a = \nu$, if $\Omega(a)$ is even. We call η_a the *type* of a and denote the sets of μ- and ν-states as A^μ and A^ν, respectively.

We say that $\mathbb{A}$ is *positive* in a proposition letter $p \in \mathtt{Prop}$ if each occurrence of p in each formula $\Theta(a)$ with $a \in A$ is positive. $\lhd$

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 $L(\mathbb{A}) = L(\mathbb{A}')$. We also use the symbol $\equiv$ to denote the equivalence between formulas and automata, i.e., we write $\varphi \equiv \mathbb{A}$ if $\mathbb{A}$ accepts exactly the pointed Kripke structures where φ holds. $\lhd$

Please cite this article in press as: S. Enqvist et al., Completeness for the modal μ-calculus: Separating the combinatorics from the dynamics, Theoret. Comput. Sci. (2018), https://doi.org/10.1016/j.tcs.2018.03.001

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

Definition 4.15. We call two modal automata $\mathbb{A}$ and $\mathbb{A}'$ *one-step equivalent* if they are of the form $\mathbb{A} = (A, \Theta, \Omega, a_I)$ and $\mathbb{A}' = (A, \Theta', \Omega, a_I)$ with $\Theta(a) \equiv_1 \Theta(a)$ for all $a \in A$. $\lhd$

It is obvious that one-step equivalence implies equivalence: in fact one can easily prove that if $\mathbb{A}$ and $\mathbb{A}'$ are one-step equivalent, then for any Kripke structure the respective acceptance games are *isomorphic*.

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 [42,70]. 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 $\Omega_C(b) := \Omega_B(b)$.

Observe that since a_C forms a degenerate cluster of the automaton $\mathbb{A} \wedge \mathbb{B}$, the choice of the value of $\Omega_C(a_C)$ is in fact arbitrary. We chose the value $k + 1$ since this fits the condition of *linearity* that plays a role in section 8.

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. Following the complementation method of Muller & Schupp [42], we involve 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

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

we inductively define the *(boolean) dual* of one-step formulas. $\lhd$

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 $\mathbb{A} = (A, \Theta, \Omega, a_I)$ 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 $(D, \Theta', \Omega', a_I)$, where $D := A \uplus B$, Θ' is given by

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

To define the priority map Ω', 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). Note that, the priority map Ω' could have been defined as $\Omega \uplus \Psi$, but we find it convenient for later proofs to define Ω' such that all states in A get higher priority than all states in B.

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.

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* $\mathtt{Prop}$*-automaton* $\mathbb{A}$ *positive in* $x \in \mathtt{Prop}$*, and any state* $a \in A$*, there are formulas* θ^a_0 *and* θ^a_1 *in which* x *does not appear, such that*

$$\Theta(a) \equiv_K (x \wedge \theta^a_0) \vee \theta^a_1$$

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

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

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

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

and so we are done. □

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^a_0, θ^a_1 such that $\Theta(a) \equiv_K (x \wedge \theta^a_0) \vee \theta^a_1$.

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 $\Theta^x(\underline{x}) := x$. Actually the definition of Θ^x is such that x appears guarded in all transitions except for $\underline{x}$. Second, we will “split” each state a into two states a_0 and a_1, taking care of the θ^a_0- and the θ^a_1-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 $\alpha \in \mathtt{1ML}(\mathtt{Prop}, A)$, each occurrence of a state a by the formula $(\underline{x} \wedge a_0) \vee a_1$. We shall denote the resulting substitution as $\kappa : A \to A^x$. 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 $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be any modal $\mathtt{Prop}$-automaton which is positive in $x \in \mathtt{Prop}$, 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 $\mathtt{Prop}$-automaton $\mathbb{A}^x = (A^x, \Theta^x, \Omega^x, a^x_I)$ as follows:

- $A^x := (A \times \{0, 1\}) \cup \{\underline{x}\}$. We write (a, i) as a_i, for $i \in \{0, 1\}$.
- $\Theta^x(a_0) := \theta^a_0[\kappa]$ and $\Theta^x(a_1) := \theta^a_1[\kappa]$,
- $\Theta^x(\underline{x}) := x$,
- $a^x_I := (a_I)_1$,
- $\Omega^x(a_i) := \Omega(a)$ 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 do not introduce any conjunctions between states a_0 and a_1, 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 $\mathtt{tr}_{\mathbb{A}} : A \to \mu\mathtt{ML}$ assigning to each state $a \in A$ an *equivalent* μ-calculus formula $\mathtt{tr}_{\mathbb{A}}(a)$ in the sense that $\mathbb{A}\langle a \rangle \equiv \mathtt{tr}_{\mathbb{A}}(a)$.

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

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

and

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

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 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:

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

and the proof is finished. □

We can now define fixpoint operations on automata in which the proposition letter x appears positively. These definitions, and in particular, those of the transition map Θ', are directly motivated by Proposition 4.21.

Definition 4.22. Let $\mathbb{A}$ be any modal $\mathtt{Prop}$-automaton which is positive in $x \in \mathtt{Prop}$. The $\mathtt{Prop} \setminus \{x\}$-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 $\Omega^x(\underline{x}) := 2 \cdot \mathtt{max}(\Omega^x[A^x]) + 1$.

Similarly, the $\mathtt{Prop} \setminus \{x\}$-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 $\Omega^x(\underline{x}) := 2 \cdot \mathtt{max}(\Omega^x[A^x]) + 2$. $\lhd$

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. Since these statements follow from the results we shall prove in section 8 (in particular, from Proposition 8.15) we leave semantic proofs as exercises for the reader. $\lhd$

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 $\mathbb{A}_\varphi$. 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:

$$\begin{array}{llll} \mathbb{A}_{\neg\varphi} & := \neg\mathbb{A}_\varphi & & \\ \mathbb{A}_{\varphi\vee\psi} & := \mathbb{A}_\varphi \vee \mathbb{A}_\psi & \mathbb{A}_{\varphi\wedge\psi} & := \mathbb{A}_\varphi \wedge \mathbb{A}_\psi \\ \mathbb{A}_{\Diamond\varphi} & := \Diamond\mathbb{A}_\varphi & \mathbb{A}_{\Box\varphi} & := \Box\mathbb{A}_\varphi \\ \mathbb{A}_{\mu x.\varphi} & := \mu x.\mathbb{A}_\varphi, & \mathbb{A}_{\nu x.\varphi} & := \nu x.\mathbb{A}_\varphi, \end{array}$$

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

We finish this section by stating the semantic correctness of the above definition. The proof of this proposition below proceeds by a routine induction on the complexity of the formula φ — we leave the details as an exercise to the reader.

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

$$\varphi \equiv \mathbb{A}_\varphi. \tag{2}$$

Remark 4.26. It may come as a surprise to the reader that this proposition does not play any role in the sequel (this is also the reason why we do not provide proof details). This does not mean that we could work with just *any* translation, however: in section 8 below we will provide a translation map $\mathtt{tr}$ in the opposite direction, and require that a formula φ is *provably* equivalent to its double translation $\mathtt{tr}(\mathbb{A}_\varphi)$. This observation, establishing a very strong (but non-semantic) link between φ and $\mathbb{A}_\varphi$, is in fact the content of one of our main results, Theorem 2. $\lhd$

5. Games for automata

5.1. Introduction

In this section we introduce two of our main tools: 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, which is based on the infinite tableau game introduced by Niwiński & Walukiewicz [44]. Our variant, 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 $\mathcal{S}(\mathbb{A})$, which is played by two players, $\forall$ and $\exists$, it helps to think of $\exists$ as defending the claim that the language $L(\mathbb{A})$ is nonempty. In fact, we may think of $\exists$'s winning strategies as blueprints for constructing (tree-based) structures that are to be accepted by $\mathbb{A}$. 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 come up with more evidence to her claims and 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 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 play of $\mathcal{S}(\mathbb{A})$ can be seen as a *bundle* of plays of the acceptance game played on exactly the structure that $\exists$ is constructing.

We will see the details further on, but to give a rough indication of how this works, positions of the satisfiability game $\mathcal{S}(\mathbb{A})$ will represent *macro-states* of $\mathbb{A}$, that is, subsets of the state space A. Given a position that represents such a macro-state B, $\exists$ is to defend the claim that there is a pointed Kripke model which is accepted by the automaton $\mathbb{A}\langle a\rangle$, for *every* position $a \in B$. Clearly, if such a pointed structure $(\mathbb{S}, s)$ exists, then each position (a, s) is winning for her in the acceptance game, and so for each a she will have a marking $m_a : R[s] \to \mathsf{P}A$ such that the formula $\Theta(a)$ holds in the one-step model $(V^{\dagger}, R[s])$. The information of these markings can be combined by defining, for each successor $t \in R[s]$, the binary relation $H_t := \{(a, b) \mid a \in B, b \in m_a(t)\}$; observe that, for each $t \in R[s]$, the set $\mathsf{Ran}H_t$ constitutes the macro-state of $\mathbb{A}$ of all states $b \in A$ for which the pair (t, b) is a possible next basic position in the acceptance game, following a basic position of the form (a, s) with $a \in B$. Now, essentially because of bisimulation invariance, for the satisfiability game we may focus on the *range* $\mathcal{R}_s := \{H_t \mid t \in R[s]\}$ of the map $H : R[s] \to A^{\sharp}$ (where $A^{\sharp}$ is the collection of binary relations over the set A).

Summarizing, and abstracting from the concrete pointed model $(\mathbb{S}, s)$, in order to defend her claim that the macro-state B is "satisfiable", the immediate evidence that $\exists$ has to come up with, is a pair $(\Upsilon, \mathcal{R}) \in \mathsf{P}\mathtt{Prop} \times \mathsf{P}A^{\sharp}$, that is, a *one-step frame* over the set $A^{\sharp}$ of binary relations on A. And the conditions that this one-step frame has to satisfy to be a legitimate move in the satisfiability game after the position representing $B \subseteq A$, is that

$$\Upsilon, \mathcal{R}, n_a \Vdash^1 \Theta(a),$$

for all $a \in B$, where $n_a : \mathcal{R} \to A$ is the *natural a-marking*

$$n_a : Q \mapsto Q[a]$$

mapping an arbitrary element $R \in \mathcal{R}$ to the set of its a-successors. In other words, pairs of the form $\Gamma = (\Upsilon, \mathcal{R})$ 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 its *range* $\mathsf{Ran}R \subseteq A$, (while its domain $\mathsf{Dom}R$ will help us to keep track of traces by remembering the predecessor of the elements of the macro states at each level). 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 $\Gamma = (\Upsilon, \mathcal{R}) \in \mathsf{K}A^{\sharp} = \mathsf{P}\mathtt{Prop} \times \mathsf{P}A^{\sharp}$ that provide legitimate moves in all associated acceptance games, in the sense that $(\Upsilon, \mathcal{R}, n_a) \Vdash^1 \Theta(a)$ for all $a \in \mathsf{Ran}R$. Positions of the form $(\Upsilon, \mathcal{R})$ are for $\forall$, and his set of admissible moves is simply given by $\mathcal{R}$ itself, that is, the moves available to him are provided by the binary relations in $\mathcal{R}$.[3] 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 $\Gamma = (\Upsilon, \mathcal{R}) \in \mathsf{K}A^{\sharp}$, followed by $\forall$ picking a next basic position $Q \in \mathcal{R}$.

Binary relation thus provide an elegant way to encode the interaction of the two players during one round of the satisfiability game, their main advantage shows up when, finally, we consider the *winning condition* of the satisfiability game. Clearly, the sequence of basic positions of an infinite $\mathcal{S}(\mathbb{A})$-play Σ provides an $A^{\sharp}$-stream, that is, an infinite sequence of binary relations over A. As we mentioned before, the idea is that Σ corresponds to a *bundle* of plays of the acceptance game for $\mathbb{A}$. The requirement that $\exists$ needs to win *all* of these plays can be nicely expressed in terms of the collection of *traces* over the $A^{\sharp}$-stream associated with Σ.

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

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. The key concept here is that of a *trace* running through an infinite sequence of binary relations. This concept appears in many papers dealing with decidability questions on fixpoint logics, going back to at least Street & Emerson [61] (where it appears under the name 'derivation sequence'); we took our terminology from Niwinński & Walukiewicz [44].

Definition 5.1. Fix a set A. Given a finite word $\Sigma = R_1R_2R_3 \dots R_k$ over the set $A^\sharp$ (see Definition 2.2), a *trace* through Σ is a finite A-word $\alpha = a_0a_1a_2 \dots a_k$ such that $a_iR_{i+1}a_{i+1}$ for all $i < k$. A *trace* through an $A^\sharp$-stream $\Sigma = R_1R_2R_3 \dots$ is an A-stream $\alpha = a_0a_1a_2 \dots$, such that $a_iR_{i+1}a_{i+1}$ for all $i < \omega$). In both cases we denote the set of traces through Σ as Tr_Σ.

Given a stream $\Sigma = R_1R_2R_3 \dots$ over $A^\sharp$ we denote by $\Sigma|_k$ the word $R_1 \dots R_k$, and for a trace $\tau = a_0a_1a_2 \dots$ through Σ we denote by $\tau|_k$ the restricted trace $a_0 \dots a_k$ through $\Sigma|_k$. We use similar notation for restrictions of words over $A^\sharp$ of length $\geq k$. $\lhd$

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* $\mathbb{G}_\Sigma$ as the directed graph with vertices $\omega \times A$ and edges $E_\mathbb{G} := \{((i,a),(j,b)) \mid j = i+1 \text{ and } R_iab\}$. $\lhd$

Definition 5.3. Fix a finite set A and a priority map $\Omega : A \to \omega$. We call a trace $\tau = a_0a_1 \dots$ through a $A^\sharp$-stream a *bad* trace if $\max(\Omega[Inf(\tau)])$, the highest priority occurring infinitely often on τ, is odd. We let NBT_Ω denote the set of $A^\sharp$-streams that contain no bad trace. $\lhd$

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 $\Omega : A \to \omega$, 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. $\square$

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 $n_a : A^\sharp \to \mathsf{P}A$ given by

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

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

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

Any object $\Gamma = (\curlyvee, \mathcal{R}) \in \mathsf{K}A^\sharp$ 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 $\alpha \in \mathtt{1ML}(\mathtt{Prop}, A)$, $a \in A$ and one-step frame $\Gamma = (\curlyvee, \mathcal{R})$, we write $\Gamma \Vdash^1_a \alpha$ or $A^\sharp, \Gamma \Vdash^1_a \alpha$ to denote that $A^\sharp, n_a, \Gamma \Vdash^1 \alpha$, and we define $[\![\alpha]\!]^1_a := \{\Gamma \in \mathsf{K}A^\sharp \mid \Gamma \Vdash^1_a \alpha\}$.

Given an object $\Gamma \in \mathsf{K}A^\sharp$, we let $\curlyvee_\Gamma$ and $\mathcal{R}_\Gamma$ denote the unique objects such that $\Gamma = (\curlyvee_\Gamma, \mathcal{R}_\Gamma)$. $\lhd$

Clearly then, we may indeed think of $\Gamma \in \mathsf{K}A^\sharp$ as the *family* $\{(\Gamma, n_a\restriction_{\mathcal{R}_\Gamma}) \mid a \in A\}$ 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. 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 play of the satisfiability game is given by the induced stream $\Sigma = R_0R_1 \dots \in (A^\sharp)^\omega$ of basic positions. This winner is $\exists$ if Σ belongs to the set NBT_Ω (also denoted by $NBT_\mathbb{A}$), 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 $\mathbb{A}$. $\lhd$

Please cite this article in press as: S. Enqvist et al., Completeness for the modal μ-calculus: Separating the combinatorics from the dynamics, Theoret. Comput. Sci. (2018), https://doi.org/10.1016/j.tcs.2018.03.001

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

Position	Player	Admissible moves
$R \in A^\sharp$	$\exists$	$\bigcap_{a \in \mathsf{Ran}R} [\![\Theta(a)]\!]^1_a$
$(\Upsilon, \mathcal{R}) \in \mathsf{K}A^\sharp$	$\forall$	$\{Q \in A^\sharp \mid Q \subseteq R \text{ for some } R \in \mathcal{R}\}$

In words, a play of the satisfiability game proceeds in *rounds*, moving from one basic position $R \in A^\sharp$ to the next. At such a basic position R, player $\exists$ has to pick an object from the set $\bigcap_{a \in \mathsf{Ran}R} [\![\Theta(a)]\!]^1_a$. That is, she has to come up with a pair $\Gamma = (\Upsilon, \mathcal{R}) \in \mathsf{K}A^\sharp$ which, taken as a family of one-step models as described above, satisfies $(\Upsilon, \mathcal{R}) \Vdash^1_a \Theta(a)$ for all $a \in \mathsf{Ran}(R)$.

Remark 5.7. An alternative and perhaps more natural version of $\mathcal{S}(\mathbb{A})$ would restrict the moves available to $\forall$ at position $(\Upsilon, \mathcal{R}) \in \mathsf{K}A^\sharp$ 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 $(\Upsilon, \mathcal{R}) \in \mathsf{K}A^\sharp$ 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. $\lhd$

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 $R = \varnothing$, that is, R is the empty relation. In this case we have that $\bigcap_{a \in \mathsf{Ran}R} [\![\Theta(a)]\!]^1_a = \mathsf{K}A^\sharp$, so that $\exists$ could pick any pair of the form $(\Upsilon, \varnothing)$ and win (as we will see in the next paragraph).

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

Finally, the situation where $\exists$ chooses $(\Upsilon, \mathcal{R})$ such that $\varnothing \in \mathcal{R}$, 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 $\forall$. But as we have just seen, should $\forall$ indeed pick the empty relation, then he would loose at the very next step of the play. $\lhd$

Convention 5.9. Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be a modal automaton. Since we will only consider plays 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. 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 play $\Sigma \in \mathrm{PM}_\exists @R_I$.

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

$$S := \{R_0 R_1 \ldots R_k \mid R_0 \Gamma_1 R_1 \ldots \Gamma_k R_k \in \mathrm{PM}^\chi_\exists @R_I\}.$$

$\exists$'s strategy can then be seen as a map $\chi : S \to \mathsf{K}A^\sharp$ and so it naturally induces a map $\sigma_\chi : S \to \mathsf{K}S$ given by

$$\sigma_\chi(\Sigma) := \Big(\Upsilon_{\chi(\Sigma)}, \{\Sigma \cdot Q \mid Q \in \mathcal{R}_{\chi(\Sigma)}\}\Big),$$

where the notations $\Upsilon_{\chi(\Sigma)}$ 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 $\mathbb{S}_\chi$. 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 $(a, s) \in \mathrm{Win}_\exists$ an A-marking $m_{a,s}$ on $R[s]$ such that the induced one-step model satisfies the following condition:

$$(V^\dagger(s), R[s], m_{a,s}) \Vdash^1 \Theta(a) \tag{3}$$

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 play $\Sigma = R_0 \dots R_k$, where $R_0 = R_I$. By a simultaneous induction, with any such play 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 play $(a_I, s_0)(a_1, s_1) \dots (a_k, s_k)$ of the acceptance game. Clearly, if we can maintain this condition indefinitely, $\exists$ will be the winner of the resulting infinite play, 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 play $\Sigma = R_0 \dots R_k$, where $R_0 = R_I$. By the inductive hypothesis it follows that all states $a \in \mathrm{Ran} R_k$ are such that $(a, s_k) \in \mathrm{Win}_\exists(\mathcal{A}(\mathbb{A}, \mathbb{S}))$. In particular, $\exists$'s winning strategy m provides a marking $m_{a,s_k} : R[s_k] \to \mathsf{P}A$ for each $a \in \mathrm{Ran} R_k$ (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

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

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 a subset of 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 play $\Sigma \cdot \chi(\Sigma) \cdot Q$ and the $\mathbb{S}$-path $s_0 \dots s_k s_{k+1}$ satisfy the condition (3) as an exercise to the reader. □

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 [9] 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. ◁

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 $\mathcal{S}(\mathbb{A})$-play 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 $(\Upsilon, \mathcal{R})$ of $\exists$ there is a legitimate move $(\Upsilon, \mathcal{R}')$ such that $\mathcal{R}' \subseteq \mathcal{N}$ and $\mathcal{R}' \overrightarrow{\mathsf{P}}{\subseteq} \mathcal{R}$. Then for any winning position in $\mathcal{S}(\mathbb{A})$ $\exists$ has a winning strategy that restricts her moves to pairs $(\Upsilon, \mathcal{R})$ with $\mathcal{R} \subseteq \mathcal{N}$.*

Proof. Assume that $\exists$ has a winning strategy χ 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 χ' that always selects moves $(\Upsilon, \mathcal{R})$ with $\mathcal{R} \subseteq \mathcal{N}$.

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

$$\Sigma = R_0 \chi_0 R_1 \chi_1 \dots R_k$$

which is χ'-guided, we will define a parallel play

$$\Sigma^* = R_0 \chi_0^* R_1 \chi_1^* \dots R_k$$

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

For the case where $k = 0$ there is nothing to prove, so assume inductively that there are plays 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 $(\Upsilon, \mathcal{R})$ be the move prescribed by $\exists$'s winning strategy χ in the partial play Σ^*, then by assumption there is a legitimate move $(\Upsilon, \mathcal{R}')$ such that $\mathcal{R}' \subseteq \mathcal{N}$ and $\mathcal{R}' \overrightarrow{\mathsf{P}}{\subseteq} \mathcal{R}$. Then we let

$$\chi'_\Sigma := (\Upsilon, \mathcal{R}')$$

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

To finish the inductive step, consider an arbitrary continuation of the play $\Sigma \cdot (\Upsilon, \mathcal{R}')$, say, where $\forall$ plays some relation Q. By definition, Q is a subset of some $Q' \in \mathcal{R}'$, while by $\mathcal{R}' \overrightarrow{\mathsf{P}}{\subseteq} \mathcal{R}$ 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 $\Sigma^* \cdot (\Upsilon, \mathcal{R})$. In other words, the two $k+1$-length plays $\Sigma \cdot (\Upsilon, \mathcal{R}') \cdot Q$ and $\Sigma^* \cdot (\Upsilon, \mathcal{R}) \cdot Q$ satisfy the required conditions. □

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 play of $\mathcal{S}(\mathbb{A})$, picks a move $\Gamma = (\Upsilon, \mathcal{R})$. Then without loss of generality we can assume that:

(1) $\mathsf{Dom}(Q) \subseteq \mathsf{Ran}(R)$, 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. Ideally, 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 $\{\sigma_b \mid b \in \mathsf{Ran}R\}$ be a family of substitutions such that the formula*

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

is consistent. Then $\exists$ has an admissible move $(\Upsilon, \mathcal{R})$ 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 $\{\sigma_b \mid b \in \mathsf{Ran}R\}$ into one single substitution over $A \times A$.

In detail, let $\lambda_b : A \to A \times A$ denote the substitution $\lambda_b : d \mapsto (b, d)$ that "tags" an arbitrary variable d with the state b, and let $\sigma : A \times A \to \mu\mathtt{ML}(\mathtt{Prop})$ 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 $\bigwedge\{\Theta(b)[\lambda_b] \mid b \in \mathsf{Ran}R\} \in \mathtt{1ML}(\mathtt{Prop}, A \times A)$ and the substitution σ, namely, that the formula

$$\big(\bigwedge\{\Theta(b)[\lambda_b] \mid b \in \mathsf{Ran}R\}\big)[\sigma]$$

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

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

As a special case of this result, consider the situation where the substitutions σ_b are all the same, given by $\sigma_b(a) = \mathtt{tr}_{\mathbb{A}}(a)$ for each $b, a \in A$, where $\mathtt{tr}_{\mathbb{A}} : A \to \mu\mathtt{ML}$ denotes the map translating states in $\mathbb{A}$ to their corresponding equivalent $\mu\mathtt{ML}$-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 $\bigwedge_{b \in \mathsf{Ran}R} \mathtt{tr}_{\mathbb{A}}(b)$.

But more importantly, Corollary 5.14 will be a key ingredient in our later construction 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 plays 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

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

Position	Player	Admissible moves
$(R, R') \in A^\sharp \times A'^\sharp$	I	$\bigcap_{a \in \mathsf{Ran}R} [\![\Theta(a)]\!]^1_a \times \{R'\}$
$(\Gamma, R') \in \mathsf{K}A^\sharp \times A'^\sharp$	II	$\{\Gamma\} \times \bigcap_{a' \in \mathsf{Ran}R'} [\![\Theta'(a')]\!]^1_{a'}$
$(\Gamma, \Gamma') \in \mathsf{K}A^\sharp \times \mathsf{K}A'^\sharp$	II	$\{\mathcal{Z} \subseteq A^\sharp \times A'^\sharp \mid (\Gamma, \Gamma') \in \overline{\mathsf{K}}\mathcal{Z}\}$
$\mathcal{Z} \subseteq A^\sharp \times A'^\sharp$	I	$\mathcal{Z}$

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 $\mathsf{Ran}R$ implies the conjunction[4] of $\mathsf{Ran}R'$. Here we take the 'conjunction of R' to be the conjunction of the set of automata $\{\mathbb{A}\langle b\rangle \mid b \in \mathsf{Ran}R\}$.

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 four 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 responds with some suitably related one-step model for R', inducing a move in the game $\mathcal{S}(\mathbb{A}')$. Concretely, for a one-step model $\Gamma = (\Upsilon, \mathcal{R})$, player II provides a one-step model $\Gamma' = (\Upsilon, \mathcal{R}')$ and a binary relation $\mathcal{Z} \subseteq A^\sharp \times A'^\sharp$ such that $(\mathcal{R}, \mathcal{R}') \in \overline{\mathsf{P}}\mathcal{Z}$. Player I then finishes the round by picking a pair (Q, Q') from $\mathcal{Z}$ as the next basic position.

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 play of $\mathcal{C}(\mathbb{A}, \mathbb{A}')$ naturally induces infinite plays 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 [54], 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 [6] or Joyal [32]). 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. $\lhd$

Definition 5.16. Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ and $\mathbb{A}' = (A', \Theta', \Omega', a'_I)$ be modal automata. 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*. Given $\Gamma = (\Upsilon, \mathcal{R}) \in \mathsf{K}A^\sharp$ and $\Gamma' = (\Upsilon', \mathcal{R}') \in \mathsf{K}A'^\sharp$, we shall write $(\Gamma, \Gamma') \in \overline{\mathsf{K}}\mathcal{Z}$ to say that $(\mathcal{R}, \mathcal{R}') \in \overline{\mathsf{P}}\mathcal{Z}$ and $\Upsilon = \Upsilon'$.

For the *winning conditions* of this game, consider an infinite play Σ of $\mathcal{C}(\mathbb{A}, \mathbb{A}')$, and let $(R_n, R'_n)_{n<\omega}$ be the induced stream of basic positions in Σ. Then player I is the winner of Σ if $(R_n)_{n<\omega} \in NBT_\Omega$ but $(R'_n)_{n<\omega} \notin NBT_{\Omega'}$; that is, if there is a bad trace on the $\mathbb{A}'$-side but not on the $\mathbb{A}$-side. If the position $\big(\{(a_I, a_I)\}, \{(a'_I, a'_I)\}\big)$ 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: $\mathbb{A} \models_{\mathsf{G}} \mathbb{A}'$. $\lhd$

A particularly simple type of strategy for Player II in the consequence game, that we will make use of a number of times, is what we call a *functional strategy*, in which the response chosen by Player II at each position (Γ, R') is as follows. Where $\Gamma = (\Upsilon, \mathcal{R})$, he picks a map $F : \mathcal{R} \to A'^\sharp$, and defines his two subsequent moves as $\Gamma' := (\mathsf{K}F)\Gamma$ and $\mathcal{Z} := \mathsf{Gr}F$ $(= \{(R, FR) \mid R \in \mathcal{R}\})$, respectively. Since these moves are completely determined by the map F, we will usually specify a functional strategy simply by defining the map F. To check legitimacy, we need to verify that $(\mathsf{K}F)\Gamma \in \bigcap_{a' \in \mathsf{Ran}R'} [\![\Theta'(a')]\!]^1_{a'}$, the fact that $(\Gamma, (\mathsf{K}F)\Gamma) \in \overline{\mathsf{K}}(\mathsf{Gr}F)$ is immediate by the definitions.

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 5.12.

Proposition 5.17. *Let $\mathbb{A} = (A, \Theta_A, \Omega_A, a_I)$ and $\mathbb{B} = (B, \Theta_B, \Omega_B, b_I)$ be modal automata, and let $\mathcal{N} \subseteq A^\sharp$ be some set of relations. Assume that for every basic position $(Q, R) \in A^\sharp \times B^\sharp$ of the consequence game, and for every legitimate move $(\Upsilon, \mathcal{Q}')$ for player I at this position, she has a legitimate move $(\Upsilon, \mathcal{Q})$ such that $\mathcal{Q} \subseteq \mathcal{N}$ and $\mathcal{Q}\, \overline{\mathsf{P}}\!\subseteq \mathcal{Q}'$. Then for any winning position in $\mathcal{C}(\mathbb{A}, \mathbb{B})$ player I has a winning strategy that restricts her moves to pairs $(\Upsilon, \mathcal{Q})$ with $\mathcal{Q} \subseteq \mathcal{N}$.*

Proof. We write $Q_0 := \{(a_I, a_I)\}$, $R_0 := \{(b_I, b_I)\}$, and abbreviate $\mathcal{C} := \mathcal{C}(\mathbb{A}, \mathbb{B})@(Q_0, R_0)$. Let f be a winning strategy for player I in $\mathcal{C}$. In the same game we will provide I with a winning strategy $\overline{f}$, that restricts her moves to pairs $(\Upsilon, \mathcal{Q})$ with $\mathcal{Q} \subseteq \mathcal{N}$. This strategy $\overline{f}$ will be defined by induction on the length of a partial $\overline{f}$-guided play, while by a simultaneous induction we will:

[4] While the consequence game definitely has a proof-theoretic flavour, this 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.

(†) associate with each $\overline{f}$-guided play $\Sigma = (Q_n, R_n)_{n\leq k}$
an f-guided shadow play $\Sigma' = (Q'_n, R_n)_{n\leq k}$ such that $Q_n \subseteq Q'_n$ for all $n \leq k$.

Clearly this holds at the start of every $\mathcal{C}$-play if we take $Q'_0 := Q_0$. For the inductive step of the definition, fix a partial $\overline{f}$-guided play $\Sigma = (Q_n, R_n)_{n\leq k}$, and let $\Sigma' = (Q'_n, R_n)_{n\leq k}$ be the inductively given shadow play. In order to provide player I with a move Γ in Σ, first consider the move $\Gamma' = (\Upsilon, \mathcal{Q}') \in \mathsf{K}A^\sharp$ provided by f in the shadow play Σ'. By assumption there is a pair $\Gamma = (\Upsilon, \mathcal{Q}) \in \mathsf{K}A^\sharp$ which is a legitimate move at position (Q'_k, R_k) and such that $\mathcal{Q} \subseteq \mathcal{N}$ and $\mathcal{Q}\,\overline{\mathsf{P}}\!\subseteq \mathcal{Q}'$. Since $Q_k \subseteq Q'_k$ it is easy to see that this move Γ is also legitimate at the last position (Q_k, R_k) of Σ. Hence we may take this Γ to be the move suggested by the strategy $\overline{f}$.

Continuing the inductive definition, suppose that player II's answers to I's move Γ are, successively, $\Delta = (\Upsilon, \mathcal{R})$, with $\mathcal{R} \in B^\sharp$, and $\mathcal{Z} \subseteq A^\sharp \times B^\sharp$. Now consider the relation $\mathcal{Z}' \subseteq A^\sharp \times B^\sharp$ defined by $\mathcal{Z}' := \supseteq; \mathcal{Z}$ (where ; denotes the relational composition). We claim that

$$\Delta \text{ and } \mathcal{Z}' \text{ are legitimate moves for II at position } (\Gamma', R) \tag{4}$$

and

$$\text{for all } (Q', R) \in \mathcal{Z}' \text{ there is a } (Q, R) \in \mathcal{Z} \text{ such that } Q \subseteq Q'. \tag{5}$$

For a proof of (4), observe that the legitimacy of Δ is obvious, while that of $\mathcal{Z}'$ follows from the fact that $(\mathcal{Q}', \mathcal{R}) \in \overline{\mathsf{P}}\supseteq; \overline{\mathsf{P}}\mathcal{Z} = \overline{\mathsf{P}}(\supseteq; \mathcal{Z}) = \overline{\mathsf{P}}\mathcal{Z}'$. The claim (5) is immediate from the definitions.

Based on the statements (4) and (5), we can finish our inductive definition: in the play $\Sigma \cdot (\Gamma, R_k) \cdot (\Gamma, \Delta) \cdot \mathcal{Z}$ we let the strategy $\overline{f}$ pick a pair $(Q, R) \in \mathcal{Z}$ as given by (5). Clearly this is a legitimate move for player I. Finally, where $\Sigma \cdot (Q, R)$ is the continuation of Σ in terms of basic positions, the associated continuation of the shadow play is $\Sigma' \cdot (Q', R)$, and so it is obvious that player I has been able to maintain the constraint (†).

It should be clear that the thus defined strategy $\overline{f}$ always picks legitimate moves of the right type. It remains to check that it is a winning strategy in $\mathcal{C}$.

It is straightforward to verify that player I will never get stuck in an $\overline{f}$-guided play, so we confine our attention to infinite plays. Let $\Sigma = (Q_n, R_n)_{n<\omega}$ be an infinite $\overline{f}$-guided play, then clearly there is an infinite f-guided shadow play $\Sigma' = (Q'_n, R_n)_{n<\omega}$ such that $Q_n \subseteq Q'_n$ for all $n < \omega$. By assumption that f is a winning strategy in $\mathcal{C}$, the play Σ' is a win for player I. That is, all traces through $(Q'_n)_{n<\omega}$ are good, while there is a bad trace through $(R_n)_{n<\omega}$. Obviously then, all traces through $(Q_n)_{n<\omega}$ are good, and so the existence of a bad trace through $(R_n)_{n<\omega}$ means that Σ as well is a win for player I. □

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

Proposition 5.18. *For any two modal automata $\mathbb{A}$ and $\mathbb{A}'$ it holds that*

$$\mathbb{A} \models_{\mathsf{G}} \mathbb{A}' \textit{ implies } \mathbb{A} \models \mathbb{A}'. \tag{6}$$

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 play with basic positions $(a'_I, s_I), (a'_1, s_1)...(a'_n, s_n)$, we shall define an f-guided shadow play $(R_I, R'_I)(R_1, R'_1)...(R_n, R'_n)$ such that the following conditions hold:
(1) $a_I a_1 ... a_n$ is a trace through $R_I R_1 ... R_n$ iff $(a_I, s_I)(a_1, s_1)...(a_n, s_n)$ is a χ-guided play; furthermore, each $b \in \mathsf{Ran}(R_n)$ is the last element of some trace on $R_I R_1 ... R_n$.
(2) $a'_I a'_1 ... a'_n$ is a trace through $R'_I R'_1 ... R'_n$.

Furthermore, we shall associate these shadow plays in a uniform manner, so that the shadow play of an initial segment of a partial play Σ is an initial segment of the shadow play associated with Σ. First, note that this means that $\exists$ wins all χ'-guided infinite plays: if $(a'_I, s_I)(a'_1, s_1)(a'_2, s_2)...$ is a loss for $\exists$ then $a'_I a'_1 a'_2 ...$ is a bad trace through $R'_I R'_1 R'_2 ...$ in the shadow play $(R_I, R'_I)(R_1, R'_1)(R_2, R'_2)...$ by condition (2). Since this play was f-guided and f is a winning strategy, this means that there must be some bad trace $a_I a_1 a_2 ...$ through $R_I R_1 R_2 ...$, and by condition (1) we get that $(a_I, s_I)(a_1, s_1)(a_2, s_2)...$ is an infinite χ-guided play, 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 play Σ consisting of positions (a'_I, s_I), (a'_1, s_1), ..., (a'_n, s_n) with a shadow play $(R_I, R'_I)(R_1, R'_1)...(R_n, R'_n)$ satisfying the conditions (1) and (2). For each $a \in \mathsf{Ran} R_n$, by (1) there is a χ-guided partial play with last position (a, s_n).

So we see that the move $\chi(a, s_n) : R[s_n] \to \mathsf{P}(A)$ prescribed by the winning strategy χ is legitimate for each such a. Define the map $H : R[s_n] \to A^\sharp$ by

$$H(v) := \{(a, b) \mid a \in \mathsf{Ran}R_n \;\&\; b \in \chi(a, s_n)(v)\}.$$

Then it follows by the one-step bisimulation invariance theorem (Proposition 4.6) that the pair $(\mathsf{K}H)(V^\dagger(s_n), R[s_n]) = (V^\dagger(s_n), H[R[s_n]])$ is a legitimate move for Player I in the consequence game at the position (R_n, R'_n). So the winning strategy f provides some $\mathcal{R} \in \mathsf{P}A'^\sharp$ and a binary relation $\mathcal{Z} \subseteq A^\sharp \times A'^\sharp$ such that $(H[R[s_n]], \mathcal{R}) \in \overline{\mathsf{P}}\mathcal{Z}$ and such that, for all $a' \in \mathsf{Ran}R'_n$:

$$V^\dagger(s_n), \mathcal{R}, m_{a'} \Vdash_1 \Theta'(a'),$$

where $m_{a'}$ is the natural marking at a'. We get $(R[s_n], \mathcal{R}) \in \overline{\mathsf{P}}(H; \mathcal{Z})$, so it follows by Proposition 4.8 (applied to the converse of the relation $H; \mathcal{Z}$) that $V^\dagger(s_n), R[s_n], h \Vdash_1 \Theta(a')$, where the marking h is defined by:

$$h(v) = \bigcup\{m_{a'}(Q) \mid (H(v), Q) \in \mathcal{Z}\}.$$

So we set $\chi'(\Sigma) = h$, and this is a legitimate move. Furthermore, if $b' \in h(v)$, then there is some Q with $(a', b') \in Q$ and $(H(v), Q) \in \mathcal{Z}$, and we can continue the shadow play with the pair $(H(v), Q)$. Then the extended shadow play

$$(R_I, R'_I)(R_1, R'_1)\ldots(R_n, R'_n)(H(v), Q)$$

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

It should be stressed that the converse direction of Proposition 5.18 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 $\mathbb{A} \not\models_{\mathsf{G}} \mathbb{A}'$, 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 $\mathbb{A}^x[\mu x.\mathbb{A}/x] \models_{\mathsf{G}} \mu x.\mathbb{A}$.*

Proof. Recall that the automaton $\mu x.\mathbb{A}$ has the same carrier as the automaton $\mathbb{A}^x$, and that the automaton $\mathbb{A}^x[\mu x.\mathbb{A}/x]$ is built from $\mu x.\mathbb{A}$ together with one disjoint copy of $\mathbb{A}^x$, so $\mathbb{A}^x[\mu x.\mathbb{A}/x]$ will contain for each state a in $\mu x.\mathbb{A}$ an extra state a' corresponding to a belonging to the disjoint copy of $\mathbb{A}^x$. With this in mind, we define a map f from states of $\mathbb{A}^x[\mu x.\mathbb{A}/x]$ 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 $\mathbb{A}^x[\mu x.\mathbb{A}/x]$ 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 (functional) strategy for Player II in the game $\mathcal{C}(\mathbb{A}^x[\mu x.\mathbb{A}/x], \mu x.\mathbb{A})$ defined by choosing the map $F_0{\restriction}_{\mathcal{R}}$ in response to the first move $(\Upsilon, \mathcal{R})$ made by Player I, and choosing the map $F{\restriction}_{\mathcal{R}}$ as a response to every other move $(\Upsilon, \mathcal{R})$ made by Player I. It can be checked that this is a winning strategy for Player II. □

Note that this proposition is stated in terms of the guardified automaton $\mathbb{A}^x$ rather than $\mathbb{A}$. It is not too hard to convince oneself that the automata $\mathbb{A}[\mu x.\mathbb{A}]$ and $\mathbb{A}^x[\mu x.\mathbb{A}]$ are *semantically* equivalent, but the consequence game is a stronger concept and in general it is somewhat surprisingly not true that $\mathbb{A}[\mu x.\mathbb{A}/x] \models_{\mathsf{G}} \mu x.\mathbb{A}$. 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 $\theta^b_0 = p$ and $\theta^b_1 = q$, and we can take $\theta^a_0 = \bot$ and $\theta^a_1 = \Theta(a)$. (The priority map Ω is irrelevant to the example and hence not specified.) We shall show that $\mathbb{A}[\mu x.\mathbb{A}] \not\models_{\mathsf{G}} \mu x.\mathbb{A}$.

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

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

State	$\Theta(\mathbb{A})$	$\Theta^x(\mathbb{A}^x)$	$\Theta'(\mu x.\mathbb{A})$	$\Theta''(\mathbb{A}[\mu x.\mathbb{A}])$
a	$\Diamond b$	–	–	$\Diamond b$
b	$(x \wedge p) \vee q$	–	–	$\big(\Diamond((\underline{x} \wedge b_0) \vee b_1) \wedge p) \vee q\big)$
a_0	–	$\bot$	$\bot$	$\bot$
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

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

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

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 family of relations $\mathcal{R}$ and a binary relation $\mathcal{Z}$ such that $(\emptyset, \mathcal{R})$ satisfies the formula $\Diamond((\underline{x} \wedge b_0) \vee b_1)$, since $\underline{x}$ is the start state of $\mu x.\mathbb{A}$ and such that $(\{\{(a,b)\}\}, \mathcal{R}) \in \overline{\mathsf{P}}\mathcal{Z}$. To satisfy the mentioned formula, there must be a relation $Q \in \mathcal{R}$ such that $\mathcal{Z}\{(a,b)\}Q$ and either (i) $(\underline{x}, b_0), (\underline{x}, \underline{x}) \in Q$ or (ii) $(\underline{x}, b_1) \in Q$.

In either case, Player I first picks the pair $(\{(a,b)\}, Q)$ belonging to $\mathcal{Z}$ as his next move. 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 Player II.

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 family $\mathcal{R}'$ that Player II can respond with such that $(\{q\}, \mathcal{R}')$ 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 family $\mathcal{R}'$ that Player II can choose such that $(\{p\}, \mathcal{R}')$ satisfies the formula $\Theta'(b_1)$, which is just q. So again Player II is stuck, and Player I wins. $\lhd$

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

Proof. Clearly, the proof of the first item is trivial. Concerning the transitivity of $\models_{\mathsf{G}}$, 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 $\mathcal{C}(\mathbb{A}, \mathbb{A}'')$. $\square$

6. Taming the traces – one step at a time

We have seen that the winner of an infinite play 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 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. What is useful about this fact for our purposes, as we shall see, is that the (deterministic) stream automaton recognizing $NBT_{\mathbb{A}}$ can in turn be used to replace an arbitrary modal automaton $\mathbb{A}$ with a "simulating" automaton $\mathbb{D}$, for which the language $NBT_{\mathbb{D}}$ is much easier to handle. While the combinatorics of the trace graph of a stream corresponding to a match of the satisfiability game for a modal automaton $\mathbb{A}$, as defined in section 5, is rather intricate, for the simulating automaton $\mathbb{D}$ it is straightforward.

Please cite this article in press as: S. Enqvist et al., Completeness for the modal μ-calculus: Separating the combinatorics from the dynamics, Theoret. Comput. Sci. (2018), https://doi.org/10.1016/j.tcs.2018.03.001

More generally, we shall be interested in special classes of modal automata for which the trace graphs of plays of the satisfiability game have a particularly simple 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' 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 plays 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 $R \in A^\sharp$ be a relation over some set A. We call R *functional* if each $a \in A$ has at most one R-successor. The set of functional relations in $A^\sharp$ will be denoted by ${A^\sharp}_f$. $\lhd$

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 $a \in \mathsf{Dom}(R_1)$ 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, after some initial phase in which traces can merge or abort, the trace graph on any stream on ${A^\sharp}_f$ consists of at most $|A|$ many, mutually disjoint, infinite traces. Thus, in some intuitive sense it is easy to check whether there is any bad trace on such a stream.

A wider class of relations that maintains some of this simplicity is the following. Recall that the notion of a cluster of an automaton, which plays a key role in the definitions of this section, was introduced in Definition 4.12.

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

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

The set of clusterwise functional relations in $A^\sharp$ will be denoted by ${A^\sharp}_{cf}$. $\lhd$

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 $\Sigma = R_1 R_2 R_3 \dots$ be any stream over ${A^\sharp}_{cf}$. Then there exists a collection F of at most $|A|$ many infinite traces over Σ, such that every infinite trace on this stream is eventually equal to some trace in F.*

Proof. Straightforward. $\square$

In particular this means that we only have to examine the $|A|$ many traces in F to find out whether there is a bad trace on $R_0 R_1 R_2 \dots$, 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 number of infinite traces of a stream over ${A^\sharp}_{cf}$ is 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 $\mathbb{A} = (A, \Theta, \Omega, a_I)$. 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 $b \lhd a$;
(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 $A^\sharp{}_{thin}$. $\lhd$

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. By the first condition on thin relations, 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 bad 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 bad 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, and for all $k \geq n$ and all i, $\tau_i(k)$ is *not* a safe state in 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 thinness 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, and so we are done. $\square$

Again, this combinatorial property greatly simplifies the problem of checking whether there is some bad trace on $R_0 R_1 R_2 \ldots$, 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 $\mathbb{A}$. $\lhd$

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 3.15. Recall that for a finite set of formulas Ψ, the formula $\nabla\Psi$ is given as $\nabla\Psi := \Box \bigvee \Psi \wedge \bigwedge \Diamond\Psi$.

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

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

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

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

where $\pi \in \mathtt{LitC}(\mathtt{Prop})$ and $B \subseteq A$. $\lhd$

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

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 $\Gamma = (\Upsilon, \mathcal{R}) \in \mathsf{K}A^\sharp$, observe that

$$(\Upsilon, \mathcal{R}) \Vdash^1_a \nabla B \text{ iff } B \subseteq \bigcup \{R[a] \mid R \in \mathcal{R}\} \text{ and } R[a] \cap B \neq \varnothing \text{ for all } R \in \mathcal{R}. \quad (7)$$

This boils down to

$$(\mathtt{Y}, \mathcal{R}) \Vdash^1_a \nabla\varnothing \text{ iff } \mathcal{R} = \varnothing \tag{8}$$

in the specific case where $B = \varnothing$. ◁

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 $\Theta(a) = \nabla\varnothing \vee \nabla\{a\}$ 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 $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be a disjunctive automaton, and let $R \in A^\sharp$ be a basic position of the satisfiability game. Then for every legitimate move $(\mathtt{Y}, \mathcal{R})$ of $\exists$ at R there is a legitimate move $(\mathtt{Y}, \mathcal{R}')$ such that $(\mathcal{R}', \mathcal{R}) \in \overrightarrow{\mathsf{P}}\subseteq$ and $\mathcal{R}' \subseteq A^\sharp{}_f$.*

Proof. We first prove the following claim:

Claim 1. Let $\alpha \in \mathtt{1ML}_d(\mathtt{Prop}, A)$ and $(\mathtt{Y}, \mathcal{R}) \in \mathsf{K}A^\sharp$ be such that $(\mathtt{Y}, \mathcal{R}) \Vdash^1_a \alpha$. Then there is some $\mathcal{R}' \in \mathsf{P}A^\sharp$ such that $(\mathcal{R}', \mathcal{R}) \in \overrightarrow{\mathsf{P}}\subseteq$, $(\mathtt{Y}, \mathcal{R}') \Vdash^1_a \alpha$ and $|R[a]| = 1$ for all $R \in \mathcal{R}'$.

Proof of Claim. Since every $\alpha \in \mathtt{1ML}_d(\mathtt{Prop}, A)$ 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 $(\mathtt{Y}, \mathcal{R}) \Vdash^1_a \nabla B$, and distinguish cases. If $B = \varnothing$ the claim holds trivially since by (8) we find $\mathcal{R} = \varnothing$ and so we may take $\mathcal{R}' := \mathcal{R}$.

Hence we may focus on the case where $B \neq \varnothing$. First of all it follows from (7) that $\mathcal{R} \neq \varnothing$ and that $R[a] \cap B \neq \varnothing$ 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 (7) 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 $(\mathcal{R}'', \mathcal{R}) \in \overrightarrow{\mathsf{P}}\subseteq$ and $(\mathtt{Y}, \mathcal{R}'') \Vdash^1_a \nabla B$.

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

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

that is, we 'split' every $R \in \mathcal{Q}$. Using (7) it is then a matter of routine to verify that $\mathcal{R}'$ meets the required conditions. □

The Proposition now follows by repeatedly applying Claim 1, once for each $a \in \mathsf{Ran}(R)$. □

From this it follows by Proposition 5.12 that we may always assume without loss of generality that $\exists$ restricts her moves to pairs $(\mathtt{Y}, \mathcal{R})$ where $\mathcal{R} \subseteq A^\sharp{}_f$, so that the stream over $A^\sharp$ consisting of the basic positions of an infinite play in the satisfiability game consists only of functional relations. In the sequel we will state a number of results that are similar to Proposition 6.11; we will omit proofs since these all follow essentially the same line of reasoning.

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 plays 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. Semi-disjunctive automata are to modal automata what the *weakly aconjunctive formulas* introduced by Walukiewicz in [69] 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 [34].

Definition 6.12. 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 $\pi \in \mathtt{LitC(Prop)}$ and $\mathcal{B}$ ranges over collection of C-safe sets, defines the set $\mathtt{1ML}_C(\mathtt{Prop}, A)$ 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). $\lhd$

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.

The definition of semi-disjunctive automata is tailored towards the following proposition, the proof of which is similar to that of Proposition 6.11.

Proposition 6.13. *Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be a semi-disjunctive automaton, and let $R \in A^\sharp$ be a basic position of the satisfiability game. Then for every legitimate move $(\Upsilon, \mathcal{R})$ of $\exists$ at R, there is a legitimate move $(\Upsilon, \mathcal{R}')$ such that $(\mathcal{R}', \mathcal{R}) \in \overrightarrow{\mathsf{P}}\subseteq$ 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 plays guided by this strategy induce streams over $A^\sharp$ consisting entirely of thin relations.

We can make this point more explicit by considering the thin satisfiability game of Definition 6.7. We now obtain the following result, as an immediate consequence of Propositions 6.13 and 5.12:

Corollary 6.14. *Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ 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 6.13 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 $(\Upsilon, \mathcal{R})$ 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 $(\Upsilon, \mathcal{R})$ 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. $\square$

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 6.15. *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 and $\mathbb{A}$ is positive in x, then $\mathbb{A}[\mathbb{B}/x]$ is semi-disjunctive, 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

$$\pi \wedge \nabla\{\bigwedge A \mid A \in \mathcal{A}\} \wedge \nabla\{\bigwedge B \mid B \in \mathcal{B}\}, \tag{9}$$

where $\pi \in \mathtt{LitC}(\mathtt{Prop} \setminus \{x\})$, $\mathcal{A}$ 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 3.14(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

$$\text{if } \alpha \in \mathtt{1ML}_d(\mathtt{Prop}, A), \text{ then } \alpha[\kappa] \text{ is equivalent to a } C_a\text{-safe formula, for any } a \in A. \tag{10}$$

Clearly we may restrict our attention to formulas α of the form ∇B with $B \subseteq A$. The key to proving (10) is the following one-step equivalence:

$$\nabla\{\gamma_0 \vee \gamma_1 \mid \gamma \in \Gamma\} \equiv_1 \bigvee \Big\{ \nabla\{\gamma_i \mid (\gamma, i) \in Z\} \;\mid\; Z \subseteq \Gamma \times \{0,1\}, \mathsf{Dom} Z = \Gamma \Big\}. \tag{11}$$

It follows from (11) 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 $a \in A$. □

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 that we have introduced for modal automata in the previous section.

Theorem 3. *Let $\mathbb{A}$ and $\mathbb{D}$ be respectively a semi-disjunctive and an arbitrary modal automaton, and assume that $\mathbb{A} \models_{\mathsf{G}} \mathbb{D}$. 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 does not 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 6.16. *Let S be any set, α any one-step formula in $\mathtt{1ML}(\mathtt{Prop}, A)$ and let $m, m' : S \to \mathsf{P}(A)$ be A-markings such that $(\Upsilon, S, m) \Vdash^1 \alpha$ and $(\Upsilon, S, m') \Vdash^1 \alpha^\partial$. 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 $\big(\{(a_I, a_I)\}, \{(d_I, d_I)\}\big)$. 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 Proposition 5.17 and another variation of Proposition 6.11. Call a relation $R \subseteq B^\sharp$ $\mathbb{A}$-*thin* if the relation $\mathsf{Res}_A(R)$ is thin with respect to $\mathbb{A}$.

Claim 1. Without loss of generality we may assume that in any play of $\mathcal{S}_{thin}(\mathbb{A} \wedge \neg\mathbb{D})$, $\exists$ only picks moves $(\Upsilon, \mathcal{R})$ 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 play $\Sigma = R_0 \ldots R_n$:

(†) R_n is thin, and for $n \geq 1$ satisfies $|\mathsf{Ran}(R_n) \cap D| = 1$;
(‡) There is a χ-guided shadow $\mathcal{C}(\mathbb{A}, \mathbb{D})$-play of the form $(S_0, S'_0)(S_1, S'_1)\ldots(S_n, S'_n)$, 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 $d \in \mathsf{Ran}(R_{i-1}) \cap D$ 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 (†) and (‡) 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 ∀ can extend the play $R_0 \dots R_n$ 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 play Σ consisting, for some $n > 1$, of the positions $R_0 R_1 \dots R_n$, with shadow play $(S_0, S'_0)(S_1, S'_1) \dots (S_n, S'_n)$. Let $\Gamma = (\Upsilon, \mathcal{R}) \in \mathsf{K}B^\sharp$ be an arbitrary move chosen by ∃ 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:

$$(\Gamma, n_b) \Vdash^1 \Theta_B(b) \text{ for all } b \in \mathsf{Ran} R_n, \tag{12}$$

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

$$(\Gamma, n_d) \Vdash^1 \Theta_D(d)^\partial, \tag{13}$$

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

Recall that $\mathsf{Res}_A : B^\sharp \to A^\sharp$ is the map sending a relation R to $R \cap (A \times A)$, so that $(\mathsf{KRes}_A)\Gamma$ is the pair $(\Upsilon, \{R \cap (A \times A) \mid R \in \mathcal{R}\})$. We write $\{R \cap (A \times A) \mid R \in \mathcal{R}\} = \mathcal{R}_A$ for short, so that $(\mathsf{KRes}_A)\Gamma = (\Upsilon, \mathcal{R}_A)$. By the one-step bisimulation invariance theorem, we may infer from (12) and (c) that

$$((\mathsf{KRes}_A)\Gamma, n_a) \Vdash^1 \Theta_A(a), \text{ for all } a \in \mathsf{Ran} S_n, \tag{14}$$

so that $(\mathsf{KRes}_A)\Gamma$ is an admissible move for player I in the consequence game at position (S_n, S'_n). Thus we find an element $\Gamma' = (\Upsilon, \mathcal{R}') \in \mathsf{K}_{\mathtt{Prop}} D^\sharp$ such that $\Gamma' \in \bigcap_{b \in \mathsf{Ran} S'_n} [\![\Theta_D(b)]\!]^1_b$, and a relation $\mathcal{Z} \subseteq \mathcal{R} \times \mathcal{R}'$ with $(\overline{\mathsf{P}}\mathcal{Z})\mathcal{R}_A\mathcal{R}'$, dictated by Player II's winning strategy χ in the consequence game. By our inductive assumptions on S'_n we get in particular that

$$(\Gamma', n_d) \Vdash^1 \Theta_D(d). \tag{15}$$

We shall prove the following claim:

Claim 2. There is some $S \in \mathcal{R}$, some $S' \in \mathcal{R}'$ and some $c \in D$ with $(\mathsf{Res}_A S, S') \in \mathcal{Z}$ and $(d, c) \in S' \cap \mathsf{Res}_D S$.

Proof of Claim 2. First, we define the marking $m : \mathcal{R} \to \mathsf{P}(D)$ by setting:

$$m(S) = \bigcup \{S'[d] \mid (\mathsf{Res}_A S, S') \in \mathcal{Z}\}.$$

We first claim that:

$$(\Gamma, m) \Vdash^1 \Theta_D(d). \tag{16}$$

Since we know that $(\Gamma', n_d) \Vdash^1 \Theta_D(d)$, by Proposition 4.8 it suffices to prove that the one-step model (Γ, m) one-step simulates (Γ', n_d). The atomic condition holds obviously. To establish the (back) condition, if $S \in \mathcal{R}$ then $\mathsf{Res}_A S \in \mathcal{R}_A$, so there is some $S' \in \mathcal{R}'$ with $(S, S') \in \mathcal{Z}$, and it immediately follows by definition of m that $n_d(S') \subseteq m(S)$. Conversely, for the (forth) condition, take an arbitrary relation $S' \in \mathcal{R}'$. Then there is some $Q \in \mathcal{R}_A$ with $(Q, S') \in \mathcal{Z}$, and Q must equal $\mathsf{Res}_A S$ for some $S \in \mathcal{R}$. Again, it immediately follows from the definition of m that $n_d(S') \subseteq m(S)$ as required.

By Proposition 6.16 it follows from (13) and (16) that there is some $c \in D$ and some $S \in \mathcal{R}$ such that $c \in n_d(S) \cap m(S)$. Then by the definitions we find that, respectively, $(d, c) \in \mathsf{Res}_D S$ and $(d, c) \in S'$ for some S' with $(\mathsf{Res}_A S, S') \in \mathcal{Z}$ as required. □

With Claim 2 in place, we define the next move for ∀ prescribed by the strategy σ to be the relation $R_{n+1} := \mathsf{Res}_A S \cup \{(d, c)\}$, where $S \in \mathcal{R}$ and $c \in D$ are as described in the claim, so that $(d, c) \in S' \cap \mathsf{Res}_D S$ for some S' with $(\mathsf{Res}_A S, S') \in \mathcal{Z}$. Note that this is a legitimate move in response to $(\Upsilon, \mathcal{R})$ since $R_{n+1} \subseteq S \in \mathcal{R}$. The shadow play is then extended by the pair $(S_{n+1}, S'_{n+1}) := (\mathsf{Res}_A S, S')$ so that condition (‡c) of the induction hypothesis holds as an immediate consequence of the claim. For condition (†), it is obvious that $|\mathsf{Ran}(R_n) \cap D| = 1$; 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 ∀, first observe that he never gets stuck, so that we may focus on infinite plays. It suffices to prove that every infinite σ-guided play contains a bad trace, so consider an arbitrary such play $\Sigma = (R_i)_{i \geq 0}$.

Clearly we may assume that all initial parts of Σ, corresponding to the partial plays $(R_i)_{0\leq i\leq n}$, satisfy the conditions (†) and (‡). From this it follows that Σ itself has an infinite χ-guided shadow play $(S_i, S_i')_{i\geq 0}$ satisfying the condition (‡a–c). In addition, it follows from (†) that Σ will contain a *unique* trace in D, which by (‡) will also be a trace on the right side of the shadow play in the consequence game. That is, the play $R_0R_1R_n\dots$ contains a unique trace of the form $b_Id_1d_2d_3\dots$ 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 bad, then we are done. If not, then given the priorities assigned to states in $\neg\mathbb{D}$ 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 $b_Ia_1a_2a_3\dots$ on the left side $S_0S_1S_2\dots$ of the shadow play in the consequence game, since this shadow play was guided by the winning strategy χ of Player II. But then this trace $b_Ia_1a_2a_3\dots$ is also a bad trace in the play $R_0R_1R_2\dots$ of the satisfiability game. Summarizing, 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 is a loss for $\exists$ as required. □

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 $\mathsf{sim}(\mathbb{A})$. 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 [42]; they also mention the simulation problem and hint at a solution, but do not provide details. Emerson & Jutla [17] 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 [17], independently from [40]). A fairly general construction, for tree automata with various acceptance conditions, was given by Muller & Schupp [43]. 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 [2]. However, we shall present a strengthened version of this simulation theorem: rather than merely showing $\mathbb{A}$ and $\mathsf{sim}(\mathbb{A})$ to be semantically equivalent, we shall prove stronger claims involving the consequence game. As one might expect, we have $\mathbb{A} \models_{\mathsf{G}} \mathsf{sim}(\mathbb{A})$, from which it follows that $\mathsf{sim}(\mathbb{A})$ accepts every model accepted by $\mathbb{A}$ and vice versa. But something stronger holds: in fact given any modal automaton $\mathbb{B}$ which is positive in the proposition letter x, we have:

$$\mathbb{B}[\mathsf{sim}(\mathbb{A})/x] \models_{\mathsf{G}} \mathbb{B}[\mathbb{A}/x].$$

and

$$\mathbb{B}[\mathbb{A}/x] \models_{\mathsf{G}} \mathbb{B}[\mathsf{sim}(\mathbb{A})/x].$$

In a sense, what this strengthened simulation says is that not only is every modal automaton equivalent in terms of the consequence game to its simulating disjunctive automaton, but this equivalence is *preserved* by substitution inside any given automaton $\mathbb{B}$. Besides strengthening the simulation theorem for modal automata, this can also be seen as an automata theoretic counterpart to and generalization of one of the key lemmas in Walukiewicz' completeness proof (labelled "Lemma 39" in [69]). This lemma states that his "tableau consequence" relation holds between a tableau for the formula $\widehat{\alpha}(\widehat{\varphi})$ and one for φ, where $\varphi := \mu p.\widehat{\alpha}(p)$ and the operation $\widehat{(\cdot)}$ produces an equivalent disjunctive normal form of a formula.

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

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

$$\bigwedge\{p \mid p \in \Upsilon\} \wedge \bigwedge\{\neg p \mid p \in X \setminus \Upsilon\}.$$

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

$$\Theta'(R) = \bigvee\{\widehat{\Upsilon} \wedge \nabla\mathcal{R} \mid (\Upsilon, \mathcal{R}) \in \bigcap_{b\in \mathsf{Ran} R} [\![\Theta(b)]\!]^1_b\},$$

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 Proposition 5.4, we may pick some deterministic parity automaton $\mathbb{Z} = (Z, \zeta, \Omega', z_I)$ that recognizes NBT_Ω. Finally we define $\mathsf{sim}(\mathbb{A})$ 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 $G_{\mathbb{A}} : D \to A^\sharp$ 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. $\lhd$

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* $\mathsf{sim}(\cdot)$ *assigns to each modal automaton* $\mathbb{A}$ *a disjunctive modal automaton* $\mathsf{sim}(\mathbb{A})$ *such that, for any modal automaton* $\mathbb{B}$ *which is positive in* p*:*

(1) $\mathbb{B}[\mathbb{A}/p] \models_{\mathsf{G}} \mathbb{B}[\mathsf{sim}(\mathbb{A})/p]$.
(2) $\mathbb{B}[\mathsf{sim}(\mathbb{A})/p] \models_{\mathsf{G}} \mathbb{B}[\mathbb{A}/p]$;

Note that by taking $\mathbb{B}$ to be a trivial automaton, with just one state assigned the variable p as its one-step formula, we get $\mathsf{sim}(\mathbb{A}) \models_{\mathsf{G}} \mathbb{A}$ and $\mathbb{A} \models_{\mathsf{G}} \mathsf{sim}(\mathbb{A})$ as a special case of Theorem 4. From this, the semantic equivalence of an automaton and its simulation follows by Proposition 5.18.

To get a hint of the proof of Theorem 4(1), one should observe that the game trees of $\mathcal{S}(\mathbb{A})$ and $\mathcal{S}(\mathsf{sim}(\mathbb{A}))$ 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 $\mathbb{A}$ *are* just 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). A full proof can be found in the appendix A. Here, we give only the simpler proof for a special case: $\mathbb{A} \models_{\mathsf{G}} \mathsf{sim}(\mathbb{A})$, which is the only case of Theorem 4(1) that is strictly needed for the completeness proof.

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 play Σ of the consequence game $\mathcal{C}(\mathbb{A}, \mathsf{sim}(\mathbb{A}))$ 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 play. If Player I continues the play Σ consisting of basic positions $(R_0, R'_0) \dots (R_k, R'_k)$ by choosing the move $(\Upsilon, \mathcal{R}) \in \mathsf{K}A^\sharp$, then we let Player II respond with the map $F : \mathcal{R} \to (A^\sharp \times Z)^\sharp$ 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 functional winning strategy for Player II, and we leave the details to the reader. $\square$

Proving clause (2) is the difficult part of Theorem 4, and 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 $\mathsf{sim}(\mathbb{A})$ that we have associated with an arbitrary automaton $\mathbb{A}$:

Definition 7.2. Let $\mathbb{A} = (A, \Theta_{\mathbb{A}}, \Omega_{\mathbb{A}}, a_I)$ and $\mathbb{D} = (D, \Theta_{\mathbb{D}}, \Omega_{\mathbb{D}}, d_I)$ 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 $G : D \to A^\sharp$ satisfying the following conditions:

(DC1) $G(d_I) = \{(a_I, a_I)\}$
(DC2) Let $\Delta = (\Upsilon, E) \in \mathsf{K}D$ be such that $(\Upsilon, E, \mathsf{sing}) \Vdash^1 \Theta_{\mathbb{D}}(d)$, where sing is the singleton map given by $e \mapsto \{e\}$. Then $(\mathsf{K}G)\delta \in \mathsf{K}A^\sharp$ satisfies $(\mathsf{K}G)\Delta \Vdash^1_a \Theta_{\mathbb{A}}(a)$ for all $a \in \mathsf{Ran}(G(d))$.
(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. $\lhd$

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 $\mathsf{sim}(\mathbb{A})$, 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 $\mathsf{Ran}G(d)$, or differently put: for each for $a \in \mathsf{Ran}G(d)$, 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 $a \in \mathsf{Ran}G(d)$ is "entailed" by d in the following sense. Given an object $\Delta = (\Upsilon, \{d_1, \dots, d_k\}) \in \mathsf{K}D$, we can see Δ as a one-step model over D by taking the singleton map $\mathsf{sing} : D \to \mathsf{P}D$ (restricted to the set $E = \{d_1, \dots, d_k\}$) as a D-marking. Similarly, applying the map $\mathsf{K}G$ to Δ, we obtain the object $(\mathsf{K}G)\Delta = (\Upsilon, \{G(d_1), \dots, G(d_k)\}) \in \mathsf{K}A^\sharp$ which, as we have seen, may be taken as an A-indexed family of one-step models $(\Upsilon, \{G(d_1), \dots, G(d_k)\}, n_a)$ over A. Now the condition (DC2) requires that if $(\Upsilon, \{d_1, \dots, d_k\}, \mathsf{sing})$ satisfies the one-step formula $\Theta_{\mathbb{D}}(d)$, then the one-step model $(\Upsilon, \{G(d_1), \dots, G(d_k)\}, n_a)$ satisfies $\Theta_{\mathbb{A}}(a)$, for each $a \in \mathsf{Ran}(G(d))$. 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 $G(d_i)_{i\in\omega}$ 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* $\mathsf{sim}(\cdot)$ *assigns a disjunctive companion to any modal automaton.*

Proof. It is fairly straightforward to check that the projection map $G_{\mathbb{A}} : D \to A^\sharp$ 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 $\mathsf{sim}(\mathbb{A})$ is a disjunctive companion of $\mathbb{A}$. □

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*

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

To warm up, we first prove a simplified version of the Proposition (which immediately yields the statement $\mathsf{sim}(\mathbb{A}) \models_{\mathsf{G}} \mathbb{A}$ in Theorem 4(1)). For the proof, and henceforth, we shall rely on the following notation: if R is a binary relation and $d \in \mathsf{Dom}R$, and there exists a *unique* d' such that $(d, d') \in R$ (i.e. $R[d]$ is a singleton), then we denote this unique d' by d^+_R, or just d^+ when R is clear from context.

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

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

Proof. Since $\mathbb{D}$ is disjunctive, and because of the link between the moves of Player I in $\mathcal{C}(\mathbb{D}, \mathbb{A})$ at position (R, R') and $\exists$'s moves in the game $\mathcal{S}(\mathbb{D})$ at position R, from the Propositions 6.11 and 5.12 we may assume without loss of generality that player I only picks moves $\Gamma = (\Upsilon, \mathcal{R})$ such that all relations in $\mathcal{R}$ are functional, and we can also assume that $\mathsf{Dom}Q \subseteq \mathsf{Ran}R$ for each $Q \in \mathcal{R}$. 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).

The functional strategy that we will give to Player II will be completely determined by the map G witnessing that $\mathbb{D}$ is a companion of $\mathbb{A}$. More specifically, suppose that at a position $(\{(e, d)\}, R)$, player I plays $(\Upsilon, \mathcal{R})$; 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 play. To show that χ is winning for II, consider an infinite χ-guided play $\Sigma = (R_n, R'_n)_{n\in\omega}$. It follows from our assumptions on player I's strategy that $R_0 R_1 \ldots$ carries only a single trace, say, $d_0 d_1 \ldots$, 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. □

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 play $\Sigma = (R_n, R'_n)_{n\in\omega}$ 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 $\Sigma_l := R_0 R_1 \ldots$ 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 $\Sigma_r := R'_0 R'_1 \ldots$ 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

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

a Σ_l-trace

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

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 play leading to final position (R_n, R'_n), a map $g_n : \mathsf{Ran}_A R'_n \to \mathsf{Ran}_D R_n$. 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 play 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 $\mathbb{D}$. ◁

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

Definition 7.7. With $\Gamma, \Gamma' \in \mathsf{K}A^\sharp$, $a, a' \in A$, and $B \subseteq A$, we write $\Gamma \rightrightarrows^B_{a,a'} \Gamma'$ to abbreviate that $(\Gamma, m) \rightrightarrows^1 (\Gamma', m')$ where $m : R \mapsto R[a] \cap B$ and $m' : R \mapsto R[a'] \cap B$. Similarly, we write $\Gamma \leftrightarrows^B_{a,a'} \Gamma'$ as an abbreviation for $(\Gamma, m) \leftrightarrows^1 (\Gamma', m')$. ◁

Proposition 7.8. *Let $\Gamma, \Gamma' \in \mathsf{K}A^\sharp$. (1) If $\Gamma \rightrightarrows^B_{a,a'} \Gamma'$ then for all $\beta \in \mathtt{1ML}(\mathtt{Prop}, B)$ we have*

$$\Gamma \Vdash^1_a \beta \text{ implies } \Gamma' \Vdash^1_{a'} \beta.$$

(2) If $\Gamma \leftrightarrows^B_{a,a'} \Gamma'$ then for all $\beta \in \mathtt{1ML}(\mathtt{Prop}, B)$ we have

$$\Gamma \Vdash^1_a \beta \text{ iff } \Gamma' \Vdash^1_{a'} \beta.$$

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

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 $G : D \to A^\sharp$ be the map witnessing that $\mathbb{D}$ is a disjunctive companion of $\mathbb{A}$. We recall our notation Ran_A and Res_A from Definition 2.2.

Our goal is to provide player II with a functional 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 play Σ ending with (R, R'), player I always picks an element $\Gamma = (\Upsilon, \mathcal{R})$ such that

(Ass1) $\mathsf{Dom}Q \subseteq \mathsf{Ran}R$ for all $Q \in \mathcal{R}$;
(Ass2) $Q \cap (D \times B) = \varnothing$, for all $Q \in \mathcal{R}$;
(Ass3) $|Q[d] \cap D| = 1$ for all $d \in \mathsf{Ran}R \cap D$ and all $Q \in \mathcal{R}$;
(Ass4) for all $b \in \mathsf{Ran}R \cap B$, either $Q[b] \cap D = \varnothing$ for all $Q \in \mathcal{R}$, or $|Q[b] \cap D| = 1$ for all $Q \in \mathcal{R}$.

Proof of Claim. We shall use Proposition 5.17. We focus on the most difficult clause, (Ass4), and leave the rest to the reader. Consider $b \in \mathsf{Ran}R \cap B$ and suppose that $\Gamma = (\Upsilon, \mathcal{R}) \in \mathsf{K}(B \cup D)$ is a legitimate move for Player I, which we may assume already satisfies (Ass2). We recall that transition map of $\mathbb{B}[\mathbb{D}/p]$, i.e. Θ_{BD} is defined by $\Theta_{BD}(b) = \Theta_B(b)[\Theta_D(d_I)/p]$. If $\Gamma \notin [\![\Theta(d_I)]\!]^1_b$ then we set $\mathcal{R}' = \{Q \setminus (\{b\} \times D) \mid Q \in \mathcal{R}\}$. If $\Gamma \in [\![\Theta(d_I)]\!]^1_b$, then pick some disjunct $\widehat{\Upsilon} \wedge \nabla E_b$ of $\Theta_D(d_I)$ such that $(B \cup D)^\sharp, \Gamma \Vdash^1_b \widehat{\Upsilon} \wedge \nabla E_b$. As a step towards defining $\mathcal{R}'$ we construct a set $\mathcal{R}_b$ as follows: if $E_b = \emptyset$ then we must have $\mathcal{R} = \emptyset$ and set $\mathcal{R}_b = \mathcal{R} = \emptyset$. Otherwise, if $E_b \neq \emptyset$, set

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

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 $Q \in \mathcal{R}_b$, $(\mathcal{R}_b, \mathcal{R}) \in \overrightarrow{\mathsf{P}}\subseteq$, and $(B \cup D)^\sharp, (\Upsilon, \mathcal{R}_b) \Vdash^1_b \Theta_{BD}(b)$. The key observation is that if $(\Upsilon, \mathcal{R})$ satisfies $\Theta_D(d_I)$, then so does $(\Upsilon, \mathcal{R}_b)$, and this is proved by simply verifying that the required back-and-forth properties hold for E_b and $\mathcal{R}_b$. By repeating the procedure for all $b \in \mathsf{Ran}R \cap B$, we will finally find some $\mathcal{R}'$ satisfying $(\mathcal{R}', \mathcal{R}) \in \overrightarrow{\mathsf{P}}\subseteq$ together with the additional required properties. □

To appreciate the above claim, consider an arbitrary partial play

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

with $R_0 = R'_0 = \{(b_I, b_I)\}$. It follows by Claim 1 that we may assume each element $c \in \mathsf{Ran} R_k$ to lie on some trace through $R_0, \ldots, R_k$, and that every trace through $R_0, \ldots, R_k$ 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 $d \in \mathsf{Ran} R_n$ 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 $R_0, \ldots, R_n$ ending at d has *exactly* one continuation through $R_{n+1}, \ldots, R_k$. (This does not imply that all plays 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 $\mathsf{Ran}_D R_k$ that we will introduce now. Intuitively, we say, for $d, d' \in \mathsf{Ran}_D R_k$, 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 the formal definition we assume as given an arbitrarily chosen injective map $\mathsf{mb} : D \to \omega$, which we will use to break "ties". We call $\mathsf{mb}(d)$ the *birth minute* of d. We let $\mathsf{tb}_\Sigma(d)$ be the smallest pair of natural numbers (j, l) in the lexicographic order on $\omega \times \omega$ such that there is some $e \in \mathsf{Ran}_D R_j$ with $\mathsf{mb}(e) = l$ and such that the unique trace on $R_j \ldots R_k$ beginning with e ends with d. Note that by (Ass1) and (Ass2) such an e is guaranteed to exist and the corresponding trace is unique because of trace functionality in D. The pair $\mathsf{tb}_\Sigma(d) = (j, l)$ is called the *time of birth* of d relative to the play Σ; we simply write $\mathsf{tb}(d)$ if Σ is clear from context.

Given a state $d \in \mathsf{Ran}_D R_k$, by (Ass1) there is a trace τ through $R_0, \ldots, R_k$ such that $\tau(k) = d$. By (Ass2), all such traces start in B and at some moment j move to the $\mathbb{D}$-part of the automaton.

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

Finally, we define a strict total ordering on $\mathsf{Ran}_D R_k$ relative to Σ by saying that d is Σ-*older* than d' if $\mathsf{tb}(d)$ is smaller than $\mathsf{tb}(d')$ (in the lexicographic order). By injectivity of tb, every subset of $\mathsf{Ran}_D R_k$ has a Σ-*oldest* member, i.e. one with the smallest time of birth in the lexicographic order. We leave it for the reader to verify that, for $d \in \mathsf{Ran} R_n$ with $n < k$, it holds that $\mathsf{tb}(d^+) \leq \mathsf{tb}(d)$.

We now turn to the definition of player II's winning strategy χ in the consequence game $\mathcal{C}$ between $\mathbb{B}[\mathbb{D}/p]$ and $\mathbb{B}[\mathbb{A}/p]$. By a simultaneous induction on the length of a partial χ-play

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

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

$$F_n : (B \cup D)^\sharp \to (B \cup A)^\sharp$$

and

$$g_n : \mathsf{Ran}_A R'_n \to \mathsf{Ran}_D R_n.$$

We let the F-maps determine a functional strategy for player II in the following sense. Suppose that in the mentioned partial play Σ, player I legitimately picks an element $\Gamma = (\Upsilon, \mathcal{R}) \in \mathsf{K}(B \cup D)^\sharp$. Then player II's response will be the map $F_{n+1}\restriction_{\mathcal{R}}$, that is, the map F_{n+1}, restricted to the set $\mathcal{R} \subseteq (B \cup D)^\sharp$. Recall that, since we are specifying a functional strategy, we only need to define one of the two moves of player II since the other move is then uniquely determined.
Inductively we will ensure that the following conditions are maintained:

(*) $F_n R_n = R'_n$,
(†0) $R'_n \cap (A \times B) = \emptyset$,
(†1) $\mathsf{Res}_B R'_n = \mathsf{Res}_B R_n$,
(†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) $\mathsf{Res}_A R'_n \subseteq \bigcup \{G(d) \mid d \in \mathsf{Ran}_D R_n\}$,
(‡) $a \in \mathsf{Ran} G(g_n a)$, for all $a \in \mathsf{Ran}_A R_n$.

For some explanation and motivation of these conditions, observe that (*) indicates that Σ itself is indeed χ-guided. For conditions (†0)–(†3), 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 $\mathcal{C}(\mathbb{B}[\mathbb{D}/p], \mathbb{B}[\mathbb{A}/p])$-play is the same, and condition (†3) states that every pair $(a, b) \in \mathsf{Res}_A \mathsf{Ran} R'_n$ is 'covered' or 'implied' by some $d \in \mathsf{Ran}_D R_n$. Finally, (‡) states that, for every $a \in \mathsf{Ran} R'_n$, the map g_n picks a specific element $d \in \mathsf{Ran}_D R_n$ such that $a \in \mathsf{Ran}(G(d))$. 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 plays.

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

In the inductive case we will define the maps F_{n+1} and g_{n+1} for a partial play Σ as above. For the definition of $F_{n+1} : (B \cup D)^\sharp \to (B \cup A)^\sharp$, first observe that by (†0)–(†3) we are only interested in relations $R \in (B \cup D)^\sharp$ that are of the form $R = \mathsf{Res}_B R \cup (R \cap (B \times D)) \cup \mathsf{Res}_D R$. We will define F_{n+1} by treating these three parts of R separately, using, respectively, the identity map on $B^\sharp$ and two auxiliary maps that we define now.

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

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

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 $L : B \times D \to \mathsf{P}(B \times A)$, given by

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

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

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

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 play 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 $a \in \mathsf{Ran}_A R'_{n+1}$, 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 element of this set with the earliest birth minute (or equivalently in this case, the earliest birth date).

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 play $\Sigma \cdot (R_{n+1}, R'_{n+1})$) of the set $\{(g_n b)^+ \mid (b, a) \in R'_{n+1}\} \subseteq D$, i.e. the one with the earliest birth date.

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 $(a_I, a) \in G(d)$. 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 time of birth.

We now turn to the second case in 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 play-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) play $\Sigma \cdot (R_{n+1}, R'_{n+1})$. 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 Fig. 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).

Claim 2. By playing according to the strategy χ, player II indeed maintains the conditions (*), (†) and (‡).

Proof of Claim. Let Σ be a partial χ-play satisfying the conditions (*), (†) and (‡), and let $(R_{n+1}, R'_{n+1}) \in \mathsf{Gr}(F_{n+1})$ be any possible next position. It suffices to show that (R_{n+1}, R'_{n+1}) also satisfies (*), (†) and (‡).

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

$$(b, a) \in F_{n+1} R_{n+1} \iff (b, a) \in G((g_n b)^+). \tag{17}$$

for all $b, a \in A$. To prove (17), consider the following chain of equivalences, which hold for all $b, a \in A$:

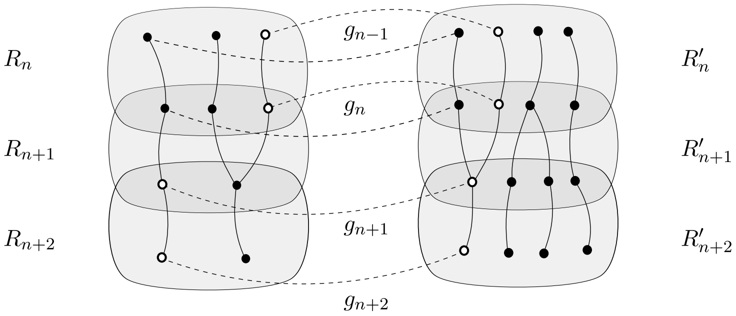

Fig. 1. A trace merge results in a trace jump.

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

Finally, for condition (‡), let $a \in \mathsf{Ran}_A R'_{n+1}$ 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 $(b,a) \in \mathsf{Res}_A R'_{n+1}$. But then $(b,a) \in G((g_n b)^+)$ by (17), so that indeed we find $a \in \mathsf{Ran}G(g_{n+1}a)$. 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 $a \in \mathsf{Ran}G(d)$; so we certainly have $a \in \mathsf{Ran}G(g_{n+1}a)$. □

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 play Σ ending with the position (R_n, R'_n) and a subsequent move $\Gamma = (\Upsilon, \mathcal{R}) \in \mathsf{K}(B \cup D)^\sharp$ by player I such that $\Gamma \Vdash^1_e \Theta_{BD}(e)$ for all $e \in \mathsf{Ran}R_n$. We need to show that

$$(\mathsf{K}F_{n+1})\Gamma \Vdash^1_c \Theta_{BA}(c) \tag{18}$$

for an arbitrary element $c \in \mathsf{Ran}R'_n$. 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 (‡) we find $c \in \mathsf{Ran}(Gd)$, where $d := g_n c$ belongs to $\mathsf{Ran}_D R_n$. Let the map $\mathsf{succ}_d : \mathcal{R} \to D$ be given by setting $\mathsf{succ}_d(Q)$ to be the unique $d' \in D$ with $(d,d') \in Q$ — this is well-defined by (Ass3) in Claim 1. We will apply the second clause of Definition 7.2 with $\Delta = (\mathsf{K}\mathsf{succ}_d)\Gamma$.

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

$$\Gamma \Vdash^1_d \Theta_D(d). \tag{19}$$

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

$$(\mathsf{K}\mathsf{sing})(\mathsf{K}\mathsf{succ}_d)\Gamma \Vdash^1_I \Theta_D(d). \tag{20}$$

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

$$(\mathsf{K}G)(\mathsf{K}\mathsf{succ}_d)\Gamma \Vdash^1_c \Theta_A(c). \tag{21}$$

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

$$(\mathsf{K}(G \circ \mathsf{succ}_d))\Gamma \Vdash^1_c \Theta_{BA}(c). \tag{22}$$

We now claim that

$$\text{for all } Q \in \mathcal{R}, a \in A:\ (c,a) \in (G \circ \mathsf{succ}_d)(Q) \Longrightarrow (c,a) \in F_{n+1}Q. \tag{23}$$

For a proof of (23), assume that $(c,a) \in (G \circ \mathsf{succ}_d)(Q) = G(d^+_Q)$. 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 (23) and the observations that (with $\Gamma = (\Upsilon, \mathcal{R})$) we have $\mathsf{K}(G \circ \mathsf{succ}_d)\Gamma = (\Upsilon, \{(G \circ \mathsf{succ}_d)(Q) \mid Q \in \mathcal{R}\})$ and $(\mathsf{K}F_{n+1})\Gamma = (\Upsilon, \{F_{n+1}(Q) \mid Q \in \mathcal{R}\})$, that

$$\mathsf{K}(G \circ \mathsf{succ}_d)\Gamma \sqsupseteq^A_{c,c} (\mathsf{K}F_{n+1})\Gamma. \tag{24}$$

But from this Proposition 7.8 yields (18), 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 $\Gamma \Vdash^1_c \Theta_B(c)[\Theta_D(d_I)/p]$, while we need to establish that $(\mathsf{K}F_{n+1})\Gamma \Vdash^1_c \Theta_B(c)[\Theta_A(a_I)/p]$. To achieve this it clearly suffices to show that

$$\Gamma \Vdash^1_c \alpha[\Theta_D(d_I)/p] \text{ implies } (\mathsf{K}F_{n+1})\Gamma \Vdash^1_c \alpha[\Theta_A(a_I)/p] \tag{25}$$

for all $\alpha \in \mathtt{1ML}(\mathtt{Prop}, B)$. We will prove (25) by induction on the one-step formula α, taken as a lattice term over the set $\{p\} \cup \mathtt{1ML}(\mathtt{Prop} \setminus \{p\}, B))$. 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

$$|Q[c] \cap D| = 1, \text{ for all } Q \in \mathcal{R}. \tag{26}$$

To see this, note that from $\Gamma \Vdash^1_c \Theta_D(d_I)$ it follows by the shape of disjunctive one-step formulas that either $\Gamma \Vdash^1_c \nabla\varnothing$ or $\Gamma \Vdash^1_c \nabla E$ for some non-empty $E \subseteq D$. In the first case we find by (8) in Remark 6.10 that $\mathcal{R} = \varnothing$, which clearly satisfies (26). In the second case we obtain by (7) in the same Remark that $Q[c] \cap D \neq \varnothing$ for some $Q \in \mathcal{R}$, so that (26) follows by (Ass4) in Claim 1.

But by (26) we may assume the existence of a map $\mathsf{succ}_c : \mathcal{R} \to D$ such that $Q[c] = \{\mathsf{succ}_c(Q)\}$ for all $Q \in \mathcal{R}$. (This also covers the case where $\mathcal{R} = \varnothing$.) It easily follows from $\Gamma \Vdash^1_c \Theta_D(d_I)$ that $(\mathsf{K}\mathsf{sing})(\mathsf{K}\mathsf{succ}_c)\Gamma \Vdash^1_I \Theta_D(d_I)$. Hence, by Definition 7.2 and functoriality of K we obtain

$$\mathsf{K}(G \circ \mathsf{succ}_c)\Gamma \Vdash^1_{a_I} \Theta_A(a_I). \tag{27}$$

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

$$\mathsf{K}(G \circ \mathsf{succ}_c)\Gamma \rightrightarrows^A_{a_I, c} (\mathsf{K}F_{n+1})\Gamma. \tag{28}$$

But from (27) and (28) it is immediate by Proposition 7.8 that

$$(\mathsf{K}F_{n+1})\Gamma \Vdash^1_c \Theta_A(a_I), \tag{29}$$

as required.

Base Case b: $\alpha \in \mathtt{1ML}(\mathtt{Prop} \setminus \{p\}, B)$, 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 (†1) Γ and $(\mathsf{K}F_{n+1})\Gamma$ coincide as one-step models over c. Formally:

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

Here we leave it for the reader to verify that $\Gamma \leftrightarrows^B_{cc} (\mathsf{K}\mathsf{Res}_B)\Gamma$ and $(\mathsf{K}\mathsf{Res}_B)(\mathsf{K}F_{n+1})\Gamma \leftrightarrows^B_{cc} (\mathsf{K}F_{n+1})\Gamma$, in the two respective steps where we apply Proposition 7.8.

Inductive case The inductive case in the proof of (25) is trivial.

This finishes the proof of Claim 3. □

Claim 4. Suppose Σ is an infinite χ-guided play with basic positions

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

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

Proof of Claim. Fix a χ-guided play $\Sigma = (R_i, R'_i)_{i \geq 0}$ 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 (†1), τ is also a trace on $R_0R_1R_2\ldots$, 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

$$\tau = b_0b_1\ldots b_na_{n+1}a_{n+2}a_{n+3}\ldots,$$

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

$$\text{there exists an index } k > n \text{ such that } g_{j+1}a_{j+1} = (g_ja_j)^+ \text{ for all } j \geq k. \tag{30}$$

In order to prove (30), recall that a *trace jump* occurs at the index $j > n$ if we have $g_{j+1}a_{j+1} \neq (g_ja_j)^+$. 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

$$\mathsf{tb}(g_ja_j) \geq \mathsf{tb}((g_ja_j)^+) = \mathsf{tb}(g_{j+1}a_{j+1}).$$

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

$$\mathsf{tb}(g_ka_k), \mathsf{tb}(g_{k+1}a_{k+1}), \mathsf{tb}(g_{k+2}a_{k+2}), \ldots$$

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 $g_{j+1}a_{j+1} \neq (g_ja_j)^+$. Recall that we defined $g_{j+1}a_{j+1}$ to be the oldest element of the set

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

But since $(a_j, a_{j+1}) \in R'_{j+1}$, it follows that $g_{j+1}a_{j+1}$ must be older than $(g_ja_j)^+$, with respect to the age relation induced by the play $(R_0, R'_0), \ldots, (R_{j+1}, R'_{j+1})$, and so $\mathsf{tb}(g_{j+1}a_{j+1})$ must be strictly smaller than $\mathsf{tb}((g_ja_j)^+) \leq \mathsf{tb}(g_ja_j)$, as required. This completes the proof of (30).

Let us finally see how (30) entails Claim 4. Suppose there exists an index k as in (30), and consider $g_ka_k \in \mathsf{Ran}_D R_k$. Pick an arbitrary initial trace $b_0\ldots b_nd_{n+1}\ldots d_k$ of $R_0\ldots R_k$ leading up to $g_ka_k = d_k$ (as mentioned already after Claim 1, the existence of such a trace follows from our assumptions on player I's strategy). Then the stream

$$b_0, \ldots, d_{k-1}, g_ka_k, g_{k+1}a_{k+1}, g_{k+2}a_{k+2}, \ldots$$

is a trace of $R_0R_1R_2\ldots$ by the property of the index k described in (30). Furthermore, it follows that $a_ka_{k+1}a_{k+2}\ldots$ is a trace of the stream

$$G(g_ka_k), G(g_{k+1}a_{k+1}), G(g_{k+2}a_{k+2}), \ldots$$

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 $d = g_ja_j$ 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 $d' = d^+ = (g_ja_j)^+ = g_{j+1}a_{j+1}$. From this we get $(a_j, a_{j+1}) \in G(g_{j+1}a_{j+1})$ as required. Note too that $a_ka_{k+1}a_{k+2}\ldots$ 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 $g_ka_k, g_{k+1}a_{k+1}, g_{k+2}a_{k+2}, \ldots$ is itself a bad trace, and so we have found a bad trace on $R_0R_1R_2\ldots$ as required. □

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 plays. But Claim 4 states that II wins all infinite plays of $\mathcal{C}(\mathbb{B}[\mathbb{D}/p], \mathbb{B}[\mathbb{A}/p])$ as well. □

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 $\mathtt{tr}$ 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. Of course, the translation will map an automaton $\mathbb{A}$ to a semantically equivalent formula, as the interested reader can verify this using routine arguments. We will leave out the proof since, just like the equivalence of a formula φ and the automaton $\mathbb{A}_\varphi$, that of $\mathbb{A}$ and its translation will actually not play a role in the main completeness proof (see also Remark 4.26).

The key property of the translation that we are after is something else, namely the following statement that we already mentioned in the introduction to the paper as one of our main lemmas:

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

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 (Definition 8.16), it is Theorem 2 that opens the door to proof theory for automata.

In order to provide the translation $\mathtt{tr}(\mathbb{A})$ of an automaton $\mathbb{A}$, we first define a map $\mathtt{tr}_\mathbb{A}$ assigning a formula to *each* state of $\mathbb{A}$. The formula $\mathtt{tr}(\mathbb{A})$ is then obtained by applying the map $\mathtt{tr}_\mathbb{A}$ 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 8.1. 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 $\Theta : A \to \mathtt{1ML}(\mathtt{Prop}, A \cup \mathtt{Prop})$. $\lhd$

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 8.2. A *generalized one-step model* is a structure (Υ, S, m) such that S is some set, $\Upsilon : S \uplus \{\star\} \to \mathsf{P}\mathtt{Prop}$ is a $\mathtt{Prop}$-marking on the set $S \uplus \{\star\}$ (with $\{\star\}$ denoting a singleton) and m is an A-marking on the set S. The one-step satisfaction relation $\Vdash^1$ for generalized one-step formulas in $\mathtt{1ML}(\mathtt{Prop}, A \cup \mathtt{Prop})$ is defined in the most obvious way: we treat a generalized one-step model (Υ, S, m) as if it were the standard one-step model $(\Upsilon(\star), S, \Upsilon{\restriction_S} \cup m)$ over $(\mathtt{Prop}, \mathtt{Prop} \cup A)$.

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: The *winning conditions* and the notion of *acceptance* are as

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

in the acceptance game for standard modal automata. $\lhd$

Remark 8.3. 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 $A^s := A \cup \{\underline{a} \mid a \in A\}$, $\Theta^s(a) := \Theta(a)[\underline{b}/b \mid b \in A]$, $\Theta^s(\underline{a}) := a$, $\Omega^s(a) := \Omega(a)$, 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. $\lhd$

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

Definition 8.4. 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 $\Theta : A \to \mathtt{1ML}(\mathtt{Prop}, A \cup \mathtt{Prop})$ and a priority function $\Omega : A \to \omega$.

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 $\mathbb{A}\langle a \rangle$ denote the initialized automaton (A, Θ, Ω, a). $\lhd$

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 $\mathtt{tr}_\mathbb{A}$ we will first confine our attention to so-called *linear* automaton structures.

Definition 8.5. An automaton structure $\mathbb{A} = (A, \Theta, \Omega)$ will be called *linear* if the relation $\sqsubset_\mathbb{A}$ is a strict linear order satisfying $(\lhd_\mathbb{A} \setminus \rhd_\mathbb{A}) \subseteq \sqsubset_\mathbb{A}$.

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 $\Omega(a)$, i.e., $\Omega'(a) = \Omega(a) \mod 2$, for all $a \in A$.

A linear refinement is called a *linearization*. $\lhd$

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 play of the acceptance game passes from one cluster to the next. Our focus on linear automaton structures is justified by the observation that all linearizations of an automaton $\mathbb{A}$ are equivalent to $\mathbb{A}$ (and hence, to one another). In the sequel we shall need a formulation of this equivalence in terms of the consequence game defined in section 5.

Proposition 8.6. *Every automaton structure $\mathbb{A}$ has a linearization $\mathbb{A}^l$ such that, for all $a \in A$*

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

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. Before giving the detailed definition, we recall that our notation for formula substitution has been given in Definition 3.3.

Definition 8.7. By induction on the size of a linear modal $\mathtt{Prop}$-automaton structure $\mathbb{A}$ we define a map $\mathtt{tr}_{\mathbb{A}} : A \to \mu\mathtt{ML}(\mathtt{Prop})$.

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

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

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 $\mathtt{Prop} \cup \{m\}$-generalised automaton structure $(A^-, \Theta^-, \Omega^-)$ with

$$A^- := A \setminus \{m\}$$
$$\Theta^- := \Theta{\restriction}_{A^-}$$
$$\Omega^- := \Omega{\restriction}_{A^-} .$$

Clearly we have $|\mathbb{A}^-| < |\mathbb{A}|$, so that inductively we may assume a map $\mathtt{tr}_{\mathbb{A}^-} : A^- \to \mu\mathtt{ML}(\mathtt{Prop} \cup \{m\})$.[5]

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

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

Second, by putting

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

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

[5] 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.

Remark 8.8. 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 $\{m_1, \ldots, m_n\}$, and the automaton structure $\mathbb{A}^-$ would remove all the maximal states. We would then get a set of "equations" $m_i := \Theta(m_i)[\mathtt{tr}_{\mathbb{A}^-}(b) \mid b \sqsubset_{\mathbb{A}} m_i]$, which is solved by a formula of the *vectorial μ-calculus* [2], and this formula can then be translated into the one-dimensional μ-calculus using the Bekič principle for simultaneous fixpoints. ◁

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 8.9. *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*

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

for all $a \in A$.

Proof. The proof of this proposition is conceptually straightforward, boiling down to the observation that, where M is the set $\{m_1, \ldots, m_n\}$ of maximal states of an automaton $\mathbb{A}$, it does not matter in which way we order the states in M to obtain a linearization of $\mathbb{A}$, in the sense that all choices provide provably equivalent translations. To prove this, it suffices to show that the Bekič principle holds in the class of modal μ-algebras, which can be done by verifying that the proofs in [2, section 1.4] in fact do not rely on completeness of the underlying lattices, but only on the existence of all least and greatest fixpoints. □

Proposition 8.9 ensures that modulo provable equivalence the following definition of $\mathtt{tr}(\mathbb{A})$ 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 8.10. To 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 $\mathtt{tr}_{\mathbb{A}} := \mathtt{tr}_{\mathbb{A}^l}$.

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

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

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

Remark 8.11. The translation given in Definition 8.10 is reasonably standard. A particular feature of the formula $\mathtt{tr}(\mathbb{A})$ is that it will always be *strongly guarded* in the sense that there is a modality between any two occurrences of a fixpoint operator. Its alternation depth will not exceed the maximal size of a cluster in the automaton $\mathbb{A}$. ◁

The following lemma gives two useful representations of the translation map $\mathtt{tr}_{\mathbb{A}}$ associated with an automaton structure $\mathbb{A}$. The point of the second result is that it displays each formula $\mathtt{tr}_{\mathbb{A}}(a)$ 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 $\mathbb{A}{\downarrow}a := (B, \Theta\restriction_B, \Omega\restriction_B)$ of which the carrier is given as $B := \{b \in A \mid b \sqsubseteq a\}$. ◁

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

$$\mathtt{tr}_{\mathbb{A}}(a) \equiv_K \Theta(a)[\mathtt{tr}_{\mathbb{A}}(b)/b \mid b \in A]. \tag{31}$$

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

$$\mathtt{tr}_{\mathbb{A}}(a) \equiv_K \eta_a a.\Theta(a)[\mathtt{tr}_{(\mathbb{A}\downarrow a)^-}(b)/b \mid b \sqsubset a][\mathtt{tr}_{\mathbb{A}}(b)/b \mid a \sqsubset b] \tag{32}$$

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

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

by definition of $\mathtt{tr}_{(\mathbb{A}\downarrow a)^-}$. Hence, we may read (32) as stating that

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

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 8.9 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:

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

For $a \neq m$, we have:

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

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:

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

as required. □

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 can be found in the appendix of the report version [19] of this paper.

Proposition 8.14. *Let $\mathbb{A} = (A, \Theta, \Omega, a_I)$ be a modal* Prop*-automaton with initial state a_I, which is positive in $x \in$* Prop*. Then we have:*

$$\mathtt{tr}(\mathbb{A}) \equiv_K (x \wedge \mathtt{tr}_{\mathbb{A}^x}((a_I)_0) \vee \mathtt{tr}_{\mathbb{A}^x}((a_I)_1) \tag{33}$$

$$\mathtt{tr}(\mu x.\mathbb{A}) \equiv_K \mu x.\mathtt{tr}_{\mathbb{A}^x}((a_I)_1) \tag{34}$$

$$\mathtt{tr}(\nu x.\mathbb{A}) \equiv_K \nu x.(\mathtt{tr}_{\mathbb{A}^x}((a_I)_0) \vee \mathtt{tr}_{\mathbb{A}^x}((a_I)_1)). \tag{35}$$

Note that we can alternatively write Proposition 8.14(34) as:

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

since we chose $(a_I)_1$ as the start state of $\mathbb{A}^x$. (The corresponding equation for $\nu x.\mathbb{A}$ does however *not* hold. To see this, take the $\{x, p, q\}$-automaton $\mathbb{A}$ to have just a single state mapped to the formula $(x \wedge p) \vee q$. We have $\mathtt{tr}(\nu x.\mathbb{A}) \equiv_K p \vee q$, but $\nu x.\mathtt{tr}(\mathbb{A}^x) \equiv_K q$.)

As mentioned, the central result of this section is the following.

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

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

Proof. A full proof can be found in the appendix of [19]. 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:

$$\begin{aligned} \mathtt{tr}(\mu x.\mathbb{A}) &\equiv_K \mu x.\mathtt{tr}_{\mathbb{A}^x}((a_I)_1) && \text{(Proposition 8.14(34))} \\ &\equiv_K \mu x.(x \wedge \mathtt{tr}_{\mathbb{A}^x}((a_I)_0)) \vee \mathtt{tr}_{\mathbb{A}^x}((a_I)_1) && \text{(Proposition 4.21)} \\ &\equiv_K \mu x.\mathtt{tr}(\mathbb{A}) && \text{(Proposition 8.14(33))} \end{aligned}$$

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

$$\begin{aligned} \mathtt{tr}(\nu x.\mathbb{A}) &\equiv_K \nu x.(\mathtt{tr}_{\mathbb{A}^x}((a_I)_0) \vee \mathtt{tr}_{\mathbb{A}^x}((a_I)_1) && \text{(Proposition 8.14(35))} \\ &\equiv_K \nu x.((x \wedge \mathtt{tr}_{\mathbb{A}^x}((a_I)_0)) \vee \mathtt{tr}_{\mathbb{A}^x}((a_I)_1) && \text{(Proposition 4.21)} \\ &\equiv_K \nu x.\mathtt{tr}(\mathbb{A}) && \text{(Proposition 8.14(33))} \end{aligned}$$

and the proof is done. □

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

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

The other cases are similar. □

We are now ready to define proof theoretic notions for automata.

Definition 8.16. A modal automaton $\mathbb{A}$ will be called *consistent* if the formula $\mathtt{tr}(\mathbb{A})$ 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 $\mathtt{tr}(\mathbb{A}) \leq_K \mathtt{tr}(\mathbb{B})$, and that $\mathbb{A}$ and $\mathbb{B}$ are provably equivalent if $\mathtt{tr}(\mathbb{A}) \equiv_K \mathtt{tr}(\mathbb{B})$. 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 $\varphi \leq_K \mathtt{tr}(\mathbb{A})$, etc. ◁

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 **K**, from which it follows that equivalent one-step formulas, seen as formulas of basic modal logic, are in fact provably equivalent.

Proposition 8.17. *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 6.15 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 ∃ 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. *∃ 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 [34]. An analogous lemma also features prominently in Walukiewicz' completeness proof [69, Theorem 31]. Observe that, in line with Kozen's approach, as an immediate consequence of our Proposition and Corollary 6.14, 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 $\mathsf{Dom} Q \subseteq \mathsf{Ran} R$.

First of all, it is immediate by the definitions (Definition 8.16 and 8.10) and the Propositions 8.6 and 8.9 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 $a \lhd_{\mathbb{A}} b$ but not $b \lhd_{\mathbb{A}} a$ (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 play 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 play played so far. More in detail, given a partial play Σ, with current position $R = \mathsf{last}(\Sigma) \in A^\sharp$, $\exists$ associates with every state $a \in \mathsf{Ran} R$ a 'private' formula $\mathtt{tr}_\Sigma(a)$ that *tightens* the 'public' formula $\mathtt{tr}_{\mathbb{A}}(a)$ in the sense that $\mathtt{tr}_\Sigma(a) \leq_K \mathtt{tr}_{\mathbb{A}}(a)$. $\exists$'s strategy will then be geared towards keeping the formula

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

consistent throughout the play. 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 [34], will be given in the appendix of [19].

Proposition 9.1. *Suppose that $\gamma \wedge \mu x.\varphi$ 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 play Σ with $R = \mathsf{last}(\Sigma)$ 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 $a \in \mathsf{Ran} R$, so that it is one of $\exists$'s tasks to avoid a being unfolded infinitely often on any continuation of τ which consists of lower priority states. The idea is now to think of the states in $\mathsf{Ran} R \setminus \{a\}$ 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

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

and at the same time ensure that the formula $\mathtt{tr}_\Sigma(a)$ is a least fixpoint formula, that is, of the form $\mu a.\varphi$ (cf. (32) 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 play Σ' extending Σ with $\mathsf{last}(\Sigma') = \mathsf{last}(\Sigma) = R$, by a Σ'-trace τ' that is a continuation of τ, then $\exists$ can guarantee that $\mathtt{tr}_{\Sigma'}(a) \leq_K \neg\gamma$. Hence, if in such a situation we would have that $\mathtt{tr}_{\Sigma'}(b) = \mathtt{tr}_\Sigma(b)$ for all $b \in \mathsf{Ran} R \setminus \{a\}$, we would find that

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

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 $\mathbb{A}$.

Please cite this article in press as: S. Enqvist et al., Completeness for the modal μ-calculus: Separating the combinatorics from the dynamics, Theoret. Comput. Sci. (2018), https://doi.org/10.1016/j.tcs.2018.03.001

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

$$\mathtt{tr}^{\delta}_{\mathbb{A}}(a) := \eta_a a.\delta(a) \wedge \Theta(a)[\mathtt{tr}_{(\mathbb{A}\downarrow a)^-}(b)/b \mid b \sqsubset a][\mathtt{tr}^{\delta}_{\mathbb{A}}(b)/b \mid a \sqsubset b] \tag{36}$$

we define the *tightening* map $\mathtt{tr}^{\delta}_{\mathbb{A}} : A \to \mu\mathtt{ML}(\mathtt{Prop})$ associated with δ. ◁

To obtain an understanding of this definition, it makes sense to compare it to the characterization of the translation map $\mathtt{tr}_{\mathbb{A}}$ in (32):

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

and to note that the map $\mathtt{tr}^{\delta}_{\mathbb{A}}$ is defined by a downward induction on the priority of states in $\mathbb{A}$. The definition of $\mathtt{tr}^{\delta}_{\mathbb{A}}$ and its comparison with (32) 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 $\mu\mathtt{ML}$ to $\mu\mathtt{ML}$-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) $\mathtt{tr}^{\delta}_{\mathbb{A}} \leq_K \mathtt{tr}_{\mathbb{A}}$*;*
(2) $\mathtt{tr}^{\delta}_{\mathbb{A}} \leq_K \delta$*;*
(3) if $\delta \leq_K \delta'$ *then* $\mathtt{tr}^{\delta}_{\mathbb{A}} \leq_K \mathtt{tr}^{\delta'}_{\mathbb{A}}$.

A key feature of our approach is that in a partial play 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 play Σ, this enables us to dynamically update the decoration by tightening it with relevant information about contexts. The 'private' formula $\mathtt{tr}_\Sigma(a)$ that, as mentioned earlier on, we want to associate with a state a in the range of the relation $\mathsf{last}(\Sigma)$, 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: $\mathsf{mst}_\Sigma(a)$. The map mst_Σ, associating a trace on Σ with each state in $\mathsf{Ran}(\mathsf{last}(\Sigma))$, 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 play are good (in the sense that the highest priority 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 $\bigwedge\{\mathtt{tr}_\Sigma(b) \mid b \in \mathsf{Ran}(\mathsf{last}\Sigma), b \neq a\}$, representing the context of a state $a \in \mathsf{Ran}(\mathsf{last}\Sigma)$ in one partial play Σ, still represents the same context in an extension Σ' of Σ with $\mathsf{Ran}(\mathsf{last}\Sigma') = \mathsf{Ran}(\mathsf{last}(\Sigma))$. 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 play 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 $\mathsf{mst}_\Sigma(a)$ 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 a continuation $\Sigma \cdot Q$ of a partial play Σ of $\mathcal{S}_{thin}(\mathbb{A})$ 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 play 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 $\mathcal{S}_{thin}(\mathbb{A})$-play 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 play, and the thinness constraint will allow us to focus on a small number of infinite traces on any infinite play, 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.

Please cite this article in press as: S. Enqvist et al., Completeness for the modal μ-calculus: Separating the combinatorics from the dynamics, Theoret. Comput. Sci. (2018), https://doi.org/10.1016/j.tcs.2018.03.001

Convention 9.4. Throughout this section we will restrict our attention to plays of the thin satisfiability game of the form $\Sigma = (R_i)_{i<\kappa}$ where $\mathsf{Dom}R_{i+1} \subseteq \mathsf{Ran}R_i$ 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 play of $\mathcal{S}_{thin}(\mathbb{A})$, player $\exists$ dynamically maintains an intricate data structure, associating with each partial play Σ the following entities (we recall from Definition 5.1 that Tr_Σ in the following clauses denotes the set of all traces through Σ):

– a set $V_\Sigma \subseteq \mathsf{Tr}_\Sigma$ of *selected* traces on Σ;
– a *most significant trace* map $\mathsf{mst}_\Sigma : \mathsf{Ran}(\mathsf{last}\Sigma) \to V_\Sigma$;
– a total *relevance order* $\prec_\Sigma$ on the set Tr_Σ;
– a total *priority order* $<_\Sigma$ on the set $\mathsf{Tr}_\Sigma \times A$.

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:) $\mathsf{mst}_\Sigma : \mathsf{Ran}(\mathsf{last}\Sigma) \to V_\Sigma$ is a bijection, with $\mathsf{mst}^{-1} = \mathsf{last}$ and, for all $\tau \in \mathsf{Tr}_\Sigma$:

$$\tau \preceq_\Sigma \mathsf{mst}_\Sigma(\mathsf{last}(\tau)).$$

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 $\mathsf{last}(\Sigma)$, with the map $\mathsf{mst}_\Sigma : \mathsf{Ran}(\mathsf{last}\Sigma) \to V_\Sigma$ selecting a *most significant trace* $\mathsf{mst}_\Sigma(a) \in \mathsf{Tr}_\Sigma$ for each state a. The compatibility condition on the map mst_Σ ensures that for each $a \in \mathsf{Ran}(\mathsf{last}\Sigma)$ there is a unique trace $\tau \in V_\Sigma$ such that $a = \mathsf{last}(\tau)$, viz., the trace $\tau = \mathsf{mst}_\Sigma(a)$.

The family of relevance relations $\prec_\Sigma$ 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 $\prec_\Sigma$. 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 play Σ. We make sure that this ranking is compatible with the priority order $\sqsubset_\mathbb{A}$ 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 $\prec_\Sigma$ 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. ◁

Before we can give the definitions of the actual structures $(V_\Sigma, \mathsf{mst}_\Sigma, \prec_\Sigma, <_\Sigma)$, we need a few auxiliary notions.

Definition 9.6. Given a trace τ on a partial $\mathcal{S}_{thin}(\mathbb{A})$-play Σ, let $C_\tau := C_{\mathsf{last}(\tau)}$ 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 $\rho = \rho|_k \cdot \mathsf{last}(\rho) = \rho|_k \cdot \rho(k+1)$. 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 $\mathsf{last}(\rho)$ is a safe state in its cluster.

Given a trace ρ, we define the *last refreshment date* of ρ, denoted $\mathsf{lrd}(\rho)$, 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 $j > k$. ◁

Note that if ρ is a trace on a partial $\mathcal{S}_{thin}(\mathbb{A})$-play Σ and $\mathsf{lrd}(\rho) > 0$, then the trace $\rho|_{\mathsf{lrd}(\rho)}$ is indeed refreshed. On the other hand, if $\mathsf{lrd}(\rho) = 0$ then ρ is the unique trace that has never been refreshed, where uniqueness is due to the assumption that Σ consists of thin relations. That is, $\rho|_k$ is not refreshed for any k. From this it is easily seen that if $\mathsf{lrd}(\sigma) = \mathsf{lrd}(\sigma')$ 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 play Σ, starting with the relevance order $\prec_\Sigma$.

Definition 9.7. Let Σ be a partial play of the thin satisfiability game for $\mathbb{A}$. We define the relation $\prec_\Sigma \subseteq \mathsf{Tr}_\Sigma \times \mathsf{Tr}_\Sigma$ by the following case distinction.

$$\sigma \prec_\Sigma \sigma' \text{ iff } \begin{cases} \text{either} & \mathsf{lrd}(\sigma) > \mathsf{lrd}(\sigma') \\ \text{or} & \mathsf{lrd}(\sigma) = \mathsf{lrd}(\sigma') \text{ and } \sigma|_k \prec_{\Sigma|_k} \sigma'|_k \\ \text{or} & \mathsf{lrd}(\sigma) = \mathsf{lrd}(\sigma') \text{ and } \sigma|_k = \sigma'|_k \text{ and } \sigma(k+1) \sqsubset \sigma'(k+1), \end{cases}$$

where in the last two cases we let k be such that $k+1 = \mathsf{lrd}(\sigma) = \mathsf{lrd}(\sigma')$. ◁

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 play 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 $\prec_\Sigma$ is a strict total order, see Proposition 9.10(1).

Definition 9.8. For any given partial play Σ and $a \in \mathsf{Ran}(\mathsf{last}\Sigma)$, set $\mathsf{mst}_\Sigma(a)$ to be the greatest trace σ with $\mathsf{last}(\sigma) = a$ according to the relevance order $\prec_\Sigma$, which clearly exists since the collection Tr_Σ is always finite and $\prec_\Sigma$ is a strict total order. With this definition in place, we simply set

$$V_\Sigma := \{\mathsf{mst}_\Sigma(b) \mid b \in \mathsf{Ran}(\mathsf{last}\Sigma)\}$$

to be the range of the map mst_Σ. ◁

Finishing the definition of the data structure associated with a partial play 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 play Σ.

Definition 9.9. For Σ being the unique initial play 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 play Σ of length k, and let $Q \subseteq A \times A$ be a thin relation. We define $<_{\Sigma\cdot Q}$ in a series of two steps.

(Step 1) First we define a new order $<^0_\Sigma$ on the set $\mathsf{Tr}_{\Sigma\cdot Q} \times A$. Basically, $<^0_{\Sigma\cdot Q}$ 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) $(\sigma, b) <^0_{\Sigma\cdot Q} (\sigma', b')$, in case σ is refreshed and σ' is not;
- (b) $(\sigma, b) <^0_{\Sigma\cdot Q} (\sigma', b')$ iff either $(\sigma|_k, b) <_\Sigma (\sigma'|_k, b')$ or $\sigma|_k = \sigma'|_k$ and $\mathsf{last}(\sigma) \sqsubset \mathsf{last}(\sigma')$, in case $\sigma \neq \sigma'$ and σ and σ' are either both refreshed or both not refreshed;
- (c) $(\sigma, b) <^0_{\Sigma\cdot Q} (\sigma, b')$ iff $b \sqsubset b'$.

(Step 2) Second, we move all erasable $(\sigma, b) \in \mathsf{Tr}_{\Sigma\cdot Q} \times A$ to the bottom of the list, where a pair (σ, b) is *erasable* if either $b \sqsubset \mathsf{last}(\sigma)$, or $b = \mathsf{last}(\sigma) \in A^\nu$. More precisely, we define the order $<_{\Sigma\cdot Q}$ by setting:
- (a) $(\sigma, b) <_{\Sigma\cdot Q} (\sigma', b')$, in case (σ, b) is erasable and (σ', b') is not;
- (b) $(\sigma, b) <_{\Sigma\cdot Q} (\sigma', b')$ iff $(\sigma, b) <^0_{\Sigma\cdot Q} (\sigma', b')$, 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 play Σ, the data structure $(V_\Sigma, \mathsf{mst}_\Sigma, \prec_\Sigma, <_\Sigma)$ is well defined and satisfies the required conditions.

Proposition 9.10. *The following hold for any partial play $\Sigma = (R_i)_{i\le k}$ of $\mathcal{S}_{thin}(\mathbb{A})$:*

(1) the relation $\prec_\Sigma$ is a strict total order on Tr_Σ;
(2) the map mst_Σ is a bijection from $\mathsf{Ran} R_k$ to V_Σ satisfying the compatibility condition;
(3) the relation $<_\Sigma$ is a strict total order on $\mathsf{Tr}_\Sigma \times A$ 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 play Σ. The case for the unique play consisting of the relation $\{(a_I, a_I)\}$ is trivial, so we focus on the induction step: let Σ be a play of a given length for which the induction hypothesis holds for all shorter plays, and let σ, τ be traces on Σ. Suppose that we have neither $\sigma \prec_\Sigma \tau$ nor $\tau \prec_\Sigma \sigma$. If $\mathsf{lrd}(\sigma) = \mathsf{lrd}(\tau) = 0$ then we must have $\tau = \sigma$, and the only other possibility is that $\mathsf{lrd}(\sigma) = \mathsf{lrd}(\tau) = k+1$ 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 $\mathsf{last}(\tau)$ is maximal among the traces ending with $\mathsf{last}(\tau)$, hence $\tau \preceq_\Sigma \mathsf{mst}_\Sigma(\mathsf{last}(\tau))$. 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. □

Next we establish some basic properties of the relevance order and most significant trace.

Proposition 9.11. *Let Σ be any partial play 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 $\sigma \cdot b \prec_{\Sigma \cdot Q} \tau \cdot b$.*

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 play Σ. Furthermore, we have $(\sigma \cdot b)|_k = \sigma$ and $(\tau \cdot b)|_k = \tau$, so since $\sigma \prec_\Sigma \tau$ we get $\sigma \cdot b \prec_{\Sigma \cdot Q} \tau \cdot b$ from the definition of the relevance order.

In the latter case, we distinguish cases:

(Case 1:) $\mathsf{lrd}(\tau) < \mathsf{lrd}(\sigma)$. Then since the traces $\sigma \cdot b$ and $\tau \cdot b$ are both not refreshed we have $\mathsf{lrd}(\tau \cdot b) = \mathsf{lrd}(\tau) < \mathsf{lrd}(\sigma) = \mathsf{lrd}(\sigma \cdot b)$, and thus we get $\sigma \cdot b \prec_{\Sigma \cdot Q} \tau \cdot b$ as required.

(Case 2:) $\mathsf{lrd}(\tau) = \mathsf{lrd}(\sigma) = k+1$ 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 $\mathsf{lrd}(\tau) = \mathsf{lrd}(\sigma) = 0$ since it would follow that $\tau = \sigma$.) In both these cases it follows that $\sigma \cdot b \prec_{\Sigma \cdot Q} \tau \cdot b$, since $\mathsf{lrd}(\tau \cdot b) = \mathsf{lrd}(\tau) = \mathsf{lrd}(\sigma) = \mathsf{lrd}(\sigma \cdot b)$, $\tau \cdot b|_k = \tau|_k$ 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. □

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 play 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 $\mathsf{mst}(\mathsf{last}(\tau|_k)) = \tau|_k$.*

Proof. Suppose for a contradiction that there is some trace σ on $\Sigma|_k$ with $\mathsf{last}(\sigma) = \mathsf{last}(\tau|_k)$ 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. □

The corresponding *upward* coherence condition does not hold: due to the occurrence of *trace merging* it is not always the case that $\sigma \cdot a \in V_{\Sigma \cdot Q}$ whenever $\sigma \in V_\Sigma$ and $(\mathsf{last}(\sigma), a) \in Q$. In case we have $\sigma \cdot a \neq \mathsf{mst}_{\Sigma \cdot Q}(a)$ (and thus $\sigma \cdot a \notin V_{\Sigma \cdot Q}$), we say that a *trace jump* occurs.

9.2.2. Infinite traces

The data structure $(V_\Sigma, \mathsf{mst}_\Sigma, \prec_\Sigma, <_\Sigma)$ 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 play in the thin satisfiability game. To see how this works, we turn to infinite plays, but first we adopt a notational convention.

Convention 9.13. Let Σ be a (full or partial) play 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 $\mathsf{mst}_{\Sigma|_k}$, etc.

Definition 9.14. Let Σ be an infinite play of $\mathcal{S}_{thin}(\mathbb{A})$. We call a trace τ on Σ *continuous* if it is infinite and $\tau|_k \in V_k$ for all $k \in \omega$. We let V_Σ denote the set of all continuous traces through Σ. ◁

Continuous traces have various nice properties; in particular, any continuous τ satisfies $\tau|_k = \mathsf{mst}_k(\tau(k))$ and thus $\tau|_k \cdot \tau(k+1) = \mathsf{mst}_{k+1}(\tau|_{k+1})$, for all $k \in \omega$. That is, continuous traces have *no trace jumps* — the property explaining the name. What makes 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 play 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 $\sigma_0, \ldots \sigma_n$, 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.

Please cite this article in press as: S. Enqvist et al., Completeness for the modal μ-calculus: Separating the combinatorics from the dynamics, Theoret. Comput. Sci. (2018), https://doi.org/10.1016/j.tcs.2018.03.001

Observe that the set $\{\mathsf{last}(\sigma_i|_k) \mid 0 \leq i \leq n\}$ is a subset of $\mathsf{Ran}(\mathsf{last}_{\Sigma|_k}) \subseteq A$, and so there must be indices $i \neq j$ such that $\mathsf{last}(\sigma_i|_k) = \mathsf{last}(\sigma_j|_k)$. But then it follows by the compatibility condition that $\sigma_i|_k = \mathsf{mst}_k(\mathsf{last}(\sigma_i|_k)) = \mathsf{mst}_k(\mathsf{last}(\sigma_j|_k)) = \sigma_j|_k$, which provides the desired contradiction. □

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 allows us to prove this is the following.

Proposition 9.16. *Let Σ be an infinite play 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. □

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 $\mathcal{S}_{thin}(\mathbb{A})$-play carries a bad trace, then it also carries a bad trace that is in addition continuous.

Proposition 9.17. *Let Σ be an infinite play of the thin satisfiability game for $\mathbb{A}$. Then for every bad trace $\tau \in \mathsf{Tr}_\Sigma$ there is a continuous trace $\widehat{\tau}$ 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 $\mathsf{mst}(\tau(k)) \cdot \tau(k+1) \neq \mathsf{mst}(\tau(k+1))$.

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 $\widehat{\tau} =_\infty \tau$.

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 $\rho = \mathsf{mst}_k(\tau(k))$. Then

$$\widehat{\tau} := \rho \cdot \sigma$$

is a continuous trace that is eventually equal to τ. Since it is obvious that $\widehat{\tau}$ is eventually equal to τ, it suffices to show that

$$\widehat{\tau}|_j = \mathsf{mst}(\widehat{\tau}(j)) \tag{37}$$

for all $j \in \omega$. First, by a simple induction we prove (37) for $j \geq k$. For $j = k$ we have $\mathsf{mst}_k(\tau(k)) = \widehat{\tau}|_k$ by definition, and assuming that the induction hypothesis holds for j we have:

$$\begin{aligned}
\widehat{\tau}|_{j+1} &= \widehat{\tau}|_j \cdot \widehat{\tau}(j+1) && \text{(obvious)}\\
&= \mathsf{mst}_j(\widehat{\tau}(j)) \cdot \widehat{\tau}(j+1) && \text{(induction hypothesis)}\\
&= \mathsf{mst}_j(\tau(j)) \cdot \tau(j+1) && \text{(definition } \widehat{\tau})\\
&= \mathsf{mst}_{j+1}(\tau(j+1)) && (j+1 \text{ not discontinuous})\\
&= \mathsf{mst}_{j+1}(\widehat{\tau}(j+1)) && \text{(definition } \widehat{\tau})
\end{aligned}$$

Second, for $j < k$ we have $\widehat{\tau}|_j = (\widehat{\tau}|_k)|_j = \rho|_j$, where $\rho = \mathsf{mst}_j(\tau(k))$, and so by Proposition 9.12 we obtain from $\rho \in V_k$ that $\widehat{\tau}|_j \in V_j$, meaning that $\widehat{\tau}|_j = \mathsf{mst}_j(\widehat{\tau}(j))$ as required. □

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 $f : \omega \to \omega$ by setting

$$f(i) := \mathsf{lrd}(\mathsf{mst}_i(\tau(i))).$$

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 $\mathsf{mst}_n(\tau(n)) \cdot \tau(n+1)$ is not refreshed either since it has the same last and next-to-last entries as $\tau|_{n+1}$. This means that:

$$\mathsf{lrd}\big(\mathsf{mst}_n(\tau(n)) \cdot \tau(n+1)\big) = \mathsf{lrd}\big(\mathsf{mst}_n(\tau(n))\big) = f(n).$$

Hence, if $f(n) < f(n+1)$, then we get

$$\mathsf{lrd}\big(\mathsf{mst}_n(\tau(n)) \cdot \tau(n+1)\big) < \mathsf{lrd}\big(\mathsf{mst}_{n+1}(\tau(n+1))\big),$$

and it immediately follows that

$$\mathsf{mst}_{n+1}(\tau(n+1)) \prec_{\Sigma|_{n+1}} \mathsf{mst}_n(\tau(n)) \cdot \tau(n+1).$$

But this directly contradicts the compatibility condition for the most significant trace. □

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 $g : \omega \to \mathsf{Tr}_q \times A$, setting:

$$g(n) := \big(\mathsf{mst}_n(\tau(n))|_q\ ,\ \mathsf{mst}_n(\tau(n))(q+1)\big)$$

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 $\mathsf{mst}_n(\tau(n))$ that will determine its place in the relevance order among traces in Tr_n: the place of the restricted trace $\mathsf{mst}_n(\tau(n))|_q$ in the relevance order at that stage, which is the last stage before the last refreshment date of $\mathsf{mst}_n(\tau(n))$, and the state visited by $\mathsf{mst}_n(\tau(n))$ at its last refreshment date. Since we already know what the last refreshment date of $\mathsf{mst}_n(\tau(n))$ is (namely $q+1$), these two pieces of information indeed suffice to determine the place of $\mathsf{mst}_n(\tau(n))$ in the relevance order.

We order the elements of $\mathsf{Tr}_q \times A$ lexicographically with respect to the relevance order and the priority order. More precisely put, we define the strict total order $\prec_q|\sqsubset$ on $\mathsf{Tr}_q \times A$ 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 $\mathsf{mst}_{n+1}(\tau(n+1)) = \mathsf{mst}_n(\tau(n)) \cdot \tau(n+1)$, and if $n > k_1 \geq q$ it follows that

$$\big(\mathsf{mst}_{n+1}(\tau(n+1))\big)|_q = \big(\mathsf{mst}_n(\tau(n))\big)|_q$$

and

$$\big(\mathsf{mst}_{n+1}(\tau(n+1))\big)(q+1) = \big(\mathsf{mst}_n(\tau(n))\big)(q+1)$$

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:

$$\mathsf{mst}_n(\tau(n)) \cdot \tau(n+1) \prec_{n+1} \mathsf{mst}_{n+1}(\tau(n+1)). \tag{38}$$

But the trace $\mathsf{mst}_n(\tau(n)) \cdot \tau(n+1)$ is not refreshed since $\tau|_{n+1}$ is not refreshed, so

$$\mathsf{lrd}\big(\mathsf{mst}_n(\tau(n)) \cdot \tau(n+1)\big) = f(n) = q+1,$$

and since $f(n+1) = q+1$ we have $\mathsf{lrd}\big(\mathsf{mst}_{n+1}(\tau(n+1))\big) = q+1$ by definition of the map f. Hence from the definition of the relevance order there are two possibilities for (38). Either we have

$$\big(\mathsf{mst}_n(\tau(n)) \cdot \tau(n+1)\big)|_q \prec_q \big(\mathsf{mst}_{n+1}(\tau(n+1))\big)|_q,$$

or else $\big(\mathsf{mst}_n(\tau(n)) \cdot \tau(n+1)\big)|_q = \big(\mathsf{mst}_{n+1}(\tau(n+1))\big)|_q$ and:

$$\big(\mathsf{mst}_n(\tau(n)) \cdot \tau(n+1)\big)(q+1) \sqsubset \big(\mathsf{mst}_{n+1}(\tau(n+1))\big)(q+1).$$

But since $n > q$ we have:

$$\big(\mathsf{mst}_n(\tau(n)) \cdot \tau(n+1)\big)|_q = \big(\mathsf{mst}_n(\tau(n))\big)|_q$$

and

$$\big(\mathsf{mst}_n(\tau(n)) \cdot \tau(n+1)\big)(q+1) = \big(\mathsf{mst}_n(\tau(n))\big)(q+1)$$

and so in each of the two cases we get $g(n) \prec_q|\sqsubset g(n+1)$ as required. □

It now follows that there can be only finitely many discontinuous points of τ above k_1, simply because the set $\mathsf{Tr}_q \times A$ 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. □

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 play; a Σ-*item* is defined as a pair $(\tau, a) \in \mathsf{Tr}_\Sigma \times A$ that is *selected* such that $\tau \in V_\Sigma$. A Σ-item is called *in focus* if $\tau = \mathsf{mst}_\Sigma(a)$, a μ-*item* if $\Omega(a)$ is odd, and a ν-*item* otherwise. The *most significant item* (MSI) of Σ, denoted msi_Σ, is defined to be the highest ranking Σ-item (τ, a) in the priority order $<_\Sigma$ which is in focus at Σ. ◁

Note that the MSI of Σ must be of the form $(\mathsf{mst}_\Sigma(a), a)$ for some $a \in \mathsf{Ran}(\mathsf{last}(\Sigma))$.

Definition 9.19. Let Σ be an infinite play 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 $S_\Sigma = \{(\sigma_1, d_1), \ldots, (\sigma_m, d_m)\}$ 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 $(\sigma_1|_j, d_1)$, ..., $(\sigma_m|_j, d_m)$ (in some fixed order). ◁

The framework of orderings $<_\Sigma$ associated with partial plays Σ is designed to make the following proposition true.

Proposition 9.20. *Let Σ be an infinite play 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 $\mathcal{S}_{thin}$-play 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 $j \geq k$:

$$(\rho|_j, b) <_j (\rho'|_j, b') \text{ iff } (\rho|_k, b) <_k (\rho'|_k, b'). \tag{39}$$

Proof of Claim. Fix two Σ-items $(\rho, b), (\rho', b')$ in F. If $\rho = \rho'$, then by the shuffle merge condition (39) 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

$$(\rho|_{k+1}, b) <_{k+1} (\rho'|_{k+1}, b') \text{ but } (\rho|_k, b) >_k (\rho'|_k, b'). \tag{40}$$

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 (40). From this the claim is immediate. □

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 plays. 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 $(\tau|_{j+1}, a) <^0_{j+1} (\sigma, d)$. 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 $(\sigma|_j \cdot \mathsf{last}(\sigma), d)$ with $(\tau|_j, a) <_{\Sigma|_j} (\sigma|_j, d)$. Furthermore by Proposition 9.12 we have that $\sigma|_j = \mathsf{mst}(\sigma(j+1))|_j = \mathsf{mst}(\sigma(j)) \in V_j$ 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 $(\sigma \cdot b_i, d) <^0_{j+1} (\tau, a)$. 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. □

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 $j \geq k$:

$$(\sigma|_j, d) <_j (\sigma'|_j, d') \text{ iff } (\sigma|_k, d) <_k (\sigma'|_k, d'). \tag{41}$$

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. □

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). □

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. □

9.3. Decorations

Our aim is now to define, by induction on the length of a partial play Σ, for every trace ρ on Σ two decorations $\delta_\rho, \delta^+_\rho : A \to \mu\mathtt{ML}$. 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 $\mathtt{tr}^{\rho}_{\mathbb{A}}(a)$ and $\mathtt{tr}^{\rho+}_{\mathbb{A}}(a)$ for $a \in A$ (check Definition 9.2) simply by:

$$\begin{aligned}\mathtt{tr}^{\rho}_{\Sigma}(a) &:= \mathtt{tr}^{\delta_\rho}_{\mathbb{A}}(a)\\ \mathtt{tr}^{\rho+}_{\Sigma}(a) &:= \mathtt{tr}^{\delta^+_\rho}_{\mathbb{A}}(a)\end{aligned}$$

In addition, the following abbreviations will help to avoid notational clutter:

$$\begin{aligned}\mathtt{tr}^{b}_{\Sigma}(a) &:= \mathtt{tr}^{\mathsf{mst}_\Sigma(b)}_{\Sigma}(a) \qquad & \mathtt{tr}_{\Sigma}(a) &:= \mathtt{tr}^{a}_{\Sigma}(a)\\ \mathtt{tr}^{b+}_{\Sigma}(a) &:= \mathtt{tr}^{\mathsf{mst}_\Sigma(b)+}_{\Sigma}(a) \qquad & \mathtt{tr}^{+}_{\Sigma}(a) &:= \mathtt{tr}^{a+}_{\Sigma}(a)\end{aligned}$$

The *context formula* for a finite partial play Σ is defined to be the formula

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

where a is the unique state such that for some trace τ, the pair (τ, a) is the MSI of Σ. ◁

We are now ready for the inductive definition of the decorations δ_ρ and δ^+_ρ, where ρ is a trace on a partial play Σ. Recall that msi_Σ is the highest priority Σ-item in focus.

Definition 9.22. For the unique trace ρ on the initial play consisting only of the relation $\{(a_I, a_I)\}$ we define $\delta_\rho(b) = \top$ for all $b \in A$. This ensures that $\mathtt{tr}^{\rho}_{\Sigma}(a) \equiv_K \mathtt{tr}_{\mathbb{A}}(a)$ for all $a \in A$.

Inductively, suppose that decorations δ_ρ and δ^+_ρ have been defined for all traces on the partial play Σ'. Let $\Sigma = \Sigma' \cdot Q$ be a continuation of Σ', and let $\rho = \sigma \cdot a$ be an arbitrary trace on Σ.

$$\delta_\rho(b) := \begin{cases} \top & \text{if } \Omega(b) \text{ is even}\\ \top & \text{if } b \sqsubset \mathsf{last}(\rho)\\ \top & \text{if } (\rho, b) <_\Sigma \mathsf{msi}_\Sigma\\ \delta^+_\sigma(b) & \text{if } \mathsf{msi}_\Sigma \leq_\Sigma (\rho, b).\end{cases}$$

In all cases, setting

$$\delta^+_\rho(b) := \begin{cases} \delta_\rho(b) \wedge \neg\gamma(\Sigma) & \text{if } (\rho, b) = \mathsf{msi}_\Sigma \text{ and } b \in A^\mu\\ \delta_\rho(b) & \text{otherwise.}\end{cases}$$

defines the decoration δ^+_ρ in terms of δ_ρ. ◁

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 play $\Sigma' \cdot Q$, and those items $(\sigma \cdot a, b)$ such that $\Omega(b)$ has a lower priority than $\mathsf{last}(\sigma \cdot a) = a$.

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 play in the (thin) satisfiability game for $\mathbb{A}$ such that the formula*

$$\bigwedge_{b \in \mathsf{Ran}(\mathsf{last}(\Sigma))} \mathtt{tr}_\Sigma(b)$$

is consistent. Then so is

$$\bigwedge_{b \in \mathsf{Ran}(\mathsf{last}(\Sigma))} \Theta(b)[\mathtt{tr}^{b}_{\Sigma}(d)^+/d \mid d \in A].$$

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:

$$\delta_\tau(a) \wedge \Theta(a)[\mathtt{tr}_{(\mathbb{A}\downarrow a)^-}(b)/b \mid b \sqsubset a][\mathtt{tr}^\tau_\Sigma(b)/b \mid a \sqsubset b].$$

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 — where we recall that we fixed the state $a \in \mathsf{Ran}(\mathsf{last}(\Sigma))$ in the beginning of the proof.

Claim 1. The state $a \in \mathsf{Ran}(\mathsf{last}(\Sigma))$ (given by $\mathsf{msi}_\Sigma = (\tau, a)$) satisfies

$$\theta[\mu a.\neg\gamma(\Sigma) \wedge \theta/a] \leq_K \Theta(a)[\mathtt{tr}^\tau_\Sigma(d)^+/d \mid d \in A].$$

Claim 2. All $b \neq a$ in the range of $\mathsf{last}(\Sigma)$ satisfy

$$\mathtt{tr}_\Sigma(b) \leq \Theta(b)[\mathtt{tr}^b_\Sigma(d)^+/d \mid d \in A].$$

The proof of both claims are given in the appendix of [19]. □

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 ∃ has a surviving strategy which maintains the consistency of the formula $\bigwedge_{b\in\mathsf{Ran}(\mathsf{last}(\Sigma))} \mathtt{tr}_\Sigma(b)$.

Proposition 9.24. *Let Σ be a partial play of length k in the (thin) satisfiability game for $\mathbb{A}$ such that the formula*

$$\bigwedge_{b\in\mathsf{Ran}(\mathsf{last}(\Sigma))} \mathtt{tr}_\Sigma(b)$$

is consistent. Then ∃ has a legitimate move $(\Upsilon, \mathcal{R})$ such that, for all $Q \in \mathcal{R}$, the formula

$$\bigwedge_{b\in\mathsf{Ran}(Q)} \mathtt{tr}_{\Sigma\cdot Q}(b)$$

is consistent.

Proof. It follows by Proposition 9.23 that the formula

$$\bigwedge_{b\in\mathsf{Ran}(\mathsf{last}(\Sigma))} \Theta(b)[\mathtt{tr}^b_\Sigma(d)^+/d \mid d \in A]$$

is consistent. By Corollary 5.14 we now find an admissible move $(\Upsilon, \mathcal{R})$ for ∃ such that, for all $Q \in \mathcal{R}$, the formula

$$(*) \quad \bigwedge_{(b,d)\in Q} \mathtt{tr}^b_\Sigma(d)^+$$

is consistent, so it suffices to show that this implies the required conjunction $\bigwedge_{b\in\mathsf{Ran}(Q)} \mathtt{tr}_{\Sigma\cdot Q}(b)$. For this, it suffices to show for all $d \in \mathsf{Ran}(Q)$ that the formula $\mathtt{tr}_{\Sigma\cdot Q}(d)$ is implied by some conjunct of $(*)$.

So let $d \in \mathsf{Ran}(Q)$. Then by Proposition 9.12 the trace $\mathsf{mst}_{\Sigma\cdot Q}(d)$ is of the form $\mathsf{mst}_\Sigma(d')\cdot d$ for some d' with $(d', d) \in Q$. So we either have $\delta_{\mathsf{mst}_{\Sigma\cdot Q}(d)} = \delta^+_{\mathsf{mst}_\Sigma(d')}$ or $\delta_{\mathsf{mst}_{\Sigma\cdot Q}(d)} = \top$. In both cases we have $\mathtt{tr}^{d'}_\Sigma(d)^+ \leq_K \mathtt{tr}_{\Sigma\cdot Q}(d)$, and so we are done. □

We are now ready for the proof of Theorem 5 itself.

Proof of Theorem 5. We shall define the winning strategy χ for ∃, and simultaneously maintain the induction hypothesis that for every partial χ-guided play Σ, the formula

$$\bigwedge\{\mathtt{tr}_\Sigma(a) \mid a \in \mathsf{Ran}(\mathsf{last}(\Sigma))\}$$

is consistent.

It is clear from Proposition 9.24 that ∃ has a strategy χ guaranteeing that she never gets stuck and that the induction hypothesis remains true for all χ-guided partial plays. It remains to show that this strategy is actually winning, i.e., that ∃ wins every infinite χ-guided play. So suppose for a contradiction that Σ is an infinite χ-guided play containing a bad trace. By Proposition 9.20, there is a continuous bad trace τ on Σ such that, with a denoting 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) $\mathsf{Ran}(\mathsf{last}(\Sigma|_j)) = \mathsf{Ran}(\mathsf{last}(\Sigma|_{j'}))$;
(2) $(\tau|_j, a) = \mathsf{msi}_{\Sigma|_j}$ and $(\tau|_{j'}, a) = \mathsf{msi}_{\Sigma|_{j'}}$;
(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 $\mathsf{Ran}(\mathsf{last}(\Sigma|_m))$ is an element of the finite set $\mathsf{P}A$, 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

$$\delta_{\sigma|_{n+1}}(d) = \delta_{\sigma|_n}(d) \tag{42}$$

for all $n \geq j$. But since we both have $\mathsf{msi}_{\Sigma|_n} \leq_{\Sigma|_n} (\tau|_n, a) <_{\Sigma|_n} (\sigma|_n, d)$ and $\mathsf{msi}_{n+1} \leq_{n+1} (\tau|_{n+1}, a) <_{n+1} (\sigma|_{n+1}, d)$, this is an immediate consequence of the definition of $\delta_{\tau|_{n+1}}$ from $\delta_{\tau|_n}$. □

On the basis of this we can prove the following key claim.

Claim 2. For any $b \in \mathsf{Ran}(\mathsf{last}(\Sigma|_j))$, we have

$$\mathtt{tr}_{\Sigma|_{j'}}(b) \leq_K \mathtt{tr}_{\Sigma|_j}(b).$$

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 $B := \mathsf{Ran}(\mathsf{last}(\Sigma|_j)) = \mathsf{Ran}(\mathsf{last}(\Sigma|_{j'}))$. Our proof of Claim 2 is based on the following observation.

$$\delta^b_{j'}(d) \leq_K \delta^b_j(d) \text{ for all } (b, d) \in B \times A, \tag{43}$$

where in order to avoid notational clutter, we abbreviate $\delta^b_j := \delta_{\mathsf{mst}_j(b)}$, $\delta^b_{j'} := \delta_{\mathsf{mst}_{j'}(b)}$, $\mathsf{mst}_j := \mathsf{mst}_{\Sigma|_j}$, and $\mathsf{mst}_{j'} := \mathsf{mst}_{\Sigma|_{j'}}$. In order to prove (43), make a case distinction.

(case 1) If $d \sqsubset b$ then (43) is immediate since $\delta^b_j(d) = \top$.
(case 2) If $b \sqsubseteq d$ and $(\mathsf{mst}_j(b), d) <_{\Sigma_j} (\tau, a)$ then (43) is immediate as well since $\delta^b_j(d) = \top$.
(case 3) If $b \sqsubseteq d$ and $(\tau, a) <_{\Sigma_j} (\mathsf{mst}_j(b), d)$ then $\mathsf{mst}_j(b)$ 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:

$$\mathsf{mst}_{j'}(b) = \sigma|_{j'}. \tag{44}$$

To see this, first observe that $\mathsf{last}(\sigma|_{j'}) = b$, since $\mathsf{last}(\sigma|_{j'}) = \sigma(j')$ by immediate unravelling of the definitions, $\sigma(j') = \sigma(j)$ by our earlier assumption on k, and $\sigma(j) = b$ by the fact that $\sigma(j) = \mathsf{last}(\sigma|_j) = \mathsf{last}(\mathsf{mst}_j(b)) = b$. But from $\mathsf{last}(\sigma|_{j'}) = b$ and the fact that σ is a continuous trace, we immediately obtain (44).
Finally we derive (43) as follows:

$$\begin{aligned} \delta^b_{j'}(d) &= \delta_{\sigma|_{j'}}(d) && \text{(immediate by (44))} \\ &= \delta_{\sigma|_j}(d) && \text{(Claim 1)} \\ &= \delta^b_j(d) && (\sigma|_j = \mathsf{mst}_j(b)) \end{aligned}$$

(case 4) If $b \sqsubseteq d$ and $(\mathsf{mst}_j(b), d) = (\tau, a)$ 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

$$\delta_{\tau|_l}(a) \leq_K \delta_{\tau|_j}(a) \text{ for all } l \geq j, \tag{45}$$

by a straightforward inductive proof, where the inductive step is taken care of by showing that

$$\delta_{\tau|_{n+1}}(a) \leq_K \delta_{\tau|_n} \text{ for all } n \geq j. \tag{46}$$

But this is not difficult to see, since $\mathsf{msi}_{\Sigma_n} \leq (\tau, a)$, for all $n \geq j$.

Finally, Claim 2 follows immediately from (43) by Proposition 9.3(3). □

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

$$\gamma(\Sigma|_j) = \bigwedge\{\mathtt{tr}_{\Sigma|_j}(b) \mid b \in \mathsf{Ran}(\mathsf{last}(\Sigma|_j)), b \neq a\}.$$

Since $\mathsf{Ran}(\mathsf{last}(\Sigma|_j)) = \mathsf{Ran}(\mathsf{last}(\Sigma|_{j'}))$, Claim 2 gives:

$$\bigwedge_{b \in \mathsf{Ran}(\mathsf{last}(\Sigma|_{j'}))\setminus\{a\}} \mathtt{tr}_{\Sigma|_{j'}}(b) \leq_K \gamma(\Sigma|_j). \tag{47}$$

On the other hand we claim that

$$\mathtt{tr}_{\Sigma|_{j'}}(a) \leq_K \neg\gamma(\Sigma|_j). \tag{48}$$

To see this, observe that since $(\tau|_j, a) = \mathsf{msi}_{\Sigma|_j}$ and since $\mathsf{msi}_{\Sigma|_{j''}}$ 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 $\mathtt{tr}_{\Sigma|_{j'}}$ is the tightened translation induced by $\delta_{\tau|_{j'}}$ and so by Proposition 9.3(2) we obtain

$$\mathtt{tr}_{\Sigma|_{j'}}(a) \leq_K \delta_{\tau|_{j'}}(a).$$

From this (48) is immediate.

Finally then, since $a \in \mathsf{Ran}(\mathsf{last}(\Sigma|_{j'}))$ we get from (47) and (48) that

$$\bigwedge_{b \in \mathsf{Ran}(\mathsf{last}(\Sigma|_{j'}))} \mathtt{tr}_{\Sigma|_{j'}}(b) \leq_K \gamma(\Sigma|_j) \wedge \neg\gamma(\Sigma|_j) \leq_K \bot,$$

directly contradicting the fact that the formula $\bigwedge_{b \in \mathsf{Ran}(\mathsf{last}(\Sigma|_{j'}))} \mathtt{tr}_{\Sigma|_{j'}}(b)$ is consistent. This finishes the proof of Theorem 5. □

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 8.16.) 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 $\mathbb{D}_\varphi$ 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. Recall that $\mathsf{sim}(\mathbb{A})$, the simulation of $\mathbb{A}$, is a disjunctive automaton that is equivalent to $\mathbb{A}$, in terms of the consequence game and, hence, also semantically.

Proposition 10.1. *Let $\mathbb{A}$ be any semi-disjunctive modal automaton. Then $\mathbb{A} \leq_K \mathsf{sim}(\mathbb{A})$.*

Proof. By Theorem 4 that there is a winning strategy for Player II in the consequence game $\mathcal{C}(\mathbb{A}, \mathsf{sim}(\mathbb{A}))$. Since $\mathbb{A}$ is semi-disjunctive it follows by Theorem 3 that $\forall$ has a winning strategy in the thin satisfiability game for $\mathbb{A} \wedge \neg\mathsf{sim}(\mathbb{A})$. By Kozen's Lemma (Theorem 5) it follows that the automaton $\mathbb{A} \wedge \neg\mathsf{sim}(\mathbb{A})$ is inconsistent. From this and the clauses 1 and 2 of Proposition 8.15, it is immediate that $\mathbb{A} \leq_K \mathsf{sim}(\mathbb{A})$. □

We are now ready for the statement and proof of our main result.

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

Proof. By Fact 3.14 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 $\mathbb{A}_\varphi$ 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 6.15(4) and 8.17 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 semi-disjunctive modulo provable equivalence by the Propositions 6.15(6) and 8.17, 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 $\mathbb{D} := \mathsf{sim}(\mu x.\mathbb{A})$. This automaton is clearly semantically equivalent to φ. We want to show that

$$\mu x.\mathbb{A} \leq_K \mathbb{D}, \tag{49}$$

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 (49) we will work with the automaton $\mathbb{A}^x$. First observe that

$$\mathbb{A}^x[\mathbb{D}/x] \models_{\mathsf{G}} \mathbb{A}^x[\mu x.\mathbb{A}/x],$$

by Theorem 4, and that

$$\mathbb{A}^x[\mu x.\mathbb{A}/x] \models_{\mathsf{G}} \mu x.\mathbb{A}$$

by Proposition 5.19. But since

$$\mu x.\mathbb{A} \models_{\mathsf{G}} \mathsf{sim}(\mu x.\mathbb{A}) = \mathbb{D}$$

by Theorem 4 again, we find by transitivity of the game consequence relation (Proposition 5.21) that

$$\mathbb{A}^x[\mathbb{D}/x] \models_{\mathsf{G}} \mathbb{D}.$$

By the Propositions 6.15(5) and 8.17 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

$$\mathtt{tr}(\mathbb{A}^x[\mathtt{tr}(\mathbb{D})/x]) \leq_K \mathtt{tr}(\mathbb{D}),$$

so that one application of the fixpoint rule yields

$$\mu x.\mathtt{tr}(\mathbb{A}^x) \leq_K \mathbb{D}.$$

By (34) in Proposition 8.14 this suffices to prove (49). □

Finally we see how Theorem 6 implies completeness.

Please cite this article in press as: S. Enqvist et al., Completeness for the modal μ-calculus: Separating the combinatorics from the dynamics, Theoret. Comput. Sci. (2018), https://doi.org/10.1016/j.tcs.2018.03.001

Theorem 10.2 *(Completeness). Every consistent formula $\varphi \in \mu\mathtt{ML}$ 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 6.14 $\exists$ 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. □

11. Future and related work

Recent related work Our hope is that the framework we have set up for reformulating and "deconstructing" Walukiewicz' completeness proof will help to fit it in a larger context, and there are many directions for future research here connected to existing work by various authors. First of all we mention that during the refereeing process of this paper, two publications presenting related work saw the light: the paper [1] by Afshari & Leigh presents a breakthrough result, proving completeness of a cut free and finitary proof system for the modal μ-calculus. This solves an important open problem for the μ-calculus, and also provides a fundamentally different proof of completeness for Kozen's system. This new development will no doubt open many paths for future research. In connection with the present work, we hope to use similar methods as we have developed here to analyze the proof methods of Afshari & Leigh, in the same spirit as what we have done with Walukiewicz' original proof. In particular we want to strengthen the connection between their (proof theoretic) approach with the game and automata theoretic perspective on the μ-calculus. We are optimistic that this can be done – while seemingly farther from our approach than Walukiewicz' proof, we would conjecture that the heart of Afshari & Leigh's proof is still about the combinatorics of traces, and that it can be reformulated in the language of automata, games, and "trace theory". Indeed, we believe that new insights of independent value will be uncovered by highlighting the trace combinatorics that is going on under the surface in their proof.

Another interesting contribution was made by Doumane [14], who gave a new completeness proof for the *linear-time* μ-calculus, i.e., the μ-calculus interpreted on models based on the natural numbers with the successor relation. Just like Afshari & Leigh's, Doumane's proof is constructive in the sense that it builds a proof for every valid formula rather than just showing existence. Her approach is congenial to ours in the sense that automata feature prominently in her proof. However, as Doumane points out herself, it does not seem to be straightforward to generalize her proof to the setting of the standard μ-calculus. Still, this clearly deserves further investigation.

The broader picture The modal μ-calculus is one system in a wide and varied family of logical formalisms that can collectively be called modal fixpoint logics. We believe that an important direction for future work is to provide complete axiomatizations of several fragment, extensions, and variations of the modal μ-calculus. An important example is Parikh's game logic, for which the problem of completeness is still open. More directly related examples are the μ-calculus with converse [63], the hybrid μ-calculus (which includes converse modalities in [57]), guarded fixpoint logic [28], and probabilistic variants as in e.g. [38]. We should mention here that a complete axiom system for the hybrid μ-calculus was recently presented by Tamura [62], 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 have already studied 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 was published in [18]. In a related paper [20] we have managed to extend this result further, covering systems like the monotone μ-calculus, i.e., the generalisation of the standard μ-calculus where formulas are interpreted in monotone neighbourhood frames.

Two additional possible directions of research in a similar spirit 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 [56] 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 that such systems can naturally turn up in certain contexts. See for example [26] for a study of the least fixpoint extension of the logic **S4** in a topological context.

12. List of symbols

In the table below we list the most important notations used in the paper, together with the location where this notation has been introduced.

Symbol	Introduction
$A^\sharp$	Definition 2.2
Dom, Ran, Res	Definition 2.2
$\overrightarrow{\mathsf{P}}, \overleftarrow{\mathsf{P}}, \overline{\mathsf{P}}$	Definition 2.3
$\exists, \forall$	Definition 2.6
$\mathsf{last}(\rho)$	Definition 2.7
$V^\dagger$	Definition 3.6
K	Remark 3.8
$[\![\varphi]\!]^{\mathbb{S}}$	Definition 3.9
$\equiv$	Definitions 3.10 and 4.14
$\mathbf{K}\mu, \vdash_K, \leq_K, \equiv_K$	Definition 3.12
∇	Definition 3.15
$\mathtt{Latt}(A)$	Definition 4.1
$\mathtt{1ML}(\mathtt{Prop}, A)$	Definition 4.1
$\Vdash^1$	Definition 4.3
$\Vdash^1_I$	Definition 4.4
$[\![\alpha]\!]^1$	Definition 4.4
$\underline{\leftrightarrow}^1$	Definition 4.5
$\underline{\to}^1$	Definition 4.7
$\mathcal{A}(\mathbb{A}, \mathbb{S}), [\![\mathbb{A}]\!]^{\mathbb{S}}$	Definition 4.11
$L(\mathbb{A})$	Definition 4.11

Symbol	Introduction
$\lhd, \bowtie$	Definition 4.12
$\sqsubset, \sqsubseteq$	Definition 4.13
$\leq$	Definition 4.14
$\equiv_1$	Definition 4.15
$\mathbb{A}^x$	Definition 4.19
Tr_Σ	Definition 5.1
$\max(\Omega[\mathit{Inf}(\tau)])$	Definition 5.3
n_a, U_a	Definition 5.5
$\Vdash^1_a$	Definition 5.5
$\mathcal{S}(\mathbb{A})$	Definition 5.6
$\mathcal{C}(\mathbb{A}, \mathbb{A}')$	Definition 5.16
$\models_{\mathsf{G}}$	Definition 5.16
$A^\sharp{}_{thin}$	Definition 6.5
$\mathcal{S}_{thin}(\mathbb{A})$	Definition 6.7
$\mathtt{1ML}_C(\mathtt{Prop}, A)$	Definition 6.12
$\mathsf{sim}(\mathbb{A})$	Definition 7.1
$\mathtt{tr}_{\mathbb{A}}$	Definition 8.7
$\mathtt{tr}$	Definition 8.10
$\mathtt{tr}^\delta_{\mathbb{A}}$	Definition 9.2
$\prec$	Definition 9.7

Appendix A. Proof details of section 7

Proof of Theorem 4(1). Fix modal automata $\mathbb{A} = (A, \Theta_A, \Omega_A, a_I)$ and $\mathbb{B} = (B, \Theta_B, \Omega_B, b_I)$ as well as the stream automaton $\mathbb{Z} = (Z, \zeta, z_I, \Omega_Z)$ that recognizes the stream language NBT_Ω over the alphabet $A^\sharp$. Furthermore, let $D = A^\sharp \times Z$ be the carrier set of the simulating automaton $\mathsf{sim}(\mathbb{A})$, then each state of $\mathsf{sim}(\mathbb{A})$ is of the shape (Q, z) where $z \in Z$ and $Q \in A^\sharp$. We construct a functional winning strategy for Player II in $\mathcal{C}(\mathbb{B}[\mathbb{A}/p], \mathbb{B}[\mathsf{sim}(\mathbb{A})/p])$ as follows.

Let Σ be a given partial play of $\mathcal{C}(\mathbb{B}[\mathbb{A}/p], \mathbb{B}[\mathsf{sim}(\mathbb{A})/p])$ ending with the basic position (R, R') belonging to Player I, and suppose that Player I plays the move $\Gamma = (\Upsilon, \mathcal{R}) \in \mathsf{K}(B \cup A)^\sharp$. Without loss of generality, this move is such that $\mathsf{Dom}\, Q \subseteq \mathsf{Ran}\, R$ for each $Q \in \mathcal{R}$. Furthermore we shall inductively maintain the following condition on basic positions (Q, Q'):

(†) $\mathsf{Ran}\, Q \supseteq (\mathsf{Ran}\, Q' \cap B) \cup \bigcup\{\mathsf{Ran}\, S \mid \exists z : (S, z) \in \mathsf{Ran}\, Q' \cap D\}$.

We assume that this condition has been maintained for all basic positions in the play Σ, including (R, R') in particular. The response by Player II following the play $\Sigma \cdot \Gamma$ will be determined by a map $F : \mathcal{R} \to (B \cup D)^\sharp$, such that for a given $Q \in \mathcal{R}$, we put $FQ := Q'_1 \cup Q'_2 \cup Q'_3$, where each these three relations Q'_1, Q'_2, Q'_3 will be defined separately. First, we set $Q'_1 := \mathsf{Res}_B\, Q$ (where we recall the definition Res from Definition 2.2). Next, we set:

$$Q'_2 := \{(b, (\{a_I\} \times (Q[b] \cap A), \zeta(a'_I, z_I)) \mid b \in \mathsf{Ran}\, R \cap B\},$$

where $a'_I = \{(a_I, a_I)\}$, the initial state of the pre-simulation of $\mathbb{A}$. It could be helpful to stop and check the types at this point: note that indeed $\{a_I\} \times (Q[b] \cap A) \in A^\sharp$, so $(\{a_I\} \times (Q[b] \cap A), \zeta(a'_I, z_I)) \in D$, hence $Q'_2 \in (B \cup D)^\sharp$ as expected. Last, we set:

$$Q'_3 := \{((S, z), (\mathsf{Res}_A\, Q, \zeta(S, z))) \in D \times D \mid (S, z) \in \mathsf{Ran}\, R' \cap D\}.$$

The intuition behind the definition of FQ is the following: Q itself can be split into four parts: a $B \times B$-part, a $B \times A$-part, an $A \times A$-part and an $A \times B$-part. The last part can be considered "junk" which is not really required for Player I's move to be legitimate, so we ignore it and consider Q to be split into the first three parts. The first part is then mapped straightforwardly to Q'_1. The second part is mimicked by considering, for each $b \in \mathsf{Ran}\, R \cap B$, the set of successors $Q[b] \cap A$ of b in A as a macro-state in D, which is formally achieved turning it into a binary relation with domain $\{a_I\}$ and tagging on the appropriate state of $\mathbb{Z}$. The idea is that for each $b \in \mathsf{Ran}\, R$ with $Q[b] \cap A \neq \emptyset$ the one-step model:

$$(\Upsilon, \{(\{a_I\} \times (Q[b] \cap A), \zeta(a'_I, z_I)) \mid Q \in \mathcal{R}\})$$

satisfies the one-step formula of the initial state d_I of $\mathsf{sim}(\mathbb{A})$. Finally, for the $A \times A$-part of Q, we note that by our condition (†) each $(S, z) \in \mathsf{Ran}\, R' \cap D$ corresponds to the subset $\mathsf{Ran}\, S \subseteq \mathsf{Ran}\, R \cap A$, so the one-step model $(\Upsilon, \mathcal{R})$ chosen by Player I satisfies all one-step formulas for elements in $\mathsf{Ran}S$ since it satisfies all the one-step formulas for elements of $\mathsf{Ran}R$. This means that for each $(S, z) \in \mathsf{Ran}R' \cap D$ the one-step model:

$$(\Upsilon, \{\{((S, z), (\mathsf{Res}_A Q, \zeta(S, z)))\} \mid Q \in \mathcal{R}\})$$

satisfies the one-step formula assigned to the state (S, z) in $\mathsf{sim}(\mathbb{A})$, and this is the motivation for the definition of Q'_3.

We leave it as an exercise for the reader to check that the map F indeed determines a legitimate move for Player II, i.e. that the one-step model $(\mathsf{K}F)\Gamma$ satisfies all the one-step formulas associated with states in $\mathsf{Ran}R'$. Furthermore, it is clear that the condition (†) is preserved for any next basic position chosen by Player I. It remains only to check that every infinite match guided by this strategy is won by Player II; so consider such an infinite match with basic positions:

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

Suppose there is a bad trace τ on the $(B \cup D)^\sharp$-stream $R'_0 R'_1 R'_2 \dots$ Then given how the maps F have been constructed, it is clear that this trace either stays in B forever, or at some point goes into the D-part of $\mathbb{B}[\mathsf{sim}(\mathbb{A})/p]$ and stays there forever. In the first case, τ is itself a bad trace on the $(B \cup A)^\sharp$-stream $R_0 R_1 R_2 \dots$ In the second case, the trace τ must be of the form:

$$b_0 b_1 \dots b_k \cdot ((\{a_I\} \times (R_{k+1}[b] \cap A)), \zeta(a'_I, z_I)) \cdot (S_1, z_1)(S_2, z_2)(S_3, z_3)\dots$$

But if this trace is bad then the stream $S_1 S_2 S_3 S_4 \dots$ has a bad trace. The strategy of Player II ensures that $S_i \subseteq R_{k+1+i}$ for each i with $1 \le i < \omega$, hence the stream $R_{k+2} R_{k+3} R_{k+4} \dots$ contains a bad trace. Since we assumed that $\mathsf{Dom}R_{j+1} \subseteq \mathsf{Ran}R_j$ for all $j < \omega$, this means that $R_0 R_1 R_2 \dots$ indeed contains a bad trace, so we are done. □

References

[1] B. Afshari, G. Leigh, Cut-free completeness for modal mu-calculus, in: Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic In Computer Science (LICS'17), IEEE Computer Society, pp. 1–12.

[2] A. Arnold, D. Niwiński, Rudiments of μ-Calculus, Studies in Logic and the Foundations of Mathematics, vol. 146, North-Holland Publishing Co., Amsterdam, 2001.

[3] J.W. de Bakker, 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] P. Blackburn, M. de Rijke, Y. Venema, Modal Logic, Cambridge Tracts in Theoretical Computer Science, vol. 53, Cambridge University Press, 2001.

[6] A. Blass, A game semantics for linear logic, Ann. Pure Appl. Logic 56 (1–3) (1992) 183–220.

[7] J. Bradfield, C. Stirling, Modal μ-calculi, in: J. van Benthem, P. Blackburn, F. Wolter (Eds.), Handbook of Modal Logic, Elsevier, 2006, pp. 721–756.

[8] J.R. Büchi, Weak second-order arithmetic and finite automata, Math. Log. Q. 6 (1–6) (1960) 66–92.

[9] J.R. Büchi, L.H. Landweber, Solving sequential conditions by finite state strategies, Trans. Amer. Math. Soc. 138 (1969) 295–311.

[10] C. Cîrstea, D. Pattinson, Modular construction of modal logics, in: Ph. Gardner, N. Yoshida (Eds.), Proceedings of the 15th International Conference on Concurrency Theory, CONCUR 2004, in: LNCS, vol. 3170, Springer, 2004, pp. 258–275.

[11] E.M. Clarke, E.A. Emerson, Design and synthesis of synchronisation skeletons using branching time temporal logic, in: D. Kozen (Ed.), Logics of Programs, Springer, 1981, pp. 52–71.

[12] G. D'Agostino, M. Hollenberg, Logical questions concerning the μ-calculus, J. Symbolic Logic 65 (2000) 310–332.

[13] M. Dam, CTL* and ECTL* as fragments of the modal μ-calculus, Theoret. Comput. Sci. 126 (1) (1994) 77–96.

[14] A. Doumane, Constructive completeness for the linear-time μ-calculus, in: Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic In Computer Science, LICS'17, IEEE Computer Society, 2017.

[15] E. Emerson, J.Y. Halpern, Decision procedures and expressiveness in the temporal logic of branching time, in: Proceedings of 14th ACM Symposium on Theory of Computing, 1982, pp. 169–180.

[16] E.A. Emerson, J. Halpern, "Sometimes" and "not never" revisited: on branching versus linear time temporal logic, J. ACM 33 (1) (1986) 151–178.

[17] E.A. Emerson, C.S. Jutla, Tree automata, mu-calculus and determinacy (extended abstract), in: Proceedings of the 32nd Symposium on the Foundations of Computer Science, IEEE Computer Society Press, 1991, pp. 368–377.

[18] S. Enqvist, F. Seifan, Y. Venema, Completeness for coalgebraic fixpoint logic, in: Proceedings of the 25th EACSL Annual Conference on Computer Science Logic, CSL 2016, in: LIPIcs, vol. 62, 2016, pp. 1–19.

[19] S. Enqvist, F. Seifan, Y. Venema, Completeness for the Modal μ-Calculus: Separating the Combinatorics from the Dynamics, Technical Report PP-2016-33, Institute for Logic, Language and Computation, Universiteit van Amsterdam, 2016.

[20] S. Enqvist, F. Seifan, Y. Venema, Completeness for Modal μ-Calculi: A Coalgebraic Approach, Technical Report PP-2017-04, Institute for Logic, Language and Computation, Universiteit van Amsterdam, 2017, submitted for publication.

[21] S. Enqvist, F. Seifan, Y. Venema, An expressive completeness theorem for coalgebraic modal mu-calculi, Log. Methods Comput. Sci. 13 (2) (2017).

[22] G. Fontaine, R. Leal, Y. Venema, Automata for coalgebras: an approach using predicate liftings, in: Automata, Languages and Programming: 37th International Colloquium, ICALP'10, in: LNCS, vol. 6199, Springer, 2010, pp. 381–392.

[23] G. Fontaine, Y. Venema, Some model theory for the modal μ-calculus: syntactic characterisations of semantic properties, Log. Methods Comput. Sci. 14 (1) (2018), https://doi.org/10.23638/LMCS-14(1:14)2018.

[24] D.M. Gabbay, A. Pnueli, S. Shelah, J. Stavi, On the temporal analysis of fairness, in: Proc. 7th ACM Symposium on Principles of Programming Languages, 1980, pp. 163–173.

[25] S. Ghilardi, An algebraic theory of normal forms, Ann. Pure Appl. Logic 71 (3) (1995) 189–245.

[26] R. Goldblatt, I. Hodkinson, Spatial logic of modal mu-calculus and tangled closure operators, arXiv preprint, arXiv:1603.01766, 2016.

[27] E. Grädel, W. Thomas, T. Wilke (Eds.), Automata, Logic, and Infinite Games, LNCS, vol. 2500, Springer, 2002.

[28] E. Grädel, I. Walukiewicz, Guarded fixed point logic, in: Proceedings. 14th Symposium on Logic in Computer Science, 1999, IEEE, 1999, pp. 45–54.

[29] G. Jäger, M. Kretz, T. Studer, Canonical completeness of infinitary mu, J. Log. Algebr. Program. 76 (2) (2008) 270–292.

[30] D. Janin, I. Walukiewicz, Automata for the modal μ-calculus and related results, in: J. Wiedermann, P. Hájek (Eds.), Mathematical Foundations of Computer Science 1995, 20th International Symposium, MFCS'95, in: LNCS, vol. 969, Springer, 1995, pp. 552–562.
[31] D. Janin, I. Walukiewicz, On the expressive completeness of the propositional μ-calculus w.r.t. monadic second-order logic, in: Proceedings CONCUR '96, 1996.
[32] A. Joyal, Free lattices, communication and money games, in: Proceedings of the 10th International Congress of Logic, Methodology, and Philosophy of Science, 1995.
[33] R. Kaivola, Axiomatising linear time mu-calculus, in: I. Lee, S.A. Smolka (Eds.), Proceedings of the 6th International Conference on Concurrency Theory, CONCUR '95, in: LNCS, vol. 962, Springer, 1995, pp. 423–437.
[34] D. Kozen, Results on the propositional μ-calculus, Theoret. Comput. Sci. 27 (1983) 333–354.
[35] D. Kozen, A finite model theorem for the propositional μ-calculus, Studia Logica 47 (1988) 233–241.
[36] D. Kozen, R. Parikh, An elementary proof of the completeness of PDL, Theoret. Comput. Sci. 14 (1981) 113–118.
[37] M. Lange, 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, IEEE Computer Society, 2001, pp. 357–365.
[38] M. Mio, Probabilistic modal μ-calculus with independent product, in: International Conference on Foundations of Software Science and Computational Structures, Springer, 2011, pp. 290–304.
[39] L. Moss, Coalgebraic logic, Ann. Pure Appl. Logic 96 (1999) 277–317, Ann. Pure Appl. Logic 99 (1999) 241–259 (Erratum).
[40] A.M. Mostowski, Games with Forbidden Positions, Technical Report 78, Instytut Matematyki, Uniwersytet Gdański, Poland, 1991.
[41] A.W. Mostowski, Regular expressions for infinite trees and a standard form of automata, in: A. Skowron (Ed.), Computation Theory, in: LNCS, Springer-Verlag, 1984, pp. 157–168.
[42] D.E. Muller, P.E. Schupp, Alternating automata on infinite trees, Theoret. Comput. Sci. 54 (1987) 267–276.
[43] D.E. Muller, P.E. Schupp, Simulating alternating tree automata by nondeterministic automata, Theoret. Comput. Sci. 141 (1995) 69–107.
[44] D. Niwínski, I. Walukiewicz, Games for the μ-calculus, Theoret. Comput. Sci. 163 (1996) 99–116.
[45] R. Parikh, The logic of games and its applications, Ann. Discrete Math. 24 (1985) 111–139.
[46] D.P.R. Park, Fixpoint induction and proof of program semantics, Mach. Intell. 5 (1970) 59–78.
[47] D. Pattinson, Coalgebraic modal logic: soundness, completeness and decidability of local consequence, Theoret. Comput. Sci. 309 (1–3) (2003) 177–193.
[48] A.M. Pitts, On an interpretation of second-order quantification in first-order intuitionistic propositional logic, J. Symbolic Logic 57 (1) (1992) 33–52.
[49] A. Pnueli, The temporal logic of programs, in: Proc. 18th Symp. Foundations of Computer Science, 1977, pp. 46–57.
[50] V.R. Pratt, Semantical considerations on Floyd-Hoare logic, in: Proc. 17th IEEE Symposium on Computer Science, 1976, pp. 109–121.
[51] V.R. Pratt, A decidable μ-calculus: preliminary report, in: Proc. 22th IEEE Symposium on Foundations of Computer Science, 1981, pp. 421–427.
[52] M.O. Rabin, Decidability of second-order theories and automata on infinite trees, Trans. Amer. Math. Soc. 141 (1969) 1–35.
[53] M. Reynolds, An axiomatization of full computation tree logic, J. Symbolic Logic (2001) 1011–1057.
[54] L. Santocanale, Free μ-lattices, J. Pure Appl. Algebra 168 (2–3) (2002) 227–264.
[55] L. Santocanale, Completions of μ-algebras, Ann. Pure Appl. Logic 154 (1) (2008) 27–50.
[56] L. Santocanale, Y. Venema, Completeness for flat modal fixpoint logics, Ann. Pure Appl. Logic 162 (2010) 55–82.
[57] U. Sattler, M.Y. Vardi, The hybrid μ-calculus, in: International Joint Conference on Automated Reasoning, Springer, 2001, pp. 76–91.
[58] L. Schröder, D. Pattinson, Pspace bounds for rank-1 modal logics, ACM Trans. Comput. Log. 10 (2) (2009) 13.
[59] L. Schröder, D. Pattinson, Rank-1 modal logics are coalgebraic, J. Logic Comput. 20 (5) (2010) 1113–1147.
[60] E. Spaan, Complexity of Modal Logics, PhD thesis, ILLC, University of Amsterdam, 1993.
[61] R.S. Street, E.A. Emerson, An automata theoretic decision procedure for the propositional mu-calculus, Inform. and Comput. 81 (1989) 249–264.
[62] K. Tamura, A small model theorem for the hybrid μ-calculus, J. Logic Comput. (2013), exs052.
[63] M.Y. Vardi, Reasoning about the past with two-way automata, in: International Colloquium on Automata, Languages, and Programming, Springer, 1998, pp. 628–641.
[64] Y. Venema, Automata and fixed point logic: a coalgebraic perspective, Inform. and Comput. 204 (2006) 637–678.
[65] Y. Venema, Lectures on the Modal μ-Calculus, Lecture Notes, Institute for Logic, Language and Computation, University of Amsterdam, 2012 (manuscript).
[66] A. Visser, Uniform interpolation and layered bisimulation, in: Gödel '96, Brno, in: Lecture Notes in Logic, vol. 6, 1996, pp. 139–164.
[67] I. Walukiewicz, On completeness of the mu-calculus, in: Proceedings of the Eighth Annual Symposium on Logic in Computer Science, LICS '93, IEEE Computer Society, 1993, pp. 136–146.
[68] 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, IEEE Computer Society Press, 1995, pp. 14–24.
[69] I. Walukiewicz, Completeness of Kozen's axiomatisation of the propositional μ-calculus, Inform. and Comput. 157 (2000) 142–182.
[70] T. Wilke, Alternating tree automata, parity games, and modal μ-calculus, Bull. Belg. Math. Soc. 8 (2001) 359–391.

Please cite this article in press as: S. Enqvist et al., Completeness for the modal μ-calculus: Separating the combinatorics from the dynamics, Theoret. Comput. Sci. (2018), https://doi.org/10.1016/j.tcs.2018.03.001