Information and Computation ••• (••••) ••••••

Contents lists available at ScienceDirect

Information and Computation

www.elsevier.com/locate/yinco

Bounded game-theoretic semantics for modal mu-calculus

Lauri Hella [a], Antti Kuusisto [a,b,*], Raine Rönnholm [a,c]

[a] Tampere University, Tampere, Finland
[b] University of Helsinki, Helsinki, Finland
[c] Université Paris-Saclay, CNRS, ENS Paris-Saclay, Gif-sur-Yvette, France

ARTICLE INFO

Article history:
Received 15 January 2021
Received in revised form 27 January 2022
Accepted 16 February 2022
Available online xxxx

Keywords:
Mu-calculus
Game-theoretic semantics
Alternative semantics
Bounded semantics
Model-checking
Fixed point approximants
Closure ordinals

ABSTRACT

We introduce a new game-theoretic semantics (GTS) for the modal mu-calculus. Our so-called bounded GTS replaces parity games with alternative evaluation games where only finite paths arise; infinite paths are not needed even when the considered transition system is infinite. The novel games offer alternative approaches to various constructions in the framework of the mu-calculus. While our main focus is introducing the new GTS, we also consider some applications to demonstrate its uses. For example, we consider a natural model transformation procedure that reduces model checking games to checking a single, fixed formula in the constructed models. We also use the GTS to identify new alternative variants of the mu-calculus, including close variants of the logic with PTime model checking; variants with iteration limited to finite ordinals; and other systems where the semantic or syntactic specification of the mu-calculus has been modified in a natural way suggested by the GTS.

1. Introduction

The modal μ-calculus [26] is a well-known formalism that plays a central role in, e.g., program verification. The standard semantics of the μ-calculus is based on fixed points, but the system has also a well-known game theoretic semantics (GTS) that makes use of parity games. The related games generally involve infinite plays, and the parity condition is used for determining the winner (see, e.g., [6] for further details and a general introduction to the μ-calculus).

The agenda and contributions of this article. In this article we present an alternative game-theoretic semantics for the modal μ-calculus. Our so-called *bounded* GTS is based on games that resemble the standard semantic games for the μ-calculus, but there is an extra feature that ensures that the plays within the novel framework always end after a finite number of rounds. Thereby only finite paths arise in related evaluation games *even when investigating infinite transition systems*. Thus there is no need to keep track of the parity condition, so in that sense the games we present in this article simplify the standard framework. Furthermore, they offer an alternative perspective on the modal μ-calculus, as we show that our semantics is equivalent to the standard one.

In the new games, the evaluation of a fixed point formula begins by one of the players declaring an ordinal number; the verifying player declares ordinals for μ-formulae and the falsifying player for ν-formulae. The declared ordinal is then lowered as the game proceeds. Since ordinals are well-founded, the game will indeed end after a finite number of game steps. In general, infinite ordinals are needed in the games, but finite ordinals suffice in finite models.

* Corresponding author.
E-mail address: antti.kuusisto@tuni.fi (A. Kuusisto).

https://doi.org/10.1016/j.ic.2022.104882
0890-5401/

Please cite this article as: L. Hella, A. Kuusisto and R. Rönnholm, Bounded game-theoretic semantics for modal mu-calculus, Information and Computation, https://doi.org/10.1016/j.ic.2022.104882

While the bounded GTS provides a new perspective on the standard modal μ-calculus, our approach also leads naturally to a range of alternative semantic systems that are not equivalent to the standard semantics. Indeed, we divide the framework of bounded semantics into subsystems dubbed Γ-*bounded semantics* for different ordinals Γ. Here Γ provides a strict upper limit for the ordinals that can be used during the game play. For each Γ-bounded semantics, we define also a corresponding bounded compositional semantics and prove the game-theoretic and compositional versions equivalent.

If only finite ordinals are allowed, meaning $\Gamma = \omega$, we obtain the *finitely bounded* GTS, which is an interesting system itself. While this semantics is equivalent to the standard case in finite models, the general expressive powers differ. Indeed, we will show that the μ-calculus under finitely bounded GTS does not have the finite model property. Furthermore, we observe that the set of validities of the μ-calculus under finitely bounded semantics is strictly contained in the set of standard validities.

We then study a range of variants of the standard bounded semantics. In our standard bounded semantics, each μ and ν-formula is associated with an ordinal of its own, while in *simple bounded semantics* this scheme is relaxed and only two ordinals are used—one for all μ-formulae and another one for all ν-formulae. We identify conditions under which this new semantics is equivalent to the standard one when restricting to the ν-free fragment of the μ-calculus. Furthermore, we prove PTime-completeness of the model checking problem for a relatively expressive variant of the μ-calculus using this new semantics.[1] The result concerns both data and combined complexity. We also introduce a system of *non-well-founded* modal logic as a canonical extension of modal logic with game-theoretic recursion and also relate the system to the simple bounded semantics of the modal μ-calculus. We call this logic *modal computation logic* MCL and ultimately prove that it has PTime-complete model checking.

As a further semantic variant of the μ-calculus, we study a version of the bounded GTS where only one of the two players is required to announce and commit to ordinal values in the semantic games. We show that these *semi-bounded* systems of semantics are essentially equivalent to the standard bounded semantics. And, in addition to semantic studies, we use GTS to identify a canonical reduction of μ-calculus model checking instances to checking a single, uniform formula in the model obtained by the reduction. Furthermore, we define a natural hierarchy of variants of the reachability game and relate our different semantic games to those variants.

Further motivation of the study. While the formal results listed above are an important part of our study, the focus of our article is mainly on the *conceptual development* of the theory of the modal μ-calculus and related systems, not so much the more technical directions. While some of our technical observations have *obvious implicit similarities* to existing results and notions, we believe the systematic, formal and conceptual study initiated in this article is justified.

Indeed, we believe the bounded GTS in general can be a fruitful framework for various further developments. The setting provides an alternative perspective to parity games, replacing infinite plays with games based on finitely many rounds only, thereby leading to a conceptually interesting territory to be explored further. The fragments with PTime-model checking we identify serve as an example of the various possibilities. Furthermore, it is worth noting here, e.g., that the difference between the standard and bounded GTS for the modal μ-calculus is analogous to the relationship between while-loops and for-loops; while-loops are iterated possibly infinitely long, whereas for-loops run for $k \in \mathbb{N}$ rounds, where k can generally be an input to the loop. Finally—while it can be called into question and of course considered a matter of opinion—we argue that the new bounded semantics often makes formulae easier to read. This relates to the semantics pointing to games that always end after a finite number of steps; the clocks concretely indicate, step-by-step, how the game approaches the set of end positions. We will illustrate the naturalness of the semantics in a concrete way in Examples 4.3 and 4.9. We note that the semantics could possibly be beneficial in teaching the modal μ-calculus.

Notes on related work. There already exist several works where simple variants of the bounded semantics have been considered in the context of temporal logics with a significantly simpler recursion mechanisms than that of the μ-calculus. The papers [13], [16] consider a bounded semantics for the Alternating-time temporal logic ATL, and [15], [12] extend the related study to the extension ATL^+ of ATL. See also [14], [17]. Part of the original motivation behind the studies in [13], [16], [15], [12] (as well as the current article) relates to work with the direct aim of understanding variants and fragments of the general, expressively Turing-complete logic first presented in [28]. It is also worth mentioning that the work in the present article has already been made essential use of in constructing a canonical formula size game for the μ-calculus in [20]. The first short draft of the current submission appeared in 2017 as the manuscript [18]. It contained only the bounded GTS presented below. That paper was expanded to the conference paper [19], and the current article is the significantly extended journal version of [19].

There is a whole range of earlier but closely related logical studies that make use of notions with similar intuitions to the ones behind the bounded semantics of this paper. Indeed, for logics with time bounds, see, e.g., the paper [2] on finitary fairness and the article [27] relating to promptness in Linear temporal logic LTL. We also mention here the work related to *bounded model checking*, see, e.g., [5], [32] and [36]. The article [8] is one example of an early work that uses explicit 'clocking' of fixed point formulae in (a variant of) the μ-calculus, thereby involving some ideas that bear a similarity to some features used also in the present paper. We note that the approach and goals of [8] are different, e.g., the paper limits

[1] It has been show that the model checking of the modal μ-calculus can be done in quasipolynomial time ([9]), but it is a famous open problem whether the model checking could be done in PTime. See [7] for more on the model checking of μ-calculus.

to finite models only and does not discuss game-theoretic semantics. However, using clocks for bounding unfolding of fixed points goes back to [34], and this idea is used to link alternating parity automata to the μ-calculus in [11].

The notion of *clock bounds* in our GTS relates naturally to the study of *closure ordinals*.[2] This is demonstrated, e.g., by the proof of Theorem 5.2 below. Closure ordinals have of course received a lot of attention within the research of μ-calculus; see e.g., [1] and [10] for some relatively recent work on the topic. It is also worth noting here that game theoretic approach has also been used for studying the satisfiability of the modal μ-calculus; see e.g. [31].

2. Preliminaries

2.1. Syntax

Let Φ be a set of *proposition symbols* and Λ a set of *label symbols*. Formulae of the modal μ-calculus are generated by the grammar

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

where $p \in \Phi$ and $X \in \Lambda$.

Let φ be a formula of the μ-calculus. The set of nodes in the syntax tree of φ is denoted by $\mathrm{Sf}(\varphi)$. All of these nodes correspond to some subformula of φ, but the same subformula may have several occurrences in the syntax tree of φ, as for example in the case of $p \vee p$. We always distinguish between different occurrences of the same subformula, and thus we always assume that the position in the syntax tree of φ is known for any given subformula of φ. We also use the following notation:

$$\mathrm{Sf}_{\mu\nu}(\varphi) := \{\theta \in \mathrm{Sf}(\varphi) \mid \theta = \mu X\psi \text{ or } \theta = \nu X\psi \text{ for some } \psi \in \mathrm{Sf}(\varphi) \text{ and } X \in \Lambda\}.$$

We write $|\varphi|$ for $|\,\mathrm{Sf}(\varphi)|$.

2.2. Standard compositional semantics

A Kripke model $\mathcal{M}$ is a tuple (W, R, V), where W is a nonempty set, R a binary relation over W and $V : \Phi \to \mathcal{P}(W)$ a valuation for proposition symbols in Φ. An *assignment* $s : \Lambda \to \mathcal{P}(W)$ for $\mathcal{M}$ maps label symbols X to subsets of W.

Definition 2.1. Let $\mathcal{M} = (W, R, V)$ be a Kripke model, $w \in W$. Let φ be a formula of the μ-calculus. *Truth of φ in $\mathcal{M}$ at w under assignment s*, denoted by $\mathcal{M}, w \vDash_s \varphi$, is defined as in standard modal logic for p, $\neg p$, $\vee$, $\wedge$, $\Diamond$, $\Box$. The truth condition for label symbols is defined with respect to the assignment s:

- $\mathcal{M}, w \vDash_s X$ iff $w \in s(X)$.

To deal with μ and ν, we define an operator $\widehat{\varphi}_{X,s} : \mathcal{P}(W) \to \mathcal{P}(W)$ such that $\widehat{\varphi}_{X,s}(A) = \{w \in W \mid \mathcal{M}, w \vDash_{s[A/X]} \varphi\}$, where $s[A/X]$ is the assignment that sends X to A and treats other label symbols the same way as s. The operators $\widehat{\varphi}_{X,s}$ are always monotone and thereby have least and greatest fixed points. The semantics for the operators μX and νX is as follows:

- $\mathcal{M}, w \vDash_s \mu X\psi$ iff w is in the least fixed point of the operator $\widehat{\psi}_{X,s}$.
- $\mathcal{M}, w \vDash_s \nu X\psi$ iff w is in the greatest fixed point of the operator $\widehat{\psi}_{X,s}$.

An occurrence of a label symbol X is said to be *free* in a formula φ if that occurrence of X is a leaf (in the syntax tree of φ) which is not a descendant of any node corresponding to a formula of the form $\mu X\psi$ or $\nu X\psi$. A formula φ is called a *sentence* if it does not contain any free occurrences of label symbols. Truth of a sentence φ is independent of assignments s, so we may simply write $\mathcal{M}, w \vDash \varphi$ instead of $\mathcal{M}, w \vDash_s \varphi$.

2.3. Reachability and safety games

All of the games defined below are played on a pointed Kripke model $(\mathcal{M}, w)$ whose vocabulary contains proposition symbols p_A, p_B and q_B. These games are played by two players, A and B, starting from the state w. In each round, one of the players moves (if possible) to some state that can be directly reached in one step from the current state via the accessibility relation of $\mathcal{M}$; if q_B holds in the current state, then B moves, and otherwise A moves. If a player cannot make the required move in some state (meaning the state is a dead end), then the game ends and that player loses and the other player wins.

[2] The closure ordinal of a formula is the smallest ordinal γ (if such an ordinal exists) such that the function defined by the formula converges in at most γ steps in every model.

In a *reachability game*, if the players reach a state where p_B holds, then the game ends and B wins. If the game does not end in a finite number of moves, then A wins. Dually, in a *safety game*, A wins if (s)he reaches a state where p_A holds, while B wins all the infinite plays. Note that reachability/safety games are considered from the "perspective of player B," and a reachability game can be seen as a safety game for player A—and vice versa. In a *double reachability game* we assume that the extensions of p_A and p_B do not intersect.[3] B wins (and the game ends) if (s)he reaches p_B. Similarly, A wins if (s)he reaches p_A. Neither of the players win infinite plays. Those double reachability games that always have a winner (by infinite plays being impossible) are called *strong reachability games*.

The *alternating reachability problem* yields the answer *yes* on the input $(\mathcal{M}, w)$ if and only if B has a winning strategy in the corresponding reachability game. This problem is well known to be PTime-complete under LogSpace reductions (see, e.g., [21] which shows the result already under reductions expressible in first-order logic with order). We let AR denote the class of all positive instances of the alternating reachability problem. The following observation is well known.

Proposition 2.2. *Let $\mathcal{M}$ be a Kripke model with the propositional vocabulary $\{p_B, q_B\}$, and let w be a state in $\mathcal{M}$. Then $(\mathcal{M}, w) \in \mathrm{AR}$ if and only if $\mathcal{M}, w \vDash \chi$, where $\chi = \mu X(p_B \vee (q_B \wedge \Diamond X) \vee (\neg q_B \wedge \Box X))$.*

We can define analogous decision problems for safety and double/strong reachability games. For safety games we can give a result corresponding to Proposition 2.2 by using the sentence $\nu X(\neg p_A \wedge ((q_B \wedge \Diamond X) \vee (\neg q_B \wedge \Box X)))$. For double/strong reachability games we can use the sentence

$$\mu X(\neg p_A \wedge (p_B \vee (q_B \wedge \Diamond X) \vee (\neg q_B \wedge \Box X))).$$

In strong reachability games we can also replace μX with νX in this sentence.

3. Technical links to the literature

In this section we link the current paper to the literature, pointing out technical similarities and differences between existing work and the current paper. We note that the related links are discussed in more depth in the subsequent sections themselves, where our results are developed.

The Sections 4 and 5 define and discuss a game-theoretic semantics (GTS) for the μ-calculus and also its variants where fixed-point iteration is truncated to some predetermined ordinal. GTS of the μ-calculus based on *parity games* is well known, but our system is the first explicit game-theoretic semantics that works also for such truncated variants of the μ-calculus. The semantics is based on the players choosing ordinals, called *clock values*, as *explicit game moves* for the operators μX and νX. The related ordinals are ordered hierarchically in a *lexicographic* manner—an approach that connects very closely to the use of *progress measures* and related notions in [24], [25], [35]. Progress measures have proved highly useful. For example, the paper [22] uses small game progress measures to provide a seminal early improvement on the μ-calculus model checking problem. The paper [23] uses progress measures to develop an algorithm for the problem that uses quasipolynomial time and also "nearly linear" space only.

However, in contrast to earlier works using such clockings, in our game-theoretic approach (1) the clock values are controlled explicitly by the players; (2) clocks are used in a dually symmetric way by both of the players; and (3) when clock values are bounded, our approach gives a natural GTS leading to nonequivalent variants of the μ-calculus. Moreover, most earlier works have considered only finite models, while our definition covers also the infinite case: by requiring the players to lower the ordinals co-inductively, it is still guaranteed that games always have a finite duration.

In Section 6 we study a nonequivalent variant of the μ-calculus semantics, where the ordinals used in the games are restricted to finite ones (also in infinite models). Proposition 6.2 shows that this logic has no *finite model property*, while Proposition 6.3 nevertheless establishes that the set of validities of the logic coincide with that of the ordinary μ-calculus. These results create an interesting contrast in relation to the classical semantics of the μ-calculus.

In Section 7 we discuss variants of *simple semantics*, an approach where we remove the lexicographic dependence of the ordinals used in the standard bounded semantics of Section 4. In Theorem 7.7 we analyse sufficiently large clock value bounds for certain fragments of μ-calculus with the simple GTS. There the case (1) on the $(\nu, \Box)$-free fragment, and the case (3) on the ν-free fragment in finite models, are closely related to folklore results on closure ordinals, but they are here given an alternative game-theoretic proof. The case (2) on the ν-free fragment in infinite models uses the new concept of a *regular branching bound*.

In Section 8 we investigate a new direction of modal logic based on game-theoretic semantics. In particular, we define the new modal logic MCL (*modal computation logic*) which is precisely the modal fragment of the logic CL which is shown Turing-complete (that is, it captures RE in the sense of descriptive complexity theory) in [28]. The logic MCL is builds on the one of the main ideas of the paper, investigating clockings in a general way. While MCL is reasonably expressive, we show that it nevertheless has PTime-complete model checking.

[3] One could also consider the more general case where there can be states where both p_A and p_B are true and thus both players win the game simultaneously. Moreover, one could also define that neither of the players wins the game if they reach a dead end state where neither p_A nor p_B is true.

In Section 9 we analyse variants of GTS, where only one of the players is required to set clock values, leading to reachability/safety games. We note when only the verifier is using clock values—which are not bounded—and when the considered models are finite, the use of such clockings is essentially equivalent to the use progress measures as done, e.g., in [22].

Section 10 gives a uniform reduction of the μ-calculus model checking games (for fixed sentences) into reachability and safety games. The naturalness of the new notion of a strong reachability game is demonstrated. While such reductions exist in the literature, the approach in the section provides a particularly uniform and simple approach, connecting directly to our semantic games. The reduction works as such for finite as well as infinite models. Concerning similar, earlier work, the paper [4] gave a related treatment for showing that finite parity games can be reduced to safety games by adding an explicit memory M to the states. Our reduction also relates to the "Measured Collapse Theorem" in [8], as explained in Section 10 itself.

4. Bounded game-theoretic semantics

The general idea of game-theoretic semantics (GTS) is that truth of a formula φ is checked in a model $\mathcal{M}$ via playing a game where a proponent player (Eloise) attempts to show that φ holds in $\mathcal{M}$ while an opponent player (Abelard) tries to establish the opposite—that φ is false. In this section we define a Γ-*bounded game-theoretic semantics* for the μ-calculus, or Γ-bounded GTS, with a given ordinal $\Gamma > 0$. The semantics shares some features with a similar GTS for the *Alternating-time temporal logic* (ATL) defined in [13].

4.1. Bounded evaluation games

Let φ be a sentence of the μ-calculus and consider an occurrence of a label symbol X in φ. The *reference formula of* X, denoted $\mathrm{Rf}(X)$, is the unique subformula of φ that *binds* X. That is, $\mathrm{Rf}(X)$ is of the form $\mu X\psi$ or $\nu X\psi$ and there is no other operator μX or νX in the syntax tree of φ on the path from $\mathrm{Rf}(X)$ to X. Since φ is a sentence, every label symbol has a reference formula (and the reference formula is by definition unique for each label symbol).

Example 4.1. Consider the sentence $\varphi^* := \nu X\Box\mu Y\big(\Diamond Y \vee (p \wedge X)\big)$. Here we have $\mathrm{Rf}(X) = \varphi^*$ and $\mathrm{Rf}(Y) = \mu Y(\Diamond Y \vee (p \wedge X))$.

Definition 4.2. Let $\mathcal{M}$ be a model, $w_0 \in W$, φ_0 a sentence and $\Gamma > 0$ an ordinal. We define the Γ-*bounded evaluation game* $\mathcal{G} = G(\mathcal{M}, w_0, \varphi_0, \Gamma)$ as follows. The game has two players, *Abelard* and *Eloise*. The *positions* of the game are of the form (w, φ, c), where $w \in W$, $\varphi \in \mathrm{Sf}(\varphi_0)$ and $c : \mathrm{Sf}_{\mu\nu}(\varphi_0) \to \{\gamma \mid \gamma \leq \Gamma\}$ is a *clock mapping*. We call the value $c(\theta)$ the *clock value of* θ (for $\theta \in \mathrm{Sf}_{\mu\nu}(\varphi_0)$).

The game begins from the *initial position* (w_0, φ_0, c_0), where $c_0(\theta) = \Gamma$ for every $\theta \in \mathrm{Sf}_{\mu\nu}(\varphi_0)$. The game is then played according to the following rules:

- In a position (w, p, c) for some $p \in \Phi$, Eloise wins if $w \in V(p)$. Otherwise Abelard wins.
- In a position $(w, \neg p, c)$ for some $p \in \Phi$, Eloise wins if $w \notin V(p)$. Otherwise Abelard wins.
- In a position $(w, \psi \vee \theta, c)$, Eloise selects whether the next position is (w, ψ, c) or (w, θ, c).
- In a position $(w, \psi \wedge \theta, c)$, Abelard selects whether the next position is (w, ψ, c) or (w, θ, c).
- In a position $(w, \Diamond\psi, c)$, Eloise selects some $v \in W$ such that wRv and the next position is (v, ψ, c). If there is no such v, then Abelard wins.
- In a position $(w, \Box\psi, c)$, Abelard selects some $v \in W$ such that wRv and the next position is (v, ψ, c). If there is no such v, then Eloise wins.
- In a position $(w, \mu X\psi, c)$, Eloise chooses an ordinal $\gamma < \Gamma$. Then the game continues from $(w, \psi, c[\gamma/\mu X\psi])$. Here $c[\gamma/\mu X\psi]$ is the clock mapping that sends $\mu X\psi$ to γ and treats other formulae as c.
- In a position $(w, \nu X\psi, c)$, Abelard chooses an ordinal $\gamma < \Gamma$. Then the game continues from the position $(w, \psi, c[\gamma/\nu X\psi])$.
- Suppose that the game is in a position (w, X, c) and let $c(\mathrm{Rf}(X)) = \gamma$.
  1. Suppose that $\mathrm{Rf}(X) = \mu X\psi$ for some ψ.
     - If $\gamma = 0$, then Abelard wins.
     - Else, Eloise must select some $\gamma' < \gamma$, and then the game continues from the position (w, ψ, c'), where
       * $c'(\mu X\psi) = \gamma'$,
       * $c'(\theta) = \Gamma$ for all $\theta \in \mathrm{Sf}_{\mu\nu}(\varphi_0)$ s.t. $\theta \in \mathrm{Sf}(\psi)$,
       * $c'(\theta) = c(\theta)$ for all other $\theta \in \mathrm{Sf}_{\mu\nu}(\varphi_0)$.
  2. Suppose that $\mathrm{Rf}(X) = \nu X\psi$ for some ψ.
     - If $\gamma = 0$, then Eloise wins.
     - Else, Abelard must select some $\gamma' < \gamma$, and then the game continues from the position (w, ψ, c'), where
       * $c'(\nu X\psi) = \gamma'$,
       * $c'(\theta) = \Gamma$ for all $\theta \in \mathrm{Sf}_{\mu\nu}(\varphi_0)$ s.t. $\theta \in \mathrm{Sf}(\psi)$,
       * $c'(\theta) = c(\theta)$ for all other $\theta \in \mathrm{Sf}_{\mu\nu}(\varphi_0)$.

The positions where one of the players wins the game are called *ending positions*. The execution of the rules related to a position of the game constitutes one *round* of the game. The number of rounds in a play of the game is called the *length of the play*. We call the ordinals $\gamma < \Gamma$ *clock values* and the ordinal Γ the *clock value bound*. (We note that only rounds with formulae of type $\mu X\psi$, $\nu X\psi$ and X affect clock values.)

We observe that in evaluation games we do not need assignments s. A label symbol in $X \in \Lambda$ is simply a marker that points to a node (that node being the formula $\mathrm{Rf}(X)$) in the syntax tree of the sentence φ_0. Hence label symbols are conceptually quite different in GTS and compositional semantics. Indeed, the operators μX (respectively νX) can be given a natural reading relating to *self-reference*. In the formula $\mu X\psi$, the operator μX is *naming* the formula ψ with the name X. The atoms X inside ψ are, in turn, *claiming* that ψ holds, i.e., referring back to the formula ψ. The difference between μ and ν is that $\mu X\psi$ relates to *verifying* the formula ψ while $\nu X\psi$ is associated with preventing the falsification of ψ, i.e., *defending* ψ. Therefore, if $N(\psi)$ denotes a natural language reading of ψ, then the natural language reading of $\mu X\psi$ states that "*we can verify the claim named X which asserts that $N(\psi)$*". An analogous reading can be given to $\nu X\psi$. This scheme of reading recursive formulae via self-reference is from [28], [29].

Example 4.3. Consider the Kripke model $\mathcal{M}^* = (W, R, V)$, where we have $W = \{w_i \mid i \in \mathbb{N}\}$, $R = \{(w_0, w_i) \mid i \geq 1\} \cup \{(w_{i+1}, w_i) \mid i \geq 0\}$ and $V(p) = \{w_0\}$.

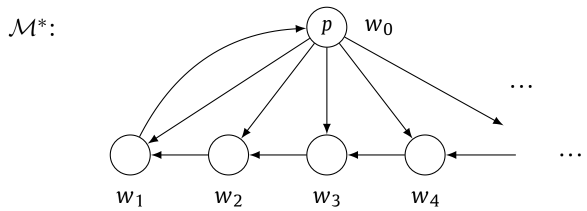

Recall the sentence $\varphi^* = \nu X\Box\mu Y(\Diamond Y \vee (p \wedge X))$ from Example 4.1 and consider the evaluation game $\mathcal{G}^* = G(\mathcal{M}^*, w_0, \varphi^*, \omega)$. In $\mathcal{G}^*$, Abelard first announces a clock value $n < \omega$ for $\mathrm{Rf}(X)$ and then makes a jump from the initial state w_0 (with a $\Box$-move). Next Eloise announces some clock value $m < \omega$ for $\mathrm{Rf}(Y)$. Then she can, by repeated $\vee$-moves, jump in the model (making a $\Diamond$-move) and loop back to the formula $\mathrm{Rf}(Y)$; each time she loops back, she needs to lower the value of m. If Eloise at some point chooses the right disjunct, Abelard can either check if p true in the current state or loop back to $\mathrm{Rf}(X)$. In the latter case, the value of n is lowered, but the value of m is reset back to ω (allowing Eloise to choose a fresh value m next time).

The game eventually ends when (1) the clock value of $\mathrm{Rf}(X)$ goes to zero, whence Abelard loses; when (2) the clock value of $\mathrm{Rf}(Y)$ goes to zero, whence Eloise loses; or when (3) Abelard chooses the left conjunct, whence Eloise wins if and only if p is true at the current state. We will return to this game in Example 4.9.

Proposition 4.4. *Let $\mathcal{G} = G(\mathcal{M}, w, \varphi, \Gamma)$ be a bounded evaluation game. Every play of $\mathcal{G}$ ends in a finite number of rounds.*

Proof. For each positive integer k, let $\prec_k$ denote the *canonical lexicographic order* of k-tuples of ordinals, that is, $(\gamma_1, \ldots, \gamma_k) \prec_k (\gamma'_1, \ldots, \gamma'_k)$ if and only if there exists some $i \leq k$ such that $\gamma_i < \gamma'_i$ and $\gamma_j = \gamma'_j$ for all $j < i$.

Consider a branch in the syntax tree of φ. Let $\psi_1, \ldots, \psi_k \in \mathrm{Sf}_{\mu\nu}(\varphi)$ be the $\mu\nu$-formulae occurring on this branch in this order (starting from the root). In each round of the game, each such sequence $(\psi_1, \ldots, \psi_k)$ is associated with the k-tuple $(c(\psi_1), \ldots, c(\psi_k))$ of clock values (that are ordinals less or equal to Γ). It is easy to see that if c and c' are clock mappings such that c' occurs later than c in the game, then we have $(c'(\psi_1), \ldots, c'(\psi_k)) \preceq_k (c(\psi_1), \ldots, c(\psi_k))$. Also, every time a transition from some label X to the reference formula $\mathrm{Rf}(X)$ is made, there is at least one branch where the k-tuple (for the relevant k) of clock values becomes strictly lowered (in relation to $\prec_k$). As ordinals are well-founded, it is clear that the game ends after finitely many rounds. $\Box$

Each evaluation game $\mathcal{G}$ can naturally be associated with a *game graph* $T(\mathcal{G}) = (P_{\mathcal{G}}, E_{\mathcal{G}})$, where $P_{\mathcal{G}}$ is the set of all possible positions (v, ψ, c) of $\mathcal{G}$ and $E_{\mathcal{G}}$ is the *successor position relation*. That is, $(p_1, p_2) \in E_{\mathcal{G}}$ if and only if the position p_2 can be reached from the position p_1 in a single turn of the game. Note that the initial position has in-degree zero, the ending positions of the game have out-degree zero and all paths from the initial position to some ending position correspond to some play of the game. Due to Proposition 4.4, game graphs of bounded evaluation games contain no infinite paths and therefore they are *acyclic*. However, they are not necessarily trees as several different positions may have a common successor.

Remark 4.5. We defined game-theoretic semantics such that ordinals can always be lowered to *any* smaller ordinal. An alternative possibility would be to require that successor ordinals are always lowered by 1. Such a semantics is called *decremental*. Using decremental semantics makes game graphs smaller, and we make this choice, e.g., in [13] when studying ATL. Notice that decremental semantics does *not* involve any choices of clock values when the clock value bound Γ is finite. This can nicely simplify, e.g., games in finite models. We note that in this article, we consider decremental semantics only in Section 8.

4.2. Game-theoretic semantics

Definition 4.6. Consider an evaluation game $\mathcal{G} = G(\mathcal{M}, w_0, \varphi_0, \Gamma)$. A *(positional) strategy* σ for Eloise in $\mathcal{G}$ is a partial mapping on the set of those positions (w, φ, c) of $\mathcal{G}$ where Eloise needs to make a move such that we have: $\sigma(w, \psi \vee \theta, c) \in \{\psi, \theta\}$, $\sigma(w, \Diamond\psi, c) \in \{v \in W \mid wRv\}$, $\sigma(w, \mu X\psi, c) \in \{\gamma \mid \gamma < \Gamma\}$, and $\sigma(w, X, c) \in \{\gamma \mid \gamma < c(\mathrm{Rf}(X))\}$ where $\mathrm{Rf}(X)$ is of the form $\mu X\psi$. We say that Eloise *plays according to* σ if she makes all her choices according to σ and that σ is a *winning strategy* if Eloise always wins when playing according to σ. Strategies for Abelard are defined analogously.

A bounded evaluation game $\mathcal{G}$ and any strategy σ of Eloise (or Abelard) can be naturally associated with a graph $T(\mathcal{G}, \sigma)$ so that $T(\mathcal{G}, \sigma)$ is *the smallest subgraph of* $T(\mathcal{G})$ *which is closed under* σ. That is, $T(\mathcal{G}, \sigma)$ is obtained inductively by starting from the initial position and proceeding as follows: (i) for positions p of Eloise we add the position $\sigma(p)$ and the edge to it, (ii) for positions p of Abelard we add all successor positions of p and the corresponding edges.

Definition 4.7. Let $\mathcal{M} = (W, R, V)$ be a model, $w \in W$, φ a sentence and $\Gamma > 0$ an ordinal. We define truth of φ in $\mathcal{M}$ and w according to Γ-*bounded game theoretic semantics*, $\mathcal{M}, w \Vdash^{\Gamma} \varphi$, as follows:

$$\mathcal{M}, w \Vdash^{\Gamma} \varphi \text{ iff Eloise has a winning strategy in } G(\mathcal{M}, w, \varphi, \Gamma).$$

Remark 4.8. Note that we define our GTS—and its variants that are studied later on—with respect to positional strategies. However, it is easy see that we would obtain equivalent systems semantics by allowing perfect recall strategies, as in our logics a player with a perfect recall winning strategy also has a positional winning strategy. For games ending in a finite number rounds this can be proven simply via backwards induction. The game variants in Sections 8 and 9 can have infinitely long plays, but it is easy to see that they can be reduced to reachability (or safety) games which are know to be positionally determined.

Example 4.9. Recall the game $\mathcal{G}^*$ from Example 4.3. We define a strategy for Eloise as follows. After Abelard has made a transition to some state w_j, Eloise chooses j for the clock value of $\mathrm{Rf}(Y)$ and jumps in the model until reaching again w_0. She then chooses the right disjunct at w_0, whence she either wins (since $w_0 \in V(p)$) or Abelard needs to lower the clock value of $\mathrm{Rf}(X)$ and the clock value of $\mathrm{Rf}(Y)$ gets reset back to ω. Clearly this is a winning strategy for Eloise and thus $\mathcal{M}^*, w_0 \Vdash^{\omega} \varphi^*$.

From the structure of the evaluation games for φ^* we find an interpretation for the meaning of φ^*: "we can infinitely repeat the process where first (1) an arbitrary transition is made, and then (2) we can reach a state where p is true and the process can be continued from (1)". Thus the clock value chosen for $\mathrm{Rf}(Y)$ is intuitively a "commitment" on *how many rounds at most it takes to reach a state where* p *holds*. The clock value for $\mathrm{Rf}(X)$, on the other hand, is a "challenge" on *how many times* p *must be reached*. Indeed, in models where p can be reached only finitely many—say n—times from the initial state, Abelard can win by choosing $n + 1$ as the initial clock value for $\mathrm{Rf}(X)$.

5. Equivalence of bounded GTS and the standard semantics

5.1. Bounded compositional semantics

In this section we define a compositional semantics based on *ordinal approximants* of fixed point operators. Let $\mathcal{M} = (W, R, V)$ be a Kripke model, $F : \mathcal{P}(W) \to \mathcal{P}(W)$ an operator and γ an ordinal. We define the sets F_μ^γ and F_ν^γ recursively as follows:

$$\begin{array}{lll} F_\mu^0 := \emptyset & \text{and} & F_\nu^0 := W. \\ F_\mu^\gamma := F\big(F_\mu^{\gamma-1}\big) & \text{and} & F_\nu^\gamma := F\big(F_\nu^{\gamma-1}\big), \ \text{ if } \gamma \text{ is a successor ordinal.} \\ F_\mu^\gamma := \bigcup_{\delta<\gamma} F_\mu^\delta & \text{and} & F_\nu^\gamma := \bigcap_{\delta<\gamma} F_\nu^\delta, \ \text{ if } \gamma \text{ is a limit ordinal.} \end{array}$$

Definition 5.1. Consider a model $\mathcal{M}$ with a state w and a related assignment s. We obtain Γ-*bounded compositional semantics* for the modal μ-calculus by defining truth of p, $\neg p$, $\vee$, $\wedge$, $\Diamond$, $\Box$ and X recursively as in the standard compositional semantics and treating the μ and ν-operators as follows:

- $\mathcal{M}, w \vDash_s^\Gamma \mu X\psi$ iff $w \in (\widehat{\psi}_{X,s,\Gamma})_\mu^\Gamma$,
- $\mathcal{M}, w \vDash_s^\Gamma \nu X\psi$ iff $w \in (\widehat{\psi}_{X,s,\Gamma})_\nu^\Gamma$,

where the operator $\widehat{\varphi}_{X,s,\Gamma} : \mathcal{P}(W) \to \mathcal{P}(W)$ is defined such that

$$\widehat{\varphi}_{X,s,\Gamma}(A) = \{w \in W \mid \mathcal{M}, w \vDash_{s[A/X]}^\Gamma \varphi\}.$$

The semantics of the μ and ν-operators can be equivalently given as follows:

- $\mathcal{M}, w \vDash_s^\Gamma \mu X\psi$ iff there exists some $\gamma < \Gamma$ s.t. $w \in (\widehat{\psi}_{X,s,\Gamma})_\mu^{\gamma+1}$.
- $\mathcal{M}, w \vDash_s^\Gamma \nu X\psi$ iff $w \in (\widehat{\psi}_{X,s,\Gamma})_\nu^{\gamma+1}$ for every $\gamma < \Gamma$.

(If Γ is a limit ordinal, we can replace the superscripts $\gamma + 1$ above by γ.)

5.2. *Equivalence of* Γ*-bounded* GTS *and* Γ*-bounded compositional semantics*

We say that a formula is in *normal form* if each label symbol in Λ occurs in the formula at most once in the μ-ν-operators (but may occur several times on the atomic level). We let φ' denote a normal form variant of φ obtained by renaming label symbols where appropriate. It is easy to show that φ is equivalent to φ' with respect to both Γ-bounded compositional semantics ($\vDash^\Gamma$) and Γ-bounded GTS ($\Vdash^\Gamma$). Thus, when proving the equivalence of these two semantics, it suffices to consider sentences that are in normal form.

Theorem 5.2. *Let Γ be an ordinal, $\mathcal{M}$ a Kripke model, $w_0 \in W$ and φ_0 a sentence of the modal μ-calculus. Now we have*

$$\mathcal{M}, w_0 \vDash^\Gamma \varphi_0 \text{ iff } \mathcal{M}, w_0 \Vdash^\Gamma \varphi_0.$$

Proof. We first present a sketch highlighting the main idea of the proof.[4] Any position (w, φ, c) in the evaluation game $\mathcal{G} = G(\mathcal{M}, w_0, \varphi_0, \Gamma)$ can be associated with the *c-generated* assignment s for which the following holds for each $X \in \mathrm{Sf}(\varphi_0)$:

1. $s(X) = (\widehat{\psi}_{X,s,\Gamma})_\mu^\gamma$ if $\mathrm{Rf}(X) = \mu X\psi$ and $c(\mathrm{Rf}(X)) = \gamma$,
2. $s(X) = (\widehat{\psi}_{X,s,\Gamma})_\nu^\gamma$ if $\mathrm{Rf}(X) = \nu X\psi$ and $c(\mathrm{Rf}(X)) = \gamma$.

Note that c-generated assignments essentially relate the clock values γ of bounded GTS to γ-approximants in the bounded compositional semantics. Moreover, note that since we may assume φ_0 to be in normal form, all occurrences of a label symbol X in φ_0 have the same reference formula and thus c-generated assignments s are uniquely defined for all X occurring in φ_0. (The values $s(Y)$ of $Y \in \Lambda \setminus \mathrm{Sf}(\varphi_0)$ may be arbitrary.)

Using the c-generated assignments and the Γ-bounded compositional semantics, each position (w, φ, c) in $\mathcal{G}$ satisfies one of the following conditions:

$E(w, \varphi, c)$: $\mathcal{M}, w, \vDash_s^\Gamma \varphi$ holds with the c-generated assignment s.
$A(w, \varphi, c)$: $\mathcal{M}, w \nvDash_s^\Gamma \varphi$ holds with the c-generated assignment s.

We then prove the following two implications for any position (w, φ, c) in $\mathcal{G}$:

1. If $E(w, \varphi, c)$ holds, Eloise either wins or she can guarantee that $E(w', \varphi', c')$ holds for the next position (w', φ', c').
2. If $A(w, \varphi, c)$ holds, Abelard either wins or he can guarantee that $A(w', \varphi', c')$ holds for the next position (w', φ', c').

From this it follows that if $\mathcal{M}, w_0 \vDash^\Gamma \varphi_0$ holds, then we can formulate a strategy for Eloise which maintains the condition $E(w, \varphi, c)$ and eventually leads to winning the game; and thus $\mathcal{M}, w_0 \Vdash^\Gamma \varphi_0$. Moreover, if $\mathcal{M}, w_0 \nvDash^\Gamma \varphi_0$, then we can dually formulate a winning strategy for Abelard, whence Eloise cannot have a winning strategy in $\mathcal{G}$ and thus $\mathcal{M}, w_0 \nVdash^\Gamma \varphi_0$.

Next we present the proof with all technicalities. We first observe that if $\mathcal{M}, w_0 \vDash^\Gamma \varphi_0$, then the condition $E(w_0, \varphi_0, c_0)$ clearly holds for the initial position of $\mathcal{G}$. Since $\mathcal{G}$ ends in a finite number of rounds, it thus suffices that we prove the claim 1 above to infer one direction of the claim of the theorem: $\mathcal{M}, w_0 \vDash^\Gamma \varphi_0$ implies $\mathcal{M}, w_0 \Vdash^\Gamma \varphi_0$. By also proving the claim 2 above, we can dually infer the contraposition: $\mathcal{M}, w_0 \nvDash^\Gamma \varphi_0$ implies $\mathcal{M}, w_0 \nVdash^\Gamma \varphi_0$.

Consider first a position (w, φ, c) with $\varphi \in \{p, \neg p \mid p \in \Phi\}$. Here the evaluation game ends and clearly Eloise wins if $E(w, \varphi, c)$ holds; and Abelard wins if $A(w, \varphi, c)$ holds.

[4] See [18] for an alternative proof, using slightly different idea, with all technical details.

Consider the position $(w, \psi \vee \theta, c)$ and let s be the c-generated assignment. Suppose first that $E(w, \psi \vee \theta, c)$ holds. Thus $\mathcal{M}, w \vDash_s^\Gamma \psi \vee \theta$, i.e., $\mathcal{M}, w \vDash_s^\Gamma \psi$ or $\mathcal{M}, w \vDash_s^\Gamma \theta$. Hence Eloise can choose $\xi \in \{\psi, \theta\}$ s.t. $\mathcal{M}, w \vDash_s^\Gamma \xi$ and thus $E(w, \xi, c)$ holds for the next position (w, ξ, c). Suppose then that $A(w, \psi \vee \theta, c)$ holds. Hence $\mathcal{M}, w \nvDash_s^\Gamma \psi \vee \theta$, i.e., $\mathcal{M}, w \nvDash_s^\Gamma \psi$ and $\mathcal{M}, w \nvDash_s^\Gamma \theta$. Thus, regardless of the choice $\xi \in \{\psi, \theta\}$ of Eloise, $A(w, \xi, c)$ holds for the next position. The case $(w, \psi \wedge \theta, c)$ is proven dually.[5]

Consider the position $(w, \Diamond\psi, c)$ and let s be the c-generated assignment. Suppose first that $\mathcal{M}, w \vDash_s^\Gamma \Diamond\psi$, i.e., there exists $v' \in W$ s.t. wRv' and $\mathcal{M}, v' \vDash_s^\Gamma \psi$. Thus Eloise can choose v' and then $E(v', \psi, c)$ holds for the next position. Suppose then that $\mathcal{M}, w \nvDash_s^\Gamma \Diamond\psi$, i.e., $\mathcal{M}, v \nvDash_s^\Gamma \psi$ for all $v \in W$ for which wRv. Thus, regardless of the choice $v' \in W$ of Eloise, $A(v', \psi, c)$ holds for the next position. The case $(w, \Box\psi, c)$ is proven dually.

Consider the position $(w, \mu X\psi, c)$ and let s' be the c-generated assignment. Suppose first that $\mathcal{M}, w \vDash_{s'}^\Gamma \mu X\psi$, i.e., there is an ordinal $\gamma' < \Gamma$ s.t. $w \in (\widehat{\psi}_{X,s',\Gamma})_\mu^{\gamma'+1}$. Let $A := (\widehat{\psi}_{X,s',\Gamma})_\mu^{\gamma'}$, whence $w \in \widehat{\psi}_{X,s',\Gamma}(A)$, i.e., $\mathcal{M}, w \vDash_{s'[A/X]}^\Gamma \psi$. Let $s := s'[A/X]$, whence $s(X) = (\widehat{\psi}_{X,s',\Gamma})_\mu^{\gamma'} = (\widehat{\psi}_{X,s,\Gamma})_\mu^{\gamma'}$ and $s(Y) = s'(Y)$ for all $Y \in \mathrm{Sf}(\varphi_0) \setminus \{X\}$. Now Eloise can choose γ' as the clock value of $\mathrm{Rf}(X)$, whence s is the $c[\gamma'/\mu X\psi]$-generated assignment and $E(w, \psi, c[\gamma'/\mu X\psi])$ holds for the next position. Suppose then that $\mathcal{M}, w \nvDash_{s'}^\Gamma \mu X\psi$, i.e., $w \notin (\widehat{\psi}_{X,s',\Gamma})_\mu^{\gamma+1}$ for any $\gamma < \Gamma$. Let $\gamma' < \Gamma$ be a value for $\mathrm{Rf}(X)$, chosen by Eloise. Now by defining A and s as above, s is the $c[\gamma'/\mu X\psi]$-generated assignment and $A(w, \psi, c[\gamma'/\mu X\psi])$ holds for the next position. The case $(w, \nu X\psi, c)$ is proven dually.

Consider the position (w, X, c) and let s' be the c-generated assignment. Moreover, suppose that $c(\mathrm{Rf}(X)) = \gamma$ and $\mathrm{Rf}(X) = \mu X\psi$, whence $s'(X) = (\widehat{\psi}_{X,s',\Gamma})_\mu^{\gamma}$. Suppose first that $\mathcal{M}, w \vDash_{s'}^\Gamma X$, i.e., $w \in s'(X) = (\widehat{\psi}_{X,s',\Gamma})_\mu^{\gamma}$. We first observe that thus the clock value γ cannot be 0. Suppose next that γ is a successor ordinal. Let $A := (\widehat{\psi}_{X,s',\Gamma})_\mu^{\gamma-1}$, whence $w \in \widehat{\psi}_{X,s',\Gamma}(A)$, i.e., $\mathcal{M}, w \vDash_{s'[A/X]}^\Gamma \psi$. Let $s := s'[A/X]$, whence $s(X) = (\widehat{\psi}_{X,s',\Gamma})_\mu^{\gamma-1} = (\widehat{\psi}_{X,s,\Gamma})_\mu^{\gamma-1}$ and $s(Y) = s'(Y)$ for all $Y \in \mathrm{Sf}(\varphi_0) \setminus \{X\}$. Now Eloise can lower the clock value of $\mathrm{Rf}(X)$ from γ to $\gamma - 1$, whence for the next position (w, ψ, c'), $E(w, \psi, c')$ holds and s is the c'-generated assignment. Finally, suppose that γ is a limit ordinal. Now $w \in \bigcup_{\delta<\gamma} (\widehat{\psi}_{X,s',\Gamma})_\mu^{\delta}$, and thus there is $\delta < \gamma$ s.t. $w \in (\widehat{\psi}_{X,s',\Gamma})_\mu^{\delta+1}$. Let $A := (\widehat{\psi}_{X,s',\Gamma})_\mu^{\delta}$, whence $w \in \widehat{\psi}_{X,s',\Gamma}(A)$. Now Eloise can lower the clock value of $\mathrm{Rf}(X)$ from γ to δ and then $E(w, \psi, c')$ holds for the next position with the c'-generated assignment $s := s'[A/X]$.

Suppose then that $\mathcal{M}, w \nvDash_{s'}^\Gamma X$, i.e., $w \notin (\widehat{\psi}_{X,s',\Gamma})_\mu^{\gamma}$. If $\gamma = 0$, then Abelard wins the evaluation game immediately. Suppose then that $\gamma \neq 0$ and let $\gamma' < \gamma$ be a clock value chosen by Eloise. Suppose that γ is a successor ordinal. Since $\gamma' \leq \gamma - 1$ and $\widehat{\psi}_{X,s',\Gamma}$ is monotone, we have $(\widehat{\psi}_{X,s',\Gamma})_\mu^{\gamma'} \subseteq (\widehat{\psi}_{X,s',\Gamma})_\mu^{\gamma-1}$. Let $A := (\widehat{\psi}_{X,s',\Gamma})_\mu^{\gamma'}$, whence $\widehat{\psi}_{X,s',\Gamma}(A) \subseteq \widehat{\psi}_{X,s',\Gamma}((\widehat{\psi}_{X,s',\Gamma})_\mu^{\gamma-1})$. Hence $w \notin \widehat{\psi}_{X,s',\Gamma}(A)$ and thus $\mathcal{M}, w \nvDash_{s'[A/X]}^\Gamma \psi$. Let $s := s'[A/X]$, whence $s(X) = (\widehat{\psi}_{X,s',\Gamma})_\mu^{\gamma'} = (\widehat{\psi}_{X,s,\Gamma})_\mu^{\gamma'}$. Thus s is the c'-generated assignment and $A(w, \psi, c')$ holds for the next position. Finally, suppose that γ is a limit ordinal, whence $\gamma' + 1 < \gamma$. Now $w \notin \bigcup_{\delta<\gamma} (\widehat{\psi}_{X,s',\Gamma})_\mu^{\delta}$, and thus, in particular, $w \notin (\widehat{\psi}_{X,s',\Gamma})_\mu^{\gamma'+1}$. Let $A := (\widehat{\psi}_{X,s',\Gamma})_\mu^{\gamma'}$, whence $w \notin \widehat{\psi}_{X,s',\Gamma}(A)$ and thus $A(w, \psi, c')$ holds for the next position with the c'-generated assignment $s := s'[A/X]$. The case when $\mathrm{Rf}(X) = \nu X\psi$ is proven dually. $\square$

Let $\mathcal{M}$ be a model. It is well-known that over $\mathcal{M}$, each operator related to a formula of the μ-calculus reaches a fixed point in less than $(\mathrm{card}(\mathcal{M}))^+$ iterations, where $(\mathrm{card}(\mathcal{M}))^+$ is the successor *cardinal* of $\mathrm{card}(\mathcal{M})$. Thus it is easy to see that the standard compositional semantics and $(\mathrm{card}(\mathcal{M}))^+$-bounded compositional semantics are equivalent in $\mathcal{M}$. Hence we obtain the following corollary.

Corollary 5.3. *Γ-bounded semantics is equivalent with the standard compositional semantics of the modal μ-calculus when $\Gamma \geq (\mathrm{card}(\mathcal{M}))^+$.*

Also note that, in the special case of *finite models*, it suffices to use finite clock values that are at most the cardinality of the model.

5.3. Stable clock value bounds and f-bounded semantics

Let $\mathcal{M}$ be a model, w be a point in $\mathcal{M}$ and φ be a formula of the modal μ-calculus. We say that a clock value bound Γ is *stable* for the triple $(\mathcal{M}, w, \varphi)$ if the following holds for all ordinals $\Gamma' \geq \Gamma$:

$$\mathcal{M}, w \Vdash^\Gamma \varphi \text{ iff } \mathcal{M}, w \Vdash^{\Gamma'} \varphi.$$

Similarly we say that Γ is stable for a model $\mathcal{M}$ if Γ is stable for $(\mathcal{M}, w, \varphi)$ with any w and φ. From Corollary 5.3 it follows that there exists a stable clock value bound for each model. Moreover, from the proof of Theorem 5.2, we can see

[5] In the case of conjunction the condition $E(w, \varphi, c)$ is maintained regardless of Abelard's choice, while Abelard can maintain $A(w, \varphi, c)$ by choosing the conjunct that is not true. The proofs for the other dual cases that we omit are similar.

that the question on which clock value bounds are stable can be reduced to the study of ordinals where fixed-points of formulae are reached (and vice versa).

Whenever we use a stable clock value bound for a given triple $(\mathcal{M}, w, \varphi)$ we obtain semantics that is equivalent to the standard semantics of the modal μ-calculus, i.e. $\mathcal{M}, w \Vdash^{\Gamma} \varphi$ if and only if $\mathcal{M}, w \vDash \varphi$. However, using smaller values of Γ, we obtain different semantic systems typically nonequivalent to the standard semantics (see the next section for examples). In order to make the Γ-bounded semantics equivalent to the standard semantics, we cannot use any *fixed* clock value bound Γ. Instead we need to set Γ based on the triple $(\mathcal{M}, w, \varphi)$. A simple way is to set $\Gamma := \mathrm{card}(\mathcal{M})^+$ in order obtain stability.[6] We refer to the semantics where Γ is set this way simply as *bounded semantics*. Corollary 5.3 can then be reformulated as follows:

Corollary 5.4. *Bounded semantics is equivalent to the standard semantics of the modal mu-calculus.*

More generally, we can obtain different variants of bounded semantics by using any function f which maps parameters of evaluation games to ordinals. Let us define this formally.

Definition 5.5. Let f be a mapping that takes as input a model $\mathcal{M}$, a point w in the domain W of $\mathcal{M}$ and a sentence φ, outputting an ordinal. Moreover, assume that if g is an isomorphism from $\mathcal{M}$ to $\mathcal{M}'$, then $f(\mathcal{M}, w, \varphi) = f(\mathcal{M}', g(w), \varphi)$.[7] Such a function f is called a *semantic clock function*. The related *f-bounded* GTS is now defined as follows:

$$\mathcal{M}, w \Vdash^{f} \varphi \text{ iff } \mathcal{M}, w \Vdash^{f(\mathcal{M}, w, \varphi)} \varphi.$$

Some natural variants of f-bounded semantics are examined in Section 8. Note also that an equivalent f-bounded compositional semantics can be defined analogously. Γ-bounded GTS is a special case of f-bounded GTS where f is a constant function.

6. Finitely bounded semantics

In this section we consider semantic systems, where a fixed value of Γ is used for all models and formulae. A particularly interesting case here is the so-called *finitely bounded semantics*, where we set $\Gamma = \omega$ for all evaluation games. In the corresponding GTS, the players can only announce *finite* clock values.[8] Finitely bounded semantics will be denoted by FBS which refers to both game-theoretic and compositional semantics with $\Gamma = \omega$. In finite models FBS is equivalent to the standard semantics, but this equivalence breaks over infinite models; see Example 6.1 below.

In the example and proofs that follow, we will consider the sentence

$$\varphi_{\mathrm{AF}p} := \mu X(p \vee \Box X)$$

which intuitively means that *on every path, p can be reached eventually*. Note that $\varphi_{\mathrm{AF}p}$ corresponds to the sentence AFp of *Computation tree logic* CTL.

Example 6.1. Let $\mathcal{M}^{\dagger}$ be the model that is otherwise identical to $\mathcal{M}^{*}$ in Example 4.3, but $V(p) = \{w_1\}$.

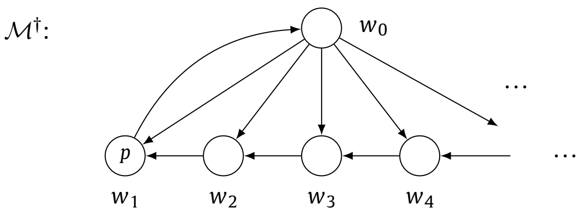

Since the state w_1 is eventually reached on every path starting from w_0, it is easy to see that $\mathcal{M}^{\dagger}, w_0 \vDash \varphi_{\mathrm{AF}p}$. However, $\mathcal{M}^{\dagger}, w_0 \nVdash^{\omega} \varphi_{\mathrm{AF}p}$ since from w_0 there is no finite upper bound on how many transitions are needed to reach w_1. Indeed, Abelard has a winning strategy in $G(\mathcal{M}^{\dagger}, w_0, \varphi_{\mathrm{AF}p}, \omega)$ since he can win by choosing a transition to w_{j+1} for any clock value $j < \omega$ for Rf(X)—chosen by Eloise.

[6] However, depending on $\mathcal{M}$, w and φ, a much smaller Γ can suffice (cf. Section 7).

[7] We note that f is too large to be a set, but this is unproblematic to our study.

[8] Note that the correspondence to for-loops is particularly natural with finitely bounded semantics: iterations can be done up to any finite bound that has to be declared in advance.

It is worth noting that $\mathcal{M}^\dagger, w_0 \Vdash^{\omega+1} \varphi_{\mathrm{AF}p}$ since if Eloise can choose ω as the initial clock value for $\mathrm{Rf}(X)$ and then lower it to $j-1$ after Abelard has made a transition to a state w_j. Moreover, we also have $\mathcal{M}^\dagger, w_0 \Vdash^{\omega} \Box\varphi_{\mathrm{AF}p}$ since Eloise will know how many transitions it takes to reach w_1 as Abelard has to make the first transition before Eloise must announce a clock value.

In the proofs that follow, we will use negations and implications of formulae of the modal μ-calculus. Such formulae are in general not included in our official syntax (in the current paper), but it is straightforward to show that they can be translated to equivalent formulae in negation normal form.

It is well known that, with standard semantics, the modal μ-calculus has the *finite model property*, i.e., every satisfiable sentence is satisfied in some finite model (see, e.g., [6]). However, with finitely bounded semantics, this property is lost.

Proposition 6.2. *The modal μ-calculus with* FBS *does not have the finite model property.*

Proof. It is easy to see that $\Box\varphi_{\mathrm{AF}p} \to \varphi_{\mathrm{AF}p}$ is valid with the standard semantics (this follows from the "fixpoint property" $\mathrm{AF}p \leftrightarrow p \vee \mathsf{AXAF}p$ of CTL). Therefore $\Box\varphi_{\mathrm{AF}p} \wedge \neg\varphi_{\mathrm{AF}p}$ is not satisfiable with the standard semantics. As the standard semantics is equivalent to FBS in finite models, $\Box\varphi_{\mathrm{AF}p} \wedge \neg\varphi_{\mathrm{AF}p}$ cannot be satisfied under FBS in any finite model. However, $\Box\varphi_{\mathrm{AF}p} \wedge \neg\varphi_{\mathrm{AF}p}$ is satisfiable with FBS in an infinite model—as demonstrated by the model $\mathcal{M}^\dagger$ in Example 6.1. $\Box$

Moreover, FBS has the following interesting connection to the standard semantics.

Proposition 6.3. *The set of validities of the modal μ-calculus with* FBS *is strictly included in the set of validities with the standard semantics.*

Proof. To prove the inclusion, let φ be a sentence valid under FBS. Then $\neg\varphi$ cannot be satisfied under FBS in any finite model. Since the standard semantics and FBS are equivalent in finite models, it follows that $\neg\varphi$ is not satisfied by the standard semantics in any finite model. Due to the finite model property of the standard semantics, $\neg\varphi$ is not satisfied by any model and thus φ is valid. The inclusion is strict as $\Box\varphi_{\mathrm{AF}p} \to \varphi_{\mathrm{AF}p}$ is valid under standard semantics but not under FBS (cf. proof of Proposition 6.2). $\Box$

We showed in [14], [17] that the claims of Propositions 6.2, 6.3 hold also for the FBS defined for CTL and ATL. There we also developed a tableau method for showing that the validity problem of CTL and ATL with FBS is decidable and has the same complexity (ExpTime) as with the standard semantics. It remains to be investigated whether the analogous ExpTime result holds also for the modal μ-calculus with FBS.

7. Simple semantics and the ν-free fragment

One key point in our Γ-bounded semantics is the use of clock mappings which utilise several clock values lexicographically. It is natural to ask if we could present the semantics in a more simple way by using just single ordinals for both players. Perhaps the most simple way to formulate such a variant goes as follows.

Definition 7.1. Let $\mathcal{M}$ be a Kripke-model, $w \in W$ and φ a sentence of the μ-calculus. The *simple Γ-bounded evaluation game* $G_s(\mathcal{M}, w, \varphi, \Gamma)$ is played the same way as the standard Γ-bounded evaluation game, but with the following differences on the way the number of remaining rounds is determined:

- Instead of using clock mappings c, Eloise is controlling an ordinal $\gamma_\exists$ and Abelard an ordinal $\gamma_\forall$. In the beginning of the game, these ordinals are set to be equal to Γ.
- Every time a transition is made from some label symbol X to the reference formula $\mu X\psi$, Eloise must lower the current value of $\gamma_\exists$. Similarly, when a transition is made from Y to the reference formula $\nu Y\psi'$, then Abelard must lower $\gamma_\forall$. (Note that the values of $\gamma_\exists$ and $\gamma_\forall$ are never increased.)

If $\gamma_\exists = 0$ and we enter a position where Eloise should lower $\gamma_\exists$, then Eloise loses the game, and similarly, if $\gamma_\forall = 0$ and we enter a position where Abelard should lower $\gamma_\forall$, Abelard loses. In positions $(\mathcal{M}, w', p)$ and $(\mathcal{M}, w', \neg p)$ where p is a proposition symbol, winning and losing is defined in the same way as in Γ-bounded games. We define truth of φ in $\mathcal{M}$ at w according to the *simple Γ-bounded* GTS such that $\mathcal{M}, w \Vdash^{\Gamma}_{sim} \varphi$ if and only if Eloise has a winning strategy in the evaluation game $G_s(\mathcal{M}, w, \varphi, \Gamma)$.

In the next example we show that the simple Γ-bounded semantics is not equivalent to the (standard) Γ-bounded semantics with any clock value bound Γ. The example thus also implies that the simple Γ-bounded semantics cannot be equivalent the standard semantics of the modal μ-calculus with any choice of Γ.

Example 7.2. Recall the formula $\varphi^* = \nu X \Box \mu Y(\Diamond Y \vee (p \wedge X))$ and the evaluation game $\mathcal{G}^* = G(\mathcal{M}^*, w_0, \varphi^*, \omega)$ from Example 4.3. Let $\Gamma > 0$ be any ordinal and consider the simple Γ-bounded evaluation game $\mathcal{G}_s^* = G_s(\mathcal{M}^*, w_0, \varphi^*, \Gamma)$. Even though $\mathcal{M}^*, w_0 \vDash \varphi^*$, Abelard can win $\mathcal{G}_s^*$ by playing as follows. When evaluating $\Box$ at w_0 he jumps to w_1, and then Eloise needs to go back to $\mathrm{Rf}(Y)$ and lower her ordinal to some $\gamma < \Gamma$. We may assume that Eloise then jumps back to w_0 and selects the disjunct $p \wedge X$ there. Then Abelard goes back to $\mathrm{Rf}(X)$ and lowers his ordinal to the same value γ which Eloise currently has. By continuing like this, we will always have $\gamma_\forall \geq \gamma_\exists$ and thus the game will eventually end at Eloise's loss by $\gamma_\exists$ becoming zero.

However, we can show that the simple Γ-bounded semantics can be made equivalent to the standard semantics when we consider the ν-free (or dually the μ-free) fragment of the modal μ-calculus. The ν-free fragment is one of the significant fragments of the μ-calculus, interesting also for its own sake. To give one example of its uses, it has been shown to correspond to *asynchronous distributed automata* in [33]. Now, in order to make simple Γ'-bounded semantics equivalent to the standard Γ-bounded semantics, we typically need to have $\Gamma' > \Gamma$. The following easy result gives a rough estimate on the required size of the clock value bound Γ'.

Proposition 7.3. *For any stable ordinal Γ and ν-free formula φ, we have:*

$$\mathcal{M}, w \Vdash^{\Gamma} \varphi \;\; \textit{iff} \;\; \mathcal{M}, w \Vdash^{\Gamma'}_{sim} \varphi, \quad \textit{where } \Gamma' = (\Gamma+1)^{|\mathrm{Sf}_{\mu\nu}(\varphi)|}$$

Proof. First note that the ordinal $(\Gamma+1)^k$, with $k = |\mathrm{Sf}_{\mu\nu}(\varphi)|$, is essentially order-isomorphic[9] to the lexicographic order of k-tuples $(\gamma_1, \dots, \gamma_k)$ of ordinals $\gamma_i \leq \Gamma$. Since the lexicographic order of such tuples corresponds to the order in which clock mappings appear in the positions of $G(\mathcal{M}, w, \varphi, \Gamma)$, it is easy to see that $\mathcal{M}, w \Vdash^{\Gamma} \varphi$ implies $\mathcal{M}, w \Vdash^{\Gamma'}_{sim} \varphi$ as φ is ν-free. For the other direction, we observe that $\mathcal{M}, w \Vdash^{\Gamma'}_{sim} \varphi$ clearly implies $\mathcal{M}, w \Vdash^{\Gamma'} \varphi$ as φ is ν-free, and since Γ is stable it thus follows that $\mathcal{M}, w \Vdash^{\Gamma} \varphi$. $\Box$

Note that when Γ is infinite, $(\Gamma+1)^{|\mathrm{Sf}_{\mu\nu}(\varphi)|}$ has the same cardinality as Γ. By doing a more fine-grained analysis on the structure of the evaluation games of the ν-free fragment, we can obtain better estimates for the size of the ordinal bounds that Eloise needs. We first make the following observation.

Lemma 7.4. *If $\mathcal{M}, w \Vdash^{\Gamma}_{sim} \varphi$ for some ν-free formula φ and ordinal $\Gamma > 0$, then $\mathcal{M}, w \vDash \varphi$.*

Proof. Let $\Gamma' \geq \Gamma$ be any ordinal that is stable for $\mathcal{M}$. Since φ is ν-free, $\mathcal{M}, w \Vdash^{\Gamma}_{sim} \varphi$ clearly implies $\mathcal{M}, w \Vdash^{\Gamma'} \varphi$. Since Γ' is stable, $\mathcal{M}, w \vDash \varphi$. $\Box$

Sizes of the ordinal bounds sufficient in the ν-free fragment can be estimated inductively accordingly to the following lemma.

Lemma 7.5. *Let φ be a ν-free formula. Consider positions (w, ψ, γ) in a simple Γ-bounded evaluation game $\mathcal{G}_s$ for φ and $\mathcal{M}$, where γ is the current value of the ordinal that Eloise controls (as φ is ν-free, the ordinal of Abelard is irrelevant). For brevity, we will write $\mathcal{M}, w \Vdash^{\gamma}_{sim} \psi$ below to denote that there exists a winning strategy for Eloise from the position (ψ, w, γ) in $\mathcal{G}_s$. The claims below hold for any subformulae ψ of φ and all states w in $\mathcal{M}$.*

- *If ψ is a literal, then $\mathcal{M}, w \vDash \psi$ implies $\mathcal{M}, w \Vdash^{0}_{sim} \psi$.*
- *If $\mathcal{M}, w \Vdash^{\gamma_1}_{sim} \psi_1$ or $\mathcal{M}, w \Vdash^{\gamma_2}_{sim} \psi_2$, then $\mathcal{M}, w \Vdash^{\min\{\gamma_1,\gamma_2\}}_{sim} \psi_1 \vee \psi_2$.*
- *If $\mathcal{M}, w \Vdash^{\gamma_2}_{sim} \psi_1$ and $\mathcal{M}, w \Vdash^{\gamma_1}_{sim} \psi_2$, then $\mathcal{M}, w \Vdash^{\max\{\gamma_1,\gamma_2\}}_{sim} \psi_1 \wedge \psi_2$.*
- *If $\mathcal{M}, w \Vdash^{\gamma_v}_{sim} \psi$ for some v for which wRv, then $\mathcal{M}, v \Vdash^{\gamma}_{sim} \Diamond\psi$, where $\gamma := \min\{\gamma_v \mid wRv$ and $\mathcal{M}, v \Vdash^{\gamma_v}_{s} \psi\}$*
- *If $\mathcal{M}, v \Vdash^{\gamma_v}_{sim} \psi$ for all v such that wRv, then $\mathcal{M}, w \Vdash^{\gamma}_{sim} \Box\psi$, where $\gamma := \sup\{\gamma_v \mid wRv$ and $\mathcal{M}, v \Vdash^{\gamma_v}_{sim} \psi\}$*
- *If $\mathcal{M}, w \Vdash^{\gamma}_{sim} \psi$ and $\mathrm{Rf}(X) = \mu X\psi$, then $\mathcal{M}, w \Vdash^{\gamma+1}_{sim} X$.*
- *If $\mathcal{M}, w \Vdash^{\gamma}_{sim} \psi$, then $\mathcal{M}, w \Vdash^{\gamma}_{sim} \mu X\psi$.*

Proof. The claims follow immediately from the rules of the simple evaluation game by using induction on φ. $\Box$

[9] Recall that finite powers of ordinals are defined based on ordinal multiplication in the natural way such that $\gamma^k := \gamma \cdot \ldots \cdot \gamma$ where γ occurs k times on the right hand side. Therefore $(\Gamma+1)^k$ is best understood via that definition as the k-fold product $(\Gamma+1) \cdot \ldots \cdot (\Gamma+1)$.

Before proving which ordinal bounds suffice for the ν-free fragment, we give the following definition. Given a model $\mathcal{M}$, the *regular branching bound* $\mathsf{RBB}(\mathcal{M})$ *of* $\mathcal{M}$ is the smallest infinite *regular*[10] cardinal that is strictly larger than all the *branchings* (i.e. out-degrees $|\{v \mid wRv\}|$) in $\mathcal{M}$. (Cf. [13], [15].))

Example 7.6. If all the branchings in $\mathcal{M}$ are finite, but there is no upper bound for their size, then $\mathsf{RBB}(\mathcal{M}) = \omega$. If for each $i < \omega$ there is a size $\aleph_i$ branching in $\mathcal{M}$, then $\mathsf{RBB}(\mathcal{M}) = \aleph_{\omega+1}$ (since $\aleph_\omega$ is a singular cardinal).

The following theorem lists cases where the simple GTS becomes equivalent with the standard semantics. The second claim of the theorem is new, but the other claims could be justified via similar known results for closure ordinals. For example, consider the fragment of the standard modal μ-calculus without ν and $\Box$. Let $\mu X\psi$ be a formula in the fragment, s an assignment and $\widehat{\psi}_{X,s}$ a related operator in a model $\mathcal{M}$. Now, it is easy to see and well known that if w is in the least fixed point of $\widehat{\psi}_{X,s}$, then $w \in (\widehat{\psi}_{X,s})^{\gamma}_{\mu}$ for some finite γ. This observation is closely related to in the claim (1) of the theorem below. Nevertheless, as the simple GTS differs from the standard compositional semantics, we shall give self-contained proofs of each claim.

Theorem 7.7. *Let $\mathcal{M}$ be a model and $\Gamma > 0$ an ordinal. The equivalence*

$$\mathcal{M}, w \Vdash^{\Gamma}_{sim} \varphi \text{ iff } \mathcal{M}, w \vDash \varphi$$

holds when at least one of the following conditions is satisfied.

(1) φ is ν-free and $\Box$-free, and $\Gamma \geq \omega$.
(2) φ is ν-free and $\Gamma \geq \mathsf{RBB}(\mathcal{M})$.
(3) φ is ν-free, $\mathcal{M}$ is finite and $\Gamma \geq \mathsf{card}(\mathcal{M}) \cdot |\varphi|$.

In a particular case of (2), it suffices to have $\Gamma = \omega$ when φ is ν-free and $\mathcal{M}$ is image finite.

Proof. As φ is ν-free in all three cases, the implication from left to right follows from Lemma 7.4. Suppose then that $\mathcal{M}, w \vDash \varphi$. Now Eloise has a winning strategy in a Γ'-bounded evaluation $\mathcal{G}$ for some Γ' which is stable for $\mathcal{M}$. Next we attach "ordinal labels" to positions in the game graph of $\mathcal{G}$. We first attach 0 to all ending positions which are winning for Eloise and then proceed inductively by using Lemma 7.5 for the respective state-subformula pairs (v, ψ)—where possible. Note that these ordinal labels are attached independently of clocks c in the positions of $\mathcal{G}$. Since Eloise has a winning strategy in $\mathcal{G}$ and Abelard does not use any clocks in $\mathcal{G}$, it is easy to see that also the initial position of $\mathcal{G}$ will now receive a label γ_0 such that $\mathcal{M}, w \Vdash^{\gamma_0}_{sim} \varphi$.

(1) By starting from zero and taking successors and (finite) maxima and minima, we cannot construct infinite ordinals. Therefore, if φ is ν-free and $\Box$-free, then γ_0 must be finite and thus $\mathcal{M}, w \Vdash^{\Gamma}_{sim} \varphi$ for all $\Gamma \geq \omega$. (2) By additionally taking suprema of such sets of ordinals, whose cardinality correspond to some branching in $\mathcal{M}$, we cannot reach the regular branching bound of $\mathcal{M}$ due to the definition of regular cardinals. Thus, if φ is ν-free, then $\gamma_0 < \mathsf{RBB}(\mathcal{M})$ and thus $\mathcal{M}, w \Vdash^{\Gamma}_{sim} \varphi$ for all $\Gamma \geq \mathsf{RBB}(\mathcal{M})$.

(3) Finally suppose that $\mathcal{M}$ is finite and φ is ν-free. Since ordinals are attached to state-subformula pairs (independently of clocks), there are at most $\mathsf{card}(\mathcal{M}) \cdot |\varphi|$ different ordinal labels attached to the positions of $\mathcal{G}$. Since ordinal labels are constructed (in the finite model) by taking successors, maxima and minima, it follows that if some ordinal γ is attached to a position, then all smaller ordinals must have already been attached too. Hence we must have $\gamma_0 \leq \mathsf{card}(\mathcal{M}) \cdot |\varphi|$ and thus $\mathcal{M}, w \Vdash^{\Gamma}_{sim} \varphi$ for all $\Gamma \geq \mathsf{card}(\mathcal{M}) \cdot |\varphi|$. $\Box$

For all the results presented for the ν-free fragment in this section, the dual results for the μ-free fragment hold as well. The proofs are obtained simply by replacing the role of Eloise with Abelard and all the operators with their duals.

8. Variants with PTime model checking

The setting of bounded GTS leads naturally to semantic variants of the μ-calculus that can quite directly be shown to have PTime complete model checking. A key idea is to make use of the intimate relationship between alternating Turing machines and semantic games. The novel systems of semantics we shall consider in this section resemble the Γ-bounded semantics but utilize a simplified way to control how many times μ and ν-formulae can be repeated in evaluation games. Indeed, recall the f-bounded and simple semantics from Definitions 5.5 and 7.1, respectively. The next definition combines these two ideas.

[10] A regular cardinal κ is equal to its own *cofinality*. This means that, for any family $(S_i)_{i\in I}$ of sets, we have $\mathsf{card}\left(\bigcup_{i\in I} S_i\right) < \kappa$ whenever $\mathsf{card}(I) < \kappa$ and $\mathsf{card}(S_i) < \kappa$ for each $i \in I$.

Definition 8.1. Let f be a semantic clock function (recall Definition 5.5). The *simple f-bounded* GTS is the generalization of the simple Γ-bounded GTS of Definition 7.1 that uses, in the game for checking whether φ holds in $\mathcal{M}$ at w, the ordinal bound $f(\mathcal{M}, w, \varphi)$ (instead of the constant bound Γ).

The naturalness and general properties of simple f-bounded semantics of course depend heavily on the properties of f. One of the first choices that naturally suggests itself is $f\big(\mathcal{M}, w, \varphi\big) = \mathsf{card}(\mathcal{M}) \cdot |\varphi|$, i.e., the model domain size times formula size. We recall that this choice of f gives precisely the semantics shown sufficient for the ν-free fragment in Theorem 7.7. Below we shall further motivate the naturalness of using the semantic clock function $f\big(\mathcal{M}, w, \varphi\big) = \mathsf{card}(\mathcal{M}) \cdot |\varphi|$, but let us first prove the following result.

Proposition 8.2. *The μ-calculus model checking problem under simple f-bounded semantics with $f(\mathcal{M}, w, \varphi) = \mathsf{card}(\mathcal{M}) \cdot |\varphi|$ is PTime-complete.*

Proof. To establish the upper bound, we define a Turing machine running in alternating logarithmic space that directly simulates the model checking game (i.e., the semantic evaluation game) with any input $\mathcal{M}, w, \varphi$. The game positions where Eloise makes a move correspond to existential machine states while Abelard's positions correspond to universal ones. We need some kind of a pointer indicating the current world of the Kripke structure and another pointer for the current subformula (i.e., node in the syntax tree). Furthermore, we keep binary representations of the clock values $\gamma_\exists$ and $\gamma_\forall$ in the memory. These binary strings are necessarily logarithmic in the input due to the choice of f. Thus it is easy to see how the required alternating Turing machine is constructed.

We obtain the lower bound via a reduction from the alternating reachability problem. Recall Proposition 2.2 and the formula χ there. Now, it follows directly from Proposition 2.2 and the third part of Theorem 7.7 that χ defines the winning set of the alternating reachability game also under simple f-bounded semantics, i.e., χ is true in $\mathcal{M}$ at w under that semantics if and only if the player B has a winning strategy in the corresponding reachability game. Note that the reduction from the alternating reachability game to the model checking problem is easy because the input to the game can also act (without modification) as an input model to the model checking problem. Putting all the above together, we now see that already with the fixed input formula χ, model checking is PTime-hard. $\square$

We note that combining the above argument with the third part of Theorem 7.7 gives a proof, based on game-theoretic semantics, of the (known) fact that the model checking problem of the ν-free fragment of the μ-calculus is PTime-complete.

We next discuss a natural generalization of modal logic. We call this logic *(propositional) modal computation logic* (MCL), as it is a fragment of the expressively Turing-complete logic defined in [28] which is studied also in, e.g., [29,30] and dubbed *computation logic* CL in [30]. The Turing-completeness of the logic means that it captures RE in the exact sense of descriptive complexity theory. This is proved in [28]. The logic CL extends the syntax of first-order logic by operators that modify the underlying structure and further operators that enable formulae to refer to themselves. Formally, self-reference is imposed simply by atoms X from where the semantic game can jump back to a formula where X occurs as a subformula.

A particular feature of MCL and CL is that their syntax allows free use of negation. The semantics of negation is precisely as in the GTS of first-order logic, i.e., the swap of roles between the verifier and falsifier (see below). In CL, this approach via GTS is actually a crucial feature, since the classical negation (of compositional semantics) would lead to CL being closed under complement and thereby not equivalent to Turing machines. This is due to the simple fact that there exist recursively enumerable classes whose complement is not recursively enumerable. Nevertheless, CL is a conservative extension of first-order logic, i.e., on first-order formulae, its semantics agrees with first-order logic. Similarly, MCL is a conservative extension of modal logic. The logic MCL can also be called *non-well-founded modal logic*, since its formulae can be regarded as non-well-founded, as we shall see.

We will later on below relate MCL to an f-bounded semantics with $f(\mathcal{M}, w, \varphi) = \mathsf{card}(\mathcal{M}) \cdot |\varphi|$. Similarly to the μ-calculus with simple f-bounded semantics with $f(\mathcal{M}, w, \varphi) = \mathsf{card}(\mathcal{M}) \cdot |\varphi|$, also MCL will prove to have PTime-complete model checking.

Definition 8.3. The syntax of MCL is given by the grammar

$$\varphi ::= p \mid X \mid \neg\varphi \mid \varphi \vee \varphi \mid \Diamond\varphi \mid X\varphi$$

where, as in the modal μ-calculus, we have $p \in \Phi$ and $X \in \Lambda$. To help separate the atomic instances of X from the instances of type $X\varphi$, it is also possible to write $((X))\varphi$ instead of just $X\varphi$. A leaf node X in the syntax tree of a formula χ is said have $((X))\psi$ as the *referent formula* if $((X))\psi$ is the first formula of type $((X))\varphi$ on the path from the leaf X to the root χ of the tree.

The semantic evaluation game for MCL does not involve any time bounds or ordinals, so the game could be described as *clock-free*. The game is played as follows in a model $\mathcal{M}$.

- The positions are of type (w, φ, Q) where the symbol $Q \in \{\exists, \forall\}$ indicates whether Eloise ($\exists$) or Abelard ($\forall$) is currently the *verifier* of the game. Here w is of course just a world of $\mathcal{M}$.

- In a position $(w, \neg\varphi, \exists)$, the game moves to the position $(w, \varphi, \forall)$, and similarly, $(w, \neg\varphi, \forall)$ moves to $(w, \varphi, \exists)$.
- In a position $(w, \Diamond\varphi, Q)$ (respectively, $(w, \varphi \vee \psi, Q)$), the currently verifying player (indicated by Q) chooses the next position which corresponds (in the usual way) to picking a $v \in W$ such that wRv (respectively, picking one of the disjuncts from $\varphi \vee \psi$). In the case for $\Diamond\varphi$, if there is no v such that wRv, the game play ends and the current verifier loses.
- In a position $(w, X\varphi, Q)$, the game moves to the position (w, φ, Q).
- In a position (w, X, Q), the game moves back up to the position $(w, X\varphi, Q)$ corresponding to the referent formula of X. If there is no such referent formula, then the game play ends and neither player wins.[11]
- The game ends in positions of type (w, p, Q), and the current verifier wins if w satisfies p. Else the other player wins.

Note indeed that no clocks are used in the game play. We write $\mathcal{M}, w \Vdash^{\mathsf{MCL}} \varphi$ if Eloise has a positional winning strategy in the game beginning from $(w, \varphi, \exists)$.

Recalling Remark 4.8, note that positional strategies suffice also for the case of MCL even though the respective evaluation games are not determined. This is because the winning condition of the game is a reachability condition.

Note that neither player wins game plays that continue for infinitely many rounds; consider, for example, the formula XX, that is, the formula $((X))X$. On the other hand, no limitations are put on the use of negation in the syntax. It is also worth noting that the semantic games of MCL are double reachability games (cf. Section 2.3). Furthermore, concerning the characterization of MCL as non-well-founded modal logic, note that the syntax trees of formulae of MCL could be replaced by non-well-founded structures where each node pointing to a leaf X would instead point directly to the corresponding formula $((X))\varphi$. In fact, we could even remove the operator $((X))$ (in addition to X) and point directly to the node corresponding to φ.

Let us next define a simple f-bounded semantics for MCL with the semantic clock function $f(\mathcal{M}, w, \varphi) = \mathit{card}(\mathcal{M}) \cdot |\varphi|$. For simplicity, we consider finite models only, so $f(\mathcal{M}, w, \varphi)$ is a finite number. The game for checking whether φ holds in $\mathcal{M}$ at w is defined in the obvious way by extending the game for MCL by two ordinal clocks—one for each of the two players—set to $f(\mathcal{M}, w, \varphi)$ in the beginning. The currently verifying player's ordinal is lowered when jumping from an atomic position with X to the position $X\varphi$ of the referent formula. Neither ordinal is raised at any stage of the game. Ordinals are always lowered by 1, so this is a decremental semantics; recall Remark 4.5. Note that thus the players do not have to make any choices relating to clock values as they always simply get decremented by 1. In a position where a clock value is already 0 but it would have to be lowered even further, the game play ends and *neither* of the players wins. Other ending positions determine the winner of the game play in the same way as in the standard, clock-free semantics of MCL. We write $\mathcal{M}, w \Vdash^{\mathsf{MCL}}_{sim} \varphi$ if Eloise has a positional winning strategy in this game for simple f-bounded semantics[12] with $f(\mathcal{M}, w, \varphi) = \mathit{card}(\mathcal{M}) \cdot |\varphi|$.

Proposition 8.4. *Let $\mathcal{M}$ be finite. Then $\mathcal{M}, w \Vdash^{\mathsf{MCL}} \varphi$ iff $\mathcal{M}, w \Vdash^{\mathsf{MCL}}_{sim} \varphi$.*

Proof. Suppose first that $\mathcal{M}, w \Vdash^{\mathsf{MCL}} \varphi$, i.e. Eloise has a winning strategy σ in the clock-free game. We will use σ to find a winning strategy σ' for Eloise in the simple f-bounded game with $f(\mathcal{M}, w, \varphi) = \mathit{card}(\mathcal{M}) \cdot |\varphi|$.

Now, since σ is positional, there exists no play of the clock-free game respecting σ where a position becomes repeated (for otherwise σ would not be winning). On the other hand, the clock-free game has at most $\mathit{card}(\mathcal{M}) \cdot |\varphi|$ states where Eloise is the current verifier. The same holds also for Abelard. We thus claim that Eloise can use the same strategy σ to also win the simple f-bounded game. To see why this holds, note first that Eloise's (respectively, Abelard's) clock can be lowered only in those positions where Eloise (respectively, Abelard) is the current verifier, and since there are at most $\mathit{card}(\mathcal{M}) \cdot |\varphi|$ such positions, the clock value never needs to drop below 0. As σ is a winning strategy, the claim holds.

The direction from $\mathcal{M}, w \Vdash^{\mathsf{MCL}}_{sim} \varphi$ to $\mathcal{M}, w \Vdash^{\mathsf{MCL}} \varphi$ is immediate. $\square$

Notice that the simple f-bounded semantics of MCL forces games to be of finite duration. As a final remark, we note the following.

Proposition 8.5. *The model checking problem of* MCL *is PTime-complete.*

Proof. The lower bound follows similarly to that of Proposition 8.2. Also the upper bound is similar (based on an alternating LogSpace machine). We simply keep a pointer to the current subformula and another pointer to the model point. We remark that we can even explicitly limit the model checking game duration to $\mathit{card}(\mathcal{M}) \cdot |\varphi|$ (by Proposition 8.4) which can of course be encoded in binary. $\square$

[11] We mention here a natural alternative semantics where the verifier can freely jump from (w, X, Q) to any position of type $(w, X\varphi, Q)$, and the game ends with neither player winning if no such positions exist.

[12] This game essentially corresponds to a reachability game from Eloise's point of view, so having a positional winning strategy is equivalent to having a general winning strategy.

9. Semi-bounded GTS

Here we introduce and discuss *semi-bounded* GTS of the μ-calculus. In this semantics, only Eloise must announce clock values while Abelard can play without setting any clocks. Thereby evaluation games can possibly take infinitely many rounds, in which case Abelard loses infinite plays. The semantics has also a dual variant where only Abelard sets clocks and Eloise plays freely, and Eloise loses infinite plays. For simplicity, we first explicitly discuss only the semi-bounded semantics where Eloise sets clock values. We let Γ-*semi-bounded* GTS ($\Vdash^{\Gamma}_{sb}$) refer to the system where the time limit bound of Eloise's clock values is Γ (and Abelard indeed uses no clocks). We note that the general idea of using clockings that concern only one player is closely related to use of progress measures in, e.g., [22].

Given a Γ-semi-bounded evaluation game $\mathcal{G}$ and a strategy σ, we can define the game graph $T(\mathcal{G})$ and its σ-closed subgraph $T(\mathcal{G},\sigma)$ analogously to the case of Γ-bounded evaluation games. Note that $T(\mathcal{G})$ may now contain infinite paths (and cycles). However, if σ is a winning strategy of the player who *does not set clock values*, then $T(\mathcal{G},\sigma)$ will not contain any infinite paths.

Theorem 9.1. *Γ-semi-bounded* GTS *and Γ-bounded* GTS *are equivalent when Γ is stable.*

Proof. Now, if Eloise has a winning strategy in a Γ-semi-bounded evaluation game, then she trivially also wins the corresponding Γ-bounded evaluation game by essentially using the same strategy—just ignoring the clock values of Abelard. Thus one direction of the equivalence is clear.

To prove the other direction, we first note that the Γ-semi-bounded games are safety games for Eloise, as in order to win, she only has to avoid those ending positions where Abelard wins. As safety games are well-known to be determined (even on infinite game arenas), we conclude that the Γ-semi-bounded games are determined.

Now, assume that Eloise has a winning strategy σ in the Γ-bounded evaluation game $\mathcal{G}$. Suppose, for the sake of contradiction, that she does not have a winning strategy in the Γ-semi-bounded evaluation game $\mathcal{G}'$. Since Γ-semi-bounded evaluation games are determined, Abelard now has a winning strategy σ' in $\mathcal{G}'$.

Consider the σ'-closed game graph $T(\mathcal{G}',\sigma')$. Since σ' is a winning strategy, $T(\mathcal{G}',\sigma')$ does not contain any infinite paths. We map the positions p in $T(\mathcal{G}',\sigma')$ to ordinals γ_p inductively as follows:

- All ending positions in $T(\mathcal{G}',\sigma')$ are mapped to 0.
- All the other positions p in $T(\mathcal{G}',\sigma')$ are mapped to the *supremum* of the ordinals corresponding to the successor positions of p in $T(\mathcal{G}',\sigma')$.

Let Γ^* be the ordinal which is a maximum of Γ and the ordinal γ_{p_0} which corresponds to initial position p_0 of $\mathcal{G}'$. Consider then the Γ^*-bounded evaluation game $\mathcal{G}^*$ which is otherwise identical to $\mathcal{G}$ but it uses Γ^* as the clock value bound. Let σ be the strategy of Abelard which is otherwise identical to σ', but it chooses clock values according to γ_p whenever Abelard needs to select (or lower) some clock value. Since σ' is a winning strategy in $\mathcal{G}'$, it follows from the inductive definition of ordinals γ_p that σ is a winning strategy in $\mathcal{G}^*$. Since $\Gamma^* \geq \Gamma$ and we assumed that Γ is stable, it follows that Abelard has a winning strategy in $\mathcal{G}$. This is a contradiction since Eloise has a winning strategy in $\mathcal{G}$. $\square$

Theorem 9.2. *The variant of Γ-semi-bounded* GTS*, where only Abelard announces clock values, is equivalent to Γ-bounded* GTS *when Γ is stable.*

Proof. If Eloise has a winning strategy in a Γ-bounded evaluation game $\mathcal{G}$, then she can use the same strategy (without setting clocks) in the corresponding Γ-semi-bounded evaluation game $\mathcal{G}'$ where only Abelard needs to announce clock values. For the other direction, suppose that she has a winning strategy in $\mathcal{G}'$. Then, proceeding exactly as in the proof above, we can first attach ordinals γ_p to positions in $\mathcal{G}'$ and then we ultimately obtain a winning strategy for Eloise in $\mathcal{G}$. $\square$

Remark 9.3. In the proofs of the theorems above we add a "clocking" to a strategy which initially does not use any clock values. Since this clocking can be described by attaching single ordinals γ_p (instead of ordinal tuples) to positions p, it may seem as if this method gives a way to replace clock mappings c with ordinals $\gamma_\exists$ and $\gamma_\forall$ as in the simple Γ-bounded semantics. However, note that the ordinals γ_p can *depend* on the clock value bound Γ (which the opponent player uses) and thus typically the ordinals γ_p become larger than Γ—even if Γ is stable. Because of this we cannot replace the clockings of *both* players with just single ordinals (as observed in Example 7.2). Instead the use of ordinals γ_p can be seen as a "meta-clocking" which only works properly when a player knows which clock values the other player is allowed to use.

Related to the remark above, we note that one could also replace clock mappings in the Γ-semi-bounded GTS with a single ordinals to obtain "*simple Γ-semi-bounded* GTS". However, such semantics cannot be equivalent to the standard semantics—regardless how large ordinals are used. This can be seen by looking at the game $\mathcal{G}^*$ in Example 4.3 one more time. Suppose that Eloise has a single ordinal $\gamma_\exists$ which she needs to lower every time when a transition from Y to $\mathrm{Rf}(Y)$ is made. Now Abelard can win the game simply by repeating $\Box$-moves and transitions from X to $\mathrm{Rf}(X)$ until Eloise loses.

Finally, recalling Section 2.3, we note that the GTS where only Eloise (resp. Abelard) announces clock values, is essentially a safety (resp. reachability) game, while our standard bounded GTS is a essentially a strong reachability game. See also the next section and Remark 10.2 for elaboration.

10. Reducing model checking to alternating reachability

In this section we study model checking of the μ-calculus for *fixed sentences*.[13] We investigate model checking separately with respect to the standard semantics and with respect to Γ-bounded semantics. The main observations link to alternating reachability in a way that has some similarities with the approach in Section 8. However, here we consider the standard semantics and reductions relating to *direct model constructions*, concentrating on data complexity only.

Given a sentence φ of the μ-calculus, we use the following notation for the corresponding model checking and bounded model checking problems:

- $\mathrm{MC}(\varphi) := \{(\mathcal{M}, w) \mid \mathcal{M}$ is finite and $\mathcal{M}, w \vDash \varphi\}$, and
- $\mathrm{BMC}(\varphi) := \{(\mathcal{M}, w, \Gamma) \mid \mathcal{M}$ is finite and $\mathcal{M}, w \vDash^{\Gamma} \varphi\}$.

Recalling the relevant notations from Section 2.3, including the formula χ given in Proposition 2.2, we note, in particular, that the alternating reachability problem AR is equal to $\mathrm{MC}(\chi)$. Our aim is to show that AR is a complete problem for model checking and bounded model checking:

Proposition 10.1. *For each formula φ of the modal μ-calculus there are LogSpace-computable model transformations J_φ and I_φ such that for any finite Kripke model $\mathcal{M}$, state w and ordinal Γ we have*

(1) $(\mathcal{M}, w, \Gamma) \in \mathrm{BMC}(\varphi)$ *iff* $J_\varphi(\mathcal{M}, w, \Gamma) \in \mathrm{AR}$, *and*
(2) $(\mathcal{M}, w) \in \mathrm{MC}(\varphi)$ *iff* $I_\varphi(\mathcal{M}, w) \in \mathrm{AR}$.

Furthermore, neither $J_\varphi(\mathcal{M}, w, \Gamma)$ nor $I_\varphi(\mathcal{M}, w)$ contain infinite paths.

Proof. Recall that the game graph of an evaluation game $\mathcal{G} = G(\mathcal{M}, w_0, \varphi, \Gamma)$ is the graph $T(\mathcal{G}) = (P_\mathcal{G}, E_\mathcal{G})$, where $P_\mathcal{G}$ is the set of positions (v, ψ, c) of $\mathcal{G}$, and $E_\mathcal{G}$ is the successor position relation. We consider the following Kripke model that is obtained from $T(\mathcal{G})$ by adding proposition symbols encoding winning end positions of Eloise and positions in which it is Eloise's turn to move: $\mathcal{M}_\mathcal{G} = (P_\mathcal{G}, E_\mathcal{G}, V_\mathcal{G})$, where $V_\mathcal{G} : \{p_B, q_B\} \to \mathcal{P}(P_\mathcal{G})$ is the valuation

- $V_\mathcal{G}(p_B) = \{(v, \psi, c) \in P_\mathcal{G} \mid \psi$ is a literal and $\mathcal{M}, v \vDash \psi\}$,
- $V_\mathcal{G}(q_B) = \{(v, \psi, c) \in P_\mathcal{G} \mid \psi$ is of the form $\theta \vee \eta$, $\Diamond\theta$, $\mu X\theta$, or X with $\mathrm{Rf}(\psi) = \mu X\theta$, or ψ is a literal and $\mathcal{M}, v \nvDash \psi\}$.

Let $r_\mathcal{G} = (w_0, \varphi, c_0)$ be the initial position of $\mathcal{G}$. Observe now that, letting Eloise play in the role of B and Abelard in the role of A, the reachability game on the Kripke-model $\mathcal{M}_\mathcal{G}$ with starting state $r_\mathcal{G}$ is essentially identical with the game $\mathcal{G}$: the positions and the rules for moves are the same, and the winning conditions are equivalent.[14] Thus, defining $J_\varphi(\mathcal{M}, w_0, \Gamma) := (\mathcal{M}_\mathcal{G}, r_\mathcal{G})$, and using Theorem 5.2, we obtain the first equivalence (1). Clearly $J_\varphi(\mathcal{M}, w_0, \Gamma)$ can be computed from the input $(\mathcal{M}, w_0, \Gamma)$ in LogSpace.

The transformation I_φ can now be defined as follows: we let $I_\varphi(\mathcal{M}, w_0) := J_\varphi(\mathcal{M}, w_0, (\mathrm{card}(\mathcal{M}))^+)$. Denote $\Gamma^* := (\mathrm{card}(\mathcal{M}))^+$ below. By Corollary 5.3 and (1) we have

$$(\mathcal{M}, w_0) \in \mathrm{MC}(\varphi) \quad \text{iff} \quad (\mathcal{M}, w_0, \Gamma^*) \in \mathrm{BMC}(\varphi) \quad \text{iff} \quad J_\varphi(\mathcal{M}, w_0, \Gamma^*) \in \mathrm{AR},$$

whence (2) holds. Clearly I_φ is LogSpace-computable.

Since game graphs of Γ-bounded evaluation games do not contain infinite paths, the same holds also for $J_\varphi(\mathcal{M}, w, \Gamma)$ and $I_\varphi(\mathcal{M}, w)$. $\Box$

Remark 10.2. Note that the transformation J_φ in the proof above actually turns a given bounded evaluation game into an *strong* reachability game (recall Section 2.3).[15] Moreover, we can use exactly the same transformation for turning those

[13] The complexities of the related problems are commonly referred to as *data complexity* as opposed to the *combined complexity* of the standard problem where the sentence is not fixed.

[14] For example, in a position $p = (v, \psi, c)$ with ψ a literal such that $\mathcal{M}, v \nvDash \psi$, B loses the alternating reachability game since p does not have any $E_\mathcal{G}$-successors.

[15] Strictly speaking we should also give an interpretation for p_A, but here e.g. $V_\mathcal{G}(p_A) = \emptyset$ works as those position in which Abelard wins the evaluation game correspond to states where B loses by not being able to move in the strong reachability game. However, perhaps the most natural choice would be to define $V_\mathcal{G}(p_A)$ as the set of those positions (v, ψ, c), where ψ is a literal and $\mathcal{M}, v \nvDash \psi$.

semi-bounded evaluation games, where only *Eloise* uses clocks, into safety games. Dually we can formulate an analogous transformation which turns those semi-bounded evaluation games, where only *Abelard* uses clocks, into reachability games (which may contain infinite paths).

Thus, checking the truth of an arbitrary sentence of the modal μ-calculus can be reduced via I_φ to checking the truth of the simple alternation free sentence χ. A related idea was used in [4] for showing that *finite* parity games can be reduced to safety games by adding an explicit memory M to the states. Proposition 10.1 resembles also the "Measured Collapse Theorem" in [8], which states that checking the truth of any sentence φ of the μ-calculus can be reduced to checking the truth of an alternation free sentence φ'. However, unlike in Proposition 10.1, the result of [8] is not a reduction to MC(ψ) for a fixed sentence ψ, as φ' depends on φ. Moreover, the sentence φ' is actually a translation of φ to a different logic, called $\mu^\sharp$-calculus, whose semantics is based on an additional domain of tuples that can be related to our clock values.

It should be noted that the existence of LogSpace-computable reductions from the model checking problems BMC and MC to AR follows directly from the well-known fact that alternating reachability is a PTime-complete problem. However, the main point here is that our reductions J_φ and I_φ arise in a natural and straightforward way from the bounded evaluation game. Moreover, except for LogSpace-computability, the proof above does not rely on any point on the assumption that the Kripke models are finite. Thus we see that the reductions J_φ and I_φ work on infinite Kripke models as well as on finite ones: for any Kripke model $\mathcal{M}$, state w and ordinal Γ we have

- $\mathcal{M}, w \vDash^\Gamma \varphi$ iff $J_\varphi(\mathcal{M}, w, \Gamma) \vDash \chi$, and
- $\mathcal{M}, w \vDash \varphi$ iff $I_\varphi(\mathcal{M}, w) \vDash \chi$.

11. Conclusion and future directions

Our study has focused on conceptual developments relating to the modal μ-calculus, the main result being the new GTS and its variants. There are many relevant future research directions; we mention here some of them. Firstly, it would be interesting to understand new *clocking patterns* in general, in addition to the finitely bounded, the simple bounded, the semi-bounded and the clock-free semantics discussed above. These investigations could naturally be pushed to involve more general logics beyond modal logic, possibly containing, e.g., operators that modify the underlying models, and thereby directly linking to the research on the general logical framework of [28] and the research program of [28] and [29].

More concretely, pinpointing the complexity of the satisfiability problem of the modal μ-calculus under finitely bounded semantics remains to be done. Finally, using ordinals to reduce arbitrary game arenas to well-founded trees is in general an interesting research direction.[16] Relating to this and the work in Section 10, it would be particularly interesting to better understand reductions of general games to strong reachability games.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgments

Antti Kuusisto and Raine Rönnholm were funded by the Academy of Finland project *Theory of Computational Logics*, grant numbers 324435 and 328987, and Antti Kuusisto also by the Academy of Finland project *Explaining AI via Logic* (XAILOG), grant number 345612. The work of Raine Rönnholm was also supported by Jenny and Antti Wihuri Foundation.

References

[1] Bahareh Afshari, Graham E. Leigh, On closure ordinals for the modal mu-calculus, in: Simona Ronchi Della Rocca (Ed.), Computer Science Logic, 2013 (CSL 2013), CSL 2013, September 2-5, 2013, Torino, Italy, in: LIPIcs, vol. 23, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2013, pp. 30–44.
[2] Rajeev Alur, Thomas A. Henzinger, Finitary fairness, ACM Trans. Program. Lang. Syst. 20 (6) (1998) 1171–1194.
[3] Benjamin Aminof, Sasha Rubin, First-cycle games, Inf. Comput. 254 (2017) 195–216.
[4] Julien Bernet, David Janin, Igor Walukiewicz, Permissive strategies: from parity games to safety games, RAIRO Theor. Inform. Appl. 36 (3) (2002) 261–275.
[5] Armin Biere, Alessandro Cimatti, Edmund M. Clarke, Yunshan Zhu, Symbolic model checking without BDDs, in: Rance Cleaveland (Ed.), TACAS'99, Proceedings, in: LNCS, vol. 1579, Springer, 1999, pp. 193–207.
[6] Julian Bradfield, Colin Stirling, Modal mu-Calculi, Elsevier, 2007, pp. 721–756.
[7] Julian C. Bradfield, Igor Walukiewicz, The mu-calculus and model checking, in: Edmund M. Clarke, Thomas A. Henzinger, Helmut Veith, Roderick Bloem (Eds.), Handbook of Model Checking, Springer, 2018, pp. 871–919.
[8] Doron Bustan, Orna Kupferman, Moshe Y. Vardi, A measured collapse of the modal μ-calculus alternation hierarchy, in: V. Diekert, M. Habib (Eds.), Proc. STACS 2004, in: LNCS, vol. 2996, Springer, 2004, pp. 522–533.

[16] The problem of finding equivalent finite duration games for infinite duration games (on finite arenas) has been studied, e.g., in [3] with an essentially different kind of method.

[9] Cristian S. Calude, Sanjay Jain, Bakhadyr Khoussainov, Wei Li, Frank Stephan, Deciding parity games in quasipolynomial time, in: Hamed Hatami, Pierre McKenzie, Valerie King (Eds.), Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC, 2017, Montreal, QC, Canada, June 19-23, 2017, ACM, 2017, pp. 252–263.

[10] Marek Czarnecki, How fast can the fixpoints in modal mu-calculus be reached?, in: Luigi Santocanale (Ed.), 7th Workshop on Fixed Points in Computer Science, FICS 2010, Brno, Czech Republic, August 21–22, 2010, Laboratoire d'Informatique Fondamentale de Marseille, 2010, pp. 35–39, https://hal.archives-ouvertes.fr/hal-00512377/document#page=36.

[11] E. Allen Emerson, Charanjit S. Jutla, Tree automata, mu-calculus and determinacy (extended abstract), in: 32nd Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 1-4 October 1991, IEEE Computer Society, 1991, pp. 368–377.

[12] Valentin Goranko, Antti Kuusisto, Raine Rönnholm, Game-theoretic semantics for ATL+ with applications to model checking, Inf. Comput. 276 (2021) 104554.

[13] Valentin Goranko, Antti Kuusisto, Raine Rönnholm, Game-theoretic semantics for alternating-time temporal logic, in: Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, AAMAS, 2016, pp. 671–679.

[14] Valentin Goranko, Antti Kuusisto, Raine Rönnholm, CTL with finitely bounded semantics, in: Sven Schewe, Thomas Schneider, Jef Wijsen (Eds.), Symposium on Temporal Representation and Reasoning, TIME 2017, in: LIPIcs, vol. 90, Schloss Dagstuhl, 2017, pp. 14:1–14:19.

[15] Valentin Goranko, Antti Kuusisto, Raine Rönnholm, Game-theoretic semantics for ATL+ with applications to model checking, in: Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS, 2017, pp. 1277–1285.

[16] Valentin Goranko, Antti Kuusisto, Raine Rönnholm, Game-theoretic semantics for alternating-time temporal logic, ACM Trans. Comput. Log. 19 (3) (2018) 17:1–17:38.

[17] Valentin Goranko, Antti Kuusisto, Raine Rönnholm, Alternating-time temporal logic ATL with finitely bounded semantics, Theor. Comput. Sci. 797 (2019) 129–155.

[18] Lauri Hella, Antti Kuusisto, Raine Rönnholm, Bounded game-theoretic semantics for modal mu-calculus, CoRR, arXiv:1706.00753 [abs], https://arxiv.org/pdf/1706.00753v1.pdf, 2017.

[19] Lauri Hella, Antti Kuusisto, Raine Rönnholm, Bounded game-theoretic semantics for modal mu-calculus and some variants, in: Jean-François Raskin, Davide Bresolin (Eds.), 11th Symp. on Games, Automata, Logics, and Formal Verification, GandALF, in: EPTCS, vol. 326, 2020, pp. 82–96.

[20] Lauri Hella, Miikka Vilander, Formula size games for modal logic and μ-calculus, J. Log. Comput. 29 (8) (2019) 1311–1344.

[21] Neil Immerman, Descriptive Complexity, Graduate Texts in Computer Science, Springer, 1999.

[22] Jurdzinski Marcin, Small progress measures for solving parity games, in: Horst Reichel, Sophie Tison (Eds.), STACS 2000, 17th Annual Symposium on Theoretical Aspects of Computer Science, Lille, France, February 2000, Proceedings, in: Lecture Notes in Computer Science, vol. 1770, Springer, 2000, pp. 290–301.

[23] Jurdzinski Marcin, Ranko Lazic, Succinct progress measures for solving parity games, in: 32nd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2017, IEEE Computer Society, 2017, pp. 1–9.

[24] Nils Klarlund, Progress measures and finite arguments for infinite computations, Ph.D. thesis, Cornell University, 1990.

[25] Nils Klarlund, Dexter Kozen, Rabin measures and their applications to fairness and automata theory, in: Proceedings of the Sixth Annual Symposium on Logic in Computer Science, LICS'91, Amsterdam, the Netherlands, July 15–18, 1991, IEEE Computer Society, 1991, pp. 256–265.

[26] Dexter Kozen, Results on the propositional mu-calculus, Theor. Comput. Sci. 27 (1983) 333–354.

[27] Orna Kupferman, Nir Piterman, Moshe Y. Vardi, From liveness to promptness, in: Werner Damm, Holger Hermanns (Eds.), Computer Aided Verification, CAV, in: LNCS, vol. 4590, Springer, 2007, pp. 406–419.

[28] Antti Kuusisto, Some Turing-complete extensions of first-order logic, in: Adriano Peron, Carla Piazza (Eds.), GandALF 2014, in: EPTCS, vol. 161, 2014, pp. 4–17.

[29] Antti Kuusisto, On games and computation, CoRR, arXiv:1910.14603 [abs], 2019.

[30] Antti Kuusisto, Game-theoretic semantics and partial specifications, https://arxiv.org/pdf/2010.05456.pdf, 2020.

[31] Damian Niwinski, Igor Walukiewicz, Games for the mu-calculus, Theor. Comput. Sci. 163 (1&2) (1996) 99–116.

[32] Wojciech Penczek, Bozena Wozna, Andrzej Zbrzezny, Bounded model checking for the universal fragment of CTL, Fundam. Inform. 51 (1–2) (2002) 135–156.

[33] Fabian Reiter, Asynchronous distributed automata: a characterization of the modal mu-fragment, in: Ioannis Chatzigiannakis, Piotr Indyk, Fabian Kuhn, Anca Muscholl (Eds.), 44th International Colloquium on Automata, Languages, and Programming, ICALP 2017, July 10-14, 2017, Warsaw, Poland, in: LIPIcs, vol. 80, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017, pp. 100:1–100:14.

[34] Robert S. Streett, E. Allen Emerson, The propositional mu-calculus is elementary, in: Jan Paredaens (Ed.), Automata, Languages and Programming, 11th Colloquium, Antwerp, Belgium, July 16–20, 1984, Proceedings, in: Lecture Notes in Computer Science, vol. 172, Springer, 1984, pp. 465–472.

[35] Robert S. Streett, E. Allen Emerson, An automata theoretic decision procedure for the propositional mu-calculus, Inf. Comput. 81 (3) (1989) 249–264.

[36] Wenhui Zhang, Bounded semantics, Theor. Comput. Sci. 564 (2015) 1–29.