

Bounded Game-Theoretic Semantics for Modal Mu-Calculus and Some Variants

Lauri Hella

Tampere University
Finland

lauri.hella@tuni.fi

Antti Kuusisto

University of Helsinki and Tampere University
Finland

antti.kuusisto@helsinki.fi

Raine Rönholm

ENS Paris-Saclay
France

raine.ronholm@tuni.fi

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. For example, they have already been successfully used as a basis for an approach leading to a natural formula size game for the logic. 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, and we also use the GTS to identify new alternative variants of the mu-calculus with PTime model checking.

1 Introduction

The modal μ -calculus [17] 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., [5] 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 μ -calculus, as we show that our semantics is equivalent to the standard one.

In the novel 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, and since ordinals are well-founded, the game will indeed end in finite time, i.e., after a finite number of game steps. In general, infinite ordinals are needed in the games, but finite ordinals suffice in finite models.

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 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 introduce yet another class of variants of the bounded GTS consisting of the systems of *f-bounded semantics*. In the Γ -bounded semantics, each μ and ν -formula is associated with an ordinal of its own, while in the *f-bounded semantics* this scheme is relaxed and only two ordinals are used, one for all μ -formulae and another one for all ν -formulae. The particular ordinals fixed in the beginning of the game depend on the particular variant of *f-bounded semantics*. We prove PTime-completeness of the model checking problem of a range of simple yet expressive systems of *f-semantics*. The result concerns both data and combined complexity. 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.

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 μ -calculus and related systems, not so much the more technical directions. While some of the technical results we obtain have straightforward and *obvious implicit similarities* to existing notions, such as finite approximants of fixed points, 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 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, we argue that the new semantics can quite often make formulae easier to read; we will illustrate this in Examples 3.3 and 3.7.

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 [8], [11] consider a bounded semantics for the Alternating-time temporal logic ATL, and [10], [7] extend the related study to the extension ATL^+ of ATL. See also [9], [12]. Part of the original motivation behind the studies in [8], [11], [10], [7] (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 presented in [19]. 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 [15]. The first short draft of the current submission appeared in 2017 as the arXiv manuscript [13]. It contained only the game-theoretic semantics presented below; the current article is the extended, full conference version of that short draft. The extended preprint of the current article is also available as the arXiv manuscript [14], containing full technical details.

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 [1] on finitary fairness and the article [18] relating to promptness in Linear temporal logic LTL. We also mention here the work related to *bounded model checking*, see, e.g., [4], [21] and [23]. The article [6] 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. However, the approach and goals of [6] are different, e.g., the paper limits to finite models only and does not discuss game-theoretic semantics at all.

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

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

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 $\text{Sf}(\varphi)$. 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:

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

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 $V : \Phi \rightarrow \mathcal{P}(W)$ a valuation for proposition symbols in Φ . An *assignment* $s : \Lambda \rightarrow \mathcal{P}(W)$ 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} and w under assignment s* , denoted by $\mathcal{M}, w \models_s \varphi$, is defined as in standard modal logic for $p, \neg p, \vee, \wedge, \diamond, \square$. The truth condition for label symbols is defined with respect to the assignment s :

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

To deal with μ and ν , we define an operator $\widehat{\varphi}_{X,s} : \mathcal{P}(W) \rightarrow \mathcal{P}(W)$ such that $\widehat{\varphi}_{X,s}(A) = \{w \in W \mid \mathcal{M}, w \models_{s[A/X]} \varphi\}$, 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:

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

A label symbol X is said to occur *free* in a formula φ if it occurs in φ but is not a subformula of any subformula of φ of the form $\mu X \psi$ or $\nu X \psi$. A formula φ is called a *sentence* if it does not contain any free label symbols. Truth of a sentence φ is independent of assignments s , so we may simply write $\mathcal{M}, w \models \varphi$ instead of $\mathcal{M}, w \models_s \varphi$.

2.3 Alternating Reachability Games

The *alternating reachability game* problem, which is well known to be PTime-complete (see, e.g., [16]), is defined as follows. The input to the problem is a finite pointed model (\mathcal{M}, w) , i.e., \mathcal{M} is Kripke model and w a state in it. We assume the vocabulary of \mathcal{M} contains the proposition symbols p_B and q_B . The game is 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; if q_B holds in the current state, then B moves, and otherwise A moves. If the players reach a state where p_B holds, then the game ends and B wins. 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. If the game does not end in a finite number of moves, then A wins. The alternating reachability game problem yields the answer *yes* on the input (\mathcal{M}, w) iff B has a winning strategy in the game. We let AR denote the class of all positive instances of the alternating reachability game problem. The following observation is well known.

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

3 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 *bounded game-theoretic semantics* for the μ -calculus, or bounded GTS. The semantics shares some features with a similar GTS for the *Alternating-time temporal logic* (ATL) defined in [8] (see also [10]).

3.1 Bounded Evaluation Games

Let φ be a sentence of the μ -calculus and $X \in \text{Sf}(\varphi)$. The *reference formula of X* , denoted $\text{Rf}(X)$, is the unique subformula of φ that *binds* X . That is, $\text{Rf}(X)$ is of the form $\mu X \psi$ or $\nu X \psi$ and there is no other operator μX or νX in the syntax tree on the path from $\text{Rf}(X)$ 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 3.1. Consider the sentence $\varphi^* := \nu X \Box \mu Y (\Diamond Y \vee (p \wedge X))$. Here we have $\text{Rf}(X) = \varphi^*$ and $\text{Rf}(Y) = \mu Y (\Diamond Y \vee (p \wedge X))$.

Definition 3.2. 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* $\mathcal{G} = (\mathcal{M}, w_0, \varphi_0, \Gamma)$ as follows. The game has two players, *Abelard* and *Eloise*. The *positions* of the game are of the form (w, φ, c) , where $w \in W$, $\varphi \in \text{Sf}(\varphi_0)$ and

$$c : \text{Sf}_{\mu\nu}(\varphi_0) \rightarrow \{\gamma \mid \gamma \leq \Gamma\}$$

is a *clock mapping*. We call the value $c(\theta)$ the *clock value of θ* (for $\theta \in \text{Sf}_{\mu\nu}(\varphi_0)$).

The game begins from the *initial position* (w_0, φ_0, c_0) , where $c_0(\theta) = \Gamma$ for every $\theta \in \text{Sf}_{\mu\nu}(\varphi_0)$. 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 $(w, \square \psi, c)$, 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 the position $(w, \psi, c[\gamma/\mu X \psi])$. 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 $c(\text{Rf}(X)) = \gamma$.
 1. Suppose that $\text{Rf}(X) = \mu X \psi$ 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 $\theta \in \text{Sf}_{\mu\nu}(\varphi_0)$ s.t. $\theta \in \text{Sf}(\psi)$,
 - * $c'(\theta) = c(\theta)$ for all other $\theta \in \text{Sf}_{\mu\nu}(\varphi_0)$.
 2. Suppose that $\text{Rf}(X) = \nu X \psi$ 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 $\theta \in \text{Sf}_{\mu\nu}(\varphi_0)$ s.t. $\theta \in \text{Sf}(\psi)$,
 - * $c'(\theta) = c(\theta)$ for all other $\theta \in \text{Sf}_{\mu\nu}(\varphi_0)$.

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 $\text{Rf}(X)$) 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 [19], [20].

Example 3.3. 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 3.1 and consider the evaluation game $\mathcal{G}^* = (\mathcal{M}^*, w_0, \varphi^*, \omega)$. In \mathcal{G}^* , Abelard first announces a clock value $n < \omega$ for $\text{Rf}(X)$ and then makes a jump from the initial state w_0 (with a \Box -move). Next Eloise announces some clock value $m < \omega$ for $\text{Rf}(Y)$. Then she can, by repeated \vee -moves, jump in the model (making a \Diamond -move) and loop back to the formula $\text{Rf}(Y)$; 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 $\text{Rf}(X)$. 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 $\text{Rf}(X)$ goes to zero, whence Abelard loses; when (2) the clock value of $\text{Rf}(Y)$ 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 3.7.

Proposition 3.4. *Let $\mathcal{G} = (\mathcal{M}, w, \varphi, \Gamma)$ 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, $(\gamma_1, \dots, \gamma_k) \prec_k (\gamma'_1, \dots, \gamma'_k)$ 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 $\psi_1, \dots, \psi_k \in \text{Sf}_{\mu\nu}(\varphi)$ 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 (ψ_1, \dots, ψ_k) is associated with the k -tuple $(c(\psi_1), \dots, c(\psi_k))$ 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 $(c'(\psi_1), \dots, c'(\psi_k)) \preceq_k (c(\psi_1), \dots, c(\psi_k))$. Also, every time a transition from some label X to the reference formula $\text{Rf}(X)$ 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 thus clear that the game must end after finitely many rounds. \square

Each evaluation game \mathcal{G} can naturally be associated with a *game tree* $T(\mathcal{G}) = (P_{\mathcal{G}}, E_{\mathcal{G}})$, where $P_{\mathcal{G}}$ is the set of positions (ν, ψ, c) of \mathcal{G} and $E_{\mathcal{G}}$ is the successor position relation. $T(\mathcal{G})$ is formed by beginning from the initial position and adding transitions to all possible successor positions. This procedure is then repeated from the successor positions until an ending position is reached. In the game tree, the initial position is of course the root and ending positions are leafs. Complete branches correspond to possible plays of the game. Due to Proposition 3.4, the game tree of any bounded evaluation game is well-founded, i.e., it does not contain infinite branches. However, if the clock value bound Γ or the model \mathcal{M} is infinite, then the out-degree of some of the nodes of the game tree can be infinite.

3.2 Game-Theoretic Semantics

Definition 3.5. Let $\mathcal{G} = (\mathcal{M}, w_0, \varphi_0, \Gamma)$ be an evaluation game. A *strategy* σ for Eloise in \mathcal{G} is a partial mapping on the set of those positions (w, φ, c) of the game where Eloise needs to make a move such that: $\sigma(w, \psi \vee \theta, c) \in \{\psi, \theta\}$, $\sigma(w, \diamond\psi, c) \in \{v \in W \mid wRv\}$, $\sigma(w, \mu X\psi, c) \in \{\gamma \mid \gamma < \Gamma\}$, and $\sigma(w, X, c) \in \{\gamma \mid \gamma < c(\text{Rf}(X))\}$ where $\text{Rf}(X)$ 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 .

Definition 3.6. 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:

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

Example 3.7. Recall the game \mathcal{G}^* from Example 3.3. 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 $\text{Rf}(Y)$ 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 $\text{Rf}(X)$ and the clock value of $\text{Rf}(Y)$ 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)”. Hence the clock value chosen for $\text{Rf}(Y)$ is intuitively a “commitment” on *how many rounds at most it will take to reach a state where p holds*. The clock value for $\text{Rf}(X)$, 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 $\text{Rf}(X)$.

4 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, $F : \mathcal{P}(W) \rightarrow \mathcal{P}(W)$ an operator and γ an ordinal. We define the sets F_μ^γ and F_ν^γ recursively as follows:

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

Definition 4.1. 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 , \square and X recursively as in the standard compositional semantics and treating the μ and ν -operators as follows:

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

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

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

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

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

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 simply 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 (\models^Γ) and Γ -bounded GTS (\Vdash^Γ). Therefore, when proving the equivalence of these two semantics, it suffices that we consider sentences that are in normal form. Indeed, henceforth we assume that all formulae are in this normal form.

Theorem 4.2. *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*

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

Proof. (Sketch.) We present here a proof sketch highlighting the main ideas. For a rigorous, fully detailed proof, please see the appendix of the full arXiv preprint version [14] of the current submission.

The key in both directions of the proof is the following condition (\star) which is a property satisfied/unsatisfied by positions (w, φ, c) in the evaluation game $\mathcal{G} = (\mathcal{M}, w_0, \varphi_0, \Gamma)$:

(\star) There is an assignment s such that $\mathcal{M}, w \models_s^\Gamma \varphi$, and for each $X \in \text{Sf}(\varphi_0)$:

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

Note that this condition essentially relates the clock values γ of bounded GTS to γ -approximants in the bounded compositional semantics.

Proving the left to right implication, we first note that (\star) holds in the initial position of \mathcal{G} by the assumption $\mathcal{M}, w_0 \models^\Gamma \varphi_0$. Then we show that whenever (\star) holds for the current position, Eloise either wins the game in the current position or she can maintain (\star) to the next position. By maintaining (\star) , we obtain a winning strategy since \mathcal{G} ends in a finite number of rounds.

For the other direction of the equivalence, we suppose that Eloise has a winning strategy σ in \mathcal{G} . Since the game tree of \mathcal{G} is well-founded, we can use well-founded (backwards) induction on the positions in the tree to prove that: if a position (w, φ, c) can be reached with σ , then (\star) holds for (w, φ, c) . Hence, in particular, (\star) holds in the initial position of \mathcal{G} and thus $\mathcal{M}, w_0 \models^\Gamma \varphi_0$. \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 at most $(\text{card}(\mathcal{M}))^+$ iterations, where $(\text{card}(\mathcal{M}))^+$ is the successor cardinal of $\text{card}(\mathcal{M})$. Thus it is easy to see that the standard compositional semantics and $(\text{card}(\mathcal{M}))^+$ -bounded compositional semantics are equivalent in \mathcal{M} . Hence we obtain the following corollary:

Corollary 4.3. *Γ -bounded GTS is equivalent with the standard compositional semantics of the modal μ -calculus when $\Gamma \geq (\text{card}(\mathcal{M}))^+$.*

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 Finitely Bounded Semantics

As stated in Corollary 4.3, the bounded semantics becomes equivalent with the standard (unbounded) semantics if we set a sufficiently large clock value bound Γ . However, using smaller values of Γ , we obtain different semantic systems typically nonequivalent to the standard semantics. We can either set some fixed bound for Γ or use a value that is determined by some parameters—such as the size of the given model and the given formula. In this section we consider the former case; systems relating to the latter case are examined in Section 6.

A particularly interesting case with a fixed value of Γ 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.¹ 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 5.1 below.

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

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

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

Example 5.1. Recall the model \mathcal{M}^* from Example 3.3. Let \mathcal{M}^\dagger be the model that is otherwise identical to \mathcal{M}^* , 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 $\mathcal{M}^\dagger, w_0 \models \varphi_{\text{AF}p}$. However, $\mathcal{M}^\dagger, w_0 \not\models^\omega \varphi_{\text{AF}p}$ 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 $(\mathcal{M}^\dagger, w_0, \varphi_{\text{AF}p}, \omega)$ since he can win by choosing a transition to w_{j+1} for any clock value $j < \omega$ for $\text{Rf}(X)$ —chosen by Eloise.

It is worth noting that $\mathcal{M}^\dagger, w_0 \models^{\omega+1} \varphi_{\text{AF}p}$ since if Eloise can choose ω as the initial clock value for $\text{Rf}(X)$ and then lower it to $j-1$ after Abelard has made a transition to a state w_j . Moreover, we also have $\mathcal{M}^\dagger, w_0 \models^\omega \Box \varphi_{\text{AF}p}$ 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., [5]). However, with finitely bounded semantics, this property is lost.

Proposition 5.2. *The modal μ -calculus with FBS does not have the finite model property.*

Proof. It is easy to see that $\Box \varphi_{\text{AF}p} \rightarrow \varphi_{\text{AF}p}$ is valid with the standard semantics (this follows from the “fixpoint property” $\text{AF}p \leftrightarrow p \vee \text{AXAF}p$ of CTL). Therefore $\Box \varphi_{\text{AF}p} \wedge \neg \varphi_{\text{AF}p}$ is not satisfiable with the standard semantics. As the standard semantics is equivalent to FBS in finite models, $\Box \varphi_{\text{AF}p} \wedge \neg \varphi_{\text{AF}p}$ cannot be satisfied under FBS in any finite model. However, $\Box \varphi_{\text{AF}p} \wedge \neg \varphi_{\text{AF}p}$ is satisfiable with FBS in an infinite model—as demonstrated by the model \mathcal{M}^\dagger in Example 5.1. \square

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

¹Note 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.

Proposition 5.3. *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 $\Box\varphi_{AFp} \rightarrow \varphi_{AFp}$ is valid under standard semantics but not under FBS (cf. proof of Proposition 5.2). \square

We showed in [9], [12] that the claims of Propositions 5.2, 5.3 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 μ -calculus with FBS.

6 Variants with PTime Model Checking

The bounded GTS leads naturally to semantic variants of the μ -calculus that can quite directly be shown to have PTime complete model checking. The main point is to make use of the intimate relationship between alternating Turing machines and semantic games. The novel systems of semantics we consider resemble the Γ -bounded semantics but utilize a simplified way to control how many times μ and ν -formulae can be repeated in semantic games.

To present the alternative semantic systems in detail, let f be a map 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. We assume that if g is an isomorphism from \mathcal{M} to \mathcal{M}' , then $f(\mathcal{M}, w, \varphi) = f(\mathcal{M}', g(w), \varphi)$.² The function f gives rise to the *simple f -bounded GTS* defined as follows.

Definition 6.1. Let \mathcal{M} be a Kripke-model, $w \in W$ and φ a sentence of the μ -calculus. The *simple f -bounded evaluation game* $\mathcal{G}_f = (\mathcal{M}, w, \varphi)$ is played the same way as the Γ -bounded evaluation game $\mathcal{G}_\Gamma = (\mathcal{M}, w, \varphi, \Gamma)$, but with the following differences on the way the number of remaining rounds is determined:

- Eloise is controlling an ordinal γ_\exists and Abelard an ordinal γ_\forall . In the beginning of the game, these ordinals are set to be equal to $f(\mathcal{M}, w, \varphi)$.
- 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 γ_\exists . Similarly, when a transition is made from Y to the reference formula $\nu Y\psi'$, then Abelard must lower γ_\forall . (Note that the values of γ_\exists and γ_\forall are never increased.)

If $\gamma_\exists = 0$ and we enter a position where Eloise should lower γ_\exists , then Eloise loses the game, and similarly, if $\gamma_\forall = 0$ and we enter a position where Abelard should lower γ_\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 f -bounded semantics* such that $\mathcal{M}, w \Vdash^f \varphi$ iff Eloise has a winning strategy in the game $\mathcal{G}_f = (\mathcal{M}, w, \varphi)$ of the simple f -bounded semantics.

Henceforth we mostly talk about f -bounded semantics instead of simple f -bounded semantics to keep the presentation simpler.

²We note that f is too large to be a set, but this is unproblematic to our study.

The naturalness and the general properties of f -bounded semantics of course depend heavily on the choice of f . One of the simpler choices is to define $f(\mathcal{M}, w, \varphi) = \text{card}(\mathcal{M}) \cdot |\varphi|$ where $|\varphi|$ is the length of φ , i.e., the number of symbol occurrences.³ This semantics has the natural property that in finite models, if the players always lower their ordinal by the minimum amount 1, then, if the game ends due to γ_{\exists} or γ_{\forall} being zero, some state-subformula pair must have been repeated. Furthermore, we can now prove the following result.

Proposition 6.2. *The μ -calculus model checking problem is PTime-complete under simple f -bounded semantics with $f(\mathcal{M}, w, \varphi) = \text{card}(\mathcal{M}) \cdot |\varphi|$.*

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, φ . The game positions where Eloise makes a move correspond to existential machine states while Abelard's positions correspond to universal states. 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 γ_{\exists} and γ_{\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 the alternating reachability game. Recall Proposition 2.2 and the formula χ there. We will show that, as in standard semantics, χ defines the winning set of the alternating reachability game also under our f -bounded semantics, i.e., χ is true in \mathcal{M} at w under our semantics if and only if the player B has a winning strategy in the corresponding alternating reachability game. Indeed, it is easy to show that when B has a winning strategy in an alternating reachability game, she can ensure a win so that no state of the game is visited more than once. Thus our choice of f for the f -bounded semantics guarantees that Eloise has a winning strategy in the corresponding the semantic game. And if Eloise has a winning strategy in a semantic game $\mathcal{G}_f(\mathcal{M}, w, \chi)$, then clearly B wins the corresponding alternating reachability game. Thus, already with the fixed input formula χ , model checking is PTime-hard. \square

It is worth noting here that in fact all the systems with $f(\mathcal{M}, w, \varphi) = \text{card}(\mathcal{M})^k \cdot |\varphi|$ (for different positive integers k) have PTime-complete model checking: the proof of Proposition 6.2 goes through with trivial modifications.

The f -bounded semantics with $f(\mathcal{M}, w, \varphi) = \text{card}(\mathcal{M}) \cdot |\varphi|$ is obviously very different in spirit from the standard semantics, and the f -bounded semantics itself changes as we modify f . Also, several further variants of the semantics immediately suggest themselves, for example the possibility of setting different limits for Eloise and Abelard, including the possibility of no limit at all. Also, letting different occurrences of μ and ν -formulae be associated with different clocks similarly to the standard semantics, but without resetting the clocks, is one of many possible interesting scenarios.

Concerning the case where we do not set clocks at all but allow the players to play indefinitely long, winning occurs only when an atomic position with a literal (e.g., p or $\neg p$) is reached. Thus the games are not determined, i.e., it is possible that neither player has a winning strategy (consider, e.g., the formula μXX). This *free semantics* for modal logic results in a system that is essentially directly a fragment of the general, Turing-complete logic \mathcal{L} of [19]. On the other hand, the different ‘‘clocking scenarios’’ described above (and further variants thereof) can be naturally imposed on \mathcal{L} , and it would indeed make sense to study related phenomena in that framework.

³Each proposition symbol p and label X counts as one symbol despite the possible subindices: for example, p_1 is one symbol, not two symbols.

7 Reducing Model Checking to Alternating Reachability

In this section we study model checking of the μ -calculus for *fixed sentences*.⁴ We investigate model checking separately with respect to the standard semantics and with respect to Γ -bounded semantics. Given a sentence φ of the μ -calculus, we use the following notation for the corresponding model checking and bounded model checking problems:

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

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

Proposition 7.1. *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) $(\mathcal{M}, w, \Gamma) \in \text{BMC}(\varphi)$ iff $J_\varphi(\mathcal{M}, w, \Gamma) \in \text{AR}$, and
- (2) $(\mathcal{M}, w) \in \text{MC}(\varphi)$ iff $I_\varphi(\mathcal{M}, w) \in \text{AR}$.

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

Proof. Recall that the game tree of an evaluation game $\mathcal{G} = (\mathcal{M}, w_0, \varphi, \Gamma)$ is the tree $T(\mathcal{G}) = (P_{\mathcal{G}}, E_{\mathcal{G}})$, 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: $\mathcal{T}_{\mathcal{G}} = (P_{\mathcal{G}}, E_{\mathcal{G}}, V_{\mathcal{G}})$, where $V_{\mathcal{G}} : \{p_B, q_B\} \rightarrow \mathcal{P}(P_{\mathcal{G}})$ is the valuation

- $V_{\mathcal{G}}(p_B) = \{(v, \psi, c) \in P_{\mathcal{G}} \mid \psi \text{ is a literal and } \mathcal{M}, v \models \psi\}$,
- $V_{\mathcal{G}}(q_B) = \{(v, \psi, c) \in P_{\mathcal{G}} \mid \psi \text{ is of the form } \theta \vee \eta, \diamond\theta, \mu X\theta, \text{ or } X \text{ with } \text{Rf}(\psi) = \mu X\theta, \text{ or } \psi \text{ is a literal and } \mathcal{M}, v \not\models \psi\}$.

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 alternating reachability game on the Kripke-model $\mathcal{T}_{\mathcal{G}}$ 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.⁵ Thus, defining $J_\varphi(\mathcal{M}, w_0, \Gamma) := (\mathcal{T}_{\mathcal{G}}, r_{\mathcal{G}})$, and using Theorem 4.2, 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 $I_\varphi(\mathcal{M}, w_0) := J_\varphi(\mathcal{M}, w_0, (\text{card}(\mathcal{M}))^+)$. Denote $\Gamma^* := (\text{card}(\mathcal{M}))^+$ below. By Corollary 4.3 and (1) we have

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

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

Since game trees of bounded evaluation games are well-founded, it is clear that $J_\varphi(\mathcal{M}, w, \Gamma)$ and $I_\varphi(\mathcal{M}, w)$ do not contain infinite paths. \square

⁴The 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.

⁵For example, in a position $p = (v, \psi, c)$ with ψ a literal such that $\mathcal{M}, v \not\models \psi$, B loses the alternating reachability game since p does not have any $E_{\mathcal{G}}$ -successors.

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 [3] for showing that *finite* parity games can be reduced to *safety games* by adding explicit memory M to the states. The elements of M are essentially the same as our clock values in the finite case, except that they are given only for one of the players. This is why the resulting game in [3] is a safety game, and this can lead to infinite plays—unlike our reachability games in $I_\varphi(\mathcal{M}, w)$.

Proposition 7.1 resembles also the “Measured Collapse Theorem” in [6], 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 7.1, the result of [6] is not a reduction to $\text{MC}(\psi)$ 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 μ^\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

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

8 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 f -bounded and the 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 [19] and the research program of [19] and [20].

More concretely, pinpointing the complexity of the satisfiability problem of the modal μ -calculus under finitely bounded semantics remains to be done. Also, it would be interesting to investigate whether the scheme of using tuples of ordinals for defining our bounded GTS can be modified to work with single ordinals in a natural way. Finally, using ordinals to reduce arbitrary game arenas to well-founded trees is in general an interesting research direction.⁶ Relating to this and the work in Section 7, it would be particularly interesting to better understand reductions of general games to (well-founded) alternating reachability games.

⁶The problem of finding equivalent finite duration games for infinite duration games (on finite arenas) has been studied, e.g., in [2] with an essentially different kind of method.

Acknowledgements

Antti Kuusisto was funded by the Academy of Finland project *Theory of Computational Logics*, grant numbers 324435 and 328987. The work of Raine Rönholm was partially supported by Jenny and Antti Wihuri Foundation. We thank the anonymous referees for comments and additional references.

References

- [1] Rajeev Alur & Thomas A. Henzinger (1998): *Finitary Fairness*. *ACM Trans. Program. Lang. Syst.* 20(6), pp. 1171–1194, doi:10.1145/295656.295659.
- [2] Benjamin Aminof & Sasha Rubin (2017): *First-cycle games*. *Information and Computation* 254, pp. 195–216, doi:10.1016/j.ic.2016.10.008.
- [3] Julien Bernet, David Janin & Igor Walukiewicz (2002): *Permissive strategies: from parity games to safety games*. *RAIRO Theor. Informatics Appl.* 36(3), pp. 261–275, doi:10.1051/ita:2002013.
- [4] Armin Biere, Alessandro Cimatti, Edmund M. Clarke & Yunshan Zhu (1999): *Symbolic Model Checking without BDDs*. In Rance Cleaveland, editor: *TACAS '99, Proceedings, LNCS 1579*, Springer, pp. 193–207, doi:10.1007/3-540-49059-0_14.
- [5] Julian Bradfield & Colin Stirling (2007): *Modal μ -calculi*, pp. 721–756. Elsevier, doi:10.1016/s1570-2464(07)80015-2.
- [6] Doron Bustan, Orna Kupferman & Moshe Y. Vardi (2004): *A Measured Collapse of the Modal μ -Calculus Alternation Hierarchy*. In Volker Diekert & Michel Habib, editors: *STACS 2004, 21st Annual Symposium on Theoretical Aspects of Computer Science, Proceedings, LNCS 2996*, Springer, pp. 522–533, doi:10.1007/978-3-540-24749-4_46.
- [7] Valentin Goranko, Antti Kuusisto & Raine Rönholm: *Game-Theoretic Semantics for ATL+ with Applications to Model Checking*. To appear in *Information and Computation*.
- [8] Valentin Goranko, Antti Kuusisto & Raine Rönholm (2016): *Game-Theoretic Semantics for Alternating-Time Temporal Logic*. In: *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, AAMAS*, pp. 671–679.
- [9] Valentin Goranko, Antti Kuusisto & Raine Rönholm (2017): *CTL with Finitely Bounded Semantics*. In Sven Schewe, Thomas Schneider & Jef Wijsen, editors: *24th International Symposium on Temporal Representation and Reasoning, TIME, LIPIcs 90*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 14:1–14:19, doi:10.4230/LIPIcs.TIME.2017.14.
- [10] Valentin Goranko, Antti Kuusisto & Raine Rönholm (2017): *Game-Theoretic Semantics for ATL+ with Applications to Model Checking*. In: *Proceedings of the 16th Conference on Autonomous Agents and Multiagent Systems, AAMAS*, pp. 1277–1285.
- [11] Valentin Goranko, Antti Kuusisto & Raine Rönholm (2018): *Game-Theoretic Semantics for Alternating-Time Temporal Logic*. *ACM Trans. Comput. Log.* 19(3), pp. 17:1–17:38, doi:10.1145/3179998.
- [12] Valentin Goranko, Antti Kuusisto & Raine Rönholm (2019): *Alternating-time temporal logic ATL with finitely bounded semantics*. *Theor. Comput. Sci.* 797, pp. 129–155, doi:10.1016/j.tcs.2019.05.029.
- [13] Lauri Hella, Antti Kuusisto & Raine Rönholm (2017): *Bounded game-theoretic semantics for modal μ -calculus*. CoRR abs/1706.00753. Available at <https://arxiv.org/pdf/1706.00753v1.pdf>.
- [14] Lauri Hella, Antti Kuusisto & Raine Rönholm (2020): *Bounded game-theoretic semantics for modal μ -calculus*. CoRR abs/1706.00753. Available at <https://arxiv.org/pdf/1706.00753v2.pdf>.
- [15] Lauri Hella & Miikka Vilander (2019): *Formula size games for modal logic and μ -calculus*. *J. Log. Comput.* 29(8), pp. 1311–1344, doi:10.1093/logcom/exz025.
- [16] Neil Immerman (1999): *Descriptive complexity*. Graduate texts in computer science, Springer, doi:10.1007/978-1-4612-0539-5.

- [17] Dexter Kozen (1983): *Results on the Propositional mu-Calculus*. *Theor. Comput. Sci.* 27, pp. 333–354, doi:10.1016/0304-3975(82)90125-6.
- [18] Orna Kupferman, Nir Piterman & Moshe Y. Vardi (2007): *From Liveness to Promptness*. In Werner Damm & Holger Hermanns, editors: *Computer Aided Verification, 19th International Conference, CAV, Proceedings, LNCS 4590*, Springer, pp. 406–419, doi:10.1007/978-3-540-73368-3_44.
- [19] Antti Kuusisto (2014): *Some Turing-Complete Extensions of First-Order Logic*. In Adriano Peron & Carla Piazza, editors: *GandALF 2014, EPTCS 161*, pp. 4–17, doi:10.4204/EPTCS.161.4.
- [20] Antti Kuusisto (2019): *On Games and Computation*. CoRR abs/1910.14603.
- [21] Wojciech Penczek, Bożena Wozna & Andrzej Zbrzezny (2002): *Bounded Model Checking for the Universal Fragment of CTL*. *Fundam. Inform.* 51(1-2), pp. 135–156.
- [22] Robert S. Streeet & E. Allen Emerson (1989): *An Automata Theoretic Decision Procedure for the Propositional Mu-Calculus*. *Inf. Comput.* 81(3), pp. 249–264, doi:10.1016/0890-5401(89)90031-X.
- [23] Wenhui Zhang (2015): *Bounded semantics*. *Theoretical Computer Science* 564, pp. 1–29, doi:10.1016/j.tcs.2014.10.026.