MODEL CHECKING GAMES FOR THE QUANTITATIVE μ-CALCULUS

DIANA FISCHER [1], ERICH GRÄDEL [1], AND ŁUKASZ KAISER [1]

[1] Mathematische Grundlagen der Informatik, RWTH Aachen
E-mail address: {fischer,graedel,kaiser}@logic.rwth-aachen.de

ABSTRACT. We investigate quantitative extensions of modal logic and the modal μ-calculus, and study the question whether the tight connection between logic and games can be lifted from the qualitative logics to their quantitative counterparts. It turns out that, if the quantitative μ-calculus is defined in an appropriate way respecting the duality properties between the logical operators, then its model checking problem can indeed be characterised by a quantitative variant of parity games. However, these quantitative games have quite different properties than their classical counterparts, in particular they are, in general, not positionally determined. The correspondence between the logic and the games goes both ways: the value of a formula on a quantitative transition system coincides with the value of the associated quantitative game, and conversely, the values of quantitative parity games are definable in the quantitative μ-calculus.

1. Introduction

There have been a number of recent proposals to extend the common qualitative, i.e. two-valued, logical formalisms for specifying the behaviour of concurrent systems, such as propositional modal logic ML, the temporal logics LTL and CTL, and the modal μ-calculus L_μ, to quantitative formalisms. In quantitative logics, the formulae can take, at a given state of a system, not just the values *true* and *false*, but quantitative values, for instance from the (non-negative) real numbers. There are several scenarios and applications where it is desirable to replace purely qualitative statements by quantitative ones, which can be of very different nature: we may be interested in the probability of an event, the value that we assign to an event may depend on how late it occurs, we can ask for the number of occurrences of an event in a play, and so on. We can consider transition structures, where already the atomic propositions take numeric values, or we can ask about the 'degree of satisfaction' of a property. There are several papers that deal with either of these topics, resulting in different specification formalisms and in different notions of transition structures. In particular, due to the prominence and importance of the modal μ-calculus in verification, there have been several attempts to define a quantitative μ-calculus. In some of these, the term quantitative refers to probability, i.e. the logic is interpreted over probabilistic transition systems [11], or used to describe winning conditions in stochastic games [5, 1, 8]. Other variants introduce quantities by allowing discounting in the respective version of

Key words and phrases: games, logic, model checking, quantitative logics.

a "next"-operator for qualitative transition systems [1], Markov decision processes and Markov chains [2], and for stochastic games [4].

While there certainly is ample motivation to extend qualitative specification formalisms to quantitative ones, there also are problems. As has been observed in many areas of mathematics, engineering and computer science where logical formalisms are applied, quantitative formalisms in general lack the clean and clear mathematical theory of their qualitative counterparts, and many of the desirable mathematical and algorithmic properties tend to get lost. Also, the definitions of quantitative formalisms are often ad hoc and do not always respect the properties that are required for logical methodologies. In this paper we have a closer look at quantitative modal logic and the quantitative μ-calculus in terms of their description by appropriate semantic games. The close connection to games is a fundamental aspect of logics. The evaluation of logical formulae can be described by model checking games, played by two players on an arena which is formed as the product of a structure $\mathcal{K}$ and a formula ψ. One player (Verifier) attempts to prove that ψ is satisfied in $\mathcal{K}$ while the other (Falsifier) tries to refute this.

For the modal μ-calculus L_μ, model checking is described by *parity games*, and this connection is of crucial importance for the model theory, the algorithmic evaluation and the applications of the μ-calculus. Indeed, most competitive model checking algorithms for L_μ are based on algorithms to solve the strategy problem in parity games [10]. Furthermore, parity games enjoy nice properties like positional determinacy and can be intuitively understood: often, the best way to make sense of a μ-calculus formula is to look at the associated game. In the other direction, winning regions of parity games (for any fixed number of priorities) are definable in the modal μ-calculus.

In this paper, we explore the question to what extent the relationship between the μ-calculus and parity games can be extended to a quantitative μ-calculus and appropriate quantitative model checking games. The extension is not straightforward, and requires that one defines the quantitative μ-calculus in the 'right' way, so as to ensure that it has appropriate closure and duality properties (such as closure under negation, De Morgan equalities, quantifier and fixed point dualities) to make it amenable to a game-based approach. Once this is done, we can indeed construct a quantitative variant of parity games, and prove that they are the appropriate model checking games for the quantitative μ-calculus. As in the classical setting the correspondence goes both ways: the value of a formula in a structure coincides with the value of the associated model checking game, and conversely, the values of quantitative parity games (with a fixed number of priorities) are definable in the quantitative μ-calculus. However, the mathematical properties of quantitative parity games are different from their qualitative counterparts. In particular, they are, in general, not positionally determined, not even up to approximation. The proof that the quantitative model checking games correctly describe the value of the formulae is considerably more difficult than for the classical case.

As in the classical case, model checking games lead to a better understanding of the semantics and expressive power of the quantitative μ-calculus. Further, the game-based approach also sheds light on the consequences of different choices in the design of the quantitative formalism, which are far less obvious than for classical logics.

2. Quantitative μ-calculus

In [3], de Alfaro, Faella, and Stoelinga introduce a quantitative μ-calculus, that is interpreted over metric transition systems, where predicates can take values in arbitrary metric spaces. Furthermore, their μ-calculus allows discounting in modalities and is studied in connection with quantitative versions of basic system relations such as bisimulation.

We base our calculus on the one proposed in [3] but modify it in the following ways.

(1) We decouple discounts from the modal operators.
(2) We allow discount factors to be greater than one.
(3) In the definition of transition systems we allow additional discounts on the edges.

These changes make the logic more robust and more general, and, as we will show in the next section, will permit us to introduce a negation operator with the desired duality properties that are fundamental to a game-based analysis.

Quantitative transition systems, similar to the ones introduced in [3] are directed graphs equipped with quantities at states and discounts on edges. In the sequel, $\mathbb{R}^+$ is the set of non-negative real numbers, and $\mathbb{R}_\infty^+ := \mathbb{R}^+ \cup \{\infty\}$.

Definition 2.1. A *quantitative transition system (QTS)* is a tuple

$$\mathcal{K} = (V, E, \delta, \{P_i\}_{i \in I}),$$

consisting of a directed graph (V, E), a discount function $\delta : E \to \mathbb{R}^+ \setminus \{0\}$ and functions $P_i : V \to \mathbb{R}_\infty^+$, that assign to each state the values of the predicates at that state.

A transition system is *qualitative* if all functions P_i assign only the values 0 or ∞, i.e. $P_i : V \to \{0, \infty\}$, where 0 stands for false and ∞ for true, and it is *non-discounted* if $\delta(e) = 1$ for all $e \in E$.

We now introduce a quantitative version of the modal μ-calculus to describe properties of quantitative transition systems.

Definition 2.2. Given a set $\mathcal{V}$ of variables X, predicate functions $\{P_i\}_{i \in I}$, discount factors $d \in \mathbb{R}^+$ and constants $c \in \mathbb{R}^+$, the formulae of *quantitative μ-calculus* (Qμ) can be built in the following way:

(1) $|P_i - c|$ is a Qμ-formula,
(2) X is a Qμ-formula,
(3) if φ, ψ are Qμ-formulae, then so are $(\varphi \wedge \psi)$ and $(\varphi \vee \psi)$,
(4) if φ is a Qμ-formula, then so are $\Box \varphi$ and $\Diamond \varphi$,
(5) if φ is a Qμ-formula, then so is $d \cdot \varphi$,
(6) if φ is a formula of Qμ, then $\mu X.\varphi$ and $\nu X.\varphi$ are formulae of Qμ.

Formulae of Qμ are interpreted over quantitative transition systems. Let $\mathcal{F}$ be the set of functions $f : V \to \mathbb{R}_\infty^+$, with $f_1 \le f_2$ if $f_1(v) \le f_2(v)$ for all v. Then $(\mathcal{F}, \le)$ forms a complete lattice with the constant functions $f = \infty$ as top element and $f = 0$ as bottom element.

Given an interpretation $\varepsilon : \mathcal{V} \to \mathcal{F}$, a variable $X \in \mathcal{V}$, and a function $f \in \mathcal{F}$, we denote by $\varepsilon[X \leftarrow f]$ the interpretation ε', such that $\varepsilon'(X) = f$ and $\varepsilon'(Y) = \varepsilon(Y)$ for all $Y \ne X$.

Definition 2.3. Given a QTS $\mathcal{K} = (V, E, \delta, \{P_i\}_{i \in I})$ and an interpretation ε, a Qμ-formula yields a valuation function $\llbracket \varphi \rrbracket_\varepsilon^\mathcal{K} : V \to \mathbb{R}_\infty^+$ defined as follows:

(1) $\llbracket |P_i - c| \rrbracket_\varepsilon^\mathcal{K}(v) = |P_i(v) - c|$,
(2) $\llbracket \varphi_1 \wedge \varphi_2 \rrbracket_\varepsilon^\mathcal{K} = \min\{\llbracket \varphi_1 \rrbracket_\varepsilon^\mathcal{K}, \llbracket \varphi_2 \rrbracket_\varepsilon^\mathcal{K}\}$ and $\llbracket \varphi_1 \vee \varphi_2 \rrbracket_\varepsilon^\mathcal{K} = \max\{\llbracket \varphi_1 \rrbracket_\varepsilon^\mathcal{K}, \llbracket \varphi_2 \rrbracket^\mathcal{K}\}$,

(3) $\llbracket \Diamond\varphi \rrbracket_{\varepsilon}^{\mathcal{K}}(v) = \sup_{v' \in vE} \delta(v, v') \cdot \llbracket \varphi \rrbracket_{\varepsilon}^{\mathcal{K}}(v')$ and $\llbracket \Box\varphi \rrbracket_{\varepsilon}^{\mathcal{K}}(v) = \inf_{v' \in vE} \frac{1}{\delta(v,v')} \llbracket \varphi \rrbracket_{\varepsilon}^{\mathcal{K}}(v')$,

(4) $\llbracket d \cdot \varphi \rrbracket_{\varepsilon}^{\mathcal{K}}(v) = d \cdot \llbracket \varphi \rrbracket_{\varepsilon}^{\mathcal{K}}(v)$,

(5) $\llbracket X \rrbracket_{\varepsilon}^{\mathcal{K}} = \varepsilon(X)$,

(6) $\llbracket \mu X.\varphi \rrbracket_{\varepsilon}^{\mathcal{K}} = \inf\{f \in \mathcal{F} : f = \llbracket \varphi \rrbracket_{\varepsilon[X \leftarrow f]}^{\mathcal{K}}\}$,

(7) $\llbracket \nu X.\varphi \rrbracket_{\varepsilon}^{\mathcal{K}} = \sup\{f \in \mathcal{F} : f = \llbracket \varphi \rrbracket_{\varepsilon[X \leftarrow f]}^{\mathcal{K}}\}$.

For formulae without free variables, we can simply write $\llbracket \varphi \rrbracket^{\mathcal{K}}$ rather than $\llbracket \varphi \rrbracket_{\varepsilon}^{\mathcal{K}}$.

We call the fragment of Qμ consisting of formulae without fixed-point operators *quantitative modal logic* QML. If Qμ is interpreted over qualitative transition systems, it coincides with the classical μ-calculus and we say that $\mathcal{K}, v$ is a model of φ, $\mathcal{K}, v \models \varphi$ if $\llbracket \varphi \rrbracket^{\mathcal{K}}(v) = \infty$. Over non-discounted quantitative transition systems, the definition above coincides with the one in [3]. For discounted systems we take the natural definition for $\Diamond$ and use the dual one for $\Box$, thus the $\frac{1}{\delta}$ factor. As we will show, this is the only definition for which there is a well-behaved negation operator and with a close relation to model checking games.

We always assume the formulae to be *well-named*, i.e. each fixed-point variable is bound only once and no variable appears both free and bound and we use the notions of *alternation level* and *alternation depth* in the usual way, as defined in e.g. [9].

Note that all operators in Qμ are monotone, thus guaranteeing the existence of the least and greatest fixed points, and their inductive definition according to the Knaster-Tarski Theorem stated below.

Proposition 2.4. *The least and greatest fixed points exist and can be computed inductively:* $\llbracket \mu X.\varphi \rrbracket_{\varepsilon}^{\mathcal{K}} = g_{\gamma}$ *with* $g_0(v) = 0$ *(and* $\llbracket \nu X.\varphi \rrbracket_{\varepsilon}^{\mathcal{K}} = g_{\gamma}$ *with* $g_0(v) = \infty$*) for all* $v \in V$ *where*

$$g_{\alpha} = \begin{cases} \llbracket \varphi \rrbracket_{\varepsilon[X \leftarrow g_{\alpha-1}]} & \text{for } \alpha \text{ successor ordinal,} \\ \lim_{\beta < \alpha} \llbracket \varphi \rrbracket_{\varepsilon[X \leftarrow g_{\beta}]} & \text{for } \alpha \text{ limit ordinal,} \end{cases}$$

and γ *is such that* $g_{\gamma} = g_{\gamma+1}$.

3. Negation and Duality

So far, the quantitative logics Qμ and QML lack a negation operator and the associated dualities between $\wedge$ and $\vee$, $\Diamond$ and $\Box$, and between least and greatest fixed points. Let us clarify in the following definition what we expect from such an operator.

Definition 3.1. A *negation operator* $f_{\neg}$ for Qμ is a function $\mathbb{R}_{\infty}^{+} \to \mathbb{R}_{\infty}^{+}$, such that when we define $\llbracket \neg\varphi \rrbracket = f_{\neg}(\llbracket \varphi \rrbracket)$, the following equivalences hold for every $\varphi \in$ Qμ:

(1) $\neg\neg\varphi \equiv \varphi$

(2) $\neg(\varphi \wedge \psi) \equiv \neg\varphi \vee \neg\psi$ and $\neg(\varphi \vee \psi) \equiv \neg\varphi \wedge \neg\psi$

(3) $\neg\Box\varphi \equiv \Diamond\neg\varphi$ and $\neg\Diamond\varphi \equiv \Box\neg\varphi$

(4) $\neg d \cdot \varphi \equiv \beta(d) \cdot \neg\varphi$ for some β independent of φ

(5) $\neg\mu X.\varphi \equiv \nu X.\neg\varphi[X/\neg X]$ and $\neg\nu X.\varphi \equiv \mu X.\neg\varphi[X/\neg X]$

A straightforward calculation shows that the function

$$f_{\frac{1}{x}} : \mathbb{R}_{\infty}^{+} \to \mathbb{R}_{\infty}^{+} : x \mapsto \begin{cases} 1/x & \text{for } x \neq 0, x \neq \infty, \\ \infty & \text{for } x = 0, \\ 0 & \text{for } x = \infty, \end{cases}$$

is a negation operator for Qμ. Hence, we can safely include negation into the definition of Qμ. If we do so, we of course have to demand that the fixed-point variables in the

definition of least and greatest fixed point formulae, see Definition 2.2, only occur under an even number of negations, so as to preserve monotonicity.

Moreover, we show that $f_{\frac{1}{x}}$ is the only negation operator with the required properties. You should note that this is the case even for non-discounted transition systems, and thus it motivates our definition of the semantics of Qμ, in particular of the modal operators, on quantitative transition systems.

Theorem 3.2. $f_{\frac{1}{x}}$ *is the only negation operator for* Qμ, *even for non-discounted systems.*

4. Quantitative Parity Games

Quantitative parity games are an extension of classical parity games. The two main differences are the possibility to assign real values in final positions to denote the payoff for Player 0 and the possibility to discount payoff values on edges.

Definition 4.1. A *quantitative parity game* is a tuple $\mathcal{G} = (V, V_0, V_1, E, \delta, \lambda, \Omega)$ where V is a disjoint union of V_0 and V_1, i.e. positions belong to either Player 0 or 1. The transition relation $E \subseteq V \times V$ describes possible moves in the game and $\delta : V \times V \to \mathbb{R}^+$ maps every move to a positive real value representing the discount factor. The payoff function $\lambda : \{v \in V : vE = \emptyset\} \to \mathbb{R}_\infty^+$ assigns values to all terminal positions and the priority function $\Omega : V \to \{0, \dots, n\}$ assigns a priority to every position.

How to play. Every play starts at some vertex $v \in V$. For every vertex in V_i, Player i chooses a successor vertex, and the play proceeds from that vertex. If the play reaches a terminal vertex, it ends. We denote by $\pi = v_0 v_1 \dots$ the (possibly infinite) play through vertices $v_0 v_1 \dots$, given that $(v_n, v_{n+1}) \in E$ for every n. The outcome $p(\pi)$ of a finite play $\pi = v_0 \dots v_k$ can be computed by multiplying all discount factors seen throughout the play with the value of the final node,

$$p(v_0 v_1 \dots v_k) = \delta(v_0, v_1) \cdot \delta(v_1, v_2) \cdot \dots \cdot \delta(v_{k-1}, v_k) \cdot \lambda(v_k).$$

The outcome of an infinite play depends only on the lowest priority seen infinitely often. We will assign the value 0 to every infinite play, where the lowest priority seen infinitely often is odd, and ∞ to those, where it is even.

Goals. The two players have opposing objectives regarding the outcome of the play. Player 0 wants to maximise the outcome, while Player 1 wants to minimise it.

Strategies. A strategy for player $i \in 0, 1$ is a function $s : V^* V_i \to V$ with $(v, s(v)) \in E$. A play $\pi = v_0 v_1 \dots$ is *consistent with a strategy* s for player i, if $v_{n+1} = s(v_0 \dots v_n)$ for every n such that $v_n \in V_i$. For strategies σ, ρ for the two players, we denote by $\pi_{\sigma, \rho}(v)$ the unique play starting at node v which is consistent with both σ and ρ.

Determinacy. A game is *determined* if, for each position v, the highest outcome Player 0 can assure from this position and the lowest outcome Player 1 can assure converge,

$$\sup_{\sigma \in \Gamma_0} \inf_{\rho \in \Gamma_1} p(\pi_{\sigma, \rho}(v)) = \inf_{\rho \in \Gamma_1} \sup_{\sigma \in \Gamma_0} p(\pi_{\sigma, \rho}(v)) =: \text{val}\mathcal{G}(v),$$

where Γ_0, Γ_1 are the sets of all possible strategies for Player 0, Player 1 and the achieved outcome is called the *value of $\mathcal{G}$ at v*.

Classical parity games can be seen as a special case of quantitative parity games when we map winning to payoff ∞ and losing to payoff 0. Formally, we say that a quantitative parity game $\mathcal{G} = (V, V_0, V_1, E, \delta, \lambda, \Omega)$ is *qualitative* when $\lambda(v) = 0$ or $\lambda(v) = \infty$ for all $v \in V$

with $vE = \emptyset$. In qualitative games, we denote by $W_i \in V$ the winning region of player i, i.e. W_0 is the region where player 0 has a strategy to guarantee payoff ∞ and W_1 is the region where player 1 can guarantee payoff 0. Note that there is no need for the discount function δ in the qualitative case as the payoff can not be changed by discounting.

Qualitative parity games have been extensively studied in the past. One of their fundamental properties is *positional determinacy*. In every parity game, the set of positions can be partitioned into the winning regions W_0 and W_1 for the two players, and each player has a positional winning strategy on her winning region (which means that the moves selected by the strategy only depend on the current position, not on the history of the play).

Unfortunately, this result does not generalise to quantitative parity games. Example 4.2 shows that there are simple quantitative games where no player has a positional winning strategy. In the depicted game there is no optimal strategy for Player 0, and even if one fixes an approximation of the game value, Player 0 needs infinite memory to reach this approximation, because she needs to loop in the second position as long as Player 1 looped in the first one to make up for the discounts. (By convention, we depict positions of Player 0 with a circle and of Player 1 with a square and the number inside is the priority for non-terminal positions and the payoff in terminal ones.)

Example 4.2.
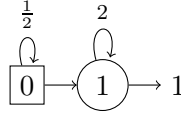
4.1. Model Checking Games for Qμ

A game $(\mathcal{G}, v)$ is a model checking game for a formula φ and a structure $\mathcal{K}, v'$, if the value of the game starting from v is exactly the value of the formula evaluated on $\mathcal{K}$ at v'. In the qualitative case, that means, that φ holds in $\mathcal{K}, v'$ if Player 0 wins in $\mathcal{G}$ from v.

Definition 4.3. For a quantitative transition system $\mathcal{K} = (S, T, \delta_S, P_i)$ and a Qμ-formula φ in negation normal form, the quantitative parity game $MC[\mathcal{K}, \varphi] = (V, V_0, V_1, E, \delta, \lambda, \Omega)$, which we call the *model checking game* for $\mathcal{K}$ and φ, is constructed in the following way.
Positions. The positions of the game are the pairs (ψ, s), where ψ is a subformula of φ, and $s \in S$ is a state of the QTS $\mathcal{K}$, and the two special positions (0) and (∞). Positions (ψ, s) where the top operator of ψ is $\square, \wedge$, or ν belong to Player 1 and all other positions belong to Player 0.
Moves. Positions of the form $(|P_i - c|, s), (0)$, and (∞) are terminal positions. From positions of the form $(\psi \wedge \theta, s)$, resp. $(\psi \vee \theta, s)$, one can move to (ψ, s) or to (θ, s). Positions of the form $(\Diamond\psi, s)$ have either a single successor (0), in case s is a terminal state in $\mathcal{K}$, or one successor (ψ, s') for every $s' \in sT$. Analogously, positions of the form $(\square\psi, s)$ have a single successor (∞), if $sT = \emptyset$, or one successor (ψ, s') for every $s' \in sT$ otherwise. Positions of the form $(d \cdot \psi, s)$ have a unique successor (ψ, s'). Fixed-point positions $(\mu X.\psi, s)$, resp. $(\nu X.\psi, s)$ have a single successor (ψ, s). Whenever one encounters a position where the fixed-point variable stands alone, i.e. (X, s'), the play goes back to the corresponding definition, namely (ψ, s').
Discounts. The discount of an edge is d for transitions from positions $(d \cdot \psi, s)$, it is $\delta_S(s, s')$ for transitions from $(\Diamond\psi, s)$ to (ψ, s'), it is $1/\delta_S(s, s')$ for transitions from $(\square\psi, s)$ to (ψ, s'), and 1 for all outgoing transitions from other positions.

Payoffs. The payoff function λ assigns $|[\![P_i]\!](s) - c|$ to all positions $(|P_i - c|, s)$, ∞ to position (∞), and 0 to position (0).

Priorities. The priority function Ω is defined as in the classical case using the alternation level of the fixed-point variables, see e.g. [9]. Positions (X, s) get a lower priority than positions (X', s') if X has a lower alternation level than X'. The priorities are then adjusted to have the right parity, so that an even value is assigned to all positions (X, s) where X is a ν-variable and an odd value to those where X is a μ-variable. The maximum priority, equal to the alternation depth of the formula, is assigned to all other positions.

It is well-known that qualitative parity games are model checking games for the classical μ-calculus, see e.g. [6] or [12]. A proof that uses the unfolding technique can be found in [9]. We generalise this connection to the quantitative setting as follows.

Theorem 4.4. *For every formula φ in $Q\mu$, a quantitative transition system $\mathcal{K}$, and $v \in \mathcal{K}$, the game $\mathrm{MC}[\mathcal{K}, \varphi]$ is determined and*

$$\mathrm{valMC}[\mathcal{K}, \varphi](\varphi, v) = [\![\varphi]\!]^{\mathcal{K}}(v).$$

4.2. Unfolding Quantitative Parity Games

To prove the model checking theorem in the quantitative case, we start with games with one priority, known as reachability and safety games. The construction of ε-optimal strategies is obtained by a generalisation of backwards induction. At first, we fix the notation and show a few basic properties.

Definition 4.5. A number $k \in \mathbb{R}_{\infty}^{+}$ is called ε-*close* to $p \in \mathbb{R}_{\infty}^{+}$, when either p is finite and $|k - p| \leq \varepsilon$ or $p = \infty$ and $k \geq \frac{1}{\varepsilon}$. A strategy σ in a determined game $\mathcal{G}$ is ε-*optimal* from v if it assures a payoff ε-close to $\mathrm{val}\mathcal{G}(v)$. Furthermore, we say that k is ε-*above* p (or ε-*below*), if $k \geq p'$ (or $k \leq p'$) for some p' that is ε-close to p.

We slightly abuse the word "close" as ε-closeness is *not* symmetric, since $\frac{1}{\varepsilon}$ is ε-close to ∞, but ∞ is not ε-close to any number $r \in \mathbb{R}^{+}$. Still, the following lemmas should convince you that our definition suits our considerations well.

Definition 4.6. For every history $h = v_0 \ldots v_\ell$ of a play, let $\Delta(h) = \Pi_{i < \ell} \delta(v_i, v_{i+1})$ be the product of all discount factors seen in h, and let $D(h) = \max(\Delta(h), \frac{1}{\Delta(h)})$. Note that for every play $\pi = v_0 v_1 \ldots$ and every k,

$$p(\pi) = \Delta(v_0 \ldots v_k) \cdot p(v_k v_{k+1} \ldots).$$

Lemma 4.7. *Let $x, y \in \mathbb{R}_{\infty}^{+}$, $\varepsilon \in (0, 1)$, $\Delta \in \mathbb{R}^{+} \setminus \{0\}$, and $D = \max\{\Delta, \frac{1}{\Delta}\}$.*
   (1) If x is ε/D-close to y, then $\Delta \cdot x$ is ε-close to $\Delta \cdot y$. This holds in particular when $\Delta = \Delta(h)$ and $D = D(h)$ for a history h.
   (2) If x is $\varepsilon/2$-close to y and y is $\varepsilon/2$-close to z, then x is ε-close to z.

This lemma remains valid if we replace the close-relation by the above- or below-relation.

Proposition 4.8. *Reachability and Safety games are determined, for every position v there exist strategies σ^{ε} and ρ^{ε} that guarantee payoffs ε-above (or respectively ε-below) $\mathrm{val}\mathcal{G}(v)$.*

The next step is to prove the determinacy of quantitative parity games. For this purpose, we present a method to unfold a quantitative parity game into a sequence of games with a smaller number of priorities. This technique is inspired by the proof of correctness

of the model checking games for L_μ in [9]. We can extend this method to prove Theorem 4.4 by showing that, as in the classical case, the unfolding of $MC[\mathcal{K}, \varphi]$ is closely related to the inductive evaluation of fixed points in φ on $\mathcal{K}$.

From now on, we assume that the minimal priority in $\mathcal{G}$ is even and call it m. This is no restriction, since, if the minimal priority is odd, we can always consider the dual game, where the roles of the players are switched and all priorities are decreased by one.

Definition 4.9. We define the *truncated game* $\mathcal{G}^- = (V, E^-, \lambda, \Omega^-)$ for a quantitative parity game $\mathcal{G} = (V, E, \lambda, \Omega)$. We assume without loss of generality that all nodes with minimal priority in $\mathcal{G}$ have unique successors with a discount of 1. In $\mathcal{G}^-$ we remove the outgoing edge from each of these nodes. Since these nodes are terminal positions in $\mathcal{G}^-$, their priority does not matter any more for the outcome of a play and Ω^- assigns them a higher priority, e.g. $m + 1$. Formally,

$$E^- = E \setminus \{(v, v') : \Omega(v) = m\}$$

$$\Omega^-(v) = \begin{cases} \Omega(v) & \text{if } \Omega(v) \neq m, \\ m + 1 & \text{if } \Omega(v) = m. \end{cases}$$

The *unfolding of* $\mathcal{G}$ is a sequence of games $\mathcal{G}_\alpha^-$, for ordinals α, which all coincide with $\mathcal{G}^-$, except for the valuation functions λ_α. Below we give the construction of the $\lambda_\alpha's$.

For all terminal nodes v of the original game $\mathcal{G}$ we have $\lambda_\alpha(v) = \lambda(v)$ for all α. For the new terminal nodes, i.e. all $v \in V$, such that $vE^- = \emptyset$ and $vE = \{w\}$, the valuation is given by:

$$\lambda_\alpha(v) = \begin{cases} \infty & \text{for } \alpha = 0, \\ \text{val}\mathcal{G}_{\alpha-1}^-(w) & \text{for } \alpha \text{ successor ordinal}, \\ \lim_{\beta < \alpha} \text{val}\mathcal{G}_\beta^-(w) & \text{for } \alpha \text{ limit ordinal}. \end{cases}$$

The intuition behind the definition of λ_α is to give an incentive for Player 0 to reach the new terminal nodes by first giving them the best possible valuation, and later by updating them to values of their successor in a previous game $\mathcal{G}_\beta^-, \beta < \alpha$.

To determine the value of the original game $\mathcal{G}$, we inductively compute the values for each game in $\mathcal{G}_\alpha$, until they do not change any more. Let γ be an ordinal for which $\text{val}\mathcal{G}_\gamma^- = \text{val}\mathcal{G}_{\gamma+1}^-$. Such an ordinal exists, since the values of the games in the unfolding are monotonically decreasing (which follows from determinacy of these games and definition). We set $g(v) = g_\gamma(v) = \text{val}\mathcal{G}_\gamma^-(v)$ and show that g is the value function of the original game $\mathcal{G}$.

To prove this, we need to introduce strategies for Player 1 and Player 0, which are inductively constructed from the strategies in the unfolding. To give an intuition for the construction, we view a play in $\mathcal{G}$ as a play in the unfolding of $\mathcal{G}$. Let us look more closely at the situation of each player.

The Strategy of Player 0

Player 0 wants to achieve the value $g_\gamma(v_0)$ or to come ε-close. To reach this goal, she imagines to play in $\mathcal{G}_\gamma^-$ and uses her ε-optimal strategies $\sigma_\gamma^\varepsilon$ for that game. Between every two occurrences of nodes of minimal priority throughout the play, she plays a strategy $\sigma_\gamma^{\varepsilon i}$.
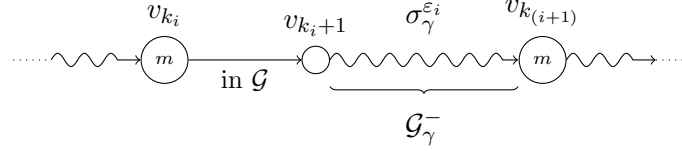
Player 0's strategy after having seen i nodes of priority m.

Initially, ε_i will be $\frac{\varepsilon}{2}$, ε being the approximation value she wants to attain in the end. Then she chooses a lower ε_{i+1} every time she passes an edge outside of $\mathcal{G}^-$. She will adjust the approximation value not only by cutting it in half every time she changes the strategy, but also according to the discount factors seen so far, since they also can dramatically alter the value of the approximation.

For a history h or a full play π, let $L(h)$ (resp. $L(\pi)$) be the number of nodes with minimal priority m occurring in h (or π).

Definition 4.10. The strategy σ^ε for Player 0 in the game $\mathcal{G}$, after history $h = v_0 \ldots v_\ell$ is given as follows. In the case that $L(h) = 0$ (i.e., no position of minimal priority has been seen), let $\varepsilon' := \varepsilon/2$, and $\sigma^\varepsilon(h) := \sigma_\gamma^{\varepsilon'}(h)$. Otherwise, let v_k be the last node of priority m in the history $h = v_0 \ldots v_\ell$,
$$\varepsilon' := \frac{\varepsilon}{2^{L(h)+1} D(v_0 \ldots v_k)}.$$
and
$$\sigma^\varepsilon(h) := \sigma_\gamma^{\varepsilon'}(v_{k+1} \ldots v_\ell).$$

Now let us consider a play $\pi = v_0 \ldots v_k v_{k+1} \ldots$, consistent with a strategy σ^ε, where v_k is the first node with minimal priority. The following property about values $g_\gamma(v_0)$ and $g_\gamma(v_{k+1})$ in such case (and an analogous, but more tedious one for Player 1) is the main technical point in proving ε-optimality.

Lemma 4.11. $\Delta(v_0 \ldots v_k) \cdot g_\gamma(v_{k+1})$ is $\frac{\varepsilon}{2}$-above $g_\gamma(v_0)$.

With the above lemma we prove the ε-optimality of the strategies σ^ε, as stated in the proposition below.

Proposition 4.12. The strategy σ^ε is ε-optimal, i.e. for every $v \in V$ and every strategy ρ for Player 1, $p(\pi_{\sigma^\varepsilon, \rho}(v))$ is ε-above $g(v)$.

The Strategy of Player 1

Now we look at the situation of Player 1. The problem of Player 1 is that he cannot just combine his strategies for $\mathcal{G}_\gamma^-$. If he did so, he would risk going infinitely often through nodes with minimal priority which is his worst case scenario. Intuitively speaking, he needs a way to count down, so that will be able to come close enough to his desired value, but will stop going through the nodes with minimal priority after a finite number of times. To achieve that, he utilises the strategy index as a counter. Like Player 0, he starts with a strategy for $\mathcal{G}_\gamma^-$, but with every strategy change at the nodes of minimal priority he not only adjusts the approximation value according to the previous one and the discount factors seen so far, but also lowers the strategy index in the following way. If the current game index is a successor ordinal, he just changes the index to its predecessor and adjusts the approximation value in the same way Player 0 does. If the current game index is a limit value, he uses the fact, that there is a game index belonging to a game which has an outcome close enough to still

reach his desired outcome. In the situation depicted below he would choose an α such that $\mathrm{val}\mathcal{G}^-_\alpha(v_{k_1+1})$ is $\frac{\varepsilon}{4}$-below $\lambda_\gamma(v_{k_1})$.
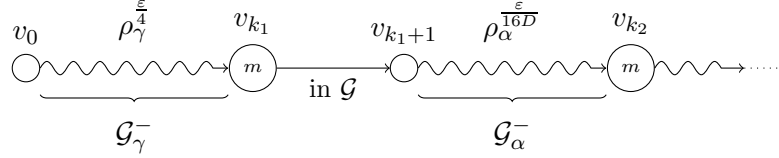
Player 1's strategy at the beginning of the play for a limit ordinal γ.

Finally, after a finite number of changes, as the ordinals are well-founded, he will be playing some version of $\rho_0^{\varepsilon_l}$ and keep on playing this strategy for the rest of the play.

Now we formally describe Player 1's strategy. Let us first fix some notation considering game indices. For a limit ordinal α, a node $v \in V$ of priority m, and for $\varepsilon \in (0,1)$, we denote by $\alpha \upharpoonright \varepsilon, v$ the index for which the value $\mathrm{val}\mathcal{G}^-_\alpha(v)$ is ε-below $\lambda_\alpha(w)$, where $\{w\} = vE$.

Definition 4.13. For a given approximation value ε', a starting ordinal ζ, and a history $h = v_0 \ldots v_l$, we define game indices $\alpha_\zeta(h, \varepsilon')$, approximation values $\varepsilon(h, \varepsilon')$, and a strategy $\rho^{\varepsilon'}$ for Player 1 in the following way.

If $L(h) = 0$, we fix $\alpha_\zeta(h, \varepsilon') = \zeta$ and $\varepsilon(h, \varepsilon') = \varepsilon'$.

For $h = v_0 \ldots v_k v_{k+1} \ldots v_l$, where v_k is the last node with minimal priority in h, let $h' = v_0 \ldots v_{k-1}$ and put

$$\alpha_\zeta(h, \varepsilon') = \begin{cases} \alpha_\zeta(h', \varepsilon') - 1 & \text{for } \alpha_\zeta(h', \varepsilon') \text{ successor ordinal,} \\ \alpha_\zeta(h', \varepsilon') \upharpoonright \left(\frac{\varepsilon'}{4^{L(h')+1}D(h')}, v_k\right) & \text{for } \alpha_\zeta(h', \varepsilon') \text{ limit ordinal,} \\ 0 & \text{for } \alpha_\zeta(h', \varepsilon') = 0, \end{cases}$$

and $\varepsilon(h, \varepsilon') = \frac{\varepsilon'}{4^{L(h)}D(v_0 \ldots v_k)}$.

The ε'-optimal strategy for Player 1 is given by:

$$\rho_\zeta^{\varepsilon'}(v_0 \ldots v_l) = \rho_{\alpha_\zeta(v_0 \ldots v_l, \varepsilon')}^{\frac{\varepsilon(v_0 \ldots v_l, \varepsilon')}{4}}.$$

Proposition 4.14. *The strategy ρ_ζ^ε is ε-optimal, i.e. for every $\varepsilon \in (0,1)$, for all $v \in V$, and strategies σ of Player 0: $p(\pi_{\sigma, \rho_\zeta^\varepsilon}(v))$ is ε-below $g_\zeta(v)$.*

Having defined the ε-optimal strategies σ^ε and ρ_γ^ε, we can formulate the conclusion.

Proposition 4.15. *For a QPG $\mathcal{G} = (V, E, \lambda, \Omega)$, for all $v \in V$,*

$$\sup_{\sigma \in \Gamma_0} \inf_{\rho \in \Gamma_1} p(\pi_{\sigma, \rho}(v)) = \inf_{\rho \in \Gamma_1} \sup_{\sigma \in \Gamma_0} p(\pi_{\sigma, \rho}(v)) = \mathrm{val}\mathcal{G}(v) = g(v).$$

4.3. Quantitative μ-calculus and Games

After establishing determinacy for quantitative parity games we are ready to prove Theorem 4.4. In the proof, we first use structural induction to show that $\mathrm{MC}[\mathcal{K}, \varphi]$ is a model checking game for QML formulae. Further, we only need to inductively consider formulae of the form $\varphi = \nu X.\psi$.

Note that in the game $\text{MC}[\mathcal{Q}, \varphi]$, the positions with minimal priority are of the form (X, v) each with a unique successor (φ, v). Our induction hypothesis states that for every interpretation g of the fixed-point variable X, it holds that:

$$\llbracket \varphi \rrbracket^{\mathcal{Q}}_{[X \leftarrow g]} = \text{valMC}[\mathcal{Q}, \psi[X/g]]. \tag{4.1}$$

By Theorem 2.4, we know that we can compute $\nu X.\psi$ inductively in the following way: $\llbracket \nu X.\psi \rrbracket^{\mathcal{K}}_{\varepsilon} = g_{\gamma}$ with $g_0(v) = \infty$ for all $v \in V$ and

$$g_{\alpha} = \begin{cases} \llbracket \psi \rrbracket_{\varepsilon[X \leftarrow g_{\alpha-1}]} & \text{for } \alpha \text{ successor ordinal,} \\ \lim_{\beta < \alpha} \llbracket \psi \rrbracket_{\varepsilon[X \leftarrow g_{\beta}]} & \text{for } \alpha \text{ limit ordinal,} \end{cases}$$

and where $g_{\gamma} = g_{\gamma+1}$.

Now we want to prove that the games $\text{MC}[\mathcal{Q}, \psi[X/g_{\alpha}]]$ coincide with the unfolding of $\text{MC}[\mathcal{Q}, \varphi]$. We say that two games coincide if the game graph is essentially the same, except for some additional moves where neither player has an actual choice and there is no discount that could change the outcome. In our case these are the moves from $\varphi = \nu X.\psi$ to ψ, which allows us to show the following lemma.

Lemma 4.16. *The games* $\text{MC}[\mathcal{Q}, \psi[X/g_{\alpha}]]$ *and* $\text{MC}[\mathcal{Q}, \varphi]^{-}_{\alpha}$ *coincide for all* α.

From the above and Proposition 4.15, we conclude that the value of the game $\text{MC}[\mathcal{Q}, \varphi]$ is the limit of the values $\text{MC}[\mathcal{Q}, \varphi]^{-}_{\alpha}$, whose value functions coincide with the stages of the fixed-point evaluation g_{α} for all α, and thus

$$\text{valMC}[\mathcal{Q}, \varphi] = \text{valMC}[\mathcal{Q}, \varphi]^{-}_{\gamma} = g_{\gamma} = \llbracket \varphi \rrbracket^{\mathcal{Q}}. \qquad \blacksquare$$

5. Describing Game Values in $\text{Q}\mu$

Having model checking games for the quantitative μ-calculus is just one direction in the relation between games and logic. The other direction concerns the definability of the winning regions in a game by formulae in the corresponding logic. For the classical μ-calculus such formulae have been constructed by Walukiewicz and it has been shown that for any parity game of fixed priority they define the winning region for Player 0, see e.g. [9]. We extend this theorem to the quantitative case in the following way. We represent quantitative parity games $(V, V_0, V_1, E, \delta_G, \lambda_G, \Omega_G)$ with priorities $\Omega(V) \in \{0, \dots d-1\}$ by a quantitative transition system $\mathcal{Q}_{\mathcal{G}} = (V, E, \delta, V_0, V_1, \Lambda, \Omega)$, where $V_i(v) = \infty$ when $v \in V_i$ and $V_i(v) = 0$ otherwise, $\Omega(v) = \Omega_G(v)$ when $vE \neq \emptyset$ and $\Omega(v) = d$ otherwise,

$$\delta(v, w) = \begin{cases} \delta_G(v, w) & \text{when } v \in V_0, \\ \frac{1}{\delta_G(v, w)} & \text{when } v \in V_1, \end{cases}$$

and payoff predicate $\Lambda(v) = \lambda_G(v)$ when $vE = \emptyset$ and $\Lambda(v) = 0$ otherwise.

We then build the formula Win_d and formulate the theorem

$$\text{Win}_d = \nu X_0.\mu X_1.\nu X_2.\dots \lambda X_{d-1} \bigvee_{j=0}^{d-1} ((V_0 \wedge P_j \wedge \Diamond X_j) \vee (V_1 \wedge P_j \wedge \Box X_j)) \vee \Lambda,$$

where $\lambda = \nu$ if d is odd, and $\lambda = \mu$ otherwise, and $P_i := \neg(\mu X.(2 \cdot X \vee |\Omega - i|))$.

Theorem 5.1. *For every* $d \in \mathbb{N}$, *the value of any quantitative parity game* $\mathcal{G}$ *with priorities in* $\{0, \dots d-1\}$ *coincides with the value of* Win_d *on the associated transition system* $\mathcal{Q}_{\mathcal{G}}$.

6. Conclusions and Future Work

In this work, we showed how the close connection between the modal μ-calculus and parity games can be lifted to the quantitative setting, provided that the quantitative extensions of the logic and the games are defined in an appropriate manner. This is just a first step in a systematic investigation of what connections between logic and games survive in the quantitative setting. These investigations should as well be extended to quantitative variants of other logics, in particular LTL, CTL, CTL*, and PDL.

Following [3] we work with games where discounts are multiplied along edges and values range over the non-negative reals with infinity. Another natural possibility is to use addition instead of multiplication and let the values range over the reals with $-\infty$ and $+\infty$. Crash games, recently introduced in [7], are defined in such a way, but with values restricted to integers. Gawlitza and Seidl present an algorithm for crash games over finite graphs which is based on strategy improvement [7]. It is possible to translate back and forth between quantitative parity games and crash games with real values by taking logarithms of the discount values on edges as payoffs for moves in the crash game. The exponent of the value of such a crash game is then equal to the value of the original quantitative parity game. This suggests that the methods from [7] can be applied to quantitative parity games as well. This could lead to efficient model-checking algorithms for $Q\mu$ and would thus further justify the game-based approach to model checking modal logics.

References

[1] Luca de Alfaro. Quantitative verification and control via the mu-calculus. In Roberto M. Amadio and Denis Lugiez, editors, *CONCUR*, volume 2761 of *LNCS*, pages 102–126. Springer, 2003.

[2] Luca de Alfaro, Marco Faella, Thomas A. Henzinger, Rupak Majumdar, and Mariëlle Stoelinga. Model checking discounted temporal properties. *Theoretical Computer Science*, 345(1):139–170, 2005.

[3] Luca de Alfaro, Marco Faella, and Mariëlle Stoelinga. Linear and branching system metrics. Technical Report ucsc-crl-05-01, School of Engineering, University of California, Santa Cruz, 2005.

[4] Luca de Alfaro, Thomas A. Henzinger, and Rupak Majumdar. Discounting the future in systems theory. In Jos C. M. Baeten, Jan Karel Lenstra, Joachim Parrow, and Gerhard J. Woeginger, editors, *ICALP*, volume 2719 of *Lecture Notes in Computer Science*, pages 1022–1037. Springer, 2003.

[5] Luca de Alfaro and Rupak Majumdar. Quantitative solution of omega-regular games. *J. Comput. Syst. Sci.*, 68(2):374–397, 2004.

[6] E. Allen Emerson, Charanjit S. Jutla, and A. Prasad Sistla. On model-checking for fragments of μ-calculus. In *CAV 93*, volume 697 of *Lecture Notes in Computer Science*, pages 385–396. Springer, 1993.

[7] Thomas Gawlitza and Helmut Seidl. Computing game values for crash games. In Kedar S. Namjoshi *et al.*, eds, *ATVA*, *Lect. Notes in Comp. Science* 4762, pp. 177-191. Springer, 2007.

[8] Hugo Gimbert and Wieslaw Zielonka. Perfect information stochastic priority games. In Lars Arge *et al.*, eds, *ICALP*, *Lect. Notes in Comp. Science* 4596, pp. 850-861. Springer, 2007.

[9] Erich Grädel. Finite model theory and descriptive complexity. In *Finite Model Theory and Its Applications*, pages 125–230. Springer-Verlag, 2007.

[10] Marcin Jurdziński. Small progress measures for solving parity games. In Horst Reichel and Sophie Tison, editors, *STACS*, volume 1770 of *Lecture Notes in Computer Science*, pages 290–301. Springer, 2000.

[11] Annabelle McIver and Carroll Morgan. Results on the quantitative μ-calculus qMμ. *ACM Trans. Comput. Log.*, 8(1), 2007.

[12] Colin Stirling. Games and modal mu-calculus. In Tiziana Margaria and Bernhard Steffen, editors, *TACAS*, volume 1055 of *Lecture Notes in Computer Science*, pages 298–312. Springer, 1996.