Model Checking Games for the Quantitative μ-Calculus

Diana Fischer, Erich Grädel and Łukasz Kaiser
Mathematische Grundlagen der Informatik, RWTH Aachen
{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 [12], or used to describe winning conditions in stochastic games [5, 1, 8]. Other variants introduce quantities by allowing discounting in the respective version of 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 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. *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 $f_1 \leq f_2$ if $f_1(v) \leq f_2(v)$ for all v. Then $(\mathcal{F}, \leq)$ 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 $Y \neq X$.

Definition 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* $[\![\varphi]\!]^{\mathcal{K}}_{\varepsilon} : V \to \mathbb{R}^+_\infty$ *defined as follows:*

(1) $[\![|P_i - c|]\!]^{\mathcal{K}}_{\varepsilon}(v) = |P_i(v) - c|$,

(2) $[\![\varphi_1 \wedge \varphi_2]\!]^{\mathcal{K}}_{\varepsilon} = \min\{[\![\varphi_1]\!]^{\mathcal{K}}_{\varepsilon}, [\![\varphi_2]\!]^{\mathcal{K}}_{\varepsilon}\}$ *and* $[\![\varphi_1 \vee \varphi_2]\!]^{\mathcal{K}}_{\varepsilon} = \max\{[\![\varphi_1]\!]^{\mathcal{K}}_{\varepsilon}, [\![\varphi_2]\!]^{\mathcal{K}}\}$,

(3) $[\![\Diamond\varphi]\!]^{\mathcal{K}}_{\varepsilon}(v) = \sup_{v' \in vE} \delta(v, v') \cdot [\![\varphi]\!]^{\mathcal{K}}_{\varepsilon}(v')$ *and* $[\![\Box\varphi]\!]^{\mathcal{K}}_{\varepsilon}(v) = \inf_{v' \in vE} \frac{1}{\delta(v,v')} [\![\varphi]\!]^{\mathcal{K}}_{\varepsilon}(v')$,

(4) $[\![d \cdot \varphi]\!]^{\mathcal{K}}_{\varepsilon}(v) = d \cdot [\![\varphi]\!]^{\mathcal{K}}_{\varepsilon}(v)$,

(5) $[\![X]\!]^{\mathcal{K}}_{\varepsilon} = \varepsilon(X)$,

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

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

For formulae without free variables, we can simply write $[\![\varphi]\!]^{\mathcal{K}}$ *rather than* $[\![\varphi]\!]^{\mathcal{K}}_{\varepsilon}$.

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 $[\![\varphi]\!]^{\mathcal{K}}(v) = \infty$. 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 4. *The least and greatest fixed points exist and can be computed inductively:* $[\![\mu X.\varphi]\!]^{\mathcal{K}}_{\varepsilon} = g_\gamma$ *with* $g_0(v) = 0$ *(and* $[\![\nu X.\varphi]\!]^{\mathcal{K}}_{\varepsilon} = g_\gamma$ *with* $g_0(v) = \infty$*) for all* $v \in V$ *where*

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

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 what we expect from such an operator. Syntactically, we want to add to the formula building rules of Qμ a new rule saying that for every formula $\varphi \in$ Qμ, also $\neg\varphi$ is a formula of Qμ. For fixed point formulae $\mu X.\varphi$ and $\nu X.\varphi$ we then have to require that X only occurs positively (i.e. in the scope of an even number of negation signs) in φ, to guarantee the monotonicity and, accordingly, the existence of the least and greatest fixed points. Semantically, the meaning of negation has to be defined by an operator $f_\neg : \mathbb{R}^+_\infty \to \mathbb{R}^+_\infty$ satisfying the properties outlined in the following definition.

Definition 5. *A* negation operator $f_\neg$ *for* Qμ *is a function* $\mathbb{R}^+_\infty \to \mathbb{R}^+_\infty$*, such that when we define* $[\![\neg\varphi]\!] = f_\neg([\![\varphi]\!])$*, 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 as carried out below shows that the function

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

is a negation operator for Qμ.

Proposition 6. *$f_{\frac{a}{x}}$ is a negation operator in* Qμ *for every $a \in \mathbb{R}^+ \backslash \{0\}$.*

Proof.

(1) $f_{\frac{a}{x}}(f_{\frac{a}{x}}(r)) = r$ for every $r \in \mathbb{R}^+_\infty$

(2) $[\![\neg(\varphi \wedge \psi)]\!]^{\mathcal{K}} = f_{\frac{a}{x}}(\min\{[\![\varphi_1]\!]^{\mathcal{K}}, [\![\varphi_2]\!]^{\mathcal{K}}\})$
$= \max\{f_{\frac{a}{x}}([\![\varphi_1]\!]^{\mathcal{K}}), f_{\frac{a}{x}}([\![\varphi_2]\!]^{\mathcal{K}})\} = [\![\neg\varphi \vee \neg\psi]\!]^{\mathcal{K}}$

(3) $[\![\neg(\varphi \vee \psi)]\!]^{\mathcal{K}} = f_{\frac{a}{x}}(\max\{[\![\varphi_1]\!]^{\mathcal{K}}, [\![\varphi_2]\!]^{\mathcal{K}}\})$
$= \min\{f_{\frac{a}{x}}([\![\varphi_1]\!]^{\mathcal{K}}), f_{\frac{a}{x}}([\![\varphi_2]\!]^{\mathcal{K}})\} = [\![\neg\varphi \wedge \neg\psi]\!]^{\mathcal{K}}$

(4) $[\![\neg\Box\varphi]\!]^{\mathcal{K}} = f_{\frac{a}{x}}(\inf_{v' \in vE} \frac{1}{\delta(v,v')} \cdot [\![\varphi]\!]^{\mathcal{K}}(v'))$
$= \sup_{v' \in vE} \delta(v,v') \cdot f_{\frac{a}{x}}([\![\varphi]\!]^{\mathcal{K}}(v')) = [\![\Diamond\neg\varphi]\!]^{\mathcal{K}}$

(5) $[\![\neg\Diamond\varphi]\!]^{\mathcal{K}} = f_{\frac{a}{x}}(\sup_{v' \in vE} \delta(v,v') \cdot [\![\varphi]\!]^{\mathcal{K}})$
$= \inf_{v' \in vE} \frac{1}{\delta(v,v')} \cdot f_{\frac{a}{x}}([\![\varphi]\!]^{\mathcal{K}}) = [\![\Box\neg\varphi]\!]^{\mathcal{K}}$

(6) $[\![\neg d \cdot \varphi]\!]^{\mathcal{K}} = f_{\frac{a}{x}}(d \cdot [\![\varphi]\!]^{\mathcal{K}}(v))$
$= \frac{1}{d} \cdot f_{\frac{a}{x}}([\![\varphi]\!]^{\mathcal{K}}(v)) = [\![\frac{1}{d} \cdot \neg\varphi]\!]^{\mathcal{K}}$

(7) $[\![\neg\mu X.\varphi]\!]^{\mathcal{K}} = [\![\nu X.\neg\varphi[X/\neg X]]\!]^{\mathcal{K}}$.
We will show this case by induction over the stages of the fixed-point evaluation as in Theorem 4. Let $[\![\mu X.\varphi]\!]^{\mathcal{K}} = \lim_n g_n$, and $[\![\nu X.\neg\varphi[X/\neg X]]\!]^{\mathcal{K}} = \lim_n h_n$. The base case $f_{\frac{a}{x}}(g_0) = h_0$, where $g_0 = 0$ and $h_0 = \infty$ as previously defined, holds by definition of $f_{\frac{a}{x}}$. Assume that $f_{\frac{a}{x}}(g_\alpha) = h_\alpha$ for stage α. The induction step $f_{\frac{a}{x}}(g_{\alpha+1}) = f_{\frac{a}{x}}([\![\varphi]\!]^{\mathcal{K}}_{\varepsilon[X \leftarrow g_\alpha]}) = [\![\neg\varphi]\!]^{\mathcal{K}}_{\varepsilon[X \leftarrow g_\alpha]}$
$= [\![\neg\varphi]\!]^{\mathcal{K}}_{\varepsilon[X \leftarrow f_{\frac{a}{x}}(h_\alpha)]} = [\![\neg\varphi[X/\neg X]]\!]^{\mathcal{K}}_{\varepsilon[X \leftarrow h_\alpha]} = h_{\alpha+1}$
follows from the induction hypothesis, the limit step follows trivially.

(8) $[\![\neg\nu X.\varphi]\!]^{\mathcal{K}} = [\![\mu X.\neg\varphi[X/\neg X]]\!]^{\mathcal{K}}$. The proof is analogous to (7) $\Box$

Hence, we can add an inductive rule for negation to the definition of Qμ. Moreover, we show that $f_{\frac{a}{x}}$ are the only negation operators for Qμ with the required properties. For this purpose we use the following technical lemma.

Lemma 7. *Let $g : \mathbb{R} \to \mathbb{R}$ be a function that satisfies the following conditions:*

(1) $g(x + y) = g(x) + g(y) + c$ for some $c \in \mathbb{R}$,

(2) $g(g(x)) = x$,

(3) $x < y \implies g(y) < g(x)$.

Then $g(x) = -x - c$.

Proof. First, we establish the following equalities:

(i) $g(0) = -c$, because $g(0+0) = g(0) + g(0) + c$ by (1).

(ii) $g(-x) = -g(x) - 2c$ as $-c = g(0) = g(x + (-x)) = g(x) + g(-x) + c$.

(iii) $g(c) = -2c$ as $g(-c) = 0$ by (2) and $g(c) = -g(-c) - 2c$ by (ii).

Let us now compare $x + c$ with $-g(x)$ for arbitrary x.

First, if $x + c > -g(x)$, then

$$\begin{aligned} g(x+c) &< g(-g(x)) && \text{by (3)} \\ g(x) + g(c) + c &< -g(g(x)) - 2c && \text{by (1) and (ii)} \\ g(x) - 2c + c &< -x - 2c && \text{by (iii) and (2),} \end{aligned}$$

so $g(x) < -x - c$, which contradicts the assumption that $x + c > -g(x)$.

The case that $x + c < -g(x)$ is treated analogously and also leads to a contradiction. Hence, $-g(x) = x + c$ and therefore also $g(x) = -x - c$ which concludes our proof. □

Note that we prove that $f_{\frac{a}{x}}$ are the only negation operators even for non-discounted transition systems. Observe that each $f_{\frac{a}{x}}$ operates on discounts so that the function $\beta(d) = \frac{1}{d}$. This motivates our definition of the semantics of Qμ, in particular it explains the $\frac{1}{\delta(v,v')}$ factor for $[\![\Box\varphi]\!]^{\mathcal{K}}$ in Definition 3.

Theorem 8. *$f_{\frac{a}{x}}$ for $a \in \mathbb{R}^+ \setminus \{0\}$ are the only negation operators for* Qμ*, even for non-discounted transition systems.*

Proof. According to property (4) in Definition 5, we require $f_\neg(d \cdot x) = \beta(d) \cdot f_\neg(x)$ for some β. If we take $x = 1$, we get $\beta(d) = \frac{f_\neg(d)}{f_\neg(1)}$. Let $a = \frac{1}{f_\neg(1)}$, so $f_\neg(d \cdot x) = f_\neg(d) \cdot f_\neg(x) \cdot a$. Now let $g(x) = \ln f_\neg(e^x)$. From our considerations above, we get

$$g(x+y) = \ln f_\neg(e^{(x+y)}) = \ln f_\neg(e^x \cdot e^y) = g(x) + g(y) + \ln(a).$$

By property (1) we require that $f_\neg(f_\neg(x)) = x$. By definition of g we have $f_\neg(x) = e^{g(\ln(x))}$ which implies that $g(g(\ln(x))) = \ln(x)$. As ln is a function onto $\mathbb{R}$, we have $g(g(x)) = x$, and as both ln and exp are monotone, g satisfies conditions (1) – (3) of Lemma 7 and thus $g(x) = -x - a$. Thus, $f_\neg(x) = \frac{a}{x}$ and $\beta(d) = \frac{1}{d}$. □

The canonical choice for negation in Qμ is $f_{\frac{1}{x}}$. The dualities between $\wedge$ and $\vee$, $\Diamond$ and $\Box$, and between least and greatest fixed points imply that Qμ has a negation normal form: every formula can be translated into one in which negation is applied only to atoms.

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 9. *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* $\Omega : V \to \{0, \ldots, n\}$ *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 $\pi = v_0 v_1 \ldots$ the (possibly infinite) play through vertices $v_0 v_1 \ldots$, given that $(v_n, v_{n+1}) \in E$ for every n. The outcome $p(\pi)$ of a finite play $\pi = v_0 \ldots v_k$ can be computed by multiplying all discount factors seen throughout the play with the value of the final node,

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

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 $\pi = v_0 v_1 \ldots$ is *consistent with a strategy* s for player i, if $v_{n+1} = s(v_0 \ldots v_n)$ 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 coincide,

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

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. There are already simple quantitative games where no player has a positional winning strategy. Consider Example 10 where, by convention, we depict positions of Player 0 with a circle and the ones of Player 1 with a square. The number inside a node stands for its priority and the one written below terminal nodes denotes the payoff. In this 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.

Example 10.

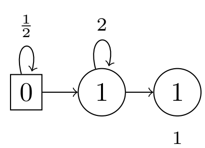

4.1 Model Checking Games for Qμ

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 11. *For a quantitative transition system $\mathcal{K} = (S, T, \delta_S, P_i)$ and a Qμ-formula φ, the quantitative parity game* $\mathrm{MC}[\mathcal{K}, \varphi] = (V, V_0, V_1, E, \delta, \lambda, \Omega)$*, 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 $\Box, \wedge,$ 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 $(\Box\psi, s)$ 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 $(\Box\psi, s)$ 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 [13]. A proof that uses the unfolding technique can be found in [9]. We generalise this connection to the quantitative setting as follows.

Theorem 12. *For every formula φ in Qμ, a quantitative transition system $\mathcal{K}$, and $v \in \mathcal{K}$, the game $\mathrm{MC}[\mathcal{K}, \varphi]$ is determined and*

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

By using the method of unfolding, we give a direct proof for both the determinacy of quantitative parity games and the connection to model checking Qμ. An alternative method to prove the determinacy of quantitative parity games would be by means of Martin's theorem [11] that guarantees the determinacy of all two-valued Borel games. To apply this result to a quantitative game $\mathcal{G}$, consider for each $x \in \mathbb{R}^+_\infty$ the two-valued game $\mathcal{G}_x$ which is identical to $\mathcal{G}$ except that Player 0 wins in $\mathcal{G}_x$ if the payoff of the corresponding play in $\mathcal{G}$ is greater than x, and Player 1 wins in the other case. As sets of paths defined by the parity condition are Borel, and quantitative payoffs appear only on finite paths, it follows from Martin's theorem that all the games $\mathcal{G}_x$ are determined. The value of $\mathcal{G}$ is then the supremum of all x such that Player 0 wins $\mathcal{G}_x$.

Example 13. *A model checking game for $\varphi = \mu X.(P \vee 2 \cdot \Diamond X)$ on the QTS $\mathcal{Q}$ shown in Figure 1 (left), with $P(a) = 0$, $P(b) = 1$, is also depicted in Figure 1 (right). The nodes are labelled with the corresponding subformulae of φ, and the state of $\mathcal{Q}$. Only the edges with discount factor different from 1 are labelled.*

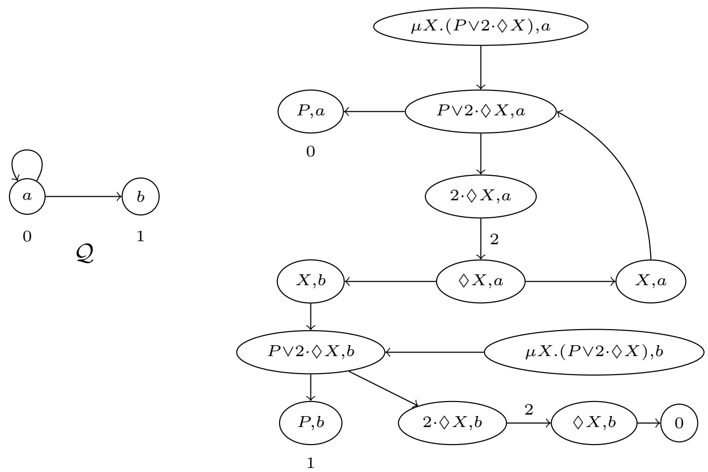

Figure 1: QTS $\mathcal{Q}$ and Model Checking Game for $\mu X.(P \vee 2 \cdot \Diamond X)$ and $\mathcal{Q}$.

Note that in this game only Player 0 *is allowed to make any choices. When we start at the top node, corresponding to an evaluation of φ at a in $\mathcal{Q}$, the only choice she has to make is either to keep playing (by looping), or to end the game by moving to a terminal position.*

4.2 Unfolding Quantitative Parity Games

To prove the model checking theorem in the quantitative case, we start with games with one priority. We call these games *reachability games* if the only priority is odd and *safety games* if it is even. As reachability and safety games are dual to each other, let us focus on reachability games where the only priority is odd and will be assigned to every node, i.e. infinite plays will have outcome 0, meaning that Player 1 wins. The construction of ε-optimal strategies is obtained by a generalisation of backwards induction and is carried out in detail below. At first, we fix the notation and show a few basic properties.

Definition 14. *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 15. *For every history $h = v_0 \dots v_\ell$ 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 $\pi = v_0 v_1 \dots$ and every k,

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

Lemma 16. *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.

To determine the value of a play in a reachability game after k steps, we use backwards induction, and we inductively define a sequence of approximate payoff functions $f_i : V \to \mathbb{R}^+_\infty$.

The first payoff function corresponds to the immediate payoff of the game.

$$f_0(v) = \begin{cases} \lambda(v) & \text{for } vE = \emptyset \\ 0 & \text{otherwise} \end{cases}$$

The payoff function f_{i+1} corresponds to the payoff that can be guaranteed for Player 0 in $i+1$ steps.

$$f_{i+1}(v) = \begin{cases} \lambda(v) & \text{for } vE = \emptyset \\ \sup_{w \in vE} \delta(v,w) \cdot f_i(w) & \text{for } v \in V_0 \\ \inf_{w \in vE} \delta(v,w) \cdot f_i(w) & \text{for } v \in V_1 \end{cases}$$

Intuitively, these functions determine the value of a game that ends after at most k steps. This means that the outcome will be 0 if Player 0 does not succeed to reach a terminal position in at most k steps.

We define $f(v)$ as $\lim_{i\to\infty} f_i(v)$ and we show that $f(v)$ is indeed the value of a game started at v. To show that $f(v)$ is well-defined, we note that the sequence f_i is monotonically increasing, i.e. $f_{i+1}(v) \geq f_i(v)$ for all v, i. This can be easily proved by induction on i. Moreover, by commutativity of lim and sup, inf, the properties in the definition of f_i are sustained in f.

Lemma 17. *For all moves $(v,w) \in E$,*

$$f(v) = \begin{cases} \lambda(v) & \textit{for } vE = \emptyset \\ \sup_{w \in vE} \delta(v,w) \cdot f(w) & \textit{for } v \in V_0 \\ \inf_{w \in vE} \delta(v,w) \cdot f(w) & \textit{for } v \in V_1 \end{cases}$$

Let us define the strategies that approximate the payoffs f_k, σ^k_ε for Player 0, corresponding to f_k, and ρ_ε for Player 1, corresponding to f. Given two non-zero real numbers ε, δ, let $\text{next}(\varepsilon, \delta) := \varepsilon/(2\max(\delta, \delta^{-1}))$. For σ^k_ε choose at v a successor node $w \in vE$ such that $\delta(v,w) \cdot f_{k-1}(w)$ is $\frac{\varepsilon}{2}$-close to $f_k(v)$. From w on, play according to $\sigma^{k-1}_{\varepsilon'}$, where $\varepsilon' = \text{next}(\varepsilon, \delta(v,w))$ For ρ_ε in v proceed in an analogous way: choose a successor node w such that $\delta(v,w) \cdot f(w)$ is $\frac{\varepsilon}{2}$-close to $f(v)$ and play $\rho_{\varepsilon'}$ from w. When the opponent makes a move, adjust ε to ε' in the same way.

Note that if the game is finitely branching, the ε-approximations are not necessary as one can choose the maximal and minimal value directly.

Lemma 18. *$p(\pi_{\sigma_\varepsilon^k,\rho}(v))$ is ε-above $f_k(v)$ for every strategy ρ of Player 1.*

Proof. We prove this lemma by induction over k. For $k = 0$, $f_0(v) \neq 0$ only in case that v is a terminal node, but then $f_0(v) = \lambda(v) = p(\pi_{\sigma_\varepsilon^0,\rho}(v))$.

By induction hypothesis $p(\pi_{\sigma_{\varepsilon'}^{k-1},\rho}(w))$ is ε'-above $f_{k-1}(w)$ for all $w \in V$ and $\varepsilon' \in (0,1)$. In particular, by Lemma 16, when $\varepsilon' = \text{next}(\varepsilon, \delta(v,w))$ for a predecessor v of w, then $\delta(v,w) \cdot p(\pi_{\sigma_{\varepsilon'}^{k-1},\rho}(w))$ is $\frac{\varepsilon}{2}$-above $\delta(v,w) \cdot f_{k-1}(w)$.

If $v \in V_0$ then σ_ε^k by definition chooses w such that $\delta(v,w) \cdot f_{k-1}(w)$ is $\frac{\varepsilon}{2}$-close to $f_k(v)$. Thus, by the above and Lemma 16, $p(\pi_{\sigma_\varepsilon^k,\rho}(v)) = \delta(v,w) \cdot p(\pi_{\sigma_{\varepsilon'}^{k-1},\rho}(w))$ is ε-above $f_k(v)$.

If $v \in V_1$ then ρ chooses any successor w', by definition

$$f_k(v) = \inf_{w \in vE} \delta(v,w) \cdot f_{k-1}(w) \leq \delta(v,w') \cdot f_{k-1}(w'),$$

and thus $p(\pi_{\sigma_\varepsilon^k,\rho}(v)) = \delta(v,w') \cdot p(\pi_{\sigma_{\varepsilon'}^{k-1},\rho}(w'))$ is even $\frac{\varepsilon}{2}$-above $f_k(w)$. ☐

Lemma 19. *The strategies $\sigma_{\frac{\varepsilon}{2}}^k$ are ε-optimal, i.e. for every ε and v there is a k such that $p(\pi_{\sigma_{\frac{\varepsilon}{2}}^k,\rho}(v))$ is ε-above $f(v)$ for every strategy ρ of Player 1.*

Proof. As $f(v) = \lim_{i\to\infty} f_i(v)$, for every ε there is a k such that $f_k(v)$ is $\frac{\varepsilon}{2}$-close to $f(v)$. By Lemma 18, $p(\pi_{\sigma_{\frac{\varepsilon}{2}}^k,\rho}(v))$ is $\frac{\varepsilon}{2}$-above $f_k(v)$ and thus by Lemma 16 it is ε-above $f(v)$. ☐

Lemma 20. *The strategy ρ_ε is ε-optimal, i.e. $p(\pi_{\sigma,\rho_\varepsilon}(v))$ is ε-below $f(v)$ for every strategy σ of Player 0.*

Proof. Towards a contradiction, assume there is a σ such that $p(\pi_{\sigma,\rho_\varepsilon}(v)) > f(v) + \varepsilon$. Then $p(\pi_{\sigma,\rho}(v)) > 0$, and thus $\pi_{\sigma,\rho}(v)$ is finite.

We show by induction on the length k of $\pi_{\sigma,\rho}(v)$ that $p(\pi_{\sigma,\rho_\varepsilon}(v))$ is ε-below $f(v)$. The case $k = 0$ means that v is a terminal position and then $p(\pi_{\sigma,\rho_\varepsilon}(v)) = \lambda(v) = f(v)$.

For the induction step, let $\pi_{\sigma,\rho_\varepsilon}(v)) = v_0 v_1 \dots v_k$. By definition of ρ_ε, the play $v_1 \dots v_k$ is played consistent with $\rho_{\varepsilon'}$ and therefore, by inductive assumption, $p(v_1 \dots v_k)$ is ε'-below $f(v_1)$. Since $\varepsilon' = \text{next}(\varepsilon, \delta(v_0, v_1))$ we have, by Lemma 16, that $p(v_0 v_1 \dots v_k) = \delta(v_0, v_1) \cdot p(v_1 \dots v_k)$ is $\frac{\varepsilon}{2}$-below $\delta(v_0, v_1) \cdot f(v_1)$.

If $v_0 \in V_0$ then $f(v_0) = \sup_{w \in v_0 E} \delta(v_0, w) \cdot f(w) \geq \delta(v_0, v_1) \cdot f(v_1)$ and so $p(v_0 v_1 \dots v_k)$ is even $\frac{\varepsilon}{2}$-below $f(v_0)$. If $v_0 \in V_1$ then, by definition, ρ_ε chooses a v_1 so that $\delta(v_0, v_1) \cdot f(v_1)$ is $\frac{\varepsilon}{2}$-close to $f(v_0)$ and thus, by Lemma 16, $p(v_0 v_1 \dots v_k)$ is ε-below $f(v_0)$. ☐

Proposition 21. *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 12 by showing that, as in the classical case, the unfolding of $\mathrm{MC}[\mathcal{K}, \varphi]$ 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 22. *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:

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

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 $\mathrm{val}\mathcal{G}^-_\gamma = \mathrm{val}\mathcal{G}^-_{\gamma+1}$. 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 $g(v) = g_\gamma(v) = \mathrm{val}\mathcal{G}^-_\gamma(v)$ 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 $\sigma_\gamma^{\varepsilon_i}$.

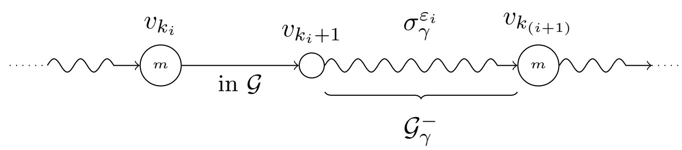

Player 0's strategy after having seen i nodes of priority m.

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 23. *The strategy σ^ε for Player 0 in the game $\mathcal{G}$, after history $h = v_0 \dots v_\ell$ 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 $h = v_0 \dots v_\ell$,*

$$\varepsilon' := \frac{\varepsilon}{2^{L(h)+1} D(v_0 \dots v_k)}.$$

and

$$\sigma^\varepsilon(h) := \sigma_\gamma^{\varepsilon'}(v_{k+1} \dots v_\ell).$$

Now let us consider a play $\pi = v_0 \dots v_k v_{k+1} \dots$, 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) allows us to prove the ε-optimality of the strategies σ^ε, as stated in the proposition below.

Lemma 24. $\Delta(v_0 \dots v_k) \cdot g_\gamma(v_{k+1})$ *is* $\frac{\varepsilon}{2}$*-above* $g_\gamma(v_0)$.

Proof. Let us look at π as played in $\mathcal{G}^-_{\gamma+1} = \mathcal{G}^-_\gamma$. By definition, $\lambda_{\gamma+1}(v_k) = g_\gamma(v_{k+1})$, hence $p(\pi) = \Delta(v_0 \dots v_k) \cdot \lambda_{\gamma+1}(v_k) = \Delta(v_0 \dots v_k) \cdot g_\gamma(v_{k+1})$. As $\sigma^{\frac{\varepsilon}{2}}_\gamma$ is $\frac{\varepsilon}{2}$-optimal in $\mathcal{G}^-_\gamma$, we know that $p(\pi)$ is $\frac{\varepsilon}{2}$-above $\mathrm{val}\mathcal{G}^-_{\gamma+1}(v_0) = g_{\gamma+1}(v_0) = g_\gamma(v_0)$. □

Proposition 25. *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 $g(v)$.*

Proof. Let us fix v and a strategy ρ for Player 1. We distinguish the following two cases.

Case 1: $\pi_{\sigma^\varepsilon,\rho}(v)$ visits nodes of minimal priority infinitely often.
In this case, the outcome of the play is ∞ and there is nothing left to show.

Case 2: $\pi_{\sigma^\varepsilon,\rho}(v)$ visits nodes of minimal priority only finitely often.
We will prove this case by induction over the number of nodes with minimal priority occurring during the play. If $L(\pi_{\sigma^\varepsilon,\rho}(v)) = 0$, then the whole play is equivalent to a play in $\mathcal{G}^-_\gamma$ and σ^ε is equivalent to the $\frac{\varepsilon}{2}$-optimal strategy $\sigma^{\frac{\varepsilon}{2}}_\gamma$ and thus gives a payoff $\frac{\varepsilon}{2}$-above $g_\gamma(v) = g(v)$.

Now let us look at a play $\pi_{\sigma^\varepsilon,\rho}(v) = v_0 \dots v_k v_{k+1} \dots$, where v_k is the first node with minimal priority and $L(\pi_{\sigma^\varepsilon,\rho}(v)) = n$. For a play suffix s, let

$$\rho'(s) = \rho(v_0 \dots v_k s) \text{ and } \sigma'(s) = \sigma^{\frac{\varepsilon}{2D(v_0 \cdots v_k)}}(s) = \sigma^\varepsilon(v_0 \dots v_k s),$$

i.e. we play a part of $\pi_{\sigma^\varepsilon,\rho}$, starting at v_{k+1}.

As $L(\pi_{\sigma',\rho'}(v_{k+1})) = n - 1$, by induction hypothesis, $p(\pi_{\sigma^\varepsilon,\rho'}(v_{k+1}))$ is $\frac{\varepsilon}{2D(v_0 \dots v_k)}$-above $g(v_{k+1})$. By Lemma 16 it follows, that $\Delta(v_0 \dots v_k) \cdot p(\pi_{\sigma^\varepsilon,\rho'}(v_{k+1}))$ is $\frac{\varepsilon}{2}$-above $\Delta(v_0 \dots v_k) \cdot g(v_{k+1})$. Using Lemma 24, we get that $\Delta(v_0 \dots v_k) \cdot g_\gamma(v_{k+1})$ is $\frac{\varepsilon}{2}$-above $g_\gamma(v_0)$. By transitivity (Lemma 16) and the above, we obtain that $p(\pi_{\sigma^{\frac{\varepsilon}{2}}_\gamma,\rho}(v_0)) = \Delta(v_0 \dots v_k) \cdot p(\pi_{\sigma^\varepsilon,\rho'}(v_{k+1}))$ is ε-above $g_\gamma(v_0)$ and so we get the conclusion. □

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.

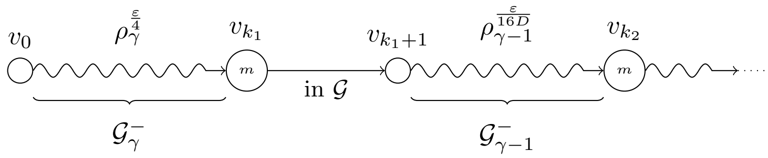

Player 1's strategy at the beginning of the play for γ a successor ordinal.

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 val$\mathcal{G}_\alpha^-(v_{k_1+1})$ is $\frac{\varepsilon}{4}$-below $\lambda_\gamma(v_{k_1})$.

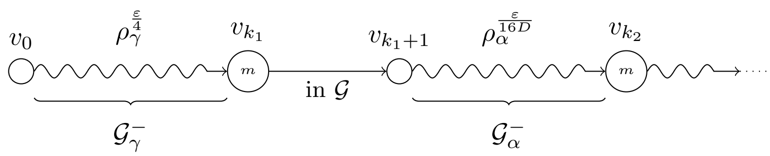

Player 1's strategy at the beginning of the play for a limit ordinal γ.

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 val$\mathcal{G}_\alpha^-(v)$ is ε-below $\lambda_\alpha(w)$, where $\{w\} = vE$.

Definition 26. *For a given approximation value ε', a starting ordinal ζ, and a history $h = v_0 \dots v_l$, 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 $h = v_0 \dots v_k v_{k+1} \dots v_l$, where v_k is the last node with minimal priority in h, let $h' = v_0 \dots v_{k-1}$ and put

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

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

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

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

We motivate the above definition with a nice property of a play consistent with such a strategy for Player 1. Let $\pi = v_0 \dots v_{k_1} v_{k_1+1} \dots v_{k_l} v_{k_l+1} \dots$ be a play that is consistent with a strategy $\rho_\zeta^{\varepsilon'}$, where v_{k_i} are the nodes with minimal priority.

To simplify the notation, let $k_0 = -1$, and $\alpha_0 = \zeta, \alpha_i = \alpha_\zeta(v_0 \dots v_{k_i}, \varepsilon')$ and $\varepsilon_i = \frac{\varepsilon'}{4^i D(v_0 \dots v_{k_i})} = \varepsilon(v_0 \dots v_{k_i}, \varepsilon')$.

Lemma 27. *$\Delta(v_{k_i+1} \dots v_{k_{(i+1)}}) \cdot g_{\alpha_{i+1}}(v_{k_{(i+1)}+1})$ is $\frac{\varepsilon_i}{2}$-below $g_{\alpha_i}(v_{k_i+1})$.*

Proof. We will show two steps from which the lemma follows by transitivity (Lemma 16).

(1) $\Delta(v_{k_i+1} \dots v_{k_{(i+1)}}) \cdot g_{\alpha_{i+1}}(v_{k_{(i+1)}+1})$ is $\frac{\varepsilon_i}{4}$-below $\Delta(v_{k_i+1} \dots v_{k_{(i+1)}}) \cdot \lambda_{\alpha_i}(v_{k_{(i+1)}})$.

(2) $\Delta(v_{k_i+1} \dots v_{k_{(i+1)}}) \cdot \lambda_{\alpha_i}(v_{k_{(i+1)}})$ is $\frac{\varepsilon_i}{4}$-below $g_{\alpha_i}(v_{k_i+1})$.

Step 1. By choice of α_{i+1}, we know that $g_{\alpha_{i+1}}(v_{k_{(i+1)}+1})$ is

$$\frac{\varepsilon'}{4^{i+1} D(v_0, \dots, v_{k_{(i+1)}})} = \frac{\varepsilon_i}{4 \cdot D(v_{k_i+1} \dots v_{k_{(i+1)}})}\text{-below } \lambda_{\alpha_i}(v_{k_{(i+1)}}).$$

In fact, in case that α_i is a successor ordinal, it holds that $g_{\alpha_{i+1}}(v_{k_{(i+1)}+1})) = \lambda_{\alpha_i}(v_{k_{(i+1)}})$.

Therefore, by Lemma 16, $\Delta(v_{k_i+1} \dots v_{k_{(i+1)}}) \cdot g_{\alpha_{i+1}}(v_{k_{(i+1)}+1})$ is $\frac{\varepsilon_i}{4}$-below $\Delta(v_{k_i+1} \dots v_{k_{(i+1)}}) \cdot \lambda_{\alpha_i}(v_{k_{(i+1)}})$ as intended.

Step 2. On the part of the play $\pi_i = v_{k_i+1} \dots v_{k_{(i+1)}}$, the strategy ρ_ζ^ε is equivalent to the $\frac{\varepsilon_i}{4}$-optimal strategy $\rho_{\alpha_i}^{\varepsilon_i}$ in $\mathcal{G}_{\alpha_i}^-$ and thus yields a payoff $p(\pi_i)$ that is $\frac{\varepsilon_i}{4}$-below $g_{\alpha_i}(v_{k_i+1})$. By definition, $p(\pi_i) = \Delta(v_{k_i+1} \dots v_{k_{(i+1)}}) \cdot \lambda_{\alpha_i}(v_{k_{(i+1)}})$, which completes the proof. □

Using the lemma above, we can now prove the ε-optimality of Player 1's strategy.

Proposition 28. *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 $g_\zeta(v)$.*

Proof. Let us fix a strategy σ for Player 0 and distinguish two cases.

Case 1: $\pi_{\sigma,\rho_\zeta^\varepsilon}(v)$ visits nodes of minimal priority infinitely often.

In this case, we show that $g_\zeta(v) = \infty$. Towards a contradiction, assume that $g_\zeta(v) < \infty$ and consider,

$$\pi_{\sigma,\rho_\zeta^\varepsilon}(v) = v_0 \dots v_{k_1} v_{k_1+1} \dots v_{k_l} v_{k_l+1} \dots,$$

where v_{k_i} are the nodes with minimal priority. Now we can use Lemma 27, i.e. we know for $k_0 = -1$, and $\alpha_0 = \zeta, \alpha_i = \alpha_\zeta(v_0 \dots v_{k_i}, \varepsilon)$, that

$$\Delta(v_{k_i+1} \dots v_{k_{(i+1)}}) \cdot g_{\alpha_{i+1}}(v_{k_{(i+1)}+1}) \text{ is } \frac{\varepsilon_i}{2}\text{-below } g_{\alpha_i}(v_{k_i+1}).$$

In particular, if $g_{\alpha_i}(v_{k_i+1})$ is finite, then $g_{\alpha_{i+1}}(v_{k_{(i+1)}+1})$ is finite as well, as ∞ is only ε-close to ∞ (see the remark after the definition of ε-closeness).

The sequence $\alpha_0 > \alpha_1 > \dots$ is a strictly decreasing sequence of ordinals and therefore, for some l, $\alpha_l = 0$. But in the game $\mathcal{G}_0^-$, node v_{k_l} is a terminal node, and thus $g_{\alpha_l}(v_{k_l}) = g_0(v_{k_l}) = \lambda_0(v_{k_l}) = \infty$, which, using the above, contradicts $g_\zeta(v) < \infty$.

Case 2: $\pi_{\sigma,\rho_\zeta^\varepsilon}(v)$ visits nodes of minimal priority only finitely often. We will prove this case by induction over the number of nodes with minimal priority.

If $L(\pi_{\sigma,\rho_\zeta^\varepsilon}(v)) = 0$ then the whole play is equivalent to a play in $\mathcal{G}_\zeta$ and ρ_ζ^ε is equivalent to the $\frac{\varepsilon}{4}$-optimal strategy $\rho_\zeta^{\frac{\varepsilon}{4}}$ and thus gives a payoff $\frac{\varepsilon}{4}$-below $g_\zeta(v)$. Now let us look at a play $\pi_{\sigma,\rho_\zeta^\varepsilon}(v) = v_0 \dots v_k v_{k+1} \dots$, where v_k is the first node with minimal priority and $L(\pi_{\sigma,\rho_\zeta^\varepsilon}(v)) = n$.

For a play suffix s, let $\sigma'(s) = \sigma(v_0 \dots v_k s)$ and

$$\rho'(s) = \rho_{\alpha_1}^{\frac{\varepsilon}{4D(v_0 \dots v_k)}}(s) = \rho_{\alpha_0=\zeta}^\varepsilon(v_0 \dots v_k s).$$

As $L(\pi_{\sigma',\rho'}(v_{k+1})) = n-1$, by induction hypothesis, $p(\pi_{\sigma',\rho'}(v_{k+1}))$ is $\frac{\varepsilon}{4D(v_0 \dots v_k)}$-below $g_{\alpha_1}(v_{k+1})$.

By Lemma 16, $\Delta(v_0 \dots v_k) \cdot p(\pi_{\sigma',\rho'}(v_{k+1}))$ is $\frac{\varepsilon}{4}$-below $\Delta(v_0 \dots v_k) \cdot g_{\alpha_1}(v_{k+1})$.

Now we can use Lemma 27, which tells us, that $\Delta(v_0 \dots v_k) \cdot g_{\alpha_1}(v_{k+1})$ is $\frac{\varepsilon}{2}$-below $g_{\alpha_0}(v_0)$.

By Lemma 16, we get that $p(\pi_{\sigma,\rho_\zeta^\varepsilon(v_0)}) = \Delta(v_0 \dots v_k) \cdot p(\pi_{\sigma',\rho'}(v_{k+1}))$ is $\frac{\varepsilon}{4} + \frac{\varepsilon}{2} < \varepsilon$-below $g_\zeta(v_0)$, which completes the proof. $\square$

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

Proposition 29. *For a quantitative parity game $\mathcal{G} = (V, E, \lambda, \Omega)$, for all $v \in V$,*

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

Note that the game $\mathcal{G}$ in the proposition above can be of any cardinality. Additionally, observe that for any quantitative parity game one can construct an infinite quantitative parity game with the same value in which all discounts are equal to one. To construct such a game it is sufficient to introduce a memory structure to store the product of discounts seen during the play. Then we take the synchronous product of the original game with such a memory structure and define the payoff in the terminal positions of the new game as the product of the original payoff and the value stored in the memory.

From the above remark follows that we only needed to prove Proposition 29 for games without discounts. Our motivation to present the proof for arbitrary discounted games is that it gives better insight into the construction of the strategies, in particular into the error bounds needed to achieve an ε-optimal strategy with one more priority.

4.3 Quantitative μ-calculus and Games

After establishing determinacy for quantitative parity games we are ready to prove Theorem 12. Let us first look at formulae without fixed-point operators and prove the following lemma.

Lemma 30. $\mathrm{MC}[\mathcal{K},\varphi]$ *is a model checking game for* $\varphi \in \mathrm{QML}$.

Proof. Note that in case of QML-model checking all plays are finite and the game graph is a tree of finite depth. The value of a the game from node v is $f(v)$ as defined in section 4.2. Note that the process of computing the sequence f_i stops after a finite number of steps, namely the depth of the game tree. We have to show by induction on the structure of φ that indeed $f(v) = [\![\varphi]\!](v)$.

In case that φ is the distance between a predicate and a constant, the value of the formula is the same as the evaluation of the terminal position. In case that $\varphi = \varphi_1 \wedge \varphi_2$, we have the evaluation $[\![\varphi]\!](v) = \min\{[\![\varphi_1]\!](v), [\![\varphi_2]\!](v)\}$. But that also means, that in the model checking game, the corresponding position belongs to Player 1 and the value $f(v)$ is computed as $\min_{w\in vE} f(w)$, where the next positions w are the subformulae φ_1, φ_2, evaluated at v. By induction hypothesis, the values $f(w)$ coincide with $[\![\varphi_1]\!](v)$ and $[\![\varphi_2]\!](v)$.

In case that $\varphi = \Box\varphi'$, we have $[\![\varphi]\!](v) = \inf_{w\in vE} \frac{1}{\delta(v,w)}[\![\varphi]\!](w)$. Hence, in the model checking game, the corresponding position belongs to Player 1 and the value $f(v)$ is computed as $\inf_{w\in vE} \frac{1}{\delta(v,w)} f(w)$, where the next positions are the subformula φ' evaluated at each of the successor nodes w in the original transition system. By induction hypothesis, these values coincide with $[\![\varphi']\!](w)$. If $\varphi = d \cdot \varphi'$, we have $[\![\varphi]\!](v) = d \cdot [\![\varphi']\!](v)$. In the model checking game this position has only one successor $w = vE$, which corresponds to the evaluation of subformula φ' at state v, and the value is computed as $f(v) = d \cdot f(w)$. Again by induction hypothesis, the value $f(w)$ coincides with $[\![\varphi']\!](v)$. $\Box$

Further, we only need to inductively consider formulae of the form $\varphi = \nu X.\psi$.

Note that in the game $\mathrm{MC}[\mathcal{Q},\varphi]$, 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:

$$[\![\varphi]\!]^{\mathcal{Q}}_{[X\leftarrow g]} = \mathrm{valMC}[\mathcal{Q}, \psi[X/g]]. \tag{1}$$

By Theorem 4, we know that we can compute $\nu X.\psi$ inductively in the following way: $[\![\nu X.\psi]\!]^{\mathcal{K}}_{\varepsilon} = g_\gamma$ with $g_0(v) = \infty$ for all $v \in V$ and

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

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

Now we want to prove that the games $\mathrm{MC}[\mathcal{Q}, \psi[X/g_\alpha]]$ coincide with the unfolding of $\mathrm{MC}[\mathcal{Q},\varphi]$. 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 31. *The games* $\mathrm{MC}[\mathcal{Q}, \psi[X/g_\alpha]]$ *and* $\mathrm{MC}[\mathcal{Q}, \varphi]^-_\alpha$ *coincide for all* α.

Proof. Considering $\alpha = 0$, note that $\mathrm{MC}[\mathcal{Q}, \psi[X/g_0]]$ coincides with $\mathrm{MC}[\mathcal{Q}, \varphi]^-_0$ by construction. The induction hypothesis is that $\mathrm{MC}[\mathcal{Q}, \psi[X/g_\alpha]]$ coincides with $\mathrm{MC}[\mathcal{Q}, \varphi]^-_\alpha$ for some α.

If α is a successor ordinal: the interesting positions in $\mathrm{MC}[\mathcal{Q}, \psi[X/g_\alpha]]$ are the terminal positions of the form (g_α, a) with valuation

$$\lambda(g_\alpha, a) = g_\alpha(a) = [\![\psi]\!]_{[X \leftarrow g_{\alpha-1}]}(a).$$

In $\mathrm{MC}[\mathcal{Q}, \varphi]^-_\alpha$, the corresponding terminal positions are of the form (X, a) with valuation

$$\lambda_\alpha(X, a) = \mathrm{valMC}[\mathcal{Q}, \varphi]^-_{\alpha-1}(\varphi, a).$$

By induction hypothesis, we have $\mathrm{MC}[\mathcal{Q}, \varphi]_{\alpha-1}$ coincides with $\mathrm{MC}[\mathcal{Q}, \psi[X/g_{\alpha-1}]]$ and therefore, by equation 1 (in section 4.3),

$$\mathrm{valMC}[\mathcal{Q}, \varphi]^-_{\alpha-1}(\varphi, a) = \mathrm{valMC}[\mathcal{Q}, \psi[X/g_{\alpha-1}]](\psi, a) = [\![\psi]\!]_{[X \leftarrow g_{\alpha-1}]}(a).$$

If α is a limit ordinal: we only consider terminal positions in $\mathrm{MC}[\mathcal{Q}, \psi[X/g_\alpha]]$ of the form (g_α, a) with valuation

$$\lambda(g_\alpha, a) = g_\alpha(a) = \lim_{\beta < \alpha} [\![\psi]\!]_{[X \leftarrow g_\beta]}(a).$$

In $\mathrm{MC}[\mathcal{Q}, \varphi]^-_\alpha$, the corresponding terminal positions are of the form (X, a) with valuation

$$\lambda_\alpha(X, a) = \lim_{\beta < \alpha} \mathrm{valMC}[\mathcal{Q}, \varphi]^-_\beta(\varphi, a).$$

By induction hypothesis, we have that for all $\beta < \alpha$ $\mathrm{MC}[\mathcal{Q}, \varphi]_\beta$ coincides with $\mathrm{MC}[\mathcal{Q}, \psi[X/g_\beta]]$ and therefore, by equation 1,

$$\lim_{\beta < \alpha} \mathrm{valMC}[\mathcal{Q}, \varphi]^-_\beta(\varphi, a) = \lim_{\beta < \alpha} \mathrm{valMC}[\mathcal{Q}, \psi[X/g_\beta]](\psi, a) = \lim_{\beta < \alpha} [\![\psi]\!]_{[X \leftarrow g_\beta]}(a).$$

□

From the above lemma and Proposition 29, we conclude that the value of the game $\mathrm{MC}[\mathcal{Q}, \varphi]$ is the limit of the values $\mathrm{MC}[\mathcal{Q}, \varphi]^-_\alpha$, whose value functions coincide with the stages of the fixed-point evaluation g_α for all α, and thus

$$\mathrm{valMC}[\mathcal{Q}, \varphi] = \mathrm{valMC}[\mathcal{Q}, \varphi]^-_\gamma = g_\gamma = [\![\varphi]\!]^{\mathcal{Q}}.$$

□

5 Describing Game Values in Qμ

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 32. *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 $\mathcal{Q}_\mathcal{G}$.*

Proof. Note that $P_i(v) = \infty$ when $\Omega(v) = i$ and $P_i(v) = 0$ otherwise. The formula Win_d is therefore analogous to the one in the qualitative case and the proof is similar, as well. We show that the model checking game for Win_d, $\text{MC}[\mathcal{Q}_\mathcal{G}, \text{Win}_d]$, coincides with $\mathcal{G}$ modulo stupid moves (moves that would lead to an immediate loss for the current player).

We consider a position $v \in V_i$ in original game with priority $\Omega(v) = k$ and distinguish two cases.

If v is a terminal position then the corresponding quantitative transition system has also a terminal position v, where all predicates $P_0, \dots P_{d-1}$ give a value of 0. In the game $\text{MC}[\mathcal{Q}_\mathcal{G}, \text{Win}_d]$ the play goes to a position (Λ, v), which gives a value of $\lambda_G(v)$ as in the original game.

If v is non-terminal of priority k then the corresponding quantitative transition system $\mathcal{Q}_\mathcal{G}$ has a position v, where $V_i(v) = \infty$ and $P_k(v) = \infty$, and the discount on the outgoing edges are $\delta_G(v, w)$ if $i = 0$, or $\frac{1}{\delta_G(v,w)}$ if $i = 1$. The game $\text{MC}[\mathcal{Q}_\mathcal{G}, \text{Win}_d]$ then gets to a position $(\Diamond X_k, v)$ if $i = 0$ or $(\Box X_k, v)$ if $i = 1$, except for the case that one of the players makes an immediately losing move. Then, player i makes a move that exactly corresponds to a move in the original game from v to some successor w, visits a position (X_k, w) with priority k and the situation repeats for w as depicted in Figure 2.

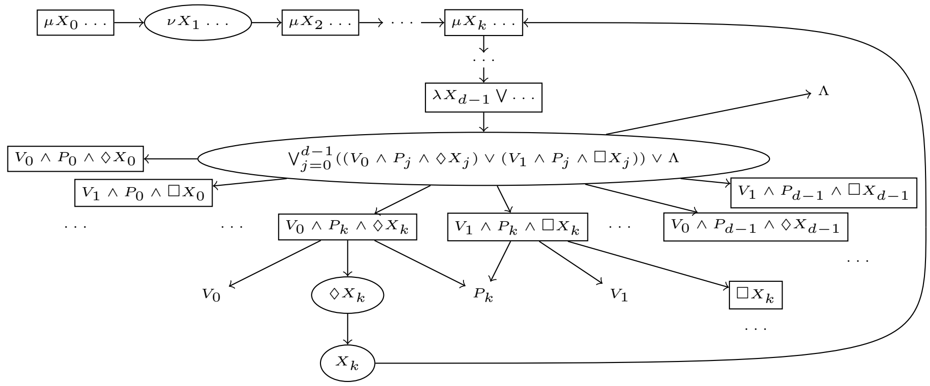

Figure 2: Configuration in the Model Checking Game for Win_d.

Let us convince you that indeed the game will proceed to position (Λ, v), $(\Diamond X_k, v)$ or $(\Box X_k, v)$, or else some player has made a stupid move. We only consider the case for $v \in V_0$ in the original game. In the positions corresponding to the series of fixed points preceding the subformula $\theta = \bigvee_{j=0}^{d-1}((V_0 \wedge P_j \wedge \Diamond X_j) \vee (V_1 \wedge P_j \wedge \Box X_j)) \vee \Lambda$ there is no choice for either of the players, so the play proceeds to the position (θ, v) and it is Player 0's turn to choose a successor. If v is terminal position, the only reasonable choice is to go to the position (Λ, v), as for all other positions Player 1 can make a move to a terminal position (P_i, v) for $i = 0, \dots d-1$, which will give a payoff of 0. If v is non-terminal and has priority k in the original game, i.e. $P_j(v) = 0$ for $j \neq k$, and $\Lambda(v) = 0$, therefore any other move than to a subformula containing the right priority predicate P_k would be pointless. Of course, Player 0 will go to the position corresponding to the disjoint $(V_0 \wedge P_k \wedge \Diamond X_k)$ as we assumed that the position belonged to her in the original game, and therefore $V_0(v) = \infty$. In this position, Player 1 may choose the next position, but again any other move than the one to $(\Diamond X_k, v)$ would be stupid, since all the other predicates will give a value of ∞ by choice of Player 0 in the last move. From position $(\Diamond X_k, v)$, Player 0 can choose a successor w, i.e. lead the game to position (X_k, w) with priority k. The discount $\delta((\Diamond X_k, v), (X_k, w))$ is the same as in the original game by construction of $\mathcal{Q}_\mathcal{G}$ and $\text{MC}[\mathcal{Q}_\mathcal{G}, \text{Win}_d]$.

The other case, $v \in V_1$, is analogous, the only difference is that now the play will go through a position $(\Box X_k, v)$, where Player 1 can choose a successor (X_k, w).

Hence, the two games $\mathcal{G}$ and $\text{MC}[\mathcal{Q}_\mathcal{G}, \text{Win}_d]$ coincide, i.e. all the relevant choices to be made by the players and priorities seen throughout the plays are essentially the same. Therefore, for all $v \in V$,

$$\text{val}\mathcal{G}(v) = \text{valMC}[\mathcal{Q}_\mathcal{G}, \text{Win}_d](\text{Win}_d, v) = [\![\text{Win}_d]\!]^{\mathcal{Q}_\mathcal{G}}(v).$$

□

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 Qμ 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 *Proceedings of the 5th International Conference on Computer Aided Verification, CAV '93*, volume 697, pages 385–396. Springer-Verlag, 1993.

[7] Thomas Gawlitza and Helmut Seidl. Computing game values for crash games. In Kedar S. Namjoshi, Tomohiro Yoneda, Teruo Higashino, and Yoshio Okamura, editors, *ATVA*, volume 4762 of *Lecture Notes in Computer Science*, pages 177–191. Springer, 2007.

[8] Hugo Gimbert and Wieslaw Zielonka. Perfect information stochastic priority games. In Lars Arge, Christian Cachin, Tomasz Jurdzinski, and Andrzej Tarlecki, editors, *ICALP*, volume 4596 of *Lecture Notes in Computer Science*, pages 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] Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102:363–371, 1975.

[12] Annabelle McIver and Carroll Morgan. Results on the quantitative μ-calculus q$M\mu$. *ACM Trans. Comput. Log.*, 8(1), 2007.

[13] 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.