

Quantitative Reductions and Vertex-Ranked Infinite Games*

Alexander Weinert

Reactive Systems Group, Saarland University, 66123 Saarbrücken, Germany

`weinert@react.uni-saarland.de`

We introduce quantitative reductions, a novel technique for structuring the space of quantitative games and solving them that does not rely on a reduction to qualitative games. We show that such reductions exhibit the same desirable properties as their qualitative counterparts and additionally retain the optimality of solutions. Moreover, we introduce vertex-ranked games as a general-purpose target for quantitative reductions and show how to solve them. In such games, the value of a play is determined only by a qualitative winning condition and a ranking of the vertices.

We provide quantitative reductions of quantitative request-response games to vertex-ranked games, thus showing EXPTIME-completeness of solving the former games. Furthermore, we exhibit the usefulness and flexibility of vertex-ranked games by showing how to use such games to compute fault-resilient strategies for safety specifications. This work lays the foundation for a general study of fault-resilient strategies for more complex winning conditions.

1 Introduction

The study of quantitative infinite games has garnered great interest lately, as they allow for a much more fine-grained analysis and specification of reactive systems than classical qualitative games [5, 6, 7, 15, 17, 26, 27]. While there exists previous work investigating quantitative games, the approaches to solving them usually rely on ad-hoc solutions that are tailor-made to the considered winning condition. Moreover, quantitative games are usually solved by reducing them to a qualitative game in a first step, hardcoding a certain value of interest during the reduction. In particular, to the best of our knowledge, no general framework for the analysis of such games that is analogous to the existing one for qualitative games has been developed. In this work, we introduce such a framework that disentangles the study of quantitative games from that of qualitative ones.

Qualitative infinite games have been applied successfully in the verification and synthesis of reactive systems [1, 2, 3, 8]. They have given rise to a multitude of algorithms that ascertain system correctness and that synthesize correct-by-construction systems. In such a game, two players, called Player 0 and Player 1, move a token in a directed graph. After infinitely many moves, the resulting sequence of vertices is evaluated and one player is declared the winner of the play. For example, in a qualitative request-response game [24], the goal for Player 0 is to ensure that every visit to a vertex denoting some request is eventually followed by a visit to a vertex denoting an answer to that request. In order to solve qualitative games, i.e., to determine a winning strategy for one player, one often reduces a complex game to a potentially larger, but conceptually simpler one. For example, in a multi-dimensional request-response game, i.e., in a request-response game in which there exist multiple conditions that can be requested and answered, one stores the set of open requests and demands that every request is closed at infinitely many positions. As this is a Büchi condition, which is much simpler than the request-response condition, one is able to reduce request-response games to Büchi games.

*Supported by the project “TriCS” (ZI 1516/1-1) of the German Research Foundation (DFG) and the Saarbrücken Graduate School of Computer Science

In recent years the focus of research has shifted from the study of qualitative games, in which one player is declared the winner of a given play, to that of quantitative games, in which the resulting play is assigned some value or cost. Such games allow modeling systems in which, for example, requests have to be answered within a certain number of steps [7, 11, 16, 17, 18], systems with one or more finite resources which may be drained and charged [4, 9, 10, 23], or scenarios in which each move incurs a certain cost for either player [15, 27].

In general, Player 0 aims to minimize the cost of the resulting play, i.e., to maximize its value, while Player 1 seeks to maximize the cost, thus minimizing the value. In a quantitative request-response game, for example, it is the goal of Player 0 to minimize the number of steps between requests and their corresponding answers. The typical questions asked in the context of such games are “Does there exist an upper bound on the time between requests and responses that Player 0 can ensure?” [11, 17, 18, 22], “Can Player 0 ensure an average cost per step greater than zero?” [27], “What is the minimal time between requests and responses that Player 0 can ensure?” [26], or “What is the minimal average level of the resource that Player 0 can ensure without it ever running out?” [4]. The former two questions can be seen as boundedness questions, while the latter two are asking for optimal solutions.

Such decision problems are usually solved by fixing some bound b on the cost of the resulting plays and subsequently reducing the quantitative game to a qualitative one, hardcoding the fixed b in the process. If the cost of the resulting play is below b , then Player 0 is declared the winner. For example, in order to determine the winner in a quantitative request-response game as described above for some bound b , we construct a Büchi game in which every time a request is opened, a counter for that request is started which counts up to the bound b and is reset if the request is answered. Once any counter exceeds the value b , we move to a terminal position indicating that Player 0 has lost. We then require that every counter is inactive infinitely often, which is again a Büchi condition and thus much simpler than the original quantitative request-response condition. Player 0 wins the resulting game if and only if she can ensure that every request is answered within at most b steps.

Such reductions are usually specific to the problem being addressed. Furthermore, they immediately abandon the quantitative aspect of the game under consideration, as the bound is hardcoded during the first step of the analysis. Thus, even when only changing the bound one is interested in, the reduction has to be recomputed and the resulting qualitative game has to be solved from scratch. In our request-response example, if one is interested in checking whether Player 0 can ensure every request to be answered within at most $b' \neq b$ steps, one constructs a new Büchi game for the bound b' . This game is then solved independently of the one previously computed for the bound b .

In this work, we lift the concept of reductions for qualitative games to quantitative games. Such quantitative reductions enable the study of a multitude of optimization problems for quantitative games in a way similar to decision problems for qualitative games. When investigating quantitative request-response games using quantitative reductions, for example, we only compute a single, simpler quantitative game and subsequently check this game for a winning strategy for Player 0 for any bound b . If she has such a strategy in the latter game, the quantitative reduction yields a strategy for her satisfying the same bound in the former one.

In general, we retain the intuitive property of reductions for qualitative games: The properties of a complex quantitative game can be studied by investigating a potentially larger, but conceptually simpler quantitative game.

Contributions We present the first framework for reductions between quantitative games and we provide vertex-ranked games as general-purpose targets for such reductions. Moreover, we show tight

bounds on the complexity of solving vertex-ranked games with respect to a given bound.

Finally, we provide two examples illustrating the use of the concepts introduced in this work: First, we define quantitative request-response games and solve them using quantitative reductions to vertex-ranked games. Second, we illustrate the versatility of vertex-ranked games by using them to compute fault-resilient strategies for safety games with faults.

After introducing qualitative and quantitative games formally in Section 2, we define quantitative reductions in Section 3 and show that they provide a mechanism to solve a quantitative game with respect to a given bound: If a game \mathcal{G} can be reduced to a game \mathcal{G}' , then we can use a strategy for Player 0 that minimizes the cost of plays in \mathcal{G}' to construct a strategy for her which minimizes the cost of plays in \mathcal{G} .

In Section 4, we define vertex-ranked games, very general classes of quantitative games that can be used as targets for quantitative reductions. The quantitative condition of such games is quite simple, as the cost of a play is determined only by a qualitative winning condition and a ranking of the vertices of the game. If the resulting play is winning according to the qualitative condition, then its cost is given by the highest rank visited at all or visited infinitely often, depending on the particular variant of vertex-ranked games. Otherwise, the value of the play is infinite. We show that solving such vertex-ranked games is as hard as solving the underlying qualitative winning condition.

Finally, in Section 5 we provide two examples of the versatility of vertex-ranked games: First, we define and solve request-response games with costs via quantitative reductions. Second, we discuss how to use vertex-ranked games to compute fault-resilient strategies in safety games with faults. In such games, after Player 0 has picked a move, say to vertex v , a fault may occur, which overrides the choice of Player 0 and the game continues in vertex $v' \neq v$ instead. By using vertex-ranked games, we are able to compute strategies that are resilient against as many faults as possible.

Proofs omitted due to space restrictions as well as an additional example can be found in the full version [25].

2 Preliminaries

We first define notions that are common to both qualitative and quantitative games. Afterwards, we recapitulate the standard notions for qualitative games before defining quantitative games and lifting the notions for qualitative games to the quantitative case.

We denote the non-negative integers by \mathbb{N} and define $[n] = \{0, 1, \dots, n-1\}$ for every $n \geq 1$. Also, we define $\infty > n$ for all $n \in \mathbb{N}$ and $\mathbb{N}_\infty = \mathbb{N} \cup \{\infty\}$. An *arena* $\mathcal{A} = (V, V_0, V_1, E, v_I)$ consists of a finite, directed graph (V, E) , a partition (V_0, V_1) of V into the positions of Player 0 and Player 1, and an initial vertex $v_I \in V$. The size of \mathcal{A} , denoted by $|\mathcal{A}|$, is defined as $|V|$. A *play* in \mathcal{A} is an infinite path $\rho = v_0 v_1 v_2 \dots$ through (V, E) starting in v_I . To rule out finite plays, we require every vertex to be non-terminal.

A *strategy* for Player i is a mapping $\sigma: V^* V_i \rightarrow V$ that satisfies $(v, \sigma(\pi v)) \in E$ for all $\pi \in V^*$, $v \in V_i$. We say that σ is *positional* if $\sigma(\pi v) = \sigma(v)$ for every $\pi \in V^*$, $v \in V_i$. A play $v_0 v_1 v_2 \dots$ is *consistent* with a strategy σ for Player i , if $v_{j+1} = \sigma(v_0 \dots v_j)$ for all j with $v_j \in V_i$.

A *memory structure* $\mathcal{M} = (M, m_I, \text{Upd})$ for an arena (V, V_0, V_1, E, v_I) consists of a finite set M of memory states, an initial memory state $m_I \in M$, and an update function $\text{Upd}: M \times V \rightarrow M$. We extend the update function to finite play prefixes in the usual way: $\text{Upd}^+(v_I) = m_I$ and $\text{Upd}^+(\pi v) = \text{Upd}(\text{Upd}^+(\pi), v)$ for play prefixes $\pi \in V^+$ and $v \in V$. A next-move function $\text{Nxt}: V_i \times M \rightarrow V$ for Player i has to satisfy $(v, \text{Nxt}(v, m)) \in E$ for all $v \in V_i$, $m \in M$. It induces a strategy σ for Player i with memory \mathcal{M} via $\sigma(v_0 \dots v_j) = \text{Nxt}(v_j, \text{Upd}^+(v_0 \dots v_j))$. A strategy is called *finite-state* if it can be implemented by a memory structure. We define $|\mathcal{M}| = |M|$. In a slight abuse of notation, the size $|\sigma|$ of a finite-state

strategy is the size of a memory structure implementing it.

An arena $\mathcal{A} = (V, V_0, V_1, E, v_I)$ and a memory structure $\mathcal{M} = (M, m_I, \text{Upd})$ for \mathcal{A} induce the expanded arena $\mathcal{A} \times \mathcal{M} = (V \times M, V_0 \times M, V_1 \times M, E', (v_I, m_I))$, where E' is defined via $((v, m), (v', m')) \in E'$ if and only if $(v, v') \in E$ and $\text{Upd}(m, v') = m'$. Every play $\rho = v_0 v_1 v_2 \dots$ in \mathcal{A} has a unique extended play $\text{ext}_{\mathcal{M}}(\rho) = (v_0, m_0)(v_1, m_1)(v_2, m_2) \dots$ in $\mathcal{A} \times \mathcal{M}$ defined by $m_0 = m_I$ and $m_{j+1} = \text{Upd}(m_j, v_{j+1})$, i.e., $m_j = \text{Upd}^+(v_0 \dots v_j)$. We omit the index \mathcal{M} if it is clear from the context. The extended play of a finite play prefix in \mathcal{A} is defined analogously.

Let \mathcal{A} be an arena and let $\mathcal{M}_1 = (M_1, m_I^1, \text{Upd}_1)$ and $\mathcal{M}_2 = (M_2, m_I^2, \text{Upd}_2)$ be memory structures for \mathcal{A} and for $\mathcal{A} \times \mathcal{M}_1$, respectively. We define $\mathcal{M}_1 \times \mathcal{M}_2 = (M_1 \times M_2, (m_I^1, m_I^2), \text{Upd})$, where $\text{Upd}((m_1, m_2), v) = (m'_1, m'_2)$ if $\text{Upd}_1(m_1, v) = m'_1$ and $\text{Upd}_2(m_2, (v, m'_1)) = m'_2$. Via a straightforward induction and in a slight abuse of notation we obtain $\text{Upd}_{\mathcal{M}_2}^+(\text{Upd}_{\mathcal{M}_1}^+(\pi)) = \text{Upd}_{\mathcal{M}_1 \times \mathcal{M}_2}^+(\pi)$ for all finite plays π , where we identify $(v, m_1, m_2) = ((v, m_1), m_2) = (v, (m_1, m_2))$.

2.1 Qualitative Games

A *qualitative game* $\mathcal{G} = (\mathcal{A}, \text{Win})$ consists of an arena \mathcal{A} with vertex set V and a set $\text{Win} \subseteq V^\omega$ of winning plays for Player 0. The set of winning plays for Player 1 is $V^\omega \setminus \text{Win}$. As our definition of games is very general, the infinite object Win may not be finitely describable. If it is, however, in a slight abuse of notation, we define $|\mathcal{G}|$ as the sum of $|\mathcal{A}|$ and the size of a description of Win .

A strategy σ for Player i is a *winning strategy* for i in $\mathcal{G} = (\mathcal{A}, \text{Win})$ if all plays consistent with σ are winning for i . If Player i has a winning strategy, then we say she wins \mathcal{G} . *Solving* a game amounts to determining its winner, if one exists. A game is determined if one player has a winning strategy.

2.2 Quantitative Games

Quantitative games extend the classical model of qualitative games. In a quantitative game, plays are not partitioned into winning and losing plays, but rather they are assigned some measure of quality. We keep this definition very general in order to encompass many of the already existing models. In Section 4, we define concrete examples of such games and show how to solve them optimally.

A quantitative game $\mathcal{G} = (\mathcal{A}, \text{Cost})$ consists of an arena \mathcal{A} with vertex set V and a cost-function $\text{Cost}: V^\omega \rightarrow \mathbb{N}_\infty$ for plays. Similarly to the winning condition in the qualitative case, Cost is an infinite object. If it is finitely describable, we, again slightly abusively, define the size $|\mathcal{G}|$ of \mathcal{G} as the sum of $|\mathcal{A}|$ and the size of a description of Cost . A play ρ in \mathcal{A} is winning for Player 0 in \mathcal{G} if $\text{Cost}(\rho) < \infty$.¹ Winning strategies, the winner of a game, and solving a game are defined as in the qualitative case.

We extend the cost-function over plays to strategies by defining $\text{Cost}(\sigma) = \sup_\rho \text{Cost}(\rho)$ and $\text{Cost}(\tau) = \inf_\rho \text{Cost}(\rho)$, where ρ ranges over the plays consistent with the strategy σ for Player 0 and over the plays consistent with the strategy τ for Player 1, respectively. Moreover, we say that a strategy σ for Player 0 (Player 1) is optimal if its cost is minimal (maximal) among all strategies for her (him).

For every strategy σ for Player 0, $\text{Cost}(\sigma) < \infty$ implies that σ is winning for Player 0. However, the converse does not hold true: Each play consistent with some strategy σ may have finite cost, while for every $n \in \mathbb{N}$ there exists a play ρ consistent with σ with $\text{Cost}(\rho) \geq n$. Dually, a strategy τ for Player 1 has $\text{Cost}(\tau) = \infty$, if and only if τ is winning for him.

¹In order to simplify the presentation, we only consider the case in which Player 0 aims to minimize the cost of a play. All concepts in this work can, however, be easily adapted to the dual case in which Player 0 aims to maximize the cost of a play.

Player 0 wins \mathcal{G} with respect to b if she has a strategy σ with $\text{Cost}(\sigma) \leq b$. Dually, if Player 1 has a strategy τ with $\text{Cost}(\tau) > b$, then he wins \mathcal{G} with respect to b . Solving a quantitative game \mathcal{G} with respect to b amounts to deciding whether or not Player 0 wins \mathcal{G} with respect to b .

If Player 0 has a strategy σ with $\text{Cost}(\sigma) \leq b$, then for all strategies τ for Player 1 we have $\text{Cost}(\tau) \leq b$. Dually, if Player 1 has a strategy τ with $\text{Cost}(\tau) > b$, then for all strategies σ for Player 0 we have $\text{Cost}(\sigma) > b$. We say that a quantitative game is determined if for each $b \in \mathbb{N}$, either Player 0 has a strategy with cost at most b , or Player 1 has a strategy with cost strictly greater than b .

We say that $b \in \mathbb{N}$ is a cap of a quantitative game \mathcal{G} if Player 0 winning \mathcal{G} implies that she has a strategy with cost at most b . A cap b for a game \mathcal{G} is tight if it is minimal.

3 Quantitative Reductions

Before defining quantitative reductions, we first recall the definition of qualitative ones. Let $\mathcal{G} = (\mathcal{A}, \text{Win})$ and $\mathcal{G}' = (\mathcal{A}', \text{Win}')$ be qualitative games. We say that \mathcal{G} is reducible to \mathcal{G}' via the memory structure \mathcal{M} for \mathcal{A} if $\mathcal{A}' = \mathcal{A} \times \mathcal{M}$ and if $\rho \in \text{Win}$ if and only if $\text{ext}(\rho) \in \text{Win}'$. Then, Player 0 wins \mathcal{G} if and only if she wins \mathcal{G}' . Moreover, if σ' is a winning strategy for Player 0 in \mathcal{G}' that is implemented by \mathcal{M}' , then a winning strategy for her in \mathcal{G} is implemented by $\mathcal{M} \times \mathcal{M}'$.

We now define quantitative reductions as an analogous technique for the study of quantitative games and show that they exhibit the same properties as qualitative reductions. Intuitively, given two quantitative games \mathcal{G} and \mathcal{G}' , we aim to say that \mathcal{G} is reducible to \mathcal{G}' if plays in one game can be translated into plays in the other, retaining their cost along the transformation. In fact, two such associated plays do not need to carry identical cost, but it suffices that the order on plays induced by their cost is retained.

To capture this intuition, we introduce b -correction functions. Let $b \in \mathbb{N}_\infty$. A function $f: \mathbb{N}_\infty \rightarrow \mathbb{N}_\infty$ is a b -correction function if

- for all $b'_1 < b'_2 < b$ we have $f(b'_1) < f(b'_2)$,
- for all $b' < b$ we have $f(b') < f(b)$, and
- for all $b' \geq b$ we have $f(b') \geq f(b)$.

For $b = \infty$ these requirements degenerate to demanding that f is strictly monotonic, which in turn implies $f(\infty) = \infty$ and $f(b) \neq \infty$ for all $b \neq \infty$. Dually, if $b = 0$, we only require that $f(0)$ bounds the values of $f(b)$ from below. As an example, the function cap_b , which is defined as $\text{cap}_b(b') = \min\{b, b'\}$ for $b' \neq \infty$ and $\text{cap}_b(\infty) = \infty$ is a b -correction function for all parameters $b \in \mathbb{N}_\infty$.

Let $\mathcal{G} = (\mathcal{A}, \text{Cost})$ and $\mathcal{G}' = (\mathcal{A}', \text{Cost}')$ be quantitative games, let \mathcal{M} be some memory structure for \mathcal{A} , let $b \in \mathbb{N}_\infty$, and let $f: \mathbb{N}_\infty \rightarrow \mathbb{N}_\infty$ be some function. We say that \mathcal{G} is b -reducible to \mathcal{G}' via \mathcal{M} and f if

- $\mathcal{A}' = \mathcal{A} \times \mathcal{M}$,
- f is a b -correction function,
- $\text{Cost}'(\text{ext}(\rho)) = f(\text{Cost}(\rho))$ for all plays ρ of \mathcal{A} with $\text{Cost}(\rho) < b$, and
- $\text{Cost}'(\text{ext}(\rho)) \geq f(b)$ for all plays ρ of \mathcal{A} with $\text{Cost}(\rho) \geq b$.

We write $\mathcal{G} \leq_{\mathcal{M}, f}^b \mathcal{G}'$ in this case. Moreover, we use cap_b as a “default” function: If $f = \text{cap}_b$, we omit stating f explicitly and write $\mathcal{G} \leq_{\mathcal{M}}^b \mathcal{G}'$. The penultimate condition implies that for each play $\text{ext}(\rho)$ in \mathcal{A}' with $\text{Cost}'(\text{ext}(\rho)) \leq f(b)$ there exists some b' such that $\text{Cost}'(\text{ext}(\rho)) = f(b')$.

Quantitative reductions are downwards-closed with respect to the parameter b : If $\mathcal{G} \leq_{\mathcal{M},f}^b \mathcal{G}'$ for some $b \in \mathbb{N}_\infty$ and $b' \leq b$, then $\mathcal{G} \leq_{\mathcal{M},f}^{b'} \mathcal{G}'$. Moreover, similarly to the case of qualitative reductions, quantitative reductions are transitive.

Theorem 1. *Let $\mathcal{G}_1, \mathcal{G}_2, \mathcal{G}_3$ be quantitative games such that $\mathcal{G}_1 \leq_{\mathcal{M}_1, f_1}^{b_1} \mathcal{G}_2 \leq_{\mathcal{M}_2, f_2}^{b_2} \mathcal{G}_3$ for some $b_1, b_2 \in \mathbb{N}_\infty$, some memory structures $\mathcal{M}_1, \mathcal{M}_2$, and some b_1 - and b_2 -correction functions f_1 and f_2 .*

Then, we have $\mathcal{G}_1 \leq_{\mathcal{M}, f}^b \mathcal{G}_3$, where $\mathcal{M} = \mathcal{M}_1 \times \mathcal{M}_2$, $f = f_2 \circ f_1$, and $b = b_1$ if $b_2 \geq f_1(b_1)$ and $b = \max\{b' \mid f_1(b') \leq b_2\}$ otherwise.

Quantitative reductions indeed retain the costs of strategies. To this end, we first demonstrate that correction functions tie the cost of plays in \mathcal{G}' to that of plays in \mathcal{G} .

Lemma 2. *Let \mathcal{G} and \mathcal{G}' be quantitative games such that $\mathcal{G} \leq_{\mathcal{M}, f}^b \mathcal{G}'$, for some $b \in \mathbb{N}_\infty$, some memory structure \mathcal{M} , and some b -correction function f . All of the following hold true for all $b' \in \mathbb{N}$ and all plays ρ in \mathcal{G} :*

1. *If $b' < b$ and $\text{Cost}'(\text{ext}(\rho)) < f(b')$, then $\text{Cost}(\rho) < b'$.*
2. *If $b' < b$ and $\text{Cost}'(\text{ext}(\rho)) = f(b')$, then $\text{Cost}(\rho) = b'$.*
3. *If $\text{Cost}'(\text{ext}(\rho)) \geq f(b)$, then $\text{Cost}(\rho) \geq b$.*

These properties of correction functions when used in quantitative reductions enable us to state and prove the main result of this section, which establishes quantitative reductions as the quantitative counterpart to qualitative reductions: If $\mathcal{G} \leq_{\mathcal{M}, f}^{b+1} \mathcal{G}'$, then all plays of cost at most b in \mathcal{G} are “tracked” precisely in \mathcal{G}' . Hence, as long as the cost of a strategy in \mathcal{G} is at most b , it is possible to construct a strategy in \mathcal{G}' with cost at most $f(b)$. This holds true for both players.

If a strategy has cost greater than b , however, we do not have a direct correspondence between costs of plays in \mathcal{G} and \mathcal{G}' anymore. If, however, b additionally is a cap of \mathcal{G} , and if \mathcal{G} is determined, then we can still show that Player 1 has a strategy of infinite cost in \mathcal{G} if he has a strategy of cost greater than $f(b)$ in \mathcal{G}' .

Theorem 3. *Let \mathcal{G} and \mathcal{G}' be determined quantitative games such that $\mathcal{G} \leq_{\mathcal{M}, f}^{b+1} \mathcal{G}'$ for some b, \mathcal{M} , and f , where $b \in \mathbb{N}$ is a cap of \mathcal{G} .*

1. *Let $b' < b + 1$. Player i has a strategy σ' in \mathcal{G}' with $\text{Cost}'(\sigma') = f(b')$ if and only if they have a strategy σ in \mathcal{G} with $\text{Cost}(\sigma) = b'$.*
2. *If Player 1 has a strategy τ' in \mathcal{G}' with $\text{Cost}'(\tau') \geq f(b + 1)$, then he has a strategy τ in \mathcal{G} with $\text{Cost}(\tau) = \infty$.*

We proved Theorem 3 by constructing optimal strategies for Player 0 in \mathcal{G} from optimal strategies for her in \mathcal{G}' . These strategies use the set of all play prefixes of \mathcal{G}' as memory states and may thus be of infinite size. If Player 0 can achieve a certain cost in \mathcal{G}' using a finite-state strategy, however, then she can achieve the corresponding cost in \mathcal{G} with a finite-state strategy as well.

Theorem 4. *Let \mathcal{G} and \mathcal{G}' be quantitative games such that $\mathcal{G} \leq_{\mathcal{M}_1, f}^b \mathcal{G}'$ for some b, \mathcal{M}_1 , and f and let $b' < b$. If Player i has a finite-state strategy σ' with $\text{Cost}'(\sigma') = f(b')$ in \mathcal{G}' that is implemented by \mathcal{M}_2 , then she has a finite-state strategy σ with $\text{Cost}(\sigma) = b'$ in \mathcal{G} that is implemented by $\mathcal{M}_1 \times \mathcal{M}_2$.*

Using quantitative reductions we are able to structure the space of quantitative games similarly to that of qualitative games. Quantitative games are, however, usually solved by reducing them to qualitative games for a fixed bound b . Thus, there exists no “foundation” of the space of quantitative winning conditions, i.e., there is no canonical simple class of quantitative games that provides a natural target for quantitative reductions. We provide such a foundation in the form of vertex-ranked games in the following section.

4 Vertex-Ranked Games

We introduce a very simple form of quantitative games, which we call vertex-ranked games. In such games, the cost of a play is determined solely by a qualitative winning condition and a ranking of the vertices of the arena by natural numbers. We show that both winning conditions only incur a polynomial overhead in the complexity of solving them with respect to the underlying winning conditions. Furthermore, we show that the memory structures implementing winning strategies for either player only incur a polynomial overhead in comparison to the memory structures implementing winning strategies for the underlying conditions. Moreover, we discuss the optimization problem for such games, i.e., the problem of determining the minimal b such that Player 0 has a strategy of cost at most b in such a game.

Let \mathcal{A} be an arena with vertex set V , let $\text{Win} \subseteq V^\omega$ be a qualitative winning condition, and let $\text{rk}: V \rightarrow \mathbb{N}$ be a ranking function on vertices. We define the quantitative vertex-ranked sup-condition

$$\text{RANK}^{\text{sup}}(\text{Win}, \text{rk}): v_0 v_1 v_2 \cdots \mapsto \begin{cases} \sup_{j \rightarrow \infty} \text{rk}(v_j) & \text{if } v_0 v_1 v_2 \cdots \in \text{Win} \text{ ,} \\ \infty & \text{otherwise ,} \end{cases}$$

as well as its prefix-independent version, the vertex-ranked limsup-condition

$$\text{RANK}^{\text{lim}}(\text{Win}, \text{rk}): v_0 v_1 v_2 \cdots \mapsto \begin{cases} \limsup_{j \rightarrow \infty} \text{rk}(v_j) & \text{if } v_0 v_1 v_2 \cdots \in \text{Win} \text{ ,} \\ \infty & \text{otherwise .} \end{cases}$$

A vertex-ranked sup- or limsup-game $\mathcal{G} = (\mathcal{A}, \text{RANK}^X(\text{Win}, \text{rk}))$ with $X \in \{\text{sup}, \text{lim}\}$ consists of an arena \mathcal{A} with vertex set V , a qualitative winning condition Win , and a vertex-ranking function $\text{rk}: V \rightarrow \mathbb{N}$.

If $\mathcal{G}_X = (\mathcal{A}, \text{RANK}^X(\text{Win}, \text{rk}))$ is a vertex-ranked sup- or limsup-game, we call the game $(\mathcal{A}, \text{Win})$ the qualitative game corresponding to \mathcal{G}_X . Moreover, if \mathcal{G}_{sup} is a vertex-ranked sup-game, we denote the vertex-ranked limsup-game with the same arena, winning condition, and rank function by \mathcal{G}_{lim} and vice versa. In either case, we denote the corresponding qualitative game by \mathcal{G} .

The remainder of this section is dedicated to providing tight bounds on the complexity of solving vertex-ranked games with respect to some given bound. In particular, we show that vertex-ranked sup-games can be solved with only an additive linear blowup compared to the complexity of solving the corresponding qualitative games. Vertex-ranked limsup-games, on the other hand, can be solved while incurring only a polynomial blowup compared to solving the corresponding qualitative games.

4.1 Solving Vertex-Ranked sup-Games

Let us start by noting that solving vertex-ranked sup-games is at least as hard as solving the underlying qualitative games. This is due to the fact that Player 0 has a winning strategy in $(\mathcal{A}, \text{Win})$ if and only if she has a strategy with cost at most zero in $(\mathcal{A}, \text{RANK}^{\text{sup}}(\text{Win}, \text{rk}))$, where rk is the constant function assigning zero to every vertex.

We now turn our attention to finding an upper bound for the complexity of solving vertex-ranked sup-games with respect to some bound. To achieve a general treatment of such games, we first introduce some notation. Let \mathfrak{G} be a class of qualitative games. We define the extension of \mathfrak{G} to vertex-ranked sup-games as

$$\mathfrak{G}_{\text{sup}}^{\text{RNK}} = \{(\mathcal{A}, \text{RANK}^{\text{sup}}(\text{Win}, \text{rk})) \mid (\mathcal{A}, \text{Win}) \in \mathfrak{G}, \text{rk is vertex-ranking function for } \mathcal{A}\} \text{ .}$$

We first show that we can use a decision procedure solving games from \mathfrak{G} to solve games from $\mathfrak{G}_{\text{sup}}^{\text{RNK}}$ with respect to a given b . To this end, we remove all vertices from which Player 1 can enforce a visit to

a vertex of rank greater than b and proclaim that Player 0 wins the quantitative game with respect to b if and only if she wins the qualitative game corresponding to the resulting quantitative game. To ensure that we are able to solve the resulting qualitative game, we assume some closure properties of \mathfrak{G} .

Let $\mathcal{A} = (V, V_0, V_1, E, v_I)$ and $\mathcal{A}' = (V', V'_0, V'_1, E', v'_I)$ be arenas. We say that \mathcal{A}' is a *sub-arena* of \mathcal{A} if $V' \subseteq V$, $V'_0 \subseteq V_0$, $V'_1 \subseteq V_1$, and $E' \subseteq E$ and write $\mathcal{A}' \sqsubseteq \mathcal{A}$ in this case. We call a class of qualitative (or quantitative) games \mathfrak{G} *proper* if for each $(\mathcal{A}, \text{Win})$ (or $(\mathcal{A}, \text{Cost})$) in \mathfrak{G} and each sub-arena $\mathcal{A}' \sqsubseteq \mathcal{A}$ the game $(\mathcal{A}', \text{Win}')$ (or $(\mathcal{A}', \text{Cost}')$), where Win' (or Cost') is the restriction of Win (or Cost) to vertices from \mathcal{A}' , is a member of \mathfrak{G} as well, if all games in \mathfrak{G} are determined and if all $\mathcal{G} \in \mathfrak{G}$ are finitely representable.

Moreover, in order to formalize the idea of removing vertices from which one player can enforce a visit to some set of vertices, we recall the attractor construction. Let $\mathcal{A} = (V, V_0, V_1, E, v_I)$ be an arena with n vertices and let $X \subseteq V$. We define $\text{Attr}_i(X) = \text{Attr}_i^n(X)$ inductively with $\text{Attr}_i^0(X) = X$ and

$$\begin{aligned} \text{Attr}_i^j(X) = \{v \in V_i \mid \exists v' \in \text{Attr}_i^{j-1}(X). (v, v') \in E\} \cup \\ \{v \in V_{1-i} \mid \forall (v, v') \in E. v' \in \text{Attr}_i^{j-1}(X)\} \cup \text{Attr}_i^{j-1}(X) . \end{aligned}$$

Intuitively, the i -attractor $\text{Attr}_i(X)$ is the set of all vertices from which Player i can enforce a visit to X . The set $\text{Attr}_i(X)$ can be computed in linear time in $|E|$ and Player i has a positional strategy σ such that each play starting in some vertex in $\text{Attr}_i(X)$ and consistent with σ eventually encounters some vertex from X [21]. We call σ the attractor strategy towards X .

Player 0 wins $\mathcal{G}_{\text{sup}} \in \mathfrak{G}_{\text{sup}}^{\text{RNK}}$ with respect to some bound b if and only if the initial vertex of \mathcal{G}_{sup} is not in the 1-attractor A of the vertices of rank greater than b , and if she is able to win the game \mathcal{G}' obtained from \mathcal{G} by removing the 1-attractor of A . Since A is computable in linear time, we obtain a generic decision procedure solving vertex-ranked sup-games in linear time and space.

Theorem 5. *Let \mathfrak{G} be a proper class of qualitative games \mathcal{G} that can be solved in time $t(|\mathcal{G}|)$ and space $s(|\mathcal{G}|)$, where t and s are monotonic functions.*

Then, the following problem can be solved in time $\mathcal{O}(n) + t(|\mathcal{G}|)$ and space $\mathcal{O}(n) + s(|\mathcal{G}|)$: “Given some game $\mathcal{G}_{\text{sup}} \in \mathfrak{G}_{\text{sup}}^{\text{RNK}}$ with n vertices and some bound $b \in \mathbb{N}$, does Player 0 win \mathcal{G}_{sup} with respect to b ?”

This theorem provides an upper bound on the complexity of solving vertex-ranked sup-games. Intuitively, we prove Theorem 5 by showing that, for any vertex-ranked sup-game \mathcal{G}_{sup} , a winning strategy for Player 0 in \mathcal{G} that never moves to the Player 1-attractor towards vertices of rank greater than b has cost at most b . Thus, an upper bound on the size of winning strategies for Player 0 for games from \mathfrak{G} provides an upper bound for strategies of finite cost in $\mathfrak{G}_{\text{sup}}^{\text{RNK}}$ as well. Moreover, if the decision procedure deciding \mathfrak{G} constructs winning strategies for one or both players, we can adapt the decision procedure deciding $\mathfrak{G}_{\text{sup}}^{\text{RNK}}$ to construct strategies of cost at most (greater than) b for Player 0 (Player 1) as well.

Corollary 6. *Let \mathfrak{G} be a proper class of qualitative games and let $\mathcal{G} \in \mathfrak{G}$. If σ is a finite-state winning strategy for Player i in \mathcal{G} , then Player i has a finite-state winning strategy σ_{sup} in \mathcal{G}_{sup} with $|\sigma_{\text{sup}}| \in \mathcal{O}(|\sigma|)$. Furthermore, if σ is effectively constructible, then σ_{sup} is effectively constructible.*

Finally, this procedure enables us to solve the optimization problem for vertex-ranked sup-games from $\mathfrak{G}_{\text{sup}}^{\text{RNK}}$: Recall that if Player 0 wins \mathcal{G}_{sup} with respect to some b , she wins it with respect to all $b' \geq b$ as well. Hence, using a binary search, $\log(n)$ invocations of the decision procedure obtained in the proof of Theorem 5 suffice to determine the minimal b such that Player 0 wins \mathcal{G}_{sup} with respect to b . Hence, it is possible to determine the minimal such b in time $\mathcal{O}(\log(n)(n + t(|\mathcal{G}|)))$ and space $\mathcal{O}(n) + s(|\mathcal{G}|)$.

4.2 Solving Vertex-Ranked limsup-Games

We now turn our attention to solving vertex-ranked limsup-games. Solving these games is again at least as hard as solving their corresponding qualitative games, due to the same reasoning as given above for vertex-ranked sup-games. Thus, we again only provide upper bounds on the complexity of solving such games. To this end, given some class \mathfrak{G} of games, we define the corresponding class of vertex-ranked limsup-games

$$\mathfrak{G}_{\text{lim}}^{\text{RNK}} = \{(\mathcal{A}, \text{RANK}^{\text{lim}}(\text{Win}, \text{rk})) \mid (\mathcal{A}, \text{Win}) \in \mathfrak{G}, \text{rk is vertex-ranking function for } \mathcal{A}\} .$$

We identify two criteria on classes of qualitative games \mathfrak{G} , each of which is sufficient for quantitative games in $\mathfrak{G}_{\text{lim}}^{\text{RNK}}$ to be solvable with respect to some given b . More precisely, we provide decision procedures for $\mathfrak{G}_{\text{lim}}^{\text{RNK}}$ for the case that

1. games from \mathfrak{G} can be solved in conjunction with coBüchi-conditions, and for the case that
2. games from \mathfrak{G} are prefix-independent.

For the first case, fix some class of games \mathfrak{G} and let $\mathcal{G}_{\text{lim}} = (\mathcal{A}, \text{RANK}^{\text{lim}}(\text{Win}, \text{rk})) \in \mathfrak{G}_{\text{lim}}^{\text{RNK}}$ with vertex set V . Moreover, recall that a play in Win has cost at most b in \mathcal{G}_{lim} if it visits vertices of rank greater than b only finitely often. In general, such behavior is formalized by the qualitative co-Büchi condition $\text{COBÜCHI}(F) = \{\rho \in V^\omega \mid \text{inf}(\rho) \cap F = \emptyset\}$, where $\text{inf}(\rho)$ denotes the set of vertices occurring infinitely often in ρ . Clearly, Player 0 has a strategy of cost at most b in \mathcal{G}_{lim} if and only if she wins $(\mathcal{A}, \text{Win} \cap \text{COBÜCHI}(\{v \in V \mid \text{rk}(v) > b\}))$. This observation gives rise to the following remark.

Remark 7. *Let \mathfrak{G} be a class of qualitative games such that the games in $\{(\mathcal{A}, \text{Win} \cap \text{COBÜCHI}(F)) \mid (\mathcal{A}, \text{Win}) \in \mathfrak{G}, F \subseteq V, V \text{ is vertex set of } \mathcal{A}\}$ can be solved in time $t(|\mathcal{G}|, |F|)$ and space $s(|\mathcal{G}|, |F|)$, where t and s are monotonic functions.*

Then, the following problem can be solved in time $t(|\mathcal{G}_{\text{lim}}|, n)$ and space $s(|\mathcal{G}_{\text{lim}}|, n)$: “Given some game $\mathcal{G}_{\text{lim}} \in \mathfrak{G}_{\text{lim}}^{\text{RNK}}$ with n vertices and some bound $b \in \mathbb{N}$, does Player 0 win \mathcal{G}_{lim} with respect to b ?”

In this case, we solve vertex-ranked limsup-games via a decision procedure for solving qualitative games as-is. Such a procedure trivially exists if the winning conditions of games from \mathfrak{G} are closed under intersection with co-Büchi conditions. Thus, we obtain solvability of a wide range of classes of vertex-ranked limsup-games, e.g., co-Büchi-, parity-, Muller-, Streett- and Rabin games.

We now turn our attention to the second case described above: We consider classes $\mathfrak{G}_{\text{lim}}^{\text{RNK}}$ where a play is only determined to be winning or losing in a game from \mathfrak{G} due to some infinite suffix. Formally, we say that a qualitative winning condition $\text{Win} \subseteq V^\omega$ is *prefix-independent* if for all infinite plays $\rho \in V^\omega$ and all play prefixes $\pi \in V^*$, we have $\rho \in \text{Win}$ if and only if $\pi\rho \in \text{Win}$. A qualitative game is prefix-independent if its winning condition is prefix-independent. A class of games is prefix-independent if every game in the class is prefix-independent.

Let \mathfrak{G} be a prefix-independent class of games and let $\mathcal{G}_{\text{lim}} \in \mathfrak{G}_{\text{lim}}^{\text{RNK}}$. Moreover, let $b \in \mathbb{N}$. In order to solve \mathcal{G}_{lim} with respect to b , we adapt the classic algorithm for solving prefix-independent qualitative games (cf., e.g., the work by Chatterjee, Henzinger, and Piterman [13]). Thereby, we repeatedly compute the set of vertices from which Player 0 has a strategy of cost at most b in the corresponding vertex-ranked sup-game \mathcal{G}_{sup} and remove their attractor from the game similarly to the construction of a decision procedure for vertex-ranked sup-games in Theorem 5. We claim that Player 0 has a strategy with cost at most b in \mathcal{G}_{lim} if and only if v_l was not removed during the above construction.

Theorem 8. *Let \mathfrak{G} be a proper prefix-independent class of qualitative games \mathcal{G} that can be solved in time $t(|\mathcal{G}|)$ and space $s(|\mathcal{G}|)$, where t and s are monotonic functions.*

Then, the following problem can be solved in time $\mathcal{O}(n^3 + n^2 \cdot t(|\mathcal{G}_{\text{lim}}|))$ and space $\mathcal{O}(n + s(|\mathcal{G}_{\text{lim}}|))$: “Given some game $\mathcal{G}_{\text{lim}} \in \mathfrak{G}_{\text{lim}}^{\text{RNK}}$ with n vertices and some bound $b \in \mathbb{N}$, does Player 0 win \mathcal{G}_{lim} with respect to b ?”

Intuitively, we prove Theorem 8 by constructing a strategy σ for Player 0 by “stitching together” the attractor-strategies towards her winning regions in the decreasing vertex-ranked sup-games and the winning strategies for her in the respective vertex-ranked sup-games. As each play consistent with that strategy descends down the hierarchy of sup-games thus constructed, we can reuse the memory states of the winning strategies in these games when implementing σ . Thus, a monotonic upper bound on the size of strategies with cost at most b in \mathcal{G}_{sup} is an upper bound on the size of such strategies in \mathcal{G}_{lim} as well.

Corollary 9. *Let \mathfrak{G} be a proper prefix-independent class of qualitative games such that, if Player 0 wins \mathcal{G} , then she has a finite-state winning strategy of size at most $m(|\mathcal{G}|)$, where m is a monotonic function.*

If Player 0 wins $\mathcal{G}_{\text{lim}} \in \mathfrak{G}_{\text{lim}}^{\text{RNK}}$ with finite-state strategies, then she has a finite-state winning strategy σ_{lim} of size at most $m(|\mathcal{G}_{\text{lim}}|)$ in \mathcal{G}_{lim} . Furthermore, if winning strategies for Player i in the games in \mathfrak{G} are effectively constructible, then σ_{lim} is effectively constructible.

Moreover, in order to find the optimal b such that Player 0 wins \mathcal{G}_{lim} with respect to b , we can again employ a binary search. Thus, we can determine the optimal such b in time $\mathcal{O}(\log(n)(n^3 + n^2 \cdot t(|\mathcal{G}|)))$ and space $\mathcal{O}(n + s(|\mathcal{G}|))$.

Having thus defined both quantitative reductions as well as a canonical target for such reductions, we now give an example of how to solve quantitative games using this tools.

5 Applications

In this section, we give examples of how to use quantitative reductions and vertex-ranked games to solve quantitative games. First, in Section 5.1, we formally introduce a quantitative variant of request-response games, which we call request-response games with costs, and show how to solve such games using quantitative reductions and vertex-ranked sup-request-response games. Subsequently, in Section 5.2, we show that vertex-ranked games are useful in their own right, by showing how to use them to synthesize controllers that are resilient against disturbances.

5.1 Reducing Request-Response Games with Costs to Vertex-Ranked Request-Response Games

Recall that a play satisfies the qualitative request-response condition if every request that is opened is eventually answered. We extend this condition to a quantitative one by equipping the edges of the arena with costs and measuring the maximal cost incurred between opening and answering a request.

Fix some arena \mathcal{A} with vertex set V and set E of edges. Formally, the qualitative request-response condition $\text{REQRES}(\Gamma)$ consists of a family of so-called request-response pairs $\Gamma = (Q_c, P_c)_{c \in [d]}$, where $Q_c, P_c \subseteq V$ for all $c \in [d]$. Player 0 wins a play according to this condition if each visit to some vertex from Q_c is answered by some later visit to a vertex from P_c , i.e., we define

$$\text{REQRES}((Q_c, P_c)_{c \in [d]}) = \{v_0 v_1 v_2 \cdots \in V^\omega \mid \forall c \in [d] \forall j \in \mathbb{N}. v_j \in Q_c \text{ implies } \exists j' \geq j. v_{j'} \in P_c\} .$$

We say that a visit to a vertex from Q_c *opens a request for condition c* and that the first visit to a vertex from P_c afterwards *answers the request for that condition*.

Proposition 10 ([24]). *Request-response games with n vertices and d request-response pairs can be solved in time $\mathcal{O}(n^2 d^2 2^d)$.*

Furthermore, let \mathcal{G} be a request-response game with d request response pairs. If Player 0 has a winning strategy in \mathcal{G} , then she has a finite-state winning strategy of size at most $d2^d$.

We extend this winning condition to a quantitative one using families of cost functions $\text{Cost} = (\text{Cost}_c)_{c \in [d]}$, where $\text{Cost}_c: E \rightarrow \mathbb{N}$ for each $c \in [d]$ and lift the cost functions Cost_c to play infixes π in \mathcal{A} by adding up the costs along π . The cost-of-response for a request for condition c at position j is then defined as

$$\text{REQRESCOR}_c(v_0 v_1 v_2 \dots, j) = \begin{cases} \min\{\text{Cost}_c(v_j \dots v_{j'}) \mid j' \geq j \text{ and } v_{j'} \in P_c\} & \text{if } v_j \in Q_c, \\ 0 & \text{otherwise,} \end{cases}$$

with $\min \emptyset = \infty$, which naturally extends to the (total) cost-of-response

$$\text{REQRESCOR}(\rho, j) = \max_{c \in [d]} \text{REQRESCOR}_c(\rho, j) .$$

Finally, we define the request-response condition with costs as

$$\text{COSTREQRES}(\Gamma, \text{Cost})(\rho) = \sup_{j \rightarrow \infty} \text{REQRESCOR}(\rho, j) ,$$

i.e., it measures the maximal cost incurred by any request in ρ .

We call a game $\mathcal{G} = (\mathcal{A}, \text{COSTREQRES}(\Gamma, \text{Cost}))$ a request-response game with costs. We denote the largest cost assigned to any edge by W . As we assume the functions Cost_c to be given in binary encoding, the largest cost W assigned to an edge may be exponential in the description length of \mathcal{G} .

If all Cost_c assign zero to every edge, then the request-response condition with costs coincides with the qualitative request-response condition. In general, however, the request-response condition with costs is a strengthening of the classical request-response condition: If some play ρ has finite cost according to the condition with costs, then it is winning for Player 0 according to the qualitative condition, but not vice versa.

Remark 11. *Let $\mathcal{G} = (\mathcal{A}, \text{COSTREQRES}(\Gamma, \text{Cost}))$ be a request-response game with costs. If a strategy σ for Player 0 in \mathcal{G} has finite cost, then σ is a winning strategy for Player 0 in $(\mathcal{A}, \text{REQRES}(\Gamma))$.*

This remark together with a detour via qualitative request-response games yield a cap for request-response games with costs.

Lemma 12. *Let \mathcal{G} be a request-response game with costs with n vertices, d request-response pairs, and largest cost of an edge W . If Player 0 has a strategy with finite cost in \mathcal{G} , then she also has a strategy with cost at most $d2^d nW$.*

Having obtained a cap for request-response games with costs, we can now turn to the main result of this section: Request-response games with costs are reducible to vertex-ranked sup-request-response games. In order to show this, we use a memory structure that keeps track of the costs incurred by the requests open at each point in the play [26].

Lemma 13. *Let \mathcal{G} be a request-response game with costs with n vertices, d request-response pairs, and highest cost of an edge W . Then $\mathcal{G} \leq_{\mathcal{M}}^{b+1} \mathcal{G}'$ for $b = d2^d nW$, some memory structure \mathcal{M} of size $\mathcal{O}(2nb^d)$, and a vertex-ranked sup-request-response game \mathcal{G}' with d request-response pairs.*

Thus, in order to solve a request-response game with costs with respect to some b , it suffices to solve a vertex-ranked sup-request-response game with respect to b . This, in turn, can be done by reducing the problem to that of solving a request-response game as shown in Theorem 5. Using this reduction together with the framework of quality-preserving reductions, we are able to provide an upper bound on the complexity of solving request-response games with respect to some bound b .

Theorem 14. *The following decision problem is in EXPTIME: “Given some request-response game with costs \mathcal{G} and some bound $b \in \mathbb{N}$, does Player 0 have a strategy σ with $\text{Cost}(\sigma) \leq b$ in \mathcal{G} ?”*

Moreover, solving request-response games is known to be EXPTIME-hard [12]. Thus, solving quantitative request-response games via quantitative reductions is asymptotically optimal.

Also, recall that Player 0 has a strategy with cost b' in some request-response game with costs if and only if she has a strategy with cost b' in the vertex-ranked sup-request-response game \mathcal{G}' constructed in the proof of Lemma 13, which has as many request-response pairs d as \mathcal{G} . Due to Proposition 10, if she has a strategy of cost at most b' in \mathcal{G}' , she has one of the same cost and of size at most $d2^d$ in \mathcal{G}' , as argued in Section 4.1. Hence, due to Theorem 4, we obtain an exponential upper bound on the size of optimal strategies for Player 0.

Corollary 15. *Let \mathcal{G} be a request-response game with costs with n vertices, d request-response pairs, and highest cost of an edge W . If Player 0 has a strategy with finite cost, then she also has a strategy with the same cost of size at most $\mathcal{O}(nb^d d2^d)$, where $b = d2^d nW$.*

Finally, the optimization problem of finding the minimal b' such that Player 0 wins a request-response game \mathcal{G} with respect to b' can be solved in exponential time as well. Recall that if Player 0 wins \mathcal{G} with respect to some b' , then she also wins it with respect to all $b'' \geq b'$. Since we can assume $b' \leq b = d2^d nW$, we can perform a binary search for b' on the interval $\{0, \dots, b\}$. Hence, the optimal b' can be found in time $\mathcal{O}(\log(b)n^3 b^{3d} d2^d)$.

5.2 Fault Resilient Strategies for Safety Games

We now demonstrate the flexibility and versatility of vertex-ranked games in their own right. To this end, we consider the problem of synthesizing a controller for a reactive system that is embedded into some environment. This setting is typically modeled as an infinite game in which Player 0 and Player 1 take the roles of the controller and of the environment, respectively. Here, we consider safety games, i.e., we assume that the specification for the controller is given as a game in which it is the aim of Player 0 to keep the play inside a safe subset of the vertices.

Dallal et al. [14] argue that this setting is not sufficiently expressive to correctly model a real-world scenario, since it assumes that Player 0 can accurately predict the effect of her actions on the state of the system. In a realistic setting, in contrast, faults may occur, i.e., an action chosen by a controller may be executed incorrectly, or it may not be executed at all.

In order to model such faults, Dallal et al introduce arenas with faults $\mathcal{A}_F = (V, V_0, V_1, E, F, v_I)$, which consist of an arena (V, V_0, V_1, E, v_I) and a set of faults $F \subseteq V_0 \times V$. In such an arena, whenever it is Player 0's turn, say at vertex v , a fault $(v, v') \in F$ may occur, resulting in the play continuing in vertex v' instead of that chosen by Player 0. Moreover, Dallal et al. consider safety conditions, i.e., qualitative winning conditions of the form $\text{Safety}(S) = \{v_0 v_1 v_2 \dots \mid \forall i \in \mathbb{N}. v_i \in S\}$ for some $S \subseteq V$. Hence, it is the aim of Player 0 to keep the play inside the “safe” set of vertices S . If the play leaves the set S , it is declared winning for Player 1. The task at hand is to compute a fault-resilient strategy for Player 0 that forces the play to remain inside S and that can “tolerate” as many faults as possible.

Safety games without faults are solved by a simple attractor construction: As soon as the play enters $W_1 = \text{Attr}_1(V \setminus S)$, Player 1 can play consistently with his attractor strategy towards $V \setminus S$ in order to win the play. Thus, it is Player 0's aim to keep the play inside $W_0 = V \setminus W_1$.

Dallal et al. solve the problem of computing fault-resilient strategies for safety games by adapting the classical algorithm for solving safety games to this setting. In doing so, they obtain a value $\text{val}(v)$ for each vertex v that denotes the minimal number of faults that need to occur in order for the play to reach W_1 , if Player 0 plays well. Furthermore, they show that val can be computed in polynomial time in $|V|$. Finally, due to the existence of positional winning strategies for both players in safety games, they obtain $\text{val}(v) \in [n] \cup \{\infty\}$ for all $v \in V$. Then, a fault-resilient strategy for Player 0 is one that maximizes the minimal value $\text{val}(v)$ witnessed during any play. Dallal et al. construct such a strategy on the fly during the computation of $\text{val}(v)$ [14].

This task can, however, easily be reframed as a vertex ranked game in the arena $\mathcal{A} = (V, V_0, V_1, E, v_I)$, which we obtain from \mathcal{A}_F by omitting the faults. In that game, we assign to each vertex the rank $\text{rk}(v) = |V| - \text{val}(v)$ if $\text{val}(v) \in [n]$ and $\text{rk}(v) = 0$ otherwise, i.e., if $\text{val}(v) = \infty$. Then, Player 0 has a strategy with cost at most b in $\mathcal{G} = (\mathcal{A}, \text{RANK}^{\text{sup}}(\text{Safety}(S), \text{rk}))$ if and only if she has a winning strategy in the original safety game with faults that tolerates at least $|V| - b$ faults.

This formulation as a vertex-ranked game enables further study of games in arenas with faults. Here, we require the winning condition to be a safety condition in order to compute $\text{val}(v)$. In recent work, we have shown how to compute this value for more complex qualitative winning conditions [20]. If $\text{val}(v)$ is effectively computable for a given qualitative winning condition, one can easily obtain fault-resilient strategies by formulating the task as a vertex-ranked game as demonstrated.

Finally, the formulation as a vertex-ranked game yields a method to compute eventually-fault-resilient strategies, i.e., strategies that are resilient to a large number of faults after a finite “start-up” phase. In order to obtain such strategies, it suffices to view the resulting vertex-ranked games as a lim sup-game instead of a sup-game and to solve it optimally as described in Section 4.2.

6 Conclusion

In this work, we have lifted the concept of reductions, which has yielded a multitude of results in the area of qualitative games, to quantitative games. We have shown that this novel concept exhibits the same useful properties for quantitative games as it does for qualitative ones and that it furthermore retains the quality of strategies.

Additionally, we have provided two very general types of quantitative games that serve as targets for quantitative reductions, namely vertex-ranked sup games and vertex-ranked lim sup-games. For both kinds of games we have shown a polynomial overhead on the complexity of solving them with respect to some bound, on the memory necessary to achieve a given cost, and on the complexity of determining the optimal cost that either player can ensure.

Finally, we have demonstrated the versatility of these tools by using them to solve quantitative request-response games and by showing how to solve the problem of computing fault-resilient strategies in safety games via vertex-ranked games. This latter formulation enables a general study of games with faults, even in the presence of more complex winning conditions than the safety condition considered by Dallal et al. [14] and in this work. We are currently investigating how to leverage vertex-ranked games for the synthesis of fault-resistant strategies in parity games.

Further research continues in two additional directions: Firstly, while the framework of quantitative reductions and vertex-ranked games yields upper bounds on the complexity of solving quantitative

games, it does not directly yield lower bounds on the complexity of the problems under investigation. Consider, for example, the problem of solving parity games with costs with respect to some bound, which is PSPACE-complete [26]. It is possible to reduce this problem to that of solving a vertex-ranked parity game of exponential size and linearly many colors similarly to the reduction presented in this work, which yields an EXPTIME-algorithm. It remains open how to use quantitative reductions to obtain an algorithm for this problem that only requires polynomial space.

Secondly, another goal for future work is the establishment of an analogue to the Borel hierarchy for quantitative winning conditions. In the qualitative case, this hierarchy establishes clear boundaries for reductions between infinite games, i.e., a game whose winning condition is in one level of the Borel hierarchy cannot be reduced to one with a winning condition in a lower level. Also, each game with a winning condition in the hierarchy is known to be determined [19]. To the best of our knowledge, it is open how to define such a hierarchy for quantitative winning conditions which exhibit similar properties.

Acknowledgements I would like to thank Martin Zimmermann for many fruitful discussions.

References

- [1] Roderick Bloem, Krishnendu Chatterjee, Karin Greimel, Thomas A. Henzinger, Georg Hofferek, Barbara Jobstmann, Bettina Könighofer & Robert Könighofer (2014): *Synthesizing robust systems*. *Acta Inf.* 51(3-4), pp. 193–220, doi:10.1007/s00236-013-0191-5.
- [2] Roderick Bloem, Krishnendu Chatterjee, Thomas A. Henzinger & Barbara Jobstmann (2009): *Better Quality in Synthesis through Quantitative Objectives*. In: *CAV 2009, LNCS 5643*, Springer, pp. 140–156, doi:10.1007/978-3-642-02658-4_14.
- [3] Patricia Bouyer, Nicolas Markey, Jörg Olschewski & Michael Ummels (2011): *Measuring Permissiveness in Parity Games: Mean-Payoff Parity Games Revisited*. In Tevfik Bultan & Pao-Ann Hsiung, editors: *ATVA 2011, LNCS 6996*, Springer, pp. 135–149, doi:10.1007/978-3-642-24372-1_11.
- [4] Patricia Bouyer, Nicolas Markey, Mickael Randour, Kim G. Larsen & Simon Laursen (2016): *Average-energy games*. *Act. Inf.*, doi:10.1007/s00236-016-0274-1.
- [5] Lubos Brim, Jakub Chaloupka, Laurent Doyen, Raffaella Gentilini & Jean-François Raskin (2011): *Faster algorithms for mean-payoff games*. *Form. Meth. in Sys. Des.* 38(2), pp. 97–118, doi:10.1007/s10703-010-0105-x.
- [6] Véronique Bruyère, Emmanuel Filiot, Mickael Randour & Jean-François Raskin (2017): *Meet your expectations with guarantees: Beyond worst-case synthesis in quantitative games*. *Inf. Comput.* 254, pp. 259–295, doi:10.1016/j.ic.2016.10.011.
- [7] Véronique Bruyère, Quentin Hautem & Mickael Randour (2016): *Window parity games: an alternative approach toward parity games with time bounds*. In Domenico Cantone & Giorgio Delzanno, editors: *GandALF 2016, EPTCS 226*, pp. 135–148, doi:10.4204/EPTCS.226.10.
- [8] Pavol Cerný, Krishnendu Chatterjee, Thomas A. Henzinger, Arjun Radhakrishna & Rohit Singh (2011): *Quantitative Synthesis for Concurrent Programs*. In Ganesh Gopalakrishnan & Shaz Qadeer, editors: *CAV 2011, LNCS 6806*, Springer, doi:10.1007/978-3-642-22110-1_20.
- [9] Krishnendu Chatterjee & Laurent Doyen (2012): *Energy parity games*. *Theo. Comp. Sci.* 458, pp. 49–60, doi:10.1016/j.tcs.2012.07.038.
- [10] Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger & Jean-François Raskin (2010): *Generalized Mean-payoff and Energy Games*. In Kamal Lodaya & Meena Mahajan, editors: *FSTTCS, LIPIcs 8*, pp. 505–516, doi:10.4230/LIPIcs.FSTTCS.2010.505.

- [11] Krishnendu Chatterjee, Thomas A. Henzinger & Florian Horn (2009): *Finitary winning in ω -regular games*. *ACM Trans. Comput. Log.* 11(1), doi:10.1145/1614431.1614432.
- [12] Krishnendu Chatterjee, Thomas A. Henzinger & Florian Horn (2011): *The Complexity of Request-Response Games*. In: *LATA 2011, LNCS 6638*, Springer, pp. 227–237, doi:10.1007/978-3-642-21254-3_17.
- [13] Krishnendu Chatterjee, Thomas A. Henzinger & Nir Piterman (2006): *Algorithms for Büchi Games*. In: *GDV*. Available at <http://arxiv.org/abs/0805.2620>.
- [14] Eric Dallal, Daniel Neider & Paulo Tabuada (2016): *Synthesis of safety controllers robust to unmodeled intermittent disturbances*. In: *CDC 2016, IEEE*, pp. 7425–7430, doi:10.1109/CDC.2016.7799416.
- [15] Andrzej Ehrenfeucht & Jan Mycielski (1979): *Positional strategies for mean payoff games*. *Internat. J. Game Theory* 8, pp. 109–113, doi:10.1007/BF01768705.
- [16] Peter Faymonville & Martin Zimmermann (2017): *Parametric Linear Dynamic Logic*. *Inf. Comput.* 253, pp. 237–256, doi:10.1016/j.ic.2016.07.009.
- [17] Nathanaël Fijalkow & Martin Zimmermann (2014): *Parity and Streett Games with Costs*. *LMCS* 10(2), doi:10.2168/LMCS-10(2:14)2014.
- [18] Orna Kupferman, Nir Piterman & Moshe Y. Vardi (2009): *From liveness to promptness*. *Form. Meth. in Sys. Des.* 34(2), pp. 83–103, doi:10.1007/s10703-009-0067-z.
- [19] Donald A. Martin (1975): *Borel determinacy*. *Ann. of Math.* 102, pp. 363–371, doi:10.2307/1971035.
- [20] Daniel Neider, Alexander Weinert & Martin Zimmermann (2018): *Synthesizing Optimally Resilient Controllers*. In Dan Ghika & Achim Jung, editors: *CSL 2018, Schloss Dagstuhl - LZI*. To appear.
- [21] Anil Nerode, Jeffrey B. Remmel & Alexander Yakhnis (1996): *McNaughton Games and Extracting Strategies for Concurrent Programs*. *Ann. Pure Appl. Logic* 78(1-3), pp. 203–242, doi:10.1016/0168-0072(95)00032-1.
- [22] Sven Schewe, Alexander Weinert & Martin Zimmermann (2018): *Parity Games with Weights*. In Dan Ghika & Achim Jung, editors: *CSL 2018, Schloss Dagstuhl - LZI*. To appear.
- [23] Yaron Velner, Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger, Alexander Moshe Rabinovich & Jean-François Raskin (2015): *The complexity of multi-mean-payoff and multi-energy games*. *Inf. Comput.* 241, pp. 177–196, doi:10.1016/j.ic.2015.03.001.
- [24] Nico Wallmeier, Patrick Hütten & Wolfgang Thomas (2003): *Symbolic Synthesis of Finite-State Controllers for Request-Response Specifications*. In: *CIAA 2003, LNCS 2759*, Springer, pp. 11–22, doi:10.1007/3-540-45089-0_3.
- [25] Alexander Weinert (2017): *Quantitative Reductions and Vertex-Ranked Infinite Games*. *arXiv abs/1704.00904*. Available at <http://arxiv.org/abs/1704.00904>.
- [26] Alexander Weinert & Martin Zimmermann (2016): *Easy to Win, Hard to Master: Optimal Strategies in Parity Games with Costs*. In Jean-Marc Talbot & Laurent Regnier, editors: *CSL 2016, LIPIcs 62*, pp. 31:1–31:17, doi:10.4230/LIPIcs.CSL.2016.31.
- [27] Uri Zwick & Mike Paterson (1996): *The complexity of mean payoff games on graphs*. *Theo. Comp. Sci.* 158(1-2), pp. 343–359, doi:10.1016/0304-3975(95)00188-3.