Quantitative Reductions and Vertex-Ranked Infinite Games*

Alexander Weinert

Reactive Systems Group, Saarland University, 66123 Saarbrücken, Germany
weinert@react.uni-saarland.de

Abstract. We introduce quantitative reductions, a novel technique for solving quantitative games that does not rely on a reduction to qualitative games. We demonstrate that such reductions exhibit the same desirable properties as qualitative ones. In addition, they retain the optimality of solutions. As a general-purpose target for quantitative reductions, we introduce vertex-ranked games, in which the value of a play is determined only by a qualitative winning condition and a ranking of the vertices. Moreover, we demonstrate how to solve such games optimally. Finally, we provide quantitative reductions to vertex-ranked games for both quantitative request-response and Muller games. Thereby, we obtain ExpTime-completeness of solving the former games, while obtaining a new proof for the membership of solving the latter games in ExpTime.

1 Introduction

The study of quantitative infinite games has garnered great interest lately. This is due to them allowing for a much more fine-grained analysis and specification of reactive systems than classical qualitative games. While there exists previous work investigating such games, the approaches to these games usually rely on ad-hoc solutions that are tailor-made to the problem under consideration. In particular, no general tools have been developed for the analysis of such games. We introduce 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. 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 nodes is evaluated and one player is declared as the winner of the play. For example, in a qualitative request-response game, 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. To solve qualitative games, i.e., to determine a winning strategy for one player, one often reduces a

* Supported by the project "TriCS" (ZI 1516/1-1) of the German Research Foundation (DFG).

complex game to a potentially larger, but conceptually simpler one. For example, in a multi-dimensional request-response game, i.e., 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.

In recent years the focus of research has shifted from the study of qualitative games, in which one player is declared as the winner, to that of quantitative games, in which the resulting play is assigned some value. Such games allow to model systems in which, for example, requests have to be answered within a certain number of steps [12,9,5,11,2], systems with one or more finite resources which may be drained and charged [3,17,4,1], or scenarios in which each move incurs a certain cost for either player [8,20].

In general, one player aims to minimize the value of the resulting play, while the other one seeks to maximize it. 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 "Can Player 0 ensure an upper bound on the time between requests and responses?" [12,5,11], "What is the minimal time between requests and responses that Player 0 can ensure?" [19], "What is the minimal average level of the resource that Player 0 can ensure without it ever running out?" [1], or "Can Player 0 ensure an average cost per step greater than zero?" [20].

Such decision problems are usually answered by reducing the quantitative game to a qualitative one, where some bound is hardcoded during the reduction. If the value of the resulting play is below the bound, then Player 0 is declared as the winner. For example, in order to determine the winner in a quantitative request-response game as described above for a given bound b, we construct a Büchi game in which every time a request is opened, a counter is started which counts up to the bound b and is reset if the request is answered. Once a counter exceeds the value b, we move to a position indicating that Player 0 has lost. We then require that every counter is not running infinitely often, which is again a Büchi condition, i.e., it is much simpler than the original quantitative request-response condition. The resulting game is won by Player 0 if and only if she can ensure that every request is answered within at most b steps.

Such reductions are usually very specific to the problem being addressed. Furthermore, they abandon the quantitative aspect of the game under consideration immediately, as the bound is hardcoded during the reduction. Thus, even when only changing the bound one is interested in, the reduction has to be recomputed and the resulting 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 would construct a new Büchi game. This game would then be solved independently of the previously computed one 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 similar way 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. We can then easily 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 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.

Our Contributions We present the first framework for reductions between quantitative games and we provide vertex-ranked games general-purpose targets for such reductions. Additionally, we show tight bounds on the complexity of solving such games with respect to a given bound. Finally, we illustrate the usefulness of this framework by using quantitative reductions to provide tight bounds on the complexity of solving quantitative request-response games with respect to some given bound and to provide an upper bound on the complexity of solving quantitative Muller games with respect to some given bound.

After introducing both qualitative and quantitative games formally in Section 2, we define quantitative reductions in Section 3. In Theorem 2 we show that they provide a mechanism to solve a quantitative game with respect to a given bound. More precisely, we show that 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 value of plays in $\mathcal{G}'$ to construct a strategy for her in $\mathcal{G}$ which minimizes the value of plays in that game as well.

In Section 4, we define very general classes of quantitative games, so-called vertex-ranked games, that can be used as targets for quantitative reductions. Such games are very simple quantitative games, in which the cost of a play is given only by a qualitative winning condition and a ranking of the vertices of the game by natural numbers. If the resulting play is winning according to the qualitative condition, then its value is given by the highest rank visited at all or visited infinitely often, depending on the specific variant of vertex-ranked games. Otherwise, the value of the play is infinite. We provide asymptotically tight bounds on the complexity of solving vertex-ranked games in Theorem 4 and Theorem 5.

Finally, we demonstrate the usefulness of quantitative reductions and vertex-ranked games by using them to solve quantitative request-response games with respect to a given bound in Section 5.1 and to solve quantitative Muller games with respect to a given bound in 5.2. No bound on the complexity of solving quantitative request-response games was known so far. We show the problem to be ExpTime-complete. Moreover, both results demonstrate that our framework is amenable to modular and conceptually simpler proofs when compared to previously existing ad-hoc solutions.

2 Preliminaries

We begin by defining 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 $[n] = \{0, 1, \ldots, n-1\}$ 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 $\rho = v_0 v_1 v_2 \cdots$ 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 $\sigma\colon V^* V_i \to V$ where $(v, \sigma(\pi v)) \in E$ for all $\pi \in V^*$, $v \in V_i$. We say that σ is *positional* if $\sigma(\pi v) = \sigma(v)$ for every $\pi \in V^*$, $v \in V_i$. We often view positional strategies as a mapping $\sigma\colon V_i \to V$. A play $v_0 v_1 v_2 \cdots$ is *consistent* with a strategy σ for Player i, if $v_{j+1} = \sigma(v_0 \cdots v_j)$ for all j with $v_j \in V_i$.

A *memory structure* $\mathcal{M} = (M, m_I, \mathrm{Upd})$ 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 $\mathrm{Upd}\colon M \times V \to M$. The update function is extended to finite play prefixes in the usual way: $\mathrm{Upd}^+(v_I) = m_I$ and $\mathrm{Upd}^+(\pi v) = \mathrm{Upd}(\mathrm{Upd}^+(\pi), v)$ for play prefixes $\pi \in V^+$ and $v \in V$. A next-move function $\mathrm{Nxt}\colon V_i \times M \to V$ for Player i has to satisfy $(v, \mathrm{Nxt}(v, m)) \in E$ for all $v \in V_i$, $m \in M$. It induces a strategy σ for Player i with memory $\mathcal{M}$ via $\sigma(v_0 \cdots v_j) = \mathrm{Nxt}(v_j, \mathrm{Upd}^+(v_0 \cdots v_j))$. A strategy is called *finite-state* if it can be implemented by a memory structure. We define $|\mathcal{M}| = |M|$. The size $|\sigma|$ of a finite-state strategy is the size of a smallest memory structure implementing it.

An arena $\mathcal{A} = (V, V_0, V_1, E, v_I)$ and a memory structure $\mathcal{M} = (M, m_I, \mathrm{Upd})$ 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 $\mathrm{Upd}(m, (v, v')) = m'$. Every play $\rho = v_0 v_1 v_2 \cdots$ in $\mathcal{A}$ has a unique extended play $\mathrm{ext}_{\mathcal{M}}(\rho) = (v_0, m_0)(v_1, m_1)(v_2, m_2) \cdots$ in $\mathcal{A} \times \mathcal{M}$ defined by $m_0 = m_I$ and $m_{j+1} = \mathrm{Upd}(m_j, v_{j+1})$, i.e., $m_j = \mathrm{Upd}^+(v_0 \cdots v_j)$. 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, let $\mathcal{M}_1 = (M_1, m_I^1, \mathrm{Upd}_1)$ be a memory structure for $\mathcal{A}$, and let $\mathcal{M}_2 = (M_2, m_I^2, \mathrm{Upd}_2)$ be a memory structure for $\mathcal{A} \times \mathcal{M}_1$. We define $\mathcal{M}_1 \times \mathcal{M}_2 = (M_1 \times M_2, (m_I^1, m_I^2), \mathrm{Upd})$, where $\mathrm{Upd}((m_1, m_2), v) = (m_1', m_2')$ if $\mathrm{Upd}_1(m_1, v) = m_1'$ and $\mathrm{Upd}_2(m_2, (v, m_1')) = m_2'$. Via a straightforward induction and in a slight abuse of notation we obtain $\mathrm{Upd}^+_{\mathcal{M}_2}(\mathrm{Upd}^+_{\mathcal{M}_1}(\pi)) = \mathrm{Upd}^+_{\mathcal{M}_1 \times \mathcal{M}_2}(\pi)$ for all finite and infinite 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, with $V' \supseteq V$.[1] 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, we define $|\mathcal{G}| = |\mathcal{A}| + |\text{Win}|$, with $|\text{Win}|$ as the description size of Win.

A strategy σ for Player i is a *winning strategy* for $\mathcal{G} = (\mathcal{A}, \text{Win})$ if all plays consistent with σ are winning for that player. 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

We define quantitative games as an extension of classical 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 $\mathcal{G} = (\mathcal{A}, \text{Cst})$ consists of an arena $\mathcal{A}$ with vertex set V and a cost-function $\text{Cst}\colon (V')^\omega \to \mathbb{N}_\infty$ for plays, where $V' \supseteq V$. Similarly to Win in the qualitative case, Cst is an infinite object. If it is finitely describable, we define the size $|\mathcal{G}|$ of $\mathcal{G}$ as the sum of $|\mathcal{A}|$ and the description length of Cst. A play in $\mathcal{A}$ is winning for Player 0 in $\mathcal{G}$ if $\text{Cst}(\rho) < \infty$. 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 $\text{Cst}(\sigma) = \sup_\rho \text{Cst}(\rho)$ and $\text{Cst}(\tau) = \inf_\rho \text{Cst}(\rho)$ for strategies σ and τ of Player 0 and Player 1, respectively. The supremum and infimum range over all plays ρ consistent with σ and τ, respectively. Moreover, we say that a strategy σ for Player i is optimal if its cost is minimal (for Player 0) or maximal (for Player 1) among all strategies for that player.

For any strategy σ for Player 0, $\text{Cst}(\sigma) < \infty$ 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 $\text{Cst}(\rho) = n$. Dually, while each winning strategy τ for Player 1 has $\text{Cst}(\tau) = \infty$, the converse does not hold true.

Player 0 wins $\mathcal{G}$ with respect to b if she has a strategy σ with $\text{Cst}(\sigma) \leq b$. Dually, if Player 1 has a strategy τ with $\text{Cst}(\tau) > b$, 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.

Note that, if Player 0 has a strategy σ with $\text{Cst}(\sigma) \leq b$, then for all strategies τ for Player 1 we have $\text{Cst}(\tau) \leq b$. Dually, if Player 1 has a strategy τ

[1] We define the winning condition over a superset of V in order to simplify the removal of parts of the arena later on.

with $\mathrm{Cst}(\tau) \geq b$, then for all strategies σ for Player 0 we have $\mathrm{Cst}(\sigma) \geq b$. A quantitative game is determined if for each $b \in \mathbb{N}$, one player has a strategy with cost at most b (for Player 0), or at least b (for Player 1).

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

In this section, we lift the concept of reductions between qualitative games to quantitative ones. Recall that a qualitative game $\mathcal{G} = (\mathcal{A}, \mathrm{Win})$ is reducible to $\mathcal{G}' = (\mathcal{A}', \mathrm{Win}')$ via the memory structure $\mathcal{M}$ for $\mathcal{A}$ if $\mathcal{A}' = \mathcal{A} \times \mathcal{M}$ and if $\rho \in \mathrm{Win}$ if and only if $\mathrm{ext}(\rho) \in \mathrm{Win}'$. 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 $\mathcal{M} \times \mathcal{M}'$.

Our main goal is to develop a notion of quantitative reductions that does not only preserve winning plays, but also retains the cost of plays, possibly with respect to some scaling. To allow for a very general notion of offset, we introduce b-correction functions. Let $b \in \mathbb{N}_\infty$. A function $f \colon \mathbb{N}_\infty \to \mathbb{N}_\infty$ 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, for $b = 0$ we only demand that $f(0)$ bounds the values of $f(b)$ from below. As an example, the function cap_b, which is defined as $\mathrm{cap}_b(b') = \min\{b, b'\}$ if $b' \neq \infty$ and $\mathrm{cap}_b(\infty) = \infty$ is a b-correction function for all parameters $b \in \mathbb{N}_\infty$.

Let $\mathcal{G} = (\mathcal{A}, \mathrm{Cst})$ and $\mathcal{G}' = (\mathcal{A}', \mathrm{Cst}')$ be quantitative games, let $\mathcal{M}$ be some memory structure for $\mathcal{A}$, let $b \in \mathbb{N}_\infty$, and let $f \colon \mathbb{N}_\infty \to \mathbb{N}_\infty$ be some function. We say that $\mathcal{G}$ is b-reducible to $\mathcal{G}'$ via $\mathcal{M}$ and f if all of the following hold true:

- $\mathcal{A}' = \mathcal{A} \times \mathcal{M}$,
- f is a b-correction function,
- $\mathrm{Cst}'(\mathrm{ext}(\rho)) = f(\mathrm{Cst}(\rho))$ for all $\rho \in \mathrm{Plays}(\mathcal{A})$ with $\mathrm{Cst}(\rho) < b$, and
- $\mathrm{Cst}'(\mathrm{ext}(\rho)) \geq f(b)$ for all $\rho \in \mathrm{Plays}(\mathcal{A})$ with $\mathrm{Cst}(\rho) \geq b$.

We write $\mathcal{G} \leq^b_{\mathcal{M},f} \mathcal{G}'$ in this case, or $\mathcal{G} \leq^b_{\mathcal{M}} \mathcal{G}'$, if $f = \mathrm{cap}_b$. Note that the penultimate condition implies that for each play $\mathrm{ext}(\rho)$ in $\mathcal{A}'$ with $\mathrm{Cst}'(\mathrm{ext}(\rho)) \leq f(b)$ there exists some b' such that $\mathrm{Cst}'(\mathrm{ext}(\rho)) = f(b')$.

Quantitative reductions are monotonic with respect to the parameter b: If $\mathcal{G} \leq^b_{\mathcal{M},f} \mathcal{G}'$ for some $b \in \mathbb{N}_\infty$ and $b' \leq b$, then $\mathcal{G} \leq^{b'}_{\mathcal{M},f} \mathcal{G}'$. 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 where $\mathcal{G}_1 \leq^{b_1}_{\mathcal{M}_1, f_1} \mathcal{G}_2 \leq^{b_2}_{\mathcal{M}_2, f_2} \mathcal{G}_3$ for some $b_1, b_2, \mathcal{M}_1, \mathcal{M}_2, f_1$, and f_2. Then we have $\mathcal{G}_1 \leq^{b}_{\mathcal{M}, f} \mathcal{G}_3$, 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.*

Proof. Let $\mathcal{G}_j = (\mathcal{A}_j, \mathrm{Cst}_j)$ for $j \in \{1, 2, 3\}$. Clearly, we have $\mathcal{A}_3 = \mathcal{A}_2 \times \mathcal{M}_2 = \mathcal{A}_1 \times \mathcal{M}_1 \times \mathcal{M}_2$ and

$$\begin{aligned} f_2 \circ f_1(\mathrm{Cst}_1(\rho)) &= f_2(\mathrm{Cst}_2(\mathrm{ext}_{\mathcal{M}_1}(\rho))) \\ &= \mathrm{Cst}_3(\mathrm{ext}_{\mathcal{M}_2}(\mathrm{ext}_{\mathcal{M}_1}(\rho))) = \mathrm{Cst}_3(\mathrm{ext}_{\mathcal{M}_1 \times \mathcal{M}_2}(\rho)) \end{aligned}$$

for all $\rho \in V^\omega$. It remains to show that $f_2 \circ f_1$ is a b-correction function, for some b as stated in the theorem.

First, assume $b_2 \geq f_1(b_1)$ and pick x and x' such that $x < x' < b_1$. As f_1 is a b_1-correction function, we obtain $f_1(x) < f_1(x')$, $f_1(x) < f_1(b_1)$, and $f_1(x') < f_1(b_1)$. Since f_2 is a b_2-correction function and as $f_1(b_1) \leq b_2$ by assumption, we furthermore obtain $f_2(f_1(b)) < f_2(f_1(x'))$. Now pick some x such that $x < b_1$. Then $f_1(x) < f_1(b_1)$ and $f_2(f_1(b_1)) \leq f_2(b_2)$. Thus, $f_2(f_1(b)) < f_2(f_1(b_1))$. Finally, pick some x such that $x \geq b_1$. Then $f_1(x) \geq f_1(b_1)$. If $f_1(x) < b_2$, then $f_2(f_1(x)) \geq f_2(f_1(b_1))$. If, however, $f_1(x) \geq b_2$, then $f_2(f_1(x)) \geq f_2(b_2) \geq f_2(f_1(b_1))$, which concludes this part of the proof.

Now assume $b_2 < f_1(b_1)$ and let b' be maximal such that $f_1(b') < b_2$. We show that $f_2 \circ f_1$ is a b'-correction function. First, pick x and x' such that $x < x' < b'$. Then $f_1(x) < f_1(x') < b_2$, i.e., $f_2(f_1(x)) < f_2(f_1(x'))$. Now, pick x such that $x < b'$. Then we have $f_1(x) < f_1(b') < b_2$, which implies $f_2(f_1(x)) < f_2(f_1(b'))$. Finally, pick x such that $x \geq b'$. If $f_1(x) < b_2$, then $f_1(x) \geq f_1(b')$ and $f_2(f_1(x)) \geq f_2(f_1(b'))$. If, however, $f_1(x) \geq b_2$, then $f_1(x) > f_1(b')$ and, as $f_1(b') < b_2$, we moreover obtain $f_2(f_1(x)) \geq f_2(f_1(b'))$. □

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

Lemma 1. *Let $\mathcal{G}$ and $\mathcal{G}'$ be quantitative games such that $\mathcal{G} \leq^{b}_{\mathcal{M}, f} \mathcal{G}'$, for some b, $\mathcal{M}$, and f. All of the following hold true for all $b' \in \mathbb{N}$ and all plays ρ in $\mathcal{A}$:*

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

Proof. 1) Let $b' < b$ and let ρ such that $\mathrm{Cst}'(\mathrm{ext}(\rho)) < f(b')$. Towards a contradiction assume $\mathrm{Cst}(\rho) = b'' \geq b'$. We have $f(b'') = f(\mathrm{Cst}(\rho)) = \mathrm{Cst}'(\mathrm{ext}(\rho))$. If $b'' < b$, then we obtain $f(b') \leq f(b'')$, which implies $f(b') \leq \mathrm{Cst}'(\mathrm{ext}(\rho))$, contradicting $\mathrm{Cst}'(\mathrm{ext}(\rho)) < f(b')$. If, however, $b'' \geq b$, then $\mathrm{Cst}'(\mathrm{ext}(\rho)) = f(b'') \geq f(b) > f(b')$, again contradicting $\mathrm{Cst}'(\mathrm{ext}(\rho)) < f(b')$.

2) Let $b' < b$ and let ρ such that $\mathrm{Cst}'(\mathrm{ext}(\rho)) = f(b')$. Towards a contradiction assume $\mathrm{Cst}(\rho) = b'' \neq b'$. We again have $f(b'') = \mathrm{Cst}'(\mathrm{ext}(\rho))$. First assume

$b'' < b'$. Then we have $b'' < b' < b$, which implies $f(b'') < f(b')$, contradicting $\mathrm{Cst}'(\mathrm{ext}(\rho)) = f(b')$. If $b' < b'' < b$, we obtain the contradiction $\mathrm{Cst}'(\mathrm{ext}(\rho)) > f(b')$ analogously. Finally, if $b \leq b''$, then $\mathrm{Cst}'(\mathrm{ext}(\rho)) = f(b'') \geq f(b) > f(b')$, which again contradicts $\mathrm{Cst}'(\mathrm{ext}(\rho)) = f(b')$.

3) Let ρ such that $\mathrm{Cst}'(\mathrm{ext}(\rho)) \geq f(b)$. Towards a contradiction assume $\mathrm{Cst}(\rho) = b' < b$. We again have $f(b') = \mathrm{Cst}'(\mathrm{ext}(\rho))$. However, we obtain $f(b') < f(b)$ due to f being a b-correction function. This contradicts $\mathrm{Cst}'(\mathrm{ext}(\rho)) = f(b') \geq f(b)$. □

These properties of correction functions when used in quantitative reductions enable us to state and prove the main result of this section. This result establishes quantitative reductions as the quantitative counterpart to qualitative reductions: If $\mathcal{G} \leq^{b+1}_{\mathcal{M},f} \mathcal{G}'$ and b is a cap of $\mathcal{G}$, 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 $\mathcal{G}'$ whose cost is at most $f(b)$. This holds true for both players. If a strategy has cost greater than $f(b)$, however, we do not have a direct correspondence between costs of plays in $\mathcal{G}$ and $\mathcal{G}'$ anymore. We are, however, still able to claim the existence of a strategy of infinite cost for Player 1 in $\mathcal{G}$ once he can ensure a cost greater than $f(b)$ in $\mathcal{G}'$, due to b being a cap of $\mathcal{G}$.

Theorem 2. *Let $\mathcal{G}$ and $\mathcal{G}'$ be determined quantitative games such that $\mathcal{G} \leq^{b+1}_{\mathcal{M},f} \mathcal{G}'$ 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 $\mathrm{Cst}'(\sigma') = f(b')$ if and only if she has a strategy σ in $\mathcal{G}$ with $\mathrm{Cst}(\sigma) = b'$.*
2. *If Player 1 has a strategy τ' in $\mathcal{G}'$ with $\mathrm{Cst}'(\tau') \geq f(b + 1)$, then he has a strategy τ in $\mathcal{G}$ with $\mathrm{Cst}(\tau) = \infty$.*

Proof. 1) Let σ' be a strategy for Player i in $\mathcal{G}'$ such that $\mathrm{Cst}'(\sigma') = f(b')$ for some $b' \leq b$. For all play prefixes π ending in a vertex in V_i in $\mathcal{G}$, we define the strategy σ for Player i in $\mathcal{G}$ via $\sigma(\pi) = v$, if $\sigma'(\mathrm{ext}(\pi)) = (v, m)$. Let the infinite play ρ be consistent with σ. A straightforward induction shows that $\rho' = \mathrm{ext}(\rho)$ is consistent with σ'. If $i = 0$, then $\mathrm{Cst}'(\rho') = \mathrm{Cst}'(\mathrm{ext}(\rho)) \leq f(b')$ and thus, $\mathrm{Cst}(\rho) \leq b'$, due to Lemma 1.1 and Lemma 1.2, as $b' < b + 1$. This in turn implies $\mathrm{Cst}(\sigma) \leq b'$. If $i = 1$, we obtain $\mathrm{Cst}(\rho) \geq b'$ using similar reasoning.

Since $b' < b+1 < \infty$, we obtain $f(b') < \infty$: If $f(b') = \infty$, we have $f(b'+1) = \infty$, which contradicts strict monotonicity of f up to and including $b+1$. Let ρ' be a play consistent with σ' such that $\mathrm{Cst}'(\rho') = f(b')$. Since $\mathrm{Cst}'(\sigma') = f(b') < \infty$, such a play exists. Moreover, let ρ be the unique play such that $\mathrm{ext}(\rho) = \rho'$. By induction we obtain that ρ is consistent with σ. Additionally, we have $\mathrm{Cst}(\rho) = b'$ due to $\mathrm{Cst}'(\rho') = \mathrm{Cst}'(\mathrm{ext}(\rho)) = f(b')$ and Lemma 1.2. Hence, $\mathrm{Cst}(\sigma) \geq b'$ if $i = 0$, and $\mathrm{Cst}(\sigma) \leq b'$ if $i = 1$.

Now let σ be a strategy in $\mathcal{G}$ with $\mathrm{Cst}(\sigma) = b' < b$. For all play prefixes $\mathrm{ext}(\pi) = (v_0, m_0) \cdots (v_j, m_j)$ ending in a vertex in $V_i \times M$ in $\mathcal{G}'$, we define the strategy σ' as $\sigma'(\mathrm{ext}(\pi)) = (v, \mathrm{Upd}(m_j, v))$ if $\sigma(\pi) = v$. We claim $\mathrm{Cst}'(\sigma') = f(b')$. Let $\mathrm{ext}(\rho)$ be a play consistent with σ'. A straightforward induction yields that ρ is consistent with σ, hence $\mathrm{Cst}(\rho) \leq b'$ and $\mathrm{Cst}(\mathrm{ext}(\rho)) \leq f(b')$ due to $b' < b+1$. Hence, $\mathrm{Cst}'(\sigma') \leq f(b')$. Now let ρ be a play consistent with σ such that $\mathrm{Cst}(\rho) = b'$. Since $b' < \infty$, such a play exists. Via another straightforward induction we obtain that $\mathrm{ext}(\rho)$ is consistent with σ'. As $\mathrm{Cst}'(\mathrm{ext}(\rho)) = f(b')$, we obtain $\mathrm{Cst}'(\sigma') \geq f(b')$.

2) Let τ' be a strategy for Player 1 in $\mathcal{G}'$ with $\mathrm{Cst}'(\tau') \geq b+1$. We define the strategy τ for Player 1 in $\mathcal{G}$ via $\tau(\pi) = v$ if $\tau'(\mathrm{ext}(\pi)) = (v, m)$ for all play prefixes π in $\mathcal{G}$. Let ρ' be a play consistent with τ' such that $\mathrm{Cst}'(\rho') \geq f(b+1)$ and let ρ be the unique play in $\mathcal{G}$ such that $\mathrm{ext}(\rho) = \rho'$. A straightforward induction yields that ρ is consistent with τ. Then we obtain $\mathrm{Cst}(\rho) \geq b+1$ due to Lemma 1.3, which in turn implies $\mathrm{Cst}(\tau) \geq b+1$. Since b is a cap of $\mathcal{G}$ and due to determinacy of $\mathcal{G}$, this implies that there exists a strategy τ'' for Player 1 in $\mathcal{G}$ such that $\mathrm{Cst}(\tau'') = \infty$. □

We prove Theorem 2, 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 are thus of infinite size. If Player 0 can play achieve a certain cost in $\mathcal{G}'$ using a finite-state strategy, however, then she can do so in $\mathcal{G}$ with a finite-state strategy as well.

Theorem 3. *Let $\mathcal{G}$ and $\mathcal{G}'$ be quantitative games such that $\mathcal{G} \leq^b_{\mathcal{M}_1,f} \mathcal{G}$ for some b, $\mathcal{M}$, and f and let $b' \leq b$. If Player i has a finite-state strategy σ' with $\mathrm{Cst}(\sigma') = f(b')$ in $\mathcal{G}'$ that is implemented by $\mathcal{M}_2$, then she has a finite-state strategy σ with $\mathrm{Cst}(\sigma) = b'$ in $\mathcal{G}'$ that is implemented by $\mathcal{M}_1 \times \mathcal{M}_2$.*

Proof. Let $\mathcal{G} = (\mathcal{A}, \mathrm{Cst})$, $\mathcal{G}' = (\mathcal{A}', \mathrm{Cst}')$, $\mathcal{M}_1 = (M_1, m^1_I, \mathrm{Upd}_1)$, and $\mathcal{M}_2 = (M_2, m^2_I, \mathrm{Upd}_2)$ such that σ' is implemented by $\mathcal{M}_2$ with the next-move function $\mathrm{Nxt}' \colon (V \times M_1) \times M_2 \to (V \times M_1)$. We define $\mathrm{Nxt}(v, (m_1, m_2)) = v^*$ if $\mathrm{Nxt}'((v, m_1), m_2) = (v^*, \mathrm{Upd}(m_1, v^*))$. Let σ be the strategy implemented by $\mathcal{M}_1 \times \mathcal{M}_2$ with the next-move function Nxt.

Let $\rho = v_0 v_1 v_2 \cdots$ be a play consistent with σ, let

$$\mathrm{ext}_{\mathcal{M}_1 \times \mathcal{M}_2}(\rho) = (v_0, m^0_1, m^0_2)(v_1, m^1_1, m^1_2)(v_2, m^2_1, m^2_2) \cdots$$

be its extension with respect to $\mathcal{M}_1 \times \mathcal{M}_2$, and let $j \in \mathbb{N}$ such that $v_j \in V_0$. Then $v_{j+1} = \sigma(v_0 \cdots v_j) = \mathrm{Nxt}(v_j, (m^1_j, m^2_j))$. Due to the definition of Nxt, this implies $\mathrm{Nxt}'((v_j, m^1_j), m^2_j) = (v_{j+1}, m^1_{j+1})$, where $m^1_{j+1} = \mathrm{Upd}_1(m_j, v_{j+1})$ due to the construction of $\mathcal{A} \times \mathcal{M}_1$. Hence, $\mathrm{ext}_{\mathcal{M}_1}(\rho)$ is consistent with σ', i.e., $\mathrm{Cst}'(\mathrm{ext}_{\mathcal{M}_1}(\rho)) \leq f(b')$, which in turn implies $\mathrm{Cst}(\rho) \leq b'$ for $i = 0$ and $\mathrm{Cst}'(\mathrm{ext}_{\mathcal{M}_1}(\rho)) \geq f(b')$ and $\mathrm{Cst}(\rho) \geq b'$ for $i = 1$.

Note that, due to similar reasoning, for each play $\mathrm{ext}_{\mathcal{M}_1}(\rho)$ consistent with σ' the play ρ is consistent with σ'. If $i = 1$ or $b' < \infty$, this concludes the proof. If, however, $i = 0$ and $b' = \infty$, then we furthermore obtain $b = \infty$ and that f is a strictly monotonic function with $f(\infty) = \infty$. Hence, if there exists a play $\mathrm{ext}_{\mathcal{M}_1}(\rho)$ consistent with σ' with $\mathrm{Cst}'(\mathrm{ext}_{\mathcal{M}_1}(\rho)) = \infty$, then $\mathrm{Cst}(\rho) = \infty$ and hence, $\mathrm{Cst}(\sigma) = \infty$. If, however, the costs of the plays consistent with σ' diverges, then the cost of the plays consistent with σ diverges as well and we obtain $\mathrm{Cst}(\sigma) = \infty$. □

After establishing quantitative reductions as the counterpart to qualitative ones, we now turn our attention to providing a "backend" for such reductions.

4 Vertex-Ranked Games

We introduce a very simple form of quantitative games, so-called 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 provide tight bounds on the complexity of solving such games with respect to a given bound and on the necessary memory for achieving such a bound. 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 $V' \supseteq V$, let $\mathrm{Win} \subseteq (V')^\omega$ be a (qualitative) winning condition, and let $\mathit{rank}\colon V' \to \mathbb{N}$ be a ranking function on vertices. We define the quantitative vertex-ranked sup-condition

$$\textsc{Rank}^{\sup}(\mathrm{Win}, \mathit{rank})(v_0v_1v_2\cdots) = \begin{cases} \sup_{j\to\infty} \mathit{rank}(v_j) & \text{if } v_0v_1v_2\cdots \in \mathrm{Win} \\ \infty & \text{otherwise} \end{cases}$$

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

$$\textsc{Rank}^{\lim}(\mathrm{Win}, \mathit{rank})(v_0v_1v_2\cdots) = \begin{cases} \limsup_{j\to\infty} \mathit{rank}(v_j) & \text{if } v_0v_1v_2\cdots \in \mathrm{Win} \\ \infty & \text{otherwise} \end{cases} .$$

A vertex-ranked sup- or lim sup-game $\mathcal{G} = (\mathcal{A}, \textsc{Rank}^X(\mathrm{Win}, \mathit{rank}))$ with $X \in \{\sup, \lim\}$ consists of an arena $\mathcal{A}$ with vertex set V, a qualitative winning condition Win, and a a vertex-ranking function $\mathit{rank}\colon V \to \mathbb{N}$. We assume the ranks to be encoded in binary.

If $\mathcal{G} = (\mathcal{A}, \textsc{Rank}^X(\mathrm{Win}, \mathit{rank}))$ is a vertex-ranked sup- or lim sup-game, we call the game $(\mathcal{A}, \mathrm{Win})$ the qualitative game corresponding to $\mathcal{G}$. Moreover, if $\mathcal{G}_{\sup}$ is a vertex-ranked sup-game, we denote the vertex-ranked lim sup-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 upper 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 lim sup-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 their corresponding qualitative games. This is due to the fact that Player 0 has a winning strategy in $(\mathcal{G}, \mathrm{Win})$ if and only if she has a strategy with cost at most 0 in $(\mathcal{A}, \textsc{Rank}^{\sup}(\mathrm{Win}, rank))$, where *rank* assigns the rank 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

$$\mathfrak{G}^{\text{RNK}}_{\sup} = \{(\mathcal{A}, \textsc{Rank}^{\sup}(\mathrm{Win}, rank)) \mid (\mathcal{A}, \mathrm{Win}) \in \mathfrak{G}, rank \text{ is vertex-ranking function for } \mathcal{A}\} \enspace .$$

We first show that we can use a decision procedure solving games from $\mathfrak{G}$ to solve games from $\mathfrak{G}^{\text{RNK}}_{\sup}$ 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 (quantitative) games $\mathfrak{G}$ *proper* if for each $(\mathcal{A}, \mathrm{Win})$ $((\mathcal{A}, \mathrm{Cst}))$ in $\mathfrak{G}$ and each sub-arena $\mathcal{A}' \sqsubseteq \mathcal{A}$ the game $(\mathcal{A}', \mathrm{Win})$ $((\mathcal{A}', \mathrm{Cst}))$ 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 and let $X \subseteq V$. We define $\mathrm{Attr}_i(X) = \mathrm{Attr}^n_i(X)$ inductively with $\mathrm{Attr}^0_i(X) = X$ and

$$\mathrm{Attr}^j_i(X) = \{v \in V_i \mid \exists v' \in \mathrm{Attr}^{j-1}_i(X).\, (v, v') \in E\} \cup \{v \in V_{1-i} \mid \forall (v, v') \in E.\, v' \in \mathrm{Attr}^{j-1}_i(X)\} \cup \mathrm{Attr}^{j-1}_i(X) \enspace .$$

Intuitively, $\mathrm{Attr}_i(X)$ is the set of all vertices from which Player i can enforce a visit to X. The set $\mathrm{Attr}_i(X)$ can be computed in linear time in $|V|$ and Player i has a positional strategy σ such that each play starting in some vertex in $\mathrm{Attr}_i(X)$ and consistent with σ eventually encounters some vertex from X [16]. We call σ the attractor strategy towards X.

Let $\mathcal{A}$ be an arena with vertex set V, let $X \subseteq V$, and let $A = \mathrm{Attr}_i(X)$. If $v_I \notin A$, then we define $\mathcal{A} \setminus A = (V \setminus A, V_0 \setminus A, V_1 \setminus A, E \setminus (A \times A), v_I)$. Note that $\mathcal{A} \setminus A$ is again an arena. We lift this notation to qualitative (quantitative) games $\mathcal{G} = (\mathcal{A}, \mathrm{Win})$ $((\mathcal{A}, \mathrm{Cst}))$ by defining $\mathcal{G} \setminus A = (\mathcal{A} \setminus A, \mathrm{Win})$ $((\mathcal{A} \setminus A, \mathrm{Cst}))$.

If $v_I \in A$, however, then both $\mathcal{A} \setminus A$ and $\mathcal{G} \setminus A$ are undefined. The game $\mathcal{G} \setminus A$ can be constructed in linear time and is of size at most $|\mathcal{G}|$.

We now show that the idea behind the construction for solving vertex-ranked sup-games described above is indeed correct.

Lemma 2. *Let $\mathfrak{G}$ be a proper class of qualitative games, let $\mathcal{G}_{\text{sup}} \in \mathfrak{G}^{\text{RNK}}_{\text{sup}}$ with vertex set V, initial vertex v_I, and ranking function rank and let $b \in \mathbb{N}$.*

Player 0 has a strategy with cost at most b in $\mathcal{G}_{\text{sup}}$ if and only if $v_I \notin A$ and if she has a winning strategy in $\mathcal{G} \setminus A$, where $A = \text{Attr}_1(\{v \in V \mid \mathit{rank}(v) > b\})$.

Proof. Let $X_b = \{v \in V \mid \mathit{rank}(v) > b\}$. We first show that, if $v_I \notin A$ and if Player 0 wins $\mathcal{G}' = \mathcal{G} \setminus A$, say with strategy σ', then she has a strategy of cost at most b in $\mathcal{G}_{\text{sup}}$. Since $\mathcal{A}' \sqsubseteq \mathcal{A}$, the strategy σ' is a strategy for Player 0 in $\mathcal{G}_{\text{sup}}$ as well and each play consistent with σ' in $\mathcal{G}'$ is consistent with σ' in $\mathcal{G}$ as well as vice versa. Let ρ be a play in $\mathcal{G}_{\text{sup}}$ consistent with σ'. Since σ' is winning for Player 0 in $\mathcal{G}'$, we have $\rho \in \text{Win}$. Moreover, since $X_b \subseteq A$, and as ρ visits only vertices occurring in $\mathcal{G}'$, we obtain $\textsc{Rank}^{\text{sup}}(\text{Win}, \mathit{rank})(\rho) \leq b$ and thus $\text{Cst}(\sigma') \leq b$, which concludes this direction of the proof.

For the other direction, first assume $v_I \in A$, let τ_A be the attractor strategy towards X_b for Player 1. We obtain $\text{Cst}(\tau_A) > b$ in $\mathcal{G}_{\text{sup}}$: By playing consistently with τ_A, Player 1 forces the play to eventually reach a vertex in X_b, i.e., a vertex v with $\mathit{rank}(v) > b$. Thus, $\text{Cst}(\tau_A) > b, \textit{i.e.},\ \text{Cst}(\sigma) > b$ for all strategies σ of Player 0.

Finally, assume $v_I \notin A$ and that Player 0 does not have a winning strategy in $\mathcal{G}'$. Towards a contradiction, additionally assume that she has a strategy σ with cost at most b in $\mathcal{G}_{\text{sup}}$. Note that no play consistent with σ visits any vertex from A. Otherwise, playing consistently with his attractor strategy towards X_b from the first visit to A, Player 1 could construct a play consistent with σ, but with cost greater than b. Thus, σ is a strategy for Player 0 in $\mathcal{G}'$ and we obtain that all plays consistent with σ in $\mathcal{A}$ are consistent with σ in $\mathcal{A}'$ and vice versa. Since $\text{Cst}(\sigma) \leq b$, we obtain $\textsc{Rank}^{\text{sup}}(\text{Win}, \mathit{rank})(\rho) < \infty$, i.e., $\rho \in \text{Win}$ for all plays ρ consistent with σ, a contradiction. □

Using this lemma, we are able to construct a decision procedure solving games from $\mathfrak{G}^{\text{RNK}}_{\text{sup}}$ using a decision procedure solving games from $\mathfrak{G}$.

Theorem 4. *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 $\mathcal{G}_{\text{sup}} \in \mathfrak{G}^{\text{RNK}}_{\text{sup}}$ with n vertices and some bound $b \in \mathbb{N}$, does Player 0 win $\mathcal{G}_{\text{sup}}$ with respect to b?"

Proof. First note that, since $\mathfrak{G}$ is proper, $\mathfrak{G}^{\text{RNK}}_{\text{sup}}$ is proper as well. Given the vertex-ranked sup-game $\mathcal{G}_{\text{sup}} = (\mathcal{A}, \textsc{Rank}^{\text{sup}}(\text{Win}, \mathit{rank}))$, let $X_b = \{v \in V \mid \mathit{rank}(v) > b\}$ and let $A = \text{Attr}_1(X_b)$. We define the decision procedure dec_{sup} deciding the given problem such that it returns false if $v_I \in A$. Otherwise, dec_{sup} returns true if and only if Player 0 wins $\mathcal{G} \setminus A$. The procedure dec_{sup} indeed decides the given decision problem due to Lemma 2.

Since we can compute and remove the Player-1-attractor A in linear time in $|\mathcal{A}| = n$ [16], the decision procedure $\mathrm{dec}_{\mathrm{sup}}$ indeed requires time $\mathcal{O}(n) + t(|\mathcal{G}|)$ and space $\mathcal{O}(n) + s(|\mathcal{G}|)$. □

This theorem provides an upper bound on the complexity of solving vertex-ranked sup-games. In the proof of Lemma 2 we show that, for any vertex-ranked sup-game $\mathcal{G}_{\mathrm{sup}}$, a winning strategy for Player 0 in $\mathcal{G} \setminus A$ 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 $\mathfrak{G}^{\mathrm{RNK}}_{\mathrm{sup}}$ 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 $\mathfrak{G}^{\mathrm{RNK}}_{\mathrm{sup}}$ to construct strategies of cost at most (greater than) b for Player 0 (Player 1) as well.

Finally, this procedure enables us to solve the optimization problem for vertex-ranked sup-games from $\mathfrak{G}^{\mathrm{RNK}}_{\mathrm{sup}}$: Recall that if Player 0 wins $\mathcal{G}_{\mathrm{sup}}$ 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 $\mathrm{dec}_{\mathrm{sup}}$ suffice to determine the minimal b such that Player 0 wins $\mathcal{G}_{\mathrm{sup}}$ 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 lim sup-Games

We now turn our attention to solving vertex-ranked lim sup-games. For the same reasons as above, solving these games is at least as hard as solving their corresponding qualitative games. Thus, we again show upper bounds on the complexity of solving these games. To this end, given some class $\mathfrak{G}$ of games, we define the corresponding vertex-ranked lim sup-games

$$\mathfrak{G}^{\mathrm{RNK}}_{\mathrm{lim}} = \{(\mathcal{A}, \textsc{Rank}^{\mathrm{lim}}(\mathrm{Win}, \mathit{rank})) \mid (\mathcal{A}, \mathrm{Win}) \in \mathfrak{G}, \mathit{rank} \text{ is vertex-ranking function for } \mathcal{A}\} \ ,$$

We identify two sufficient criteria on classes of qualitative games $\mathfrak{G}$ for quantitative games in $\mathfrak{G}^{\mathrm{RNK}}_{\mathrm{lim}}$ being solvable with respect to some given b. More precisely, we provide decision procedures for $\mathfrak{G}^{\mathrm{RNK}}_{\mathrm{lim}}$ if games from $\mathfrak{G}$ can be solved in conjunction with CoBüchi-conditions and if games from $\mathfrak{G}$ are prefix-independent.

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

Remark 1. Let $\mathfrak{G}$ be a class of qualitative games such that games from $\{(\mathcal{A}, \mathrm{Win} \cap \textsc{CoBüchi}(F)) \mid (\mathcal{A}, \mathrm{Win}) \in \mathfrak{G}, F \subseteq V, V \text{ is vertex set of } \mathcal{A}\}$ 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 $t(|\mathcal{G}|, n)$ and space $s(|\mathcal{G}|, n)$: "Given some game $\mathcal{G}_{\lim} \in \mathfrak{G}_{\lim}^{\text{RNK}}$ with n vertices and some bound $b \in \mathbb{N}$, does Player 0 win $\mathcal{G}_{\lim}$ with respect to b?"

In this first case, we use a decision procedure for solving qualitative games as-is, which requires the existence of such a specific decision procedure. 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 classes of vertex-ranked lim sup-games. We now turn our attention to solving another wide range of such games. More precisely, we consider games from $\mathfrak{G}_{\lim}^{\text{RNK}}$ where the winning conditions of the qualitative games from $\mathfrak{G}$ are insensitive to finite prefixes.

Formally, we call a qualitative winning condition $\text{Win} \subseteq V^\omega$ *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.

We now turn our attention to games $\mathcal{G}_{\lim}$ from $\mathfrak{G}_{\lim}^{\text{RNK}}$ where $\mathfrak{G}$ is prefix-independent. In order to do so, we adapt the classical algorithm for solving prefix-independent qualitative games (cf., e.g., [7]). Thereby, we repeatedly compute the set of vertices from which Player 0 has a strategy of cost at most b in $\mathcal{G}_{\sup}$ and remove their attractor from the game. Once the obtained games stabilize, we proclaim that Player 0 has a strategy with cost at most b in $\mathcal{G}_{\lim}$ if and only if v_I was removed during the above construction.

In order to formalize this approach, let $\mathcal{G}$ be a quantitative game with vertex set V. For each $v \in V$, we write $\mathcal{G}_v$ to denote the game $\mathcal{G}$ with its initial vertex replaced by v. All other components, i.e., the structure of the arena and the cost-function, remain unchanged. We write $\text{W}_i^b(\mathcal{G})$ to denote the set of all vertices v such that Player i has a strategy of cost at most b, if $i = 0$, or greater than b, if $i = 1$, in $\mathcal{G}_v$.

The approach described above is sound, as the sup-game is harder to win than the lim sup-game for Player 0. Also, if Player 0 does not win the sup-game from any vertex, then she also does not win the lim sup-game from any vertex.

Lemma 3. *Let* $\mathcal{G}_{\lim} = (\mathcal{A}, \text{RANK}^{\lim}(\text{Win}, \mathit{rank}))$ *be some vertex-ranked* lim sup*-game such that* Win *is prefix-independent and such that* $\mathcal{G}_{\sup}$ *is determined. If* $\text{W}_0^b(\mathcal{G}_{\sup}) = \emptyset$*, then* $\text{W}_0^b(\mathcal{G}_{\lim}) = \emptyset$.

Proof. Let V be the vertex set of $\mathcal{G}_{\sup}$ and $\mathcal{G}_{\lim}$. Since $\text{W}_0^b(\mathcal{G}_{\sup}) = \emptyset$ and since $\mathcal{G}_{\sup}$ is determined, we obtain $\text{W}_1^b(\mathcal{G}_{\sup}) = V$. For each $v \in V$, let τ'_v be a strategy for Player 1 in $(\mathcal{G}_{\sup})_v$ with cost greater than b. We now define a single strategy τ for Player 1 in $\mathcal{G}_{\lim}$ with cost greater than b. For each $\pi = v_0 \cdots v_j \in V^*$ we define $\tau(\pi) = \tau'_{v_k}(v_k \cdots v_j)$, where $k = \max\{k' \mid \mathit{rank}(v_{k'-1}) > b\}$, with $\max \emptyset = 0$. We claim that τ has cost greater than b in all $(\mathcal{G}_{\lim})_v$. As all $(\mathcal{G}_{\lim})_v$ have the cost-function Cst, we formally claim $\text{Cst}(\tau) > b$.

Let $\rho = v_0 v_1 v_2 \cdots$ be a play of $\mathcal{G}_v$ consistent with τ. If there are infinitely many positions j with $\mathit{rank}(v_j) > b$, then $\text{Cst}(\rho) > b$. Thus, assume the opposite and let j be the maximal position with $rank(v_j) > b$. Then the suffix $\rho' = v_{j+1}v_{j+2}v_{j+3}\cdots$ of ρ is consistent with $\tau'_{v_{j+1}}$. Since ρ' does not encounter any vertices of rank greater than b, while $\mathrm{Cst}(\rho') > b$ due to ρ' being consistent with a strategy of cost greater than b, we obtain $\rho' \notin \mathrm{Win}$. This implies $\rho \notin \mathrm{Win}$ due to prefix-independence of Win. Hence, $\mathrm{Cst}(\rho) = \infty$, which, together with the statement above, implies $\mathrm{Cst}(\tau) > b$. □

We formalize the construction described above in the following theorem.

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 $\mathcal{O}(n^3 + n^2 \cdot t(|\mathcal{G}|))$ and space $\mathcal{O}(n) + s(|\mathcal{G}|)$: "Given some game $\mathcal{G}_{\mathrm{lim}} \in \mathfrak{G}^{\mathrm{RNK}}_{\mathrm{lim}}$ with n vertices and some bound $b \in \mathbb{N}$, does Player 0 win $\mathcal{G}_{\mathrm{lim}}$ with respect to b?"

Proof. Given $\mathcal{G}_{\mathrm{lim}} = (\mathcal{A}, \textsc{Rank}^{\mathrm{lim}}(\mathrm{Win}, rank))$ with vertex set V of size n, we define $\mathcal{G}_0 = \mathcal{G}_{\mathrm{sup}}$, as well as $X_j = \mathrm{W}^b_0(\mathcal{G}_j)$, $A_j = \mathrm{Attr}_0(X_j)$, which is computed in the arena of $\mathcal{G}_j$, and $\mathcal{G}_{j+1} = \mathcal{G}_j \setminus A_j$ for all $j \in \mathbb{N}$. As we only remove vertices from the $\mathcal{G}_j$, we obtain $\mathcal{G}_{j+1} \sqsubseteq \mathcal{G}_j$. Thus, the series of games stabilizes at $j = n$ at the latest, i.e., $\mathcal{G}_j = \mathcal{G}_n$ for all $j \geq n$. We define $A = \bigcup_{j \leq n} A_j$ and $\mathcal{G}' = \mathcal{G}_n$ and claim that Player 0 has a strategy with cost at most b in $\mathcal{G}$ if and only if $v_I \in A$. We first argue that this suffices to show the desired result.

Let $\mathrm{dec}_{\mathrm{sup}}$ be the decision procedure deciding whether or not Player 0 has a strategy with cost at most b in games from $\mathfrak{G}^{\mathrm{RNK}}_{\mathrm{sup}}$, as constructed in the proof of Theorem 4. The decision procedure $\mathrm{dec}_{\mathrm{sup}}$ can be easily modified to return $W^b_0(\mathcal{G}_j)$ instead of a yes/no-answer by applying it to each $(\mathcal{G}_j)_v$ individually. This modified procedure $\mathrm{dec}'_{\mathrm{sup}}$ runs in time at most $\mathcal{O}(n^2 + n \cdot t(|\mathcal{G}|))$ and space $\mathcal{O}(n) + s(|\mathcal{G}|)$, where $t(|\mathcal{G}|)$ and $s(|\mathcal{G}|)$ are the time and space required to solve $\mathcal{G}$, respectively.

For $j \in \{0, \ldots, n\}$, the decision procedure $\mathrm{dec}_{\mathrm{lim}}$ first computes $\mathcal{G}_j$ in linear time in n and reusing the space used for solving $\mathcal{G}_{j-1}$. It then computes X_j requiring a single call to the modified $\mathrm{dec}_{\mathrm{sup}}$. It subsequently computes A_j in time $\mathcal{O}(n)$ and space $\mathcal{O}(n)$. Finally, it returns false if and only if v_I is in the arena of $\mathcal{G}_n$. In total, we obtain a runtime of $\mathrm{dec}_{\mathrm{lim}}$ of $\mathcal{O}(n^3 + n^2 \cdot t(|\mathcal{G}|))$. The only additional memory required by $\mathrm{dec}_{\mathrm{lim}}$ is that for storing the sets X_j and A_j, the size of which is bounded from above by n. The games $\mathcal{G}_j$ can be stored by reusing the memory occupied by $\mathcal{G}$, due to $\mathcal{G}_j \sqsubseteq \mathcal{G}_{j-1}$. Hence, the procedure $\mathrm{dec}_{\mathrm{lim}}$ requires space $\mathcal{O}(n) + s(|\mathcal{G}|)$.

It remains to show that Player 0 has a strategy with cost at most b in $\mathcal{G}$ if and only if $v_I \notin \mathcal{A}_n$, i.e., if $v \in A$. To this end, first assume $v_I \in A$ and note that we have $A_j \supseteq X_j$. However, for each two $j \neq j'$, we have $A_j \cap A_{j'} = \emptyset$ and, in particular, $X_j \cap X_{j'} = \emptyset$. Hence, for each $v \in A$ there exists a unique j such that $v \in A_j$.

We define the strategy σ for Player 0 in $\mathcal{G}$ inductively such that any play consistent with σ only descends through the X_j. Formally, we construct σ such that it satisfies the following invariant:

Let $\rho = v_0 v_1 v_2 \cdots$ be a play consistent with σ and let $k \in \mathbb{N}$. If $v_k \in A_j \setminus X_j$, then $v_{k+1} \in \bigcup_{j' \leq j} A_{j'} \cup X_{j'}$. Moreover, if $v_k \in (A_j \setminus X_j) \cap V_0$, then the move to v_{k+1} is the move prescribed by the attractor strategy of Player 0 towards X_j. If $v_k \in X_j$, then $v_{k+1} \in X_j \cup \bigcup_{j' < j} A_{j'} \cup X_{j'}$.

Clearly, this invariant holds true for $\pi = v_I$. Thus, let $\pi = v_0 \cdots v_k$ be a play prefix consistent with σ. If $v_k \in V_1$, let v^* be an arbitrary successor of v_k in $\mathcal{G}$ and assume towards a contradiction that πv^* violates the invariant. If $v_k \in A_j \setminus X_j$, then in $\mathcal{G}_j$ there exists an edge from v_k leading to some vertex $v^* \notin A_j$, a contradiction to the definition of the attractor. If, however, $v_k \in X_j$ and $v^* \notin X_j \cup \bigcup_{j' < j} A_{j'} \cup X_{j'}$, then Player 1 has a strategy τ in $(\mathcal{G}_j)_{v^*}$ with cost greater than b. Thus, a play that begins in v_k, moves to v^* and is consistent with τ afterwards has cost greater b, i.e., Player 0 does not have a strategy with cost at most b in $(\mathcal{G}_j)_{v_k}$, a contradiction to $v_k \in X_j = \mathrm{W}_0^b(\mathcal{G}_j)$. Hence, πv^* satisfies the invariant for each successor v^* of $v_k \in V_1$.

Now assume $v_k \in V_0$ and first let $v_0 \in A_j \cup X_j$ for some $j \in \mathbb{N}$. Let σ_j^A be the attractor strategy for Player 0 towards X_j. If $v_k \in A_j \setminus X_j$, we define $\sigma(\pi) = \sigma_j^A(v_k)$, which satisfies the invariant due to the definition of the attractor strategy. If, however, $v_k \in X_j$, let k' be minimal such that $v_{k''} \in X_j$ for all k'' with $k' \leq k'' \leq k$. Moreover, let σ_j^v be a strategy for Player 0 such that every play consistent with σ_j^v in $\mathcal{G}_j$ with initial vertex v has cost at most b. Such a strategy exists due to $X_j = \mathrm{W}_0^b(\mathcal{G}_j)$. We define $\sigma(\pi) = \sigma_j^{v_{k'}}(v_{k'} \cdots v_k)$, which satisfies the invariant to similar reasoning as above.

In order to show $\mathrm{Cst}(\sigma) \leq b$, let $\rho = v_0 v_1 v_2 \cdots$ be a play consistent with σ. Due to the invariant of σ and since $v_0 \in A$, the play ρ descends through the A_j and the X_j, i.e., once it encounters some X_j, it never moves to any $A_{j'} \setminus X_{j'}$ with $j' \geq j$ nor to any $X_{j'}$ with $j' > j$. Also, ρ stabilizes in some X_j, i.e., there exists a $k \in \mathbb{N}$ such that $v_{k'} \in X_j$ for all $k' \geq k$, as σ prescribes moves according to the attractor strategy towards X_j when in $A_j \setminus X_j$. Moreover, due to the definition of σ, the suffix $\rho' = v_k v_{k+1} v_{k+2} \cdots$ is consistent with $\sigma_j^{v_k}$, i.e., we obtain $\rho' \in \mathrm{Win}$ and that the maximal vertex-rank encountered in ρ is at most b. As Win is prefix-independent, we obtain $\rho \in \mathrm{Win}$ as well as $\limsup_{k \to \infty} \mathit{rank}(v_k) \leq b$. Hence, $\textsc{Rank}^{\lim}(\mathrm{Win}, \mathit{rank})(\rho) \leq b$, which concludes this direction of the proof.

Now assume $v_I \notin A$ and consider $\mathcal{G}'$ with vertex set $V \setminus A$. Since the construction of the $\mathcal{G}_j$ stabilized, we have $A_j = X_j = W_0^b(\mathcal{G}') = \emptyset$, i.e., Player 1 has a strategy with cost greater than b from any starting vertex in $\mathcal{G}'$. Due to Lemma 3, this implies that he has such a strategy from every vertex in $\mathcal{G}_{\sup} \setminus A$, call it τ. Note that there exists no Player-0-vertex in $V \setminus A$ that has an outgoing edge leading into A, as this would contradict the definition of the Player-0-attractors A_j. Hence, τ is a strategy for Player 1 in $\mathcal{G}$ as well and we retain $\mathrm{Cst}(\tau) > b$. □

The strategy σ for Player 0 constructed in the proof of Theorem 5 works by "stitching together" the attractor-strategies leading her to the X_j and the strategies for her in the respective vertex-ranked sup-games. As no play consistent with σ ever returns to earlier X_j or A_j, we can reuse the memory states of the winning strategies in the $\mathcal{G}_j$ when implementing σ. Thus, a monotonic upper bound on the size of strategies with cost at most b in $\mathcal{G}_{\text{sup}}$ is an upper bound on the size of such strategies in $\mathcal{G}_{\text{lim}}$ as well.

Moreover, in order to find the optimal b such that Player 0 wins $\mathcal{G}_{\text{lim}}$ 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 $\mathcal{O}(n) + s(|\mathcal{G}|)$.

5 Applications

Having defined the framework of quantitative reductions in Section 3 and vertex-ranked games as general-purpose targets for such reductions in Section 4, we now turn to applications of both concepts. In particular, we first reduce request-response games with costs to vertex-ranked request-response games, thereby establishing ExpTime-membership of the problem of solving the former games with respect to a given bound. Moreover, we reduce quantitative Muller games to quantitative safety games, thus providing a novel proof of ExpTime-membership of the problem of solving the former games with respect to a given bound.

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.

Formally, the qualitative request-response condition $\textsc{ReqRes}(\Gamma)$ consists of a family of so-called request-response pairs $\Gamma = (Q_c, P_c)_{c \in [d]}$. 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

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

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.*

Theorem 6.

1. *Request-response games with* n *vertices and* d *request-response pairs can be solved in time* $\mathcal{O}(n^2 d^2 2^d)$. *[18]*
2. *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* $d2^d$ *[18].*

We extend this winning condition to a quantitative one using families of cost functions $\mathrm{Cst} = (\mathrm{Cst}_c)_{c \in [d]}$, where $\mathrm{Cst}_c \colon E \to \mathbb{N}$ for each $c \in [d]$. The cost-of-response for a request for condition c at position j is defined as

$$\textsc{ReqResCor}_c(v_0 v_1 v_2 \cdots, j) = \begin{cases} \min\{\mathrm{Cst}_c(v_j \cdots v_{j'}) \mid j' \geq j \text{ and } v_{j'} \in P_c\} & \text{if } v_j \in Q_c \\ 0 & \text{otherwise} \end{cases} ,$$

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

$$\textsc{ReqResCor}(\rho, j) = \max_{c \in [d]} \textsc{ReqResCor}_c(\rho, j) \enspace .$$

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

$$\textsc{CostReqRes}(\mathit{\Gamma}, \mathrm{Cst})(\rho) = \sup_{j \to \infty} \textsc{ReqResCor}(\rho, j) \enspace ,$$

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

A game $\mathcal{G} = (\mathcal{A}, \textsc{CostReqRes}(\mathit{\Gamma}, \mathrm{Cst}))$ is called a request-response game with costs. We denote the largest cost assigned to any edge by W. As we assume the functions Cst_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 Cst_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 2. Let $\mathcal{G} = (\mathcal{A}, \textsc{CostReqRes}(\mathit{\Gamma}, \mathrm{Cst}))$ 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 $(\mathcal{A}, \textsc{ReqRes}(\mathit{\Gamma}))$.

Using this remark and a detour via qualitative request-response games, we provide a cap for request-response games with costs.

Lemma 4. *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 in $\mathcal{G}$, then she also has a strategy with cost at most $d2^d nW$.*

Proof. Let $\mathcal{G} = (\mathcal{A}, \textsc{CostReqRes}(\mathit{\Gamma}, \mathrm{Cst}))$ and let $\mathcal{G}' = (\mathcal{A}, \textsc{ReqRes}(\mathit{\Gamma}))$ be a qualitative request-response game obtained by disregarding the cost functions of $\mathcal{G}$. Moreover, let σ be a strategy with finite cost for $\mathcal{G}$. Due to Remark 2, σ is winning for Player 0 in $\mathcal{G}'$ as well, i.e., Player 0 wins $\mathcal{G}'$. Thus, due to Theorem 6.2, she has a winning strategy σ' of size at most $d2^d$ in $\mathcal{G}'$. Let σ' be implemented by the memory structure $\mathcal{M}$ and let $b = d2^d nW$. We show $\mathrm{Cst}(\sigma') \leq b$.

Let $\rho = v_0 v_1 v_2 \cdots$ be a play consistent with σ' and assume towards a contradiction $\textsc{CostReqRes}(\mathit{\Gamma}, \mathrm{Cst})(\rho) > b$. Then there exist $c \in [d]$ and $j \in \mathbb{N}$

such that $\textsc{ReqResCor}_c(\rho, j) > b$. As each edge has cost at most W, the request for condition c opened at position j is not answered for at least $d2^dn$ steps, i.e., we obtain $v_{j'} \notin P_c$ for all j' with $j \le j' \le j + d2^dn$. Let $\mathrm{ext}(\rho) = (v_0, m_0)(v_1, m_1)(v_2, m_2)\cdots$. Since $|\mathcal{M}| \le d2^d$, there exists a vertex repetition on the play infix $(v_j, m_j)\cdots(v_{j+d2^dn}, m_{j+d2^dn})$ of $\mathrm{ext}(\rho)$, say at positions k and k' with $j \le k < k' \le j + d2^dn$. Then the play $\rho' = v_0 \cdots v_k(v_{k+1}\cdots v_{k'})^\omega$ is consistent with σ'.

In ρ', however, a request for condition c is opened at position $j \le k$. Since $j \le k' \le j + d2^dn$, this request is not answered in the play infix $v_j \cdots v_k \cdots v_{k'}$, i.e., it is never answered. Hence, $\rho' \notin \textsc{ReqRes}(\Gamma)$, which contradicts σ' being a winning strategy for Player 0 in $\mathcal{G}'$. □

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 [19].

Lemma 5. *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 $\mathcal{G} \le^{b+1}_{\mathcal{M}} \mathcal{G}'$ for $b = d2^dnW$, 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.*

Proof. Let $\mathcal{G} = (\mathcal{A}, \textsc{CostReqRes}(\Gamma, \mathrm{Cst}))$ with initial vertex v_I. Recall that $b = d2^dnW$, is a cap of $\mathcal{G}$ due to Lemma 4. We first define the memory structure $\mathcal{M}$. Intuitively, we use it to keep track of the currently open requests and the costs they have incurred up to the cap b. Once the cost of a single request incurs a cost greater than b, the memory structure raises a Boolean flag, which indicates that Player 1 can unbound the cost of that request.

Let $r\colon [d] \to \{\bot\} \cup [b+1] = \{\bot, 0, \ldots, b\}$ be a function mapping conditions c to the cost $r(c) \in [b+1]$ they have incurred so far, or to $r(c) = \bot$ if no request for that condition is pending. We call such a function a *request-function* and denote the set of all request functions by R. We define the initial request function r_I such that $r_I(c) = 0$ if $v_I \in Q_c$ and $r_I(c) = \bot$ otherwise. In order to be able to access the current vertex during the update of the memory structure, we store it in the memory structure as well. By accessing the current vertex together with the vertex that we move to, we are thus able to obtain the cost of the traversed edge. Finally, we store a flag that indicates whether or not the bound b has been exceeded. Hence, we define the set of memory states $M = V \times R \times \{0, 1\}$ with the initial memory state $m_I = (v_I, r_I, 0)$.

We define the update function $\mathrm{Upd}((v, r, f), v') = (v', r', f')$ by performing the following steps in order:

- For each $c \in [d]$, if $r(c) \neq \bot$, set $r'(c) = r(c) + \mathrm{Cst}_c((v, v'))$. Otherwise, set $r'(c) = \bot$.
- Now, if there exists a condition c such that $r'(c) > f$, then set $r'(c) = \bot$ for all c and set f' to 1. Otherwise, set f' to f.

- For each $c \in [d]$, if $v' \in Q_c$, set $r(c')$ to $\max\{r(c'), 0\}$ where $\bot < 0$.
- For each $c \in [d]$, if $v' \in P_c$, set $r(c')$ to $\bot$.

We obtain $\mathcal{M} = (M, m_I, \mathrm{Upd})$. Note that $|\mathcal{M}| \in \mathcal{O}(2nb^d)$, i.e., $\mathcal{M}$ is of exponential size in d, but only of polynomial size in n and W.

Using this definition, we obtain that if we have $\textsc{CostReqRes}(\Gamma, \mathrm{Cst})(\rho) \leq b$, then $\mathrm{ext}(\rho)$ remains in vertices of the form $(v, v, r, 0)$. Dually, if we have $\textsc{CostReqRes}(\Gamma, \mathrm{Cst})(\rho) > b$, then $\mathrm{ext}(\rho)$ eventually moves to vertices of the form $(v, v, r, 1)$ and remains there ad infinitum.

Let $\Gamma = (Q_c, P_c)_{c \in [d]}$. In order to obtain $\mathcal{G}'$, it remains to define the vertex-ranking function $\mathit{rank}\colon V \times \mathcal{M} \to \mathbb{N}$, as well as the family of request-response pairs Γ' for $\mathcal{G}'$. We define the former as $\mathit{rank}(v, v, r, 0) = \max_{c \in [d]} r(c)$ and $\mathit{rank}(v, v, r, 1) = b + 1$ and the latter as $\Gamma' = (Q'_c, P'_c)_{c \in [d]}$, where $Q'_c = Q_c \times Q_c \times R \times \{0, 1\}$ and $P'_c = P_c \times P_c \times R \times \{0, 1\}$ for all $c \in [d]$. Note that $\rho \in \textsc{ReqRes}(\Gamma)$ if and only if $\mathrm{ext}(\rho) \in \textsc{ReqRes}(\Gamma')$. Moreover, let $\rho = v_0 v_1 v_2$ be some play in $\mathcal{G}$ and let $\mathrm{ext}(\rho) = (v_0, v_0, r_0, f_0)(v_1, v_1, r_1, f_1)(v_2, v_2, r_2, f_2) \cdots$ be its extension. Note that the current vertex v_j is replicated in the memory state in order to be able to access it in the update of the memory state during the move to v_{j+1}, thereby attaining access to the traversed edge (v_j, v_{j+1}). If $\textsc{CostReqRes}(\Gamma, \mathrm{Cst})(\rho) \leq b$ and if, for some $c \in [d]$ and some $j \in \mathbb{N}$, we have $\textsc{ReqResCor}_c(\rho, j) = b'$, then $r_{j'}(c) = b'$, where j' is the earliest position at which the request for c opened at position j is answered. Dually, if $r_{j'}(c) = b'$ for some $j' \in \mathbb{N}$ and some $c \in [d]$, then $\textsc{ReqResCor}_c(\rho, j) = b'$, where j is the earliest position at which the request for condition c is opened without being answered prior to position j'.

We define $\mathcal{G}' = (\mathcal{A} \times \mathcal{M}, \textsc{Rank}^{\mathrm{sup}}(\textsc{ReqRes}(\Gamma'), \mathit{rank}))$. Moreover, since Γ' is the extension of Γ to the vertices of $\mathcal{A} \times \mathcal{M}$, $\mathcal{G}'$ contains d many request-response pairs.

It remains to show $\mathcal{G} \leq^{b+1}_{\mathcal{M}} \mathcal{G}'$. Recall that the $(b+1)$-correction-function $\cap_{b+1}$ is given implicitly. Clearly, the first and second condition of the definition of the quantitative reduction hold true, i.e., the arena of $\mathcal{G}'$ is $\mathcal{A} \times \mathcal{M}$ and cap_{b+1} is a $(b + 1)$-correction function. It remains to show the two latter conditions. To this end, let $\rho = v_0 v_1 v_2 \cdots \in V^\omega$ and let

$$\mathrm{ext}(\rho) = (v_0, v_0, r_0, f_0)(v_1, v_1, r_1, f_1)(v_2, v_2, r_2, f_2) \cdots \ .$$

We introduce the shorthands $\mathrm{Cst}_{\mathcal{G}} = \textsc{CostReqRes}(\Gamma, \mathrm{Cst})$ and $\mathrm{Cst}_{\mathcal{G}'} = \textsc{Rank}^{\mathrm{sup}}(\textsc{ReqRes}(\Gamma'), \mathit{rank})$.

We first show $\mathrm{Cst}_{\mathcal{G}}(\rho) = \mathrm{Cst}_{\mathcal{G}'}(\mathrm{ext}(\rho))$ for all ρ with $\mathrm{Cst}_{\mathcal{G}}(\rho) < b + 1$. Let $\mathrm{Cst}_{\mathcal{G}}(\rho) = b' < b + 1$ and note that this implies $\rho \in \textsc{ReqRes}(\Gamma)$ and $\mathrm{ext}(\rho) \in \textsc{ReqRes}(\Gamma')$. As argued above, we obtain $\mathit{rank}(v_j, v_j, r_j, f_j) \leq b'$ for all j, which implies $\mathrm{Cst}_{\mathcal{G}'}(\mathrm{ext}(\rho)) \leq b'$. Moreover, let $c \in [d]$ and $j \in \mathbb{N}$ such that $\textsc{ReqResCor}_c(\rho, j) = b'$. Since $b' < \infty$, such c and j exist. The play $\mathrm{ext}(\rho)$ visits a vertex of rank b' at the position at which the request for condition c opened at position j is answered for the first time. Thus, $\mathrm{Cst}_{\mathcal{G}'}(\mathrm{ext}(\rho)) \geq b'$, which concludes this part of the proof.

It remains to show that $\mathrm{Cst}_{\mathcal{G}'}(\mathrm{ext}(\rho)) \geq \mathrm{cap}_{b+1}(b+1) = b+1$ holds true for all ρ with $\mathrm{Cst}_{\mathcal{G}}(\rho) \geq b+1$. To this end, let $\mathrm{Cst}_{\mathcal{G}}(\rho) = b' \geq b+1$. As argued above, the extended play $\mathrm{ext}(\rho)$ eventually moves to vertices of the form $(v, v, r, 1)$ and remains there. Hence, $\mathrm{Cst}_{\mathcal{G}'}(\mathrm{ext}(\rho)) = b+1$ if $\rho \in \textsc{ReqRes}(\Gamma)$, i.e., if $\mathrm{ext}(\rho) \in \textsc{ReqRes}(\Gamma')$. If, however, $\rho \notin \textsc{ReqRes}(\Gamma)$, then $\mathrm{ext}(\rho) \notin \textsc{ReqRes}(\Gamma')$ and hence, $\mathrm{Cst}_{\mathcal{G}'}(\rho) = \infty > b+1$. □

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 4. 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 7. *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 $\mathrm{Cst}(\sigma) \leq b$ in $\mathcal{G}$?"*

Proof. Let $\mathcal{G}$ contain n vertices, d request-response pairs, and let W be the largest cost assigned to any edge. We first construct the vertex-ranked sup-request-response game $\mathcal{G}'$ as shown in Lemma 5. Recall that the game $\mathcal{G}'$ contains $\mathcal{O}(n(d2^d nW)^d)$ vertices and d request-response pairs. Due to the instantiation of Theorem 4 with the decision procedure for qualitative request-response games from Theorem 6.1, the game $\mathcal{G}'$ can be solved with respect to b in time $\mathcal{O}(n^3(d2^d nW)^{3d} d^2 2^d)$, which is exponential in the description length of $\mathcal{G}$, due to $W^{3d} \in \mathcal{O}((2^{|\mathcal{G}|})^{|\mathcal{G}|}) = \mathcal{O}(2^{|\mathcal{G}|^2})$. □

Moreover, request-response games are known to be ExpTime-hard [6]. 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 5, which has as many request-response pairs d as $\mathcal{G}$. Due to Theorem 6.2, 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 3, we obtain an exponential upper bound on the size of optimal strategies for Player 0.

Corollary 1. *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 Reducing Quantitative Muller Games to Vertex-Ranked Safety Games

Having shown how our framework can be used to find optimal strategies in request-response games with costs in a structured and modular way, we now show how it can be used to greatly simplify such existing proofs. To this end, we show how to reduce quantitative Muller games to vertex-ranked safety games [15].

Let $\mathcal{A}$ be some arena with vertex set V and recall the qualitative Muller condition, which is defined by a partition of 2^V into $(\mathcal{F}_0, \mathcal{F}_1)$ as $\textsc{Muller}(\mathcal{F}_0, \mathcal{F}_1) = \{\rho \in V^\omega \mid \inf(\rho) \in \mathcal{F}_0\}$, where $\inf(\rho)$ denotes the set of vertices that are visited infinitely often by ρ.

McNaughton introduced a quantitative characterization of the Muller condition by assigning a score to each prefix of a play and each subset of the set of vertices [14]. In order to characterize the infinity-set of a play, the score of a subset F measures how often F has been visited completely without leaving it. For a play ρ, the limes inferior of the score of $\inf(\rho)$ tends towards infinity, while the limes inferior of the score for all other sets is zero [14].

Formally, for any set $F \subseteq V$ with $F \neq \emptyset$, the score $\mathrm{Score}_F(\pi)$ is defined inductively using an accumulator that stores the vertices of F that have already been visited, as $(\mathrm{Acc}_F(\varepsilon), \mathrm{Score}_F(\varepsilon)) = (\emptyset, 0)$ and

$$(\mathrm{Acc}_F(\pi v), \mathrm{Score}_F(\pi v)) = \begin{cases} (\emptyset, 0) & \text{if } v \notin F \\ (\emptyset, \mathrm{Score}_F(\pi) + 1) & \text{if } \mathrm{Acc}_F(\pi) = F \setminus \{v\} \\ (\mathrm{Acc}_F(\pi) \cup \{v\}, \mathrm{Score}_F(\pi)) & \text{otherwise} \end{cases}$$

We generalize the score-function to families $\mathcal{F}$ of subsets of vertices, i.e., $\mathcal{F} \subseteq 2^V$, by defining $\mathrm{Score}_{\mathcal{F}}(\pi) = \max_{F \in \mathcal{F}}(\mathrm{Score}_F(\pi))$ and to infinite plays by defining $\mathrm{Score}_{\mathcal{F}}(v_0 v_1 v_2 \cdots) = \sup_{j \to \infty} \mathrm{Score}_{\mathcal{F}_1}(v_0 \cdots v_j)$. This definition inspires the quantitative Muller condition, which is defined as $\textsc{QuantMuller}(\mathcal{F}_0, \mathcal{F}_1)(\rho) = \mathrm{Score}_{\mathcal{F}_1}(\rho)$. We obtain a cap for such games using a result by Fearnley and Zimmermann [10].

Lemma 6. *Let $\mathcal{G} = (\mathcal{A}, \mathrm{Cst})$ be a quantitative Muller game. If Player 0 has a strategy with finite cost in $\mathcal{G}$, then she has a strategy σ with $\mathrm{Cst}(\sigma) \leq 2$.*

Proof. Let $\mathcal{G} = (\mathcal{A}, \textsc{QuantMuller}(\mathcal{F}_0, \mathcal{F}_1))$. Since $\mathrm{Cst}(\sigma) < \infty$, for every play ρ consistent with σ and every prefix π of ρ, we have that there exists an upper bound on $\mathrm{Score}_F(\pi)$ for all $F \in \mathcal{F}_1$. Moreover, as the score of $\inf(\rho)$ tends towards ∞, this implies $\inf(\rho) \in \mathcal{F}_0$, i.e., σ is a winning strategy for the qualitative Muller game $\mathcal{G}' = (\mathcal{A}, \textsc{Muller}(\mathcal{F}_0, \mathcal{F}_1))$.

Since Player 0 wins $\mathcal{G}'$, she has a strategy σ' with $\text{Score}_{\mathcal{F}_1}(\pi) \leq 2$ for all prefixes π of all plays consistent with σ' [10] . Thus $\text{QuantMuller}(\mathcal{F}_0, \mathcal{F}_1)(\sigma') \leq 2$ indeed holds true. □

We now show how to reduce quantitative Muller games to vertex-ranked sup-safety games based on previous work by Neider et al. [15]. Recall that a safety game is a very simple qualitative game, in which it is Player 0's goal to avoid a certain set of undesirable vertices. The constructed safety game uses play prefixes of cost at most 3 as vertices and mimics plays π of cost at most 3 in the Muller game by moving to some vertex π' such that the score and the accumulator are equal in π and π' for all $F \in \mathcal{F}_1$. We show how to lift this qualitative construction to the setting of quantitative games by providing a quantitative reduction from quantitative Muller games to vertex-ranked sup-safety games.

Lemma 7. *Let $\mathcal{G}$ be a quantitative Muller game with vertex set V. There exists a memory structure $\mathcal{M}$ of size at most $(|V|!)^3$ and a vertex-ranked* sup*-safety game $\mathcal{G}'$ such that $\mathcal{G} \leq^3_{\mathcal{M}} \mathcal{G}'$.*

Proof. Let $\mathcal{G} = (\mathcal{A}, \text{Cst})$ with vertex set V and $\text{Cst} = \text{QuantMuller}(\mathcal{F}_0, \mathcal{F}_1)$. We say that two play prefixes π and π' are $\mathcal{F}_1$-equivalent if they end in the same vertex and if, for each $F \in \mathcal{F}_1$, we have $\text{Acc}_F(\pi) = \text{Acc}_F(\pi')$ and $\text{Score}_F(\pi) = \text{Score}_F(\pi')$. In this case, we write $\pi \approx_{\mathcal{F}_1} \pi'$. For each play π, we denote the $\mathcal{F}_1$-equivalence-class of π by $[\pi]_{\approx_{\mathcal{F}_1}}$. Let $\text{Plays}_{\leq 2} = \{\pi \in V^* \mid \text{Score}_{\mathcal{F}_1}(\pi) \leq 2\}$. For the sake of readability, we omit the $\mathcal{F}_1$ for the remainder of this proof wherever unambiguously possible.

We define the set of memory states $M = \text{Plays}_{\leq 2}/\approx \cup \{\bot\}$, the initial memory state $m_I = [v_I]$, and the update function Upd as $\text{Upd}(\bot, v) = \bot$, $\text{Upd}(\pi, v) = [\pi v]$ if $\text{Score}_{\mathcal{F}_1}(\pi v) \leq 2$ and $\text{Upd}(\pi, v) = \bot$ otherwise. We obtain $|M| \in \mathcal{O}(|\text{Plays}_{\leq 2}/\approx|)$. Since $|\text{Plays}_{\leq 2}/\approx| \leq (|V|!)^3$, the memory structure $\mathcal{M}$ is indeed of exponential size in $|V|$ [15].

A straightforward induction shows that this memory structure tracks the score of a play precisely as long as it does not exceed the value two on any prefix. More formally, it satisfies the following invariant:

> Let $\pi = v_0 \cdots v_j$ be a play prefix in $\mathcal{G}$ such that $\text{Score}_{\mathcal{F}_1}(v_0 \cdots v_k) \leq 2$ for all k with $0 \leq k \leq j$. Moreover, let $\text{Upd}^+(\pi) = \pi'$. Then $\pi \approx \pi'$.

Recall $\text{Safety}(S) = \{v_0 v_1 v_2 \cdots \in V^\omega \mid \forall j \in \mathbb{N}.\, v_j \notin S\}$. We define the vertex-ranked sup-safety game $\mathcal{G}' = (\mathcal{A} \times \mathcal{M}, \text{Rank}^{\text{sup}}(\text{Safety}(V \times \{\bot\}), \mathit{rank}))$, where $\mathit{rank}(v, \pi) = \text{Score}_{\mathcal{F}_1}(\pi)$ for all $\pi \in \text{Plays}_{\leq 2}$, and $\mathit{rank}(v, \bot) = 3$.

Let $\text{Cst}' = \text{Rank}^{\text{sup}}(\text{Safety}(V \times \{\bot\}), \mathit{rank})$. Clearly, the first two items of the definition of $\mathcal{G} \leq^3_{\mathcal{M}} \mathcal{G}'$ hold true. Moreover, $\mathcal{G}'$ is indeed of exponential size in $|\mathcal{G}|$, as argued above. It remains to show $\text{Cst}(\rho) = \text{Cst}'(\text{ext}(\rho))$ for all ρ with $\text{Cst}(\rho) < 3$ and $\text{Cst}'(\text{ext}(\rho)) \geq 3$ for all other ρ.

First, let $\rho = v_0 v_1 v_2 \cdots$ be some play with $\text{Cst}(\rho) \leq 2$ and let $\text{ext}(\rho) = (v_0, m_0)(v_1, m_1)(v_2, m_2) \cdots$. Then $\text{Score}_{\mathcal{F}_1}(v_0 \cdots v_j) \leq 2$ for all $j \in \mathbb{N}$. Thus, due to the invariant above and the definition of rank, we obtain $\mathit{rank}(v_j, m_j) = \text{Score}(v_0 \cdots v_j)$ for all $j \in \mathbb{N}$, which implies $\text{Cst}'(\text{ext}(\rho)) = \text{Cst}(\rho)$.

Towards a proof of the latter statement, let $\rho = v_0 v_1 v_2 \cdots$ be a play such that $\mathrm{Cst}(\rho) \geq 3$ and let j be the minimal position such that $\mathrm{Cst}(v_0 \cdots v_j) = 3$. Since $\mathrm{Cst}(v_0 \cdots v_j) = \mathrm{Score}_F(v_0 \cdots v_j)$ for some $F \in \mathcal{F}_1$ and since the score is at most incremented during each step, we obtain $\mathrm{Score}_F(v_0 \cdots v_{j-1}) = 2$ and $\mathrm{Acc}_F(v_0 \cdots v_{j-1}) = F \setminus \{v_j\}$. Let $\mathrm{Upd}^+(v_0 \cdots v_{j-1}) = \pi'$. Due to the invariant we obtain $\mathrm{Score}_F(\pi') = 2$ and $\mathrm{Acc}_F(\pi') = F \setminus \{v_j\}$. Thus, $\mathrm{Score}_F(\pi' v_j) = 3$, hence $\mathrm{ext}(v_0 \cdots v_j) = (v_0, m_I) \cdots (v_j, \bot)$, which implies $\mathrm{ext}(\rho) \notin \textsc{Safety}(V \times \{\bot\})$, which in turn yields $\mathrm{Cst}'(\mathrm{ext}(\rho)) = \infty > 3$. □

Thus, in order to solve a quantitative Muller game with respect to some b, it suffices to solve a vertex-ranked sup-safety game $\mathcal{G}'$ with respect to b. Recall that this is only constructive if Player 0 wins $\mathcal{G}'$ with respect to $b < 3$, i.e., only in this case are we able to construct a strategy with cost at most b for her in $\mathcal{G}$. Otherwise, Theorem 2 yields the existence of a strategy of cost ∞ for Player 1 in $\mathcal{G}$, but we cannot construct such a strategy from his strategy of cost greater than two in $\mathcal{G}'$. This is consistent with results of Neider et al. [15] and with the fact that Muller conditions are in a higher level of the Borel hierarchy than safety conditions, i.e., qualitative Muller games cannot be reduced to safety games.

We can, however, solve the resulting vertex-ranked safety game with respect to a given bound by solving a qualitative safety game as shown in Theorem 4. Using this reduction together with the framework of quality-preserving reductions, we obtain an upper bound on the complexity of solving quantitative Muller games with respect to some bound b.

Theorem 8. *The following problem is in* ExpTime*: "Given some quantitative Muller game $\mathcal{G}$ and some bound $b \in \mathbb{N}$, does Player 0 win $\mathcal{G}$ with respect to b?"*

Proof. Given $\mathcal{G}$, we first construct the vertex-ranked sup-safety game $\mathcal{G}'$ as shown in Lemma 7. Recall that $\mathcal{G}'$ contains at most $(|V|!)^3$ vertices. Due to Theorem 4 and the fact that safety games can be solved in linear time in the number of vertices, $\mathcal{G}'$ can indeed be solved in time at most $(|V|!)^3$ with respect to a given bound b, which is exponential in the number of vertices $|V|$ of $\mathcal{G}$. □

As discussed in Section 4.1, since both players have positional winning strategies in safety games, Lemma 7 yields that if Player 0 has a strategy with cost at most 3 in a quantitative Muller game $\mathcal{G}$, then she also has a strategy in $\mathcal{G}$ with the same cost and of size at most exponential in $|V|$. Moreover, finding the minimal b such that Player 0 has a strategy of cost at most b in $\mathcal{G}$ requires solving at most three safety games of exponential size in $|V|$. Thus, the optimization problem for quantitative Muller games can be solved in exponential time.

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, moreover, that it retains the quality of strategies.

Additionally, we have provided two very general types of quantitative games that serve as targets for quantitative reductions, i.e., vertex-ranked sup games and vertex-ranked lim sup-games. For both kinds of games we have provided tight bounds 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 usefulness of quantitative reductions and vertex-ranked games by providing reductions from quantitative request-response games to vertex-ranked request-response games and from quantitative Muller games to vertex-ranked safety games. Thereby, we have obtained tight bounds on the complexity of solving the former games optimally and an upper bound on the complexity of solving the latter games optimally. For quantitative request-response games, no such bound was known previously. Both proofs show that our framework enables well-structured and modular analyses of quantitative games.

The reduction of request-response games with costs to vertex-ranked request-response games yields a tight upper bound on the runtime complexity of solving such games optimally. The reduction of Muller games to safety games, however, shows that an approach via reductions may not always yield optimal runtime complexity. Consider also, for example, the problem of solving parity games with costs with respect to some bound, which is PSpace-complete [19]. 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.

Finally, 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 [13]. To the best of our knowledge, it is open how to define such a hierarchy for quantitative winning conditions which exhibits similar properties.

References

1. Bouyer, P., Markey, N., Randour, M., Larsen, K.G., Laursen, S.: Average-energy games. Acta Informatica (2016)
2. Bruyère, V., Hautem, Q., Randour, M.: Window parity games: an alternative approach toward parity games with time bounds. In: Cantone, D., Delzanno, G. (eds.) GandALF. EPTCS, vol. 226, pp. 135–148 (2016)
3. Chatterjee, K., Doyen, L.: Energy parity games. TCS 458, 49–60 (2012)

4. Chatterjee, K., Doyen, L., Henzinger, T.A., Raskin, J.F.: Generalized Mean-payoff and Energy Games. In: Lodaya, K., Mahajan, M. (eds.) FSTTCS. LIPIcs, vol. 8, pp. 505–516 (2010)
5. Chatterjee, K., Henzinger, T.A., Horn, F.: Finitary winning in ω-regular games. ACM Trans. Comput. Log. 11(1) (2009)
6. Chatterjee, K., Henzinger, T.A., Horn, F.: The complexity of request-response games. In: Language and Automata Theory and Applications, pp. 227–237. Springer Nature (2011)
7. Chatterjee, K., Henzinger, T.A., Piterman, N.: Algorithms for Büchi games. In: GDV (2006)
8. Ehrenfeucht, A., Mycielski, J.: Positional strategies for mean payoff games. Internat. J. Game Theory 8, 109–113 (1979)
9. Faymonville, P., Zimmermann, M.: Parametric linear dynamic logic. In: Peron, A., Piazza, C. (eds.) GandALF. EPTCS, vol. 161, pp. 60–73 (2014)
10. Fearnley, J., Zimmermann, M.: Playing Muller Games in a Hurry. Internat. J. Found. Comput. Sci. 23, 649–668 (2012), special Issue GandALF 2010.
11. Fijalkow, N., Zimmermann, M.: Parity and Streett Games with Costs. LMCS 10(2) (2014)
12. Kupferman, O., Piterman, N., Vardi, M.Y.: From liveness to promptness. Formal Methods in System Design 34(2), 83–103 (2009)
13. Martin, D.A.: Borel determinacy. Annals of Mathematics 102, 363–371 (1975)
14. McNaughton, R.: Playing infinite games in finite time. In: Salomaa, A., Wood, D., Yu, S. (eds.) A Half-Century of Automata Theory. pp. 73–91. World Scientific (2000)
15. Neider, D., Rabinovich, R., Zimmermann, M.: Down the Borel hierarchy: Solving Muller games via safety games. TCS 560(3), 219 – 234 (2014), special Issue GandALF 2012
16. Nerode, A., Remmel, J.B., Yakhnis, A.: McNaughton games and extracting strategies for concurrent programs. Ann. Pure Appl. Logic 78(1-3), 203–242 (1996)
17. Velner, Y., Chatterjee, K., Doyen, L., Henzinger, T.A., Rabinovich, A., Raskin, J.F.: The complexity of multi-mean-payoff and multi-energy games. Inf. and Comp. 241, 177 – 196 (2015)
18. Wallmeier, N., Hütten, P., Thomas, W.: Symbolic synthesis of finite-state controllers for request-response specifications. In: Ibarra, O.H., Dang, Z. (eds.) CIAA. LNCS, vol. 2759, pp. 11–22. Springer (2003)
19. Weinert, A., Zimmermann, M.: Easy to Win, Hard to Master: Optimal Strategies in Parity Games with Costs. In: Talbot, J.M., Regnier, L. (eds.) CSL '16. LIPIcs, vol. 62, pp. 31:1–31:17 (2016)
20. Zwick, U., Paterson, M.: The complexity of mean payoff games on graphs. TCS 158(1-2), 343–359 (1996)