Fixed-Dimensional Energy Games are in Pseudo-Polynomial Time *

Marcin Jurdziński[1], Ranko Lazić[1], and Sylvain Schmitz[1,2]

[1] DIMAP, Department of Computer Science, University of Warwick, UK
[2] LSV, ENS Cachan & CNRS & INRIA, France

Abstract. We generalise the hyperplane separation technique (Chatterjee and Velner, 2013) from multi-dimensional mean-payoff to energy games, and achieve an algorithm for solving the latter whose running time is exponential only in the dimension, but not in the number of vertices of the game graph. This answers an open question whether energy games with arbitrary initial credit can be solved in pseudo-polynomial time for fixed dimensions 3 or larger (Chaloupka, 2013). It also improves the complexity of solving multi-dimensional energy games with given initial credit from non-elementary (Brázdil, Jančar, and Kučera, 2010) to 2EXPTIME, thus establishing their 2EXPTIME-completeness.

1 Introduction

Multi-Dimensional Energy Games are played turn-by-turn by two players on a finite *multi-weighted* game graph, whose edges are labelled with integer vectors modelling discrete energy consumption and refuelling. Player 1's objective is to keep the accumulated energy non-negative in every component along infinite plays. This setting is relevant to the synthesis of resource-sensitive controllers balancing the usage of various resources like fuel, time, money, or items in stock, and finding optimal trade-offs; see [4, 10, 3, 11] for some examples. Maybe more importantly, energy games are the key ingredient in the study of several related resource-conscious games, notably multi-dimensional mean-payoff games [6] and games played on vector addition systems with states (VASS) [4, 2, 9].

The main open problem about these games has been to pinpoint the complexity of deciding whether Player 1 has a winning strategy when starting from a particular vertex and given an initial energy vector as part of the input. This particular *given initial credit* variant of energy games is also known as *Z-reachability* VASS games [4, 5]. The problem is also equivalent via logarithmic-space reductions to deciding *single-sided* VASS games with a non-termination objective [2], and to deciding whether a given VASS (or, equivalently, a Petri net) simulates a given finite state system [9, 1]. As shown by Brázdil, Jančar, and Kučera [4], all these problems can be solved in $(d-1)$EXPTIME where $d \geq 2$ is the number

* Work funded in part by the ANR grant 11-BS02-001-01 ReacHard, the Leverhulme Trust Visiting Professorship VP1-2014-041, and the EPSRC grant EP/M011801/1.

of energy components, i.e. a TOWER of exponentials when d is part of the input. The best known lower bound for this problem is 2EXPTIME-hardness [9], leaving a substantial complexity gap. So far, the only tight complexity bounds are for the case $d = 2$: Chaloupka [5] shows the problem to be PTIME-complete when using unit updates, i.e. when the energy levels can only vary by -1, 0, or 1. However, quoting Chaloupka, 'since the presented results about 2-dimensional VASS are relatively complicated, we suspect this [general] problem is difficult.'

When inspecting the upper bound proof of Brázdil et al. [4], it turns out that the main obstacle to closing the gap and proving 2EXPTIME-completeness lies in the complexity upper bounds for energy games with an *arbitrary initial credit*—which is actually the variant commonly assumed when talking about energy games. Given a multi-weighted game graph and an initial vertex v, we now wish to decide whether there exists an initial energy vector $\mathbf{b}$ such that Player 1 has a winning strategy starting from the pair $(v, \mathbf{b})$. As shown by Chatterjee, Doyen, Henzinger, and Raskin [6], this variant is simpler: it is coNP-complete. However, the parameterised complexity bounds in the literature [4, 7] for this simpler problem involve an exponential dependency on the number $|V|$ of vertices in the input game graph, which translates into a tower of exponentials when solving the given initial credit variant.

Contributions. We show in this paper that the arbitrary initial credit problem for d-dimensional energy games can be solved in time $O(|V| \cdot \|E\|)^{O(d^4)}$ where $|V|$ is the number of vertices of the input multi-weighted game graph and $\|E\|$ the maximal value that labels its edges, i.e. in pseudo-polynomial time (see Thm. 3.3). We then deduce that the given initial credit problem for general multi-dimensional energy games is 2EXPTIME-complete, and also in pseudo-polynomial time when the dimension is fixed (see Thm. 3.5), thus closing the gap left open in [4, 9]. Our parameterised bounds are of practical interest because typical instances of energy games would have small dimension but might have a large number of vertices. By the results of Chatterjee et al. [6], another consequence is that we can decide the existence of a *finite-memory* winning strategy for fixed-dimensional *mean-payoff* games in pseudo-polynomial time. The existence of a finite-memory winning strategy is the most relevant problem for controller synthesis, but until now, solving fixed-dimensional mean-payoff games in pseudo-polynomial time required infinite memory strategies [8].

Overview. We prove our upper bounds on the complexity of the arbitrary initial credit problem for d-dimensional energy games by reducing them to *bounding games*, where Player 1 additionally seeks to prevent arbitrarily high energy levels (Sec. 2.3). We further show these games to be equivalent to *first-cycle bounding games* in Sec. 5, where the total effect of the first simple cycle defined by the two players determines the winner. More precisely, first-cycle bounding games rely on a hierarchically-defined colouring of the game graph by *perfect half-spaces* (see Sec. 4), and the two players strive respectively to avoid or produce cycles in those perfect half-spaces.

First-cycle bounding games coloured with perfect half-spaces can be seen as generalising quite significantly both the 'local strategy' approach of Chaloupka

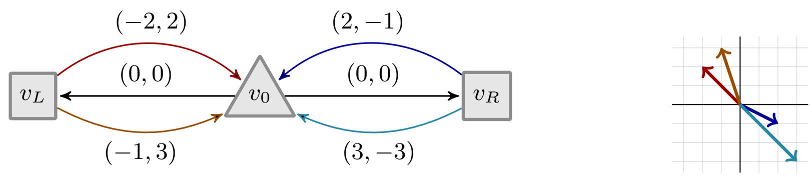

Fig. 1: A 2-dimensional multi-weighted game graph.

[5] for 2-dimensional energy games, and the 'separating hyperplane technique' of Chatterjee and Velner [8] for multi-dimensional mean-payoff games.

The reduction to first-cycle bounding games has several important corollaries: the *determinacy* of bounding games, and the existence of a *small hypercube property*, which in turn allow to derive the announced complexity bounds on energy games (see Sec. 3). In fact, we found with first-cycle bounding games a highly versatile tool, which we use extensively in our proofs on energy games.

We start by presenting the necessary background on energy and bounding games in Sec. 2. Some omitted material can be found in the full paper available from http://arxiv.org/abs/1502.06875.

2 Multi-Weighted Games

We define in this section the various games we consider in this work. We start by defining multi-weighted game graphs, which provide a finite representation for the infinite arenas over which our games are played. We then define energy games in Sec. 2.2, and their generalisation as bounding games in Sec. 2.3.

2.1 Multi-Weighted Game Graphs

We consider game graphs whose edges are labelled by vectors of integers. They are tuples of the form (V, E, d), where d is the dimension in $\mathbb{N}$, $V \stackrel{\text{def}}{=} V_1 \uplus V_2$ is a finite set of vertices, which is partitioned into Player 1 vertices (V_1) and Player 2 vertices (V_2), and E is a finite set of edges included in $V \times \mathbb{Z}^d \times V$, and such that every vertex has at least one outgoing edge; we call the labels in $\mathbb{Z}^d$ 'weights'.

Example 2.1. Figure 1 shows on its left-hand-side an example of a 2-dimensional multi-weighted game graph. Throughout this paper, Player 1 vertices are depicted as triangles and Player 2 vertices as squares.

Norms. For a vector $\mathbf{a}$, we denote the maximum absolute value of its entries by $\|\mathbf{a}\| \stackrel{\text{def}}{=} \max_{1\leq i\leq d} |\mathbf{a}(i)|$, and we call it the *norm* of $\mathbf{a}$. By extension, for a set of edges E, we let $\|E\| \stackrel{\text{def}}{=} \max_{(v,\mathbf{u},v')\in E} \|\mathbf{u}\|$. We assume, without loss of generality, that $\|E\| > 0$ in our multi-weighted game graphs. Regarding complexity, we encode vectors of integers in binary, hence $\|E\|$ may be exponential in the size of the multi-weighted game graph.

Paths and Cycles. Given a multi-weighted game graph (V, E, d), a *configuration* is a pair $(v, \mathbf{a})$ with v in V and $\mathbf{a}$ in $\mathbb{Z}^d$. A *path* is a finite sequence of configurations

$\pi = (v_0, \mathbf{a}_0)(v_1, \mathbf{a}_1) \cdots (v_n, \mathbf{a}_n)$ in $(V \times \mathbb{Z}^d)^*$ such that for every $0 \leq j < n$ there exists an edge $(v_j, \mathbf{a}_{j+1} - \mathbf{a}_j, v_{j+1})$ in E (where addition is performed componentwise). The *total weight* of such a path π is $w(\pi) \stackrel{\text{def}}{=} \sum_{0 \leq j < n} \mathbf{a}_{j+1} - \mathbf{a}_j = \mathbf{a}_n - \mathbf{a}_0$. A *cycle* is a path $(v_0, \mathbf{a}_0)(v_1, \mathbf{a}_1) \cdots (v_n, \mathbf{a}_n)$ with $v_0 = v_n$. Such a cycle is *simple* if $v_j = v_k$ for some $0 \leq j < k \leq n$ implies $j = 0$ and $k = n$. We assume, without loss of generality, that every cycle contains at least one Player 1 vertex. We often identify simple cycles with their respective weights; the weights of the four simple cycles of the game graph in Fig. 1 are displayed on its right-hand-side.

Proposition 2.2. *In any game graph (V, E, d), the total weight of any simple cycle has norm at most $|V| \cdot \|E\|$.*

Plays and Strategies. Let v_0 be a vertex from V. A *play from* v_0 is an infinite configuration sequence $\rho = (v_0, \mathbf{a}_0)(v_1, \mathbf{a}_1) \cdots$ such that $\mathbf{a}_0 = \mathbf{0}$ is the null vector and every finite prefix $\rho|_n \stackrel{\text{def}}{=} (v_0, \mathbf{a}_0) \cdots (v_n, \mathbf{a}_n)$ is a path. Note that, because $\mathbf{a}_0 = \mathbf{0}$, the total weight of this prefix is $w(\rho|_n) = \mathbf{a}_n$. We define the *norm* of a play ρ as the supremum of the norms of total weights of its prefixes: $\|\rho\| \stackrel{\text{def}}{=} \sup_n \|w(\rho|_n)\|$. A *strategy* for Player p, $p \in \{1, 2\}$, is a function σ_p taking as input a non-empty path $\pi \cdot (v, \mathbf{a})$ ending in a Player p vertex $v \in V_p$, and returning an edge $\sigma_p(\pi \cdot (v, \mathbf{a})) = (v, \mathbf{u}, v')$ from E. We employ the usual notions of plays *consistent* with strategies, and given some winning condition on plays, of *winning* strategies for a player.

Example 2.1 (continued). For instance, in the game graph depicted in Fig. 1, a strategy for Player 1 could be to move to v_L whenever the current energy level on the first coordinate is non-negative, and to v_R otherwise—note that this is an *infinite-memory* strategy—:

$$\sigma_1(\pi \cdot (v_0, \mathbf{a})) \stackrel{\text{def}}{=} \begin{cases} (v_0, (0,0), v_L) & \text{if } \mathbf{a}(1) \geq 0, \\ (v_0, (0,0), v_R) & \text{otherwise,} \end{cases} \tag{1}$$

and one for Player 2 could be to always select one particular edge in every vertex, regardless of the current energy vector—this is called a *counterless* strategy [4]—:

$$\sigma_2(\pi \cdot (v, \mathbf{a})) \stackrel{\text{def}}{=} \begin{cases} (v_L, (-2,2), v_0) & \text{if } v = v_L \\ (v_R, (2,-1), v_0) & \text{otherwise.} \end{cases} \tag{2}$$

2.2 Multi-Dimensional Energy Games

Suppose (V, E, d) is a multi-weighted game graph, v_0 an initial vertex, and $\mathbf{b}$ is a vector from $\mathbb{N}^d$. A play ρ from v_0 is *winning* for Player 1 in the *energy game* $\Delta_{\mathbf{b}}(V, E, d)$ with *initial credit* $\mathbf{b}$ if, for all n, $\mathbf{b} + w(\rho|_n) \geq \mathbf{0}$, using the product ordering over $\mathbb{Z}^d$. Otherwise, Player 2 wins the play. An immediate property of energy games is *monotonicity*: if σ_1 is winning for Player 1 with some initial credit $\mathbf{b}$, and $\mathbf{b}' \geq \mathbf{b}$, then it is also winning for Player 1 with initial credit $\mathbf{b}'$.

Example 2.1 (continued). For example, one may observe that the strategy (1) for Player 1 is winning for the game graph of Fig. 1 with initial credit $(2, 2)$ (or larger). A geometric intuition comes from the directions of the total weights of simple cycles in Fig. 1: by choosing alternatively edges to v_L or v_R, Player 1 is able to balance the energy levels above the '$x + y = 0$' line.

2.3 Multi-Dimensional Bounding Games

A generalisation of energy games sometimes considered in the literature is to further impose a maximal *capacity* $\mathbf{c} \in \mathbb{N}^d$ (also called an upper bound) on the energy levels during the play [10, 11]. Player 1 then wins a play ρ if $0 \leq \mathbf{b} + w(\rho|_n) \leq \mathbf{c}$ for all n.

In the spirit of the arbitrary initial credit variant of energy games, we also quantify $\mathbf{c}$ existentially. This defines the *bounding game* $\Gamma(V, E, d)$ over a multi-weighted game graph (V, E, d), where a play ρ is winning for Player 1 if its norm $\|\rho\|$ is finite, i.e. if the set $\{\|w(\rho|_n)\| : n \in \mathbb{N}\}$ of norms of total weights of all finite prefixes of ρ is bounded, and Player 2 wins otherwise, if it is unbounded. In other words, Player 1 strives to contain the current vector within some d-dimensional hypercube, while Player 2 attempts to escape.

Example 2.1 (continued). Note that Player 2 is now winning the bounding game defined by the game graph of Fig. 1 from any of the three vertices, for example using the strategy (2). Indeed, this strategy ensures that the only simple cycles that can be played have weights $(-2, 2)$ and $(2, -1)$. Because these vectors belong to an open half-plane, the total energy will drift deeper and deeper inside that open half-plane and its norm will grow unbounded.

Lossy Game Graphs. If Player 1 wins the bounding game, then there exists some initial credit for which she also wins the energy game. For a converse, let $\mathsf{lossy}(V, E, d)$ denote the *lossy* multi-weighted game graph obtained from (V, E, d) by inserting, at each Player 1 vertex and for each $1 \leq i \leq d$, a self-loop labelled by the negative unit vector $-\mathbf{e}_i$. In a bounding game played over a lossy game graph, it turns out that Player 1 can always bound the current vector from above by playing these unit decrements, hence she only has to ensure that the current vector remains bounded from below, i.e. she has to win the energy game for some initial credit. Formally (see the full paper for a proof):

Proposition 2.2. *From any vertex in any multi-weighted game graph (V, E, d):*
1. *Player 1 wins the energy game $\Delta_\mathbf{b}(V, E, d)$ for some $\mathbf{b} \in \mathbb{N}^d$ if and only if Player 1 wins the bounding game $\Gamma(\mathsf{lossy}(V, E, d))$.*
2. *Player 2 wins the energy game $\Delta_\mathbf{b}(V, E, d)$ for all $\mathbf{b} \in \mathbb{N}^d$ if and only if Player 2 wins the bounding game $\Gamma(\mathsf{lossy}(V, E, d))$.*

Our task in the following will be therefore to prove an upper bound on the time complexity required to solve bounding games.

Example 2.3. By Prop. 2.2, because she was winning the energy game of Fig. 1 with initial credit $(2, 2)$, Player 1 is now winning the bounding game played on the lossy version of the multi-weighted game graph of Fig. 1.

Example 2.4. As a rather different example, consider the multi-weighted game graph of Fig. 2. Although Player 2 does not control any vertex, and Player 1 controls the 'direction of divergence', Player 2 wins the associated bounding

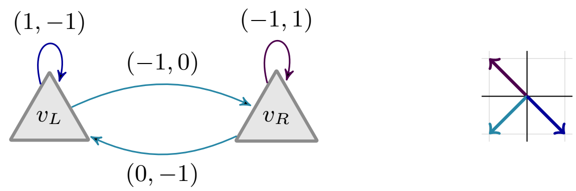

Fig. 2: A 2-dimensional game graph with only Player 1 vertices.

game. Indeed, Player 1 can either eventually stay forever at one of the two vertices, or visit both vertices infinitely often. In any case, she loses.

3 Complexity Upper Bounds

Our main results are new parameterised complexity upper bounds for deciding whether Player 1 has a winning strategy in a given energy game. In turn, we rely for these results on a *small hypercube property* of bounding games, which we introduce next, and which will be a consequence of the study of first-cycle bounding games in Sec. 5.

3.1 Small Hypercube Property

In a bounding game, if Player 1 is winning, then by definition she has a winning strategy σ_1 such that for all plays ρ consistent with σ_1 there exists some bound B_ρ with $\|\rho\| \leq B_\rho$. We considerably strengthen this statement in Sec. 5 where we construct an explicit winning strategy, which yields an explicit *uniform* bound B for all consistent plays:

Lemma 3.1. *Let (V, E, d) be a multi-weighted game graph. If Player 1 wins the bounding game $\Gamma(V, E, d)$, then she has a winning strategy which ensures $\|\rho\| \leq (4|V| \cdot \|E\|)^{2(d+2)^3}$ for all consistent plays ρ.*

Note that our bound is polynomial in $|V|$ the number of vertices, unlike the bounds found in comparable statements by Brázdil et al. [4], Lem. 7 and Chatterjee et al. [7], Lem. 3, which incur an exponential dependence on $|V|$. This entails pseudo polynomial complexity bounds when d is fixed:

Corollary 3.2. *Bounding games on multi-weighted graphs (V, E, d) are solvable in time $(|V| \cdot \|E\|)^{O(d^4)}$.*

Proof. By Lem. 3.1, the bounding game is equivalent to a reachability game where Player 2 attempts to see the norm of the total weight exceed $B \stackrel{\text{def}}{=} (4|V| \cdot \|E\|)^{2(d+2)^3}$. This can be played within a finite arena of size $(2B+1)^d$ and solved in time linear in that size using the usual attractor computation algorithm. □

3.2 Energy Games with Arbitrary Initial Credit

The *arbitrary initial credit problem* for energy games takes as input a multi-weighted game graph and an initial vertex v_0 and asks whether there exists a vector $\mathbf{b}$ in $\mathbb{N}^d$ such that Player 1 wins $\Delta_{\mathbf{b}}(V, E, d)$ from v_0:

Theorem 3.3. *The arbitrary initial credit problem for energy games on multi-weighted game graphs* (V, E, d) *is solvable in time* $(|V| \cdot \|E\|)^{O(d^4)}$.

Proof. This follows from Prop. 2.2, and Cor. 3.2 applied to the game graph $\Gamma(\mathsf{lossy}(V, E, d))$. □

3.3 Energy Games with Given Initial Credit

The *given initial credit problem* for energy games takes as input a multi-weighted game graph (V, E, d), an initial vertex v_0, and a credit $\mathbf{b}$ in $\mathbb{N}^d$ and asks whether Player 1 wins the energy game $\Delta_{\mathbf{b}}(V, E, d)$ from v_0. Thanks to Lem. 3.1, a proof of the upcoming Thm. 3.5 could be obtained using the work of Brázdil et al. [4], and more generally the techniques of Rackoff [12]. As usual in this work, we rather proceed by transferring that setting to that of bounding games (and thus to that of first-cycle bounding games). Our key lemma shows that any energy game with a given initial credit played over some multi-weighted game graph is equivalent to some bounding game played over a double-exponentially larger game graph:

Lemma 3.4. *Let* $\mathbf{b} \in \mathbb{N}^d$, (V, E, d) *be a multi-weighted game graph, and* $v \in V$. *One can construct in time* $O(|V'| \cdot |E| + d \cdot \log \|\mathbf{b}\|)$ *a multi-weighted game graph* (V', E', d) *and a vertex* $v_{\mathbf{b}}$ *in* V' *with* $|V'| \leq (4|V| \cdot \|E\|)^{2^d (d+3)^{3d}}$ *and* $\|E'\| = \|E\|$ *s.t., for all* $p \in \{1, 2\}$, *Player* p *wins the energy game* $\Delta_{\mathbf{b}}(V, E, d)$ *from* v *iff Player* p *wins the bounding game* $\Gamma(V', E', d)$ *from* $v_{\mathbf{b}}$.

By applying Cor. 3.2 to the game graph (V', E', d) and since $|E| \leq |V|^2 \cdot \|E\|^d$, we obtain a $\mathsf{2EXPTIME}$ upper bound on the given initial credit problem, which is again pseudo-polynomial when d is fixed:

Theorem 3.5. *The given initial credit problem with credit* $\mathbf{b}$ *for energy games on multi-weighted game graphs* (V, E, d) *is solvable in time* $O(|V| \cdot \|E\|)^{2^{O(d \cdot \log d)}} + O(d \cdot \log \|\mathbf{b}\|)$.

This matches the $\mathsf{2EXPTIME}$ lower bound from [9], and encompasses Chaloupka's PTIME upper bound in dimension $d = 2$ with unit updates, i.e. with $\|E\| = 1$. Because the given initial credit problem for energy games of fixed dimension $d \geq 4$ is $\mathsf{EXPTIME}$-hard [9], the bound in terms of $\|E\|$ in Thm. 3.5 cannot be improved.

4 Perfect Half-Spaces

We recall in this section the definition of subsets of $\mathbb{Q}^d$ called *perfect half-spaces*. They will be used next in Sec. 5 to define a condition for Player 2 to win bounding games, which relies on Player 2's ability to force cycles inside perfect half-spaces. This can be understood as a generalisation of Chatterjee and Velner's approach for solving multi-dimensional *mean-payoff* games [8], which as we recall in the full paper relies on a similar ability to force cycles inside open half-spaces. We employ perfect half-spaces in Sec. 5 to colour the edges in *first-cycle bounding games*, which determine the winner using both the colours and the weight of the first cycle formed along a play.

4.1 Definitions from Linear Algebra

Given a subset $\mathbf{A}$ of $\mathbb{Q}^d$, we write $\mathsf{span}(\mathbf{A})$ (resp., $\mathsf{cone}(\mathbf{A})$) for the *vector space* (resp., the *cone*) *generated* by $\mathbf{A}$, i.e., the closure of $\mathbf{A}$ under addition and under multiplication by all (resp., nonnegative) rationals. A *k-perfect half-space* of $\mathbb{Q}^d$, where $k \in \{1, 2, \ldots, d\}$, is a (necessarily disjoint) union $H_d \cup \cdots \cup H_k$ such that:

- H_d is an open half-space of $\mathbb{Q}^d$;
- for all $j \in \{k, \ldots, d-1\}$, $H_j \subseteq \mathbb{Q}^d$ is an open half-space of the boundary of H_{j+1}.

Whenever we write a k-perfect half-space in form $H_d \cup \cdots \cup H_k$, we assume that each H_j is j-dimensional. We additionally define the $(d+1)$-perfect half-space as the empty set; a *partially-perfect half-space* is then a k-perfect half-space for some k in $\{1, \ldots, d+1\}$. A *perfect half-space* is a 1-perfect half-space.

4.2 Generated Perfect Half-Spaces

In order to pursue effective and parsimonious strategy constructions, we consider perfect half-spaces generated by particular sets of vectors, which will correspond to the total weights of simple cycles in multi-weighted game graphs. Given a norm M in $\mathbb{N}$, we say that an open half-space H is *M-generated* if its boundary equals $\mathsf{span}(\mathbf{B})$ for some set $\mathbf{B}$ of vectors of norm at most M. By extension, a partially-perfect half-space is *M-generated* if each of its open half-spaces is M-generated.

Proposition 4.1. *Any k-dimensional vector space of $\mathbb{Q}^d$ has at most $\mathcal{L}(k) \stackrel{\text{def}}{=} 2(2M+1)^{d(k-1)}$ open half-spaces that are M-generated.*

Example 4.2. In the game graph of Fig. 2, there are three 1-generated open half-spaces of interest: the half-plane $H_2 \stackrel{\text{def}}{=} \{(x, y) : x + y < 0\}$ with boundary $\mathsf{span}((-1, 1), (1, -1))$ and containing $(-1, -1)$, and the two half-lines $H_1 \stackrel{\text{def}}{=} \{(x, y) : x + y = 0 \land x < 0\}$ and $H_1' \stackrel{\text{def}}{=} \{(x, y) : x + y = 0 \land x > 0\}$ with boundary $\mathsf{span}(\mathbf{0})$ and containing, respectively, $(-1, 1)$ and $(1, -1)$. Those open half-spaces define two perfect half-spaces: $H_2 \cup H_1$ and $H_2 \cup H_1'$.

4.3 Hierarchy of Perfect Half-Spaces

Finally, we fix a ranked tree-like structure on all M-generated partially-perfect half-spaces, which provide a scaffolding on which we will build strategies in multi-dimensional bounding games. Observe that an M-generated partially-perfect half-space $H_d \cup \cdots \cup H_k$ for $k > 1$ can be extended using any of the M-generated open half-spaces H of the boundary of H_k; note that this boundary then equals $\mathsf{span}(H)$. In Example 4.2, H_2 can be extended using H_1 or H_1', and $\mathsf{span}(H_1) = \mathsf{span}(H_1') = \{(x, y) : x + y = 0\}$.

The set of M-generated perfect half-spaces can be totally ordered by positing a linear ordering $<$ between all M-generated open half-spaces. We write $\prec$ for the lexicographically induced linear ordering between all M-generated perfect half-spaces of $\mathbb{Q}^d$: if $\mathcal{H} = H_d \cup \cdots \cup H_1$ and $\mathcal{H}' = H_d' \cup \cdots \cup H_1'$, we define $\mathcal{H} \prec \mathcal{H}'$ to hold iff $H_j = H_j'$ for all $j \in \{k+1, \ldots, d\}$ and $H_k < H_k'$ for some $k \in \{1, 2, \ldots, d\}$.

5 First-Cycle Bounding Games

We define in this section *first-cycle bounding games*, which provide the key technical arguments for most of our results. Such games end as soon as a cycle is formed along a play, and the weight of this cycle determines the winner, along with a colouring information chosen by Player 2. In sections 5.2 and 5.3, we are going to show that first-cycle bounding games and infinite bounding games are equivalent, by translating winning strategies for each Player p, $p \in \{2, 1\}$, from first-cycle bounding games to bounding games. This yields in particular the small hypercube property of Lem. 3.1.

5.1 Definition

We define the *first-cycle bounding game* $G(V, E, d)$ on a multi-weighted game graph (V, E, d):

- at any Player-1 vertex, Player 2 chooses a $|V| \cdot \|E\|$-generated perfect half-space $\mathcal{H}$ of $\mathbb{Q}^d$, and then Player 1 chooses an outgoing edge, whose occurrence in the play becomes coloured by $\mathcal{H}$;
- at any Player-2 vertex, he chooses an outgoing edge;
- the game finishes as soon as a vertex is visited twice, which produces a simple cycle C with coloured Player-1 edges;
- Player 2 wins if $w(C)$, the total weight of the cycle, is in the largest partially-perfect half-space of $\mathbb{Q}^d$ that is contained in all the colours in C, i.e. the least common ancestor of all the colours in C; Player 1 wins otherwise.

Example 5.1. Player 2 wins the first-cycle bounding game played in Fig. 1 (but loses in its lossy version). For example, strategy (2) is winning for Player 2 if he colours the edges outgoing from v_0 by the perfect half-space $H_2' \cup H_1$ where $H_2' \stackrel{\text{def}}{=} \{(x, y) : x + y > 0\}$ and $H_1 \stackrel{\text{def}}{=} \{(x, y) : x + y = 0 \wedge x < 0\}$.

Example 5.2. Player 2 wins the first-cycle bounding game played in Fig. 2. Indeed, he can choose the colour $H_2 \cup H_1$ in v_L and the colour $H_2 \cup H_1'$ in v_R. Then Player 1 cannot avoid forming a simple cycle in either $H_2 \cup H_1$ (if cycling on v_L), in $H_2 \cup H_1'$ (if cycling on v_R), or in H_2 (if cycling between v_L and v_R).

Observe that first-cycle bounding games are finite perfect information games, and are thus *determined*: from any vertex, either Player 1 wins or Player 2 wins.

5.2 Winning Strategies for Player 2

Suppose σ is a strategy of Player 2 from a vertex v_0 in a first-cycle bounding game $G(V, E, d)$. Let $\widetilde{\sigma}$ be the following strategy of Player 2 in the infinite bounding game $\Gamma(V, E, d)$:

- at any Player-2 vertex, $\widetilde{\sigma}$ chooses the edge specified by σ;
- whenever a cycle is formed, $\widetilde{\sigma}$ cuts it out of its memory, and continues playing according to σ.

Lemma 5.3. *If σ is winning for Player 2 in $G(V, E, d)$ from some vertex v_0, then $\widetilde{\sigma}$ is winning for Player 2 in $\Gamma(V, E, d)$ from the same vertex v_0.*

Proof idea. Consider any infinite play $\widetilde{\rho}$ consistent with $\widetilde{\sigma}$, and let:

- ρ be obtained from $\widetilde{\rho}$ by colouring all Player 1's edges with the $|V| \cdot \|E\|$-generated perfect half-spaces of $\mathbb{Q}^d$ as specified by σ;
- $C_1, C_2, \ldots$ be the cycle decomposition of ρ, and for each n, ρ_n be the simple path that remains after removing C_n;
- $\mathcal{H}_n$ be the largest partially-perfect half-space of $\mathbb{Q}^d$ that is contained in all the colours in C_n, for each n.

Since σ is winning for Player 2 in the first-cycle game, each cycle weight $w(C_n)$ belongs to the partially-perfect half-space $\mathcal{H}_n$. The bulk of the proof consists in extracting a 'direction of divergence' of the total energy, notwithstanding that the $\mathcal{H}_n$'s may keep varying.

In short, by distinguishing those n's for which the length of the simple path ρ_n is the smallest one that occurs infinitely often, we are able to show that the set of $\mathcal{H}_n$'s that occur infinitely often has a unique smallest element $\mathcal{H} = H_d \cup \cdots \cup H_k$ with respect to inclusion. Further linear-algebraic reasoning then shows that one of the component half-spaces $H_{k'}$ of $\mathcal{H}$ provides the desired direction of divergence: after some $N > 0$, all the sums of cycle weights $w(C_N) + w(C_{N+1}) + \cdots + w(C_n)$ belong to the topological closure $\overline{H_{k'}}$ and their distances from the boundary of $H_{k'}$ diverge. See the full paper for details. □

5.3 Winning Strategies for Player 1

If there is no winning strategy for Player 2 in the first-cycle bounding game $G(V, E, d)$ from a vertex v_0, then by determinacy of first-cycle bounding games, there is a winning strategy σ for Player 1 in $G(V, E, d)$ from v_0.

Example 5.4. Consider the lossy version of the game graph in Fig. 1. Because Player 1 wins the energy game with initial credit $(2, 2)$, by Prop. 2.2 and Lem. 5.3, she wins the first-cycle bounding game. One winning strategy, whose moves depend only on the latest visited vertex (here only v_0) and colour $\mathcal{H}$ chosen by Player 2 in v_0, is as follows:

(i) if $(-2, 2)$ and $(-1, 3)$ are both outside $\mathcal{H}$, move to v_L, and
(ii) if $(2, -1)$ and $(3, -3)$ are both outside $\mathcal{H}$, move to v_R, and
(iii) otherwise perform the self-loop labelled $(-1, 0)$.

Observe that the first two cases (i) and (ii) are disjoint. Since there is no perfect half-space that contains $(-1, 0)$ and intersects both $\{(-2, 2), (-1, 3)\}$ and $\{(2, -1), (3, -3)\}$, this strategy is indeed winning for Player 1—the same would apply if she were to choose the other self-loop $(0, -1)$ instead.

The proof of our main result consists in constructing from σ a finite-memory winning strategy $\widetilde{\sigma}$ for Player 1 in the infinite bounding game $\Gamma(V, E, d)$ from v_0, which ensures the small hypercube property stated in Lem. 3.1. Let us outline this construction. The memory of $\widetilde{\sigma}$ consists of:

a simple path γ from the initial vertex v_0 to the current vertex v, in which Player 1's edges are coloured by $|V| \cdot \|E\|$-generated perfect half-spaces of $\mathbb{Q}^d$ (this can be represented concretely by a sequence of coloured edges from E);

a colour i.e. a $|V| \cdot \|E\|$-generated perfect half-space $\mathcal{H} = H_d \cup \cdots \cup H_1$ of $\mathbb{Q}^d$ (initially the $\prec$-minimal one);

counters $\mathsf{c}(k, W)$ for every $k \in \{1, 2, \ldots, d\}$ and for every nonzero total weight W of a simple cycle, which are natural numbers (initially 0).

Strategy $\widetilde{\sigma}$ copies its moves from strategy σ for the first-cycle bounding game, based on the coloured simple path and the colour it has in its memory. Whenever a cycle is formed it is removed from the simple path, and provided its weight W is nonzero, all the counters $\mathsf{c}(k, W)$ are incremented.

Together with the current path, the counters provide the current energy level, which equals $w(\gamma) + \sum_W \mathsf{c}(d, W) \cdot W$ throughout the play, where W ranges over all simple cycle weights. To keep the counters and thus the total energy bounded, $\widetilde{\sigma}$ may perform one of the following operations after a counter increment:

- a *k-shift to $H'_k > H_k$* changes the current colour $\mathcal{H}$ to the $\prec$-minimal perfect half-space of the form $H_d \cup \cdots \cup H_{k+1} \cup H'_k \cup \cdots \cup H'_1$, and resets to 0 all the counters $\mathsf{c}(k', W)$ with $k' < k$;
- a *k-cancellation* changes the current colour $\mathcal{H}$ to the $\prec$-minimal perfect half-space of the form $H_d \cup \cdots \cup H_{k+1} \cup H'_k \cup \cdots \cup H'_1$. Simultaneously, given some simple cycle weights $W_1, \ldots, W_n$ and a positive integral solution $\mathbf{x}$ to $\sum_{i=1}^{n} \mathbf{x}(i) W_i = \mathbf{0}$, it subtracts $\mathbf{x} \cdot u(k)$ where $u(k) \stackrel{\text{def}}{=} (4|V| \cdot \|E\|)^{(2k-1)(d+2)^2}$ from all the tuples $(\mathsf{c}(k', W_1), \ldots, \mathsf{c}(k', W_n))$ with $k' \geq k$, and resets to 0 all the counters $\mathsf{c}(k', W)$ with $k' < k$.

These operations allow to maintain two main invariants, from which the small hypercube property of Lem. 5.5 is derived. For all $1 \leq k \leq d$ and simple path weights W in the span of H_k:

- initially, after any $>k$-shift, and after any $\geq k$-cancellation, $\mathsf{c}(k, W) < \mathcal{U}(k) \stackrel{\text{def}}{=} (4|V| \cdot \|E\|)^{2k(d+2)^2}$, the so-called *$k$-soft bound*;
- at all times, $\mathsf{c}(k, W) < \mathcal{U}(k) + u(k)$, the so-called *$k$-hard bound*.

To ensure those invariants, strategy $\widetilde{\sigma}$ further maintains that, whenever $\mathsf{c}(k, W) \geq \mathcal{U}(k)$ and W is in $\mathsf{span}(H_k)$, then W is in $\overline{H_k}$. When this new invariant cannot be preserved by any k-shift, then a version of the Farkas-Minkowski-Weyl Theorem implies that it can be enforced through a k-cancellation, in which a small positive integral solution can be found for the associated system of equations where $W_1, \ldots, W_n$ are the offending cycle weights.

This strategy shows a statement dual to Lem. 5.3, and thereby entails both the equivalence of infinite bounding games with first-cycle bounding games and the small hypercube property of Lem. 3.1 (see the full paper for a proof):

Lemma 5.5. *If σ is winning for Player 1 in $G(V, E, d)$ from some vertex v_0, then $\widetilde{\sigma}$ is winning for Player 1 in $\Gamma(V, W, d)$ from v_0, and ensures energy levels of norm at most* $(4|V| \cdot \|E\|)^{2(d+2)^3}$.

6 Concluding Remarks

In this paper, we have shown in Thm. 3.3 and Thm. 3.5 that fixed-dimensional energy games can be solved in pseudo-polynomial time, regardless of whether the initial credit is arbitrary or fixed. For the variant with given initial credit, this closes a large complexity gap between the TOWER upper bounds of Brázdil, Jančar, and Kučera [4] and the lower bounds of Courtois and Schmitz [9], and also settles the complexity of simulation problems between VASS and finite state systems [9]:

Corollary 6.1. *The given initial credit problem for energy games is* 2EXPTIME-*complete, and* EXPTIME-*complete in fixed dimension* $d \geq 4$.

The main direction for extending these results is to consider a *parity* condition on top of the energy condition. Abdulla, Mayr, Sangnier, and Sproston [2] show that multi-dimensional energy parity games with given initial credit are decidable. They do not provide any complexity upper bounds—although one might be able to show TOWER upper bounds from the memory bounds on winning strategies shown by Chatterjee et al. [7], Lem. 3—, leaving a large complexity gap with 2EXPTIME-hardness. This gap also impacts the complexity of *weak simulation* games between VASS and finite state systems [2].

Acknowledgements. The authors thank Dmitry Chistikov for his assistance in proving Lem. 5.3, the anonymous reviewers for their insightful comments, and Christoph Haase, Jérôme Leroux, and Claudine Picaronny for helpful discussions.

References

1. Abdulla, P.A., Atig, M.F., Hofman, P., Mayr, R., Kumar, K.N., Totzke, P.: Infinite-state energy games. CSL-LICS 2014. ACM (2014)
2. Abdulla, P.A., Mayr, R., Sangnier, A., Sproston, J.: Solving parity games on integer vectors. Concur 2013. LNCS, vol. 8052, pp. 106–120. Springer (2013)
3. Brázdil, T., Chatterjee, K., Kučera, A., Novotný, P.: Efficient controller synthesis for consumption games with multiple resource types. CAV 2012. LNCS, vol. 7358, pp. 23–38. Springer (2012)
4. Brázdil, T., Jančar, P., Kučera, A.: Reachability games on extended vector addition systems with states. ICALP 2010. LNCS, vol. 6199, pp. 478–489. Springer (2010)
5. Chaloupka, J.: Z-reachability problem for games on 2-dimensional vector addition systems with states is in P. Fund. Inform. 123(1), 15–42 (2013)
6. Chatterjee, K., Doyen, L., Henzinger, T.A., Raskin, J.F.: Generalized mean-payoff and energy games. FSTTCS 2010. LIPIcs, vol. 8, pp. 505–516. LZI (2010)
7. Chatterjee, K., Randour, M., Raskin, J.F.: Strategy synthesis for multi-dimensional quantitative objectives. Acta Inf. 51(3–4), 129–163 (2014)
8. Chatterjee, K., Velner, Y.: Hyperplane separation technique for multidimensional mean-payoff games. Concur 2013. LNCS, vol. 8052, pp. 500–515. Springer (2013)
9. Courtois, J., Schmitz, S.: Alternating vector addition systems with states. MFCS 2014, Part I. LNCS, vol. 8634, pp. 220–231. Springer (2014)
10. Fahrenberg, U., Juhl, L., Larsen, K.G., Srba, J.: Energy games in multiweighted automata. ICTAC 2011. LNCS, vol. 6916, pp. 95–115. Springer (2011)
11. Juhl, L., Larsen, K.G., Raskin, J.: Optimal bounds for multiweighted and parametrised energy games. Theories of Programming and Formal Methods—Essays Dedicated to Jifeng He on the Occasion of his 70th Birthday. LNCS, vol. 8051, pp. 244–255. Springer (2013)
12. Rackoff, C.: The covering and boundedness problems for vector addition systems. Theor. Comput. Sci. 6(2), 223–231 (1978)