Perfect Half Space Games

Thomas Colcombet
IRIF
CNRS & Université Paris-Diderot

Marcin Jurdziński
Ranko Lazić
DIMAP & Dpt. of Computer Science
University of Warwick

Sylvain Schmitz
LSV
ENS Paris-Saclay & CNRS & INRIA
Université Paris-Saclay

***Abstract*—We introduce perfect half space games, in which the goal of Player 2 is to make the sums of encountered multi-dimensional weights diverge in a direction which is consistent with a chosen sequence of perfect half spaces (chosen dynamically by Player 2). We establish that the bounding games of Jurdziński et al. (ICALP 2015) can be reduced to perfect half space games, which in turn can be translated to the lexicographic energy games of Colcombet and Niwiński, and are positionally determined in a strong sense (Player 2 can play without knowing the current perfect half space). We finally show how perfect half space games and bounding games can be employed to solve multi-dimensional energy parity games in pseudo-polynomial time when both the numbers of energy dimensions and of priorities are fixed, regardless of whether the initial credit is given as part of the input or existentially quantified. This also yields an optimal 2-EXPTIME complexity with given initial credit, where the best known upper bound was non-elementary.**

I. Introduction

A d-dimensional energy game [5, 14] sees two players compete in a finite game graph, whose edges are decorated with vectors of weights in $\mathbb{Z}^d$. The d weights represent various discrete resources that can be consumed or replenished by the actions of the game. The objective of Player 1, given an initial credit in $\mathbb{N}^d$, is to play indefinitely without depleting any of the resources—more precisely to keep the current sum of encountered weights plus initial credit non-negative in every dimension—while Player 2 attempts to foil this. The primary motivation for these games is controller synthesis for resource-sensitive reactive systems, where they are also closely related to multi-dimensional mean-payoff games—and actually equivalent if finite-memory strategies are sought for the latter [14, Lemma 6]. But they appear in diverse settings: for example, in process algebra, they are equivalent to the simulation problem between a finite state system and a Petri net or a basic parallel process [10, Propositions 6.2 and 6.4]; in artificial intelligence, they allow to solve the model-checking problem for the resource-bounded logic RB$\pm$ATL [3, 2].

The algorithmic issues surrounding multi-dimensional energy games have come under considerable scrutiny. Deciding whether there exists an initial credit that would allow Player 1 to win is coNP-complete [14, Theorem 3], while the complexity when the initial credit is given as part of the input becomes 2-EXPTIME-complete [10, 12]. Finally, both decision problems are in pseudo-polynomial time when d is fixed [12].

Open Questions: However, these recent advances do not settle the case of multi-dimensional energy parity games [7], where Player 1 must ensure that, in addition to the quantitative energy objective (specifying resource consumption and replenishment), she also complies with a qualitative ω-regular objective in the form of a parity condition (specifying functional requirements). These games with arbitrary initial credit are still coNP-complete as a consequence of [7, Lemma 4]. With given initial credit, they were first proven decidable by Abdulla, Mayr, Sangnier, and Sproston [1], and used to decide both the model-checking problem for a suitable fragment of the μ-calculus against Petri net executions and the *weak* simulation problem between a finite state system and a Petri net; they also allow to decide the model-checking problem for the resource logic RB$\pm$ATL* [2]. As shown by Jančar [11], d-dimensional energy games using $2p$ priorities can be reduced to 'extended' multi-dimensional energy games of dimension $d' \stackrel{\text{def}}{=} d+p$, with complexity upper bounds shown earlier by Brázdil, Jančar, and Kučera [5] to be in $(d'-1)$-EXPTIME when $d' \geq 2$ is fixed, and in TOWER when d' is part of the input, leaving a substantial complexity gap with the 2-EXPTIME-hardness shown in [10].

Contributions: We introduce in Section II *perfect half space games*, both

- as intermediate objects in a chain of reductions from multi-dimensional energy parity games to mean-payoff games (see Figure 1), allowing us to derive new tight complexity upper bounds based on recent advances by Comin and Rizzi [9] on the complexity of mean-payoff games, and
- as a means to gain a deeper understanding of how winning strategies in energy games are structured.

More precisely, in perfect half-space games, positions are pairs: a vertex from a d-dimensional game graph as above, together with a d-dimensional *perfect half space*. The latter is a maximal salient blunt cone in $\mathbb{Q}^d$: a union of open half spaces of dimensions d, $d-1$, …, 1, where each is contained in the boundary of the previous one. In these games, Player 1 may not change the current perfect half space, but Player 2 may change it arbitrarily at any move. However, the goal of Player 2 is to make the sums of encountered weights diverge in a direction which is consistent with the chosen perfect half spaces; thus the greater the dimension of the component open half spaces that Player 2 varies infinitely often, the harder it is for him to win. For example, with $d = 2$, if Player 2 eventually settles on the perfect half space that consists of the half plane $x < 0$ and the half line $x = 0 \wedge y < 0$, then he wins provided the sequence of total weights is such that either

978-1-5090-3018-7/17/$31.00

their x-coordinates diverge to $-\infty$, or their x-coordinates do not diverge to $+\infty$ and their y-coordinates diverge to $-\infty$; if however Player 2 switches between the two half lines of $x = 0$ infinitely often, then he can only win in the former manner.

Firstly, we show that perfect half space games can be easily translated to the *lexicographic energy games* of Colcombet and Niwiński [8]. The translation amounts to normalising the edge weights with respect to the current perfect half spaces, and inserting another d dimensions in which we encode appropriate penalties for Player 2 that are imposed whenever he changes the perfect half space (cf. Section III-B). We deduce that perfect half space games are positionally determined, and moreover that Player 2 has winning strategies that are oblivious to the current perfect half space. Along the way, we provide in Section III-A a proof of the positional determinacy of lexicographic energy games, along with pseudo-polynomial complexity upper bounds for their decision problem when d is fixed, based on the recent results of Comin and Rizzi [9] for mean-payoff games.

Secondly, we establish that perfect half space games capture *bounding games* (cf. Section IV). The latter were central to obtaining the tight complexity upper bounds for multi-dimensional energy games [12]. They are played purely on the d-dimensional game graphs and have a simple winning condition: the goal of Player 1 is to keep the total absolute value of weights bounded (i.e., contained in some d-dimensional hypercube). One reading of this reduction is that whenever Player 2 has a winning strategy in a bounding game, he has one that 'announces' at every move some perfect half space and succeeds in forcing the total weights to be unbounded in a direction consistent with the infinite sequence of his announcements. The proof is difficult, and relies on a construction from the previous paper [12] of a winning strategy for Player 1 in the bounding game given her winning strategy in a first-cycle game featuring perfect half spaces. Composing this with our complexity bounds for lexicographic energy games gives us a new approach to solving bounding games, improving the time complexity from the previously best $(|V| \cdot \|E\|)^{O(d^4)}$ [12, Corollary 3.2] to $(|V| \cdot \|E\|)^{O(d^3)}$, where V is the set of vertices and $\|E\|$ the maximal absolute value over the weights in the input multi-dimensional game graph (cf. Corollary IV.6).

Thirdly, building on Jančar's reduction, we show how multi-dimensional energy parity games can be solved by means of bounding games (cf. Section V). For the given initial credit problem, we obtain 2-EXPTIME-completeness, closing the aforementioned complexity gap. When the dimension d and the number of priorities $2p$ are fixed, we obtain that, for both arbitrary and given initial credits, the winner is decidable in pseudo-polynomial time. With arbitrary initial credit, our new bound $(|V| \cdot \|E\|)^{O((d+p)^3 \log(d+p))}$ improves when $p = 0$ over the previously best $(|V| \cdot \|E\|)^{O(d^4)}$ [12, Theorem 3.3].

Structure of the Paper: The chain of reductions we use in this paper is depicted in Figure 1, and we shall essentially work our way up through it. In Section II we introduce multi-dimensional game graphs and perfect half space games. In Section III we show how to employ lexicographic energy games for solving perfect half space games. We apply these results to bounding games in Section IV and multi-dimensional energy parity games in Section V, before concluding.

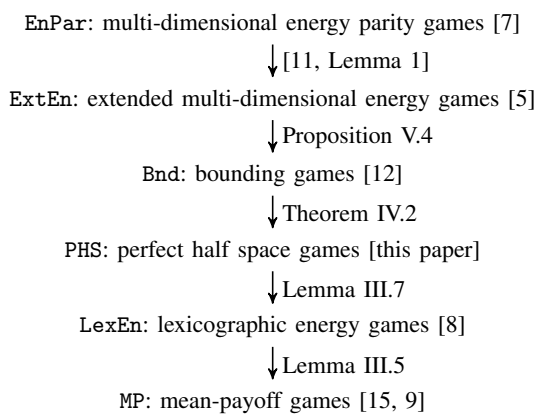

Fig. 1. The reductions between the various games in this paper.

II. Perfect Half Space Games

A. Multi-Weighted Game Graphs

We consider *multi-dimensional game graphs* whose edges are labelled by *multi-weights*, which are vectors of integers. They are tuples of the form (V, E, d), where d is the dimension in $\mathbb{N}_{>0}$, $V \stackrel{\text{def}}{=} V_1 \uplus V_2$ is a finite set of vertices partitioned into Player 1 vertices and Player 2 vertices, and E is a finite set of edges included in $V \times \mathbb{Z}^d \times V$, such that every vertex has at least one outgoing edge. We may write just '*weight*' instead of 'multi-weight' when there is no risk of confusion, and also $v \xrightarrow{\mathbf{w}} v'$ to denote an edge $(v, \mathbf{w}, v')$. Given a path P in the game, we denote by $\mathbf{w}(P)$ the sum of the weights encountered.

For a vector $\mathbf{w}$ in $\mathbb{Z}^d$, we let $\|\mathbf{w}\| \stackrel{\text{def}}{=} \max_{1 \leq i \leq d} |\mathbf{w}(i)|$ denote its infinity norm; we define the norm $\|E\| \stackrel{\text{def}}{=} \max_{v \xrightarrow{\mathbf{w}} v' \in E} \|\mathbf{w}\|$ as the maximum of the norms of edge weights. We assume all our integers to be encoded in binary, hence $\|E\|$ might be exponential in the size of the multi-weighted game graph.

Without loss of generality, we assume that the players strictly alternate ($v \xrightarrow{\mathbf{w}} v'$ in E implies v in V_i and v' in V_{3-i} for some i in $\{1, 2\}$), the weight of every edge is determined by its vertices ($v \xrightarrow{\mathbf{w}} v'$ and $v \xrightarrow{\mathbf{w}'} v'$ in E implies $\mathbf{w} = \mathbf{w}'$), and not all weights are zero ($\|E\| > 0$).

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

B. Perfect Half Spaces

We *represent* partially perfect half spaces by tuples $\mathbf{H} = (\mathbf{h}_1, \ldots, \mathbf{h}_k)$ of $k \leq d$ mutually normal nonzero d-dimensional integer vectors, which are normal to the represented half spaces. For this, let $\prec$ denote the (strict) lexicographic ordering, and for any d-dimensional vector $\mathbf{a}$, let $\mathbf{a} \cdot \mathbf{H}$ denote the pointwise dot-product $(\mathbf{a} \cdot \mathbf{h}_1, \ldots, \mathbf{a} \cdot \mathbf{h}_k)$. The *partially perfect half space* denoted by $\mathbf{H}$ is then $\{\mathbf{a} \in \mathbb{Q}^d : \mathbf{a} \cdot \mathbf{H} \prec \mathbf{0}\}$.

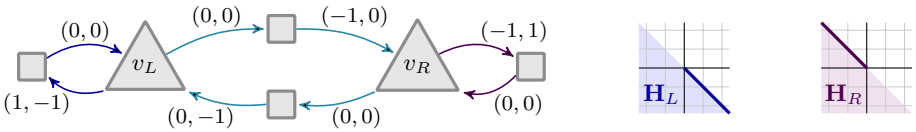

Fig. 2. A 2-dimensional game graph $(V, E, 2)$ and two perfect half spaces.

Let $|\mathbf{H}| \stackrel{\text{def}}{=} k$. When $|\mathbf{H}| = d$, the representation is a (full) *perfect half space*; when $|\mathbf{H}| = 0$, it is the empty set since there is only one 0-dimensional vector and the ordering $\prec$ is strict.

We define the *norm* $\|\mathbf{H}\|$ as the maximum of $\|\mathbf{h}_1\|$, …, $\|\mathbf{h}_k\|$.

Example II.2. The two perfect half spaces of interest on the right-hand side of Figure 2 are $\{(x, y) : x+y < 0\} \cup \{(x, y) : x + y = 0 \wedge x > 0\}$ denoted by $\mathbf{H}_L \stackrel{\text{def}}{=} ((1, 1), (-1, 1))$, and $\{(x, y) : x + y < 0\} \cup \{(x, y) : x + y = 0 \wedge x < 0\}$ denoted by $\mathbf{H}_R \stackrel{\text{def}}{=} ((1, 1), (1, -1))$. They have the half-plane $\{(x, y) : x+y < 0\}$ with normal vector $(1, 1)$ in common, but differ on which half-line of its boundary $\{(x, y) : x + y = 0\}$ they contain. □

We shall reason sometimes directly on the representations of partially perfect half spaces through the *prefix ordering*. We write $\mathbf{H} \leq^{\text{pref}} \mathbf{H}'$ when $\mathbf{H}$ is a prefix of $\mathbf{H}'$, and $\operatorname{lcp}_i \mathbf{H}_i$ for the longest common prefix of a finite or infinite set of partially perfect half spaces $\mathbf{H}_1, \mathbf{H}_2, \ldots$. Observe that, if $\mathbf{a} \cdot \mathbf{H}' \prec \mathbf{0}$ and $\mathbf{H} \leq^{\text{pref}} \mathbf{H}'$, then $\mathbf{a} \cdot \mathbf{H} \preceq \mathbf{0}$.

C. Perfect Half Space Games

We write $(\widehat{V}, \widehat{E}, d)$ for the weighted game graph obtained from (V, E, d) by pairing vertices in V with perfect half spaces of appropriately bounded norms, which may be changed only by Player 2:

- for both $i \in \{1, 2\}$, $\widehat{V}_i \stackrel{\text{def}}{=} V_i \times \mathcal{H}$ where $\mathcal{H}$ is the set of all perfect half spaces of norm at most $|V| \cdot \|E\|$;
- $\widehat{E}$ is the set of all $(v, \mathbf{H}) \xrightarrow{\mathbf{w}} (v', \mathbf{H}')$ such that $v \xrightarrow{\mathbf{w}} v'$ is in E and if $v \in V_1$ then $\mathbf{H} = \mathbf{H}'$.

Let $\texttt{PHS}(\widehat{V}, \widehat{E}, d)$ denote the *perfect half space game* in which the goal of Player 2 is for the total weight to diverge in a direction consistent with the chosen perfect half spaces:

Definition II.3 (Winning Condition for Perfect Half-Space Games)**.** *An infinite play* $(v_0, \mathbf{H}_0) \xrightarrow{\mathbf{w}_1} (v_1, \mathbf{H}_1) \xrightarrow{\mathbf{w}_2} (v_2, \mathbf{H}_2) \cdots$ *is* winning for Player 2 *if there exists a partially perfect half space* $(\mathbf{g}_1, \ldots, \mathbf{g}_k)$ *with* $k > 0$ *that is a prefix of* $\mathbf{H}_i$ *for all sufficiently large* i*, s.t.* $\limsup_n \sum_{j=1}^{n} \mathbf{w}_j \cdot \mathbf{g}_k = -\infty$ *and, for all* $1 \leq \ell < k$*,* $\liminf_n \sum_{j=1}^{n} \mathbf{w}_j \cdot \mathbf{g}_\ell < +\infty$.

Observe that whether Player 2 wins from $(v, \mathbf{H})$ does not depend on $\mathbf{H}$, hence we say that Player 2 *wins from* v if there exists $\mathbf{H} \in \mathcal{H}$ such that he wins from $(v, \mathbf{H})$—equivalently, he wins from $(v, \mathbf{H})$ for all $\mathbf{H} \in \mathcal{H}$—, and similarly for Player 1.

Given a finite path

$$P \stackrel{\text{def}}{=} (v_0, \mathbf{H}_0) \xrightarrow{\mathbf{w}_1} (v_1, \mathbf{H}_1) \cdots (v_{n-1}, \mathbf{H}_{n-1}) \xrightarrow{\mathbf{w}_n} (v_n, \mathbf{H}_n)$$

in a perfect half space game, we denote by $\operatorname{lcp}(P) \stackrel{\text{def}}{=} \operatorname{lcp}_{0 \leq i \leq n} \mathbf{H}_i$ the longest partially perfect half space that agrees with all the perfect half spaces seen along the path. We also inherit the notation $\mathbf{w}(P) \stackrel{\text{def}}{=} \sum_{i=1}^{n} \mathbf{w}_i$ that accounts for the sum of the weights in P. We say that P *is winning for Player* 1 if $\mathbf{w}(P) \cdot \operatorname{lcp}(P) \succeq \mathbf{0}$. Similarly, P *is winning for Player* 2 if $\mathbf{w}(P) \cdot \operatorname{lcp}(P) \prec \mathbf{0}$. Note that when P is in fact a cycle, then its infinite iteration is winning for a player if and only if the cycle is winning for them according to this definition.

Example II.4. Player 2 wins the perfect half space game on the graph of Example II.1 from any vertex by choosing the perfect half space $\mathbf{H}_L$ from Example II.2 when going to v_L and $\mathbf{H}_R$ when going to v_R. Indeed, either Player 1 eventually only uses the left (blue) cycle, in which case $(\mathbf{g}_1, \mathbf{g}_2) \stackrel{\text{def}}{=} \mathbf{H}_L$ itself can be used as witness in Definition II.3, or she eventually only uses the right (violet) cycle, in which case $(\mathbf{g}_1, \mathbf{g}_2) \stackrel{\text{def}}{=} \mathbf{H}_R$ fits, or she alternates infinitely often between v_R and v_L (using the cyan cycle), in which case the partially perfect half space $(\mathbf{g}_1) \stackrel{\text{def}}{=} ((1, 1))$ is a witness of his victory. □

III. Solving Perfect Half Space Games

As an intermediate step towards the proof of our determinacy and complexity results for perfect half space games (Theorem III.8), we employ another winning condition introduced in [8]: that of lexicographic energy games. We start by presenting a proof of their positional determinacy, and an upper bound for their decision problem using the state-of-the-art results of Comin and Rizzi [9] for mean-payoff games. We then proceed to show how perfect half space games can be reduced to lexicographic energy ones in Section III-B.

A. Solving Lexicographic Energy Games

1) Lexicographic Energy Games [8] are played on multi-weighted game graphs (V, E, d), as described in Section II. An infinite play $v_0 \xrightarrow{\mathbf{w}_1} v_1 \xrightarrow{\mathbf{w}_2} \cdots$ is *winning for Player* 2 if there exists $1 \leq k \leq d$ s.t. $\limsup_n \sum_{j=1}^{n} \mathbf{w}_j(k) = -\infty$ and, for all $1 \leq \ell < k$, $\liminf_n \sum_{j=1}^{n} \mathbf{w}_j(\ell) < +\infty$.[1]

Put differently, lexicographic energy games are akin to perfect half space games, except that the same full perfect half space $(-\mathbf{e}_1, \ldots, -\mathbf{e}_d)$ is associated to every vertex of the game graph, where $\mathbf{e}_i$ for $1 \leq i \leq d$ denotes the unit vector with 1 in coordinate i and 0 everywhere else.

[1] Lexicographic energy games bear a superficial resemblance to two different definitions of lexicographic mean-payoff games, due respectively to Bloem et al. [4] and to Bruyère et al. [6]. However, the definition that would best match lexicographic energy games would be multi-dimensional 'pointwise' lexicographic mean-payoff games, which do *not* enjoy positional determinacy, and all these definitions are unfit for our purposes.

Example III.1. Let us consider the multi-weighted game graph of Example II.1. Player 1 wins the lexicographic energy game from any initial vertex, by moving to v_L and looping on the left (blue) loop. □

2) Strategies: A strategy for a player is *positional* if, from each of her vertices, the player using it always chooses the same outgoing edge, no matter where the play started or how it evolved so far. We say that a game is *positionally determined* if the two players have positional strategies σ and τ, respectively, such that for every vertex $v \in V$, either σ is winning for Player 1 from v, or τ is winning for Player 2 from v.

3) Reduction to Mean-Payoff Games: A *mean-payoff game* is played on a *weighted game graph*, i.e. a 1-dimensional weighted game graph $(V, E, 1)$, and is denoted $\mathtt{MP}(V, E)$. From an infinite play $v_0 \xrightarrow{u_1} v_1 \xrightarrow{u_2} v_2 \cdots$, Player 1 ('Max') gains a payoff $\liminf_{n\to\infty}(u_1+\cdots+u_n)/n$, whereas Player 2 ('Min') loses a payoff $\limsup_{n\to\infty}(u_1+\cdots+u_n)/n$. A strategy for Max is *optimal* for her if by following it she is guaranteed to gain at least as much as when using any other strategy, and optimal strategies for Min are defined symmetrically. By the positional determinacy of mean-payoff games [15], there exist positional optimal strategies for both players, yielding the same payoff for both from each initial vertex, called the *value* of the vertex.

A strategy for Max is *winning* from some initial vertex if by following it she is guaranteed to gain at least ≥ 0, and a strategy for Max is winning if by following it he is guaranteed to lose at least < 0. Note that not every winning strategy for Min needs to be optimal, but that if she wins then any optimal strategy is winning: Min wins the game if and only if the value of the initial vertex is ≥ 0, and Max wins otherwise.

For a multi-weighted game graph (V, E, d), and for every i, $1 \leq i \leq d$, let the set $E(i)$ consist of the edges $v \xrightarrow{\mathbf{w}(i)} v'$ where $v \xrightarrow{\mathbf{w}} v' \in E$.

Theorem III.2. *(i) Lexicographic energy games are positionally determined.*

(ii) There is an algorithm for solving lexicographic energy games whose running time is in $O\left(|V|^{d+1} \cdot |E| \cdot \prod_{i=1}^{d} \|E(i)\|\right)$.

We start by describing a translation from lexicographic energy games to mean-payoff games, similar to the classical translation from parity games [13]: the idea is to write the d-dimensional weights into a single weight by shifting the most significant components by appropriate amounts. We define accordingly the sets of weighted edges $E^{(i)}$ for $i = d, d-1, \ldots, 1$ as follows:

- $E^{(d)} \stackrel{\text{def}}{=} E(d)$;
- for all $i = d-1, d-2, \ldots, 1$, and for all $v \xrightarrow{\mathbf{w}} v' \in E$, if $v \xrightarrow{r_{i+1}} v' \in E^{(i+1)}$ then $v \xrightarrow{r_i} v' \in E^{(i)}$, where $r_i \stackrel{\text{def}}{=} \mathbf{w}(i) \cdot \left(|V| \cdot \|E^{(i+1)}\| + 1\right) + r_{i+1}$.

We will argue directly that positional optimal strategies for the two players in the mean-payoff game $\mathtt{MP}(V, E^{(1)})$ witness positional determinacy of the lexicographic energy game $\mathtt{LexEn}(V, E, d)$.

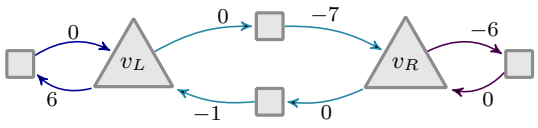

Fig. 3. The weighted game graph $(V, E^{(1)})$ constructed from the graph of Figure 2.

Example III.3. The weighted game graph obtained from the multi-weighted game graph of Example II.1 is depicted in Figure 3 (indeed $|V| \cdot \|E^{(2)}\| + 1 = 7$). Max has a positional optimal strategy consisting in moving to v_L and using the left (blue) loop; every vertex has value 6. □

The outcome of this encoding of d-dimensional weights in $E^{(1)}$ is the following, easy to establish, proposition.

Proposition III.4. *The total weight of a simple cycle in the multi-weighted game graph (V, E, d) is $\prec \mathbf{0}$ (or $= \mathbf{0}$, or $\succ \mathbf{0}$, respectively) if and only if the total weight of the cycle in the weighted game graph $(V, E^{(1)})$ is negative (or zero, or positive, respectively).*

In order to show the positional determinacy of lexicographic energy games, we rely on the following lemma proven in the appendix.

Lemma III.5. *If the value of the mean-payoff game $\mathtt{MP}(V, E^{(1)})$ is non-negative (negative, resp.) at a vertex v, then by using a positional optimal strategy from that mean-payoff game, Player* 1 *(Player* 2*, resp.) wins the corresponding lexicographic energy game $\mathtt{LexEn}(V, E, d)$ from v.*

Proof of Theorem III.2. By Lemma III.5, in order to compute a positional winning strategy for one the players in a lexicographic energy game $\mathtt{LexEn}(V, E, d)$, it suffices to find a positional optimal strategy in the corresponding mean-payoff game $\mathtt{MP}(V, E^{(1)})$. This entails the positional determinacy of lexicographic energy games (cf., e.g., [15]). Regarding complexity, the state-of-the-art algorithm for solving mean-payoff games due to Comin and Rizzi [9] runs in time $O\left(|V|^2 \cdot |E| \cdot \|E\|\right)$. Observe that $|E^{(1)}| = |E|$ and $\|E^{(1)}\| = O\left(|V|^{d-1} \cdot \prod_{i=1}^{d} \|E(i)\|\right)$, and hence the algorithm of Comin and Rizzi can be used to solve lexicographic energy games in time $O\Big(|V|^{d+1} \cdot |E| \cdot \prod_{i=1}^{d} \|E(i)\|\Big)$. □

B. Translation to Lexicographic Energy Games

We now reduce perfect half space games to lexicographic energy games. Given a perfect half space game played on a d-dimensional multi-weighted game graph, the idea is to play a lexicographic energy game on a $2d$-dimensional game graph, where the extra dimensions are used to penalise Player 2 for changing of perfect half space.

1) Flag Vectors and Interleavings: For any d-dimensional perfect half spaces $\mathbf{H}$ and $\mathbf{H}'$, let the *flag vector* $\mathbf{e}_{\mathbf{H},\mathbf{H}'}$ be defined for all $i = 1, \ldots, d$ by $\mathbf{e}_{\mathbf{H},\mathbf{H}'}(i) \stackrel{\text{def}}{=} 0$ if the i-th coordinates of $\mathbf{H}$ and $\mathbf{H}'$ are identical, and $\mathbf{e}_{\mathbf{H},\mathbf{H}'}(i) \stackrel{\text{def}}{=} 1$

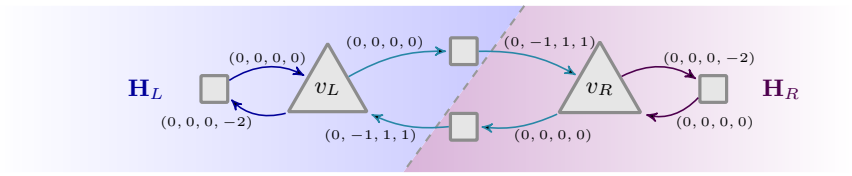

Fig. 4. The translation of the graph from Figure 2 to lexicographic energy games.

otherwise. For any d-dimensional vectors $\mathbf{a}$ and $\mathbf{b}$, let $\mathbf{a} \Delta \mathbf{b}$ be their *interleaving* $(\mathbf{a}(1), \mathbf{b}(1), \ldots, \mathbf{a}(d), \mathbf{b}(d))$.

2) Translation: We write $(\widehat{V}, \widetilde{E}, 2d)$ for the weighted game graph obtained from $(\widehat{V}, \widehat{E}, d)$ by doubling the dimension, where the even indices of weights in $\widetilde{E}$ contain the corresponding weights from $\widehat{E}$ but normalised with respect to the current perfect half space, and the odd indices are occupied by flag vectors that penalise Player 2 for changing the perfect half spaces. More precisely, $\widetilde{E}$ is the set of all $(v, \mathbf{H}) \xrightarrow{\mathbf{e}_{\mathbf{H},\mathbf{H}'} \Delta (\mathbf{w} \cdot \mathbf{H})} (v', \mathbf{H}')$ such that $(v, \mathbf{H}) \xrightarrow{\mathbf{w}} (v', \mathbf{H}')$ is in $\widehat{E}$.

Let $\texttt{LexEn}(\widehat{V}, \widetilde{E}, 2d)$ denote the lexicographic energy game played on the multi-weighted game graph $(\widehat{V}, \widetilde{E}, 2d)$.

Example III.6. We depict in Figure 4 a fragment of the translated game graph $(\widehat{V}, \widetilde{E}, 2d)$ for the perfect half space game from Example II.4. The vertices on the left of the median dashed line are all paired with $\mathbf{H}_L$, while those on the right are paired with $\mathbf{H}_R$. The flag vector $\mathbf{e}_{\mathbf{H}_L,\mathbf{H}_R} = (0, 1) = \mathbf{e}_{\mathbf{H}_R,\mathbf{H}_L}$ is interleaved with the normalised vectors on the two middle edges entering v_R and v_L.

In contrast to Example III.1, Player 1 now loses the lexicographic energy game in Figure 4. Indeed, if she plays the middle simple cycle (in cyan) infinitely often, then the energy on the first coordinate converges to 0 and the energy in the second coordinate diverges to $-\infty$. Otherwise (i.e., if the number of occurrences of the middle cycle is bounded), the energy in the first three coordinates does not diverge and the energy in the fourth coordinate diverges to $-\infty$. □

The correctness of this translation is a direct consequence of the definitions, as shown in the following lemma proven in the appendix.

Lemma III.7. *The winning strategies of Player i, $i \in \{1, 2\}$, are the same in* $\texttt{PHS}(\widehat{V}, \widehat{E}, d)$ *and* $\texttt{LexEn}(\widehat{V}, \widetilde{E}, 2d)$.

Define a strategy τ for Player 2 in the perfect half space game $\texttt{PHS}(\widehat{V}, \widehat{E}, d)$ to be *(perfect half space) oblivious at* v for $v \in V_2$ if it chooses the same move in $(v, \mathbf{H})$ for all $\mathbf{H}$. It is *perfect half space oblivious* if it is oblivious at all vertices $v \in V_2$. We are now ready to prove the main theorem of this section.

Theorem III.8. *(i) There is an algorithm for solving perfect half space games whose running time is in* $O\left((3|V| \cdot \|E\|)^{2(d+1)^3}\right)$.

(ii) If Player 2 has a winning strategy in the perfect half space game $\texttt{PHS}(\widehat{V}, \widehat{E}, d)$*, then he has one that is perfect half space oblivious.*

Proof of Theorem III.8(i). The upper bound on the running time is a consequence of Lemma III.7 and Theorem III.2.ii. Observe that the vertex set is of size $|\widehat{V}| = |V| \cdot |\mathcal{H}| \leq |V| \cdot (2|V| \cdot \|E\| + 1)^{d^2} \leq (3|V| \cdot \|E\|)^{d^2+1}$. Regarding the norms, $\|\widetilde{E}\| = \max\{\|\mathbf{w} \cdot \mathbf{H}\| : v \xrightarrow{\mathbf{w}} v' \in E, \mathbf{H} \in \mathcal{H}\}$, hence $\|\widetilde{E}\| \leq d \cdot |V| \cdot \|E\|^2 \leq (3|V| \cdot \|E\|)^{2+\log d}$. Hence a time bound in $O\left((3|V| \cdot \|E\|)^m\right)$ where $m = (d^2 + 1)(2d + 3) + 2d(2 + \log d) \leq 2(d + 1)^3$. □

Proof of Theorem III.8(ii). The idea of the following proof is to show that, for any vertex of the weighted game graph winning for Player 2, there is a 'good' perfect half space $\mathbf{H}$ such that following a positional strategy $\tau_{\mathbf{H}}$ winning from $(v, \mathbf{H})$ will also win from any $(v, \mathbf{H}')$.

More formally, we prove by induction on $k \leq |V_2|$ that there exists a winning positional strategy τ for Player 2 which is perfect half space oblivious at k distinct vertices in V_2.

The induction hypothesis obviously holds for $k = 0$ by using a positional strategy in $\texttt{PHS}(\widehat{V}, \widehat{E}, d)$, which exists by Theorem III.2 and Lemma III.7. For the induction step, let us suppose that τ is a winning positional strategy for Player 2 oblivious at $k < |V_2|$ distinct vertices $v_1, \ldots, v_k \in V_2$. Let v be another $v \in V_2$ distinct from $v_1, \ldots, v_k$ vertices; τ and v are now fixed for the remainder of the proof.

For all perfect half spaces $\mathbf{H}$, let us denote by $\tau_{\mathbf{H}}$ the strategy τ modified in such a way that it behaves in $(v, \mathbf{H}')$ as in $(v, \mathbf{H})$ for all $\mathbf{H}' \neq \mathbf{H}$. The result is still a valid strategy (by definition of the perfect half space game) and is of course oblivious at v as well as at $v_1, \ldots, v_k$. We want to show that there exists $\mathbf{H}$ such that $\tau_{\mathbf{H}}$ fulfils the induction hypothesis. This is the case for any $\mathbf{H}$ if v is not in the winning region. We shall therefore assume that v is in the winning region for Player 2; thus τ is winning from every $(v, \mathbf{H})$ but might use different moves depending on $\mathbf{H}$.

a) Good Perfect Half-Spaces: Let us call a perfect half space $\mathbf{H}$ *good (for τ and v)* if $\tau_{\mathbf{H}}$ is winning for Player 2 starting in $(v, \mathbf{H})$, and bad otherwise. As shown in the appendix, there must exist a good perfect half space, as otherwise Player 2 would not win from v.

Claim III.9. *There exists a good perfect half space.*

b) A Winning Strategy $\tau_{\mathbf{H}}$: Let $\mathbf{H}$ be a good perfect half space that exists according to Claim III.9. Let us show that $\tau_{\mathbf{H}}$ fulfils the condition of the induction hypothesis. As already mentioned, it is oblivious at $\{v, v_1, \ldots, v_k\}$. We have to prove that it is winning. For this, let us consider any play consistent with $\tau_{\mathbf{H}}$ starting from some $(v', \mathbf{H}')$ in the winning region for

Player 2. Two cases can happen. Either this play does not visit the vertex v, and in this case it was already a run consistent with τ, and hence it is winning for Player 2. Otherwise it visits v, and after that point it continues in a way consistent with $\tau_{\mathbf{H}}$ starting from $(v, \mathbf{H})$, and hence is winning for Player 2 since $\mathbf{H}$ is good. This establishes the induction hypothesis, and thus completes the proof of Theorem III.8(ii). □

IV. Bounding Games

In this section, we define bounding games (as introduced in [12]) and show how these can be reduced to perfect half space games (Theorem IV.2 below). Corollary IV.6 then summarises our knowledge about bounding games.

For a weighted game graph (V, E, d), we denote by $\texttt{Bnd}(V, E, d)$ the *bounding game* in which Player 1 ('Guard') strives to contain the total weight within some d-dimensional hypercube, while Player 2 ('Fugitive') attempts to escape. More precisely, an infinite play $v_0 \xrightarrow{\mathbf{w}_1} v_1 \xrightarrow{\mathbf{w}_2} v_2 \cdots$ is winning for Player 1 if and only if the set $\{\|\sum_{i=1}^{n} \mathbf{w}_i\| \,:\, n \in \mathbb{N}\}$ of norms of total weights of all finite prefixes of the play is bounded.

Example IV.1. Consider again the multi-weighted game graph of Example II.1. Observe that Player 1 cannot choose to play solely in the left (blue) cycle, as the accumulated weights would drift towards $(+\infty, -\infty)$; a similar argument holds with the right (violet) cycle. Hence, she must somehow balance the effect of the two cycles by switching infinitely often between v_L and v_R, but the effect of the middle (cyan) cycle then makes the simulated weights drift towards $(-\infty, -\infty)$. In fact, by the upcoming Theorem IV.2 and as seen in Example II.4, Player 2 wins this game. □

Theorem IV.2. *Let (V, E, d) be a multi-weighted game graph, v be a vertex in V, and $i \in \{1, 2\}$. Player i wins the bounding game $\texttt{Bnd}(V, E, d)$ from v if and only if Player i wins the perfect half space game $\texttt{PHS}(\widehat{V}, \widehat{E}, d)$ from v.*

By Theorem III.8, perfect half space games are determined, hence we can focus on Player 2. One implication is straightforward: a winning strategy for Player 2 in $\texttt{PHS}(\widehat{V}, \widehat{E}, d)$ also wins $\texttt{Bnd}(V, E, d)$ when ignoring the perfect half spaces. Note that this translates an oblivious strategy in $\texttt{PHS}(\widehat{V}, \widehat{E}, d)$ into a positional one in $\texttt{Bnd}(V, E, d)$.

Lemma IV.3. *If Player 2 wins $\texttt{PHS}(\widehat{V}, \widehat{E}, d)$ from v, then he wins $\texttt{Bnd}(V, E, d)$ from v with the same strategy (where perfect half spaces are projected away).*

Proof sketch. Let Player 2 follow a winning strategy for the perfect half space game, projected onto the arena of the bounding game, and consider any resulting play. By the winning condition of the former game, the total weights have unbounded distances from some hyperplane, and so have unbounded norms. □

It remains therefore to establish the converse implication in order to complete the proof of Theorem IV.2.

Lemma IV.4. *If Player 2 wins $\texttt{Bnd}(V, E, d)$ from v, then he wins $\texttt{PHS}(\widehat{V}, \widehat{E}, d)$ from v.*

The proof of this lemma relies on [12, Lemma 5.5]—the most involved result in that paper—, which shows how to construct a winning strategy for Player 1 from v in $\texttt{Bnd}(V, E, d)$ from a winning strategy in a *first-cycle* variant $\texttt{FC}(\widehat{V}, \widehat{E}, d)$ of $\texttt{PHS}(\widehat{V}, \widehat{E}, d)$ from v. As these first-cycle games are determined, this entails that, if Player 2 wins from v in $\texttt{Bnd}(V, E, d)$, then he also wins from v in the first-cycle game $\texttt{FC}(\widehat{V}, \widehat{E}, d)$, and it remains to show how to build a winning strategy for him in $\texttt{PHS}(\widehat{V}, \widehat{E}, d)$. The reasoning itself is surprisingly subtle, and similar to the one employed in the proof of [12, Lemma 5.3].

Proof of Lemma IV.4. By [12, Lemma 5.5], there exists a winning strategy σ for Player 2 from some $(v, \mathbf{H})$ in the following *first-cycle game* $\texttt{FC}(\widehat{V}, \widehat{E}, d)$:

1) the game finishes as soon as the play has a suffix $C = (v_0, \mathbf{H}_0) \xrightarrow{\mathbf{w}_1} (v_1, \mathbf{H}_1) \xrightarrow{\mathbf{w}_2} \cdots \xrightarrow{\mathbf{w}_n} (v_n, \mathbf{H}_n)$ such that $v_0 = v_n \in V_1$;
2) Player 2 wins if and only if $\mathbf{H}_0 = \mathbf{H}_n$ and C is winning for him: the total weight $\mathbf{w}(C) \stackrel{\text{def}}{=} \mathbf{w}_1 + \cdots + \mathbf{w}_n$ of the cycle is in the partially perfect half space defined by the longest common prefix, i.e. $\mathbf{w}(C) \cdot \mathrm{lcp}(C) \prec \mathbf{0}$ (recall that $\mathrm{lcp}(C) \stackrel{\text{def}}{=} \mathrm{lcp}_{1 \leq i \leq n} \mathbf{H}_i$).

Let σ^* denote the strategy for Player 2 from $(v, \mathbf{H})$ in $\texttt{PHS}(\widehat{V}, \widehat{E}, d)$ that amounts to following σ and repeatedly cutting out the winning cycles. We want to show that σ^* is winning: consider for this a play $(v_0, \mathbf{H}_0) \xrightarrow{\mathbf{w}_1} (v_1, \mathbf{H}_1) \xrightarrow{\mathbf{w}_2} \cdots$ consistent with σ^* starting from $v_0 = v$.

Let us consider the V_1 *cycle decomposition* of this play: the latter is the infinite sequence of 'V_1-simple' cycles C obtained by pushing the triples of visited vertices and perfect half spaces and indices $(v_0, \mathbf{H}_0, 0), (v_1, \mathbf{H}_1, 1), \ldots$ onto a stack, and as soon as we push a pair $(v_e, \mathbf{H}_e, e)$ with an element $(v_s, \mathbf{H}_s, s)$ with $v_s = v_e \in V_1$ already present in the stack, we pop the cycle C thus formed from the stack and push $(v_e, \mathbf{H}_e, e)$ back on top. We call the indices $s(C) \stackrel{\text{def}}{=} s$ and $e(C) \stackrel{\text{def}}{=} e$ the *start* and *end* of the cycle, and denote by $\mathrm{lcp}(C)$ and $\mathbf{w}(C)$ the longest common prefix of its perfect half spaces and total weight respectively. Because σ is winning in $\texttt{FC}(\widehat{V}, \widehat{E}, d)$, all the cycles C formed in the cycle decomposition satisfy condition 2 above, hence $\mathbf{H}_{s(C)} = \mathbf{H}_{e(C)}$ and $\mathbf{w}(C) \cdot \mathrm{lcp}(C) \prec \mathbf{0}$.

Let us now consider the longest $\mathbf{P}$ such that there exists a sufficiently large index $i_{\mathbf{P}}$ such that $\mathbf{P} = \mathrm{lcp}_{s(C) \geq i_{\mathbf{P}}}(\mathrm{lcp}(C))$. We call a partially perfect half space representation $\mathbf{H}$ *recurring* if $\mathbf{H} = \mathrm{lcp}(C)$ for infinitely many cycles C in the V_1 cycle decomposition of our play; such a vector $\mathbf{H}$ is necessarily non-empty.

Claim IV.5. *$\mathbf{P}$ is recurring.*

Proof of Claim IV.5. We reason on the height of the stack used to construct the V_1 cycle decomposition of the play. Let us call ρ_i the stack at step i. Since its height $|\rho_i|$ is bounded by $2|V_1|$, there is a smallest height h that occurs infinitely often, and a minimal index i_h such that h is the minimal height in

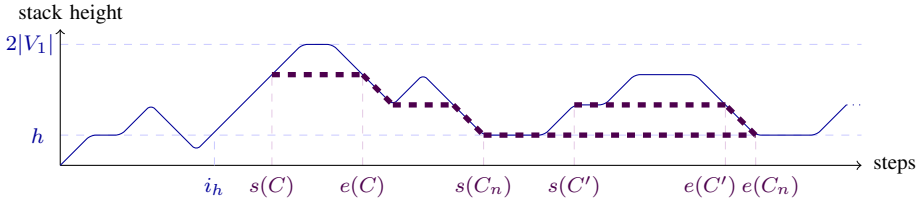

Fig. 5. Stack heights in the proof of Claim IV.5.

the infinite suffix starting from i_h. We depict the stack heights along the play in blue in Figure 5.

Let us call a *downward path* a sequence of cycles $C_1, \dots, C_n$ such that, for all $1 \leq i < n$, C_i and C_{i+1} are either two successive cycles with $|\rho_{s(C_i)}| > |\rho_{s(C_{i+1})}|$ or two cycles (not necessarily successive) with $e_{C_i} = s_{C_{i+1}}$. Observe that in both cases, they visit a common perfect half space $\mathbf{H}_{e_{C_i}}$, hence $\mathrm{lcp}(C_i)$ and $\mathrm{lcp}(C_{i+1})$ are comparable for the prefix ordering.

Assume there are two recurring representations of partially perfect half spaces $\mathbf{H}$ and $\mathbf{H}'$. Let us show that they have a common prefix that is also recurring. For this, consider two occurrences $\mathrm{lcp}(C) = \mathbf{H}$ and $\mathrm{lcp}(C') = \mathbf{H}'$ of $\mathbf{H}$ and $\mathbf{H}'$ with $i_h < s(C) < s(C')$. As shown by the thick dashed violet line in Figure 5, and since a stack height of h occurs infinitely often, there must be two downward paths $C = C_1, \dots, C_n$ resp. $C' = C_{n+m}, \dots, C_n$ from C resp. C' to a single cycle C_n. Thus the sequence $C = C_1, \dots, C_n, \dots, C_{n+m} = C'$ is such that, for all $1 \leq i < n+m$, $\mathrm{lcp}(C_i)$ and $\mathrm{lcp}(C_{i+1})$ are comparable for the prefix ordering. The set $\{\mathrm{lcp}(C_i) : 1 \leq i \leq n+m\}$ is a finite meet-semilattice for the prefix ordering, thus with a bottom element $\mathbf{G} \leq^{\text{pref}} \mathbf{H}, \mathbf{H}'$. As there are infinitely many such pairs of occurrences of the recurring $\mathbf{H}$ and $\mathbf{H}'$ and finitely many different such $\mathbf{G}$ with $\|\mathbf{G}\| \leq |V| \cdot \|E\|$, one of the latter must be recurring.

To conclude the proof, assume now that $\mathbf{P}$ is not recurring and let us derive a contradiction. Note that, for all cycles C with $s(C) \geq i_{\mathbf{P}}$, $\mathbf{P} \leq^{\text{pref}} \mathrm{lcp}(C)$. Since $\mathbf{P}$ is not recurring, there must be two incomparable recurring $\mathbf{H}$ and $\mathbf{H}'$, such that $\mathbf{P} <^{\text{pref}} \mathbf{H}$ and $\mathbf{P} <^{\text{pref}} \mathbf{H}'$; we shall further assume that $\mathbf{H}$ is minimal in length with this property. By the previous argument, they have a common prefix $\mathbf{G} <^{\text{pref}} \mathbf{H}, \mathbf{H}'$, which is also recurring, and which we shall also assume minimal in length. Since $\mathbf{H}$ was chosen minimal, there is no recurring $\mathbf{G}'$ incomparable with $\mathbf{G}$, and since $\mathbf{G}$ is minimal, there is no recurring $\mathbf{G}' <^{\text{pref}} \mathbf{G}$ either, hence there exists an index i such that $\mathbf{G} = \mathrm{lcp}_{s(C) \geq i}(\mathrm{lcp}(C))$. As $\mathbf{P} \leq^{\text{pref}} \mathbf{G}$ and $\mathbf{P}$ was chosen of maximal length with this property, $\mathbf{P} = \mathbf{G}$ is recurring. □

Let us conclude the proof of Lemma IV.4. Write $\mathbf{P}$ as $(\mathbf{p}_1, \dots, \mathbf{p}_{|\mathbf{P}|})$. For all cycles C with $s(C) \geq i_{\mathbf{P}}$, $\mathbf{P} \leq^{\text{pref}} \mathrm{lcp}(C)$ shows that $\mathbf{w}(C) \cdot \mathbf{P} \preceq \mathbf{0}$. There are then $|\mathbf{P}|+1$ cases for such cycles C: either there is $1 \leq k \leq |\mathbf{P}|$ with $\mathbf{w}(C) \cdot \mathbf{p}_k < 0$ and $\mathbf{w}(C) \cdot \mathbf{p}_\ell = 0$ for all $1 \leq \ell \leq k$, or $\mathbf{w}(C) \cdot \mathbf{P} = \mathbf{0}$. By Claim IV.5, the $|\mathbf{P}|$ first cases occur (cumulatively) infinitely often; let k^* with $1 \leq k^* \leq |\mathbf{P}|$ be the smallest that does. Then, as there are only finitely many occurrences of cases $k < k^*$, and finitely many $\mathbf{w}_j$ and $\mathbf{H}_j$ not taken into account in the set of cycles C with $s(C) \geq i_{\mathbf{P}}$, $(\mathbf{p}_1, \dots, \mathbf{p}_{k^*})$ is a witness for Definition II.3: Player 2 wins the play. □

By Theorem III.8, Theorem IV.2 and the proof of Lemma IV.3, we now have the following improvement over [12, Corollary 3.2].

Corollary IV.6. *(i) There is an algorithm for solving bounding games whose running time is in* $(|V| \cdot \|E\|)^{O(d^3)}$.
(ii) Player 2 has positional winning strategies for bounding games.

V. Multi-Dimensional Energy Parity Games

In this section, we define multi-dimensional energy parity games (as introduced in [7]) as well as extended multi-dimensional energy games (from [5]), and show how to solve them with an arbitrary (Corollary V.5) or a given (Corollary V.7) initial credit.

A. Multi-Dimensional Energy Parity Games

The *multi-dimensional energy parity games* are played on finite multi-weighted game graphs (V, E, d) enriched with a *priority function* $\pi : V \to \mathbb{N}_{>0}$; we let p be the number of distinct even priorities. Given an *initial credit* $\mathbf{c} \in \mathbb{N}^d$, we denote by $\mathtt{EnPar}_{\mathbf{c}}(V, E, d, p)$ the multi-dimensional energy parity game where Player 1 wins a play $v_0 \xrightarrow{\mathbf{w}_1} v_1 \xrightarrow{\mathbf{w}_2} v_2 \cdots$ if it satisfies

- the *energy objective*: for all $i > 0$, her energy level at step i is non-negative on all components: $\mathbf{c} + \sum_{j \leq i} \mathbf{w}_j \geq \mathbf{0}$, where comparisons are taken componentwise, and
- the *parity objective*: the least priority $\pi(v_i)$ that appears infinitely often is odd;

Player 2 wins the play otherwise. A multi-dimensional energy game ignores the parity condition—equivalently $\pi(v) = 1$ for all $v \in V$.

Example V.1. Let us consider once more the graph of Example II.1. Player 2 wins the energy game with any initial credit: if Player 1 eventually uses only the left (blue) loop, then the second component will eventually become negative, and similarly for the right (violet) loop and the first component. Hence she must switch infinitely often between her two vertices using the middle (cyan) loop, but this decreases the 1-norm of her current energy level. □

Example V.2. Consider the 1-weighted game graph with priorities of Figure 6. Player 1 is losing for all initial credits c

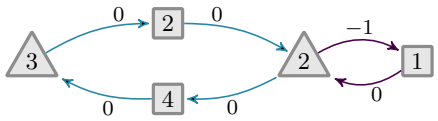

Fig. 6. A 1-weighted game graph with priorities.

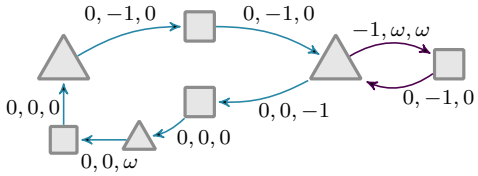

Fig. 7. An extended 3-weighted game graph encoding Figure 6.

in this game: due to the energy objective the violet loop on the right can be played at most c times, and eventually only the cyan loop on the left will be played, but then the parity objective is not satisfied by the play. □

B. Extended Multi-Dimensional Energy Games

Extended multi-dimensional energy games allow special weights (denoted by 'ω') that let Player 1 choose any value she wants for the component. Formally, let $\mathbb{Z}_\omega \stackrel{\text{def}}{=} \mathbb{Z} \uplus \{\omega\}$; in infinity norms of the extended multi-weights, ω is treated as 1. An *extended (finite) multi-dimensional weighted game graph* (V, E, d), where $E \subseteq (V_1 \times \mathbb{Z}_\omega^d \times V_2) \cup (V_2 \times \mathbb{Z}^d \times V_1)$. A play on such a graph is an infinite sequence $v_0 \xrightarrow{\mathbf{w}_1} v_1 \xrightarrow{\mathbf{w}_2} v_2 \cdots$ such that $v_i \xrightarrow{\mathbf{u}_{i+1}} v_{i+1} \in E$ for all $0 \leq i$ and $\mathbf{w}_{i+1}$ instantiates $\mathbf{u}_{i+1}$ by replacing ω's with values from $\mathbb{N}$; strategies for Player 1 now have to specify how to instantiate ω's to form plays. Using the energy objective as before to determine winners of plays, we obtain the *extended multi-dimensional energy game* $\texttt{ExtEn}_{\mathbf{c}}(V, E, d)$ where $\mathbf{c}$ is the initial credit [5].

The following proposition shows how to get rid of priorities in multi-dimensional energy parity games at the price of extra dimensions and the use of extended games: each even priority is associated with an extra dimension, which is decremented by one upon entering a vertex with this priority, and incremented by ω upon entering a vertex with a smaller odd priority (a pair of additional vertices might need to be introduced if the originating vertex was a Player 2 vertex); see Figure 7 for the extended multi-weighted game thus constructed from Figure 6.

Fact V.3 (Jančar [11, Lemma 1])**.** *Let (V, E, d) be a weighted game graph, $v \in V$ an initial vertex, π a priority function with p distinct even priorities, and $\mathbf{c} \in \mathbb{N}^d$ an initial credit. We can construct in logarithmic space an extended weighted game graph $(V', E', d+p)$ with $V \subseteq V'$, $|V'| \leq 3|V|$, $|E'| \leq |E| + 2|V|$, and $\|E'\| = \|E\|$ such that:*

(i) Player 1 wins $\texttt{EnPar}_{\mathbf{c}}(V, E, d, p)$ from v if and only if there exists $\mathbf{c}' \in \mathbb{N}^p$ such that she wins $\texttt{ExtEn}_{\mathbf{cc}'}(V', E', d+p)$ from v, and

(ii) Player 2 wins $\texttt{EnPar}_{\mathbf{c}}(V, E, d, p)$ from v if and only if for all $\mathbf{c}' \in \mathbb{N}^p$ he wins $\texttt{ExtEn}_{\mathbf{cc}'}(V', E', d+p)$ from v.

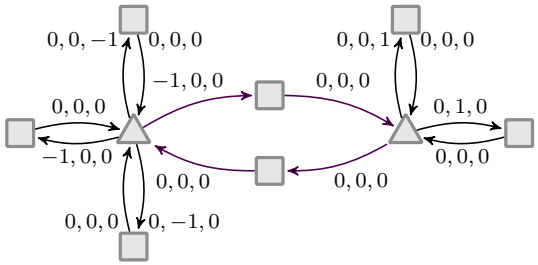

Fig. 8. Part of the translation of Figure 7 for bounding games.

C. Arbitrary Initial Credit

We show how extended multi-dimensional energy games, in case the initial credit is existentially quantified for Player 1 (this is the *arbitrary initial credit problem*), can be solved efficiently by translating them to bounding games. The ideas behind the translation are simple: enable Player 1 to keep the energy bounded at all times by artificial decreasing self-loops, and to instantiate ω weights arbitrarily by encoding them as increasing self-loops. However, the proof of correctness (cf. Proposition V.4) is unexpectedly non-trivial and makes use of perfect half space games.

The translation is an extension of the translation in [12, Section 2.3], which did not handle ω weights, and performs the following:

- at every vertex owned by Player 1 and for every coordinate i, a self-loop is inserted whose weight is the negative unit vector $-\mathbf{e}_i$ (these make use of new dummy Player 2 vertices, to meet the requirements of player alternation and weight determinacy);
- for every edge whose weight $\mathbf{u}$ is not in $\mathbb{Z}^d$, all ω coordinates in $\mathbf{u}$ are replaced by 0, and then a dummy Player 2 vertex is inserted succeeded by a new Player 1 vertex that has a self-loop of weight $\mathbf{e}_i$ for each coordinate i that was ω in $\mathbf{u}$ (the latter make use of further dummy Player 2 vertices as before).

Figure 8 illustrates this construction on the right violet loop of the graph of Figure 7.

Proposition V.4. *Let (V, E, d) be an extended multi-dimensional weighted game graph and $v \in V$ an initial vertex. We can construct in logarithmic space a multi-dimensional weighted game graph $(V^\dagger, E^\dagger, d)$ with $V \subseteq V^\dagger$, $|V^\dagger| \leq (d+1)|V| + (d+2)|E|$, $|E^\dagger| \leq 2(d+1)|E| + 2d|V|$, and $\|E^\dagger\| = \|E\|$ such that:*

(i) Player 1 wins $\texttt{ExtEn}_{\mathbf{c}}(V, E, d)$ from v for some $\mathbf{c} \in \mathbb{N}^d$ if and only if she wins $\texttt{Bnd}(V^\dagger, E^\dagger, d)$ from v, and

(ii) Player 2 wins $\texttt{ExtEn}_{\mathbf{c}}(V, E, d)$ from v for all $\mathbf{c} \in \mathbb{N}^d$ if and only if he wins $\texttt{Bnd}(V^\dagger, E^\dagger, d)$ from v.

Proof. Regarding the right-to-left implication in item (i), by [12, Lemma 3.1], Player 1 has a winning strategy σ in $\texttt{Bnd}(V^\dagger, E^\dagger, d)$ from v that ensures that all total weights are at most $B \stackrel{\text{def}}{=} (4|V^\dagger| \cdot \|E^\dagger\|)^{2(d+2)^3}$. In particular, σ does not get stuck in any of the artificial self-loops. Hence, σ gives rise to a winning strategy in $\texttt{ExtEn}_{(B,\ldots,B)}(V, E, d)$ from v.

By the determinacy of bounding games (cf. Theorem IV.2, Theorem III.8 and Theorem III.2), it now suffices to establish the right-to-left implication in item (ii). Let τ be a winning strategy of Player 2 in the perfect half space game $\mathtt{PHS}(\widehat{V^\dagger}, \widehat{E^\dagger}, d)$, that is positional and perfect half space oblivious, and let $\overline{\tau}$ be its projection onto the extended graph (V, E, d).

Consider any $\mathbf{c} \in \mathbb{N}^d$, any play π in $\mathtt{ExtEn}_{\mathbf{c}}(V, E, d)$ from v that is consistent with $\overline{\tau}$, and let $\pi^\dagger$ be a play in $\mathtt{PHS}(\widehat{V^\dagger}, \widehat{E^\dagger}, d)$ from v that corresponds to π (i.e., where the instantiations of ω weights in π are reproduced by the corresponding increasing self-loops). Observe that any Player 2 vertex v' in π also occurs in $\widehat{\pi^\dagger}$, and so the perfect half space chosen by τ at v' must contain every negative unit vector $-\mathbf{e}_i$ (otherwise, Player 1 could proceed to win by repeating forever one of the artificial self-loops at the successor of v'), i.e., be disjoint from the non-negative orthant $\mathbb{Q}^d_{\geq 0}$.

Since τ is winning, there exists a partially perfect half space $(\mathbf{g}_1, \ldots, \mathbf{g}_k)$ which is a prefix of all perfect half spaces that are chosen by τ along a suffix of $\widehat{\pi^\dagger}$, and there exist a_1, b_1, $\ldots$, a_{k-1}, b_{k-1} such that:

- the dot products of the total weights along $\widehat{\pi^\dagger}$ with $\mathbf{g}_k$ are unbounded below, and
- for every $\ell = 1, \ldots, k-1$, the dot products of the total weights along $\widehat{\pi^\dagger}$ with $\mathbf{g}_\ell$ are in the interval $[a_\ell, b_\ell]$.

Hence, for the sequence of total weights along π with $\mathbf{c}$ subtracted, the same holds. But, by the observation above, the denotation of $(\mathbf{g}_1, \ldots, \mathbf{g}_k)$ is disjoint from the non-negative orthant, implying that

$$\{\mathbf{x} \cdot \mathbf{g}_k \,:\, \mathbf{c} + \mathbf{x} \geq \mathbf{0} \text{ and } \forall 1 \leq \ell < k,\, \mathbf{x} \cdot \mathbf{g}_\ell \in [a_\ell, b_\ell]\}$$

is bounded below. We conclude that the total weights along π are not contained in $\mathbb{N}^d - \mathbf{c}$, showing as required that $\overline{\pi}$ is a winning strategy in $\mathtt{ExtEn}_{\mathbf{c}}(V, E, d)$. □

From Corollary IV.6, Fact V.3 and Proposition V.4, we obtain our first improved upper bound.

Corollary V.5. *The arbitrary initial credit problem for multi-dimensional energy parity games on (V, E, d) with p even priorities is solvable in time $(|V| \cdot \|E\|)^{O((d+p)^3 \log(d+p))}$.*

We also deduce that Player 2 has positional winning strategies in multi-dimensional energy parity games with arbitrary initial credit; this could already be derived by Fact V.3 from the case of extended energy games with arbitrary initial credit, shown in Lemma 19 in the arXiv version of [5].

D. Given Initial Credit

The *given initial credit problem* for multi-dimensional energy parity games takes as input a multi-weighted game graph (V, E, d), a priority function π, an initial vertex v, and an initial credit $\mathbf{c}$ in $\mathbb{N}^d$ and asks whether Player 1 wins the multi-dimensional energy parity game $\mathtt{EnPar}_{\mathbf{c}}(V, E, d, \pi)$ from v.

Following [12, Lemma 3.4], we show that any multi-dimensional energy parity game with a given initial credit is equivalent to a bounding game played over a doubly-exponentially larger graph in terms of d, and exponentially larger in terms of p.

Lemma V.6. *be a multi-weighted game graph, π a priority function with p distinct even priorities, and $v \in V$. One can construct in time $O(|V^\ddagger| \cdot |E| + d \cdot \log \|\mathbf{c}\|)$ a multi-weighted game graph $(V^\ddagger, E^\ddagger, d+p)$ and a vertex $v_{\mathbf{c}}$ in $V^\ddagger$, where $|V^\ddagger|$ is in $(|V| \cdot \|E\|)^{2^{O(d \log(d+p))}}$ and $\|E^\ddagger\| = \|E\|$ such that, for all $i \in \{1, 2\}$, Player i wins the multi-dimensional energy parity game $\mathtt{EnPar}_{\mathbf{c}}(V, E, d, \pi)$ from v if and only if Player i wins the bounding game $\mathtt{Bnd}(V^\ddagger, E^\ddagger, d+p)$ from $v_{\mathbf{c}}$.*

Proof sketch. We use the same arguments as in the proof of [12, Lemma 3.4]. The only difference is that we need to handle the parity condition, and thus to go through extended multi-dimensional energy games and replace [12, Proposition 2.2] with the combination of Fact V.3 and Proposition V.4. These only incur a polynomial overhead in the size of the weighted game graphs, hence a bound for $|V^\ddagger|$ in $(|V| \cdot \|E\|)^{2^{O(d^\ddagger \log d^\ddagger)}}$ with $d^\ddagger \stackrel{\text{def}}{=} d+p$ can be deduced directly from [12, Lemma 3.4].

We refine this bound by observing that only the first d components of the $(d+p)$-dimensional bounding game we construct should be treated as initialised, while the p remaining ones in Fact V.3 are arbitrary, hence the blowing-up construction of [12, Lemma 3.4] only needs to be applied d times, yielding instead a bound in $(|V| \cdot \|E\|)^{2^{O(d \log d^\ddagger)}}$; see Eq. (9) in the arXiv version of [12]. □

By applying Corollary IV.6 to the game graph $(V^\ddagger, E^\ddagger, d+p)$ and since $|E| \leq |V|^2$, we obtain a 2-EXPTIME upper bound on the given initial credit problem, which is again pseudo-polynomial when d and p are fixed.

Corollary V.7. *The given initial credit problem with initial credit $\mathbf{c}$ for multi-dimensional energy parity games on (V, E, d) with p even priorities is solvable in time*

$$(|V| \cdot \|E\|)^{2^{O(d \cdot \log(d+p))}} + O(d \cdot \log \|\mathbf{c}\|)\ .$$

This matches the 2-EXPTIME lower bound from [10], and generalises [12, Theorem 3.5] to multi-dimensional energy parity games. Because the given initial credit problem for energy games of fixed dimension $d \geq 4$ and number of even priorities $p = 0$ is already EXPTIME-hard [10], there is no hope of improving the pseudo-polynomial bound in Corollary V.7 to a polynomial one.

VI. Concluding Remarks

In this paper, we have shown a chain of reductions and strategy transfers from multi-dimensional energy parity games to perfect half space games and lexicographic energy games, see Figure 1.

There are two main outcomes. On the complexity side, we obtain tighter upper bounds for multi-dimensional energy parity games, both with arbitrary and given initial credit. In particular, in addition to closing the complexity gap with given initial credit, our 2-EXPTIME upper bound in Corollary V.7 also

closes complexity gaps for several problems already mentioned in the introduction:

- deciding extended multi-dimensional energy games with given initial input [5],
- deciding whether a Petri net weakly simulates a finite state system, or satisfies a formula of the μ-calculus fragment defined in [1], and
- deciding the model-checking problem for RB$\pm$ATL [2].

The second outcome is a rather precise description of the winning strategies for Player 2 in these games. Here, the perfect half space viewpoint is especially enlightening: Player 2 can win by 'announcing' in which perfect half spaces it is attempting to escape.

APPENDIX

A. Proof of Theorem III.2

Lemma III.5. *If the value of the mean-payoff game* $\mathtt{MP}(V, E^{(1)})$ *is non-negative (negative, resp.) at a vertex* v*, then by using a positional optimal strategy from that mean-payoff game, Player* 1 *(Player* 2*, resp.) wins the corresponding lexicographic energy game* $\mathtt{LexEn}(V, E, d)$ *from* v*.*

Proof. We prove the lemma for Player 2 (in mean-payoff terminology, Min); the argument for Player 1 (Max) is analogous. For this, let us fix ourselves a positional optimal strategy for Min in the mean-payoff game $(V, E^{(1)})$. We show that this strategy is also winning for Player 2 in the lexicographic energy game $\mathtt{LexEn}(V, E, d)$. Hence, in the rest of the proof, we consider a play P consistent with this strategy in $\mathtt{LexEn}(V, E, d)$, and we aim at showing that it is winning.

Let $C_1, C_2, \ldots$ be the infinite sequence of simple cycles obtained by the '*cycle decomposition*' of the play P: we start with an empty sequence of cycles, we then push successive vertices of the play on a stack, and each time we push a vertex that is already present on the stack, we pop the resulting simple cycle from the top of the stack and add it to the sequence of simple cycles. Observe ($\star$) that every simple cycle $C_1, C_2, \ldots$ has total multi-weight $\prec \mathbf{0}$. Indeed, as a cycle in the strategy subgraph of an optimal strategy for Min in the mean-payoff game $\mathtt{MP}(V, E^{(1)})$ with a negative value, it has a negative total weight [15], and hence the observation ($\star$) follows by Proposition III.4.

For a cycle C in the multi-weighted game graph (V, E, d), call the *leading dimension* the least $k = 1, \ldots, d$ such that $\mathbf{w}(C)(k) \neq 0$ (recall that $\mathbf{w}(C)$ is the total multi-weight of the edges in the cycle). The leading dimension k^* of the play P is the smallest dimension that is the leading dimension of infinitely many cycles $C_1, C_2, \ldots$; note that in the proof for Player 1, k^* can equal $d + 1$.

The core of the proof is now contained in the following claims A.1 and A.2.

Claim A.1. *For all* $1 \leq \ell < k^*$, $\liminf_n \sum_{i=1}^n \mathbf{w}(C_i)(\ell) < +\infty$.

Proof of Claim A.1. Indeed, by the definition of k^*, for all $\ell < k^*$, we have that $\mathbf{w}(C_i)(\ell) = 0$ for all sufficiently large i. Hence the sequence of sums $\sum_{i=1}^n \mathbf{w}(C_i)(\ell)$ is eventually constant, so its inferior limit is finite. □

Claim A.2. *If* $k^* \leq d$*, then* $\limsup_n \sum_{i=1}^n \mathbf{w}(C_i)(k^*) = -\infty$.

Proof of Claim A.2. Indeed, from the definition of k^* and the fact that every cycle in the decomposition has total weight $\prec \mathbf{0}$, we have that $\mathbf{w}(C_i)(k^*)$ is:

- ≤ -1 for infinitely many i;
- ≤ 0 for all sufficiently large i.

The sequence of sums $\sum_{i=1}^n \mathbf{w}(C_i)(k)$ therefore has limit superior $-\infty$. □

From the two above claims, we get that the play P is won by Min. The reader may worry that the expressions of the form '$\sum_{i=1}^n \mathbf{w}(C_i)(\ell)$' differ from those in the definition of lexicographic energy games: there might be a non-empty simple path remaining indefinitely 'on the stack' of the cycle decomposition, and thus not taken into account. However, if we want just to determine whether the corresponding limit inferior (superior, resp.) is less than $+\infty$ (equal to $-\infty$, resp.), then the discrepancy is benign because, for every simple path P', we have $|\mathbf{w}(P')(k)| \leq |V| \cdot \|E\|$. □

B. Proof of Lemma III.7

Lemma III.7. *The winning strategies of Player* i*,* $i \in \{1, 2\}$*, are the same in* $\mathtt{PHS}(\widehat{V}, \widehat{E}, d)$ *and* $\mathtt{LexEn}(\widehat{V}, \widetilde{E}, 2d)$.

Proof. Consider any infinite play

$$P \stackrel{\text{def}}{=} (v_1, \mathbf{H}_1) \xrightarrow{\mathbf{w}_1} (v_2, \mathbf{H}_2) \xrightarrow{\mathbf{w}_2} \cdots$$

in the perfect half space game $\mathtt{PHS}(\widehat{V}, \widehat{E}, d)$ along with its corresponding play

$$\widetilde{P} \stackrel{\text{def}}{=} (v_1, \mathbf{H}_1) \xrightarrow{\mathbf{e}_{\mathbf{H}_1,\mathbf{H}_2} \Delta(\mathbf{w}_1 \cdot \mathbf{H}_1)} (v_2, \mathbf{H}_2) \xrightarrow{\mathbf{e}_{\mathbf{H}_2,\mathbf{H}_3} \Delta(\mathbf{w}_2 \cdot \mathbf{H}_2)} \cdots$$

in the lexicographic energy game $\mathtt{LexEn}(\widehat{V}, \widetilde{E}, 2d)$.

We show that P is winning for Player i in $\mathtt{PHS}(\widehat{V}, \widehat{E}, d)$ if and only if $\widetilde{P}$ is winning for the same player in $\mathtt{LexEn}(\widehat{V}, \widetilde{E}, 2d)$. This in turn entails that winning strategies for each player can be transferred between the two games.

It suffices to show this for Player 2. If $\widetilde{P}$ is winning for Player 2, then there exists $1 \leq k \leq 2d$ such that $\limsup_n \sum_{j=1}^n (\mathbf{e}_{\mathbf{H}_j,\mathbf{H}_{j+1}} \Delta \mathbf{w}_j \cdot \mathbf{H}_j)(k) = -\infty$ and for all $1 \leq \ell < k$, $\liminf_n \sum_{j=1}^n (\mathbf{e}_{\mathbf{H}_j,\mathbf{H}_{j+1}} \Delta\, \mathbf{w}_j \cdot \mathbf{H}_j)(\ell) < +\infty$. Since the coefficients $\mathbf{e}_{\mathbf{H}_j,\mathbf{H}_{j+1}}$ are all non-negative, k cannot correspond to one of these dimensions. Hence k is even; let $\hat{k} \stackrel{\text{def}}{=} k/2$. Because $\liminf_n \sum_{j=1}^n (\mathbf{e}_{\mathbf{H}_j,\mathbf{H}_{j+1}} \Delta\, \mathbf{w}_j \cdot \mathbf{H}_j)(\ell) < +\infty$ for all *odd* $1 \leq \ell < k$, we deduce that the visited perfect half spaces $\mathbf{H}_1, \mathbf{H}_2, \ldots$ differ on their first $\hat{k}$ coordinates only finitely many times. Hence there is an infinite suffix of the play where all the perfect half spaces share a common prefix $(\mathbf{g}_1, \ldots, \mathbf{g}_{\hat{k}})$. Then $\limsup_n \sum_{j=1}^n \mathbf{w}_j \cdot \mathbf{g}_{\hat{k}} = \limsup_n \sum_{j=1}^n (\mathbf{e}_{\mathbf{H}_j,\mathbf{H}_{j+1}} \Delta\, \mathbf{w}_j \cdot \mathbf{H}_j)(2\hat{k}) = -\infty$ and for all $1 \leq \hat{\ell} < \hat{k}$, $\liminf_n \sum_{j=1}^n \mathbf{w}_j \cdot \mathbf{g}_{\hat{\ell}} = \liminf_n \sum_{j=1}^n (\mathbf{e}_{\mathbf{H}_j,\mathbf{H}_{j+1}} \Delta\, \mathbf{w}_j \cdot \mathbf{H}_j)(2\hat{\ell}) < +\infty$, hence P is also winning for Player 2 in the perfect half space game.

Conversely, if P is winning for Player 2 in the perfect half space game $\mathtt{PHS}(\widehat{V}, \widehat{E}, d)$, then there is an infinite suffix starting at some index i with $\mathbf{G} \preceq \mathrm{lcp}_{j \geq i} \mathbf{H}_j$ satisfying Definition II.3, and let $k \stackrel{\text{def}}{=} |\mathbf{G}|$. The k first odd coordinates of the weights in the corresponding infinite suffix in $\widetilde{P}$ are thus all 0, hence the energy will not diverge on these coordinates. Furthermore, the k first even coordinates in the same suffix are such that $\limsup_n \sum_{j=i}^{n} (\mathbf{e}_{\mathbf{H}_j, \mathbf{H}_{j+1}} \Delta \mathbf{w}_j \cdot \mathbf{H}_j)(2k) = \limsup_n \sum_{j=i}^{n} \mathbf{w}_j \cdot \mathbf{g}_k = -\infty$ and, for all $1 \leq \ell < k$, $\liminf_n \sum_{j=i}^{n} (\mathbf{e}_{\mathbf{H}_j, \mathbf{H}_{j+1}} \Delta \mathbf{w}_j \cdot \mathbf{H}_j)(2\ell) = \liminf_n \sum_{j=i}^{n} \mathbf{w}_j \cdot \mathbf{g}_\ell < +\infty$. Thus $\widetilde{P}$ is also winning for Player 2 in $\mathtt{LexEn}(\widehat{V}, \widetilde{E}, 2d)$. □

C. Proof of Theorem III.8(ii)

Key Remark: For a strategy τ' of Player 2, we say that a path is (τ', v)-*elementary path* if it is consistent with τ', it starts in some $(v, \mathbf{H})$, ends in some $(v, \mathbf{H}')$, and does not visit the vertex v in between. Consider a (τ, v)-elementary path P starting in $(v, \mathbf{H})$ and ending in $(v, \mathbf{H}')$. Then there is a $(\tau_{\mathbf{H}}, v)$-elementary path $P^{\mathbf{H}'}$ that is exactly like P but for the fact that it begins in $(v, \mathbf{H}')$. This one happens to be a *cycle* consistent with τ_H. Then we clearly have $\mathrm{lcp}(P) \leq^{\text{pref}} \mathrm{lcp}(P^{\mathbf{H}'})$, since every perfect half space that occurs in $P^{\mathbf{H}'}$ already occurs in P. Since furthermore $\mathbf{w}(P^{\mathbf{H}'}) = \mathbf{w}(P)$, this means that if P is winning for Player 2, then the same holds for $P^{\mathbf{H}'}$.

Claim A.3. *If a perfect half space* $\mathbf{H}$ *is bad then there exists a* (τ, v)*-elementary path starting in* $(v, \mathbf{H})$ *and losing for Player* 2*.*

Proof of Claim A.3. Indeed, if $\mathbf{H}$ is bad, there exists a play resulting from playing a strategy for Player 1 from $(v, \mathbf{H})$ against $\tau_{\mathbf{H}}$, which is winning for Player 1; by Theorem III.2 and Lemma III.7 we can assume this strategy to be positional. Two cases may happen: Either this play never visits v (except at the initial position). In this case, this play was already a play consistent with τ, contradicting the fact that the strategy τ was winning from $(v, \mathbf{H})$. Otherwise, the infinite play encounters at least once more some vertex $(v, \mathbf{H}')$. Let P be the prefix of the play from $(v, \mathbf{H})$ to $(v, \mathbf{H}')$. This is a (τ, v)-elementary path. Since P has been obtained from the fight of a positional strategy for Player 1 against $\tau_{\mathbf{H}}$, the infinite play ultimately repeats the cycle $P^{\mathbf{H}'}$. Thus $P^{\mathbf{H}'}$ is losing for Player 2. According to the above key remark, P was thus already losing for Player 2. □

Claim III.9. *There exists a good perfect half space.*

Proof of Claim III.9. Assume for the sake of contradiction that all perfect half spaces are bad. We shall prove that in this case τ was losing from $(v, \mathbf{H})$. Let us fix for all perfect half spaces $\mathbf{H}$ a (τ, v)-elementary path $P(\mathbf{H})$ starting from $(v, \mathbf{H})$ ending in some $(v, f(\mathbf{H}))$ and losing for Player 2 (it exists according to Claim A.3). Let us now construct a play consistent with τ starting from $(v, \mathbf{H})$ as follows: assuming the partial play constructed so far ends in $(v, \mathbf{H})$, we extend it by concatenating the path $P(\mathbf{H})$ to it, yielding a longer play ending in $(v, f(\mathbf{H}))$. We iterate this process and, going to the limit, we obtain an infinite play P consistent with τ. However, this play is an infinite concatenation of finitely many $P(\mathbf{H})$ paths, which are all losing for Player 2. Hence P is losing for Player 2. This contradicts the fact that τ was assumed to be winning from $(v, \mathbf{H})$. The claim is proved: there has to be a good perfect half space. □

Acknowledgements

Work funded in part by the EPSRC grants EP/M011801/1 and EP/P020992/1, and the ERC grant 259454 (GALE). The authors thank Joël Ouaknine and Prakash Panangaden for organising and hosting the Infinite-State Systems workshop at Bellairs Research Institute in March 2015, where this work started.

References

[1] P. A. Abdulla, R. Mayr, A. Sangnier, and J. Sproston, "Solving parity games on integer vectors," in *Concur 2013*, ser. LNCS, vol. 8052. Springer, 2013, pp. 106–120.

[2] N. Alechina, N. Bulling, S. Demri, and B. Logan, "On the complexity of resource-bounded logics," in *RP 2016*, ser. LNCS, vol. 9899. Springer, 2016, pp. 36–50.

[3] N. Alechina, N. Bulling, B. Logan, and H. N. Nguyen, "The virtues of idleness: A decidable fragment of resource agent logic," *Artif. Intell.*, 2017, to appear.

[4] R. Bloem, K. Chatterjee, T. A. Henzinger, and B. Jobstmann, "Better quality in synthesis through quantitative objectives," in *CAV 2009*, ser. LNCS, vol. 5643. Springer, 2009, pp. 140–156.

[5] T. Brázdil, P. Jančar, and A. Kučera, "Reachability games on extended vector addition systems with states," in *ICALP 2010*, ser. LNCS, vol. 6199. Springer, 2010, pp. 478–489, arXiv version available from http://arxiv.org/abs/1002.2557.

[6] V. Bruyère, N. Meunier, and J. Raskin, "Secure equilibria in weighted games," in *CSL-LICS 2014*, 2014, pp. 26:1–26:26.

[7] K. Chatterjee, M. Randour, and J.-F. Raskin, "Strategy synthesis for multi-dimensional quantitative objectives," *Acta Inf.*, vol. 51, no. 3–4, pp. 129–163, 2014.

[8] T. Colcombet and D. Niwiński, "Lexicographic energy games," Manuscript, 2017.

[9] C. Comin and R. Rizzi, "Improved pseudo-polynomial bound for the value problem and optimal strategy synthesis in mean payoff games," *Algorithmica*, 2016, to appear.

[10] J. Courtois and S. Schmitz, "Alternating vector addition systems with states," in *MFCS 2014*, ser. LNCS, vol. 8634. Springer, 2014, pp. 220–231.

[11] P. Jančar, "On reachability-related games on vector addition systems with states," in *RP 2015*, ser. LNCS, vol. 9328. Springer, 2015, pp. 50–62.

[12] M. Jurdziński, R. Lazić, and S. Schmitz, "Fixed-dimensional energy games are in pseudo-polynomial time," in *ICALP 2015*, ser. LNCS, vol. 9135. Springer, 2015, pp. 260–272, arXiv version available from https://arxiv.org/abs/1502.06875.

[13] A. Puri, "Theory of hybrid and discrete systems," Ph.D. dissertation, University of California, Berkeley, 1995.

[14] Y. Velner, K. Chatterjee, L. Doyen, T. A. Henzinger, A. Rabinovich, and J. Raskin, "The complexity of multi-mean-payoff and multi-energy games," *Inform. and Comput.*, vol. 241, pp. 177–196, 2015.

[15] U. Zwick and M. Paterson, "The complexity of mean payoff games on graphs," *Theor. Comput. Sci.*, vol. 158, no. 1, pp. 343–359, 1996.