



HAL
open science

Reachability games with relaxed energy constraints

Loïc Hélouët, Nicolas Markey, Ritam Raha

► **To cite this version:**

Loïc Hélouët, Nicolas Markey, Ritam Raha. Reachability games with relaxed energy constraints. 2020. hal-02964729

HAL Id: hal-02964729

<https://inria.hal.science/hal-02964729>

Preprint submitted on 12 Oct 2020

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Reachability games with relaxed energy constraints

Loïc Hélouët^a, Nicolas Markey^a, Ritam Raha^{b,c}

^a*Inria & CNRS & Univ. Rennes, France*

^b*Univ. Antwerpen, Belgium*

^c*LaBRI, Univ. Bordeaux, France*

Abstract

We study games with reachability objectives under energy constraints. We first prove that under strict energy constraints (either only lower-bound constraint or interval constraint), those games are LOGSPACE-equivalent to energy games with the same energy constraints but without reachability objective (i.e., for infinite runs). We then consider two relaxations of the upper-bound constraints (while keeping the lower-bound constraint strict): in the first one, called *weak upper bound*, the upper bound is *absorbing*, i.e., when the upper bound is reached, the extra energy is not stored; in the second one, we allow for *temporary violations* of the upper bound, imposing limits on the number or on the amount of violations.

We prove that when considering weak upper bound, reachability objectives require memory, but can still be solved in polynomial-time for one-player arenas; we prove that they are in coNP in the two-player setting. Allowing for bounded violations makes the problem PSPACE-complete for one-player arenas and EXPTIME-complete for two players. We then address the problem of existence of bounds for a given arena. We show that with reachability objectives, existence can be a simpler problem than the game itself, and conversely that with infinite games, existence can be harder.

Keywords: Games, Energy

1. Introduction

Weighted graphs are a convenient formalism to address questions related to consumption, production and storage of resources: vertices represent the possible configurations of the system under study, and transitions carry positive or negative integers to represent the evolution of resources. When dealing with open systems, weighted graphs are used as the arenas of two-player turn-based game, modelling the interactions of the system with its hostile environment. Various objectives have been considered for such games, such as optimizing the total or average amount of resources that have been collected along the play, or maintaining the total amount within given bounds. The latter kind

Email addresses: loic.helouet@inria.fr (Loïc Hélouët), nicolas.markey@irisa.fr (Nicolas Markey), ritam.raha@uantwerpen.be (Ritam Raha)

of objectives, usually called *energy objectives* [CdAHS03, BFL⁺08], has been widely studied in the untimed setting [CD12, CDHR10, DDG⁺10, FJLS11, JLR13, JLS15, VCD⁺15, BMR⁺15, BHM⁺17, DM18], and to a lesser extent in the timed setting [BFLM10, BLM12]. Energy objectives can be used to model the evolution of the available energy in an autonomous system that has to achieve its tasks, but also recharge batteries to avoid running out of power. Energy objectives were used to model moulding machines: such machines inject molten plastic into a mould, using pressure obtained by storing liquid in a tank [CJL⁺09]; the level of liquid has to be controlled to maintain sufficient pressure, but excessive pressure in the tank reduces the lifetime of the valve.

Energy games impose strict constraints on the energy level at all stages of the play. Two kinds of constraints have been mainly considered in the literature: lower-bound constraints (a.k.a. L-energy constraints) impose a strict lower bound (usually 0), but no upper bound; on the other hand, lower- and upper-bound constraints (a.k.a. LU-energy constraints) require that the energy level always remains within a bounded interval $[L; U]$. Finding strategies that realize L-energy objectives along infinite runs is in PTIME in the one-player setting, and in $\text{NP} \cap \text{coNP}$ for two players; for LU-energy objectives, it is respectively PSPACE-complete and EXPTIME-complete [BFL⁺08]. Some works also consider the existence of an initial energy level for which a winning strategy exists [CDHR10].

In this paper, we focus on weighted games combining energy objectives with reachability objectives. Our first result is the (expected) proof that L-energy games with or without reachability objectives are interreducible; the same holds for LU-energy games. We then focus on relaxations of the energy constraints, in two different directions. In both cases, the lower bound remains unchanged, as it corresponds to running out of energy, which we always want to avoid. We thus only relax the upper-bound constraint. The first direction concerns *weak upper bounds*, already introduced in [BFL⁺08]: in that setting, hitting the upper bound is allowed, but there is no overload: trying to exceed the upper bound simply maintains the energy level at this maximal level. Following [BFL⁺08], we name these objectives LW-energy objectives. They could be used e.g. as a (simplified) model for batteries. When considered alone, LW-energy objectives are not much different from L-energy objectives, in the sense that the aim is to find a reachable *positive loop*. LW-energy games (with no other objectives besides maintaining the energy above L) are in PTIME for one-player games, and in $\text{NP} \cap \text{coNP}$ for two players [BFL⁺08]. When combining LW-energy and reachability objectives, we prove in this paper that the situation changes: different loops may have different effects on the energy level, and we have to keep track of the final energy level reached when iterating those loops.

We introduce and study a second way of relaxing upper bounds, which we call *soft upper bound*: it consists in allowing a limited number (or amount) of violations of the soft upper bound (possibly within an additional strong upper-bound): when modeling a pressure tank, the lower-bound constraint is strict (pressure should always be available) but the upper bound is soft (excessive pressure may be temporarily allowed if needed). We consider different kinds of restrictions (on the number or amount of violations), and prove that deciding

whether Player 1 has a strategy to keep violations below a given bound is PSPACE-complete for one-player arenas, and EXPTIME-complete for two-player ones. Then we consider the apparently simpler *existence* problem, which aim is to decide whether an upper bound allows to win an energy game. We give complexities of this problem for reachability and infinite games. Surprisingly, depending on the nature of the game, existence can be easier or harder than the corresponding game with a given upper bound.

Related Work: Quantitative games have been considered in many articles since the 1970s, with various kinds of objectives, such as ultimately optimizing the total payoff, mean-payoff [EM79, ZP96], or discounted sum [ZP96, And06]. Energy objectives, which are a kind of safety objectives on the total payoff, were introduced in [CdAHS03] and rediscovered in [BFL⁺08]. Several works have extended those works by combining quantitative conditions together, e.g. multi-dimensional energy conditions [FJLS11, JLS15] or conjunctions of energy- and mean-payoff objectives [CDHR10]. Combinations with qualitative objectives (e.g. reachability [CDH17] or parity objectives [CD12, CRR14, DM18]) were also considered. Similar objectives have been considered in slightly different settings e.g. Vector Addition Systems with States [Rei16] and one-counter machines [HKOW09, GHOW10, Hun15]. This work is an extended version of [HMR19] containing full proofs of our results. The part of the paper addressing the existence problem (Sec. 6) contains new material w.r.t the conference paper.

2. Preliminaries

Definition 1. A two-player turn-based arena is a 3-tuple $G = (Q_1, Q_2, E)$ where $Q = Q_1 \uplus Q_2$ is a set of states, $E \subseteq Q \times \mathbb{Z} \times Q$ is a set of weighted edges. For $q \in Q$, we let $qE = \{(q, w, q') \in E \mid w \in \mathbb{Z}, q' \in Q\}$, which we assume is non-empty for any $q \in Q$. A one-player arena is a two-player arena where $Q_2 = \emptyset$.

Consider a state $q_0 \in Q$. A *finite path* in an arena G from an initial state q_0 is an finite sequence of edges $\pi = (e_i)_{0 \leq i < n}$ such that for every $0 \leq i < n$, writing $e_i = (q_i, w_i, q'_i)$, it holds $q'_i = q_{i+1}$. Fix a path $\pi = (e_i)_{0 \leq i < n}$. Using the notations above, we write $|\pi|$ for the size n of π , $\hat{\pi}_i$ for the i -th state q_i of π (with the convention that $q_n = q'_{n-1}$), and $\text{first}(\pi) = \hat{\pi}_0$ for its first state and $\text{last}(\pi) = \hat{\pi}_n$ for its last state. The empty path is a special finite path from q_0 ; its length is zero, and q_0 is both its first and last state. Given two finite paths $\pi = (e_i)_{0 \leq i < n}$ and $\pi' = (e'_j)_{0 \leq j < n'}$ such that $\text{last}(\pi) = \text{first}(\pi')$, the concatenation $\pi \cdot \pi'$ is the finite path $(f_k)_{0 \leq k < n+n'}$ such that $f_k = e_k$ if $0 \leq k < n$ and $f_k = e'_{k-n}$ if $n \leq k < n+n'$. For $0 \leq k \leq n$, the k -th prefix of π is the finite path $\pi_{<k} = (e_i)_{0 \leq i < k}$. We write $\text{FPaths}(G, q_0)$ for the set of finite paths in G issued from q_0 (we may omit G when it is clear from the context). We define analogously $\text{Paths}(G, q_0)$ for the set of infinite paths from q_0 .

A *strategy* for Player 1 (resp. Player 2) from q_0 is a function $\sigma: \text{FPaths}(q_0) \rightarrow E$ associating with any finite path π with $\text{last}(\pi) \in Q_1$ (resp. $\text{last}(\pi) \in Q_2$) an edge originating from $\text{last}(\pi)$. A strategy is said *memoryless* when $\sigma(\pi) = \sigma(\pi')$ whenever $\text{last}(\pi) = \text{last}(\pi')$. A finite path $\pi = (e_i)_{0 \leq i < n}$ *conforms* to a strategy σ of

Player 1 (resp. of Player 2) from q_0 if $\text{first}(\pi) = q_0$ and for every $0 \leq k < n$, either $e_k = \sigma(\pi_{<k})$, or $\text{last}(\pi_{<k}) \in Q_2$ (resp. $\text{last}(\pi_{<k}) \in Q_1$). This is extended to infinite paths in the natural way. Given a strategy σ of Player 1 (resp. of Player 2) from q_0 , the outcomes of σ is the set of all (finite or infinite) paths π issued from q_0 that conform to σ .

A game is a triple (G, q_0, O) where G is a two-player arena, q_0 is an initial state in Q , and $O \subseteq \text{Paths}(G, q_0)$ is a set of infinite paths, also called *objective* (for Player 1). A strategy for Player 1 from q_0 is winning in (G, q_0, O) if its infinite outcomes all belong to O .

Given a set $R \subseteq Q$ of states, the reachability objective defined by R is the set of all paths containing some state in R , while the safety objective defined by R is the set of all infinite paths never visiting any state in R . In this paper, we also focus on *energy objectives* [CdAHS03, BFL⁺08], which we now define.

Definition 2. Fix a finite-state two-player arena $G = (Q_1, Q_2, E)$, and an extra state $q_{\text{err}} \notin Q_1 \cup Q_2$. Let $L \in \mathbb{Z}$. The L -energy arena associated with G is the infinite-state arena $G_L = (C_1, C_2, T)$ where $C_1 = \{q_{\text{err}}\} \cup Q_1 \times [L; +\infty)$ and $C_2 = Q_2 \times [L; +\infty)$ are sets of configurations, and $T \subseteq C_1 \times \mathbb{Z} \times C_2$ is such that

- for any (q, l) and (q', l') in $Q \times [L; +\infty)$ and any $w \in \mathbb{Z}$, we have $((q, l), w, (q', l')) \in T$ if, and only if, $(q, w, q') \in E$ and $l' = l + w \geq L$; we also impose a loop $(q_{\text{err}}, 0, q_{\text{err}}) \in T$.
- for any $(q, l) \in Q \times [L; +\infty)$, we have $((q, l), w, q_{\text{err}}) \in T$ if, and only if, there is a transition $(q, w, q') \in E$ such that $l + w < L$

Similarly, given $L \in \mathbb{Z}$ and $U \in \mathbb{Z}$, the LU -energy arena associated with G is the finite-state arena $G_{LU} = (C_1, C_2, T)$ where $C_1 = (Q_1 \times [L; U]) \cup \{q_{\text{err}}\}$ and $C_2 = Q_2 \times [L; U]$, and $T \subseteq C_1 \times \mathbb{Z} \times C_2$ is such that

- for any (q, l) and (q', l') in $Q \times [L; U]$ and any $w \in \mathbb{Z}$, we have $((q, l), w, (q', l')) \in T$ if, and only if, $(q, w, q') \in E$ and $l' = l + w \in [L; U]$; we also impose a loop $(q_{\text{err}}, 0, q_{\text{err}}) \in T$.
- for any $(q, l) \in Q \times [L; U]$, we have $((q, l), w, q_{\text{err}}) \in T$ if, and only if, there is a transition $(q, w, q') \in E$ such that $l + w < L$ or $l + w > U$;

Finally, given $L \in \mathbb{Z}$ and $W \in \mathbb{Z}$, the LW -energy arena associated with G is the finite-state arena $G_{LW} = (C_1, C_2, T)$ where $C_1 = \{q_{\text{err}}\} \cup Q_1 \times [L; W]$ and $C_2 = Q_2 \times [L; W]$, and $T \subseteq C_1 \times \mathbb{Z} \times C_2$ is such that

- for any (q, l) and (q', l') in $Q \times [L; W]$ and any $w \in \mathbb{Z}$, we have $((q, l), w, (q', l')) \in T$ if, and only if, $(q, w, q') \in E$ and $l' = \min(W, l + w) \geq L$; we also impose a loop $(q_{\text{err}}, 0, q_{\text{err}}) \in T$.
- for any $(q, l) \in Q \times [L; W]$, we have $((q, l), w, q_{\text{err}}) \in T$ if, and only if, there is a transition $(q, w, q') \in E$ such that $l + w < L$;

An L -run (resp. LU -run, LW -run) ρ in G from q with initial energy level l is a path in G_L (resp. G_{LU} , G_{LW}) from (q, l) never visiting q_{err} . With such a run $\rho = (t_i)_i$ in G , writing $t_i = ((q_i, l_i), w_i, (q'_i, l'_i))$, we associate the path $\pi = (e_i)_i$ such that $e_i = (q_i, w_i, q'_i)$. We define $\hat{\rho}_i = (q_i, l_i)$, corresponding to the i -th configuration along ρ , and $\tilde{\rho}_i = l_i$, which we name the energy level in that configuration.

Similarly, a path π is said L -feasible (resp. LU -feasible, LW -feasible) from initial energy level L if there exists an L -run (resp. LU -run, LW -run) from $(\text{first}(\pi), L)$ whose associated path is π . Notice that if such a run exists, it is unique (because paths are defined as sequences of transitions).

The L -energy (resp. LU -energy, LW -energy) objective is the set of infinite paths that are L -feasible (resp. LU -feasible, LW -feasible) (from initial energy level L). Similarly, given a target set $R \subseteq Q$, the L -energy- (resp. LU -energy-, LW -energy-) reachability objective is the set of L -feasible (resp. LU -feasible, LW -feasible) paths visiting R .

Remark 3. Taking L as the initial energy level results in no loss of generality, since any energy level can be obtained by adding an initial transition from (q_0, L) .

In many cases, strong upper bounds are too strict, as many system do not break as soon as their maximal energy level is reached. Imposing a *weak* upper bound is a way to relax these constraints. We introduce another way to relax energy constraints, by allowing for (limited) violations of the upper bound: given two strict bounds L and U in \mathbb{Z} , a soft upper bound $S \in \mathbb{Z}$ with $L \leq S \leq U$, and an LU -run ρ , the set of violations along ρ is the set $\mathbf{V}(\rho) = \{i \in [0; |\rho|] \mid \tilde{\rho}_i > S\}$ of positions along ρ where the energy level exceeds the soft upper bound S . There are many ways to quantify violations along a run. We consider three of them in this paper, namely the total number of violations, the maximal number of consecutive violations, and the sum of the violations. We thus define the following three quantities: $\#\mathbf{V}(\rho) = |\mathbf{V}(\rho)|$, $\overline{\#\mathbf{V}}(\rho) = \max\{i - j + 1 \mid \forall k \in [i, j]. k \in \mathbf{V}(\rho)\}$, and $\Sigma\mathbf{V}(\rho) = \sum_{i \in \mathbf{V}(\rho)} (\tilde{\rho}_i - S)$.

Figure 1 is an arena for an $LSU^\#$ -energy game, and Figure 2 shows the evolution of $\#\mathbf{V}$ along a winning run in this arena. One can notice that with a strong upper bound of 3, state q_t would not be reachable. On the other hand, if the strong upper bound is set to $U = 6$, and the soft upper bound is set to $S = 3$, then there exists a run from q_0 to q_t , but that requires 3 violations of S , and a total amount of violations of 6.

Given three values $L \leq S \leq U$, and a threshold $V \in \mathbb{N}$, the $LSU^\#$ -energy (resp. $LSU^{\overline{\#}}$ -energy, LSU^Σ -energy) objective is the set of LU -feasible infinite paths π such that, along their associated runs ρ from (q_0, L) , the number $\#\mathbf{V}(\rho)$ of violations (resp. maximal number of consecutive violations $\overline{\#\mathbf{V}}(\rho)$, sum $\Sigma\mathbf{V}(\rho)$ of violations) of the soft upper bound S is at most V . Similarly, for a set of states R , the $LSU^\#$ -energy-reachability (resp. $LSU^{\overline{\#}}$ -energy-reachability, LSU^Σ -energy-reachability) objective is the set of LU -feasible paths π reaching R such that along their associated run from (q_0, L) , the number $\#\mathbf{V}(\rho)$ of violations (resp. maximal number of consecutive violations $\overline{\#\mathbf{V}}(\rho)$, sum $\Sigma\mathbf{V}(\rho)$ of violations) of the soft upper bound S is at most V .




Figure 1: An arena for a $\text{LSU}^\#$ -energy reachability game.




Figure 2: Energy level and $\#V$ along a winning run in a $\text{LSU}^\#$ -energy reachability game.

We study the complexity of deciding the existence of a winning strategy for the objectives defined above, in both the one- and two-player settings. In addition, given L , we address the bound-existence problem, which asks to decide if there exists an upper bound $U \in \mathbb{Z}$ (resp, a weak upper bound $W \in \mathbb{Z}$) for which Player 1 wins the LU-energy (resp. LW-energy, LSU^* -energy) game. Tables 1 and 2 summarize known results, and the results obtained in this paper (where LSU^* -energy gathers all three energy constraints with violations).

	1 player	2 players
L-energy	in PTIME (Thm. 4, [CDH17])	in $\text{NP} \cap \text{coNP}$ (Thm. 4, [CDH17])
LU-energy	PSPACE-c. (Thm. 7)	EXPTIME-c. (Thm. 7)
LW-energy	in PTIME (Thm. 23)	in coNP (Coro. 27)
LSU^* -energy	PSPACE-c. (Thm. 28)	EXPTIME-c. (Thm. 28)

Table 1: Summary of our results : Reachability

	Reachability		Infinite runs	
	1 player	2 players	1 player	2 players
LU-energy (Given $L, \exists U?$)	in PTIME (Prop. 30)	in $\text{NP} \cap \text{coNP}$ (Prop. 30)	in NP (Thm. 31)	in \mathcal{L} -EXPTIME ([JLR13])
LW-energy (Given $L, \exists W?$)	in PTIME (Prop. 30)	in $\text{NP} \cap \text{coNP}$ (Prop. 30)	in PTIME (Thm 32)	in \mathcal{L} -EXPTIME ([JLR13])
LSU^* -energy (Given $L, S, V, \exists U$)	PSPACE-c (remark 29)	EXPTIME-c (remark 29)	PSPACE-c (remark 29)	EXPTIME-c (remark 29)

Table 2: Summary of our results: (weak) upper bound existence

3. Energy reachability games with strict bounds

In this section, we focus on the L-energy-reachability and LU-energy-reachability problems. We first prove that L-energy-reachability problems are inter-reducible with L-energy problems, which entails:

Theorem 4. *Two-player L-energy-reachability games are decidable in $\text{NP} \cap \text{coNP}$. The one-player version is in PTIME.*

Remark 5. *Notice that these results are not a direct consequence of the results of [CD12] about energy parity games: that paper focuses on the existence of an initial energy level for which Player 1 has a winning strategy with energy-parity objectives (which encompass our energy-reachability objectives). When the answer is positive, one can compute the minimal initial energy level for which a winning strategy exists, but the (deterministic) algorithm runs in exponential time.*

Remark 6. *These results were already proven in [CDH17]: for one-player arenas, the authors develop a PTIME algorithm, while they prove LOGSPACE-equivalence with L-energy games for the two-player setting (the result then follow from [BFL⁺08]). Our proof uses similar arguments as in the latter proof, but with a uniform, full and direct reduction back and forth both for the one and two-player cases.*

Proof. We prove that L-energy-reachability and L-energy games are interreducible. The theorem then follows from the results of [BFL⁺08].

First consider a two-player arena $G = (Q_1, Q_2, E)$, an initial state q_0 , and an L-energy objective. We define a new arena $G' = (Q_1 \cup Q_c \cup \{q_t\}, Q_2, E')$ (assuming $q_t \notin Q$) where $Q_c = \{q_c \mid q \in Q\}$ is a copy of all the vertices of G . Note that, q_c is always a Player 1 vertex; intuitively, states in Q_c are used to allow Player 1 to stop the game and reach the target state q_t , if enough energy has been stored. The set of transitions E' is obtained from E as follows (where the (positive) rational¹ value of ε will be fixed later):

- for each $(q, w, q') \in E$, there is a transition $(q, w + \varepsilon, q'_c)$ and $(q'_c, 0, q')$ in E' ;
- for each $q_c \in Q_c$, there is a transition $(q_c, -\delta, q_t)$ in E' , where $\delta = 1 + \sum_{(q, w, q') \in E} |w|$;
- finally, E' contains an edge $(q_t, 0, q_t)$.

We claim that Player 1 has a winning strategy from q_0 for the L-energy-reachability objective in G' if, and only if, she has a winning strategy from q_0 for the L-energy objective in G .

¹Our definition of arenas do not allow for rational weights, but by scaling up all constants (including the energy bounds) we get an equivalent instance of our problem with only integer bounds.




Figure 3: Schema of the reduction from L-energy to L-energy-reachability objectives

First assume that Player 1 has a winning strategy σ in G for the L-energy objective; then we can assume that this strategy is memoryless [BFL⁺08]; we define the strategy σ' as follows: for any state q of G , letting $q' = \sigma(q)$, we define $\sigma'(\pi \cdot q) = q'_c$, and

$$\sigma'(\pi \cdot q \cdot q'_c) = \begin{cases} q' & \text{if } |\pi| \leq \frac{\delta}{2\varepsilon} - 1 \\ q_t & \text{otherwise.} \end{cases}$$

Obviously, any outcome μ' of σ' from q_0 reaches q_t . First note that, by construction of σ' , the prefix ν' of μ' just before reaching q_t has odd length, say $2n - 1$. Also note that it corresponds to an outcome ν of σ in G of length n . Since σ is assumed winning, ν must be L-feasible; moreover, we have $\tilde{\nu}'_{2i} = \tilde{\nu}'_{2i-1} = \tilde{\nu}_i + i \cdot \varepsilon$ for all $0 \leq i < n$. Now, $\tilde{\nu}_i \geq L$ for all i , since ν is an outcome of σ , so that also $\tilde{\nu}'_i \geq L$ for all i . Moreover, $|\nu'| = \delta/2\varepsilon - 1$ implies that, $|\nu| = \delta/\varepsilon$, so that $\tilde{\nu}'_{2n-1} \geq L + \delta$, and $\tilde{\mu}'_{2n} \geq L$. It follows that σ' is winning in G' for the L-energy-reachability objective.

Conversely, assume that Player 1 wins the L-energy-reachability game G' , and write σ' for a winning strategy in G' from q_0 . We may assume that no negative cycle occurs along any outcome of σ' : indeed, consider the (finite) execution tree of σ' , and assume that it involves a negative cycle starting and ending at some state q ; then there must exist a subtree rooted at q which contains no other occurrences of q ; by redefining σ' so as to play as in this subtree after any occurrence of q , we remove all occurrences of our negative cycle, while preserving reachability of q_t and still satisfying the energy constraint (since removing negative cycles may only increase the energy level).

Now, take any outcome ρ' of σ' from q_0 , it must eventually reach q_t . First note that, any prefix of ρ' is of the form $q_0 q_1^c q_1 \dots q_t$. Hence, if we take any prefix π' of ρ' before reaching q_t and drop the alternate vertices, we get a corresponding path in G . Now, as ρ' eventually reaches q_t and since the edge leading to q_t has weight $-\delta$, a positive cycle must have been visited along ρ' in G' . From σ' , we can then build a strategy σ that repeats the first positive cycle visited (after dropping the alternate vertices). Formally, $\sigma(\pi \cdot q) = q'$ if $\sigma'(\pi' \cdot q) = q'_c$ where π is obtained dropping alternate vertices from π' and π' contains no positive cycle. When π is a run of the form $\pi = \rho_1 \cdot \beta_1 \dots \beta_{k-1} \cdot \rho_k$, where each β_i is a positive cycle, we take $\sigma(\rho_1 \cdot \beta_1) = \sigma(\rho_1)$. The resulting strategy σ then never takes the edge to q_t , since it only plays moves returned by σ' along outcomes that do not contain positive cycles. Moreover, all simple cycles generated by σ in G' are positive cycles; by taking $\varepsilon < \frac{1}{|Q|+1}$, these cycles still are positive cycles in G , so that σ is winning in G for the L-energy objective.

We now prove the converse reduction, which relies on similar ideas: we consider a two-player arena $G = (Q_1, Q_2, E)$, an initial state q_0 , and an L-energy-reachability objective; we assume without loss of generality that there is a unique target state q_t , and write $Attr_1(q_t)$ for the Player 1-attractor of q_t in G . We build (in polynomial time) a two-player arena $G' = (Q'_1, Q'_2, E')$ from G as follows:

- $Q'_1 = (Q_1 \cap Attr_1(q_t)) \cup \{q'_0, q_s\}$ and $Q'_2 = Q_2$. State q_0 will serve as the new initial state, and q_s is a sink state that guarantees existence of an infinite run;
- letting $E_0 = \{(q, w - \varepsilon, q') \mid (q, w, q') \in E \text{ and } q \in Q'_1 \cup Q'_2 \setminus \{q_t\}\} \cup \{(q_t, 0, q_t), (q_s, -1, q_s)\} \cup \{(q'_0, q) \mid q \in \{q_0\} \cap Attr_1(q_t)\}$, we define $E' = E_0 \cup \{(q, 0, q_s) \mid qE_0 = \emptyset\}$. This way, all states have an outgoing edge, possibly to the sink state q_s if no other transition exists. As for the first reduction, the exact value of ε will be fixed later.

We prove that Player 1 wins the L-energy-reachability game in G from q_0 if, and only if, she wins the L-energy game in G' from q'_0 .




Figure 4: Schema of the reduction from L-energy-reachability to L-energy objectives

For the first direction, if Player 1 has a winning strategy to reach q_t from q_0 in G while maintaining the energy level above L , then she has such a strategy σ along whose outcomes the energy level is bounded above by $L + 2\delta$ (where $\delta = 1 + \sum_{(q,w,q') \in E} |w|$): indeed, if energy level $L + \delta$ is reached along some outcome, then Player 1 can achieve the reachability objective by playing her memoryless *attractor strategy*. Choosing the attractor strategy ensures reaching q_t , and will decrease the energy level by at most δ along any outcome. Similarly, following the attractor strategy can increase the energy level by no more than δ . Strategy σ can be chosen to contain no negative cycles, so we can bound the length of the outcomes by $(\delta + 1) \cdot |Q|$. Now, by taking $\varepsilon < \frac{1}{(\delta+1) \cdot |Q|}$, we can mimic strategy σ in G' : all outcomes only visits states in the attractor of q_t , and reach q_t in at most $(\delta + 1) \cdot |Q| + 1$ steps (the extra step is the transition from q'_0 to q_0). The ε difference in the weights is compensated by the initial credit 1 harvested when moving from q'_0 to q_0 , so that all outcomes satisfy the L-energy constraint.

Conversely, if Player 1 has a winning strategy σ' from q'_0 in G' , then we can assume that this strategy is memoryless [BFL⁺08]. Some of the outcomes may reach q_t , some may not. Since σ' is memoryless, it cannot take any negative cycle, as this would yield an outcome whose energy level tends to $-\infty$. Hence it

may only take nonnegative cycles in G' , which correspond to positive cycles in G (since $\varepsilon > 0$). As a consequence, when mimicking σ' in G , for those outcomes that do not reach q_t , the energy level will grow arbitrarily high; when it exceeds δ , Player 1 can play her attractor-strategy to reach q_t . This concludes our proof for two-player games. \square

Similarly, for LU-energy-reachability objectives, we prove the same complexities as with classical LU-energy objectives:

Theorem 7. *One-player LU-energy-reachability games are PSPACE-complete. Two-player LU-energy-reachability games are EXPTIME-complete.*

Proof. Membership in PSPACE and EXPTIME is proved by considering the *expanded game* G_{LU} , in which energy is represented explicitly in states: it can be used to check reachability of a target state for both the one- and the two-player cases. Recall that G_{LU} is of exponential size, and that reachability in a graph is in NLOGSPACE. Hence, for the one-player case, a strategy for an LU-energy-reachability game can be found on-the-fly in PSPACE, without explicitly building the expanded arena. For the two-player case, recall that reachability in two-player games is in PTIME w.r.t the size of the arena (by building an attractor for the target states). We solve reachability in the exponential-size arena G_{LU} , which results in an EXPTIME algorithm.

For both the one- and the two-player settings, the hardness proofs for LU-energy objectives are readily adapted to LU-energy-reachability objectives, since they are based on reachability-like problems (reachability in bounded one-counter automata [FJ13] and countdown games [JLS07], respectively).

Let us first show PSPACE-hardness of the LU-energy-reachability games for the one-player setting. We reduce the question of reachability in bounded one-counter automata to existence of a winning strategy in one-player LU-energy-reachability games. A bounded one-counter automaton is a machine $M = (Q, q_0, q_f, \delta_M)$ with one counter c that stores values between 0 and an integer upper bound b . The elements of the machine are its set of states Q , an initial state q_0 , and a target state q_f . Transitions in δ are tuples of the form (q, p, g_1, g_2, q') where q and q' are states in Q , $p \in [-b, b]$ is an increment or decrement of the counter, and $g_1, g_2 \in [0, b]$ are lower- and upper bounds on the value of c . The semantics of bounded one-counter automata is as follows: the execution starts from configuration $(q_0, 0)$. A transition (q, p, g_1, g_2, q') can be fired from some configuration (q, c) if $g_1 \leq c \leq g_2$; it leads to configuration $(q', c + p)$. The reachability question asks whether, starting from configuration $(q_0, 0)$, the one-counter automaton can reach state q_f .

We build a one-player LU-energy-reachability game with arena G_M with set of states $Q' = Q \cup \{q^{t,1}, q^{t,2}, q^{t,3}, q^{t,4} \mid \exists t = (q, p, g_1, g_2, q') \in \delta_M\}$ and bounds $[0, U]$ with $U = b$. The energy level in this game will simulate the value of the counter. For every transition $t = (q, p, g_1, g_2, q') \in \delta_M$, we create five transitions:

- two transitions $(q, -g_1, q^{t,1})$ and $(q^{t,1}, g_1, q^{t,2})$ to test that $c \geq g_1$;
- two transitions $(q^{t,2}, U - g_2, q^{t,3})$ and $(q^{t,3}, g_2 - U, q^{t,4})$ to test that $c \leq g_2$;

- one transition $(q^{t,4}, p, q')$ to move to the next state q' with an update of the energy level.

Notice that the intermediate states have only one incoming- and one outgoing transitions; we write $(q, v) \xrightarrow{[g_1, g_2], p} (q', v + p)$ for the sequence

$$(q, v) \xrightarrow{-g_1} (q^{t,1}, v - g_1) \xrightarrow{g_1} (q^{t,2}, v) \xrightarrow{U - g_2} (q^{t,3}, v + U - g_2) \xrightarrow{g_2 - U} (q^{t,4}, v) \xrightarrow{p} (q', v + p),$$

provided that all intermediary configurations fulfill the LU-energy constraint.

We start the LU-energy-reachability game on G_M with energy level 0. Assume that there exists a winning run ρ from q_0 to q_f of the form

$$\rho = (q_0, 0) \xrightarrow{[g_1, g_2], p} (q', p) \dots (q_i, v_i) \xrightarrow{[g_{i,1}, g_{i,2}], p_i} (q_{i+1}, v_i + p_i) \dots \\ (q_{n-1}, v_{n-1}) \xrightarrow{[g_{n-1,1}, g_{n-1,2}], p_{n-1}} (q_f, v_n).$$

As ρ is a winning run for the LU-energy-reachability game with bounds $[0, U = b]$, we have that for every i , $v_i - g_{i,1} \geq 0$, so $v_i \geq g_{i,1} \geq 0$. Similarly, we have $v_i + b - g_{i,2} \leq b$, so $v_i \leq g_{i,2} \leq b$. The bounded character of LU-energy-reachability game guarantees that q_f is reached with a final energy level $v_n \in [0, U]$. Now, if an energy level is such that $v_i \leq g_{i,1}$ but $v_i \geq g_{i,2}$, then one can play moves $(q_i, v_i) \xrightarrow{-g_{i,1}} (q^{i,1}, v_i - g_{i,1}) \xrightarrow{g_1} (q^{i,2}, v_i)$, but the system will get stuck in configuration $(q^{i,2}, v_i)$. So there exists a winning run of M starting from $(q_0, 0)$ and reaching q_f iff there exists a winning run for the LU game. Notice that G_M has a size linear in the size of M . As one-counter games are PSPACE-complete [FJ13], this shows PSPACE-hardness of one-player LU-energy-reachability games.

We can now show EXPTIME-hardness of two-player LU-energy-reachability games. We can show that existence of a winning strategy in countdown games can be brought back to a LU-energy-reachability game. A countdown game is a pair (S, E) where S is a set of states, and $E \subseteq S \times \mathbb{N} \times S$ is a set of edges. The game starts from a configuration (s_0, c_0) , where s_0 is an initial state and c_0 the value of the unique counter in the game. Then, for every turn of the game, from configuration (s, c) , Player 1 chooses a value d such that $0 < d \leq c$, and such that there exists at least one transition of the form $(s, d, s') \in E$. This choice is immediately followed by a choice of Player 2 to move to a state s'' such that $(s, d, s'') \in E$. The configuration reached after this decision becomes $(s'', c - d)$. A play is winning for Player 1 if it reaches a configuration of the form $(s, 0)$. On the other hand, Player 2 wins a play if it ends in a configuration (s, c) that is deadlocked, i.e. such that $c < d$ for every $(s, d, s') \in E$. The problem whether Player 1 has a winning strategy in a countdown game is EXPTIME-complete, as shown in [JLS07].

We recast the winning strategy question in LU-energy-reachability games the following way. We build an arena $G = (Q_1 \cup Q_2, E_{LU})$. For every state $s \in S$, we create a state $q_s \in Q_1$. Further, we add two particular states q_f and q_t

to Q_1 . Then, for every edge of the form (s, d, s') , we create a state $q_{s,d}$ in Q_2 . For every edge (s, d, s') , we create an edge $(q_s, -d, q_{s,d})$ and an edge $(q_{s,d}, 0, q'_s)$ in E_{LU} . Further, we create an edge $(q_s, 0, q_t)$ from every state $q_s \in Q_1$, and an edge $(q_t, +U, q_t)$. We set as initial value of the arena c_0 , we take as lower bound $L = 0$ and as upper bound $U = c_0$. A winning play in arena G is a play ρ of the form

$$(q_{s_0}, c_0) \xrightarrow{-d_1} (q_{s_0,d_1}, c_0) \xrightarrow{0} (q_{s_1}, c_1) \xrightarrow{-d_2} \dots (q_{s_k}, c_k) \xrightarrow{0} (q_t, c_k) \xrightarrow{U} (q_t, c_k + U).$$

This play is winning if, and only if, all weights c_0, c_1, \dots, c_k and $c_k + U$ belong to $[0, U]$, hence if, and only if, $c_k = 0$. Then, ρ is winning if, and only if, $(s_0, c_0) \xrightarrow{-d_1} (s_1, c_1) \dots (s_k, c_k = 0)$ is winning. So, if Player 1 has a strategy to win the LU-energy-reachability game starting from energy level c_0 , then the countdown game is winning for Player 1 starting from configuration (s_0, c_0) . \square

4. Energy reachability games with weak upper bound

Finding a strategy that satisfies an LW-energy constraint along an infinite run is conceptually easy: it suffices to find a cycle that can be iterated once with a positive effect. It follows that memoryless strategies are enough, and the LW-energy problem was shown to be in PTIME for one-player arenas, and in $NP \cap \text{coNP}$ for two-player arenas [BFL⁺08].

The situation is different when we have a reachability condition: players may have to keep track of the exact energy level in order to find their way to the target state. Obviously, considering the expanded arena G_{LW} , we easily get exponential-time algorithms for LW-energy-reachability objectives. However, as proved below, in the one-player case, a PTIME algorithm exists.

Example 8. Consider the one-player arena of Fig. 5, where the lower bound is $L = 0$, the weak-upper bound is $W = 5$, and the target state is q_t . Starting from q_0 with initial credit 0, we first have to move to q_1 , and then iterate the positive cycle $\beta_1 = (q_1, +2, q_2) \cdot (q_2, -2, q_3) \cdot (q_3, +1, q_1)$ three times, ending up in q_1 with energy level 3. We then take the cycle $\beta_2 = (q_1, +2, q_2) \cdot (q_2, -5, q_4) \cdot (q_4, +5, q_1)$, which raises the energy level to 5 when we come back to q_1 , so that we




Figure 5: A one-player arena with LW-energy-reachability objective

can reach q_t . Notice that β_1 has to be repeated three times before taking cycle β_2 , and that repeating β_1 more than three times maintains the energy level at 4, which is not sufficient to reach q_t . This suggests that Player 1 needs memory and cannot rely on a single cycle to win LW-energy-reachability games.

This example shows that winning strategies for Player 1 may have to monitor the exact energy level all along the computation, thereby requiring exponential memory (assuming that all constants are encoded in binary; with unary encoding, the expanded game G_{LW} would have polynomial size and directly give a polynomial-time algorithm).

Proposition 9. *In LW-energy-reachability games, exponential memory may be necessary for Player 1 (assuming binary-encoded constants).*

Notice that this does not prevent from having PTIME algorithms: the strategy in Example 8 is not very involved, but it depends on the energy level (up to W).

4.1. One-player case

In this section, we focus on the one-player case. In order to get a polynomial-time algorithm, we begin with analyzing cycles in the graph, and prove that a path witnessing LW-energy reachability can have a special form, which can be represented compactly in polynomial size. We then characterize interesting cycles to iterate in order to efficiently search for interesting cycles to iterate.

4.1.1. Reshaping winning paths

In this section, we prove the following result:

Proposition 10. *Let π be a finite path in a one-player arena G . There exists a path π' of the form $\alpha_1 \cdot \varphi_1^{n_1} \cdot \alpha_2 \cdot \varphi_2^{n_2} \cdot \dots \cdot \varphi_k^{n_k} \cdot \alpha_{k+1}$, where φ_j are distinct cycles of length less than $|Q|$, α_j are acyclic paths, and n_j are integers, such that if $(q, u) \xrightarrow{\pi}_{\text{LW}} (q', u')$, then also $(q, u) \xrightarrow{\pi'}_{\text{LW}} (q', u'')$ with $u'' \geq u'$.*

We begin with a series of simple lemmas. Our first lemma states that starting with higher energy level can only be beneficial:

Lemma 11. *Let π be a finite path in a one-player arena G . If $(q, u) \xrightarrow{\pi}_{\text{LW}} (q', u')$, then for any $v \geq u$, $(q, v) \xrightarrow{\pi}_{\text{LW}} (q', v')$ for some $v' \geq u'$.*

Proof. Write $\pi = (e_i)_{0 \leq i < n}$, with $e_i = (q_i, p_i, q'_i)$ for each i . The sequence defined as

$$u_0 = u \qquad u_{i+1} = \min(W, u_i + p_i)$$

is the sequence of energy levels along the run $(q, u) \xrightarrow{\pi}_{\text{LW}} (q', u')$, so that $u_n = u'$. For $v \geq u$, letting

$$v_0 = v \qquad v_{i+1} = \min(W, v_i + p_i),$$

we easily prove by induction that for all i , $u_i \leq v_i \leq W$, which entails that $(q, v) \xrightarrow{\pi}_{\text{LW}} (q', v')$ with $v' = v_n \geq u_n = u'$. \square

Notice that, even if we add condition $u' > u$ in the hypotheses of Lemma 11, it need not be the case that $v' > v$. In other terms, a sequence of transitions may have a positive effect on the energy level from some configuration, and a negative effect from some other configuration, due to the weak upper bound. Below, we prove a series of results related to this issue, and that will be useful for the rest of the proof. First, the effect of a given path (i.e., the net amount of energy that is harvested) decreases when the initial energy level increases:

Lemma 12. *Let π be a finite path in a one-player arena G , and consider two LW-runs $(q, u) \xrightarrow{\pi}_{LW} (q', u')$ and $(q, v) \xrightarrow{\pi}_{LW} (q', v')$ with $u \leq v$. Then $u' - u \geq v' - v$, and if the inequality is strict, then the energy level along the run $(q, v) \xrightarrow{\pi}_{LW} (q', v')$ must have hit W .*

Proof. The first statement is proven by induction: we again write $\pi = (e_i)_{0 \leq i < n}$, with $e_i = (q_i, p_i, q'_i)$ for each i , and

$$\begin{aligned} u_0 &= u & u_{i+1} &= \min(W, u_i + p_i) \\ v_0 &= v & v_{i+1} &= \min(W, v_i + p_i). \end{aligned}$$

Then $u_{i+1} - u_i = \min(W - u_i, p_i)$ and $v_{i+1} - v_i = \min(W - v_i, p_i)$. Since $u_i \leq v_i$ for all i , we also have $W - u_i \geq W - v_i$, and $u_{i+1} - u_i \geq v_{i+1} - v_i$. By summing up these inequalities, we get $u_{i+1} - u_0 \geq v_{i+1} - v_0$. Now, as long as $W - v_i \geq p_i$ (then also $W - u_i \geq p_i$), the inequalities above are equalities. It follows that if the inequality is strict, then the run $(q, v) \xrightarrow{\pi}_{LW} (q', v')$ must have hit W . \square

The next lemma is more precise about the effect of following a path when starting from the maximal energy level W :

Lemma 13. *Let π be a finite path in a one-player arena G , for which there is an LW-run $(q, u) \xrightarrow{\pi}_{LW} (q', u')$. If u' is the maximal energy level along that run, then $(q, W) \xrightarrow{\pi}_{LW} (q', W)$; if u is the maximal energy level along the run above, then $(q, W) \xrightarrow{\pi}_{LW} (q', W + u' - u)$.*

Proof. Write $\pi = (e_i)_{0 \leq i < n}$, with $e_i = (q_i, p_i, q'_i)$ for each i . If u' is the maximal energy level, then for all i , we have $\sum_{j=i}^{n-1} p_j \geq 0$. Now, define

$$v_0 = W \quad v_{i+1} = \min(W, v_i + p_i).$$

If $v_n < W$, then by induction we also have $v_i < W$ for all i , contradicting the fact that $v_0 = W$. This proves our first result.

Similarly, if u is the maximal energy level, then for all i , we have $\sum_{j=0}^i p_j \leq 0$. Then for all i , $v_{i+1} = v_i + p_i \leq W$, so that $v_n - v_0 = u' - u$. Our second result follows. \square

From Lemma 11, it follows that any run witnessing LW-energy reachability can be assumed to contain no cycles with nonpositive effect. Formally:

Lemma 14. *Let π be a finite path in a one-player arena G . If $(q, u) \xrightarrow{\pi}_{LW} (q', u')$ and π can be decomposed as $\pi_1 \cdot \pi_2 \cdot \pi_3$ in such a way that $(q, u) \xrightarrow{\pi_1}_{LW} (s, v) \xrightarrow{\pi_2}_{LW} (s, v')$ and $(s, v') \xrightarrow{\pi_3}_{LW} (q', u')$ with $v' \leq v$, then $(q, u) \xrightarrow{\pi_1 \cdot \pi_3}_{LW} (q', u'')$ with $u'' \geq u'$.*

Proof. Since $(s, v') \xrightarrow{\pi_3}_{\text{LW}} (q', u')$ and $v' \leq v$, by Lemma 11 we also have $(s, v) \xrightarrow{\pi_3}_{\text{LW}} (q', u'')$ for some $u'' \geq u'$. The result follows. \square

The following lemmas show that several occurrences of a cycle having positive effect along a path can be gathered together. This will be useful to prove the existence of a short path witnessing LW-energy reachability.

Lemma 15. *Let π be a finite path in a one-player arena G . If $(q, u) \xrightarrow{\pi}_{\text{LW}} (q', u')$ with $u' > u$ and $(q, w) \xrightarrow{\pi}_{\text{LW}} (q', w')$ with $w' > w$, then for any $u \leq v \leq w$, it holds that $(q, v) \xrightarrow{\pi}_{\text{LW}} (q', v')$ with $v' > v$.*

Proof. Using Lemma 11, we immediately have $(q, v) \xrightarrow{\pi}_{\text{LW}} (q', v')$. As in the previous proof, we define sequences

$$\begin{aligned} u_0 &= u & u_{i+1} &= \min(W, u_i + p_i) \\ v_0 &= v & v_{i+1} &= \min(W, v_i + p_i) \\ w_0 &= w & w_{i+1} &= \min(W, w_i + p_i). \end{aligned}$$

We still have $u_i \leq v_i \leq w_i$ for all i . Moreover, if $v_j < W$ for all $j \leq i$, then $v_i - u_i = v - u$. As a consequence, if $v' \leq v$, then it must be the case that $v_j = W$ for some j ; but then $w_j = v_j$, since $v_j \leq w_j \leq W$. It follows that $w_k = v_k$ for all $k \geq j$, so at the end of π we have $w' = v'$. Assuming $v' \leq v$ raises a contradiction since we have $v' = w' > w \geq v$. Hence $v' > v$. \square

Lemma 16. *Let π be a cycle on q such that $(q, u) \xrightarrow{\pi}_{\text{LW}} (q, v)$ for some $u \leq v$. Then $(q, u) \xrightarrow{\pi^{W-L}}_{\text{LW}} (q, v')$ for some $v' \geq v$, and $(q, v') \xrightarrow{\pi}_{\text{LW}} (q, v')$.*

Proof. The case where $u = v$ is trivial. We assume $u < v$. Applying Lemma 11 inductively, we get that the cycle can be iterated arbitrarily many times; this also proves that the sequence of energy levels reached at the end of each iteration is non-decreasing.

Now, assume that $(q, v') \xrightarrow{\pi}_{\text{LW}} (q, v'')$ for some $v'' \neq v'$. Then $v'' > v'$. Lemma 15 then entails that the sequence of energy levels reached at the end of each iteration is increasing. Since the loop has been iterated $W - L$ times, the energy level in v'' would exceed W , which is impossible. This proves our result. \square

We now prove Prop. 10. Fix a path π in G , and assume that some cycle φ appears (at least) twice along π (if not, then π is already in the desired form): the first time from some configuration (q, v) to some configuration (q, v') , and the second time from (q, w) to (q, w') . First, we may assume that φ has length at most $|Q|$, since otherwise we can take an inner subcycle. We may also assume that $w > v'$, as otherwise we can apply Lemma 14 to get rid of the resulting nonpositive cycle between (q, v') and (q, w) . For the same reason we may assume $v' > v$ and $w' > w$. As a consequence, by Lemmas 15 and 11, by repeatedly iterating φ from (q, v) , we eventually reach some configuration (q, x) with $x \geq w'$, from which we can follow the suffix of π after the second occurrence

of φ . It follows that all occurrences of φ along π can be grouped together, and we can restrict our attention to runs of the form $\alpha_1 \cdot \varphi_1^{n_1} \cdot \alpha_2 \cdot \varphi_2^{n_2} \cdots \varphi_k^{n_k} \cdot \alpha_{k+1}$ where the cycles φ_j are distinct, and have size at most $|Q|$, and the finite runs α_j are acyclic. Notice that each occurrence of any cycle φ_j can be assumed to have positive effect, and by Lemma 16, we may assume $n_j = W - L$ for all j .

4.1.2. Characterizing interesting cycles

While Prop. 10 allows us to only consider paths of a special form, it does not provide *short* witnesses, since there may be exponentially many cycles of length less than or equal to $|Q|$, and the witnessing run may need to iterate several cycles looping on the same state (see Example 8). In order to circumvent this problem, we have to show that all cycles need not be considered, and that one can compute the "useful" cycles efficiently. For this, we introduce *universal* cycles, which are cycles that can be iterated from any initial energy level (above L).

Definition 17. *Let G be a one-player arena, and q be a state of G . Let W be a weak-upper bound and $L \leq W$ be a lower bound. A universal cycle on q in G is a cycle φ with $\text{first}(\varphi) = \text{last}(\varphi) = q$ such that $(q, L) \xrightarrow{\varphi}_{LW} (q, v_{\varphi,L})$ for some $v_{\varphi,L}$ (i.e., the energy level never drops below the lower bound L when following φ with initial energy level L). A universal cycle is positive if $v_{\varphi,L} > L$.*

When a cycle φ is iterated $W - L$ times in a row, then some universal cycle σ is also iterated $W - L - 1$ times (by considering the state with minimal energy level along φ). As a consequence, iterating only universal cycles is enough: we may now only look for runs of the form $\beta_1 \cdot \sigma_1^{n_1} \cdot \beta_2 \cdot \sigma_2^{n_2} \cdots \sigma_k^{n_k} \cdot \beta_{k+1}$ where σ_j 's are *universal* cycles of length at most $|Q|$. Now, assume that some state q admits two universal cycles σ and σ' , and that both cycles appear along a given run π . Write h (resp. h') for the energy levels reached after iterating σ (resp. σ') $W - L$ times. We define a preorder on universal cycles of q by letting $\sigma \triangleright \sigma'$ when $h \geq h'$. Then if $\sigma \triangleright \sigma'$, each occurrence of σ' along π can be replaced with σ , yielding a run π' that still satisfies the LW-energy condition (and has the same first and last states). Generalizing this argument, each state that admits universal cycles has an optimal universal cycle of length at most $|Q|$, and it is enough to iterate only this universal cycle to find a path witnessing reachability. This provides us with a *small witness*, of the form $\gamma_1 \cdot \tau_1^{W-L} \cdot \gamma_2 \cdot \tau_2^{W-L} \cdots \tau_k^{W-L} \cdot \gamma_{k+1}$ where τ_j are optimal universal cycles of length at most $|Q|$ and γ_j are acyclic paths. Since it suffices to consider at most one universal cycle per state, we have $k \leq |Q|$. From this, we immediately derive an NP algorithm for solving LW-energy reachability for one-player arenas: it suffices to non-deterministically select each portion of the path, and compute that each portion is LW-feasible (notice that there is no need for checking universality nor optimality of cycles; those properties were only used to prove that small witnesses exist). Checking LW-feasibility requires computing the final energy level reached after iterating a cycle $W - L$ times; this can be performed by detecting the highest energy level along that cycle, and computing how much the energy level decreases from that point on until the end of the cycle. This provides us with a way of *accelerating* the computation of the effect of iterating cycles.

We now prove that optimal universal cycles of length at most $|Q|$ can be computed for a given state q_0 . For this we unwind the graph from q_0 as a DAG of depth $|Q|$, so that it includes all cycles of length at most $|Q|$. We name the states of this DAG $[q', d]$, where q' is the name of a state of the arena and d is the depth of this state in the DAG (using square brackets to avoid confusion with configurations (q, l) where l is the energy level); hence there are transitions $([q', d], w, [q'', d + 1])$ in the DAG as soon as there is a transition (q', w, q'') in the arena.

We then explore this DAG from its initial state $[q_0, 0]$, looking for (paths corresponding to) universal cycles. Our aim is to keep track of all runs from $[q_0, 0]$ to $[q', d]$ that are prefixes of universal cycles starting from q_0 . Actually, we do not need to keep track of those runs explicitly, and it suffices for each such run to remember the following two values:

- the maximal energy level M that has been observed along the run so far (starting from energy level L , with weak upper bound W);
- the difference m between the maximal energy level M and the final energy level in $[q', d]$. Notice that $m \geq 0$, and that the final energy level in $[q', d]$ is $M - m$.

Example 18. Figure 6 shows two universal cycles from q_0 in an LW-energy game with $L = 0$ and $W = 5$. The first cycle, going via q_2 , ends with $M_1 = 5$ (reached in q_2) and $m_1 = 4$, thus with a final energy level of 1 (when starting from energy level 0); actually, iterating this cycle will not improve this final energy level. The second cycle, via q_4 and q_5 , has a maximal energy level $M_2 = 4$ (reached in q_3) and ends with $m_2 = 1$. Hence, after one iteration of this cycle, one can end in state q_0 with energy level $W - m_2 = 4$.




Figure 6: Two cycles with upper bound $W = 5$

If we know the values (M, m) of some path from $[q_0, 0]$ to $[q', d]$, we can decide if a given transition with weight w from $[q', d]$ to $[q'', d + 1]$ can be taken (the resulting path can still be a prefix of a universal cycle if $M - m + w \geq L$), and how the values of M and m have to be updated: if $w > m$, the run will reach a new maximal energy level, and the new pair of values is $(\min(W; M - m + w), 0)$;

if $m + L - M \leq w \leq m$, then the transition can be taken: the new energy level $M - m + w$ will remain between L and M , and we update the pair of values to $(M, m - w)$; finally, if $w < m + L - M$, the energy level would go below L , and the resulting run would not be a prefix of a universal cycle.

Following these ideas, we inductively attach labels to the states of the DAG: initially, $[q_0, 0]$ is labelled with $(M = L, m = 0)$; then if a state $[q', d]$ is labelled with (M, m) , and if there is a transition from $[q', d]$ to $[q'', d + 1]$ with weight w :

- if $w > m$, then we label $[q'', d + 1]$ with the pair $(\min(W; M - m + w), 0)$;
- if $m + L - M \leq w \leq m$, we label $[q'', d + 1]$ with $(M, m - w)$.

The following lemma makes a link between runs in the DAG and labels computed by our algorithm:

Lemma 19. *Let $[q, d]$ be a state of the DAG, and M and m be two integers such that $0 \leq m \leq M - L$. Upon termination of this algorithm, state $[q, d]$ of the DAG is labelled with (M, m) if, and only if, there is an LW-run of length d from (q_0, L) to $(q, M - m)$ along which the energy level always remains in the interval $[L, M]$ and equals M at some point.*

Proof. The proof is by induction on d . The result is trivial for $d = 0$. Now, assume it holds for some depth $d - 1$, and pick a state $[q, d]$. For the first direction, if $[q, d]$ is labelled with (M, m) , then this label was added using some transition $([q', d - 1], w, [q, d])$ and some label (M', m') of $[q', d - 1]$. By induction, there is an LW-run ρ of length $d - 1$ from (q_0, L) to $(q', M' - m')$ in G along which the energy level remains in the interval $[L, M']$. We consider two cases, corresponding to the two ways of updating the pair of values:

- if $w > m'$, then we have $M = \min(W, M' - m' + w)$ and $m = 0$. Now, the transition $([q', d - 1], w, [q, d])$ in the DAG originates from a transition (q', w, q) in G ; taking this transition after ρ provides us with the run of length d from (q_0, L) to $(q, M - m)$ along which the energy level remains in $[L, M]$, as required;
- if $m' + L - M' \leq w \leq m'$, then $M = M'$ and $m = m' - w$. Again, taking transition (q', w, q) after ρ provides us with the LW-run we are looking for.

Conversely, if there is an LW-run ρ of length d from (q_0, L) to $(q, M - m)$ along which the energy level always remains in the interval $[L, M]$, then we write $\rho = \rho' \cdot ((q', l'), w, (q, M - m))$, distinguishing its last transition. By induction, $[q', d - 1]$ must have been labelled with a pair (M', m') such that $l' = M' - m'$ and the energy level along ρ' remained within $[L, M']$. Now, from the existence of a transition $((q', l'), w, (q, M - m))$, we know that there is a transition $([q', d - 1], w, [q, d])$ in the DAG, which will generate the required label of $[q, d]$. \square

Lemma 20. *Let $[q_0, d]$ be a state of the DAG, with $d > 0$. Let m be a nonnegative integer such that $L + m < W$. Upon termination of this algorithm, state $[q_0, d]$ is labelled with (M, m) for some $M > L + m$ if, and only if, there is a universal cycle φ on q_0 of length d such that $(q_0, L) \xrightarrow{\varphi^{W-L}}_{LW} (q_0, W - m)$.*

Proof. First assume that $[q_0, d]$ is labelled with (M, m) for some M such that $M - m > L$. From Lemma 19, there is a cycle φ on q_0 of length d generating a run $(q_0, L) \xrightarrow{\varphi}_{\text{LW}} (q_0, M - m)$ along which the energy level is within $[L, M]$. Then $M - m \geq L$, so that Lemma 16 applies: we then get $(q_0, L) \xrightarrow{\varphi^{W-L}}_{\text{LW}} (q_0, E)$ with $(q_0, E) \xrightarrow{\varphi}_{\text{LW}} (q_0, E)$ and $E \geq L$. Write $(p_i)_{0 \leq i < |\varphi|}$ for the sequence of weights along φ . Also write ρ for the run $(q_0, L) \xrightarrow{\varphi}_{\text{LW}} (q_0, M - m)$, and σ for the run $(q_0, E) \xrightarrow{\varphi}_{\text{LW}} (q_0, E)$.

As $L < M - m$, then by Lemma 12, it must be the case that energy level W is reached along σ . Write i_0 for the first position along ρ for which the energy level is M . Assume $\tilde{\sigma}_{i_0} \neq W$: by Lemma 11, we must have $M = \tilde{\rho}_{i_0} \leq \tilde{\sigma}_{i_0} < W$. Then for all $k \geq i_0$, $\sum_{l=i_0}^k p_l \leq 0$. Since $\tilde{\sigma}_{i_0} < W$, then also $\tilde{\sigma}_k < W$ for all $k \geq i_0$. According to Lemma 12, energy level W is reached in σ , so there exists some $k_0 < i_0$ such that $\tilde{\sigma}_{k_0} = W$. However, as i_0 is the index of the first maximal value in ρ , we have $\tilde{\rho}_{k_0} < M$, and the energy level increases in run ρ between k_0 and i_0 . So according to Lemma 13, we should have $\tilde{\sigma}_{i_0} = W$, which raises a contradiction. Hence we proved $\tilde{\sigma}_{i_0} = W$; applying the second result of Lemma 13, we get $E = W - m$.

Conversely, if there is a universal cycle φ satisfying the conditions of the lemma, then it must have positive effect when run from energy level L . Let F be such that $(q_0, L) \xrightarrow{\varphi}_{\text{LW}} (q_0, F)$, and M be the maximal energy level encountered along the run $(q_0, L) \xrightarrow{\varphi}_{\text{LW}} (q_0, F)$. By Lemma 19, state $[q_0, d]$ is labelled with (M, m') for some $m' \geq 0$ such that $F = M - m'$. By Lemma 16, we must have $(q_0, L) \xrightarrow{\varphi^{W-L}}_{\text{LW}} (q_0, W - m')$. \square

The algorithm above computes optimal universal cycles, but it still runs in exponential time (in the worst case) since it may generate exponentially many different labels in each state $[q, d]$ (one per path from $[q_0, 0]$ to $[q, d]$). We now explain how to only generate polynomially-many pairs (M, m) . This is based on the following partial order on labels: we let $(M, m) \preceq (M', m')$ whenever $M - m \leq M' - m'$ and $m' \leq m$. Notice in particular that

- if $M = M'$, then $(M, m) \preceq (M', m')$ if, and only if, $m' \leq m$;
- if $m = m'$, then $(M, m) \preceq (M', m')$ if, and only if, $M \leq M'$.

The following lemma entails that it suffices to store maximal labels w.r.t. \preceq :

Lemma 21. *Consider two paths π and π' such that $\text{first}(\pi) = \text{first}(\pi')$ and $\text{last}(\pi) = \text{last}(\pi')$, and with respective values (M, m) and (M', m') such that $(M, m) \preceq (M', m')$. If π is a prefix of a universal cycle φ , then π' is a prefix of a universal cycle φ' with $\varphi' \triangleright \varphi$.*

Proof. Let $q = \text{first}(\pi)$ and $q' = \text{last}(\pi)$. We write ψ for the path such that $\varphi = \pi \cdot \psi$ (i.e., ψ is a path from q' to q). Then $(q, L) \xrightarrow{\pi}_{\text{LW}} (q', M - m) \xrightarrow{\psi}_{\text{LW}} (q, F)$ for some $F \geq L$. Also, $(q, L) \xrightarrow{\pi'}_{\text{LW}} (q', M' - m')$. Since $M - m \leq M' - m'$,

we have $(q', M' - m') \xrightarrow{\psi}_{\text{LW}} (q, F')$, and $F \leq F'$. We can thus let $\varphi' = \pi' \cdot \psi$: by Lemma 11, the final energy level reached after iterating φ' is higher than or equal to the energy level reached after iterating φ , since $m' \leq m$. Hence $\varphi' \triangleright \varphi$. \square

It remains to prove that by keeping only maximal labels, we only store a polynomial number of labels:

Lemma 22. *If the algorithm labelling the DAG only keeps maximal labels (for \preceq), then it runs in polynomial time.*

Proof. We prove that, when attaching to each node $[q, d]$ of the DAG only the maximal labels (w.r.t \preceq) reached for a path of length d ending in state q , the number of values for the first component of the different labels that appear at depth $d > 0$ in the DAG is at most $d \cdot |Q|$. Since it only stores optimal labels, our algorithm will never associate to a state $[q, d]$ two labels having the same value on their first component. So, any state at depth d will have at most $d \cdot |Q|$ labels.

So we prove, by induction on d , that the number of different values for the first component among the labels appearing at depth $d > 0$ is at most $d \cdot |Q|$. This is true for $d = 1$ since the initial state $(q, 0)$ only contains $(M = 0, m = 0)$, and each transition with nonnegative weight w will create one new label $(w, 0)$ (transitions with negative weight are not prefixes of universal cycles). Now, since all those labels have value 0 as their second component, each state $[q, 1]$ in the DAG will be attached at most one label. Hence, the total number of labels (and the total number of different values for their first component) is at most $|Q|$ at depth 1 in the DAG.

Now, assume that labels appearing at depth $d > 1$ are all drawn from a set of labels $\{(M_i, m_i) \mid 1 \leq i \leq n\}$ in which the number of different values of M_i is at most $d \cdot |Q|$. Consider a state $[q', d]$, labelled with $\{(M_i, m_i) \mid 1 \leq i \leq n_{q', d}\}$ (even if it means reindexing the labels). Pick a transition from $[q', d]$ to $[q'', d+1]$, with weight w . For each pair (M_i, m_i) associated with $[q', d]$, it creates a new label in $[q'', d+1]$, which is

- either $(\min(W; M_i - m_i + w), 0)$ if $m_i < w$ (maximal energy level increases);
- or $(M_i, m_i - w)$ if $m_i + L - M_i \leq w \leq m_i$ (maximal energy level in unchanged).

Now, for a state $(q'', d+1)$, the set of labels created by all incoming transitions can be grouped as follows:

- labels having zero as their second component; among those, our algorithm only stores the one with maximal first component, as $(M_i, 0) \preceq (M_j, 0)$ as soon as $M_i \leq M_j$;
- for each M_i appearing at depth d , labels having M_i as their first component; again, we only keep the one with minimal second component, as $(M, m_i) \preceq (M, m_j)$ when $m_j \leq m_i$.

Last, for this state $[q'', d + 1]$, we keep at most one label for each distinct value among the first components M_i of labels appearing at depth d , and possibly one extra label with second value 0. In other terms, at depth $d + 1$ the values that appear as first component of labels are obtained from values at depth d , plus possibly one per state; Hence, at depth $d + 1$, there exists at most $(d + 1) \cdot |Q|$ labels, which completes the proof of the induction step. \square

Using the algorithm above, we can compute, for each state q of the original arena, the smallest value m_q for which there exists a universal cycle on q that, when iterated sufficiently many times, leads to configuration $(q, W - m_q)$. Since universal cycles can be iterated from any energy level, if q is reachable, then it is reachable with energy level $W - m_q$. We make this explicit by adding to our arena a special self-loop on q , labelled with $\text{set}(W - m_q)$, which sets the energy level to $W - m_q$ (in the same way as *recharge transitions* of [EF13]).

In the resulting arena, we can restrict to paths of the form $\gamma_1 \cdot \nu_1 \cdot \gamma_2 \cdot \nu_2 \cdots \nu_k \cdot \gamma_{k+1}$, where ν_i are newly added transitions labelled with $\text{set}(W - m)$, and γ_i are acyclic paths. Such paths have length at most $(|Q| + 1)^2$. We can then inductively compute the maximal energy level that can be reached (under our LW-energy constraint) in any state after paths of length less than or equal to $(|Q| + 1)^2$. This can be performed by unwinding (as a DAG) the modified arena from the source state q_0 up to depth $(|Q| + 1)^2$, and labelling the states of this DAG (and in particular target states if they are reachable) by the maximal energy level with which that state can be reached from (q_0, L) ; this is achieved in a way similar to our algorithm for computing the effect of universal cycles, but this time only keeping the maximal energy level that can be reached (under LW-energy constraint). As there are at most $|Q|$ states per level in this DAG of depth at most $(|Q| + 1)^2$, we get:

Theorem 23. *The existence of a winning path in one-player LW-energy-reachability games can be decided in PTIME.*

Example 24. *Consider the one-player arena of Fig. 7. We assume $L = 0$, and fix an even weak upper bound W . The state s has $W/2$ disjoint cycles: for each odd integer i in $[0; W - 1]$, the cycle c_i is made of three consecutive edges with weights $-i, +W$ and $-W + i + 1$. Similarly, the state s' has $W/2$ disjoint cycles: for even integers i in $[0; W - 1]$, the cycle c'_i has weights $-i, +W$ and $-W + i + 1$. Finally, there are: two sequences of k edges of weight 0 from s to s' and from s' to s ; an edge from the initial state to s of weight 1, and from s' to target state q_t of weight $-W$. The total number of states then is $2W + 2k + 2$.*

In order to go from the initial state q_0 , with energy level 0, to the final state q_t , we have to first take the cycle c_1 (with weights $-1, +W, -W + 2$) on s (no other cycles c_i can be taken). We then reach configuration $(s, 2)$. Iterating c_1 has no effect, and the only next interesting cycle is c_2 , for which we have to go to s' . After running c_2 , we end up in $(s', 3)$. Again, iterating c_2 has no effect, and we go back to s , take c_3 , and so on. We have to take each cycle c_i (at least) once, and take the sequences of k edges between s and s' $W/2$ times each. In the end, we have a run of length $3W + Wk + 2$.




Figure 7: An example showing that more than one cycle per state can be needed.

Let us look at the universal cycles that we have in this arena: besides the cycles made of the $2k$ edges with weight zero between s and s' , the only possible universal cycles have to depart from the first state of each cycle c_i (as they are the only states having a positive outgoing edge). Following Lemma 16, such cycles can be iterated arbitrarily many times, and set the energy level to some value in $[L; W]$. Since the only edge available at the end of a universal cycle has weight $+W$, the exact value of the universal cycles is irrelevant: the energy level will be W anyway when reaching the second state of each cycle c_i . As a consequence, using **set**-edges in this example does not shorten the witnessing run, which then cannot be shorter than $3W + Wk + 2$.

4.2. Two-player case

We now move to the two-player setting. We begin with proving a result similar to Lemma 11:

Lemma 25. *Let G be a two-player arena, equipped with an LW -energy-reachability objective. Let q be a state of G , and $u \leq u'$ in $[L; W]$. If Player 1 wins the game from (q, u) , then she also wins from (q, u') .*

Proof. Let σ be a winning strategy for Player 1 from (q, u) . If she plays the same strategy from (q, u') , then for any strategy of Player 2, the resulting outcome from (q, u') follows the same transitions as the outcome of the same strategies from u , with higher energy level. Since σ is winning from (q, u) , it is also winning from (q, u') . \square

By Martin's theorem [Mar75], our games are determined. It follows that if Player 2 wins from some configuration (q, v) , she also wins from (q, v') for all $L \leq v' \leq v$ (assuming the contrary, i.e. (q, v') winning for Player 1, would lead to the contradictory statement that (q, v) is both winning for Player 1 and Player 2). Using classical techniques [CD12], we prove that Player 2 can be restricted to play memoryless strategies:

Proposition 26. *For two-player LW -energy-reachability games, memoryless strategies are sufficient for Player 2.*

Proof. According to Lemma 25, for each state q , there is an integer $v_q \in [L; W+1]$ such that Player 1 wins the game from any configuration (q, v) satisfying $v_q \leq v \leq W$, while Player 2 wins the game from any configuration (q, v) with $L \leq v < v_q$.

Assume that Player 2 wins the game from some state (q, v) , with $L \leq v < v_q$. Denote with $(q, p_i, q_i)_{1 \leq i \leq m}$ the set of outgoing transitions from q . By definition of v_{q_i} , Player 1 wins the game from any configuration of the form (q_i, v) with $v \geq v_{q_i}$. Since Player 2 wins from (q, v) , there must exist an index $1 \leq i \leq m$ such that $v + p_i < v_{q_i}$. This defines a winning move for Player 2 from (q, v) . The same argument applies in all states, and yields a memoryless winning strategy for Player 2. \square

A direct consequence of this result and of Theorem 23 is the following:

Corollary 27. *Two-player LW-energy-reachability games are in coNP.*

Whether those games can be solved in NP or are coNP-hard remains an open question.

5. Energy reachability games with soft upper bound

We now consider games with limited violations, i.e. (reachability) games with $LSU^\#$ -energy, $LSU^{\overline{\#}}$ -energy and LSU^Σ -energy objectives. We address the problems of deciding the winner in the one-player and two-player settings, and consider the existence and minimization questions.

Theorem 28. *$LSU^\#$ -energy, $LSU^{\overline{\#}}$ -energy and LSU^Σ -energy (reachability) games are PSPACE-complete for one-player arenas, and EXPTIME-complete for two-player arenas.*

Proof. Membership in PSPACE and EXPTIME can be obtained by building a variant G_{LSU} of the G_{LU} arena: besides storing the energy level in each state, we can also store the amount of violations (for any of the three measures we consider). More precisely, given an arena G , lower and upper bounds L and U on the energy level, a soft bound S , and a bound V on the measure of violations, for any of our three measures of violations, the maximal energy level that can be reached along a path with violations smaller than or equal to V is $S + V \cdot w_{\max}$, where w_{\max} is the maximal weight in our arena; for this reason, we may assume that $U \leq S + V \cdot w_{\max}$. We then define a new arena² G_{LSU} with set of states $(Q \times ([L; U] \cup \{\perp\}) \times ([0; V] \cup \{\perp\}))^3$, and each transition (q, w, q') of the original arena generates a transition from state $(q, l, (n, c, s))$ to state $(q', l', (n', c', s'))$ whenever

- l' correctly encodes the evolution of the energy level:
 - $l' = l + w$ if l and $l + w$ are in $[L; U]$;

²In order to factor our proof, we store all three measures of violations in one single arena.

- $l' = \perp$ if either $l = \perp$ or $l + w < L$ or $l + w > U$;
- n' correctly stores the number of violations:
 - $n' = \perp$ if $l' = \perp$ or $n = \perp$;
 - $n' = n$ if $l' \in [L; S]$;
 - $n' = n + 1$ if $l' \in (S; U]$ and $n + 1 \leq V$;
 - $n' = \perp$ if $l' \in (S; U]$ and $n + 1 > V$.
- c' is updated to count the current number of consecutive violations:
 - $c' = \perp$ if $l' = \perp$ or $c = \perp$;
 - $c' = 0$ if $l' \in [L; S]$;
 - $c' = c + 1$ if $l' \in (S; U]$ and $c + 1 \leq V$;
 - $c' = \perp$ if $l' \in (S; U]$ and $c + 1 > V$.
- s' encodes the sum of all violations:
 - $s' = \perp$ if $l' = \perp$ or $s = \perp$;
 - $s' = s$ if $l' \in [L; S]$;
 - $s' = s + (l' - S)$ if $l' \in (S; U]$ and $s + (l' - S) \leq V$;
 - $s' = \perp$ if $l' \in (S; U]$ and $s + (l' - S) > V$.

In this arena, n , c and s keep track of the number of violations, number of consecutive violations, and sum of violations; their values are set of \perp as soon as they exceed the violation bound V , or if the energy level has left its range of allowed values $[L; U]$. The arena G_{LSU} is of exponential size, and our LSU^* -energy-reachability problems can be reduced to solving reachability of the relevant set of states in that arena (e.g., Player 1 wins the $\text{LSU}^\#$ -energy reachability game if, and only if, she wins in the modified game G_{LSU} for the objective of reaching the target set without visiting states where $n = \perp$).

Our hardness results are proved by setting the number/amount of allowed violations to zero, thereby recovering the classical LU-energy-reachability games, which were proved PSPACE-complete and EXPTIME-complete for one-player and two-player arenas, respectively.

Solving $\text{LSU}^\#$ -energy, $\text{LSU}^{\bar{\#}}$ -energy, LSU^Σ -energy infinite-duration games can be performed with arena G_{LSU} built above. Now, the objective in $\text{LSU}^\#$ -energy, $\text{LSU}^{\bar{\#}}$ -energy, LSU^Σ -energy games is to enforce infinite runs, that avoid states with $l = \perp$ and with $n = \perp$, $c = \perp$ or $s = \perp$, depending on the chosen criterion on violation. Again, these strategies can be found in PSPACE for the one-player case, and in EXPTIME in the two-player case. For the hardness part, reduction from LU-energy games obtained by setting $V = 0$ still applies. \square

Remark 29. *In our proof, the strong upper bound U is given as input. We could also assume that no such upper bound is given: if the amount of violations is bounded by V along a run, then the maximal energy level that can be reached in $LSU^\#$ -energy, $LSU^\#$ -energy games is bounded by $S + V \cdot w_{\max}$, where w_{\max} is the maximal weight of transitions in G , and by $S + V$ in LSU^Σ -energy games. As games that are winning for a strong upper bound U are also winning with strong upper bound $U+1$, we can prove existence of a bound by solving a game with this maximal energy levels as strong upper bound. We can even compute the smallest U for which Player 1 wins a given LSU^* -energy game using binary search. LSU^* -energy games and their bound existence problems hence have the same PSPACE and EXPTIME complexities.*

6. Existence of bounds for Energy Games

In this section we address the problems of existence and computation of upper bounds U under which Player 1 has a winning strategy in energy games.

6.1. Existence of bounds for reachability games

The case of energy-reachability is trivial: as soon as there exists a winning strategy for L-energy-reachability, the maximal energy level along any outcome can serve as a bound for LW-energy- and LU-energy-reachability games. As a consequence:

Proposition 30. *The upper-bound- and weak-upper-bound existence problems for energy-reachability games are decidable in PTIME for one-player arenas and in $NP \cap \text{coNP}$ for two-player arenas.*

Noticing that from a configuration $(q, L + |Q| \cdot w_{\max})$ where w_{\max} is the maximal absolute value of a weight in the arena, one can reach any state q_t connected to q , we can give a bound on the maximal value of strong and weak upper bounds in energy-reachability games. Even if the path from q to q_t is composed only of negative transitions, an energy level of $L + |Q| \cdot w_{\max}$ allows to reach q_t without reaching an energy level below L . Similarly, if the path from q to q_t is composed only of positive transitions, the value of energy can reach a level up to $L + 2 \cdot |Q| \cdot w_{\max}$. Using binary search in interval $[L, L + 2 \cdot |Q| \cdot w_{\max}]$, we can hence find minimal strong upper bounds for energy-reachability games with one-player arenas in PSPACE, and weak upper bounds in PTIME.

A similar reasoning can be applied in the two-player setting, based on the doubly-exponential bound given by [BHM⁺17, Lemma 2]. Again using binary search, the optimal bounds can then be computed in EXPTIME.

6.2. Existence of bounds for infinite-duration games

In this section, we focus on the existence of upper bounds for infinite-duration games.

6.2.1. LU-energy games

We begin with LU-energy constraints, for which we prove the following result:

Theorem 31. *Given an arena G and a lower bound, deciding whether there exists an upper bound U such that the LU-energy game is winning for Player 1 is decidable in NP for one-player arenas, and in 2-EXPTIME for two-player arenas.*

Proof. The existence of an upper-bound for two-player multiweighted LU-energy games where energy levels and bounds are defined as a k -dimensional vector of energy levels, has been shown decidable in $2k$ -EXPTIME in [JLR13]. Using this result we immediately get that the existence of an upper-bound for two-player LU-energy games is decidable in 2-EXPTIME.

We now focus on the one-player case, and prove that it is decidable in NP. As a first step, assuming that an infinite path along which the energy level remains bounded exists, we prove that it can be detected by looking for adequate cycles (and whether they are reachable).

Indeed, assume that such a witness path exists from some given configuration (q_0, l_0) ; then there must be a configuration (q, l) occurring at least twice along that path, so that there is a witness path of the form $\rho_1 \cdot \rho_2^\omega$, where ρ_2 is a zero-cycle. Then two cases may occur:

- (i) either ρ_2 contains a simple zero-cycle: in that case, that simple cycle can be iterated after some initial prefix (and with some initial energy level). So in that case, there is a witness path of the form $\rho_1 \cdot \rho_2^\omega$ where ρ_2 is a simple zero-cycle; we can additionally assume that ρ_2 is *universal* (see Def. 17), i.e., that the energy level along ρ_2 is always above the initial energy level.
- (ii) or, if ρ_2 does not contain a simple zero-cycle, it must contain at least one simple positive cycle σ_1 and at least one simple negative cycle σ_2 . Then clearly enough two such cycles (and some acyclic paths $\alpha_{1 \rightarrow 2}$ and $\alpha_{2 \rightarrow 1}$ connecting between) suffice to get an infinite run in which energy level remains bounded (and above L): σ_1 can be iterated to gain energy, and when the energy level is high enough the path goes through $\alpha_{1 \rightarrow 2}$ to σ_2 , possibly cycling along σ_2 as long as there remains enough energy to take $\alpha_{2 \rightarrow 1}$ and reach σ_1 again. Then again, there exists a witness path of the form $\rho_1 \cdot \rho_2^\omega$ where ρ_2 is a universal 0-cycle starting with a positive simple cycle σ_1 .

Our NP algorithm will then non-deterministically select

- a candidate state q_0 to be the first state of ρ_2 ;
- either a candidate universal simple 0-cycle starting in q_0 , in order to witness Condition (i), or a candidate universal simple positive cycle starting at q_0 and a candidate negative simple cycle reachable and co-reachable from q_0 , in order to witness Condition (ii),

and check that those candidate indeed satisfy the requirements, including the fact that q_0 is reachable from the initial state. Those verifications can be performed in polynomial time, so that our algorithm is indeed in NP. \square

6.2.2. LW games

In [JLR13], the authors show how to find a minimal bound for multi-dimensional energy games with weak upper bounds. Multi-dimensional energy games are weighted games in which transitions are labeled by k -vectors of weights. Accordingly, configurations are pairs of the form (q, \vec{v}) where \vec{v} is a vector of integers of dimension k . The paper first recalls that the set Min_W of k -dimensional vectors that are winning for LW-energy games and are minimal (w.r.t. componentwise vector comparison) is upward-closed, and effectively computable, using a coverability tree.

The technique obviously applies to our LW-energy games, which are games of dimension 1. However, with games of dimension 1, one needs not build a coverability tree, and we propose below a more efficient technique to check existence of a weak upper energy bound allowing a winning strategy.

Theorem 32. *Existence of a bound for an LW-energy game can be decided in PTIME for single player, and in $NP \cap coNP$ for two players.*

Proof. We proceed in two steps. We first prove that if a weak bound exists for an LW-energy game, then there is a bound smaller or equal to $K_W = w_{max}^2 \cdot \frac{|Q| \cdot |Q-1|}{2}$. We then show that knowing K_W , one can apply relaxation techniques to find whether a bound lower than K_W exists.

Lemma 33. *Let G be an arena and L be a lower bound. There, if there exists a weak upper bound W such that an LW-energy game is winning for player 1, then there exists a bound smaller than $K_W = w_{max}^2 \cdot \frac{|Q| \cdot |Q-1|}{2}$, where w_{max} is the maximal absolute value of weights in G .*

Proof. We establish the bound K_W by analyzing the shape of a coverability tree of the arena, following the construction proposed by [JLR13]. A coverability tree is a structure $T = (N, \lambda, \rightarrow)$ where N is a set of nodes, λ is a labeling (a pair of state and vector of dimension k). The construction of a coverability tree is performed inductively, and starts from the initial node n_0 with $\lambda(n_0) = (q_0, 0^k)$. Each node n with $\lambda(n) = (q, \mathbf{v})$ has a successor n' with $\lambda(n') = (q', \mathbf{v}')$ if there exists a transition (q, p, q') and $\mathbf{v} = \mathbf{v}' + p$ (where v is a k -dimensional vector of energy levels, and p a k -dimensional vector of weights). The construction stops as soon as one component in \mathbf{v}' is negative (this node is losing, and witnesses a violation of the lower bound), or if (q', \mathbf{v}') has a predecessor n'' in its path from n_0 such that $\lambda(n'') = (q', \mathbf{v}'')$ with $\mathbf{v}'' \leq \mathbf{v}'$ (this node is said winning). Configurations form a well-quasi-order, so the construction of this coverability tree eventually terminates. Then Player 1 wins the game iff it has a strategy to guide a run to a winning configuration from n_0 . Indeed, if Player 1 can enforce a run ρ from q_0 to some state q and then a cycle φ on q , then he can repeat it an arbitrary number of times, as the energy levels on all components increase. Now, it remains to fix the weak upper bounds. Taking as upper bounds $W[1..k]$ the maximal energy levels reached along the winning path from q_0 and during the first iteration of the cycle does not prevent any transition of path $\rho.\varphi$, nor

of any path of the form $\rho.\varphi^k$ for any value of k . So W gives safe upper bounds on all components. Conversely, if Player 1 cannot avoid a losing configuration, then he loses the game.

For the single player case, it is hence sufficient to find a path to a winning configuration in the tree, hence find a run which size is at most the height of the coverability tree. Similarly, in the two-player setting, Player 1 has to enforce winning paths, that have the same maximal height.

One can notice that in dimension 1, along a run, if a configuration (q, v) appears, then there can be at most v occurrences of configurations of the form (q, v') that are incomparable with their predecessors. Indeed, every configuration of the form (q, v') such that $v' \geq v$ is greater than (q, v) . The first nodes immediately after n_0 are of the form n_1 with $\lambda(n_1) = (q_1, v_1)$, where $v_1 \leq w_{\max}$. Along a path $n_0.n_1 \dots$, the next occurrence of a node n_k referring to a state different from q_1 occurs at latest w_{\max} steps after n_1 , i.e. in a node with label $\lambda(N_k) = (q_k, v_k)$ with $v_k \leq 2.w_{\max}$, and so on. Hence the maximal size of path in the coverability tree is bounded by the size of a shuffle of sequences $(q_i, w_i)(q_i, w_i - 1) \dots$ of similar states with decreasing weight (one sequence per state), where $w_i \leq i.w_{\max}$ and bounded by $L_{T_G} = \sum_{i=1}^{|Q|} i.w_{\max} = w_{\max} \cdot \frac{|Q| \cdot (|Q| - 1)}{2}$. So $K_W = w_{\max}^2 \cdot \frac{|Q| \cdot (|Q| - 1)}{2}$. \square

Now building the coverability tree for an energy game with a single weight is not useful. In the worst case, checking for existence of a winning path (or of a winning strategy) may require to explore a tree of size in $O(Q^{L_{T_G}})$, which can be very inefficient. We now show how to check existence of a weak upper energy bound smaller than K_W in PTIME.

We recall that if Player 1 wins a LW-energy game with a weak bound W , she wins for every bound $W' \geq W$. Similarly, if Player 1 loses a LW-energy game with a weak bound W , she loses for every bound $W' \leq W$. For a single player, the algorithm of [BFL⁺08] decides whether a LW-energy game is winning for a given bound B in PTIME (for a single player). It uses a modified version of the Bellman-Ford algorithm, to detect lassos where the cyclic part is a positive cycle. Interestingly, in this algorithm, the bound B is only used during relaxation : the maximal energy $E(q, q')$ needed to go from q to q' via transition (q'', p, q') is $\min(B, E(q, q'') + p)$ if $E(q, q') \leq E(q, q'') + p$. The Bellman-Ford Algorithm runs performs $O(|Q|^3)$ relaxations, hence runs in $O(|Q|^3 \cdot r)$ where r is the cost of a relaxation. However, the number of relaxations does not depend on bound B . Relaxation computes a minimal value between two integers smaller than B , which can be done in $O(\log B)$. As in our case we run the algorithm with $B = K_W$, our algorithm runs in $O(|Q|^3 \cdot \log(w_{\max} \cdot Q))$. Hence checking existence of a weak bound for a single player can be done in PTIME w.r.t. the size of G . If the answer is negative, then following the results of [JLR13], no bounds exist.

For the two-player case, strategies are simple: if the energy level is too low to reach a winning configuration, then Player 1 should play to accumulate energy, and otherwise she should play to move to a winning configuration. Hence, finding a strategy simply consists in deciding for each states q a pair of states (q_e for the energy accumulation strategy, q_p for the game progress) and a threshold value

v_t allowing to choose to move to q_e or to q_t . One can hence guess a strategy $\sigma : Q \times [L, K_W] \rightarrow Q$ for Player 1 in PTIME, and then verify that its opponent has no strategy to violate the lower bound (i.e that no configuration (q, v) with $v < L$ is reachable. This can be done again using the Bellman-Ford algorithm, that in addition to computing shortest paths, detects negative cycles. Conversely, one can guess a strategy for Player 2, and use the algorithm above for a single player. The two-player case is hence in $\text{NP} \cap \text{coNP}$.

Remark 34. *Since we can derive an exponential upper bound U_{\max} on the values of U that allow winning outcomes in LU-energy games or LSU^* -energy games, and W_{\max} on the values of W in LW-energy games, we can give algorithms to characterize optimal values for those bounds: checking that a given value B is optimal is achieved by checking that B is indeed a valid bound, and that $B - 1$ is not. Computing such optimal bounds can be performed using binary search, which requires a polynomial number of such verifications.*

In the one-player LW-energy games setting, checking that a given value W is the optimal weak-upper bound can be achieved in polynomial time. For one-player LU-energy games, checking that a given value U is the optimal strong upper bound is in DP. In two-player LU-energy games, checking optimality of a strong upper bound is in EXPTIME. For weak upper bounds in two-player LW-energy games, since checking the existence of an infinite run for a fixed weak upper bound is in $\text{NP} \cap \text{coNP}$ (see [BFL⁺08]), we can guess and check polynomial certificates that a given bound B is valid and that $B - 1$ is not, and symmetrically. So checking whether a given bound is the optimal weak upper bound is in $\text{NP} \cap \text{coNP}$.

7. Conclusion

This paper has considered several variants of energy games. The first variant defines games with upper and lower bound constraints, combined with reachability objectives. The second variant defines games with a strong lower bound and a soft upper bound, which can be temporarily exceeded. In the one player case, complexities ranges from PTIME to PSPACE-complete, and in the two-player case from $\text{NP} \cap \text{coNP}$ to EXPTIME-complete. In general, the complexity is the same for a reachability and for an infinite run objective. Interestingly, for LW-energy games, the complexity of the single player case is PTIME, but reachability objectives require exponential memory (in the size of the weak upper bound) while strategies are memoryless for infinite run objectives. The associated bound existence problems range from PTIME to PSPACE-complete in the one-player setting, and from PTIME to 2-EXPTIME for two-player. However, proving optimality of a given bound is not harder than the energy game itself.

A possible extension of this work is to consider energy games with mean-payoff functions and discounted total payoff, both for the energy level and for the violation constraints.

References

- [And06] Daniel Andersson. An improved algorithm for discounted payoff games. In Janneke Huitink and Sophia Katrenko, editors, *Proceedings of the 11th ESSLLI Student Session*, pages 91–98, August 2006.
- [BFL⁺08] Patricia Bouyer, Uli Fahrenberg, Kim Guldstrand Larsen, Nicolas Markey, and Jiří Srba. Infinite runs in weighted timed automata with energy constraints. In Franck Cassez and Claude Jard, editors, *Proceedings of the 6th International Conferences on Formal Modelling and Analysis of Timed Systems (FORMATS'08)*, volume 5215 of *Lecture Notes in Computer Science*, pages 33–47. Springer-Verlag, September 2008.
- [BFLM10] Patricia Bouyer, Uli Fahrenberg, Kim Guldstrand Larsen, and Nicolas Markey. Timed automata with observers under energy constraints. In Karl Henrik Johansson and Wang Yi, editors, *Proceedings of the 13th International Workshop on Hybrid Systems: Computation and Control (HSCC'10)*, pages 61–70. ACM Press, April 2010.
- [BHM⁺17] Patricia Bouyer, Piotr Hofman, Nicolas Markey, Mickael Randour, and Martin Zimmermann. Bounding average-energy games. In Javier Esparza and Andrzej Murawski, editors, *Proceedings of the 20th International Conference on Foundations of Software Science and Computation Structure (FoSSaCS'17)*, volume 10203 of *Lecture Notes in Computer Science*, pages 179–195. Springer-Verlag, April 2017.
- [BLM12] Patricia Bouyer, Kim Guldstrand Larsen, and Nicolas Markey. Lower-bound constrained runs in weighted timed automata. In *Proceedings of the 9th International Conference on Quantitative Evaluation of Systems (QEST'12)*, pages 128–137. IEEE Comp. Soc. Press, September 2012.
- [BMR⁺15] Patricia Bouyer, Nicolas Markey, Mickael Randour, Kim Guldstrand Larsen, and Simon Laursen. Average-energy games. In Javier Esparza and Enrico Tronci, editors, *Proceedings of the 6th International Symposium on Games, Automata, Logics and Formal Verification (GandALF'15)*, volume 193 of *Electronic Proceedings in Theoretical Computer Science*, pages 1–15, September 2015.
- [CD12] Krishnendu Chatterjee and Laurent Doyen. Energy parity games. *Theoretical Computer Science*, 458:49–60, November 2012.
- [CdAHS03] Arindam Chakrabarti, Luca de Alfaro, Thomas A. Henzinger, and Mariëlle Stoelinga. Resource interfaces. In Rajeev Alur and Insup Lee, editors, *Proceedings of the 3rd International Conference on Embedded Software (EMSOFT'03)*, volume 2855 of *Lecture Notes in Computer Science*, pages 117–133. Springer-Verlag, October 2003.

- [CDH17] Krishnendu Chatterjee, Laurent Doyen, and Thomas A. Henzinger. The cost of exactness in quantitative reachability. In Luca Aceto, Giorgio Bacci, Giovanni Bacci, Anna Ingólfssdóttir, Axel Legay, and Radu Mardare, editors, *Models, Algorithms, Logics and Tools: Essays Dedicated to Kim Guldstrand Larsen on the Occasion of His 60th Birthday*, volume 10460 of *Lecture Notes in Computer Science*, pages 367–381. Springer-Verlag, August 2017.
- [CDHR10] Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger, and Jean-François Raskin. Generalized mean-payoff and energy games. In Kamal Lodaya and Meena Mahajan, editors, *Proceedings of the 30th Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'10)*, volume 8 of *Leibniz International Proceedings in Informatics*, pages 505–516. Leibniz-Zentrum für Informatik, December 2010.
- [CJL⁺09] Franck Cassez, Jan J. Jensen, Kim Guldstrand Larsen, Jean-François Raskin, and Pierre-Alain Reynier. Automatic synthesis of robust and optimal controllers – an industrial case study. In Rupak Majumdar and Paulo Tabuada, editors, *Proceedings of the 12th International Workshop on Hybrid Systems: Computation and Control (HSCC'09)*, volume 5469 of *Lecture Notes in Computer Science*, pages 90–104. Springer-Verlag, April 2009.
- [CRR14] Krishnendu Chatterjee, Mickael Randour, and Jean-François Raskin. Strategy synthesis for multi-dimensional quantitative objectives. *Acta Informatica*, 51(3-4):129–163, June 2014.
- [DDG⁺10] Aldric Degorre, Laurent Doyen, Raffaella Gentilini, Jean-François Raskin, and Szymon Toruńczyk. Energy and mean-payoff games with imperfect information. In Anuj Dawar and Helmut Veith, editors, *Proceedings of the 24th International Workshop on Computer Science Logic (CSL'10)*, volume 6247 of *Lecture Notes in Computer Science*, pages 260–274. Springer-Verlag, August 2010.
- [DM18] Dario Della Monica and Aniello Murano. Parity-energy ATL for qualitative and quantitative reasoning in MAS. In Elisabeth André, Sven Koenig, Mehdi Dastani, and Gita Sukthankar, editors, *Proceedings of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS'18)*, pages 1441–1449. International Foundation for Autonomous Agents and Multiagent Systems, 2018.
- [EF13] Daniel Ejsing-Dunn and Lisa Fontani. Infinite runs in recharge automata. Master’s thesis, Computer Science Department, Aalborg University, Denmark, June 2013.

- [EM79] Andrzej Ehrenfeucht and Jan Mycielski. Positional strategies for mean payoff games. *International Journal of Game Theory*, 8(2):109–113, June 1979.
- [FJ13] John Fearnley and Marcin Jurdziński. Reachability in two-clock timed automata is PSPACE-complete. In Fedor V. Fomin, Rusins Freivalds, Marta Kwiatkowska, and David Peleg, editors, *Proceedings of the 40th International Colloquium on Automata, Languages and Programming (ICALP'13) – Part II*, volume 7966 of *Lecture Notes in Computer Science*, pages 212–223. Springer-Verlag, July 2013.
- [FJLS11] Uli Fahrenberg, Line Juhl, Kim Guldstrand Larsen, and Jiří Srba. Energy games in multiweighted automata. In Antonio Cerone and Pekka Pihlajasaari, editors, *Proceedings of the 8th International Colloquium on Theoretical Aspects of Computing (ICTAC'11)*, volume 6916 of *Lecture Notes in Computer Science*, pages 95–115. Springer-Verlag, August-September 2011.
- [GHOW10] Stefan Göller, Christoph Haase, Joël Ouaknine, and James Worrell. Model checking succinct and parametric one-counter automata. In Samson Abramsky, Cyril Gavoille, Claude Kirchner, Friedhelm Meyer auf der Heide, and Paul G. Spirakis, editors, *Proceedings of the 37th International Colloquium on Automata, Languages and Programming (ICALP'10) – Part II*, volume 6199 of *Lecture Notes in Computer Science*, pages 575–586. Springer-Verlag, July 2010.
- [HKOW09] Christoph Haase, Stephan Kreutzer, Joël Ouaknine, and James Worrell. Reachability in succinct and parametric one-counter automata. In Mario Bravetti and Gianluigi Zavattaro, editors, *Proceedings of the 20th International Conference on Concurrency Theory (CONCUR'09)*, volume 5710 of *Lecture Notes in Computer Science*, pages 369–383. Springer-Verlag, September 2009.
- [HMR19] Loïc Hélouët, Nicolas Markey, and Ritam Raha. Reachability games with relaxed energy constraints. In Jérôme Leroux and Jean-François Raskin, editors, *Proceedings Tenth International Symposium on Games, Automata, Logics, and Formal Verification, GandALF 2019, Bordeaux, France, 2-3rd September 2019*, volume 305 of *EPTCS*, pages 17–33, 2019.
- [Hun15] Paul Hunter. Reachability in succinct one-counter games. In Mikołaj Bojańczyk, Sławomir Lasota, and Igor Potapov, editors, *Proceedings of the 9th Workshop on Reachability Problems in Computational Models (RP'15)*, volume 9328 of *Lecture Notes in Computer Science*, pages 37–49. Springer-Verlag, September 2015.
- [JLR13] Line Juhl, Kim Guldstrand Larsen, and Jean-François Raskin. Optimal bounds for multiweighted and parametrised energy games. In

- Zhiming Liu, Jim Woodcock, and Yunshan Zhu, editors, *Theories of Programming and Formal Methods – Essays Dedicated to Jifeng He on the Occasion of His 70th Birthday*, volume 8051 of *Lecture Notes in Computer Science*, pages 244–255. Springer-Verlag, 2013.
- [JLS07] Marcin Jurdziński, François Laroussinie, and Jeremy Sproston. Model checking probabilistic timed automata with one or two clocks. In Orna Grumberg and Michael Huth, editors, *Proceedings of the 13th International Conference on Tools and Algorithms for Construction and Analysis of Systems (TACAS’07)*, volume 4424 of *Lecture Notes in Computer Science*, pages 170–184. Springer-Verlag, March 2007.
- [JLS15] Marcin Jurdziński, Ranko Lazić, and Sylvain Schmitz. Fixed-dimensional energy games are in pseudo-polynomial time. In Magnús M. Halldórsson, Kazuo Iwana, Naoki Kobayashi, and Bettina Speckmann, editors, *Proceedings of the 42nd International Colloquium on Automata, Languages and Programming (ICALP’15) – Part II*, volume 9135 of *Lecture Notes in Computer Science*, pages 260–272. Springer-Verlag, July 2015.
- [Mar75] Donald A. Martin. Borel determinacy. *Annals of Mathematics*, 102(2):363–371, September 1975.
- [Rei16] Julien Reichert. On the complexity of counter reachability games. *Fundamenta Informaticae*, 143(3-4):415–436, 2016.
- [VCD⁺15] Yaron Velner, Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger, Alexander Rabinovich, and Jean-François Raskin. The complexity of multi-mean-payoff and multi-energy games. *Information and Computation*, 241:177–196, April 2015.
- [ZP96] Uri Zwick and Mike Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158(1-2):343–359, May 1996.