

Limit Your Consumption!

Finding Bounds in Average-energy Games*

Kim G. Larsen

Department of Computer Science, Aalborg University
Aalborg, Denmark

kgl@cs.aau.dk

Simon Laursen

simlau@cs.aau.dk

Martin Zimmermann

Reactive Systems Group, Saarland University
Saarbrücken, Germany

zimmermann@react.uni-saarland.de

Energy games are infinite two-player games played in weighted arenas with quantitative objectives that restrict the consumption of a resource modeled by the weights, e.g., a battery that is charged and drained. Typically, upper and/or lower bounds on the battery capacity are part of the problem description. Here, we consider the problem of determining upper bounds on the average accumulated energy or on the capacity while satisfying a given lower bound, i.e., we do not determine whether a given bound is sufficient to meet the specification, but if there exists a sufficient bound to meet it.

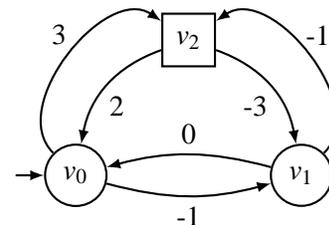
In the classical setting with positive and negative weights, we show that the problem of determining the existence of a sufficient bound on the long-run average accumulated energy can be solved in doubly-exponential time. Then, we consider recharge games: here, all weights are negative, but there are recharge edges that recharge the energy to some fixed capacity. We show that bounding the long-run average energy in such games is complete for exponential time. Then, we consider the existential version of the problem, which turns out to be solvable in polynomial time: here, we ask whether there is a recharge capacity that allows the system player to win the game.

We conclude by studying tradeoffs between the memory needed to implement strategies and the bounds they realize. We give an example showing that memory can be traded for bounds and vice versa. Also, we show that increasing the capacity allows to lower the average accumulated energy.

1 Introduction

Quantitative games provide a natural framework for synthesizing controllers with resource restrictions and for performance requirements for reactive systems with an uncontrollable environment. In a traditional two-player graph game of infinite duration (see [18]), two players, Player 0 (who represents the system to be synthesized) and Player 1 (representing the antagonistic environment), construct an infinite path by moving a pebble through a graph, which describes the interaction between the system and its environment. The objective, which encodes the controller's specification, determines the winner of such a play. Quantitative games extend classical ones by having weights on edges for modeling costs, consumption or rewards, and a quantitative objective to encode the specification in terms of the weights.

Consider the game depicted to the right: we interpret negative weights as energy consumption and correspondingly positive weights as recharges. Then, Player 0 (who moves the pebble at the circled vertices) can always maintain an energy level (the sum of the weights seen along a play prefix starting with energy 0) between zero and five using the following strategy: when at vertex v_0 with non-zero energy level go to vertex v_1 , otherwise go to vertex v_2 in



*Supported by the DFG project TriCS (ZI 1516/1-1), the ERC Advanced Grant LASSO and the EU FET projects SENSATION and CASSTING.

order to satisfy the lower bound. At vertex v_1 she moves to v_0 if the energy level is zero, otherwise to v_2 . It is straightforward to verify that the strategy has the desired property when starting at the initial vertex v_0 with initial energy 0. However, this strategy requires memory to implement, as its choices depend on the current energy level.

Quantitative games [2, 9, 24] and objectives such as mean-payoff [7, 26, 28], energy [4, 10, 19], and their combination [12] have attracted considerable attention recently. The focus has been on establishing the computational complexity of deciding whether Player 0 wins the game and on memory requirements. In mean-payoff games, Player 0's goal is to optimize the long-run average gain per edge taken, whereas in energy games the goal is to keep the accumulated energy within given bounds. Recently, the average-energy objective was introduced [5] to capture the specification in an industrial case study [8]. In this study, the authors synthesize a controller to operate an oil pump using timed games and UPPAAL TiGA. The controller has to keep the amount of oil in an accumulator within given bounds while minimizing the average amount of oil in the accumulator in the long run. A discrete version of this problem is exactly an average-energy game, where the goal for Player 0 is to optimize the long-run average accumulated energy during a play while keeping the accumulated energy within given bounds.

Recall the introductory example above. The strategy for Player 0 described there realizes the long-run average 4: the consistent play $v_0(v_2v_0v_1)^\omega$ with energy levels $0, (3, 5, 4)^\omega$ has average 4, obtained by dividing the sum of the levels in the period by the length of the period. Every other consistent play has a smaller or equal average.

The computational complexity of these quantitative objectives are typically studied with respect to given bounds on the energy level or given thresholds on the mean-payoff or on the average accumulated energy. In this work, we consider the variants where the bounds and thresholds are existentially quantified instead of given as part of the input, i.e., we ask if there exist bounds and thresholds such that Player 0 has a winning strategy. This question is natural for models with bounds and thresholds as it is desirable to know if a given model is realizable for some bounds. In a second step, one would then determine the minimal bounds for which Player 0 is able to win.

In particular, we study existential questions on two different game models, average-energy games and average-bounded recharge games. Average-energy games are defined as in [5] with both positive and negative weights on edges whereas in average-bounded recharge games all weights are negative, but there are designated recharge-edges that recharge the energy to some fixed capacity.

Our contribution. For average-energy games, we show that the problem of deciding whether there exists a threshold to which Player 0 can bound the long-run average accumulated energy while keeping the accumulated energy non-negative can be solved in doubly-exponential time. To this end, we show that the problem is equivalent to determining whether the maximal energy level can be uniformly bounded by a strategy. The latter problem is known to be in 2EXPTIME [19]. The challenging part is to construct a strategy that uniformly bounds the energy from the strategy that only bounds the long-run average accumulated energy, but might reach arbitrarily high energy levels. But whenever the energy level increases above the given threshold, it has to drop below it at some later point. Thus, we can always play like in a situation where the peak between these two threshold crossings is as small as possible. This yields a new strategy that bounds the energy level. Our result is one step towards solving the open problem of solving lower-bounded average-energy games with a given threshold [5].

For average-bounded recharge games, we show that given a bound on the long-run average energy, deciding the winner is EXPTIME-complete. For the existential versions of the problem, we show that it remains EXPTIME-hard when the recharge capacity is quantified and the average threshold is given. The problem becomes solvable in polynomial time when only the recharge capacity is considered: here, we

ask whether there is a recharge capacity such that Player 0 wins the game with respect to this capacity.

Finally, we study tradeoffs between the different bounds and the memory requirements of winning strategies, and show that increasing the upper bound on the maximal energy level allows to improve the average energy level and memory can be traded for smaller upper bounds and vice versa.

Related work. The average energy objective was first introduced in [25] under the name total-reward but has until recently not undergone a systematic study. Independently, it was studied (under the name total-payoff) for Markov decision processes and stochastic games [3], and [5] presented a comprehensive investigation into the problem in the deterministic case. The latter also considered extensions where the average-energy objective is combined with bounds on the energy, which is the model we consider here.

Several other games with combined objectives have been introduced such as mean-payoff parity [11], energy-parity [10], multi-dimensional energy [16], multi-dimensional mean-payoff [26] and the combination of multi-dimensional energy, mean-payoff and parity [12]. In [6], consumption games are studied where edges only have negative weights, and some distinguished edges recharge the energy to a level determined by Player 0. This model is related to recharge games, but in recharge games the recharge capacity is given and we consider average-bounded objectives. Existential questions in games have been studied before in the form of determining the emptiness of a set of bounds that allow Player 0 to win a quantitative game, e.g., for multi-dimensional energy games with upper bounds [19] and for games with objectives in parameterized generalizations of LTL [1, 17, 22, 27].

2 Definitions

An *arena* $\mathcal{A} = (V, V_0, V_1, E, v_I)$ consists of a finite directed graph (V, E) without terminal vertices, a partition $V = V_0 \uplus V_1$ of the vertices, and an initial vertex $v_I \in V$. Vertices in V_0 are under Player 0's control and are drawn as circles, whereas vertices in V_1 are under Player 1's control and drawn as rectangles. A play in \mathcal{A} is an infinite path $\rho = v_0 v_1 v_2 \dots$ with $v_0 = v_I$. A *game* $\mathcal{G} = (\mathcal{A}, \text{Win})$ consists of an arena \mathcal{A} , and a set $\text{Win} \subseteq V^\omega$ of winning plays for Player 0, the *objective* of \mathcal{G} . The objectives we consider are induced by weight functions, assigning integer weights to edges, which are encoded in binary. We say an algorithm runs in *pseudo-polynomial time*, if it runs in polynomial time in the number of vertices and in the largest absolute weight. An algorithm runs in *polynomial time*, if it runs in polynomial time in the number of vertices and in the size of the encoding of the largest absolute weight.

A *strategy* for Player $i \in \{0, 1\}$ is a mapping $\sigma_i: V^*V_i \rightarrow V$ such that $(v, \sigma_i(wv)) \in E$ for all $wv \in V^*V_i$. A play $v_0 v_1 v_2 \dots$ is *consistent* with a strategy σ_i for Player i if $v_{n+1} = \sigma_i(v_0 v_1 \dots v_n)$ for every n with $v_n \in V_i$. A strategy σ_0 for Player 0 is winning for the game $\mathcal{G} = (\mathcal{A}, \text{Win})$ if every play that is consistent with σ_0 is in Win . We say that Player 0 wins \mathcal{G} if she has a winning strategy for \mathcal{G} . We define $\text{Prefs}(\sigma)$ to denote the set of finite play prefixes that are consistent with σ . We denote the last vertex of a non-empty word w by $\text{Last}(w)$.

A *memory structure* $\mathcal{M} = (M, m_I, \text{Upd})$ for an arena (V, V_0, V_1, E, v_I) consists of a finite set M of memory states, an initial memory state $m_I \in M$, and an update function $\text{Upd}: M \times E \rightarrow M$. The update function can be extended to $\text{Upd}^+: V^+ \rightarrow M$ in the usual way: $\text{Upd}^+(v_0) = m_I$ and $\text{Upd}^+(v_0 \dots v_n v_{n+1}) = \text{Upd}(\text{Upd}^+(v_0 \dots v_n), (v_n, v_{n+1}))$. A next-move function (for Player i) $\text{Nxt}: V_i \times M \rightarrow V$ has to satisfy $(v, \text{Nxt}(v, m)) \in E$ for all $v \in V_i$ and all $m \in M$. It induces a strategy σ for Player i via $\sigma(v_0 \dots v_n) = \text{Nxt}(v_n, \text{Upd}^+(v_0 \dots v_n))$. A strategy is called *finite-state (positional)* if it can be implemented by a memory structure (with a single state). Intuitively, the next move of a positional strategy only depends on the last vertex of the play prefix. An arena $\mathcal{A} = (V, V_0, V_1, E, v_I)$ and a memory structure $\mathcal{M} = (M, m_I, \text{Upd})$ for \mathcal{A} induce the expanded arena $\mathcal{A} \times \mathcal{M} = (V \times M, V_0 \times M, V_1 \times M, E', (v_I, m_I))$

where $((v, m), (v', m')) \in E'$ if and only if $(v, v') \in E$ and $\text{Upd}(m, (v, v')) = m'$. Each play $v_0 v_1 v_2 \dots$ in \mathcal{A} has a unique extended play $(v_0, m_0)(v_1, m_1)(v_2, m_2) \dots$ in $\mathcal{A} \times \mathcal{M}$ defined by $m_0 = m_I$ and $m_{n+1} = \text{Upd}(m_n, (v_n, v_{n+1}))$, i.e., $m_n = \text{Upd}^+(v_0 \dots v_n)$. A game $\mathcal{G} = (\mathcal{A}, \text{Win})$ is *reducible* to $\mathcal{G}' = (\mathcal{A}', \text{Win}')$ via \mathcal{M} , written $\mathcal{G} \leq_{\mathcal{M}} \mathcal{G}'$, if $\mathcal{A}' = \mathcal{A} \times \mathcal{M}$ and every play ρ in \mathcal{G} is won by the player who wins the extended play ρ' in \mathcal{G}' , i.e., $\rho \in \text{Win}$ if, and only if, $\rho' \in \text{Win}'$.

Lemma 1. *If $\mathcal{G} \leq_{\mathcal{M}} \mathcal{G}'$ and Player i has a positional winning strategy for \mathcal{G}' , then she has a finite-state winning strategy for \mathcal{G} which is implemented by \mathcal{M} .*

3 Finding Bounds in Average-energy Games

In this section, we study average-energy games with existentially quantified bounds on the average accumulated energy: our main theorem shows that these games are solvable in doubly-exponential time.

A weight function for an arena (V, V_0, V_1, E, v_I) is a function $w: E \rightarrow \mathbb{Z}$ mapping every edge to an integer weight. The energy level of a play prefix is the accumulated weight of its edges, i.e., $\text{EL}(v_0 \dots v_n) = \sum_{i=0}^{n-1} w(v_i, v_{i+1})$. We consider several objectives obtained by specifying upper and lower bounds on the energy level and on the long-run average accumulated energy.

- The lower-bounded energy objective requires Player 0 to keep the energy level non-negative:

$$\text{Energy}_L(w) = \{v_0 v_1 v_2 \dots \in V^\omega \mid \forall n. 0 \leq \text{EL}(v_0 \dots v_n)\}$$

- The lower- and upper-bounded energy objective requires Player 0 to keep the energy level always between 0 and some given upper bound cap , the so-called capacity:

$$\text{Energy}_{LU}(w, cap) = \{v_0 v_1 v_2 \dots \in V^\omega \mid \forall n. 0 \leq \text{EL}(v_0 \dots v_n) \leq cap\}$$

- The average-energy objective requires Player 0 to keep the long-run average of the accumulated energy below a given threshold t :

$$\text{AvgEnergy}(w, t) = \{v_0 v_1 v_2 \dots \in V^\omega \mid \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} \text{EL}(v_0 \dots v_i) \leq t\}$$

- Also, we consider conjunctions of objectives, i.e., the lower-bounded average-energy objective

$$\text{AvgEnergy}_L(w, t) = \text{Energy}_L(w) \cap \text{AvgEnergy}(w, t)$$

and the lower- and upper-bounded average-energy objective

$$\text{AvgEnergy}_{LU}(w, cap, t) = \text{Energy}_{LU}(w, cap) \cap \text{AvgEnergy}(w, t).$$

Note that we always assume the initial energy level to be zero. This is not a restriction, as one can always add a fresh initial vertex with an edge to the old initial vertex that is labeled by the desired initial energy level. Similarly, one can reduce arbitrary non-zero lower bounds to the case of the lower bound being zero, which is the one we consider here.

Decidability of determining the winner of a game with lower-bounded average-energy objective with a given threshold t is an open problem [5]. To take a step towards solving this problem, we consider the existential variant of the problem, i.e., we ask whether there exists some threshold t such that Player 0 wins the game with objective $\text{AvgEnergy}_L(w, t)$:

Problem 1. *Existence of a threshold in a lower-bounded average-energy game.*

Input: Arena $\mathcal{A} = (V, V_0, V_1, E, v_I)$ and $w: E \rightarrow \mathbb{Z}$

Question: *Exists a threshold $t \in \mathbb{N}$ s.t. Player 0 wins $(\mathcal{A}, \text{AvgEnergy}_L(w, t))$?*

We show that this problem is reducible to asking for the existence of an upper bound on the capacity cap . Note that such an upper bound also bounds the average accumulated energy. However, the converse is non-trivial as the average can be bounded while the energy level is unbounded. Formally, we consider the following problem:

Problem 2. *Existence of an upper bound in a lower- and upper-bounded energy game.*

Input: Arena $\mathcal{A} = (V, V_0, V_1, E, v_I)$ and $w: E \rightarrow \mathbb{Z}$

Question: *Exists a capacity $cap \in \mathbb{N}$ s.t. Player 0 wins $(\mathcal{A}, \text{Energy}_{LU}(w, cap))$?*

The main theorem of this section shows that the existence of a threshold in a lower-bounded average-energy game can be checked in doubly-exponential time. Our choice of encoding the weights influences the complexity of the problem: if the weights are encoded in unary, then the complexity drops to EXPTIME. Furthermore, the problem is trivially at least as hard as solving mean-payoff games.

Theorem 1. *The threshold problem for lower-bounded average-energy games is in 2EXPTIME.*

To prove this theorem, it suffices to show that Problem 1 and Problem 2 are equivalent, as the latter problem was shown to be in 2EXPTIME [19].

Lemma 2. *Let \mathcal{A} be an arena and let w be a weight function for \mathcal{A} . Player 0 wins $(\mathcal{A}, \text{AvgEnergy}_L(w, t))$ for some $t \in \mathbb{N}$ if, and only if, Player 0 wins $(\mathcal{A}, \text{Energy}_{LU}(w, cap))$ for some $cap \in \mathbb{N}$.*

Proof. It is clear that a winning strategy σ for $(\mathcal{A}, \text{Energy}_{LU}(w, cap))$ for some $cap \in \mathbb{N}$ is a winning strategy for $(\mathcal{A}, \text{AvgEnergy}_L(w, cap))$, as if the energy level is always below some cap , then the average energy is also bounded by cap .

For the other direction, assume that σ is a winning strategy for Player 0 in $(\mathcal{A}, \text{AvgEnergy}_L(w, t))$ for some $t \in \mathbb{N}$. Now, we want to construct a strategy σ' that is winning for Player 0 in $(\mathcal{A}, \text{Energy}_{LU}(w, cap))$ for some $cap \in \mathbb{N}$. Note that σ might bound the average to some value while the energy level might be unbounded. But whenever the energy level increases above t , it has to drop below t at some point. We use this property to construct a strategy σ' that bounds the energy level.

First, we need to introduce some notation. Fix a play prefix $w \in \text{Prefs}(\sigma)$ with $\text{EL}(w) > t$ and define

$$\text{Peak}(w) = \sup\{\text{EL}(wx) \mid wx \in \text{Prefs}(\sigma) \text{ and } \text{EL}(wx') > t \text{ for all } x' \sqsubseteq x\},$$

i.e., $\text{Peak}(w)$ is the supremum of the energy levels of prolongations of w that are consistent with σ and have not yet had an energy level below t . Applying König's Lemma [21] and the fact that σ is a winning strategy implies that the peak is always bounded.

Remark 1. *We have $\text{Peak}(w) \in \mathbb{N}$ for every $w \in \text{Prefs}(\sigma)$.*

For an energy level $c \in \mathbb{N}$ and a vertex $v \in V$ we define the set of possible ways to end up in vertex v with the energy level c playing consistently with σ as

$$\text{Real}(v, c) = \{w \in \text{Prefs}(\sigma) \mid \text{Last}(w) = v \text{ and } \text{EL}(w) = c\}.$$

For every combination (v, c) with $c > t$, we pick a representative from $\text{Real}(v, c)$ that minimizes the peak height among all such realizations, i.e., we define $\text{Rep}(v, c)$ to be an element w from $\text{Real}(v, c)$ with minimal peak-value $\text{Peak}(w)$ among the play prefixes in $\text{Real}(v, c)$. Note that $\text{Rep}(v, c)$ might be undefined, i.e., if there is no play prefix ending in v with energy level c .

Intuitively, we construct a new strategy that mimics the behavior of σ until the energy level increases above t . At this point, the history is replaced by the representative for the last vertex and the current energy level. Then, our new strategy mimics the behavior of σ with this history until the threshold t is again crossed from below. Then, the next representative is picked. This strategy satisfies an upper bound, as only a finite number of representatives, each with a bounded peak-value, are considered when mimicking σ . To formalize this, we recursively define $h: V^+ \rightarrow \text{Prefs}(\sigma)$ via $h(v_I) = v_I$ and

$$h(wv) = \begin{cases} \text{Rep}(v, \text{EL}(h(w)v)) & \text{if } \text{EL}(h(w)) \leq t \text{ and } \text{EL}(h(w)v) > t \\ h(w)v & \text{otherwise} \end{cases}$$

for a play prefix $wv \in V^+$ ending in a vertex v , i.e., $h(w)$ is the play prefix that simulates w . Now, we define the new strategy σ' via $\sigma'(w) = \sigma(h(w))$. The following remark implies that this is well-defined, although Rep and therefore h and σ might be undefined for certain inputs.

Remark 2. *Let w be consistent with σ' . Then, $h(w)$ is defined and consistent with σ , $\text{Last}(w) = \text{Last}(h(w))$, and $\text{EL}(w) = \text{EL}(h(w))$.*

Applying the remark inductively we conclude that $h(w)$ is defined for every play prefix w that is consistent with σ' . This implies that $\sigma'(w)$ is well-defined for every such w that ends in a vertex from V_0 . Furthermore, this also implies that σ' still satisfies the lower bound on the energy level.

Thus, it remains to prove that an upper bound exists. Let $\rho = v_0v_1v_2 \cdots$ be consistent with σ' and let n be such that $\text{EL}(v_0 \cdots v_n) \leq t$ and $\text{EL}(v_0 \cdots v_nv_{n+1}) > t$. If there is no such n , then σ bounds the energy level by t and we are done. Furthermore, define n' to be minimal with $n' > n + 1$ and $\text{EL}(v_0 \cdots v_{n'}) \leq t$ and $\text{EL}(v_0 \cdots v_{n'}v_{n'+1}) > t$ (if no such n' exists the reasoning is analogous). As the energy level between the positions $n + 1$ and n' never crosses the threshold t from below, we are always in the second case of the definition of h . Thus, after the play prefix $v_0 \cdots v_{n+1}$, the strategy σ' mimics the behavior of σ after the prefix $h(v_0 \cdots v_{n+1}) = \text{Rep}(v_{n+1}, \text{EL}(v_0 \cdots v_{n+1}))$. Therefore, the energy level between these two positions is bounded by $\text{Peak}(\text{Rep}(v_{n+1}, \text{EL}(v_0 \cdots v_{n+1})))$. As we only take those representatives into account that have an energy level between $t + 1$ and $t + W$, where W is the largest positive weight in the image of w , the energy level of the play is bounded by the maximal peak of one of these representatives. Finally, this bound is uniform for all plays that are consistent with σ' . Thus, σ' is winning in the game $(\mathcal{A}, \text{AvgEnergy}_L(w, \text{cap}))$ for some cap . \square

Note that we do not obtain any upper bounds on the energy level or on the long-run average energy realized by σ' , as they depend on properties of σ . One can even construct examples that show these values to be arbitrarily large by starting with a *bad* winning strategy σ for the energy game.

4 Finding Bounds in Average-bounded Recharge Games

In this section, we study a variation of energy games called recharge games (the name is inspired by recharge automata, first introduced in [14]). In such games, there are designated recharge edges that recharge the energy to some given capacity. All other edges have non-positive cost, i.e., they only decrease the energy level or leave it unchanged. This is reminiscent of so-called consumption games [6], where Player 0 picks the new energy level while traversing a recharge edge. There, one is interested in which initial energy levels allow Player 0 to win and to compute upper bounds on the recharge levels picked by Player 0.

In this section, we go beyond just bounding the energy level by also considering bounds on the average accumulated energy, as we have done for average-energy games. However, the resulting games are intractable, as soon as the threshold on the average is part of the input. These results are presented in Subsection 4.1. To overcome the high complexity, in Subsection 4.2 we consider the problem where the recharge capacity is existentially quantified: this problem is solvable in polynomial time by a reduction to three-color parity games.

Here, we consider weight functions with only non-positive weights and a special recharge action R , i.e., $w: E \rightarrow -\mathbb{N} \cup \{R\}$. The recharge action R returns the energy level to some given upper bound capacity cap . The recharge energy level is the energy left since the last recharge action, which is defined as $EL_{cap}(v_0 \cdots v_n) = cap + EL(x)$, where x is the longest suffix of $v_0 \cdots v_n$ without an R -edge, i.e., $w(v_j, v_{j+1}) \neq R$ for all (v_j, v_{j+1}) in x , which implies that a play starts with energy level cap . We define the objective of a recharge game as

$$\text{Recharge}(w, cap) = \{v_0 v_1 v_2 \cdots \in V^\omega \mid \forall n. EL_{cap}(v_0 \cdots v_n) \geq 0\}$$

and the average-bounded version as

$$\text{AvgRecharge}(w, cap, t) = \{v_0 v_1 v_2 \cdots \in V^\omega \mid \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} EL_{cap}(v_0 \cdots v_i) \leq t\} \cap \text{Recharge}(w, cap).$$

4.1 Solving Average-bounded Recharge Games

First, we show that solving average-bounded recharge games for a given threshold t and a given recharge capacity cap is EXPTIME-complete and that the problem is still EXPTIME-hard, if the capacity is existentially quantified and only the threshold is given. Formally, we are interested in the following problems:

Problem 3. *Solving Average-bounded recharge games*

Input: Arena $\mathcal{A} = (V, V_0, V_1, E, v_I)$, $w: E \rightarrow -\mathbb{N} \cup \{R\}$, $cap \in \mathbb{N}$, and $t \in \mathbb{N}$.

Question: Does Player 0 win $(\mathcal{A}, \text{AvgRecharge}(w, cap, t))$?

Problem 4. *Solving Average-bounded recharge games with existentially quantified capacity*

Input: Arena $\mathcal{A} = (V, V_0, V_1, E, v_I)$, $w: E \rightarrow -\mathbb{N} \cup \{R\}$, and $t \in \mathbb{N}$.

Question: Exists $cap \in \mathbb{N}$ s.t. Player 0 wins $(\mathcal{A}, \text{AvgRecharge}(w, cap, t))$?

First, we consider Problem 3.

Theorem 2. *Solving average-bounded recharge games is EXPTIME-complete.*

We begin the proof by presenting an exponential time algorithm for solving average-bounded recharge games by reducing them to mean-payoff games, similarly to the reduction from lower- and upper-bounded energy games to mean-payoff games [5]. The mean-payoff objective is given by

$$\text{MeanPayoff}(w, t) = \{v_0 v_1 v_2 \cdots \in V^\omega \mid \limsup_{n \rightarrow \infty} \frac{1}{n} EL(v_0 \cdots v_{n-1}) \leq t\}.$$

Lemma 3. *Average-bounded recharge games can be solved in exponential time.*

Proof. Fix an arena $\mathcal{A} = (V, V_0, V_1, E, v_I)$, $w: E \rightarrow -\mathbb{N} \cup \{R\}$, $cap \in \mathbb{N}$, and $t \in \mathbb{N}$. We construct a memory structure $\mathcal{M} = (M, m_I, \text{Upd})$ to reduce the average-bounded recharge game to a mean-payoff game. To this end, let $M = \{0, \dots, cap\} \cup \{\perp\}$, $m_I = cap$, $\text{Upd}(\perp, (v, v')) = \perp$, and

$$\text{Upd}(c, (v, v')) = \begin{cases} cap & \text{if } w(v, v') = R, \\ c + w(v, v') & \text{if } c + w(v, v') \geq 0, \\ \perp & \text{if } c + w(v, v') < 0. \end{cases}$$

Intuitively, the memory structure keeps track of the energy level as long as it is non-negative. If it is negative, then a sink state is reached. Finally, we define a new weight function w' by $w'((v, c), (v', m)) = c$ for every $c \in M \setminus \{\perp\}$ and $m \in M$ and $w'((v, \perp), (v', \perp)) = t + 1$.

Remark 3. Let $\rho = v_0 v_1 v_2 \dots$ and $\rho' = (v_0, m_0)(v_1, m_1)(v_2, m_2) \dots$ be such that ρ is a play in \mathcal{A} and ρ' is the corresponding extended play in $\mathcal{A} \times \mathcal{M}$.

1. If there is no $s \leq n$ such that $\text{EL}_{\text{cap}}(v_0 \dots v_s) < 0$, then $m_n = \text{EL}_{\text{cap}}(v_0 \dots v_n)$.
2. If there is an $s \leq n$ such that $\text{EL}_{\text{cap}}(v_0 \dots v_s) < 0$, then $m_n = \perp$.
3. If there is no s such that $\text{EL}_{\text{cap}}(v_0 \dots v_s) < 0$, then

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \text{EL}((v_0, m_0) \dots (v_{n-1}, m_{n-1})) = \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} \text{EL}_{\text{cap}}(v_0 \dots v_i).$$

4. If there is an s such that $\text{EL}_{\text{cap}}(v_0 \dots v_s) < 0$, then

$$\limsup_{n \rightarrow \infty} \frac{1}{n} \text{EL}((v_0, m_0) \dots (v_{n-1}, m_{n-1})) = t + 1.$$

5. $\rho \in \text{AvgRecharge}(w, \text{cap}, t)$ if, and only if, $\rho' \in \text{MeanPayoff}(w', t)$.

Thus, we have $(\mathcal{A}, \text{AvgRecharge}(w, \text{cap}, t)) \leq_{\mathcal{M}} (\mathcal{A} \times \mathcal{M}, \text{MeanPayoff}(w', t))$. Hence, positional determinacy of mean-payoff games [13], Lemma 1, and mean-payoff games being solvable in pseudo-polynomial time [28] yield the exponential time algorithm. \square

An application of Lemma 1 additionally yields an upper bound on the necessary memory states to implement a winning strategy.

Corollary 1. If Player 0 wins an average-bounded recharge game with capacity cap , then she also wins it with a finite-state strategy of size $\text{cap} + 2$.

Conversely, it is straightforward to show that this bound is tight: consider the average-bounded recharge game depicted in Figure 1 with some fixed even capacity cap and threshold $t = \frac{\text{cap}}{2}$. With cap memory states, Player 0 can implement a strategy whose unique consistent play has the form $(v_0 v_1^{\text{cap}})^\omega$ which has the energy levels $(\text{cap}, \text{cap} - 1, \dots, 1, 0)^\omega$, which results in a long-run average of t . However, with $n < \text{cap}$ memory states, the best Player 0 is able to do is to implement a strategy whose unique consistent play has the form $(v_0 v_1^n)^\omega$ which has the energy levels $(\text{cap}, \text{cap} - 1, \dots, \text{cap} - n)^\omega$, which results in a long-run average of $(\text{cap} - n) + \frac{n}{2} = \text{cap} - \frac{n}{2} > \text{cap} - \frac{\text{cap}}{2} = t$. Every other play that is implementable with n memory states has an even higher average. Thus, Player 0 needs cap memory states to meet the bound on the average.

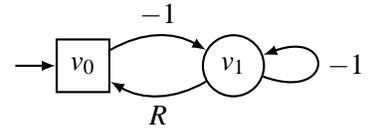


Figure 1: The arena for the lower bound on memory requirements in average-bounded recharge games.

Next, we give an EXPTIME lower bound by a reduction from countdown games. The arena $\mathcal{A} = (V, V_0, V_1, E, v_\perp)$ and the weight function w of such a game are subject to some restrictions:

1. The initial vertex is in V_0 and there is a designated sink vertex $v_\perp \in V_1$ with a self loop,
2. every vertex in V_0 has an edge to v_\perp and all other edges are in $V_0 \times (V_1 \setminus \{v_\perp\}) \cup (V_1 \setminus \{v_\perp\}) \times V_0$,
3. all edges in $V_0 \times (V_1 \setminus \{v_\perp\})$ have negative weight and there are no two outgoing transitions from a vertex in V_0 with the same weight, and

4. all other edges have weight zero.

The objective is given as

$$\text{Countdown}(w, c) = \{v_0 v_1 v_2 \dots \in V^\omega \mid \exists n. v_n = v_\perp \text{ and } c + EL(v_0 \dots v_n) = 0\}.$$

Intuitively, Player 0 picks negative weights that are subtracted from the initial energy c and Player 1 picks the next vertex to continue at (vertices of the countdown game are in V_0 , V_1 only contains auxiliary vertices). Player 0 wins if the energy level is exactly zero at some point, at which she has to move to the sink vertex. Otherwise, Player 1 wins. Solving countdown games is EXPTIME-complete [20]. Our reduction is a straightforward adaption of the reduction from countdown to average-energy games [5].

Lemma 4. *Solving average-bounded recharge games is EXPTIME-hard.*

Proof. Fix $\mathcal{A} = (V, V_0, V_1, E, v_I)$ and w satisfying the requirements of a countdown game and some initial energy c . We add a fresh vertex v'_I to V_1 , add an edge from v'_I to v_I and label it with the recharge action R to obtain the arena \mathcal{A}' and the weight function w' . As every play that does not reach the sink vertex traverses infinitely many edges with negative weight, we have $\rho \in \text{Countdown}(w, c)$ if, and only if, $v'_I \cdot \rho \in \text{AvgRecharge}(w', c, 0)$. Thus, Player 0 wins $(\mathcal{A}', \text{AvgRecharge}(w', c, 0))$ if, and only if, she wins $(\mathcal{A}, \text{Countdown}(w, c))$. Hence, solving average-bounded recharge games is EXPTIME-hard. \square

Note that the hardness depends on the requirement to bound the average. Recharge games without average-bound are solvable in pseudo-polynomial time, as such a game can be expressed as a one-dimensional consumption game [6]. Determining the minimal cover (the analogue of our capacity in consumption games, see [6] for a formal definition) for the initial vertex and comparing it to the given capacity yields the desired result, as the minimal cover in a one-dimensional consumption game can be computed in pseudo-polynomial-time [6]. Whether recharge games can be solved in polynomial time is open. In the next subsection, we present a variant that is solvable in polynomial time.

Also, the previous hardness proof can be adapted to recharge games with a given threshold and existentially quantified capacity (Problem 4). To this end, we add the initial gadget presented in Figure 2 to a countdown game \mathcal{G} . In order to win this game, Player 0 has to reach the Player 1 vertex with energy level c . If the energy level is larger then Player 1 can take the edge with weight $-c$ and reach the sink with a positive energy level. Hence, the average accumulated energy will be non-zero, too. Conversely, if the energy level is smaller than c , then taking the same edge yields a negative energy level. Hence, in both cases the objective $\text{AvgRecharge}(w, \text{cap}, 0)$ is violated, independently of the value of c . However, if Player 0 reaches the Player 1 vertex with energy level c , then she wins from there, if and only if, she has a winning strategy for the countdown game \mathcal{G} with initial value c . Thus, she wins the recharge game with objective $\text{AvgRecharge}(w, \text{cap}, 0)$ for some cap if, and only if, she wins the countdown game \mathcal{G} with objective $\text{Countdown}(w, c)$.

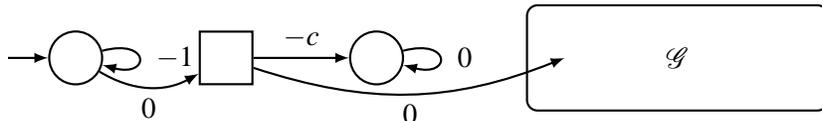


Figure 2: The gadget for showing Problem 4 EXPTIME-hard.

Theorem 3. *Solving average-bounded recharge games with existentially quantified capacity and a given threshold is EXPTIME-hard.*

However, it is an open problem whether these games can be solved in exponential time. The reduction to mean-payoff games presented above depends on the capacity being part of the input. This is related to the absence of good upper bounds on the necessary capacity to achieve a given threshold.

4.2 Finding a Sufficient Capacity in Recharge Games

To tackle the high complexity of solving average-bounded recharge games, we consider the problem where the recharge capacity cap and the threshold t are existentially quantified. As the energy level is always bounded from above by cap , which implies that the average accumulated energy is also bounded by cap , it suffices to consider the objective $\text{Recharge}(w, cap)$, analogous results hold for the objective $\text{AvgRecharge}(w, cap, t)$. We show that the following problem can be solved in polynomial time.

Problem 5. *Existence of a sufficient recharge level in recharge games*

Input: Arena $\mathcal{A} = (V, V_0, V_1, E, v_I)$ and $w: E \rightarrow -\mathbb{N} \cup \{\mathbf{R}\}$

Question: *Exists a capacity cap s.t. Player 0 wins $(\mathcal{A}, \text{Recharge}(w, cap))$?*

One attempt to prove this result is to again encode the game as a one-dimensional consumption game as described above. However, this only yields a pseudo-polynomial time algorithm. In the following, we present a truly polynomial time algorithm by a reduction to three-color parity games. Given a coloring $\Omega: V \rightarrow \mathbb{N}$, $\text{Parity}(\Omega)$ denotes the (max)-parity objective, which contains all plays $v_0 v_1 v_2 \dots \in V^\omega$ such that the maximal color appearing infinitely often in $\Omega(v_0)\Omega(v_1)\Omega(v_2)\dots$ is even.

Theorem 4. *The existence of a sufficient recharge level in a recharge game can be determined in polynomial time.*

Proof. Fix an arena $\mathcal{A} = (V, V_0, V_1, E, v_I)$ and $w: E \rightarrow -\mathbb{N} \cup \{\mathbf{R}\}$. We construct a three-color parity game with the following property: Player 0 wins the parity game if, and only if, there is a cap such that Player 0 wins $(\mathcal{A}, \text{Recharge}(w, cap))$. We assume w.l.o.g. that every vertex of \mathcal{A} either only has incoming edges labeled with \mathbf{R} , only has incoming edges labeled with 0, or only has incoming edges labeled with a negative weight. This can always be achieved by tripling the set of vertices, one copy for each type of incoming edge. The new initial vertex is some fixed copy of the original initial vertex. This transformation does not change the winner and only results in a linear increase in the number of states.

Now, we can speak of recharge-vertices, zero-vertices, and of decrement-vertices and define the coloring Ω such that it assigns color 2 to the recharge-vertices, color 1 to the decrement-vertices, and color 0 to the zero-vertices. We claim that Player 0 has a winning strategy for the induced parity game if, and only if, there is a cap such that Player 0 wins $(\mathcal{A}, \text{Recharge}(w, cap))$.

First, assume Player 0 has a winning strategy for the parity game, which we can assume w.l.o.g. to be positional [15, 23]. Let W be the largest absolute weight in the image of w and define $cap = (|V| - 1) \cdot W$. We claim that σ is a winning strategy for Player 0 in $(\mathcal{A}, \text{Recharge}(w, cap))$. Assume it is not: then, there is a play prefix $v_0 \dots v_n$ that is consistent with σ such that $\text{EL}_{cap}(v_0 \dots v_n) < 0$. Let $v_i \dots v_n$ be the suffix since the last recharge edge was traversed, i.e., $-\text{EL}(v_i \dots v_n) > cap$. By the choice of cap , there are positions j and j' satisfying $i < j < j' \leq n$ such that $v_j = v_{j'}$ and $\text{EL}(v_j \dots v_{j'}) < 0$, i.e., there is a cycle with negative cost and without recharge edge. As σ is positional, the play $v_0 \dots v_{j-1}(v_j \dots v_{j-1})^\omega$ obtained by reaching and then repeating this cycle is consistent with σ as well. However, in the parity game, this cycle visits no recharge-vertex, but at least one decrement-vertex. Hence, it is losing for Player 0, which contradicts σ being a winning strategy. Hence, σ is indeed also a winning strategy for $(\mathcal{A}, \text{Recharge}(w, cap))$.

Now, assume there is some cap and a strategy σ that is winning for Player 0 in $(\mathcal{A}, \text{Recharge}(w, cap))$. We claim that this strategy is also winning for her in the parity game. Assume, it is not, i.e., there is a

play that is consistent with σ , but losing for Player 0 in the parity game. By our choice of colors, this implies that this play visits only finitely many recharge-vertices, but infinitely many decrement-vertices. Thus, it has a prefix whose recharge energy level is negative. But this contradicts the fact that σ is a winning strategy for the recharge game.

To conclude, it remains to remark that three-color parity games can be solved in polynomial time. \square

By applying both directions of the equivalence, we obtain the following corollary.

Corollary 2. *If there is a cap such that Player 0 wins $(\mathcal{A}, \text{Recharge}(w, \text{cap}))$, then she also wins $(\mathcal{A}, \text{Recharge}(3 \cdot (n - 1) \cdot W, w))$, where n is the number of vertices of \mathcal{A} and W is the largest absolute weight in the domain of w . Player 0 wins the latter game with a finite-state strategy of size three.*

Note that this can be improved slightly by a finer analysis: the factor $(n - 1)$ can be replaced by the number of decrement-vertices. Conversely, it is straightforward to construct examples that prove these bounds to be tight, e.g., a cycle of n edges, one being a recharge edge and all others having weight $-W$.

5 Tradeoffs in Recharge Games

In this section, we illustrate two different tradeoff scenarios between different quality measures for winning strategies that occur in average-bounded recharge games, i.e., tradeoffs between capacity and long-run average and between memory size and long-run average. Note that increasing the recharge capacity in such a game has a (possibly negative) influence on the long-run average, as every recharge returns the energy level to the capacity. All games we consider here are solitary games for Player 0, i.e., every vertex belongs to Player 0. Thus, a strategy can be identified with the unique play consistent with it.

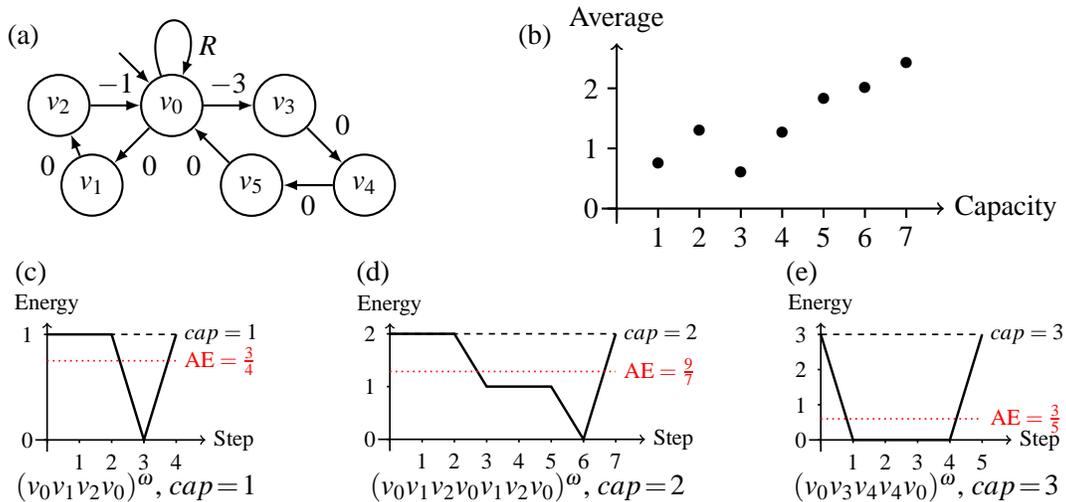


Figure 3: (a) An average-bounded recharge game with tradeoff between capacity and long-run average. (b) A plot of the tradeoff. (c) - (e) Energy progressions of different plays in the average-bounded recharge game for different capacities.

First, we study the tradeoff between the capacity and the long-run average energy level. Consider the game in Figure 3(a): Player 0 wins the game for $\text{cap} = 1$ and $t = 1$ by realizing the long-run average $\frac{3}{4}$ with the play $(v_0v_1v_2v_0)^\omega$ (Figure 3(c)). But, by increasing the capacity to $\text{cap} = 2$, it is no longer possible for her to win for $t = 1$, as the best long-run average she can realize is $\frac{9}{7}$ by playing $(v_0v_1v_2v_0v_1v_2v_0)^\omega$

(Figure 3(d)). However, for $cap = 3$, she can again win for $t = 1$, and it is possible to realize the long-run average $\frac{3}{5}$ by playing $(v_0v_3v_4v_5v_0)^\omega$ (Figure 3(e)). Again, with $cap = 4$ Player 0 loses for $t = 1$.

This example shows that higher capacity can be traded for a lower long-run average and that the tradeoff is non-monotonic. Figure 3(b) shows a plot of the tradeoff for capacities ranging from 1 to 7.

Another tradeoff scenario is between the number of memory states required to implement a strategy and the long-run average energy level it realizes. Consider the recharge game from Figure 1: as discussed below Corollary 1, Player 0 can win for the threshold $t = \frac{cap}{2}$ with cap memory states. However, with $n < cap$ memory states, she can only guarantee the long-run average $(cap - n) + \frac{n}{2}$. In particular, the best long-run average that is realizable by a positional strategy (which requires one memory state to implement) is $cap - \frac{1}{2}$ (see Fig. 4).

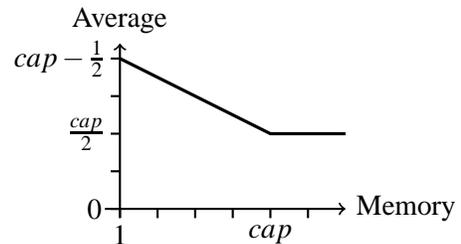


Figure 4: A plot of the tradeoff between memory size and long-run average in the game in Figure 1.

6 Conclusion

We continued the study of average-energy games by considering problems where the bound on the average is existentially quantified instead of given as part of the input. We showed that solving this problem is equivalent to determining whether the maximal energy level can be uniformly bounded by a strategy. The latter problem is known to be decidable in doubly-exponential time, which therefore also holds for our original problem. Then, we considered a different type of energy evolution where energy is only consumed or reset to some fixed capacity. Solving the average-bounded variants of these games is shown to be complete for exponential time. Due to this high complexity, we again considered a variant where the bounds are existentially quantified. This problem turns out to be solvable in polynomial time. Finally, we studied tradeoffs between the different bounds and the memory requirements of winning strategies: increasing the upper bound on the maximal energy level is shown to allow to improve the average energy level and memory can be traded for smaller upper bounds and vice versa.

For future work, it would be interesting to extend our results to a multi-dimensional setting. Also, the exact complexity of determining the existence of an upper bound in average-energy games is open. Finally, the decidability of average-energy games with a given threshold, but without an upper bound on the energy level is open [5]. In current work, we study whether our approach presented in Section 3 can be adapted to solve these problems, e.g., by not picking representatives by minimizing peak height but some other measure. These questions are also related to the complexity of recharge games with a given threshold where the capacity is existentially quantified. Finally, we are studying upper bounds on the tradeoffs presented in Section 5.

Acknowledgements. The work presented here was carried out while the third author visited the Distributed and Embedded Systems Unit at Aalborg University and while the second author visited the Reactive Systems Group at Saarland University. We thank these institutions for their hospitality.

References

- [1] Rajeev Alur, Kousha Etessami, Salvatore La Torre & Doron Peled (2001): *Parametric temporal logic for "model measuring"*. *ACM Trans. Comput. Log.* 2(3), pp. 388–407, doi:10.1145/377978.377990.

- [2] R. Bloem, K. Chatterjee, T.A. Henzinger & B. Jobstmann (2009): *Better Quality in Synthesis through Quantitative Objectives*. In Ahmed Bouajjani & Oded Maler, editors: *CAV 2009, LNCS 5643*, Springer, pp. 140–156, doi:10.1007/978-3-642-02658-4_14.
- [3] Endre Boros, Khaled M. Elbassioni, Vladimir Gurvich & Kazuhisa Makino (2015): *Markov Decision Processes and Stochastic Games with Total Effective Payoff*. In Ernst W. Mayr & Nicolas Ollinger, editors: *STACS 2015, LIPIcs 30*, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 103–115, doi:10.4230/LIPIcs.STACS.2015.103.
- [4] Patricia Bouyer, Uli Fahrenberg, Kim G. Larsen, Nicolas Markey & Jiří Srba (2008): *Infinite runs in weighted timed automata with energy constraints*. In Franck Cassez & Claude Jard, editors: *Formats 2008, LNCS 5215*, Springer, pp. 33–47, doi:10.1007/978-3-540-85778-5_4.
- [5] Patricia Bouyer, Nicolas Markey, Mickael Randour, Kim G. Larsen & Simon Laursen (2015): *Average-energy games*. In Javier Esparza & Enrico Tronci, editors: *GandALF 2015, EPTCS 193*, Open Publishing Association, pp. 1–15, doi:10.4204/EPTCS.193.1.
- [6] Tomáš Brázdil, Krishnendu Chatterjee, Antonín Kucera & Petr Novotný (2012): *Efficient Controller Synthesis for Consumption Games with Multiple Resource Types*. In P. Madhusudan & Sanjit A. Seshia, editors: *CAV 2012, LNCS 7358*, Springer, pp. 23–38, doi:10.1007/978-3-642-31424-7_8.
- [7] Luboš Brim, Jakub Chaloupka, Laurent Doyen, Raffaella Gentilini & Jean-François Raskin (2011): *Faster algorithms for mean-payoff games*. *Formal Methods in System Design* 38(2), pp. 97–118, doi:10.1007/s10703-010-0105-x.
- [8] Franck Cassez, Jan J. Jensen, Kim G. Larsen, Jean-François Raskin & Pierre-Alain Reynier (2009): *Automatic Synthesis of Robust and Optimal Controllers – An Industrial Case Study*. In Rupak Majumdar & Paulo Tabuada, editors: *HSCC 2009, LNCS 5469*, Springer, pp. 90–104, doi:10.1007/978-3-642-00602-9_7.
- [9] Arindam Chakrabarti, Luca de Alfaro, Thomas A. Henzinger & Mariëlle Stoelinga (2003): *Resource Interfaces*. In Rajeev Alur & Insup Lee, editors: *EMSOFT 2003, LNCS 2855*, Springer, pp. 117–133, doi:10.1007/978-3-540-45212-6_9.
- [10] Krishnendu Chatterjee & Laurent Doyen (2012): *Energy parity games*. *Theor. Comput. Sci.* 458, pp. 49–60, doi:10.1016/j.tcs.2012.07.038.
- [11] Krishnendu Chatterjee, Thomas A. Henzinger & Marcin Jurdziński (2005): *Mean-Payoff Parity Games*. In: *LICS 2005*, IEEE Computer Society, pp. 178–187, doi:10.1109/LICS.2005.26.
- [12] Krishnendu Chatterjee, Mickael Randour & Jean-François Raskin (2013): *Strategy synthesis for multi-dimensional quantitative objectives*. *Acta Informatica* 51(3), pp. 129–163, doi:10.1007/s00236-013-0182-6.
- [13] A. Ehrenfeucht & J. Mycielski (1979): *Positional strategies for mean payoff games*. *International Journal of Game Theory* 8, pp. 109–113, doi:10.1007/BF01768705.
- [14] Daniel Ejsing-Dunn & Lisa Fontani (2013): *Infinite Runs in Recharge Automata*. Master’s thesis, Computer Science Department, Aalborg University, Denmark. Available at <http://www.cassting-project.eu/wp-content/uploads/master13-EF.pdf>.
- [15] E. Allen Emerson & Charanjit S. Jutla (1991): *Tree Automata, Mu-Calculus and Determinacy (Extended Abstract)*. In: *FOCS 1991*, IEEE, pp. 368–377, doi:10.1109/SFCS.1991.185392.
- [16] Uli Fahrenberg, Line Juhl, Kim G. Larsen & Jiří Srba (2011): *Energy Games in Multiweighted Automata*. In Antonio Cerone & Pekka Pihlajasaari, editors: *ICTAC 2011, LNCS 6916*, Springer, pp. 95–115, doi:10.1007/978-3-642-23283-1_9.
- [17] Peter Faymonville & Martin Zimmermann (2014): *Parametric Linear Dynamic Logic*. In Adriano Peron & Carla Piazza, editors: *GandALF 2014, EPTCS 161*, pp. 60–73, doi:10.4204/EPTCS.161.8.
- [18] Erich Grädel, Wolfgang Thomas & Thomas Wilke, editors (2002): *Automata, Logics, and Infinite Games: A Guide to Current Research*. *LNCS 2500*, Springer, doi:10.1007/3-540-36387-4.

- [19] Line Juhl, Kim G. Larsen & Jean-François Raskin (2013): *Optimal Bounds for Multiweighted and Parametrised Energy Games*. In Zhiming Liu, Jim Woodcock & Huibiao Zhu, editors: *Theories of Programming and Formal Methods - Essays Dedicated to Jifeng He on the Occasion of His 70th Birthday*, LNCS 8051, Springer, pp. 244–255, doi:10.1007/978-3-642-39698-4_15.
- [20] Marcin Jurdziński, Jeremy Sproston & François Laroussinie (2008): *Model Checking Probabilistic Timed Automata with One or Two Clocks*. *LMCS* 4(3), doi:10.2168/LMCS-4(3:12)2008.
- [21] Dénes König (1927): *Über eine Schlussweise aus dem Endlichen ins Unendliche*. *Acta Litt. ac. sci. Szeged* 3, pp. 121–130.
- [22] Orna Kupferman, Nir Piterman & Moshe Y. Vardi (2009): *From liveness to promptness*. *Formal Methods in System Design* 34(2), pp. 83–103, doi:10.1007/s10703-009-0067-z.
- [23] Andrzej Mostowski (1991): *Games with Forbidden Positions*. Technical Report 78, University of Gdańsk.
- [24] Mickael Randour (2013): *Automated Synthesis of Reliable and Efficient Systems Through Game Theory: A Case Study*. In Thomas Gilbert, Markus Kirkilionis & Gregoire Nicolis, editors: *ECCS 2012, Springer Proceedings in Complexity*, Springer, pp. 731–738, doi:10.1007/978-3-319-00395-5_90.
- [25] F. Thuijsman & O.J. Vrieze (1987): *The bad match; A total reward stochastic game*. *OR Spektrum* 9(2), pp. 93–99, doi:10.1007/BF01732644.
- [26] Yaron Velner, Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger, Alexander Moshe Rabinovich & Jean-François Raskin (2015): *The complexity of multi-mean-payoff and multi-energy games*. *Inf. Comput.* 241, pp. 177–196, doi:10.1016/j.ic.2015.03.001.
- [27] Martin Zimmermann (2015): *Parameterized Linear Temporal Logics Meet Costs: Still not Costlier than LTL*. In Javier Esparza & Enrico Tronci, editors: *GandALF 2015, EPTCS* 193, Open Publishing Association, pp. 144–157, doi:10.4204/EPTCS.193.11.
- [28] Uri Zwick & Mike Paterson (1996): *The complexity of mean payoff games on graphs*. *TCS* 158(1-2), pp. 343–359, doi:10.1016/0304-3975(95)00188-3.