

Average-energy games*

Patricia Bouyer Nicolas Markey Mickael Randour

LSV – CNRS & ENS Cachan – France

Kim G. Larsen Simon Laursen

Aalborg University – Denmark

Two-player quantitative zero-sum games provide a natural framework to synthesize controllers with performance guarantees for reactive systems within an uncontrollable environment. Classical settings include mean-payoff games, where the objective is to optimize the long-run average gain per action, and energy games, where the system has to avoid running out of energy.

We study *average-energy* games, where the goal is to optimize the long-run average of the accumulated energy. We show that this objective arises naturally in several applications, and that it yields interesting connections with previous concepts in the literature. We prove that deciding the winner in such games is in $\text{NP} \cap \text{coNP}$ and at least as hard as solving mean-payoff games, and we establish that memoryless strategies suffice to win. We also consider the case where the system has to minimize the average-energy *while* maintaining the accumulated energy within predefined bounds at all times: this corresponds to operating with a finite-capacity storage for energy. We give results for one-player and two-player games, and establish complexity bounds and memory requirements.

1 Introduction

Quantitative games. Game-theoretic formulations are a standard tool for the synthesis of provably-correct controllers for reactive systems [22]. We consider two-player (system vs. environment) turn-based games played on finite graphs. Vertices of the graph are called *states* and partitioned into states of player 1 and states of player 2. The game is played by moving a pebble from state to state, along *edges* in the graph, and starting from a given initial state. Whenever the pebble is on a state belonging to player i , player i decides where to move the pebble next, according to his *strategy*. The infinite path followed by the pebble is called a *play*: it represents one possible behavior of the system. A *winning objective* encodes acceptable behaviors of the system and can be seen as a set of winning plays. The goal of player 1 is to ensure that the outcome of the game will be a winning play, whatever the strategy played by his adversary.

To reason about resource constraints and the performance of strategies, *quantitative games* have been considered in the literature. See for example [10, 3, 29], or [30] for an overview. Those games are played on *weighted* graphs, where edges are fitted with integer weights modeling rewards or costs. The performance of a play is evaluated via a *payoff function* that maps it to the numerical domain. The objective of player 1 is then to ensure a sufficient payoff with regard to a given threshold value. Seminal classes of quantitative games include mean-payoff (*MP*), total-payoff (*TP*) and energy games (*EG*). In *MP* games [15, 33, 24], player 1 has to optimize his long-run average gain per edge taken whereas, in *TP* games [20, 19], player 1 has to optimize his long-run sum of weights. Energy games [10, 5, 23] model safety-like properties: the goal is to ensure that the running sum of weights never drops below zero and/or that it never exceeds a given upper bound $U \in \mathbb{N}$. All three classes share common properties. First, *MP* games, *TP* games, and *EG* games with only a lower bound (EG_L) are memoryless determined

*Work partially supported by European project CASSTING (FP7-ICT-601148) and ERC project EQualIS (StG-308087).

(given an initial state, either player 1 has a strategy to win, or player 2 has one, and in both cases no memory is required to win). Second, deciding the winner for those games is in $\text{NP} \cap \text{coNP}$ and no polynomial algorithm is known despite many efforts (e.g., [8, 12]). Energy games with both lower and upper bounds (EG_{LU}) are more complex: they are EXPTIME-complete and winning requires memory in general [5].

While those classes are well-known, it is sometimes necessary to go beyond them to accurately model practical applications. For example, multi-dimensional games and conjunctions with a parity objective model trade-offs between different quantitative aspects [11, 14, 32]. Similarly, window objectives address the need for strategies ensuring good quantitative behaviors within reasonable time frames [12].

Average-energy games. We study the *average-energy* (AE) payoff function: in AE games, the goal of player 1 is to optimize the *long-run average accumulated energy* over a play. We introduce this objective to formalize the specification desired in a practical application [9], which we detail in the following as a motivating example. Interestingly, it turns out that this payoff first appeared long ago [31], but it was not subject to a systematic study until very recently: see related work for more discussion.

In addition to being meaningful w.r.t. practical applications, AE games also have theoretical interest. In [13], Chatterjee and Prabhu define the *average debit-sum level* objective, which can be seen as a variation of the *average-energy* where the accumulated energy is taken to be zero in any point where it is actually positive (hence, it focuses on the average debt). They use the corresponding games to compute the values of quantitative timed simulation functions. In particular, they provide a pseudo-polynomial-time algorithm to solve those games, but the complexity of deciding the winner as well as the memory requirements are open. Here, we solve those questions for the very similar average-energy objective.

Motivating example. Our example is a simplified version of the industrial application studied by Cassez *et al.* [9]. Consider a machine that consumes oil, stored in a connected accumulator. We want to synthesize an appropriate controller to operate the oil pump that fills the accumulator, and by the effect of pressure, that releases oil from the accumulator into the machine with a (time-varying) rate according to desired production. In order to ensure safety, the oil level in the accumulator should be maintained at all times between a minimal and a maximal level. This part of the specification can be encoded as an energy objective with both lower and upper bounds (EG_{LU}). At the same time, the more oil (thus pressure) in the accumulator, the faster the whole apparatus wears out. Hence, an ideal controller should minimize the average level of oil in the long run. This desire can be formalized through the average-energy payoff (AE). Overall, the specification is thus to minimize the average-energy under the strong energy constraints: we denote the corresponding objective by AE_{LU} .

Contributions. Our main results are summarized in Table 1.

A) We establish that the average-energy objective can be seen as a *refinement* of total-payoff, in the same sense as total-payoff is seen as a refinement of mean-payoff [19]: it allows to distinguish strategies yielding identical mean-payoff and total-payoff.

B) We show that deciding the winner in two-player AE games is in $\text{NP} \cap \text{coNP}$ whereas it is in P for one-player games. In both cases, memoryless strategies suffice (Thm. 5). Those complexities match the state-of-the-art for MP and TP games [33, 24, 19, 8]. Furthermore we prove that AE games are at least as hard as mean-payoff games (Thm. 7). Therefore, the $\text{NP} \cap \text{coNP}$ -membership can be considered optimal w.r.t. our knowledge of MP games. Technically, the crux of our approach is as follows. First, we show that memoryless strategies suffice in one-player AE games (Thm. 3): this requires to prove important properties of the AE payoff as classical sufficient criteria for memoryless determinacy present in the literature fail to apply directly. Second, we establish a polynomial-time algorithm for the one-player case: it exploits the structure of winning strategies and mixes graph techniques with local linear program

Game objective	1-player	2-player	memory
MP	in P [26]	in $NP \cap coNP$ [33]	memoryless [15]
TP	in P [17]	in $NP \cap coNP$ [19]	memoryless [20]
EG_L	in P [5]	in $NP \cap coNP$ [10, 5]	memoryless [10]
EG_{LU}	PSPACE-complete [16]	EXPTIME-complete [5]	pseudo-polynomial
AE	in P	in $NP \cap coNP$	memoryless
AE_{LU} , polynomial U	in P	in $NP \cap coNP$	polynomial
AE_{LU} , arbitrary U	in EXPTIME / PSPACE-hard	EXPTIME-complete	pseudo-polynomial
AE_L	EXPTIME-easy / NP-hard	<i>open</i> / EXPTIME-hard	<i>open</i> (\geq pseudo-p.)

Table 1: Complexity of deciding the winner and memory requirements for quantitative games: MP stands for mean-payoff, TP for total-payoff, EG_L (resp. EG_{LU}) for lower-bounded (resp. lower- and upper-bounded) energy, AE for average-energy, and AE_L (resp. AE_{LU}) for average-energy under a lower bound (resp. and upper bound $U \in \mathbb{N}$) on the energy. Results without reference are proved in this paper.

solving (Thm. 4). Finally, we lift memoryless determinacy to the two-player case using results by Gimbert and Zielonka [21] and obtain the $NP \cap coNP$ -membership as a corollary (Thm. 6).

C) We establish an EXPTIME algorithm to solve two-player AE games with lower- and upper-bounded energy (AE_{LU}) with an arbitrary upper bound $U \in \mathbb{N}$ (Thm. 8). It relies on a reduction of the AE_{LU} game to a pseudo-polynomially larger AE game where the energy constraints are encoded in the graph structure. Applying straightforwardly the AE algorithm on this game would only give us $NEXPTIME \cap coNEXPTIME$ -membership, hence we avoid this blowup by further reducing the problem to a particular MP game and applying a pseudo-polynomial algorithm, with some care to ensure that overall the algorithm only requires pseudo-polynomial time in the original AE_{LU} game. Since the simpler EG_{LU} games (i.e., AE_{LU} with a trivial AE constraint) are already EXPTIME-hard [5], the EXPTIME-membership result is optimal. We also prove that pseudo-polynomial memory is both sufficient and in general necessary to win in AE_{LU} games, for both players (Thm. 9). Whether one-player AE_{LU} games belong to PSPACE is an open question. For polynomial (in the size of the game graph) values of the upper bound U —or if it is given in unary—the complexity of the two-player AE_{LU} problem collapses to $NP \cap coNP$ with the same approach, and polynomial memory suffices for both players.

D) We provide partial answers for the AE_L objective— AE under a lower bound constraint on energy but no upper bound. We provide an EXPTIME algorithm for the one-player case, by reducing the problem to an AE_{LU} game with a sufficiently large upper bound. That is, we prove that if the player can win for the AE_L objective, then he can do so without ever increasing its energy above a well-chosen bound. We also prove the AE_L problem to be at least NP-hard in one-player games and EXPTIME-hard in two-player games (Lem. 12) via reductions from the subset-sum problem and countdown games respectively. Finally, we show that memory is required for both players in two-player AE_L games (Lem. 13), and that pseudo-polynomial memory is both sufficient and necessary in the one-player case (Thm. 10). The decidability status of two-player AE_L games remains open as we only provide a correct but incomplete incremental algorithm (Lem. 11). We conjecture that the two-player AE_L problem is decidable and sketch a potential approach to solve it. We highlight the key remaining questions and discuss some connections with related models that are known to be difficult.

Observe that in many applications, the energy must be stocked in a finite-capacity storage for which an upper bound is provided. Hence, the model of choice in this case is AE_{LU} .

Related work. The *average-energy* payoff—Eq. (1)—appeared in a paper by Thuijsman and Vrieze in the late eighties [31], under the name *total-reward*. This definition is different from the classical *total-payoff*—see Sect. 2—commonly studied in the formal methods community (see for example [20, 19]), which, despite that, has been referred in many papers as either total-payoff or total-reward equivalently. We will see in this paper that both definitions are *indeed* different and exhibit different behaviors.

Maybe due to this confusion, the payoff of Eq. (1)—which we call *average-energy* thus avoiding misunderstandings—was not studied extensively until recently. Nothing was known about memoryless determinacy and complexity of deciding the winner. Independently to our work, Boros et al. recently studied the same payoff (under the name *total-payoff*). In [4], they study Markov decision processes and stochastic games with the payoff of Eq. (1) and solve both questions. Their results overlap with ours for *AE* games (Table 1). Let us first mention that our results were obtained independently. Second, and *most importantly*, our approach and *techniques are different*, and we believe our take on the problem yields some interest for our community. Indeed, the algorithm of Boros *et al.* entirely relies on linear programming in the one-player case, and resorts to approximation by discounted games in the two-player one. Our techniques are arguably more constructive and based on inherent properties of the payoff. In that sense, it is closer to what is usually deemed important in our field. For example, we provide an extensive comparison with classical payoffs. We base our proof of memoryless determinacy on *operational understanding* of the *AE* which is crucial in order to formalize proper specifications. Our technique then benefits from seminal works [21] to bypass the reduction to discounted games and obtain a direct proof, thanks to our more constructive approach. Lastly, while [4] considers the *AE* problem in the stochastic context, we focus on the deterministic one but consider multi-criteria extensions by adding bounds on the energy (*AE_{LU}* and *AE_L* games). Those extensions are *completely new*, exhibit theoretical interest and are adequate for practical applications in constrained energy systems, as witnessed by the case study of [9].

Recent work of Brázdil *et al.* [7] considers the optimization of a payoff under energy constraint. They study mean-payoff in consumption systems, i.e., simplified one-player energy games where all edges consume energy but some states can atomically produce a reload of the energy up to the allowed capacity.

Full details and proofs of the results presented here can be found in the extended paper [6].

2 Preliminaries

Graph games. We consider turn-based games played on graphs between two players denoted by \mathcal{P}_1 and \mathcal{P}_2 . A *game* is a tuple $G = (S_1, S_2, E, w)$ where (i) S_1 and S_2 are disjoint finite sets of *states* belonging to \mathcal{P}_1 and \mathcal{P}_2 , with $S = S_1 \uplus S_2$, (ii) $E \subseteq S \times S$ is a finite set of *edges*, and (iii) $w: E \rightarrow \mathbb{Z}$ is an integer *weight function*. Given edge $(s_1, s_2) \in E$, we write $w(s_1, s_2)$ as a shortcut for $w((s_1, s_2))$. We denote by W the largest absolute weight assigned by function w . A game is called 1-player if $S_1 = \emptyset$ or $S_2 = \emptyset$.

A *play* from an initial state $s_{\text{init}} \in S$ is an infinite sequence $\pi = s_0 s_1 \dots s_n \dots$ such that $s_0 = s_{\text{init}}$ and for all $i \geq 0$ we have $(s_i, s_{i+1}) \in E$. The (finite) *prefix* of π up to position n gives the sequence $\pi(n) = s_0 s_1 \dots s_n$, the last element s_n is denoted $\text{last}(\pi(n))$. The set of all plays in G is denoted by $\text{Plays}(G)$ and the set of all prefixes is denoted by $\text{Pref}(G)$. We say that a prefix $\rho \in \text{Pref}(G)$ belongs to \mathcal{P}_i , $i \in \{1, 2\}$, if $\text{last}(\rho) \in S_i$. The set of prefixes that belong to \mathcal{P}_i is denoted by $\text{Pref}_i(G)$. The classical concatenation between prefixes (resp. prefix and play) is denoted by the \cdot operator. The length of a non-empty prefix $\rho = s_0 \dots s_n$ is defined as the number of edges and denoted by $|\rho| = n$.

Payoffs of plays. Given a play $\pi = s_0 s_1 \dots s_n \dots$ we define

- its *energy level* at position n as $EL(\pi(n)) = \sum_{i=0}^{n-1} w(s_i, s_{i+1})$;
- its *mean-payoff* as $\overline{MP}(\pi) = \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} w(s_i, s_{i+1}) = \limsup_{n \rightarrow \infty} \frac{1}{n} EL(\pi(n))$;

- its *total-payoff* as $\overline{TP}(\pi) = \limsup_{n \rightarrow \infty} \sum_{i=0}^{n-1} w(s_i, s_{i+1}) = \limsup_{n \rightarrow \infty} EL(\pi(n))$;
- and its *average-energy* as

$$\overline{AE}(\pi) = \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n \left(\sum_{j=0}^{i-1} w(s_j, s_{j+1}) \right) = \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n EL(\pi(i)). \quad (1)$$

We will sometimes consider those measures defined with \liminf instead of \limsup , in which case we write \underline{MP} , \underline{TP} and \underline{AE} respectively. Finally, we also consider those measures over prefixes: we naturally define them by dropping the $\limsup_{n \rightarrow \infty}$ operator and taking $n = |\rho|$ for a prefix $\rho \in \text{Pref}_i(G)$. In this case, we simply write $MP(\rho)$, $TP(\rho)$ and $AE(\rho)$ to denote the fact that we consider *finite* sequences.

Strategies. A strategy for \mathcal{P}_i , $i \in \{1, 2\}$, is a function $\sigma_i: \text{Pref}_i(G) \rightarrow S$ such that for all $\rho \in \text{Pref}_i(G)$ we have $(\text{last}(\rho), \sigma_i(\rho)) \in E$. A strategy σ_i for \mathcal{P}_i is *finite-memory* if it can be encoded by a deterministic finite-state Moore machine. A strategy is *memoryless* if it does not depend on the history but only on the current state of the game. We denote by $\Sigma_i(G)$, the sets of strategies for player \mathcal{P}_i . We drop G when the context is clear.

A play $\pi = s_0 s_1 \dots$ is *consistent* with a strategy σ_i of \mathcal{P}_i if, for all $n \geq 0$ where $\text{last}(\pi(n)) \in S_i$, we have $\sigma_i(\pi(n)) = s_{n+1}$. Given an initial state $s_{\text{init}} \in S$ and strategies σ_1 and σ_2 for the two players, we denote by $\text{Outcome}(s_{\text{init}}, \sigma_1, \sigma_2)$ the unique play that starts in s_{init} and is consistent with both σ_1 and σ_2 . When fixing the strategy of only \mathcal{P}_i , we denote the set of consistent outcomes by $\text{Outcomes}(s_{\text{init}}, \sigma_i)$.

Objectives. An *objective* in G is a set $\mathcal{W} \subseteq \text{Plays}(G)$ that is declared winning for \mathcal{P}_1 . Given a game G , an initial state s_{init} , and an objective \mathcal{W} , a strategy $\sigma_1 \in \Sigma_1$ is winning for \mathcal{P}_1 if for all strategy $\sigma_2 \in \Sigma_2$, we have that $\text{Outcome}(s_{\text{init}}, \sigma_1, \sigma_2) \in \mathcal{W}$. Symmetrically, a strategy $\sigma_2 \in \Sigma_2$ is winning for \mathcal{P}_2 if for all strategy $\sigma_1 \in \Sigma_1$, we have that $\text{Outcome}(s_{\text{init}}, \sigma_1, \sigma_2) \notin \mathcal{W}$. That is, we consider *zero-sum* games.

We consider the following objectives and combinations of those objectives.

- Given an initial energy level $c_{\text{init}} \in \mathbb{N}$, the **lower-bounded energy** (EG_L) objective $\text{Energy}_L(c_{\text{init}}) = \{\pi \in \text{Plays}(G) \mid \forall n \geq 0, c_{\text{init}} + EL(\pi(n)) \geq 0\}$ requires non-negative energy at all times.
- Given an upper bound $U \in \mathbb{N}$ and an initial energy level $c_{\text{init}} \in \mathbb{N}$, the **lower- and upper-bounded energy** (EG_{LU}) objective $\text{Energy}_{LU}(U, c_{\text{init}}) = \{\pi \in \text{Plays}(G) \mid \forall n \geq 0, c_{\text{init}} + EL(\pi(n)) \in [0, U]\}$ requires that the energy always remains non-negative and below the upper bound U along a play.
- Given a threshold $t \in \mathbb{Q}$, the **mean-payoff** (MP) objective $\text{MeanPayoff}(t) = \{\pi \in \text{Plays}(G) \mid \overline{MP}(\pi) \leq t\}$ requires that the mean-payoff is at most t .
- Given a threshold $t \in \mathbb{Z}$, the **total-payoff** (TP) objective $\text{TotalPayoff}(t) = \{\pi \in \text{Plays}(G) \mid \overline{TP}(\pi) \leq t\}$ requires that the total-payoff is at most t .
- Given a threshold $t \in \mathbb{Q}$, the **average-energy** (AE) objective $\text{AvgEnergy}(t) = \{\pi \in \text{Plays}(G) \mid \overline{AE}(\pi) \leq t\}$ requires that the average-energy is at most t .

For the MP , TP and AE objectives, note that \mathcal{P}_1 aims to *minimize* the payoff value while \mathcal{P}_2 tries to maximize it. The reversed convention is also often used in the literature but both are equivalent. For our motivating example, seeing \mathcal{P}_1 as a minimizer is more natural. Note that we define the objectives using the \limsup variants of MP , TP and AE , but similar results are obtained for the \liminf variants.

Decision problem. Given a game G , an initial state $s_{\text{init}} \in S$, and an objective $\mathcal{W} \subseteq \text{Plays}(G)$ as defined above, the associated *decision problem* is to decide if \mathcal{P}_1 has a winning strategy for this objective.

We recall classical results in Table 1. Memoryless strategies suffice for both players for EG_L [10, 5], MP [15] and TP [17, 20] objectives. Since all associated problems can be solved in polynomial time

for 1-player games, it follows that the 2-player decision problem is in $\text{NP} \cap \text{coNP}$ for those three objectives [5, 33, 19]. For the EG_{LU} objective, memory is in general needed and the associated decision problem is EXPTIME-complete [5] (PSPACE-complete for one-player games [16]).

Game values. Given a game with an objective $\mathcal{W} \in \{\text{MeanPayoff}, \text{TotalPayoff}, \text{AvgEnergy}\}$ and an initial state s_{init} , we refer to the *value* from s_{init} as $v = \inf\{t \in \mathbb{Q} \mid \exists \sigma_1 \in \Sigma_1, \text{Outcomes}(s_{\text{init}}, \sigma_1) \subseteq \mathcal{W}(t)\}$. For both MP and TP objectives, it is known that the value can be achieved by an optimal memoryless strategy; for the AE objective it follows from our results (Thm. 5).

3 Average-Energy

In this section, we consider the problem of ensuring a *sufficiently low* average-energy.

Problem 1 (AE). *Given a game G , an initial state s_{init} , and a threshold $t \in \mathbb{Q}$, decide if \mathcal{P}_1 has a winning strategy $\sigma_1 \in \Sigma_1$ for the objective $\text{AvgEnergy}(t)$.*

3.1 Relation with classical objectives

Several links between EG_L , MP and TP objectives can be established. Intuitively, \mathcal{P}_1 can only ensure a lower bound on energy if he can prevent \mathcal{P}_2 from enforcing strictly-negative cycles (otherwise the initial energy is eventually exhausted). This is the case if and only if \mathcal{P}_1 can ensure a non-negative mean-payoff in G (here, he wants to maximize the MP), and if this is the case, \mathcal{P}_1 can prevent the running sum of weights from ever going too far beyond zero along a play, hence granting a lower bound on total-payoff.

The TP objective is sometimes seen as a *refinement* of MP for the case where \mathcal{P}_1 —as a minimizer—can ensure MP equal to zero but not lower, i.e., the MP game has value zero [19]. Indeed, one may use the TP to further discriminate between strategies that guarantee MP zero. In the same philosophy, the average-energy can help in distinguishing strategies that yield identical total-payoffs. See Fig. 1. The AE values in both examples can be computed easily using the upcoming technical lemmas (Sect. 3.2).

In these examples, the average-energy is clearly comprised between the infimum and supremum total-payoffs. This remains true for any play. In particular, *if the mean-payoff value from a state is not zero*, its total-payoff value is infinite and the following holds: either \mathcal{P}_1 can force AE equal to $-\infty$ or \mathcal{P}_2 can force AE equal to $+\infty$.

3.2 Useful properties of the average-energy

Classical sufficient criteria. Various sufficient criteria—or connected approaches—to deduce memoryless determinacy appear in the literature [15, 2, 1, 20, 27]. Unfortunately, they cannot be applied straight out of the box to the AE payoff. Intuitively, a common requirement is for winning objectives to be closed under *cyclic permutation* and under *concatenation*. Without further assumptions, the AE objective satisfies neither. Indeed, consider cycles represented by sequences of *weights* $\mathcal{C}_1 = \{-1\}$, $\mathcal{C}_2 = \{1\}$ and $\mathcal{C}_3 = \{1, -2\}$. We see that $AE(\mathcal{C}_1\mathcal{C}_2) = (-1 + 0)/2 = -1/2 < AE(\mathcal{C}_2\mathcal{C}_1) = (1 - 0)/2 = 1/2$, hence AE is not closed under cyclic permutations. Intuitively, the order in which the weights are seen *does* matter, in contrast to most classical payoffs. For concatenation, see that $AE(\mathcal{C}_3) = 0$ while $AE(\mathcal{C}_3\mathcal{C}_3) = -1/2 < 0$. Here the intuition is that the overall AE is impacted by the energy of the first cycle which is strictly negative (-1). In a sense, the AE of a cycle can only be maintained through repetition if this cycle is neutral with regard to the total energy level, i.e., if it has energy level zero: we will formalize this intuition in Lem. 2.




Fig. 1: Both plays have identical mean-payoff and total-payoff: $\overline{MP}(\pi_1) = \underline{MP}(\pi_1) = \overline{MP}(\pi_2) = \underline{MP}(\pi_2) = 0$, $\overline{TP}(\pi_1) = \overline{TP}(\pi_2) = 5$, and $\underline{TP}(\pi_1) = \underline{TP}(\pi_2) = 1$. But play π_1 has a lower average-energy: $\overline{AE}(\pi_1) = \underline{AE}(\pi_1) = 3 < \overline{AE}(\pi_2) = \underline{AE}(\pi_2) = 11/3$.

Extraction of prefixes. We establish two useful properties of the average-energy that help us to prove memoryless determinacy. The following lemma describes the impact of adding a finite prefix to an infinite play: it will help us in decomposing plays when needed.

Lemma 1. [Average-energy prefix] Let $\rho \in \text{Pref}_s(G)$, $\pi \in \text{Plays}(G)$. Then, $\overline{AE}(\rho \cdot \pi) = EL(\rho) + \overline{AE}(\pi)$. The same equality holds for \underline{AE} .

Extraction of a best cycle. The next lemma is crucial to prove that memoryless strategies suffice: under well-chosen conditions, one can always select a best cycle in a play—hence, there is no interest in mixing different cycles and no use for memory. It holds only for sequences of cycles *that have energy level zero*: since they do not change the energy, they do not modify the AE of the following suffix of play, and one can decompose the AE as a weighted average over zero cycles. The concatenation of cycles $\mathcal{C}_a = s s' \dots s$ and $\mathcal{C}_b = s s'' \dots s$ is to be understood as $\mathcal{C}_a \cdot \mathcal{C}_b = s s' \dots s s'' \dots s$.

Lemma 2. [Repeated zero cycles of bounded length] Let $\mathcal{C}_1, \mathcal{C}_2, \mathcal{C}_3, \dots$ be an infinite sequence of cycles $\mathcal{C}_i \in \text{Pref}_s(G)$ such that (i) $\pi = \mathcal{C}_1 \cdot \mathcal{C}_2 \cdot \mathcal{C}_3 \dots \in \text{Plays}(G)$, (ii) $\forall i \geq 1, EL(\mathcal{C}_i) = 0$ and (iii) $\exists \ell \in \mathbb{N}_{>0}$ such that $\forall i \geq 1, |\mathcal{C}_i| \leq \ell$. Then the following properties hold.

1. The average-energy of π is the weighted average of the average-energies of the cycles:

$$\overline{AE}(\pi) = \limsup_{k \rightarrow \infty} \left[\frac{\sum_{i=1}^k |\mathcal{C}_i| \cdot \overline{AE}(\mathcal{C}_i)}{\sum_{i=1}^k |\mathcal{C}_i|} \right]. \quad (2)$$

2. For any cycle $\mathcal{C} \in \text{Pref}_s(G)$ such that $EL(\mathcal{C}) = 0$, we have that $\overline{AE}(\mathcal{C}^\omega) = \overline{AE}(\mathcal{C})$.
3. Repeating the best cycle gives the lowest AE : $\inf_{i \in \mathbb{N}_{>0}} \overline{AE}(\mathcal{C}_i) = \inf_{i \in \mathbb{N}_{>0}} \overline{AE}((\mathcal{C}_i)^\omega) \leq \overline{AE}(\pi)$.

Similar properties hold for \underline{AE} .

3.3 One-player games

We assume that the unique player is \mathcal{P}_1 , hence that $S_2 = \emptyset$. The proofs are similar for the case where all states belong to \mathcal{P}_2 (i.e., $S_1 = \emptyset$). Similarly, we present our results for the \overline{AE} variant, but they carry over to the \underline{AE} one. Actually, since we show that we can restrict ourselves to *memoryless* strategies, all consistent outcomes will be periodic and thus both variants will be equal over those outcomes.




Fig. 2: The best cycle $\mathcal{C}_{s,2}$ is computed by looking for a path from $(s, 2)$ to $(s, 0)$ with sum zero in the first dimension (zero cycle) and minimal sum in the second dimension (minimal AE). Here, the cycle via s' is clearly better, with AE equal to $-1/2$ in contrast to $1/2$ via s'' .

Memoryless determinacy. Intuitively, we use Lem. 1 and Lem. 2 to transform any arbitrary path in a simple *lasso path*, repeating a unique simple cycle, and yielding an at least as good AE , thus proving that any threshold achievable with memory can also be achieved without it.

Theorem 3. *Memoryless strategies are sufficient to win one-player AE games.*

Polynomial-time algorithm. We know the form of optimal memoryless strategies: an optimal lasso path $\pi = \rho \cdot \mathcal{C}^\omega$ w.r.t. the AE . We establish a polynomial-time algorithm to solve one-player AE games.

The crux is computing, for each state s , the best—w.r.t. the AE —zero cycle \mathcal{C}_s starting and ending in s (if any). This is achieved through linear programming (LP) over expanded graphs. For each state s and length $k \in \{1, \dots, |S|\}$, we compute the best cycle $\mathcal{C}_{s,k}$ by considering a graph (Fig. 2) that models all cycles of length k from s and that uses $k + 1$ levels and two-dimensional weights on edges of the form $(c, l \cdot c)$ where c is the weight in the original game and $l \in \{k, k - 1, \dots, 1\}$ is the level of the edge. In the LP, we look for cycles $\mathcal{C}_{s,k}$ of length k on s such that (a) the sum of weights in the first dimension is zero (thus $\mathcal{C}_{s,k}$ is a zero cycle), and (b) the sum in the second one is minimal. Fortunately, this sum is exactly equal to $AE(\mathcal{C}) \cdot k$ thanks to the l factors used in the weights of the expanded graph. Hence, we obtain the optimal cycle $\mathcal{C}_{s,k}$ (in polynomial time). Doing this $|S|$ times for each state s , we obtain for each of them the optimal cycle \mathcal{C}_s (if one zero cycle exists). Then, by Lem. 1, it remains to compute the least EL with which each state s can be reached using classical graph techniques (e.g., Bellman-Ford), and to pick the optimal combination to obtain an optimal memoryless strategy, in polynomial time.

Theorem 4. *The AE problem for one-player games is in P.*

3.4 Two-player games

Memoryless determinacy. We now prove that memoryless strategies still suffice in two-player games. As discussed in Sect. 3.2, classical criteria do not apply. There is, however, one result that proves particularly useful. Consider any payoff function such that memoryless strategies suffice for *both one-player* versions ($S_1 = \emptyset$, resp. $S_2 = \emptyset$). In [21, Cor. 7], Gimbert and Zielonka establish that memoryless strategies also suffice in *two-player* games with the same payoff. Thanks to Thm. 3, this entails the next theorem.

Theorem 5. *Average-energy games are determined and both players have memoryless optimal strategies.*

Solving average-energy games. By Thm. 5, one can guess an optimal memoryless strategy for \mathcal{P}_2 and solve the remaining one-player game for \mathcal{P}_1 , in polynomial time (by Thm. 4). The converse is also true: one can guess the strategy of \mathcal{P}_1 and solve the remaining game where $S_1 = \emptyset$ in polynomial time.

Theorem 6. *The AE problem for two-player games is in $\text{NP} \cap \text{coNP}$.*




Fig. 3: Example of a one-player AE_{LU} game ($U = 3$) and the evolution of energy under different strategies that maintain it within $[0, 3]$ at all times. The minimal average-energy is obtained with **play π_3** : alternating in order between the $+1$, $+2$ and -3 cycles.

We prove that MP games can be encoded into AE ones in polynomial time. The former are known to be in $NP \cap coNP$ but whether they belong to P is a long-standing open question (e.g., [33, 24, 8, 12]). Hence, w.r.t. current knowledge, the $NP \cap coNP$ -membership of the AE problem can be considered optimal. The key of the construction is to double each edge of the original game and modify the weight function such that each pair of successive edges corresponding to such a doubled edge now has a total energy level of zero, and an average-energy that is exactly equal to the weight of the original edge. Then we apply decomposition techniques as in Lem. 2 to establish the equivalence.

Theorem 7. *Mean-payoff games can be reduced to average-energy games in polynomial time.*

4 Average-Energy with Lower- and Upper-Bounded Energy

We extend the AE framework with constraints on the running energy level of the system. Such constraints are natural in many applications where the energy capacity is bounded (e.g., fuel tank, battery charge). We first study the case where the energy is subject to *both* a lower bound (here, zero) *and* an upper bound ($U \in \mathbb{N}$). We study the problem for the *fixed initial energy level* $c_{init} := 0$.

Problem 2 (AE_{LU}). *Given a game G , an initial state s_{init} , an upper bound $U \in \mathbb{N}$, and a threshold $t \in \mathbb{Q}$, decide if \mathcal{P}_1 has a winning strategy $\sigma_1 \in \Sigma_1$ for the objective $Energy_{LU}(U, c_{init} := 0) \cap AvgEnergy(t)$.*

Illustration. Consider the one-player game in Fig. 3. The energy constraints force \mathcal{P}_1 to keep the energy in $[0, 3]$ at all times. Hence, only three strategies can be followed safely, respectively inducing plays π_1 , π_2 and π_3 . Due to the bounds on energy, it is natural that strategies need to alternate between both a positive and a negative cycle to satisfy objective $Energy_{LU}(U, c_{init} := 0)$ (since no simple zero cycle exists). It is yet interesting that to play optimally (play π_3), \mathcal{P}_1 actually has to use *both* positive cycles, and in the *appropriate order* (compare plays π_2 and π_3). This type of alternation is more intricate than for other classical objectives [11, 14, 32]. This gives a hint of the complexity of AE_{LU} games.

4.1 Pseudo-polynomial algorithm and complexity bounds

We first reduce the AE_{LU} problem to the AE problem over a *pseudo-polynomial expanded game*, i.e., polynomial in the size of the original AE_{LU} game and in $U \in \mathbb{N}$. By Thm. 6 and Thm. 4, this reduction induces $NEXPTIME \cap coNEXPTIME$ -membership of the two-player AE_{LU} problem, and $EXPTIME$ -membership of the one-player one. We improve the complexity for two-player games by further reducing the AE game to an MP game. This yields $EXPTIME$ -membership, which is optimal (Thm. 8).




Fig. 4: Reduction from the AE_{LU} game in Fig. 3a to an AE game and further reduction to an MP game over the same expanded graph. For the sake of succinctness, the weights are written as $c | c'$ with c the weight used in the AE game and c' the one used in the MP game. We use the upper bound $U = 3$ and the average-energy threshold $t = 1$ (the optimal value in this case). The optimal play $\pi_3 = (acaab)^\omega$ of the original game corresponds to an optimal memoryless play in the expanded graph.


Observe that if U is encoded in unary or if U is polynomial in the size of the original game, the complexity of the AE_{LU} problem collapses to $NP \cap coNP$ for two-player games and to P for one-player games thanks to our reduction to an AE problem and the results of Thm. 6 and Thm. 4.

The reductions. Given a game $G = (S_1, S_2, E, w)$, an initial state s_{init} , an upper bound $U \in \mathbb{N}$, and a threshold $t \in \mathbb{Q}$, we reduce the AE_{LU} problem to an AE problem as follows. If at any point along a play, the energy drops below zero or exceeds U , the play will be losing for the $Energy_{LU}(U, c_{\text{init}} := 0)$ objective, hence also for its conjunction with the AE one. So we build a new game G' over the state space $(S \times \{0, 1, \dots, U\}) \cup \{\text{sink}\}$. The idea is to include the energy level within the state labels, with sink as an absorbing state reached only when the energy constraint is breached. We now consider the AE problem for threshold t on G' . By putting a self-loop of weight 1 on sink, we ensure that if the energy constraint is not guaranteed in G , the answer to the AE problem in G' will be No as the average-energy will be infinite due to reaching this positive loop and repeating it forever. Hence, we show that the AE_{LU} objective can be won in G if and only if the AE one can be won in G' (thus avoiding the sink state). The result of the reduction for the game in Fig. 3a is presented in Fig. 4.

We then show that the AE game G' can be further reduced to an MP game G'' by modifying the weight structure of the graph. Essentially, all edges leaving a state (s, c) of G' are given weight c in G'' , i.e., the current energy level, and the self-loop on sink is given weight $(\lceil t \rceil + 1)$. This modification is depicted in Fig. 4. We claim that the AE problem for threshold $t \in \mathbb{Q}$ in G' is equivalent to the MP problem for the same threshold in G'' . Indeed, we show that with our change of weight function, reaching sink implies losing, both in G' for AE and in G'' for MP , and all plays that *do not* reach sink have the same value for their average-energy in G' as for their mean-payoff in G'' .

Illustration. Consider the AE_{LU} game G in Fig. 3a. The optimal strategy is $\pi_3 = (acaab)^\omega$. Now consider the reduction to the AE game, and further to the MP game, depicted in Fig. 4. The optimal (memoryless) strategy in both the AE game G' and the MP game G'' is to create the play $\pi' = ((a,0)(c,1)(a,1)(a,3)(b,0))^\omega$, which corresponds to the optimal play π_3 in the original game. It can be checked that $AE_G(\pi_3) = \overline{AE}_{G'}(\pi') = \overline{MP}_{G''}(\pi')$.

Complexity. The reduction from the AE_{LU} game to the AE one induces a pseudo-polynomial blow-up in the number of states. Thanks to the second reduction and the use of a pseudo-polynomial algorithm for the MP game [33, 8], we get EXPTIME-membership, which is optimal for two-player games thanks to the lower bound proved for EG_{LU} [5].



(a) \mathcal{P}_1 needs to take U times (s, s') before taking (s, s) once and repeating. (b) \mathcal{P}_2 needs to increase the energy up to U using (a, c) to force \mathcal{P}_1 to take (g, d) then make him lose by taking (a, b) .

Fig. 5: Families of games witnessing the need for pseudo-polynomial-memory strategies for EG_{LU} (and AE_{LU}) objectives. The goal of \mathcal{P}_1 is to keep the energy in $[0, U]$ at all times, for $U \in \mathbb{N}$. The left game is won by \mathcal{P}_1 and the right one by \mathcal{P}_2 but both require memory polynomial in the value U to be won.

Theorem 8. *The AE_{LU} problem is EXPTIME-complete for two-player games and at least PSPACE-hard for one-player games. If the upper bound $U \in \mathbb{N}$ is polynomial in the size of the game or encoded in unary, the AE_{LU} problem collapses to $\text{NP} \cap \text{coNP}$ and P for two-player and one-player games respectively.*

4.2 Memory requirements

We prove pseudo-polynomial lower and upper bounds on memory for the two players in AE_{LU} games. The upper bound follows from the reduction to a pseudo-polynomial AE game and the memoryless determinacy of AE games proved in Thm. 5. The lower bound can be witnessed in two families of games asking for strategies using memory polynomial in the energy upper bound $U \in \mathbb{N}$ to be won by \mathcal{P}_1 (Fig. 5a) or \mathcal{P}_2 (Fig. 5b) respectively. It is interesting to observe that those families already ask for such memory when considering the simpler EG_{LU} objective.

Theorem 9. *Pseudo-polynomial-memory strategies are both sufficient and necessary to win in EG_{LU} and AE_{LU} games with arbitrary energy upper bound $U \in \mathbb{N}$, for both players. Polynomial memory suffices when U is polynomial in the size of the game or encoded in unary.*

5 Average-Energy with Lower-Bounded Energy

We conclude with the conjunction of an AE objective with a lower bound (again equal to zero) constraint on the running energy, but no upper bound. This corresponds to an *hypothetical* unbounded energy storage. Hence, its applicability is limited, but it may prove interesting on the theoretical standpoint.

Problem 3 (AE_L). *Given a game G , an initial state s_{init} and a threshold $t \in \mathbb{Q}$, decide if \mathcal{P}_1 has a winning strategy $\sigma_1 \in \Sigma_1$ for objective $\text{Energy}_L(c_{\text{init}} := 0) \cap \text{AvgEnergy}(t)$.*

This problem proves to be challenging to solve: we provide partial answers in the following, with a proper algorithm for one-player games but only a correct but incomplete method for two-player games.

Illustration. Consider the game in Fig. 3. Recall that for AE_{LU} with $U = 3$, the optimal play is π_3 , and it requires alternation between all three different simple cycles. Now consider AE_L . One may think that relaxing the objective would allow for simpler winning strategies. This is not the case. Some new plays are now acceptable w.r.t. the energy constraint, such as $\pi_4 = (abaaba)^\omega$, with $\overline{AE}(\pi_4) = 11/7$ and $\pi_5 = (aababa)^\omega$, with $\overline{AE}(\pi_5) = 18/7$. Yet, the optimal play w.r.t. the AE (under the lower-bound energy constraint) is still π_3 , hence still requires to use all the available cycles, in the appropriate order.

5.1 One-player games

We assume that the unique player is \mathcal{P}_1 . Indeed, the opposite case is easy as for \mathcal{P}_2 , the objective is a disjunction and \mathcal{P}_2 can choose beforehand which sub-objective he will transgress, and do so with a simple memoryless strategy (both AE and EG_L games admit memoryless optimal strategies as seen before). We show how to solve a one-player AE_L problem in *pseudo-polynomial time* by reduction to an AE_{LU} problem for a well-chosen upper bound $U \in \mathbb{N}$ and then application of the algorithm of Sect. 4.1.

The reduction. Given a game $G = (S_1, S_2 = \emptyset, E, w)$, an initial state s_{init} , and a threshold $t \in \mathbb{Q}$, we reduce the AE_L problem to an AE_{LU} problem with an upper bound $U \in \mathbb{N}$ which is pseudo-polynomial in the original problem. Precisely, $U := t + N^2 + N^3$, with $N = W \cdot (|S| + 2)$. The intuition is that if \mathcal{P}_1 can win a one-player AE_L game, he can win it without ever reaching energy levels higher than the chosen bound U , even if he is technically allowed to do so. Essentially, the interest of increasing the energy is making more cycles available (as they become safe to take w.r.t. the lower bound constraint), but increasing the energy further than necessary is not a good idea as it will negatively impact the average-energy. To prove this reduction, we start from an arbitrary winning path in the AE_L game, and build a witness path that is still winning for the AE_L objective, but also keeps the energy below U at all times. Our construction exploits a result of Lafourcade *et al.* that bounds the value of the counter along a path in a one-counter automaton [28]. We build upon it to define an appropriate transformation leading to the witness path and derive a sufficiently large upper bound $U \in \mathbb{N}$ for the AE_{LU} problem.

Complexity. Plugging this bound U in the pseudo-polynomial-time algorithm for AE_{LU} games yields an algorithm for one-player AE_L games that is overall also pseudo-polynomial. We prove that no truly-polynomial-time algorithm can be obtained unless $P = NP$ as the one-player AE_L problem is NP-hard. We show it by reduction from the subset-sum problem [18].

Memory requirements. Recall that for \mathcal{P}_2 , the situation is simpler and memoryless strategies suffice. By the reduction to AE_{LU} , we know that pseudo-polynomial memory suffices for \mathcal{P}_1 . This bound is tight as witnessed by the family of games already presented in Fig. 5a. To ensure the lower bound on energy, \mathcal{P}_1 has to play edge (s, s') at least U times before taking the (s, s) self-loop. But to minimize the average-energy, edge (s, s') should never be played more than necessary. The optimal strategy is the same as for the AE_{LU} problem: playing (s, s') exactly U times, then (s, s) once, then repeating, forever.

Theorem 10. *Pseudo-polynomial-memory strategies are both sufficient and necessary to win for \mathcal{P}_1 in one-player AE_L games. Memoryless strategies suffice for \mathcal{P}_2 in such games.*

5.2 Two-player games

Decidability. Assume that there exists some $U \in \mathbb{N}$ such that \mathcal{P}_1 has a winning strategy for the AE_{LU} problem with upper bound U and average-energy threshold t . Then, this strategy is trivially winning for the AE_L problem as well. This observation leads to an incremental algorithm that is correct (no false positives) but incomplete (it is not guaranteed to stop). In [6], we draw the outline of a potential approach to obtain completeness hence decidability.

Lemma 11. *There is an algorithm that takes as input an AE_L problem and iteratively solves corresponding AE_{LU} problems for incremental values of $U \in \mathbb{N}$. If a winning strategy is found for some $U \in \mathbb{N}$, then it is also winning for the original AE_L problem. If no strategy is found up to value $U \in \mathbb{N}$, then no strategy of \mathcal{P}_1 can simultaneously win the AE_L problem and prevent the energy from exceeding U at all times.*

While an incomplete algorithm clearly seems limiting from a theoretical standpoint, it is worth noting that in practice, such approaches are common and often necessary restrictions, even for problems where a

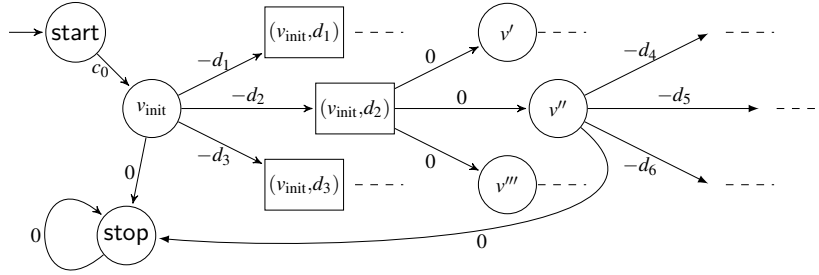


Fig. 6: Reduction from a countdown game $\mathcal{C} = (\mathcal{V}, \mathcal{E})$ with initial configuration (v_{init}, c_0) to a two-player AE_L problem for average-energy threshold $t := 0$.

complete algorithm is known to exist — because theoretical bounds granting completeness are too large to be tackled efficiently by software synthesis tools (e.g., [14]). In our case, we have already seen that if such a bound exists for the two-player AE_L problem, it needs to be at least exponential in the encoding of problem (cf. one-player AE_L games). Hence it seems likely that a prohibitive bound would be necessary, rendering the incremental algorithm of Lem. 11 more appealing in practice.

Complexity lower bound. We now prove that the two-player AE_L problem would require at least exponential time to solve. Our proof is by reduction from *countdown games*. A countdown game \mathcal{C} is a weighted graph $(\mathcal{V}, \mathcal{E})$, where \mathcal{V} is the finite set of states, and $\mathcal{E} \subseteq \mathcal{V} \times \mathbb{N} \setminus \{0\} \times \mathcal{V}$ is the edge relation. Configurations are of the form (v, c) , $v \in \mathcal{V}$, $c \in \mathbb{N}$. The game starts in an initial configuration (v_{init}, c_0) and transitions from a configuration (s, c) are performed as follows. First, \mathcal{P}_1 chooses a duration d , $0 < d \leq c$ such that there exists $e = (v, d, v') \in \mathcal{E}$ for some $v' \in \mathcal{V}$. Second, \mathcal{P}_2 chooses a state $v' \in \mathcal{V}$ such that $e = (v, d, v') \in \mathcal{E}$. Then the game advances to $(v', c - d)$. Terminal configurations are reached whenever no legitimate move is available. If such a configuration is of the form $(v, 0)$, \mathcal{P}_1 wins the play, otherwise \mathcal{P}_2 wins. Deciding the winner given an initial configuration (v_{init}, c_0) is EXPTIME-complete [25].

Our reduction is depicted in Fig. 6. The EL is initialized to c_0 , then it is decreasing along any play. Consider the AE_L objective for AE threshold $t := 0$. To ensure that the energy always stays non-negative, \mathcal{P}_1 has to switch to stop while the EL is no less than zero. In addition, to ensure an AE no more than $t = 0$, \mathcal{P}_1 has to obtain an EL at most equal to zero before switching to stop (as the AE will be equal to this EL thanks to Lem. 1 and the zero self-loop on stop). Hence, \mathcal{P}_1 wins the AE_L objective only if he can ensure a total sum of chosen durations that is *exactly* equal to c_0 , i.e., if he can reach a winning terminal configuration for the countdown game. The converse also holds.

Lemma 12. *The AE_L problem is EXPTIME-hard for two-player games.*

Memory requirements. We establish that memory is needed for both players.

Lemma 13. *Pseudo-polynomial-memory strategies are necessary to win for \mathcal{P}_1 in two-player AE_L games. Memory is also required for \mathcal{P}_2 in such games.*

References

- [1] B. Aminof & S. Rubin (2014): *First Cycle Games*. In: *Proc. of SR, EPTCS* 146, pp. 83–90.
- [2] H. Björklund, S. Sandberg & S. Vorobyov (2004): *Memoryless Determinacy of Parity and Mean Payoff Games: A Simple Proof*. *Theoretical Computer Science* 310(1-3), pp. 365–378, doi:[10.1016/S0304-3975\(03\)00427-4](https://doi.org/10.1016/S0304-3975(03)00427-4).
- [3] R. Bloem, K. Chatterjee, T.A. Henzinger & B. Jobstmann (2009): *Better Quality in Synthesis through Quantitative Objectives*. In: *Proc. of CAV, LNCS* 5643, Springer, pp. 140–156.

- [4] E. Boros, K. Elbassioni, V. Gurvich & K. Makino (2015): *Markov Decision Processes and Stochastic Games with Total Effective Payoff*. In: *Proc. of STACS, LIPIcs* 30, Schloss Dagstuhl - LZI, pp. 103–115.
- [5] P. Bouyer, U. Fahrenberg, K.G. Larsen, N. Markey & J. Srba (2008): *Infinite Runs in Weighted Timed Automata with Energy Constraints*. In: *Proc. of FORMATS, LNCS* 5215, Springer, pp. 33–47.
- [6] P. Bouyer, N. Markey, M. Randour, Larsen K.G. & S. Laursen (2015): *Average-energy games*. Research Report, Laboratoire Spécification et Vérification, ENS Cachan, France. Available at <http://www.lsv.ens-cachan.fr/~randour/papers/BMRL15.pdf>.
- [7] T. Brázdil, D. Klaška, A. Kučera & P. Novotný (2014): *Minimizing Running Costs in Consumption Systems*. In: *Proc. of CAV, LNCS* 8559, Springer, pp. 457–472.
- [8] L. Brim, J. Chaloupka, L. Doyen, R. Gentilini & J.-F. 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](https://doi.org/10.1007/s10703-010-0105-x).
- [9] F. Cassez, J.J. Jensen, K.G. Larsen, J.-F. Raskin & P.-A. Reynier (2009): *Automatic Synthesis of Robust and Optimal Controllers – An Industrial Case Study*. In: *Proc. of HSCC, LNCS* 5469, Springer, pp. 90–104.
- [10] A. Chakrabarti, L. de Alfaro, T.A. Henzinger & M. Stoelinga (2003): *Resource Interfaces*. In: *Proc. of EMSOFT, LNCS* 2855, Springer, pp. 117–133.
- [11] K. Chatterjee & L. Doyen (2010): *Energy Parity Games*. In: *Proc. of ICALP, LNCS* 6199, Springer, pp. 599–610.
- [12] K. Chatterjee, L. Doyen, M. Randour & J.-F. Raskin (2015): *Looking at mean-payoff and total-payoff through windows*. *Information and Computation* 242, pp. 25 – 52, doi:[10.1016/j.ic.2015.03.010](https://doi.org/10.1016/j.ic.2015.03.010).
- [13] K. Chatterjee & V.S. Prabhu (2013): *Quantitative timed simulation functions and refinement metrics for real-time systems*. In: *Proc. of HSCC, ACM*, pp. 273–282, doi:[10.1145/2461328.2461370](https://doi.org/10.1145/2461328.2461370).
- [14] K. Chatterjee, M. Randour & J.-F. Raskin (2014): *Strategy synthesis for multi-dimensional quantitative objectives*. *Acta Informatica* 51(3-4), pp. 129–163, doi:[10.1007/s00236-013-0182-6](https://doi.org/10.1007/s00236-013-0182-6).
- [15] A. Ehrenfeucht & J. Mycielski (1979): *Positional strategies for mean payoff games*. *International Journal of Game Theory* 8(2), pp. 109–113, doi:[10.1007/BF01768705](https://doi.org/10.1007/BF01768705).
- [16] J. Fearnley & M. Jurdziński (2013): *Reachability in two-clock timed automata is PSPACE-complete*. In: *Proc. of ICALP, LNCS* 7966, Springer, pp. 212–223.
- [17] J. Filar & K. Vrieze (1997): *Competitive Markov decision processes*. Springer.
- [18] M.R. Garey & D.S. Johnson (1979): *Computers and intractability: a guide to the Theory of NP-Completeness*. Freeman New York.
- [19] T. Gawlitza & H. Seidl (2009): *Games through Nested Fixpoints*. In: *Proc. of CAV, LNCS* 5643, Springer, pp. 291–305.
- [20] H. Gimbert & W. Zielonka (2004): *When Can You Play Positionnaly?* In: *Proc. of MFCS, LNCS* 3153, Springer, pp. 686–697.
- [21] H. Gimbert & W. Zielonka (2005): *Games Where You Can Play Optimally Without Any Memory*. In: *Proc. of CONCUR, LNCS* 3653, Springer, pp. 428–442.
- [22] E. Grädel, W. Thomas & T. Wilke, editors (2002): *Automata, Logics, and Infinite Games: A Guide to Current Research*. LNCS 2500, Springer.
- [23] L. Juhl, K.G. Larsen & J.-F. Raskin (2013): *Optimal Bounds for Multiweighted and Parametrised Energy Games*. In: *Theories of Programming and Formal Methods, LNCS* 8051, Springer, pp. 244–255, doi:[10.1007/978-3-642-39698-4_15](https://doi.org/10.1007/978-3-642-39698-4_15).
- [24] M. Jurdziński (1998): *Deciding the Winner in Parity Games is in $UP \cap co-UP$* . *Information Processing Letters* 68(3), pp. 119–124, doi:[10.1016/S0020-0190\(98\)00150-1](https://doi.org/10.1016/S0020-0190(98)00150-1).
- [25] M. Jurdziński, J. Sproston & F. Laroussinie (2008): *Model Checking Probabilistic Timed Automata with One or Two Clocks*. *Logical Methods in Computer Science* 4(3), doi:[10.2168/LMCS-4\(3:12\)2008](https://doi.org/10.2168/LMCS-4(3:12)2008).

- [26] R.M. Karp (1978): *A characterization of the minimum cycle mean in a digraph*. *Discrete Mathematics* 23(3), doi:[10.1016/0012-365X\(78\)90011-0](https://doi.org/10.1016/0012-365X(78)90011-0).
- [27] E. Kopczynski (2006): *Half-Positional Determinacy of Infinite Games*. In: *Proc. of ICALP*, LNCS 4052, Springer, pp. 336–347.
- [28] P. Lafourcade, D. Lugiez & R. Treinen (2005): *Intruder Deduction for AC-Like Equational Theories with Homomorphisms*. In: *Proc. of RTA*, LNCS 3467, Springer, pp. 308–322.
- [29] M. Randour (2013): *Automated Synthesis of Reliable and Efficient Systems Through Game Theory: A Case Study*. In: *Proceedings of the European Conference on Complex Systems 2012*, Springer Proceedings in Complexity XVII, Springer, pp. 731–738, doi:[10.1007/978-3-319-00395-5_90](https://doi.org/10.1007/978-3-319-00395-5_90).
- [30] M. Randour (2014): *Synthesis in Multi-Criteria Quantitative Games*. Ph.D. thesis, University of Mons, Belgium.
- [31] F. Thuijsman & O.J. Vrieze (1987): *The bad match; A total reward stochastic game*. *OR Spektrum* 9(2), doi:[10.1007/BF01732644](https://doi.org/10.1007/BF01732644).
- [32] Y. Velner, K. Chatterjee, L. Doyen, T.A. Henzinger, A.M. Rabinovich & J.-F. 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](https://doi.org/10.1016/j.ic.2015.03.001).
- [33] U. Zwick & M. Paterson (1996): *The Complexity of Mean Payoff Games on Graphs*. *Theoretical Computer Science* 158(1-2), pp. 343–359, doi:[10.1016/0304-3975\(95\)00188-3](https://doi.org/10.1016/0304-3975(95)00188-3).