



HAL
open science

Average-energy games

Patricia Bouyer, Nicolas Markey, Mickaël Randour, Kim G. Larsen, Simon Laursen

► **To cite this version:**

Patricia Bouyer, Nicolas Markey, Mickaël Randour, Kim G. Larsen, Simon Laursen. Average-energy games. Acta Informatica, 2018, 55 (2), pp.91 - 127. <10.1007/s00236-016-0274-1>. <hal-01889005>

HAL Id: hal-01889005

<https://hal.science/hal-01889005v1>

Submitted on 5 Oct 2018

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

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



HAL Authorization

Average-energy games

Patricia Bouyer · Nicolas Markey ·
Mickael Randour · Kim G. Larsen ·
Simon Laursen

the date of receipt and acceptance should be inserted later

Abstract 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.

Work partially supported by European project CASSTING (FP7-ICT-601148) and ERC project EQualIS (StG-308087). Mickael Randour is an F.R.S.-FNRS Postdoctoral Researcher.

Patricia Bouyer · E-mail: bouyer@lsv.ens-cachan.fr
Nicolas Markey · E-mail: markey@lsv.ens-cachan.fr
LSV – CNRS & ENS Cachan
Avenue du Président Wilson, 61, 94230 Cachan, France

Mickael Randour · E-mail: mickael.randour@ulb.ac.be
Computer Science Department, Université libre de Bruxelles (ULB)
Avenue Franklin Roosevelt, 50, 1050 Bruxelles, Belgium

Kim G. Larsen · E-mail: kgl@cs.aau.dk
Simon Laursen · E-mail: simlau@cs.aau.dk
Aalborg University
Fredrik Bajers Vej 5, 9100 Aalborg, Denmark

1 Introduction

Quantitative games. Game-theoretic formulations are a standard tool for the synthesis of provably-correct controllers for reactive systems [24]. 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 [11, 3, 33], or [34] 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 [17, 38, 26], player 1 has to optimize his long-run average gain per edge taken whereas, in *TP* games [22, 21], player 1 has to optimize his long-run sum of weights. Energy games [11, 6, 25] 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 (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., [9, 13]). Energy games with both lower and upper bounds (*EG_{LU}*) are more complex: they are EXPTIME-complete and winning requires memory in general [6].

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 [12, 15, 37]. Similarly, window objectives address the need for strategies ensuring good quantitative behaviors within reasonable time frames [13].

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 [10], which we detail in the following as a motivating example. Interestingly, it turns out that this payoff first appeared long ago [36], 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 [14], Chatterjee and Prabhu define the *average debit-sum level* objective, which can be seen as a variation of the *average-energy* where

| Game objective | 1-player | 2-player | memory |
|----------------------------|-------------------|------------------------------|--------------------------|
| MP | P [28] | $NP \cap \text{coNP}$ [38] | memoryless [17] |
| TP | P [19] | $NP \cap \text{coNP}$ [21] | memoryless [22] |
| EG_L | P [6] | $NP \cap \text{coNP}$ [11,6] | memoryless [11] |
| EG_{LU} | PSPACE-c. [18] | EXPTIME-c. [6] | pseudo-polynomial |
| AE | P | $NP \cap \text{coNP}$ | memoryless |
| AE_{LU} , polynomial U | P | $NP \cap \text{coNP}$ | polynomial |
| AE_{LU} , arbitrary U | PSPACE-c. | EXPTIME-c. | pseudo-polynomial |
| AE_L | PSPACE-e. / NP-h. | open / EXPTIME-h. | open (\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, AE_L (resp. AE_{LU}) for average-energy under a lower bound (resp. and upper bound $U \in \mathbb{N}$) on the energy, c. for complete, e. for easy, and h. for hard. Results without reference are proved in this paper.

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. in [10]. 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 [21]: 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 $NP \cap \text{coNP}$ whereas it is in P for one-player games. In both cases, memoryless strategies suffice (Thm. 8). Those complexity bounds match the state-of-the-art for MP and TP games [38,26,21,9]. Furthermore we prove that AE games are at least as hard as mean-payoff games (Thm. 10). Therefore, the $NP \cap \text{coNP}$ -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. 6): 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 solving (Thm. 7). Finally, we lift memoryless determinacy to the two-player case using results by Gimbert and Zielonka [23] and obtain the $NP \cap \text{coNP}$ -membership as a corollary (Thm. 9).

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. 13). 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 [6], 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. 14). We show that one-player AE_{LU} games are PSPACE-complete via the on-the-fly construction of a witness path based on the aforementioned reduction, answering a question left open in [7]. 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 (resp. one-player) AE_{LU} problem collapses to $NP \cap \text{coNP}$ (resp. P) 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 show PSPACE-membership for the one-player case (Thm. 17), 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 (Thm. 17) and EXPTIME-hard in two-player games (Lem. 20) 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. 21), and that pseudo-polynomial memory is both sufficient and necessary in the one-player case (Thm. 18). The decidability status of two-player AE_L games remains open as we only provide a correct but incomplete incremental algorithm (Lem. 19). 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. This paper extends previous work presented in a conference [7]: it gives a full presentation of the technical details, along with additional results and improved complexities.

The *average-energy* payoff—Eq. (1)—appeared in a paper by Thuijsman and Vrieze in the late eighties [36], 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 [22, 21]), which, despite that, has been referred in many papers as either total-payoff or total-reward equivalently. We will see 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 [5], 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 [23] to bypass the reduction to discounted games and obtain a direct proof, thanks to our more constructive approach. Lastly, while [5] 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 [10].

Recent work of Brázdil et al. [8] 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.

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* such that for all $s \in S$, there exists $s' \in S$ such that $(s, s') \in E$ (i.e., no deadlock), 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 first (resp. last) element s_0 (resp. s_n) is denoted $\text{first}(\pi(n))$ (resp. $\text{last}(\pi(n))$). The set of all plays

in G is denoted by $Plays(G)$ and the set of all prefixes is denoted by $Prefs(G)$. We say that a prefix $\rho \in Prefs(G)$ belongs to \mathcal{P}_i , $i \in \{1, 2\}$, if $last(\rho) \in S_i$. The set of prefixes that belong to \mathcal{P}_i is denoted by $Prefs_i(G)$. 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 $\rho \in Prefs(G)$. 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 $\sigma_i: Prefs_i(G) \rightarrow S$ such that for all $\rho \in Prefs_i(G)$ we have $(last(\rho), \sigma_i(\rho)) \in E$. A strategy σ_i for \mathcal{P}_i is *finite-memory* if it can be encoded by a deterministic Moore machine $(M, m_0, \alpha_u, \alpha_n)$ where M is a finite set of states (the memory of the strategy), $m_0 \in M$ is the initial memory state, $\alpha_u: M \times S \rightarrow M$ is an update function, and $\alpha_n: M \times S_i \rightarrow S$ is the next-action function. If the game is in $s \in S_i$ and $m \in M$ is the current memory value, then the strategy chooses $s' = \alpha_n(m, s)$ as the next state of the game. When the game leaves a state $s \in S$, the memory is updated to $\alpha_u(m, s)$. Formally, $(M, m_0, \alpha_u, \alpha_n)$ defines the strategy σ_i such that $\sigma_i(\rho \cdot s) = \alpha_n(\hat{\alpha}_u(m_0, \rho), s)$ for all $\rho \in S^*$ and $s \in S_i$, where $\hat{\alpha}_u$ extends α_u to sequences of states as expected. A strategy is *memoryless* if $|M| = 1$, i.e., 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 $last(\pi(n)) \in S_i$, we have $\sigma_i(\pi(n)) = s_{n+1}$. Given an initial state $s_{init} \in S$ and strategies σ_1 and σ_2 for the two players, we denote by $Outcome(s_{init}, \sigma_1, \sigma_2)$ 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 $Outcomes(s_{init}, \sigma_i)$.

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 lim sup variants of MP , TP and AE , but similar results are obtained for the lim inf 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 [11,6], MP [17] and TP [19,22] 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 [6,38,21]. For the EG_{LU} objective, memory is in general needed and the associated decision problem is EXPTIME-complete [6] (PSPACE-complete for one-player games [18]).

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. 8).

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 $\text{AvgEnergy}(t)$.*

We first compare the *AE* objective with traditional quantitative objectives and study how they can be connected (Sect. 3.1). Then we want to establish that in *AE* games, memoryless strategies are always sufficient to play optimally, for *both* players. Interestingly, this result cannot be obtained by straightforward application of many well-known sufficient criteria for memoryless determinacy existing in the literature. We thus introduce some technical lemmas that highlight the inherent features of the *AE* payoff function (Sect. 3.2) and permit to prove the result for *one-player AE* games (Sect. 3.3). We then prove that one-player *AE* games can be solved in polynomial-time via an algorithm combining graph analysis techniques with linear programming. Finally, we consider the *two-player* case (Sect. 3.4). Applying a result by Gimbert and Zielonka [23], combined with our results on the one-player case, we derive memoryless determinacy of two-player *AE* games. This also induces $\text{NP} \cap \text{coNP}$ -membership of the *AE* problem by the P algorithm of Sect. 3.3. We conclude by proving that *AE* games are at least as hard as *MP* games, hence indicating that the $\text{NP} \cap \text{coNP}$ upper bound is essentially optimal with regard to our current knowledge of *MP* games (whose membership to P is a long-standing open problem [38, 26, 9, 13]).

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 below zero along a play, hence granting a lower bound on total-payoff. We introduce the sign-reversed game G' in the next lemma, which is consistent with our view of \mathcal{P}_1 as a minimizer with regard to payoffs (as discussed in Sect. 2).


Lemma 1 *Let $G = (S_1, S_2, E, w)$ be a game and $s_{\text{init}} \in S$ be the initial state. The following assertions are equivalent.*

- A. *There exists $c_{\text{init}} \in \mathbb{N}$ such that \mathcal{P}_1 has a (memoryless) winning strategy for objective $\text{Energy}_L(c_{\text{init}})$.*
- B. *Player \mathcal{P}_1 has a (memoryless) winning strategy for objective $\text{MeanPayoff}(0)$ in the game G' defined by reversing the sign of the weight function, i.e., for all $(s_1, s_2) \in E$, $w'(s_1, s_2) = -w(s_1, s_2)$.*
- C. *Player \mathcal{P}_1 has a (memoryless) winning strategy for objective $\text{TotalPayoff}(t)$, with $t = 2 \cdot (|S| - 1) \cdot W$, in the game G' defined by reversing the sign of the weight function.*
- D. *There exists $t \in \mathbb{Z}$ such that \mathcal{P}_1 has a (memoryless) winning strategy for objective $\text{TotalPayoff}(t)$, in the game G' defined by reversing the sign of the weight function.*

Proof Proof of $A \Leftrightarrow B$ is given in [6, Proposition 12]. Proof of $B \Leftrightarrow C \Leftrightarrow D$ is in [13, Lem. 1]. \square

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 [21]. 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).



(c) Play π_1 sees energy levels $(1, 3, 5, 3)^\omega$. (d) Play π_2 sees energy levels $(1, 3, 5, 5, 3)^\omega$.

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: $\overline{AE}(\pi_1) = \underline{AE}(\pi_1) = 3 < \overline{AE}(\pi_2) = \underline{AE}(\pi_2) = 11/3$.

In these examples, the average-energy is clearly comprised between the infimum and supremum total-payoffs. This remains true for any play.

Lemma 2 For any play $\pi \in \text{Plays}(G)$, we have that

$$\underline{AE}(\pi), \overline{AE}(\pi) \in [\underline{TP}(\pi), \overline{TP}(\pi)] \subseteq \mathbb{R} \cup \{-\infty, \infty\}.$$

Proof Consider a play $\pi \in \text{Plays}(G)$. By definition of the total-payoff and thanks to weights taking integer values, we have that there exists some index $m \in \mathbb{N}_0$ such that, for all $n \geq m$, $EL(\pi(n)) \in [\underline{TP}(\pi), \overline{TP}(\pi)]$. By definition, the average-energy \overline{AE} (resp. \underline{AE}) measures the supremum (resp. infimum) limit of the averages of those partial sums, hence it holds that $\underline{AE}(\pi), \overline{AE}(\pi) \in [\underline{TP}(\pi), \overline{TP}(\pi)]$. \square

In particular, if the mean-payoff value from a state is not zero, its total-payoff value is infinite and the following lemma holds.

Lemma 3 Let $G = (S_1, S_2, E, w)$ be a game and $s_{\text{init}} \in S$ be the initial state.

1. If there exists $t < 0$ such that \mathcal{P}_1 has a (memoryless) winning strategy for objective $\text{MeanPayoff}(t)$, then \mathcal{P}_1 has a memoryless strategy that is winning for $\text{AvgEnergy}(t')$ for all $t' \in \mathbb{Q}$, i.e., this strategy ensures that any consistent outcome π is such that $\underline{AE}(\pi) = \overline{AE}(\pi) = -\infty$.

2. If \mathcal{P}_1 has no (memoryless) winning strategy for $\text{MeanPayoff}(0)$, then, for any $t' \in \mathbb{Q}$, \mathcal{P}_1 has no winning strategy for $\text{AvgEnergy}(t')$. In particular, \mathcal{P}_2 has a memoryless strategy ensuring that any consistent outcome π is such that $\underline{AE}(\pi) = \overline{AE}(\pi) = \infty$.

Proof Consider the first implication. Assume \mathcal{P}_1 has a memoryless strategy σ_1 ensuring that all outcomes $\pi \in \text{Outcomes}(s_{\text{init}}, \sigma_1)$ are such that $\overline{MP}(\pi) < 0$. For any such outcome, it is guaranteed that all simple cycles have a strictly negative energy level. Thus, we have that $\overline{TP}(\pi) = -\infty$, and by Lem. 2, it implies that $\overline{AE}(\pi) = -\infty$, as claimed. Since $\underline{AE}(\pi) \leq \overline{AE}(\pi)$ by definition, the property holds.

Now consider the second implication. Assume there exists no winning strategy for \mathcal{P}_1 for the mean-payoff objective. By equivalence **B** \Leftrightarrow **D** of Lem. 1, and memoryless determinacy of total-payoff games (see for example [22]), it follows that \mathcal{P}_2 has a memoryless strategy σ_2 ensuring that all consistent outcomes $\pi \in \text{Outcomes}(s_{\text{init}}, \sigma_2)$ are such that $\underline{TP}(\pi) = \infty$. By Lem. 2, this induces the claim. \square

3.2 Useful properties of the average-energy

In this subsection, we will first review some classical criteria that usually prove sufficient to deduce memoryless determinacy in quantitative games and discuss why they cannot be applied straight out of the box to the average-energy payoff. We will then prove two useful properties of this payoff that will later help us to prove the desired result.

Classical sufficient criteria. We briefly discuss traditional approaches to prove memoryless determinacy in quantitative games. The first one is to study a variant of the infinite-duration game where the game halts as soon as a cycle is closed and then to relate the properties of this variant to the infinite-duration game. This technique was used in the original proof of memoryless determinacy for mean-payoff games by Ehrenfeucht and Mycielski [17], and in a following simpler proof by Björklund et al. [2]. The connection between infinite-duration games and so-called *first cycle games* was recently streamlined by Aminof and Rubin [1], identifying sufficient conditions to prove that first cycle games and their infinite-duration counterparts admit optimal memoryless strategies for both players. Among those conditions is the need for winning objectives to be closed under *cyclic permutation* (intuitively, swapping cycles in a play should not induce a better payoff) and under *concatenation* (intuitively, concatenating two prefixes should not result in a payoff better than the best of the two prefixes). Without further assumptions, the average-energy objective satisfies neither. Indeed, consider individual 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. 5.

Other criteria for memoryless determinacy or half-memoryless determinacy (i.e., holding only for one of the two players) respectively appear in works by Gimbert and Zielonka [22] and by Kopczynski [29]. They involve checking that the payoff is *fairly mixing*, or *concave*. Again, both are false for arbitrary sequences of weights in the case of the average-energy, for essentially the same reasons as above. Nevertheless, we will be able to prove that memoryless strategies suffice for both players using similar ideas but first taking care of the problematic cases. Intuitively, when those cases are dealt with, we will regain a payoff that satisfies the above conditions. We also obtain *monotonicity* and *selectivity* of the payoff function as defined in [23].

Extraction of prefixes. The following lemma describes the impact of adding a finite prefix to an infinite play. We prove that the average-energy over a play can be decomposed w.r.t. to the energy level of any of its prefixes and the average-energy of the remaining suffix.

Lemma 4 (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} .

Proof Let $\rho = s_0 \dots s_k \in \text{Pref}_s(G)$ and $\pi \in \text{Plays}(G)$ be a prefix and a play over a game G . We prove the property for \overline{AE} . By definition and decomposition, we have that

$$\begin{aligned} \overline{AE}(\rho \cdot \pi) &= \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=1}^n EL((\rho \cdot \pi)(i)) \\ &= \limsup_{n \rightarrow \infty} \left[\frac{1}{n} \cdot \sum_{i=1}^k EL(\rho(i)) + \frac{1}{n} \cdot \sum_{i=k+1}^n EL(\rho) + \frac{1}{n} \cdot \sum_{i=k+1}^n EL(\pi(i-k)) \right]. \end{aligned}$$

For clarity, we rewrite this expression as $\overline{AE}(\rho \cdot \pi) = \limsup_{n \rightarrow \infty} [X_1(n) + X_2(n) + X_3(n)]$, maintaining the same order.

Since k is fixed and finite, and $EL(\rho(i))$ is bounded for all $i \leq k$, we have that $\limsup_{n \rightarrow \infty} X_1(n) = \lim_{n \rightarrow \infty} X_1(n) = 0$. Furthermore, for $n \geq k+1$, we rewrite the second term as $X_2(n) = (n-k-1) \cdot EL(\rho)/n$, and it follows that $\limsup_{n \rightarrow \infty} X_2(n) = \lim_{n \rightarrow \infty} X_2(n) = EL(\rho)$. Since both sequences $X_1(n)$ and $X_2(n)$ converge, we can write

$$\begin{aligned} \liminf_{n \rightarrow \infty} X_1(n) + \liminf_{n \rightarrow \infty} X_2(n) + \limsup_{n \rightarrow \infty} X_3(n) &\leq \overline{AE}(\rho \cdot \pi) \\ &\leq \limsup_{n \rightarrow \infty} X_1(n) + \limsup_{n \rightarrow \infty} X_2(n) + \limsup_{n \rightarrow \infty} X_3(n). \end{aligned}$$

Hence, by a small change of variable,

$$\begin{aligned} \overline{AE}(\rho \cdot \pi) &= EL(\rho) + \limsup_{n \rightarrow \infty} X_3(n) = EL(\rho) + \limsup_{n \rightarrow \infty} \left[\frac{1}{n} \cdot \sum_{i=1}^{n-k-1} EL(\pi(i)) \right] \\ &= EL(\rho) + \overline{AE}(\pi), \end{aligned}$$

as, in the limit, the $(k+1)$ missing terms in the sum are negligible. The proof for \underline{AE} is similar. \square

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.

Lemma 5 (Repeated zero cycles of bounded length) *Let C_1, C_2, C_3, \dots be an infinite sequence of cycles $C_i \in \text{Prefs}(G)$ such that (i) $\pi = C_1 \cdot C_2 \cdot C_3 \dots \in \text{Plays}(G)$,¹ (ii) $\forall i \geq 1, EL(C_i) = 0$ and (iii) $\exists \ell \in \mathbb{N}_{>0}$ such that $\forall i \geq 1, |C_i| \leq \ell$. Then the following properties hold.*

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

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

2. *For any cycle $C \in \text{Prefs}(G)$ such that $EL(C) = 0$, we have that $\overline{AE}(C^\omega) = AE(C)$.*
3. *Repeating the best cycle gives the lowest AE :*

$$\inf_{i \in \mathbb{N}_{>0}} AE(C_i) = \inf_{i \in \mathbb{N}_{>0}} \overline{AE}((C_i)^\omega) \leq \overline{AE}(\pi)$$

Similar properties hold for \underline{AE} .

Observe that since we assume a bound $\ell \in \mathbb{N}_{>0}$ on the length of cycles, and the game is played on a finite graph, Point 3 of Lem. 5 does actually allow to select a *best* cycle: the set of possible cycles of length at most ℓ is finite and the infimum is reached, hence can be replaced by the minimum.

Proof We prove the three points for \overline{AE} , similar arguments can be applied for \underline{AE} . Consider Point 1. Let $\pi = s_0^1 \dots s_{|C_1|}^1 s_1^2 \dots s_{|C_2|}^2 s_1^3 \dots$ where s_j^i denotes the j -th state of cycle C_i , with $C_1 = s_0^1 \dots s_{|C_1|}^1$ and for all $i > 1, C_i = s_{|C_{i-1}|}^{i-1} s_1^i \dots s_{|C_i|}^i$. Essentially, $s_{|C_{i-1}|}^{i-1}$ is both the last state of C_{i-1} and the first one of C_i : it can also be seen as s_0^i and we later use both notations depending on the role we consider for this state. Given index $k \in \mathbb{N}$ of a state s_k in the classical formulation $\pi = s_0 s_1 s_2 \dots$ such that s_k denotes state s_j^i in our new formulation $\pi = s_0^1 \dots s_{|C_1|}^1 s_1^2 \dots s_{|C_2|}^2 s_1^3 \dots$, we define $c(k) = i$ and $p(k) = j$, respectively denoting the index of the corresponding cycle and the position of state s_k within this cycle. We can rewrite the definition of the average-energy of π as

$$\begin{aligned} \overline{AE}(\pi) &= \limsup_{n \rightarrow \infty} \left[\frac{1}{n} \sum_{k=1}^n EL(\pi(k)) \right] \\ &= \limsup_{n \rightarrow \infty} \left[\frac{1}{n} \left(\sum_{i=1}^{c(n)-1} \sum_{j=1}^{|C_i|} EL(s_0^1 \dots s_j^i) + \sum_{j=1}^{p(n)} EL(s_0^1 \dots s_j^{c(n)}) \right) \right]. \quad (3) \end{aligned}$$

¹ We slightly abuse the notation as we see cycles as sequences of *edges*. The concatenation of cycles $C_a = s s' \dots s$ and $C_b = s s'' \dots s$ is to be understood as its natural interpretation $C_a \cdot C_b = s s' \dots s s'' \dots s$: the origin state s only appears *once* in the middle and not twice as it would with C_a and C_b seen as true sequences of states.

Observe that since all cycles are such that $EL(\mathcal{C}_i) = 0$, we have that $EL(s_0^1 \dots s_j^i) = EL(s_0^i \dots s_j^i)$ for all indices $i \in \mathbb{N}_{>0}$, $j \in \{1, \dots, |\mathcal{C}_i|\}$. In other words, the energy level in a given position only depends on the current cycle. Hence, for all $i \in \mathbb{N}_{>0}$,

$$\sum_{j=1}^{|\mathcal{C}_i|} EL(s_0^1 \dots s_j^i) = \sum_{j=1}^{|\mathcal{C}_i|} EL(s_0^i \dots s_j^i) = |\mathcal{C}_i| \cdot AE(\mathcal{C}_i)$$

where the second equality follows by definition of $AE(\mathcal{C}_i)$. Therefore, Eq. (3) becomes

$$\overline{AE}(\pi) = \limsup_{n \rightarrow \infty} \left[\frac{1}{n} \left(\sum_{i=1}^{c(n)-1} |\mathcal{C}_i| \cdot AE(\mathcal{C}_i) + \sum_{j=1}^{p(n)} EL(s_0^{c(n)} \dots s_j^{c(n)}) \right) \right].$$

Recall that, by hypothesis, there exists $\ell \in \mathbb{N}_{>0}$ such that for all $i \geq 1$, $|\mathcal{C}_i| \leq \ell$. Observe that the boundedness of cycles length implies that

- (a) $p(n) \leq \ell$,
- (b) $\sum_{j=1}^{p(n)} EL(s_0^{c(n)} \dots s_j^{c(n)})$ is bounded,
- (c) $\sum_{i=1}^{c(n)-1} |\mathcal{C}_i| \leq n = \sum_{i=1}^{c(n)-1} |\mathcal{C}_i| + p(n) \leq \sum_{i=1}^{c(n)-1} |\mathcal{C}_i| + \ell$.

Combining those three arguments, we obtain that

$$\limsup_{n \rightarrow \infty} \left[\frac{\sum_{i=1}^{c(n)-1} |\mathcal{C}_i| \cdot AE(\mathcal{C}_i)}{\sum_{i=1}^{c(n)-1} |\mathcal{C}_i| + \ell} \right] \leq \overline{AE}(\pi) \leq \limsup_{n \rightarrow \infty} \left[\frac{\sum_{i=1}^{c(n)-1} |\mathcal{C}_i| \cdot AE(\mathcal{C}_i)}{\sum_{i=1}^{c(n)-1} |\mathcal{C}_i|} \right]$$

Hence,

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

as claimed by Point 1.

Now consider Point 2. For any cycle $\mathcal{C} \in \text{Prefs}(G)$ such that $EL(\mathcal{C}) = 0$, all three hypotheses (i), (ii), and (iii) are clearly satisfied, with $\ell = |\mathcal{C}|$. Hence by Point 1, we have that

$$\overline{AE}(\mathcal{C}^\omega) = \limsup_{k \rightarrow \infty} \left[\frac{k \cdot |\mathcal{C}| \cdot AE(\mathcal{C})}{k \cdot |\mathcal{C}|} \right] = AE(\mathcal{C}).$$

Finally, we prove Point 3. The equality straightforwardly follows from Point 2. It remains to consider the inequality. By definition of the infimum, we have that, for all $k \geq 1$,

$$\inf_{i \in \mathbb{N}_{>0}} AE(\mathcal{C}_i) = \frac{\sum_{i=1}^k |\mathcal{C}_i| \cdot \inf_{i \in \mathbb{N}_{>0}} AE(\mathcal{C}_i)}{\sum_{i=1}^k |\mathcal{C}_i|} \leq \frac{\sum_{i=1}^k |\mathcal{C}_i| \cdot AE(\mathcal{C}_i)}{\sum_{i=1}^k |\mathcal{C}_i|}.$$

Hence by taking the limit, we obtain

$$\inf_{i \in \mathbb{N}_{>0}} AE(\mathcal{C}_i) = \limsup_{k \rightarrow \infty} \left[\inf_{i \in \mathbb{N}_{>0}} AE(\mathcal{C}_i) \right] \leq \limsup_{k \rightarrow \infty} \left[\frac{\sum_{i=1}^k |\mathcal{C}_i| \cdot AE(\mathcal{C}_i)}{\sum_{i=1}^k |\mathcal{C}_i|} \right] = \overline{AE}(\pi).$$

This concludes our proof. \square

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.

Memoryless determinacy. Intuitively, we use Lem. 4 and Lem. 5 to transform any arbitrary path into a simple *lasso path*, repeating a unique simple cycle, yielding an AE at least as good, thus proving that any threshold achievable with memory can also be achieved without it.

Theorem 6 *Memoryless strategies are sufficient to win one-player AE games.*

Proof As a preliminary step, we check whether the graph contains a reachable strictly negative cycle, e.g., using the Bellman-Ford algorithm in $\mathcal{O}(|S| \cdot |T|)$ -time. If so, then \mathcal{P}_1 can ensure a strictly negative mean-payoff, and by Point 1 of Lem. 3, a memoryless strategy exists to make the average-energy be $-\infty$: such a strategy consists in reaching and repeating the negative simple cycle forever.

Now, assume that the graph contains no (reachable) strictly negative cycle. If the graph also contains no zero cycle, then the energy level necessarily diverges to $+\infty$, and the average-energy is $+\infty$ along any run. Indeed, we are in the case of Point 2 of Lem. 3. Any strategy is optimal in that case: in particular, any memoryless strategy is.

For the rest of this proof, we consider the remaining case of graphs that contain no reachable strictly negative cycle, but that do contain zero cycles. We will prove that memoryless strategies suffice for \mathcal{P}_1 in those games, by induction on the number of choices of \mathcal{P}_1 . Given a game $G = (S_1, S_2 = \emptyset, E, w)$, we define $d_G = |E| - |S|$. Since we assume graphs to be deadlock-free, we have that $d_G \geq 0$ for any game G . We consider induction on the value d_G . For every game G such that $d_G = 0$ and initial state $s_{\text{init}} \in S$, \mathcal{P}_1 wins for the AE objective for threshold $t \in \mathbb{Q}$ iff he wins with a memoryless strategy: indeed, \mathcal{P}_1 actually has no choice at all in G , which is reduced to a unique outcome from s_{init} .

Now assume that memoryless strategies suffice for \mathcal{P}_1 in every game G such that $d_G \leq m$ for some $m \in \mathbb{N}$. We claim that they also suffice in every game G such that $d_G = m + 1$. Observe that if this holds, we are done as it proves that memoryless strategies suffice for \mathcal{P}_1 in all one-player AE games. Let G be such a game with $d_G = m + 1$. Recall that in $G = (S_1, S_2 = \emptyset, E, w)$ there is no strictly negative cycle by hypothesis. Let s be a state of G such that s has at least two outgoing edges. Such a state necessarily exists since $d_G \geq 1$. Consider a partition of the outgoing edges of s in two non-empty sets A, B such that $A \uplus B = \{(s_1, s_2) \in E \mid s_1 = s\}$. According to this partition, we can define in the natural way two sub-games $G_A = (S_1, S_2 = \emptyset, E \setminus B, w)$ and $G_B = (S_1, S_2 = \emptyset, E \setminus A, w)$ such that $d_{G_A} \leq m$ and $d_{G_B} \leq m$. By induction hypothesis, we know that memoryless strategies suffice to play optimally for the AE objective in those two sub-games. First, observe that if \mathcal{P}_1 has a memoryless winning strategy σ in either G_A or G_B for threshold $t \in \mathbb{Q}$, then this strategy remains winning in G . What we need to show is that if \mathcal{P}_1 cannot win in both G_A and G_B , then he also cannot win in G , even using memory in s : in the following, we assume that \mathcal{P}_1 is memoryless in any other state $s' \neq s$ (following the induction hypothesis) and we show that mixing cycles in s does not help him.

By contradiction, assume that \mathcal{P}_1 cannot win in both G_A and G_B , but he has a winning strategy σ in G , for the same threshold t . Let π be the outcome consistent with σ . Two cases are possible.

First, state s is seen *finitely often* along π . In this case, we apply Lemma 4 repeatedly on π to iteratively remove all cycles on s . Since there is no strictly negative cycle in G , we know that removing one cycle cannot increase the average-energy of the play (it either stays the same if the cycle is a zero cycle, or decreases if it is a strictly positive one). Since s is seen finitely often, we eventually obtain a play π' that sees s at most once. Therefore, this play either belongs to G_A or G_B (both if s is never visited). Furthermore, it has average-energy at most t by construction. This contradicts the claim that \mathcal{P}_1 has no winning strategy in both sub-games and concludes the proof in this case.

Second, state s is seen *infinitely often* along π . Since \mathcal{P}_1 is memoryless outside s , π only contains simple cycles and can be written as $\pi = \rho \cdot C_1 \cdot C_2 \cdot C_3 \cdots$ where ρ is an acyclic prefix ending in s and for all $i \geq 1$, C_i is a simple cycle on s . Observe that every cycle C_i belongs either to G_A or to G_B . Furthermore, since π is winning and there is no strictly negative cycle in G , only finitely many indices i_1, \dots, i_k may correspond to a strictly positive cycle. With the same reasoning as above (repeated application of Lemma 4), we have that the play $\pi' = \rho \cdot C_{i_k+1} \cdot C_{i_k+2} \cdots$, obtained by removing the first cycles up to index i_k , necessarily has a lower or equal average-energy: hence it is also winning. Now observe that the sequence of cycles $\pi'' = C_{i_k+1} \cdot C_{i_k+2} \cdots$ may still involve simple cycles from both G_A and G_B . Still, as all cycles are of length at most $|S|$, and are zero cycles, we can apply Lemma 5 to extract one best cycle C_j , $j > i_k$. Putting all this together, we have that $\pi''' = \rho \cdot (C_j)^\omega$ is such that $AE(\pi''') \leq AE(\pi)$. Furthermore, π''' is a simple lasso path that belongs either to G_A or to G_B (as it now uses a unique outgoing edge from s). Consequently, π''' describes a winning strategy in one of the sub-games, which contradicts our hypothesis and concludes our proof in this case too. \square

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

The crux of our algorithm consists in computing, for each state s , the best — w.r.t. the AE — zero cycle 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 $C_{s,k}$ 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 $C_{s,k}$ of length k on s such that (a) the sum of weights in the first dimension is zero (thus $C_{s,k}$ is a zero cycle), and (b) the sum in the second one is minimal. Fortunately, this sum is exactly equal to $AE(C) \cdot k$ thanks to the l factors used in the weights of the expanded graph. Hence, we obtain the optimal cycle $C_{s,k}$ (in polynomial time). Doing this $|S|$ times for each state s , we obtain for each of them the optimal cycle C_s (if one zero cycle exists). Then, by Lem. 4, 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 7 *The AE problem for one-player games is in P.*




Fig. 2: The best cycle $C_{s,2}$ 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'' .

Proof Let s_{init} be the initial state and $t \in \mathbb{Q}$ be the threshold. From Thm. 6, we can restrict our search to *memoryless* strategies achieving average-energy less than or equal to t . As noted in the proof of Thm. 6, if a strictly negative simple cycle exists and can be reached from s_{init} , then the answer to the AE problem is clearly Yes, as average-energy $-\infty$ is achievable. Checking if such a cycle exists and is reachable can be done in cubic time in the number of states (e.g., using Bellman-Ford to detect negative cycles).

Hence, we now assume that no negative cycle exists. The main part of our algorithm consists in computing, for each state s , the least average-energy that can be achieved along a simple *zero* cycle starting and ending in s (if any). Indeed, strictly positive cycles should be avoided as there is no negative cycle to counteract them. Applying Lem. 4, it then remains to compute the least energy level with which each state s can be reached (simple paths are sufficient as there are no negative cycles), and to pick the optimal combination. Again, this last part can be solved by using classical graph algorithms in cubic time in $|S|$.

We now focus on computing the best zero cycle from a state s . This is achieved by enumerating the possible lengths, from 1 to $|S|$ (*simple* cycles suffice). For a fixed length k , we consider a new graph $\mathcal{G}_{s,k}$, made of $k+1$ copies of the original game G . The states of $\mathcal{G}_{s,k}$ are pairs (u, l) with $u \in S$ and $0 \leq l \leq k$. The new graph is arranged in levels, indexed from $l = k$ for the top one to $l = 0$ for the bottom one: l represents the number of steps remaining to close the cycle of length k . For each edge (u, u') of G , with $w(u, u') = c$, and for each $1 \leq l \leq k$, except if both $u' = s$ and $l < k$ (in order to rule out intermediary visits to s), there is an edge from (u, l) to $(u', l-1)$. This edge carries a pair of weights $(c, l \cdot c)$. Our aim is to find a path in this graph from (s, k) to $(s, 0)$ (hence this is a simple cycle of length k) such that the sum of the weights on the first dimension is zero (hence this is a zero cycle) and the sum on the second dimension is minimized (when divided by k , this sum is precisely the average-energy, if starting from energy level zero).

This problem can be expressed as a linear program, with variables $x_{u,u',l}$ for each edge $u \rightarrow u'$ and each $1 \leq l \leq k$. While they are not required to take integer values, these variables are intended to represent the number of times the edge from (u, l) to $(u', l-1)$ is taken along a “path” in $\mathcal{G}_{s,k}$. The linear program is as follows:

minimize $\sum x_{u,u',l} \cdot l \cdot w(u,u')$ subject to

1. $0 \leq x_{u,u',l} \leq 1$ for all $x_{u,u',l}$;
2. for all (u,l) with $1 \leq l \leq k-1$, $\sum_{u'} x_{u',u,l+1} = \sum_{u'} x_{u,u',l}$;
3. $\sum_{u'} x_{s,u',k} = \sum_{u'} x_{u',s,1} = 1$;
4. $\sum x_{u,u',l} \cdot w(u,u') = 0$;
5. $\sum x_{u,u',l} \geq 1$.

Condition (2) states that each state has the same amount of “incoming” and “outgoing” flow. Condition (3) expresses the fact that we start and end up in state s . Condition (4) encodes the fact that we are looking for zero cycles, and Condition (5) rules out the (possible) trivial solution where all variables are zero.

First observe that if this LP has no solution, then there is no zero cycle of length k from s . Now, assume it has a solution $(x_{u,u',l}^0)$: this solution minimizes $\sum x_{u,u',l} \cdot l \cdot w(u,u')$. Consider a sequence of edges $s = u_k \rightarrow u_{k-1} \rightarrow \dots \rightarrow u_1 \rightarrow u_0 = s$ for which $x_{u_l, u_{l-1}, l} > 0$ for all l . The existence of such a sequence easily follows from Conditions (2) and (3). Assume that this is not a zero cycle. As there are no negative cycles, then this must be a positive cycle. But in order to fulfill Condition (4), we would need a negative cycle to compensate for this positive cycle, hence implying contradiction. We conclude that any sequence of consecutive edges as selected above is a zero cycle. Similarly, there cannot be a zero cycle of length k from s with better average-energy, as this would contradict the optimality of this solution. We thus have obtained an average-energy-optimal simple zero cycle of length k from s , in polynomial time. Indeed, the LP is polynomial in the size of $\mathcal{G}_{s,k}$, itself polynomial in the size of the original game: the expanded graph has its size bounded by $|S| \cdot (k+1)$ and all weights are bounded by $k \cdot W$ with $k \leq |S|$ and W the largest absolute weight in the original game.

As discussed above, this process can be repeated for each state s and each length k , $1 \leq k \leq |S|$, hence at most $|S|^2$ times. For each state, we select the best cycle among the $|S|$ possible ones (one for each length). Therefore, in polynomial time, we get a description of the best cycles w.r.t. the average-energy, for each $s \in S$. Clearly if no such cycle exists, then the answer to the *AE* problem is **No**, as all cycles are strictly positive and the average-energy of any play will be $+\infty$. If some exist, we can find an optimal strategy by picking the best combination between such a cycle from a state s and a corresponding prefix from s_{init} to s of minimal energy level. As presented before, this is achieved in polynomial time. Then the answer to the *AE* problem is **Yes** if and only if this optimal combination yields average-energy at most equal to t . This concludes our proof. \square

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, most 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 [23, Cor. 7], Gimbert and Zielonka establish that memoryless strategies also suffice in *two-player* games with the same payoff. Thanks to Thm. 6, this entails the next theorem.

Theorem 8 *Average-energy games are determined and both players have memoryless optimal strategies.*

Observe that this result is true for both variants of the average-energy payoff function, namely \overline{AE} and \underline{AE} . When both players play optimally, they can restrict themselves to memoryless strategies and both variants thus coincide as mentioned earlier.

Solving average-energy games. Finally, consider the complexity of deciding the winner in a two-player AE game. By Thm. 8, 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. 7). 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. Thus, we obtain the following result.

Theorem 9 *The AE problem for two-player games is in $\text{NP} \cap \text{coNP}$.*

We complete our study by proving that MP games can be encoded into AE ones in polynomial time. The former are known to be in $\text{NP} \cap \text{coNP}$ but whether they belong to P is a long-standing open question (e.g., [38, 26, 9, 13]). Hence, w.r.t. current knowledge, the $\text{NP} \cap \text{coNP}$ -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. 5 to establish the equivalence.

Theorem 10 *Mean-payoff games can be reduced to average-energy games in polynomial time.*

Proof Let $G = (S_1, S_2, E, w)$ be a game, and $t \in \mathbb{Q}$ be the threshold for the mean-payoff problem. From G , we build another game $G' = (S'_1, S'_2, E', w')$ such that

- $S'_1 = S_1 \cup E$ and $S'_2 = S_2$;
- E' contains two types of edges:
 - $(s, e) \in E'$ iff there exists s' such that $e = (s, s') \in E$. Then $w'(s, e) = 2 \cdot w(e)$.
 - $(e, s') \in E'$ for any $e = (s, s') \in E$. Then $w'(e, s') = -2 \cdot w(e)$.

We claim that \mathcal{P}_1 has a strategy ensuring objective $\text{MeanPayoff}(t)$ in G if and only if the answer for the AE problem in G' is **Yes** for the same threshold t . A similar construction is used in [5].

With a prefix $\rho = (s_i)_{i \leq n}$ in G , we can associate a prefix $\rho' = (s'_i)_{i \leq 2n}$ in G' as follows: for all $k \leq n$, $s'_{2k} = s_k$, and for all $k < n$, $s'_{2k+1} = (s_k, s_{k+1})$. The mean-payoff along ρ then equals the average energy along ρ' (assuming initial energy 0 for ρ'). Indeed, applying the same decomposition arguments as for Lem. 5 and by definition of the weight function w' , we have that

$$\begin{aligned} AE(\rho') &= \frac{1}{n} \sum_{i=0}^{n-1} \frac{2 \cdot w'(s_i, (s_i, s_{i+1})) + w'((s_i, s_{i+1}), s_{i+1})}{2} \\ &= \frac{1}{n} \sum_{i=0}^{n-1} \frac{4 \cdot w(s_i, s_{i+1}) - 2 \cdot w(s_i, s_{i+1})}{2} = \frac{1}{n} \sum_{i=0}^{n-1} w(s_i, s_{i+1}) = MP(\rho). \end{aligned}$$

Conversely, with a prefix $\rho' = (s'_i)_{i \leq 2n}$ in G' starting and ending in a state in $S_1 \cup S_2$, we can associate a prefix $\rho = (s_i)_{i \leq n}$ in G such that $s_k = s'_{2k}$ for all $k \leq n$. Again, assuming the initial energy is zero in ρ' , the average energy along ρ' equals the mean payoff along ρ .

Now, assume that \mathcal{P}_1 has a winning strategy σ in G from some state $s \in S_1 \cup S_2$, achieving mean-payoff less than or equal to t . Consider the strategy σ' for G' defined as $\sigma'(\rho') = \sigma(\rho)$ if ρ' ends in S_1 . If ρ' ends in a T -state of the form (s, s') , then we let $\sigma'(\rho') = s'$, which is the only possible outgoing edge. We see that the outcomes of σ' correspond to the outcomes of σ , so that, assuming that the initial energy level is zero, σ' enforces that the average-energy is below t for any infinite outcome. Conversely, given a strategy σ' for G' whose outcomes have average-energy below t , the strategy defined by $\sigma(\rho) = \sigma'(\rho')$ for all finite paths ρ in G secures a mean-payoff below t . Observe that the equivalence holds both between \overline{AE} and \overline{MP} , and between \underline{AE} and \underline{MP} . Indeed, we have seen that for both MP and AE games, memoryless strategies suffice and decision problems for both variants coincide. \square

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* $c_{\text{init}} := 0$. In this case, the range of acceptable energy levels is by definition constrained to the interval $[0, U]$. Our approach benefits from this: we solve the AE_{LU} problem by considering an AE problem (and subsequently, an MP problem) over an expanded game that explicitly accounts for the lower and upper bounds on the energy.

Formally, we want to decide if \mathcal{P}_1 can ensure a *sufficiently low AE* while keeping the EL within the allowed range.

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 $Energy_{LU}(U, c_{\text{init}} := 0) \cap AvgEnergy(t)$.*

Again, we present results for the supremum variant \overline{AE} but they also hold for the infimum one \underline{AE} .

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 $Energy_{LU}(U, c_{\text{init}} := 0)$ (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 alternating behavior is more intricate than for other classical conjunctions of objectives. Consider for example energy parity [12] or multi-dimensional energy games [15, 37]. It is usually necessary to use different cycles in such games: intuitively, one needs one “good” cycle for each dimension and one for the parity objective, and a winning strategy needs to alternate between those cycles.

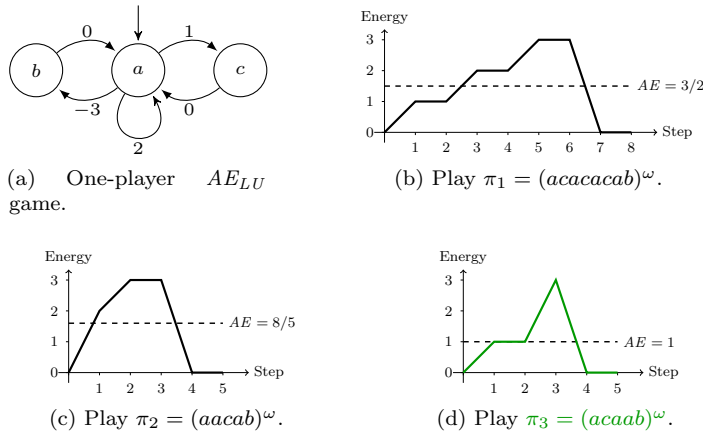


Fig. 3: Example of a one-player AE_{LU} game ($U = 3$) and the evolution of energy under different strategies that maintain it within $[0, 3]$ at all times. The minimal average-energy is obtained with **play π_3** : alternating in order between the $+1$, $+2$ and -3 cycles.

However, there is no need to use *two* different cycles that are “good” w.r.t. the same part of the objective. In the case of AE_{LU} games, we see that it is sometimes necessary to use two (or more) different cycles even though they impact the sum of weights in the same direction (e.g., several positive cycles). 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. 9 and Thm. 7, this reduction induces $\text{NEXPTIME} \cap \text{coNEXPTIME}$ -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. 13). We also improve the one-player case by observing that a witness lasso path in the MP game can be built on-the-fly, and the mean-payoff of this path can be computed using only polynomial space in the original game, hence we end up with PSPACE -membership which we also prove optimal in Thm. 13.

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 $\text{NP} \cap \text{coNP}$ for two-player games and to P for one-player games thanks to our reduction to an AE problem and the results of Thm. 9 and Thm. 7.

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 $\text{Energy}_{LU}(U, c_{\text{init}} := 0)$ 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.




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.

Lemma 11 *The AE_{LU} problem over a game $G = (S_1, S_2, E, w)$, with an initial state s_{init} , an upper bound $U \in \mathbb{N}$, and a threshold $t \in \mathbb{Q}$, is reducible to an AE problem for the same threshold $t \in \mathbb{Q}$ over a game $G' = (S'_1, S'_2, E', w')$ such that $|S'| = (U + 1) \cdot |S| + 1$ and $W' = \max\{\min\{W, U\}, 1\}$, i.e., the largest absolute weight in G' is at most the same as in G , or equal to constant 1.*

Proof Consider the game $G = (S_1, S_2, E, w)$, with initial state s_{init} , upper bound $U \in \mathbb{N}$ and threshold $t \in \mathbb{Q}$. We define the expanded game $G' = (S'_1, S'_2, E', w')$ as follows.

- $S'_1 = (S_1 \times \{0, 1, \dots, U\}) \cup \{\text{sink}\}$.
- $S'_2 = S_2 \times \{0, 1, \dots, U\}$.
- For all $(u, v) \in E$, $(u, c) \in S'$, we have that:
 1. if $d = c + w(u, v) \in [0, U]$, then $e = ((u, c), (v, d)) \in E'$ and $w'(e) = w(u, v)$,
 2. else $e = ((u, c), \text{sink}) \in E'$ and $w'(e) = 1$.
- $(\text{sink}, \text{sink}) \in E'$ and $w(\text{sink}, \text{sink}) = 1$.

The game G' starts in state $(s_{\text{init}}, 0)$ and edges are built naturally to reflect the changes in the energy level. Whenever the energy drops below zero or exceeds U , we redirect the edge to sink , where a self-loop of weight 1 is repeated forever.

We claim that \mathcal{P}_1 has a winning strategy σ_1 for the AE_{LU} objective in G if and only if he has a winning strategy σ'_1 for the AE objective in G' , for the very same average-energy threshold t .

First, consider the left-to-right implication. Assume σ_1 is winning for objective $Energy_{LU}(U, c_{\text{init}} := 0) \cap AvgEnergy(t)$ in G . The very same strategy can be followed in G' , ignoring the additional information on the energy in the state labels. Precisely, for any prefix $\rho' = (s_0, c_0)(s_1, c_1) \dots (s_n, c_n)$ in G' , we define $\sigma'_1(\rho') = (s, c)$ where $s = \sigma_1(\rho)$ for $\rho = s_0 s_1 \dots s_n$ and $c = c_n + w(s_n, s)$. Obviously, playing this strategy ensures that the special state sink is never reached, as otherwise it would not be winning for $Energy_{LU}(U, c_{\text{init}} := 0)$ in G , by construction of G' . Since all weights are identical in both games except on edges entering the sink state, we have that any consistent outcome π' of σ'_1 in G' corresponds to a consistent outcome π of σ_1 in G such that $\overline{AE}(\pi') = \overline{AE}(\pi)$, and conversely. Therefore, σ' is clearly winning for objective $AvgEnergy(t)$ in G' .

Second, consider the right-to-left implication. Assume σ'_1 is winning for objective $AvgEnergy(t)$ in G' . Then this strategy ensures that sink is avoided forever. Otherwise, there would exist a consistent outcome π' reaching sink, and such that $\overline{AE}(\pi') = \infty > t$ because of the strictly positive self-loop. Thus the strategy would not be winning. Hence by construction of G' , this strategy trivially ensures $Energy_{LU}(U, c_{\text{init}} := 0)$ in G' . From σ'_1 , we build a strategy σ_1 in G in the natural way, potentially integrating the information on the energy within the memory of σ_1 . Again, there is a bijection between plays avoiding sink in G' and plays in G , such that σ_1 is winning for $Energy_{LU}(U, c_{\text{init}} := 0) \cap AvgEnergy(t)$ in G .

Hence we have shown the claimed reduction. For the sake of completeness, observe that the reduction holds both for \overline{AE} and \underline{AE} variants of the average-energy. It remains to discuss the size of the expanded game. Observe that $|S'| = (U + 1) \cdot |S| + 1$. Furthermore, if W is the largest absolute weight in G , then $W' = \max\{\min\{W, U\}, 1\}$ is the largest one in G' . Indeed, W' is upper-bounded by U by construction (as all edges of absolute weight larger than U can be redirected directly to sink) and it is lower-bounded by 1 due to edges leading to sink. So the state space of G' is polynomial in the state space of G and in the *value* of the upper bound U , while its weights are bounded by either the largest weight W , the upper bound U or constant 1. \square

We now 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'' .

Lemma 12 *The AE problem over the game $G' = (S'_1, S'_2, E', w')$ defined in Lem. 11 is reducible to an MP problem for the same threshold $t \in \mathbb{Q}$ over a game $G'' = (S'_1, S'_2, E', w'')$ sharing the same state space but with largest absolute weight $W'' = \max\{U, \lceil t \rceil + 1\}$, where U is the energy upper bound of the original AE_{LU} problem.*

Proof Let $G' = (S'_1, S'_2, E', w')$ be the game defined in Lem. 11, as a reduction from the original game G for the AE_{LU} problem with upper bound $U \in \mathbb{N}$ and

average-energy threshold $t \in \mathbb{Q}$. We now build the game $G'' = (S'_1, S'_2, E', w'')$ by simply modifying the weight function of G' . The changes are as follows:

- For all edge $e = ((s, c), (s', c')) \in E'$, its weight in G' is $w'(e) = c' - c$ and we now set it to $w''(e) = c$ in G'' . Recall that by construction of G' , the value c represents the current energy level for any prefix ending in (s, c) . This is the value we now use for the outgoing edge. Also, this value is constrained in $[0, U]$ by definition of G' .
- For all edge $e = ((s, c), \text{sink}) \in E'$, its weight in G' is $w'(e) = 1$ and we now set it to $w''(e) = c$ in G'' for the sake of consistency (the actual value over this type of edges will not matter eventually).
- For the self-loop $e = (\text{sink}, \text{sink}) \in E'$, its weight in G' is $w'(e) = 1$ and we now set it to $w''(e) = \lceil t \rceil + 1$ in G'' . That is, reaching sink will imply a mean-payoff higher than the threshold.

Before proving the claim, we show that for all plays $\pi \in \text{Plays}(G') = \text{Plays}(G'')$ that *do not* reach sink, we have that $\overline{AE}_{G'}(\pi) = \overline{MP}_{G''}(\pi)$, where the subscript naturally refers to the change of weight function. Let

$$\pi = s'_0 s'_1 s'_2 \dots = (s_0, c_0)(s_1, c_1)(s_2, c_2) \dots$$

be such a play, where for all $i \geq 0$, $s'_i \in S'$ and $(s_i, c_i) \in S \times [0, U] \cap \mathbb{N}$ is its corresponding label. By definition of G'' , we have that,

$$\forall n \geq 0, \quad w''(s'_n, s'_{n+1}) = c_n = EL_{G'}(\pi(n)).$$

Hence by definition of the mean-payoff and the average-energy,

$$\begin{aligned} \overline{MP}_{G''}(\pi) &= \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} w''(s'_i, s'_{i+1}) \\ &= \limsup_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} EL_{G'}(\pi(i)) = \overline{AE}_{G'}(\pi). \end{aligned} \quad (4)$$

For the sake of completeness, observe that this equality does not hold for plays reaching sink, as they have infinite average-energy in G' but finite mean-payoff in G'' .

We proceed by proving the claim that \mathcal{P}_1 has a winning strategy σ'_1 for the *AE* objective in G' if and only if he has a winning strategy σ''_1 for the *MP* objective in G'' , for the very same threshold t .

First, consider the left-to-right implication. Assume σ'_1 is winning for objective *AvgEnergy*(t) in G' . We apply the same strategy in G'' straightforwardly as the underlying graph is not modified. Since this strategy is winning for the *AE* objective in G' , it necessarily avoids sink both in G' and G'' (as otherwise the *AE* would be infinite). Hence by Eq. (4), we have that σ'_1 is also winning for *MeanPayoff*(t) in G'' .

Second, consider the right-to-left implication. Assume σ''_1 is winning for objective *MeanPayoff*(t) in G'' . Since the self-loop on sink has weight $\lceil t \rceil + 1$, it is necessary that σ''_1 never reaches sink otherwise it would not be winning. Hence we apply the same strategy in G' and by Eq. (4), we have that σ''_1 is also winning for *AvgEnergy*(t) in G' .

This proves correctness of the reduction. The same reasoning can be followed for \underline{AE} (thus using \underline{MP}) instead of \overline{AE} . We end by discussing the size of G'' . Clearly, the state space S'' is identical to S' , hence $|S''| = (U + 1) \cdot |S| + 1$. However, the largest absolute weight in G'' is $W'' = \max\{U, \lceil t \rceil + 1\}$. Indeed, the self-loop on sink has weight $(\lceil t \rceil + 1)$ and all other edges have weight bounded by the energy upper bound U by construction. \square

Illustration. Consider the AE_{LU} game G depicted in Fig. 3a. We have seen that 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 $\overline{AE}_G(\pi_3) = \overline{AE}_{G'}(\pi') = \overline{MP}_{G''}(\pi')$.

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 [38,9], we get EXPTIME-membership, which is optimal for two-player games thanks to the lower bound proved for EG_{LU} [6]. The complexity is reduced when the bound U is given in unary or is polynomial in the size of the game, matching the one obtained for AE games without energy constraints.

For the one-player case, we also use the reduction to an MP game. By [17], optimal memoryless strategies exist, hence it suffices to non-deterministically build a simple lasso path in G'' , and to check that it satisfies the mean-payoff constraint. It can be done using only polynomial space through on-the-fly computation.

Theorem 13 *The AE_{LU} problem is EXPTIME-complete for two-player games and PSPACE-complete 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.*

Proof Let $G = (S_1, S_2, E, w)$ be the original AE_{LU} game, $W \in \mathbb{N}$ its largest absolute weight, $U \in \mathbb{N}$ the upper bound for energy and $t \in \mathbb{Q}$ the threshold for the AE_{LU} problem. By Lem. 11, this AE_{LU} problem is reducible to an AE problem for the same threshold t over a game $G' = (S'_1, S'_2, E', w')$ such that $|S'| = (U+1) \cdot |S| + 1$ and $W' = \max\{\min\{W, U\}, 1\}$. By Lem. 12, the AE_{LU} problem can be further reduced to an MP problem for the same threshold t over a game $G'' = (S'_1, S'_2, E', w'')$ sharing the same state space as G' but with largest absolute weight $W'' = \max\{U, \lceil t \rceil + 1\}$. We start by proving the complexity upper bounds.

First, consider the one-player case. Combining Thm. 7 and the reduction to an AE game, we obtain that one-player AE_{LU} games can be solved in pseudo-polynomial time, i.e., polynomial in $|S|$ but also in the value of U (hence exponential in the size of its binary encoding). This both gives EXPTIME-membership of one-player AE_{LU} games with arbitrary upper bounds, and P -membership of the same games with polynomial or unary upper bounds. For arbitrary bounds, we improve the complexity from EXPTIME to PSPACE. To do so, we consider the further reduction to an MP game, but we *do not* completely build the MP game G'' which is known to be of exponential size. Instead, we build non-deterministically a witness lasso path (thanks to memoryless determinacy [17], they are sufficient) and check on-the-fly that the path is winning or not, using only polynomial space. Recall that

we consider a game G'' such that $S'_2 = \emptyset$. Our non-deterministic algorithm answers YES if \mathcal{P}_1 has a winning strategy in G'' (and hence in G thanks to Lem. 11 and Lem. 12), No otherwise, and is as follows:

1. Guess a state $s'_r \in S'_1 = (S_1 \times \{0, 1, \dots, U\}) \cup \{\text{sink}\}$ that will be the starting (and ending) state of the cycling part of the lasso path. For the following, we assume that $s'_r \neq \text{sink}$ otherwise the lasso path that we are trying to build is clearly losing (see proof of Lem. 12) and the algorithm answers No. Thus, store state $s'_r = (s_r, m)$ for some $m \in \{0, \dots, U\}$.
2. Check that s'_r is reachable from the initial state $(s_{\text{init}}, 0)$. This can be done in NLOGSPACE w.r.t. the size of G'' (see e.g., [35]), hence NPSPACE w.r.t. the original problem. If it is not, then the answer is No.
3. Build step by step² a lasso path by constructing a simple cycle in G'' starting in s'_r . This construction is non-deterministic: if at any point, the sink state is reached, the algorithm returns No. The construction stops as soon as s'_r is reached, or after $|S'| + 1$ steps if s'_r is not reached: in the latter case, the answer is also No (after $|S'| + 1$ steps, we know for certain that a cycle was created hence our lasso path is complete). While constructing the cycle, we make on-the-fly computations: at each step, the next state is chosen non-deterministically and the only information that is stored — except from state s'_r used to determine the end of the cycle — is the number of steps from leaving s'_r , and the sum of the weights seen along the cycle.
4. Assume s'_r is reached (otherwise we have seen that the answer is No). Let $s'_0 s'_1 \dots s'_l$ be the sequence of states visited along the construction, with $s'_0 = s'_l = s'_r$. We have stored the length l and the sum of weights

$$\gamma = \sum_{i=0}^{l-1} w''(s'_i, s'_{i+1}).$$

Now, we check if $\frac{\gamma}{l} \leq t$: this quantity is the mean-payoff of the lasso path we have constructed. If yes, then the answer is YES, thanks to Lem. 11 and Lem. 12: the lasso path describes a winning strategy. Otherwise, the answer is No as this lasso path represents a losing strategy, by the same lemmas.

The correctness of this algorithm is guaranteed by Lem. 11 and Lem. 12. It remains to argue that it only uses polynomial space in the original AE_{LU} problem. Observe that our on-the-fly computations only need to record the state s'_r , the current state, the current length and the current sum. We have that both states belong to $S_1 \times \{0, 1, \dots, U\}$, that $l < |S'| = (U + 1) \cdot |S| + 1$ and that the sum is bounded by $l \cdot W'' = l \cdot \max\{U, \lceil t \rceil + 1\}$. Hence, encoding those values only requires a polynomial number of bits w.r.t. the input of the AE_{LU} problem (i.e., logarithmic in the upper bound U , the largest weight W and the threshold t). This proves that our algorithm lies in NPSPACE, and by Savitch's theorem [35] we know that NPSPACE = PSPACE: hence we proved the upper bound for the one-player AE_{LU} problem.

Second, consider two-player AE_{LU} games. In this case, we solve the MP problem over G'' using a pseudo-polynomial algorithm such as the one presented in [9],

² Observe that given a state in G'' , it is indeed possible to build any neighboring state using only E and w from the original game: one can effectively build the graph G'' on-the-fly.

whose complexity is $\mathcal{O}(|S^*|^3 \cdot W^*)$ for a game with $|S^*|$ states and largest absolute weight $W^* \in \mathbb{N}$. Therefore, the complexity of solving the original AE_{LU} problem is

$$\mathcal{O}(|S'|^3 \cdot W'') = \mathcal{O}\left(\left((U+1) \cdot |S| + 1\right)^3 \cdot \max\{U, \lceil t \rceil + 1\}\right),$$

which is clearly pseudo-polynomial. Hence we obtain EXPTIME-membership for two-player AE_{LU} games. If the upper bound $U \in \mathbb{N}$ is polynomial in the size of the game or encoded in unary, it is sufficient to solve the polynomially-larger AE game G' using Thm. 9 to obtain $\text{NP} \cap \text{coNP}$ -membership.

Now consider lower bounds. The AE_{LU} problem trivially encompasses the lower- and upper-bounded energy problem EG_{LU} , i.e., the AE_{LU} without consideration of the average-energy. Indeed, consider a game G with an objective $\text{Energy}_{LU}(U, c_{\text{init}} := 0)$, for some $U \in \mathbb{N}$. Assume \mathcal{P}_1 has a winning strategy for this objective. Then this strategy ensures that along any consistent outcome π , the running energy at any point is at most equal to U . By definition, this implies that $\underline{AE}(\pi) \leq \overline{AE}(\pi) \leq U$. Hence this strategy is also winning for the AE_{LU} objective written as the conjunction $\text{Energy}_{LU}(U, c_{\text{init}} := 0) \cap \text{AvgEnergy}(t := U)$. The converse is also trivially true. Ergo, any lower bound on the complexity of the EG_{LU} problem also holds for the AE_{LU} one. The EXPTIME-hardness of the two-player EG_{LU} problem was proved in [6], the PSPACE-hardness of the one-player version was proved in [18] (in the equivalent setting of reachability in bounded one-counter automata). Note that those results clearly rely on having an upper bound U larger than polynomial (w.r.t. the size of the game) and encoded in binary, as we have already shown that in the opposite case the complexity of the problem is reduced.

Finally, observe that the same reduction and complexities also hold if we use \underline{AE} instead of \overline{AE} to define the AE_{LU} problem. This concludes our proof. \square

Remark 1 One could argue that the reduction from AE games to MP games presented in Lem. 12 could be used to solve AE games without resorting to the specific analysis of Sect. 3. Indeed, in the case where the mean-payoff value is zero, any memoryless strategy (which we know to suffice) that is winning should only create zero cycles: the energy can be constrained in the range $[-2 \cdot |S| \cdot W, 2 \cdot |S| \cdot W]$ along any winning play. However, applying a pseudo-polynomial MP algorithm on this new game would only grant EXPTIME-membership for AE games (because of the polynomial dependency on W), in contrast to the $\text{NP} \cap \text{coNP}$ and P results obtained with the refined analysis for two-player and one-player AE 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. 8. Observe that winning strategies obtained via our reductions have a natural form: they are memoryless w.r.t. configurations (s, c) denoting the current state and the current energy level. As noted before, when the upper bound on energy $U \in \mathbb{N}$ is polynomial or given in unary, the expanded game is only polynomial in size, and the memory needs are also reduced.

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 (i.e., bounded energy only). Sufficiency of pseudo-polynomial memory for EG_{LU} games follows from [6] but to the best of our knowledge, it was not proved in the literature that such memory is also necessary.

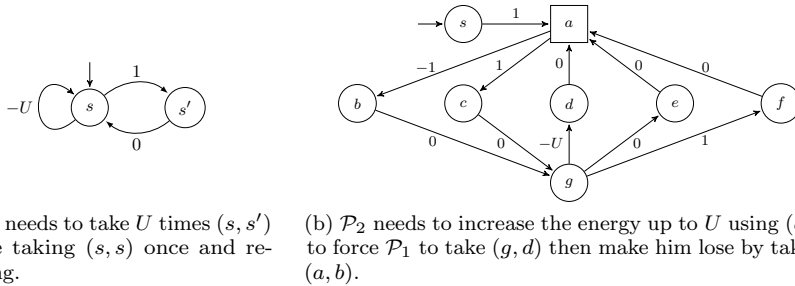


Fig. 5: Families of games witnessing the need for pseudo-polynomial-memory strategies for EG_{LU} (and AE_{LU}) objectives. The goal of \mathcal{P}_1 is to keep the energy in $[0, U]$ at all times, for $U \in \mathbb{N}$. The left game is won by \mathcal{P}_1 and the right one by \mathcal{P}_2 but both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

Theorem 14 *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.*

Proof We first prove the *upper bound on memory*. The expanded game G' built in the reduction from the AE_{LU} to the AE problem (Lem. 11) has a state space of size $|S'| = (U + 1) \cdot |S| + 1$, over which memoryless strategies suffice, by Thm. 8. Thus, winning for the AE_{LU} objective only requires memory that is polynomial in the original number of states and the upper bound value $U \in \mathbb{N}$. The same reduction holds for EG_{LU} games with an even simpler *safety* objective (never reaching sink) instead of the AE one (or equivalently with the AE objective for threshold $t = U$). Thus, with regard to the *binary encoding* of U , strategies require exponential memory in general. For the special cases of unary encoding or polynomially bounded value U , polynomial memory suffices. Note that as usual, these arguments are true for both the \overline{AE} and the \underline{AE} versions of the objective.

We now discuss the two families of games witnessing that pseudo-polynomial memory is also a *lower bound* for both players.

First, consider the one-player game depicted in Fig. 5a and parametrized by the value $U \in \mathbb{N}$. Assume the objective is EG_{LU} , asking for the energy to remain within $[0, U]$ at all times. Recall that the initial energy level is fixed to $c_{\text{init}} := 0$. It is easy to see that there is only one acceptable strategy for \mathcal{P}_1 : playing (s, s') exactly U times, then playing the self-loop (s, s) once, and repeating this forever. Indeed, any other strategy eventually leads the energy outside the allowed range. Hence, to win this game, \mathcal{P}_1 needs a strategy described by a Moore machine whose memory

contains at least $(U + 1)$ states. This proves that pseudo-polynomial memory is required for \mathcal{P}_1 in EG_{LU} games. Furthermore, the same argument can be applied on this game with objective AE_{LU} by considering the average-energy threshold $t := U$ which is trivially ensured by strategies satisfying the EG_{LU} objective.

Second, consider the two-player³ EG_{LU} game depicted in Fig. 5b. Again this game is parametrized by the energy upper bound $U \in \mathbb{N}$ and the initial energy level is fixed to $c_{\text{init}} := 0$. This game can be won by \mathcal{P}_2 using the following strategy: if the energy level is in $[1, U]$, play (a, c) , otherwise play (a, b) . Note that this strategy again requires at least $(U + 1)$ states of memory in its Moore machine (to keep track of the energy level).

This strategy is indeed winning. Observe that \mathcal{P}_1 can only decrease the energy by using edge (g, d) of weight $-U$, and this edge can only be used safely if the energy level is exactly U . In addition, the energy is bound to reach or exceed U eventually (as it will increase by 1 or 2 between each visit of a). If it exceeds U , then \mathcal{P}_2 wins directly. Otherwise, assume that the energy is U when the game is in state g . If \mathcal{P}_1 plays (g, f) , he loses (the energy reaches $U + 1$). If he plays (g, e) , \mathcal{P}_2 wins by playing (a, c) (the energy also reaches $U + 1$). And if \mathcal{P}_1 plays (g, d) , \mathcal{P}_2 wins by playing (a, b) (the energy reaches -1). Hence, \mathcal{P}_2 wins the game against all strategies of \mathcal{P}_1 .

Now, observe that \mathcal{P}_2 cannot win if he uses a strategy with less memory states in its Moore machine. Indeed, any such strategy cannot keep track of all the energy levels between 0 and U and play (a, c) a sufficient number of times in a row before switching to the appropriate choice (depending on the energy being 0 or U). Therefore, if \mathcal{P}_2 uses such a strategy, \mathcal{P}_1 can maintain the energy in the allowed range by simply reacting to edge (a, b) with (g, f) and to edge (a, c) by choosing between (g, d) (if the energy is U) and (g, e) (otherwise). Such choices are safe for \mathcal{P}_1 as the strategy of \mathcal{P}_2 does not have enough memory to distinguish the resulting energy levels from the intermediate ones.

This proves that \mathcal{P}_2 also needs pseudo-polynomial memory in EG_{LU} games. Finally, we remark that this reasoning also holds for the AE_{LU} objective with threshold $t := U$, as for the previous game. \square

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. As usual, we present our results for the supremum variant \overline{AE} .

³ In EG_{LU} games with only \mathcal{P}_2 (i.e., $S_1 = \emptyset$), \mathcal{P}_2 does not need memory to play as he can pick beforehand which of the energy bounds (lower or upper) he will transgress, and then do so with a memoryless strategy.

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 $\pi_4 = (aabaaba)^\omega$, with $\overline{AE}(\pi_4) = 11/7$ and $\pi_5 = (aaababa)^\omega$, 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. This indicates that AE_L games also require complex solutions.

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 that one-player AE_L problems lie in PSPACE by reduction to AE_{LU} problems for a well-chosen upper bound $U \in \mathbb{N}$ and then application of Thm. 13.

The reduction. Given a game $G = (S_1, S_2 = \emptyset, E, w)$ with largest weight $W \in \mathbb{N}$, 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 $U \in \mathbb{N}$ defined as $U := t + N^2 + N^3$, with $N = W \cdot (|S| + 2)$. Observe that the length of the binary encoding of U is polynomial in the size of the game, the encoding of the largest weight W and the encoding of the threshold t . 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 (stated in [31] and proved in [30, Lem. 42]). We slightly adapt it to our framework in the next lemma. The technique is identical, but the statement is more precise. In the following, we call an *expanded configuration* of the game G a couple (s, c) where $s \in S$ is a state and $c \in \mathbb{Z}$ a level of energy.

Lemma 15 *Let $g \in \mathbb{Z}$. Let (s, c) and (s', c') be two expanded configurations of the game G such that there exists an expanded path $\rho_{\text{exp}} = (s_0, c_0) \dots (s_m, c_m)$ in G from (s, c) to (s', c') with $c_i \geq g$ for every $0 \leq i \leq m$. Then, there is a path $\rho'_{\text{exp}} = (s'_0, c'_0)(s'_1, c'_1) \dots (s'_n, c'_n)$ in G from (s, c) to (s', c') such that:*

- for every $0 \leq i \leq n$, $g \leq c'_i \leq \max\{c, c', g\} + N^2 + N^3$, where $N = W \cdot (|S| + 2)$, with W the maximal absolute weight in G ;
- there is an (injective) increasing mapping $\iota: \{1, \dots, n\} \rightarrow \{1, \dots, m\}$ such that for every $1 \leq i \leq n$, $s'_i = s_{\iota(i)}$ and $c'_i \leq c_{\iota(i)}$.

Furthermore, for any two expanded paths ρ^1 and ρ^2 , with $\text{last}(\rho^1) = (s, c)$ and $\text{first}(\rho^2) = (s', c')$, if $AE(\rho^1 \cdot \rho_{\text{exp}} \cdot \rho^2) \leq g$, then it also holds that $AE(\rho^1 \cdot \rho'_{\text{exp}} \cdot \rho^2) \leq AE(\rho^1 \cdot \rho_{\text{exp}} \cdot \rho^2) \leq g$.

Proof We write $\alpha = W \cdot (|S| + 1)$, $\beta = (\alpha + W) \cdot (\alpha + W - 1) - 1$ and $K = \max\{c, c', g\} + (\alpha + W)^2$. We apply inductively a transformation that removes similar ascending and descending segments of the path. The segments are selected such that their composition is neutral w.r.t. the energy.

Pick a subpath $\rho_{\text{exp}}[k, k+h] = (s_k, c_k) \dots (s_{k+h}, c_{k+h})$ of ρ_{exp} , if it exists, such that:

- (a) $c_k \leq K$ and $c_{k+h} \leq K$;
- (b) for every $0 < \ell < h$, $c_{k+\ell} > K$;
- (c) there is $0 < \ell < h$ such that $c_{k+\ell} > K + W \cdot (|S| + 1) \cdot \beta$.

If such a subpath does not exist, then this means that the cost along ρ_{exp} is overall bounded by $K + W \cdot (|S| + 1) \cdot \beta$ (since condition (a) is not restrictive – $c, c' \leq K$), which then concludes the proof. Hence, assume such a subpath exists for the following steps.

Ascent part. Let $k \leq \ell_0 \leq \dots \leq \ell_\beta \leq k+h$ be indices such that:

- $c_{\ell_i} > K + i \cdot W \cdot (|S| + 1)$;
- for every $k \leq \ell < \ell_i$, $c_\ell \leq K + i \cdot W \cdot (|S| + 1)$.

Fix $0 \leq i \leq \beta$. Then it holds that $c_{\ell_i} \leq K + i \cdot W \cdot (|S| + 1) + W$ and thus $c_{\ell_{i+1}} - c_{\ell_i} > K + (i+1) \cdot W \cdot (|S| + 1) - (K + i \cdot W \cdot (|S| + 1) + W) = W \cdot (|S| + 1) - W = W \cdot |S|$. Let J_i be a subset of $[\ell_i; \ell_{i+1}]$ defined by $\ell_i \in J_i$, and if $j \in J_i$, then let $j' \leq \ell_{i+1}$ be the smallest index larger than j (if it exists) such that $c_{j'} > c_j$. Obviously we have $c_j < c_{j'} \leq c_j + W$. Hence the cardinal of J_i is at least $1 + \frac{W \cdot |S|}{W} \geq |S| + 1$. Hence there is a state $\tilde{s}^{(i)}$ and two indices $j_{i,1} < j_{i,2} \in J_i$ with $(s_{j_{i,1}}, c_{j_{i,1}}) = (\tilde{s}^{(i)}, \alpha_1)$ and $(s_{j_{i,2}}, c_{j_{i,2}}) = (\tilde{s}^{(i)}, \alpha_2)$ with $c_{\ell_i} \leq \alpha_1 < \alpha_2 \leq c_{\ell_{i+1}}$, hence using previous computed bounds, $0 < \alpha_2 - \alpha_1 \leq c_{\ell_{i+1}} - c_{\ell_i} < W \cdot (|S| + 2) = \alpha + W$. We write $\tilde{d}^{(i)} = \alpha_2 - \alpha_1$. The segment between indices $j_{i,1}$ and $j_{i,2}$ is a candidate for being removed. Due to the value of β , there is $d \in \{\tilde{d}^{(i)} \mid 0 \leq i \leq \beta\}$ that appears $(\alpha + W)$ times in that set.

Descent part. We do a similar reasoning for the “descent” part. There must exist indices $k \leq m_0 \leq \dots \leq m_\beta \leq k+h$ such that:

- $c_{m_i} > K + (\beta - i) \cdot W \cdot (|S| + 1)$;
- for every $m_i < m \leq k+h$, $c_m \leq K + (\beta - i) \cdot W \cdot (|S| + 1)$.

Note that we obviously have $\ell_\beta < m_0$.

Then we apply the same combinatorics as for the ascent part. There is some value $0 < d' < \alpha + W$ which appears at least $\alpha + W$ times in potential cycles within the segment $\rho_{\text{exp}}[k, k+h]$.

Transformation. The algorithm then proceeds by removing d' segments that increase the cost by d within $\rho_{\text{exp}}[\ell_0, \ell_\beta]$ and d segments that decrease the cost by d' within $\rho_{\text{exp}}[m_0, m_\beta]$. This yields another path ρ'_{exp} and an obvious injection of ρ'_{exp} into ρ_{exp} which satisfies all the mentioned constraints. The sum of all energy levels along ρ'_{exp} is smaller than that along ρ_{exp} , and any energy level along ρ'_{exp} is obtained from that along ρ_{exp} by decreasing by at most $0 < d \cdot d' < (\alpha + W)^2$.

By assumption on segment $\rho_{\text{exp}}[k, k+h]$ and bound K , we get that the cost along ρ'_{exp} is always larger than or equal to g, c and c' .

We iterate this transformation to get a uniform upper bound. We finally notice that the obtained upper bound $K + W \cdot (|S| + 1) \cdot \beta$ is bounded itself by $\max\{c, c', g\} + N^2 + N^3$, where $N = W \cdot (|S| + 2)$. This implies the expected result. \square

We build upon this lemma 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.

Lemma 16 *The AE_L problem over a one-player game $G = (S_1, S_2 = \emptyset, E, w)$, with an initial state s_{init} and a threshold $t \in \mathbb{Q}$, is reducible to an AE_{LU} problem over the same game G , for the same threshold t and upper bound $U := t + N^2 + N^3$, with $N = W \cdot (|S| + 2)$.*

Proof We prove that we can bound the energy along a witness of the one-player AE_L problem. Let σ be a winning strategy of \mathcal{P}_1 for objective $Energy_L(c_{\text{init}} := 0) \cap AvgEnergy(t)$ and $\pi = s_0 s_1 \dots s_n \dots$ be the corresponding outcome.

We build another strategy $\tilde{\sigma}$ with corresponding play $\tilde{\pi}$ such that for every n , $0 \leq c_{\text{init}} + EL(\tilde{\pi}(n)) \leq c_{\text{init}} + t + N^2 + N^3$, where $N = W \cdot (|S| + 2)$ (W is the maximal absolute weight in G), and such that $\overline{AE}(\tilde{\pi}) \leq \overline{AE}(\pi)$. We actually build the play $\tilde{\pi}$ directly, and infer strategy $\tilde{\sigma}$.

From π , we build the expanded play $\pi_{\text{exp}} = (s_0, c_0)(s_1, c_1) \dots (s_n, c_n) \dots$ such that $c_i = EL(\pi(i))$ for every $i \geq 0$. Since π is a witness satisfying the objective $Energy_L(c_{\text{init}}) \cap AvgEnergy(t)$, it holds that $c_i + c_{\text{init}} \geq 0$ for every $i \geq 0$. We now show that some pair (s, c) is visited infinitely often along π_{exp} . Toward a contradiction, assume that it is not the case. Then since energy levels are bounded from below along π , this means that $\liminf_{n \rightarrow \infty} c_n = \underline{TP}(\pi) = +\infty$, and by Lem. 2, that $\overline{AE}(\pi) = +\infty$ which contradicts the play being winning for the AE objective with threshold $t \in \mathbb{Q}$. Now select the smallest energy c and state s such that (s, c) is visited infinitely often along π_{exp} . Pick n_0 such that (1) $(s_{n_0}, c_{n_0}) = (s, c)$, (2) $\pi[\geq n_0] = s_{n_0} s_{n_0+1} \dots$ only visits states that are visited infinitely often along π , and (3) for every (s', c') along $\pi_{\text{exp}}[\geq n_0]$, it holds that $c' \geq c$.

We can then write π_{exp} as $\pi_{\text{exp}}[\leq n_0] \cdot \mathcal{C}_1 \cdot \mathcal{C}_2 \dots$ where each \mathcal{C}_i ends at configuration (s, c) (hence \mathcal{C}_i forms a cycle), and each configuration (s', c') along some \mathcal{C}_i satisfies $c' \geq c$. We write γ_i for the projection of \mathcal{C}_i on states (without energy level) — it forms a cycle as well. We obviously have

$$\overline{AE}(\pi) = EL(\pi(n_0)) + \overline{AE}(\pi[> n_0]) = c + \overline{AE}(\pi[> n_0])$$

by Lem. 4, and since $\overline{AE}(\pi) \leq t$, there must be some cycle \mathcal{C}_i such that $AE(\gamma_i) \leq t - c$. We write γ for such a γ_i , and we define $\varpi = \pi(n_0) \cdot \gamma^\omega$: it is a lasso-shaped play which also satisfies the objective $Energy_L(c_{\text{init}}) \cap AvgEnergy(t)$.

We will now modify the play ϖ , so that the energy does not grow too much along it. We write ϖ_{exp} for the expanded version of ϖ : it is of the form

$$\varpi_{\text{exp}}[\leq n_0] \cdot (\varpi_{\text{exp}}[n_0 + 1, n_0 + p])^\omega,$$

where $\varpi_{\text{exp}}[n_0 + 1, n_0 + p]$ projects onto γ when the energy information is removed (note that the last configurations of $\varpi_{\text{exp}}[\leq n_0]$ and of $\varpi_{\text{exp}}[n_0 + 1, n_0 + p]$ are

(s, c)). We will do two things: (i) first we will work on the cycle γ ; and (ii) then we will work on the prefix $\varpi[\leq n_0]$, to build a witness with a fixed upper bound on the energy. For the rest of the proof, we assume that $\varpi_{\text{exp}} = (s_0, c_0)(s_1, c_1) \dots$ so that $(s_n, c_n) = (s, c)$ for every $n = n_0 + b \cdot p$ for some integer b .

First consider point (i). Let us notice that $c \leq t$, otherwise the average-energy along ϖ could not be at most t (remember that the cost along the expanded version of γ starting at (s, c) is always larger than or equal to c by construction). We pick the first maximal subpath $\varpi_{\text{exp}}[k, k+h]$ of ϖ_{exp} with $[k, k+h] \subseteq (n_0, n_0+p)$, such that $c_{k+\ell} > t$ for every $0 \leq \ell \leq h$. By maximality of $\varpi_{\text{exp}}[k, k+h]$, it is the case that $c_{k-1} \leq t$ and $c_{k+h+1} \leq t$. We infer that $t < c_k \leq t+W$ and $t < c_{k+h} \leq t+W$, where W is the maximal absolute weight in the game G . We apply Lem. 15 to the path $\varpi_{\text{exp}}[k, k+h]$ with $g = t$, and we get that we can build an expanded path $\varpi_{\text{exp}}^{(k)}$ which is shorter than $\varpi_{\text{exp}}[k, k+h]$ and such that:

- at all positions of $\varpi_{\text{exp}}^{(k)}$, the energy is in the interval $[t, t + N^2 + N^3]$, where $N = W \cdot (|S| + 2)$;
- there is an injective increasing mapping $\iota: [0, |\varpi_{\text{exp}}^{(k)}|] \rightarrow [k, k+h]$ such that for every index $1 \leq i \leq |\varpi_{\text{exp}}^{(k)}|$, the state of $\varpi_{\text{exp}}^{(k)}[= i]$ coincides with that of $\varpi_{\text{exp}}[= \iota(i)]$ and the energy at position i of $\varpi_{\text{exp}}^{(k)}$ is smaller than or equal to $c_{\iota(i)}$.

In particular, we have a new witness for the objective $\text{Energy}_L(c_{\text{init}}) \cap \text{AvgEnergy}(t)$, which is the play $\varpi[< n_0] \cdot (\varpi[n_0, k-1] \cdot \varpi^{(k)} \cdot \varpi[k+h+1, n_0+|\gamma|-1])^\omega$, where $\varpi^{(k)}$ is the projection of $\varpi_{\text{exp}}^{(k)}$ over the states of the game G . We iterate this transformation over all relevant segments of γ (this will happen only a finite number of times), and we end up with a new lasso-play $\varpi' = \varpi[\leq n_0] \cdot (\gamma')^\omega$ such that:

- ϖ' satisfies the objective $\text{Energy}_L(c_{\text{init}}) \cap \text{AvgEnergy}(t)$;
- for every $1 \leq \ell \leq |\gamma'|$, $-c_{\text{init}} \leq \text{EL}(\varpi'(n_0 + \ell)) \leq t + N^2 + N^3$.

Now, consider point (ii). It remains to work on the prefix $\varpi[\leq n_0]$ (which is still a prefix of ϖ'). We apply Lem. 15 to the prefix $\varpi[\leq n_0]$ with $g = 0$, and we get an appropriately bounded witness.

Summing up, our construction proves that if there exists a winning play for $\text{Energy}_L(c_{\text{init}} := 0) \cap \text{AvgEnergy}(t)$ in the one-player game G , then there exists one for $\text{Energy}_{LU}(U, c_{\text{init}} := 0) \cap \text{AvgEnergy}(t)$, with $U := t + N^2 + N^3$. Since the converse implication is obvious (as the second objective is strictly stronger), this concludes the proof of the reduction to an AE_{LU} game. \square

Complexity. Plugging this bound U in the PSPACE algorithm for one-player AE_{LU} games (Thm. 13) implies PSPACE-membership for one-player AE_L games also. In terms of time complexity, we saw that this problem can thus be solved in pseudo-polynomial time. We prove that no truly-polynomial-time algorithm can be obtained unless $\text{P} = \text{NP}$ as the one-player AE_L problem is NP-hard. We show it by reduction from the subset-sum problem [20]: given a finite set of naturals $A = \{a_1, \dots, a_n\}$ and a target natural v , decide if there exists a subset $B \subseteq A$ such that $\sum_{a_i \in B} a_i = v$. The reduction is sketched in Fig. 6: a play corresponds to a choice of subset. In order to keep a positive energy level, \mathcal{P}_1 has to pick a subset that achieves a sum *at least* equal to v , but in order to satisfy the AE threshold, this sum must be *at most* v : hence \mathcal{P}_1 must be able to pick a subset whose sum is *exactly* the target v .




Fig. 6: Reduction from the subset-sum problem for target $v \in \mathbb{N}$ to a one-player AE_L problem for average-energy threshold $t := v$.

Theorem 17 *The AE_L problem is in PSPACE and at least NP-hard for one-player games.*

Proof First, consider the claim of PSPACE-membership. Let $G = (S_1, S_2 = \emptyset, E, w)$ be a game with initial state s_{init} . Consider the AE_L problem for a given average-energy threshold $t \in \mathbb{Q}$. By Lem. 16, this problem is reducible to the AE_{LU} problem with upper bound $U := t + N^2 + N^3$, with $N = W \cdot (|S| + 2)$. Hence, U is of order $\mathcal{O}(t + W^3 \cdot |S|^3)$, and its encoding is polynomial in the encoding of the original AE_L problem (including thresholds and weights, not only in the number of states of the original game!). Following the complexity analysis presented in Thm. 13, we thus conclude that the one-player AE_L problem is indeed in PSPACE. In terms of time, by using the MP reduction and the pseudo-polynomial algorithm, we have an algorithm for the one-player AE_L problem that takes time of order

$$\mathcal{O}\left(\left((U + 1) \cdot |S| + 1\right)^3 \cdot \max\{U, \lceil t \rceil + 1\}\right) = \mathcal{O}\left(\left(t + W^3 \cdot |S|^3\right)^4 \cdot |S|^3\right),$$

which is still pseudo-polynomial in the size of the original AE_L problem (i.e., polynomial in the number of states and in the values of the largest absolute weight and of the average-energy threshold).

Second, we prove that the one-player AE_L problem is NP-hard. Consider the subset-sum problem for the set $A = \{a_1, \dots, a_n\}$ such that for all $i \in \{1, \dots, n\}$, $a_i \in \mathbb{N}$, and target $v \in \mathbb{N}$. Deciding if there exists a subset $B \subseteq A$ such that $\sum_{a_i \in B} a_i = v$ is well-known to be NP-complete [20]. We reduce this problem to an AE_L problem over the game G depicted in Fig. 6. Observe that this game has polynomially as many states as the size of A , and that its largest absolute weight is equal to the maximum between the largest element of A and the target v . It is clear that there is a bijection between choices of subsets of A and plays in G . Let us fix threshold $t := v$ for the average-energy. Recall that Lem. 4 implies that the average-energy of any play is exactly its energy level at the first visit of end (because afterwards the zero self-loop is repeated forever). Hence, we have that

1. a play π in G is winning for $Energy_L(c_{\text{init}} := 0)$ if and only if the corresponding subset B is such that $\sum_{a_i \in B} a_i \geq v$;
2. a play π in G is winning for $AvgEnergy(t := v)$ if and only if the corresponding subset B is such that $\sum_{a_i \in B} a_i \leq v$.

Therefore, \mathcal{P}_1 has a winning strategy for the AE_L objective $Energy_L(c_{\text{init}} := 0) \cap AvgEnergy(t := v)$ in G if and only if there exists a subset B for which the sum of elements is exactly equal to the target v .

This proves the reduction from the subset-sum problem and the NP-hardness result. Observe two things. First, the hardness proof relies on having set elements and a target value that are not polynomial in the size of the input set A . Indeed, the subset-sum problem is solvable with a pseudo-polynomial algorithm, hence in P for polynomial values. Second, our reduction also holds for the \underline{AE} variant of the average-energy. \square

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. As shown in Thm. 14, this strategy requires pseudo-polynomial memory.

Theorem 18 *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

For the two-player AE_L problem, we only provide partial answers, as open questions remain. We first discuss decidability: we present an incremental algorithm that is correct but incomplete (Lem. 19) and we draw the outline of a potential approach to obtain completeness hence decidability. Then, we prove that the two-player AE_L problem is at least EXPTIME-hard (Lem. 20). Finally, we show that, in contrast to the one-player case, \mathcal{P}_2 also requires memory in two-player AE_L games (Lem. 21).

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).

Lemma 19 *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 complete algorithm is known to exist. For example, the existence of an initial energy level sufficient to win in multi-dimensional energy games can be decided [15] but practical implementations resort to an incremental scheme that is in practice incomplete because the theoretical bound granting completeness is too large to be tackled efficiently by software synthesis tools [4]. 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 algorithm of Lem. 19 more appealing in practice.

Nevertheless, we *conjecture* that the AE_L problem is decidable for two-player games and that, similarly to the one-player case, an upper bound on the energy can be obtained. Unfortunately, this claim is much more challenging to prove for two-player games. Clearly, the approach of Lem. 16 has to be generalized: while in one-player games we could pick a witness winning play and transform it, we now have to deal with *tree unfoldings* — describing sets of plays — because of the uncontrollable choices made by \mathcal{P}_2 .

A potentially promising approach is to define a notion close to the *self-covering trees* used in [15] for energy games. Roughly, take any winning strategy of \mathcal{P}_1 in a two-player AE_L game. Without further assumption, this strategy could be infinite-memory. It can be represented by its corresponding infinite tree unfolding where in nodes of \mathcal{P}_1 , a unique child is given by the strategy, and in nodes of \mathcal{P}_2 , all possible successors yield different branches. Every rooted branch of this tree is infinite and describes a winning play. Then, we would like to achieve the following steps.

1. Prove that all branches of this unfolding can be cut in such a way that the resulting finite tree describes a *finite-memory* strategy that is still winning for the AE_L objective.
2. Reduce the height of this finite tree by compressing parts of the branches: deleting embedded zero cycles seems to be a good candidate for the transformation to apply.
3. Derive an *upper bound* on the height of the compressed tree and, consequently, on the maximal energy level reached along any play consistent with the corresponding strategy.
4. Use this upper bound to reduce the AE_L problem to an AE_{LU} problem.

Sadly, some challenges appear on the technical side when trying to implement this approach, mainly for items 1 and 3. Intuitively, the additional difficulty (when compared to the approach developed in [15] and similar works) arises from the fact that describing what is a good cycle pattern for the AE_L objective is much more intricate than it is for a simple EG_L objective (in which case we simply look for zero cycles). This makes the precise definition of an appropriate transformation of branches, and the resulting tree height analysis, more tedious to achieve.

We also mention that the AE_L problem could be reduced, following a construction similar to the one given in Sect. 4.1, to a mean-payoff threshold problem over an infinite arena, where states of the expanded graph are arranged respectively to their energy level, ranging from zero to infinity, and where weights would also take values inside $\mathbb{N} \cup \{\infty\}$ (as they reflect the possible energy levels). To the best of our knowledge, it is not known if mean-payoff games over such particular structures are decidable. If so, an algorithm would have to fully exploit the peculiar form of those arenas, as it is for example known that general models such as pushdown games are undecidable for the mean-payoff [16].

Finally, one could envision to fill the gap between one-player and two-player AE_L games by using a general result similar to [23, Cor. 7]. Recall that we used it to derive memoryless determinacy in the two-player case from memoryless determinacy

of both one-player versions ($S_1 = \emptyset$ and $S_2 = \emptyset$). However, we here have that in one-player games, \mathcal{P}_1 requires pseudo-polynomial memory. Therefore, it is necessary to extend the result of Gimbert and Zielonka to finite-memory strategies: that is, to show that if we have a bound on memory valid in both one-player versions of a game, then this bound, or a derived one, is also valid in the two-player version. This is not known to be the case in general, and establishing it for a sufficiently general class of games seems challenging.

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 [27].

Our reduction is depicted in Fig. 7. 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. 4 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.




Fig. 7: 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 $t := 0$.

Lemma 20 *The AE_L problem is EXPTIME-hard for two-player games.*

Proof Given a countdown game $\mathcal{C} = (\mathcal{V}, \mathcal{E})$ and an initial configuration (v_{init}, c_0) , we build a game $G = (S_1, S_2, E, w)$ with initial state s_{init} such that \mathcal{P}_1 has a winning strategy in G for the AE_L objective for threshold $t := 0$ if and only if he has a

winning strategy in \mathcal{C} to reach a terminal configuration with counter value zero. The construction is depicted in Fig. 7. Formally, the game G is built as follows.

- $S_1 = \mathcal{V} \cup \{\text{start}, \text{stop}\}$.
- $S_2 = \{(v, d) \in \mathcal{V} \times \mathbb{N} \setminus \{0\} \mid \exists v' \in \mathcal{V}, (v, d, v') \in \mathcal{E}\}$.
- $s_{\text{init}} = \text{start}$.
- For each $(v, d, v') \in \mathcal{E}$, we have that $(v, (v, d)) \in E$ with $w(v, (v, d)) = -d$ and $((v, d), v') \in E$ with $w((v, d), v') = 0$.
- Additionally, $(\text{start}, v_{\text{init}}) \in E$ with $w(\text{start}, v_{\text{init}}) = c_0$, $(\text{stop}, \text{stop}) \in E$ with $w(\text{stop}, \text{stop}) = 0$ and for all $v \in \mathcal{V}$, $(v, \text{stop}) \in E$ with $w(v, \text{stop}) = 0$.

First, consider the left-to-right direction of the claim. Assume \mathcal{P}_1 has a winning strategy for the AE_L objective in G . As noted before, such a strategy necessarily reaches the energy level zero then switches to **stop** directly. Hence, applying this strategy in the countdown game ensures that the sum of durations will be exactly equal to c_0 (recall that we start our AE_L game by initializing the energy to c_0 then decrease it at every step by the duration chosen by \mathcal{P}_1). Thus, this strategy is winning in the countdown game \mathcal{C} .

Second, consider the right-to-left direction. Assume that \mathcal{P}_1 has a winning strategy in the countdown game \mathcal{C} . Playing this strategy in G ensures to reach a state $v \in S_1$ with energy level exactly equal to zero. Thus a winning strategy for the AE_L objective is to play the countdown strategy up to this point then to immediately take the edge (v, stop) . Indeed, any consistent outcome will satisfy the lower bound on energy (as the energy will never go below zero), and it will have an average-energy equal to $t = 0$ (because the energy level when reaching **stop** will be zero).

This shows both directions of the claim and concludes our proof. Observe that this reduction is also true if we consider the AE variant of the average-energy. \square

Memory requirements. We close our study of two-player AE_L games by discussing the memory needs. First note that we cannot provide upper bounds: if we had such bounds, we could derive a bound on the energy along any consistent play and reduce the AE_L problem to an AE_{LU} one as discussed before, hence proving its decidability. Second, we already know by Thm. 18 that pseudo-polynomial memory is necessary for \mathcal{P}_1 . Finally, we present a simple game (Fig. 8) where \mathcal{P}_2 needs to use memory in order to prevent \mathcal{P}_1 from winning.




Fig. 8: Simple two-player AE_L game witnessing the need for memory even for \mathcal{P}_2 .

Lemma 21 *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.*

Proof We only have to prove that \mathcal{P}_2 needs memory in the game of Fig. 8. Consider the AE_L objective for the average-energy threshold $t := 1$ on this game. Assume that \mathcal{P}_2 is restricted to *memoryless* strategies. Then, there are only two possible strategies for \mathcal{P}_2 . If \mathcal{P}_2 always takes the self-loop (s_2, s_2) , then the only consistent play is $s_1(s_2)^\omega$: it has AE equal to 1, and satisfies the lower bound constraint on energy, thus \mathcal{P}_1 wins. If \mathcal{P}_2 always takes (s_2, s_3) , then \mathcal{P}_1 can win by producing the following play: $s_1s_2(s_3s_2s_3)^\omega$. It also has AE equal to 1, and satisfies the energy constraint. Hence \mathcal{P}_2 cannot win this game with a memoryless strategy. Nonetheless, he has a winning strategy that uses memory. Let this strategy be the one that plays (s_2, s_3) once then chooses the self-loop (s_2, s_2) forever. When this strategy is used by \mathcal{P}_2 , \mathcal{P}_1 has to pick (s_3, s_2) in the first visit of s_3 otherwise he loses because the energy goes below zero. But if \mathcal{P}_1 picks this edge, the unique outcome becomes $s_1s_2s_3(s_2)^\omega$, whose average-energy is $2 > t$, hence also losing for \mathcal{P}_1 . Thus, the defined strategy is winning for \mathcal{P}_2 . \square

6 Conclusion

We presented a thorough study of the *average-energy* payoff. We showed that average-energy games belong to the same intriguing complexity class as mean-payoff, total-payoff and energy games and that they are similarly memoryless determined. We then solved average-energy games with lower- and upper-bounded energy: such a conjunction is motivated by previous case studies in the literature [10]. Lastly, we provided preliminary results for the case of average-energy with a lower bound but no upper bound on the energy. Following the publication of [7], Larsen et al. addressed a different problem in [32]: they proved that deciding if there exists a threshold $t \in \mathbb{Q}$ such that \mathcal{P}_1 can win a two-player game for objective $Energy_L(c_{\text{init}} := 0) \cap AvgEnergy(t)$ can be done in doubly-exponential time. This is indeed equivalent to deciding if there exists an upper-bound $U \in \mathbb{N}$ such that \mathcal{P}_1 can win for the objective $Energy_{LU}(U, c_{\text{init}} := 0)$, which is known to be in 2EXPTIME [25]. Unfortunately, this approach does not help in solving Problem 3, where the threshold $t \in \mathbb{Q}$ for the average-energy is part of the input: solving two-player AE_L games is still an open question.

We believe that the average-energy objective and its variations model relevant aspects of systems in practical applications as hinted by the aforementioned case study. Hence, we would like to extend our knowledge of this objective to more general models such as stochastic games, or games with multi-dimensional weights. Of course, the open questions regarding the AE_L objective are intriguing. Finally, we would like to implement our techniques in synthesis tools and assess their applicability through proper case studies.

References

1. B. Aminof and S. Rubin. First cycle games. In *Proc. of SR*, EPTCS 146, pages 83–90, 2014.
2. H. Björklund, S. Sandberg, and S. Vorobyov. Memoryless determinacy of parity and mean payoff games: A simple proof. *Theoretical Computer Science*, 310(1-3):365–378, 2004.
3. R. Bloem, K. Chatterjee, T.A. Henzinger, and B. Jobstmann. Better quality in synthesis through quantitative objectives. In *Proc. of CAV*, LNCS 5643, pages 140–156. Springer, 2009.

4. A. Bohy, V. Bruyère, E. Filiot, and J.-F. Raskin. Synthesis from LTL specifications with mean-payoff objectives. In *Proc. of TACAS*, LNCS 7795, pages 169–184. Springer, 2013.
5. E. Boros, K. Elbassioni, V. Gurvich, and K. Makino. Markov decision processes and stochastic games with total effective payoff. In *Proc. of STACS*, LIPIcs 30, pages 103–115. Schloss Dagstuhl - LZI, 2015.
6. P. Bouyer, U. Fahrenberg, K.G. Larsen, N. Markey, and J. Srba. Infinite runs in weighted timed automata with energy constraints. In *Proc. of FORMATS*, LNCS 5215, pages 33–47. Springer, 2008.
7. P. Bouyer, N. Markey, M. Randour, K.G. Larsen, and S. Laursen. Average-energy games. In *Proc. of GandALF*, EPTCS 193, pages 1–15, 2015.
8. T. Brázdil, D. Kláška, A. Kučera, and P. Novotný. Minimizing running costs in consumption systems. In *Proc. of CAV*, LNCS 8559, pages 457–472. Springer, 2014.
9. L. Brim, J. Chaloupka, L. Doyen, R. Gentilini, and J.-F. Raskin. Faster algorithms for mean-payoff games. *Formal Methods in System Design*, 38(2):97–118, 2011.
10. F. Cassez, J.J. Jensen, K.G. Larsen, J.-F. Raskin, and P.-A. Reynier. Automatic synthesis of robust and optimal controllers – an industrial case study. In *Proc. of HSCC*, LNCS 5469, pages 90–104. Springer, 2009.
11. A. Chakrabarti, L. de Alfaro, T.A. Henzinger, and M. Stoelinga. Resource interfaces. In *Proc. of EMSOFT*, LNCS 2855, pages 117–133. Springer, 2003.
12. K. Chatterjee and L. Doyen. Energy parity games. In *Proc. of ICALP*, LNCS 6199, pages 599–610. Springer, 2010.
13. K. Chatterjee, L. Doyen, M. Randour, and J.-F. Raskin. Looking at mean-payoff and total-payoff through windows. *Information and Computation*, 242:25–52, 2015.
14. K. Chatterjee and V.S. Prabhu. Quantitative timed simulation functions and refinement metrics for real-time systems. In *Proc. of HSCC*, pages 273–282. ACM, 2013.
15. K. Chatterjee, M. Randour, and J.-F. Raskin. Strategy synthesis for multi-dimensional quantitative objectives. *Acta Informatica*, 51(3-4):129–163, 2014.
16. K. Chatterjee and Y. Verner. Mean-payoff pushdown games. In *Proc. of LICS*, pages 195–204. IEEE, 2012.
17. A. Ehrenfeucht and J. Mycielski. Positional strategies for mean payoff games. *International Journal of Game Theory*, 8(2):109–113, 1979.
18. J. Fearnley and M. Jurdziński. Reachability in two-clock timed automata is PSPACE-complete. In *Proc. of ICALP*, LNCS 7966, pages 212–223. Springer, 2013.
19. J. Filar and K. Vrieze. *Competitive Markov decision processes*. Springer, 1997.
20. M.R. Garey and D.S. Johnson. *Computers and intractability: a guide to the Theory of NP-Completeness*. Freeman New York, 1979.
21. T. Gawlitza and H. Seidl. Games through nested fixpoints. In *Proc. of CAV*, LNCS 5643, pages 291–305. Springer, 2009.
22. H. Gimbert and W. Zielonka. When can you play positionnaly? In *Proc. of MFCS*, LNCS 3153, pages 686–697. Springer, 2004.
23. H. Gimbert and W. Zielonka. Games where you can play optimally without any memory. In *Proc. of CONCUR*, LNCS 3653, pages 428–442. Springer, 2005.
24. E. Grädel, W. Thomas, and T. Wilke, editors. *Automata, Logics, and Infinite Games: A Guide to Current Research*, LNCS 2500. Springer, 2002.
25. L. Juhl, K.G. Larsen, and J.-F. Raskin. Optimal bounds for multiweighted and parametrised energy games. In *Theories of Programming and Formal Methods*, LNCS 8051, pages 244–255. Springer, 2013.
26. M. Jurdziński. Deciding the winner in parity games is in $UP \cap co-UP$. *Information Processing Letters*, 68(3):119–124, 1998.
27. M. Jurdziński, J. Sproston, and F. Laroussinie. Model checking probabilistic timed automata with one or two clocks. *Logical Methods in Computer Science*, 4(3), 2008.
28. R.M. Karp. A characterization of the minimum cycle mean in a digraph. *Discrete Mathematics*, 23(3), 1978.
29. E. Kopczynski. Half-positional determinacy of infinite games. In *Proc. of ICALP*, LNCS 4052, pages 336–347. Springer, 2006.
30. P. Lafourcade, D. Lugiez, and R. Treinen. Intruder deduction for AC-like equational theories with homomorphisms. Research Report LSV-04-16, Laboratoire Spécification et Vérification, ENS Cachan, France, 2004.
31. P. Lafourcade, D. Lugiez, and R. Treinen. Intruder deduction for AC-like equational theories with homomorphisms. In *Proc. of RTA*, LNCS 3467, pages 308–322. Springer, 2005.

32. K.G. Larsen, S. Laursen, and M. Zimmermann. Limit your consumption! Finding bounds in average-energy games. In *Proc. of QAPL*, EPTCS, 2016.
33. M. Randour. 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, pages 731–738. Springer, 2013.
34. M. Randour. *Synthesis in Multi-Criteria Quantitative Games*. PhD thesis, Université de Mons, Belgium, 2014.
35. M. Sipser. *Introduction to the theory of computation*. PWS Publishing Company, 1997.
36. F. Thuijsman and O.J. Vrieze. The bad match; a total reward stochastic game. *OR Spektrum*, 9(2), 1987.
37. Y. Velner, K. Chatterjee, L. Doyen, T.A. Henzinger, A.M. Rabinovich, and J.-F. Raskin. The complexity of multi-mean-payoff and multi-energy games. *Information and Computation*, 241:177–196, 2015.
38. U. Zwick and M. Paterson. The complexity of mean payoff games on graphs. *Theoretical Computer Science*, 158(1-2):343–359, 1996.