Bounding Average-Energy Games

Patricia Bouyer[1] Piotr Hofman[1,2] Nicolas Markey[1,3]
Mickael Randour[4] Martin Zimmermann[5]

[1]LSV - CNRS & ENS Cachan [2]University of Warsaw
[3]IRISA - CNRS & INRIA & U. Rennes [4]ULB - Université libre de Bruxelles
[5]Saarland University

April 25, 2017 — *FoSSaCS 2017, Uppsala*

20th International Conference on Foundations of Software Science and Computation Structures

The talk in one slide

- Study of **average-energy** games: quantitative two-player games where the goal is to *minimize the average energy level in the long-run*.
- *AE* games studied in [BMR+16], also in conjunction with energy constraints: EG_L or EG_{LU} (lower bound only, or lower + upper bounds).

Goal of this work

Solving a problem left open in [BMR+16]: two-player games with conjunction of an *AE* constraint and an EG_L one, i.e., AE_L games.

- ▷ To solve them, we make a detour by mean-payoff games on **infinite arenas**.
- ▷ We also consider multi-dimensional extensions of AE games.

Advertisement

Full paper available on arXiv [BHM+16]: abs/1610.07858

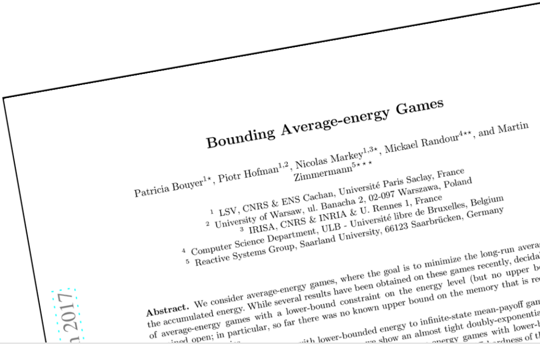

Bounding Average-energy Games

Patricia Bouyer[1⋆], Piotr Hofman[1,2], Nicolas Markey[1,3⋆], Mickael Randour[4⋆⋆], and Martin Zimmermann[5⋆⋆⋆]

[1] LSV, CNRS & ENS Cachan, Université Paris Saclay, France
[2] University of Warsaw, ul. Banacha 2, 02-097 Warszawa, Poland
[3] IRISA, CNRS & INRIA & U. Rennes 1, France
[4] Computer Science Department, ULB - Université libre de Bruxelles, Belgium
[5] Reactive Systems Group, Saarland University, 66123 Saarbrücken, Germany

Abstract. We consider average-energy games, where the goal is to minimize the long-run avera[ge of] the accumulated energy. While several results have been obtained on these games recently, decida[bility] of average-energy games with a lower-bound constraint on the energy level (but no upper b[ound]) ... open; in particular, so far there was no known upper bound on the memory that is re[quired] ...

1. Average-energy games
2. Average-energy games with lower-bounded energy
3. Multi-dimensional extensions
4. Conclusion

General context: strategy synthesis in quantitative games

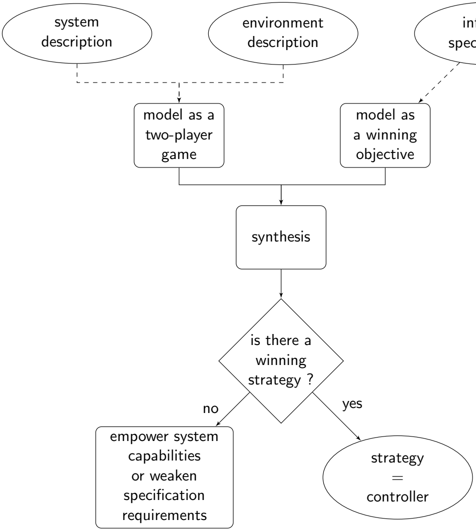

1. How complex is it to decide if a winning strategy exists?
2. How complex such a strategy needs to be? **Simpler is better**.
3. Can we synthesize one efficiently?

⇒ Depends on the *winning objective*.

Motivating example for average-energy

HYDAC oil pump industrial case study [CJL[+]09] (Quasimodo research project).

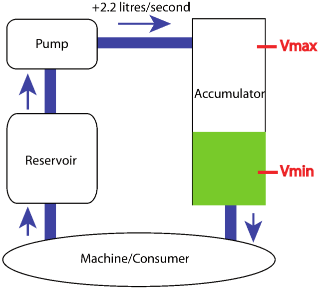

Goals:

1. Keep the oil level in the safe zone.
   ↪ **Energy objective with lower and upper bounds:** EG_{LU}
2. Minimize the average oil level.
   ↪ **Average-energy objective:** AE

⇒ **Conjunction:** AE_{LU}

Average-energy: illustration

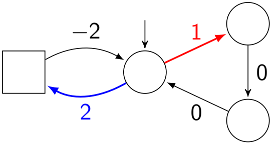

- Two-player turn-based games with integer weights.
- Focus on two *memoryless* strategies.

$\Longrightarrow$ We look at the **energy level** (EL) along a play.

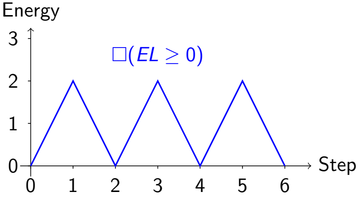

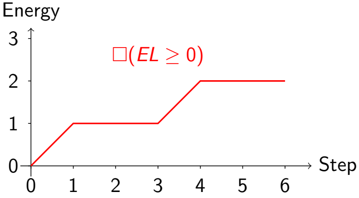

Energy objective (EG_L/EG_{LU}): e.g., always maintain $EL \geq 0$.

Average-energy: illustration

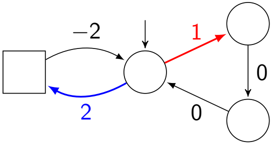

- Two-player turn-based games with integer weights.
- Focus on two *memoryless* strategies.

$\Longrightarrow$ We look at the **energy level** (*EL*) along a play.

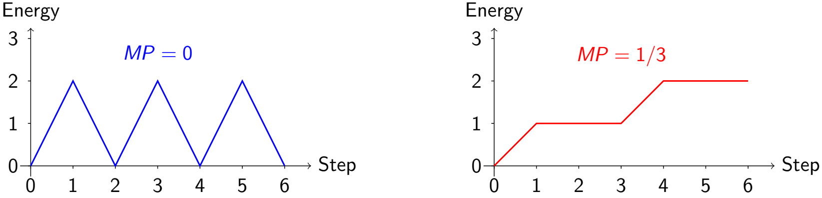

Mean-payoff (*MP*): long-run average payoff per transition.

Average-energy: illustration

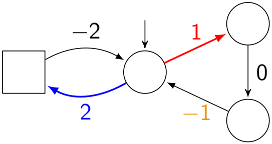

- Two-player turn-based games with integer weights.
- Focus on two *memoryless* strategies.

$\Longrightarrow$ We look at the **energy level** (*EL*) along a play.

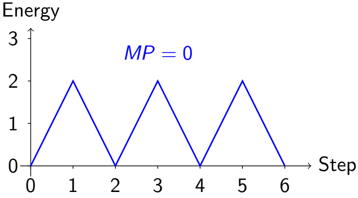

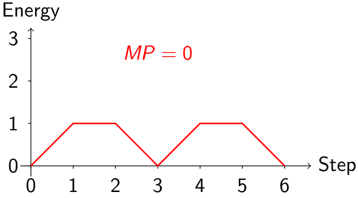

Mean-payoff (*MP*): long-run average payoff per transition.

$\Longrightarrow$ **Let's change the weights of our game.**

Average-energy: illustration

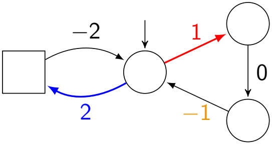

- Two-player turn-based games with integer weights.
- Focus on two *memoryless* strategies.

$\implies$ We look at the **energy level** (*EL*) along a play.

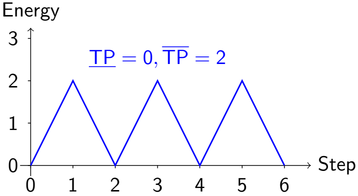

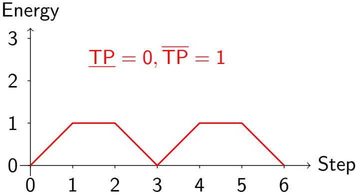

Total-payoff (TP) *refines MP* in the case $MP = 0$ by looking at high/low points of the sequence.

Average-energy: illustration

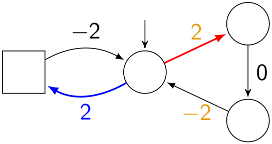

- Two-player turn-based games with integer weights.
- Focus on two *memoryless* strategies.

$\implies$ We look at the **energy level** (*EL*) along a play.

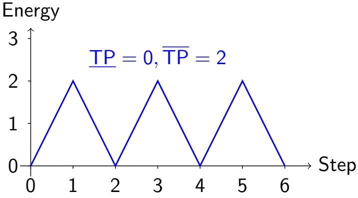

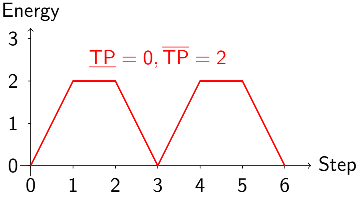

Total-payoff (TP**)** *refines MP* in the case $MP = 0$ by looking at high/low points of the sequence.

$\implies$ **Let's change the weights again.**

Average-energy: illustration

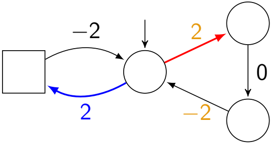

- Two-player turn-based games with integer weights.
- Focus on two *memoryless* strategies.

$\Longrightarrow$ We look at the **energy level** (*EL*) along a play.

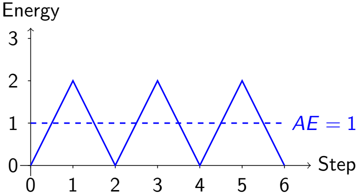

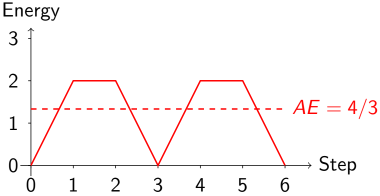

Average-energy (AE**)** *further refines* TP: average *EL* along a play.

$\Longrightarrow$ **Natural concept (cf. case study).**

Formal definitions

- We consider games $G = (S_0, S_1, E)$ between players P_0 and P_1, such that each edge $e \in E$ has an integer weight $w(e)$.
- For a prefix $\rho = (e_i)_{1 \leq i \leq n}$, we define
  - its *energy level* as $EL(\rho) = \sum_{i=1}^{n} w(e_i)$;
  - its *mean-payoff* as $MP(\rho) = \frac{1}{n} \sum_{i=1}^{n} w(e_i) = \frac{1}{n} EL(\rho)$;
  - its *average-energy* as $\mathsf{AE}(\rho) = \frac{1}{n} \sum_{i=1}^{n} EL(\rho_{\leq i})$.
- **Natural extensions to plays** by taking the upper-limit, e.g.,

$$\overline{\mathsf{AE}}(\pi) = \limsup_{n \to \infty} \frac{1}{n} \sum_{i=1}^{n} EL(\pi_{\leq i}).$$

Overview of known results

Objective	**1-player**	**2-player**	**memory**
MP	P [Kar78]	NP ∩ coNP [ZP96]	memoryless [EM79]
TP	P [FV97]	NP ∩ coNP [GS09]	memoryless [GZ04]
EG_L	P [BFL+08]	NP ∩ coNP [CdAHS03, BFL+08]	memoryless [CdAHS03]
EG_{LU}	PSPACE-c. [FJ15]	EXPTIME-c. [BFL+08]	exponential
AE	P	NP ∩ coNP	memoryless
AE_{LU}	PSPACE-c.	EXPTIME-c.	exponential
AE_L	PSPACE-e./NP-h.	*open*/EXPTIME-h.	*open* ($\geq$ exp.)

▷ Results without references are proved in [BMR+16].

▷ The one-player AE_L case is solved by reduction to an AE_{LU} game for a sufficiently large upper bound U, obtained through results on one-counter automata that permit to bound the counter value along a path.

$\Longrightarrow$ **Let's first recall how we solve AE_{LU} games.**

With energy constraints, memory is needed!

$AE_{LU} \rightsquigarrow$ minimize AE while keeping $EL \in [0, 3]$ (init. $EL = 0$).

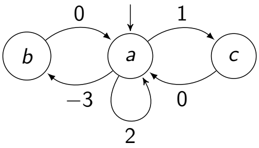

(a) One-player AE_{LU} game.

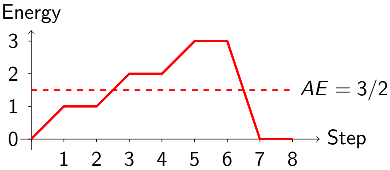

(b) Play $\pi_1 = (acacacab)^\omega$.

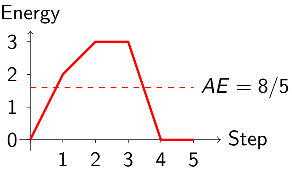

(c) Play $\pi_2 = (aacab)^\omega$.

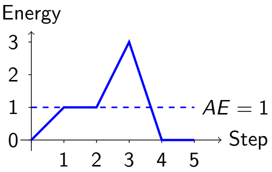

(d) Play $\pi_3 = (acaab)^\omega$.

Minimal AE with π_3: alternating between the $+1$, $+2$ and -3 cycles.

With energy constraints, memory is needed!

$AE_{LU} \rightsquigarrow$ minimize AE while keeping $EL \in [0, 3]$ (init. $EL = 0$).

Non-trivial behavior in general!
$\hookrightarrow$ **Need to choose carefully which cycles to play.**

The AE_{LU} problem is EXPTIME-complete.

$\hookrightarrow$ Reduction from AE_{LU} to AE on pseudo-polynomial game ($\implies AE_{LU} \in$ NEXPTIME $\cap$ coNEXPTIME).

$\hookrightarrow$ Reduction from this AE game to MP game + pseudo-poly. algorithm.

AE_{LU} problem: reduction to AE

↪ Expanded graph constraining the game within the energy bounds $[0, U]$. **Pseudo-polynomial size**: $\mathcal{O}(|S| \cdot (U+1))$.

↪ If we go out, AE $= \infty$.

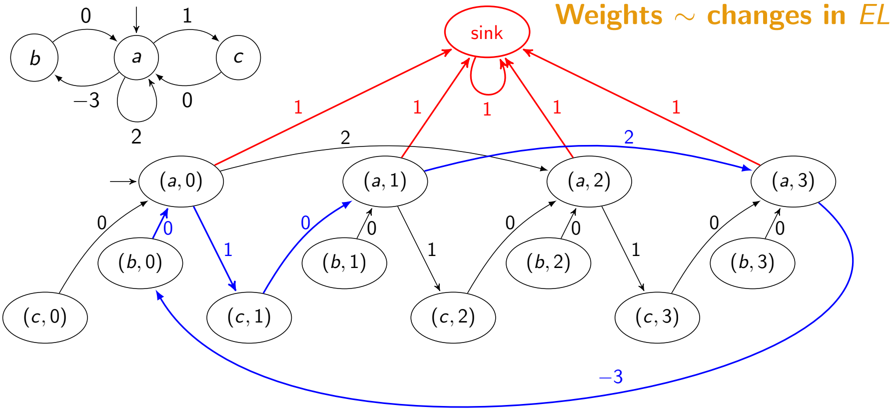

minimal AE $\wedge$ $EL \in [0, 3]$ in G $\iff$ minimal AE in G'

AE_{LU} problem: further reduction to MP

↪ Expanded graph of pseudo-poly. size: $\mathcal{O}(|S| \cdot (U+1))$.
Threshold for AE: $t = 1$.

↪ If we go out, $MP = \lceil t \rceil + 1 > t \Rightarrow$ losing.

Weights ∼ EL of prefix

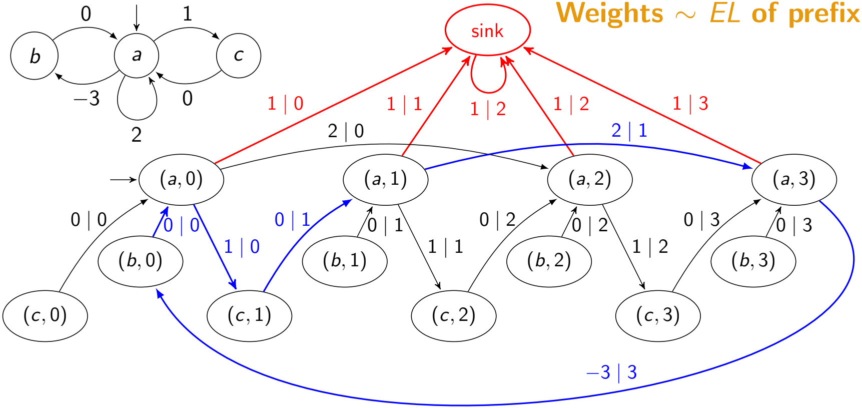

If $\neg(\Diamond \text{sink})$: $\overline{\text{AE}}(\pi)$ in $G' = \overline{\text{MP}}(\pi)$ in G''

1 Average-energy games

2 Average-energy games with lower-bounded energy

3 Multi-dimensional extensions

4 Conclusion

Tackling the two-player AE_L case

Aim of our approach

Obtain an energy upper bound U sufficient to reduce two-player AE_L games to AE_{LU} games.

- ▷ The approach used for one-player games does not suffice: we cannot modify plays directly because of P_1, the adversary.
- ▷ Defining an appropriate notion of *self-covering tree* (e.g., [CRR14]) and using it directly is difficult due to the complexity of the AE payoff (w.r.t. mean-payoff for example).

Idea

As in the AE_{LU} case, we will transform the AE_L game to an MP game on an expanded graph, with a similar construction.

$\Longrightarrow$ **Problem: this graph will be infinite!**

From an AE_L game to an infinite *MP* one

Given $G = (S_0, S_1, E)$, $s_{\text{init}} \in S$ and *AE* threshold $t \in \mathbb{Q}$, we define the *MP* game $G' = (\Gamma_0, \Gamma_1, \Delta)$:

- $\Gamma_0 = S_0 \times \mathbb{N}$ and $\Gamma_1 = S_1 \times \mathbb{N} \cup \{\bot\}$;
- Δ is given by:
  - $((s, c), c', (s', c')) \in \Delta$ if $\exists (s, w, s') \in E$ with $c' = c + w \geq 0$,
  - $((s, c), \lceil t \rceil + 1, \bot) \in \Delta$ if $\exists (s, w, s') \in E$ with $c + w < 0$,
  - $(\bot, \lceil t \rceil + 1, \bot) \in \Delta$.

$\Longrightarrow$ **Essentially the same construction as before, but with energy only bounded from below.**

Equivalence

P_0 has a winning strategy in G for AE_L with threshold t iff P_0 has a winning strategy in G' for *MP* with threshold t.

$\Longrightarrow$ **From now on, we consider the *MP* game.**

Solving the infinite *MP* game

So, it suffices to solve the *MP* game...

- ▷ Not much is known about *infinite* *MP* game.
- ▷ Our game has a *special structure*: its graph can be seen as the configuration graph of a one-counter pushdown system, where the stack height corresponds to the *EL* and the weight of an edge is given by the stack height of the target configuration.

$\Longrightarrow$ **Problem:** *MP* **games on pushdown systems with bounded weight functions are already undecidable [CV12], and our weight function is unbounded...**

$\Longrightarrow$ **We need to use the special structure!**

Sketch of our approach (1/2)

Goal

Prove that if a winning strategy exists, there exists one that wins while keeping the energy below a given bound U.

1. Along a winning play for MP, configurations below threshold t must be visited *frequently*.

   $\Longrightarrow$ **Proved through a density argument.**

2. Refining the analysis, we give an exponential (in the encoding) upper-bound on the length of the shortest *good cycle* along a winning play.

   Good cycle: $MP \leq t$ **and from a configuration below** t**.**

Sketch of our approach (2/2)

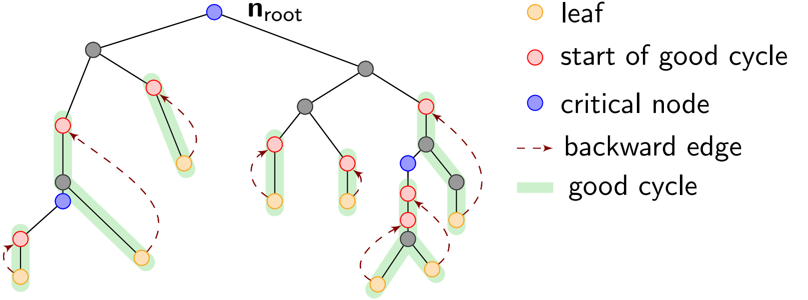

3. We define *finite good strategy trees*, which induce finite-memory winning strategies.
4. We prove that *any* winning strategy induces a finite good strategy tree.

⟹ **We need to bound the energy level in such a good strategy tree.**

Sketch of our approach (2/2)

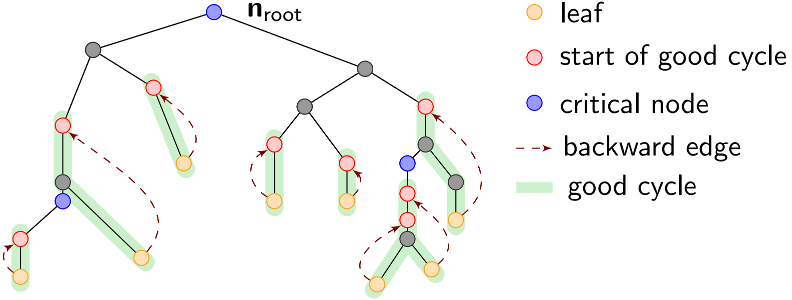

5. We build the strategy tree for a strategy σ by considering the shortest good cycles, hence the good cycles are already of bounded length (exponential) by Item 2.

$\Longrightarrow$ **We need to bound the remaining (i.e., gray) parts.**

Sketch of our approach (2/2)

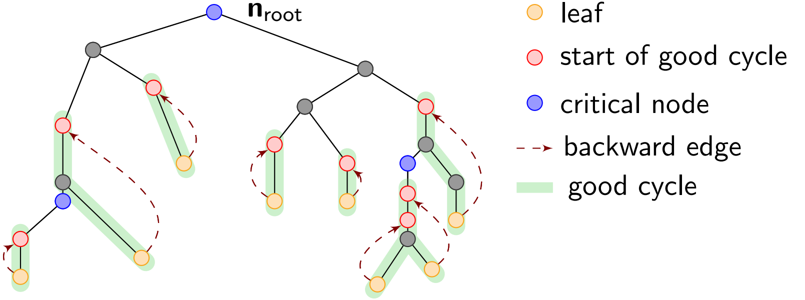

6 We consider reachability on our graph (a particular pushdown game) and show that we can bound the energy needed by strategies going from a critical node to the starting nodes of good cycles (by a double-exponential in the encoding).

$\Longrightarrow$ **We “replace” the strategy described by our tree in those gray parts by one with bounded energy.**

Sketch of our approach (2/2)

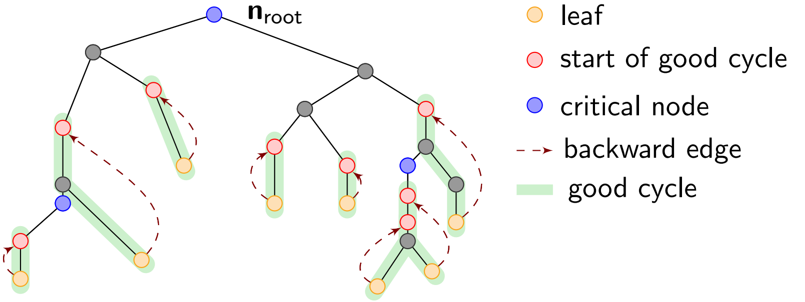

$\Longrightarrow$ Overall: we obtain that a doubly-exponential bound on the energy suffices to win the *MP* game.

$\Longrightarrow$ Applying the AE_{LU} reduction for this bound, we obtain 2-EXPTIME membership of AE_L games.

AE_L games: summary

Objective	**1-player**	**2-player**	**memory**
MP	P [Kar78]	NP ∩ coNP [ZP96]	memoryless [EM79]
TP	P [FV97]	NP ∩ coNP [GS09]	memoryless [GZ04]
EG_L	P [BFL+08]	NP ∩ coNP [CdAHS03, BFL+08]	memoryless [CdAHS03]
EG_{LU}	PSPACE-c. [FJ15]	EXPTIME-c. [BFL+08]	pseudo-polynomial
AE	P	NP ∩ coNP	memoryless
AE_{LU}	PSPACE-c.	EXPTIME-c.	exponential
AE_L	PSPACE-e./NP-h.	2-EXPTIME-e./EXPSPACE-h.	doubly-exp./super-exp.

- EXPTIME for unary encoding or polynomial weights and thresholds.
- Memory upper bound follows from our reduction, lower bound is by encoding of a succinct one-counter game [Hun14].
- EXPSPACE-hardness is also through reduction from succinct one-counter games [Hun15].

Multi-dimensional variants of *AE* games

We considered extensions to multiple dimensions (i.e., vectors of weights, bounds and thresholds) of three classes of games:

1. *AE* games (without energy bounds),
2. AE_{LU} games,
3. AE_L games.

$\Longrightarrow$ **We give a quick overview here.**

Multi-dimensional AE games

Reminder: one-dimensional version is in NP $\cap$ coNP and memoryless strategies suffice.

Undecidability

AE games with 3 or more dimensions are undecidable.

$\Longrightarrow$ We prove it via *two-dimensional robot games* [NPR16].

Robot game

$R = (\{q_0\}, \{q_1\}, T)$ where $T \subseteq Q \times [-V, V]^2 \times Q$ for some $V \in \mathbb{N}$, and q_i belongs to P_i. The game starts in q_0 with counter values $(x_0, y_0) \in \mathbb{Z}^2$ and P_0 tries to reach $(q_0, (0, 0))$.

Multi-dimensional AE_{LU} games

Reminder: one-dimensional version is EXPTIME-c. and exponential-memory strategies suffice.

Decidability

Multi-dim. AE_{LU} games are in NEXPTIME ∩ coNEXPTIME.

We generalize the construction seen before: *reduction to MP game over an expanded graph*. Two differences:

- graph is now exponential in the number of dimensions,
- multi-dim. *limsup MP* games are in NP ∩ coNP [VCD+15].

Multi-dimensional AE_L games

Reminder: one-dimensional version is in 2-EXPTIME and doubly-exponential-memory strategies suffice.

Undecidability

AE_L games with 2 or more dimensions are undecidable.

$\Longrightarrow$ We prove it via *two-counter machines*, with a proof similar to the one for total-payoff games [CDRR15].

Wrap-up

- We solved the open case from [BMR+16]: two-player AE_L games. We proved:
  - ▷ 2-EXPTIME membership,
  - ▷ EXPSPACE-hardness,
  - ▷ almost-tight memory bounds (doubly-exp. vs. super exp.).
- As a by-product, we solved a specific class of mean-payoff (one-counter) pushdown game with unbounded weight function.

  $\Longrightarrow$ Could be interesting to investigate if we can solve larger classes with similar techniques.
- In the multi-dimensional case, we proved that only AE_{LU} games remain decidable.

Thank you! Any question?

References I

Patricia Bouyer, Uli Fahrenberg, Kim G. Larsen, Nicolas Markey, and Jiří Srba.
Infinite runs in weighted timed automata with energy constraints.
In Franck Cassez and Claude Jard, editors, Proceedings of the 6th International Conferences on Formal Modelling and Analysis of Timed Systems, (FORMATS'08), volume 5215 of Lecture Notes in Computer Science, pages 33–47. Springer-Verlag, September 2008.

Patricia Bouyer, Piotr Hofman, Nicolas Markey, Mickael Randour, and Martin Zimmermann.
Bounding average-energy games.
CoRR, abs/1610.07858, 2016.

Patricia Bouyer, Piotr Hofman, Nicolas Markey, Mickael Randour, and Martin Zimmermann.
Bounding average-energy games.
In Javier Esparza and Andrzej Murawski, editors, Foundations of Software Science and Computation Structures - 20th International Conference, FOSSACS 2017, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2017, Uppsala, Sweden, April 22-29, 2017, Proceedings, Lecture Notes in Computer Science. Springer, 2017.
To appear.

Patricia Bouyer, Nicolas Markey, Mickael Randour, Kim G. Larsen, and Simon Laursen.
Average-energy games.
Acta Informatica, 2016.
Article in Press.

Arindam Chakrabarti, Luca de Alfaro, Thomas A. Henzinger, and Mariëlle Stoelinga.
Resource interfaces.
In Rajeev Alur and Insup Lee, editors, Embedded Software, Third International Conference, EMSOFT 2003, Philadelphia, PA, USA, October 13-15, 2003, Proceedings, volume 2855 of Lecture Notes in Computer Science, pages 117–133. Springer, 2003.

References II

Krishnendu Chatterjee, Laurent Doyen, Mickael Randour, and Jean-François Raskin.
Looking at mean-payoff and total-payoff through windows.
Information and Computation, 242:25–52, 2015.

Franck Cassez, Jan Jakob Jessen, Kim Guldstrand Larsen, Jean-François Raskin, and Pierre-Alain Reynier.
Automatic synthesis of robust and optimal controllers - an industrial case study.
In Rupak Majumdar and Paulo Tabuada, editors, Hybrid Systems: Computation and Control, 12th International Conference, HSCC 2009, San Francisco, CA, USA, April 13-15, 2009. Proceedings, volume 5469 of Lecture Notes in Computer Science, pages 90–104. Springer, 2009.

Krishnendu Chatterjee, Mickael Randour, and Jean-François Raskin.
Strategy synthesis for multi-dimensional quantitative objectives.
Acta Informatica, 51(3-4):129–163, 2014.

Krishnendu Chatterjee and Yaron Velner.
Mean-payoff pushdown games.
In Proceedings of the 27th Annual IEEE Symposium on Logic in Computer Science, LICS 2012, Dubrovnik, Croatia, June 25-28, 2012, pages 195–204. IEEE Computer Society, 2012.

Andrzej Ehrenfeucht and Jan Mycielski.
Positional strategies for mean payoff games.
International Journal of Game Theory, 8(2):109–113, 1979.

John Fearnley and Marcin Jurdzinski.
Reachability in two-clock timed automata is PSPACE-complete.
Information and Computation, 243:26–36, 2015.

References III

Jerzy Filar and Koos Vrieze.
Competitive Markov decision processes.
Springer-Verlag, 1997.

Thomas Gawlitza and Helmut Seidl.
Games through nested fixpoints.
In Ahmed Bouajjani and Oded Maler, editors, Proceedings of the 21st International Conference on Computer Aided Verification (CAV'09), volume 5643 of Lecture Notes in Computer Science, pages 291–305. Springer-Verlag, June 2009.

Hugo Gimbert and Wiesław Zielonka.
When can you play positionnaly?
In Jiří Fiala, Václav Koubek, and Jan Kratochíl, editors, Proceedings of the 29th International Symposium on Mathematical Foundations of Computer Science (MFCS'04), volume 3153 of Lecture Notes in Computer Science, pages 686–697. Springer-Verlag, August 2004.

Paul Hunter.
Reachability in succinct one-counter games.
arXiv, 1407.1996, 2014.

Paul Hunter.
Reachability in succinct one-counter games.
In Mikolaj Bojanczyk, Slawomir Lasota, and Igor Potapov, editors, Reachability Problems - 9th International Workshop, RP 2015, Warsaw, Poland, September 21-23, 2015, Proceedings, volume 9328 of Lecture Notes in Computer Science, pages 37–49. Springer, 2015.

References IV

Richard M. Karp.
A characterization of the minimum cycle mean in a digraph.
Discrete Mathematics, 23(3), 1978.

Reino Niskanen, Igor Potapov, and Julien Reichert.
Undecidability of two-dimensional robot games.
In Piotr Faliszewski, Anca Muscholl, and Rolf Niedermeier, editors, 41st International Symposium on Mathematical Foundations of Computer Science, MFCS 2016, August 22-26, 2016 - Kraków, Poland, volume 58 of LIPIcs, pages 73:1–73:13. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.

Yaron Velner, Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger, Alexander Rabinovich, and Jean-François Raskin.
The complexity of multi-mean-payoff and multi-energy games.
Information and Computation, 241:177–196, 2015.

Uri Zwick and Mike Paterson.
The complexity of mean payoff games on graphs.
Theoretical Computer Science, 158(1-2):343–359, 1996.