Average-Energy Games

Patricia Bouyer-Decitre[1] Kim G. Larsen[2] Simon Laursen[2]
Nicolas Markey[1] Mickael Randour[1]

[1]LSV - CNRS & ENS Cachan [2]Aalborg University

12.03.2015 - *ERC Workshop* - IST Austria

AALBORG UNIVERSITY
DENMARK

General context: strategy synthesis in quantitative games

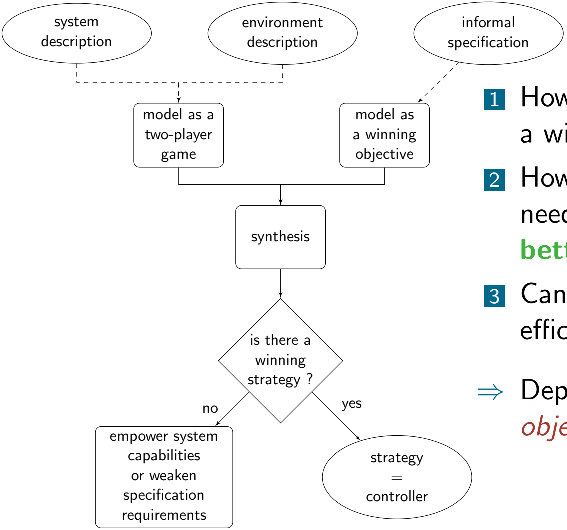

1. How complex is it to decide if a winning strategy exists?
2. How complex such a strategy needs to be? **Simpler is better**.
3. Can we synthesize one efficiently?

⇒ Depends on the *winning objective*.

The talk in one slide

- **New quantitative objective**
  - ▷ Total-payoff (TP) "refines" mean-payoff (MP) (MP value = 0)
  - ▷ **Average-energy (AE)** "refines" TP

The talk in one slide

- **New quantitative objective**
  - ▷ Total-payoff (TP) "refines" mean-payoff (MP) (MP value = 0)
  - ▷ **Average-energy (AE)** "refines" TP
  - ↪ characterizes the **average energy level** along an infinite play

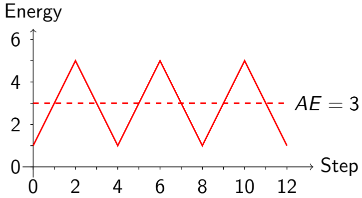

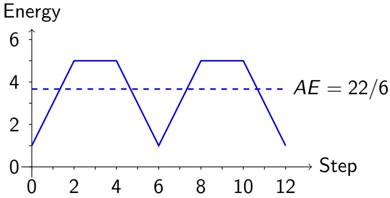

The talk in one slide

- **New quantitative objective**
  - ▷ Total-payoff (TP) "refines" mean-payoff (MP) (MP value = 0)
  - ▷ **Average-energy (AE)** "refines" TP
  - ↪ characterizes the **average energy level** along an infinite play

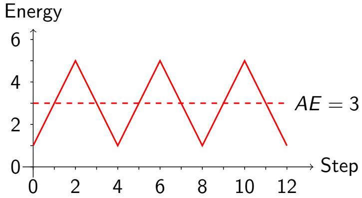

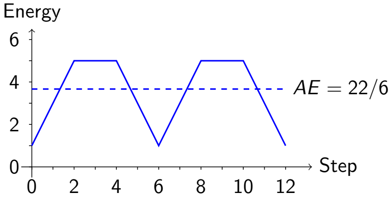

- Conjunction with **energy constraints**: lower and/or upper bounds on the energy level (e.g., fuel tank)

The talk in one slide

- **New quantitative objective**
  - ▷ Total-payoff (TP) “refines” mean-payoff (MP) (MP value = 0)
  - ▷ **Average-energy (AE)** “refines” TP
  - ↪ characterizes the **average energy level** along an infinite play

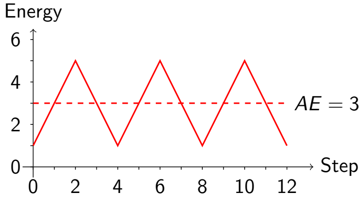

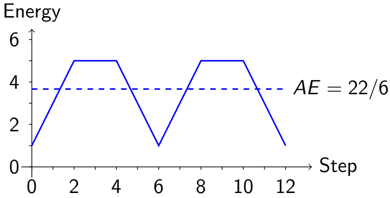

- Conjunction with **energy constraints**: lower and/or upper bounds on the energy level (e.g., fuel tank)

Ongoing work!

1 Context & Definitions

2 Average-Energy Games

3 Average-Energy with Energy Constraints

4 Conclusion

Two-player turn-based games on graphs

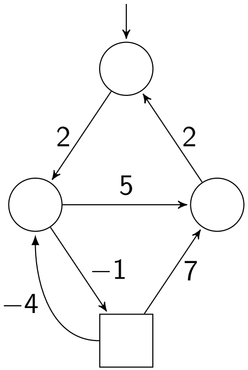

- $G = (S_1, S_2, T, w)$
- $S = S_1 \cup S_2$, $S_1 \cap S_2 = \emptyset$, $T \subseteq S \times S$, $w\colon T \to \mathbb{Z}$
- $\mathcal{P}_1$ states = ○
- $\mathcal{P}_2$ states = □
- Plays have values
  - ▷ $f : \mathit{Plays}(G) \to \mathbb{R} \cup \{-\infty, \infty\}$
- Players follow **pure** strategies
  - ▷ $\sigma_i\colon \mathit{Prefs}_i(G) \to S$

Energy, total-payoff, mean-payoff

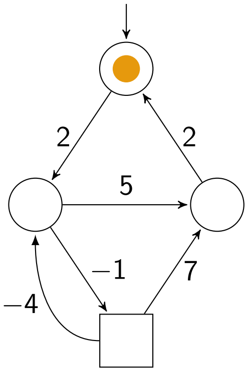

- $EL(\pi(n)) = \sum_{i=0}^{n-1} w(s_i, s_{i+1})$
- $\overline{TP}(\pi) = \limsup_{n\to\infty} EL(\pi(n))$
- $\overline{MP}(\pi) = \limsup_{n\to\infty} \frac{1}{n} EL(\pi(n))$

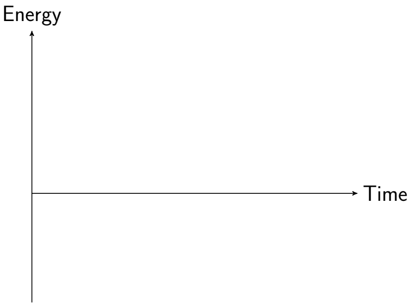

Energy, total-payoff, mean-payoff

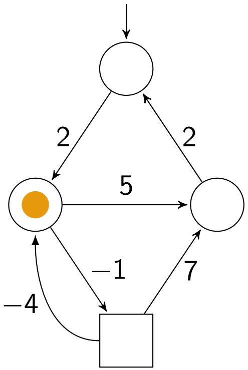

- $EL(\pi(n)) = \sum_{i=0}^{n-1} w(s_i, s_{i+1})$
- $\overline{TP}(\pi) = \limsup_{n\to\infty} EL(\pi(n))$
- $\overline{MP}(\pi) = \limsup_{n\to\infty} \frac{1}{n} EL(\pi(n))$

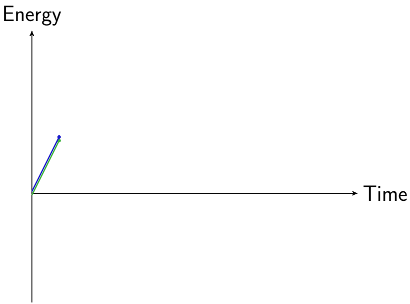

Energy, total-payoff, mean-payoff

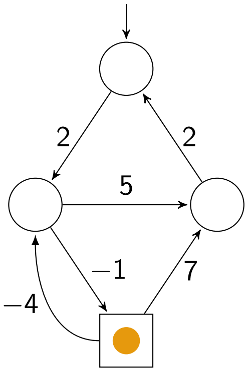

- $EL(\pi(n)) = \sum_{i=0}^{n-1} w(s_i, s_{i+1})$
- $\overline{TP}(\pi) = \limsup_{n\to\infty} EL(\pi(n))$
- $\overline{MP}(\pi) = \limsup_{n\to\infty} \frac{1}{n} EL(\pi(n))$

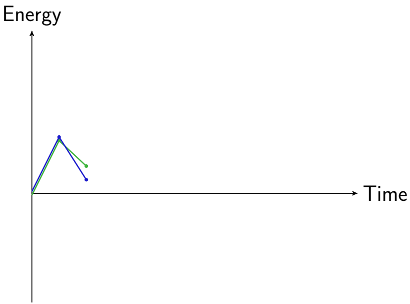

Energy, total-payoff, mean-payoff

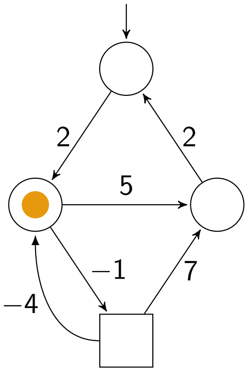

- $EL(\pi(n)) = \sum_{i=0}^{n-1} w(s_i, s_{i+1})$
- $\overline{TP}(\pi) = \limsup_{n\to\infty} EL(\pi(n))$
- $\overline{MP}(\pi) = \limsup_{n\to\infty} \frac{1}{n} EL(\pi(n))$

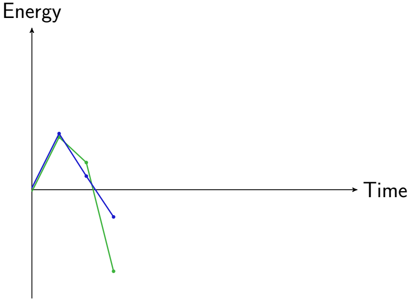

Energy, total-payoff, mean-payoff

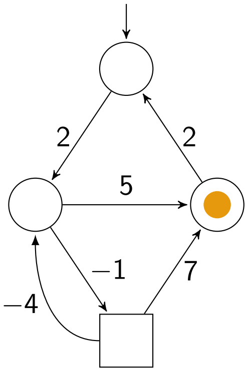

- $EL(\pi(n)) = \sum_{i=0}^{n-1} w(s_i, s_{i+1})$
- $\overline{TP}(\pi) = \limsup_{n\to\infty} EL(\pi(n))$
- $\overline{MP}(\pi) = \limsup_{n\to\infty} \frac{1}{n} EL(\pi(n))$

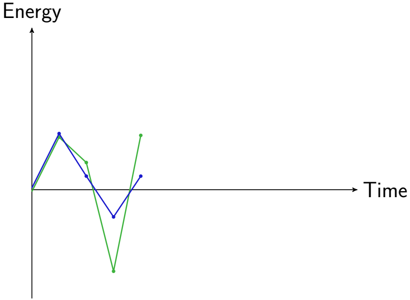

Energy, total-payoff, mean-payoff

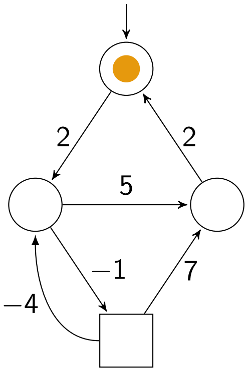

- $EL(\pi(n)) = \sum_{i=0}^{n-1} w(s_i, s_{i+1})$
- $\overline{TP}(\pi) = \limsup_{n\to\infty} EL(\pi(n))$
- $\overline{MP}(\pi) = \limsup_{n\to\infty} \frac{1}{n} EL(\pi(n))$

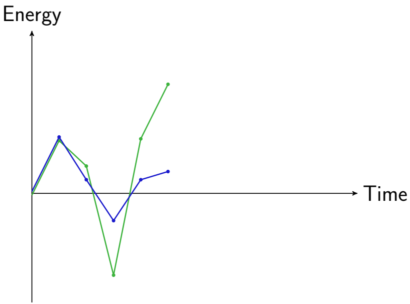

Energy, total-payoff, mean-payoff

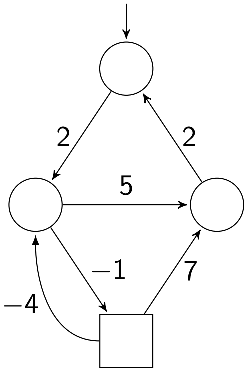

Then, $(2,5,2)^\omega$

- $EL(\pi(n)) = \sum_{i=0}^{n-1} w(s_i, s_{i+1})$
- $\overline{TP}(\pi) = \limsup_{n\to\infty} EL(\pi(n))$
- $\overline{MP}(\pi) = \limsup_{n\to\infty} \frac{1}{n} EL(\pi(n))$

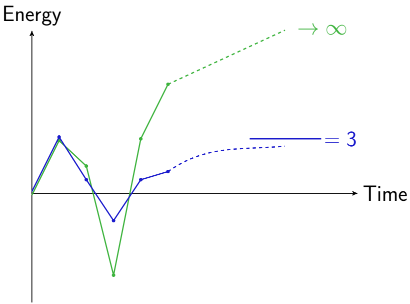

Decision problems

- **TP (MP) threshold problem**
  - ▷ Given $t \in \mathbb{Q}$ and $s_{\text{init}} \in S$, $\exists? \sigma_1 \in \Sigma_1$ s.t. $\forall \sigma_2 \in \Sigma_2$,
    $\overline{TP}(Outcome(s_{\text{init}}, \sigma_1, \sigma_2)) \leq t$
  - ↪ we take the **minimizer** point of view

- **Lower-bounded energy problem**
  - ▷ Given $c_{\text{init}} \in \mathbb{N}$ and $s_{\text{init}} \in S$, $\exists? \sigma_1 \in \Sigma_1$ s.t. $\forall \sigma_2 \in \Sigma_2$,
    $\forall n \geq 0,\ c_{\text{init}} + EL(Outcome(s_{\text{init}}, \sigma_1, \sigma_2)(n)) \geq 0$
  - ↪ **fixed** initial credit

- **Lower- and upper-bounded energy problem**
  - ▷ Given $c_{\text{init}} \in \mathbb{N}$, $U \in \mathbb{N}$ and $s_{\text{init}} \in S$, $\exists? \sigma_1 \in \Sigma_1$ s.t. $\forall \sigma_2 \in \Sigma_2$,
    $\forall n \geq 0,\ c_{\text{init}} + EL(Outcome(s_{\text{init}}, \sigma_1, \sigma_2)(n)) \in [0, U]$

Known results

Game objective	**1-player**	**2-player**
MP	P [Kar78]	NP ∩ coNP [ZP96]
TP	P [FV97]	NP ∩ coNP [GS09]
EG_L	P [BFL+08]	NP ∩ coNP [CdAHS03, BFL+08]
EG_{LU}	PSPACE-c. [FJ13]	EXPTIME-c. [BFL+08]

▷ For all objectives but EG_{LU}, *memoryless* strategies suffice for both players.

Average-energy: motivating example

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

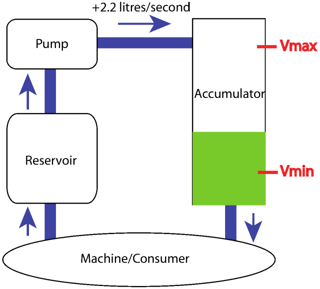

Goals:

1. Keep the oil level in the safe zone.
   $\hookrightarrow$ EG_{LU}
2. Minimize the average oil level.
   $\hookrightarrow$ AE

$\Rightarrow$ Conjunction: AE_{LU}

Average-energy: definition

Recall

- $EL(\pi(n)) = \sum_{i=0}^{n-1} w(s_i, s_{i+1})$
- $\overline{TP}(\pi) = \limsup_{n\to\infty} EL(\pi(n))$
- $\overline{MP}(\pi) = \limsup_{n\to\infty} \frac{1}{n} EL(\pi(n))$

+ infimum variants
$\underline{TP}$, $\underline{MP}$, $\underline{AE}$

Average-energy (AE)

Describes the **average energy level** along a play:

$$\overline{AE}(\pi) = \limsup_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n} EL(\pi(i))$$

TP "refines" MP

- If $\mathcal{P}_1$ (minimizer) can ensure $\underline{MP} = \overline{MP} < 0$ (memoryless), he can ensure $\underline{TP} = \overline{TP} = -\infty$.
- If $\mathcal{P}_2$ (maximizer) can ensure $\underline{MP} = \overline{MP} > 0$ (memoryless), he can ensure $\underline{TP} = \overline{TP} = \infty$.

TP "refines" MP

- If $\mathcal{P}_1$ (minimizer) can ensure $\underline{MP} = \overline{MP} < 0$ (memoryless), he can ensure $\underline{TP} = \overline{TP} = -\infty$.
- If $\mathcal{P}_2$ (maximizer) can ensure $\underline{MP} = \overline{MP} > 0$ (memoryless), he can ensure $\underline{TP} = \overline{TP} = \infty$.

$\Rightarrow$ **TP discriminates "MP-zero" strategies** depending on the high points ($\overline{TP}$) or low points ($\underline{TP}$) of cycles.

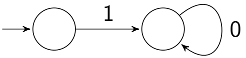

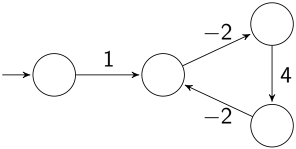

$\overline{MP} = \underline{MP} = 0$

$\overline{TP} = \underline{TP} = 1$

$\overline{MP} = \underline{MP} = 0$

$\overline{TP} = 3, \ \underline{TP} = -1$

AE "refines" TP

AE describes the **long-run average** *EL*

↪ By definition, $\underline{AE}(\pi), \overline{AE}(\pi) \in [\underline{TP}(\pi), \overline{TP}(\pi)]$.

AE "refines" TP

AE describes the **long-run average** *EL*

↪ By definition, $\underline{AE}(\pi), \overline{AE}(\pi) \in [\underline{TP}(\pi), \overline{TP}(\pi)]$.

⇒ **AE discriminates strategies with identical high/low points**.

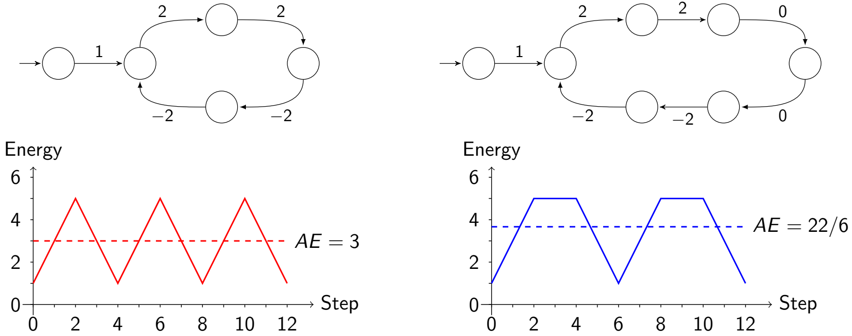

Identical MP and TP, but AE lower in the first one.

Memoryless determinacy (1/2)

Classical criteria from the literature **cannot be applied out-of-the-box** [EM79, BSV04, AR14, GZ04, Kop06].

↪ Common approach: connect *first cycle* games and infinite-duration ones.

↪ Requires e.g., closure under *cyclic permutation* and *concatenation* [AR14].

Intuitively: ability to **mix and shuffle** good cycles and stay good.

Memoryless determinacy (1/2)

Classical criteria from the literature **cannot be applied out-of-the-box** [EM79, BSV04, AR14, GZ04, Kop06].

↪ Common approach: connect *first cycle* games and infinite-duration ones.

↪ Requires e.g., closure under *cyclic permutation* and *concatenation* [AR14].

Intuitively: ability to **mix and shuffle** good cycles and stay good.

Not true in general for AE!

$$\mathcal{C}_1 = \{-1\},\ \mathcal{C}_2 = \{1\},\ \mathcal{C}_3 = \{1, -2\}$$

$$AE(\mathcal{C}_1\mathcal{C}_2) = (-1+0)/2 = -1/2 < AE(\mathcal{C}_2\mathcal{C}_1) = (1-0)/2 = 1/2$$

$$AE(\mathcal{C}_3) = 0 \text{ but } AE(\mathcal{C}_3\mathcal{C}_3) = -1/2 < 0$$

Memoryless determinacy (1/2)

Classical criteria from the literature **cannot be applied out-of-the-box** [EM79, BSV04, AR14, GZ04, Kop06].

↪ Common approach: connect *first cycle* games and infinite-duration ones.

↪ Requires e.g., closure under *cyclic permutation* and *concatenation* [AR14].

Intuitively: ability to **mix and shuffle** good cycles and stay good.

We can only shuffle/repeat cycles that are neutral w.r.t. the energy level!

↪ **zero-cycles**

Memoryless determinacy (2/2)

Two key properties:

1. **Extraction of prefixes**
   - Let $\rho \in \mathit{Prefs}(G)$, $\pi \in \mathit{Plays}(G)$. Then,

$$\overline{AE}(\rho \cdot \pi) = EL(\rho) + \overline{AE}(\pi).$$

2. **Extraction of a best cycle**
   - Given an infinite sequence of *zero-cycles*, one can select and repeat a *best cycle* to minimize the average-energy.

One-player games: strategy

Sketch (minimizer)

1. If you can ensure $MP < 0$, do it.
   - ▷ Memoryless [EM79], implies $AE = -\infty$.
2. If you *cannot* ensure $MP = 0$, forget it.
   - ▷ You are doomed, $AE = \infty$.
3. Play the strategy that minimizes

$$\overline{AE}(\rho \cdot \mathcal{C}^{\omega}) = EL(\rho) + \overline{AE}(\mathcal{C}),$$

where $\mathcal{C}$ is a simple zero-cycle.

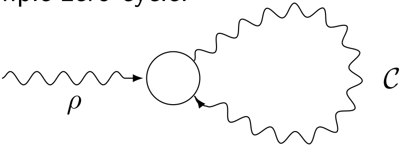

↪ Picking the best combination can be done **without memory**.

One-player games: P algorithm (1/2)

- Case $MP < 0$ is easy
  - ▷ Look for a negative cycle (e.g., Bellman-Ford, $\mathcal{O}(|S|^3)$)

One-player games: P algorithm (1/2)

- Case $MP < 0$ is easy
  - ▷ Look for a negative cycle (e.g., Bellman-Ford, $\mathcal{O}(|S|^3)$)
- Assume $MP = 0$: **pick the best combination** of ρ and $\mathcal{C}$
  - ▷ Computing the best ρ for each state is easy with classical graph algorithms (e.g., Bellman-Ford).
  - ▷ **Main task**: computing the best $\mathcal{C}$ (AE-wise) for each state.

One-player games: P algorithm (1/2)

- Case $MP < 0$ is easy
  - ▷ Look for a negative cycle (e.g., Bellman-Ford, $\mathcal{O}(|S|^3)$)
- Assume $MP = 0$: **pick the best combination** of ρ and $\mathcal{C}$
  - ▷ Computing the best ρ for each state is easy with classical graph algorithms (e.g., Bellman-Ford).
  - ▷ **Main task**: computing the best $\mathcal{C}$ (AE-wise) for each state.
- For each state, we compute the best cycle of length k, for all $k \in \{1, \ldots, |S|\}$, then pick the best one.
  - ▷ Need to compute $\mathcal{C}_{s,k}$ in polynomial time.

One-player games: P algorithm (2/2)

Computing $\mathcal{C}_{s,k}$: build a **new graph** $\mathcal{G}_{s,k}$ of size $|S| \cdot (k+1)$.

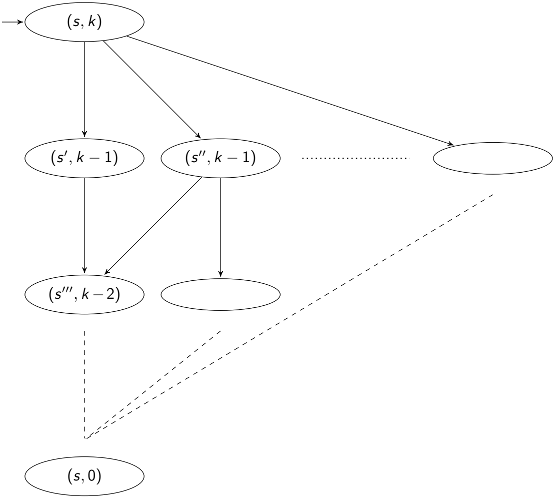

One-player games: P algorithm (2/2)

Computing $\mathcal{C}_{s,k}$: build a **new graph** $\mathcal{G}_{s,k}$ of size $|S| \cdot (k+1)$.

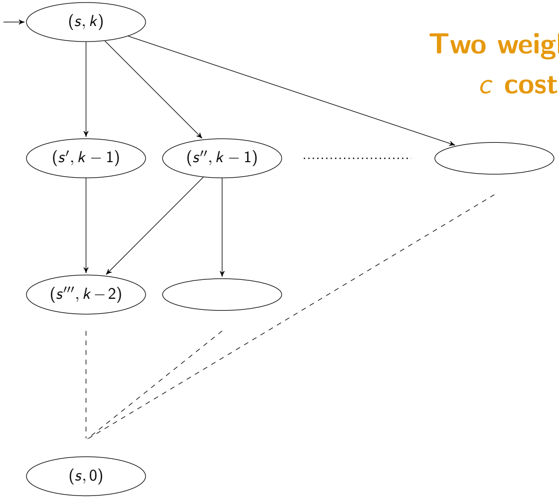

Two weights by edge: $(c, l \cdot c)$
c cost, l origin level

One-player games: P algorithm (2/2)

Computing $\mathcal{C}_{s,k}$: build a **new graph** $\mathcal{G}_{s,k}$ of size $|S| \cdot (k+1)$.

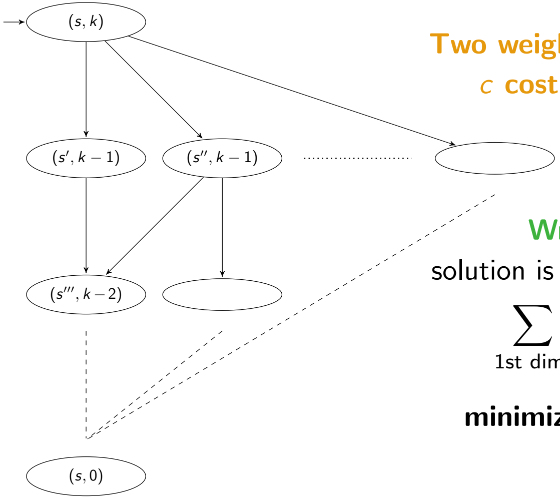

Two weights by edge: $(c, l \cdot c)$

c cost, l origin level

Write an LP s.t.

solution is a path from (s, k) to $(s, 0)$

$$\sum_{\text{1st dim.}} = 0 \text{ (zero cycle)}$$

$$\textbf{minimize} \sum_{\text{2nd dim.}} = AE(\mathcal{C}) \cdot k$$

One-player games: P algorithm (2/2)

Computing $\mathcal{C}_{s,k}$: build a **new graph** $\mathcal{G}_{s,k}$ of size $|S| \cdot (k+1)$.

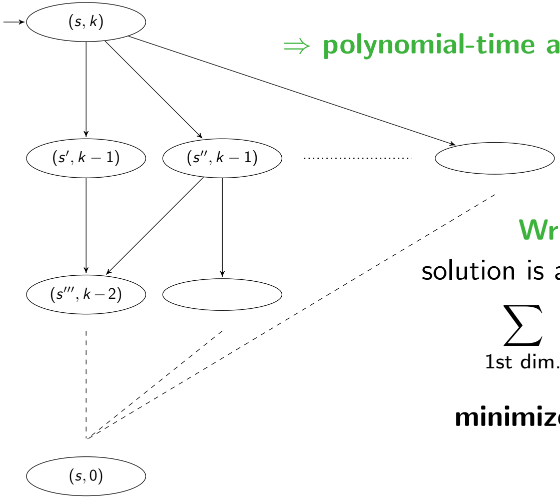

$\Rightarrow$ **polynomial-time algorithm for 1-p. games**

Write an LP s.t.

solution is a path from (s,k) to $(s,0)$

$$\sum_{\text{1st dim.}} = 0 \text{ (zero cycle)}$$

$$\textbf{minimize} \sum_{\text{2nd dim.}} = AE(\mathcal{C}) \cdot k$$

Two-player games

- **Memoryless determinacy**
  - ▷ Follows from the 1-p. results (minimizer *and* maximizer) using Gimbert and Zielonka [GZ05].
- **Threshold problem in NP ∩ coNP.**
  - ▷ Memoryless determinacy + P for one-player games.
- **"Mean-payoff" hard.**
  - ▷ Replace any edge of weight c by two consecutive edges of values $2 \cdot c$ and $-2 \cdot c$.
  - ▷ $MP(\pi)$ in $G = AE(\pi)$ in G'.

Wrap-up

Game objective	**1-player**	**2-player**
MP	P [Kar78]	NP ∩ coNP [ZP96]
TP	P [FV97]	NP ∩ coNP [GS09]
EG_L	P [BFL+08]	NP ∩ coNP [CdAHS03, BFL+08]
EG_{LU}	PSPACE-c. [FJ13]	EXPTIME-c. [BFL+08]
AE	P	NP ∩ coNP

▷ For all objectives but EG_{LU}, *memoryless* strategies suffice for both players.

Two settings

1. AE_{LU}: AE with lower (0) and upper ($U \in \mathbb{N}$) bounds.
2. AE_L: AE with only the lower bound (0).

$\hookrightarrow$ Fixed initial credit $c_{\text{init}} = 0$.

Memory is needed!

Example: $AE_L \rightsquigarrow$ minimize AE while keeping $EL \geq 0$.

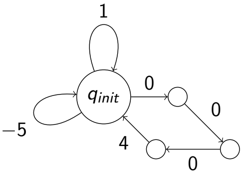

(a) 1-Player AE_L game

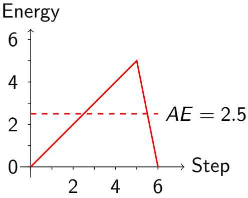

(b) Taking the +1 loop

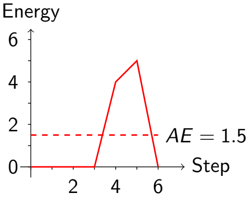

(c) Taking +4 then +1

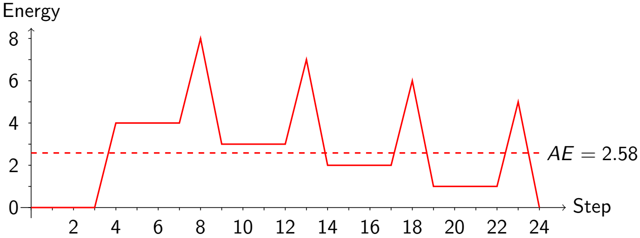

(d) Taking the +4 loop

Memory is needed!

Example: $AE_L \rightsquigarrow$ minimize AE while keeping $EL \geq 0$.

Non trivial behavior in general!

$\hookrightarrow$ **Need to choose carefully which cycles to play.**

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, $AE = \infty$.

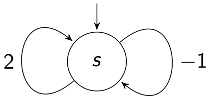

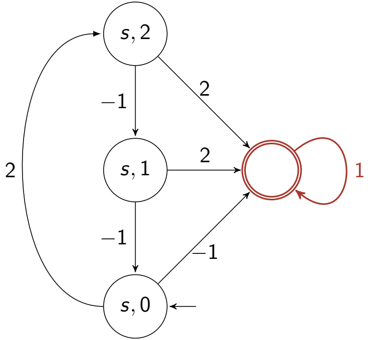

$\mathcal{P}_1$ minimizes AE and maintains $EL \in [0, 2]$ in the left game iff $\mathcal{P}_1$ minimizes AE in the right game.

AE_{LU} problem: complexity

Game objective	**1-player**	**2-player**
MP	P [Kar78]	NP ∩ coNP [ZP96]
TP	P [FV97]	NP ∩ coNP [GS09]
EG_L	P [BFL+08]	NP ∩ coNP [CdAHS03, BFL+08]
EG_{LU}	PSPACE-c. [FJ13]	EXPTIME-c. [BFL+08]
AE	P	NP ∩ coNP
AE_{LU} (poly. U)	P	NP ∩ coNP
AE_{LU} (arbitrary)	EXPTIME /PSPACE-h.	NEXPTIME ∩ coNEXPTIME /EXPTIME-h.

- ▷ Lower bounds follow from EG_{LU}.
- ▷ Memory is required (at most exponential).

AE_L problem: one-player case

Key argument: (upper) bounding the value of the energy over a witness winning path.

- ↪ It is not necessary to accumulate *too much* energy. Intuitively, otherwise we cannot keep the AE sufficiently low.
- ↪ Bound U polynomial in $|S|$, the largest absolute weight W and the threshold t for the AE constraint.
- ↪ **Reduction to an AE_{LU} problem.**
- ↪ EXPTIME-algorithm.

AE_L problem: one-player case

Key argument: (upper) bounding the value of the energy over a witness winning path.

- ↪ It is not necessary to accumulate *too much* energy. Intuitively, otherwise we cannot keep the AE sufficiently low.
- ↪ Bound U polynomial in $|S|$, the largest absolute weight W and the threshold t for the AE constraint.
- ↪ **Reduction to an AE_{LU} problem.**
- ↪ EXPTIME-algorithm.

Lower bound: NP-hard via *subset sum problem* [GJ79].

- ↪ Find a subset of a set of naturals s.t. the sum of its elements is exactly equal to target $T \in \mathbb{N}$.
- ↪ The energy LB can be used to ensure a sum $\geq T$ and the AE to ensure $\leq T$.

AE_L problem: two-player case

- Algorithm: work still in progress.
  - ↪ We believe we can apply a similar approach, upper bounding the energy.
  - ↪ Non-trivial alternations between carefully chosen cycles is required (see previous example).

AE_L problem: two-player case

- Algorithm: work still in progress.
  - ↪ We believe we can apply a similar approach, upper bounding the energy.
  - ↪ Non-trivial alternations between carefully chosen cycles is required (see previous example).
- Problem is EXPTIME-hard via *countdown games* [JSL08].

AE_L problem: complexity

Game objective	**1-player**	**2-player**
MP	P [Kar78]	NP ∩ coNP [ZP96]
TP	P [FV97]	NP ∩ coNP [GS09]
EG_L	P [BFL$^+$08]	NP ∩ coNP [CdAHS03, BFL$^+$08]
EG_{LU}	PSPACE-c. [FJ13]	EXPTIME-c. [BFL$^+$08]
AE	P	NP ∩ coNP
AE_{LU} (poly. U)	P	NP ∩ coNP
AE_{LU} (arbitrary)	EXPTIME /PSPACE-h.	NEXPTIME ∩ coNEXPTIME /EXPTIME-h.
AE_L	EXPTIME/NP-h.	???/EXPTIME-h.

▷ Memory is required (at most exponential).

Wrap-up

New quantitative objective.

- Practical motivations (e.g., HYDAC).
- "Refines" *TP* (and *MP*).
- Same complexity class as EG_L, *MP* and *TP* games.
- Still some open questions.
  - ↪ Complexity gaps.
  - ↪ Algorithm for 2-player AE_L games.

Thanks!

Do not hesitate to discuss with us!

References I

Benjamin Aminof and Sasha Rubin.
First cycle games.
In Proceedings of the 2nd International Workshop on Strategic Reasoning (SR'14), volume 146 of Electronic Proceedings in Theoretical Computer Science, pages 83–90, March 2014.

Patricia Bouyer, Uli Fahrenberg, Kim Gulstrand 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.

Henrik Björklund, Sven Sandberg, and Sergei Vorobyov.
Memoryless determinacy of parity and mean payoff games: A simple proof.
Theoretical Computer Science, 310(1-3):365–378, January 2004.

Arindam Chakrabarti, Luca de Alfaro, Thomas A. Henzinger, and Mariëlle Stoelinga.
Resource interfaces.
In Rajeev Alur and Insup Lee, editors, Proceedings of the 3rd International Conference on Embedded Software (EMSOFT'03), volume 2855 of Lecture Notes in Computer Science, pages 117–133. Springer-Verlag, October 2003.

Krishnendu Chatterjee, Laurent Doyen, Mickael Randour, and Jean-François Raskin.
Looking at mean-payoff and total-payoff through windows.
In Dang Van Hung and Mizuhito Ogawa, editors, ATVA, volume 8172 of Lecture Notes in Computer Science, pages 118–132. Springer, 2013.

References II

Franck Cassez, Jan J. Jensen, Kim Gulstrand Larsen, Jean-François Raskin, and Pierre-Alain Reynier.
Automatic synthesis of robust and optimal controllers – an industrial case study.
In Rupak Majumdar and Paulo Tabuada, editors, Proceedings of the 12th International Workshop on Hybrid Systems: Computation and Control (HSCC'09), volume 5469 of Lecture Notes in Computer Science, pages 90–104. Springer-Verlag, April 2009.

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

John Fearnley and Marcin Jurdziński.
Reachability in two-clock timed automata is PSPACE-complete.
In Fedor V. Fomin, Rusins Freivalds, Marta Kwiatkowska, and David Peleg, editors, Proceedings of the 40th International Colloquium on Automata, Languages and Programming (ICALP'13) – Part II, volume 7966 of Lecture Notes in Computer Science, pages 212–223. Springer-Verlag, July 2013.

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

Michael R. Garey and David S. Johnson.
Computers and intractability: a guide to the Theory of NP-Completeness.
Freeman New York, 1979.

References III

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.

Hugo Gimbert and Wiesław Zielonka.
Games where you can play optimally without any memory.
In Martín Abadi and Luca de Alfaro, editors, Proceedings of the 16th International Conference on Concurrency Theory (CONCUR'05), volume 3653 of Lecture Notes in Computer Science, pages 428–442. Springer-Verlag, August 2005.

Marcin Jurdziński, Jeremy Sproston, and François Laroussinie.
Model checking probabilistic timed automata with one or two clocks.
Logical Methods in Computer Science, 4(3), 2008.

Richard M. Karp.
A characterization of the minimum cycle mean in a digraph.
Discrete Mathematics, 23(3):309–311, September 1978.

References IV

Eryk Kopczynski.
Half-positional determinacy of infinite games.
In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP'06)) – Part II, volume 4052 of Lecture Notes in Computer Science, pages 336–347. Springer-Verlag, July 2006.

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