Average-Energy Games and Beyond

Patricia Bouyer

Based on joint works with:

Nicolas Markey
Mickael Randour
Kim G. Larsen
Simon Laursen
GandALF'15 / Acta Informatica

Piotr Hofman
Nicolas Markey
Mickael Randour
Martin Zimmermann
FoSSaCS'17

Thanks to Mickael for his slides!

General context: strategy synthesis in quantitative games

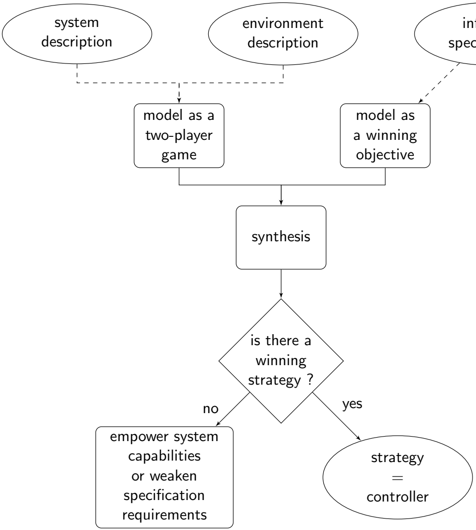

1. How complex is it to decide if a winning strategy exists?
2. How complex such a strategy needs to be? **Simpler is better**.
3. Can we synthesize one efficiently?

$\Rightarrow$ Depends on the *winning objective*.

Motivating example

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

Goals:

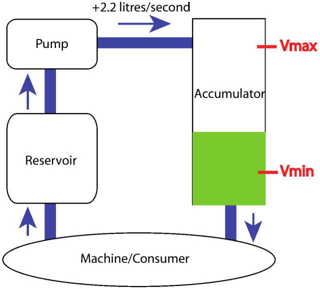

Motivating example

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

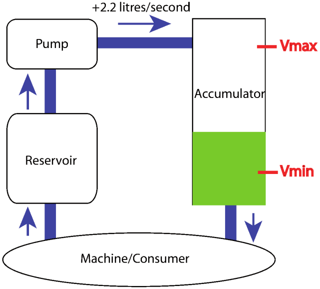

Goals:

1. Keep the oil level in the safe zone.
   ↪ **Energy objective with lower and upper bounds:** EG_{LU}

Motivating example

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

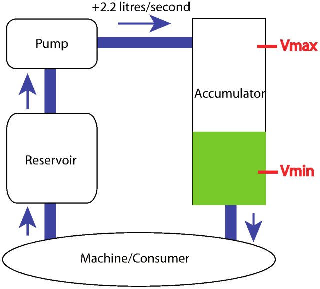

Goals:

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

Motivating example

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

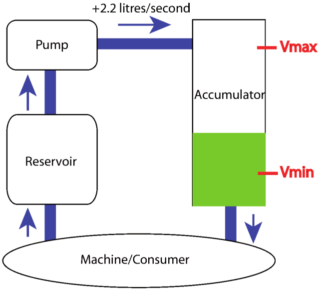

Goals:

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

$\Rightarrow$ **Conjunction:** AE_{LU}

Average-energy: illustration

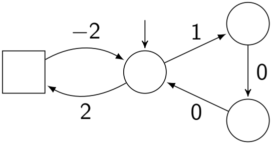

- Two-player turn-based games with integer weights.

Average-energy: illustration

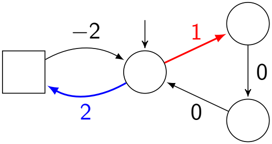

- Two-player turn-based games with integer weights.
- Focus on two *memoryless* strategies.

Average-energy: illustration

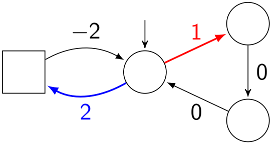

- Two-player turn-based games with integer weights.
- Focus on two *memoryless* strategies.

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

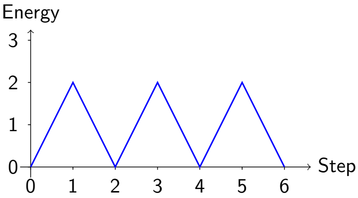

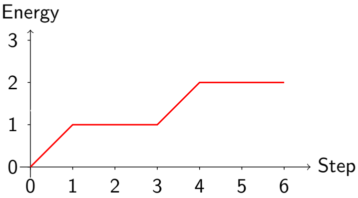

Average-energy: illustration

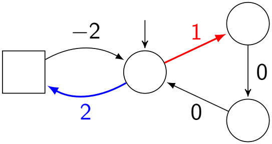

- Two-player turn-based games with integer weights.
- Focus on two *memoryless* strategies.

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

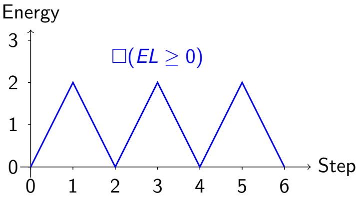

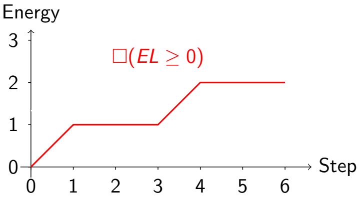

Energy objective (EG_L/EG_{LU}): e.g., always maintain $EL \geq 0$.

Average-energy: illustration

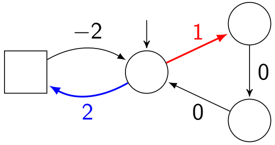

- Two-player turn-based games with integer weights.
- Focus on two *memoryless* strategies.

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

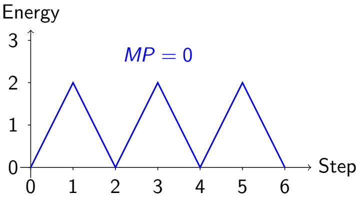

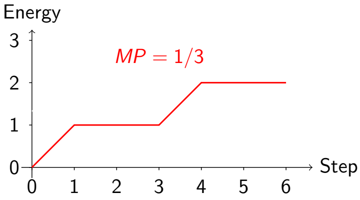

Mean-payoff (MP): long-run average payoff per transition.

Average-energy: illustration

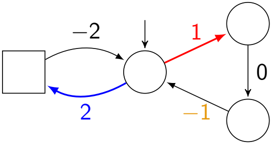

- Two-player turn-based games with integer weights.
- Focus on two *memoryless* strategies.

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

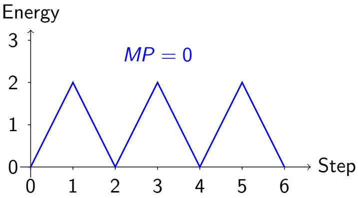

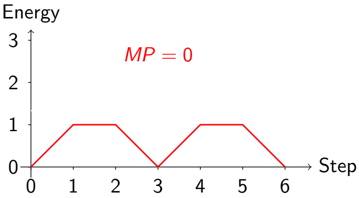

Mean-payoff (*MP*): long-run average payoff per transition.

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

Average-energy: illustration

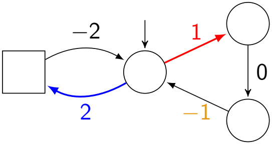

- Two-player turn-based games with integer weights.
- Focus on two *memoryless* strategies.

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

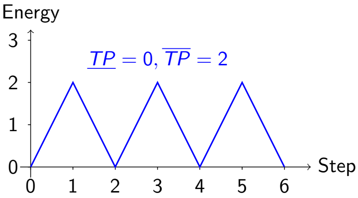

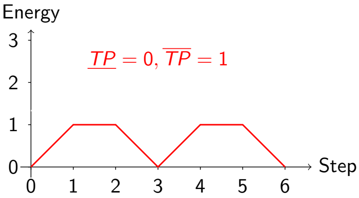

Total-payoff (*TP*) *refines MP* in the case $MP = 0$ by looking at high/low points of the sequence.

Average-energy: illustration

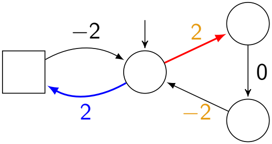

- Two-player turn-based games with integer weights.
- Focus on two *memoryless* strategies.

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

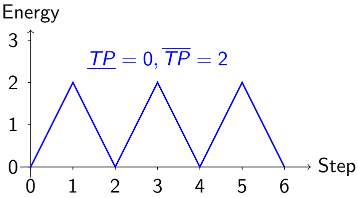

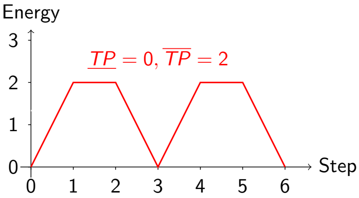

Total-payoff (TP) *refines MP* in the case $MP = 0$ by looking at high/low points of the sequence.

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

Average-energy: illustration

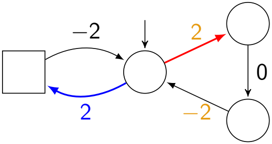

- Two-player turn-based games with integer weights.
- Focus on two *memoryless* strategies.

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

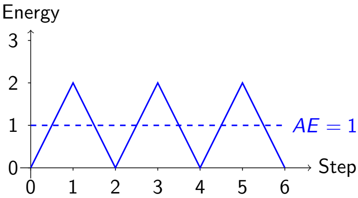

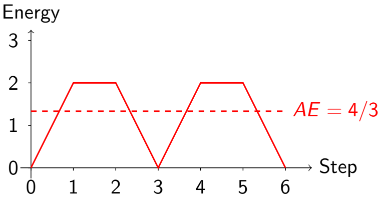

Average-energy (*AE*) *further refines TP*: average EL along a play.

Average-energy: illustration

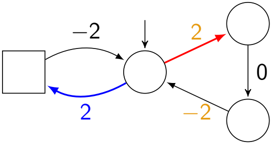

- Two-player turn-based games with integer weights.
- Focus on two *memoryless* strategies.

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

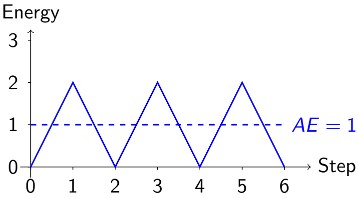

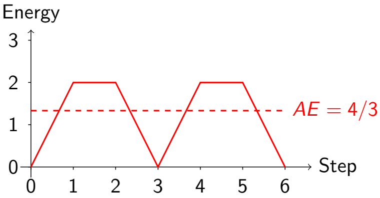

Average-energy (AE) *further refines TP*: average EL along a play.

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

Average-energy: overview

Objective	**1-player**	**2-player**	**memory**
MP	P [Kar78]	NP $\cap$ coNP [ZP96]	memoryless [EM79]
TP	P [FV97]	NP $\cap$ coNP [GS09]	memoryless [GZ04]
EG_L	P [BFL$^+$08]	NP $\cap$ coNP [CdAHS03, BFL$^+$08]	memoryless [CdAHS03]
EG_{LU}	PSPACE-c. [FJ13]	EXPTIME-c. [BFL$^+$08]	pseudo-polynomial
AE	P	NP $\cap$ coNP	memoryless

▷ For all but EG_{LU}, *memoryless* strategies suffice.

Average-energy: overview

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. [FJ13]	EXPTIME-c. [BFL+08]	pseudo-polynomial
AE	P	NP ∩ coNP	memoryless

▷ For all but EG_{LU}, *memoryless* strategies suffice.

Techniques:

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

Average-energy: overview

Objective	**1-player**	**2-player**	**memory**
MP	P [Kar78]	NP $\cap$ coNP [ZP96]	memoryless [EM79]
TP	P [FV97]	NP $\cap$ coNP [GS09]	memoryless [GZ04]
EG_L	P [BFL+08]	NP $\cap$ coNP [CdAHS03, BFL+08]	memoryless [CdAHS03]
EG_{LU}	PSPACE-c. [FJ13]	EXPTIME-c. [BFL+08]	pseudo-polynomial
AE	P	NP $\cap$ coNP	memoryless

- For all but EG_{LU}, *memoryless* strategies suffice.

Techniques:

- Classical criteria **cannot be applied out-of-the-box** [EM79, BSV04, AR14, GZ04, Kop06].
- 1-player: memorylessness proof and polynomial-time LP-based algorithm.

Average-energy: overview

Objective	**1-player**	**2-player**	**memory**
MP	P [Kar78]	NP $\cap$ coNP [ZP96]	memoryless [EM79]
TP	P [FV97]	NP $\cap$ coNP [GS09]	memoryless [GZ04]
EG_L	P [BFL+08]	NP $\cap$ coNP [CdAHS03, BFL+08]	memoryless [CdAHS03]
EG_{LU}	PSPACE-c. [FJ13]	EXPTIME-c. [BFL+08]	pseudo-polynomial
AE	P	NP $\cap$ coNP	memoryless

▷ For all but EG_{LU}, *memoryless* strategies suffice.

Techniques:

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

▷ 1-player: memorylessness proof and polynomial-time LP-based algorithm.

▷ 2-player: extension thanks to Gimbert and Zielonka [GZ05].

Average-energy: overview

Objective	**1-player**	**2-player**	**memory**
MP	P [Kar78]	NP $\cap$ coNP [ZP96]	memoryless [EM79]
TP	P [FV97]	NP $\cap$ coNP [GS09]	memoryless [GZ04]
EG_L	P [BFL$^+$08]	NP $\cap$ coNP [CdAHS03, BFL$^+$08]	memoryless [CdAHS03]
EG_{LU}	PSPACE-c. [FJ13]	EXPTIME-c. [BFL$^+$08]	pseudo-polynomial
AE	P	NP $\cap$ coNP	memoryless

- For all but EG_{LU}, *memoryless* strategies suffice.

Techniques:

- Classical criteria **cannot be applied out-of-the-box** [EM79, BSV04, AR14, GZ04, Kop06].
- 1-player: memorylessness proof and polynomial-time LP-based algorithm.
- 2-player: extension thanks to Gimbert and Zielonka [GZ05].
- MP-hardness.

With LU energy constraints, memory is needed!

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

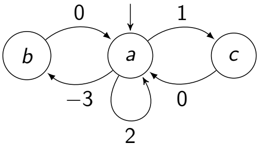

(a) One-player AE_{LU} game.

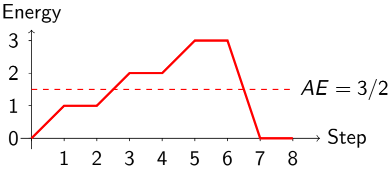

(b) Play $\pi_1 = (acacacab)^\omega$.

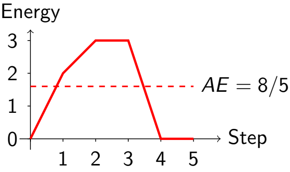

(c) Play $\pi_2 = (aacab)^\omega$.

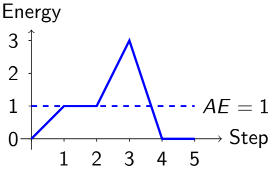

(d) Play $\pi_3 = (acaab)^\omega$.

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

With LU 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.

With LU 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 ($\Rightarrow$ $AE_{LU} \in$ NEXPTIME $\cap$ coNEXPTIME).

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

What about L constraints?

One-player case only!

- Upper bound on the energy level, thanks to Lafourcade et al [LLT04]
- Results for AE_{LU} apply!

- Unfortunately hard to extend to two-player games

With energy constraints: results overview

Objective	**1-player**	**2-player**	**memory**
MP	P [Kar78]	NP $\cap$ coNP [ZP96]	memoryless [EM79]
TP	P [FV97]	NP $\cap$ coNP [GS09]	memoryless [GZ04]
EG_L	P [BFL$^+$08]	NP $\cap$ coNP [CdAHS03, BFL$^+$08]	memoryless [CdAHS03]
EG_{LU}	PSPACE-c. [FJ13]	EXPTIME-c. [BFL$^+$08]	pseudo-polynomial
AE	P	NP $\cap$ coNP	memoryless
AE_{LU} (poly. U)	P	NP $\cap$ coNP	polynomial
AE_{LU} (arbitrary)	PSPACE-c.	EXPTIME-c.	pseudo-polynomial
AE_L	PSPACE-e./NP-h.	/EXPTIME-h.	

With energy constraints: results overview

Objective	**1-player**	**2-player**	**memory**
MP	P [Kar78]	NP $\cap$ coNP [ZP96]	memoryless [EM79]
TP	P [FV97]	NP $\cap$ coNP [GS09]	memoryless [GZ04]
EG_L	P [BFL$^+$08]	NP $\cap$ coNP [CdAHS03, BFL$^+$08]	memoryless [CdAHS03]
EG_{LU}	PSPACE-c. [FJ13]	EXPTIME-c. [BFL$^+$08]	pseudo-polynomial
AE	P	NP $\cap$ coNP	memoryless
AE_{LU} (poly. U)	P	NP $\cap$ coNP	polynomial
AE_{LU} (arbitrary)	PSPACE-c.	EXPTIME-c.	pseudo-polynomial
AE_L	PSPACE-e./NP-h.	*open*/EXPTIME-h.	*open* ($\geq$ pseudo-poly.)

Status at Barbizon last year!

With energy constraints: results overview

Objective	**1-player**	**2-player**	**memory**
MP	P [Kar78]	NP $\cap$ coNP [ZP96]	memoryless [EM79]
TP	P [FV97]	NP $\cap$ coNP [GS09]	memoryless [GZ04]
EG_L	P [BFL+08]	NP $\cap$ coNP [CdAHS03, BFL+08]	memoryless [CdAHS03]
EG_{LU}	PSPACE-c. [FJ13]	EXPTIME-c. [BFL+08]	pseudo-polynomial
AE	P	NP $\cap$ coNP	memoryless
AE_{LU} (poly. U)	P	NP $\cap$ coNP	polynomial
AE_{LU} (arbitrary)	PSPACE-c.	EXPTIME-c.	pseudo-polynomial
AE_L	PSPACE-e./NP-h.	2-EXPTIME-e./EXPSPACE-h.	*super-exp.* (doubly exp.)

Thanks to a fresh idea by Piotr Hofman

Two-player games with L-constraints

The crux idea

If the average is low (smaller than t), then there must be a large number of configurations with energy level smaller than t!

Two-player games with L-constraints

The crux idea

If the average is low (smaller than t), then there must be a large number of configurations with energy level smaller than t!

Assume $AE(\pi) \leq t$:

Two-player games with L-constraints

The crux idea

If the average is low (smaller than t), then there must be a large number of configurations with energy level smaller than t!

Assume $AE(\pi) \leq t$:

- $\exists \gamma \in \Gamma^{\leq t}$ s.t. $\forall n$, $\exists^\infty n' \geq n$,

$$density(\gamma, \pi_{[n,n']}) \geq \frac{\widetilde{t}}{4(t+1)^2|S|}$$

Two-player games with L-constraints

The crux idea

If the average is low (smaller than t), then there must be a large number of configurations with energy level smaller than t!

Assume $AE(\pi) \leq t$:

- $\exists \gamma \in \Gamma^{\leq t}$ s.t. $\forall n$, $\exists^\infty n' \geq n$,

$$density(\gamma, \pi_{[n,n']}) \geq \frac{\widetilde{t}}{4(t+1)^2|S|}$$

- There is a reachable cycle with average $\leq t$ (called good)

Two-player games with L-constraints

The crux idea

If the average is low (smaller than t), then there must be a large number of configurations with energy level smaller than t!

Assume $AE(\pi) \leq t$:

- $\exists \gamma \in \Gamma^{\leq t}$ s.t. $\forall n$, $\exists^\infty n' \geq n$,

$$density(\gamma, \pi_{[n,n']}) \geq \frac{\widetilde{t}}{4(t+1)^2|S|}$$

- There is a reachable cycle with average $\leq t$ (called good)
- All (reachable) good cycles with no strict good sub-cycles have length bounded by $8t_1t_2(t+1)^3|S|^2$

Seeing strategies as (finite) trees

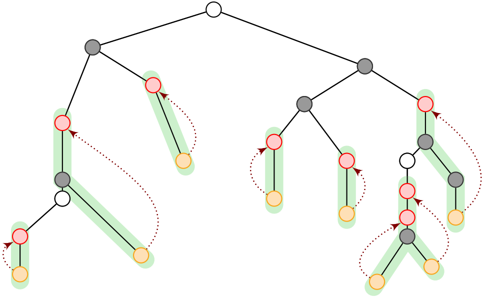

Seeing strategies as (finite) trees

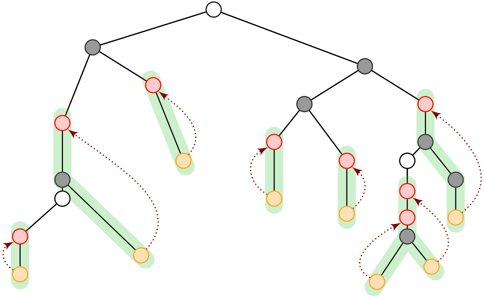

- A finite good tree represents a winning strategy

Seeing strategies as (finite) trees

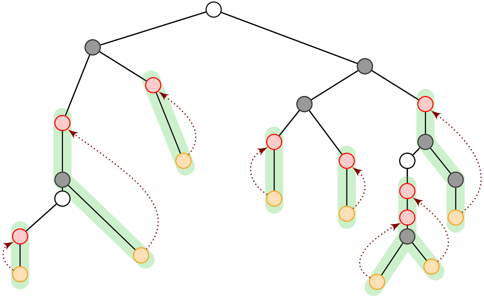

- A finite good tree represents a winning strategy
- Fix a winning strategy, and build a finite tree by “closing” minimal good cycles. We then have a finite good tree, hence a winning strategy!

What is missing?

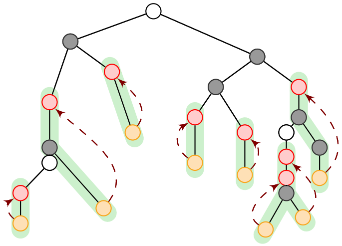

What is missing?

- energy level is bounded in green parts

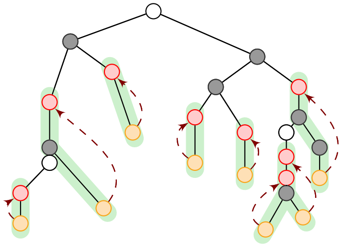

What is missing?

- energy level is bounded in green parts
- what about outside green parts?

  $\rightsquigarrow$ better understand winning strategies in pushdown games

  an original idea by [Wal01], revisited in [FZ12], from which we can derive a doubly-exponential upper bound on the energy level!

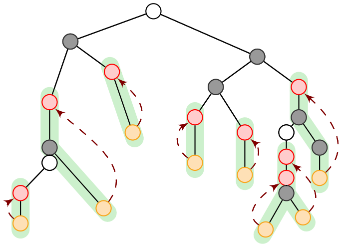

What is missing?

- energy level is bounded in green parts
- what about outside green parts?

  $\rightsquigarrow$ better understand winning strategies in pushdown games

  an original idea by [Wal01], revisited in [FZ12], from which we can derive a doubly-exponential upper bound on the energy level!
- reduced to AE_{LU} problem!

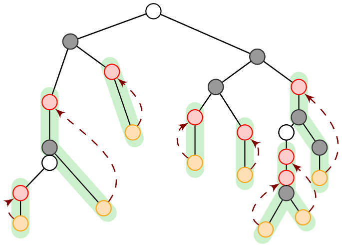

With energy constraints: results overview

Objective	**1-player**	**2-player**	**memory**
MP	P [Kar78]	NP $\cap$ coNP [ZP96]	memoryless [EM79]
TP	P [FV97]	NP $\cap$ coNP [GS09]	memoryless [GZ04]
EG_L	P [BFL+08]	NP $\cap$ coNP [CdAHS03, BFL+08]	memoryless [CdAHS03]
EG_{LU}	PSPACE-c. [FJ13]	EXPTIME-c. [BFL+08]	pseudo-polynomial
AE	P	NP $\cap$ coNP	memoryless
AE_{LU} (poly. U)	P	NP $\cap$ coNP	polynomial
AE_{LU} (arbitrary)	PSPACE-c.	EXPTIME-c.	pseudo-polynomial
AE_L	PSPACE-e./NP-h.	2-EXPTIME-e./EXPSPACE-h.	*super-exp.* (doubly exp.)

Wrap-up

"New" quantitative objective.[1]

- Yields natural payoff functions.
- AE "refines" TP (and MP).
- Same complexity class as EG_L, MP and TP games.
- AE_{LU} and AE_L now well-understood.
- Interesting proofs techniques.

[1] Appeared in [TV87] as an alternative *total reward* definition but not studied until recently. See also [CP13, BEGM15].

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.

Endre Boros, Khaled M. Elbassioni, Vladimir Gurvich, and Kazuhisa Makino.
Markov decision processes and stochastic games with total effective payoff.
In Ernst W. Mayr and Nicolas Ollinger, editors, 32nd International Symposium on Theoretical Aspects of Computer Science, STACS 2015, March 4-7, 2015, Garching, Germany, volume 30 of LIPIcs, pages 103–115. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2015.

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.

Patricia Bouyer, Nicolas Markey, Mickael Randour, Kim Guldstrand Larsen, and Simon Laursen.
Average-energy games.
In Javier Esparza and Enrico Tronci, editors, Proceedings Sixth International Symposium on Games, Automata, Logics and Formal Verification, GandALF 2015, Genoa, Italy, 21-22nd September 2015., volume 193 of EPTCS, pages 1–15, 2015.

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.

References II

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.

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.

Krishnendu Chatterjee and Vinayak S. Prabhu.
Quantitative timed simulation functions and refinement metrics for real-time systems.
In Calin Belta and Franjo Ivancic, editors, Proceedings of the 16th international conference on Hybrid systems: computation and control, HSCC 2013, April 8-11, 2013, Philadelphia, PA, USA, pages 273–282. ACM, 2013.

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.

References III

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

Wladimir Fridman and Martin Zimmermann.
Playing pushdown parity games in a hurry.
In GandALF 2012, pages 183–196, 2012.

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.

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.

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.

Frank Thuijsman and Okko Jan Vrieze.
The bad match; a total reward stochastic game.
OR Spektrum, 9(2):93–99, 1987.

Igor Walukiewicz.
Pushdown processes: Games and model-checking.
Information and Computation, 164(2):234–263, 2001.

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