Strategy Synthesis for Multi-dimensional Quantitative Objectives

Krishnendu Chatterjee[1] Mickael Randour[2]
Jean-François Raskin[3]

[1]IST Austria [2]UMONS [3]ULB

04.09.2012

CONCUR 2012: 23rd International Conference on Concurrency Theory

UMONS Université de Mons · IST AUSTRIA Institute of Science and Technology · ULB UNIVERSITÉ LIBRE DE BRUXELLES

Aim of this work

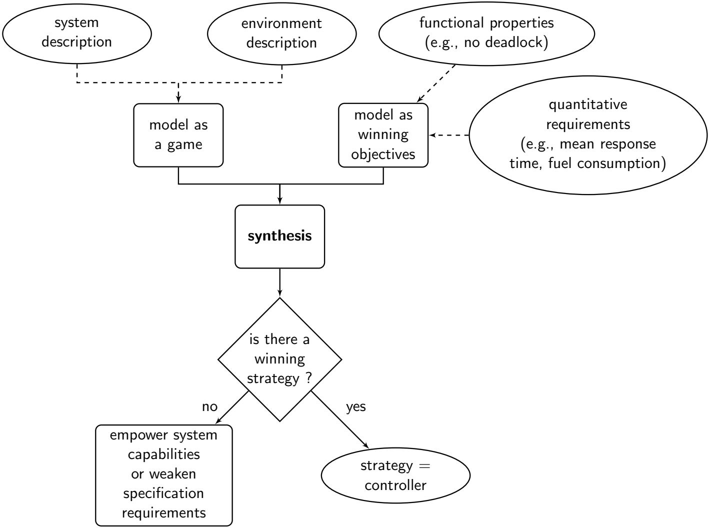

Aim of this work

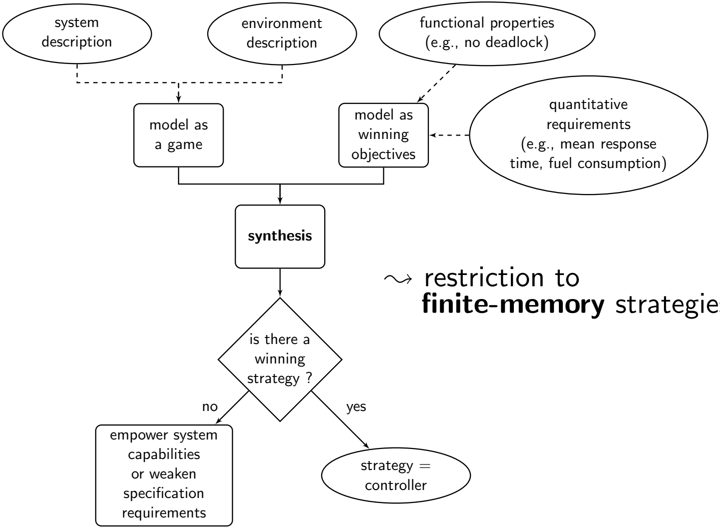

$\rightsquigarrow$ restriction to **finite-memory** strategies.

Aim of this work

- Study games with
  - ▷ multi-dimensional quantitative objectives (energy and mean-payoff)
  - ▷ *and* a parity objective.

$\rightsquigarrow$ First study of such a conjunction.

- Address questions that revolve around *strategies*:
  - ▷ bounds on memory,
  - ▷ synthesis algorithm,
  - ▷ randomness $\stackrel{?}{\sim}$ memory.

Results Overview

- Memory bounds

MEPGs	MMPPGs	
optimal	finite-memory optimal	optimal
exp.	**exp.**	infinite [CDHR10]

- Strategy synthesis (finite memory)

MEPGs	MMPPGs
EXPTIME	**EXPTIME**

- Randomness as a substitute for finite memory

	MEGs	EPGs	MMP(P)Gs	MPPGs
one-player	×	×	√	√
two-player	×	×	×	√

1. Multi energy and mean-payoff parity games
2. Memory bounds
3. Strategy synthesis
4. Randomization as a substitute to finite-memory
5. Conclusion

Turn-based games

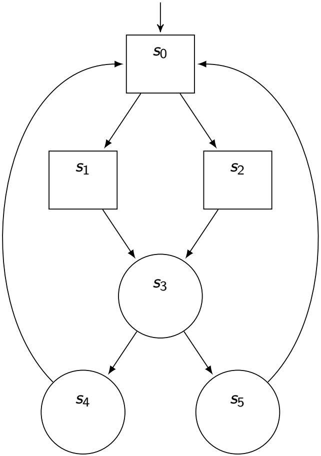

- $G = (S_1, S_2, s_{init}, E)$
- $S = S_1 \cup S_2, S_1 \cap S_2 = \emptyset, E \subseteq S \times S$
- $\mathcal{P}_1$ states = ○
- $\mathcal{P}_2$ states = □
- Plays, prefixes, **pure** strategies.

Integer k-dim. payoff function

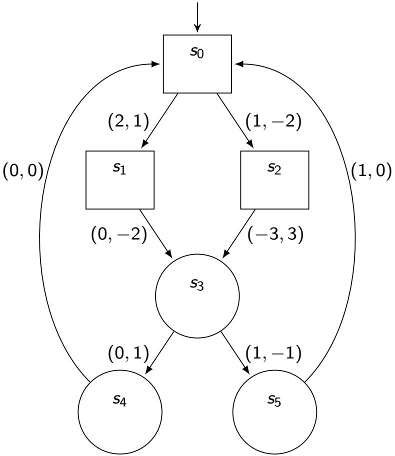

- $G = (S_1, S_2, s_{init}, E, \underline{w})$
- $w : E \to \mathbb{Z}^k$, model changes in quantities
- *Energy level* $\mathrm{EL}(\rho) = v_0 + \sum_{i=0}^{i=n-1} w(s_i, s_{i+1})$
- *Mean-payoff* $\mathrm{MP}(\pi) = \liminf_{n\to\infty} \frac{1}{n}\mathrm{EL}(\pi(n))$

Energy and mean-payoff problems

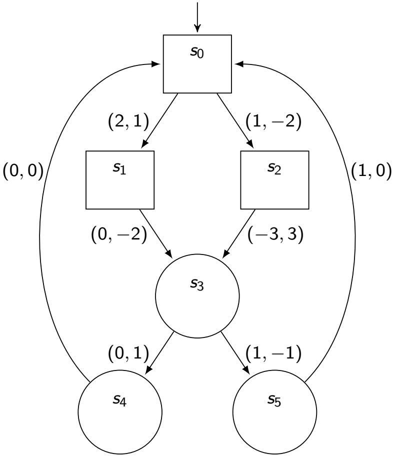

- **Unknown initial credit**

  $\exists?\ v_0 \in \mathbb{N}^k, \lambda_1 \in \Lambda_1$ s.t.

- **Mean-payoff threshold**

  Given $v \in \mathbb{Q}^k$, $\exists?\ \lambda_1 \in \Lambda_1$ s.t.

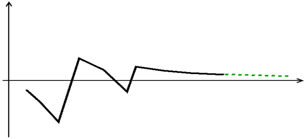

Parity problem

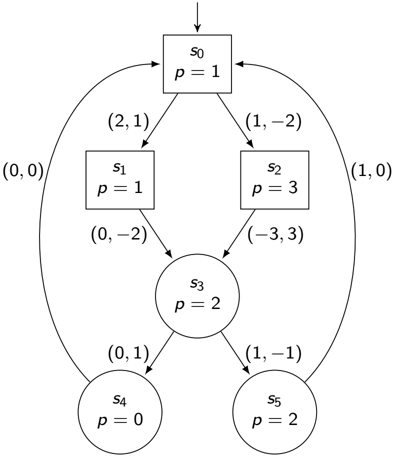

- $G_p = (S_1, S_2, s_{init}, E, w, \underline{p})$
- $p : S \to \mathbb{N}$
- $\mathsf{Par}(\pi) = \min \{p(s) \mid s \in \mathsf{Inf}(\pi)\}$
- **Even parity**

  $\exists? \, \lambda_1 \in \Lambda_1$ s.t. the parity is even

▷ canonical way to express ω-regular objectives

Known results

		Memory ($\mathcal{P}_1$)	Decision problem
Energy	1-dim [CdAHS03, BFL+08]	memoryless	NP ∩ coNP
	k-dim [CDHR10]	finite	coNP-c
	1-dim + parity [CD10]	exponential	NP ∩ coNP
Mean-payoff	1-dim [EM79, LL69]	memoryless	NP ∩ coNP
	k-dim [CDHR10]	infinite	coNP-c (fin.)
	1-dim + parity [CHJ05, BMOU11]	infinite	NP ∩ coNP

Infinite memory?

Example for MMPGs, even with only one player! [CDHR10]

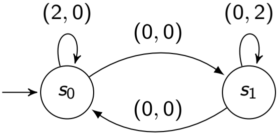

- ▷ To obtain $\mathsf{MP}(\pi) = (1, 1)$ (with $\limsup$, $(2, 2)$!), $\mathcal{P}_1$ has to visit s_0 and s_1 for longer and longer intervals before jumping from one to the other.
- ▷ Any finite-memory strategy alternating between these edges induces an ultimately periodic play s.t. $\mathsf{MP}(\pi) = (x, y)$, $x + y < 2$.

Restriction to finite memory

- Infinite memory:
  - ▷ needed for MMPGs & MPPGs,
  - ▷ practical implementation is unrealistic.

Restriction to finite memory

- Infinite memory:
  - ▷ needed for MMPGs & MPPGs,
  - ▷ practical implementation is unrealistic.
- Finite memory:
  - ▷ preserves game determinacy,
  - ▷ provides equivalence between energy and mean-payoff settings,
  - ▷ the way to go for strategy synthesis.

Obtained results

MEPGs	MMPPGs	
optimal	finite-memory optimal	optimal
exp.	**exp.**	infinite [CDHR10]

By [CDHR10], we only have to consider MEPGs. Recall that the unknown initial credit decision problem for MEGs (without parity) is coNP-complete.

Upper memory bound: even-parity SCTs

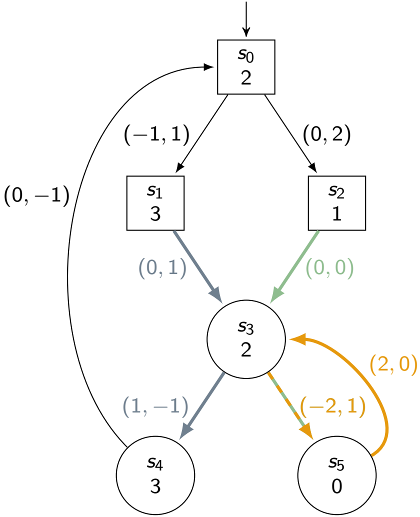

- A winning strategy λ_1 for initial credit $v_0 = (2, 0)$ is
  - $\lambda_1(*s_1 s_3) = s_4$,
  - $\lambda_1(*s_2 s_3) = s_5$,
  - $\lambda_1(*s_5 s_3) = s_5$.

Upper memory bound: even-parity SCTs

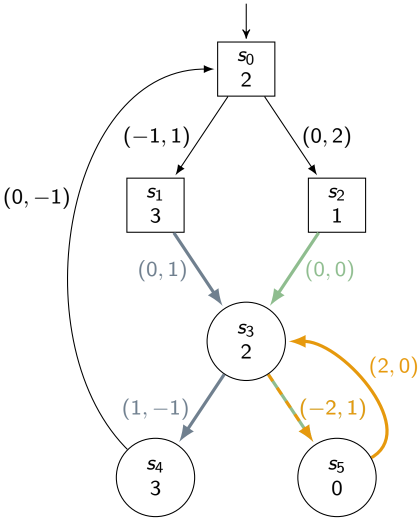

- A winning strategy λ_1 for initial credit $v_0 = (2, 0)$ is
  - $\lambda_1(*s_1 s_3) = s_4$,
  - $\lambda_1(*s_2 s_3) = s_5$,
  - $\lambda_1(*s_5 s_3) = s_5$.
- Lemma: To win, $\mathcal{P}_1$ must be able to enforce positive cycles of even parity.
  - Self-covering paths on VASS [Rac78, RY86].
  - *Self-covering trees (SCTs)* on reachability games over VASS [BJK10].

Upper memory bound: even-parity SCTs

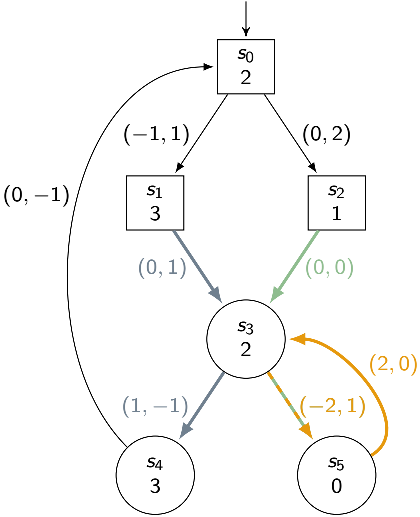

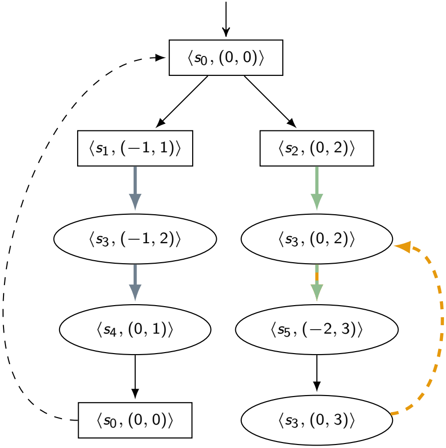

Pebble moves $\Rightarrow$ strategy.

Upper memory bound: even-parity SCTs

$T = (Q, R)$ is an epSCT for s_0, $\Theta : Q \mapsto S \times \mathbb{Z}^k$ is a labeling function.

- Root labeled $\langle s_0, (0, \dots, 0)\rangle$.
- Non-leaf nodes have
  - ▷ unique child if $\mathcal{P}_1$,
  - ▷ all possible children if $\mathcal{P}_2$.
- Leafs have *even-descendance energy ancestors*: ancestors with lower label and minimal priority even on the downward path.

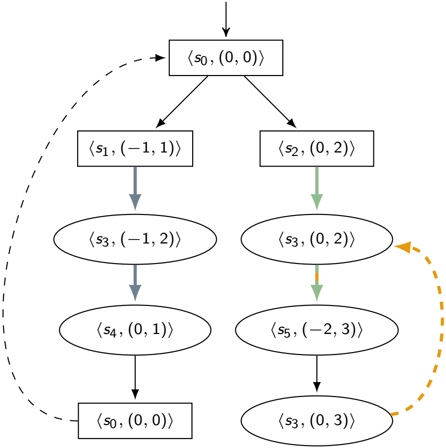

Pebble moves $\Rightarrow$ strategy.

Upper memory bound: SCTs for VASS games

$\mathcal{P}_1$ wins $\Rightarrow \exists$ SCT of depth at most exponential [BJK10].

$\leadsto$ If there exists a winning strategy, there exists a "compact" one.

$\leadsto$ Idea is to eliminate unnecessary cycles.

Limits:

- weights in $\{-1, 0, 1\}$,
- no parity,
- depth only.

Upper memory bound: SCTs for MEGs (no parity)

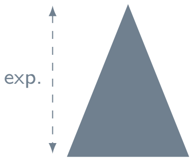

Exp. depth

Upper memory bound: SCTs for MEGs (no parity)

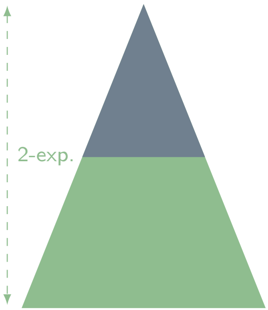

Exp. depth

⇓

Arbitrary weights, no parity

Upper memory bound: SCTs for MEGs (no parity)

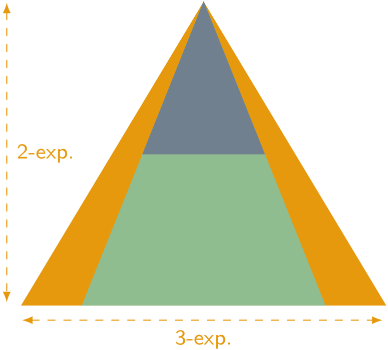

Exp. depth

⇓

Arbitrary weights, no parity

⇓

Width exp. in depth

Upper memory bound: epSCTs for MEPGs

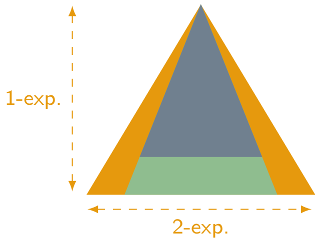

Exp. depth

⇓

Arbitrary weights, ~~no parity~~

$\rightsquigarrow$ *preserve branching, add parity*

⇓

Width exp. in depth

Upper memory bound: epSCTs for MEPGs

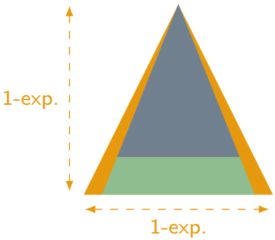

Exp. depth

⇓

Arbitrary weights, ~~no parity~~

$\rightsquigarrow$ *preserve branching, add parity*

⇓

~~Width exp. in depth~~

$\rightsquigarrow$ *encode parity as additionnal energy dimensions*

$\rightsquigarrow$ *merge nodes based on energy levels*

Lower memory bound

Lemma: *There exists a family of multi energy games* $(G(K))_{K \geq 1,} = (S_1, S_2, s_{init}, E, k = 2 \cdot K, w : E \rightarrow \{-1, 0, 1\})$ *s.t. for any initial credit,* $\mathcal{P}_1$ *needs exponential memory to win.*

Lower memory bound

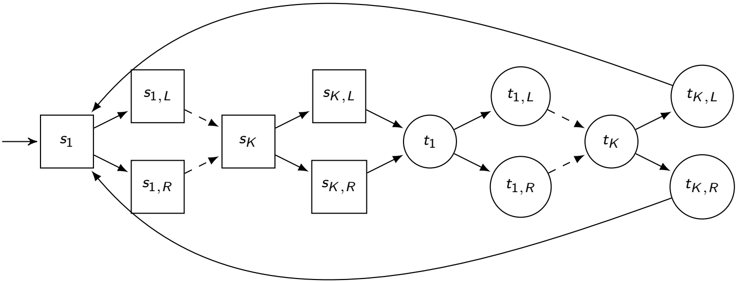

$$\forall\, 1 \leq i \leq K, w((\circ, s_i)) = w((\circ, t_i)) = (0, \ldots, 0),$$

$$w((s_i, s_{i,L})) = -w((s_i, s_{i,R})) = w((t_i, t_{i,L})) = -w((t_i, t_{i,R})),$$

$$\forall\, 1 \leq j \leq k,\ w((s_i, s_{i,L}))(j) = \begin{cases} = 1 \text{ if } j = 2 \cdot i - 1 \\ = -1 \text{ if } j = 2 \cdot i \\ = 0 \text{ otherwise} \end{cases}.$$

Lower memory bound

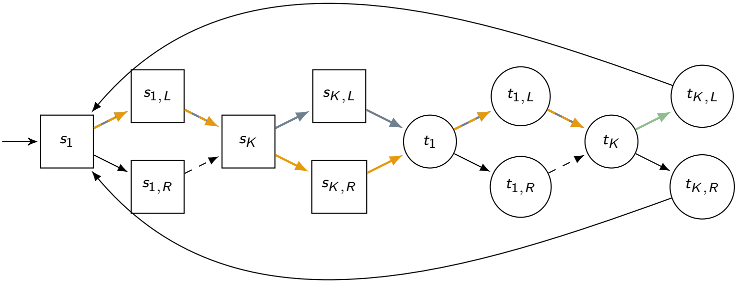

If $\mathcal{P}_1$ plays according to a Moore machine with less than 2^K states, he takes the same decision in some state t_x for the two highlighted prefixes (let $x = K$ w.l.o.g.).

$\Rightarrow \mathcal{P}_2$ can force a decrease by 2 on some dimension every visit.

$\Rightarrow \mathcal{P}_1$ loses for any $v_0 \in \mathbb{N}^k$.

Symbolic synthesis algorithm

Algorithm CpreFP for MEPGs and MMPPGs:

- ▷ symbolic (antichains) and incremental,
- ▷ winning strategy of at most exponential size,
- ▷ worst-case exponential time.

Symbolic synthesis algorithm

Algorithm CpreFP for MEPGs and MMPPGs:

- ▷ symbolic (antichains) and incremental,
- ▷ winning strategy of at most exponential size,
- ▷ worst-case exponential time.

Idea: greatest fixed point of a $\mathsf{Cpre}_{\mathbb{C}}$ operator.

- ▷ Compute for each state the set of winning initial credits, represented by the minimal elements of upper closed sets.
- ▷ Parameter $\mathbb{C}$: range of energy levels to consider.
  $\rightsquigarrow$ incremental, ensures convergence.

Symbolic synthesis algorithm: Cpre

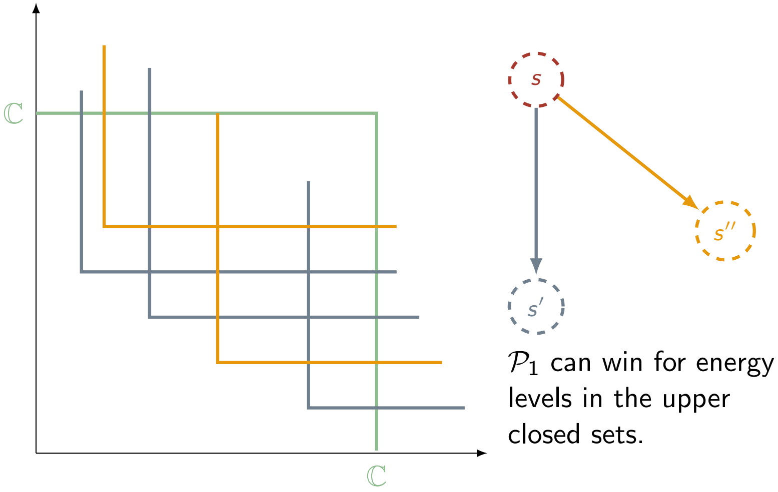

$\mathcal{P}_1$ can win for energy levels in the upper closed sets.

Symbolic synthesis algorithm: Cpre

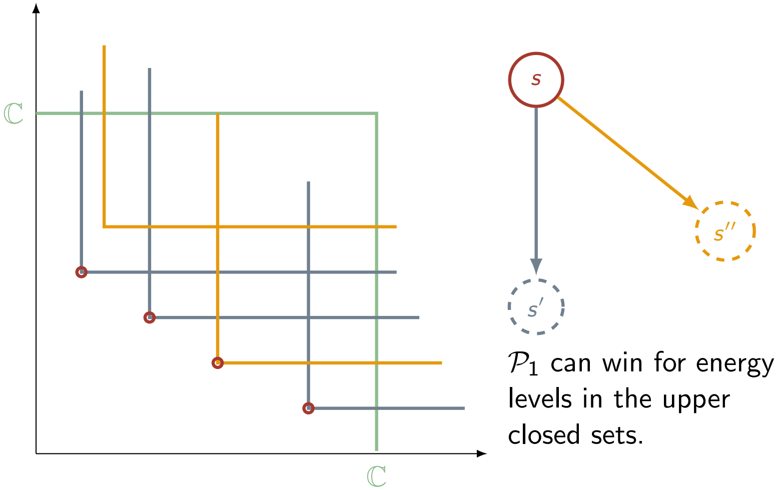

$\mathcal{P}_1$ can win for energy levels in the upper closed sets.

Symbolic synthesis algorithm: Cpre

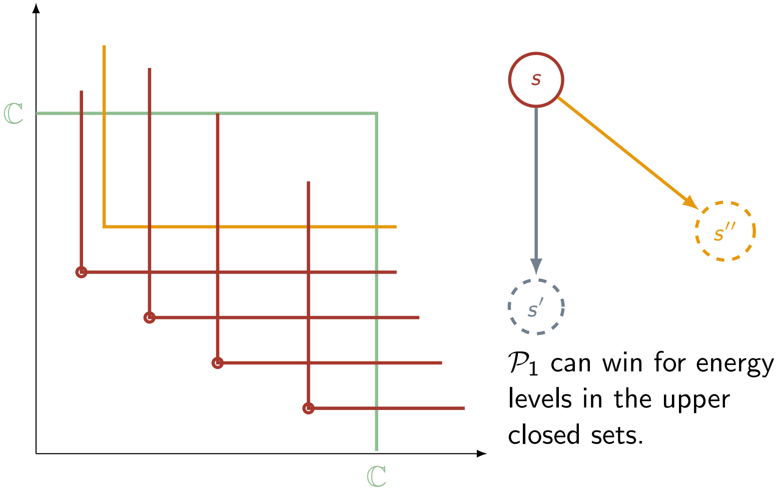

$\mathcal{P}_1$ can win for energy levels in the upper closed sets.

Symbolic synthesis algorithm: Cpre

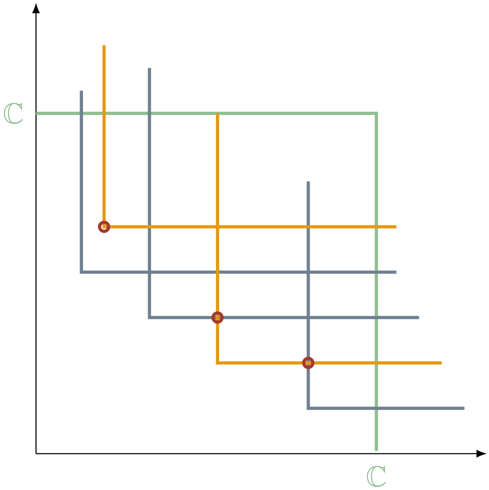

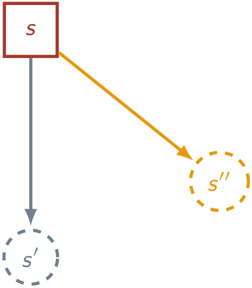

$\mathcal{P}_1$ can win for energy levels in the upper closed sets.

Symbolic synthesis algorithm: Cpre

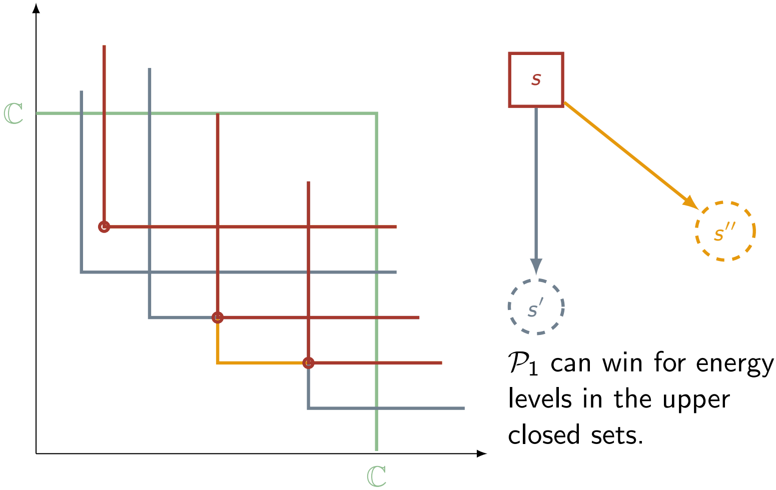

$\mathcal{P}_1$ can win for energy levels in the upper closed sets.

Symbolic synthesis algorithm: CpreFP

- **Correctness**
  - $(s_{init}, (c_1, \ldots, c_k)) \in \mathrm{Cpre}^*_{\mathbb{C}} \rightsquigarrow$ winning strategy for initial credit $(c_1, \ldots, c_k)$.
- **Completeness**
  - Winning strategy and SCT of depth $l \rightsquigarrow$ $(s_{init}, (\mathbb{C}, \ldots, \mathbb{C})) \in \mathrm{Cpre}^*_{\mathbb{C}}$ for $\mathbb{C} = 2 \cdot l \cdot W$ (cf. max init. credit).

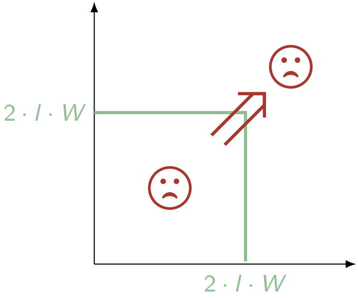

Question

When and how can $\mathcal{P}_1$ trade his pure finite-memory strategy for an equally powerful randomized memoryless one ?

- Sure semantics $\rightsquigarrow$ almost-sure semantics (i.e., probability 1).
- Illustration on **single mean-payoff Büchi games**.

Mean-payoff Büchi games

Remark. MPBGs require infinite memory for optimality.

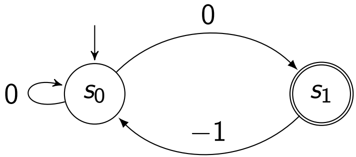

▷ $\mathcal{P}_1$ has to delay his visits of s_1 for longer and longer intervals.

Mean-payoff Büchi games

Remark. MPBGs require infinite memory for optimality.

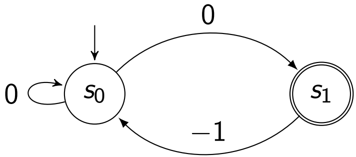

▷ $\mathcal{P}_1$ has to delay his visits of s_1 for longer and longer intervals.

Lemma: *In MPBGs, ε-optimality can be achieved surely by pure finite-memory strategies and almost-surely by randomized memoryless strategies.*

MPBGs: key idea

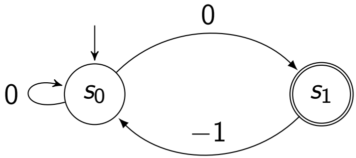

1. Uniform memoryless strategies:

MPBGs: key idea

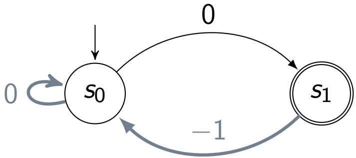

1. Uniform memoryless strategies:
   - λ_1^{gfe} ensures any cycle c has $\mathsf{EL}(c) \geq 0$ [CD10],

MPBGs: key idea

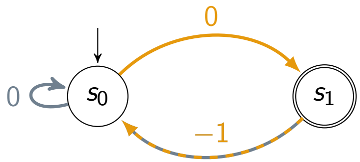

1. Uniform memoryless strategies:
   - λ_1^{gfe} ensures any cycle c has $\mathsf{EL}(c) \geq 0$ [CD10],
   - $\lambda_1^{\Diamond F}$ ensures reaching F in at most n steps (attractor).

MPBGs: key idea

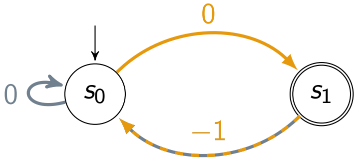

1. Uniform memoryless strategies:
   - λ_1^{gfe} ensures any cycle c has $\mathsf{EL}(c) \geq 0$ [CD10],
   - $\lambda_1^{\Diamond F}$ ensures reaching F in at most n steps (attractor).
2. Alternate using *pure memory* or *probability distributions*.
   - ▷ Frequency of $\lambda_1^{gfe} \rightarrow 1 \Rightarrow \mathsf{MP} \rightarrow \mathsf{MP}^*$.

Obtained results

	MEGs	EPGs	MMP(P)Gs	MPPGs
one-player	$\times$	$\times$	$\surd$	$\surd$
two-player	$\times$	$\times$	$\times$	$\surd$

Conclusion

- Quantitative objectives
- Parity
- Restriction to finite memory (practical interest)
- Exponential memory bounds
- EXPTIME symbolic and incremental synthesis
- Randomness instead of memory

Results Overview

- Memory bounds

MEPGs	MMPPGs	
optimal	finite-memory optimal	optimal
exp.	**exp.**	infinite [CDHR10]

- Strategy synthesis (finite memory)

MEPGs	MMPPGs
EXPTIME	**EXPTIME**

- Randomness as a substitute for finite memory

	MEGs	EPGs	MMP(P)Gs	MPPGs
one-player	×	×	√	√
two-player	×	×	×	√

Thanks. Questions ?

Patricia Bouyer, Ulrich Fahrenberg, Kim Guldstrand Larsen, Nicolas Markey, and Jirí Srba.
Infinite runs in weighted timed automata with energy constraints.
In Proc. of FORMATS, volume 5215 of LNCS, pages 33–47. Springer, 2008.

Tomás Brázdil, Petr Jancar, and Antonín Kucera.
Reachability games on extended vector addition systems with states.
In Proc. of ICALP, volume 6199 of LNCS, pages 478–489. Springer, 2010.

Patricia Bouyer, Nicolas Markey, Jörg Olschewski, and Michael Ummels.
Measuring permissiveness in parity games: Mean-payoff parity games revisited.

In Proc. of ATVA, volume 6996 of LNCS, pages 135–149. Springer, 2011.

Krishnendu Chatterjee and Laurent Doyen.
Energy parity games.
In Proc. of ICALP, volume 6199 of LNCS, pages 599–610. Springer, 2010.

Arindam Chakrabarti, Luca de Alfaro, Thomas A. Henzinger, and Mariëlle Stoelinga.
Resource interfaces.
In Proc. of EMSOFT, volume 2855 of LNCS, pages 117–133. Springer, 2003.

Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger, and Jean-François Raskin.
Generalized mean-payoff and energy games.
In Proc. of FSTTCS, volume 8 of LIPIcs, pages 505–516. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2010.

Krishnendu Chatterjee, Thomas A. Henzinger, and Marcin Jurdzinski.
Mean-payoff parity games.
In Proc. of LICS, pages 178–187. IEEE Computer Society, 2005.

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

T.M. Liggett and S.A. Lippman.
Short notes: Stochastic games with perfect information and time average payoff.
Siam Review, 11(4):604–607, 1969.

Charles Rackoff.
The covering and boundedness problems for vector addition systems.
Theor. Comput. Sci., 6:223–231, 1978.

Louis E. Rosier and Hsu-Chun Yen.
A multiparameter analysis of the boundedness problem for vector addition systems.
J. Comput. Syst. Sci., 32(1):105–135, 1986.

Upper memory bound: SCTs for MEGs (no parity)

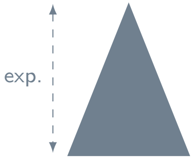

$$w : E \to \{-1, 0, 1\}^k$$
$$l = 2^{(d-1)\cdot|S|} \cdot (|S| + 1)^{c\cdot k^2}$$

Depth bound from [BJK10].

Upper memory bound: SCTs for MEGs (no parity)

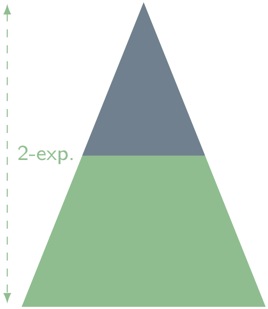

$$w : E \to \{-1, 0, 1\}^k$$
$$l = 2^{(d-1)\cdot|S|} \cdot (|S| + 1)^{c\cdot k^2}$$
$$\Downarrow$$
$w : E \to \mathbb{Z}^k$, W max absolute weight,
V bits to encode W
$$l = 2^{(d-1)\cdot W\cdot|S|} \cdot (W \cdot |S| + 1)^{c\cdot k^2}$$
$$= 2^{(d-1)\cdot 2^V\cdot|S|} \cdot (W \cdot |S| + 1)^{c\cdot k^2}$$

Naive approach: blow-up by W in the size of the state space.

Upper memory bound: SCTs for MEGs (no parity)

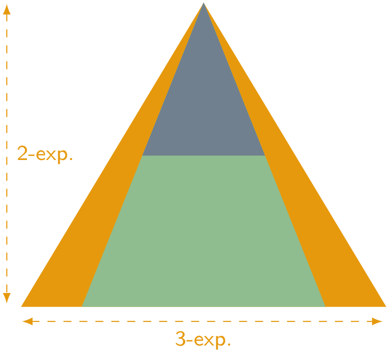

$$w : E \to \{-1, 0, 1\}^k$$
$$l = 2^{(d-1)\cdot|S|} \cdot (|S| + 1)^{c\cdot k^2}$$
$$\Downarrow$$
$w : E \to \mathbb{Z}^k$, W max absolute weight,
V bits to encode W
$$l = 2^{(d-1)\cdot W\cdot|S|} \cdot (W \cdot |S| + 1)^{c\cdot k^2}$$
$$= 2^{(d-1)\cdot 2^V\cdot|S|} \cdot (W \cdot |S| + 1)^{c\cdot k^2}$$
$$\Downarrow$$
Width bounded by $L = d^l$

Naive approach: width increases exponentially with depth.

Upper memory bound: SCTs for MEGs (no parity)

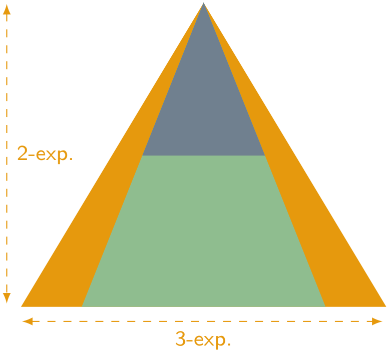

$$w : E \to \{-1, 0, 1\}^k$$
$$l = 2^{(d-1)\cdot|S|} \cdot (|S| + 1)^{c\cdot k^2}$$
$$\Downarrow$$
$w : E \to \mathbb{Z}^k$, W max absolute weight,
V bits to encode W
$$l = 2^{(d-1)\cdot W\cdot|S|} \cdot (W \cdot |S| + 1)^{c\cdot k^2}$$
$$= 2^{(d-1)\cdot 2^V\cdot|S|} \cdot (W \cdot |S| + 1)^{c\cdot k^2}$$
$$\Downarrow$$
Width bounded by $L = d^l$

Naive approach: overall, 3-exp. memory $\leq L \cdot l$, without parity.

Upper memory bound: epSCTs for MEPGs

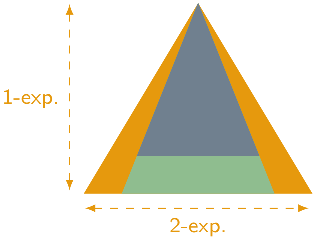

$$w : E \to \{-1, 0, 1\}^k$$
$$l = 2^{(d-1)\cdot|S|} \cdot (|S| + 1)^{c\cdot k^2}$$
$$\Downarrow$$
$w : E \to \mathbb{Z}^k$, W max absolute weight,
$$l = 2^{(d-1)\cdot|S|} \cdot (W \cdot |S| + 1)^{c\cdot k^2}$$

$$\Downarrow$$
Width bounded by $L = d^l$

Refined approach: no blow-up in exponent as branching is preserved, extension to parity.

Upper memory bound: epSCTs for MEPGs

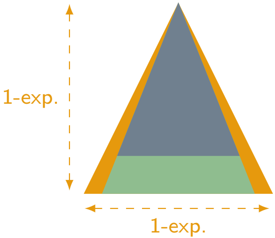

$$w : E \to \{-1, 0, 1\}^k$$
$$l = 2^{(d-1)\cdot|S|} \cdot (|S| + 1)^{c\cdot k^2}$$
$$\Downarrow$$
$w : E \to \mathbb{Z}^k$, W max absolute weight,
$$l = 2^{(d-1)\cdot|S|} \cdot (W \cdot |S| + 1)^{c\cdot k^2}$$

$$\Downarrow$$
Width bounded by $L = |S| \cdot (2 \cdot l \cdot W + 1)^k$

Refined approach: merge equivalent nodes, width is bounded by number of incomparable labels (see next slide).

Upper memory bound: epSCTs for MEPGs

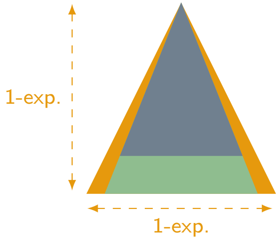

$$w : E \to \{-1, 0, 1\}^k$$
$$l = 2^{(d-1)\cdot|S|} \cdot (|S| + 1)^{c\cdot k^2}$$
$$\Downarrow$$
$w : E \to \mathbb{Z}^k$, W max absolute weight,
$$l = 2^{(d-1)\cdot|S|} \cdot (W \cdot |S| + 1)^{c\cdot k^2}$$

$$\Downarrow$$
Width bounded by $L = |S| \cdot (2 \cdot l \cdot W + 1)^k$

Refined approach: overall, **single exp. memory** $\leq L \cdot l$, for multi energy *along with* parity. Initial credit bounded by $l \cdot W$.

Upper memory bound: from MEPGs to MEGs

- Thanks to the bound on depth for MEPGs, encode parity ($2 \cdot m$ priorities) as m additional energy dimensions.
  - ▷ For each odd priority, add one dimension.
  - ▷ Decrease by 1 when this odd priority is visited.
  - ▷ Increase by l each time a smaller even priority is visited.
- $\mathcal{P}_1$ maintains the energy positive on all additional dimensions iff he wins the original parity objective.

Upper memory bound: merging nodes in SCTs

- Key idea to reduce width to single exp.
  - ▷ $\mathcal{P}_1$ only cares about the energy level.
  - ▷ If he can win with energy v, he can win with energy $\geq v$.

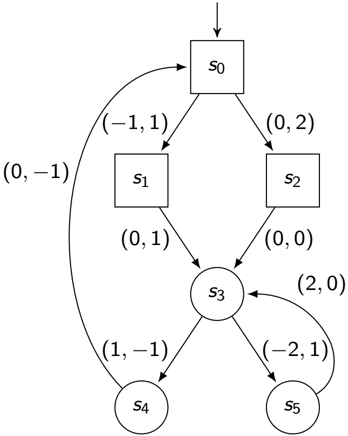

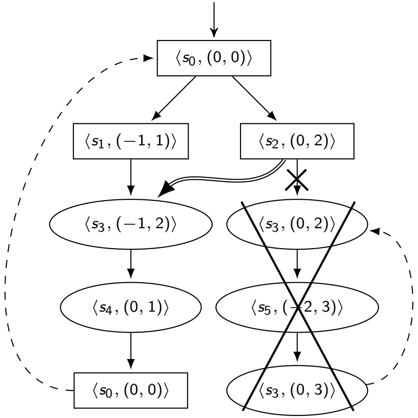

Symbolic synthesis algorithm: Cpre

- $\mathbb{C} = 2 \cdot l \cdot W \in \mathbb{N}$, $U(\mathbb{C}) = (S_1 \cup S_2) \times \{0, 1, \ldots, \mathbb{C}\}^k$,
- $\mathcal{U}(\mathbb{C}) = 2^{U(\mathbb{C})}$, the powerset of $U(\mathbb{C})$,
- $\mathsf{Cpre}_{\mathbb{C}} : \mathcal{U}(\mathbb{C}) \to \mathcal{U}(\mathbb{C})$, $\mathsf{Cpre}_{\mathbb{C}}(V) =$

$$\{(s_1, e_1) \in U(\mathbb{C}) \mid s_1 \in S_1 \wedge \exists (s_1, s) \in E, \exists (s, e_2) \in V : e_2 \leq e_1 + w(s_1, s)\}$$
$$\cup$$
$$\{(s_2, e_2) \in U(\mathbb{C}) \mid s_2 \in S_2 \wedge \forall (s_2, s) \in E, \exists (s, e_1) \in V : e_1 \leq e_2 + w(s_2, s)\}$$

- ▷ Exponential bound on the size of manipulated sets ($\sim$ width).
- ▷ Exponential bound on the number of iterations if a winning strategy exists ($\sim$ depth).

MPBGs: sketch of proof

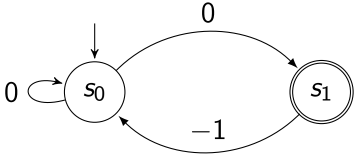

1. Let $G = (S_1, S_2, s_{init}, E, w, F)$, with F the set of Büchi states. Let $n = |S|$. Let Win be the set of winning states for the MPB objective with threshold 0 (w.l.o.g.). For all $s \in$ Win, $\mathcal{P}_1$ has two uniform memoryless strategies λ_1^{gfe} and $\lambda_1^{\Diamond F}$ s.t.

MPBGs: sketch of proof

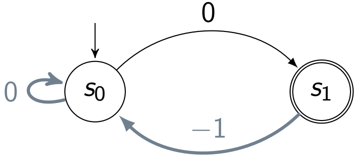

1. Let $G = (S_1, S_2, s_{init}, E, w, F)$, with F the set of Büchi states. Let $n = |S|$. Let Win be the set of winning states for the MPB objective with threshold 0 (w.l.o.g.). For all $s \in \text{Win}$, $\mathcal{P}_1$ has two uniform memoryless strategies λ_1^{gfe} and $\lambda_1^{\Diamond F}$ s.t.
   - λ_1^{gfe} ensures that any cycle c of its outcome has $\text{EL}(c) \geq 0$ [CD10],

MPBGs: sketch of proof

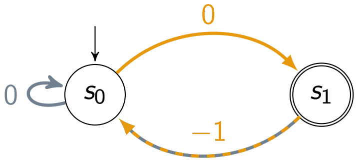

1. Let $G = (S_1, S_2, s_{init}, E, w, F)$, with F the set of Büchi states. Let $n = |S|$. Let Win be the set of winning states for the MPB objective with threshold 0 (w.l.o.g.). For all $s \in$ Win, $\mathcal{P}_1$ has two uniform memoryless strategies λ_1^{gfe} and $\lambda_1^{\Diamond F}$ s.t.

   - λ_1^{gfe} ensures that any cycle c of its outcome has $\mathrm{EL}(c) \geq 0$ [CD10],
   - $\lambda_1^{\Diamond F}$ ensures reaching F in at most n steps, while staying in Win.

MPBGs: sketch of proof

2. For $\varepsilon > 0$, we build a pure finite-memory λ_1^{pf} s.t.
   (a) it plays λ_1^{gfe} for $\dfrac{2 \cdot W \cdot n}{\varepsilon} - n$ steps, then
   (b) it plays $\lambda_1^{\Diamond F}$ for n steps, then again (a).

MPBGs: sketch of proof

2. For $\varepsilon > 0$, we build a pure finite-memory λ_1^{pf} s.t.
   (a) it plays λ_1^{gfe} for $\dfrac{2 \cdot W \cdot n}{\varepsilon} - n$ steps, then
   (b) it plays $\lambda_1^{\Diamond F}$ for n steps, then again (a).

   This ensures that
   - F is visited infinitely often,
   - the total cost of phases (a) + (b) is bounded by $-2 \cdot W \cdot n$, and thus the mean-payoff is at least $-\varepsilon$.

MPBGs: sketch of proof

3. Based on λ_1^{gfe} and $\lambda_1^{\Diamond F}$, we obtain almost-surely ε-optimal *randomized memoryless* strategies, i.e.,

$$\forall\, \varepsilon > 0,\ \exists\, \lambda_1^{rm} \in \Lambda_1^{RM},\ \forall\, \lambda_2 \in \Lambda_2,$$
$$\mathbb{P}_{s_{init}}^{\lambda_1^{rm},\lambda_2}(\mathsf{Par}(\pi) \bmod 2 = 0) = 1\ \wedge\ \mathbb{P}_{s_{init}}^{\lambda_1^{rm},\lambda_2}(\mathsf{MP}(\pi) \geq -\varepsilon) = 1.$$

MPBGs: sketch of proof

3. Based on λ_1^{gfe} and $\lambda_1^{\Diamond F}$, we obtain almost-surely ε-optimal *randomized memoryless* strategies, i.e.,

$$\forall\, \varepsilon > 0,\ \exists\, \lambda_1^{rm} \in \Lambda_1^{RM},\ \forall\, \lambda_2^{pm} \in \Lambda_2^{PM},$$
$$\mathbb{P}_{s_{init}}^{\lambda_1^{rm},\lambda_2^{pm}}(\mathrm{Par}(\pi) \bmod 2 = 0) = 1 \ \wedge\ \mathbb{P}_{s_{init}}^{\lambda_1^{rm},\lambda_2^{pm}}(\mathrm{MP}(\pi) \geq -\varepsilon) = 1.$$

MPBGs: sketch of proof

3. Based on λ_1^{gfe} and $\lambda_1^{\Diamond F}$, we obtain almost-surely ε-optimal *randomized memoryless* strategies, i.e.,

$$\forall\, \varepsilon > 0,\ \exists\, \lambda_1^{rm} \in \Lambda_1^{RM},\ \forall\, \lambda_2^{pm} \in \Lambda_2^{PM},$$
$$\mathbb{P}_{s_{init}}^{\lambda_1^{rm},\lambda_2^{pm}}(\text{Par}(\pi) \bmod 2 = 0) = 1 \ \wedge\ \mathbb{P}_{s_{init}}^{\lambda_1^{rm},\lambda_2^{pm}}(\text{MP}(\pi) \geq -\varepsilon) = 1.$$

Strategy:

$$\forall s \in S,\ \lambda_1^{rm}(s) = \begin{cases} \lambda_1^{gfe}(s) \text{ with probability } 1-\gamma, \\ \lambda_1^{\Diamond F}(s) \text{ with probability } \gamma, \end{cases}$$

for some *well-chosen* $\gamma \in\,]0,1[$.

MPBGs: sketch of proof

Büchi

- Probability of playing as $\lambda_1^{\Diamond F}$ for n steps in a row and ensuring visit of F strictly positive at all times.
- Thus λ_1^{rm} almost-sure winning for the Büchi objective.

MPBGs: sketch of proof

Mean-payoff

- Consider
  - all end components
  - in all MCs induced by pure memoryless strategies of $\mathcal{P}_2$.
- Choose γ so that all ECs have expectation $> -\varepsilon$.
- Put more probability on lengthy sequences of *gfe* edges.