Strategy Synthesis for Quantitative Objectives

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

[1] IST Austria

[2] UMONS

[3] ULB

30.11.2011

GASICS

Aim of this work

- Study games with (multi-dimensional) quantitative objectives: energy and mean-payoff.
- Address questions that revolve around *strategies*:
  - ▷ bounds on memory,
  - ▷ synthesis algorithm,
  - ▷ randomness $\stackrel{?}{\sim}$ memory.

Results Overview

- Strategy synthesis

	MEGs	MMPGs	
	optimal	finite-memory optimal	optimal
Memory	**exp.**	**exp.**	infinite [CDHR10]
Synthesis	**EXPTIME**	**EXPTIME**	/

- Randomness as a substitute for finite-memory

	MEGs	EPGs	MMPGs	MPBGs	MPPGs
1-player	×	×	√	√	√ (conj.)
2-player	×	×	×	√	√ (conj.)

Turn-based games

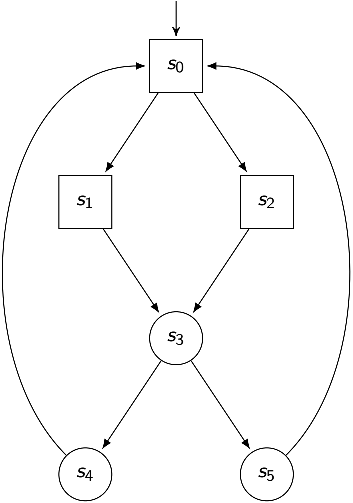

- $G = (S_1, S_2, s_{init}, E)$
  $S = S_1 \cup S_2, S_1 \cap S_2 = \emptyset, E \subseteq S \times S$
- $\mathcal{P}_1$ states $= \bigcirc$
- $\mathcal{P}_2$ states $= \square$

Turn-based games

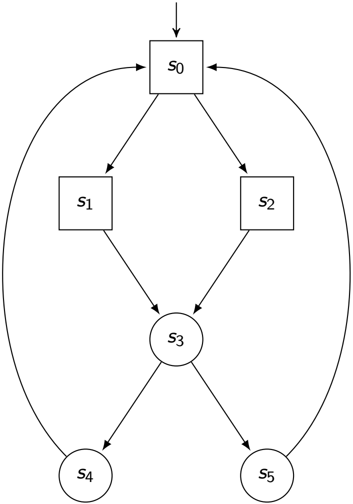

- $G = (S_1, S_2, s_{init}, E)$
  $S = S_1 \cup S_2, S_1 \cap S_2 = \emptyset, E \subseteq S \times S$
- $\mathcal{P}_1$ states = ○
- $\mathcal{P}_2$ states = □
- Play $\pi = s^0 s^1 s^2 \dots s^n \dots$ s.t. $s^0 = s_{init}$
- Prefix $\rho = \pi(n) = s^0 s^1 s^2 \dots s^n$

Pure strategies

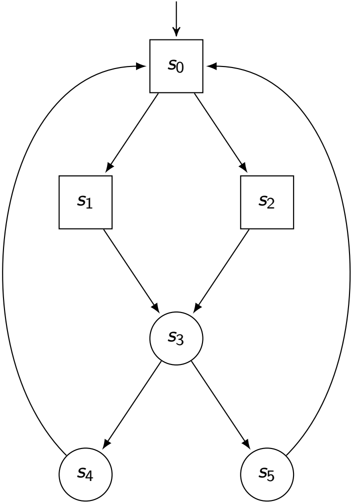

- Pure strategy for $\mathcal{P}_i$

  $\lambda_i \in \Lambda_i : \mathrm{Prefs}_i(G) \to S$ s.t. for all $\rho \in \mathrm{Prefs}_i(G)$, $(\mathrm{Last}(\rho), \lambda_i(\rho)) \in E$
- Memoryless strategy

  $\lambda_i^{pm} \in \Lambda_i^{PM} : S_i \to S$
- Finite-memory strategy

  $\lambda_i^{fm} \in \Lambda_i^{FM} : \mathrm{Prefs}_i(G) \to S$, and can be encoded as a deterministic Moore machine

Integer payoff function

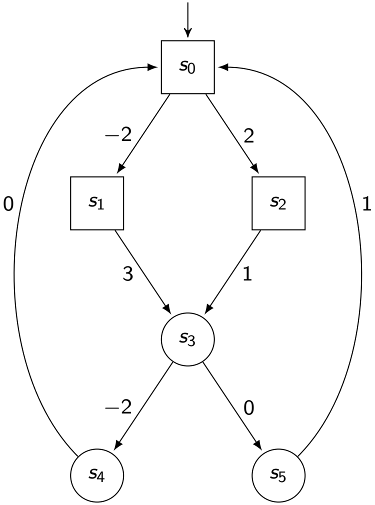

- $G = (S_1, S_2, s_{init}, E, \underline{w})$
- $w : E \to \mathbb{Z}$

Integer payoff function

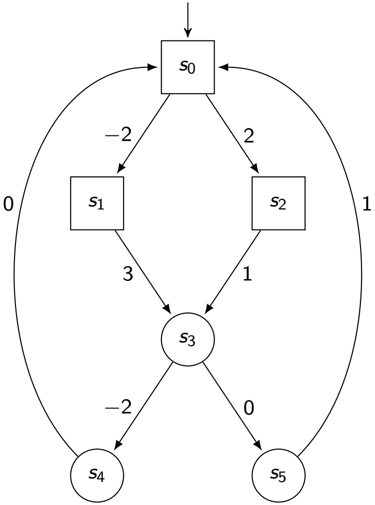

- $G = (S_1, S_2, s_{init}, E, \underline{w})$
- $w : E \to \mathbb{Z}$
- *Energy level* $\mathsf{EL}(\rho) = \sum_{i=0}^{i=n-1} w(s_i, s_{i+1})$
- *Mean-payoff* $\mathsf{MP}(\pi) = \liminf_{n\to\infty} \frac{1}{n}\mathsf{EL}(\pi(n))$

Energy and mean-payoff objectives

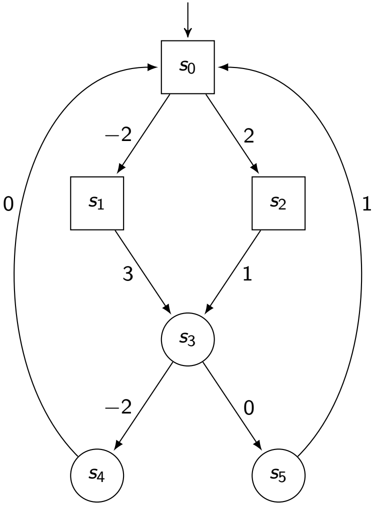

- **Energy objective**

  Given initial credit $v_0 \in \mathbb{N}$,
  $\mathsf{PosEnergy}_G(v_0) = \{\pi \in \mathsf{Plays}(G) \mid \forall\, n \geq 0 : v_0 + \mathsf{EL}(\pi(n)) \in \mathbb{N}\}$

- **Mean-payoff objective**

  Given threshold $v \in \mathbb{Q}$,
  $\mathsf{MeanPayoff}_G(v) = \{\pi \in \mathsf{Plays}(G) \mid \mathsf{MP}(\pi) \geq v\}$

Energy and mean-payoff objectives

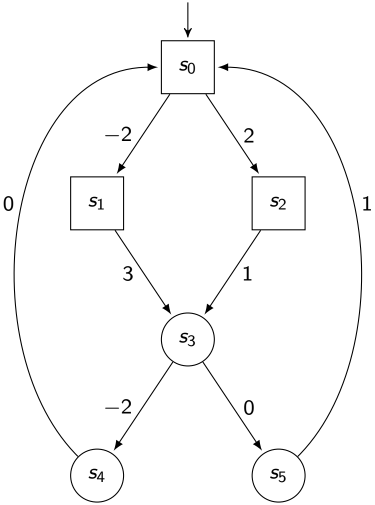

- $\lambda_1(s_3) = s_4$
  - λ_1 wins for $\text{MeanPayoff}_G(\frac{-1}{4})$
  - λ_1 loses for $\text{PosEnergy}_G(v_0)$, for any arbitrary high initial credit
- $\lambda_1(s_3) = s_5$
  - λ_1 wins for $\text{MeanPayoff}_G(\frac{1}{2})$
  - λ_1 wins for $\text{PosEnergy}_G(v_0)$, with $v_0 = 2$

Decision problems

- *Unknown initial credit problem*:

  $\exists?\, v_0 \in \mathbb{N}, \lambda_1 \in \Lambda_1$ s.t. λ_1 wins for $\mathsf{PosEnergy}_G(v_0)$
- *Mean-payoff threshold problem*:

  Given $v \in \mathbb{Q}$, $\exists?\, \lambda_1 \in \Lambda_1$ s.t. λ_1 wins for $\mathsf{MeanPayoff}_G(v)$

MPG threshold v problem equivalent to EG$-v$ unknown initial credit problem [BFL+08].

Complexity of EGs and MPGs

	EGs	MPGs
Memory to win	memoryless [CdAHS03, BFL+08]	memoryless [EM79, LL69]
Decision problem	NP ∩ coNP	NP ∩ coNP

Multi-dimensional weights

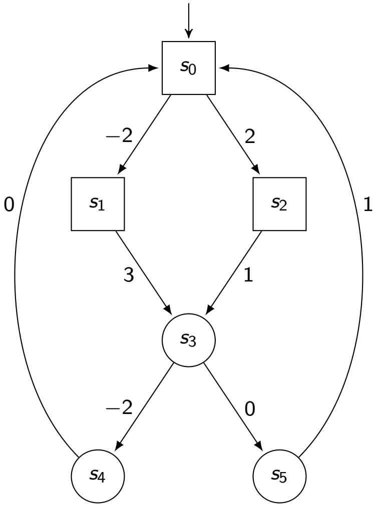

- $G = (S_1, S_2, s_{init}, E, w)$
- $w : E \to \mathbb{Z}$

Multi-dimensional weights

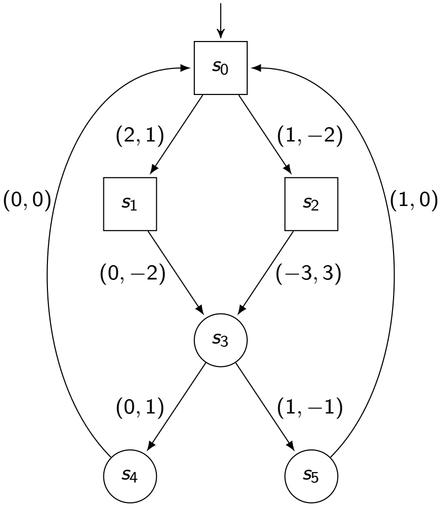

- $G = (S_1, S_2, s_{init}, E, \underline{k}, w)$
- $w : E \to \mathbb{Z}^k$

▷ multiple quantitative aspects

▷ natural extensions of energy and mean-payoff objectives and associated decision problems

MEGs & MMPGs

- Finite memory suffice for MEGs [CDHR10].

MEGs & MMPGs

- Finite memory suffice for MEGs [CDHR10].
- However, infinite memory is needed for MMPGs, even with only one player! [CDHR10]

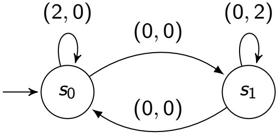

▷ To obtain $\mathsf{MP}(\pi) = (1, 1)$, $\mathcal{P}_1$ has to visit s_0 and s_1 for longer and longer intervals before jumping from one to the other.

▷ Any finite-memory strategy induces an ultimately periodic play s.t. $\mathsf{MP}(\pi) = (x, y)$, $x + y < 2$.

▷ With $\limsup$ as MP the gap is huge : $(2, 2)$.

MEGs & MMPGs

- If players are restricted to finite memory [CDHR10],
  - MEGs and MMPGs are still determined and they are log-space equivalent,
  - the unknown initial credit and the mean-payoff threshold problems are coNP-complete,
  - no clue on memory bounds for $\mathcal{P}_1$ (for $\mathcal{P}_2$, we know it is memoryless).

MEGs & MMPGs

- If players are restricted to finite memory [CDHR10],
  - ▷ MEGs and MMPGs are still determined and they are log-space equivalent,
  - ▷ the unknown initial credit and the mean-payoff threshold problems are coNP-complete,
  - ▷ no clue on memory bounds for $\mathcal{P}_1$ (for $\mathcal{P}_2$, we know it is memoryless).
- Other interesting results on decision problems on MEGs are proved in [FJLS11]. Surprisingly, given a fixed initial vector, the problem becomes EXPSPACE-hard.

Parity

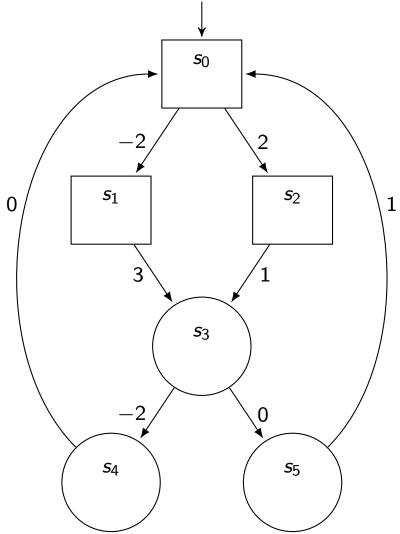

- $G = (S_1, S_2, s_{init}, E, w)$

Parity

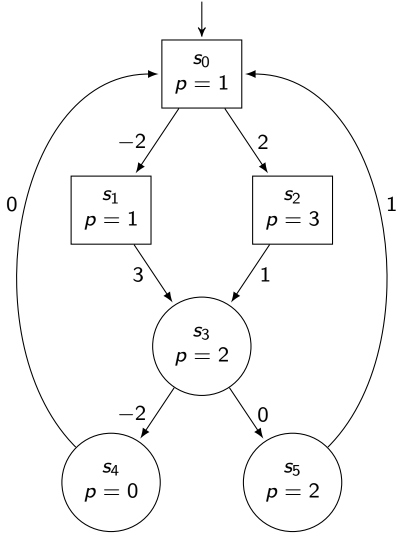

- $G_p = (S_1, S_2, s_{init}, E, w, \underline{p})$
- $p : S \to \mathbb{N}$
- $\mathsf{Par}(\pi) = \min \{p(s) \mid s \in \mathsf{Inf}(\pi)\}$

▷ $\mathsf{Parity}_{G_p} = \{\pi \in \mathsf{Plays}(G_p) \mid \mathsf{Par}(\pi) \bmod 2 = 0\}$

▷ canonical way to express ω-regular objectives

▷ achieve the energy or mean-payoff objective while satisfying the parity condition

Parity

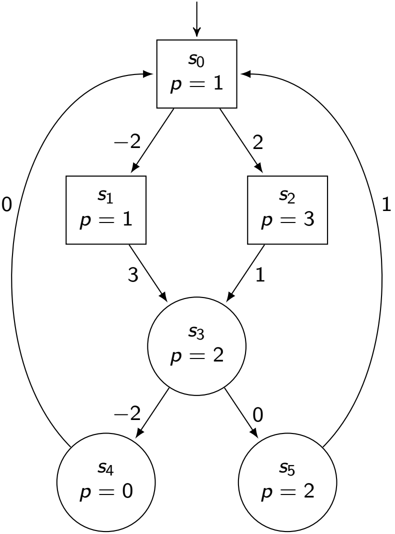

- To win the energy parity objective, $\mathcal{P}_1$ must
  - visit s_4 infinitely often,
  - alternate with visits of s_5 to fund future visits of s_4.

Parity

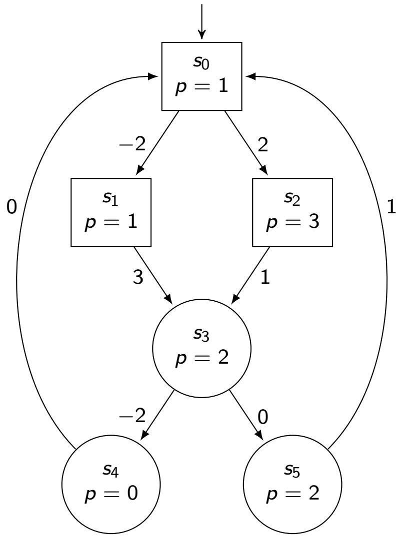

- To win the energy parity objective, $\mathcal{P}_1$ must
  - visit s_4 infinitely often,
  - alternate with visits of s_5 to fund future visits of s_4.
- To achieve optimality for the mean-payoff parity objective, $\mathcal{P}_1$ should wait longer and longer between visits of s_4.

EPGs & MPPGs

- Exponential memory suffice for EPGs and deciding the winner is in NP $\cap$ coNP [CD10].

EPGs & MPPGs

- Exponential memory suffice for EPGs and deciding the winner is in NP ∩ coNP [CD10].
- Infinite memory is needed for MPPGs and deciding the winner is in NP ∩ coNP [CHJ05, BMOU11].

EPGs & MPPGs

- Exponential memory suffice for EPGs and deciding the winner is in NP ∩ coNP [CD10].
- Infinite memory is needed for MPPGs and deciding the winner is in NP ∩ coNP [CHJ05, BMOU11].
- Finite-memory ε-strategies for MPPGs always exist [BCHJ09].
- $\mathcal{P}_1$ has a winning strategy for the MPPG $\langle G, p, w \rangle$ iff $\mathcal{P}_1$ has a winning strategy for the EPG $\langle G, p, w + \varepsilon \rangle$, with $\varepsilon = \frac{1}{|S|+1}$ [CD10].

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.

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.
- **Our goals:**
  - ▷ bounds on memory,
  - ▷ strategy synthesis algorithm,
  - ▷ encoding of memory as randomness.

Obtained results

	MEGs	MMPGs	
	optimal	finite-memory optimal	optimal
Memory	**exp.**	**exp.**	infinite [CDHR10]
Synthesis	**EXPTIME**	**EXPTIME**	/

By [CDHR10], we only have to consider MEGs. Recall that the unknown initial credit decision problem is coNP-complete.

Upper memory bound: SCTs

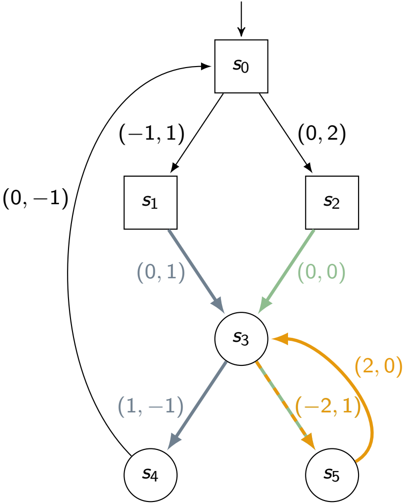

- A winning strategy λ_1 for initial credit $v_0 = (2, 0)$ is
  - $\lambda_1(*s_1s_3) = s_4$,
  - $\lambda_1(*s_2s_3) = s_5$,
  - $\lambda_1(*s_5s_3) = s_5$.

Upper memory bound: SCTs

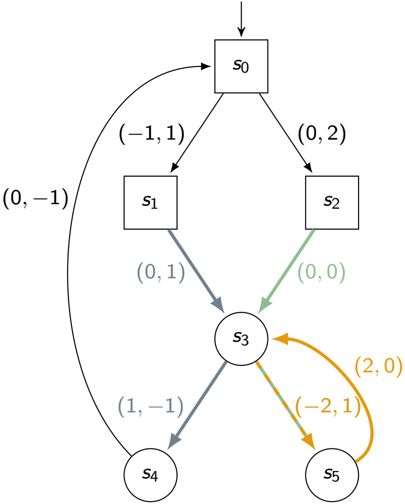

- A winning strategy λ_1 for initial credit $v_0 = (2, 0)$ is
  - $\lambda_1(*s_1 s_3) = s_4$,
  - $\lambda_1(*s_2 s_3) = s_5$,
  - $\lambda_1(*s_5 s_3) = s_5$.
- Lemma: To win, $\mathcal{P}_1$ must be able to enforce positive cycles.
  - Self-covering paths on VASS [Rac78, RY86].
  - *Self-covering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

Upper memory bound: SCTs

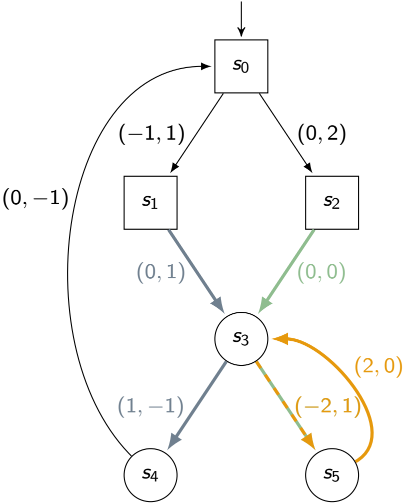

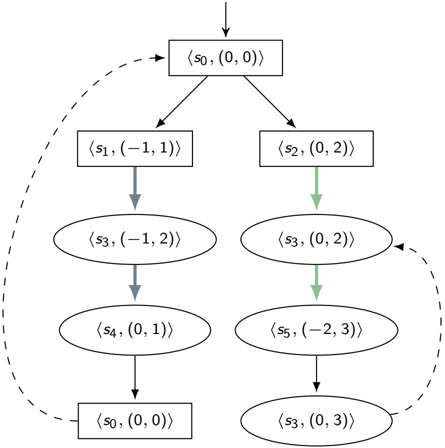

Pebble moves $\Rightarrow$ strategy.

Upper memory bound: SCTs

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

- Root labeled $\langle s_0, (0, \ldots, 0) \rangle$.
- Non-leaf nodes have
  - ▷ unique child if $\mathcal{P}_1$,
  - ▷ all possible children if $\mathcal{P}_2$.
- Leafs have *energy ancestors*: ancestors with lower label.

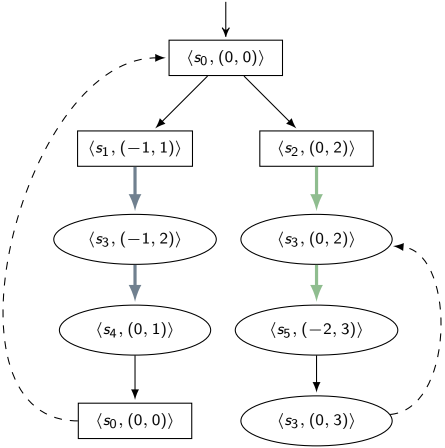

Pebble moves ⇒ strategy.

Upper memory bound: SCTs for VASS games

Theorem (application of [BJK10]): *On a VASS game with weights in $\{-1, 0, 1\}^k$, if state s is winning for $\mathcal{P}_1$, there is a SCT for s whose depth is at most $l = 2^{(d-1)\cdot|S|} \cdot (|S|+1)^{c\cdot k^2}$, with c a constant independent of the considered VASS game and d its branching degree.*

$\rightsquigarrow$ If there exists a winning strategy, there exists a “compact” one.

Upper memory bound: SCTs for MEGs

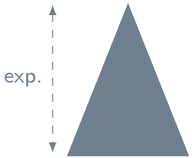

$$w : E \to \{-1, 0, 1\}^k$$
$$l = 2^{(d-1)\cdot|S|} \cdot (|S|+1)^{c\cdot k^2}$$

Depth bound from [BJK10].

Upper memory bound: SCTs for MEGs

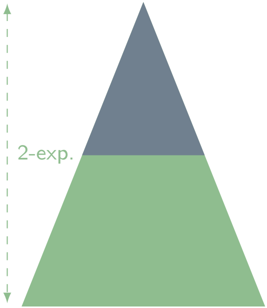

$$w : E \rightarrow \{-1, 0, 1\}^k$$
$$I = 2^{(d-1)\cdot|S|} \cdot (|S| + 1)^{c\cdot k^2}$$
$$\Downarrow$$
$w : E \rightarrow \mathbb{Z}^k$, W max absolute weight, V bits to encode W

$$I = 2^{(d-1)\cdot W \cdot |S|} \cdot (W \cdot |S| + 1)^{c\cdot k^2}$$
$$= 2^{(d-1)\cdot 2^V \cdot |S|} \cdot (W \cdot |S| + 1)^{c\cdot k^2}$$

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

Upper memory bound: SCTs for MEGs

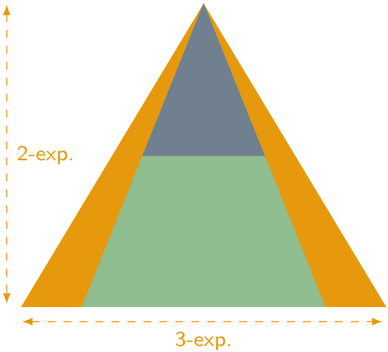

$$w : E \to \{-1, 0, 1\}^k$$
$$l = 2^{(d-1)\cdot|S|} \cdot (|S| + 1)^{c\cdot k^2}$$
$$\Downarrow$$
$w : E \to \mathbb{Z}^k$, W max absolute weight,
V bits to encode W
$$l = 2^{(d-1)\cdot W\cdot|S|} \cdot (W \cdot |S| + 1)^{c\cdot k^2}$$
$$= 2^{(d-1)\cdot 2^V\cdot|S|} \cdot (W \cdot |S| + 1)^{c\cdot k^2}$$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

Naive approach: width increases exponentially with depth.

Upper memory bound: SCTs for MEGs

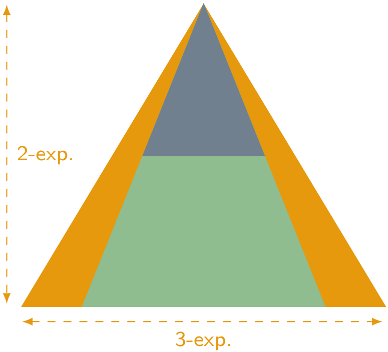

$$w : E \to \{-1, 0, 1\}^k$$
$$l = 2^{(d-1)\cdot|S|} \cdot (|S| + 1)^{c\cdot k^2}$$
$$\Downarrow$$
$w : E \to \mathbb{Z}^k$, W max absolute weight,
V bits to encode W
$$l = 2^{(d-1)\cdot W\cdot|S|} \cdot (W \cdot |S| + 1)^{c\cdot k^2}$$
$$= 2^{(d-1)\cdot 2^V\cdot|S|} \cdot (W \cdot |S| + 1)^{c\cdot k^2}$$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

Naive approach: overall, 3-exp. memory $\leq L \cdot l$.

Upper memory bound: SCTs for MEGs

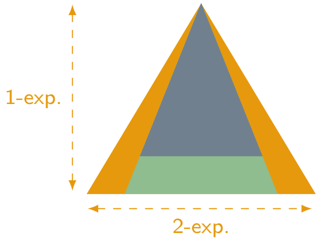

$$w : E \to \{-1, 0, 1\}^k$$
$$l = 2^{(d-1)\cdot|S|} \cdot (|S| + 1)^{c\cdot k^2}$$
$$\Downarrow$$
$w : E \to \mathbb{Z}^k$, W max absolute weight,
$$l = 2^{(d-1)\cdot|S|} \cdot (W \cdot |S| + 1)^{c\cdot k^2}$$
$$\Downarrow$$
Width bounded by $L = d^l$

Refined approach: no blow-up in exponent as branching is preserved.

Upper memory bound: SCTs for MEGs

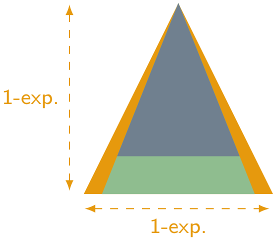

$$w : E \to \{-1, 0, 1\}^k$$
$$l = 2^{(d-1)\cdot|S|} \cdot (|S| + 1)^{c\cdot k^2}$$
$$\Downarrow$$
$w : E \to \mathbb{Z}^k$, W max absolute weight,
$$l = 2^{(d-1)\cdot|S|} \cdot (W \cdot |S| + 1)^{c\cdot k^2}$$

$$\Downarrow$$
Width bounded by
$$L = \binom{2 \cdot l \cdot W + k - 1}{k - 1}$$

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

Upper memory bound: SCTs for MEGs

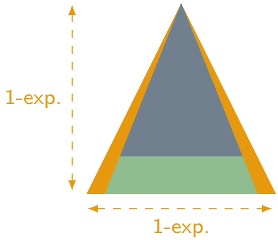

$$w : E \to \{-1, 0, 1\}^k$$
$$l = 2^{(d-1)\cdot|S|} \cdot (|S| + 1)^{c \cdot k^2}$$
$$\Downarrow$$
$w : E \to \mathbb{Z}^k$, W max absolute weight,
$$l = 2^{(d-1)\cdot|S|} \cdot (W \cdot |S| + 1)^{c \cdot k^2}$$

$$\Downarrow$$

Width bounded by
$$L = \binom{2 \cdot l \cdot W + k - 1}{k - 1}$$

Refined approach: overall, **single exp. memory** $\leq L \cdot l$.

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 $\geq v$.

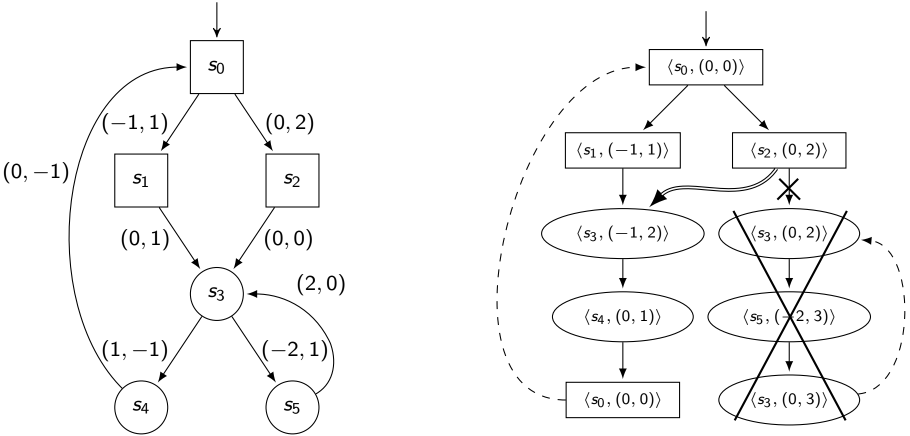

Upper memory bound

Theorem: *The size of memory needed for a finite-memory winning strategy in an energy game $G = (S_1, S_2, s_{init}, E, k, w)$ is upper bounded by an exponential*

$$memSize(|S|, k, d, W) = l \cdot |S| \cdot \binom{2 \cdot l \cdot W + k - 1}{k - 1},$$

with $l = 2^{(d-1)\cdot|S|} \cdot (W \cdot |S| + 1)^{c \cdot k^2}$, d the branching degree of the game, W the largest weight on any edge and c a constant independant of the game.

Note that given l, it is easy to see that the needed initial credit is bounded by $l \cdot W$.

Lower memory bound

Theorem: *There exists a family of games* $(G(K))_{K \geq 1,} = (S_1, S_2, s_{init}, E, k = 2 \cdot K, w)$ *such that for any initial credit,* $\mathcal{P}_1$ *needs exponential memory to win.*

Lower memory bound

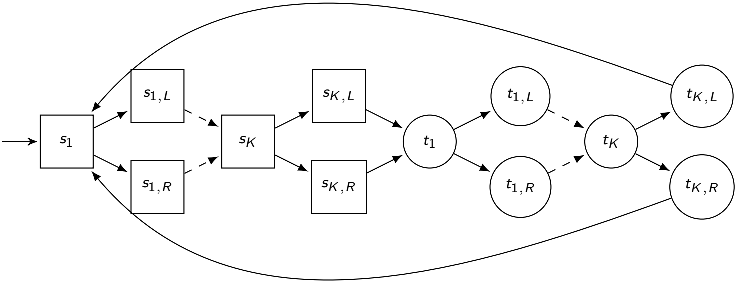

$$\forall\, 1 \leq i \leq K, w((\circ, s_i)) = w((\circ, t_i)) = (0, \ldots, 0),$$

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

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

Lower memory bound

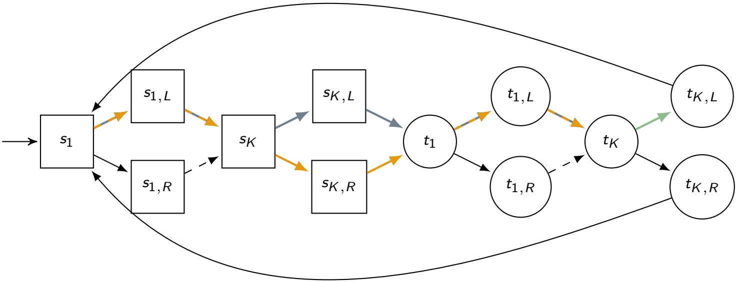

If $\mathcal{P}_1$ plays according to a Moore machine with less than 2^K states, he takes the same decision in some state t_x for the two highlighted prefixes (let $x = K$ w.n.l.o.g.).
$\Rightarrow$ $\mathcal{P}_2$ can alternate and enforce decrease by 1 every two visits $\Rightarrow$ $\mathcal{P}_1$ loses for any $v_0 \in \mathbb{N}^k$.

Symbolic strategy synthesis algorithm

Theorem: *Let $G = (S_1, S_2, s_{init}, E, k, w)$ be a multi energy game. If Player 1 has a winning strategy in G, a Moore machine whose size is at most exponential in G can be constructed in time bounded by an exponential in G.*

Symbolic strategy synthesis algorithm

Theorem: *Let $G = (S_1, S_2, s_{init}, E, k, w)$ be a multi energy game. If Player 1 has a winning strategy in G, a Moore machine whose size is at most exponential in G can be constructed in time bounded by an exponential in G.*

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

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

Symbolic strategy synthesis algorithm: Cpre

- $\mathbb{C} = I \cdot W \in \mathbb{N}$, $U(\mathbb{C}) = (S_1 \cup S_2) \times [0..\mathbb{C}]^k$,
- $\mathcal{U}(\mathbb{C}) = 2^{U(\mathbb{C})}$, the powerset of $U(\mathbb{C})$,
- $\mathsf{Cpre}_{\mathbb{C}} : \mathcal{U}(\mathbb{C}) \rightarrow \mathcal{U}(\mathbb{C})$, $\mathsf{Cpre}_{\mathbb{C}}(V) =$

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

▷ Intuitively, compute for each state the sets of winning initial credits, represented by minimal elements of upper closed sets.

Symbolic strategy synthesis algorithm: Cpre

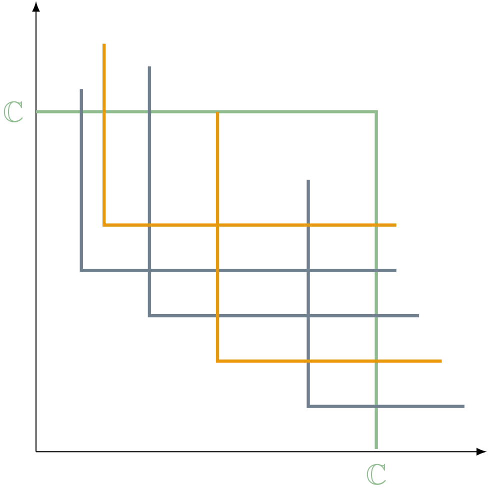

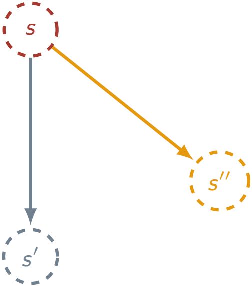

$\mathcal{P}_1$ can win for energy levels in the upper closed sets.

Symbolic strategy synthesis algorithm: Cpre

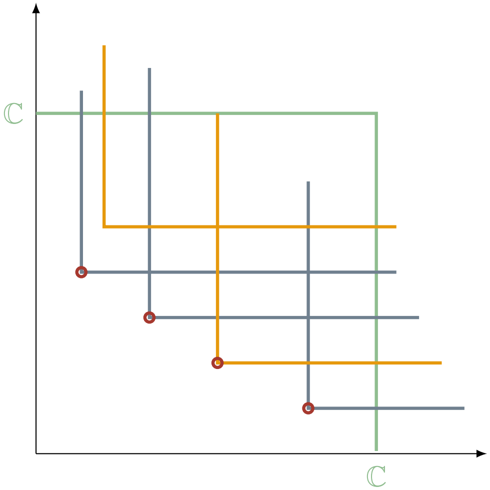

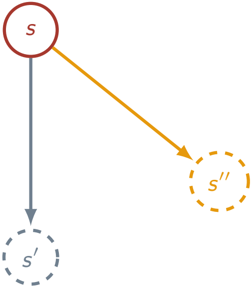

$\mathcal{P}_1$ can win for energy levels in the upper closed sets.

Symbolic strategy synthesis algorithm: Cpre

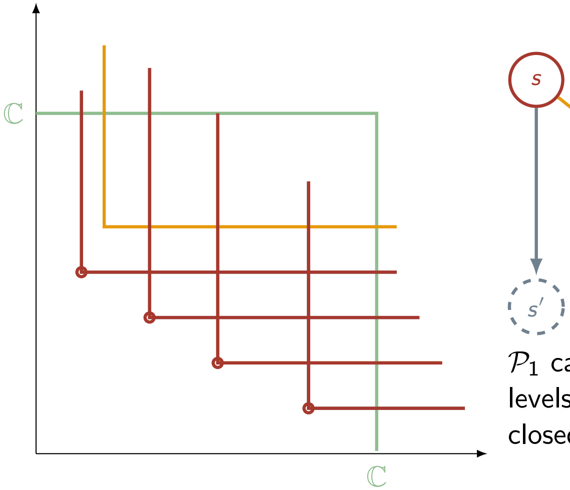

$\mathcal{P}_1$ can win for energy levels in the upper closed sets.

Symbolic strategy synthesis algorithm: Cpre

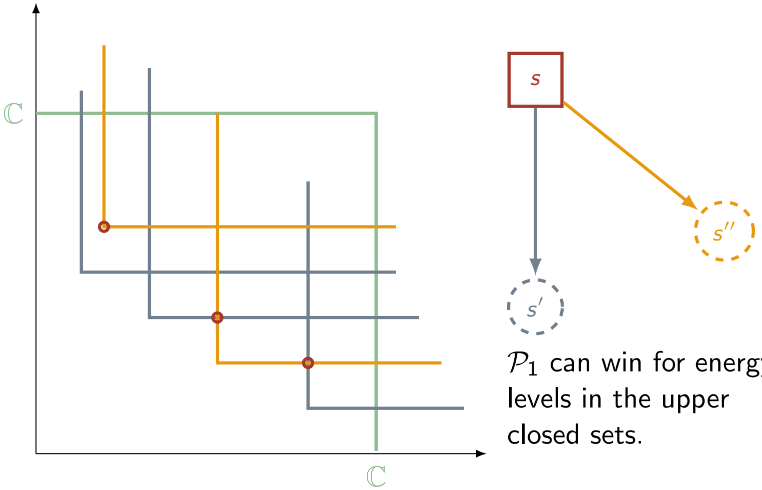

$\mathcal{P}_1$ can win for energy levels in the upper closed sets.

Symbolic strategy synthesis algorithm: Cpre

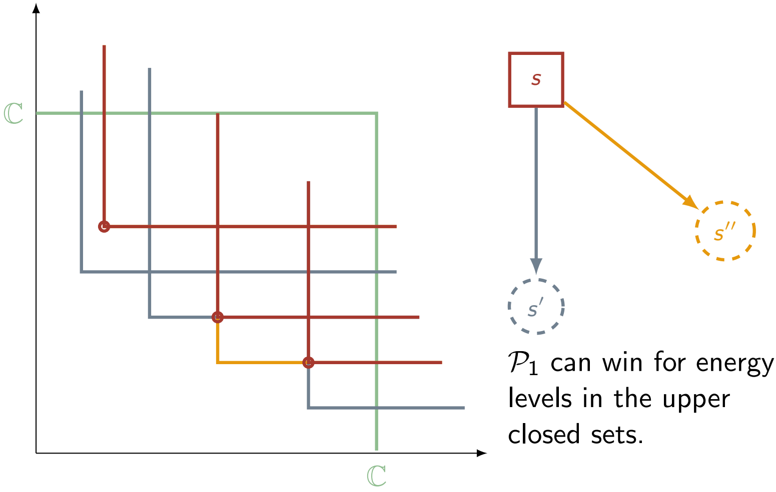

$\mathcal{P}_1$ can win for energy levels in the upper closed sets.

Symbolic strategy synthesis algorithm: Cpre

Lemma: *Let* $G = (S_1, S_2, s_{init}, E, k, w)$ *be a multi energy game, in which all absolute values of weights are bounded by* W*, if Player 1 has a winning strategy in* G *and* $T = (Q, R)$ *is a self-covering tree for* G *of depth* l*, then* $(s_{init}, \mathbb{C}) \in \mathsf{Cpre}^*_{\mathbb{C}}$ *for* $\mathbb{C} = W \cdot l$*.*

Lemma: *Let* $G = (S_1, S_2, s_{init}, E, k, w)$ *be a multi energy game, let* $\mathbb{C} \in \mathbb{N}$*, if there exists* $c \in \mathbb{N}$ *such that* $(s_{init}, c) \in \mathsf{Cpre}^*_{\mathbb{C}}$ *then Player 1 has a winning strategy in* G *for initial credit* c *and a memory used by Player 1 can be bounded by* $|\mathsf{Min}_{\preceq}(\mathsf{Cpre}_{\mathbb{C}})|$*.*

Symbolic strategy synthesis algorithm: Cpre

Lemma: *Let $G = (S_1, S_2, s_{init}, E, k, w)$ be a multi energy game, in which all absolute values of weights are bounded by W, if Player 1 has a winning strategy in G and $T = (Q, R)$ is a self-covering tree for G of depth l, then $(s_{init}, \mathbb{C}) \in \mathsf{Cpre}^*_{\mathbb{C}}$ for $\mathbb{C} = W \cdot l$.*

Lemma: *Let $G = (S_1, S_2, s_{init}, E, k, w)$ be a multi energy game, let $\mathbb{C} \in \mathbb{N}$, if there exists $c \in \mathbb{N}$ such that $(s_{init}, c) \in \mathsf{Cpre}^*_{\mathbb{C}}$ then Player 1 has a winning strategy in G for initial credit c and a memory used by Player 1 can be bounded by $|\mathsf{Min}_{\preceq}(\mathsf{Cpre}_{\mathbb{C}})|$.*

- Incremental approach can be used, by increasing the value of $\mathbb{C}$ inch by inch.
- Efficient implementation seems within reach.

Corollary for MMPGs and summary

Corollary (thanks to [CDHR10]): *Exponential memory is both sufficient and, in general, necessary for finite-memory winning on MMPGs. Finite-memory strategy synthesis is in EXPTIME.*

	MEGs	MMPGs	
	optimal	finite-memory optimal	optimal
Memory	**exp.**	**exp.**	infinite [CDHR10]
Synthesis	**EXPTIME**	**EXPTIME**	/

1. Classical energy and mean-payoff games
2. Extensions to multi-dimensions and parity
3. Strategy synthesis
4. **Randomization as a substitute to finite-memory**
5. Conclusion and ongoing work

Obtained results

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

	MEGs	EPGs	MMPGs	MPBGs	MPPGs
1-player	×	×	√	√	√ (conj.)
2-player	×	×	×	√	√ (conj.)

Obtained results

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

	MEGs	EPGs	MMPGs	**MPBGs**	MPPGs
1-player	$\times$	$\times$	$\surd$	$\boldsymbol{\surd}$	$\surd$ (conj.)
2-player	$\times$	$\times$	$\times$	$\boldsymbol{\surd}$	$\surd$ (conj.)

$\Rightarrow$ Mean-payoff Büchi games.

Probabilistic semantics

- **Büchi**: sure $\rightsquigarrow$ almost-sure.

Probabilistic semantics

- **Büchi**: sure $\rightsquigarrow$ almost-sure.
- **Mean-payoff**:
  - α-*satisfaction*. Given $\alpha \in [0,1]$, $\forall\, \lambda_2 \in \Lambda_2$, $\mathbb{P}^{\lambda_1,\lambda_2}_{s_{init}}(\mathsf{MP}_{\geq v}) \geq \alpha$.
  - β-*expectation*. Given $\beta \in \mathbb{Q}^k$, $\forall\, \lambda_2 \in \Lambda_2$, $\mathbb{E}^{\lambda_1,\lambda_2}_{s_{init}}(\mathsf{MP}) \geq \beta$.
  - 1-satisfaction of $\mathsf{MP}_{\geq v}$ $\Rightarrow$ v-expectation for MP.

Probabilistic semantics

- **Büchi**: sure $\rightsquigarrow$ almost-sure.
- **Mean-payoff**:
  - ▷ *α-satisfaction*. Given $\alpha \in [0, 1]$, $\forall \lambda_2 \in \Lambda_2$, $\mathbb{P}^{\lambda_1,\lambda_2}_{s_{init}}(\mathsf{MP}_{\geq v}) \geq \alpha$.
  - ▷ *β-expectation*. Given $\beta \in \mathbb{Q}^k$, $\forall \lambda_2 \in \Lambda_2$, $\mathbb{E}^{\lambda_1,\lambda_2}_{s_{init}}(\mathsf{MP}) \geq \beta$.
  - ▷ 1-satisfaction of $\mathsf{MP}_{\geq v}$ $\Rightarrow$ v-expectation for MP.

$\Rightarrow$ Almost-sure semantics.

Mean-payoff Büchi games

Remark. MPBGs require infinite memory for optimality.

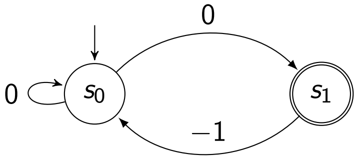

▷ $\mathcal{P}_1$ has to delay his visits of s_1 for longer and longer intervals.

Mean-payoff Büchi games

Remark. MPBGs require infinite memory for optimality.

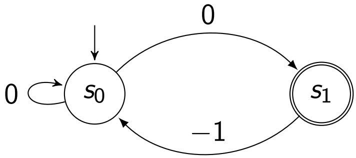

▷ $\mathcal{P}_1$ has to delay his visits of s_1 for longer and longer intervals.

Theorem: *In MPBGs, ε-optimality can be achieved using randomized memoryless strategies, both for satisfaction and expectation semantics.*

MPBGs: sketch of proof

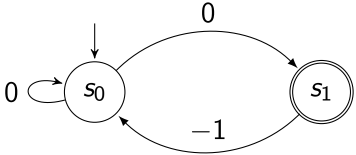

1. Let $G = (S_1, S_2, s_{init}, E, w, F)$, with F the set of Büchi states. Let $n = |S|$. Let Win be the set of winning states for the MPB objective with threshold 0 (w.n.l.o.g.). For all $s \in$ Win, $\mathcal{P}_1$ has two uniform memoryless strategies λ_1^{gfe} and $\lambda_1^{\Diamond F}$ s.t.

MPBGs: sketch of proof

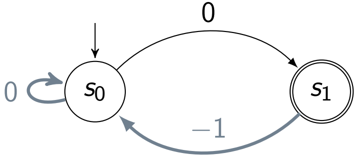

1. Let $G = (S_1, S_2, s_{init}, E, w, F)$, with F the set of Büchi states. Let $n = |S|$. Let Win be the set of winning states for the MPB objective with threshold 0 (w.n.l.o.g.). For all $s \in$ Win, $\mathcal{P}_1$ has two uniform memoryless strategies λ_1^{gfe} and $\lambda_1^{\Diamond F}$ s.t.
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

MPBGs: sketch of proof

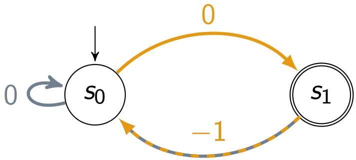

1. Let $G = (S_1, S_2, s_{init}, E, w, F)$, with F the set of Büchi states. Let $n = |S|$. Let Win be the set of winning states for the MPB objective with threshold 0 (w.n.l.o.g.). For all $s \in \text{Win}$, $\mathcal{P}_1$ has two uniform memoryless strategies λ_1^{gfe} and $\lambda_1^{\Diamond F}$ s.t.
   - λ_1^{gfe} ensures that any cycle of its outcome have MP ≥ 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 $\frac{2 \cdot W \cdot n}{\varepsilon} - n$ 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^{pf}, we build a randomized memoryless strategy λ_1^{rm} s.t. in each state,
   (a) it plays as λ_1^{gfe} with probability at least $1 - \dfrac{\varepsilon}{2 \cdot W \cdot n}$,
   (b) it plays as $\lambda_1^{\Diamond F}$ with the remaining probability.

MPBGs: sketch of proof

3. Based on λ_1^{pf}, we build a randomized memoryless strategy λ_1^{rm} s.t. in each state,
   (a) it plays as λ_1^{gfe} with probability at least $1 - \frac{\varepsilon}{2 \cdot W \cdot n}$,
   (b) it plays as $\lambda_1^{\Diamond F}$ with the remaining probability.

   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

- Long-term frequencies of transitions within a given state maintained.
- $\mathcal{P}_2$ may use the same strategy on the graph induced by λ_1^{pf} and the MDP induced by λ_1^{rm} to achieve the same overall transition probabilities.
- Achieving plays π with $\mathsf{MP}(\pi) < -\varepsilon$ with strictly positive probability on the MDP would induce that $\mathcal{P}_2$ can enforce such a play on the graph and lead to contradiction.
- Thus λ_1^{rm} almost-sure winning for the MP objective with threshold $-\varepsilon$.

Summary

	MEGs	EPGs	MMPGs	MPBGs	MPPGs
1-player	$\times$	$\times$	$\surd$	$\surd$	$\surd$ (conj.)
2-player	$\times$	$\times$	$\times$	$\surd$	$\surd$ (conj.)

Conclusion

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

Results Overview

- Strategy synthesis

	MEGs	MMPGs	
	optimal	finite-memory optimal	optimal
Memory	**exp.**	**exp.**	infinite [CDHR10]
Synthesis	**EXPTIME**	**EXPTIME**	/

- Randomness as a substitute for finite-memory

	MEGs	EPGs	MMPGs	MPBGs	MPPGs
1-player	$\times$	$\times$	$\surd$	$\surd$	$\surd$ (conj.)
2-player	$\times$	$\times$	$\times$	$\surd$	$\surd$ (conj.)

Ongoing and future work

- Extend results on MEGs/MMPGs to MEPGs/MMPPGs.
- Consider alternative, more natural definition of MP-like objective, with good synthesis properties.

Thanks. Questions ?

Roderick Bloem, Krishnendu Chatterjee, Thomas A. Henzinger, and Barbara Jobstmann.
Better quality in synthesis through quantitative objectives.
In *Proc. of CAV*, volume 5643 of *LNCS*, pages 140–156. Springer, 2009.

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.

Uli Fahrenberg, Line Juhl, Kim G. Larsen, and Jirí Srba.
Energy games in multiweighted automata.

In *Proc. of ICTAC*, volume 6916 of *LNCS*, pages 95–115. Springer, 2011.

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.