Decidable classes of mean-payoff games with imperfect information

Paul Hunter, Guillermo Pérez, Jean-François Raskin

Université Libre de Bruxelles
Richmodels SVARM Meeting @ Malta

June, 2013

Outline

1. The setting
   - Mean-payoff games
   - MPGs with imperfect information
2. Visible games
3. Pure games
   - Class description
   - Strategy transfer
   - Relevant problems
4. Results
5. Other subclasses
6. Conclusions & future work

MPGs imperfect information: example

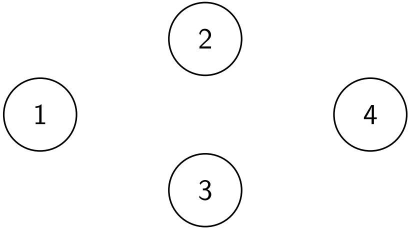

MPGs imperfect information: example

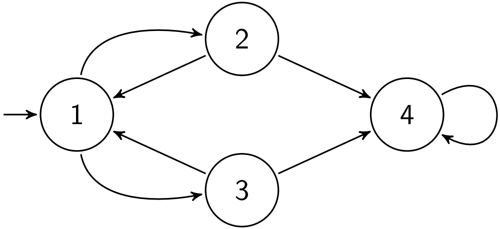

MPGs imperfect information: example

- $\Sigma = \{a, b\}$ and weights on the edges

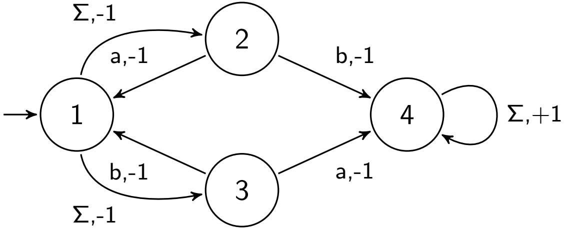

UNIVERSITÉ LIBRE DE BRUXELLES

MPGs imperfect information: example

- $\Sigma = \{a, b\}$ and weights on the edges
- Game
  - to move token: Player 1 chooses σ and Player 2 chooses edge
  - to win (P1): keep average weight of edges traversed above 0

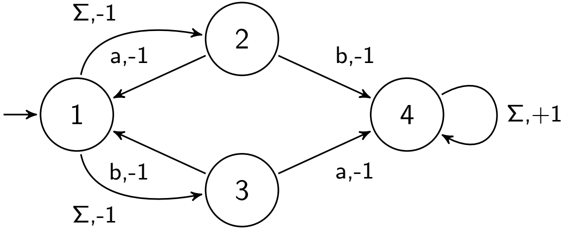

MPGs imperfect information: example

- $\Sigma = \{a, b\}$ and weights on the edges
- Game
  - to move token: Player 1 chooses σ and Player 2 chooses edge
  - to win (P1): keep average weight of edges traversed above 0
- Example: choices = a,$(1, 2)$, payoff = -1

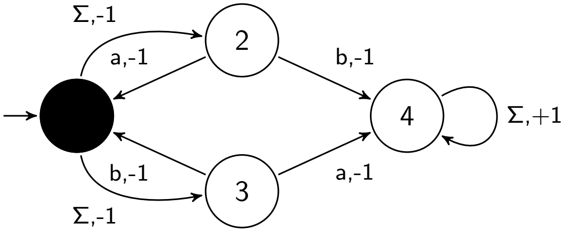

UNIVERSITÉ LIBRE DE BRUXELLES

MPGs imperfect information: example

- $\Sigma = \{a, b\}$ and weights on the edges
- Game
  - to move token: Player 1 chooses σ and Player 2 chooses edge
  - to win (P1): keep average weight of edges traversed above 0
- Example: choices = a,$(1, 2)$, payoff = -1

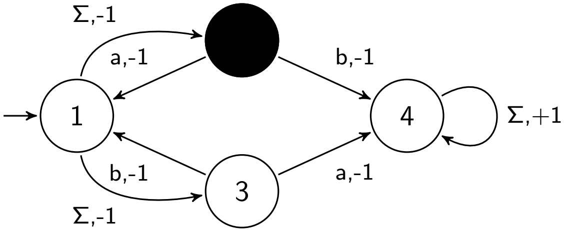

MPGs imperfect information: example

- $\Sigma = \{a, b\}$ and weights on the edges
- Game
  - to move token: Player 1 chooses σ and Player 2 chooses edge
  - to win (P1): keep average weight of edges traversed above 0

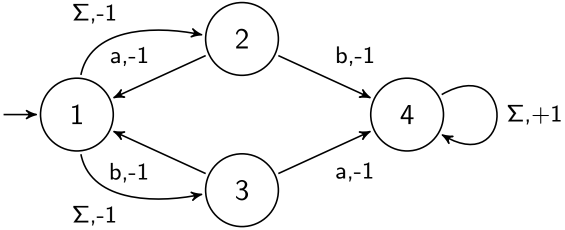

MPGs imperfect information: example

- $\Sigma = \{a, b\}$ and weights on the edges
- Game
  - to move token: Player 1 chooses σ and Player 2 chooses edge
  - to win (P1): keep average weight of edges traversed above 0
- Player 1 only sees colors, Player 2 sees everything

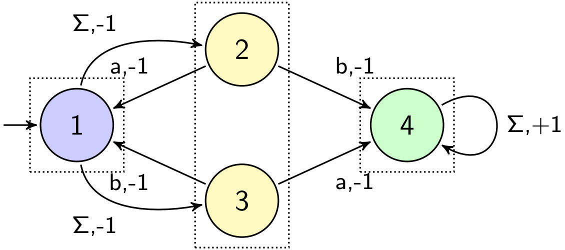

Mean-payoff game

Definition (MPGs)

- Mean-payoff games are 2-player games of infinite duration played on (directed) weighted graphs. ∃ve chooses an action, and ∀dam resolves non-determinism by choosing the next state.
- ∃ve wants to maximize the average weight of the edges traversed (i.e. MP value)
- ∀dam wants to minimize the same value (zero-sum)

UNIVERSITÉ LIBRE DE BRUXELLES

Mean-payoff value

Definition (MP value)

Given the transition relation Δ and the weight function $w : \Delta \mapsto \mathbb{Z}$ of a MPG, the MP value is either:

- $\overline{MP} = lim\ sup_{n \to \infty} \frac{1}{n} \sum_{i=0}^{n-1} w(q_i, \sigma_i, q_{i+1})$ or
- $\underline{MP} = lim\ inf_{n \to \infty} \frac{1}{n} \sum_{i=0}^{n-1} w(q_i, \sigma_i, q_{i+1})$

Problem (Winner of an MPG)

Given a threshold $\nu \in \mathbb{N}$, the MPG is won by $\exists$ve iff $\underline{MP} \geq \nu$. W.l.o.g assume $\nu = 0$.

UNIVERSITÉ LIBRE DE BRUXELLES

MPGs

Theorem

- *MPGs are determined, i.e. if $\exists$ve doesn't have a winning strategy then $\forall$dam does (and viceversa).*
- *Positional strategies suffice for either $\forall$dam or $\exists$ve to win a MPG.*

UNIVERSITÉ LIBRE DE BRUXELLES

MPGs

Theorem

- *MPGs are determined, i.e. if $\exists$ve doesn't have a winning strategy then $\forall$dam does (and viceversa).*
- *Positional strategies suffice for either $\forall$dam or $\exists$ve to win a MPG.*

$\Sigma = \{a, b\}$

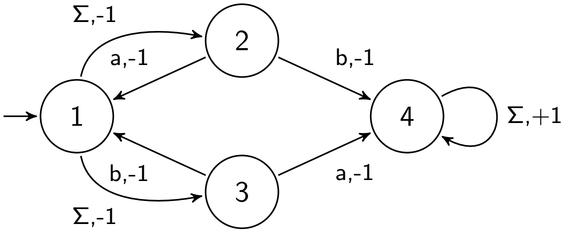

MPGs

Theorem

- *MPGs are determined, i.e. if ∃ve doesn't have a winning strategy then ∀dam does (and viceversa).*
- *Positional strategies suffice for either ∀dam or ∃ve to win a MPG.*

$\Sigma = \{a, b\}$ ∃ve has a winning strat: play b in 2 and a in 3

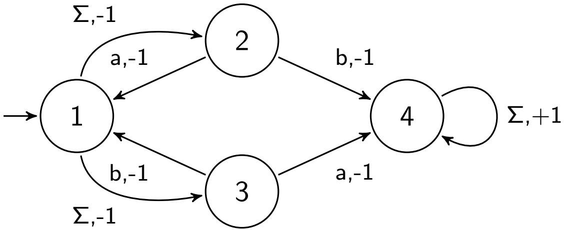

UNIVERSITÉ LIBRE DE BRUXELLES

MPG with imperfect information

Definition (MPGs with II)

A MPG with imperfect information is played on a weighted graph given with a coloring of the state space that defines equivalence classes of indistinguishable states (observations).

UNIVERSITÉ LIBRE DE BRUXELLES

MPG with imperfect information

Definition (MPGs with II)

A MPG with imperfect information is played on a weighted graph given with a coloring of the state space that defines equivalence classes of indistinguishable states (observations).

$\Sigma = \{a, b\}$

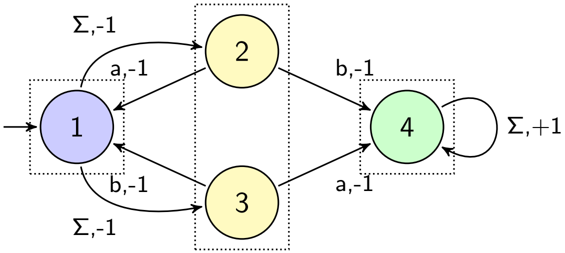

MPG with imperfect information

Definition (MPGs with II)

A MPG with imperfect information is played on a weighted graph given with a coloring of the state space that defines equivalence classes of indistinguishable states (observations).

$\Sigma = \{a, b\}$ Neither $\exists$ve nor $\forall$dam have a winning strategy anymore

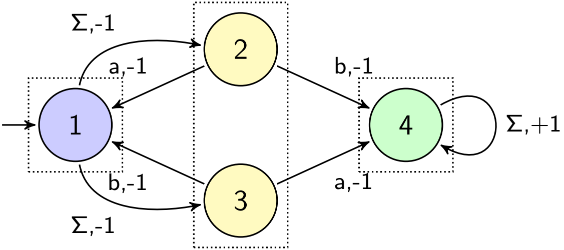

Motivation and properties

Why consider such a model?

- MPGs are natural models for systems where we want to optimize the limit-average usage of a resource.
- Imperfect information arises from the fact that most systems have a limited amount of sensors and input data.

Motivation and properties

Why consider such a model?

- MPGs are natural models for systems where we want to optimize the limit-average usage of a resource.
- Imperfect information arises from the fact that most systems have a limited amount of sensors and input data.

Theorem (Degorre et al. [2010])

- *MPGs with II are no longer "determined"*
- *MPGs with II may require infinite memory to be won by* $\exists ve$
- *The problem on MPGs with II is undecidable*

Visible MPGs with II

Definition (Degorre et al. [2010])

A visible game has weight function w s.t. $w(q_1, \sigma, q_2) = w(q'_1, \sigma', q'_2) = x$ for all transitions $(q_1, \sigma, q_2), (q'_1, \sigma', q'_2) \in \Delta$ having $obs(q_1) = obs(q'_1)$, $obs(q_2) = obs(q'_2)$ and $\sigma = \sigma'$.

UNIVERSITÉ LIBRE DE BRUXELLES

Visible MPGs with II

Definition (Degorre et al. [2010])

A visible game has weight function w s.t. $w(q_1, \sigma, q_2) = w(q'_1, \sigma', q'_2) = x$ for all transitions $(q_1, \sigma, q_2), (q'_1, \sigma', q'_2) \in \Delta$ having $obs(q_1) = obs(q'_1)$, $obs(q_2) = obs(q'_2)$ and $\sigma = \sigma'$.

Theorem (Degorre et al. [2010])

Deciding if $\exists$ve has a winning strategy in a visible MPG with II is EXPTIME-complete.

The knowledge of ∃ve

Definition (Knowledge-based subset construction)

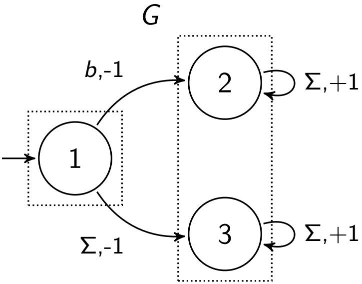

The knowledge of ∃ve

Definition (Knowledge-based subset construction)

- Δ^K based on where ∃ve might be

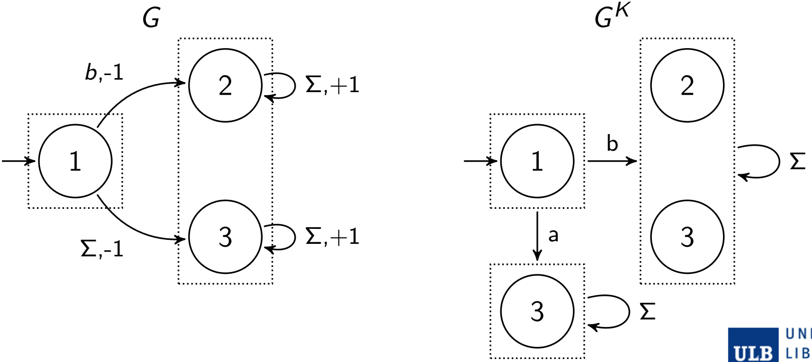

The knowledge of ∃ve

Definition (Knowledge-based subset construction)

- Δ^K based on where ∃ve might be
- w^K makes sense only in the context of visible games

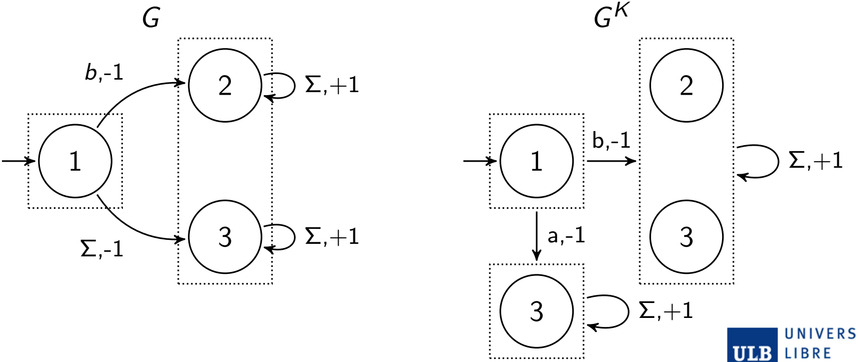

Pure games

Definition

- Simple cycles in G^K are of the form $\rho_K = K_0\sigma_0K_1\sigma_1\cdots K_n$ where $K_0 = K_n$ and $K_i \neq K_j$ for all $0 < i < j < n$.
- Let $\gamma(\rho_K) = \{\pi \mid \pi = q_0\sigma_0q_1\sigma_1\cdots q_n \text{ s.t. } q_i \in K_i \text{ for all } i \geq 0\}$
- A cycle ρ_K is <u>positive</u> if:

$$\forall\pi \in \gamma(\rho_K) : w(\rho) \geq 0$$

Definition

A pure MPG with imperfect information induces a knowledge-based subset construction G^K with all simple cycles being either <u>positive</u> or <u>negative</u>.

UNIVERSITÉ LIBRE DE BRUXELLES

Unfolding a pure game

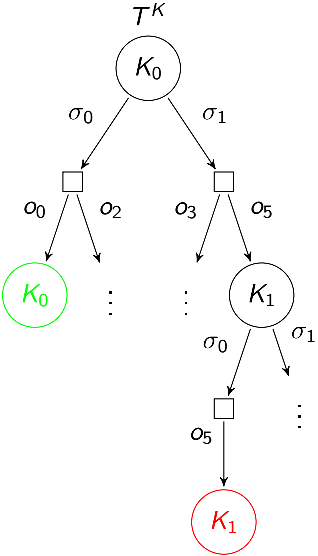

1. unfold G^K, stop when a repeated knowledge set is seen
2. label leaves as good or bad

Unfolding a pure game

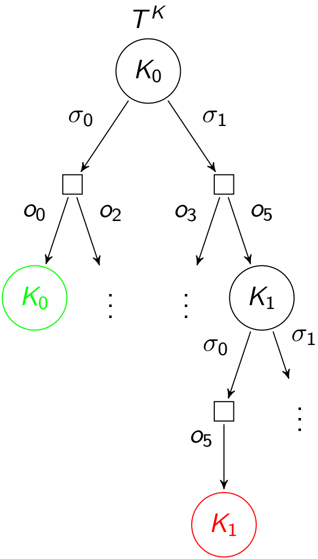

1. unfold G^K, stop when a repeated knowledge set is seen
2. label leaves as good or bad
3. T^K is finite
4. Reachability game on T^K where ∃ve (∀dam) wants to reach good (bad) leaves. Determined!

Strategy transfer

Theorem

If $\exists$*ve has a winning strategy in the reachability game on* T^K *then she also has a winning strategy in* G.

K_0

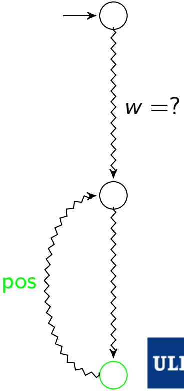

Strategy transfer

Theorem

If $\exists$*ve has a winning strategy in the reachability game on* T^K *then she also has a winning strategy in* G.

K_1
K_0

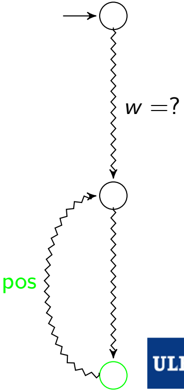

UNIVERSITÉ LIBRE DE BRUXELLES

Strategy transfer

Theorem

If ∃ve has a winning strategy in the reachability game on T^K then she also has a winning strategy in G.

K_1
K_1
K_0

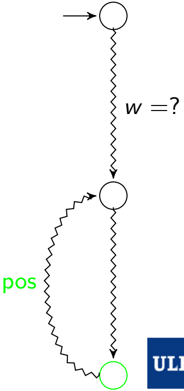

Strategy transfer

Theorem

If ∃*ve has a winning strategy in the reachability game on* T^K *then she also has a winning strategy in* G.

K_0

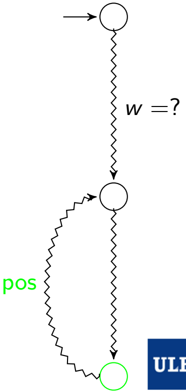

Strategy transfer

Theorem

If ∃*ve has a winning strategy in the reachability game on* T^K *then she also has a winning strategy in* G.

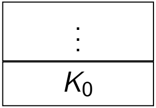

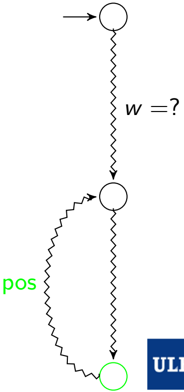

Strategy transfer

Theorem

If ∃*ve has a winning strategy in the reachability game on* T^K *then she also has a winning strategy in* G.

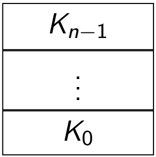

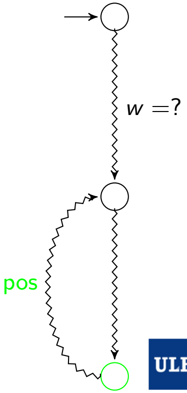

Strategy transfer

Theorem

If ∃ve has a winning strategy in the reachability game on T^K then she also has a winning strategy in G.

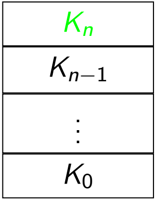

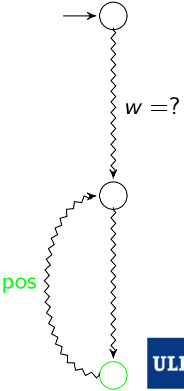

Strategy transfer

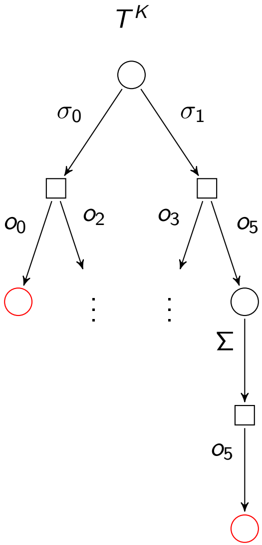

For ∀dam we can only claim a weaker statement:

Theorem

If ∀dam has a winning strategy (WS) in the reachability game on T^K then he can spoil any strategy (S) played by ∃ve .

Strategy transfer

Outcome(T^K ,WS,S)

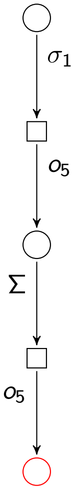

For $\forall$dam we can only claim a weaker statement:

Theorem

If $\forall$dam has a winning strategy (WS) in the reachability game on T^K then he can spoil any strategy (S) played by $\exists$ve .

1. fix S for $\exists$ve and WS for $\forall$dam

Strategy transfer

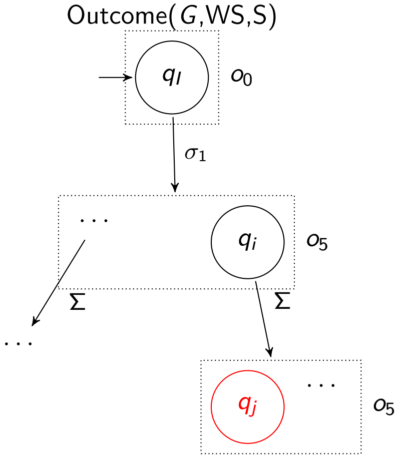

For ∀dam we can only claim a weaker statement:

Theorem

If ∀dam has a winning strategy (WS) in the reachability game on T^K then he can spoil any strategy (S) played by ∃ve .

1. fix S for ∃ve and WS for ∀dam
2. map WS to a quasistrategy in G

Relevant problems

Problem (Deciding the winner)

Does $\exists$ve have a winning strategy in a given pure MPG with II?

Problem (Class membership)

Is a given MPG with II “pure”?

ULB UNIVERSITÉ LIBRE DE BRUXELLES

Does ∃ve win pure G?

Theorem

Deciding if ∃ve has a winning strategy in a given pure MPG with II is EXPTIME-complete.

Proof.

- Hardness follows from pure games being a generalization of visible games.
- For EXPTIME membership we outline an EXPTIME algorithm to decide if ∃ve has a winning strategy in the game.

□

UNIVERSITÉ LIBRE DE BRUXELLES

Solving a pure G

Remark

$\exists$ve being able to avoid bad leaves in T^K implies she has a winning strategy for a stronger condition.

Definition (Energy Games)

$\exists$ve wins an energy game played on a weighted graph if, given an initial credit, she can keep her energy level above zero at all times.

Solving a pure G via a safety game

Definition (Safety game H)

$\mathcal{F}$ is a set of functions $f : Q \mapsto [0, 2W \cdot |Q^K|] \cup \{\bot\}$ which give the current possible states and energy level. $H = \langle\ \mathcal{F}, f_I, \Sigma,\ \Delta^H\ \rangle$ where $f_I(q_I) = W \cdot |Q^K|$ and $f_I(q) = \bot$ for all $q \neq q_I$. All functions with $f_i(q) = 0$, for some $q \in Q$ are not safe.

H

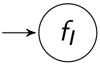

UNIVERSITÉ LIBRE DE BRUXELLES

Solving a pure G via a safety game

Definition (Safety game H)

$\mathcal{F}$ is a set of functions $f : Q \mapsto [0, 2W \cdot |Q^K|] \cup \{\bot\}$ which give the current possible states and energy level. $H = \langle\ \mathcal{F}, f_I, \Sigma,\ \Delta^H\ \rangle$ where $f_I(q_I) = W \cdot |Q^K|$ and $f_I(q) = \bot$ for all $q \neq q_I$. All functions with $f_i(q) = 0$, for some $q \in Q$ are not safe.

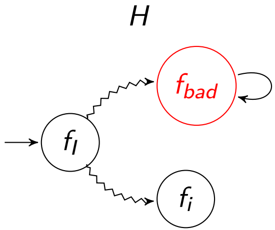

Solving a pure G via a safety game

Definition (Safety game H)

$\mathcal{F}$ is a set of functions $f : Q \mapsto [0, 2W \cdot |Q^K|] \cup \{\bot\}$ which give the current possible states and energy level. $H = \langle\ \mathcal{F}, f_I, \Sigma,\ \Delta^H\ \rangle$ where $f_I(q_I) = W \cdot |Q^K|$ and $f_I(q) = \bot$ for all $q \neq q_I$. All functions with $f_i(q) = 0$, for some $q \in Q$ are not safe.

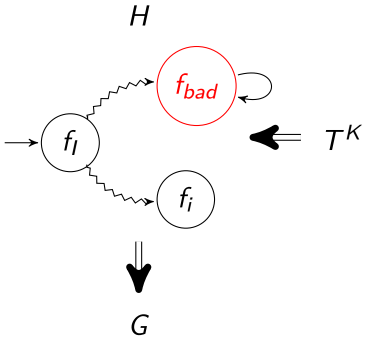

Solving a pure G via a safety game

Definition (Safety game H)

$\mathcal{F}$ is a set of functions $f : Q \mapsto [0, 2W \cdot |Q^K|] \cup \{\bot\}$ which give the current possible states and energy level. $H = \langle\, \mathcal{F}, f_I, \Sigma, \Delta^H \,\rangle$ where $f_I(q_I) = W \cdot |Q^K|$ and $f_I(q) = \bot$ for all $q \neq q_I$. All functions with $f_i(q) = 0$, for some $q \in Q$ are not safe.

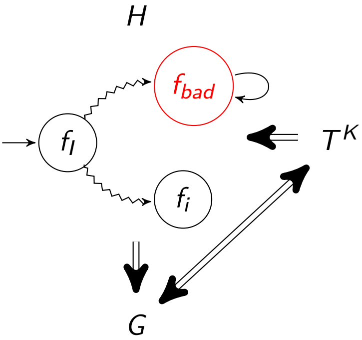

Is G pure?

Theorem

Deciding if an MPG with II is pure is coNEXPTIME-complete.

Proof.

- Membership is straightforward: non-deterministically guess a cycle in G^K, check that it is a simple cycle and that it is neither positive nor negative.
- For hardness we reduce from the SUCCINCT HAMILTONIAN-CYCLE problem.

UNIVERSITÉ LIBRE DE BRUXELLES

SUCCINCT HAM-CYCLE

Definition (Galperin and Wigderson [1983])

$G = \langle V, E \rangle$ with $m \geq 2^n$ vertices, each labelled with a distinct n-bit string. A circuit C_G receives two n-bit inputs and outputs 1 if there is an edge. C_G has $r = O(n^k)$ gates.

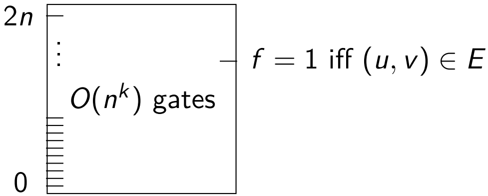

SUCCINCT HAM-CYCLE

Theorem (Exponential blow-up)

Most problems (reducible as a "projection") have an exponential blow-up when the graph is represented succinctly. SUCCINCT HAM-CYCLE is NEXPTIME-complete.

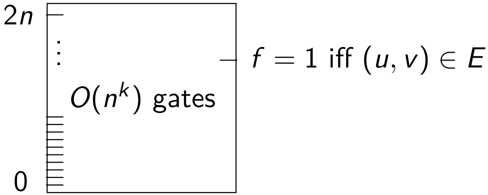

Is G pure?

Theorem

Deciding if an MPG with II is pure is coNEXPTIME-complete.

Proof.

- Membership is straightforward: non-deterministically guess a cycle in G^K, check that it is a simple cycle and that it is neither positive nor negative.
- For hardness we reduce from the SUCCINCT HAMILTONIAN-CYCLE problem.
  - "mixed" simple cycle $\Rightarrow$ simulated 2^N transitions

□

ULB UNIVERSITÉ LIBRE DE BRUXELLES

coNEXP-hardness proof

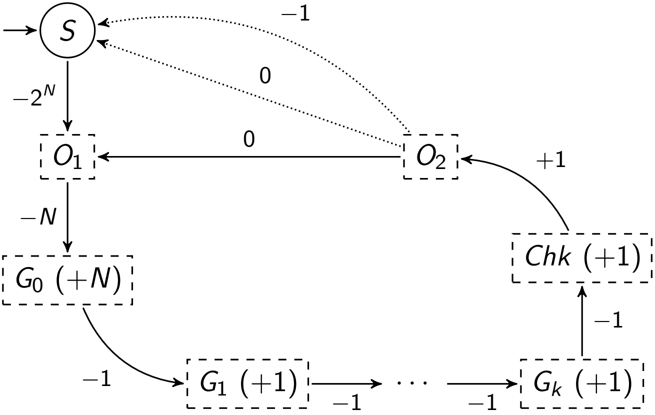

Relaxing the restriction

More general subclasses of MPGs with II for which...

1. $\exists$ve can still force good leaves in T^K, even if unfolding G^K yields good, bad and "undecided" leaves.
   - Class membership: NEXP-h

Relaxing the restriction

More general subclasses of MPGs with II for which...

1. ∃ve can still force good leaves in T^K, even if unfolding G^K yields good, bad and “undecided” leaves.
   - Class membership: NEXP-h
2. all cycles in G^K become positive or negative after unfolding them finitely many times.
   - Class membership: ??

UNIVERSITÉ LIBRE DE BRUXELLES

Relaxing the restriction

More general subclasses of MPGs with II for which...

1. ∃ve can still force good leaves in T^K, even if unfolding G^K yields good, bad and "undecided" leaves.
   - Class membership: NEXP-h
2. all cycles in G^K become positive or negative after unfolding them finitely many times.
   - Class membership: ??
3. the root of T^K can be considered "good" without having all cycles being eventually positive or negative.
   - Class membership: Undecidable

UNIVERSITÉ LIBRE DE BRUXELLES

Wrap-up

We have...

1. A class for which deciding if $\exists$ve has a winning strategy is EXPTIME-c and determining class membership is coNEXPTIME-c
2. Considered extensions of this class and the related problems.

Wrap-up

We have...

1. A class for which deciding if ∃ve has a winning strategy is EXPTIME-c and determining class membership is coNEXPTIME-c
2. Considered extensions of this class and the related problems.

We're still working on...

1. Applying window MPG objectives to the II setting.
2. Games with bounded imperfect information. [Puchala and Rabinovich, 2010]
3. Related interesting subclasses of MPGs with imperfect information.

UNIVERSITÉ LIBRE DE BRUXELLES

References I

Degorre, A., Doyen, L., Gentilini, R., Raskin, J.-F., and Torunczyk, S. (2010). Energy and mean-payoff games with imperfect information. In Dawar, A. and Veith, H., editors, CSL, volume 6247 of Lecture Notes in Computer Science, pages 260–274. Springer.

Galperin, H. and Wigderson, A. (1983). Succinct representations of graphs. Information and Control, 56(3):183–198.

Puchala, B. and Rabinovich, R. (2010). Parity games with partial information played on graphs of bounded complexity. In Mathematical Foundations of Computer Science 2010, pages 604–615. Springer.

coNEXP-hardness proof

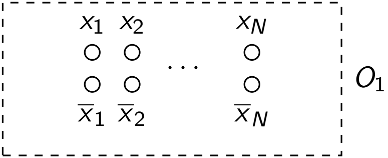

coNEXP-hardness proof

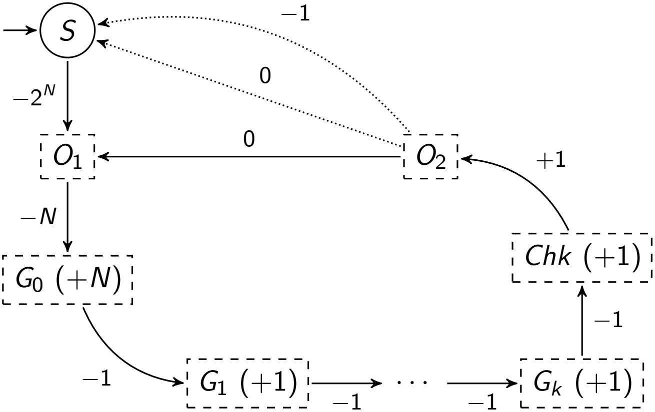

coNEXP-hardness proof

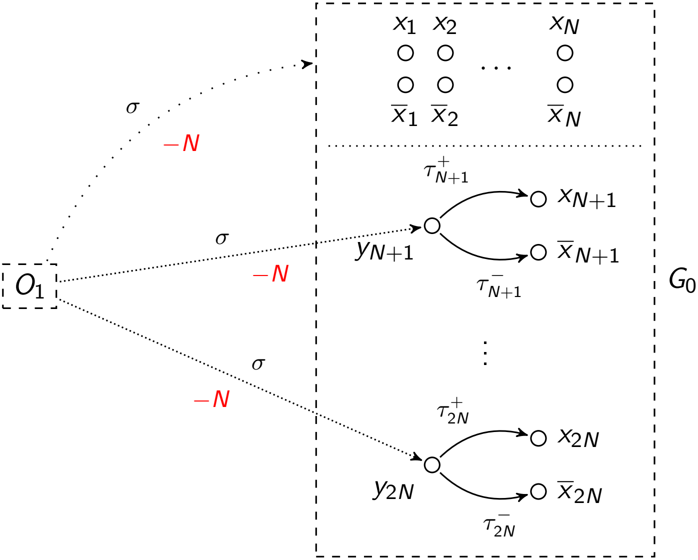

coNEXP-hardness proof

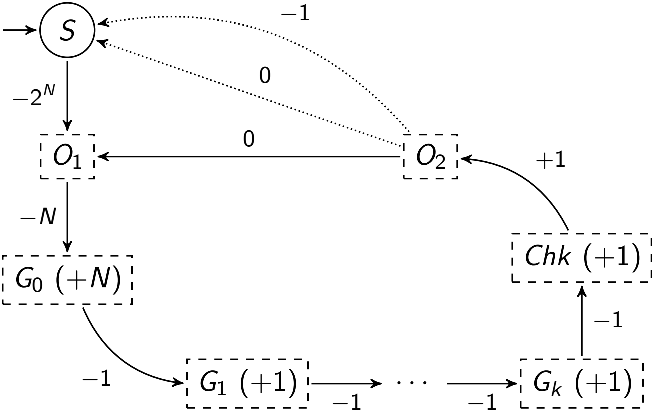

UNIVERSITÉ LIBRE DE BRUXELLES

coNEXP-hardness proof

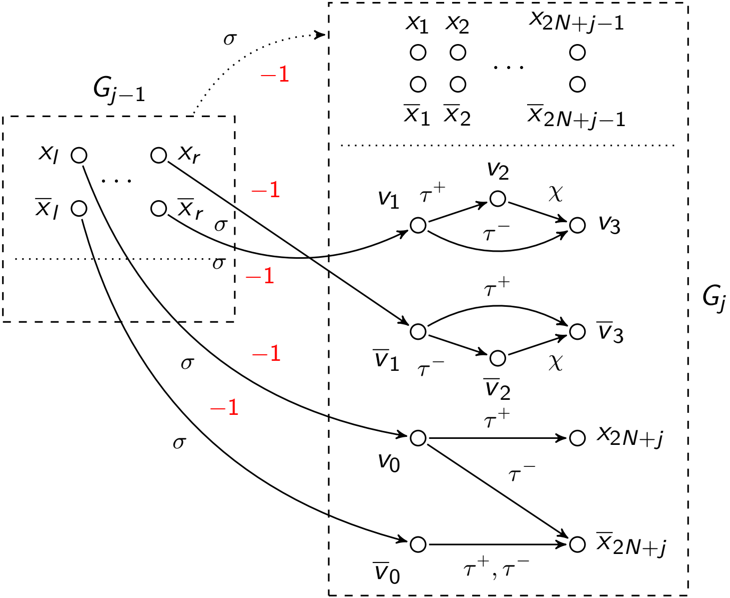

UNIVERSITÉ LIBRE DE BRUXELLES

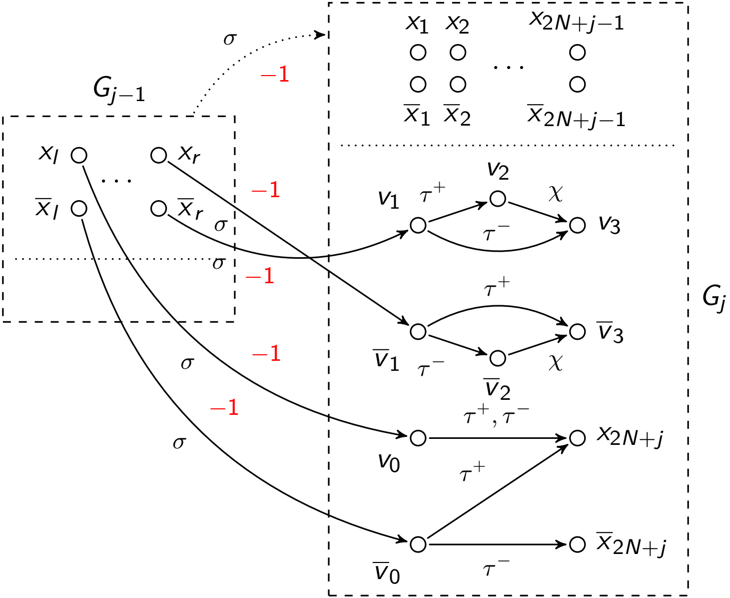
σ
−1
G_{j-1}
x_1 x_2 x_{2N+j-1}
$\overline{x}_1$ $\overline{x}_2$ $\overline{x}_{2N+j-1}$
x_l
x_r
$\overline{x}_l$
$\overline{x}_r$
−1
σ
σ
−1
v_1 τ^+ v_2 χ v_3
τ^-
G_j
τ^+
$\overline{v}_1$ τ^- $\overline{v}_2$ χ $\overline{v}_3$
σ
−1
−1
σ
τ^+, τ^-
v_0
x_{2N+j}
τ^+
$\overline{v}_0$
τ^-
$\overline{x}_{2N+j}$

coNEXP-hardness proof

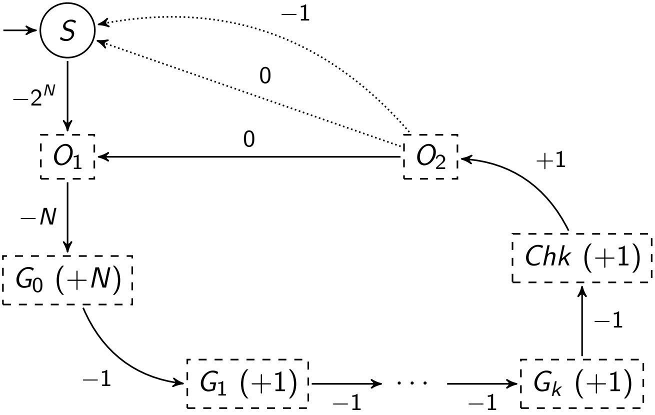

Undecidability of ATU

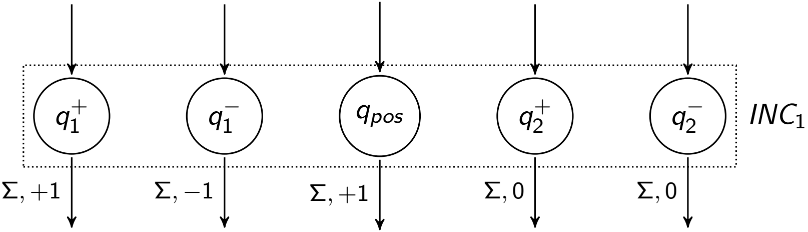

Undecidability of ATU

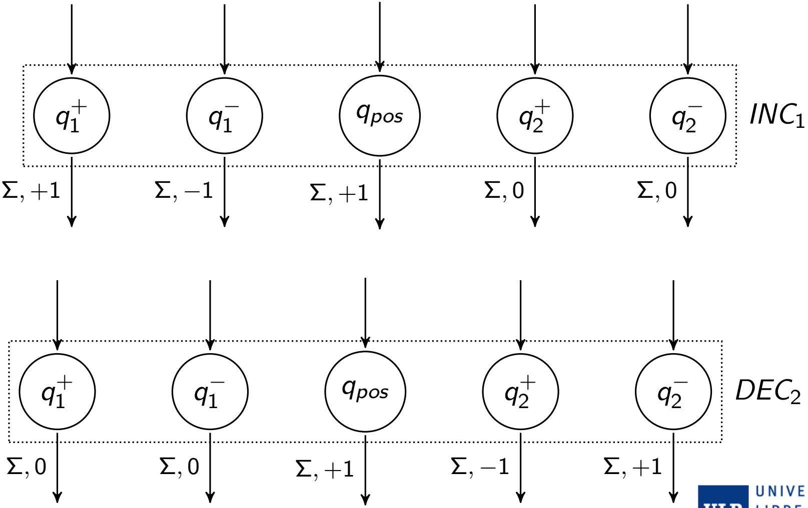

Undecidability of ATU