Quantitative Games

Jean-François Raskin
Université Libre de Bruxelles

Movep 2014
Nantes

Controller synthesis as a game

☞ support the design process with **automatic synthesis**

$$\text{Env} \;\|\; \textbf{?} \;\models\; \psi$$

- Sys is constructed by an **algorithm**
- Sys is **correct** by construction
- Underlying theory: **2-player zero-sum games**
- Env is **adversarial** (worst-case assumption)

Winning strategy = Correct Sys

Controller synthesis with quantitative objectives

Embedded Control

Communication Protocols

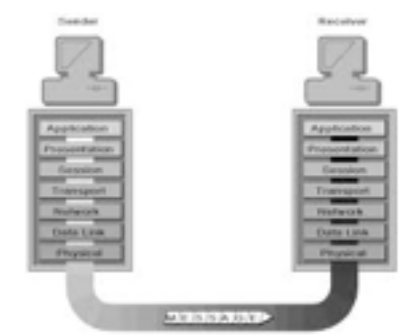

Security Protocols

Parts of OS/Chipset

In most of those examples, **quantitative** measures
of performances are important : not only a matter or correctness !!!

Plan of the talk

- 2-player zero-sum games played on weighted graphs
- Mean-payoff and energy games
    - Determinacy of MPG: an elementary proof
    - Fixpoint algorithm: a pseudo-polynomial time solution
    - Memoryless determinacy (corollary of the FP algorithm)
- Multi-dimensional mean-payoff and energy games
- Summary and conclusion

Two-player zero sum games played on weighted graphs

Game played on weighted graph

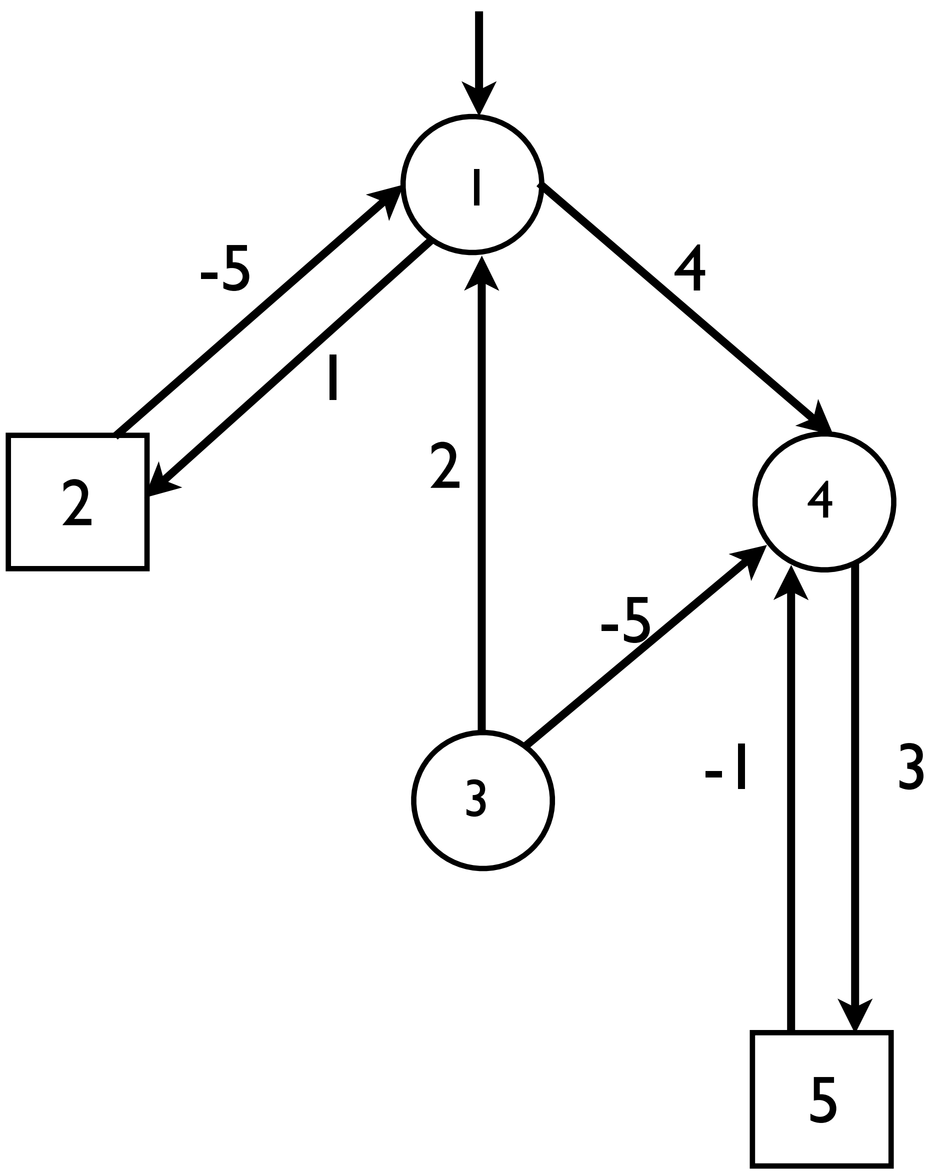

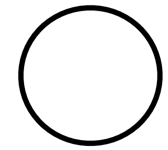 : positions of Player 1 - **maximizer=system**

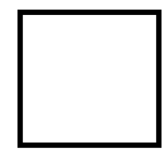 : positions of Player 2 **minimizer=environment**

Directed graph with **weights** on edges
and a partition of the states $G=(S_1,S_2,E,w)$

The game is played in **rounds**:
-**initially** a token is on state s,
-**rounds**: the player owning the current state
 chooses an outgoing edge to move the token

The outcome is an **infinite path**=**play**

Winning condition: the set of infinite paths
is partitioned into

W_1=winning for player 1

$W_2=S^\omega \setminus W_1$=winning for player 2

Players play according to strategies

(Player 1) strategy:

λ_1: $V^*. V_1 \rightarrow$edge.
Λ_1=set of strategies of Pl.1

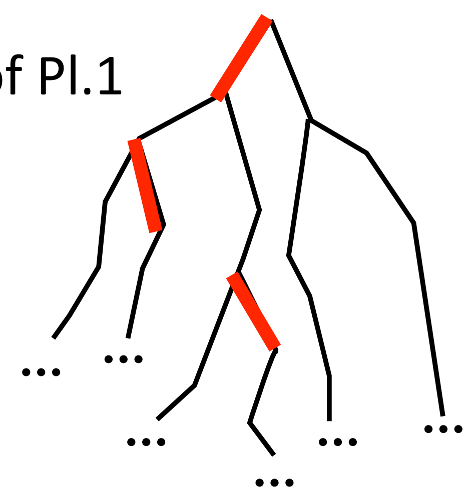

Memoryless strategy:

$\lambda_{1,\mathbf{m}}$:$V_1 \rightarrow$edge.
$\Lambda_{1,\mathbf{m}}$=set of memoryless strategies of Player1

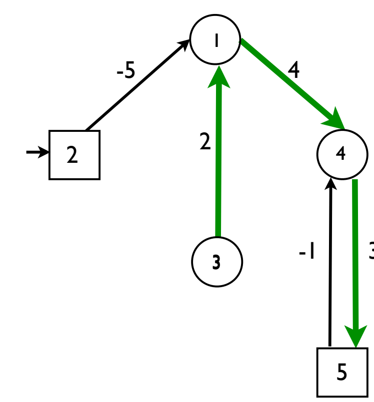

Finite-memory strategy:

$\lambda_{1,\mathbf{f}}$: $V^*. V_1 \rightarrow$edge but regular (Moore machine)
$\Lambda_{1,\mathbf{f}}$=set of finite memory strategies of Player I

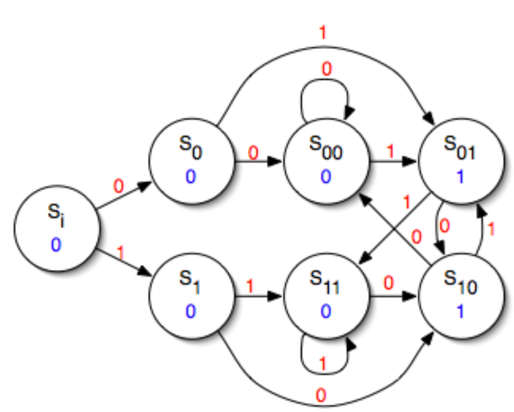

Randomized strategy:

$\lambda_{1,\mathbf{r}}$: $V^*. V_1 \rightarrow$**Dist**(edge).
$\Lambda_{1,\mathbf{r}}$=set of randomized strategies of Player1

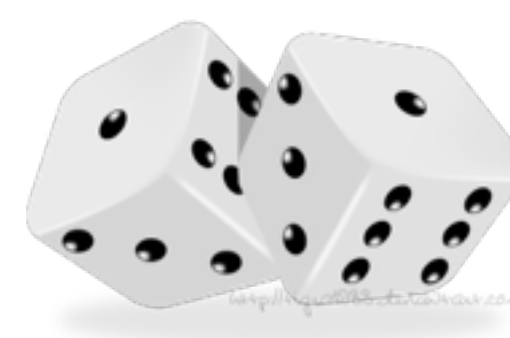

Winning strategies

- If Player 1 plays λ_1 and Player 2 plays λ_2 from s, then the **outcome** of the interaction is a **play** noted **Outcome**(s,λ_1,λ_2)

- **Outcome**$(s,\lambda_1)=\{ \pi \mid \exists\lambda_2 \text{ s.t. } \pi=$ **Outcome**$(s,\lambda_1,\lambda_2) \}$
  Outcome$(s,\lambda_2)=\{ \pi \mid \exists\lambda_1 \text{ s.t. } \pi=$ **Outcome**$(s,\lambda_1,\lambda_2) \}$

- λ_1 is a **winning strategy** for Player 1 from s
  if **for all** strategies λ_2 of Player 2: $\text{Outcome}(s,\lambda_1,\lambda_2) \in \text{Win}_1$
  equivalently, if **Outcome**$(s,\lambda_1) \subseteq \text{Win}_1$

- λ_2 is a **winning strategy** for Player 2 from s
  if **for all** strategies λ_1 of Player 1: $\text{Outcome}(s,\lambda_1,\lambda_2) \in \text{Win}_2=S^\omega\setminus\text{Win}_1$

- Player 1 **wins** the game from s if there exists a **winning strategy** for Player 1 from s (symmetrically for Player 2)

- A game is **determined** (from a state s) if either Player 1 has a winning strategy (from s) or Player 2 has a winning strategy (from s)

- A class of games is **determined** if all the games in the class are determined

Mean-payoff games
(Ehrenfeucht-Mycielski 79)

Mean-payoff games [EM79]

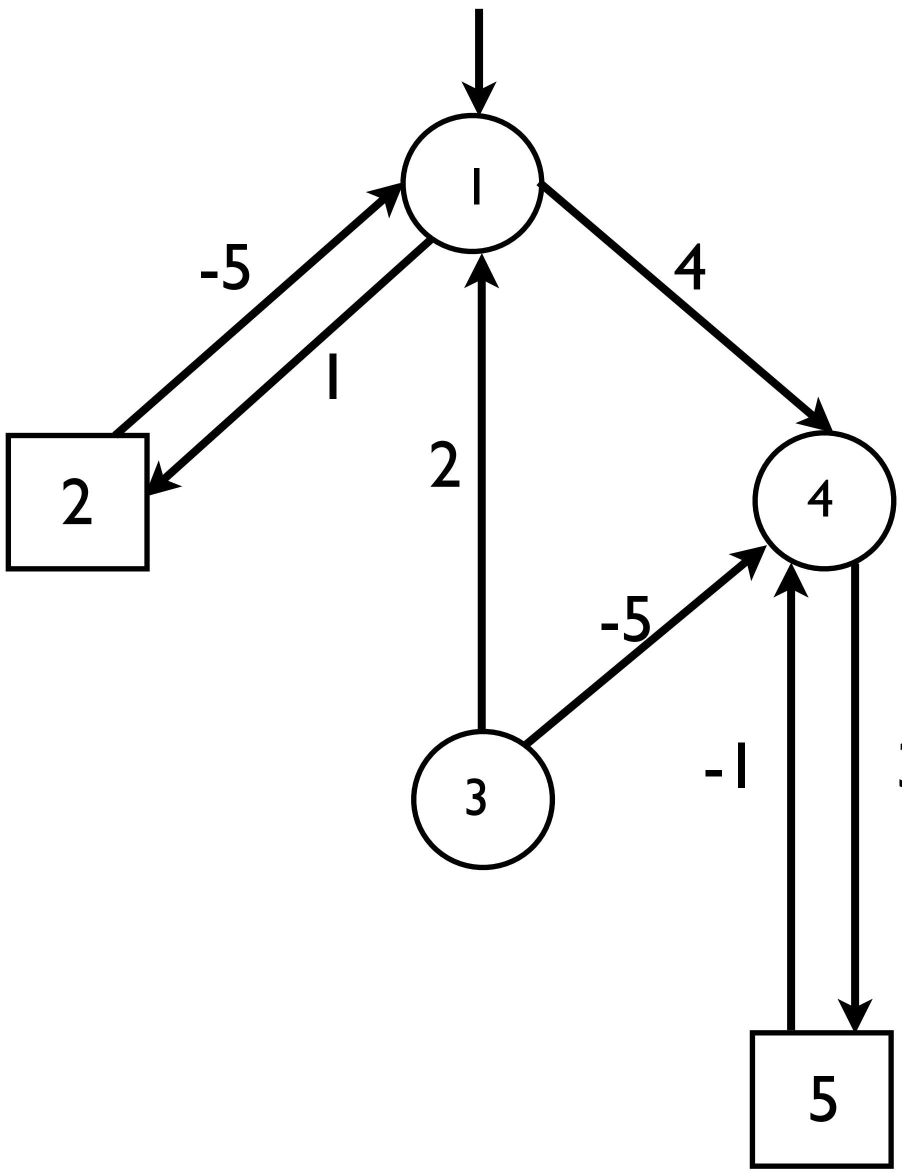

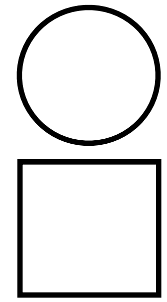

: positions of Player 1 - **maximizer=system**

: positions of Player 2 **minimizer=environment**

Edges are labelled with **rewards**

(1,4) (4,5) (5,4) ... (4,5) (5,4) .... =play

4 3 -1 3 -1 ....

= **Lim Inf**$_{n \to +\infty}$ $\Sigma_{i=1,i=n}$ r_i / n

= **MP**((1,4) (4,5) (5,4) ... (4,5) (5,4) ....)=1

Win$_1$={ play π | MP(π) ≥ **v** }

Note: **not** ω-regular.

Mean-payoff games [EM79]

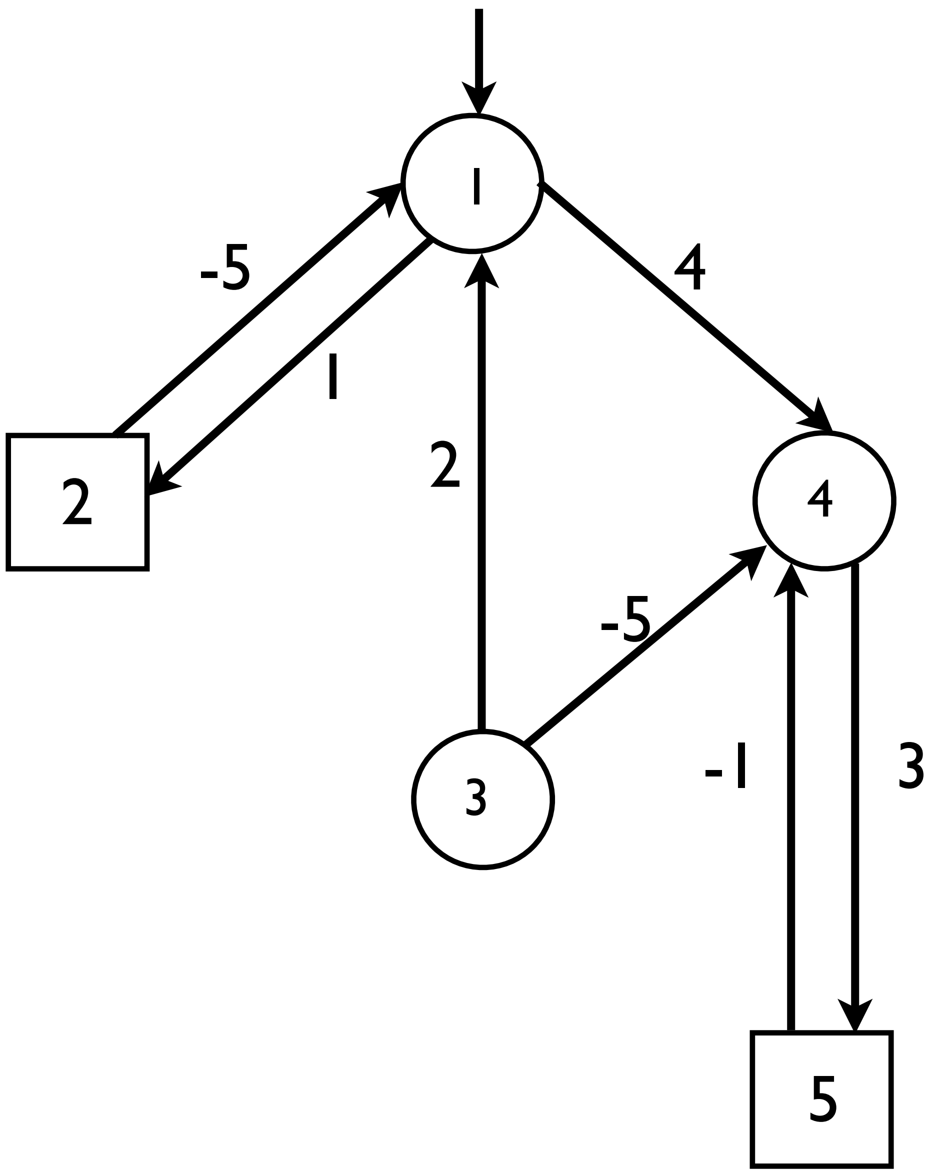

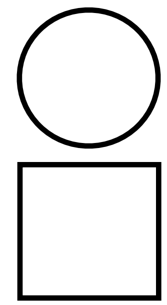

: positions of Player 1 - **maximizer=system**

: positions of Player 2 **minimizer=environment**

Edges are labelled with **rewards**

(1,4) (4,5) (5,4) ... (4,5) (5,4) .... =play
4 3 -1 3 -1 ....

= $\textbf{Lim Inf}_{n \to +\infty} \Sigma_{i=1,i=n} r_i / n$

= **MP**((1,4) (4,5) (5,4) ... (4,5) (5,4) ....)=1

☞ **Win₁**={ play π | MP(π) ≥ **v** }
Note: **not** ω-regular.

Mean-payoff games [EM79]

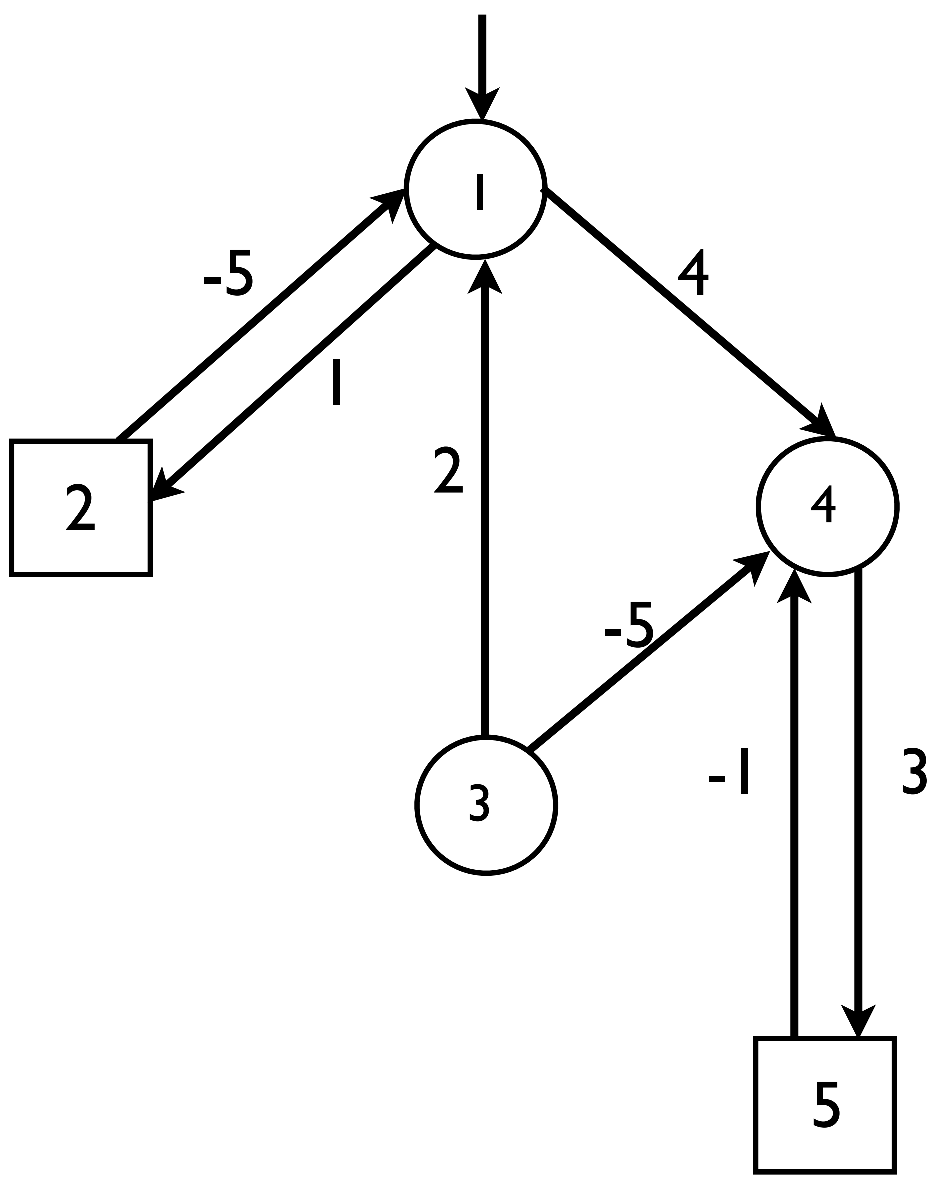

: positions of Player 1 - **maximizer=system**

: positions of Player 2 **minimizer=environment**

Edges are labelled with **rewards**

(1,4) (4,5) (5,4) ... (4,5) (5,4) .... =play

4 3 -1 3 -1 ....

= **Lim Inf**$_{n\rightarrow+\infty}$ $\Sigma_{i=1,i=n}$ r_i / n

= **MP**((1,4) (4,

W.l.o.g., **v=0**

Win₁={ play π | MP(π) ≥ **v** }

Note: **not** ω-regular.

Mean-payoff games [EM79]

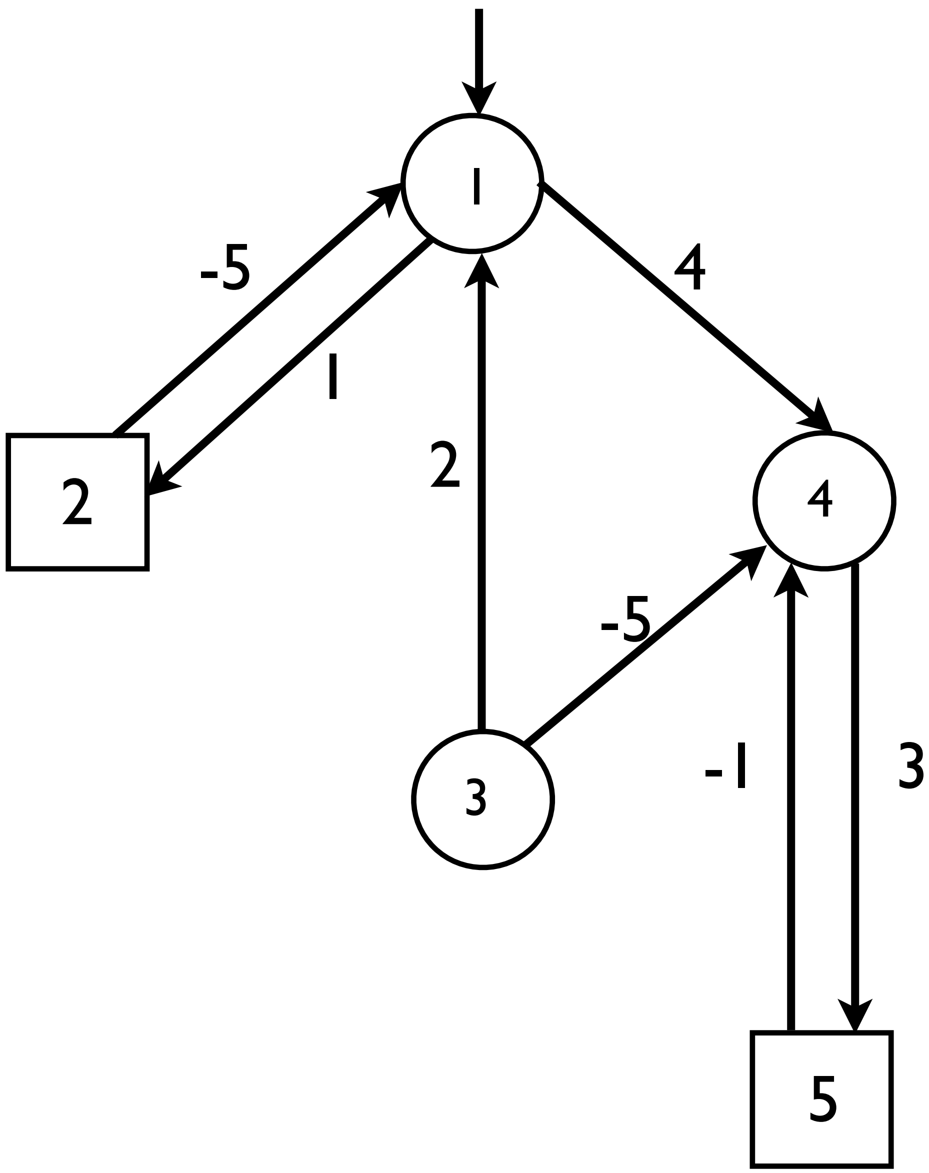

○ : positions of Player 1 - **maximizer=system**

□ : positions of Player 2 **minimizer=environment**

Edges are labelled with **rewards**

(1,4) (4,5) (5,4) ... (4,5) (5,4) .... =play

4 3 -1 3 -1 ....

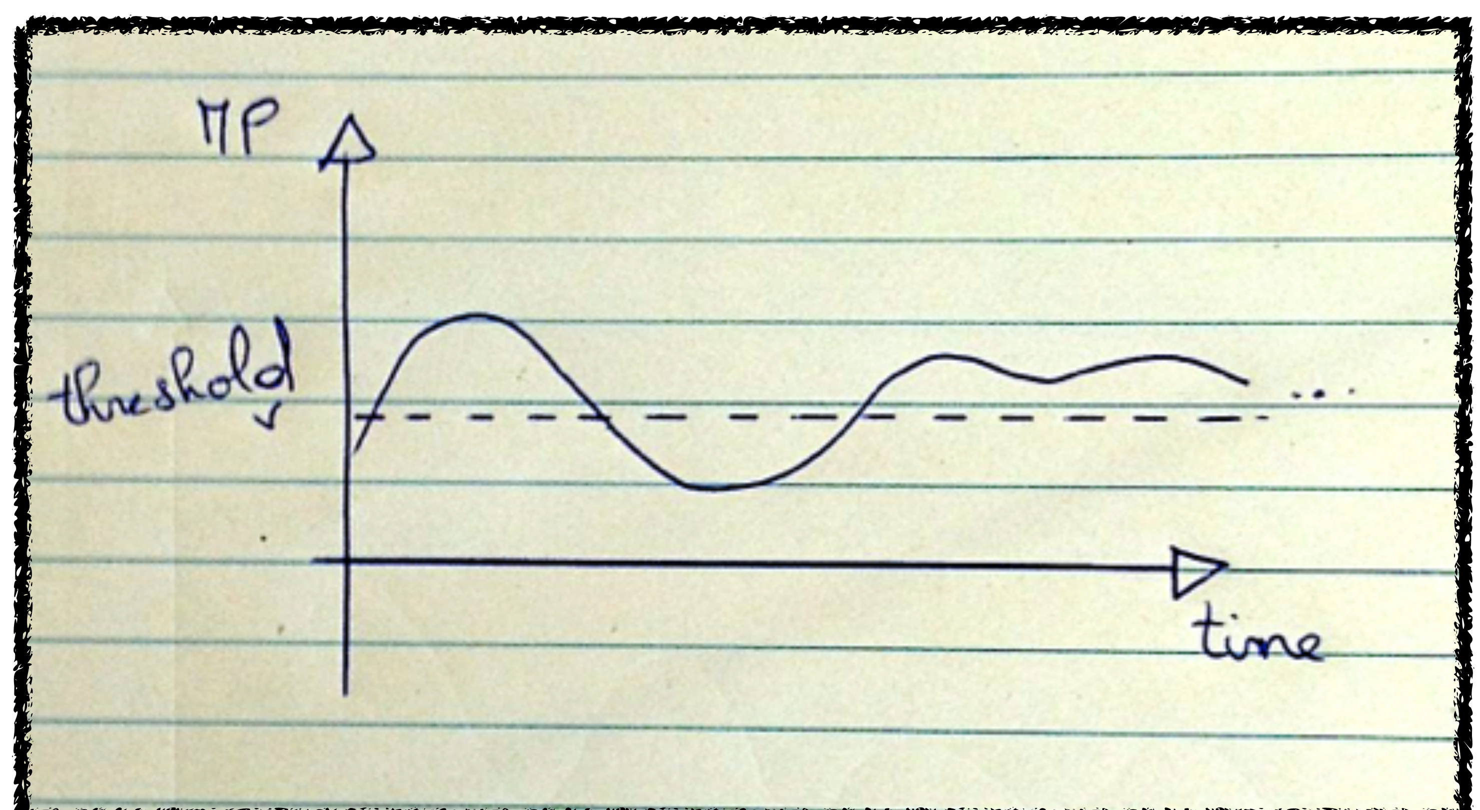

Mean-payoff games [EM79]

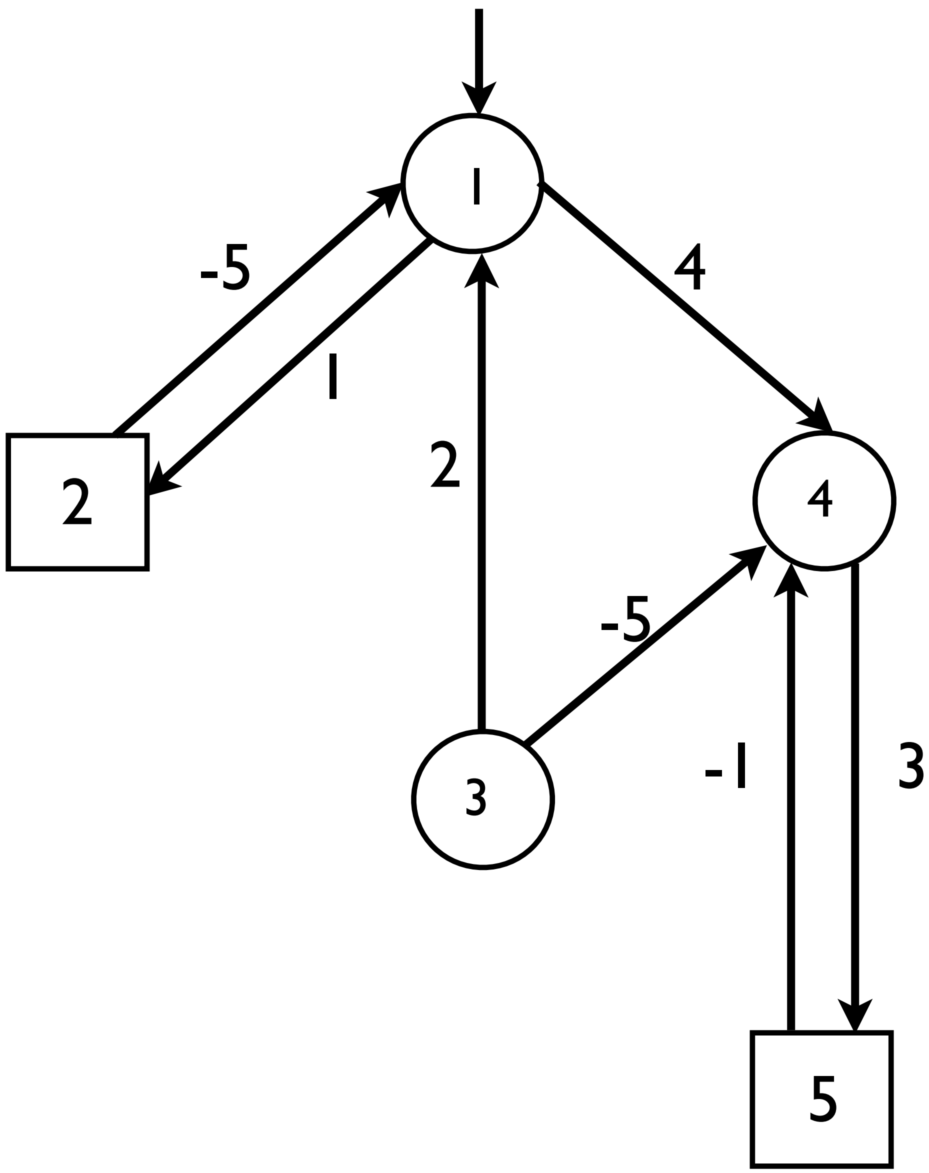

The decision problem for MPG asks:

Given an state **s**, if Player 1 has **a strategy** λ_1 s.t. **Outcome**$(s,\lambda_1) \subseteq Win_1=\{$ plays $\pi \mid MP(\pi) \geq 0 \}$

Energy games

[CdAHS03,BFLMS08]

Energy games

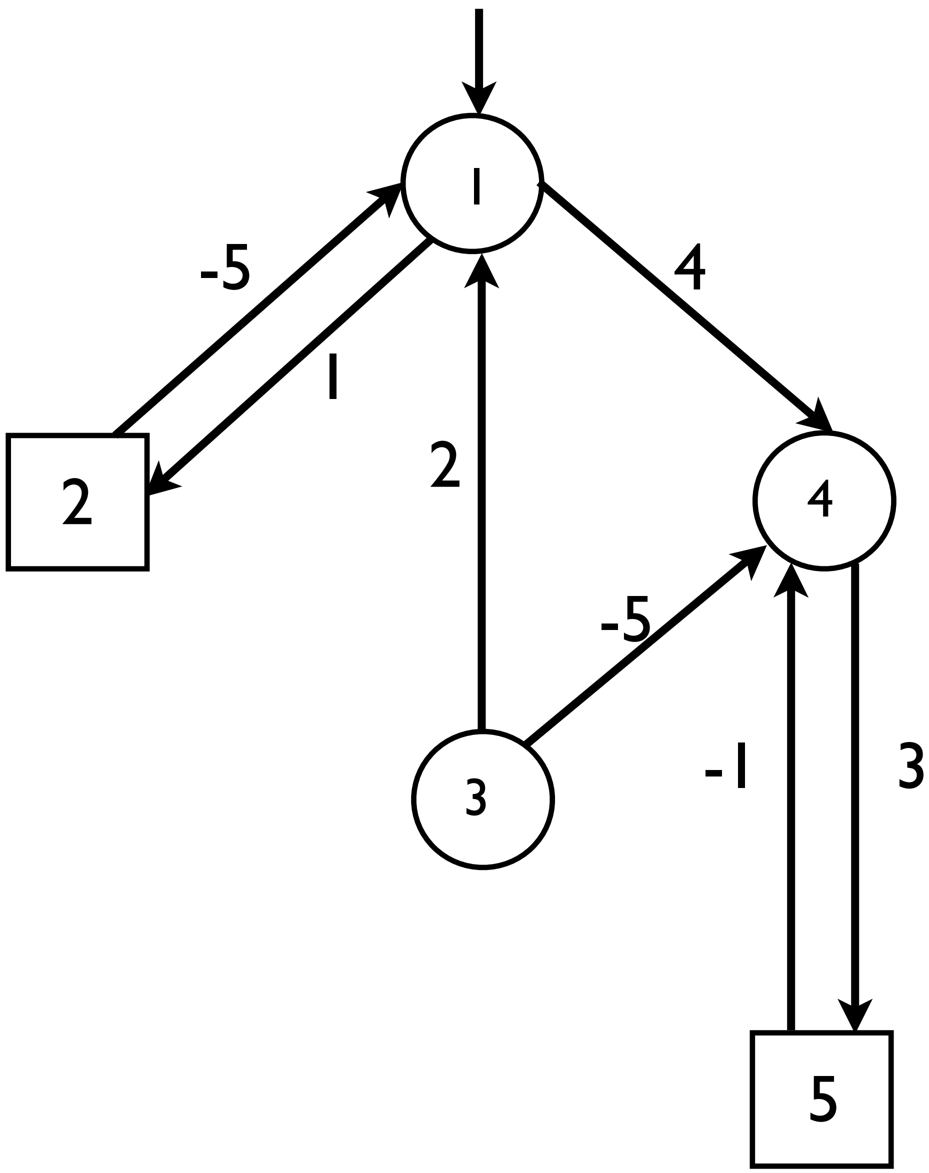

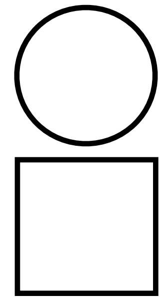

: positions of Player 1 - **maximizer=system**

: positions of Player 2 **minimizer=environment**

Edges are labelled with energy *consumptions* or energy *gains*.

Strategies are defined as for MPGs.

Initial energy level : 7

Play : (1,2) (2,1) (1,4) (4,5) (5,4) (4,5) (5,4) ...

EL : 7 8 3 7 10 9 10 9 ...

$\models \Box$ **EL**≥ 0

Energy games

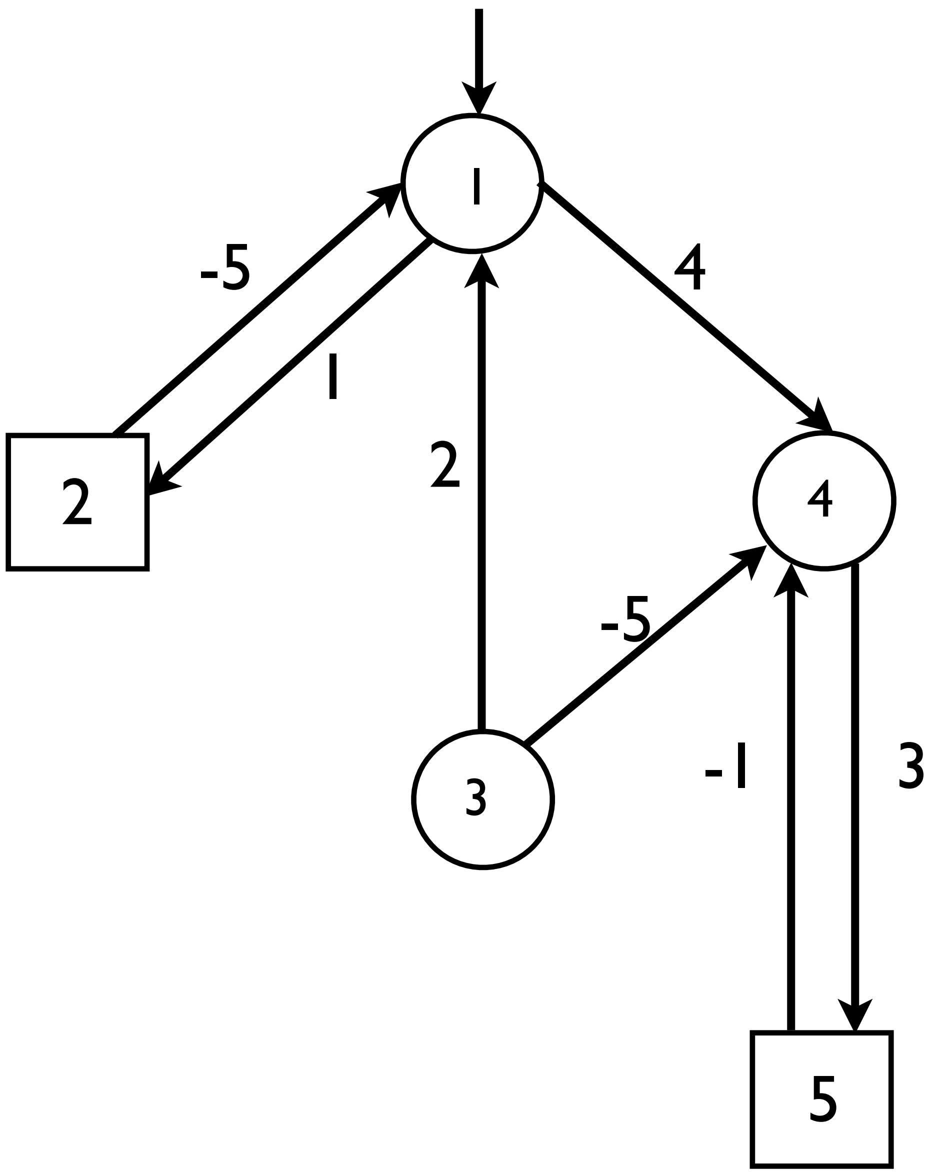

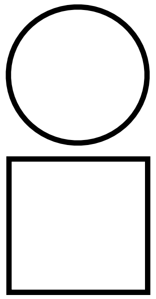

: positions of Player 1 - **maximizer=system**

: positions of Player 2 **minimizer=environment**

Edges are labelled with energy *consumptions* or energy *gains*.

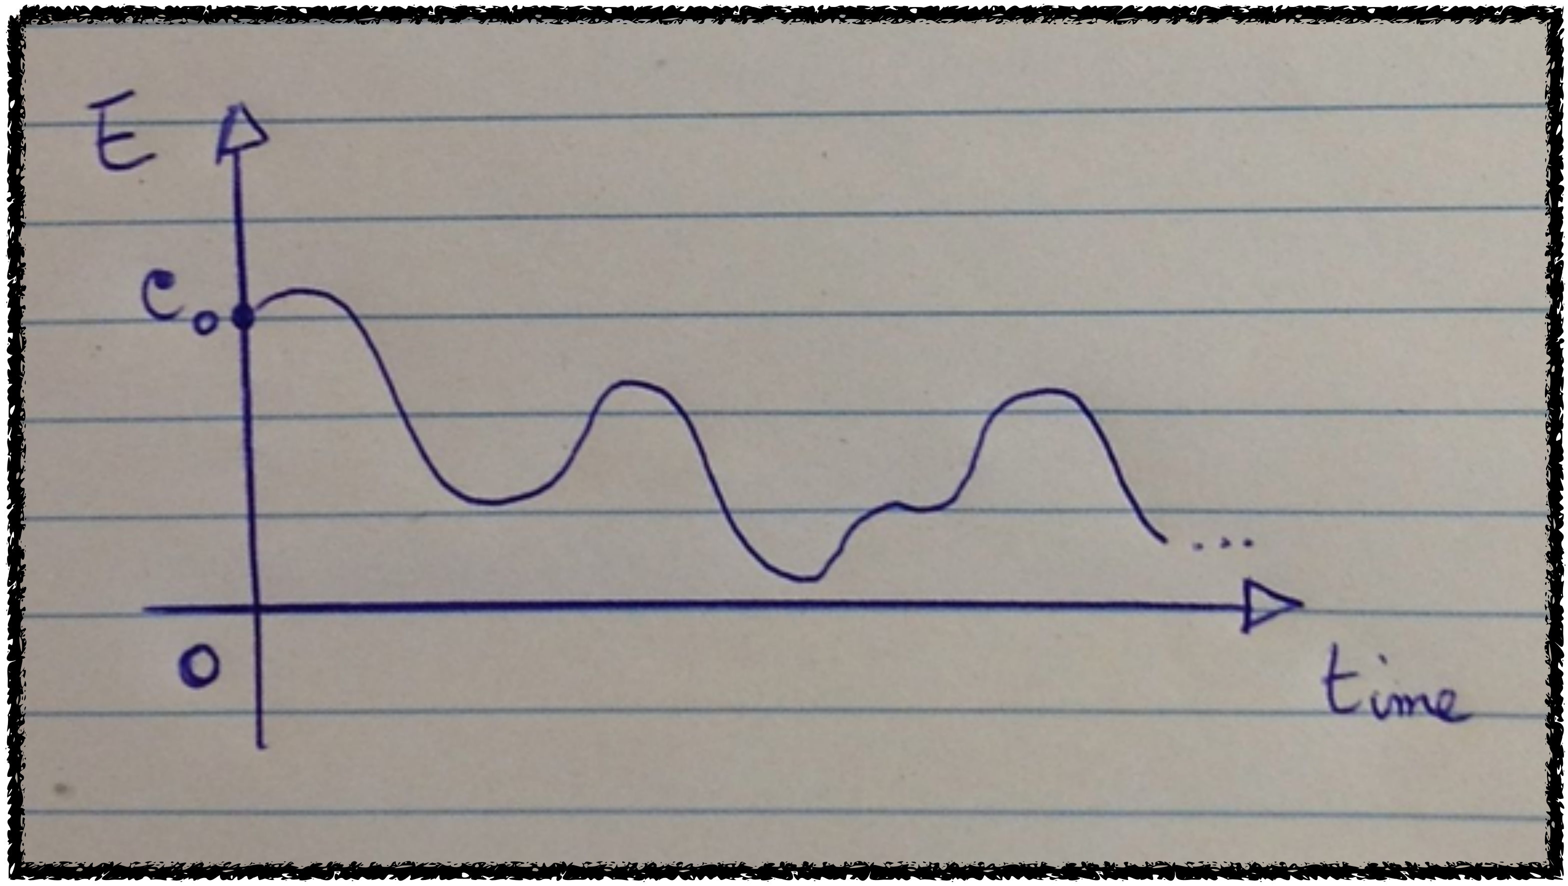

Energy games

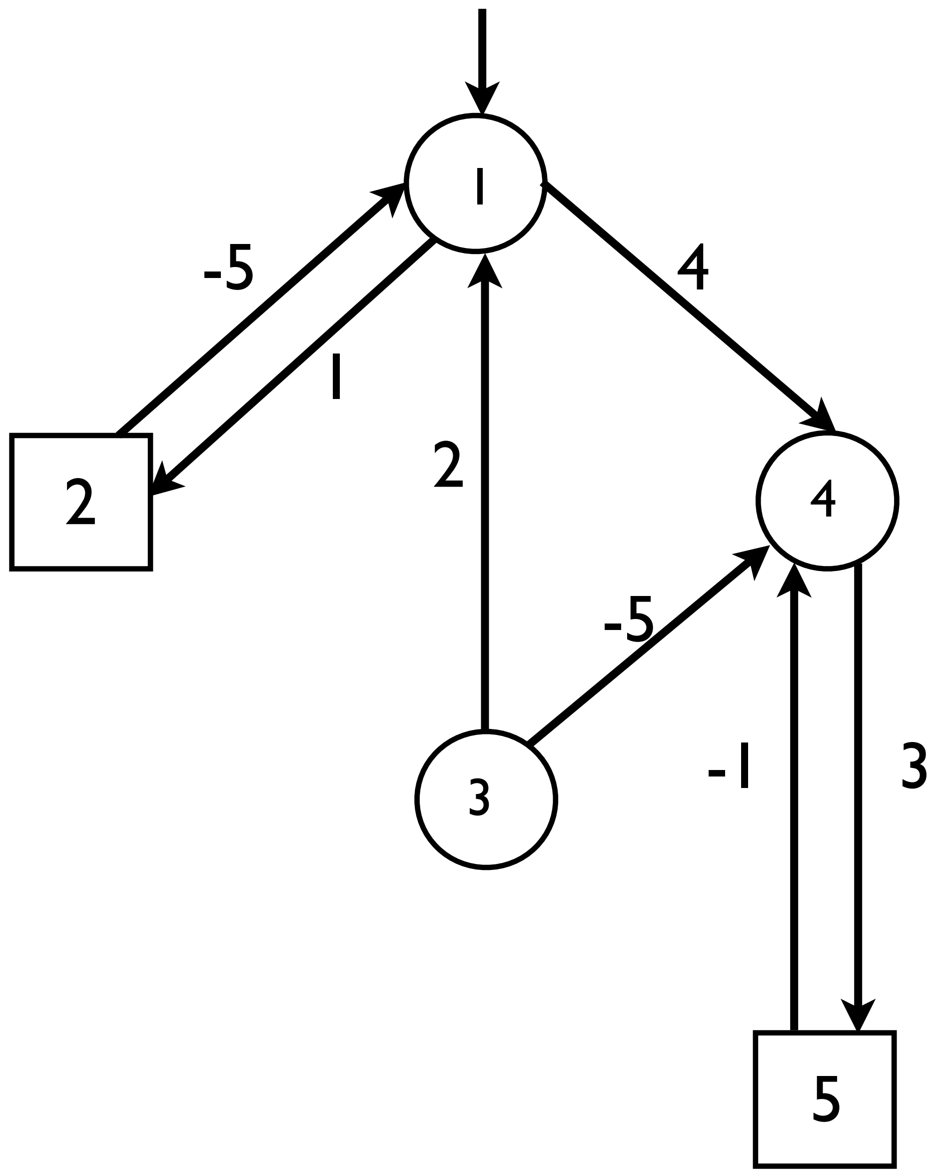

The **decision problem for EG** asks:

given state **s,** decide if there exist

✶ an **initial energy level** $c_0 \in \mathbb{N}$, and

✶ a strategy λ_1 **for Player 1** to maintain a positive energy level from c_0 at all time (i.e. $\square$ **EL**≥0), no matter what Player 2 plays.

Properties that we will prove

We will prove the following results:

- MPG and EG are **inter reducible**
- MPG and EG are **memoryless determined**, as a consequence their decision problems are in **NP∩coNP**
- There is an elegant **pseudo-polynomial** time algorithm to solve EG, and so MPG (remark: no (truly) polynomial time algorithm is known)

MPG and EG

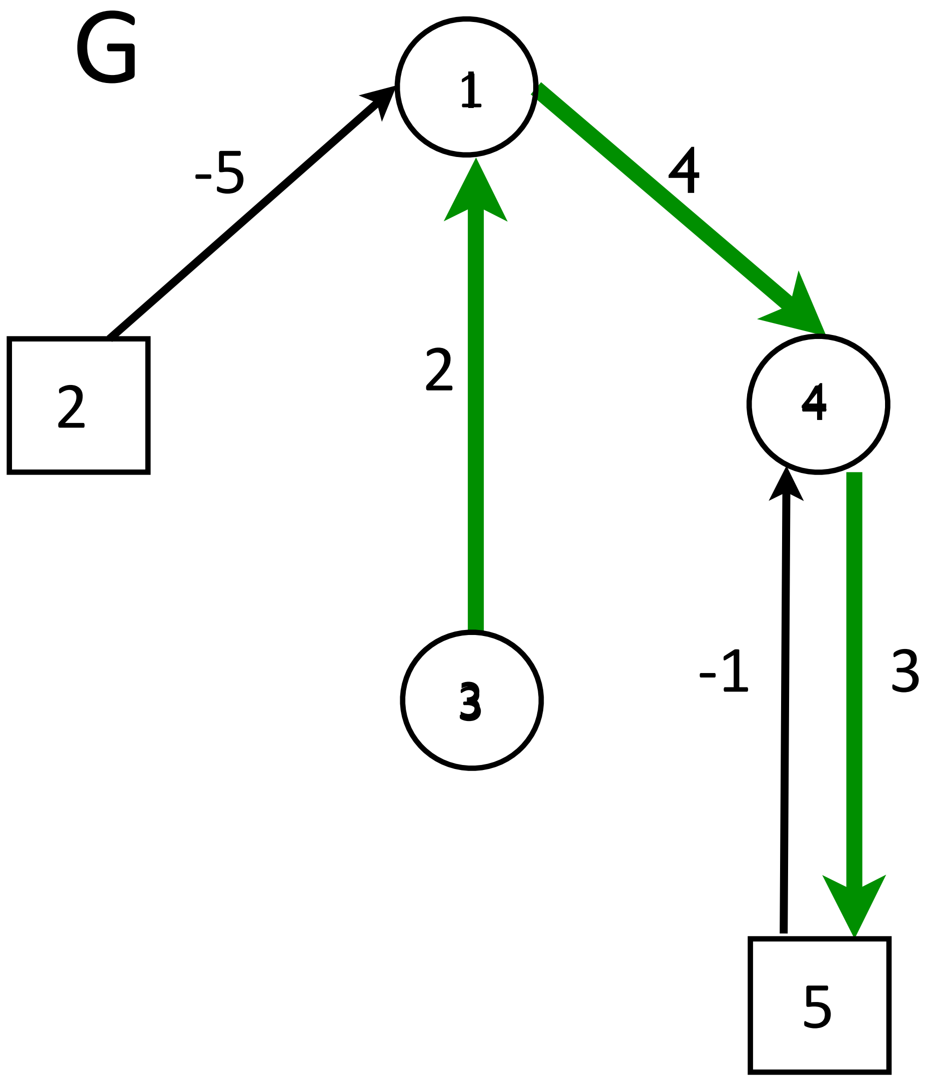

Player 1 has a
memoryless strategy
to force MP ≥ 1

MPG and EG

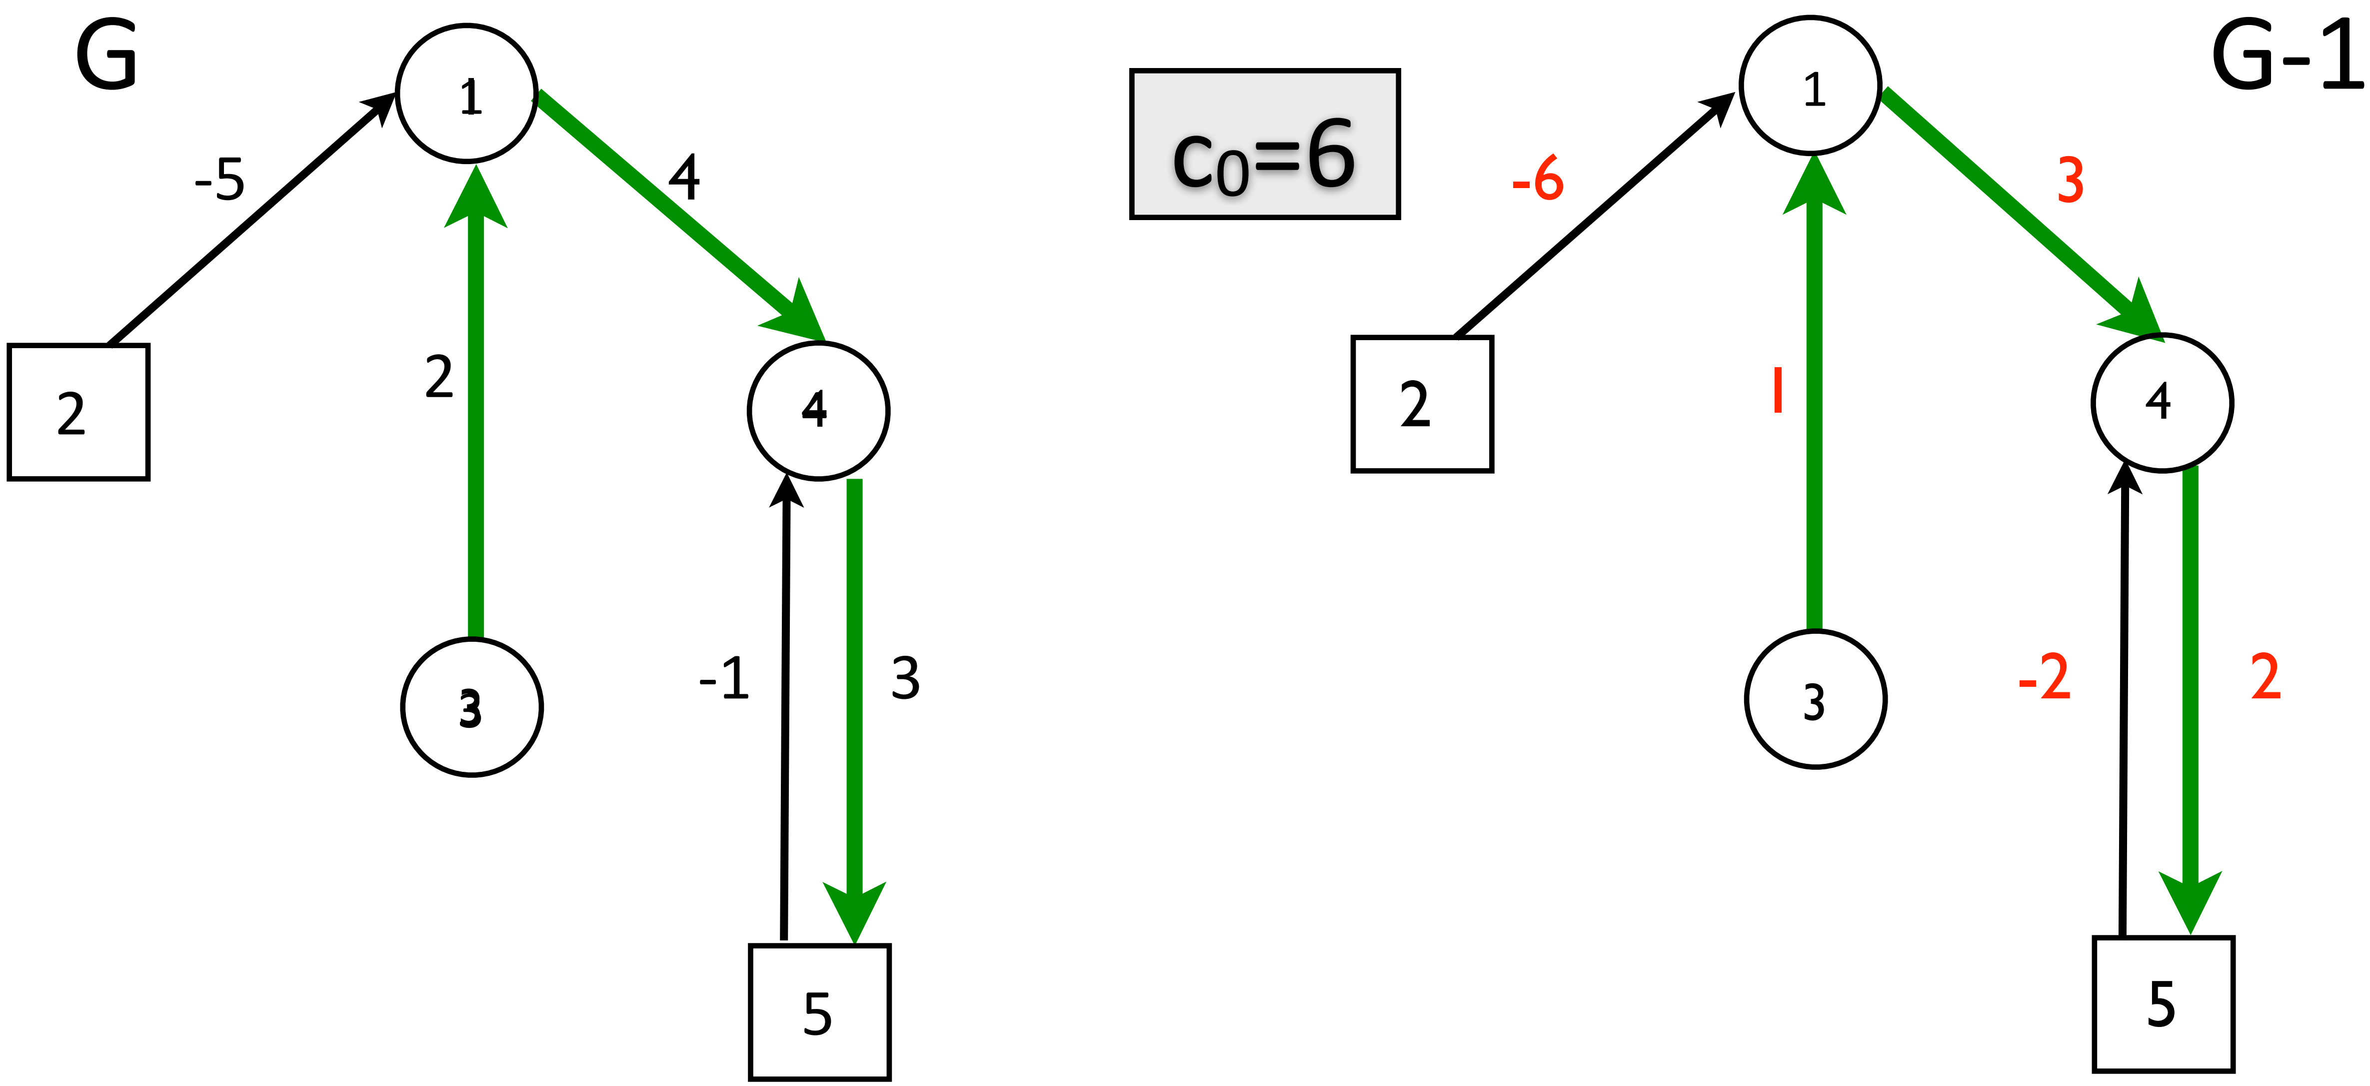

Player 1 ensures value 1 in MPG **iff** Player 1 wins EG G-1.

MPGs and EGs are **memoryless determined** and in **NP∩coNP**

Determinacy and equivalence for MPG and EG An elementary proof

Determinacy of MPG

Theorem [**Determinacy**] For all MPG G, for all states s:

- **either** $\exists \lambda_1$ for Player 1 s.t. Outcome(s, λ_1) $\subseteq \{ \pi \mid MP(\pi) \geq 0 \}$,
- **or** $\exists \lambda_2$ for Player 2 s.t. Outcome(s, λ_2) $\subseteq \{ \pi \mid MP(\pi) < 0 \}$.

Starting point: determinacy for finite tree games

Theorem [**Zermelo 1913**]. Every finite tree reachability game is determined.

Starting point: determinacy for finite tree games

Theorem [**Zermelo 1913**]. Every finite tree reachability game is determined.

- Every finite duration turn-based game can be represented as a game tree of bounded depth
- Each branch represents a play
- The winning condition is defined by a partition of the leaves of the tree: plays that are winning for Player 1 and those that are winning for Player 2
- A corollary: in **chess**, either black or white (one of the two players) is able to force win or draw

Determinacy of finite tree reachability games

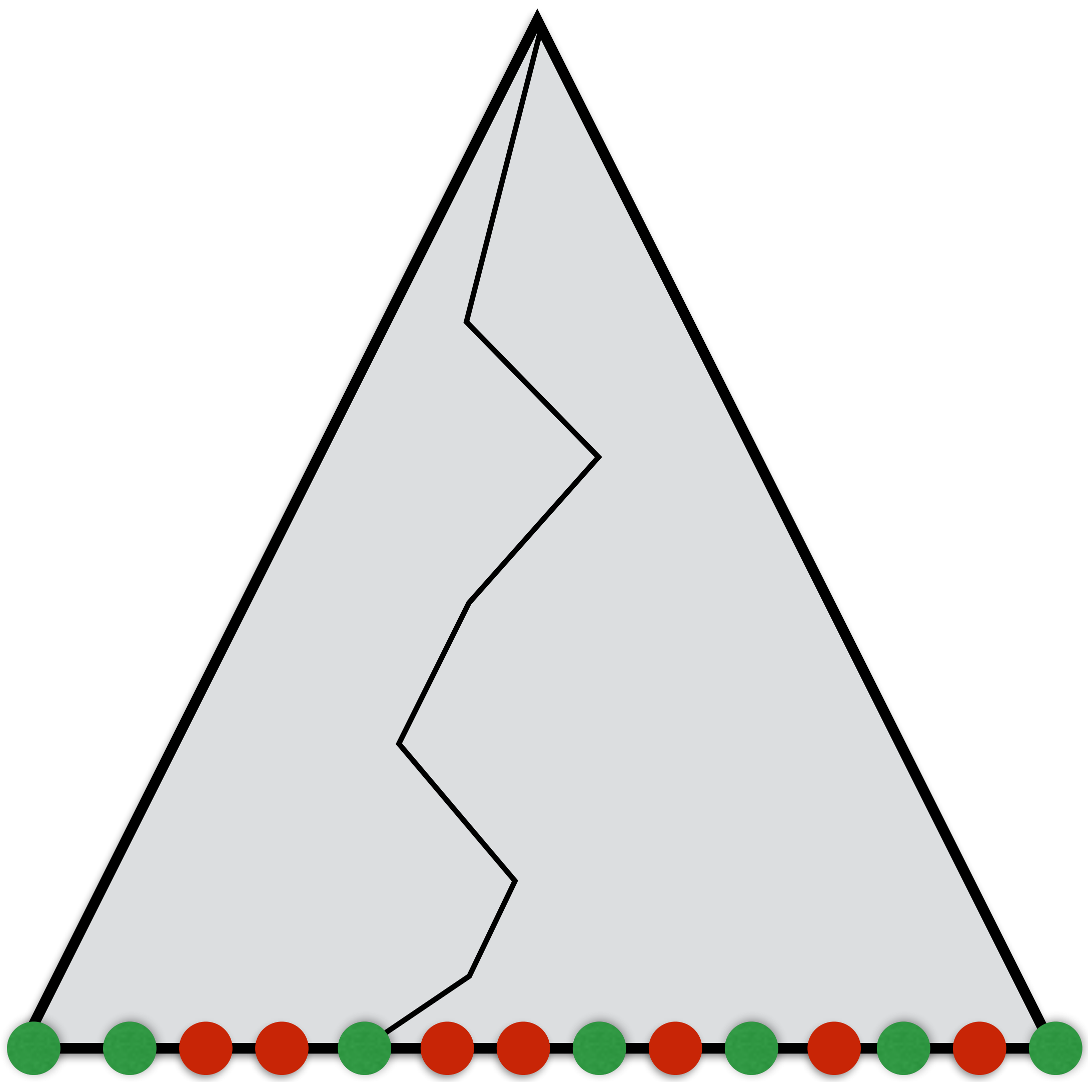

Zermelo theorem says that at the root

either Player 1 can force a **green** leaf

or Player 2 can force a **red** leaf

Determinacy of finite tree reachability games

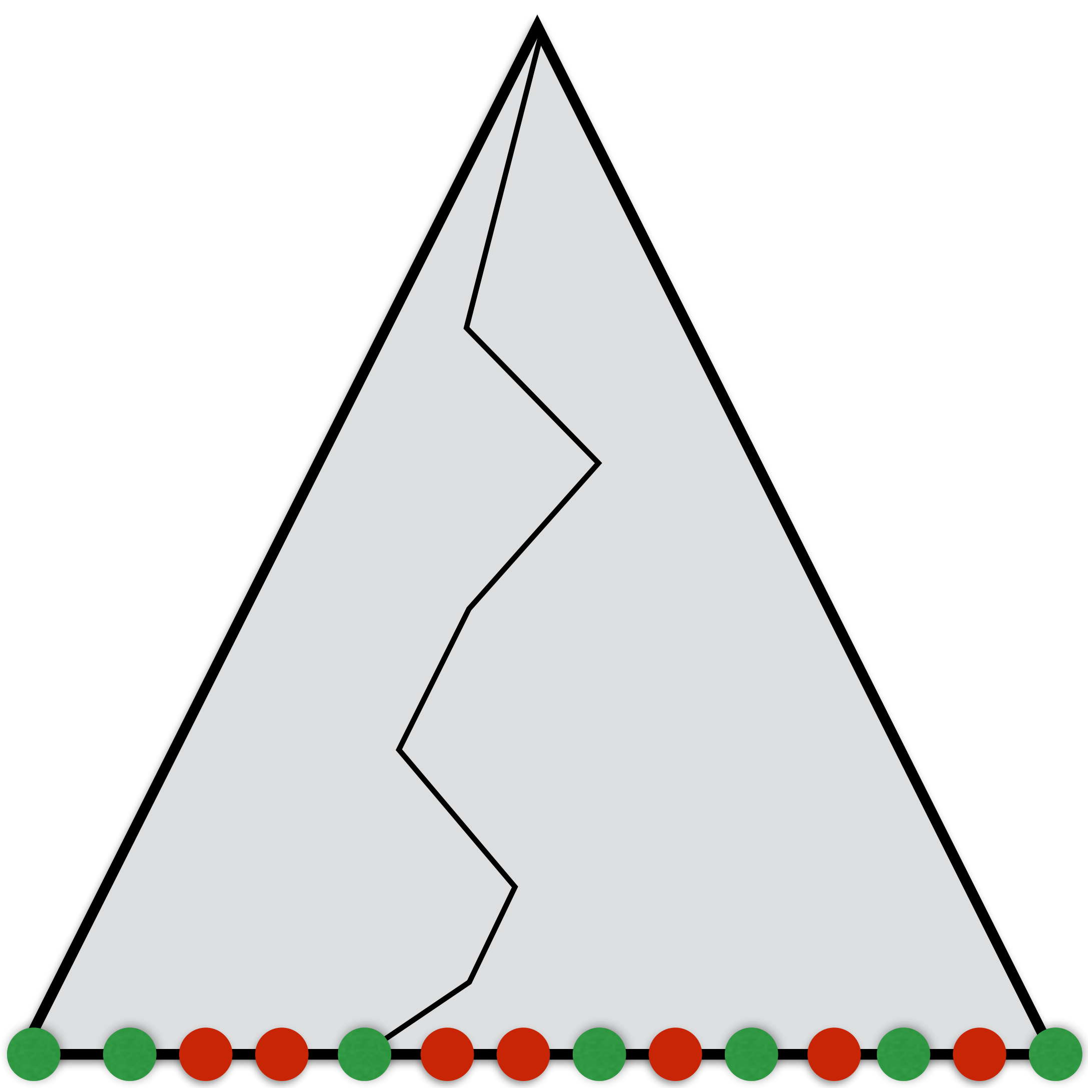

The proof is by **induction** on the depth of the tree.

Each node of the tree can be labelled in green or red so that:

1) **Player 1** can force a green leaf from any **green node**

2) **Player 2** can force a red leaf from any **red node**

So, as the root is either red or green, one of the players has a winning strategy for his objective.

Determinacy of finite tree reachability games

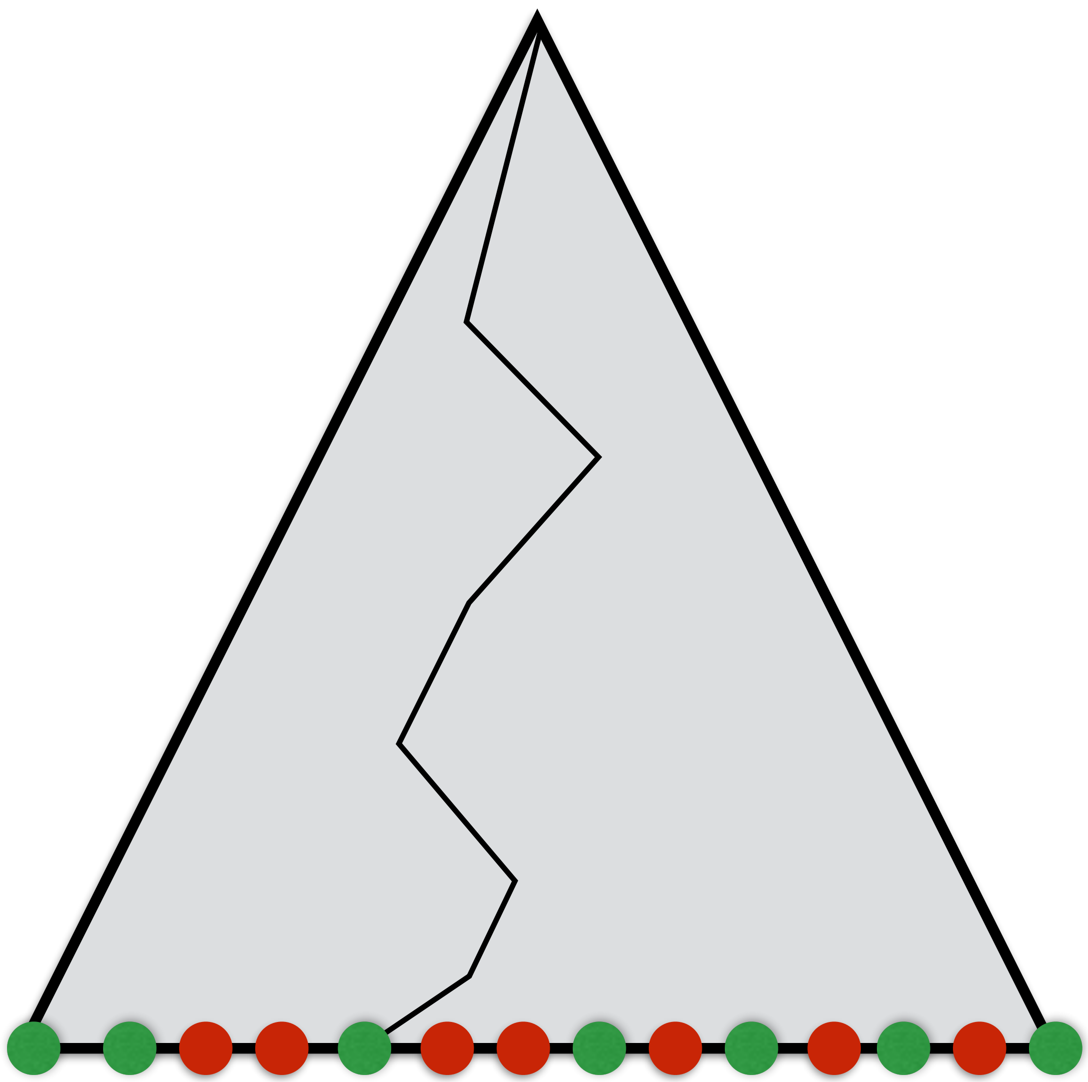

Determinacy of finite tree reachability games

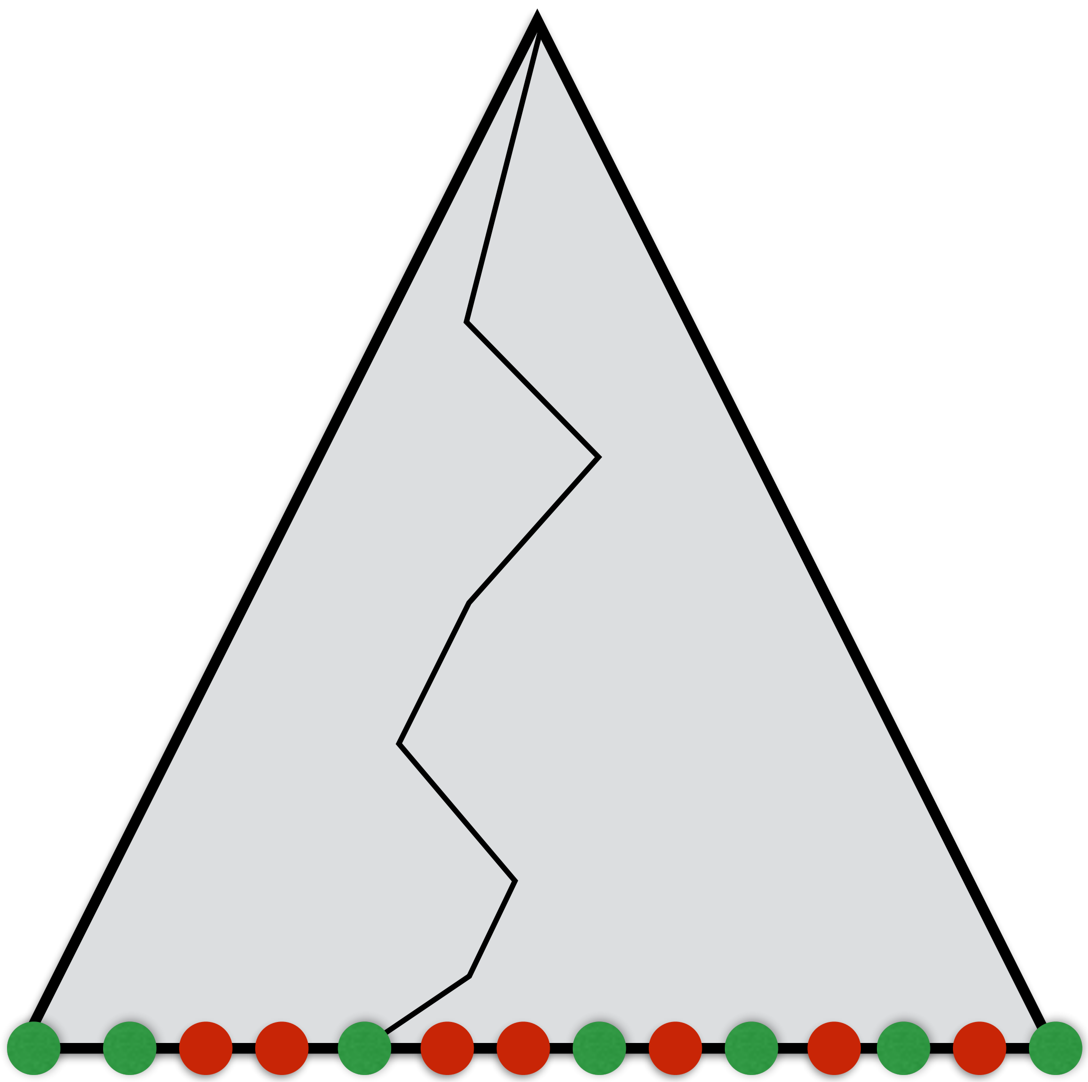

Base case : tree=one leaf. Trivial.

Determinacy of finite tree reachability games

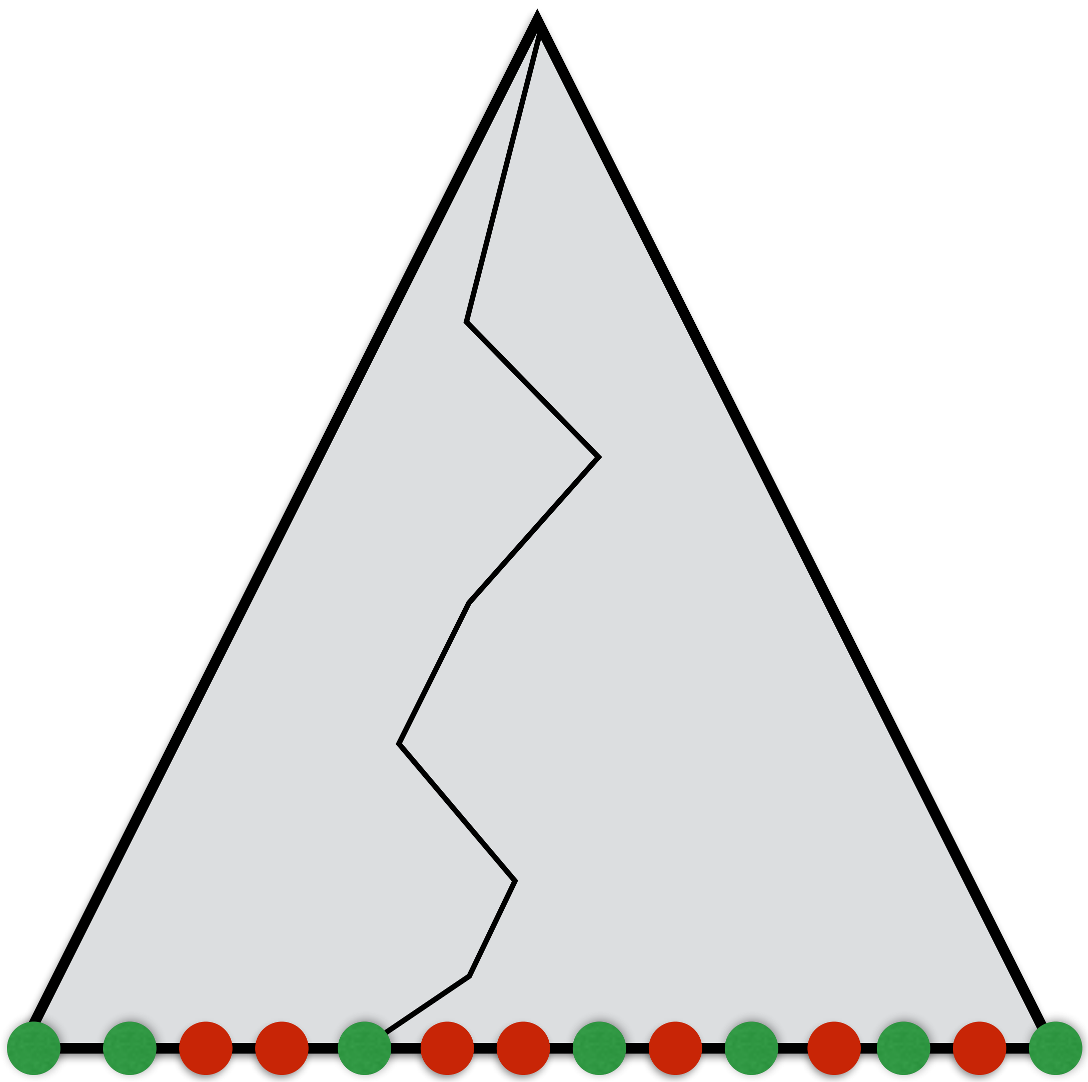

Induction.

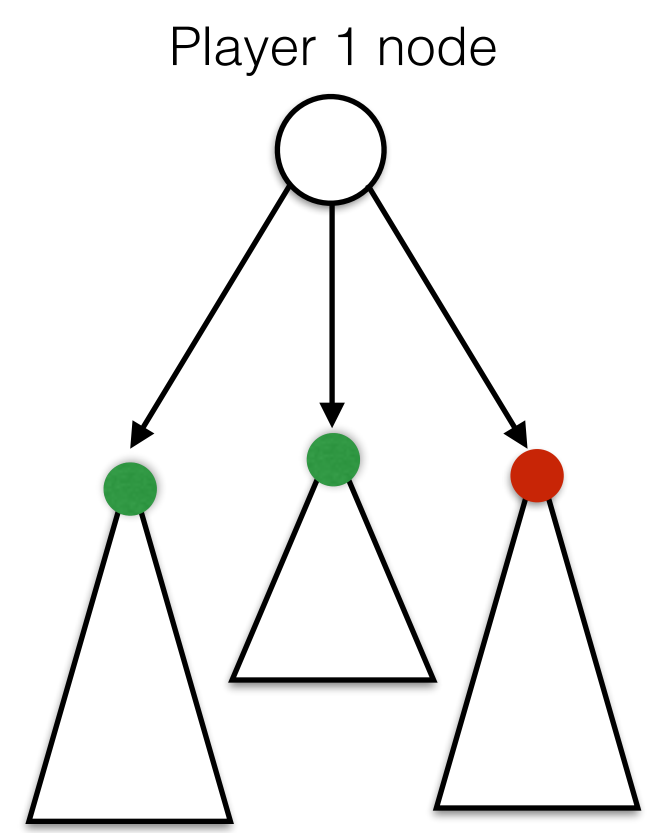

The node is green if and only if there exists one green successor

Determinacy of finite tree reachability games

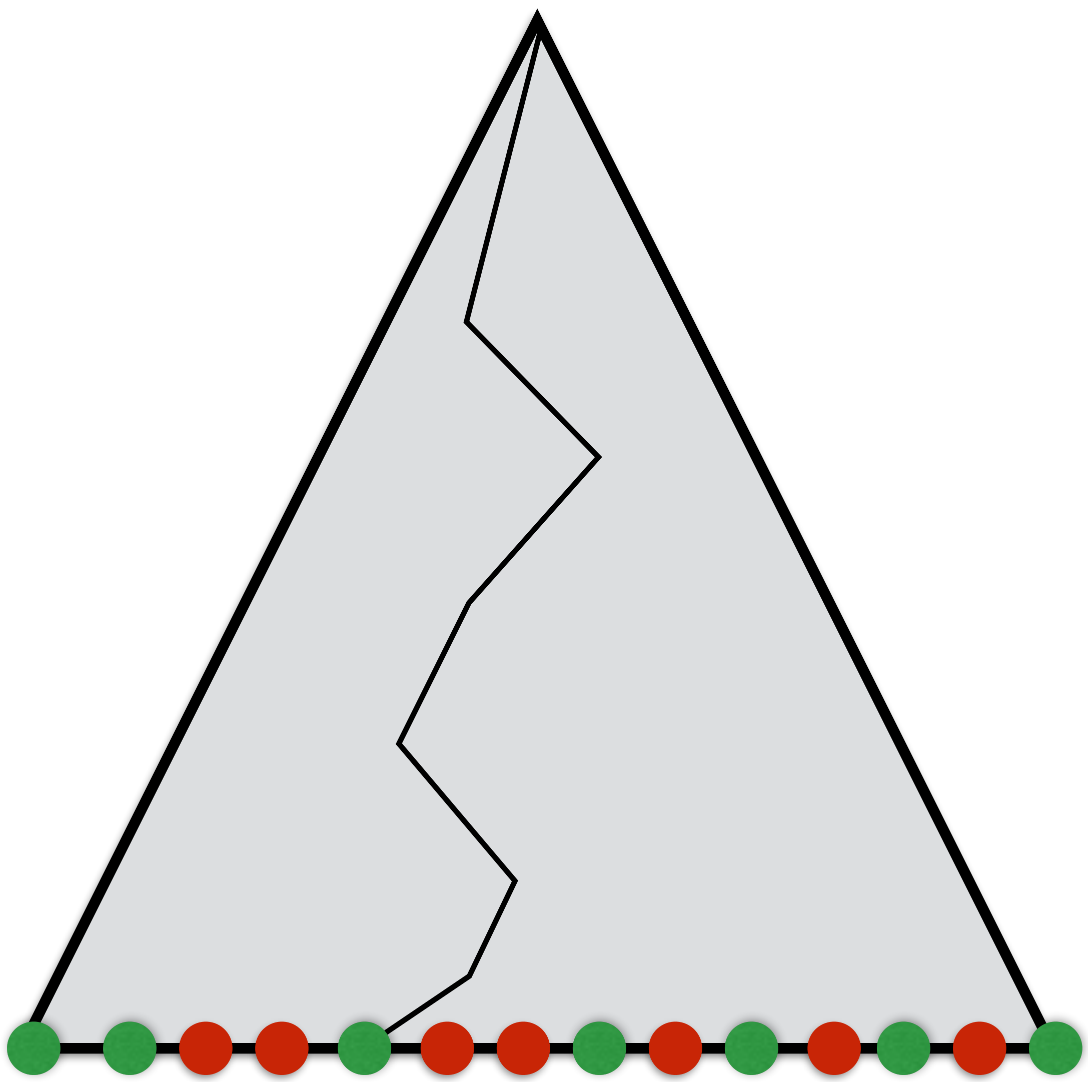

Induction.

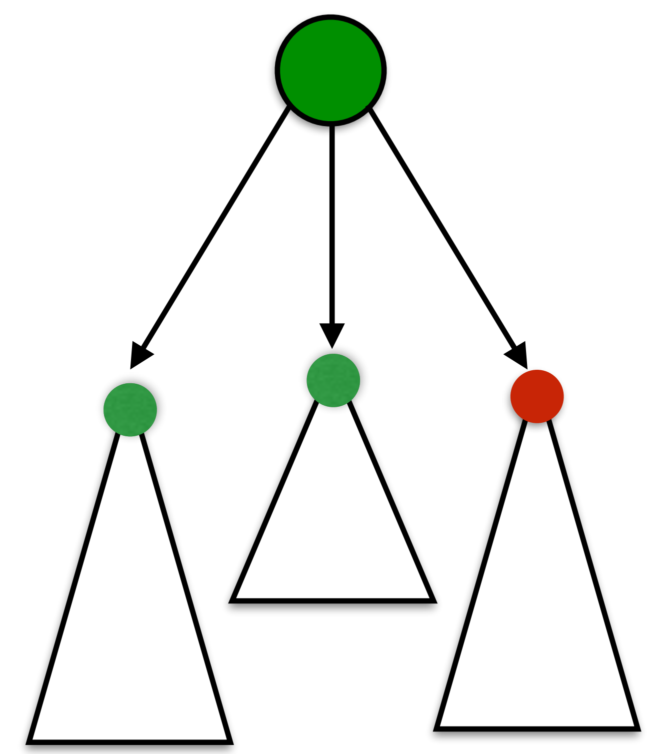

The node is green if and only if there exists one green successor

Determinacy of finite tree reachability games

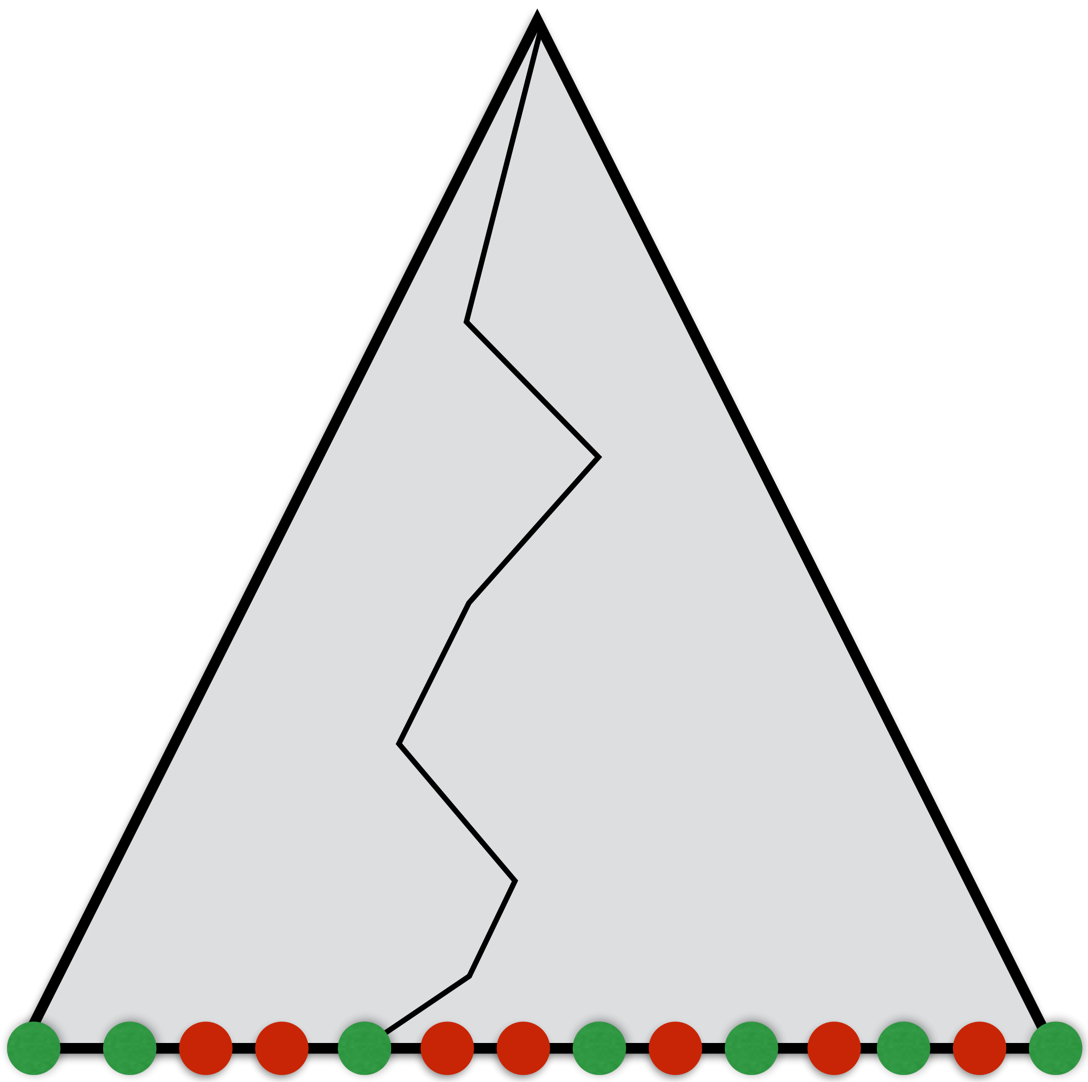

Induction.

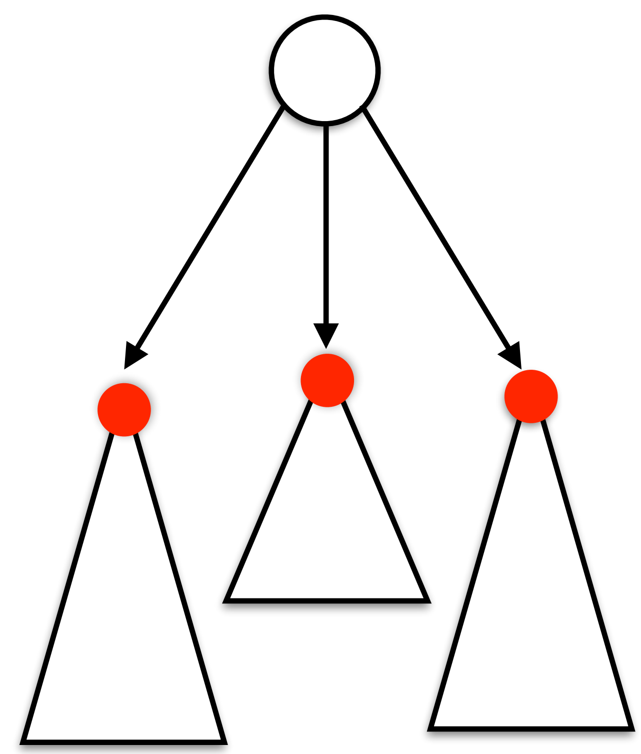

The node is green if and only if there exists one green successor

Determinacy of finite tree reachability games

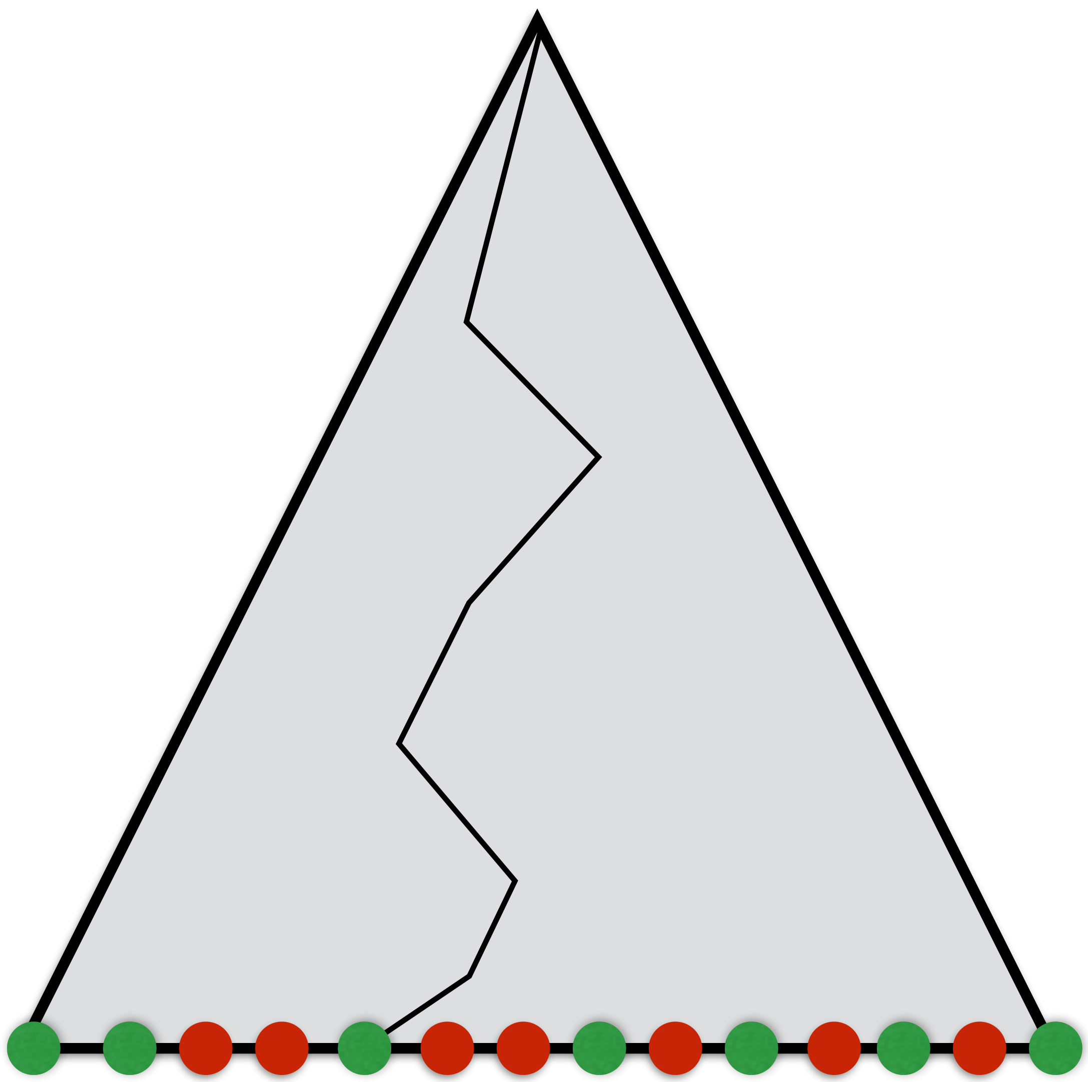

Induction.

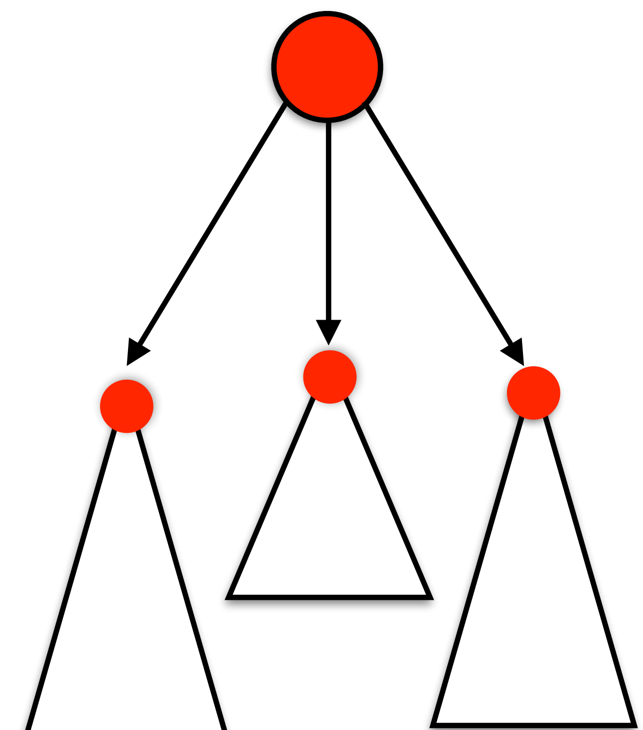

The node is green if and only if there exists one green successor

Determinacy of finite tree reachability games

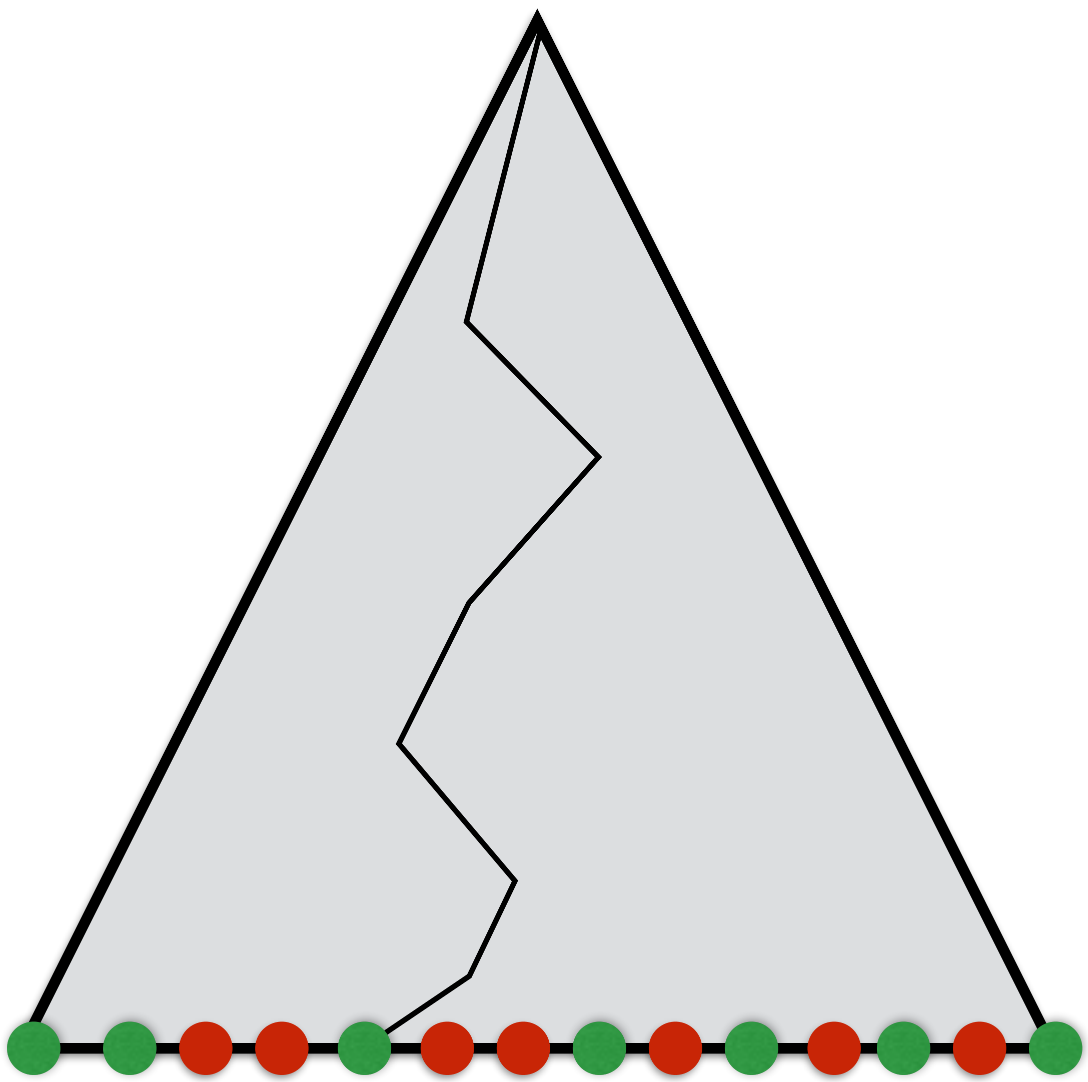

Induction.

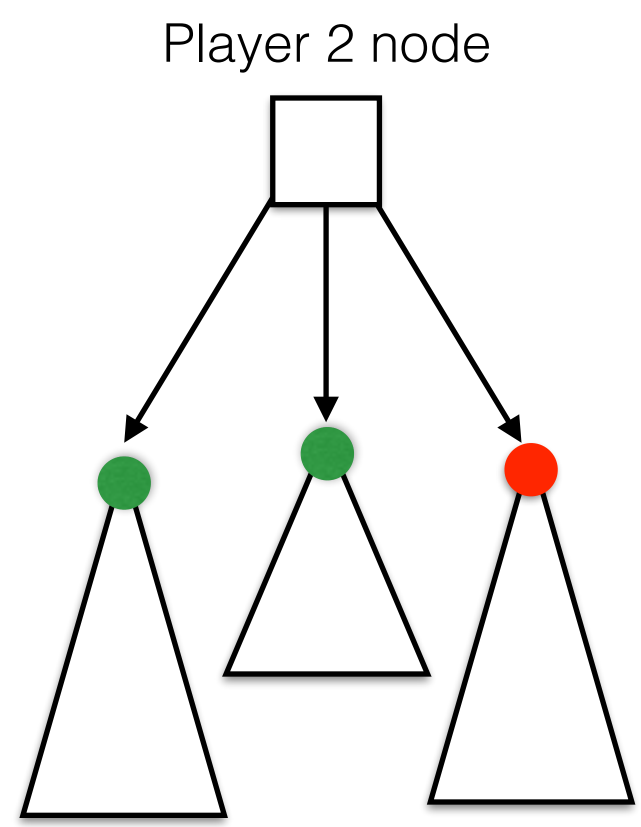

The node is red if there exists one red successor
otherwise it is green

Determinacy of finite tree reachability games

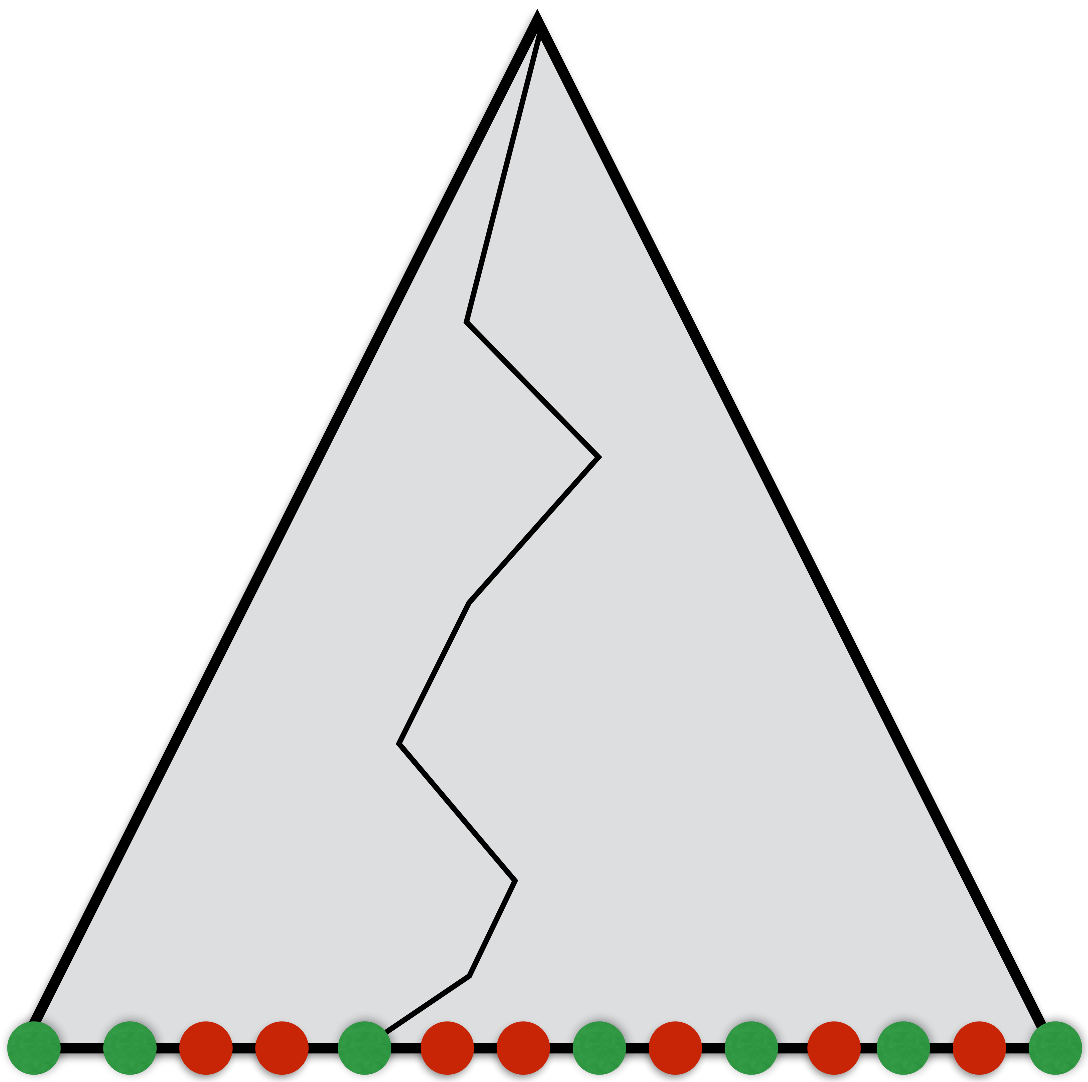

Induction.

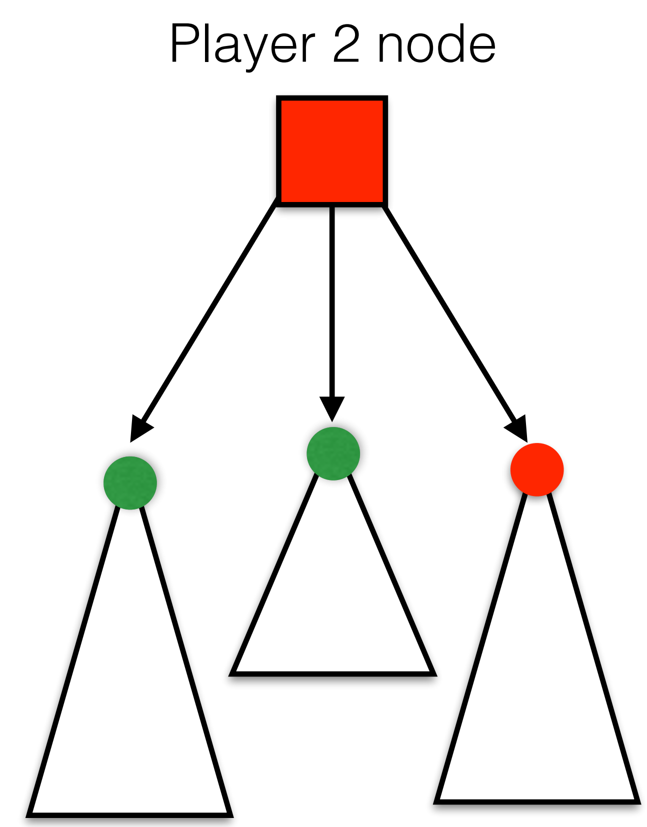

The node is red if there exists one red successor
otherwise it is green

Determinacy of finite tree reachability games

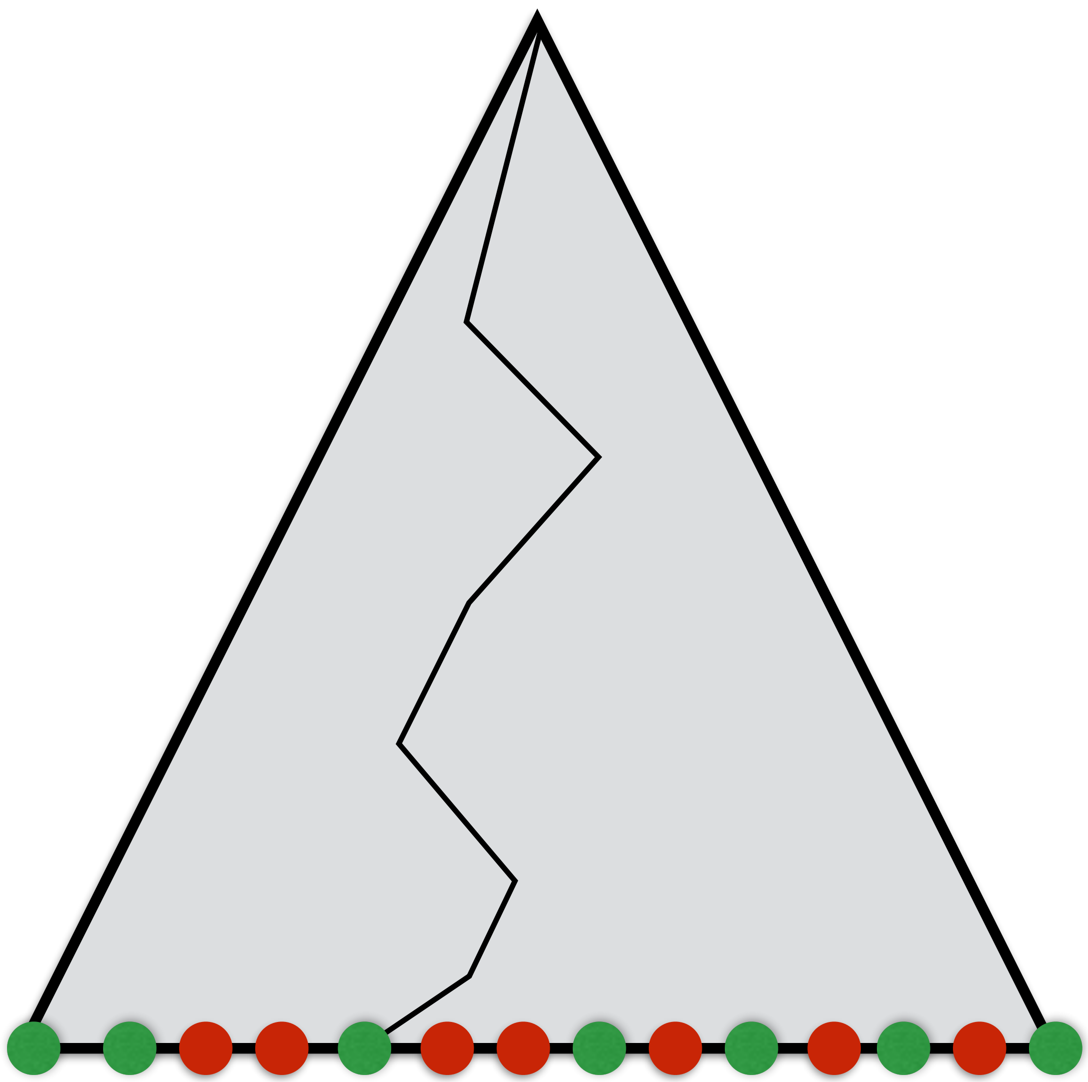

Induction.

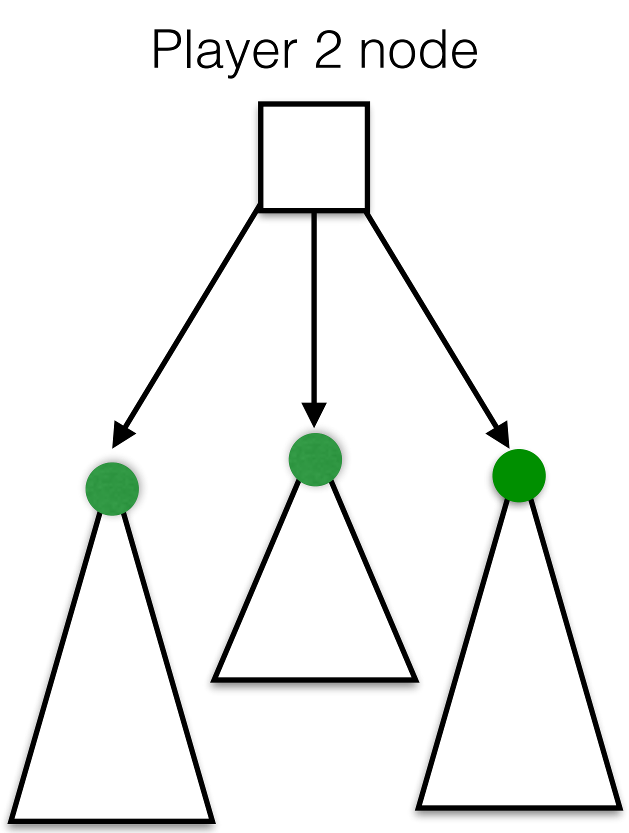

The node is red if there exists one red successor
otherwise it is green

Determinacy of finite tree reachability games

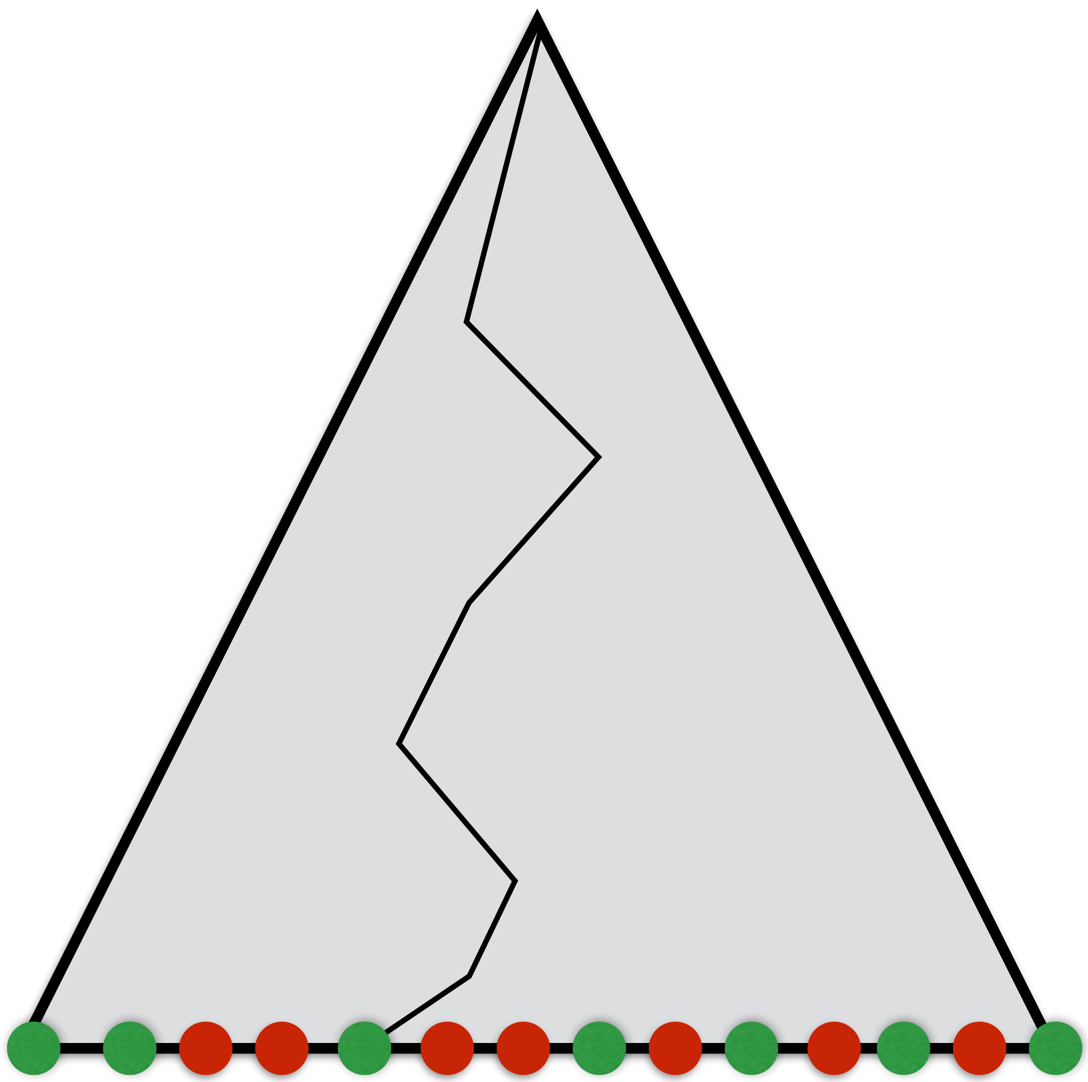

Induction.

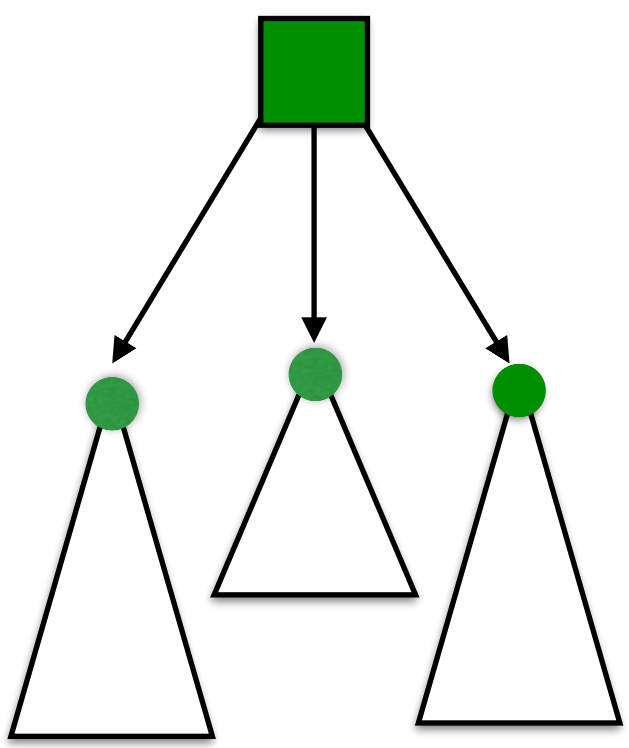

The node is red if there exists one red successor
otherwise it is green

Determinacy of finite tree reachability games

We have established:

Theorem [**Zermelo 1913**]. Every finite tree reachability game is determined.

Unfoldings for solving MPG

Unfold G up to a first repetition of a vertex:

- a leaf is winning for Pl. 1 if the cycle has a nonnegative sum
- a leaf is winning for Pl. 2 if the cycle has a negative sum

☛ By Zermelo theorem:
either Pl. 1 can force positive cycles
or Pl. 2 can force negative cycles

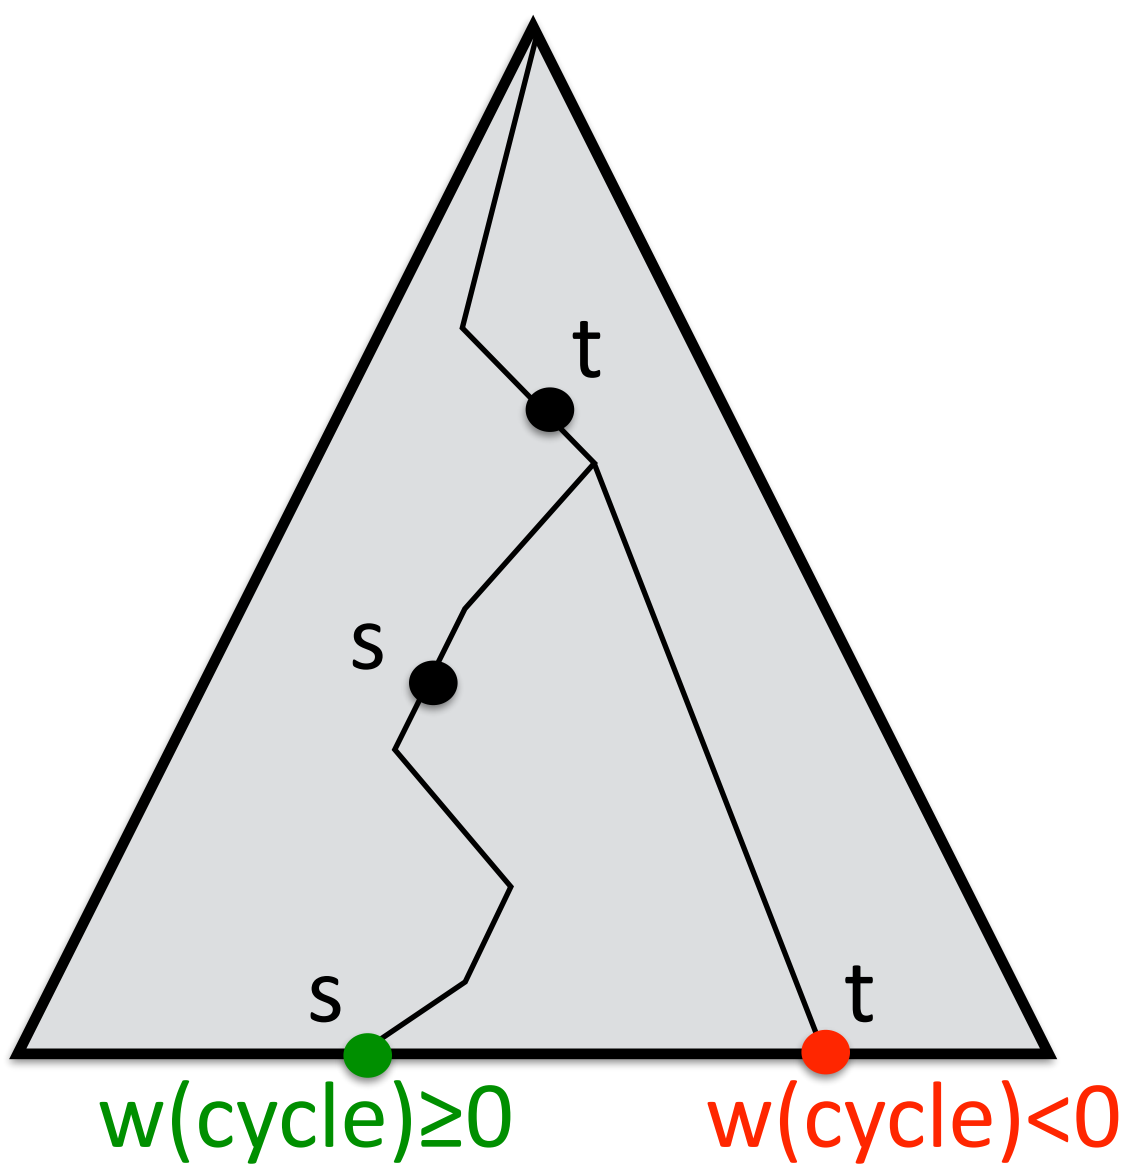

Unfolding - an example

G

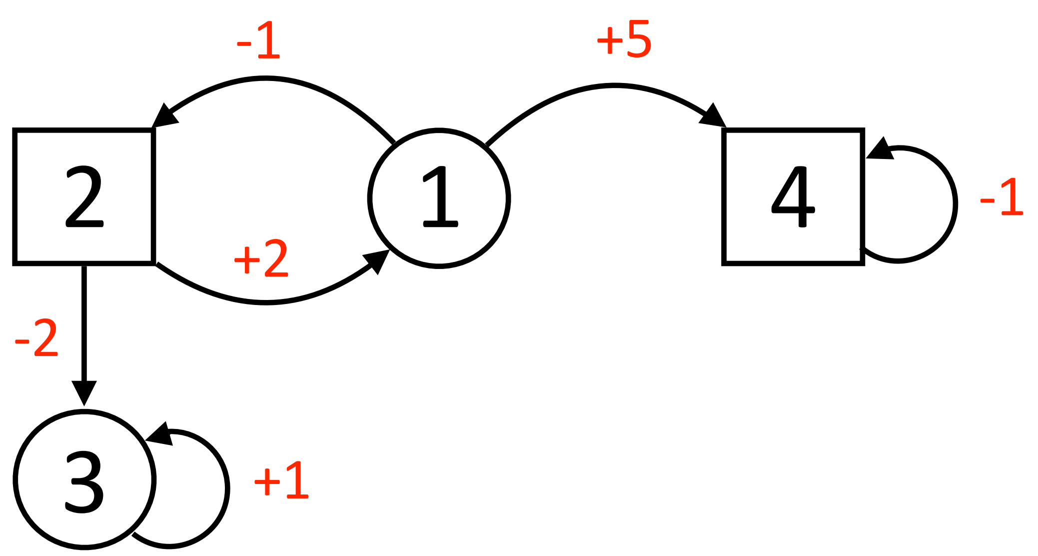

Unfolding of G starting in (1)

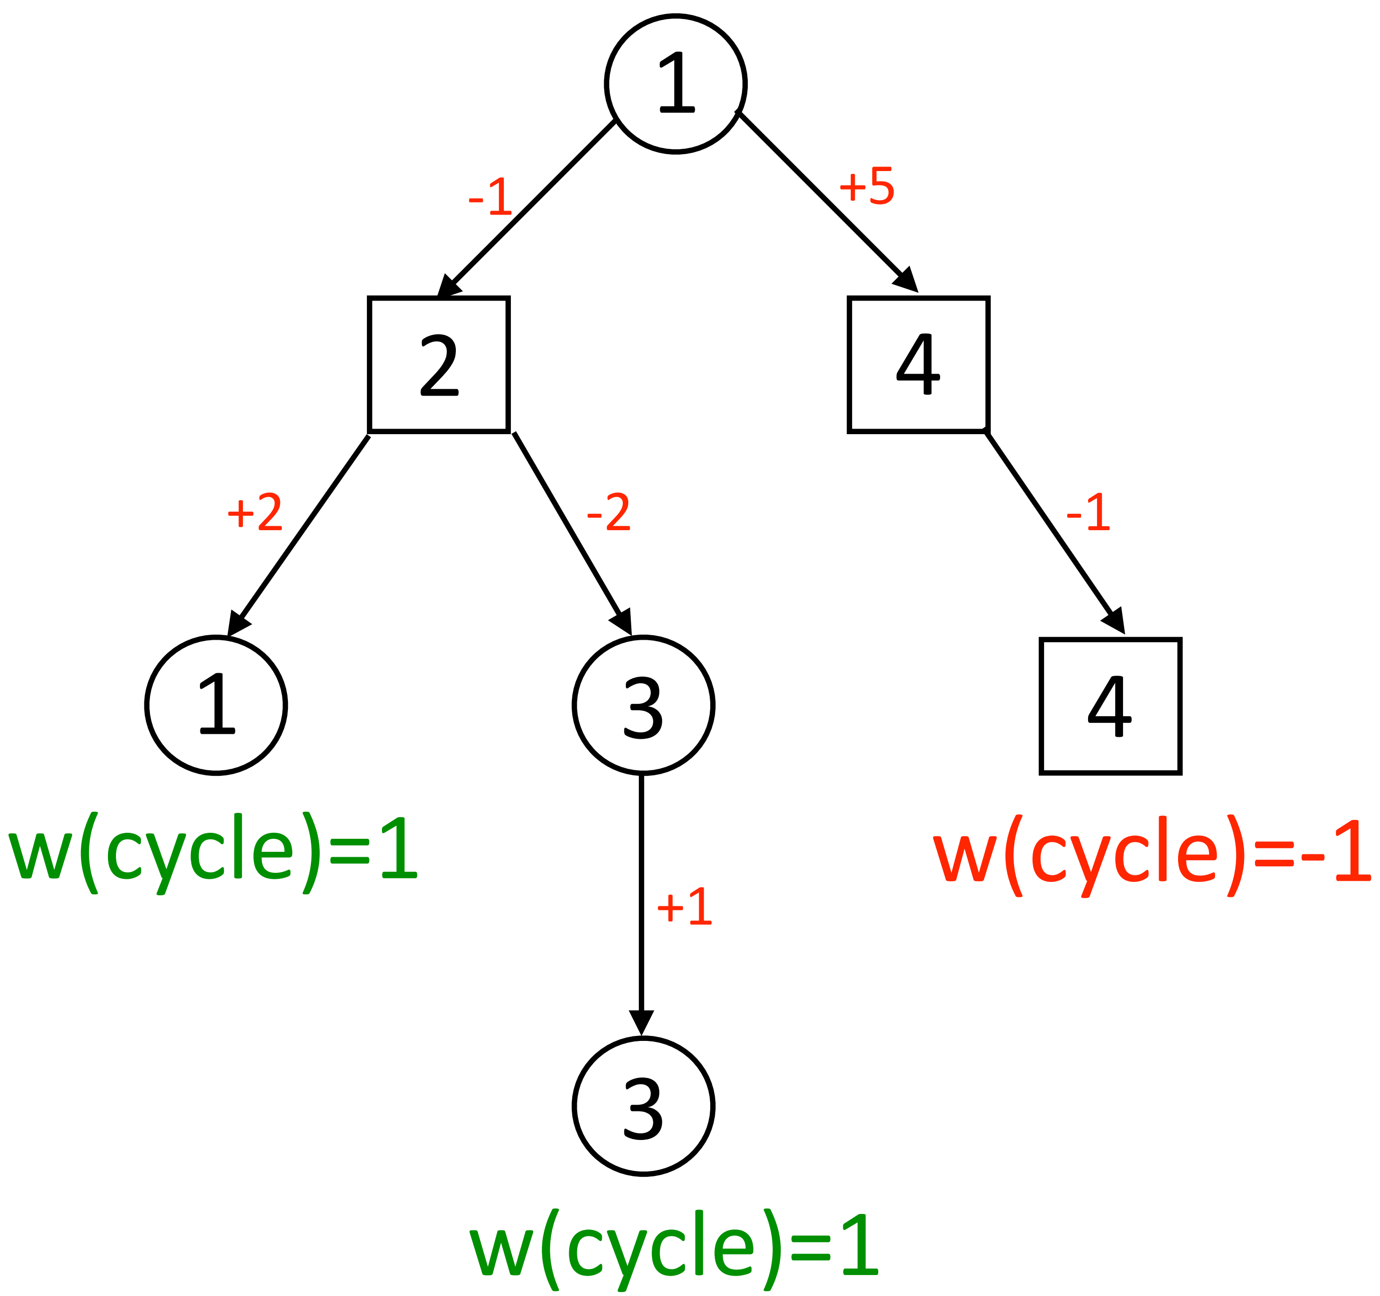

Unfolding - an example

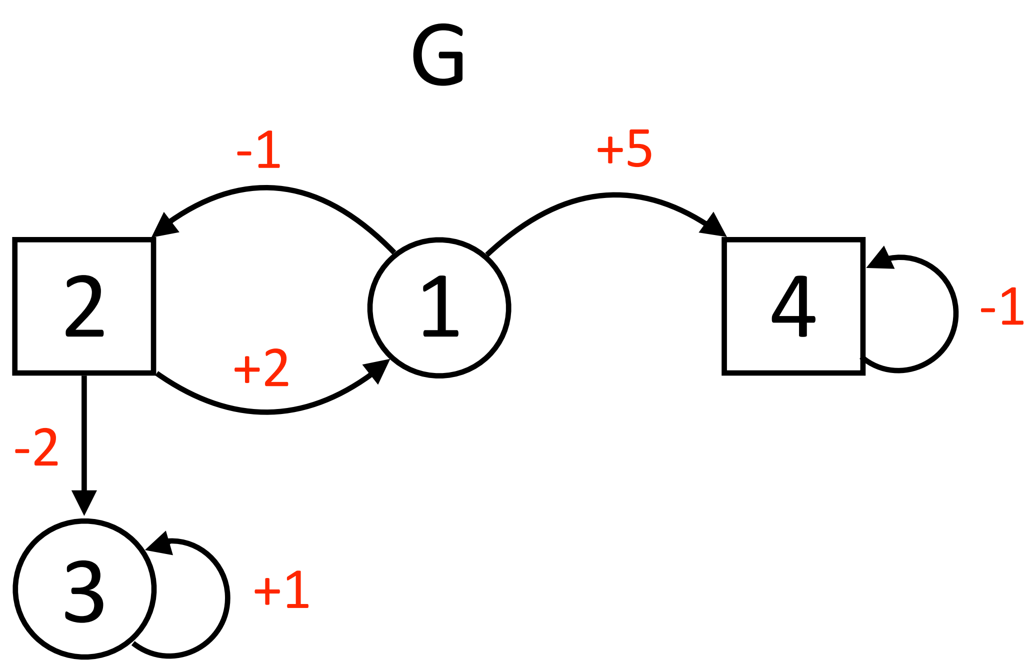

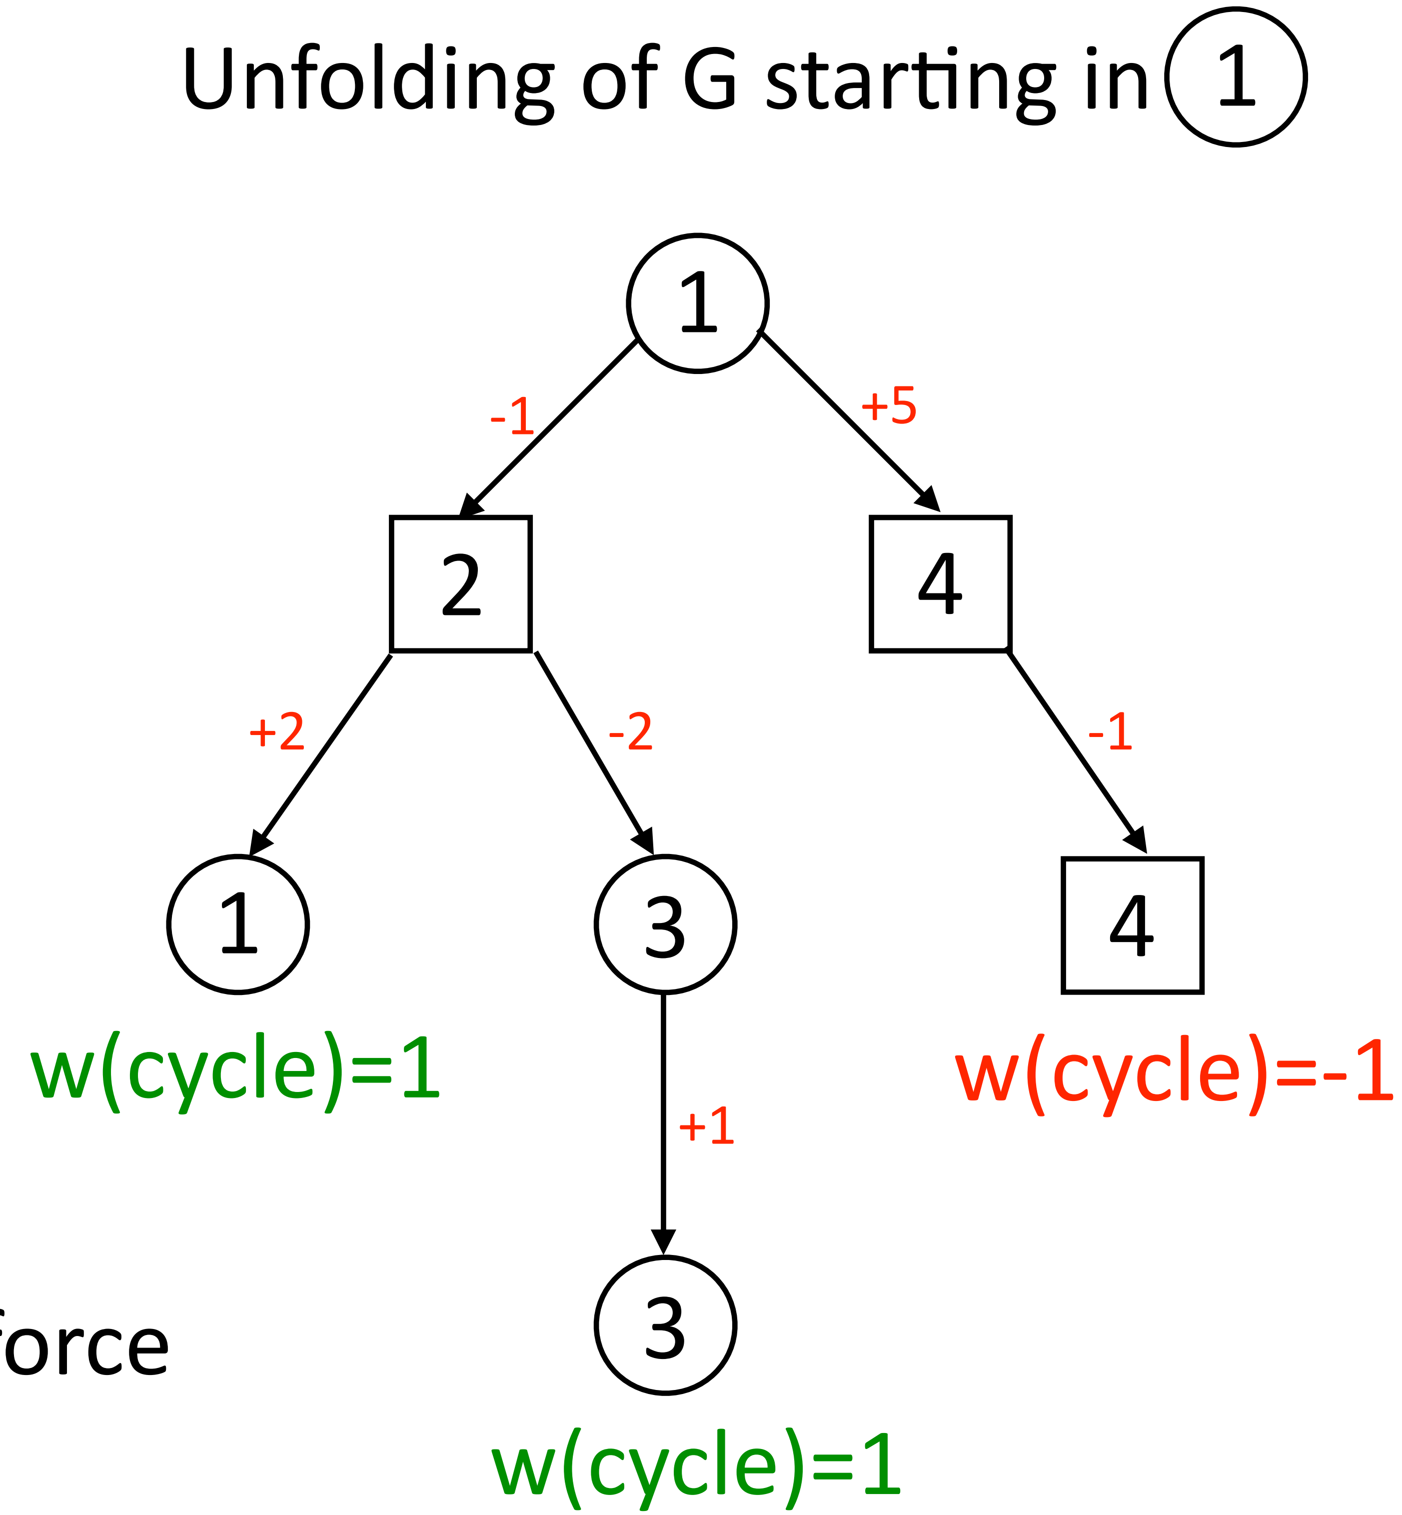

In this example, Player 1 can force nonnegative cycles !

From the tree to the game graph

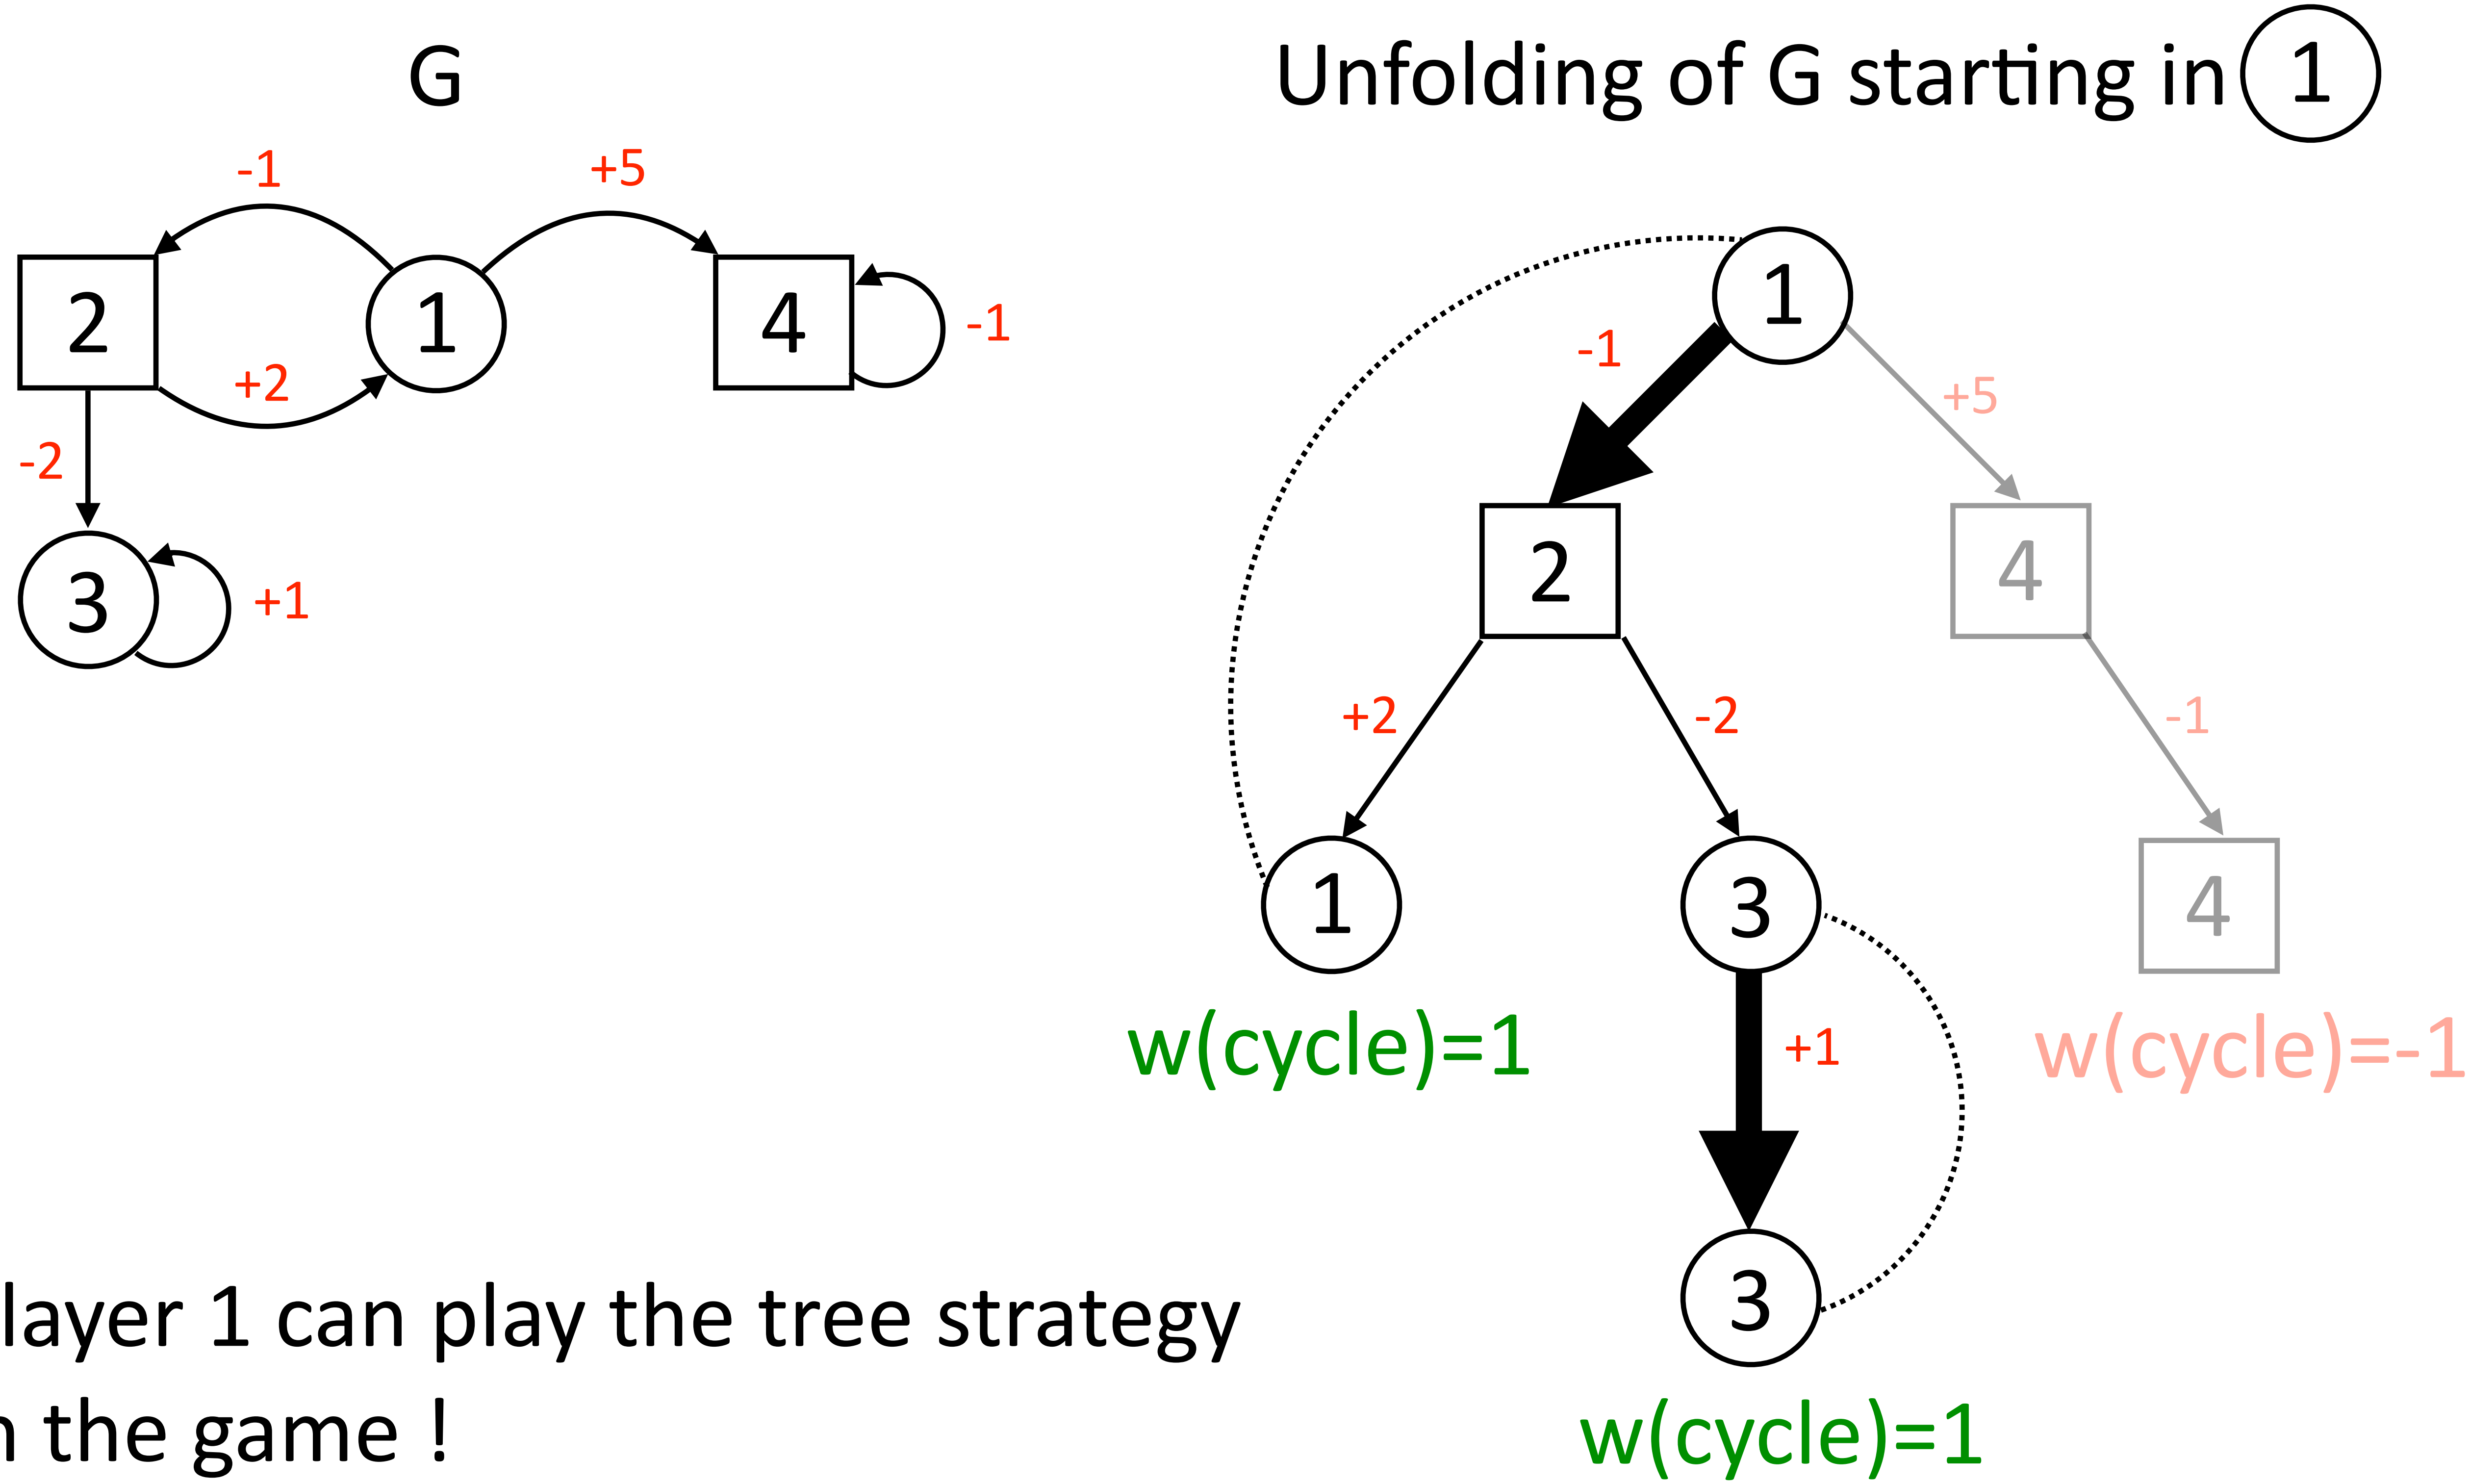

Player 1 can play the tree strategy in the game !

Transfer of strategies - MPG-EG

Lemma [**strategy transfer**] Winning strategies in the tree can be transferred into winning strategies in the **MP/EG** game:

- If Player 1 can force green leaves in the unfolding of G then Player 1 has a winning strategy in G for the MP ≥ 0 objective **and** in the EG;
- If Player 2 can force red leaves in the unfolding of G then Player 2 has a winning strategy in G for the MP < 0 objective **and** in the EG.

To establish this lemma, we rely on the notion of **cycle decomposition** of a play…

… and we will get:

Corollary [**MPG≈EG**] MPG G •≥v and EG G-v are equivalent !

Decomposition of a play into simple cycles

$s_1=s_6$

$\pi = s_0\ e_0\ s_1\ e_1\ s_2\ e_2\ s_3\ e_3\ s_4\ e_4\ s_5\ e_5\ s_6\ e_6\ s_7\ e_7\ldots$

$s_3=s_5$

$st(\pi[..0])=s_0$	$dec(\pi[..0])=\{\}$
$st(\pi[..1])=s_0s_1$	$dec(\pi[..1])=\{\}$
$st(\pi[..2])=s_0s_1s_2$	$dec(\pi[..2])=\{\}$
$st(\pi[..3])=s_0s_1s_2s_3$	$dec(\pi[..3])=\{\}$
$st(\pi[..4])=s_0s_1s_2s_3s_4$	$dec(\pi[..4])=\{\}$
$st(\pi[..5])=s_0s_1s_2s_3s_4s_5$	$dec(\pi[..5])=\{s_3s_4s_5\}$
$st(\pi[..6])=s_0s_1s_2s_5s_6$	$dec(\pi[..6])=\{s_3s_4s_5, s_1s_2s_5s_6\}$
$st(\pi[..7])=s_0s_6s_7$	$dec(\pi[..7])=\{s_3s_4s_5, s_1s_2s_5s_6\}$
…	

Proof ideas for transfer of strategies

- If Player 1 plays a winning tree strategy on the game graph then **all (simple) cycles** obtained during the cycle decomposition of any outcome **have sum of weights ≥ 0**. So, the **running sum** of all prefixes is **bounded from below by -nW** (n=number of states in G, W=absolute value of the largest weight in G).
    - This implies that the **EG** is won by Player 1 **from energy level nW**
    - The **MP** of the play is nonnegative and Player 1 wins **MP ≥ 0**
- If Player 2 plays a winning tree strategy on the game graph then **all (simple) cycles** obtained during the cycle decomposition of any outcome **have a sum of weights ≤ -1**. So, the **running sum** of all prefixes **tends to -∞** and each cycle has a **MP ≤ -1/n**
    - The energy game is won by Player 2 **no matter what is the initial energy level**
    - The MP of the play ≤ **-1/n** (the finite residue on the stack can be neglected) and is won by Player 2

Determinacy of MPG-EG and equivalence

Theorem [**MPG strong determinacy**] For all MPG G, for all states s:

- **either** $\exists \lambda_1$ for Player 1 s.t. Outcome(s, λ_1) $\subseteq \{ \pi \mid MP(\pi)$ **≥ 0** $\}$,
- **or** $\exists \lambda_2$ for Player 2 s.t. Outcome(s, λ_2) $\subseteq \{ \pi \mid MP(\pi)$ **$\leq -1/n$** $\}$.

Corollary [**MPG determinacy**] For all MPG G, for all states s:

- **either** $\exists \lambda_1$ for Player 1 s.t. Outcome(s, λ_1) $\subseteq \{ \pi \mid MP(\pi) \geq 0 \}$,
- **or** $\exists \lambda_2$ for Player 2 s.t. Outcome(s, λ_2) $\subseteq \{ \pi \mid MP(\pi) < 0 \}$.

Theorem [**Determinacy-EG**] For all EG games G, for all states s:

- **either** there exists an initial EL and a strategy for Player 1 from s to win the energy game,
- **or** Player 2 has a strategy from s to win the energy game, no matter what is the initial EL.

Theorem [**Equivalence MPG-EG**] For all games G, for all states s, Player 1 wins for **MP ≥ 0** from s if and only if Player 1 wins the **energy game** from s.

‼ Strategies in the tree are **not** guaranteed to be **memoryless** !

‼ We **need additional arguments** to prove memoryless determinacy …

Need for memory in the game tree

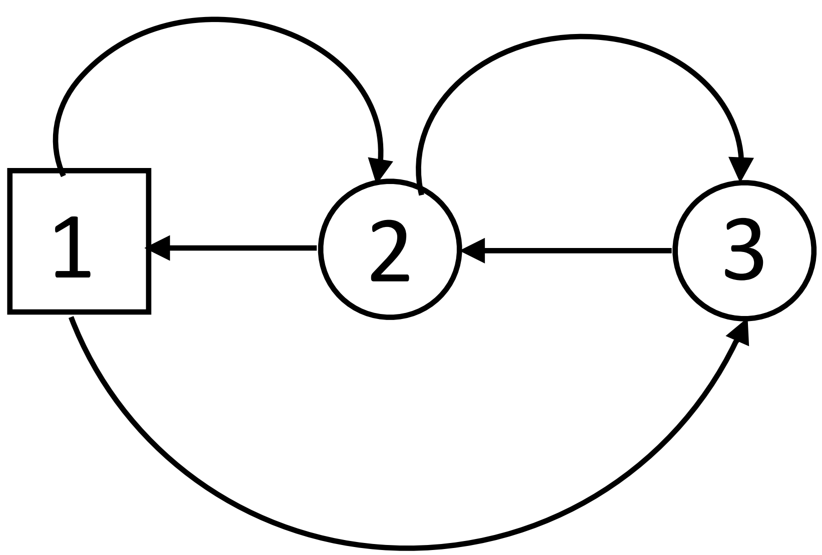

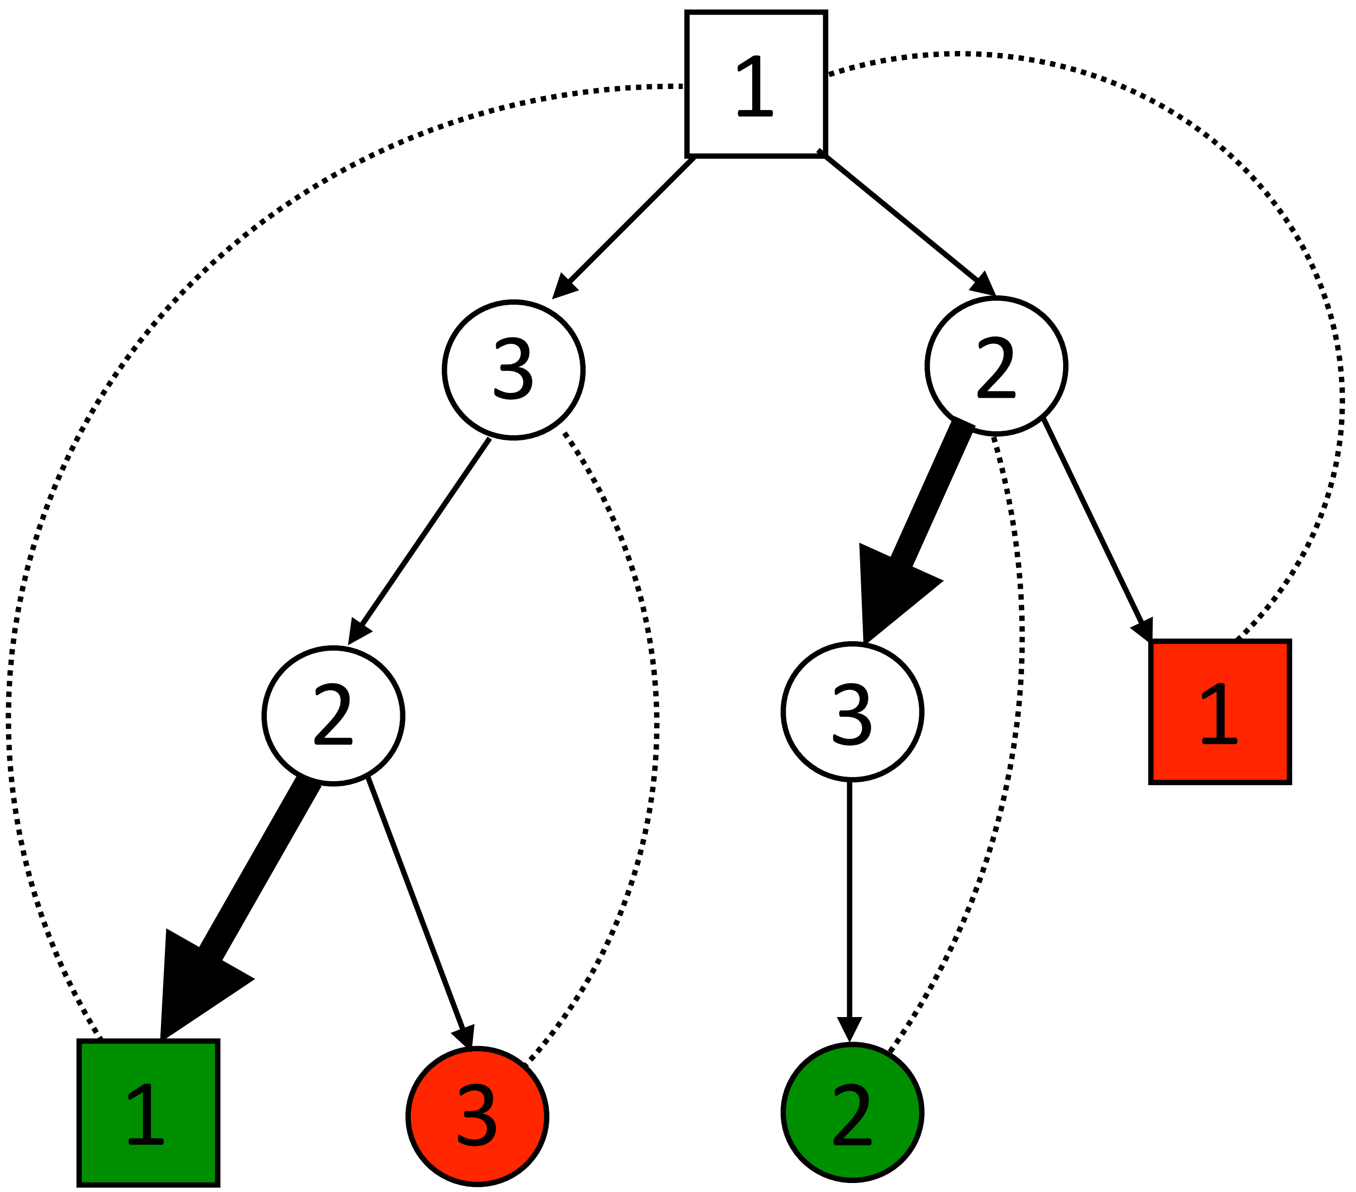

Choice in 2 is not uniform:
it depends on the history
=need for memory

A pseudo-polynomial time algorithm for energy games

EG and safety

$SAFE_i$=set of $(s,c) \in S_1 \times \mathbb{N} \cup S_2 \times \mathbb{N}$ s.t. from (s,c),
Player 1 can maintain energy level nonnegative for **i steps**

What are the **controllable predecessors** of $SAFE_i$?

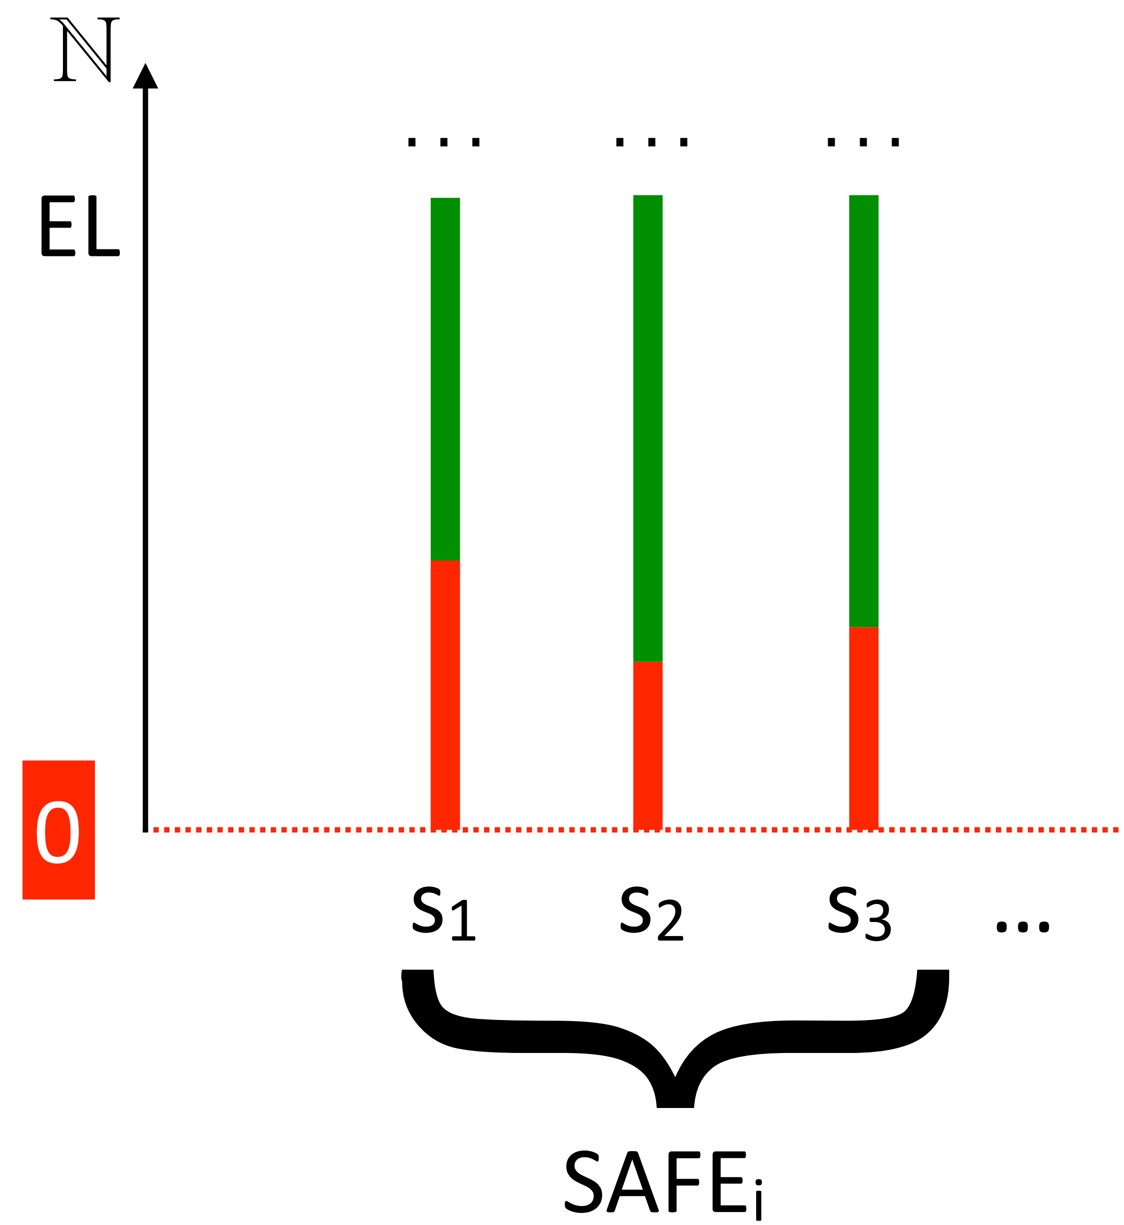

CPRE(X)

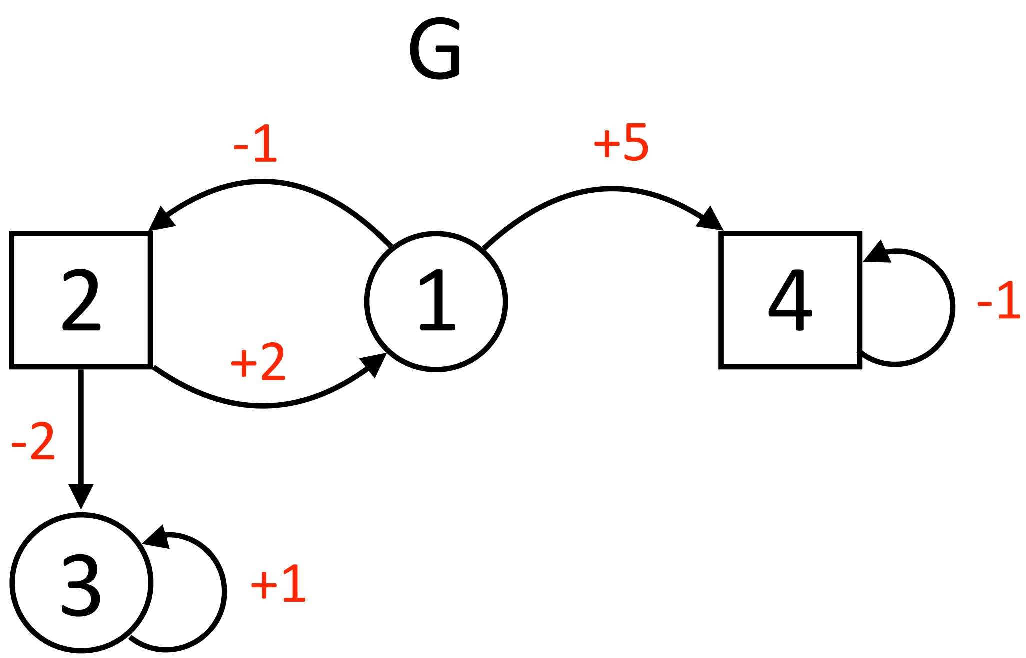

CPRE(X) where $X \subseteq (S_1 \times \mathbb{N} \cup S_2 \times \mathbb{N})$ is the set

$$\{ (s_1,c) \mid \exists\, (s_1,w,s') \in E : (s',c+w) \in X \}$$

$$\cup \quad \{ (s_2,c) \mid \forall\, (s_2,w,s') \in E : (s',c+w) \in X \}$$

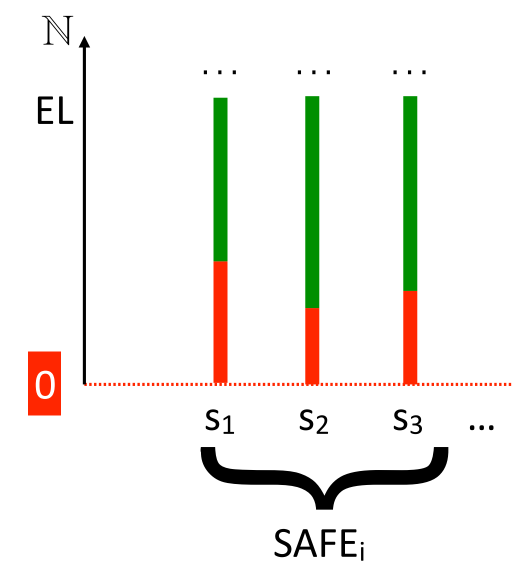

- We define $\preccurlyeq$ as $(s,c) \preccurlyeq (s',c')$ **iff** $s=s'$ and $c \leq c'$
- CPRE(X) transforms $\preccurlyeq$-upper-closed sets into $\preccurlyeq$-upper-closed sets
- Give a set X, we write $\uparrow X$ for its $\preccurlyeq$-upper-closure

CPRE(X)

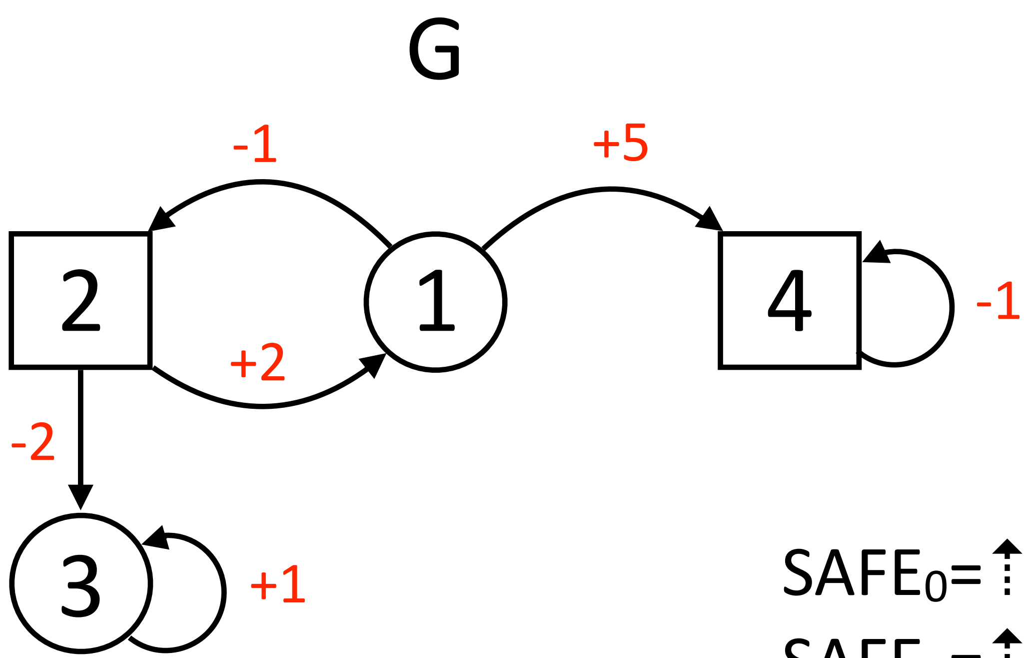

CPRE(X) where $X \subseteq (S_1 \times \mathbb{N} \cup S_2 \times \mathbb{N})$ is the set

$\{ (s_1,c) \mid \exists\, (s_1,w,s') \in E : (s',c+w) \in X \}$

$\cup \quad \{ (s_2,c) \mid \forall\, (s_2,w,s') \in E : (s',c+w) \in X \}$

$SAFE_0 = \uparrow\{(s_1,0),(s_2,0),(s_3,0),(s_4,0)\}$
$SAFE_1 = \uparrow\{(s_1,0),(s_2,2),(s_3,0),(s_4,1)\}$
$SAFE_2 = \uparrow\{(s_1,0),(s_2,2),(s_3,0),(s_4,2)\}$
$SAFE_3 = \uparrow\{(s_1,0),(s_2,2),(s_3,0),(s_4,3)\}$
$SAFE_4 = \uparrow\{(s_1,0),(s_2,2),(s_3,0),(s_4,4)\}$
$SAFE_5 = \uparrow\{(s_1,0),(s_2,2),(s_3,0),(s_4,5)\}$
$SAFE_6 = \uparrow\{(s_1,1),(s_2,2),(s_3,0),(s_4,6)\}$
$SAFE_7 = \uparrow\{(s_1,2),(s_2,2),(s_3,0),(s_4,7)\}$
$SAFE_8 = \uparrow\{(s_1,3),(s_2,2),(s_3,0),(s_4,8)\}$
$SAFE_9 = \uparrow\{(s_1,3),(s_2,2),(s_3,0),(s_4,9)\}$
…
$SAFE_k = \uparrow\{(s_1,3),(s_2,2),(s_3,0),(s_4,k)\}$

no stabilisation !

CPRE[C](X) to force termination

- Above energy requirement $C \in \mathbb{N}$, we consider the game as lost ! (conservative approximation)
- Let $C \in \mathbb{N}$, define $U(C)=\mathbb{P}((S_1 \cup S_2)\times\{0 \ldots C\})$
- CPRE[C](X) where $X \in U(C)$ is the set

$$\{ (s_1,c) \in S_1\times\{0 \ldots C\} \mid \exists\ (s_1,w,s') \in E : (s',c+w) \in X \}$$

$$\cup \quad \{ (s_2,c) \in S_2\times\{0 \ldots C\} \mid \forall\ (s_2,w,s') \in E : (s',c+w) \in X \}$$

CPRE[C](X) - properties

CPRE*[C](X) is monotone

- so it has a greatest fixed point, noted CPRE*[C]
- computed iteratively from $\top = \mathbb{P}((S_1 \cup S_2) \times \{0 \ldots C\})$
- convergence is ensured now as the lattice is finite
- The greatest fixpoint can be computed in O(|V|.|E|.W), where W is the largest weight in absolute value in G. So the complexity is **pseudo-polynomial**

CPRE[C](X)

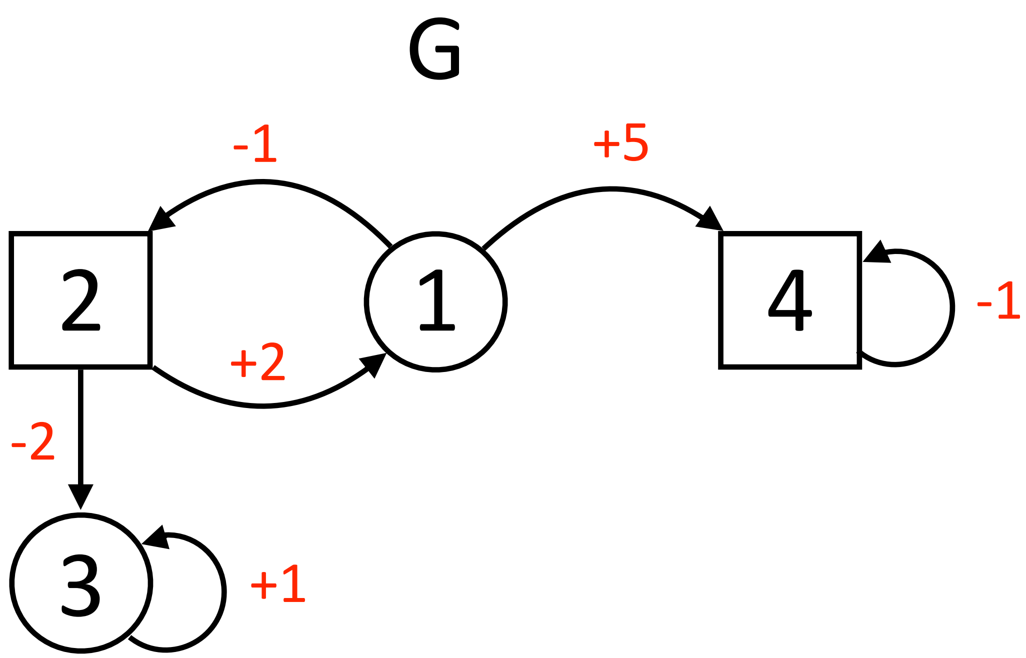

CPRE[3](X) where $X \in U(C)$ is the set

$$\{ (s_1,c) \in S_1 \times \{0 \ldots 3\} \mid \exists\, (s_1,w,s') \in E : (s',c+w) \in X \}$$

$$\cup \quad \{ (s_2,c) \in S_2 \times \{0 \ldots 3\} \mid \forall\, (s_2,w,s') \in E : (s',c+w) \in X \}$$

$SAFE_0 = \uparrow\{(s_1,0),(s_2,0),(s_3,0),(s_4,0)\}$

$SAFE_1 = \uparrow\{(s_1,0),(s_2,2),(s_3,0),(s_4,1)\}$

$SAFE_2 = \uparrow\{(s_1,0),(s_2,2),(s_3,0),(s_4,2)\}$

$SAFE_3 = \uparrow\{(s_1,0),(s_2,2),(s_3,0),(s_4,3)\}$

$SAFE_4 = \uparrow\{(s_1,3),(s_2,2),(s_3,0)\}$

$SAFE_5 = \uparrow\{(s_1,3),(s_2,2),(s_3,0)\} = \mathbf{SAFE_\infty}$

Stabilisation !

Greatest fixpoint

Correctness of the FP algorithm

Theorem [**correctness**] ∀ C ∈ ℕ, ∀ (s,c) ∈ ⇡CPRE*[C], Player 1 wins EG from s.

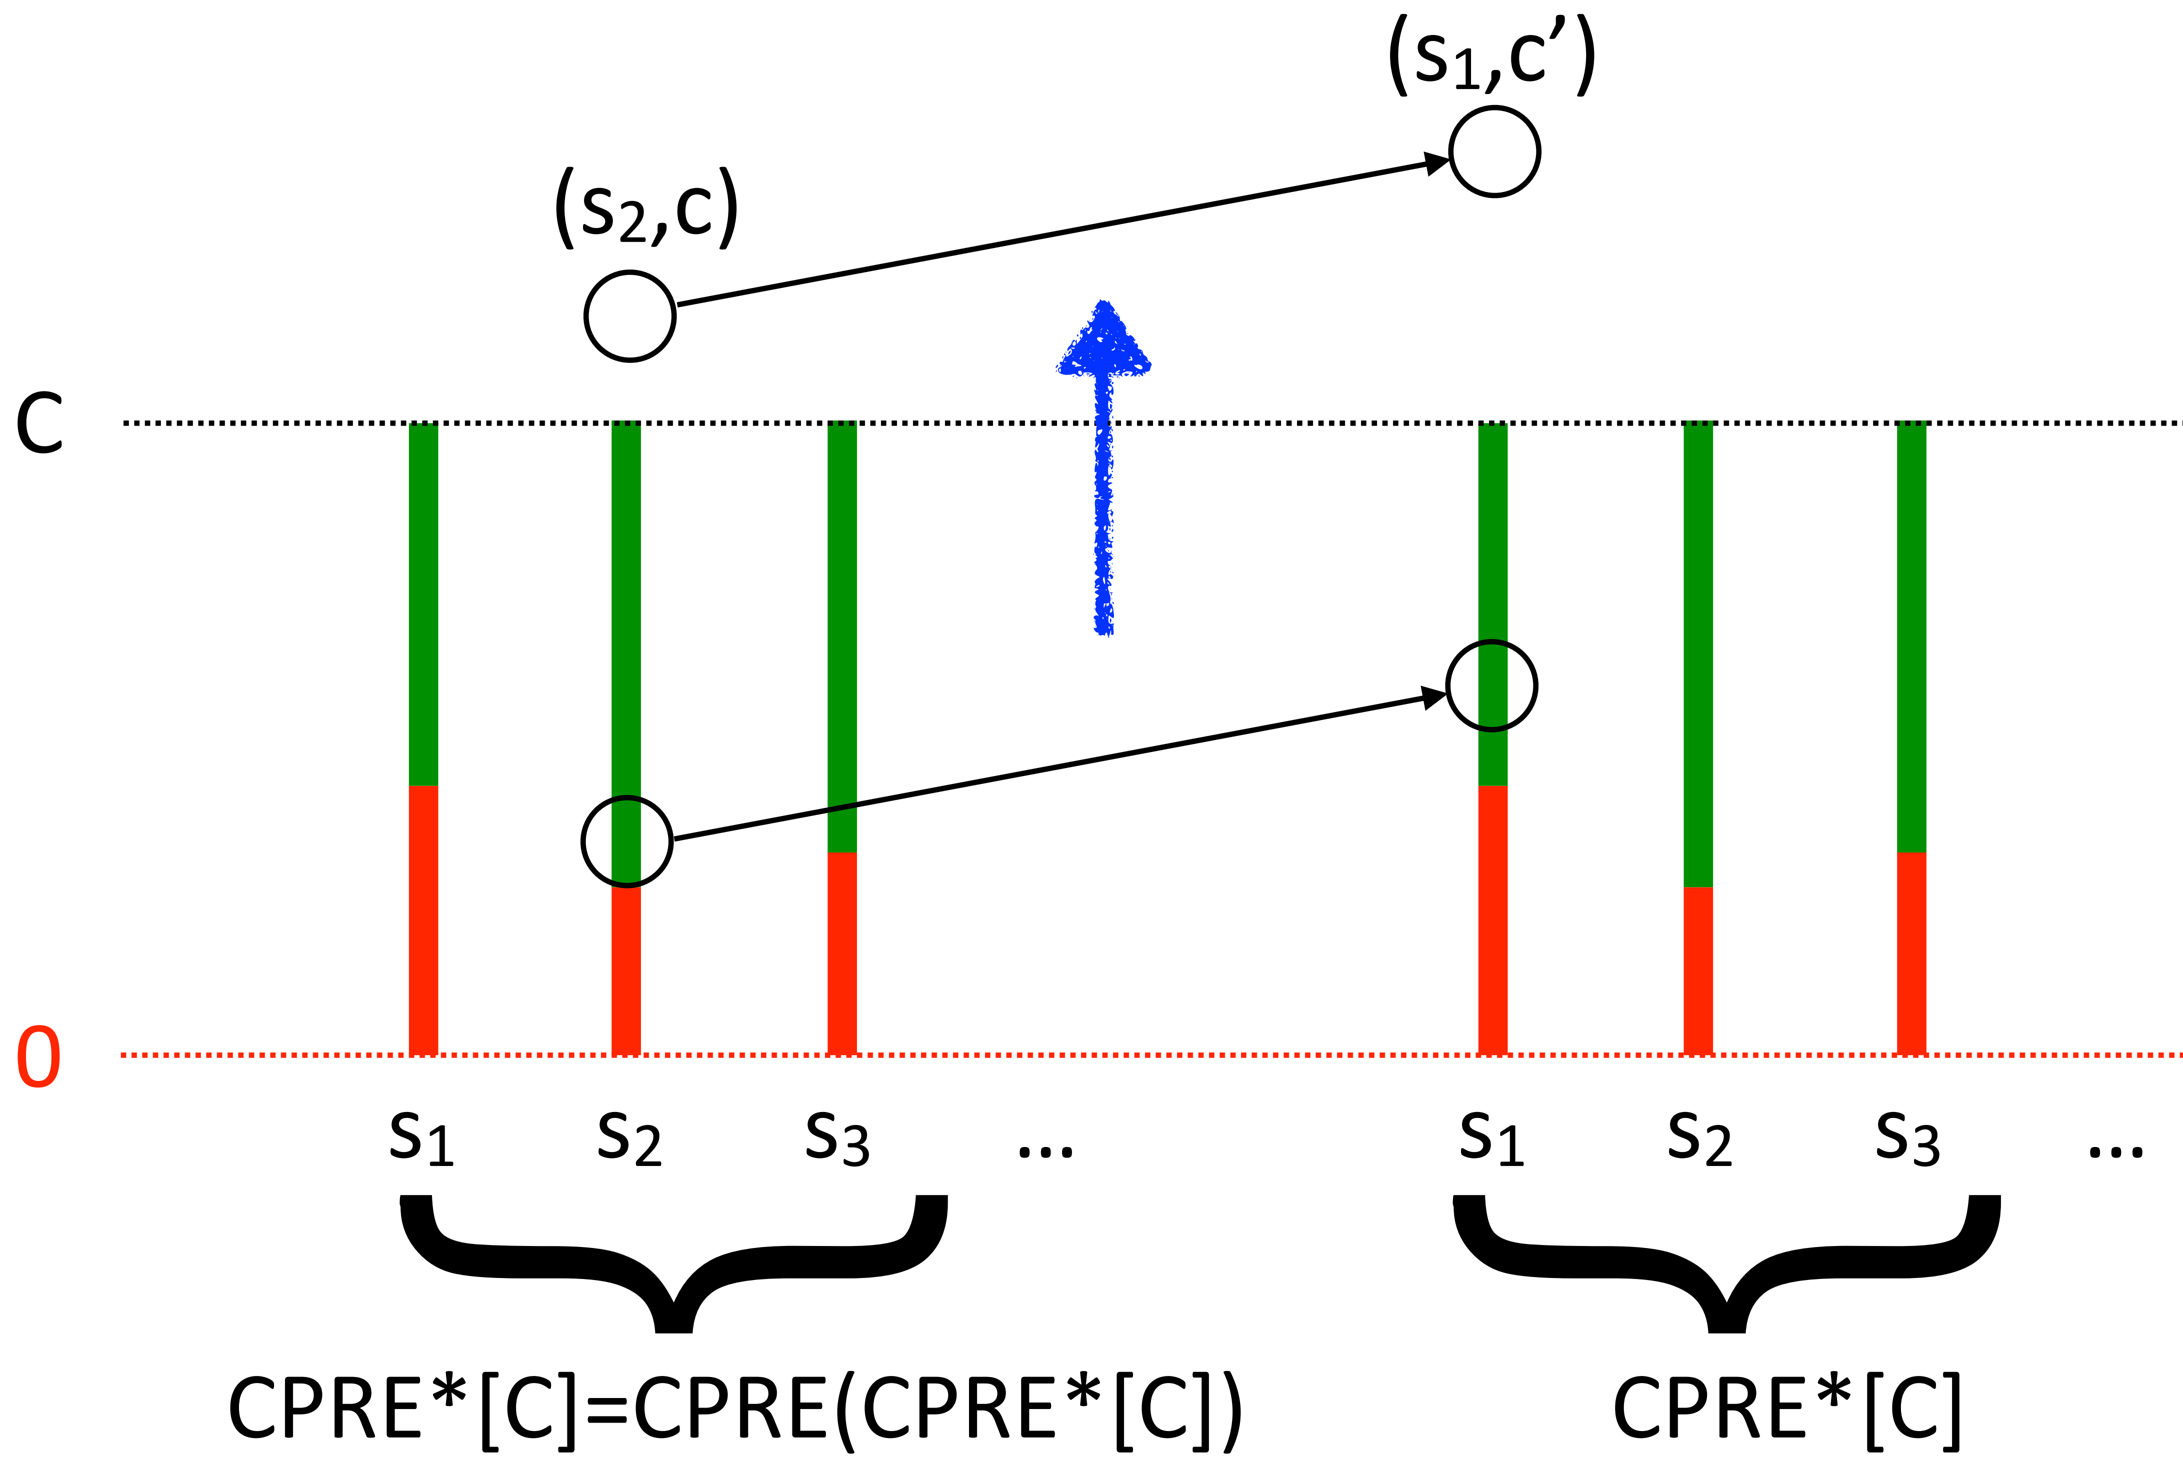

Correctness of the FP algorithm

Theorem [**correctness**] $\forall$ C $\in \mathbb{N}$, $\forall$ (s,c) $\in$ ⇡CPRE*[C], Player 1 wins EG from s.

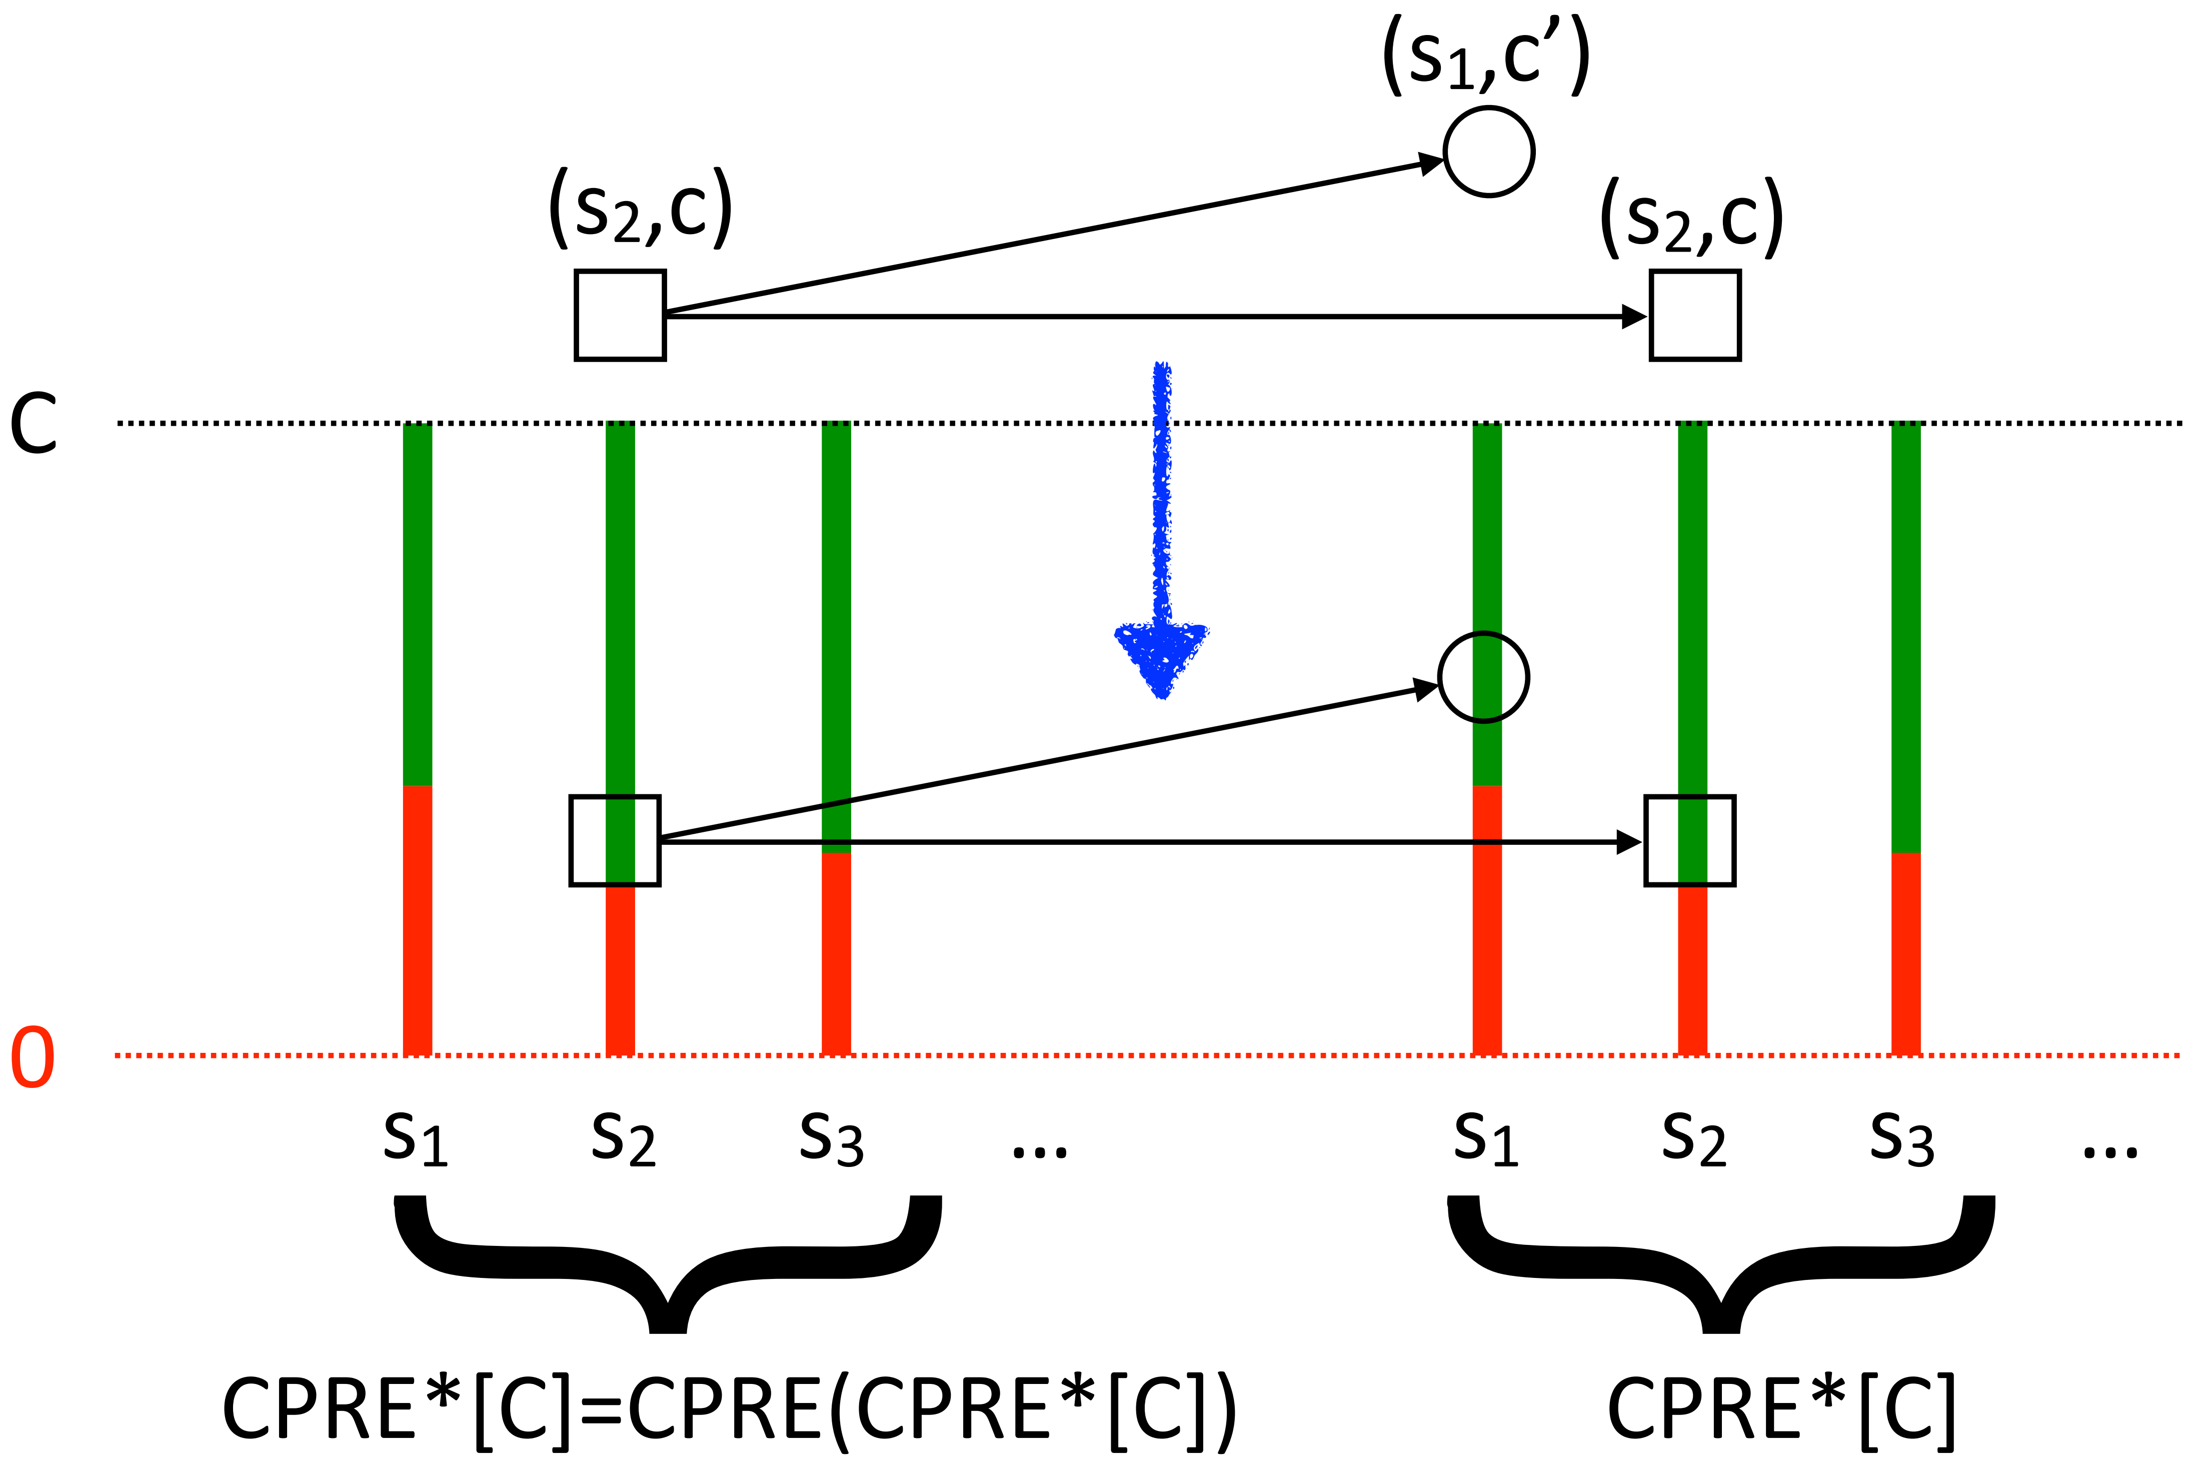

Correctness of the FP algorithm

Theorem [**correctness**] $\forall$ C $\in \mathbb{N}$, $\forall$ (s,c) $\in$ ⇡CPRE*[C], Player 1 wins EG from s with initial energy level c.

Proof. Assume we start in state s with energy level c. We construct a strategy for Player 1 s.t. if $(s_1,c_1)(s_2,c_2),\ldots,(s_n,c_n),\ldots$ is an outcome then for all positions $i \geq 0,(s_i,c_i) \in$ ⇡CPRE*[C]. So the energy level always stays nonnegative.

The proof is by induction. We consider two cases.

(1) $(s_{i-1},c_{i-1}) \in$ ⇡CPRE*[C] and $s_{i-1} \in S_1$, consider (s,c)$\in$ CPRE*[C], with s_{i-1}=s and $c_{i-1} \geq c$. From (s_{i-1},c_{i-1}) Player 1 chooses e=(s,w,s')$\in$E such that there exists (s',c') $\in$ CPRE*[C] such that c+w ≥ c'. As CPRE*[C] is a FP, such an edge exists. So, we have that (s_i,c_i) is such that s_i=s', $c_i=c_{i-1}+w \geq c+w \geq c'$ and so $(s_i,c_i) \in$ ⇡CPRE*[C].

(2) $(s_{i-1},c_{i-1}) \in$ ⇡CPRE*[C] and $s_{i-1} \in S_2$, Player 2 has chosen the edge (s_{i-1},w,s_i). Let (s,c) $\in$ CPRE*[C] be s.t. s_{i-1}=s and $c_{i-1} \geq c$. By definition of CPRE*[C], there exists $(s_i,c') \in$ CPRE*[C] such that c+w ≥ c'. So, we have that $c_i=c_{i-1}+w \geq c+w \geq c'$ and we are done.

Completeness of the FP algorithm

Theorem [**completeness**] Let C=2nW. If Player 1 has a winning strategy from s in EG, then there exists $(s,c) \in CPRE^*[C]$.

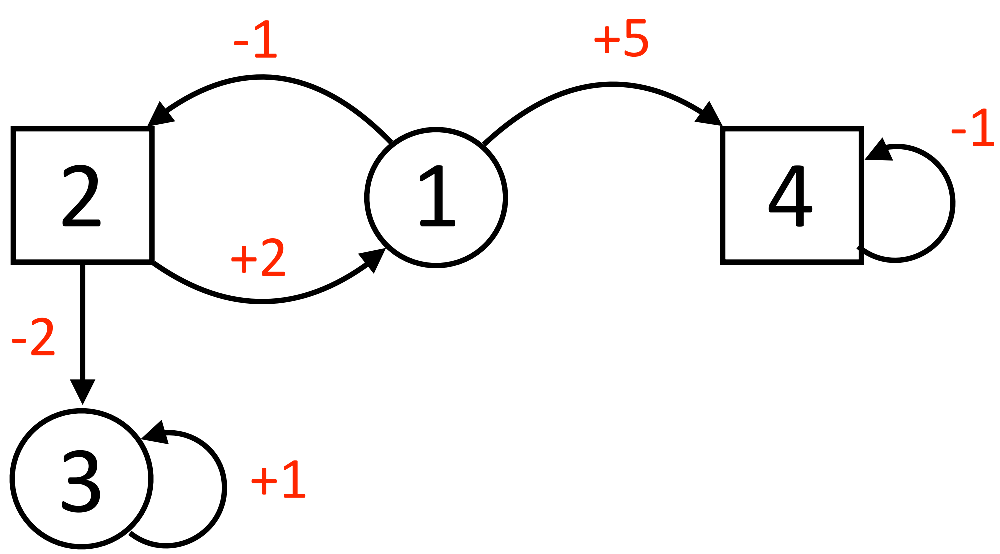

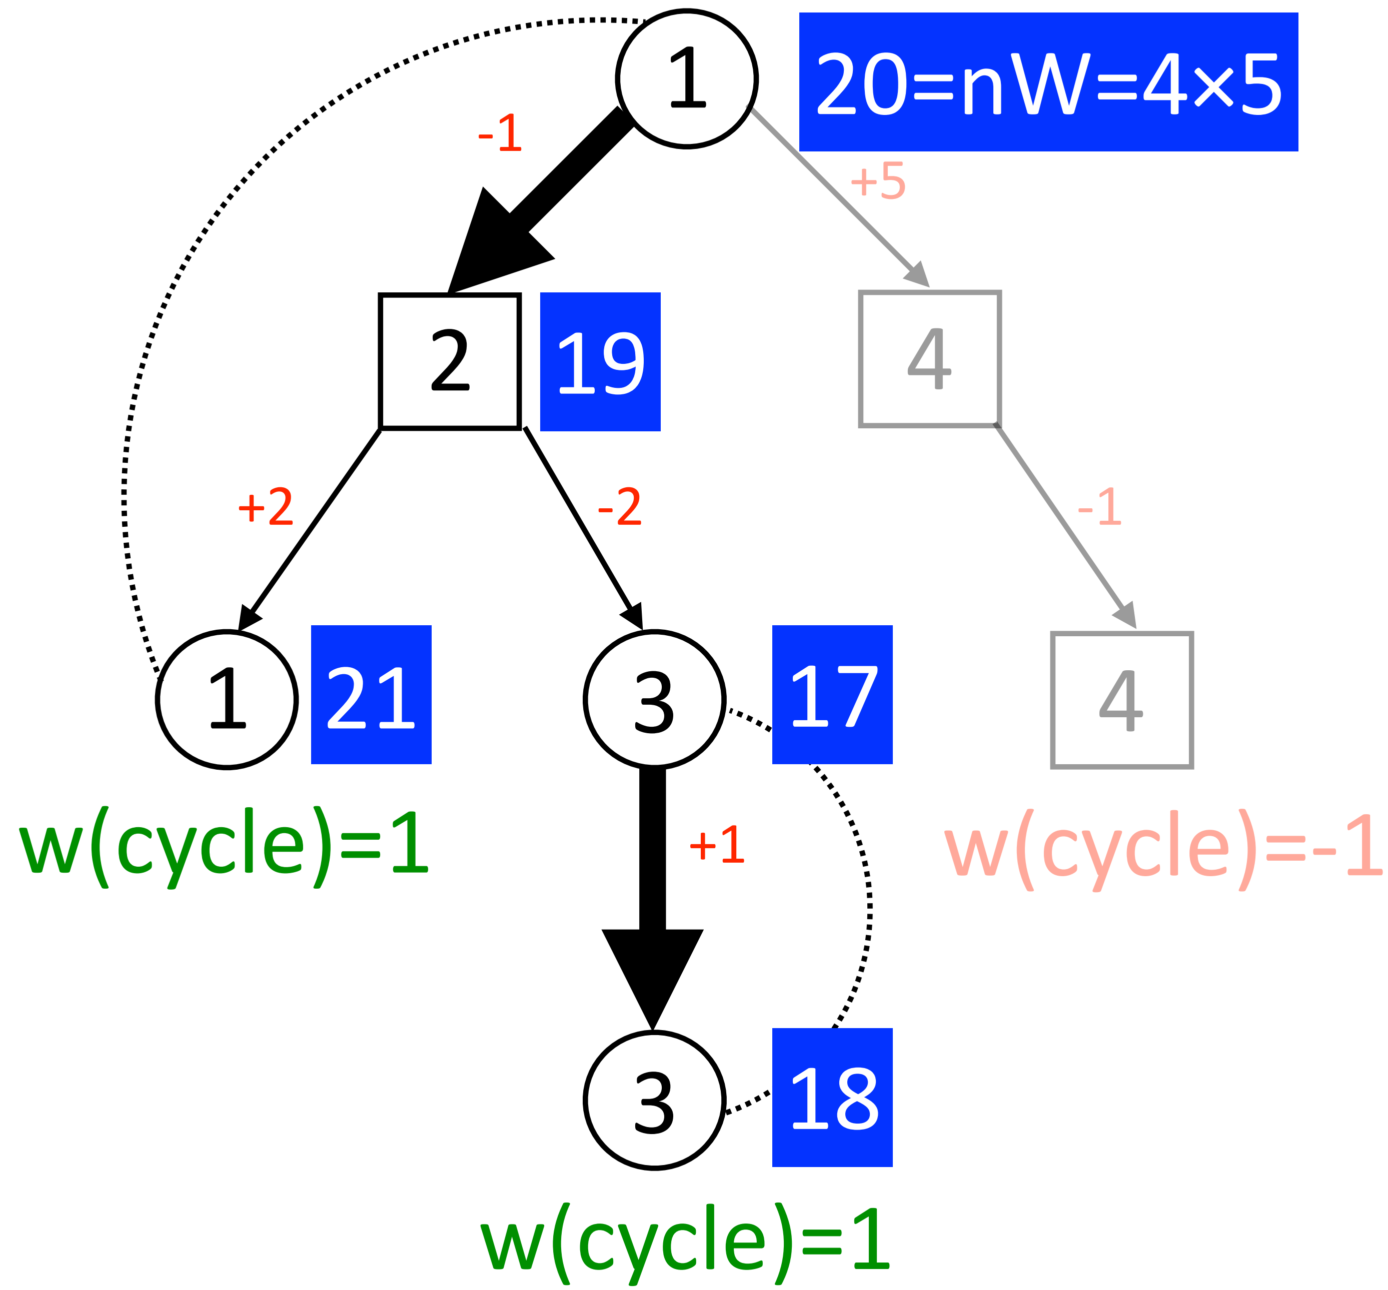

Claim:

$F = \uparrow\{(s_1,20),(s_2,19),(s_3,17)\}$

is a FP of CPRE[20]

So, $F \subseteq CPRE^*[20]$!

Completeness of the FP algorithm

Theorem [**completness**] Let C=2nW. If Player 1 has a winning strategy from s_0 in EG, then there exists $(s_0,c) \in$ CPRE*[C].

Proof. Consider the tree unfolding (up to a first repetition) of G from s_0 and the equivalent reachability game. If Player 1 wins the EG from s_0, then Player 2 cannot win the tree reachability game (because of strategy transfer) and so by determinacy Player 1 has a winning strategy in the tree.

Consider any strategy of Player 1 in the tree and the subtree induced by that strategy. We annotate the subtree as follows. The root is labelled with weight nW. Then we label the other nodes starting from the root by maintaining the energy level on each history.

It is easy to see that this tree only contains energy levels c such that $0 \leq c \leq 2nW$, indeed each branch is of length at most n and so from energy level nW, we can gain at most nW and lose at most nW.

Let $F=\{ (s,c) \in S\times\{0,..,2nW\} \mid$ there is a node n in the tree labelled with s and c' and $c \geq c'$ $\}$. Clearly, $(s_0,nW) \in F$, and F is a fix point for the operator CPRE[2nW]. So $F \subseteq$ CPRE*[2nW] (as CPRE*[2nW] is the greatest fixpoint), and $(s_0,nW) \in$ CPRE*[2nW].

Memoryless strategies for Player 1 in EG

Fixpoint and
good actions
for Player 1

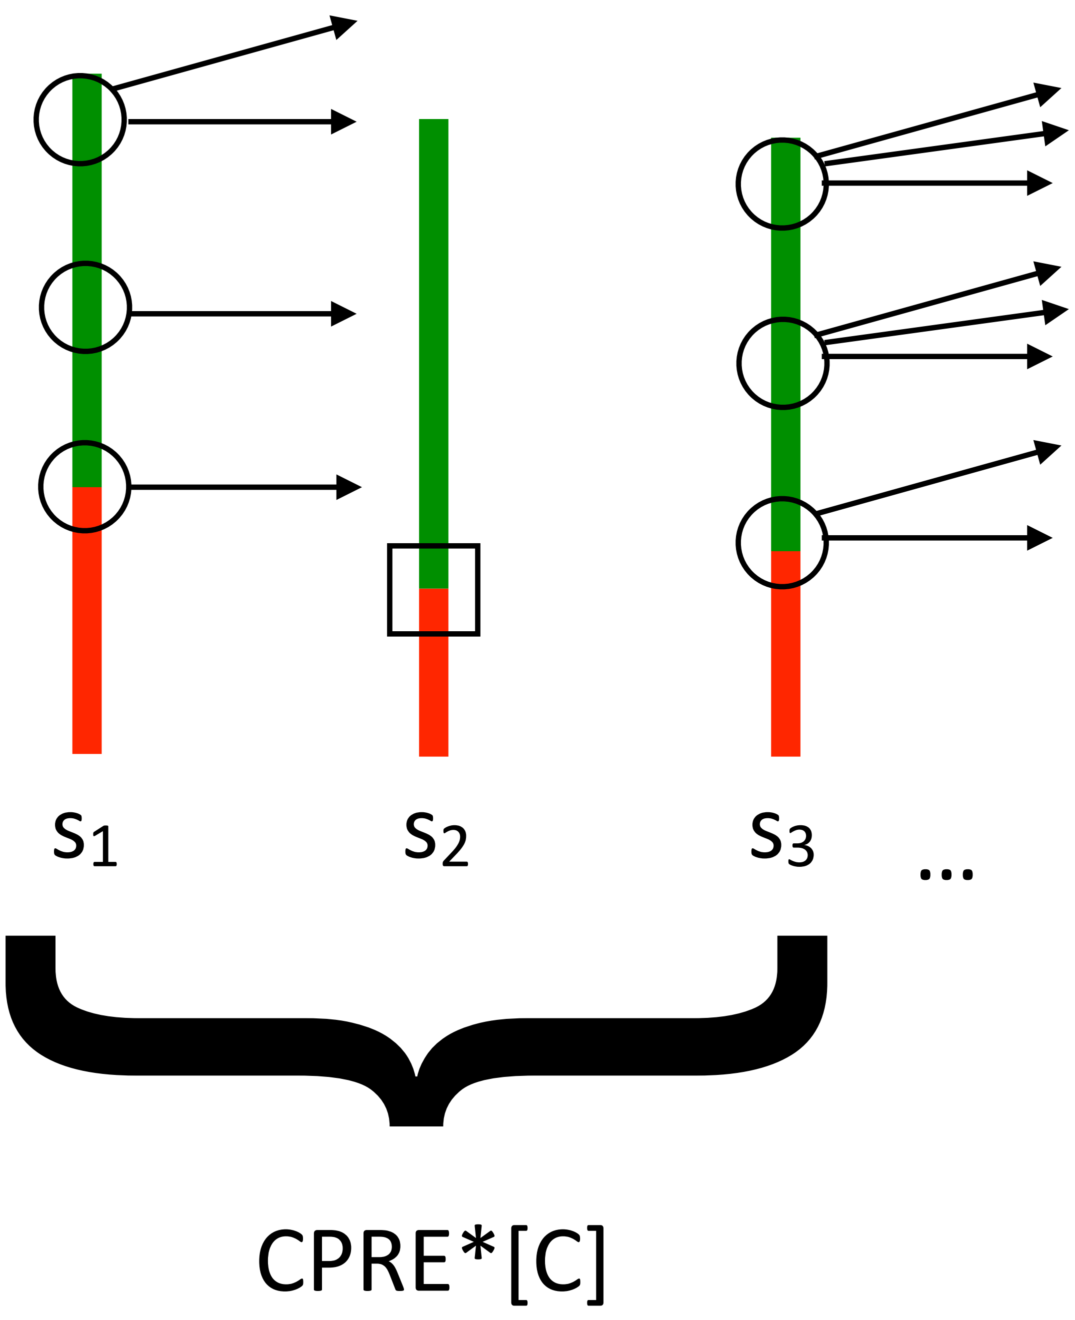

Memoryless strategies for Player 1 in EG

Fixpoint and
good actions
for Player 1

Important property:

actions that are good for EL c
are also good for all EL c’>c

=Monotonicity
implies
Memoryless

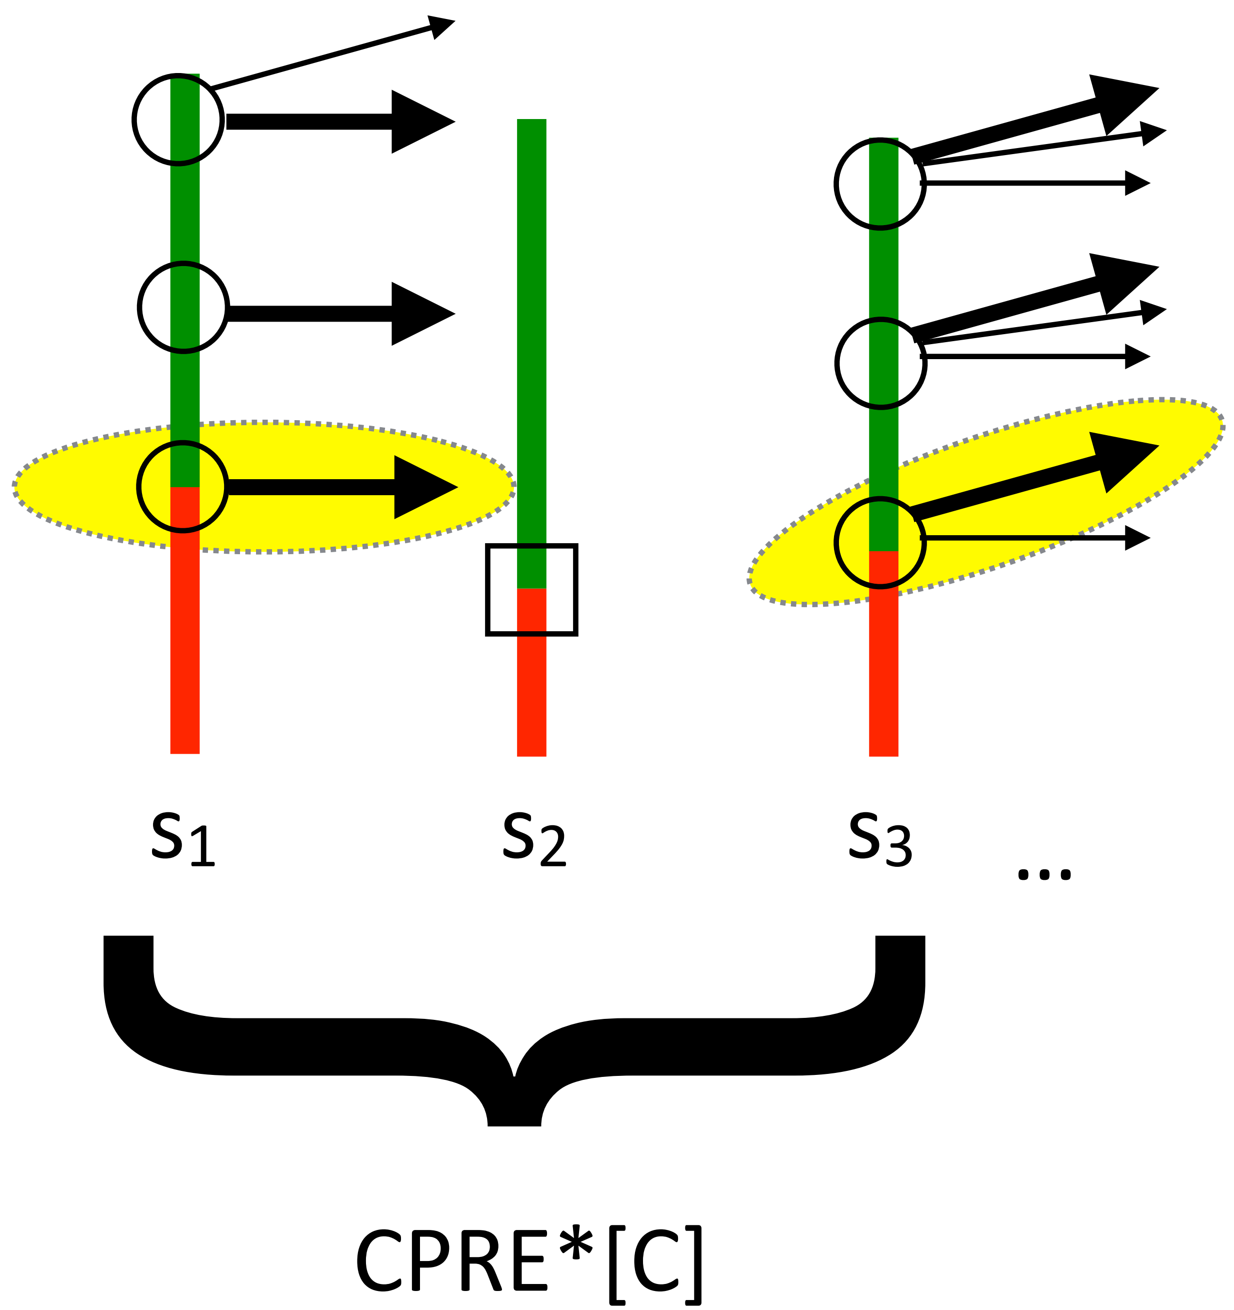

Memoryless strategies for Player 1 in EG

- Player 1 wins the EG from $Win_1 = \{ s \mid \exists (s,c) \in CPRE^*[C] \}$.
- For each $s \in Win_1 \cap S_1$, consider (s,c) where c is minimal (worst-case situation) in $CPRE^*[C]$.
- From each minimal pair (s,c), fix for s an edge (s,w,s') such $(s,c+w) \in CPRE^*[C]$.

Theorem **[memoryless]**. Strategy λ_1 is a memoryless (uniform) strategy λ_1 which is winning from each Player 1 winning state of G in the EG.

Memoryless strategies for Player 1 and 2 in MPG

As a corollary of **MPG≈EG**, and strong determinacy of MPG, we get:

Theorem [**memoryless determinacy of MPG**]
Mean-payoff games are memoryless determined, i.e. both Player 1 and Player 2 can play optimally with memoryless strategies.

Proof. Player 1 can play memoryless as it is the case in EG. For Player 2, we do the following reasoning: Player 2 can enforce **MP ≤ -1/n** in G if and only if Player 2 can enforce **MP ≥ 0** in G' where G' is equal to G in which weight **w** in G is replaced by **-w-1/n** in G'. So Player 2 can play optimally with a memoryless in G.

An we can go back to EG:

Corollary [**memoryless strategies for Player 2 in EG**].
If Player 2 wins EG from s then he has a memoryless winning strategy from s.

Multi-dim. mean-payoff and energy games

Multi-dim. Energy Games (MEGs)

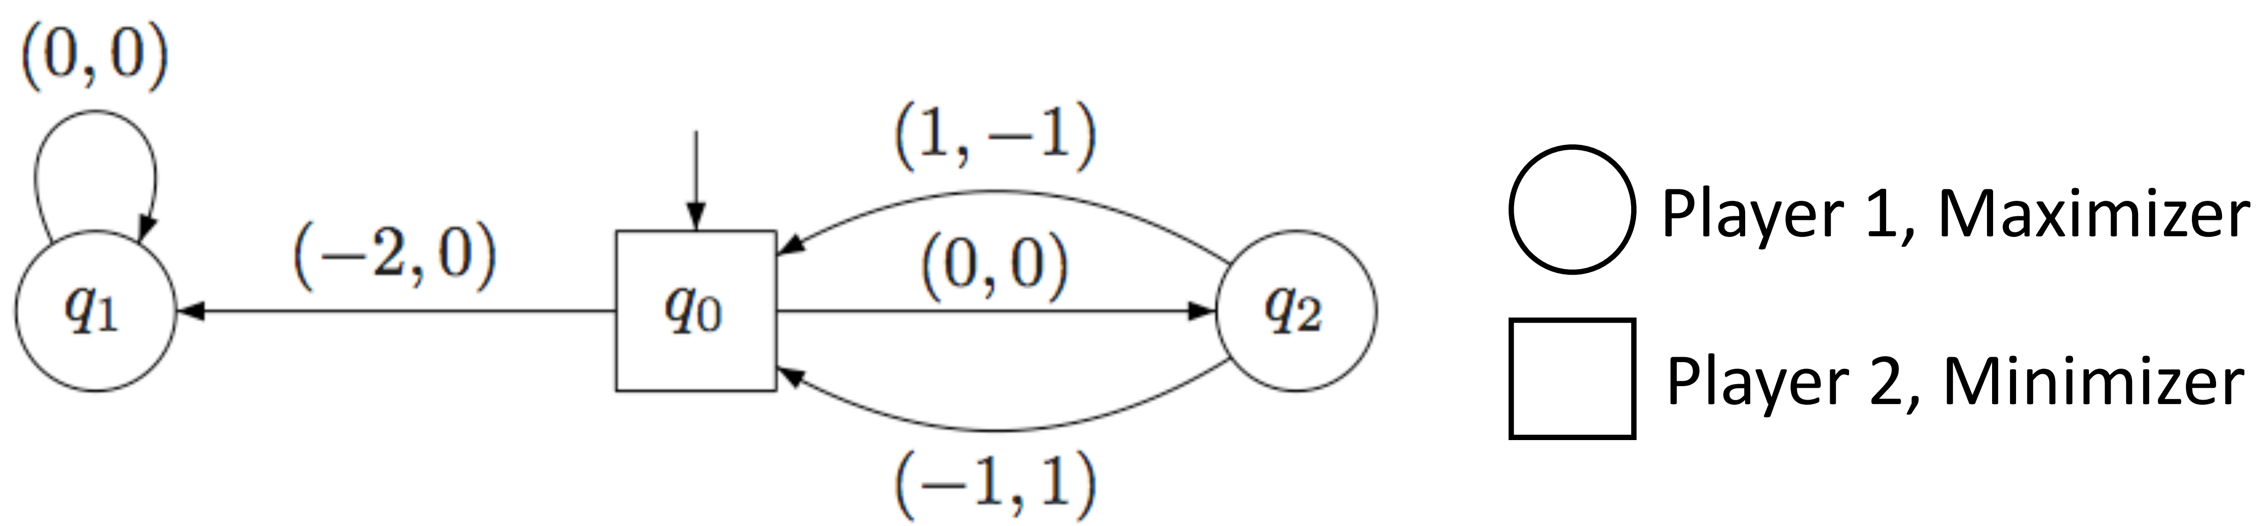

? $\exists (C_1,C_2) \in \mathbb{N}^2$ and λ_1 s. t. **Outcome**$(q_0,\lambda_1,(C_1,C_2)) \models \square$ $\mathbf{EL_1} \geq 0 \wedge \mathbf{EL_2} \geq 0$.

Multi-dim. Energy Games (MEGs)

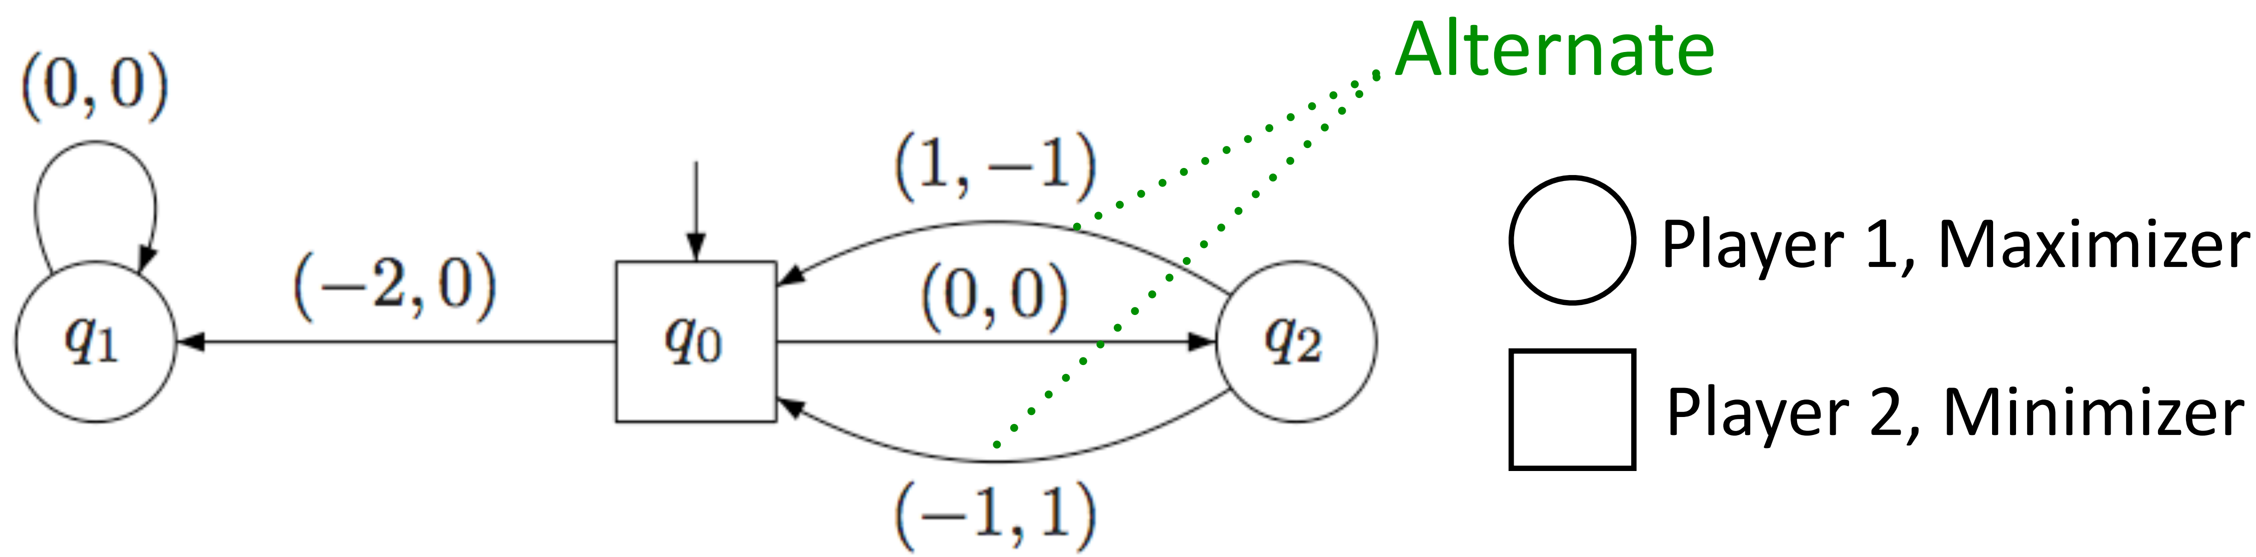

? $\exists (C_1,C_2) \in \mathbb{N}^2$ and λ_1 s. t. **Outcome**$(q_0,\lambda_1,(C_1,C_2)) \models \Box$ $\mathbf{EL_1} \geq 0 \wedge \mathbf{EL_2} \geq 0$.

- For any $(C_1,C_2) \geq (2,1)$, Player 1 has a winning strategy.
- Player 1 needs memory ! How much ?

Complexity of multi-dim. energy games [CDHR10]

Player 2 - Memoryless Strategies

MEGs equivalent to “zero-games” played on **vector addition systems extended with states** (VASS a.k.a. Petri nets) when the initial marking is (ω,ω,…,ω).

For such games, [BJK10] establishes that **memoryless** strategies are sufficient for Player 2.

Lemma[BJK10]. Memoryless strategies are sufficient for Player 2 to win in “zero-games” played on VASS.

Theorem[BJK10]. Zero-game played on vector addition systems extended with states can be solved in **Pspace** when weights are in {-1,0,1} (i.e. in **ExpSpace** for general weights).

Corollary. The unknown initial credit problem in MEGs is in **ExpSpace**.

Corollary. Player 2 plays optimally in MEGs with **memoryless strategies**.

Complexity of MEGs

- Memoryless strategies are sufficient for Player 2 to win a MEG G.
- Let $\lambda_2 \in \Sigma_{2,\mathbf{m}}$, $G(\lambda_2)$ is a multi-weighted graph.
- λ_2 is losing **iff** $G(\lambda_2)$ contains a reachable cycle (not necessarily simple) with positive effect on all dimensions.

Theorem[Kosaraju,Sullivan88]. Given a multi-weighted graph G, it is decidable in deterministic polynomial time if G contains a state s which is reachable from itself with a (not necessarily simple) path with zero effect on all dimensions.

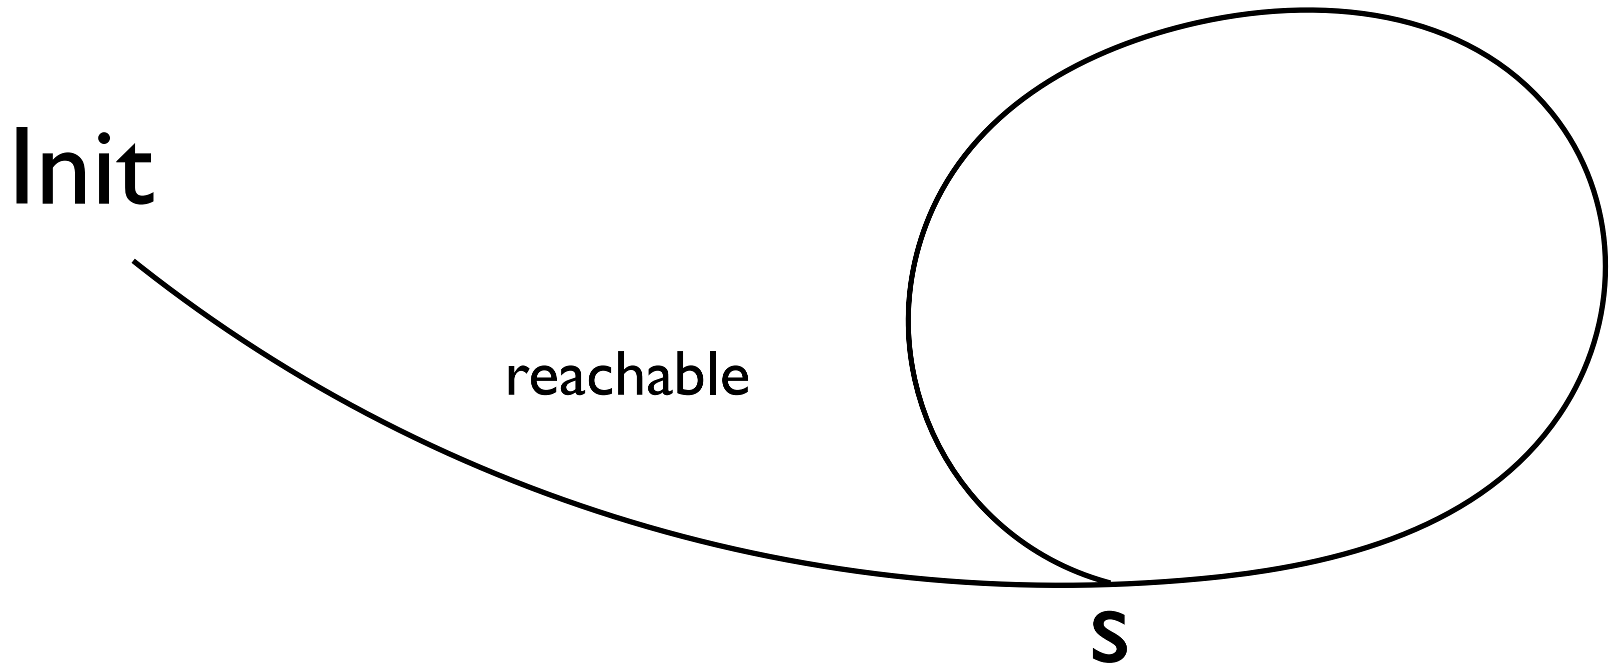

$G(\lambda_2)$+self loops

positive effect
=zero effect
if decreasing self loops added

Complexity of MEGs

- Memoryless strategies are sufficient for Player 2 to win a MEG G.
- Let $\lambda_2 \in \Sigma_{2,\mathbf{m}}$, $G(\lambda_2)$ is a multi-weighted graph.
- λ_2 is losing **iff** $G(\lambda_2)$ contains a reachable cycle (not necessarily simple) with positive effect on all dimensions.

☞ **Lemma.** The unknown initial credit problem in MEGs belongs to coNP.

Complexity of MEGs

Lemma. The unknown initial credit problem in MEGs is coNP-Hard.

Proof. We show that deciding whether **Player 1 has a winning strategy** is as hard as deciding if a 3CNF formula is **unsatisfiable**.

Let ψ be a 3CNF formula with clauses $C_1, C_2, \ldots, C_k$ over variables $\{x_1, x_2, \ldots, x_n\}$. We construct from ψ the following game structure with weight in $\mathbb{Z}^{2n}$:

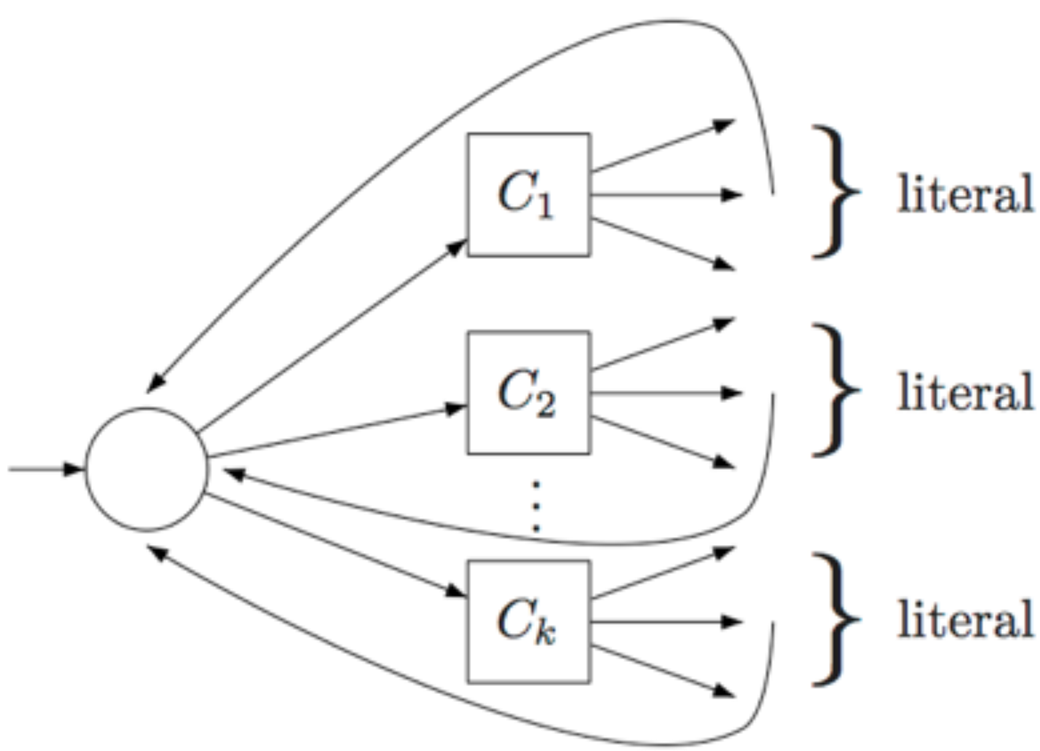

Complexity of MEGs

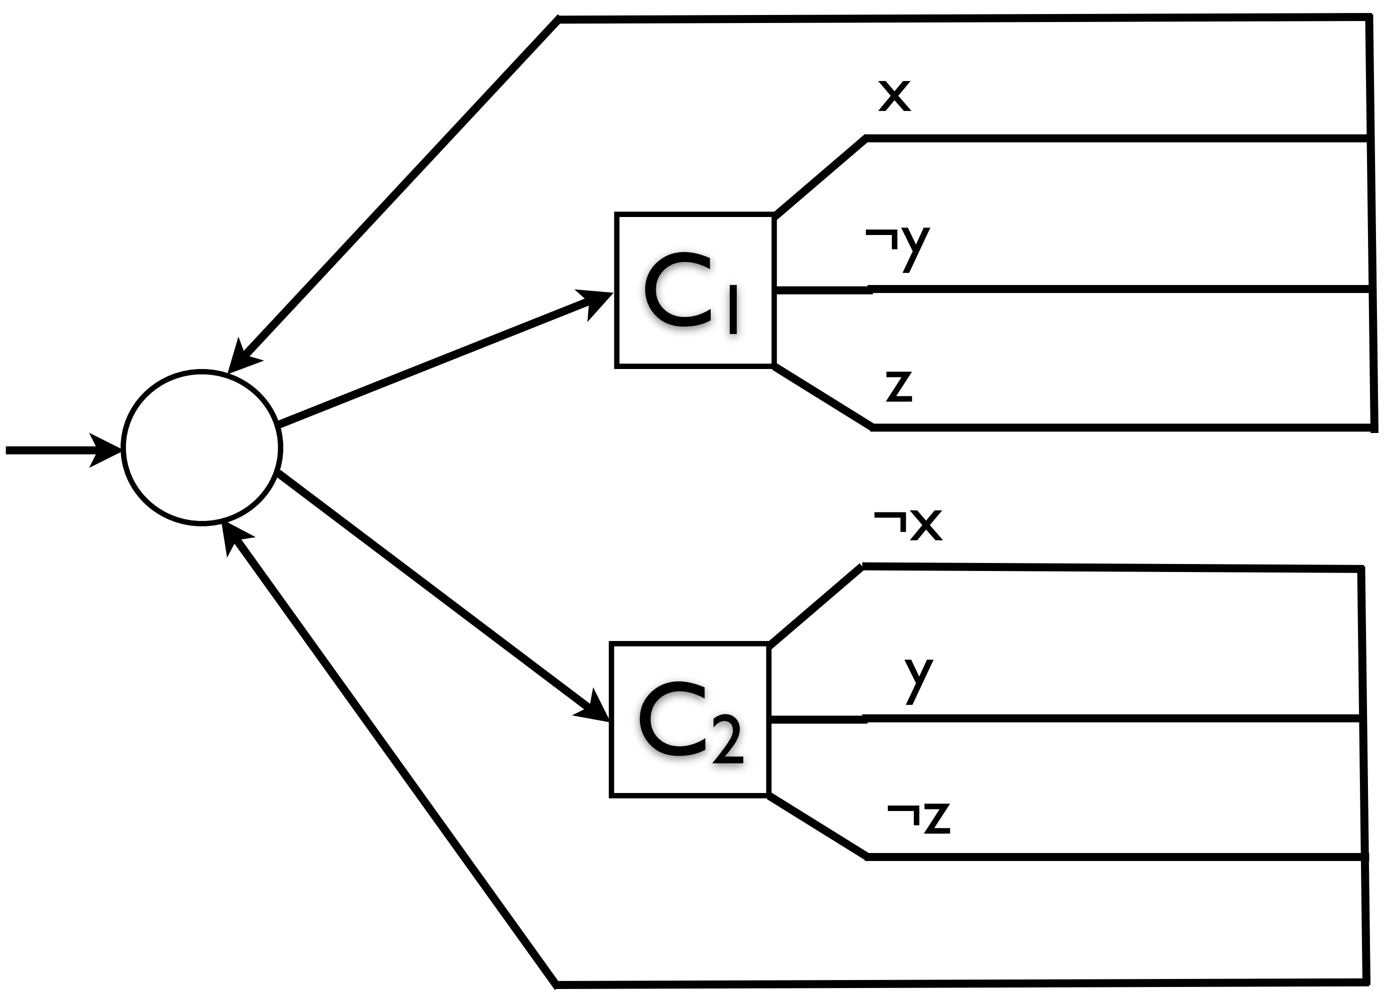

Ex: $\Phi = (x \vee \neg y \vee z) \wedge (\neg x \vee y \vee \neg z)$

Complexity of MEGs

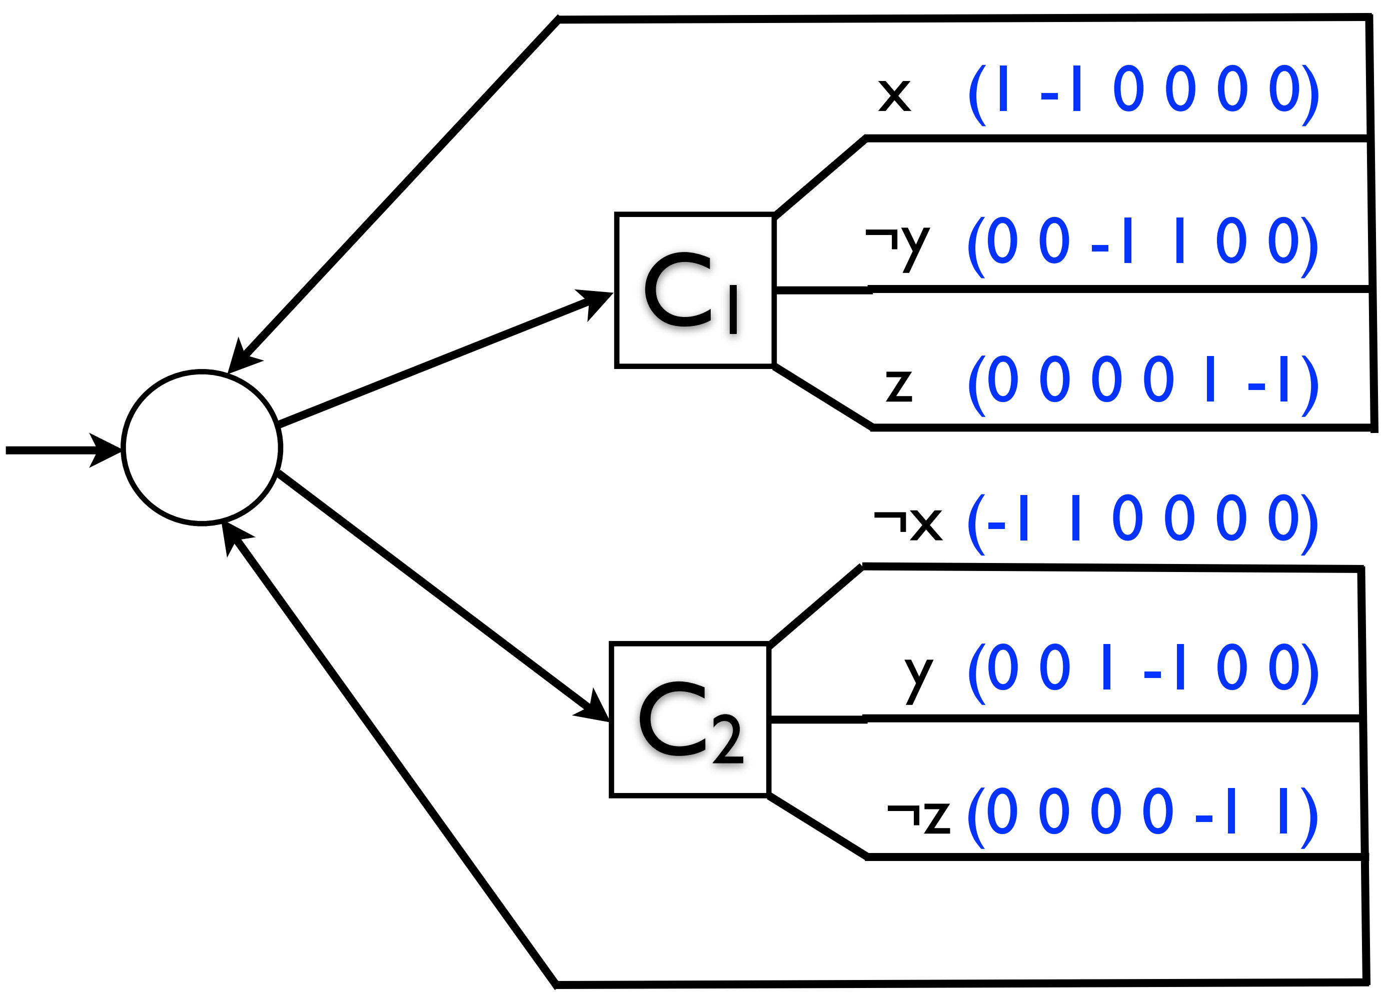

Ex: $\Phi = (x \vee \neg y \vee z) \wedge (\neg x \vee y \vee \neg z)$

We define the weight labelling as follows:

-every edge is labeled by $\{0\}^{2n}$ with the exception of edges going from literals back to initial state.
-for a literal y and an edge back to the initial state, the weight vector contains:

- 1 in the dimension of y
- -1 in the dimension of the complement of y
- 0 otherwise

Complexity of MEGs

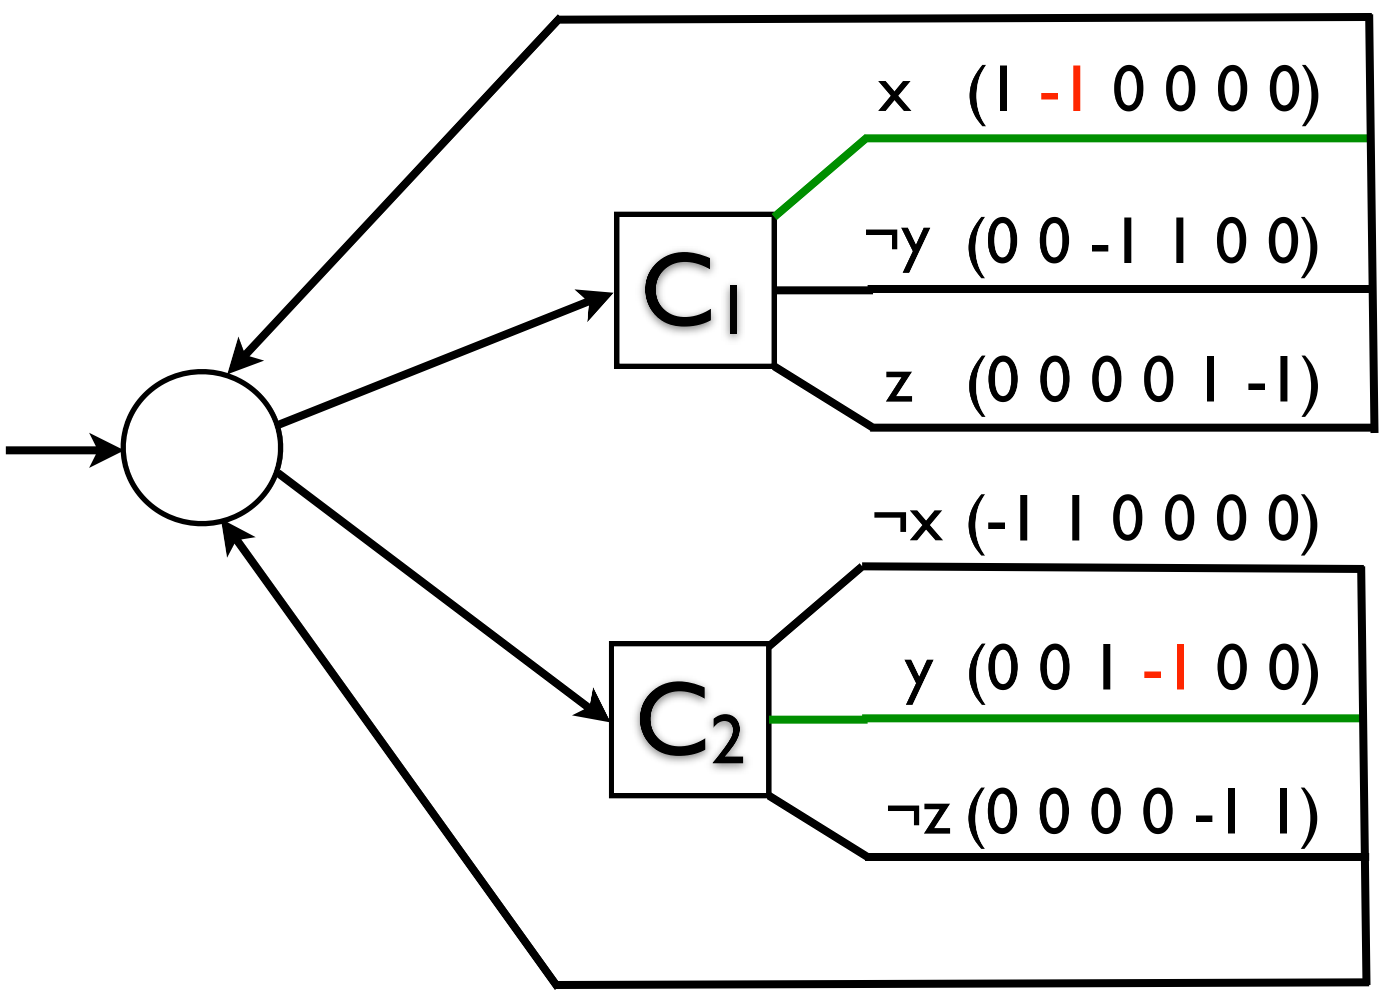

Ex: $\Phi = (x \vee \neg y \vee z) \wedge (\neg x \vee y \vee \neg z)$

$v(x)=1, v(y)=1, v(z)=0$

$v \models \Phi$

$\lambda_2(C_1)=x$

$\lambda_2(C_2)=y$

Φ is satisfiable implies Player 2 wins

Let v be s.t. $v \models \Phi$. We construct λ_2 as follows: in each clause C_i, λ_2 chooses l_{ij} s. t. $v \models l_{ij}$. Now, take any λ_1 and consider the play consistent with λ_1 and λ_2. There must exist C_i that appears ∞-often along this play: the dimension that correspond to $\neg l_{ij}$ is decreased ∞-often without ever being increased ! There is no initial credit that can help Player 1 !

Complexity of MEGs

Φ is unsatisfiable implies Player 1 is winning
or equivalently **Φ is unsatisfiable implies Player 2 is not winning**

As Φ is unsatisfiable, when Player 2 chooses one literal per clause (we know that he can play optimally without memory), it has to choose two literals that are **complementary**. Let assume that the choice of Player 2 are complementary for clauses C_i and C_j. In that case, the winning strategy for Player 1 is to alternate between C_i and C_j. This strategy is winning for a initial credit of 1 in all dimension.

$$\Phi = (\neg y) \wedge (x \vee y) \wedge (\neg x)$$

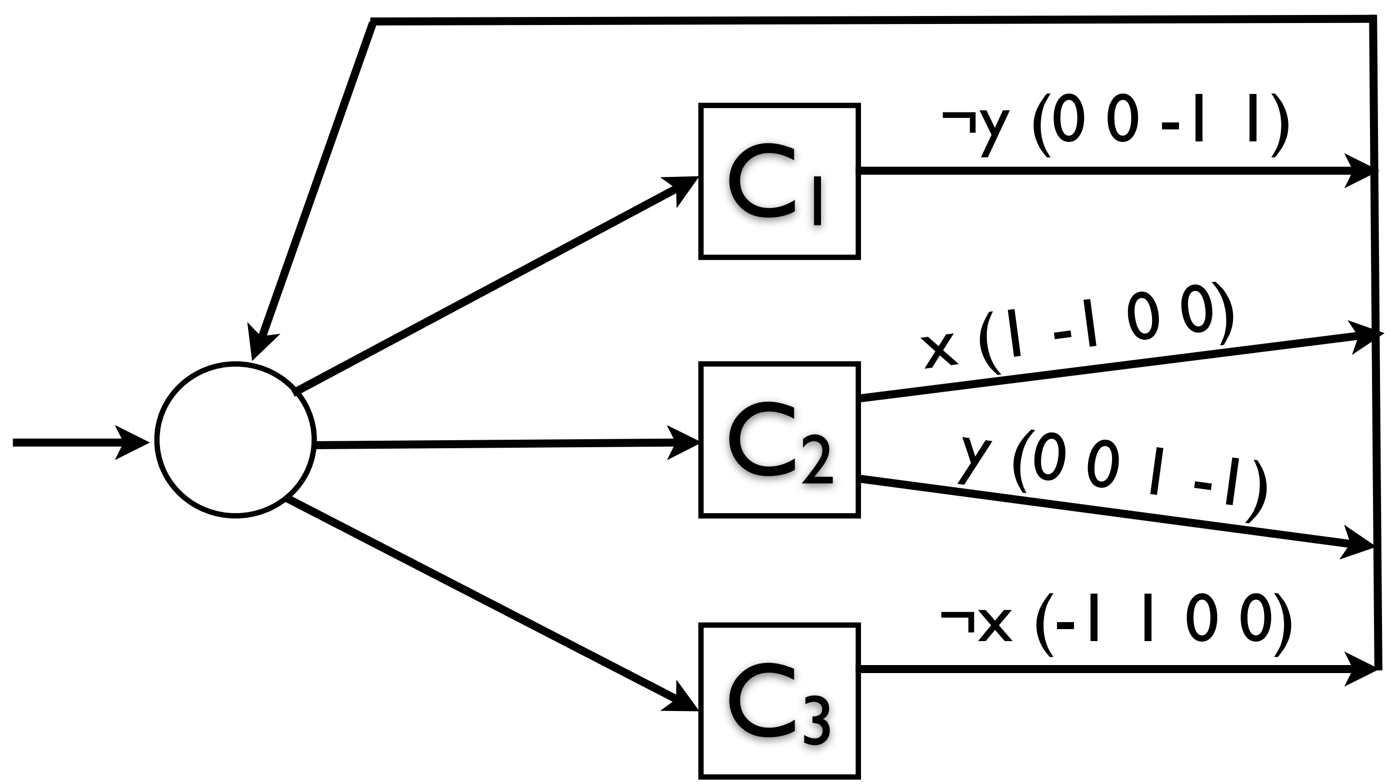

Complexity of MEGs

Φ is unsatisfiable implies Player 1 is winning

or equivalently **Φ is unsatisfiable implies Player 2 is not winning**

As Φ is unsatisfiable, when Player 2 chooses one literal per clause (we know that he can play optimally without memory), it has to choose two literals that are **complementary**. Let assume that the choice of Player 2 are complementary for clauses C_i and C_j. In that case, the winning strategy for Player 1 is to alternate between C_i and C_j. This strategy is winning for a initial credit of 1 in all dimension.

$$\Phi = (\neg y) \wedge (x \vee y) \wedge (\neg x)$$

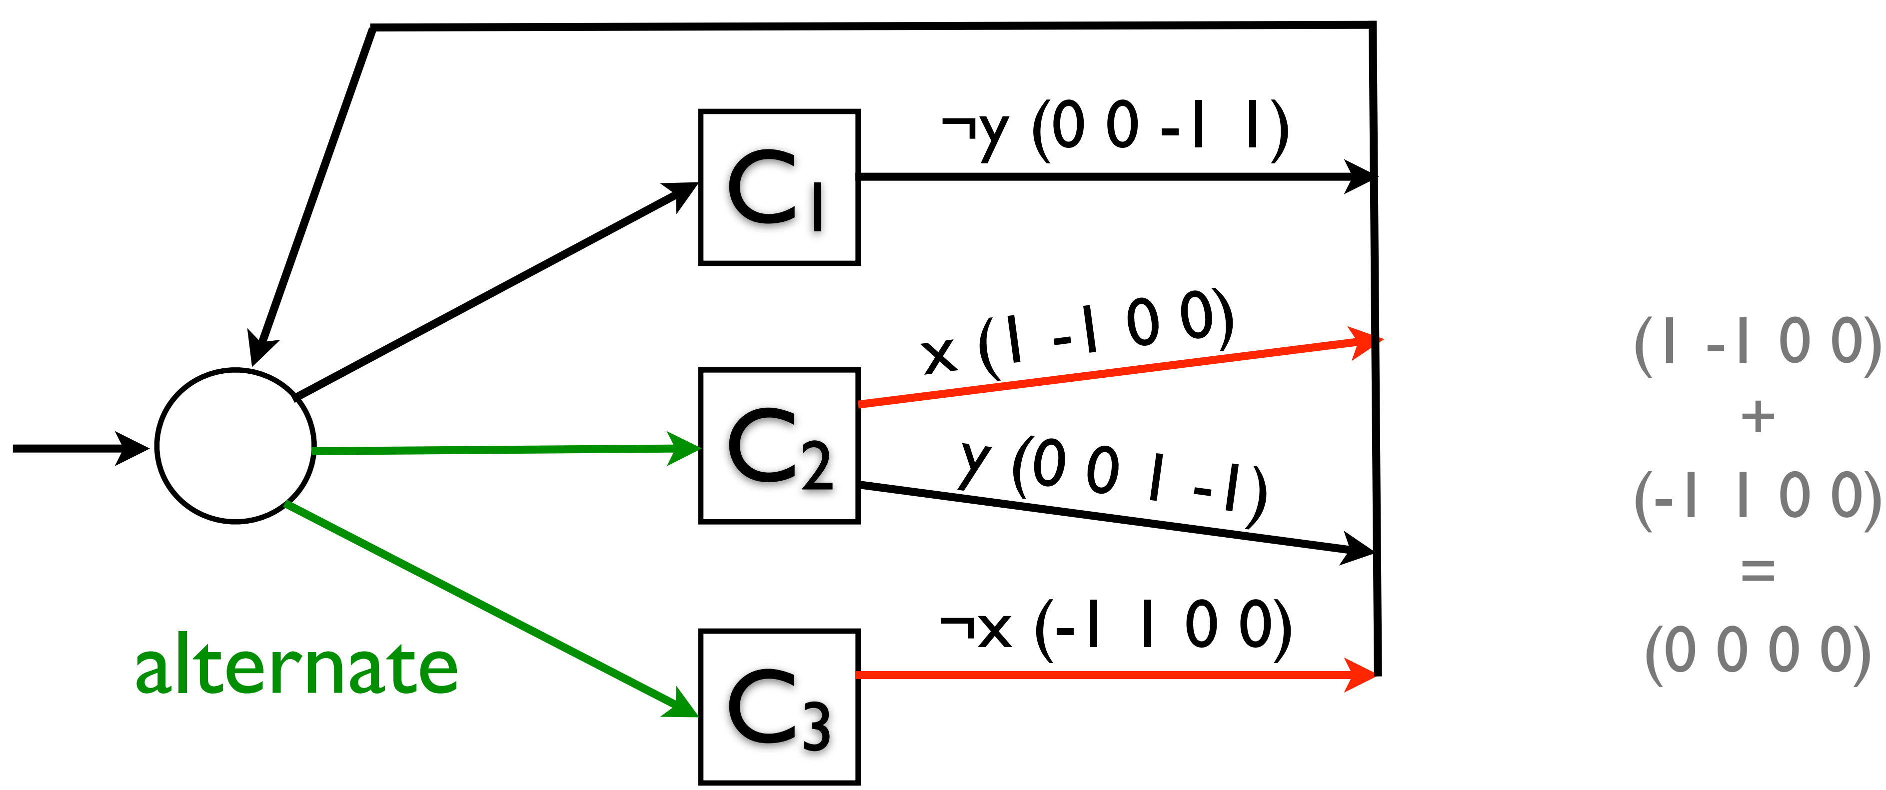

Complexity of MEGs

Φ is unsatisfiable implies Player 1 is winning
or equivalently **Φ is unsatisfiable implies Player 2 is not winning**

As Φ is unsatisfiable, when Player 2 chooses one literal per clause (we know that he can play optimally without memory), it has to choose two literals that are **complementary**. Let assume that the choice of Player 2 are complementary for clauses C_i and C_j. In that case, the winning strategy for Player 1 is to alternate between C_i and C_j. This strategy is winning for a initial credit of 1 in all dimension.

$$\Phi = (\neg y) \wedge (x \vee y) \wedge (\neg x)$$

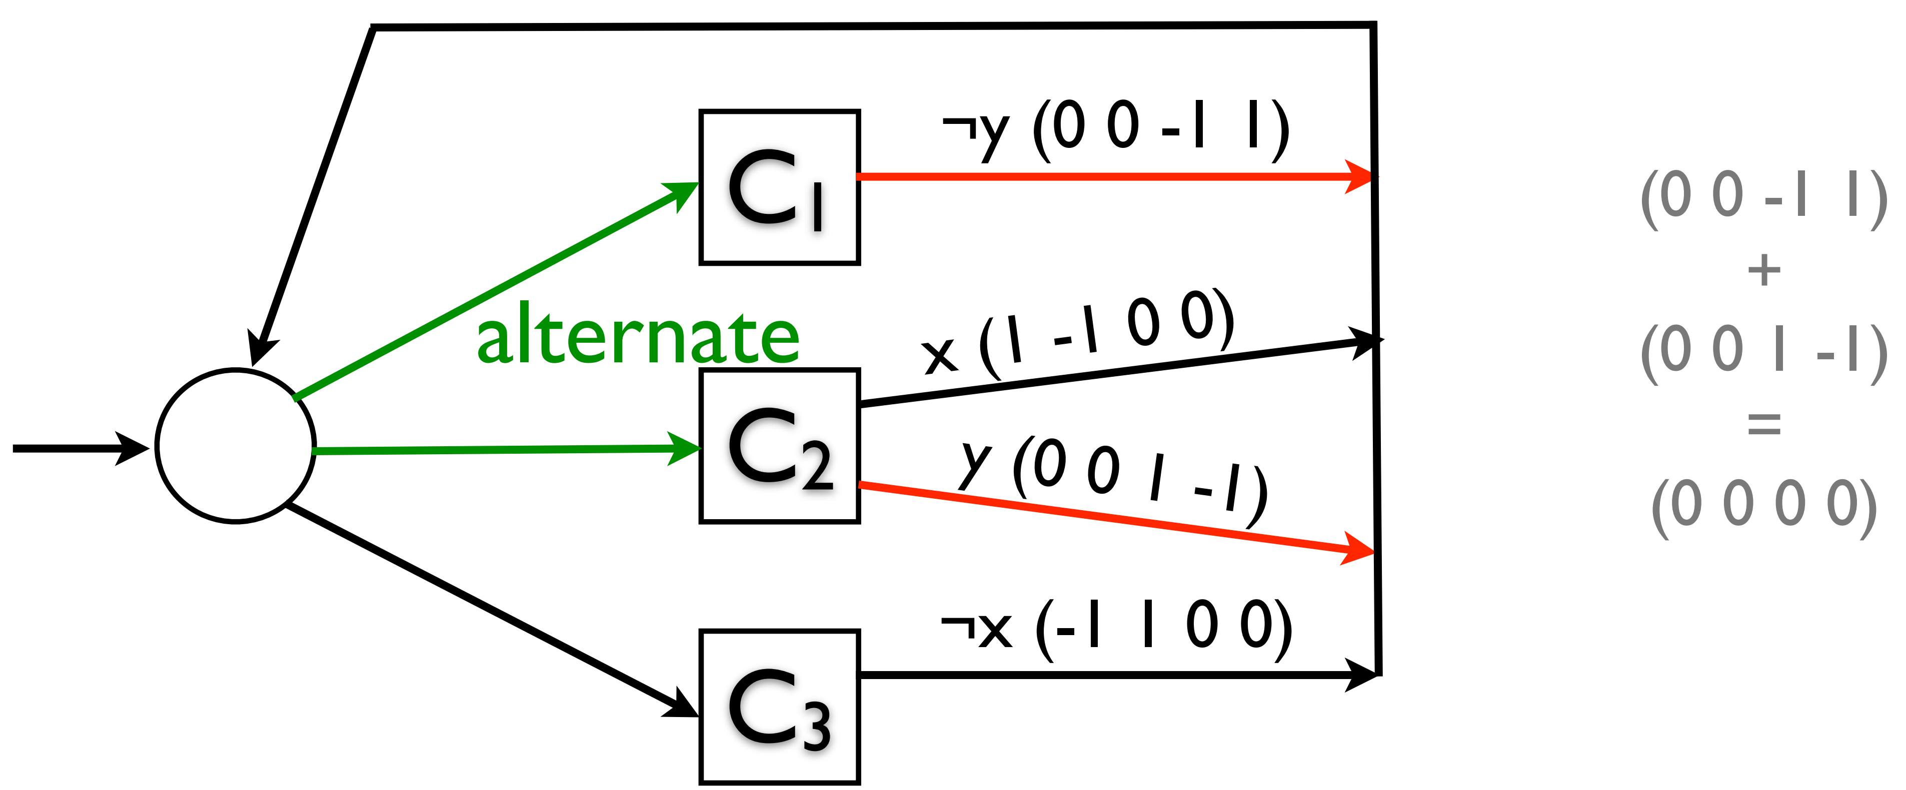

Complexity of MEGs

Theorem. The unknown initial credit problem in MEGs is **coNP-C**.

Exponential Memory is Sufficient for Player 1 in Multi-dim. Energy Games [CRR12]

Player 1 - Finite Memory Strategies

Lemma. Finite memory strategies are sufficient for Player 1 to win in MEGs.

Player 1 - Finite Memory Strategies

Lemma. Finite memory strategies are sufficient for Player 1 to win in MEGs.

Proof. First, $(\mathbb{N}^k,\leq)$ is a well-quasi ordered set, i.e.:

1. $\leq$ is a partial order (so a pre-order)
2. for all infinite sequences of elements $m_0\ m_1\ m_2 \dots m_n \dots$ in $(\mathbb{N}^k)^\omega$,

   there exists $i<j$ such that $m_i \leq m_j$

Player 1 - Finite Memory Strategies

Lemma. Finite memory strategies are sufficient for Player 1 to win in MEGs.

Proof. First, $(\mathbb{N}^k, \leq)$ is a well-quasi ordered set.

Let λ_I be winning

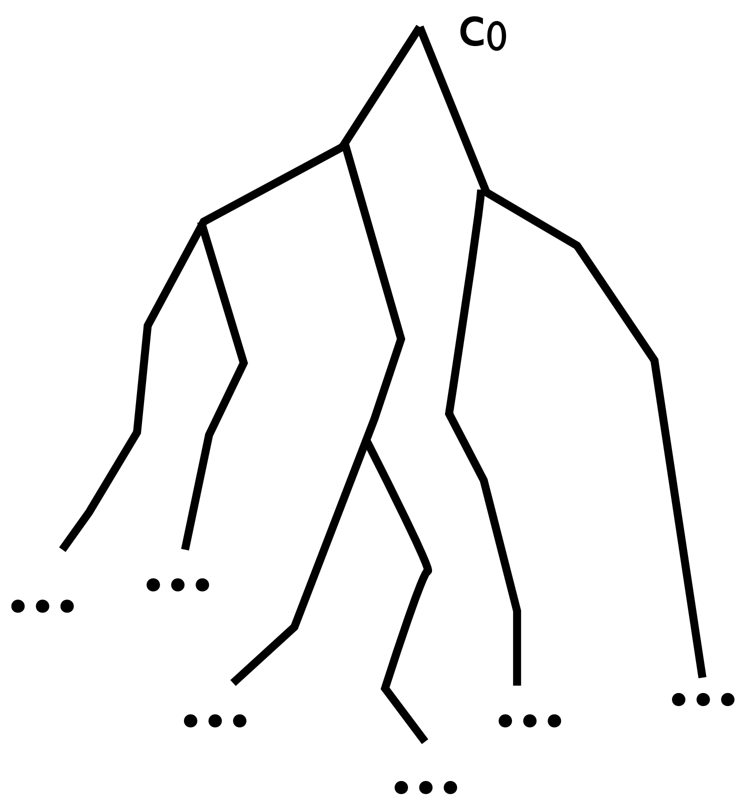

Player 1 - Finite Memory Strategies

Lemma. Finite memory strategies are sufficient for Player 1 to win in MEGs.

Proof. First, $(\mathbb{N}^k, \leq)$ is a well-quasi ordered set.

Let λ_I be winning

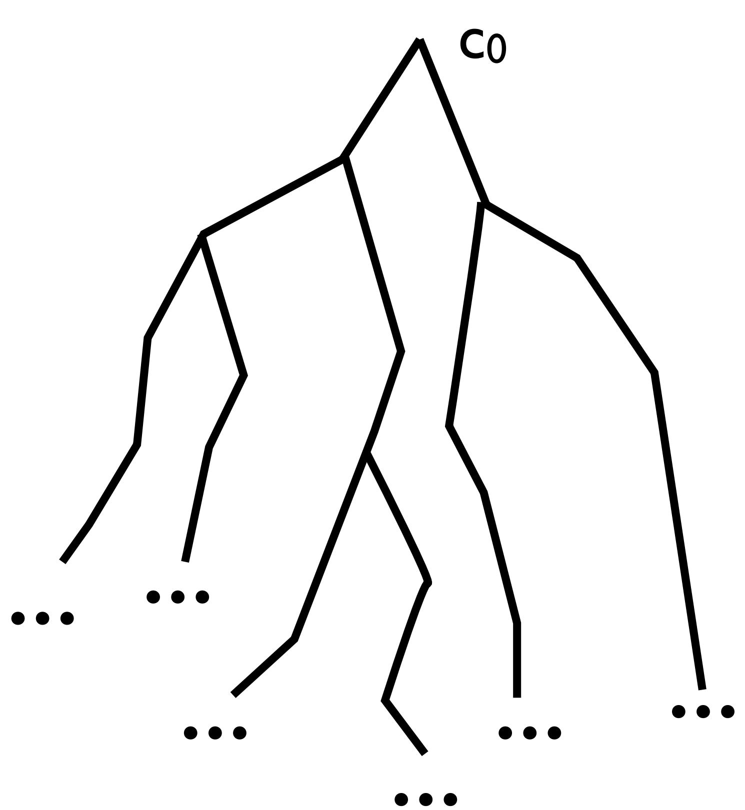

On each branch

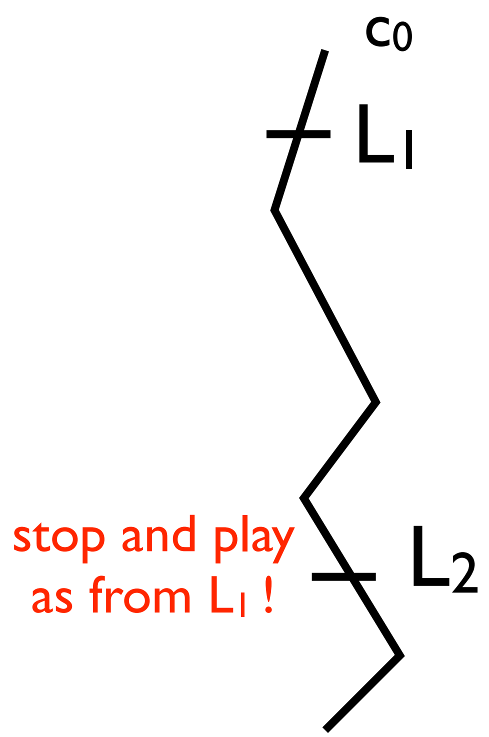

With $L_I \leq L_2$

Player 1 - Finite Memory Strategies

Lemma. Finite memory strategies are sufficient for Player 1 to win in MEGs.

Proof. First, $(\mathbb{N}^k,\leq)$ is a well-quasi ordered set.

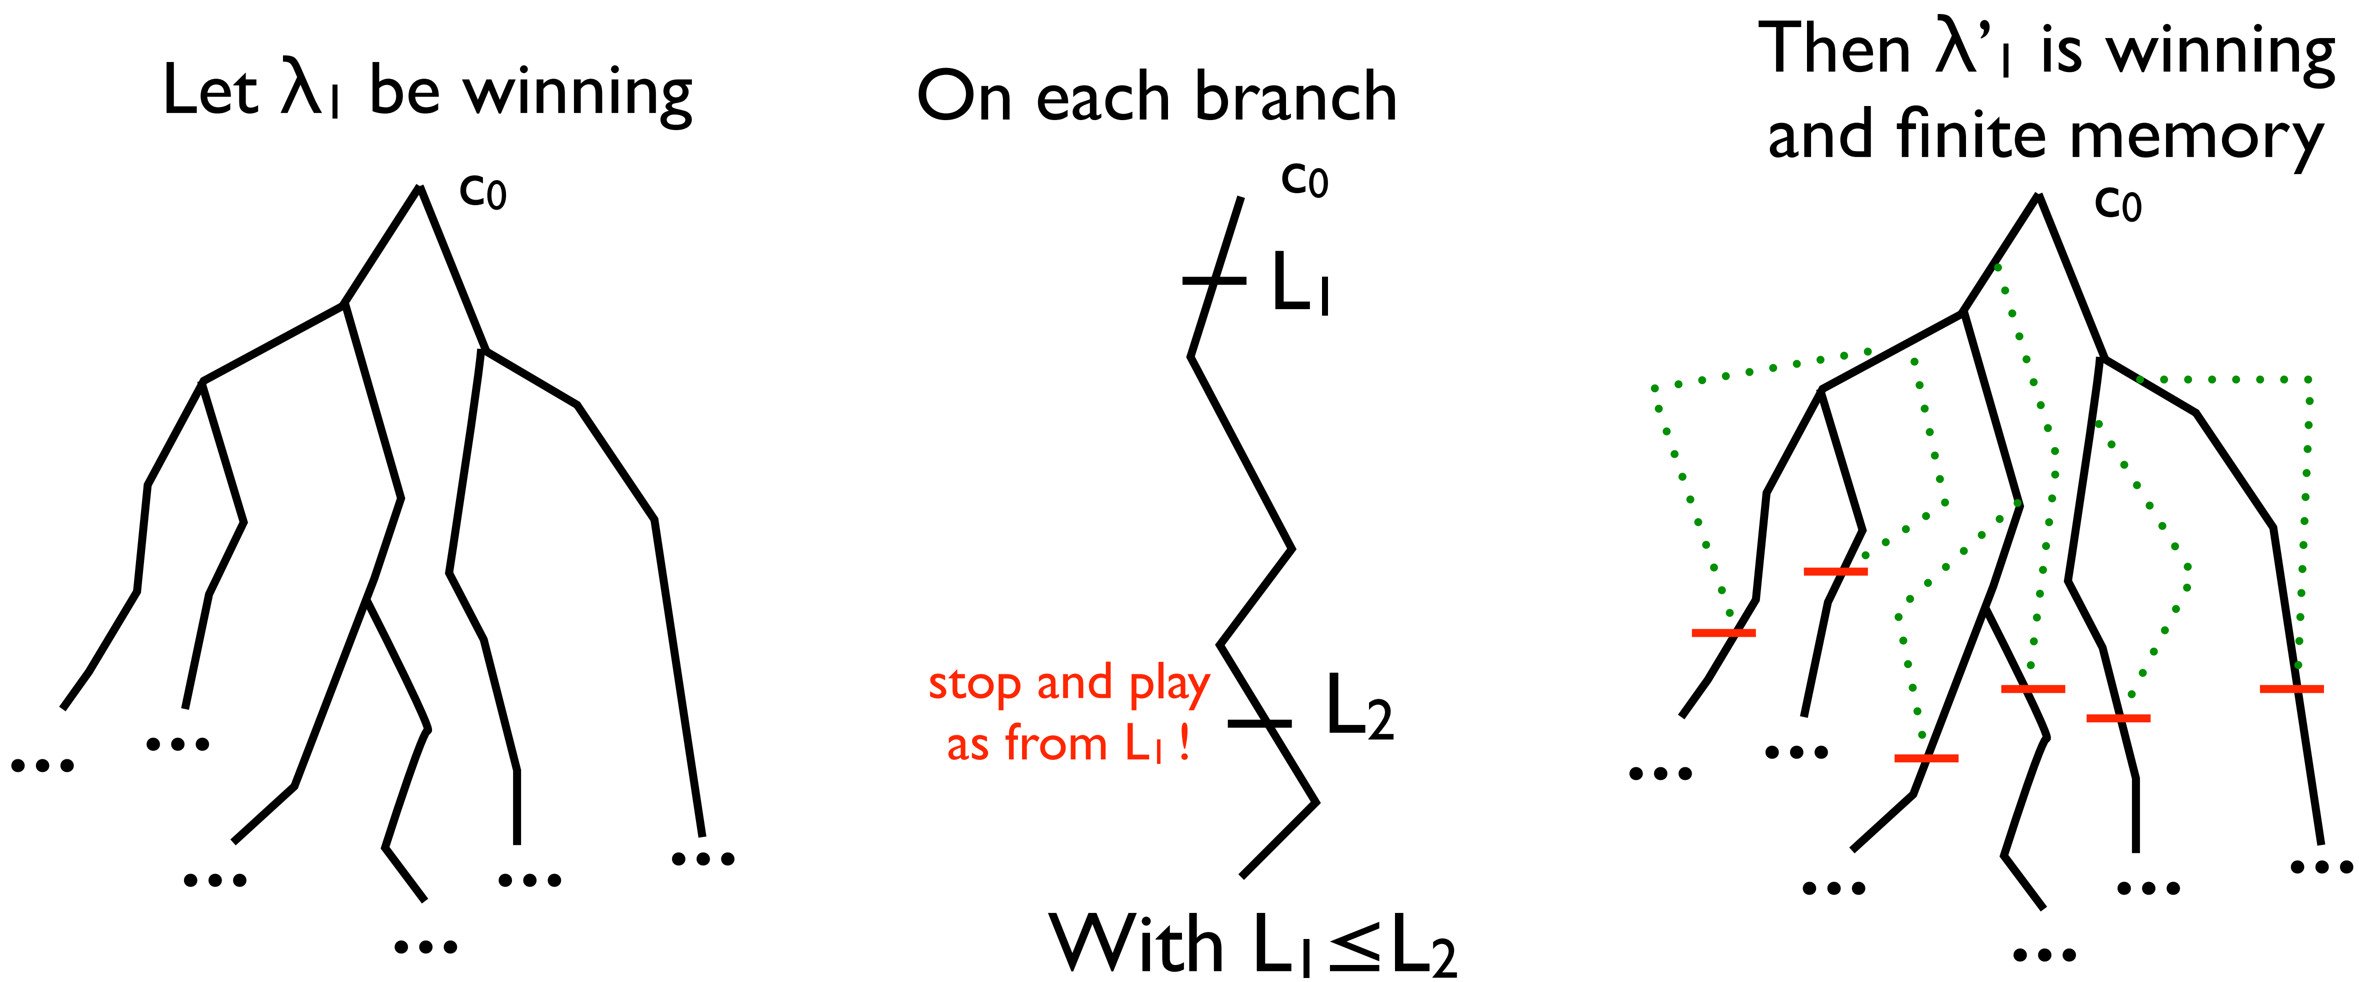

wqo+Koenig's lemma

Player 1 - Finite Memory Strategies

Lemma. Finite memory strategies are sufficient for Player 1 to win in MEGs.

Proof. First, $(\mathbb{N}^k, \leq)$ is a well-quasi ordered set.

Then λ'_1 is winning
and finite memory

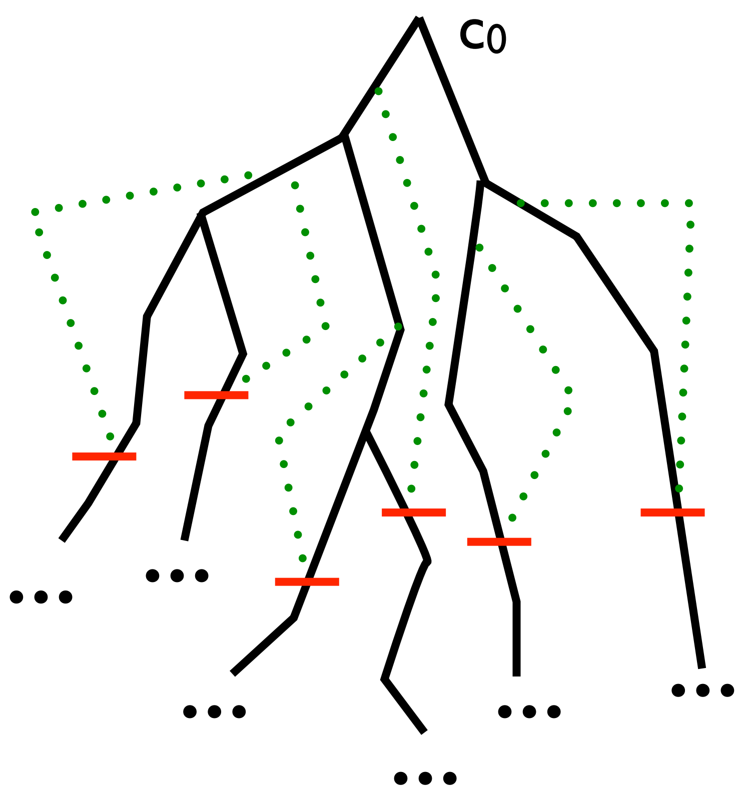

wqo+Koenig's lemma

Finite tree=winning strategy:

① play according to the choices made in tree

② in leaf, go to ancestor with lower or equal energy

Finite memory → Exponential memory

Then λ'ᴵ is winning
and finite memory

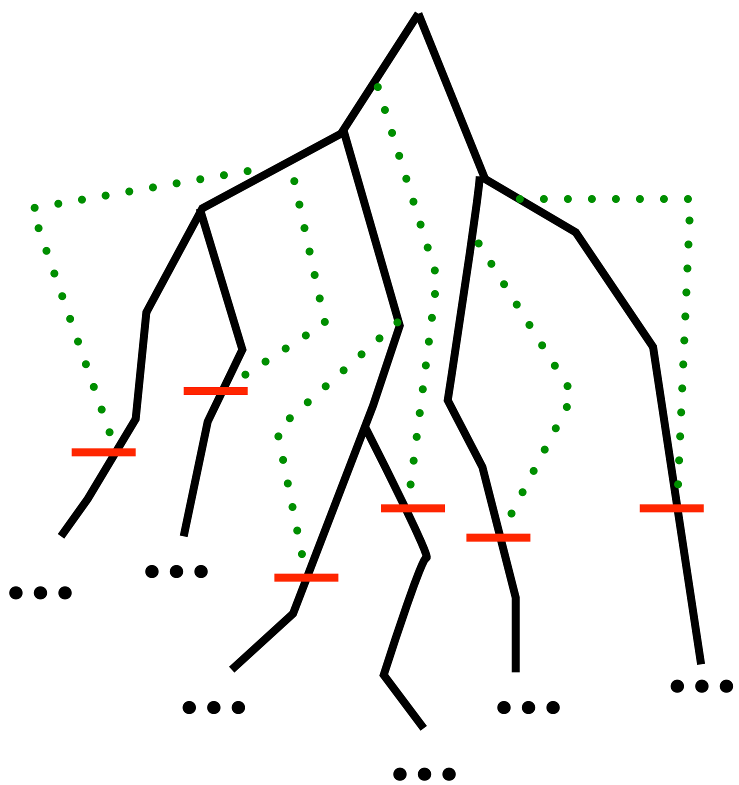

wqo+Koenig's lemma

① Exponential memory is sufficient.

- Use extensions of technics à la Rackoff (Petri nets) - refinements of [BJK10]

② Exponential memory is needed

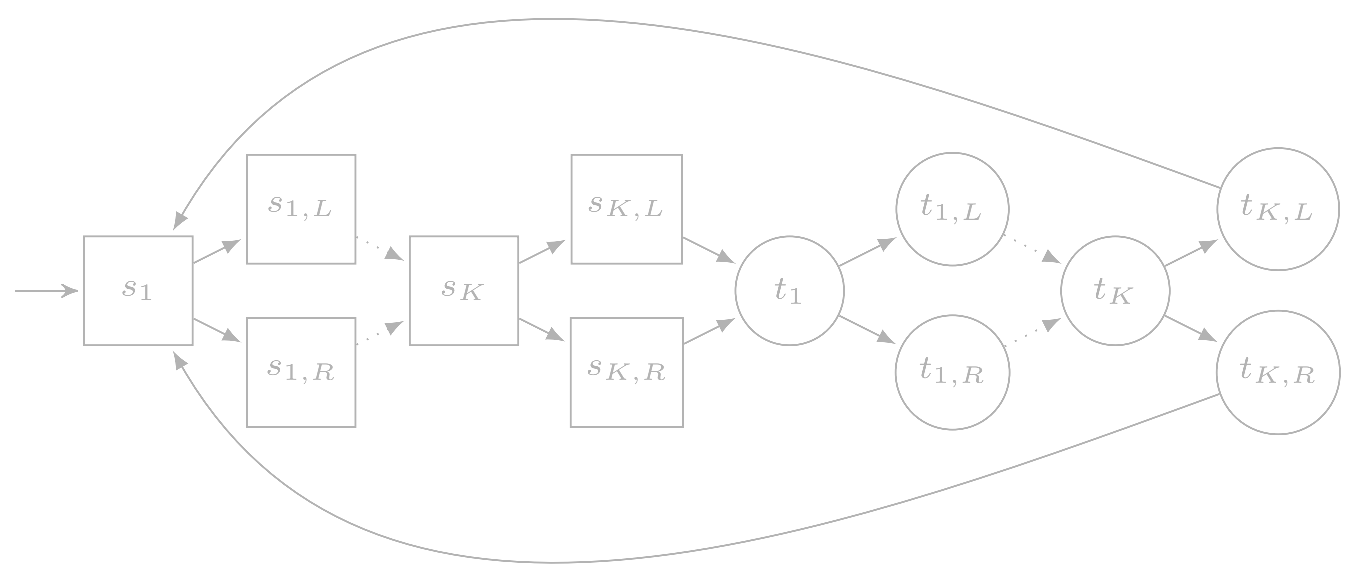

③ Leads to symbolic and incremental algorithms

Finite ... ponential memory

...fficient.

...nics à la Rackoff

...ts of [BJK10]

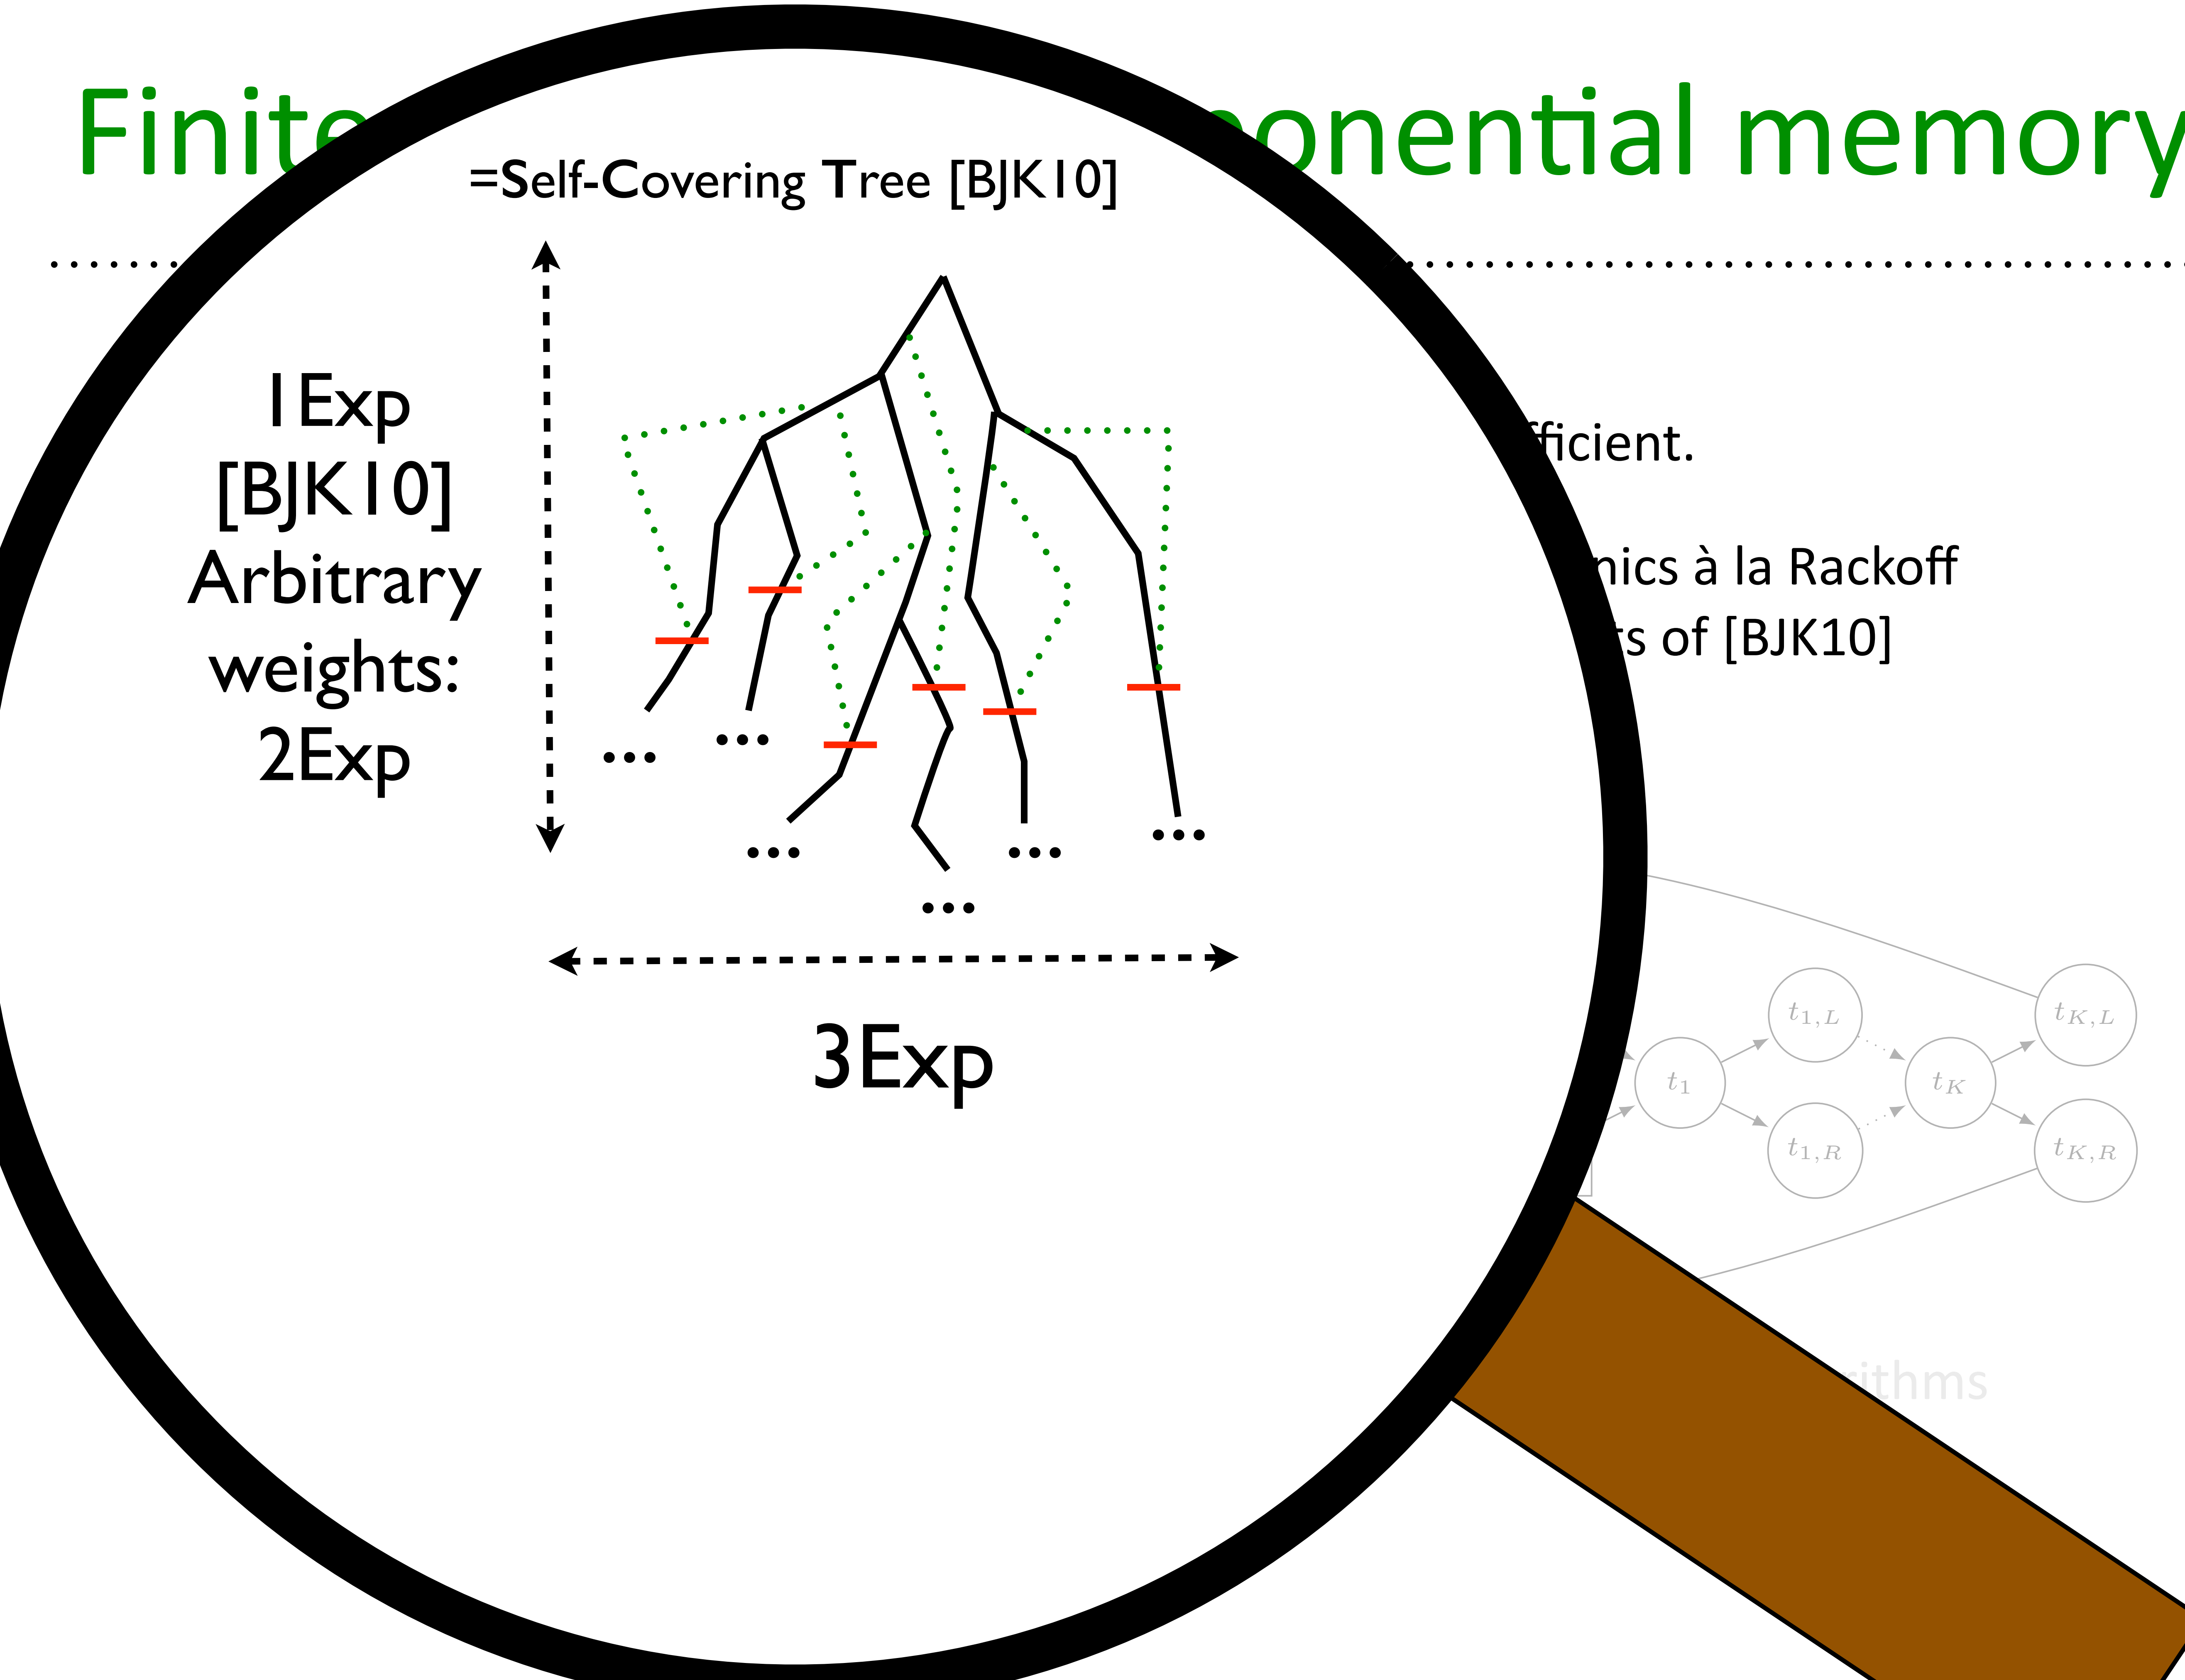

...ithms

Finite … ponential memory

=Self-Covering Tree [BJK10]

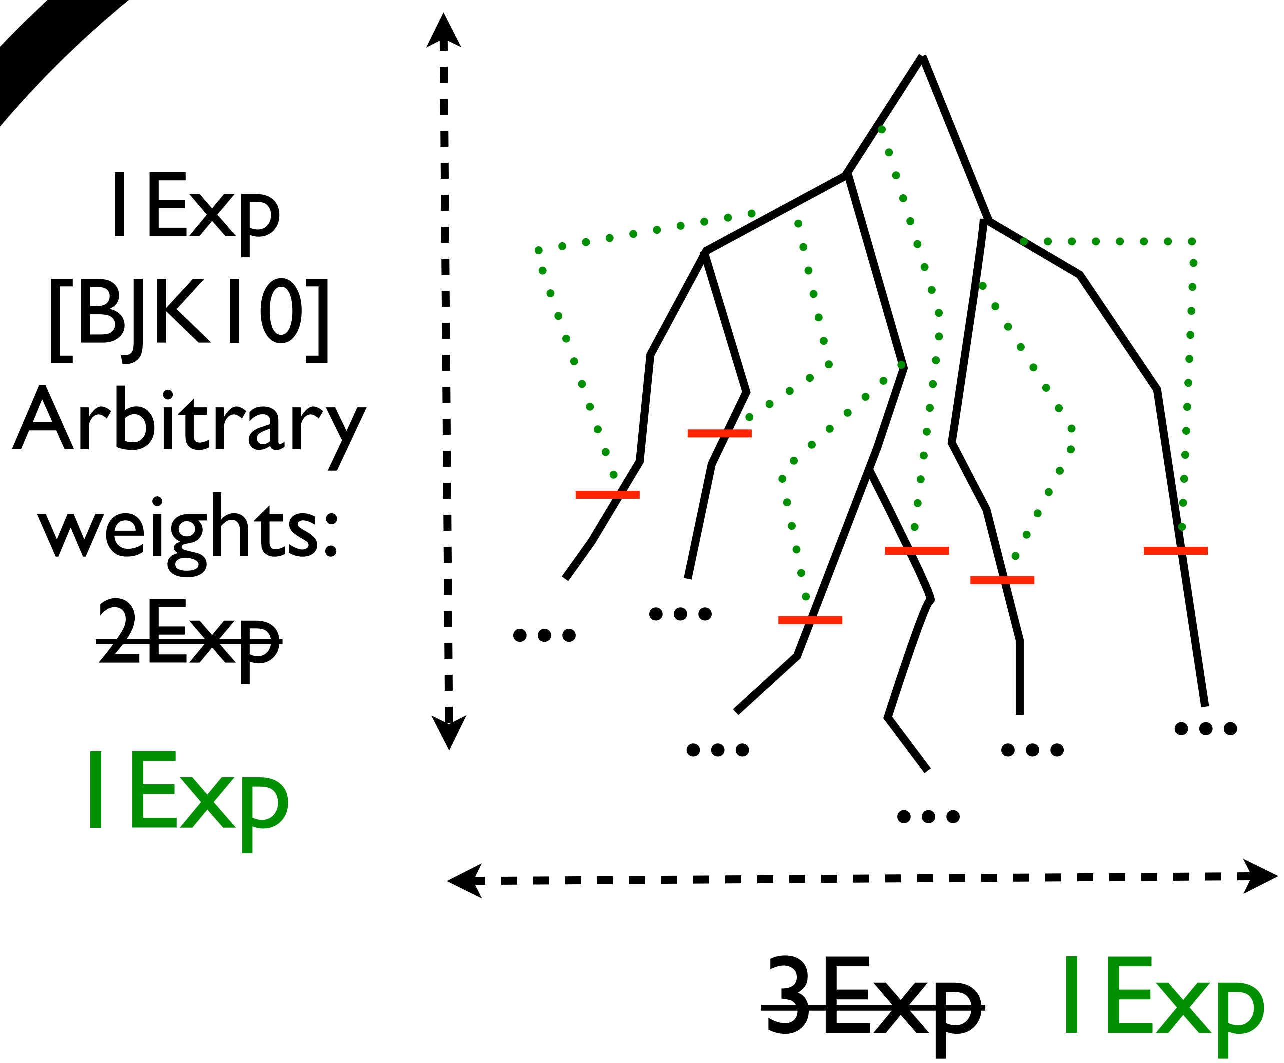

Depth: single exponential - encoding of arbitrary weights into {-1,0,1} does not add choices to the adversary.
Width: only energy level important (DAG).

…fficient.

…mics à la Rackoff
…ts of [BJK10]

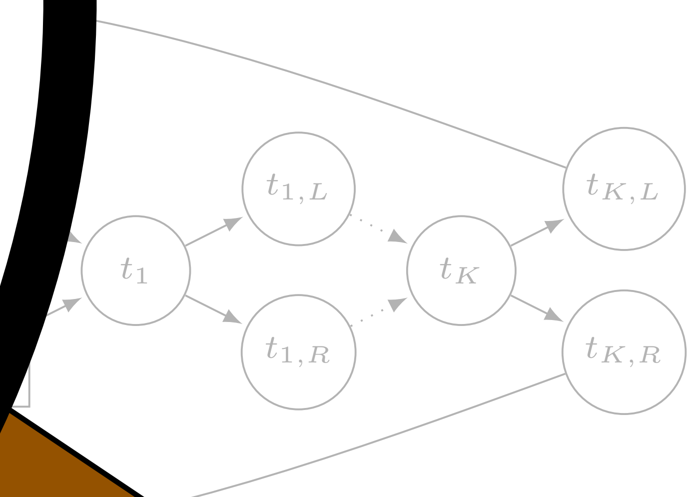

…rithms

Finite … ponential memory

fficient.

nics à la Rackoff

ts of [BJK10]

=Self-Covering Tree [BJK10]

1Exp [BJK10]
Arbitrary weights: ~~2Exp~~ **1Exp**

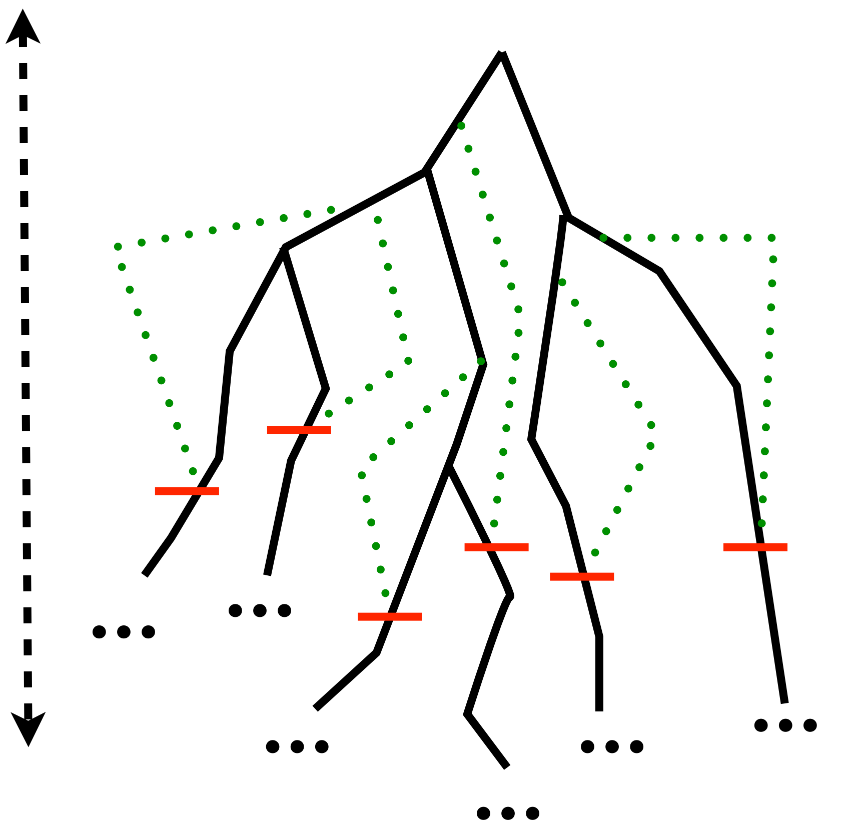

~~3Exp~~ 1Exp

Depth: single exponential - encoding of arbitrary weights into {-1,0,1} does not add choices to the adversary.
Width: only energy level important (DAG).

Works also with parity

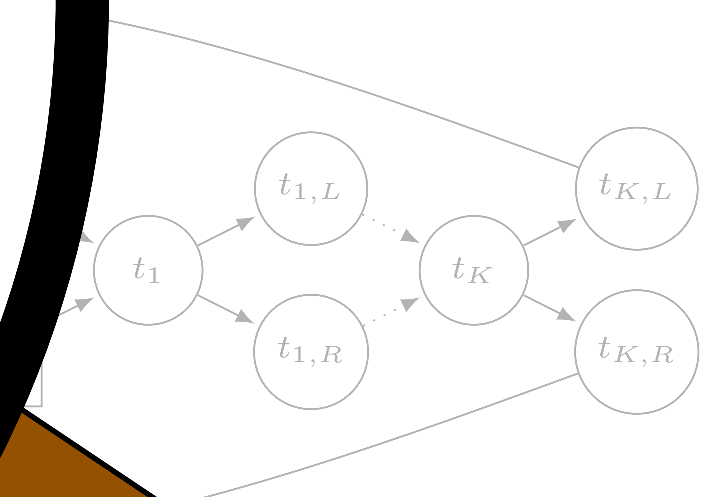

ithms

Finite memory → Exponential memory

Then λ'ı is winning
and finite memory

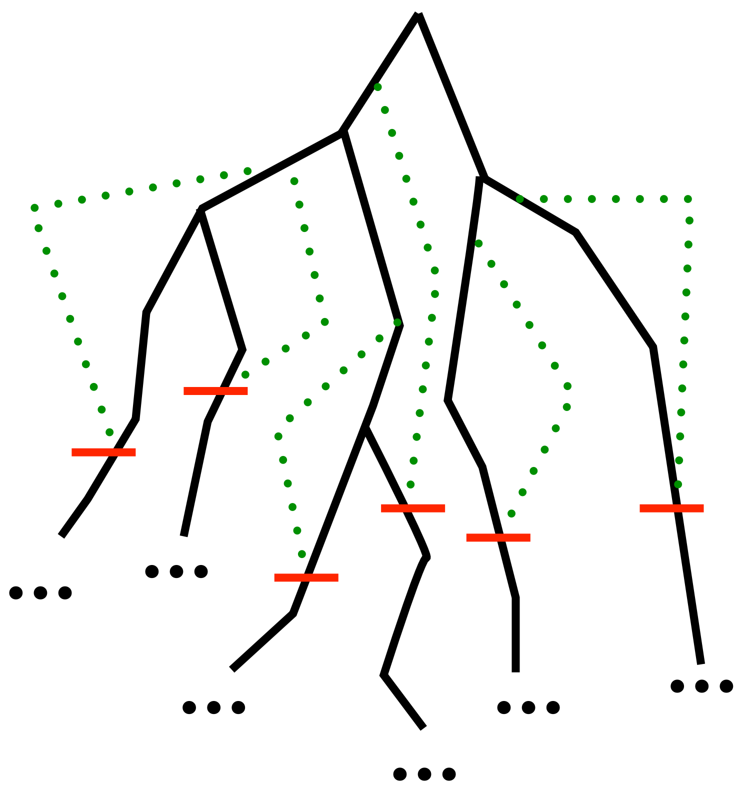

wqo+Koenig's lemma

① Exponential memory is sufficient.

- Use extensions of technics à la Rackoff (Petri nets) - refinements of [BJK10]

② Exponential memory is needed

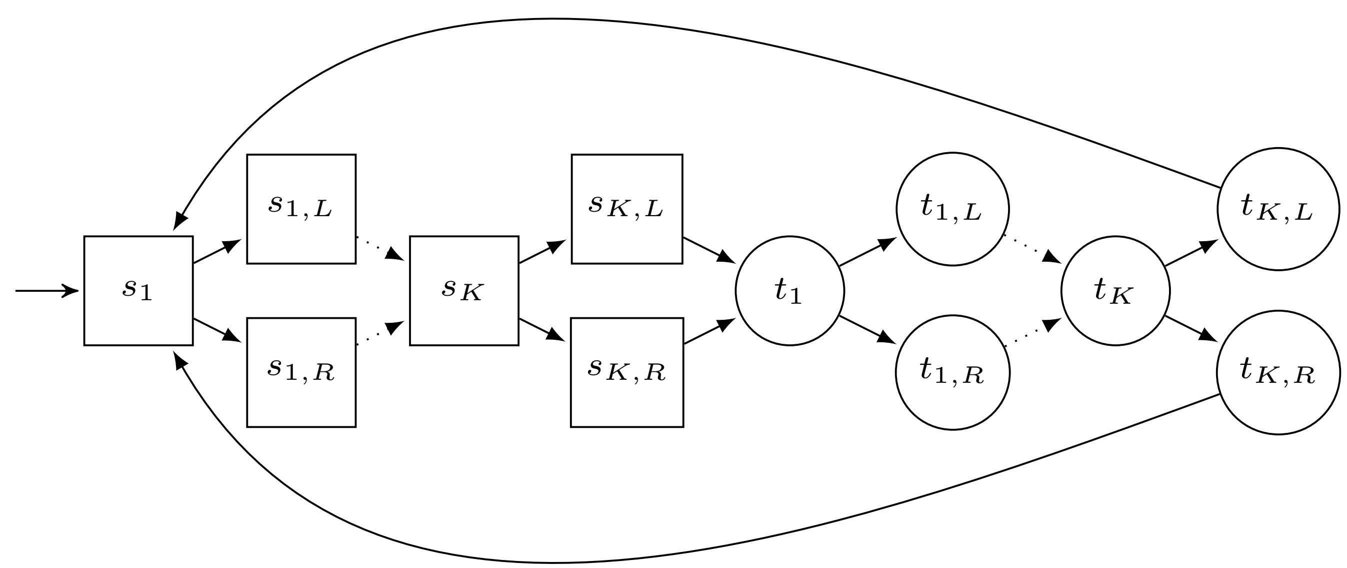

③ Leads to symbolic and incremental algorithms

Finite … onential memory

…fficient.

…nics à la Rackoff

…s of [BJK10]

…rithms

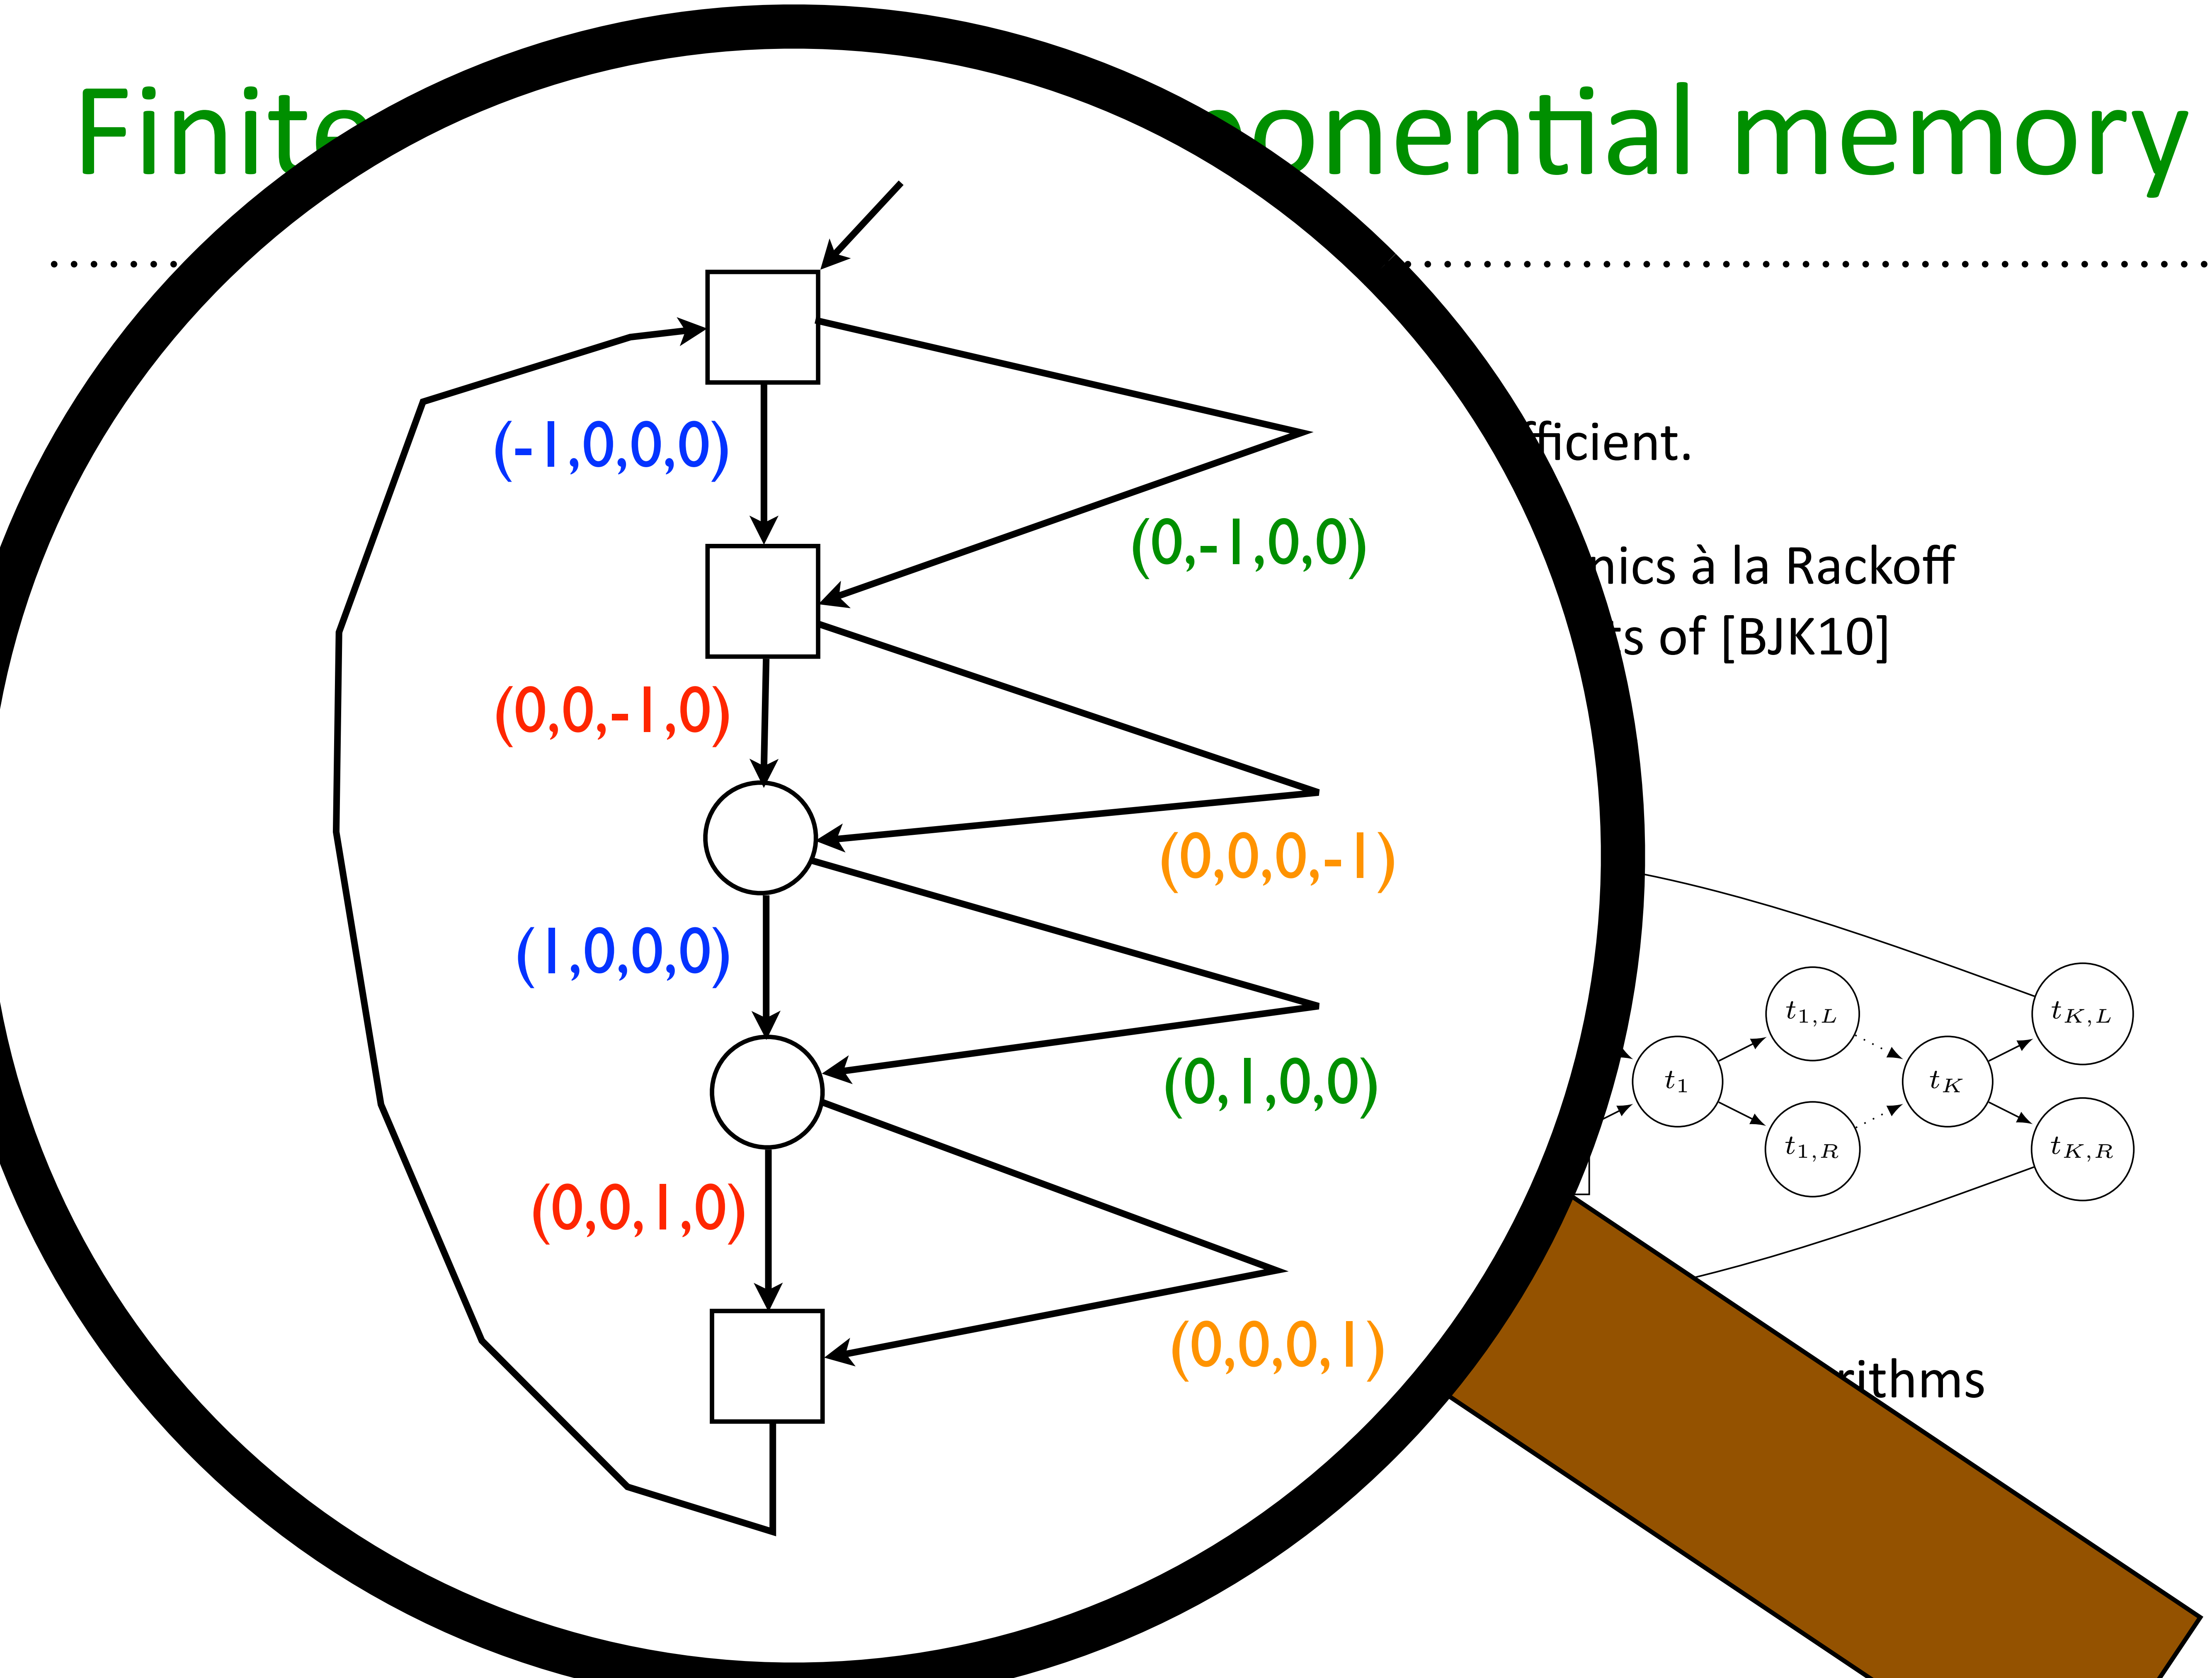

Finite memory → Exponential memory

Then λ'$_I$ is winning and finite memory

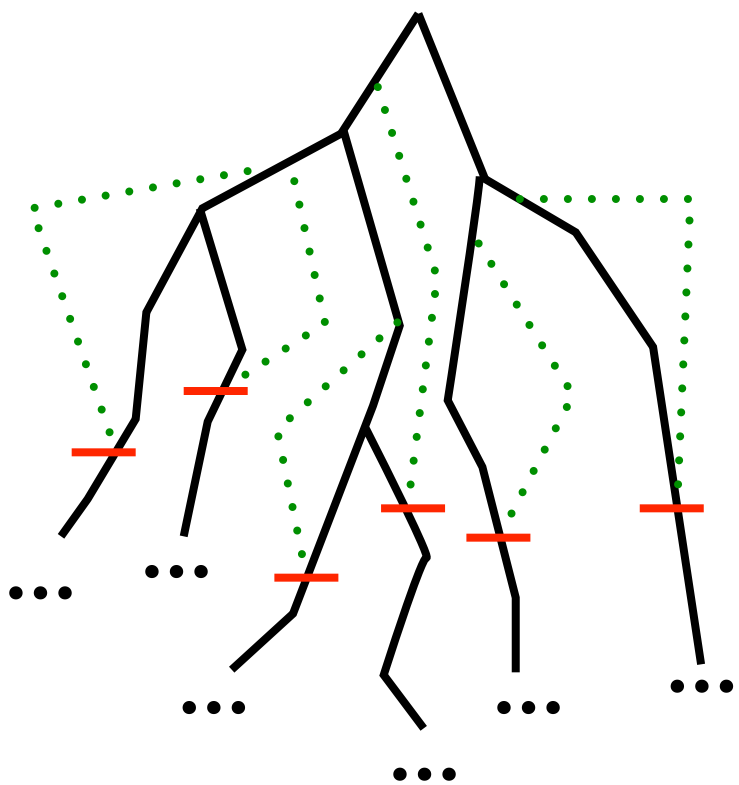

wqo+Koenig's lemma

① Exponential memory is sufficient.

- Use extensions of technics à la Rackoff (Petri nets) - refinements of [BJK10]

② Exponential memory is needed

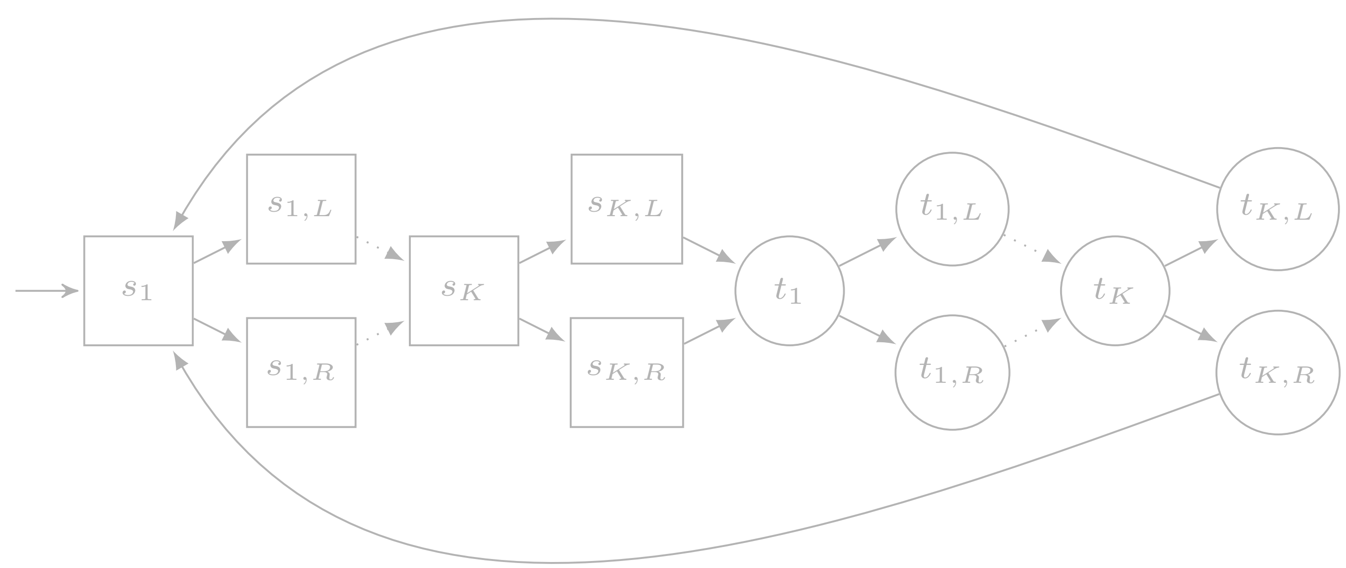

③ Leads to symbolic and incremental algorithms

Finite onential memory

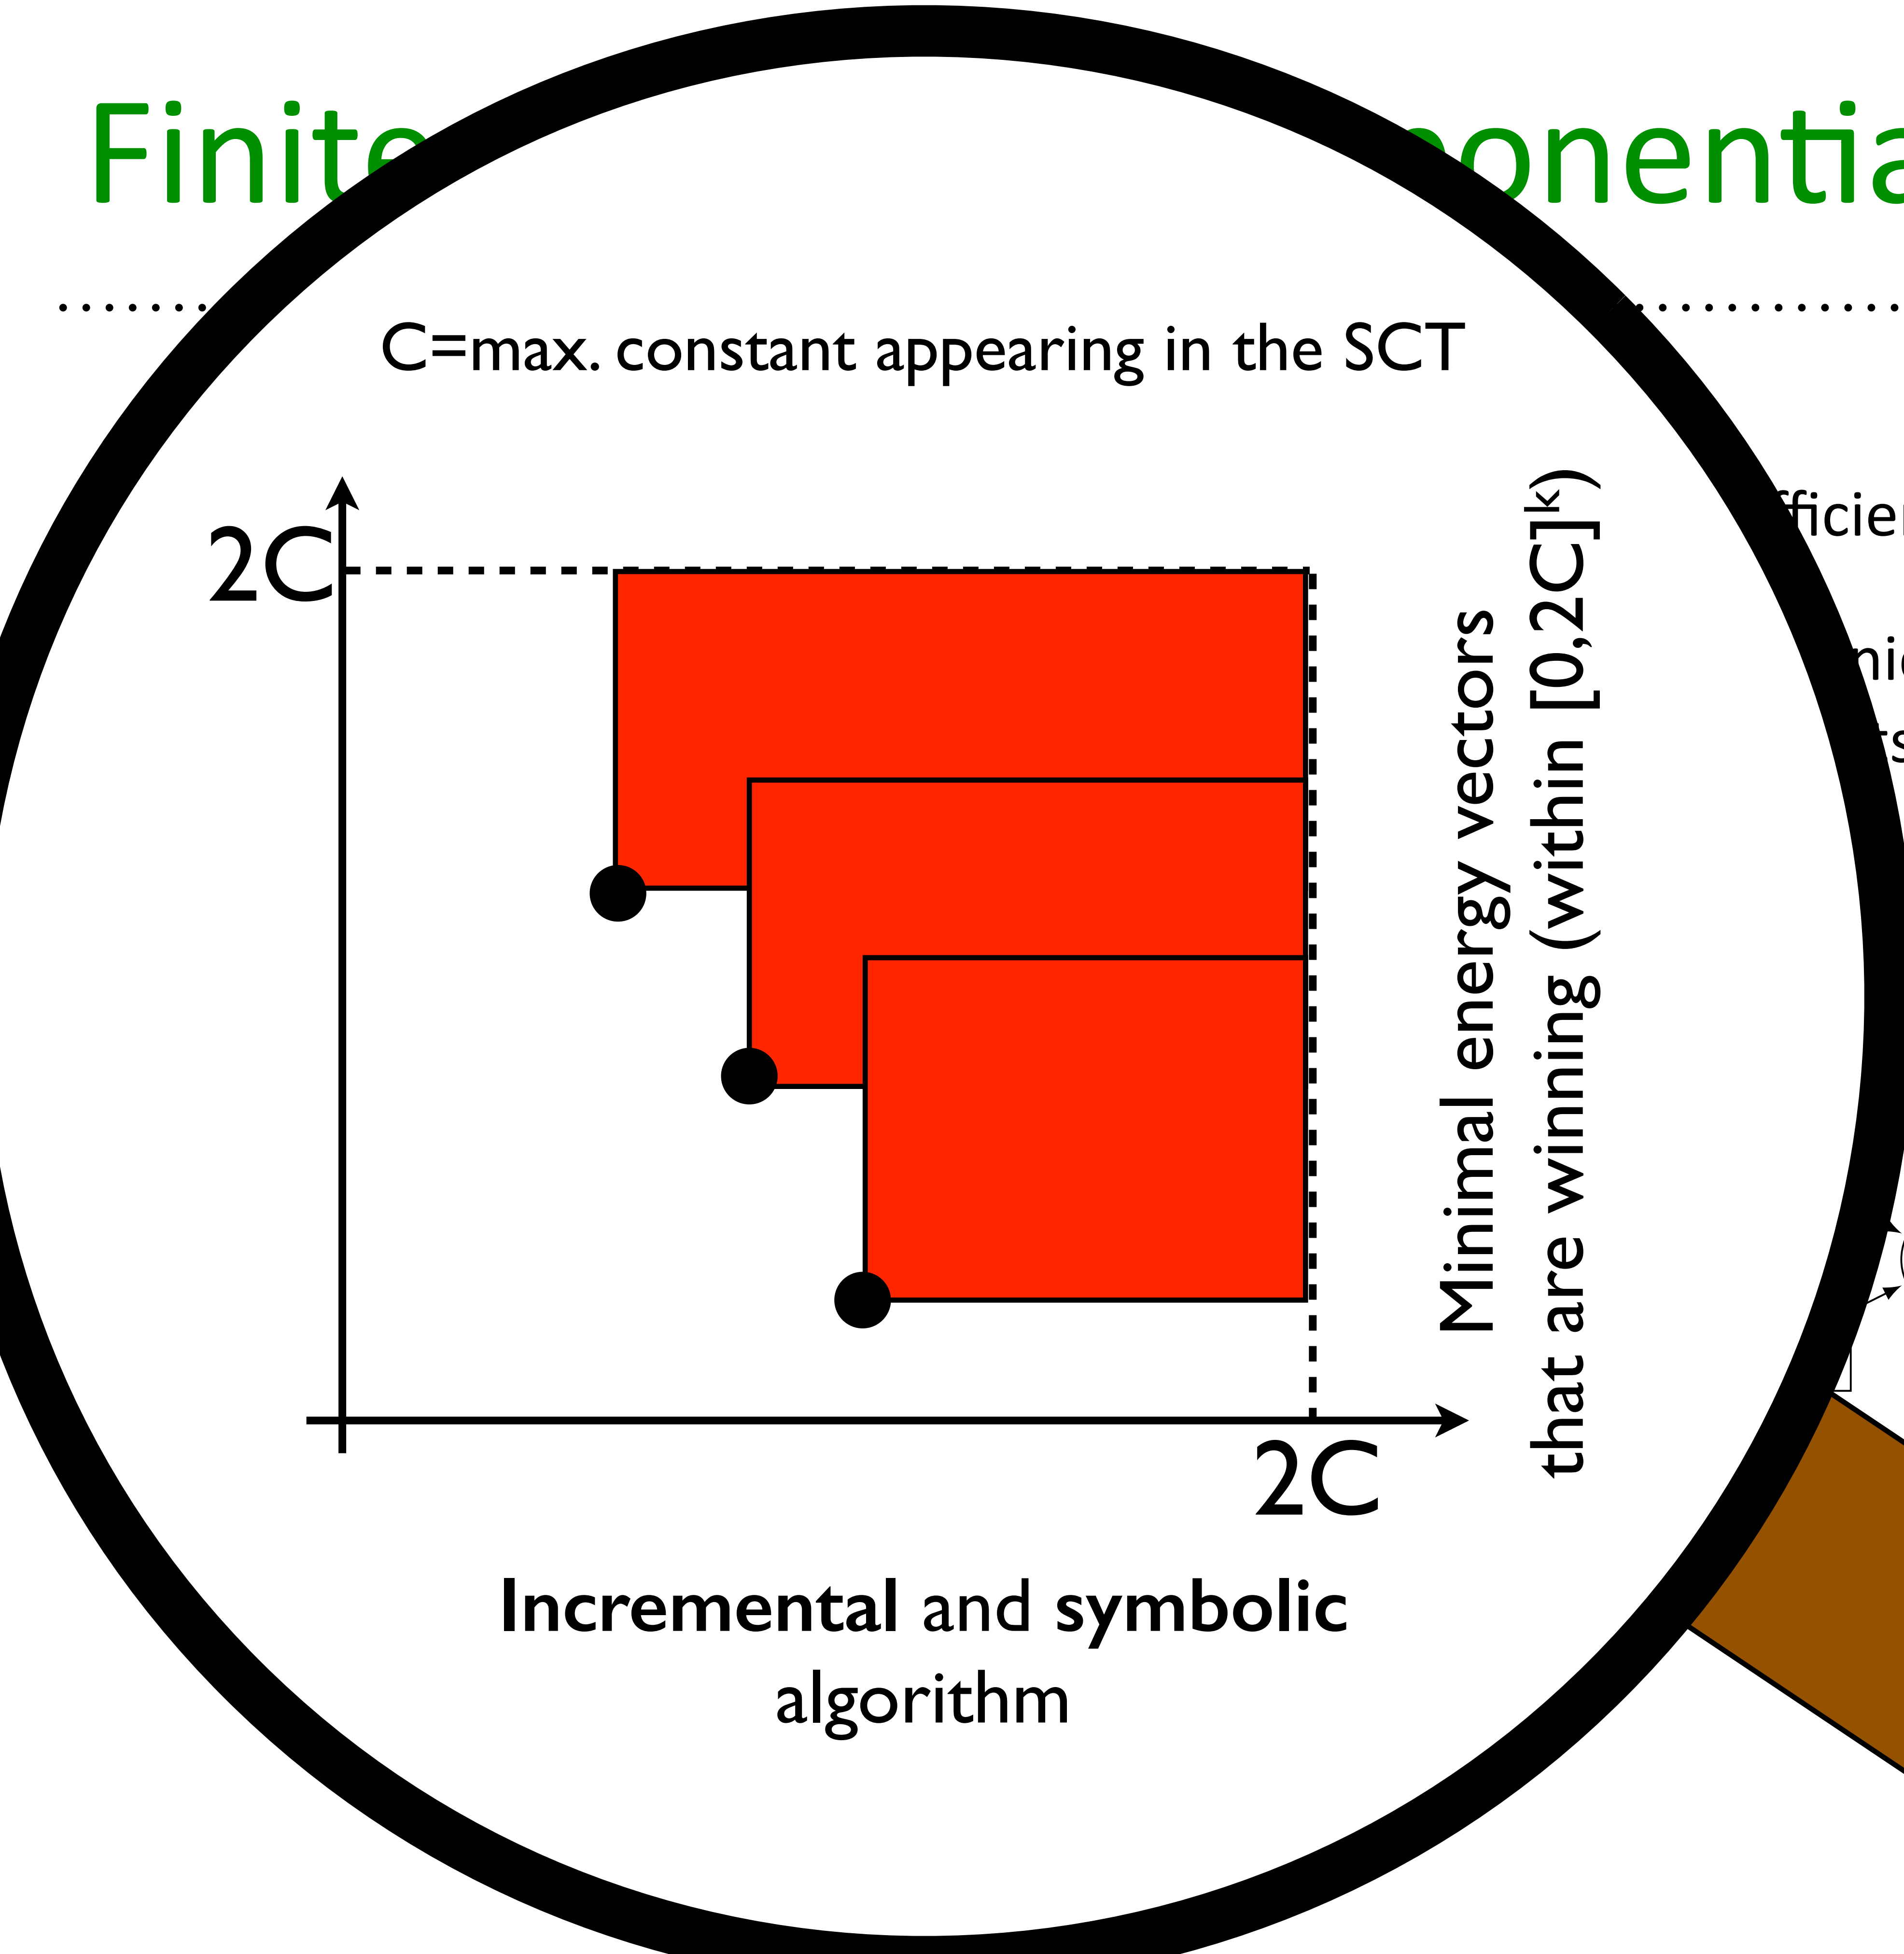

fficient.

nics à la Rackoff
s of [BJK10]

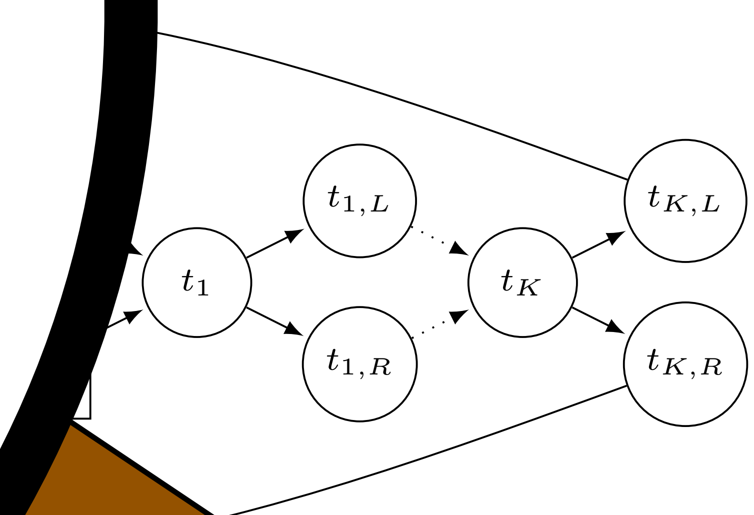

rithms

Complexity of deciding the existence of **memoryless** winning strategies for Player I in MEGs

When Player I plays memoryless

Theorem. The unknown initial credit problem in 2D-MEGs for memoryless strategies is **NP-C**.

Proof.

(i) **Easyness**: Guess a memoryless strategy for Player I and then search (in det. polynomial time) for the existence of a cycle which is negative on at least one dimension (e.g. using Karp's algorithm). If no such cycle exists, the memoryless strategy of Player I is winning.

When Player I plays memoryless

Theorem. The unknown initial credit problem in 2D-MEGs for memoryless strategies is **NP-C**.

Proof.

(i) **Hardness**: Reduction from PARTITION.

Let $A=\{a_1,a_2,\ldots,a_n\}$ and $s : A \rightarrow \mathbb{N}_0$, and let $B=\Sigma_{a\in A}\, s(a)$ (and assume B is even).

PARTITION asks if A can be partitioned into A_L and A_R such that $\Sigma_{a\in}A_L\, s(a)=\Sigma_{a\in}A_R\, s(a)$

When Player I plays memoryless

Theorem. The unknown initial credit problem in 2D-MEGs for memoryless strategies is **NP-C**.

Proof.

(i) **Hardness**: Reduction from PARTITION.

Let $A=\{a_1,a_2,\ldots,a_n\}$ and $s : A \to \mathbb{N}_0$, and let $B=\Sigma_{a\in A}\ s(a)$ (and assume B is even).

We construct the following one player MEG:

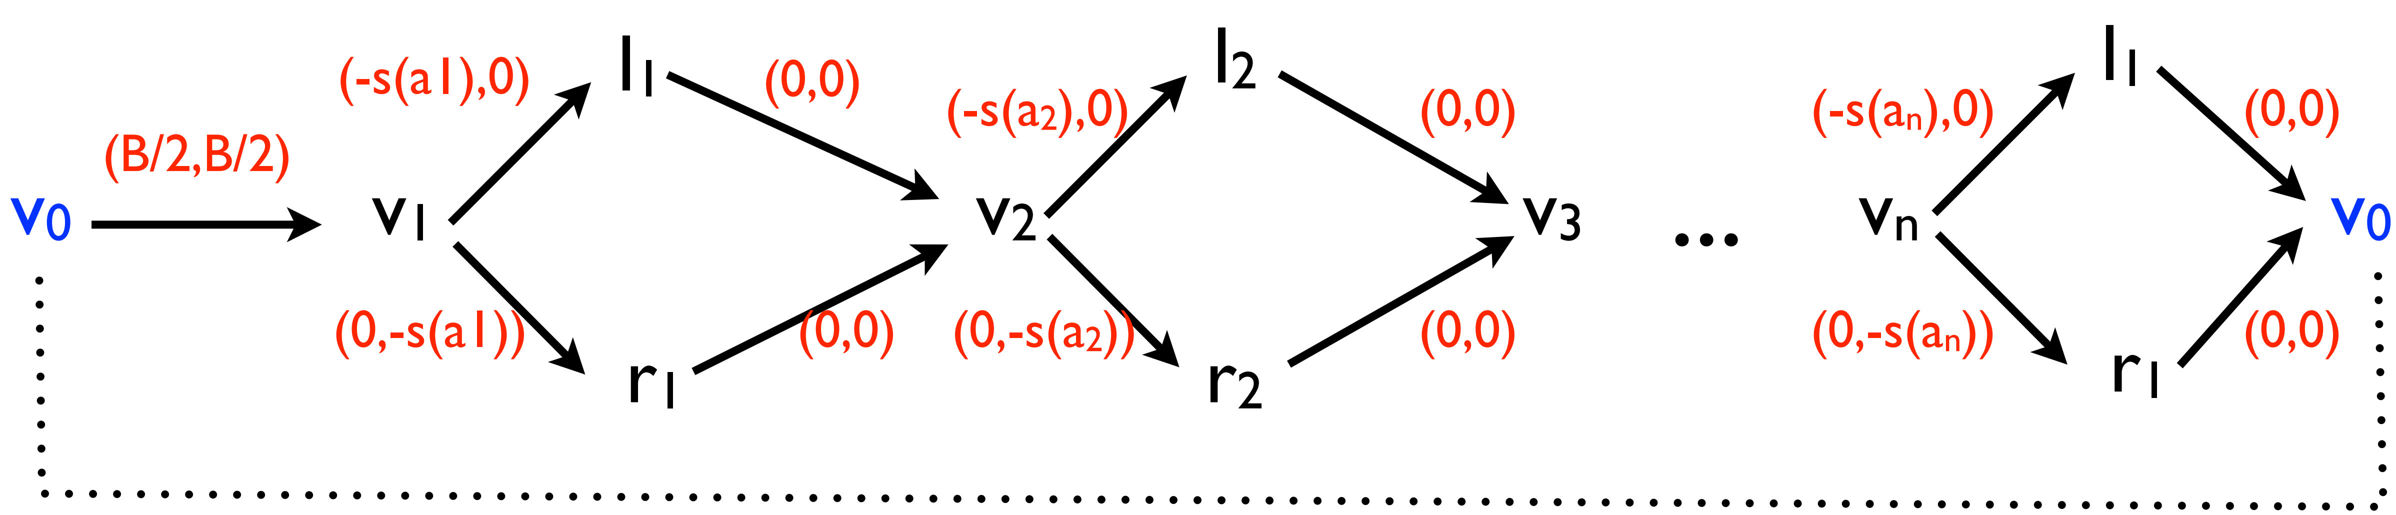

Clearly, a memoryless strategy is winning **iff** it corresponds to a valid right-left partition.

Multi-dim. Mean-payoff Games [RV11]

Two variants: LimSup - LimInf

- Lim Inf
- Lim Sup

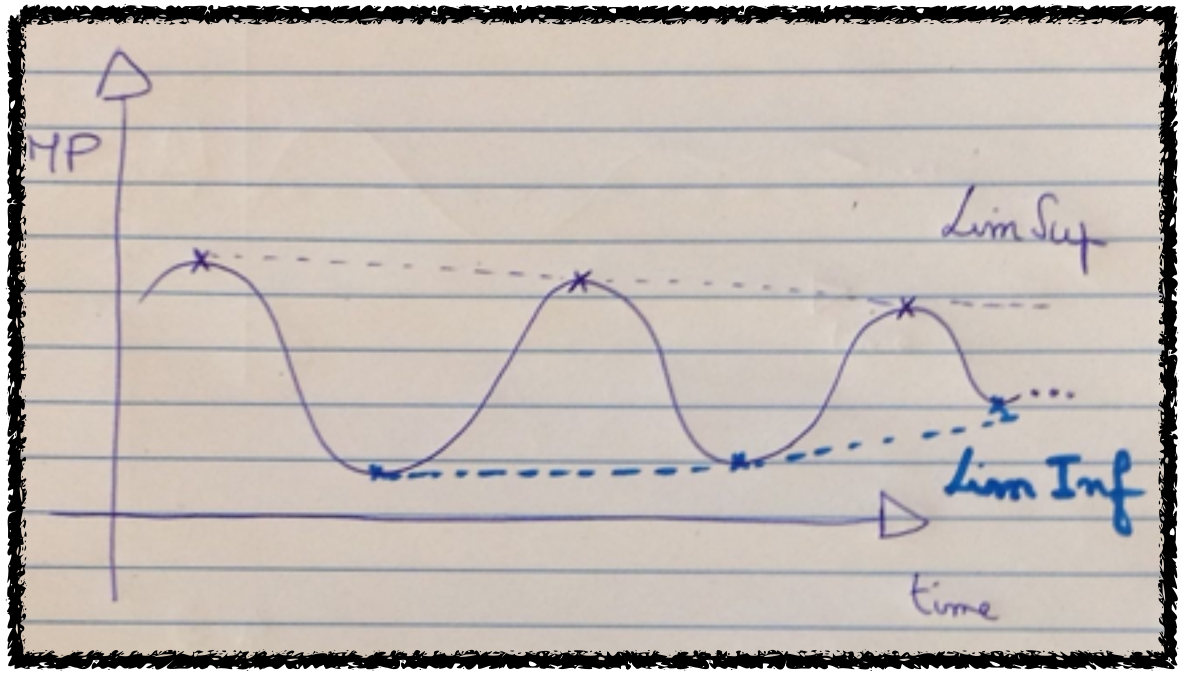

- In the one dimension case, it does **not** make a difference because there exist memoryless optimal strategies, so outcomes can be considered as **ultimately periodic** (and the two limits coincide)

- In the multi-dim. case, it **makes a difference** because optimal strategies may require **infinite memory**

MMPGs - Infinite Memory

Lim-inf MP: define the mean-payoff in each dimension as follows:

Let $\pi : \mathbb{N} \to \mathbb{Z}^2$, we associate to π the pair (u,v) where:

- u=**lim inf** $_{i\to\infty}$ $1/n \times \Sigma_{i=0..n}\ \pi(i)\downarrow 1$ %MP on first dim.
- v=**lim inf** $_{i\to\infty}$ $1/n \times \Sigma_{i=0..n}\ \pi(i)\downarrow 2$ %MP on second dim.

Consider the strategy that alternates visits to q_a and q_b such that after the **n**th alternation, the self-loop on the visited state q ($q \in \{q_a, q_b\}$) is taken **n** times. This strategy achieves threshold (1, 1) for **Lim-inf MP**, as the frequency of edges with (0,0) goes to 0.

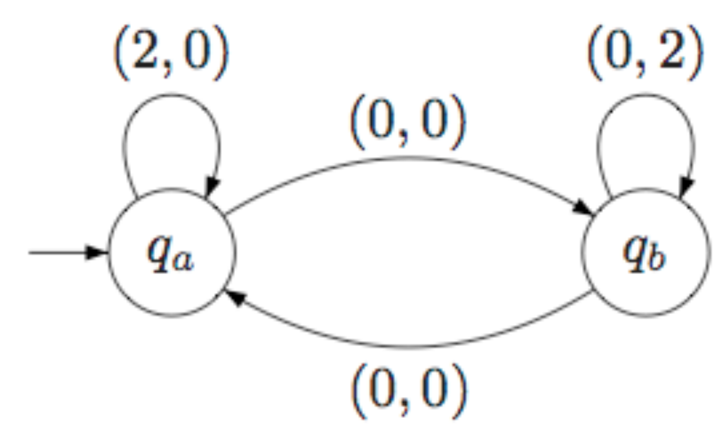

MMPGs - Infinite Memory

Lim-inf MP: define the mean-payoff in each dimension as follows:

Let π : $\mathbb{N} \to \mathbb{Z}^2$, we associate to π the pair (u,v) where:

- u=**lim inf** $_{i \to \infty}$ 1/n × $\Sigma_{i=0..n}$ π(i)↓1 %MP on first dim.
- v=**lim inf** $_{i \to \infty}$ 1/n × $\Sigma_{i=0..n}$ π(i)↓2 %MP on second dim.

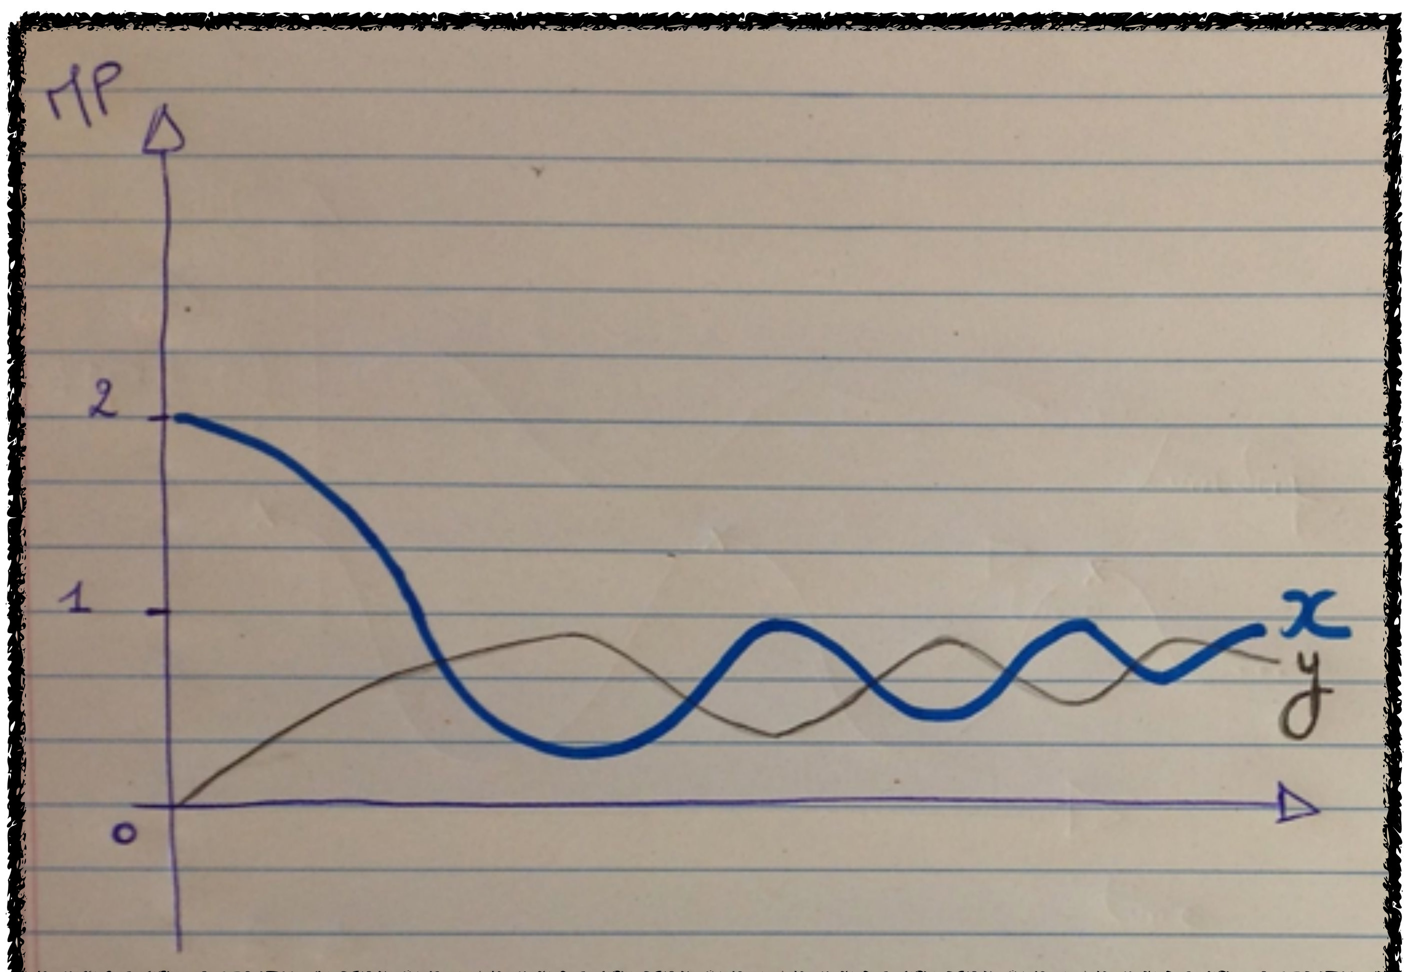

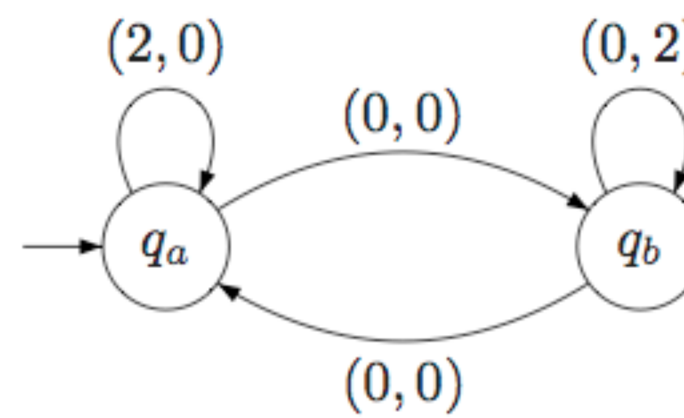

MMPGs - Infinite Memory

Lim-sup MP: define the mean-payoff in each dimension as follows:

Let $\pi : \mathbb{N} \rightarrow \mathbb{Z}^2$, we associate to π the pair (u,v) where:

- u=**lim sup** $_{i\rightarrow\infty}$ $1/n \times \Sigma_{i=0..n}\, \pi(i)\downarrow 1$ %MP on first dim.
- v=**lim sup** $_{i\rightarrow\infty}$ $1/n \times \Sigma_{i=0..n}\, \pi(i)\downarrow 2$ %MP on second dim.

Consider the strategy that alternates visits to q_a and q_b such that after the n^{th} alternation, the self-loop on the visited state q ($q \in \{q_a, q_b\}$) is taken **so many times** that the average frequency of q gets larger than (n-1)/n in the current finite prefix of the play. This is always possible and achieves threshold (2, 2) for **Lim-sup MP**.

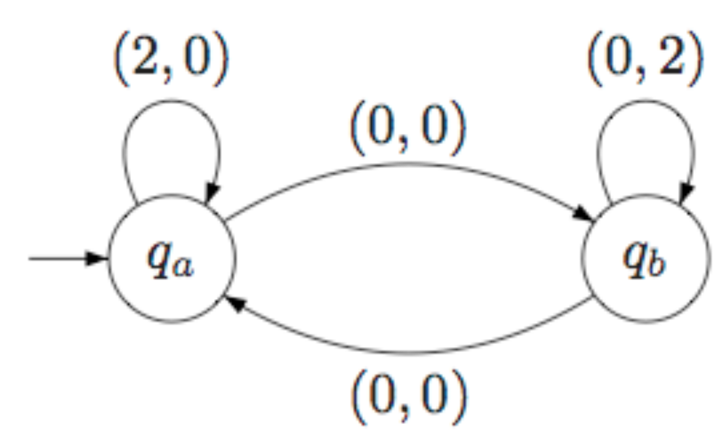

MMPGs - Infinite Memory

Lim-sup MP: define the mean-payoff in each dimension as follows:

Let $\pi : \mathbb{N} \rightarrow \mathbb{Z}^2$, we associate to π the pair (u,v) where:

- u=**lim sup** $_{i \rightarrow \infty}$ $1/n \times \Sigma_{i=0..n}\, \pi(i)\!\downarrow\! 1$ %MP on first dim.
- v=**lim sup** $_{i \rightarrow \infty}$ $1/n \times \Sigma_{i=0..n}\, \pi(i)\!\downarrow\! 2$ %MP on second dim.

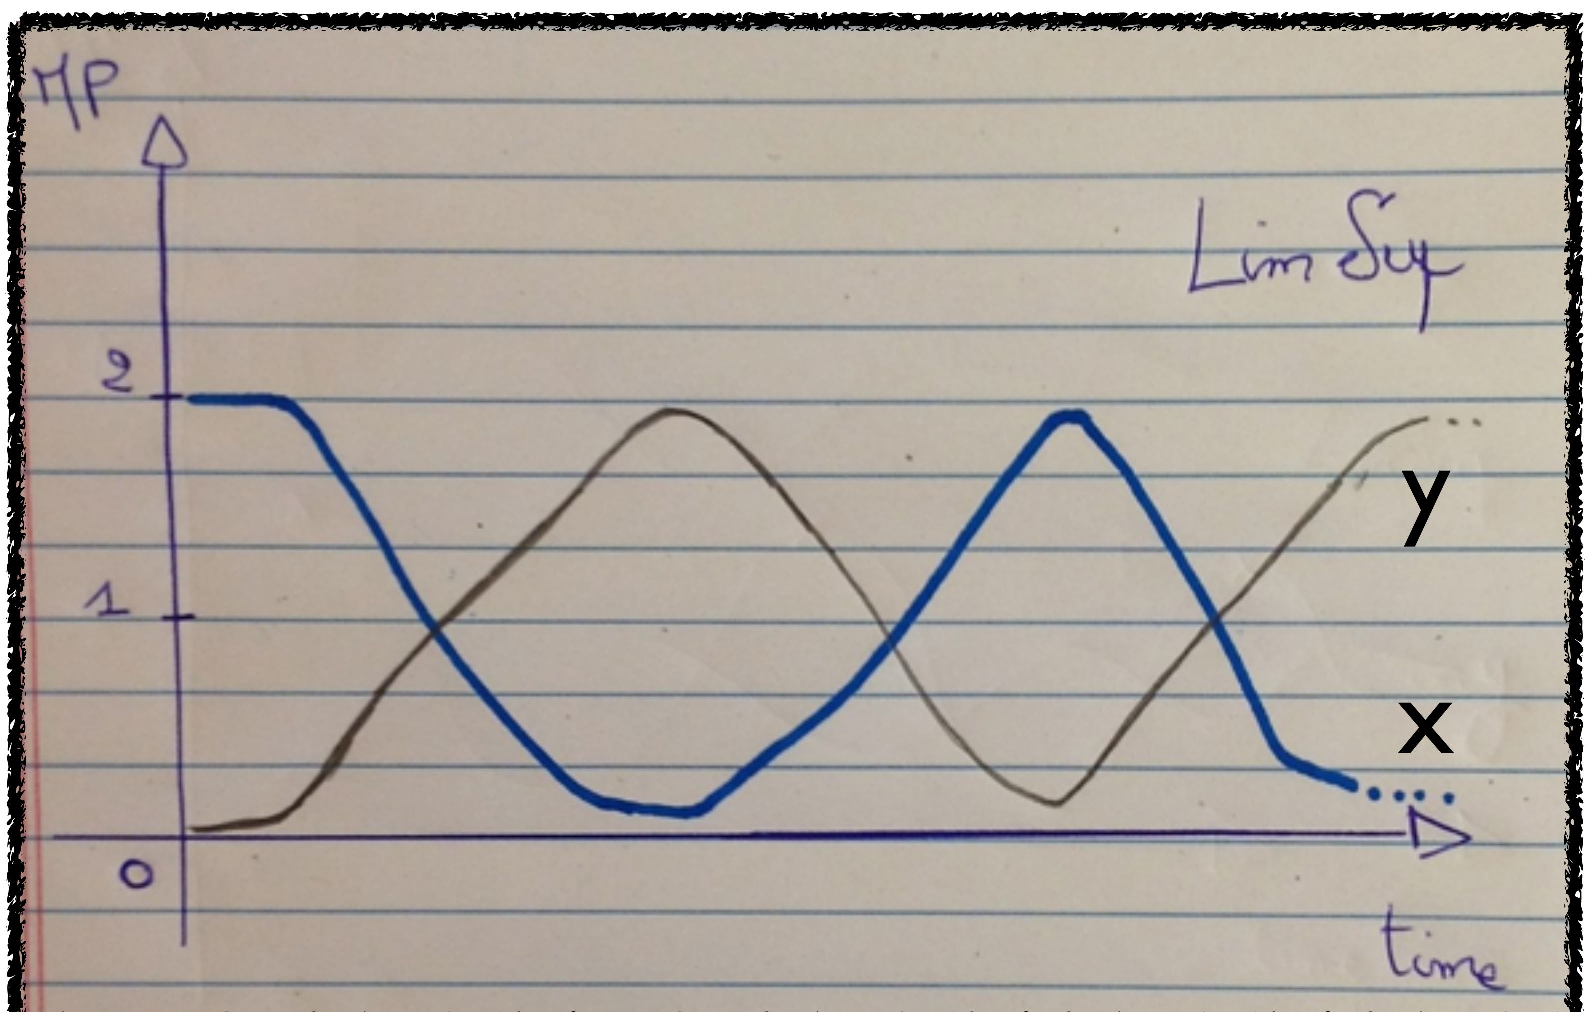

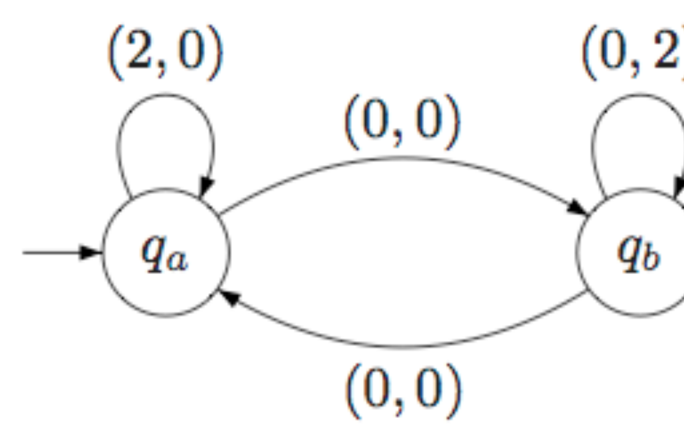

MMPGs - Infinite Memory

	Complexity	Memory Pl. 1	Synthesis Pl. 1	Memory Pl.2
Lim-sup	NP∩coNP	∞	✔	Memoryless
Lim-inf	coNP-C	∞	?	Memoryless
Lim-sup and Lim-inf	coNP-C	∞	?	Memoryless

MMPGs - Lim-sup

Lemma. If for all states $v \in V_1 \cup V_2$, for all i, $1 \leq i \leq k$, Player 1 has a winning strategy for winning the mean-payoff sup. for dimension i, then for all states $v \in V_1 \cup V_2$, Player I has a winning strategy from v for the conjunction of all k mean-payoff objectives.

Intuition: play each of the k winning strategies one after the other for longer and longer time intervals.

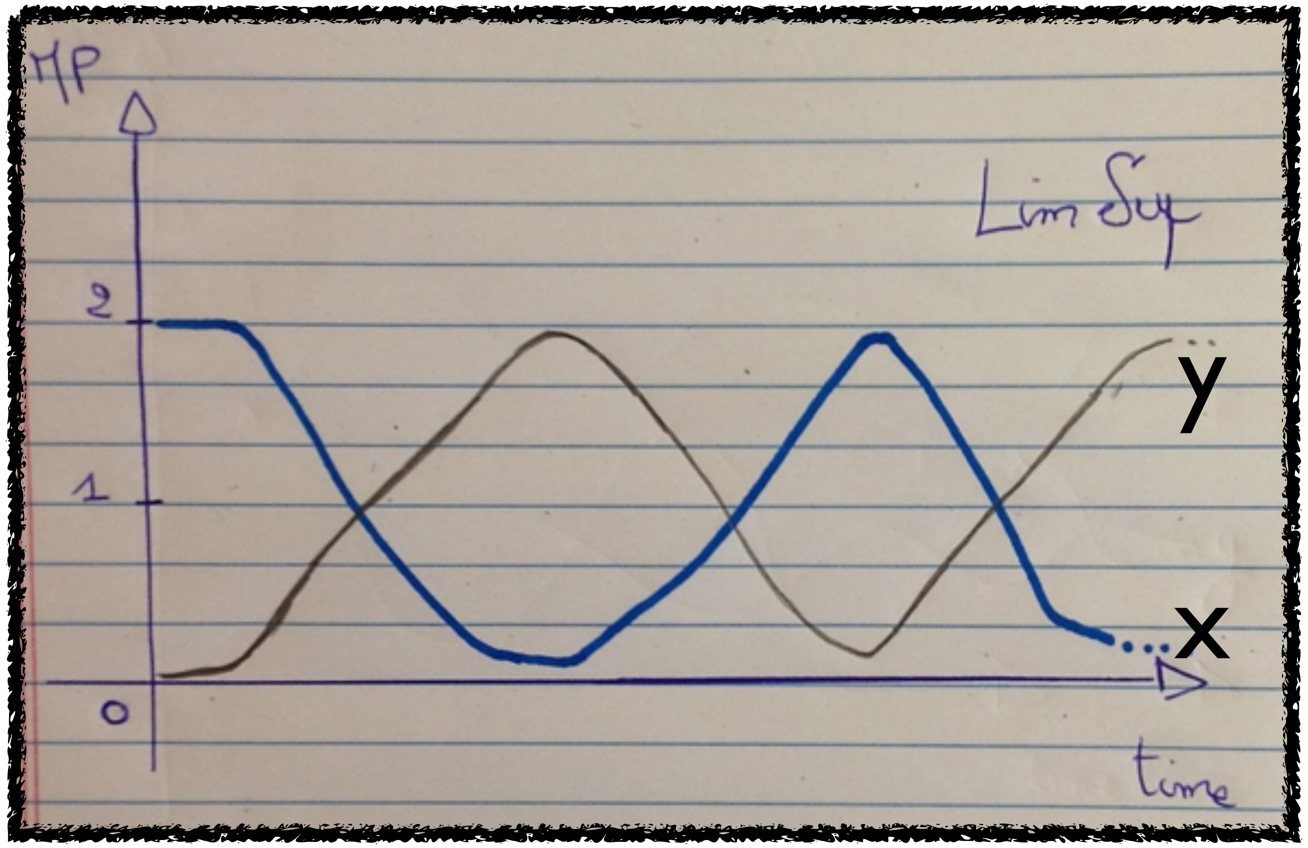

MMPGs - Lim-sup - Algorithm

Consider the following algorithm:

1. Compute W_i=set of states where Pl. I wins the 1 dim. game defined by dim. i

2. Let W be the intersection of all W_i's

3. Remove states that are not in W

Repeat until no states are removed

Let Win be the states that survived this process

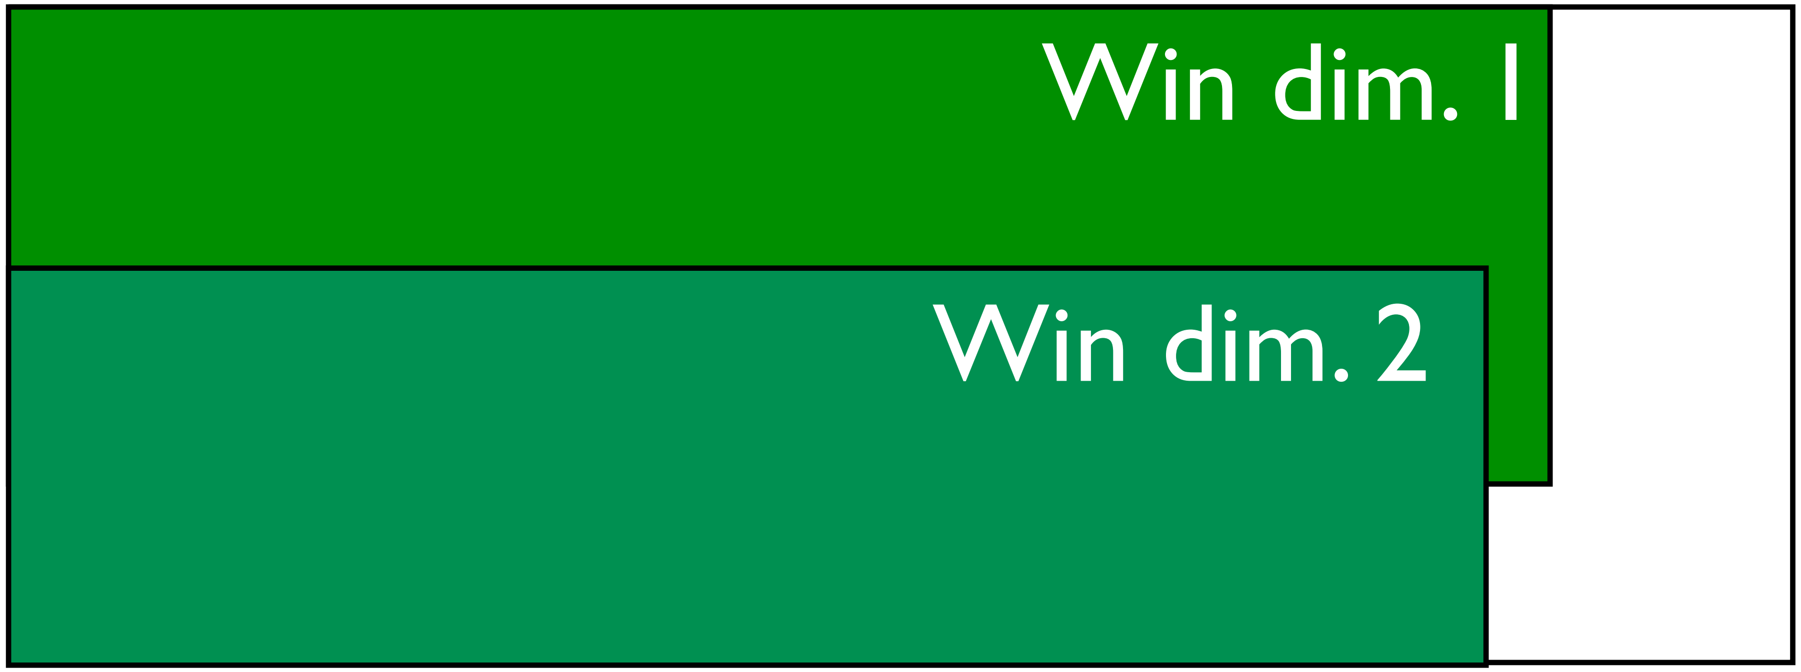

MMPGs - Lim-sup - Algorithm

Consider the following algorithm:

1. Compute W_i=set of states where Pl. I wins the 1 dim. game defined by dim. i
2. Let W be the intersection of all W_i's
3. Remove states that are not in W

Repeat until no states are removed

Let Win be the states that survived this process

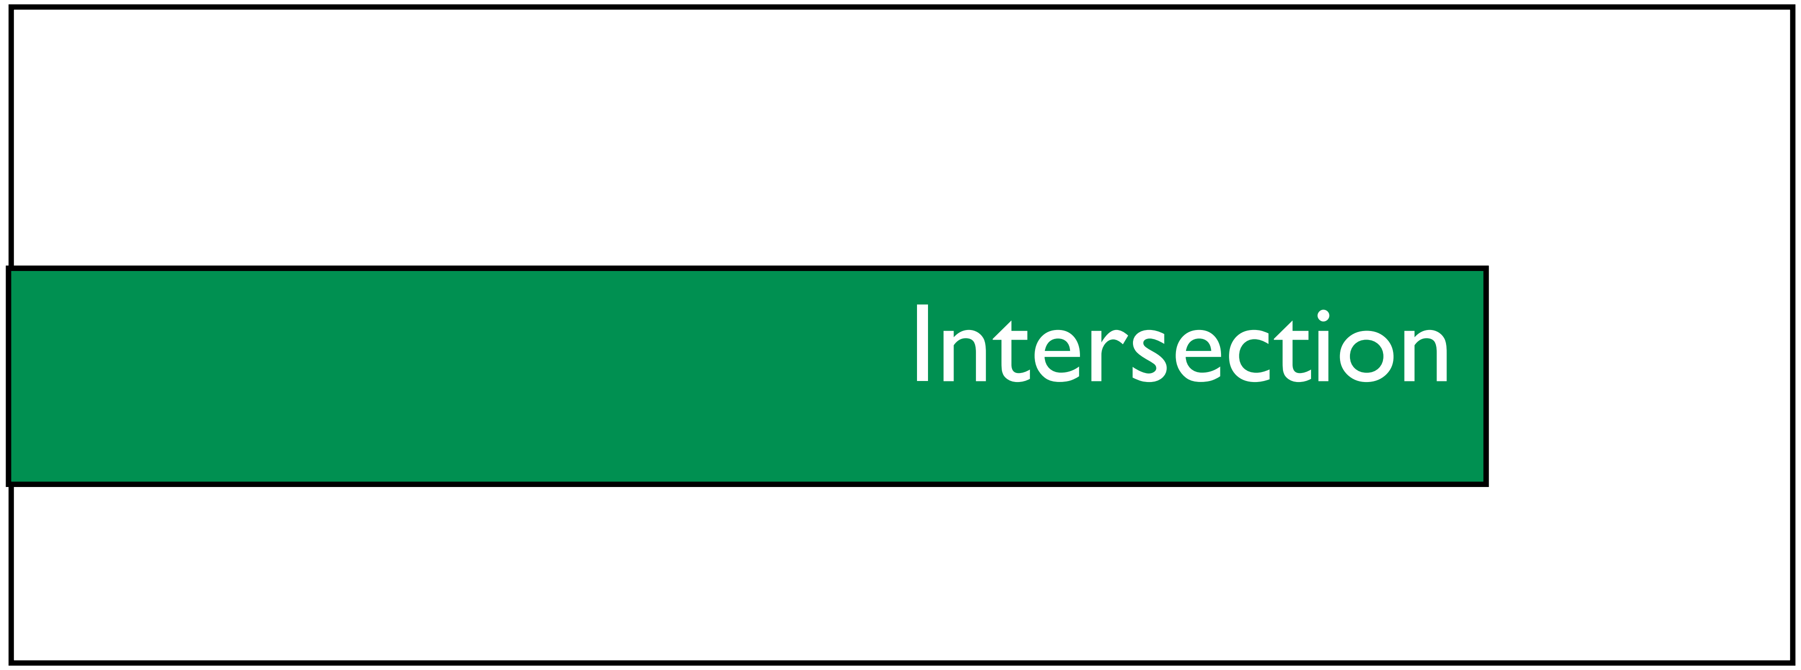

MMPGs - Lim-sup - Algorithm

Consider the following algorithm:

1. Compute W_i=set of states where Pl. I wins the 1 dim. game defined by dim. i
2. Let W be the intersection of all W_i's
3. Remove states that are not in W

Repeat until no states are removed

Let Win be the states that survived this process

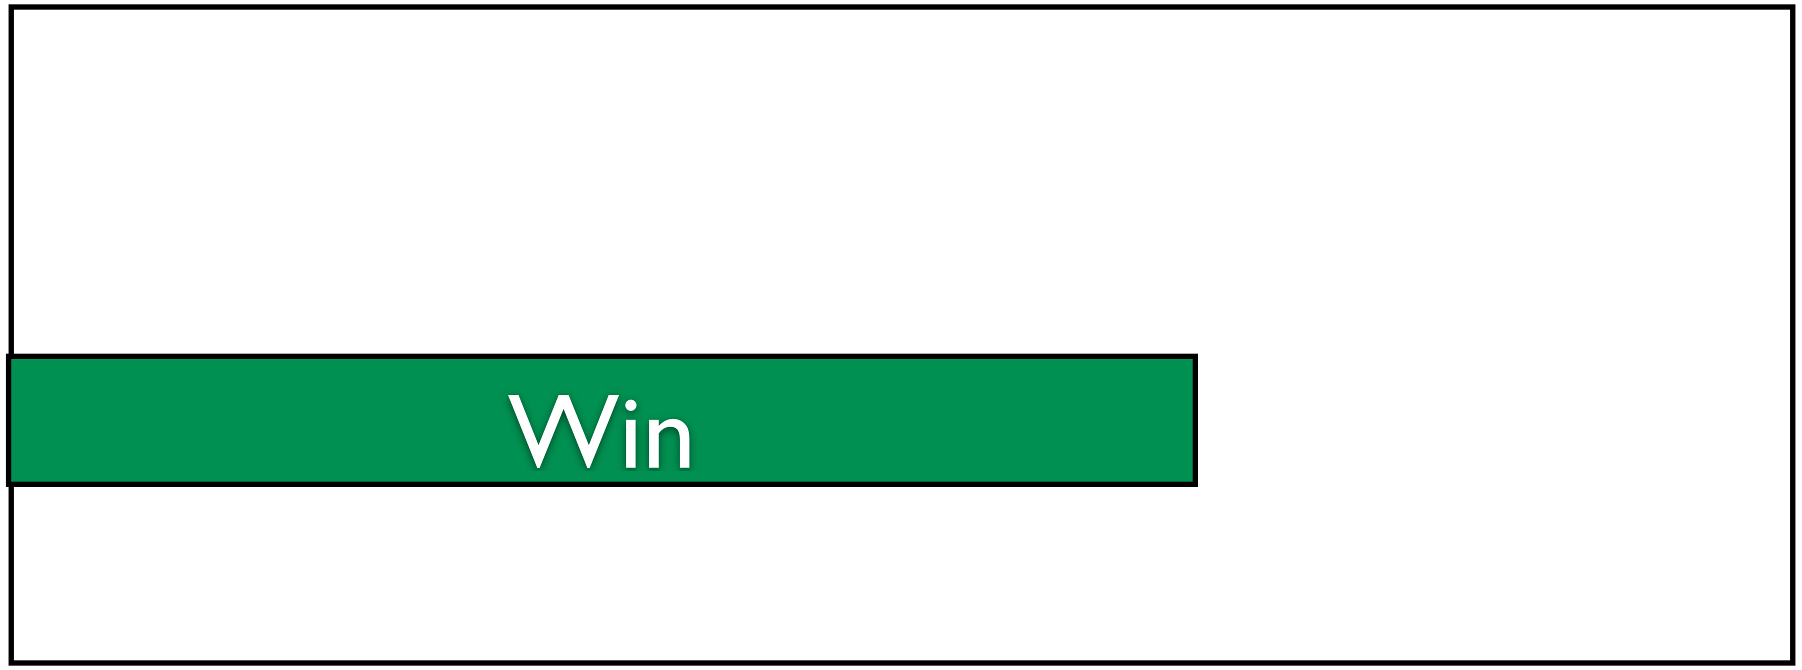

MMPGs - Lim-sup - Algorithm

Lemma. From all states in Win, Player 1 has a winning strategy for each dimension. From all states that are not in Win, Player II has a winning strategy for at least one dimension.

Theorem [RV11]. From all states in Win, Player 1 has a winning strategy for all the dimensions (for Lim Sup).

Corollary. Deciding MMPGs with Lim-sup is in NP$\cap$coNP.

Summary

	Opt. Stg. Player 1	Synthesis	Opt. Stg. Player 2	Complexity Decision
EG	Memoryless	OK	Memoryless	NP∩coNP
MP	Memoryless	OK	Memoryless	NP∩coNP
MEG	Exponential	OK	Memoryless	coNP-C
MEG-finite	-	OK	Memoryless	NP-C
MMPG - Sup	Infinite	OK	Memoryless	NP∩coNP
MMPG - Inf	Infinite	?	Memoryless	coNP-C
MMPG - Mix	Infinite	?	Memoryless	coNP-C

References

Mean-payoff games

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

Energy games

Patricia Bouyer, Ulrich Fahrenberg, Kim Guldstrand Larsen, Nicolas Markey, and Jiri Srba. I**nfinite runs in weighted timed automata with energy constraints**. In *FORMATS*, volume 5215 of *Lecture Notes in Computer Science*. Springer, 2008.

Fixpoint algorithm

Lubos Brim, Jakub Chaloupka, Laurent Doyen, Raffaella Gentilini, and Jean-François Raskin. F**aster algorithms for mean-payoff games**. *Formal Methods in System Design*, 38(2):97–118, 2011.

Multi-dimension MP and EG

Yaron Velner, Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger, Alexander Rabinovich, and Jean-François Raskin. T**he complexity of multi-mean-payoff and multi-energy games**. *CoRR*, abs/1209.3234, 2012.

Krishnendu Chatterjee, Laurent Doyen, Thomas A. Henzinger, and Jean-François Raskin. G**eneralized mean- payoff and energy games**. In *FSTTCS*, volume 8 of *LIPIcs*. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2010.

Tomás Brázdil, Petr Jancar, Antonín Kucera: **Reachability Games on Extended Vector Addition Systems with States.** ICALP (2) 2010: 478-489