Energy and Mean-payoff Games

Laurent Doyen
LSV, ENS Cachan & CNRS

joint work with

Aldric Degorre, Raffaella Gentilini,
Jean-François Raskin, Szymon Torunczyk

ACTS 2010, Chennai

Synthesis problem

Specification

avoid failure,
ensure progress,
etc.

$\models$

Correctness
relation

Synthesis problem

System - Model

$\models$

Correctness relation

Specification

avoid failure,
ensure progress,
etc.

Solved as a game – system vs. environment

➡ solution = winning strategy

This talk: quantitative games (resource-constrained systems)

Energy games

(staying alive)

Energy games (CdAHS03,BFLM08)

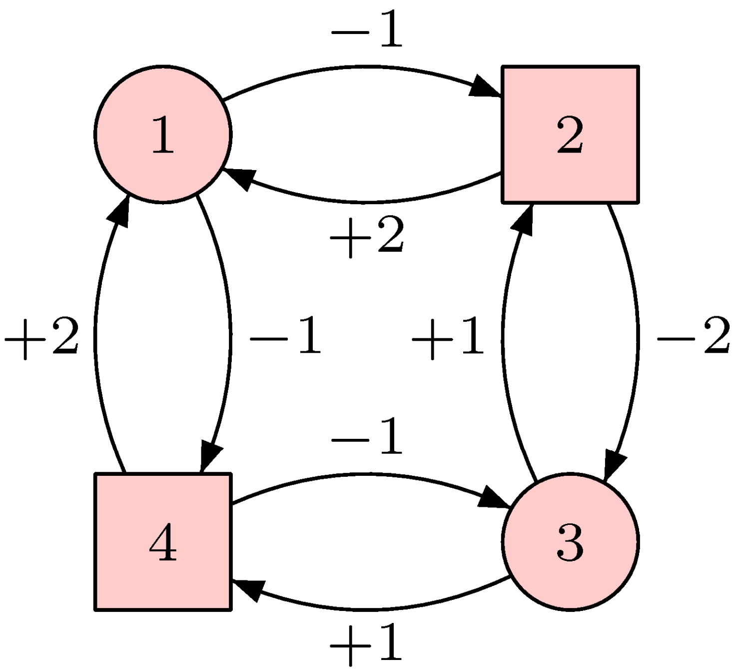

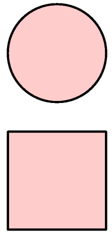

Maximizer

Minimizer

positive weight = reward

play: (1,4) (4,1) (1,4) (4,1) …

weights: -1 +2 -1 +2 …

energy level: **1** 0 2 1 3 2 4 3 …

Energy games (CdAHS03,BFL+08)

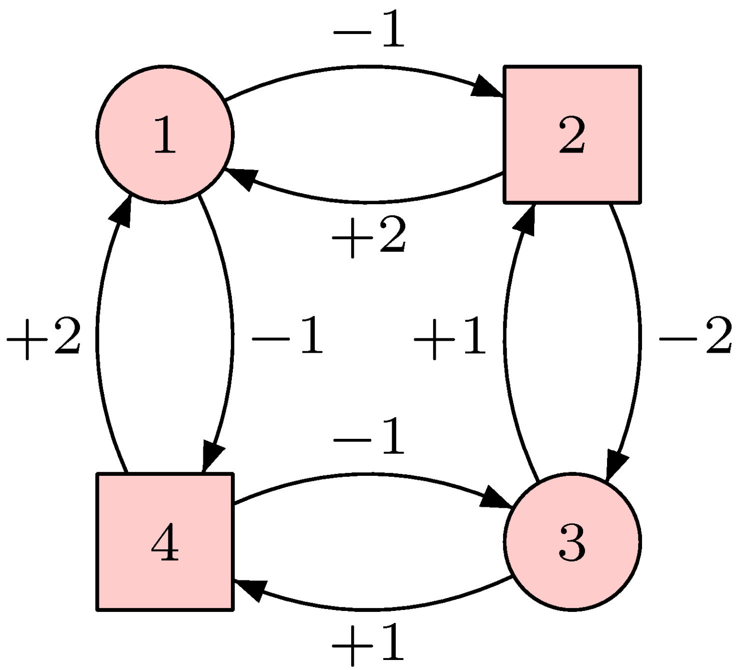

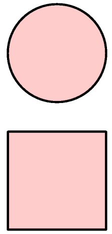

Maximizer

Minimizer

positive weight = reward

play: (1,4) (4,1) (1,4) (4,1) …

weights: -1 +2 -1 +2 …

energy level: **1** 0 2 1 3 2 4 3 …

Initial credit

Energy games

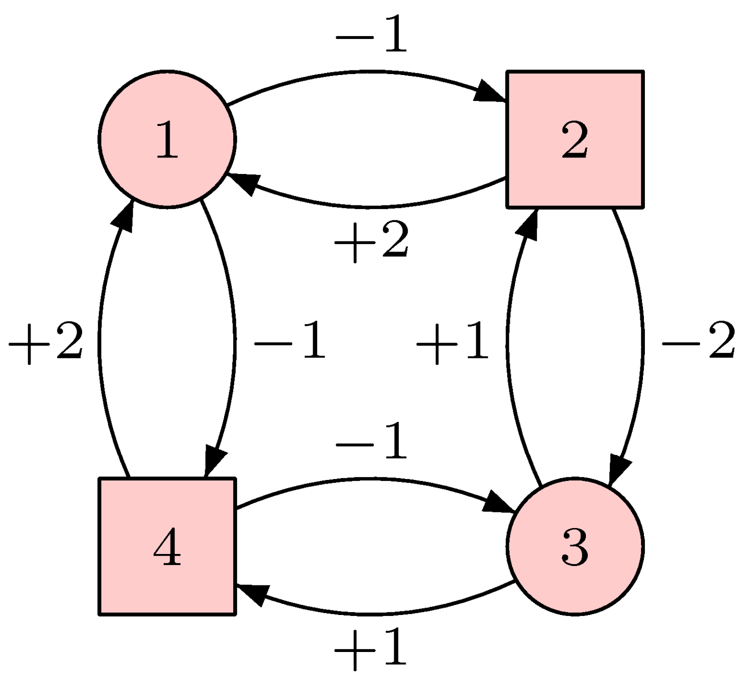

Strategies:

Maximizer $\sigma : Q^* \cdot Q_\circ \to Q$

Minimizer $\pi : Q^* \cdot Q_\Box \to Q$

play: outcome(q, σ, π)

Infinite sequence of edges consistent with strategies σ and π

outcome is winning if:

Energy level $\rightarrow$ $c_0 + \sum_{i=0}^{n-1} w_i \geq 0$ for all $n \geq 0$

Energy games

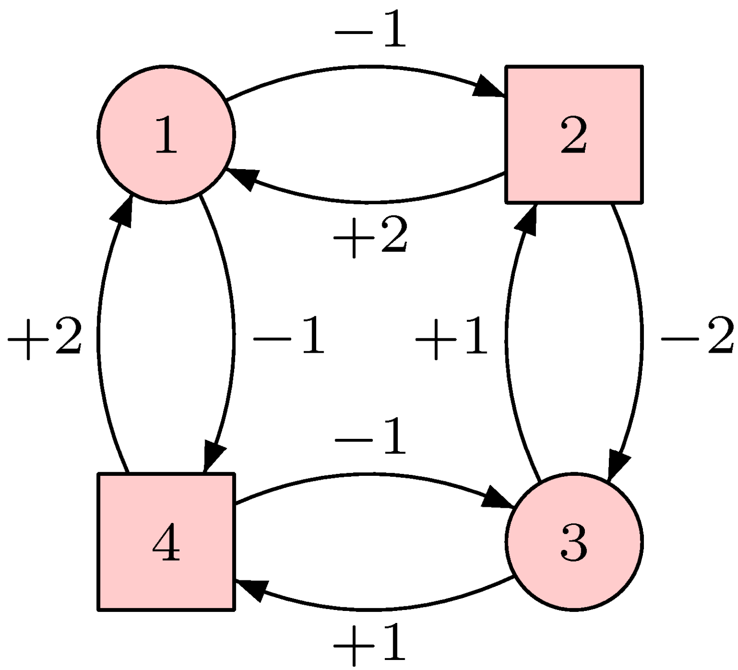

Decision problem:

Decide if there exist an initial credit c_0 and a strategy of the maximizer to maintain the energy level always nonnegative.

Energy games

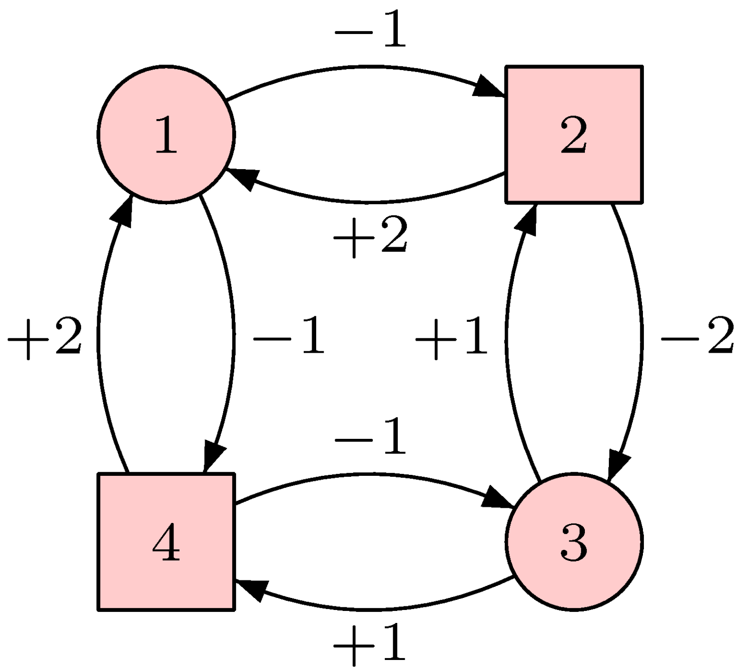

Decision problem:

Decide if there exist an initial credit c_0 and a strategy of the maximizer to maintain the energy level always nonnegative.

For energy games,
memoryless strategies suffice.

$$\sigma : Q_{\Box} \rightarrow Q$$

$$\pi : Q_{\circ} \rightarrow Q$$

Energy games

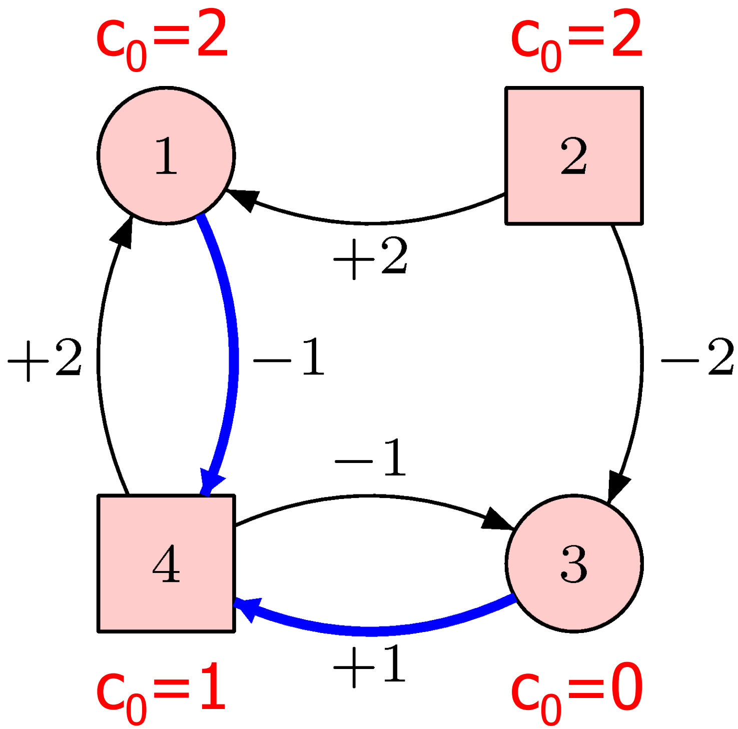

Decision problem:

Decide if there exist an initial credit c_0 and a strategy of the maximizer to maintain the energy level always nonnegative.

For energy games,
memoryless strategies suffice.

A memoryless strategy σ is winning if all cycles are nonnegative when σ is fixed.

Energy games

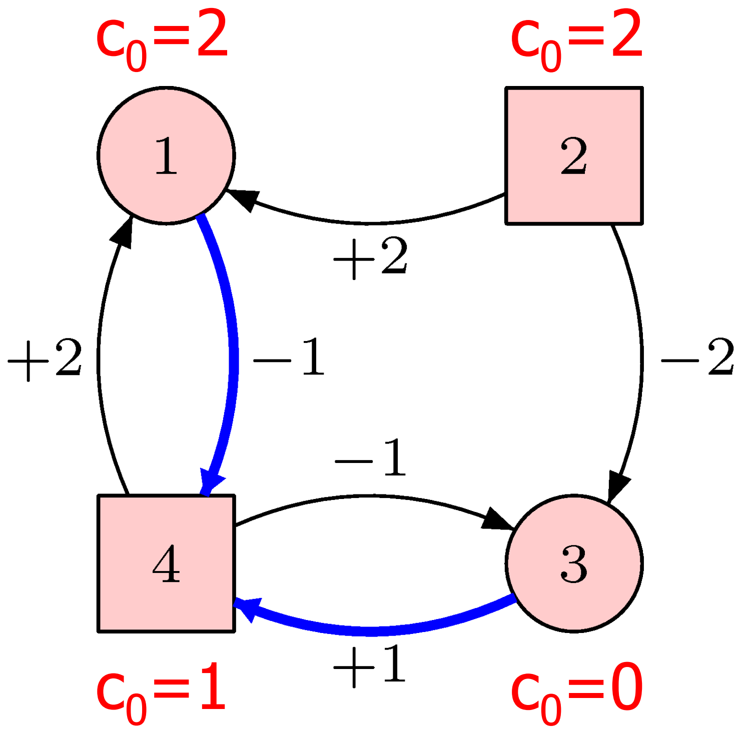

Decision problem:

Decide if there exis… …redit c_0
and a strat… …to
maintai… …always
nonnega…

NP ∩ coNP

For energy games,
memoryless strategies suffice.

A memoryless strategy σ is winning if all cycles are nonnegative when σ is fixed.

Algorithm

Algorithm for energy games

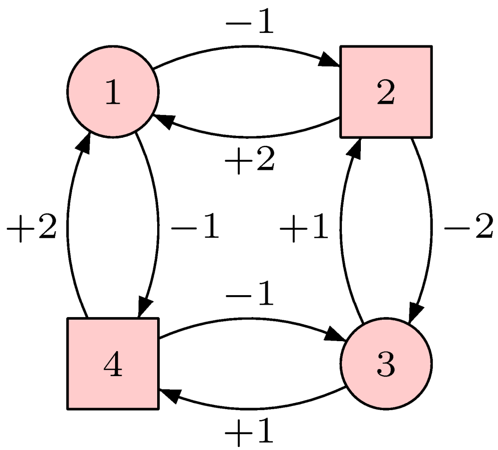

Initial credit is useful to survive before a cycle is formed

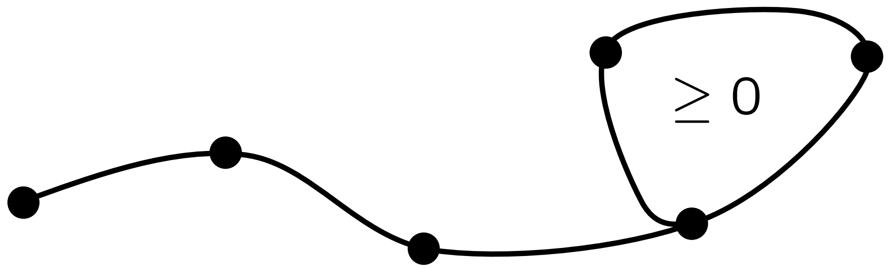

Length(AcyclicPath) ≤ Q

Q: #states
E: #edges
W: maximal weight

Algorithm for energy games

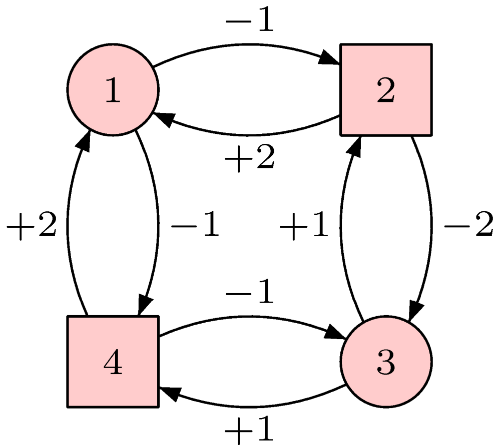

Initial credit is useful to survive before a cycle is formed

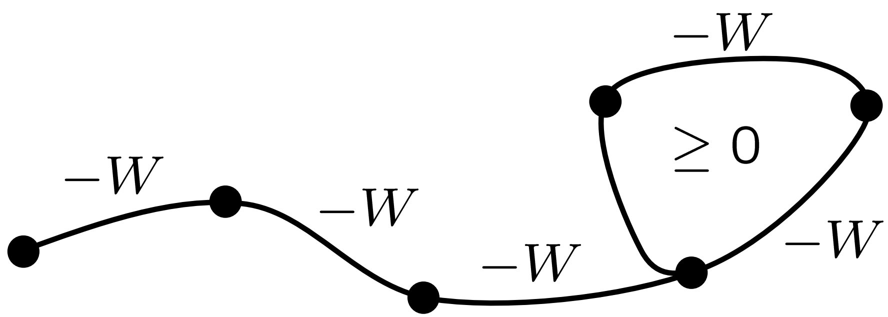

Length(AcyclicPath) ≤ Q

Q: #states
E: #edges
W: maximal weight

Minimum initial credit is at most Q·W

Algorithm for energy games

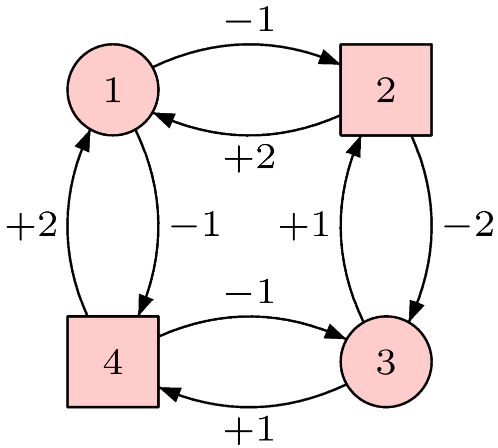

The minimum initial credit $v(\cdot)$ is such that:

in Maximizer state q:

$$v(q) + w(q, q') \geq v(q') \text{ for some } (q, q') \in E$$

in Minimizer state q:

$$v(q) + w(q, q') \geq v(q') \text{ for all } (q, q') \in E$$

➡ Compute successive under-approximations of the minimum initial credit.

Algorithm for energy games

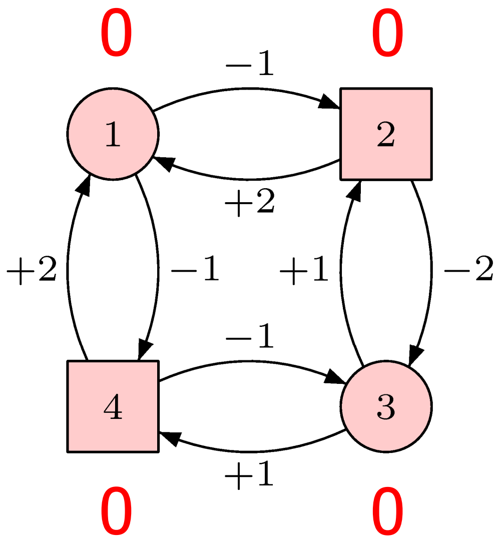

Fixpoint algorithm:

- start with $v(q) = 0$

Algorithm for energy games

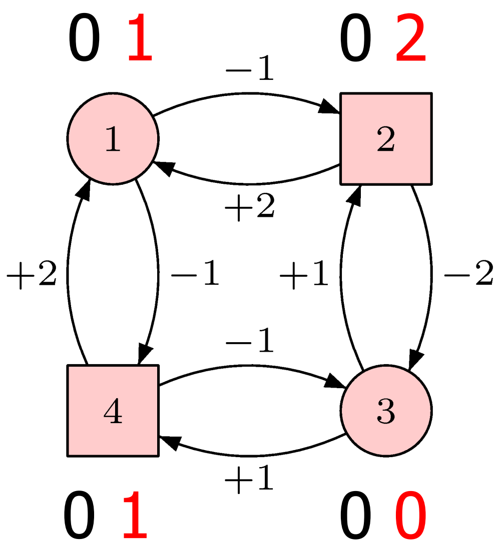

Fixpoint algorithm:

- start with $v(q) = 0$

- iterate

  at Maximizer states:

  $$v(q) \leftarrow \min\{v(q') - w(q,q') \mid (q,q') \in E\}$$

  at Minimizer states:

  $$v(q) \leftarrow \max\{v(q') - w(q,q') \mid (q,q') \in E\}$$

Algorithm for energy games

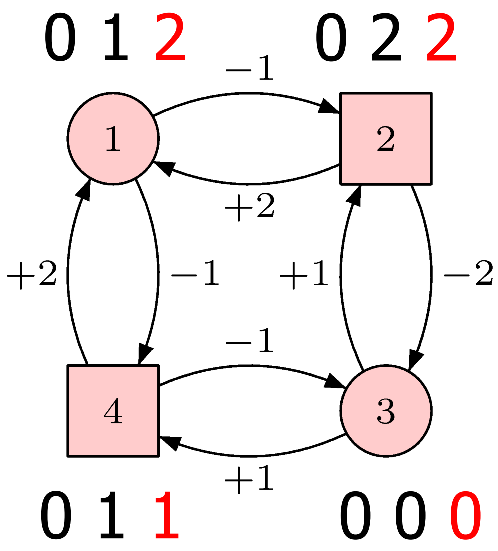

Fixpoint algorithm:

- start with $v(q) = 0$

- iterate

  at Maximizer states:

  $$v(q) \leftarrow \min\{v(q') - w(q,q') \mid (q,q') \in E\}$$

  at Minimizer states:

  $$v(q) \leftarrow \max\{v(q') - w(q,q') \mid (q,q') \in E\}$$

Algorithm for energy games

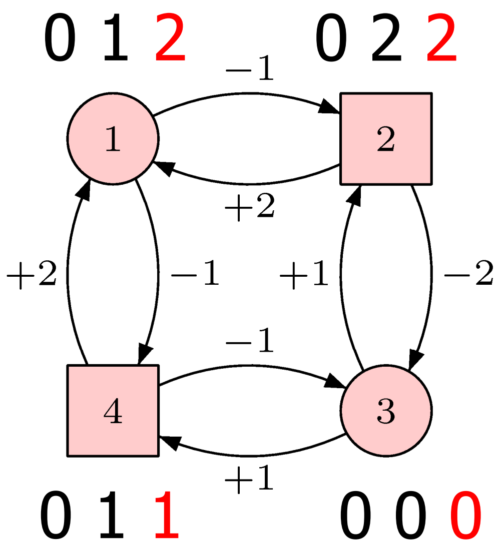

Fixpoint algorithm:

- start with $v(q) = 0$

- iterate

  at Maximizer states:

  $$v(q) \leftarrow \min\{v(q') - w(q,q') \mid (q,q') \in E\}$$

  at Minimizer states:

  $$v(q) \leftarrow \max\{v(q') - w(q,q') \mid (q,q') \in E\}$$

Termination argument: monotonic operators,

and finite codomain $v(q) \in \{0, 1, \ldots, Q \cdot W\}$

Complexity: O(E·Q·W)

Mean-payoff games

Mean-payoff games (EM79)

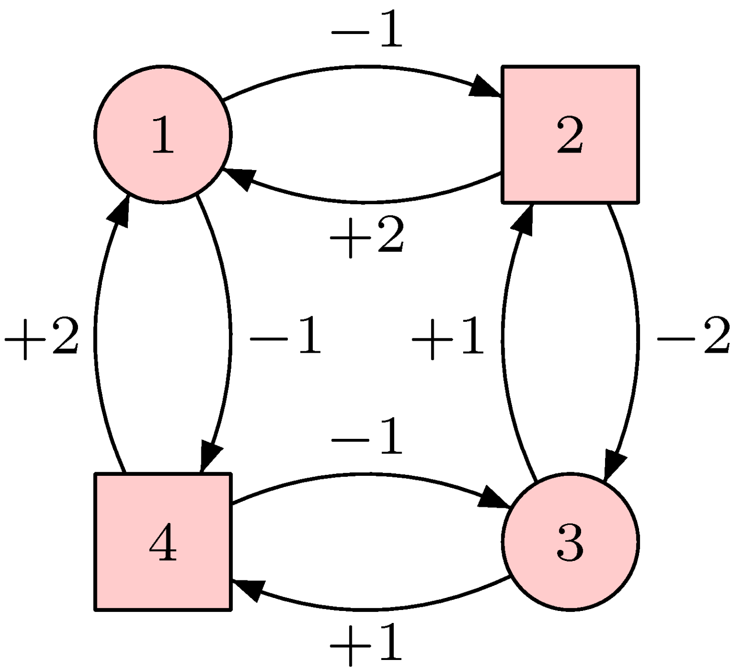

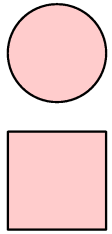

Maximizer

Minimizer

positive weight = reward

play: (1,4) (4,1) (1,4) (4,1) …

weights: -1 +2 -1 +2 …

mean-payoff value: $\frac{1}{2}$
(limit of weight average)

Mean-payoff games (EM79)

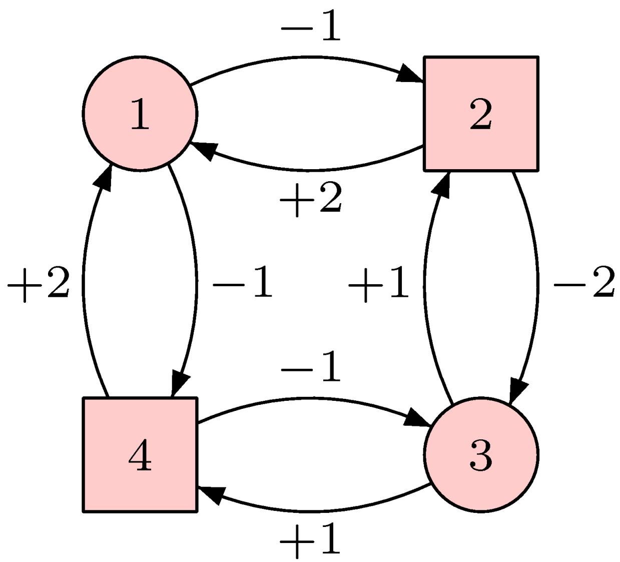

Mean-payoff value:

either $\liminf_{n\to\infty} \frac{1}{n} \cdot \sum_{i=0}^{n-1} w_i$ or $\limsup_{n\to\infty} \frac{1}{n} \cdot \sum_{i=0}^{n-1} w_i$

Decision problem:

Given a rational threshold ν, decide if there exists a strategy of the maximizer to ensure mean-payoff value at least ν.

Note: we can assume $\nu = 0$
e.g. by shifting all weights by ν.

Mean-payoff games

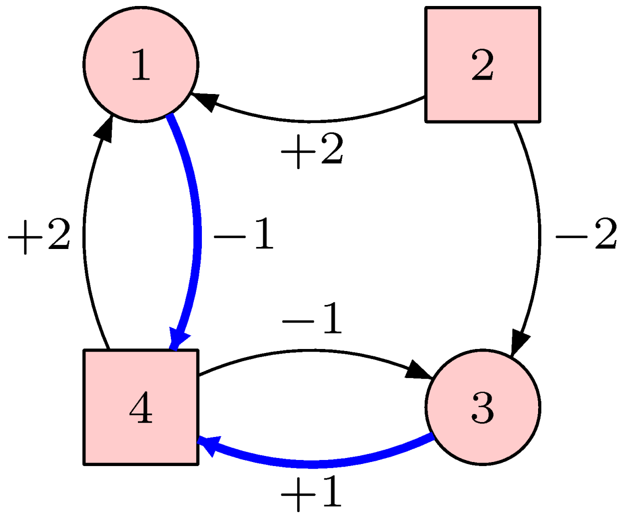

Mean-payoff value:

either $\liminf_{n\to\infty} \frac{1}{n} \cdot \sum_{i=0}^{n-1} w_i$ or $\limsup_{n\to\infty} \frac{1}{n} \cdot \sum_{i=0}^{n-1} w_i$

Decision problem:

Given a rational threshold ν, decide if there exists a strategy of the maximizer to ensure mean-payoff value at least ν.

Assuming $\nu = 0$

A memoryless strategy σ is winning if all cycles are nonnegative when σ is fixed.

Mean-payoff games

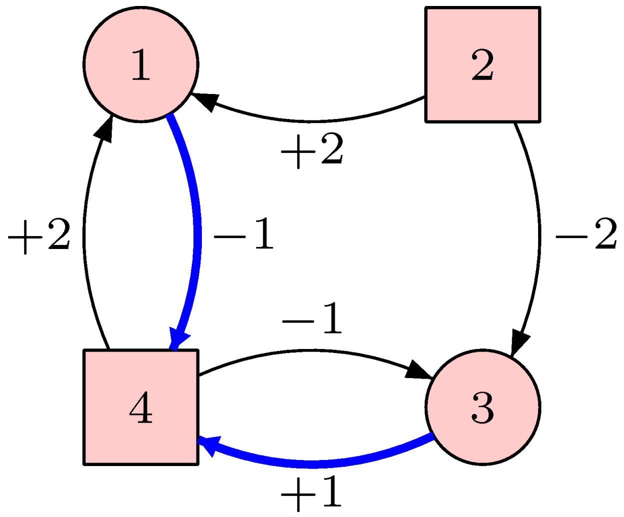

Mean-payoff value:

either $\liminf_{n\to\infty} \frac{1}{n} \cdot \sum_{i=0}^{n-1} w_i$ or $\limsup_{n\to\infty} \frac{1}{n} \cdot \sum_{i=0}^{n-1} w_i$

Decision problem:

Given a rational thres[...]ide
if there exists [...]e
maxim[...]nean-payoff
value at [...].

log-space equivalent to energy games [BFL+08]

Assuming $\nu = 0$

A memoryless strategy σ is winning if all cycles are nonnegative when σ is fixed.

Complexity

	Energy games	Mean-payoff games
Decision problem	$O(E \cdot Q \cdot W)$	$O(E \cdot Q \cdot W)$ (this talk) $O(E \cdot Q^2 \cdot W)$ [ZP96]

Deterministic Pseudo-polynomial algorithms

Outline

- Perfect information
  - Mean-payoff games
  - Energy games
  - Algorithms
- Imperfect information
  - Energy with fixed initial credit
  - Energy with unknown initial credit
  - Mean-payoff

Imperfect information

(staying alive in the dark)

Imperfect information – Why ?

System - Model

$\models$

Correctness relation

Specification

avoid failure,
ensure progress,
etc.

- Private variables/internal state
- Noisy sensors

⇒ Strategies should not rely on hidden information

Imperfect information – How ?

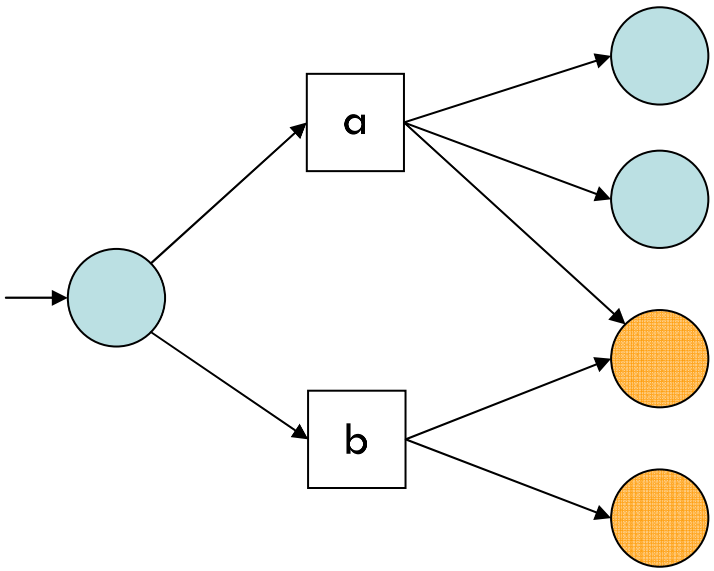

- Coloring of the state space

  observations = set of states with the same color

Imperfect information – How ?

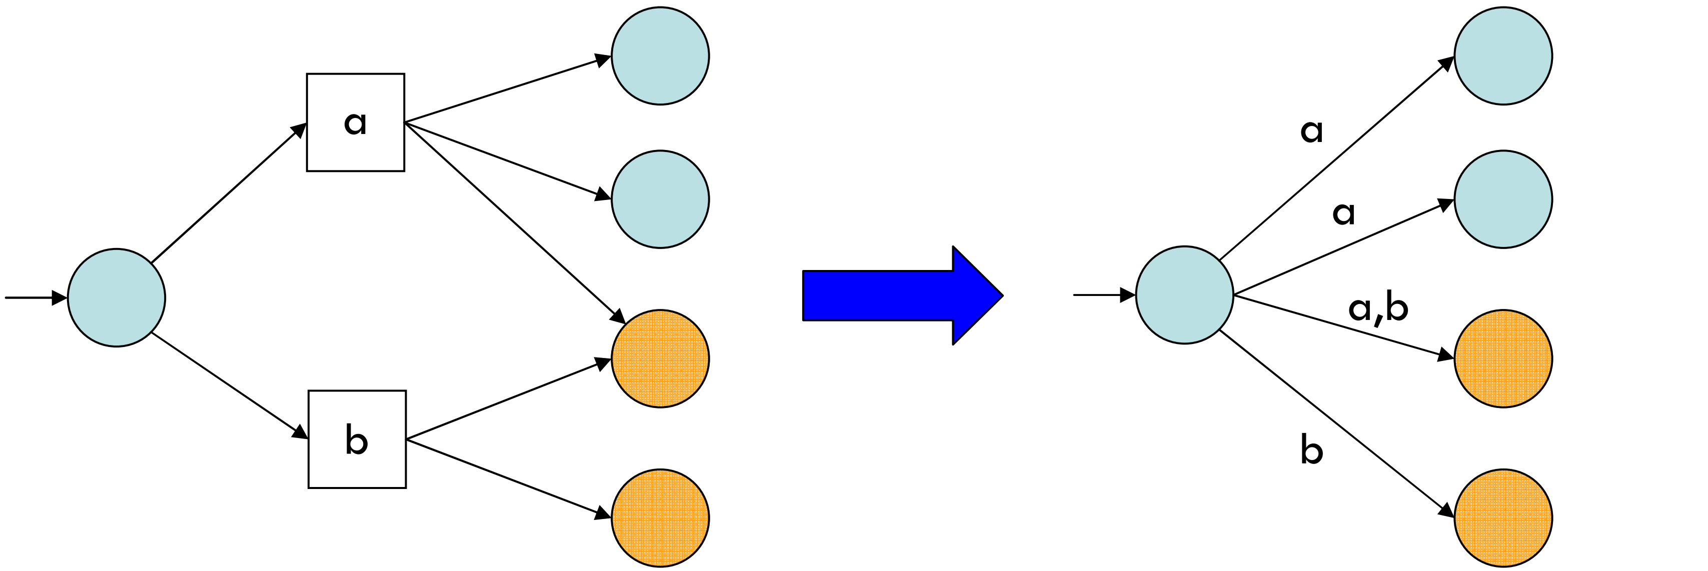

Maximizer states only

Playing the game:

1. Maximizer chooses an action (a or b)
2. Minimizer chooses successor state
   (compatible with Maximizer's action)
3. The color of the next state is visible to Maximizer

Imperfect information – How ?

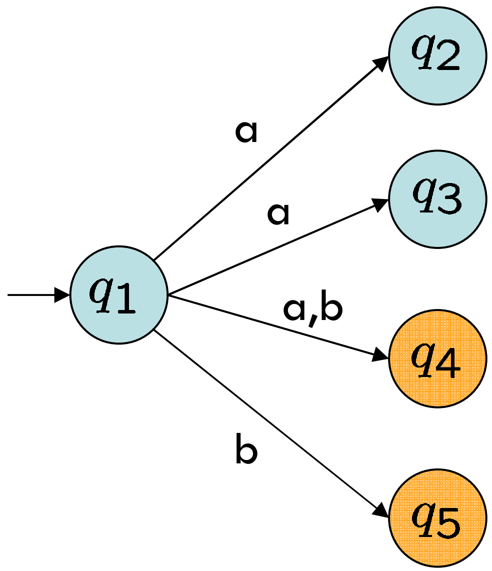

Actions $\Sigma = \{a, b\}$

Observations $\text{Obs} = \{\{q_1, q_2, q_3\}, \{q_4, q_5\}\}$

Imperfect information – How ?

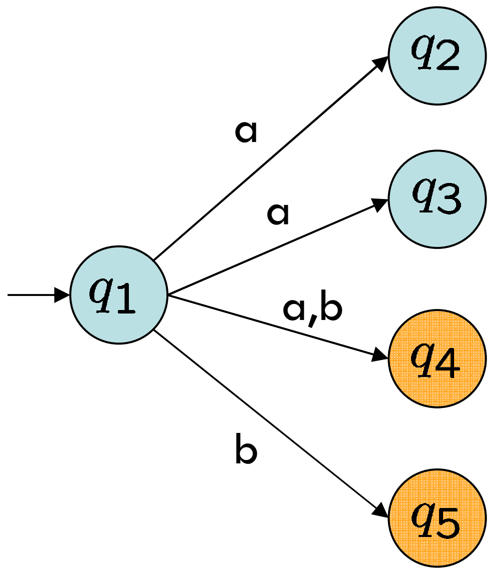

Observation-based strategies

$$\sigma : \mathsf{Obs}^+ \to \Sigma$$

Goal: all outcomes have

- nonnegative energy level,
- or nonnegative mean-payoff value

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

Observations $\mathsf{Obs} = \{\{q_1, q_2, q_3\}, \{q_4, q_5\}\}$

Complexity

	Energy games	Mean-payoff games
Perfect information	O(E·Q·W)	O(E·Q·W) (this talk) O(E·Q^2·W) [ZP96]
Imperfect information	?	?

Imperfect information

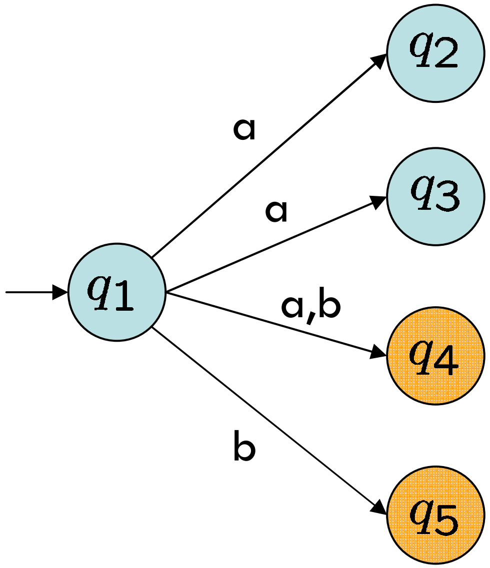

Observation-based strategies

$$\sigma : \mathsf{Obs}^+ \to \Sigma$$

Goal: all outcomes have

- nonnegative energy level,
- or nonnegative mean-payoff value

Two variants for Energy games:

- fixed initial credit
- unknown initial credit

Fixed initial credit

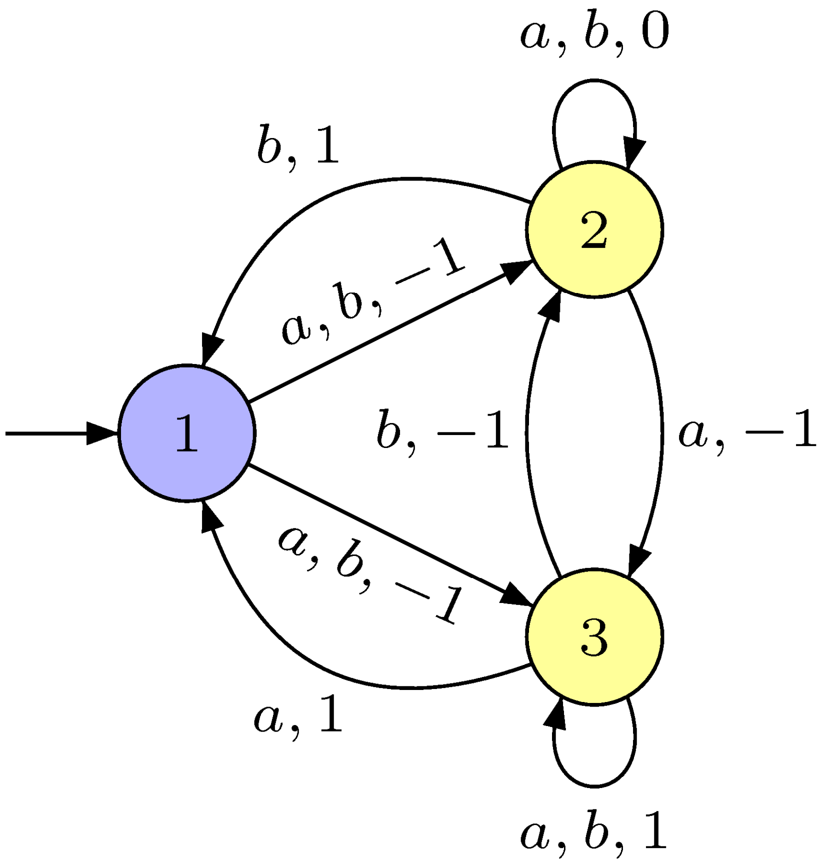

Can you win with initial credit = 3 ?

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

Observations $\mathsf{Obs} = \{\{1\}, \{2, 3\}\}$

Fixed initial credit

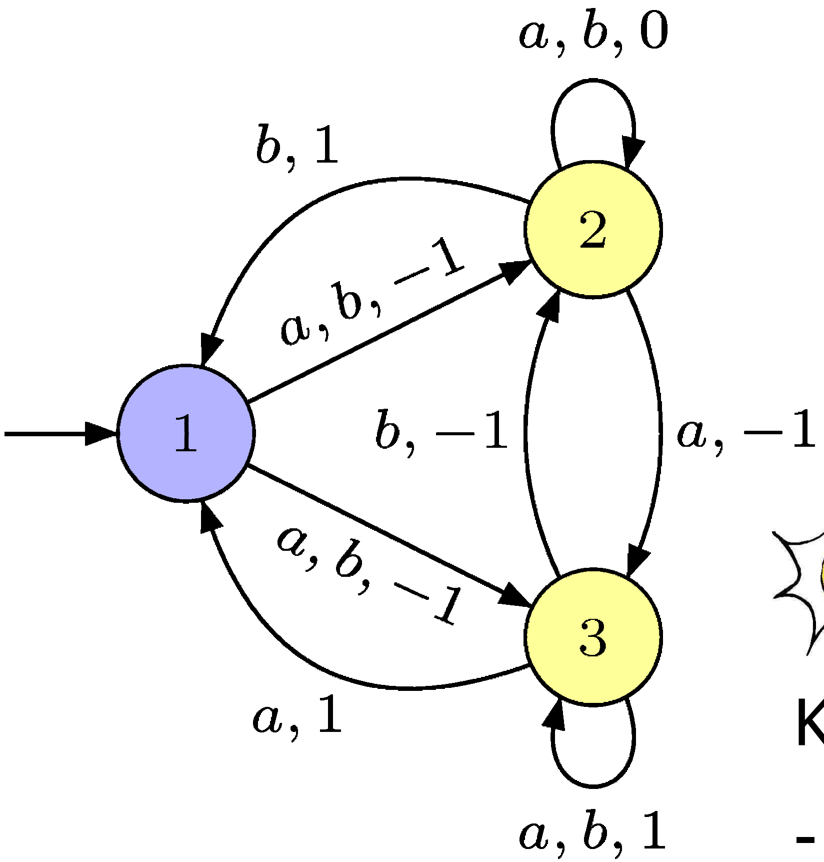

Can you win with initial credit = 3 ?

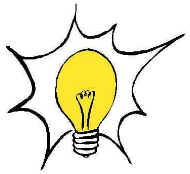

Keep track of

- which can be the current state, and
- what is the worst-case energy level

Initially: (3,⊥,⊥)

Example

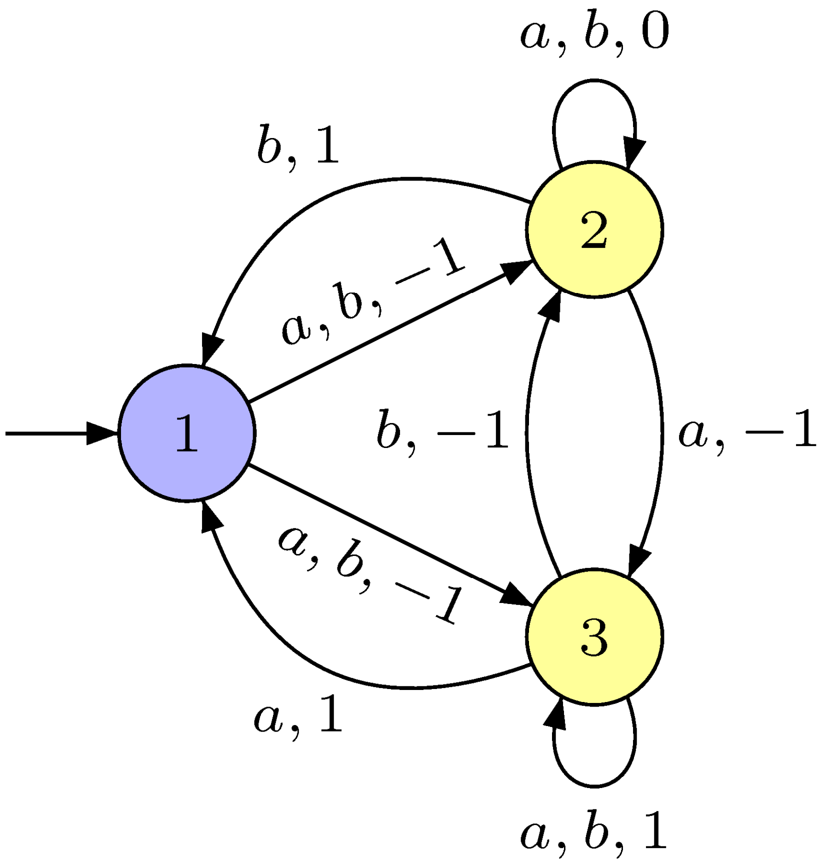

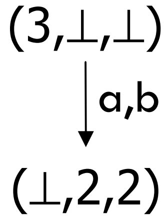

$$v(q) \leftarrow \min\{v(q') + w(q', q) \mid v(q') \neq \bot\}$$

Example

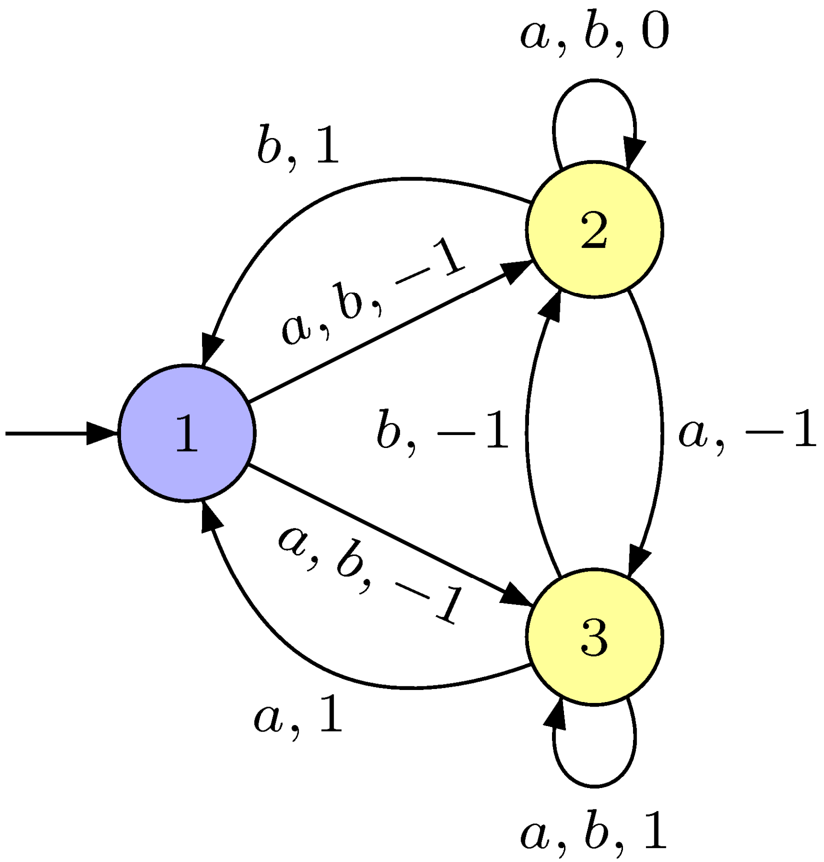

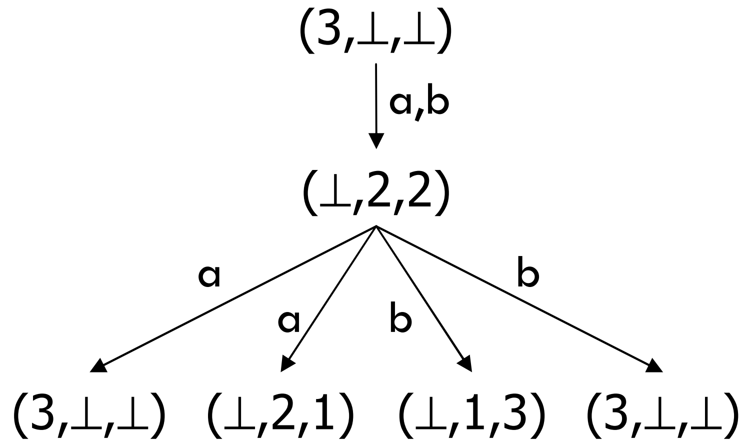

$$v(q) \leftarrow \min\{v(q') + w(q', q) \mid v(q') \neq \bot\}$$

Example

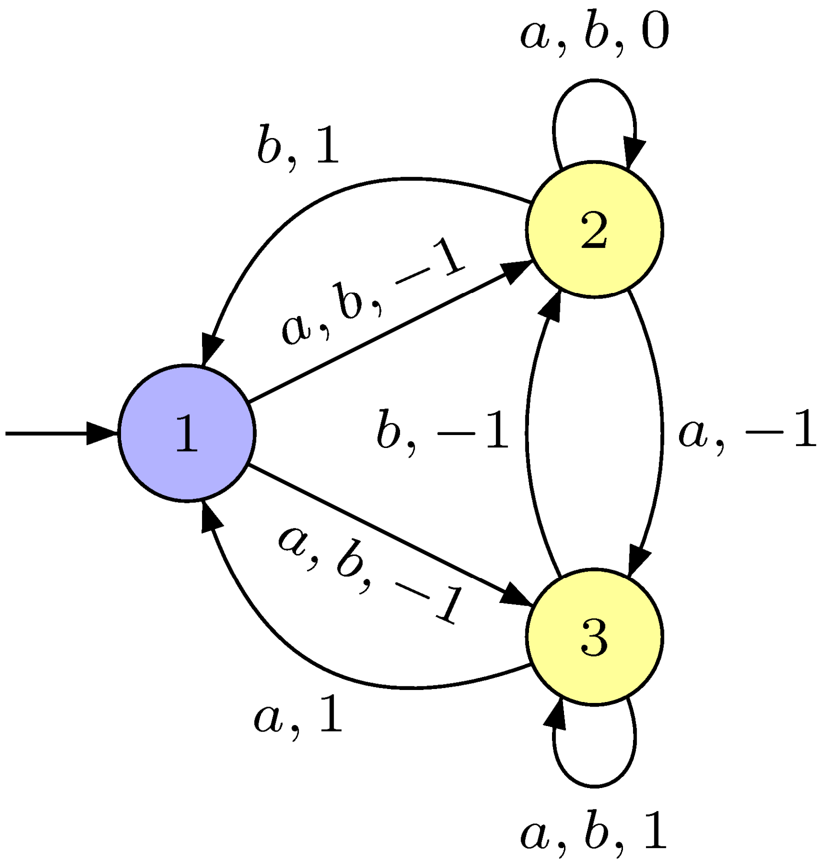

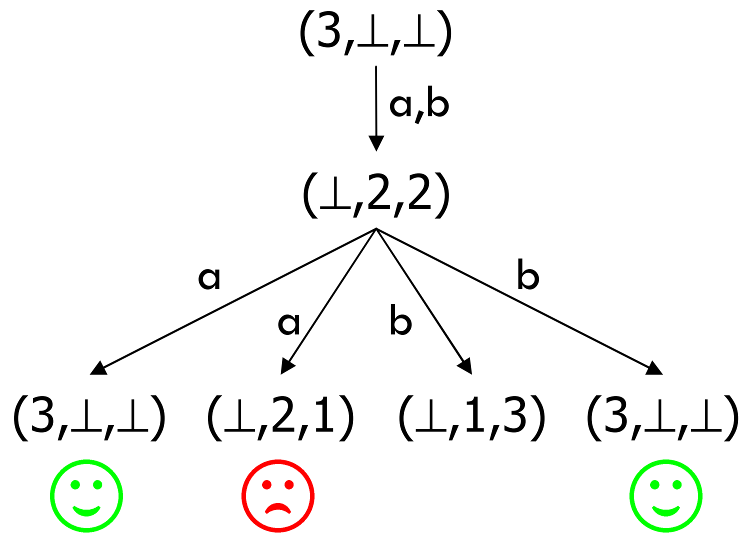

Stop search whenever

- negative value, or

- comparable ancestor

$$v(q) \leftarrow \min\{v(q') + w(q', q) \mid v(q') \neq \bot\}$$

Example

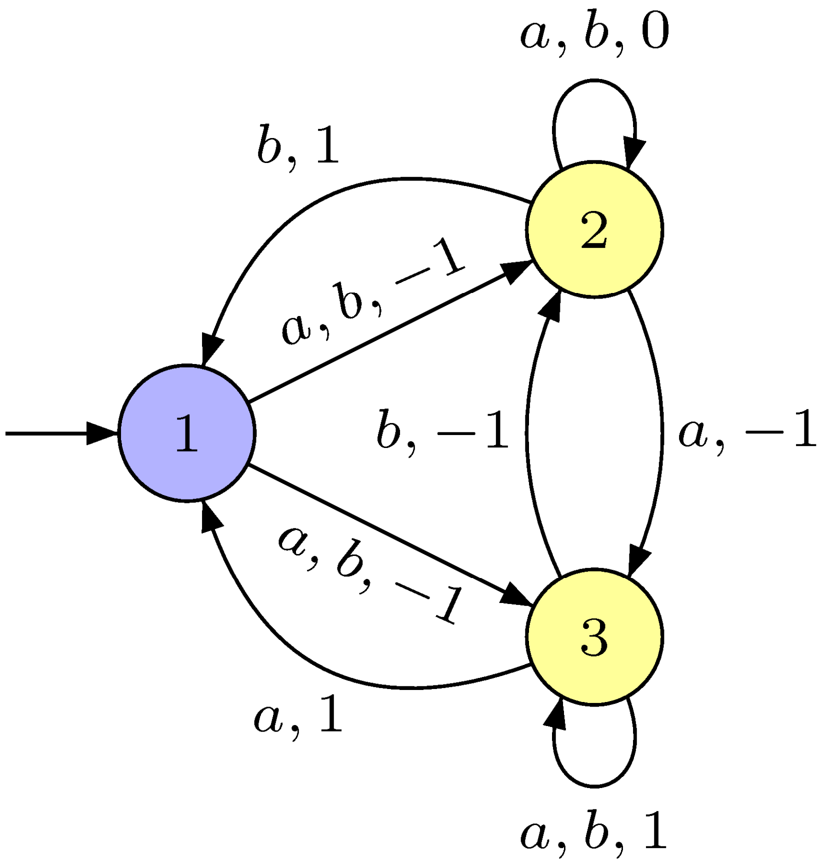

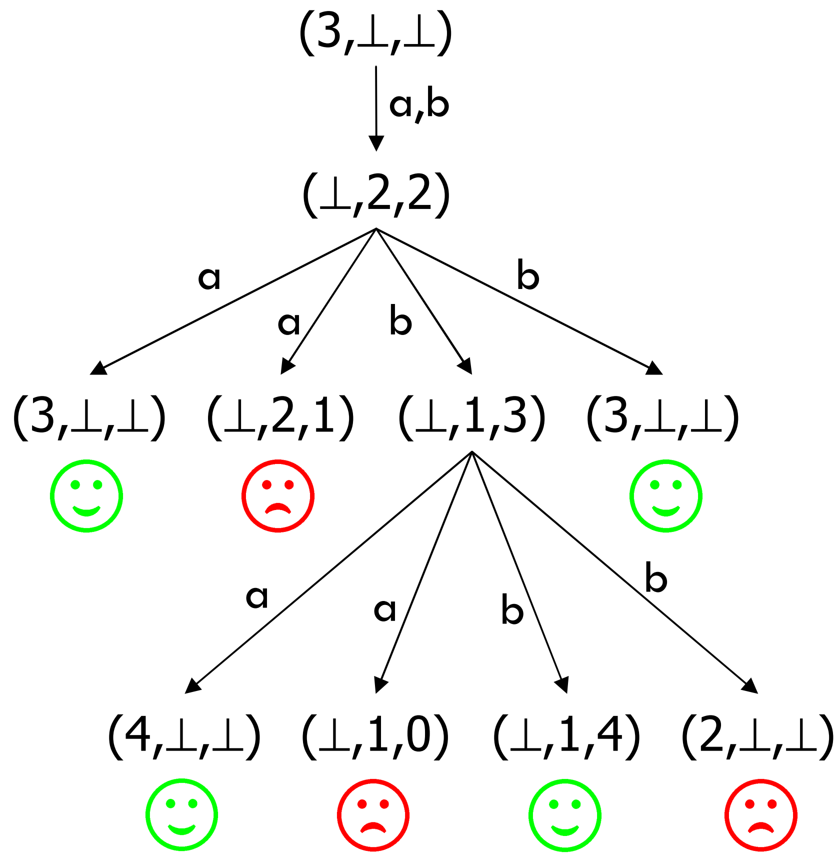

Stop search whenever:

- negative value, or

- comparable ancestor

Example

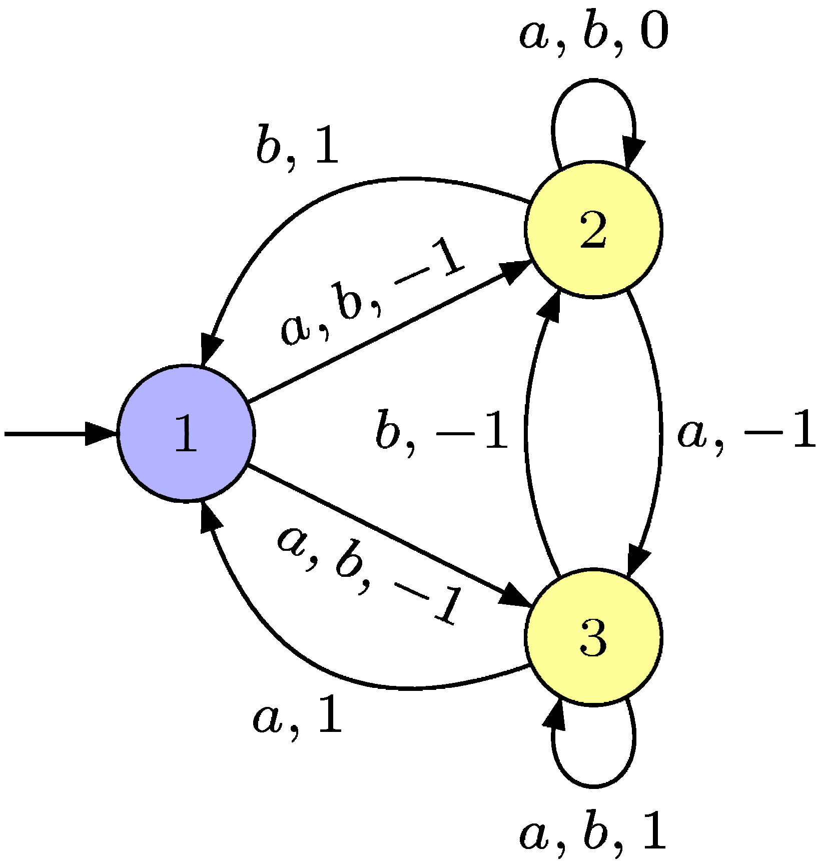

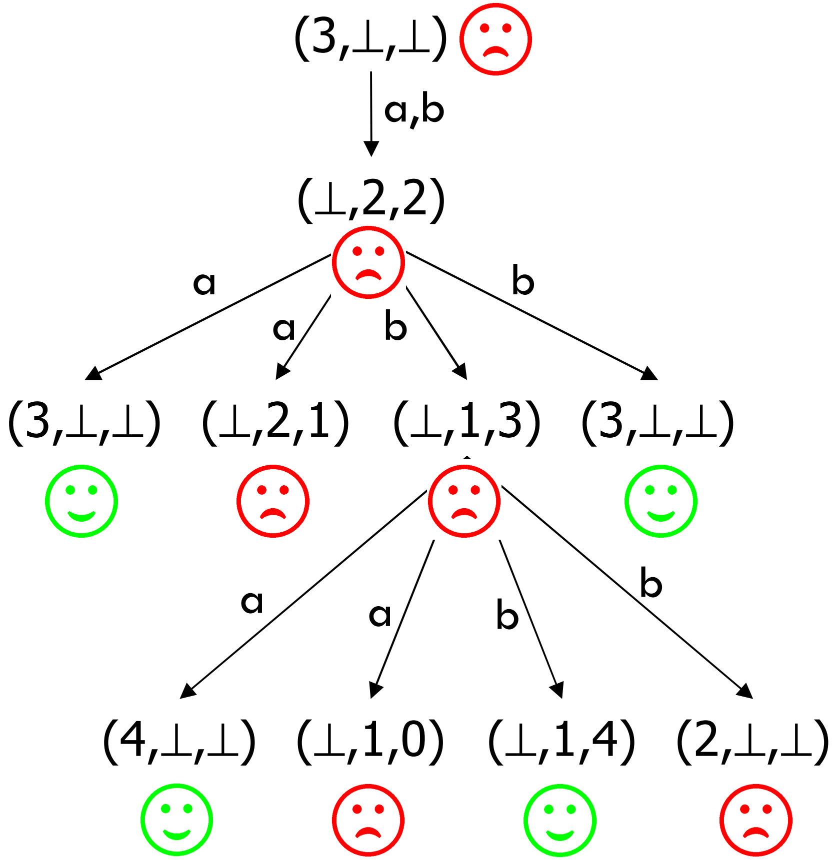

Initial credit = 3 is not sufficient !

Example

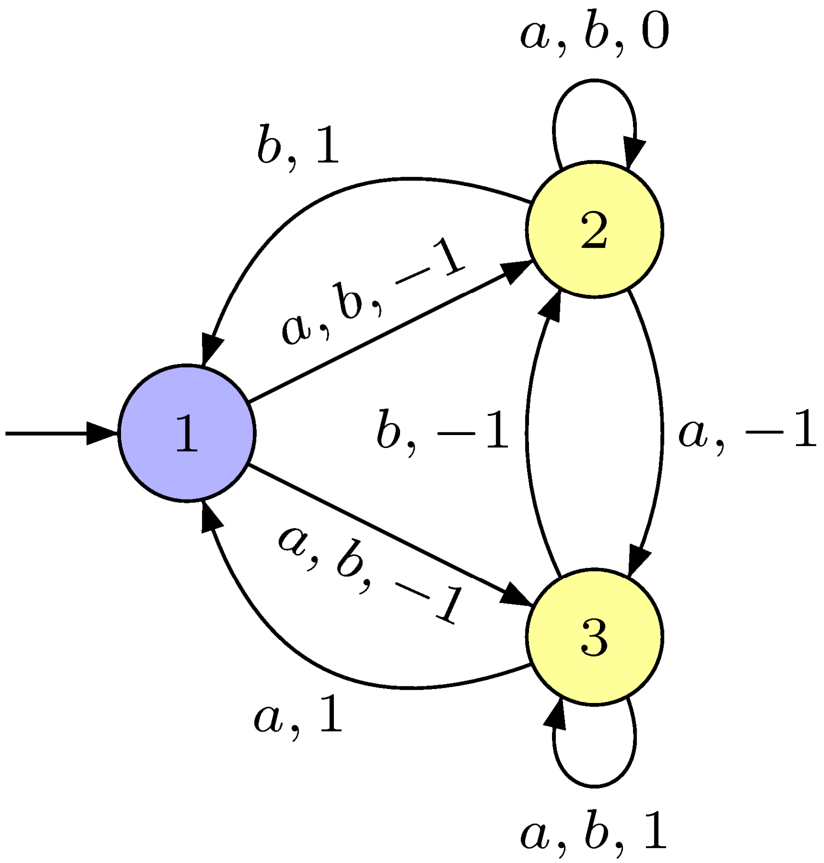

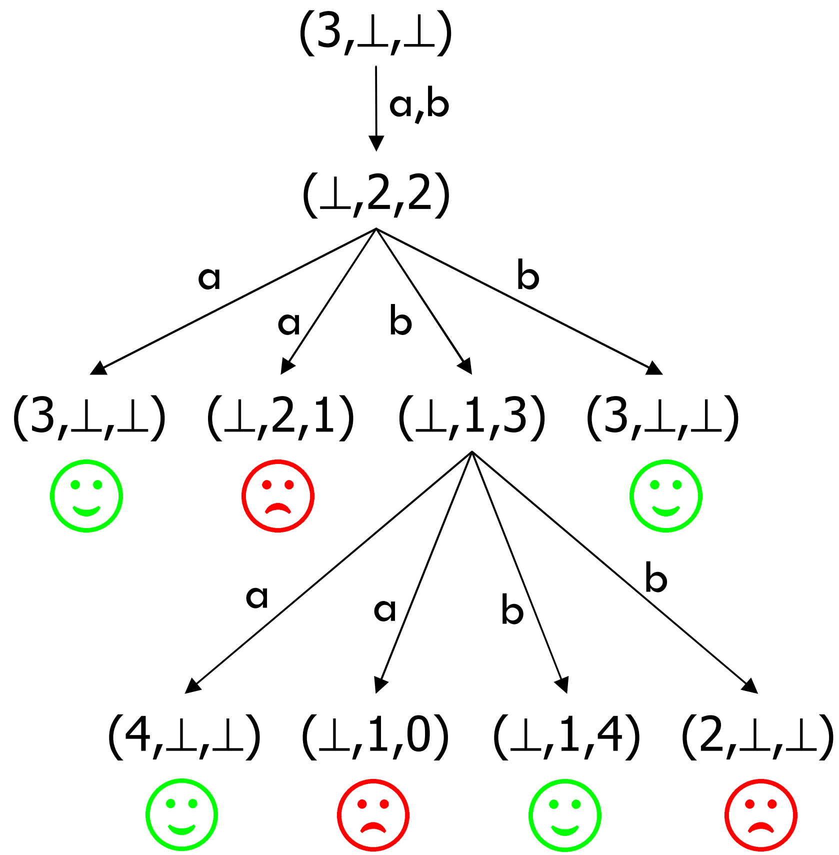

Search will terminate because $\mathbb{N}^d$ is well-quasi ordered.

Example

Upper bound: non-primitive recursive

Lower bound: EXPSPACE-hard

Proof (not shown in this talk): reduction from the infinite execution problem of Petri Nets.

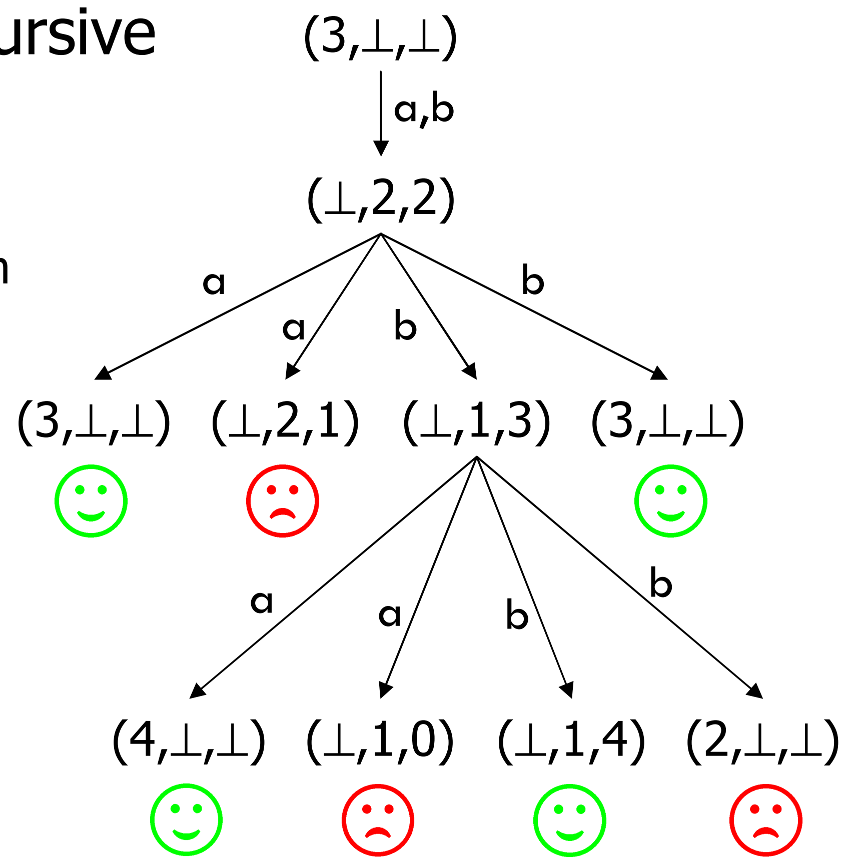

Search will terminate because $\mathbb{N}^d$ is well-quasi ordered.

Complexity

	Energy games (unknown initial credit)	Mean-payoff games
Perfect information	$O(E \cdot Q \cdot W)$	$O(E \cdot Q \cdot W)$ (this talk) $O(E \cdot Q^2 \cdot W)$ [ZP96]
Imperfect information	r.e.	?

Memory requirement

With imperfect information:

Corollary: Finite-memory strategies suffice in energy games

Memory requirement

With imperfect information:

Corollary: Finite-memory strategies suffice in energy games

In mean-payoff games:

- **infinite** memory may be required
- limsup vs. liminf definition do **not** coincide

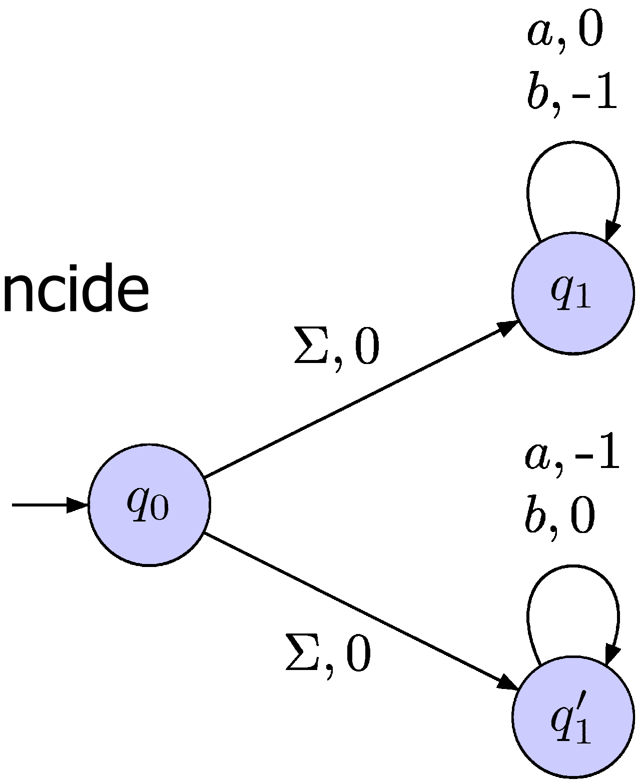

Memory requirement

	Energy games	Mean-payoff games
Perfect information	memoryless	memoryless
Imperfect information	finite memory	infinite memory

Unknown initial credit

Theorem

The unknown initial credit problem for energy games is undecidable.

(even for blind games)

Proof:

Using a reduction from the halting problem of 2-counter machines.

2-counter machines

- 2 counters c_1, c_2
- increment, decrement, zero test

q1: inc c_1 goto q2

q2: inc c_1 goto q3

q3: if c_1 == 0 goto q6
    else dec c_1 goto q4

q4: inc c_2 goto q5

q5: inc c_2 goto q3

q6: halt

2-counter machines

- 2 counters c_1, c_2
- increment, decrement, zero test

```
q1: inc c1 goto q2
q2: inc c1 goto q3
q3: if c1 == 0 goto q6
    else dec c1 goto q4
q4: inc c2 goto q5
q5: inc c2 goto q3
q6: halt
```

$$q : \text{ inc } c \textbf{ goto } q'$$

$$q : \textbf{ if } c = 0 \textbf{ then goto } q' \textbf{ else } \text{dec } c \textbf{ goto } q''.$$

$$(q, inc, c, q')$$

$$(q, 0?, c, q') \text{ and } (q, dec, c, q'')$$

$$\begin{pmatrix} q_1 \\ 0 \\ 0 \end{pmatrix} \xrightarrow{(q_1, inc, c_1, q_2)} \begin{pmatrix} q_2 \\ 1 \\ 0 \end{pmatrix} \xrightarrow{(q_2, inc, c_1, q_3)} \begin{pmatrix} q_3 \\ 2 \\ 0 \end{pmatrix} \xrightarrow{(q_3, dec, c_1, q_4)} \begin{pmatrix} q_4 \\ 1 \\ 0 \end{pmatrix} \xrightarrow{(q_4, inc, c_2, q_5)} \begin{pmatrix} q_5 \\ 1 \\ 1 \end{pmatrix} \xrightarrow{\cdots}$$

Reduction

Halting problem:

Given M and state q_{halt}, decide if q_{halt} is reachable (i.e., M halts).

q1: inc c_1 goto q2

q2: inc c_1 goto q3

q3: if c_1 == 0 goto q6
else dec c_1 goto q4

q4: inc c_2 goto q5

q5: inc c_2 goto q3

q6: halt

Reduction:

Given M, construct G_M such that M halts iff there exists a winning strategy in G_M (with some initial credit).

!

- Deterministic machine
- Nonnegative counters

Reduction

q1: inc c_1 goto q2

q2: inc c_1 goto q3

q3: if c_1 == 0 goto q6
  else dec c_1 goto q4

q4: inc c_2 goto q5

q5: inc c_2 goto q3

q6: halt

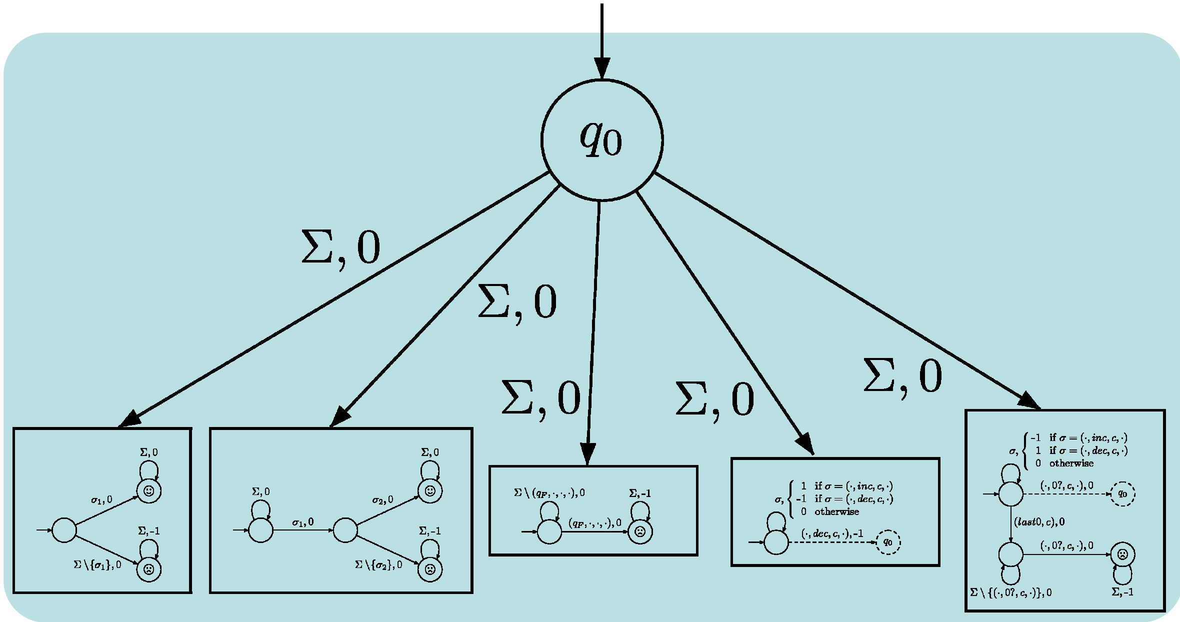

- Blind game (unique observation)
- Initial nondeterministic jump to several gadgets
- Winning strategy = (#AcceptingRun)$^{\omega}$

Gadgets

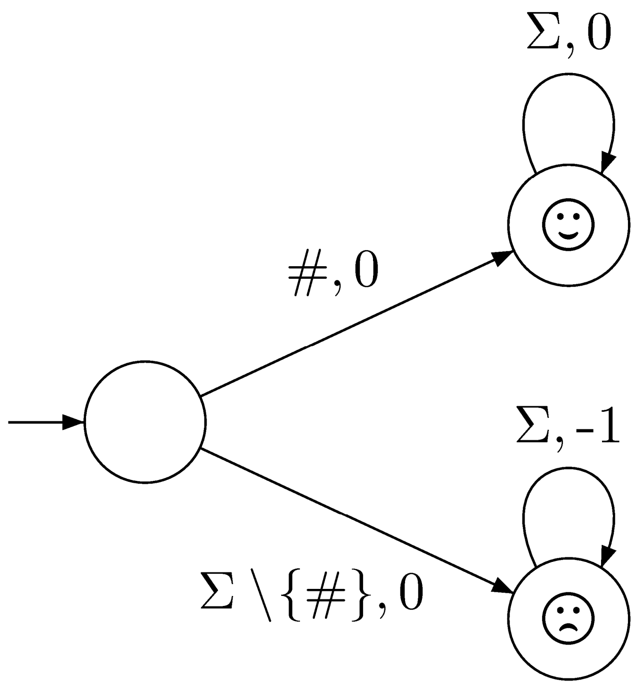

Gadget 1:

« First symbol is # »

Reminder: Winning strategy = #AcceptingRun#AcceptingRun#...

Gadgets

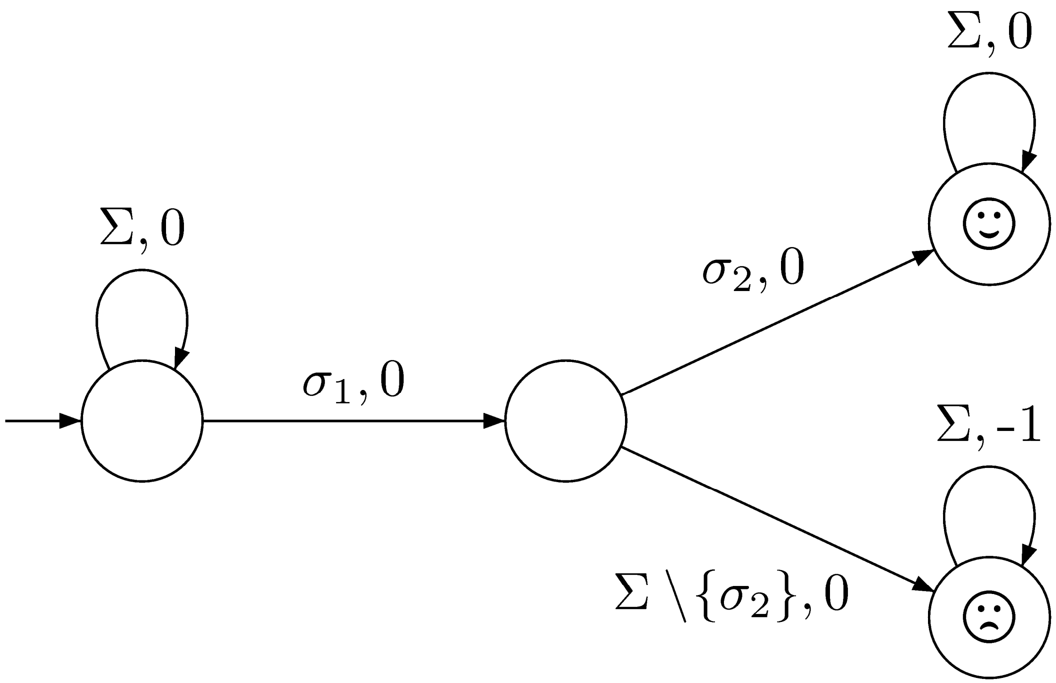

Gadget 2:

« Every σ_1 is followed by σ_2 »

E.g., $\sigma_1 = (q, \cdot, \cdot, q')$ and $\sigma_2 = (q', \cdot, \cdot, q'')$

Reminder: Winning strategy = #AcceptingRun#AcceptingRun#...

Gadgets

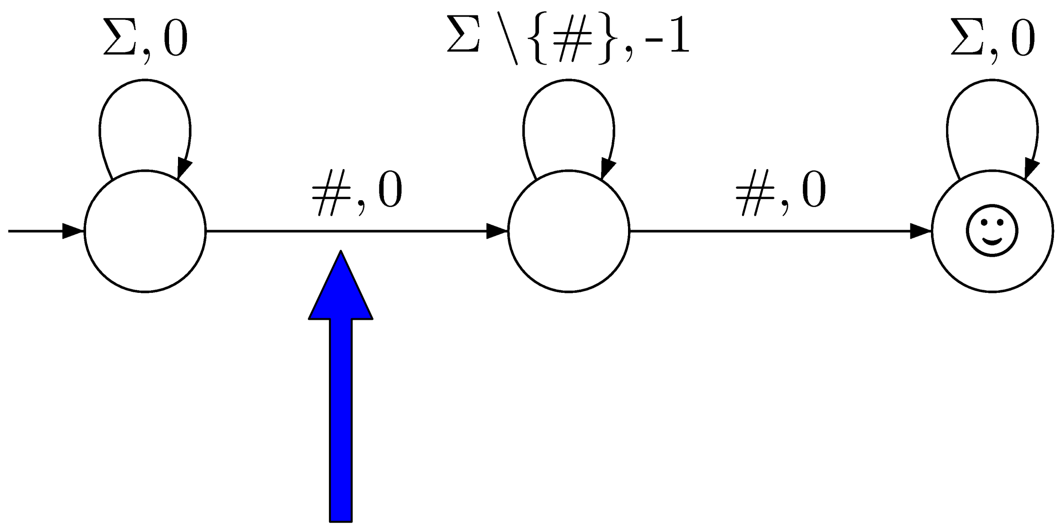

Gadget 3:

« Infinitely many # »

(and a bit more…)

Guess: this is the last #

Reminder: Winning strategy = #AcceptingRun#AcceptingRun#...

Gadgets

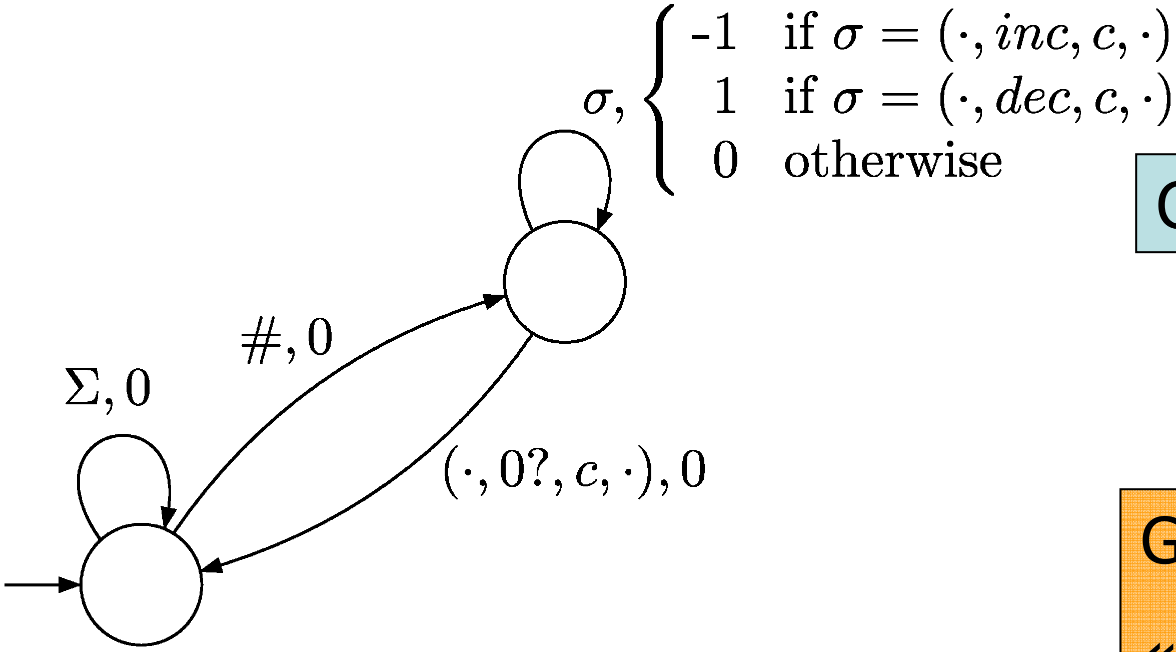

Check zero tests on c

Gadget 4:

« Counter correctness »

Gadgets

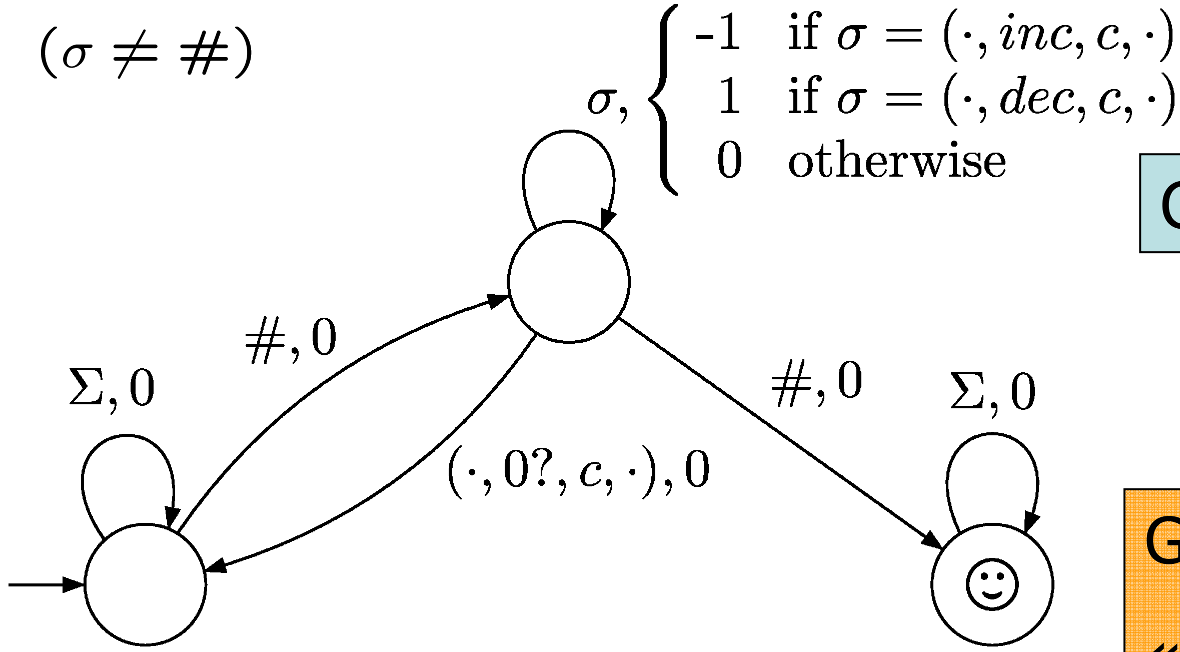

Check zero tests on c

Gadget 4:

« Counter correctness »

Gadgets

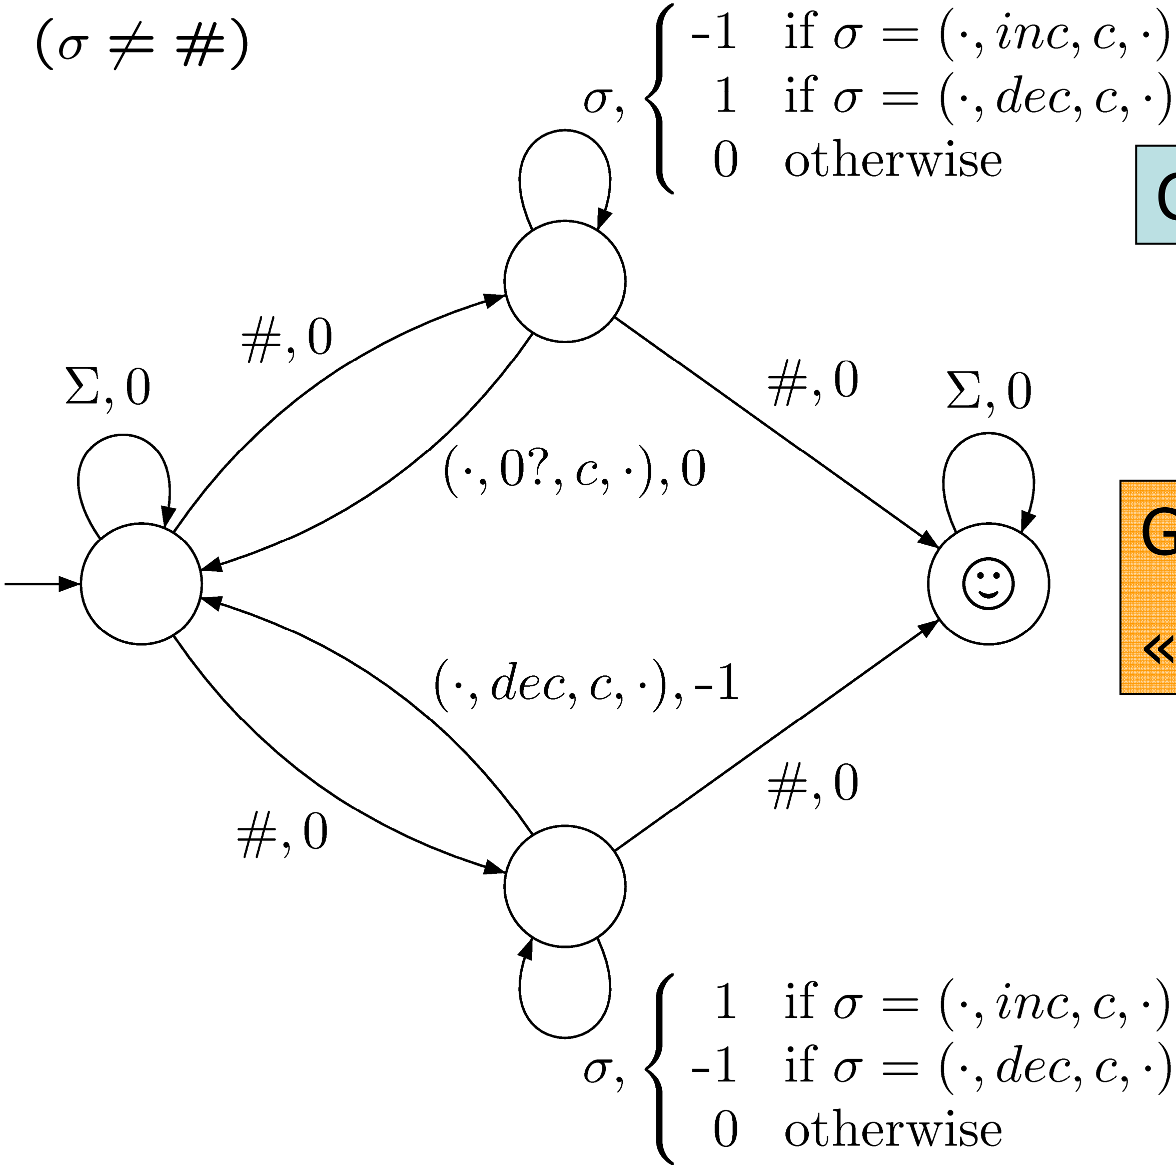

Check zero tests on c

Gadget 4:

« Counter correctness »

Check non-zero test on c

Correctness

q1: inc c_1 goto q2

q2: inc c_1 goto q3

q3: if c_1 == 0 goto q6
      else dec c_1 goto q4

q4: inc c_2 goto q5

q5: inc c_2 goto q3

q6: halt

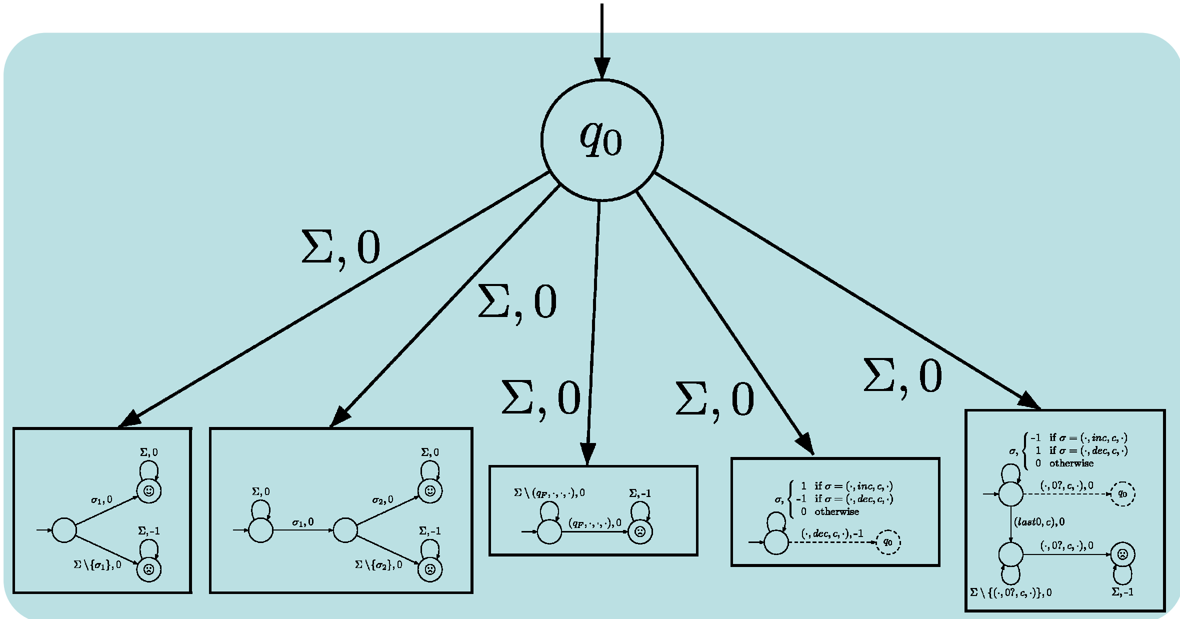

- If M halts, then (#AcceptingRun)$^{\omega}$ is a winning strategy with initial credit Length(AcceptingRun).
- If there exists a winning strategy with finite initial credit, then # occurs infinitely often, and finitely many cheats occur. Hence, M has an accepting run.

Mean-payoff games

Theorem

Mean-payoff games are undecidable (not co-r.e.).

(even blind games)

Proof:

Using a reduction from the halting problem of 2-counter machines.

Nota: the proof works for both limsup and liminf, but only for strict mean-payoff objective (i.e., MP > ν)

Mean-payoff games

Theorem

Mean-payoff games are undecidable (not co-r.e.).

(even blind games)

Proof:

Using a reduction from the halting problem of 2-counter machines.

Reduction:

Given M, construct G_M such that M halts iff there exists a strategy to ensure strictly positive mean-payoff value.

Gadgets

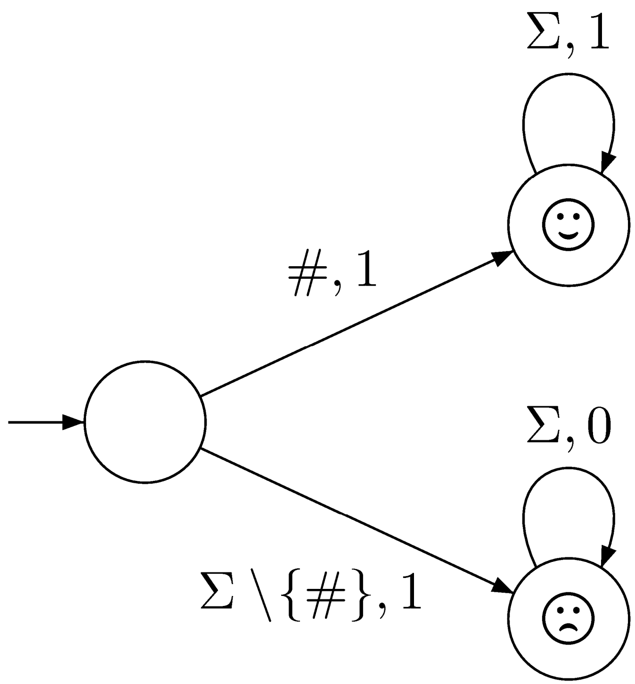

Gadget 1:

« First symbol is # »

Reminder: Winning strategy = #AcceptingRun#AcceptingRun#...

Gadgets

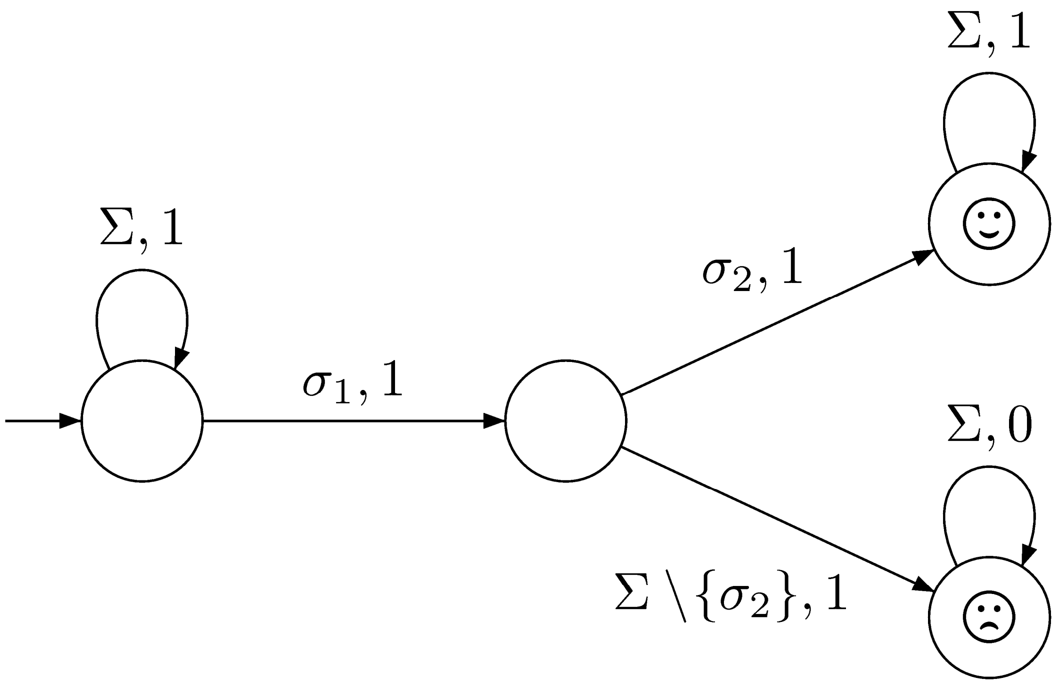

Gadget 2:

« Every σ_1 is followed by σ_2 »

E.g., $\sigma_1 = (q, \cdot, \cdot, q')$ and $\sigma_2 = (q', \cdot, \cdot, q'')$

Reminder: Winning strategy = #AcceptingRun#AcceptingRun#...

Gadgets

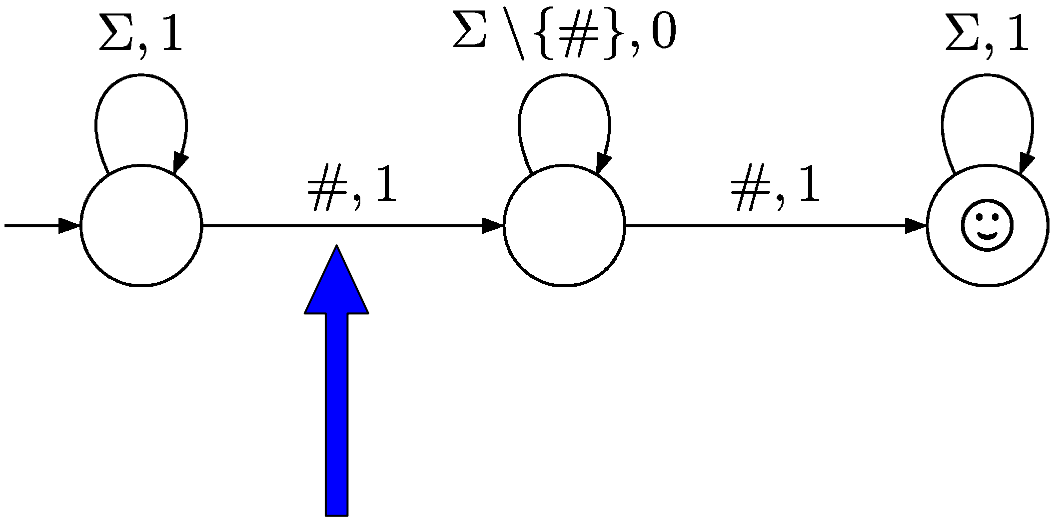

Guess: this is the last #

Gadget 3:

« Infinitely many # »

(and a bit more…)

Reminder: Winning strategy = #AcceptingRun#AcceptingRun#...

Gadgets

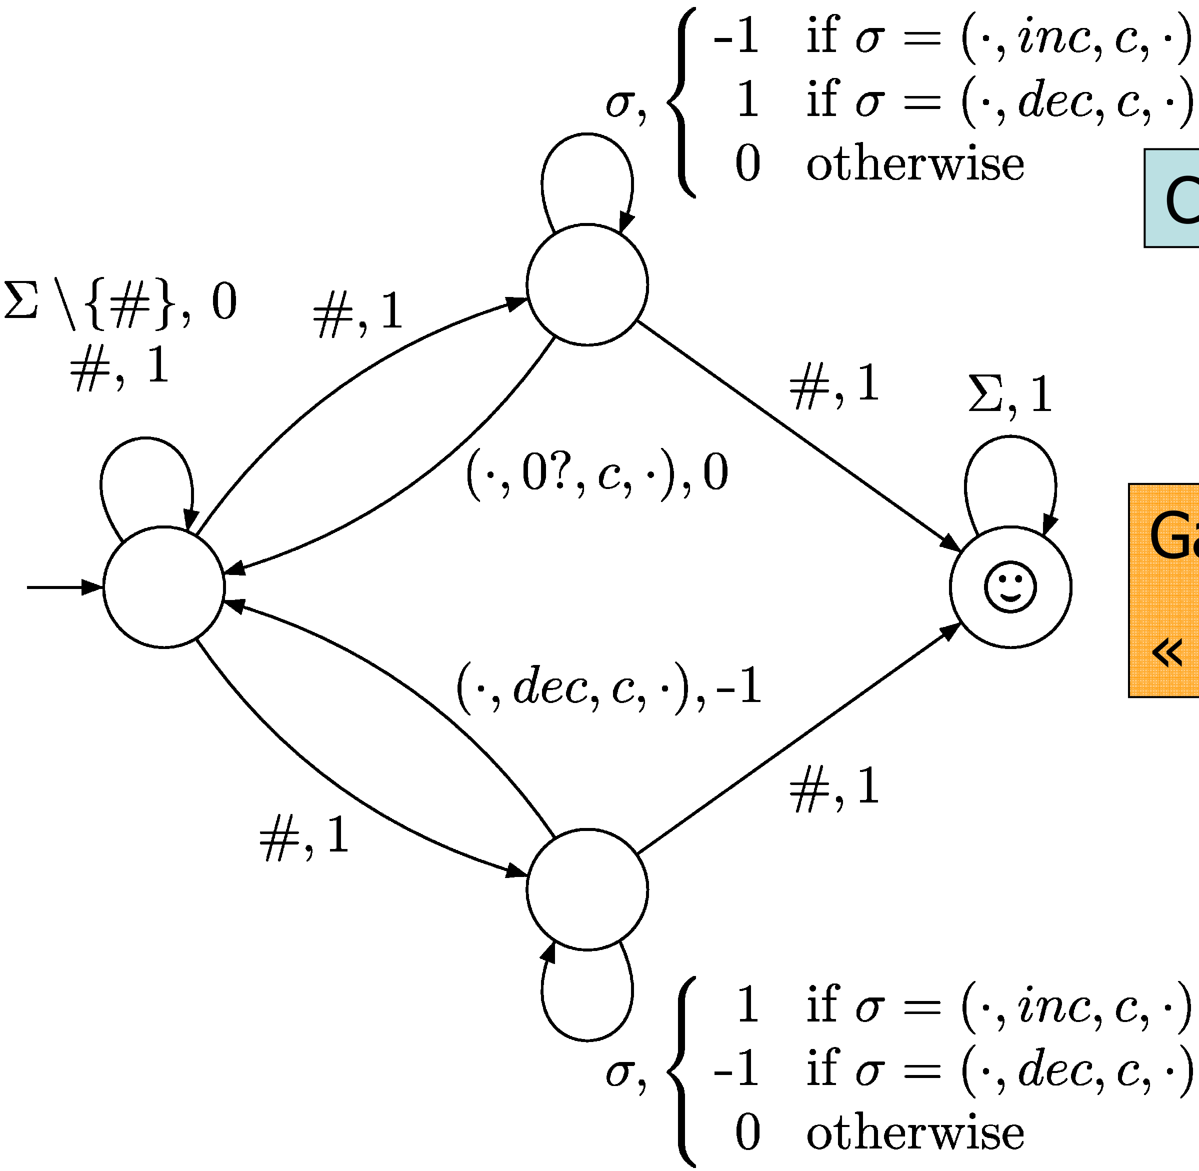

Check zero tests on c

Gadget 4:

« Counter correctness »

Check non-zero test on c

Complexity

	Energy games (unknown initial credit)	Mean-payoff games
Perfect information	$O(E \cdot Q \cdot W)$	$O(E \cdot Q \cdot W)$ (this talk) $O(E \cdot Q^2 \cdot W)$ [ZP96]
Imperfect information	r.e. not co-r.e.	? not co-r.e.

Mean-payoff games

Theorem

Mean-payoff games are undecidable (not r.e.).

(for games with at least 2 observations)

Proof:

Using a reduction from the non-halting problem of 2-counter machines.

Nota: the proof works only for limsup and non-strict mean-payoff objective (i.e., MP $\geq \nu$)

Mean-payoff games

Theorem

Mean-payoff games are undecidable (not r.e.).

(for games with at least 2 observations)

Proof:

Using a reduction from the non-halting problem of 2-counter machines.

Reduction:

Given M, construct G_M such that M **does not halt** iff there exists a strategy to ensure strictly nonnegative mean-payoff value.

Reduction

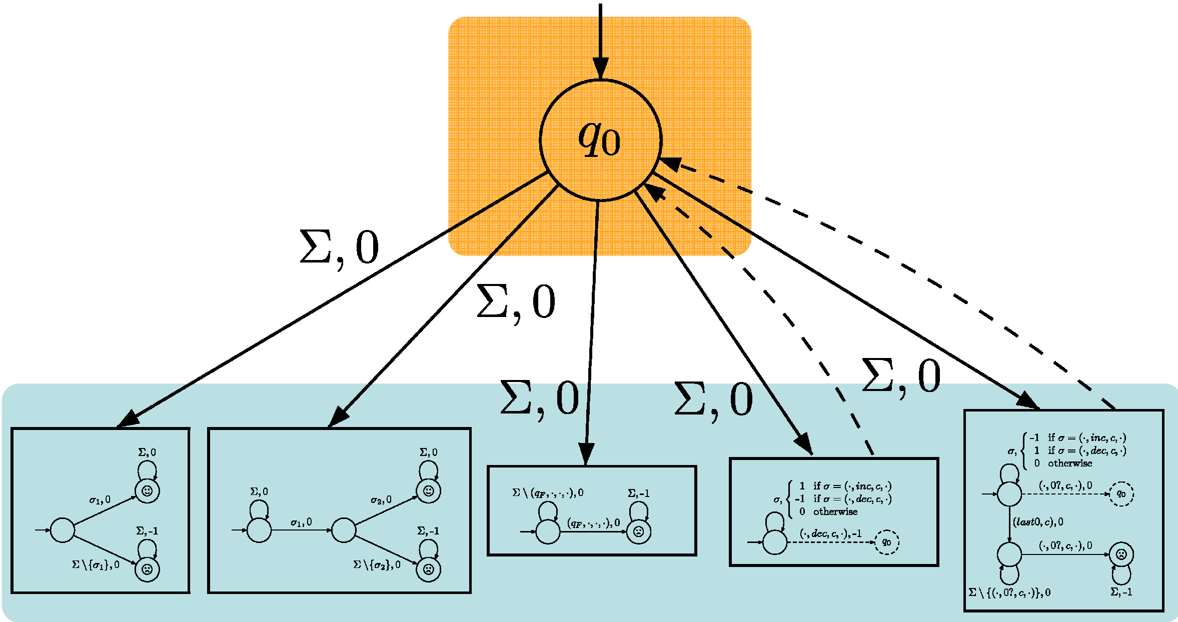

- 2-observation game
- Initial nondeterministic jump to several gadgets (+ back-edges)
- Winning strategy = Non-terminatingRun

Gadgets

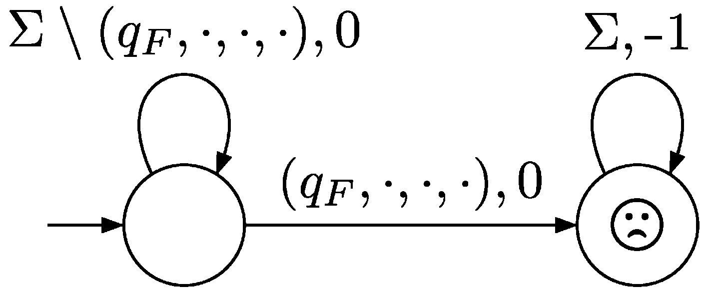

Gadget 3:

« avoid halting state »

Reminder: Winning strategy = Non-terminatingRun

Gadgets

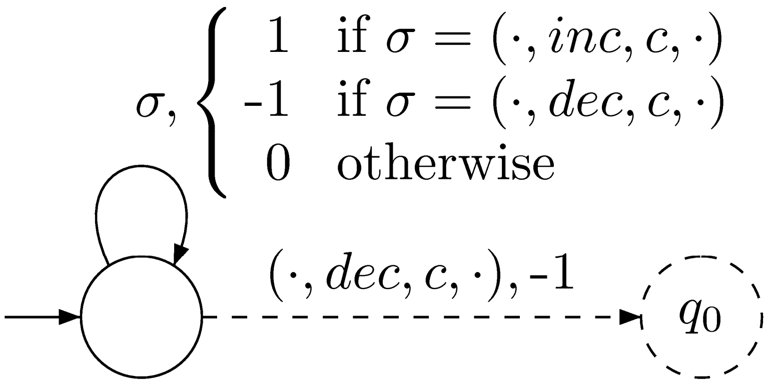

Check non-zero test on c

Gadget 5 and 6:

« Counter correctness »

Gadgets

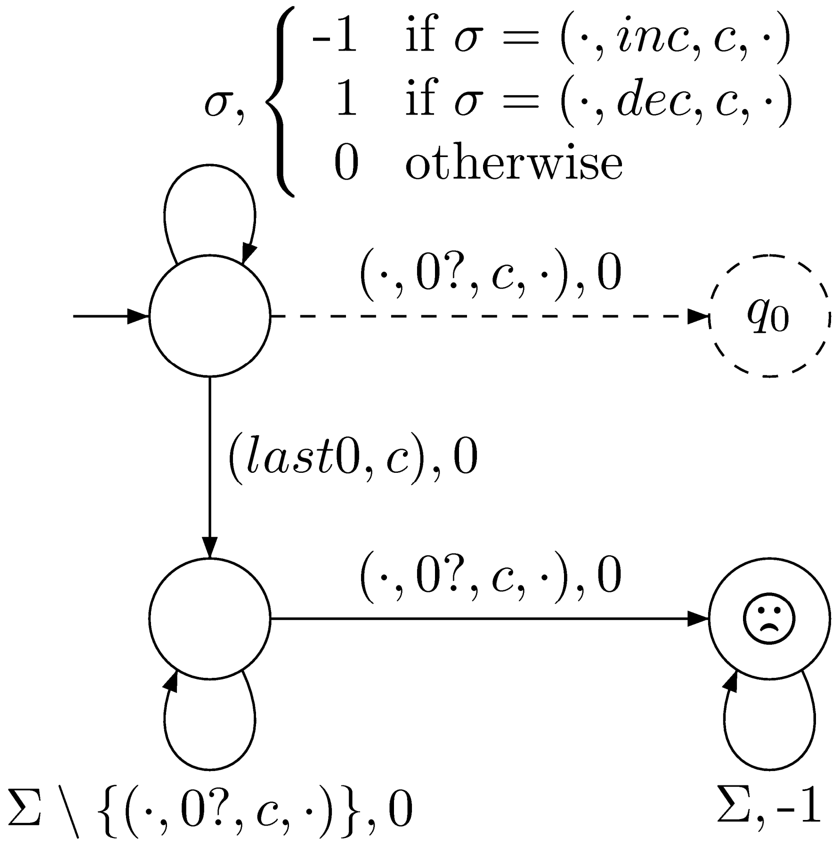

Gadget 5 and 6:

« Counter correctness »

Check zero tests on c

Correctness

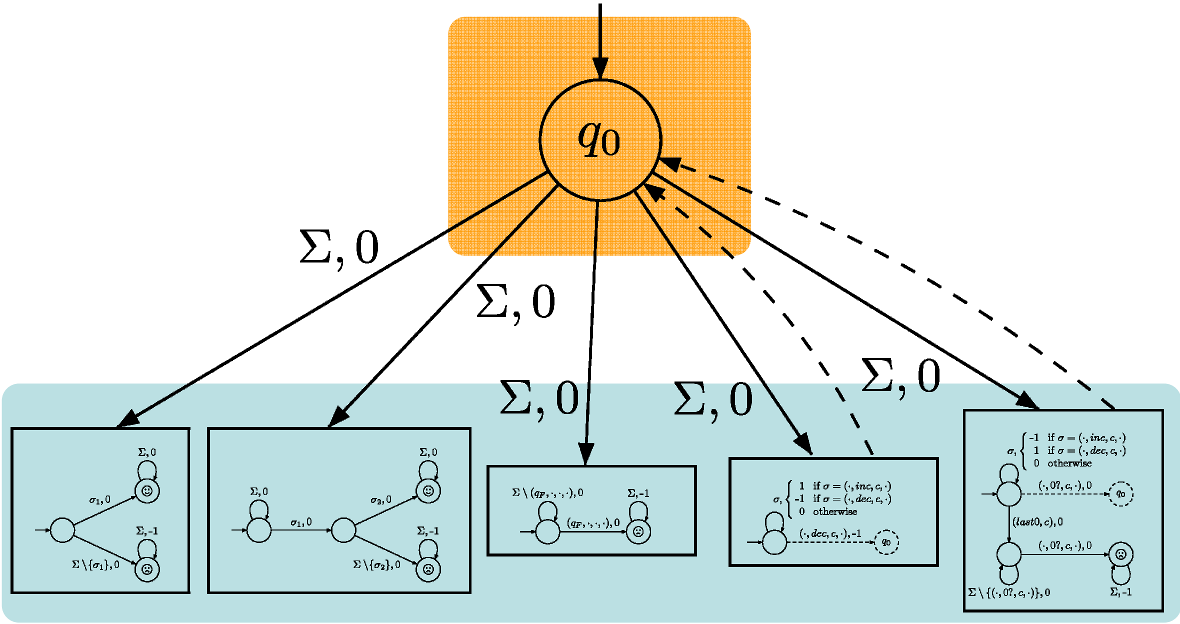

- If M does not halt, then Non-terminatingRun is a winning strategy.

- If M halts, then Maximizer has to cheat within L steps where L = Size(AcceptingRun), or reaches halting state, thus he ensures mean-payoff at most -1/L.

Complexity

	Energy games (unknown initial credit)	Mean-payoff games
Perfect information	$O(E \cdot Q \cdot W)$	$O(E \cdot Q \cdot W)$ (this talk) $O(E \cdot Q^2 \cdot W)$ [ZP96]
Imperfect information	r.e. not co-r.e.	not r.e. not co-r.e.

Nota: whether there exists a finite-memory winning strategy in mean-payoff games is also undecidable.

Decidability result

Energy and mean-payoff games with **visible weights** are decidable (EXPTIME-complete).

Weights are **visible** if

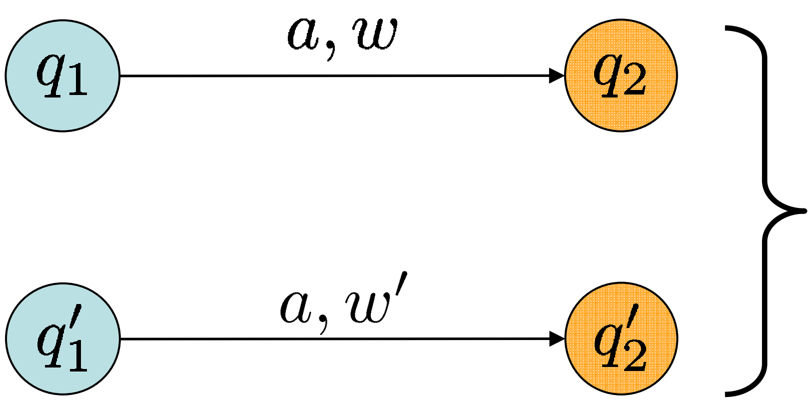

implies $w = w'$

Weighted subset construction is finite

Complexity

	Energy games (unknown initial credit)	Mean-payoff games
Perfect information	O(E·Q·W)	O(E·Q·W) (this talk) O(E·Q^2·W) [ZP96]
Imperfect information	r.e. not co-r.e.	not r.e. not co-r.e.
Visible weights	EXPTIME-complete	EXPTIME-complete

Conclusion

- Quantitative games with imperfect information
  - Undecidable in general
  - Energy with fixed initial credit decidable
  - Visible weights decidable
- Open questions
  - Strict vs. non-strict mean-payoff
  - Liminf vs. Limsup
  - Blind mean-payoff games
- Related work
  - Incorporate liveness conditions (e.g. parity)

The end

Thank you !

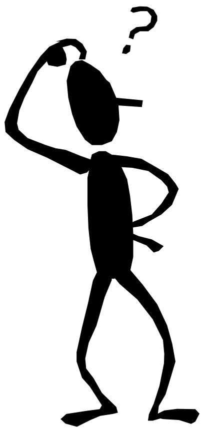

Questions ?

References

[CdAHS03] A.Chakrabarti, L. de Alfaro, T.A. Henzinger, and M. Stoelinga. **Resource interfaces**, Proc. of EMSOFT: Embedded Software, LNCS 2855, Springer, pp.117-133, 2003

[EM79] A. Ehrenfeucht, and J. Mycielski, **Positional Strategies for Mean-Payoff Games**, International Journal of Game Theory, vol. 8, pp. 109-113, 1979

[BFL+08] P. Bouyer, U. Fahrenberg, K.G. Larsen, N. Markey, and J. Srba, **Infinite Runs in Weighted Timed Automata with Energy Constraints**, Proc. of FORMATS: Formal Modeling and Analysis of Timed Systems, LNCS 5215, Springer, pp. 33-47, 2008