UNIVERSITÉ DE VERSAILLES ST-QUENTIN-EN-YVELINES

Laboratoire d'informatique
CNRS - UMR 8144
PRISM

David Auger
Université de Versailles Saint-Quentin-en-Yvelines

Simple Stochastic Games : a state of the art

jeudi 21 Mars 2013, LIRMM

A Simple Stochastic Game (Condon 1989) is defined by a directed graph with :

- three sets of vertices V_{MAX}, V_{MIN}, V_{AVE}, all of which have outdegree 2
- two 'sink' vertices 0 and 1
- a start vertex

2 1/2 players : MAX and MIN, and a 'chance' player

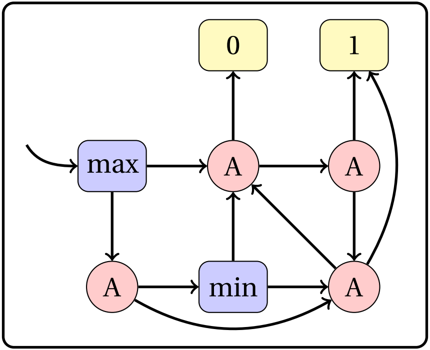

- player MAX wants to reach the 1 sink
- player MIN wants to prevent him from doing so

A play consists in moving a pebble on the graph :

- on a MAX (resp. MIN) node player MAX (resp. MIN) decides where to go next ;
- on a AVE node the next vertex is randomly determined (simple coin toss)

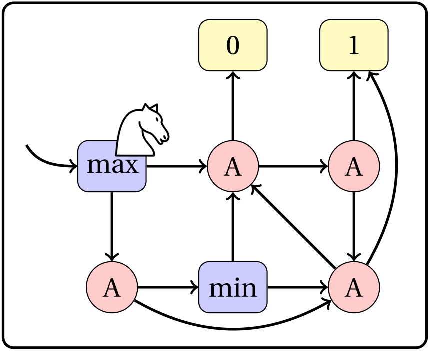

- player MAX wants to reach the 1 sink
- player MIN wants to prevent him from doing so

A play consists in moving a pebble on the graph :

- on a MAX (resp. MIN) node player MAX (resp. MIN) decides where to go next ;
- on a AVE node the next vertex is randomly determined (simple coin toss)

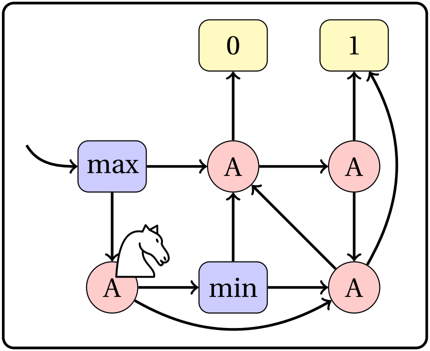

- player MAX wants to reach the 1 sink
- player MIN wants to prevent him from doing so

A play consists in moving a pebble on the graph :

- on a MAX (resp. MIN) node player MAX (resp. MIN) decides where to go next ;
- on a AVE node the next vertex is randomly determined (simple coin toss)

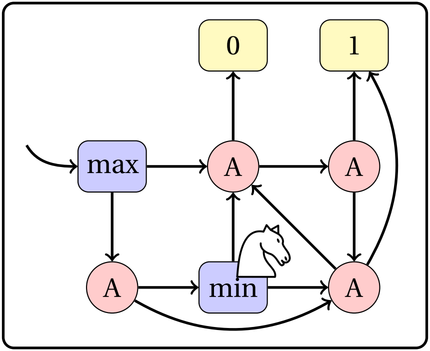

- player MAX wants to reach the 1 sink
- player MIN wants to prevent him from doing so

A play consists in moving a pebble on the graph :

- on a MAX (resp. MIN) node player MAX (resp. MIN) decides where to go next ;
- on a AVE node the next vertex is randomly determined (simple coin toss)

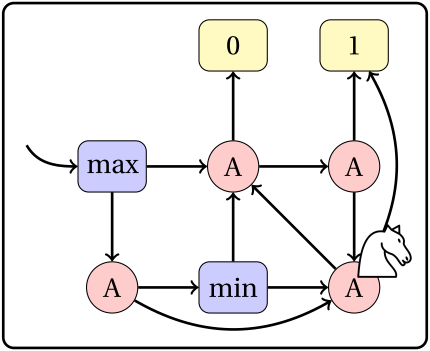

- player MAX wants to reach the 1 sink
- player MIN wants to prevent him from doing so

A play consists in moving a pebble on the graph :

- on a MAX (resp. MIN) node player MAX (resp. MIN) decides where to go next ;
- on a AVE node the next vertex is randomly determined (simple coin toss)

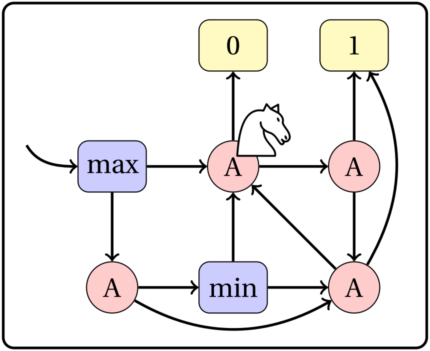

- player MAX wants to reach the 1 sink
- player MIN wants to prevent him from doing so

A play consists in moving a pebble on the graph :

- on a MAX (resp. MIN) node player MAX (resp. MIN) decides where to go next ;
- on a AVE node the next vertex is randomly determined (simple coin toss)

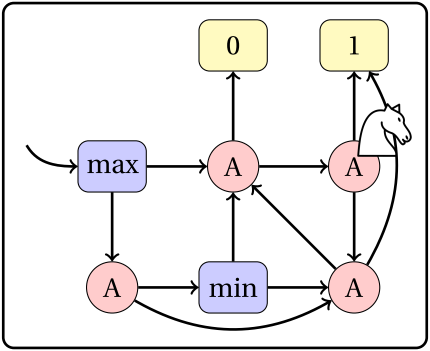

- player MAX wants to reach the 1 sink
- player MIN wants to prevent him from doing so

A play consists in moving a pebble on the graph :

- on a MAX (resp. MIN) node player MAX (resp. MIN) decides where to go next ;
- on a AVE node the next vertex is randomly determined (simple coin toss)

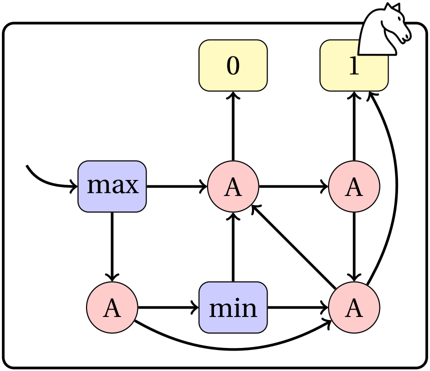

- player MAX wants to reach the 1 sink
- player MIN wants to prevent him from doing so

- General definition of a **strategy** σ for a player *MAX* :

$$\sigma : \text{history of play ending in } V_{MAX} \longmapsto \text{probability distribution on outneighbours}$$

- The **value** of a vertex x is

$$v(x) = \sup_{\substack{\sigma \text{ strategy} \\ for MAX}} \inf_{\substack{\tau \text{ strategy} \\ \text{for MIN}}} \underbrace{\mathbb{P}_{\sigma,\tau}(1 \text{ is reached} \quad | \quad \text{game starts in } x)}_{v_{\sigma,\tau(x)}}$$

- to compute values we can restrict our attention to *pure, stationnary, memoriless* strategies (**positional strategies** for short) :

$$\sigma : V_{MAX} \longrightarrow V, \quad \tau : V_{MIN} \longrightarrow V$$

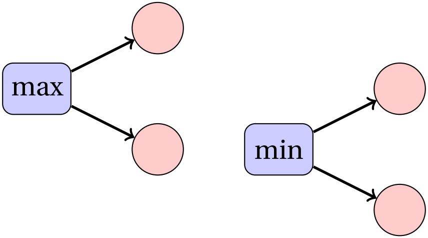

Theorem (Condon 89)

For all vertices x,

$$v(x) = \max_{\substack{\sigma \text{ positional strategy} \\ \text{for MAX}}} \min_{\substack{\tau \text{ positional strategy} \\ \text{for MIN}}} v_{\sigma,\tau}(x)$$

$$= \min_{\substack{\tau \text{ positional strategy} \\ \text{for MIN}}} \max_{\substack{\sigma \text{ positional strategy} \\ \text{for MAX}}} v_{\sigma,\tau}(x)$$

main lines of a proof ...

1. **sups** and **infs** are **maxs** and **mins** : optimal strategies and best responses exists (compacity and continuity arguments)

main lines of a proof ...

1. **sups** and **infs** are **maxs** and **mins** : optimal strategies and best responses exists (compacity and continuity arguments)
2. against a **positional** strategy σ, MIN might as well respond positional :

$$\sigma \text{ positional} \Rightarrow \quad \min_{\tau \text{ general}} v_{\sigma,\tau}(x) = \min_{\tau \text{ positional}} v_{\sigma,\tau}(x)$$

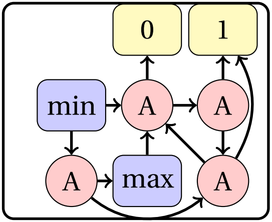

When reaching any x MIN plays the first move of any optimal strategy starting in x

idea of proof ...

3. so

$$\max_{\text{pos}} \min_{\text{pos}} = \max_{\text{pos}} \min_{\text{gen}} \leq \max_{\text{gen}} \min_{\text{gen}} \leq \min_{\text{gen}} \max_{\text{gen}} \leq \min_{\text{pos}} \max_{\text{gen}} = \min_{\text{pos}} \max_{\text{pos}}$$

idea of proof ...

3. so

$$\max_{pos}\min_{pos} = \max_{pos}\min_{gen} \leq \max_{gen}\min_{gen} \leq \min_{gen}\max_{gen} \leq \min_{pos}\max_{gen} = \min_{pos}\max_{pos}$$

4. However

$$\max_{pos}\min_{pos} = \min_{pos}\max_{pos}$$

finite number of strategies → **zero-sum matrix game** (exponentially sized)

$$\left\{\begin{array}{rl} \max & t \\ \text{for all pure } \tau, & v_{\sigma,\tau} \geq t \\ \sigma \text{ prob. on} & \text{pure strategies} \end{array}\right. = \left\{\begin{array}{rl} \min & t \\ \text{for all pure } \sigma, & v_{\sigma,\tau} \leq t \\ \tau \text{ prob. on} & \text{pure strategies} \end{array}\right.$$

by strong duality theorem

idea of proof ...

3. so

$$\max_{pos}\min_{pos} = \max_{pos}\min_{gen} \leq \max_{gen}\min_{gen} \leq \min_{gen}\max_{gen} \leq \min_{pos}\max_{gen} = \min_{pos}\max_{pos}$$

4. However

$$\max_{pos}\min_{pos} \quad = \quad \min_{pos}\max_{pos}$$

finite number of strategies → **zero-sum matrix game** (exponentially sized)

$$\left\{\begin{array}{rl} \max & t \\ \text{for all pure } \tau, & v_{\sigma,\tau} \geq t \\ \sigma \text{ prob. on} & \text{pure strategies} \end{array}\right. = \left\{\begin{array}{rl} \min & t \\ \text{for all pure } \sigma, & v_{\sigma,\tau} \leq t \\ \tau \text{ prob. on} & \text{pure strategies} \end{array}\right.$$

by strong duality theorem

5. Finally, random strategies are useless since the game is positional

Computing values

Fix σ, τ positional strategies.

- if $x \in V_{MAX}$, $v_{\sigma,\tau}(x) = v_{\sigma,\tau}(\sigma(x))$
- if $x \in V_{MIN}$, $v_{\sigma,\tau}(x) = v_{\sigma,\tau}(\tau(x))$
- if $x \in V_{AVE}$, $v_{\sigma,\tau}(x) = \frac{1}{2} v_{\sigma,\tau}(x_1) + \frac{1}{2} v_{\sigma,\tau}(x_2))$

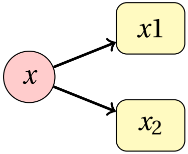

Let $S = \{$ vertices having a directed path to a sink $\}$

Computing values

Fix σ, τ positional strategies.

- if $x \in V_{MAX}$, $v_{\sigma,\tau}(x) = v_{\sigma,\tau}(\sigma(x))$
- if $x \in V_{MIN}$, $v_{\sigma,\tau}(x) = v_{\sigma,\tau}(\tau(x))$
- if $x \in V_{AVE}$, $v_{\sigma,\tau}(x) = \frac{1}{2} v_{\sigma,\tau}(x_1) + \frac{1}{2} v_{\sigma,\tau}(x_2))$

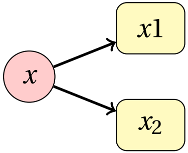

Let $S = \{$ vertices having a directed path to a sink $\}$

- if $x \notin S$ then $v_{\sigma,\tau}(x) = 0$

Computing values

Fix σ, τ positional strategies.

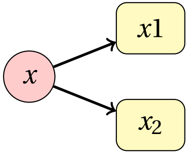

- if $x \in V_{MAX}$, $v_{\sigma,\tau}(x) = v_{\sigma,\tau}(\sigma(x))$
- if $x \in V_{MIN}$, $v_{\sigma,\tau}(x) = v_{\sigma,\tau}(\tau(x))$
- if $x \in V_{AVE}$, $v_{\sigma,\tau}(x) = \frac{1}{2} v_{\sigma,\tau}(x_1) + \frac{1}{2} v_{\sigma,\tau}(x_2))$

Let $S = \{$ vertices having a directed path to a sink $\}$

- if $x \notin S$ then $v_{\sigma,\tau}(x) = 0$
- previous system :

$$v_S = Q v_S + b$$

with $I - Q$ nonsingular so

$$v_S = (I - Q)^{-1} b$$

Computing values

Fix σ, τ positional strategies.

- if $x \in V_{MAX}$, $v_{\sigma,\tau}(x) = v_{\sigma,\tau}(\sigma(x))$
- if $x \in V_{MIN}$, $v_{\sigma,\tau}(x) = v_{\sigma,\tau}(\tau(x))$
- if $x \in V_{AVE}$, $v_{\sigma,\tau}(x) = \frac{1}{2} v_{\sigma,\tau}(x_1) + \frac{1}{2} v_{\sigma,\tau}(x_2))$

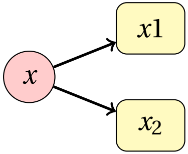

Let $S = \{$ vertices having a directed path to a sink $\}$

- if $x \notin S$ then $v_{\sigma,\tau}(x) = 0$
- previous system :

$$v_S = Q v_S + b$$

  with $I - Q$ nonsingular so

$$v_S = (I - Q)^{-1} b$$

- $I - Q$ and b have entries in $\{0, \pm 1, \pm \frac{1}{2}\}$
  $v_{\sigma,\tau}$ has rational entries with denominator at most 4^n.

stopping SSGs

A SSG is stopping if for all strategies, the game reaches a sink vertex almost surely.

Theorem (Condon 89)

For every SSG G, there is a polynomial-time computable SSG G' such that

- *G' is stopping*
- *size of G' = poly(size of G)*
- *for all vertices x, $v_{G'}(x) > \frac{1}{2}$ if and only if $v_G(x) > \frac{1}{2}$*

stopping SSGs

A SSG is stopping if for all strategies, the game reaches a sink vertex almost surely.

Idea of proof

1. $v_G(x) > \frac{1}{2} \iff v_G(x) \geq \frac{1}{2} + 4^{-n}$
2. values are **stable under perturbations**,

stopping SSGs

A SSG is stopping if for all strategies, the game reaches a sink vertex almost surely.

Idea of proof

1. $v_G(x) > \frac{1}{2} \iff v_G(x) \geq \frac{1}{2} + 4^{-n}$
2. values are **stable under perturbations**,
3. replace all arcs

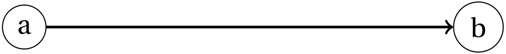

by

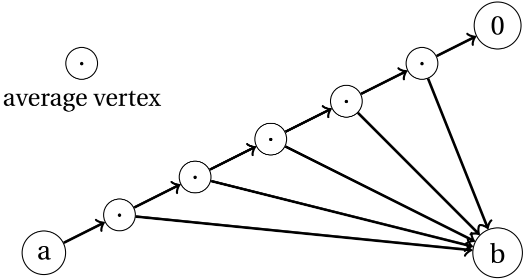

giving a small probability to every vertex to go **reach the 0 sink**

From now on we suppose SSGs **stopping**

(even if I forget to write / say it)

the switch operation

Let x be a MIN vertex.
Suppose $v_{\sigma,\tau}(x) = v_{\sigma,\tau}(x_1) > v_{\sigma,\tau}(x_2))$

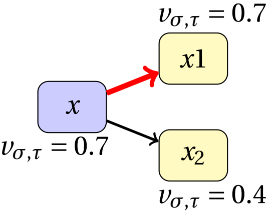

switching τ at x :
$\tau'(x) = x_2$ and equal to $\tau' = \tau$ elsewhere.

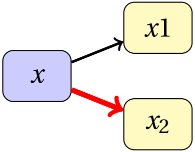

Such a switch is **profitable** for MIN : $\tau' < \tau$

- for all y, $v_{\sigma,\tau'}(y) \leq v_{\sigma,\tau}(y)$
- in particular $v_{\sigma,\tau'}(x) < v_{\sigma,\tau}(x)$

the switch operation

τ_k = time-dependent strategy equal to

- τ' at times $0, 1, \cdots k-1$
- τ thereafter.

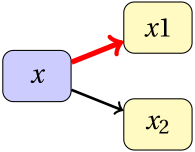

Then against σ : (following Gimbert & Horn)

- $\tau_0 = \tau$

the switch operation

τ_k = time-dependent strategy equal to

- τ' at times $0, 1, \cdots k-1$
- τ thereafter.

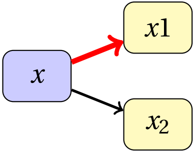

Then against σ : (following Gimbert & Horn)

- $\tau_0 = \tau$
- $\tau_1(x) < \tau(x)$

the switch operation

τ_k = time-dependent strategy equal to

- τ' at times $0, 1, \cdots k-1$
- τ thereafter.

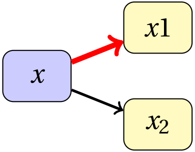

Then against σ : (following Gimbert & Horn)

- $\tau_0 = \tau$
- $\tau_1(x) < \tau(x)$
- for all $k \geq 0$: $\tau_{k+1} \leq \tau_k$
  - conditionnal on token not in x at time k same probability of reaching 1
  - conditionnal on token in x at time k the probability of reaching 1 is smaller

the switch operation

τ_k = time-dependent strategy equal to

- τ' at times $0, 1, \cdots k-1$
- τ thereafter.

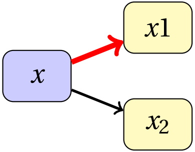

Then against σ : (following Gimbert & Horn)

- $\tau_0 = \tau$
- $\tau_1(x) < \tau(x)$
- for all $k \geq 0$: $\quad \tau_{k+1} \leq \tau_k$
  - conditionnal on token not in x at time k same probability of reaching 1
  - conditionnal on token in x at time k the probability of reaching 1 is smaller
- $$\tau = \tau_0 > \tau_1 \geq \tau_2 \geq \cdots \lim_{\infty} \tau_k = \tau'$$

optimality conditions

Suppose σ fixed, we want to compute a best-response $\tau(\sigma)$.

Lemma

Let G be a stopping SSG, and σ a positional strategy for MAX. *Then τ is a best-response to σ if and only*

$$\text{for all } x \in V_{MIN}, \quad v_{\sigma,\tau}(x) = \min(v_{\sigma,\tau}(x_1), v_{\sigma,\tau}(x_2))$$

proof : if not, switch.

optimality conditions

Suppose σ fixed, we want to compute a best-response $\tau(\sigma)$.

Lemma

Let G be a stopping SSG, and σ a positional strategy for MAX. *Then τ is a best-response to σ if and only*

$$\text{for all } x \in V_{MIN}, \quad v_{\sigma,\tau}(x) = \min(v_{\sigma,\tau}(x_1), v_{\sigma,\tau}(x_2))$$

proof : if not, switch.

Lemma

G stopping SSG, and σ, τ are optimal strategies if and only if

$$\text{for all } x \in V_{MIN}, \quad v_{\sigma,\tau}(x) = \min(v_{\sigma,\tau}(x_1), v_{\sigma,\tau}(x_2))$$

$$\text{for all } x \in V_{MAX}, \quad v_{\sigma,\tau}(x) = \max(v_{\sigma,\tau}(x_1), v_{\sigma,\tau}(x_2))$$

SSG $\Longleftrightarrow$ max / min / average systems

computing a best response

- Suppose G is an SSG and σ is fixed.
- Define

$$F_\sigma : \begin{cases} [0,1]^V & \longrightarrow & [0,1]^V \\ v_x & \longmapsto & \begin{cases} \min(v_{x_1}, v_{x_2}) \text{ if } x \in V_{MIN} \\ v_{\sigma(x)} \text{ if } x \in V_{MAX} \\ \frac{1}{2} v_{x_1} + \frac{1}{2} v_{x_2} \text{ if } x \in V_{AVE} \end{cases} \end{cases}$$

where the values of sinks are replaced by 0 or 1.

computing a best response

- Suppose G is an SSG and σ is fixed.
- Define

$$F_\sigma : \begin{cases} [0,1]^V & \longrightarrow & [0,1]^V \\ v_x & \longmapsto & \begin{cases} \min(v_{x_1}, v_{x_2}) \text{ if } x \in V_{MIN} \\ v_{\sigma(x)} \text{ if } x \in V_{MAX} \\ \frac{1}{2} v_{x_1} + \frac{1}{2} v_{x_2} \text{ if } x \in V_{AVE} \end{cases} \end{cases}$$

  where the values of sinks are replaced by 0 or 1.
- Operator F_σ is **contracting** (sup norm)
  $\rightarrow$ single fixed point = value vector of σ (values vs best response)

computing a best response

- Suppose G is an SSG and σ is fixed.
- Define

$$F_\sigma : \left\{ \begin{array}{ccc} [0,1]^V & \longrightarrow & [0,1]^V \\ v_x & \longmapsto & \left\{ \begin{array}{c} \min(v_{x_1}, v_{x_2}) \text{ if } x \in V_{MIN} \\ v_{\sigma(x)} \text{ if } x \in V_{MAX} \\ \frac{1}{2} v_{x_1} + \frac{1}{2} v_{x_2} \text{ if } x \in V_{AVE} \end{array} \right. \end{array} \right.$$

  where the values of sinks are replaced by 0 or 1.
- Operator F_σ is **contracting** (sup norm)
  → single fixed point = value vector of σ (values vs best response)
- solving $F_\sigma v = v$ by linear programing

$$\max \sum_i v_i$$

$$F_\sigma(v) \leq v$$

algorithmic complexity

Value computation problem : given a SSG and a vertex x, does

$$v(x) > \frac{1}{2}\,?$$

Theorem

The value complexity problem for SSG lies in complexity class $NP \cap co-NP$.

Guess a couple (σ, τ) of positional strategies, compute the values (linear system) and check optimality conditions.

Theorem

The value complexity problem for SSG lies in complexity class $UP \cap co-UP$.

strategy improvement algorithms

The strategy improvement algorithm a.k.a Hoffman-Karp algorithm (1966, MDP context) is

0. choose σ_0 and let $\tau_0 = \tau(\sigma_0)$ (best response)
1. while (σ_k, τ_k) is not optimal, obtain σ_{k+1} by switch σ_k; let $\tau_{k+1} = \tau(\sigma_{k+1})$

based on :

Lemma

$v_{\sigma_{k+1},\tau_{k+1}} > v_{\sigma_k,\tau_k}$

strategy improvement algorithms

The strategy improvement algorithm a.k.a Hoffman-Karp algorithm (1966, MDP context) is

0. choose σ_0 and let $\tau_0 = \tau(\sigma_0)$ (best response)
1. while (σ_k, τ_k) is not optimal, obtain σ_{k+1} by switch σ_k ; let $\tau_{k+1} = \tau(\sigma_{k+1})$

based on :

Lemma

$v_{\sigma_{k+1},\tau_{k+1}} > v_{\sigma_k,\tau_k}$

Theorem

The HK algorithm makes at most $O(2^n/n)$ iterations

Unfortunately : this can take exponential time :

- Friedmann (2009) gives a counter-example for parity game ($2^{\sqrt{n}}$ iterations, claimed 2^{cn})
- Andersson (2009) shows that this counterexample survives the reduction (to come on last slides)

the 'counter-example' of Friedman

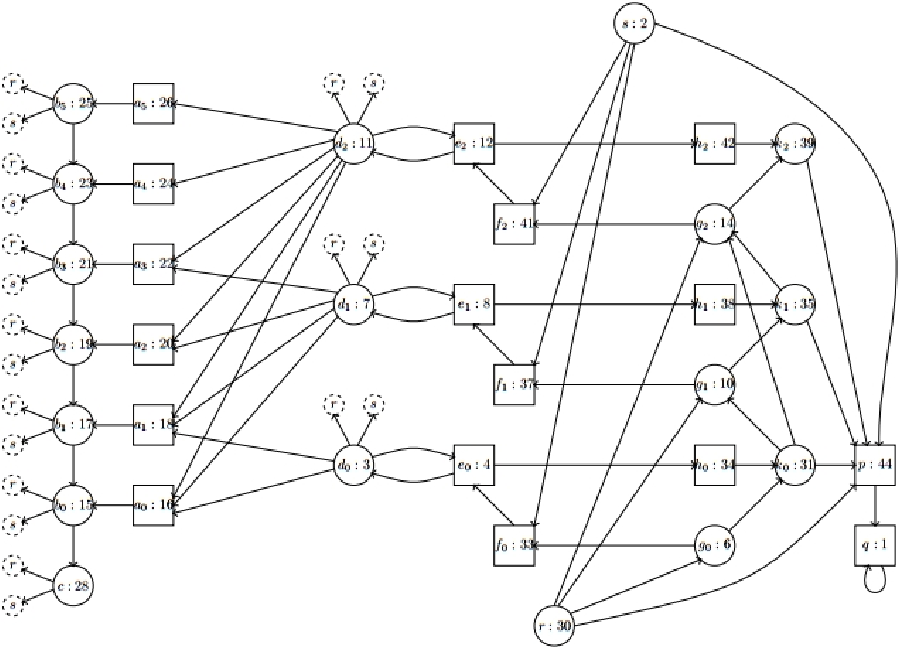

SSG without average vertices

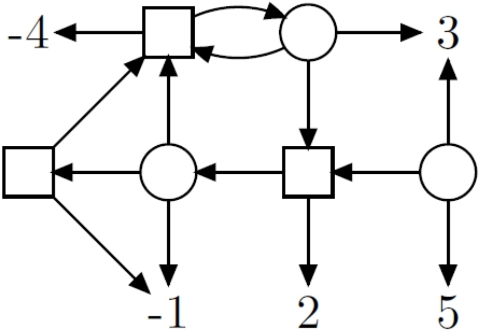

a.k.a. **deterministic graphical games** (Washburn 1966, Andersson et al. 2012)

Definition = SSG without average vertices, but allow sinks with arbitrary payoffs

SSG without average vertices

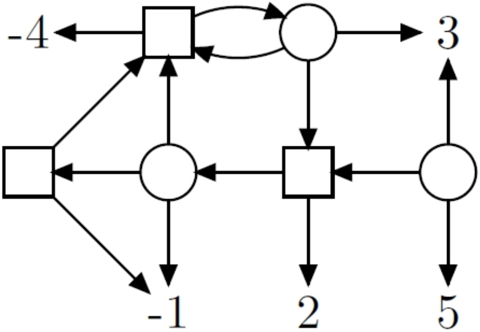

a.k.a. **deterministic graphical games** (Washburn 1966, Andersson et al. 2012)

Definition = SSG without average vertices, but allow sinks with arbitrary payoffs

Solving DGG in linear time by backtracking

While possible :

1. sink *s* with maximal payoff : if an incoming MIN arcs never go there if they have a choice : **delete arc** or **merge**
2. Do the opposite for the minimum payoff sink.

In the end remain vertices with no connection to sinks, their value is 0.

an FPT algorithm on the number of average nodes (Gimbert & Horn 2009)

Theorem

There is an algorithm which computes values and optimal strategies of SSGs with n vertices and k average vertices in time $O((k! \cdot n)$.

(Moreover the outdegree of nodes is unlimited)

an FPT algorithm on the number of average nodes (Gimbert & Horn 2009)

Theorem

There is an algorithm which computes values and optimal strategies of SSGs with n vertices and k average vertices in time $O((k! \cdot n)$.

(Moreover the outdegree of nodes is unlimited)

- a strategy consists in choosing among nodes. Hence an preference order on all nodes yields a strategy.

an FPT algorithm on the number of average nodes (Gimbert & Horn 2009)

Theorem

There is an algorithm which computes values and optimal strategies of SSGs with n vertices and k average vertices in time $O((k! \cdot n)$.

(Moreover the outdegree of nodes is unlimited)

- a strategy consists in choosing among nodes. Hence an preference order on all nodes yields a strategy.
- but an order on V_{AVE} is enough

$$0 < a_1 < a_2 \cdots a_k < 1$$

MIN tries to force the next average vertex to be great
MIN tries to force the next average vertex to be small

an FPT algorithm on the number of average nodes (Gimbert & Horn 2009)

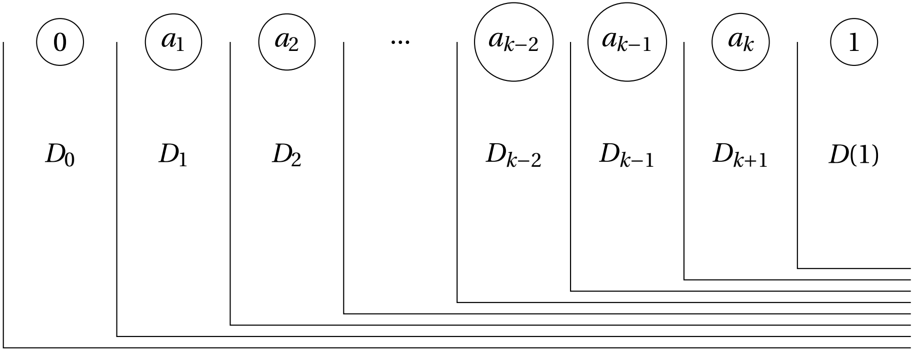

$$D_i = \text{Deterministic Attractor of } \{a_i, a_{i+1}, \cdots, a_k, 1\}$$

The deterministic attractor $D(X)$ of X is the set of *MAX*, *MIN* vertices from where MAX has a strategy forcing X to be reached.

an FPT algorithm on the number of average nodes (Gimbert & Horn 2009)

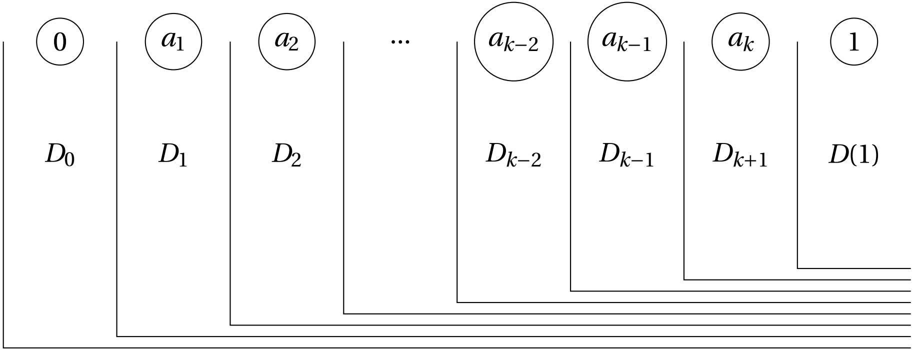

For every order f on AVE vertices, two strategies σ_f, τ_f such that
game is in $D_i \setminus D_{i+1}$ at any time $\Rightarrow$ next average vertex is a_i

Theorem

If the order f is coherent with the real values of the game (+small condition if some values are equal) then strategies σ_f, τ_f are optimal.

an FPT algorithm on the number of average nodes (Gimbert & Horn 2009)

The $O((k! \cdot n)$ was improved to :

- $O\left((4^k k^c n^c\right)$ (Chaterjee et al 2009)
- $O\left((k 2^k (k \log k + n)\right)$ (Ibsen-Jensen et al 2012), using involved extremal combinatorics to establish the bound.

an FPT algorithm on the number of average nodes (Gimbert & Horn 2009)

The $O((k! \cdot n)$ was improved to :

- $O\left((4^k k^c n^c\right)$ (Chaterjee et al 2009)
- $O\left((k 2^k (k \log k + n)\right)$ (Ibsen-Jensen et al 2012), using involved extremal combinatorics to establish the bound.

Question : here is my simple idea for $O\left(2^k n^2\right)$, what do you think ? (oral only, sorry)

parity games

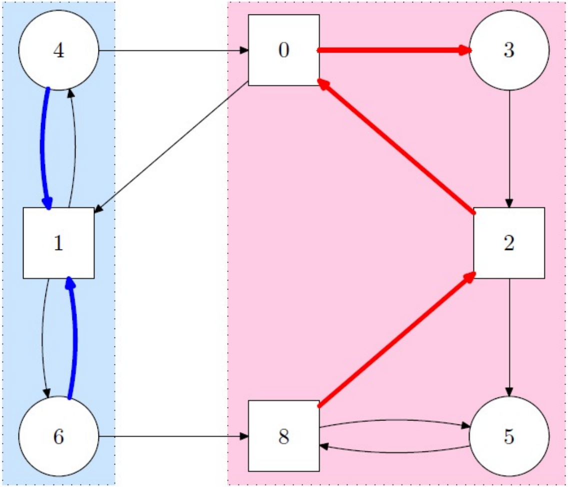

- two player game on a graph (no random)
- Play goes on forever
- every vertex has a priority
- strategies fixed, moves are determistic
- a cycle is repeated

If the greatest priority on the cycle is even, player 0 wins if it is odd player 1 wins.

Every vertex is either a win for 0 or 1

parity games

Theorem

Determining the winner of a parity game for a given start vertex is in $NP \cap co-NP$ (in fact $UP \cap co-UP$)

Open Question : Is it in P ?

parity games

Theorem

Determining the winner of a parity game for a given start vertex is in $NP \cap co-NP$ (in fact $UP \cap co-UP$)

Open Question : Is it in P?

Theorem

There is a Karp reduction from parity games to stochastic parity games, such that a vertex is winning for 1 in the PG if the corresponding vertex has value $> \frac{1}{2}$ in the SSG

idea :

- add two sinks 0 and 1
- assign for every transition a small probability to go to sink 0 (nodes of player 0) or sink 1 (nodes of player 1)

parity games

Theorem

Determining the winner of a parity game for a given start vertex is in $NP \cap co-NP$ (in fact $UP \cap co-UP$)

Open Question : Is it in P ?

Theorem

There is a Karp reduction from parity games to stochastic parity games, such that a vertex is winning for 1 in the PG if the corresponding vertex has value $> \frac{1}{2}$ in the SSG

idea :

- add two sinks 0 and 1
- assign for every transition a small probability to go to sink 0 (nodes of player 0) or sink 1 (nodes of player 1)

Open Question : is there a polynomial reduction in the other direction ?

thank you!

selective bibliography

D. Andersson and P. Miltersen.
The complexity of solving stochastic games on graphs.
Algorithms and Computation, pages 112–121, 2009.

D. Andersson.
Extending friedmann's lower bound to the hoffman-karp algorithm.
Preprint (June 2009), 2009.

M. Blum, B. Juba, and R. Williams.
Non-monotone behaviors in min/max/avg circuits and their relationship to simple stochastic games.
2008.

A. Condon et al.
On algorithms for simple stochastic games.
Advances in computational complexity theory, 13 :51–73, 1993.

K. Chatterjee, L. De Alfaro, and T.A. Henzinger.
Termination criteria for solving concurrent safety and reachability games.
In Proceedings of the twentieth Annual ACM-SIAM Symposium on Discrete Algorithms, pages 197–206. Society for Industrial and Applied Mathematics, 2009.

A. Condon.
The complexity of stochastic games.
Information and Computation, 96(2) :203–224, 1992.

H. Gimbert and F. Horn.
Simple stochastic games with few random vertices are easy to solve.
Foundations of Software Science and Computational Structures, pages 5–19, 2008.

N. Halman.
Simple stochastic games, parity games, mean payoff games and discounted payoff games are all lp-type problems.
Algorithmica, 49(1) :37–50, 2007.

R. Ibsen-Jensen and P. Miltersen.
Solving simple stochastic games with few coin toss positions.
Algorithms–ESA 2012, pages 636–647, 2012.

W. Ludwig.
A subexponential randomized algorithm for the simple stochastic game problem.
1995.

R. Somla.
New algorithms for solving simple stochastic games.
Electronic Notes in Theoretical Computer Science, 119(1) :51–65, 2005.

R. Tripathi, E. Valkanova, and VS Anil Kumar.
On strategy improvement algorithms for simple stochastic games.
Journal of Discrete Algorithms, 9(3) :263–278, 2011.