Introduction to Game Theory

(From a CS Point of View)

Olivier Serre
Serre@irif.fr

IRIF (CNRS & Université Paris Diderot – Paris 7)

10th of September 2019
Master Parisien de Recherche en Informatique

Who is who?

- Olivier Serre (IRIF, Univ. P7 & CNRS) will be teaching: Games on finite graphs; Games, tree automata and logic;
- Nathanaël Fijalkow (LaBRI, Univ. Bordeaux & CNRS) will give a lecture on universal graphs and parity games.
- Dietmar Berwanger (LSV, ÉNS Cachan & CNRS) will be teaching: Mean-payoff / Simple Stochastic games; imperfect information (single / multiplayer) games.
- Wiesław Zielonka (IRIF, Univ. P7 & CNRS) is also part of this course but he is not teaching this year.

To contact us:

`serre@irif.fr` / `nathanael.fijalkow@labri.fr` / `dwb@lsv.ens-cachan.fr` / `zielonka@irif.fr`

Purpose of introduction lecture

- Present several examples of games to give intuition of the various questions considered in this course.
- Not a formal course, based on intuitions rather than on formal reasoning.
- Of course, the next lectures will be much more formal...

Oh, I forgot!

Games are based on interactions

So does this course. Please:

- Ask questions whenever something is not clear enough.
- Answer my questions even if you are not sure to be right.
- Read the notes and ask questions by email or at the beginning / end of the course if necessary.

Chomp game (David Gale, 1974)

(D. Gale, A curious Nim-type game, Amer. Math. Monthly 81 (1974) 876-879)

The game of Chomp is like Russian Roulette for chocolate lovers :-)

A move consists of chomping a square out of the chocolate bar along with any squares to the right and below. Players alternate moves. The upper left square is poisoned though and the player forced to chomp it loses (and actually dies. . .).

Chomp game (David Gale, 1974)

(D. Gale, A curious Nim-type game, Amer. Math. Monthly 81 (1974) 876-879)

The game of Chomp is like Russian Roulette for chocolate lovers :-)

A move consists of chomping a square out of the chocolate bar along with any squares to the right and below. Players alternate moves. The upper left square is poisoned though and the player forced to chomp it loses (and actually dies. . .).

Characteristics of this game:

- Zero sum (one player wins, the other looses)
- Finite duration
- Turn based
- Perfect information
- Deterministic

For us, this will be the simplest kind of game (however very few is known about this "simple" game. . .).

Modeling of the game 3×2

Possible configurations:

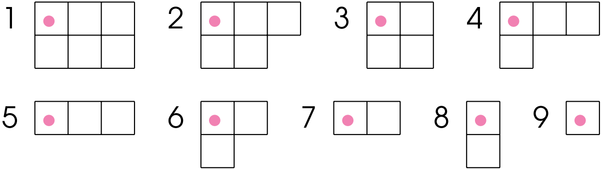

Modeling of the game **3** × **2**

Associated arena:

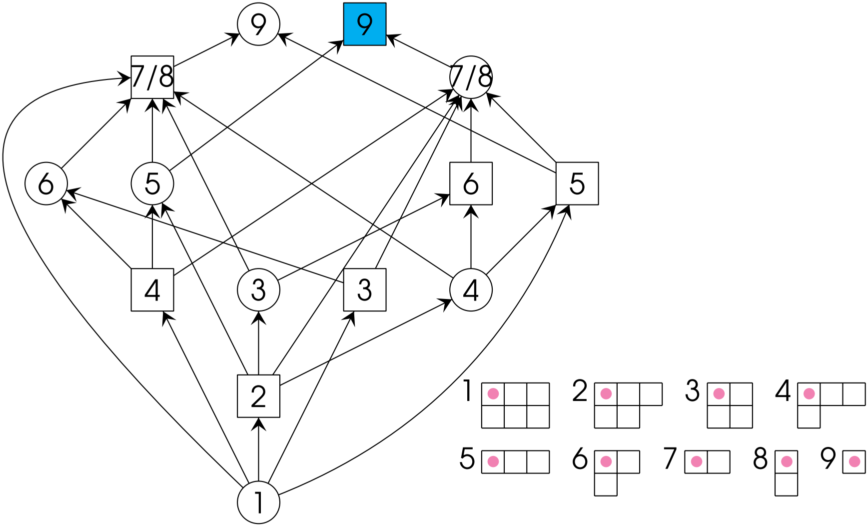

Associated arena:

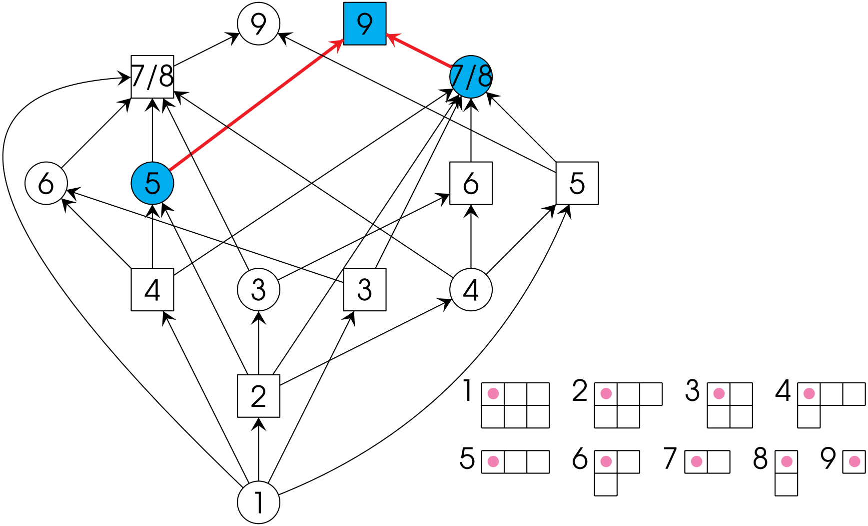

Associated arena:

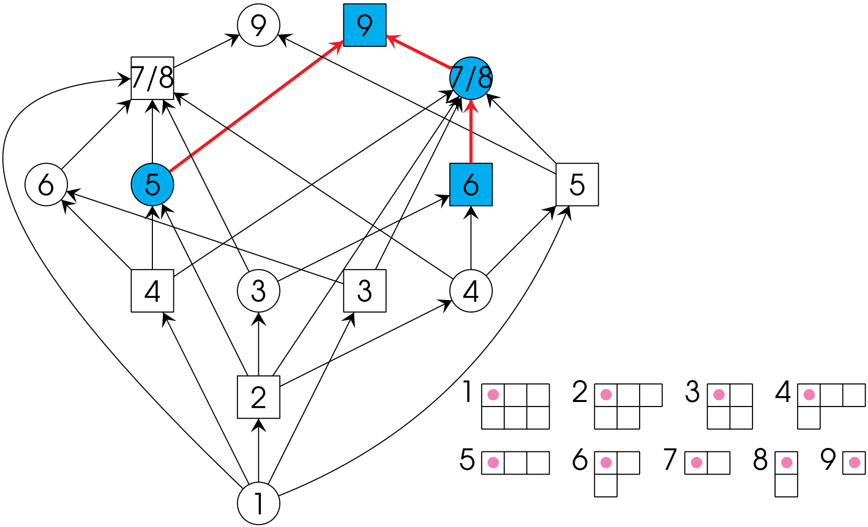

Associated arena:

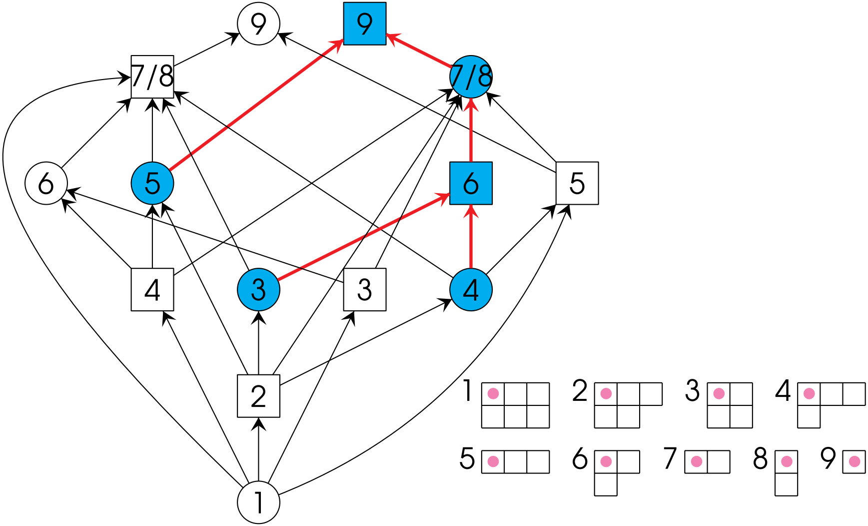

Modeling of the game **3** × **2**

Associated arena:

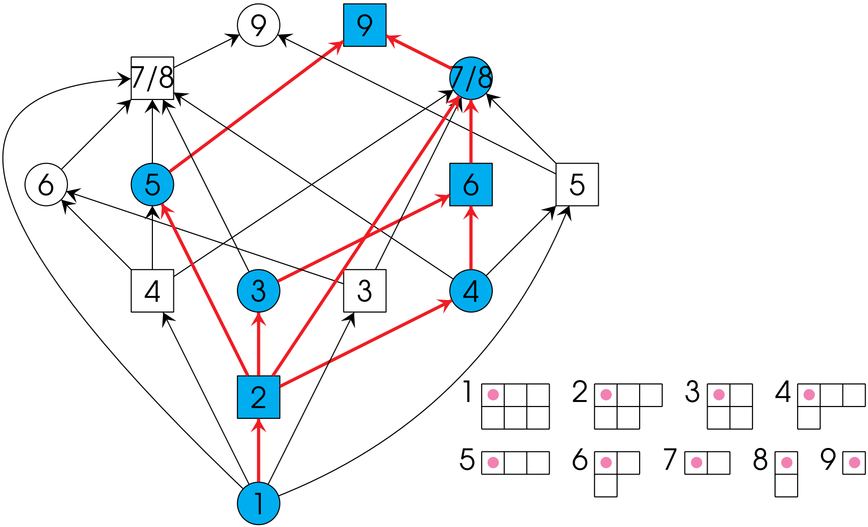

Exercice: generalization

Exercice

Given a chocolate bar of size $n \times m$ who has a winning strategy?

Exercice: generalization

Exercice

Given a chocolate bar of size $n \times m$ who has a winning strategy?

I need some help...

Designing algorithms for games on (finite) arena will be the topic of courses # 1 and # 2 (Olivier Serre).

Prisoner's dilemma

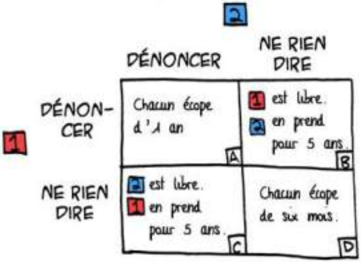

Two suspects are arrested by the police. The police have insufficient evidence for a conviction, and, having separated the prisoners, visit each of them to offer the same deal. If one testifies for the prosecution against the other (defects) and the other remains silent (cooperates), the defector goes free and the silent accomplice receives the full five-year sentence. If both remain silent, both prisoners are sentenced to only six months in jail for a minor charge. If each betrays the other, each receives a one-year sentence. Each prisoner must choose to betray the other or to remain silent. Each one is assured that the other would not know about the betrayal before the end of the investigation. How should the prisoners act?

Prisoner's dilemma

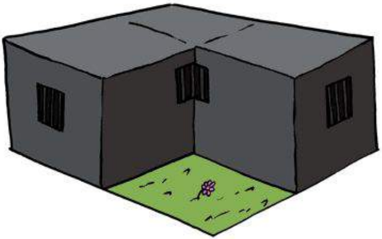

Characteristics of this game:

- Non zero sum
- Finite duration
- Concurrent
- Perfect information
- Deterministic

Nash equilibria & Prisoner's dilemma

Nash equilibrium

A strategy profile (*i.e.* a choice of action per each player) is a Nash Equilibrium if no player has anything to gain by changing only his own strategy unilaterally.

Nash equilibrium

A strategy profile (*i.e.* a choice of action per each player) is a Nash Equilibrium if no player has anything to gain by changing only his own strategy unilaterally.

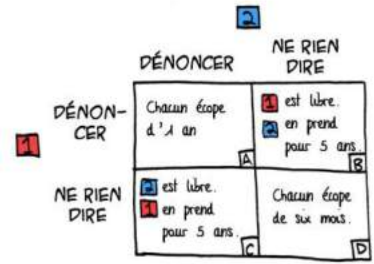

Is there a Nash Equilibrium here?

Nash equilibrium

A strategy profile (*i.e.* a choice of action per each player) is a Nash Equilibrium if no player has anything to gain by changing only his own strategy unilaterally.

You want to know more?

Strategic Games will be the topic of courses of Dietmar Berwanger.

Rock Paper Scissors (Japan, in the late 19th)

Rock Paper Scissors (Japan, in the late 19th)

Characteristics of this game:

- Zero sum
- Finite duration
- Concurrent
- Perfect information
- Deterministic

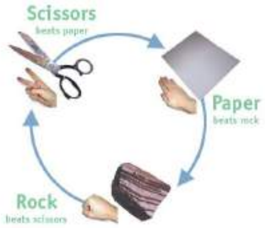

	R	P	S
R	0 / 0	-1 / +1	+1 / -1
P	+1 / -1	0 / 0	-1 / +1
S	-1 / 1	+1 / -1	0 / 0

Nash equilibria & Rock Paper Scissors

	R	P	S
R	0 / 0	-1 / +1	+1 / -1
P	+1 / -1	0 / 0	-1 / +1
S	-1 / 1	+1 / -1	0 / 0

Is there a Nash Equilibrium here?

Is there a Nash Equilibrium here?

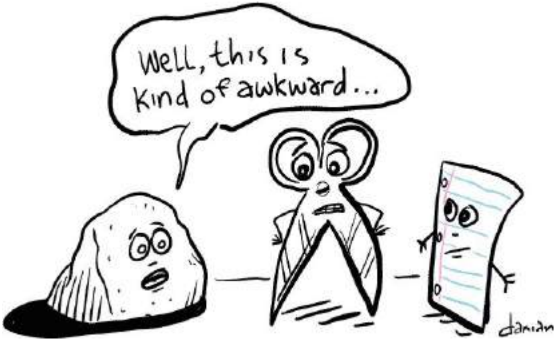

Nash equilibria & Rock Paper Scissors

Randomized strategies may be useful for concurrent games!

	R	P	S
R	0 / 0	-1 / +1	+1 / -1
P	+1 / -1	0 / 0	-1 / +1
S	-1 / 1	+1 / -1	0 / 0

Is there a Nash Equilibrium here?

Randomized strategies may be useful for concurrent games!

Theorem (Nash's theorem (1950))

Every finite game has a mixed strategy equilibrium.

You want to know more?

Strategic Games will be the topic of courses of Dietmar Berwanger.

Ice cream seller

Ice cream seller

Consider a beach (segment) with a uniformly distributed infinite set of people. Everyone buy an ice cream every day.
There are n ice cream sellers that in the morning choose simultaneously where to stay for the whole day. Of course, you buy your ice cream to the closest seller. If two sellers are sitting at the same place, they uniformly share their clients.

Characteristics of this game:

- Zero sum
- Finite duration
- Concurrent
- Perfect information
- Deterministic
- Infinite set of choices

Ice cream seller

Consider a beach (segment) with a uniformly distributed infinite set of people. Everyone buy an ice cream every day.
There are n ice cream sellers that in the morning choose simultaneously where to stay for the whole day. Of course, you buy your ice cream to the closest seller. If two sellers are sitting at the same place, they uniformly share their clients.

Exercise

- Assume $n = 2$. Can you model this game as previously? Is there a Nash equilibrium with non randomized strategies?
- Same for $n = 3$.
- Same for $n = 5$.

Stochastic games

You want to know more?

Stochastic games will be discussed in courses of Dietmar Berwanger.

Imperfect information (possibly stochastic) games

You want to know more?

Imperfect information games will be discussed in courses of Dietmar Berwanger.

Infinite duration games

Infinite duration can come from:

- The game itself, having an unbounded number of rounds or having loops.
- The winning condition (*e.g.* "go infinitely often through a good state", "never visit a bad state", "whenever a blue state is visited a red state should be visited later", "average payoff should be positive"...).

Infinite duration games

Infinite duration can come from:

- The game itself, having an unbounded number of rounds or having loops.
- The winning condition (*e.g.* "go infinitely often through a good state", "never visit a bad state", "whenever a blue state is visited a red state should be visited later", "average payoff should be positive"...).

Why is this making sense?

- To check validity of logical formulas; to deal with problem from automata on infinite trees (see courses # 9 and #10).
- To model systems that are not supposed to stop after a fixed amount of time.

You want to know more?

Most of the games in this course will have infinite duration!

Algorithmic issues

Definition (Game — informal)

A game consists of

- a state space,
- a set of actions for each player,
- a transition function describing the dynamic of the game,
- a winning condition / payoff function.

Definition (Game — informal)

A game consists of

- a state space,
- a set of actions for each player,
- a transition function describing the dynamic of the game,
- a winning condition / payoff function.

A play is a sequence of states induced by a sequence of actions chosen by the players. For each play, there is a winner / a payoff for every player.

Definition (Game — informal)

A game consists of

- a state space,
- a set of actions for each player,
- a transition function describing the dynamic of the game,
- a winning condition / payoff function.

A play is a sequence of states induced by a sequence of actions chosen by the players. For each play, there is a winner / a payoff for every player.
A strategy for some player is a function that, with every history, associate an action to play (could be a distribution of actions for randomized strategy. . .).

Definition (Game — informal)

A game consists of

- a state space,
- a set of actions for each player,
- a transition function describing the dynamic of the game,
- a winning condition / payoff function.

A play is a sequence of states induced by a sequence of actions chosen by the players. For each play, there is a winner / a payoff for every player.

A strategy for some player is a function that, with every history, associate an action to play (could be a distribution of actions for randomized strategy. . .).

Once one fixes a strategy for every player, if induces a unique play.

Definition (Game — informal)

A game consists of

- a state space,
- a set of actions for each player,
- a transition function describing the dynamic of the game,
- a winning condition / payoff function.

A strategy for some player is a function that, with every history, associate an action to play (could be a distribution of actions for randomized strategy...).
Once one fixes a strategy for every player, if induces a unique play.
A strategy is winning for some player if it only induces winning plays.

Algorithmic issues

Definition (Game — informal)

A game consists of

- a state space,
- a set of actions for each player,
- a transition function describing the dynamic of the game,
- a winning condition / payoff function.

A strategy for some player is a function that, with every history, associate an action to play (could be a distribution of actions for randomized strategy...).
Once one fixes a strategy for every player, if induces a unique play.
A strategy is winning for some player if it only induces winning plays.
A player wins a game if he has a winning strategy.

Algorithmic issues

Definition (Game — informal)

A game consists of

- a state space,
- a set of actions for each player,
- a transition function describing the dynamic of the game,
- a winning condition / payoff function.

A strategy is winning for some player if it only induces winning plays.
A player wins a game if he has a winning strategy.

Algorithmically, one wants to decide who wins a given game and how.

Sure Winning vs Almost Sure Winning

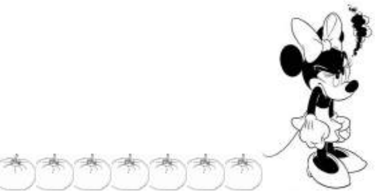

Sure Winning vs Almost Sure Winning

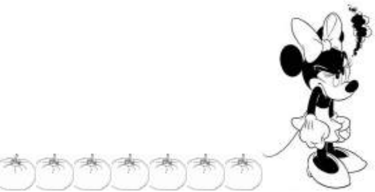

Sure Winning vs Almost Sure Winning

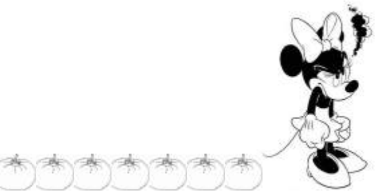

Sure Winning vs Almost Sure Winning

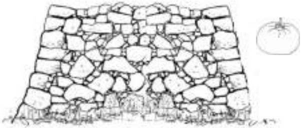

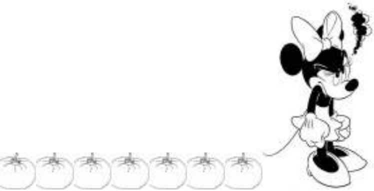

Sure Winning vs Almost Sure Winning

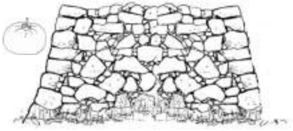

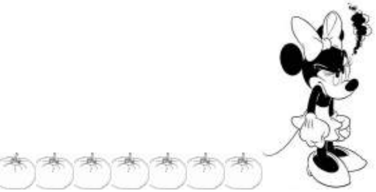

Sure Winning vs Almost Sure Winning

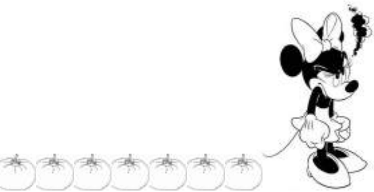

Eve wins a play iff she eventually hits Adam

Sure Winning vs Almost Sure Winning

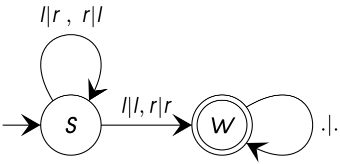

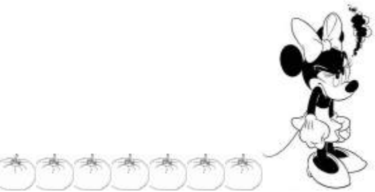

Eve wins a play iff she eventually hits Adam

Sure Winning vs Almost Sure Winning

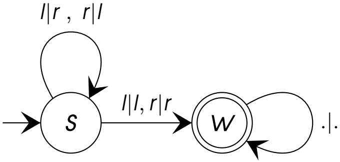

	l	r
l	w	s
r	s	w

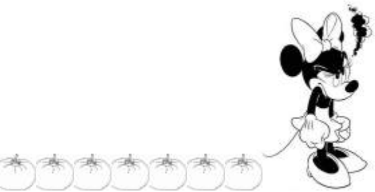

Eve wins a play iff she eventually hits Adam

Sure Winning vs Almost Sure Winning

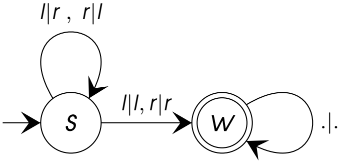

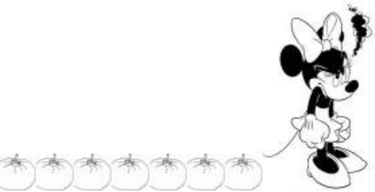

Eve does not have a surely winning strategy but she has an almost surely one.

Almost Sure Winning vs Limit Sure Winning

Almost Sure Winning vs Limit Sure Winning

Almost Sure Winning vs Limit Sure Winning

Almost Sure Winning vs Limit Sure Winning

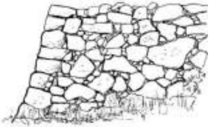

Almost Sure Winning vs Limit Sure Winning

Almost Sure Winning vs Limit Sure Winning

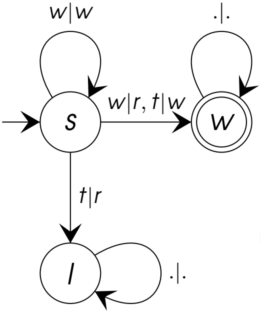

Adam wins iff he eventually reaches the castle.

Almost Sure Winning vs Limit Sure Winning

Adam does not have an almost surely winning strategy but he has a limit surely one.

Adam wins iff he eventually reaches the castle.

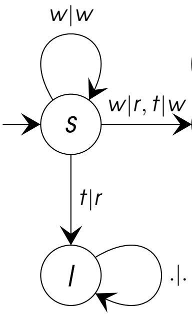

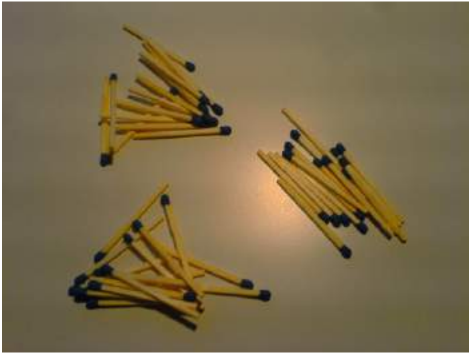

Players remove an arbitrary number of matches but all from the same group. The player that removes the last match looses.