Verification and Game Theory

Patricia Bouyer

LSV, CNRS & ENS Paris-Saclay
Université Paris-Saclay, Cachan, France

Thanks to:

my co-authors Nicolas Markey, Romain Brenguier,
Michael Ummels, Nathan Thomasset
Stéphane Le Roux for recent discussions on the subject
Thomas Brihaye for some of the slides

Outline

1. Verification and game theory

2. Games on graphs
   - The general model
   - Focus on a simple scenario
   - Adding probabilities to the setting?
   - Concurrent games

3. Conclusion

Computer programming

Computer programming is a difficult task

- understand deeply the initial problem;
- find a solution;
- write the program correctly.

Software bugs

- It is an error, a failure in a computer program or system that induces an incorrect result.
- It may have catastrophic consequences.

Software bugs

Bug example

In August 2005, a Malaysian Airlines MH124 (Boeing 777) that was on autopilot suddenly ascended 2,000 feet.

Bug consequences

- loss of confidence from users' point of view,
- loss of credibility from institutions' point of view,
- large financial loss,
- human loss,...

$\Rightarrow$ Real need to **verify** the correctness of a program!

The model-checking approach to verification

Real system *plane,...*	$\stackrel{?}{\models}$	**Specification** *arrive safely,...*

The model-checking approach to verification

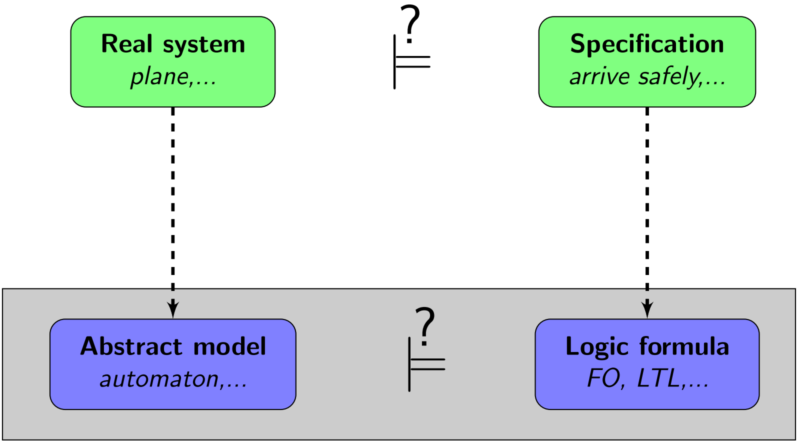

The model-checking approach to verification

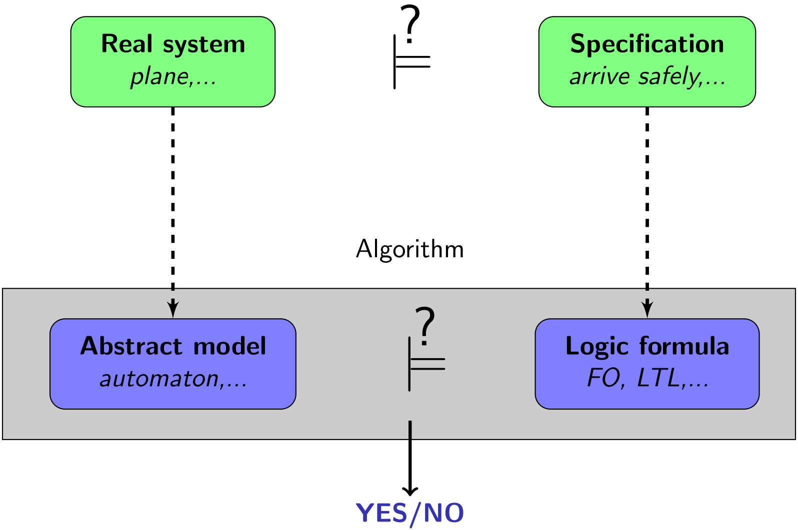

The autopilot case

Requirement: to arrive safely **in every weather condition**,

The autopilot case

Requirement: to arrive safely **in every weather condition**, **while minimising the fuel consumption**.

Controlling computer systems

Real system
plane,...

Environment weather,...

(Quant.) Spec.
arrive safely,
energy cons.,...

Controlling computer systems

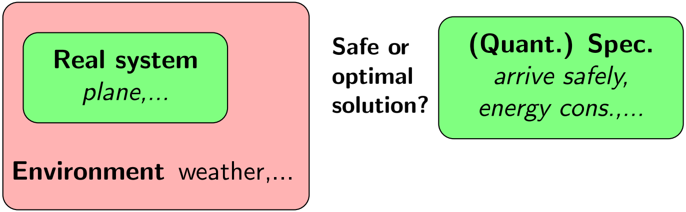

Controlling computer systems

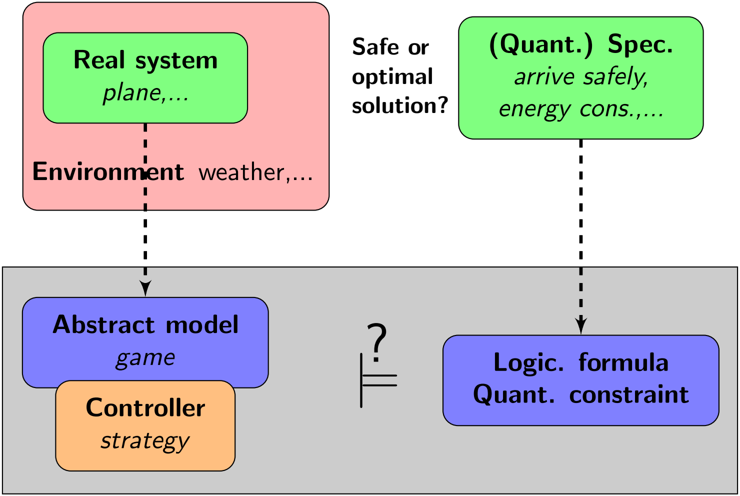

Controlling computer systems

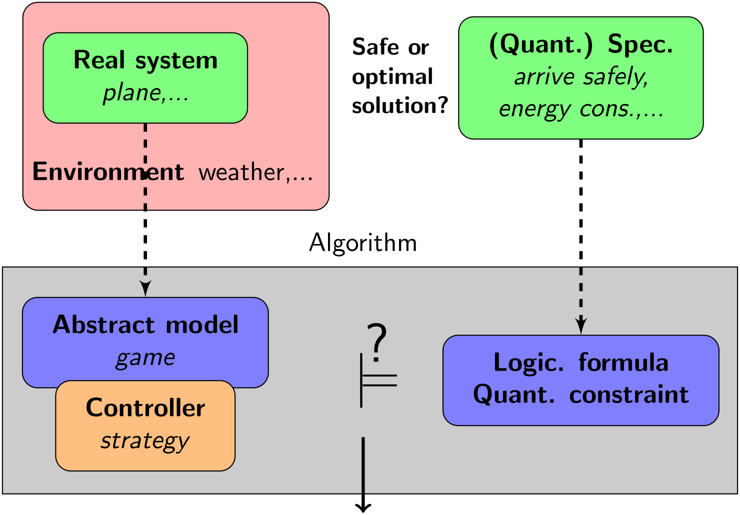

NO/YES + A controller

The autopilot case

Requirement: to arrive safely **in every weather condition**,

The autopilot case

Requirement: to arrive safely **in every weather condition**, **taking into account the other planes**,

The autopilot case

Requirement: to arrive safely **in every weather condition**, **taking into account the other planes**, **while minimising the fuel consumption**.

Controlling complex interactive computer systems

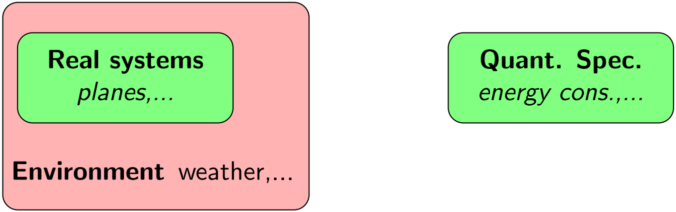

Controlling complex interactive computer systems

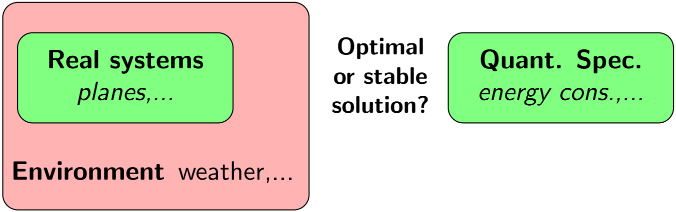

Controlling complex interactive computer systems

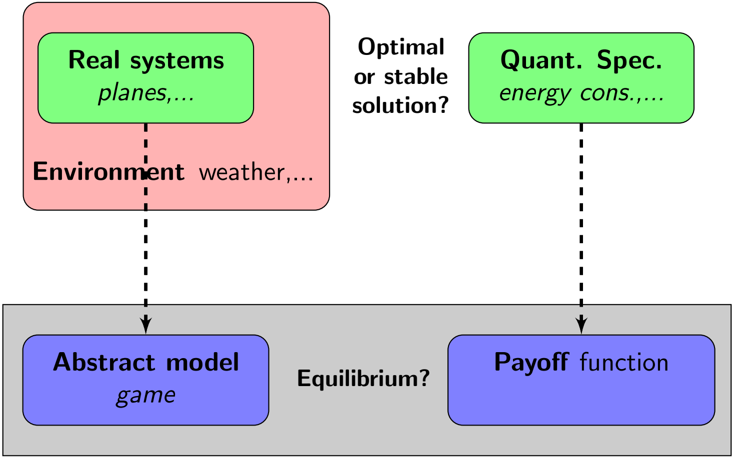

Controlling complex interactive computer systems

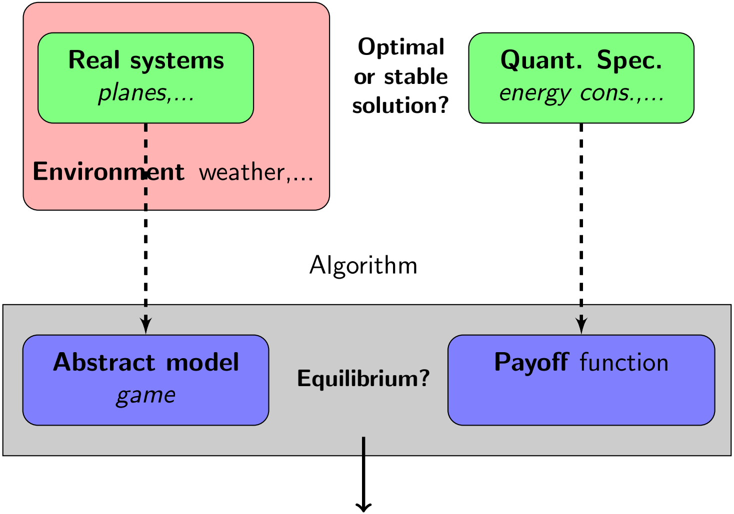

NO/YES + An equilibrium

Controlling complex interactive computer systems

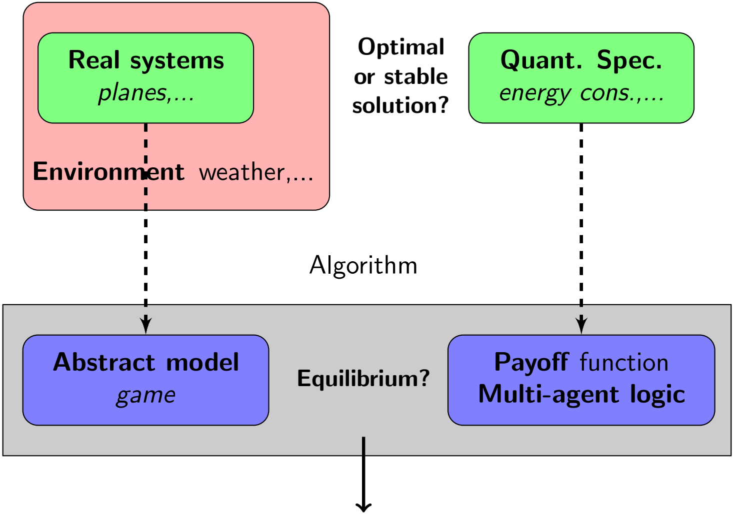

NO/YES + An equilibrium

Outline

Which games do we need for verification?

Methodology

- Pick standard models used in model-checking
- Expand them with interaction capabilities

$\rightsquigarrow$ Games played on graphs

  - Several features in the graph: stochastic or deterministic
  - Several options for interaction: turn-based *vs* concurrent, pure *vs* mixed strategies

Which games do we need for verification?

Methodology

- Pick standard models used in model-checking
- Expand them with interaction capabilities

$\rightsquigarrow$ Games played on graphs

  - Several features in the graph: stochastic or deterministic
  - Several options for interaction: turn-based *vs* concurrent, pure *vs* mixed strategies

The Nim game modelled as a turn-based game

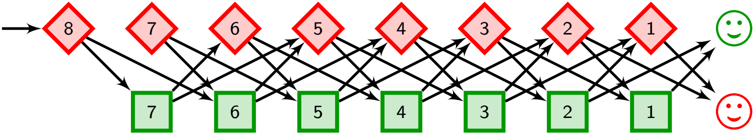

Which games do we need for verification?

Methodology

- Pick standard models used in model-checking
- Expand them with interaction capabilities

$\rightsquigarrow$ Games played on graphs

- Several features in the graph: stochastic or deterministic
- Several options for interaction: turn-based *vs* concurrent, pure *vs* mixed strategies

The Nim game modelled as a turn-based game

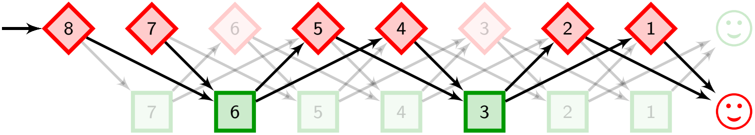

This is then just a matter of computing winning states (controller synthesis)

Outline

Multiplayer stochastic concurrent games

- Graph with stochastic nodes
- Multiple players: $\mathsf{Agt} = \{A_1, A_2, A_3, \dots\}$
- Concurrent moves: $\mathtt{a}_1\mathtt{a}_2\mathtt{a}_3 \cdots \in \Sigma^{\mathsf{Agt}}$ means that player A_1 played $\mathtt{a}_1$, player A_2 played $\mathtt{a}_2$ and player A_3 played $\mathtt{a}_3$, ...
- Payoff functions $\mathsf{payoff}_A : V^\omega \to \mathbb{R}$ for every $A \in \mathsf{Agt}$

A simple model for the medium access control problem [KNPS19]

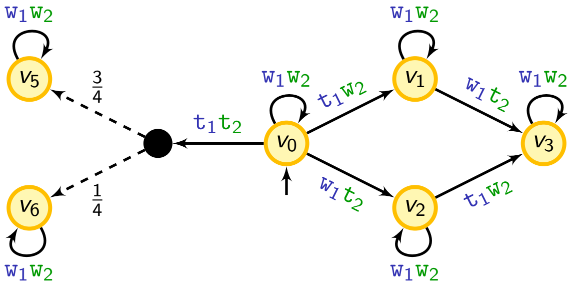

[KNPS19] Kwiatkowska, Norman, Parker, Santos. Equilibria-based probabilistic model checking for concurrent stochastic games *(FM'19)*.

How do we play those games?

According to strategies!

How do we play those games?

According to strategies!

What kind of strategies?

How do we play those games?

According to strategies!

What kind of strategies?

Mixed strategies

$\sigma_A : V^* \to \mathsf{Dist}(\Sigma)$

After history $h \in V^*$, player A will play each action $b \in \Sigma$ with probability $\sigma_A(h)(b)$.

How do we play those games?

According to strategies!

What kind of strategies?

Mixed strategies
$\sigma_A : V^* \to \mathsf{Dist}(\Sigma)$

Deterministic strategies
$\sigma_A : V^* \to \Sigma$

For every $h \in V^*$, $\sigma_A(h)$ is a Dirac measure.

How do we play those games?

According to strategies!

What kind of strategies?

Mixed strategies
$\sigma_A : V^* \to \mathrm{Dist}(\Sigma)$

Deterministic strategies
$\sigma_A : V^* \to \Sigma$

Stationary strategies
$\sigma_A : V \to \mathrm{Dist}(\Sigma)$

If $h, h' \in V^*$ are s.t. $last(h) = last(h')$, then $\sigma_A(h) = \sigma_A(h')$.

How do we play those games?

According to strategies!

What kind of strategies?

Mixed strategies
$\sigma_A : V^* \to \mathrm{Dist}(\Sigma)$

Deterministic strategies
$\sigma_A : V^* \to \Sigma$

Stationary strategies
$\sigma_A : V \to \mathrm{Dist}(\Sigma)$

Memoryless strategies
$\sigma_A : V \to \Sigma$

How do we play those games?

According to strategies!

What kind of strategies?

Mixed strategies
$\sigma_A : V^* \to \mathsf{Dist}(\Sigma)$

Deterministic strategies
$\sigma_A : V^* \to \Sigma$

Stationary strategies
$\sigma_A : V \to \mathsf{Dist}(\Sigma)$

Memoryless strategies
$\sigma_A : V \to \Sigma$

Strategy profile $\sigma = (\sigma_A)_{A \in \mathsf{Agt}}$

An example

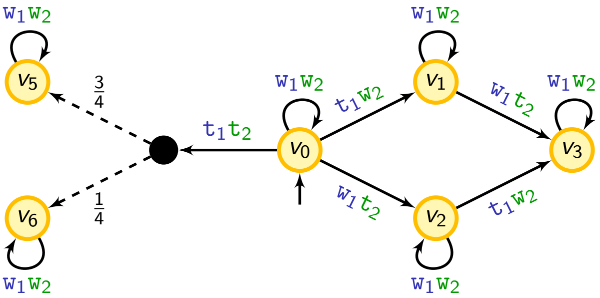

w1w2
v5
3/4
1/4
v6
w1w2
t1t2
w1w2
v0
t1w2
w1t2
v1
w1w2
v2
w1w2
w1t2
t1w2
v3
w1w2

An example

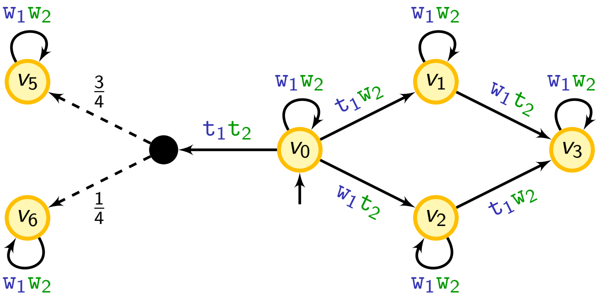

Strategy for player A_i: $\sigma_{A_i}(h) = \begin{cases} \frac{1}{3} \cdot \mathtt{t}_i + \frac{2}{3} \cdot \mathtt{w}_i & \text{if } \mathtt{t}_i \text{ available} \\ \mathtt{w}_i & \text{otherwise} \end{cases}$

An example

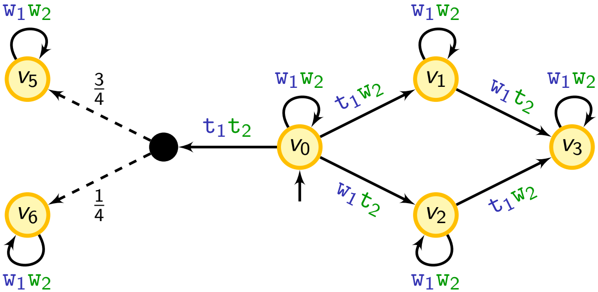

Strategy for player A_i: $\sigma_{A_i}(h) = \begin{cases} \frac{1}{3}\cdot \mathtt{t}_i + \frac{2}{3}\cdot \mathtt{w}_i & \text{if } \mathtt{t}_i \text{ available} \\ \mathtt{w}_i & \text{otherwise} \end{cases}$

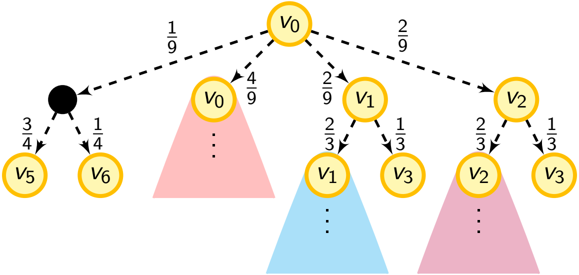

Payoffs

Given strategy profile $\sigma = (\sigma_A)_{A \in \mathsf{Agt}}$, the benefit $p_A(\sigma)$ of player A from v_0 is given by:

$$p_A(\sigma) = \mathbb{E}^{\sigma}_{v_0}(\mathsf{payoff}_A)$$

Payoffs

Given strategy profile $\sigma = (\sigma_A)_{A \in \mathsf{Agt}}$, the benefit $p_A(\sigma)$ of player A from v_0 is given by:

$$p_A(\sigma) = \mathbb{E}^{\sigma}_{v_0}(\mathsf{payoff}_A)$$

Examples

- $\phi_A \subseteq V^\omega$, and for $\rho \in V^\omega$,

$$\mathsf{payoff}_A(\rho) = \begin{cases} 1 & \text{if } \rho \models \phi_A \\ 0 & \text{otherwise} \end{cases}$$

  Then, $p_A(\sigma) = \mathbb{P}^{\sigma}_{v_0}(\phi_A)$.

Payoffs

Given strategy profile $\sigma = (\sigma_A)_{A \in \mathsf{Agt}}$, the benefit $p_A(\sigma)$ of player A from v_0 is given by:

$$p_A(\sigma) = \mathbb{E}^{\sigma}_{v_0}(\text{payoff}_A)$$

Examples

- $\phi_A \subseteq V^\omega$, and for $\rho \in V^\omega$,

$$\text{payoff}_A(\rho) = \begin{cases} 1 & \text{if } \rho \models \phi_A \\ 0 & \text{otherwise} \end{cases}$$

  Then, $p_A(\sigma) = \mathbb{P}^{\sigma}_{v_0}(\phi_A)$.

- payoff_A is a quantitative function on V^ω, for instance:
  - a mean-payoff function
  - a terminal-reward function

Subclasses of interest

- Turn-based games: V partitioned into all V_{A_i}'s

Subclasses of interest

- Turn-based games: V partitioned into all V_{A_i}'s

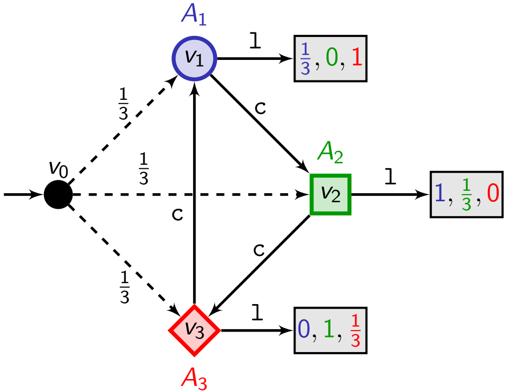

Subclasses of interest

- Turn-based games: V partitioned into all V_{A_i}'s

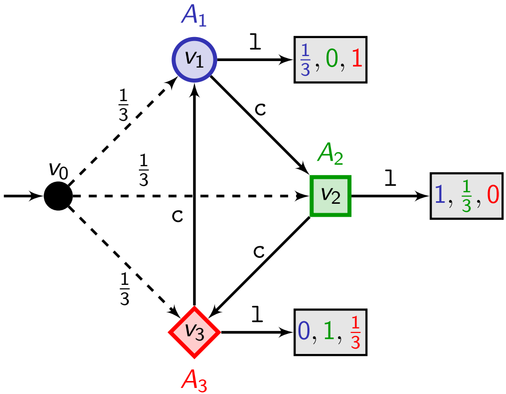

- Deterministic games

Subclasses of interest

- Turn-based games: V partitioned into all V_{A_i}'s

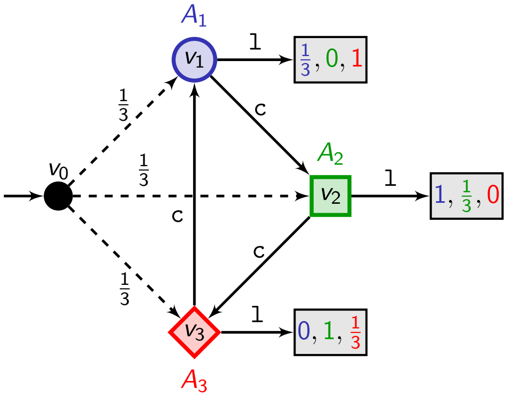

- Deterministic games
  If σ is pure and the game is deterministic, then profile σ has a single outcome $\text{out}(\sigma)$, and

$$p_A(\sigma) = \text{payoff}_A(\text{out}(\sigma))$$

Nash equilibrium in this setting

Nash equilibrium

A mixed (resp. pure) strategy profile $\sigma = (\sigma_A)_{A \in \mathsf{Agt}}$ is a mixed (resp. pure) Nash equilibrium if no player can improve her payoff by unilaterally changing her strategy, that is, for every $A \in \mathsf{Agt}$, for every mixed (resp. pure) deviation σ'_A,

$$\mathbb{E}^{\sigma}_{v_0}(\mathsf{payoff}_A) \geq \mathbb{E}^{\sigma[A/\sigma'_A]}_{v_0}(\mathsf{payoff}_A)$$

Nash equilibrium in this setting

Nash equilibrium

A mixed (resp. pure) strategy profile $\sigma = (\sigma_A)_{A \in \mathsf{Agt}}$ is a mixed (resp. pure) Nash equilibrium if no player can improve her payoff by unilaterally changing her strategy, that is, for every $A \in \mathsf{Agt}$, for every mixed (resp. pure) deviation σ'_A,

$$\mathbb{E}^{\sigma}_{v_0}(\mathsf{payoff}_A) \geq \mathbb{E}^{\sigma[A/\sigma'_A]}_{v_0}(\mathsf{payoff}_A)$$

Example

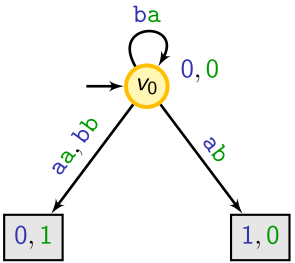

aa (that is, $\sigma_{A_i}(v_0) = \mathtt{a}$) is a (pure) Nash equilibrium

Nash equilibrium in this setting

Nash equilibrium

A mixed (resp. pure) strategy profile $\sigma = (\sigma_A)_{A \in \mathsf{Agt}}$ is a mixed (resp. pure) Nash equilibrium if no player can improve her payoff by unilaterally changing her strategy, that is, for every $A \in \mathsf{Agt}$, for every mixed (resp. pure) deviation σ'_A,

$$\mathbb{E}^{\sigma}_{v_0}(\mathsf{payoff}_A) \geq \mathbb{E}^{\sigma[A/\sigma'_A]}_{v_0}(\mathsf{payoff}_A)$$

Example – Matching penny

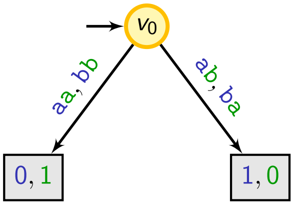

$\sigma_{A_i}(v_0) = \frac{1}{2} \cdot \mathsf{a} + \frac{1}{2} \cdot \mathsf{b}$ is the unique (mixed) Nash equilibrium

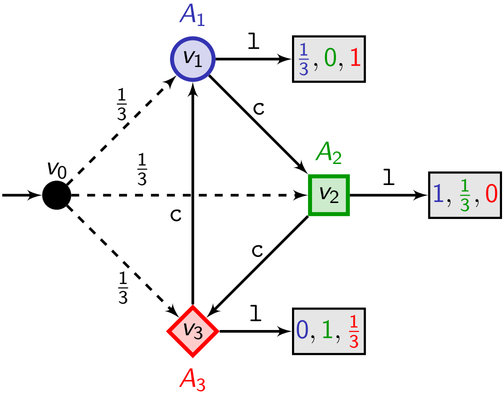

A_1
v_1
l
$\frac{1}{3}, 0, 1$
c
$\frac{1}{3}$
v_0
$\frac{1}{3}$
A_2
v_2
l
$1, \frac{1}{3}, 0$
c
c
$\frac{1}{3}$
v_3
l
$0, 1, \frac{1}{3}$
A_3

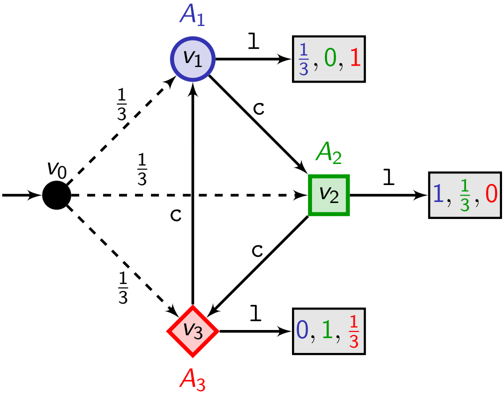

- There is no stationary Nash equilibrium

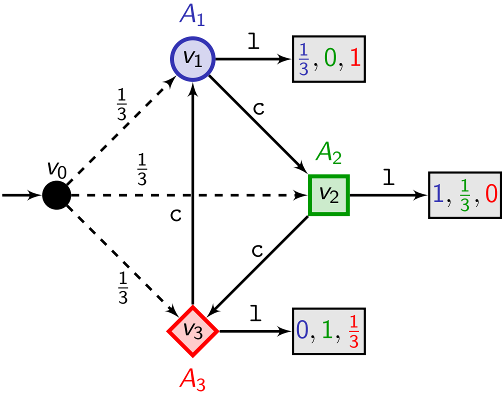

- There is no stationary Nash equilibrium
- There is a pure Nash equilibrium:
  - $v_0 v_i \mapsto$ c
  - $v_0 v_{i+1} \mapsto$ l
  - $v_0 v_i h \mapsto$ c

  It has payoff $(\frac{4}{9}, \frac{4}{9}, \frac{4}{9})$.

Problems of interest

- Universal existence:

  Does there exist a Nash equilibrium in all games?

Problems of interest

- Universal existence:

  Does there exist a Nash equilibrium in all games?

- Existence problem:

  Given a game, does there exist a Nash equilibrium?

Problems of interest

- Universal existence:

  Does there exist a Nash equilibrium in all games?

- Existence problem:

  Given a game, does there exist a Nash equilibrium?

- Constrained existence problem:

  Given a game, does there exist a Nash equilibrium which satisfies some constraint?

Problems of interest

- Universal existence:

  Does there exist a Nash equilibrium in all games?

- Existence problem:

  Given a game, does there exist a Nash equilibrium?

- Constrained existence problem:

  Given a game, does there exist a Nash equilibrium which satisfies some constraint?

- Synthesis of witness (simple?) profiles?

  Do strategy profiles require randomness? Memory?

Problems of interest

- Universal existence:

  Does there exist a Nash equilibrium in all games?

- Existence problem:

  Given a game, does there exist a Nash equilibrium?

- Constrained existence problem:

  Given a game, does there exist a Nash equilibrium which satisfies some constraint?

- Synthesis of witness (simple?) profiles?

  Do strategy profiles require randomness? Memory?

Does the standard theory apply?

- Nash theorem does not apply (requires a finite number of pure strategies)

Does the standard theory apply?

- Nash theorem does not apply (requires a finite number of pure strategies)
- But do the related fixed point theorems apply?

Does the standard theory apply?

- Nash theorem does not apply (requires a finite number of pure strategies)
- But do the related fixed point theorems apply?

Kakutani's fixpoint theorem

Let X be a non-empty, compact and convex subset of $\mathbb{R}^n$. Let $f : X \to 2^X$ be a set-valued function on X with a closed graph and the property that $f(x)$ is non-empty and convex for all $x \in X$. Then f has a fixpoint.

Does the standard theory apply?

- Nash theorem does not apply (requires a finite number of pure strategies)
- But do the related fixed point theorems apply?

Kakutani's fixpoint theorem

Let X be a non-empty, compact and convex subset of $\mathbb{R}^n$. Let $f : X \to 2^X$ be a set-valued function on X with a closed graph and the property that $f(x)$ is non-empty and convex for all $x \in X$. Then f has a fixpoint.

- Usually it applies to the best-response operator: if $\sigma \in \mathbb{S}$ ($\mathbb{S}$ is for stationary profiles), then

$$\mathsf{BR}(\sigma) = \left\{ \sigma' \in \mathbb{S} \mid \forall A \in \mathsf{Agt},\ \sigma'_A \in \mathrm{argmax}_{\sigma''_A \in \mathbb{S}_A} \mathbb{E}_{v_0}^{\sigma[A/\sigma''_A]}(\mathsf{payoff}_A) \right\}$$

Does the standard theory apply?

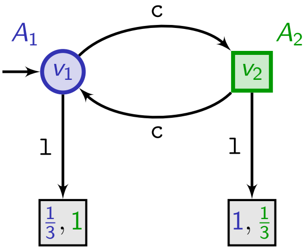

Does the standard theory apply?

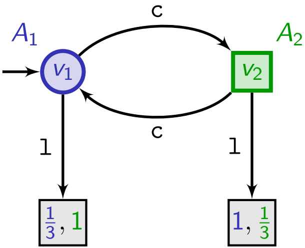

The first who leaves the loop loses!

Does the standard theory apply?

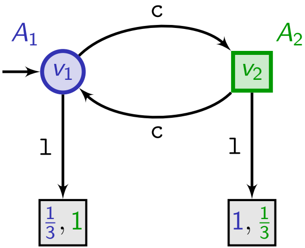

We note $(x_1, x_2) \in [0,1]^2$ for the profile σ s.t.

$$\begin{cases} \sigma_{A_1}(v_1) & = x_1 \cdot \mathtt{l} + (1 - x_1) \cdot \mathtt{c} \\ \sigma_{A_2}(v_2) & = x_2 \cdot \mathtt{l} + (1 - x_2) \cdot \mathtt{c} \end{cases}$$

The first who leaves the loop loses!

Does the standard theory apply?

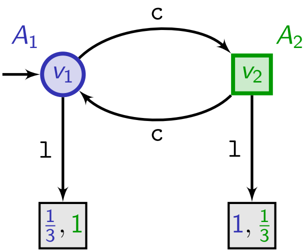

We note $(x_1, x_2) \in [0,1]^2$ for the profile σ s.t.

$$\begin{cases} \sigma_{A_1}(v_1) & = x_1 \cdot \mathtt{l} + (1 - x_1) \cdot \mathtt{c} \\ \sigma_{A_2}(v_2) & = x_2 \cdot \mathtt{l} + (1 - x_2) \cdot \mathtt{c} \end{cases}$$

The first who leaves the loop loses!

- For every $x_1, x_2 > 0$, $\mathrm{BR}((x_1, x_2)) = (0, 0)$

Does the standard theory apply?

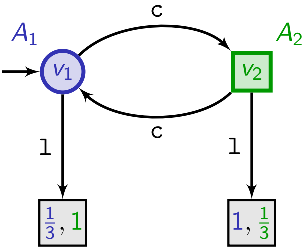

We note $(x_1, x_2) \in [0,1]^2$ for the profile σ s.t.

$$\begin{cases} \sigma_{A_1}(v_1) & = x_1 \cdot \mathtt{l} + (1 - x_1) \cdot \mathtt{c} \\ \sigma_{A_2}(v_2) & = x_2 \cdot \mathtt{l} + (1 - x_2) \cdot \mathtt{c} \end{cases}$$

The first who leaves the loop loses!

- For every $x_1, x_2 > 0$, $\mathrm{BR}((x_1, x_2)) = (0, 0)$
- $\mathrm{BR}((0, 0)) = \{(x_1, x_2) \mid x_1, x_2 > 0\}$

Does the standard theory apply?

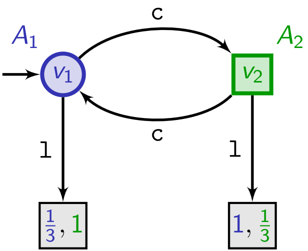

We note $(x_1, x_2) \in [0,1]^2$ for the profile σ s.t.

$$\begin{cases} \sigma_{A_1}(v_1) & = x_1 \cdot \mathtt{l} + (1 - x_1) \cdot \mathtt{c} \\ \sigma_{A_2}(v_2) & = x_2 \cdot \mathtt{l} + (1 - x_2) \cdot \mathtt{c} \end{cases}$$

The first who leaves the loop loses!

- For every $x_1, x_2 > 0$, $\mathrm{BR}((x_1, x_2)) = (0, 0)$
- $\mathrm{BR}((0, 0)) = \{(x_1, x_2) \mid x_1, x_2 > 0\}$
- The graph of BR is not closed

Does the standard theory apply? No!

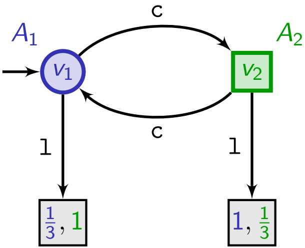

We note $(x_1, x_2) \in [0,1]^2$ for the profile σ s.t.

$$\begin{cases} \sigma_{A_1}(v_1) & = x_1 \cdot \mathtt{l} + (1 - x_1) \cdot \mathtt{c} \\ \sigma_{A_2}(v_2) & = x_2 \cdot \mathtt{l} + (1 - x_2) \cdot \mathtt{c} \end{cases}$$

The first who leaves the loop loses!

- For every $x_1, x_2 > 0$, $\mathrm{BR}((x_1, x_2)) = (0, 0)$
- $\mathrm{BR}((0, 0)) = \{(x_1, x_2) \mid x_1, x_2 > 0\}$
- The graph of BR is not closed
- Kakutani's theorem does not apply

Does the standard theory apply? No!

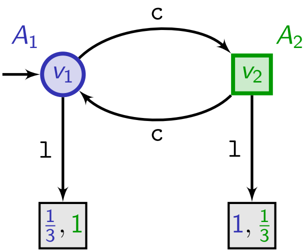

We note $(x_1, x_2) \in [0,1]^2$ for the profile σ s.t.

$$\begin{cases} \sigma_{A_1}(v_1) & = x_1 \cdot \mathtt{l} + (1 - x_1) \cdot \mathtt{c} \\ \sigma_{A_2}(v_2) & = x_2 \cdot \mathtt{l} + (1 - x_2) \cdot \mathtt{c} \end{cases}$$

The first who leaves the loop loses!

- For every $x_1, x_2 > 0$, $\mathrm{BR}((x_1, x_2)) = (0, 0)$
- $\mathrm{BR}((0,0)) = \{(x_1, x_2) \mid x_1, x_2 > 0\}$
- The graph of BR is not closed
- Kakutani's theorem does not apply

However there are infinitely many Nash equilibria:

all $(0, x_2)$ and $(x_1, 0)$ with $x_1, x_2 > 0$

No universal existence in general!

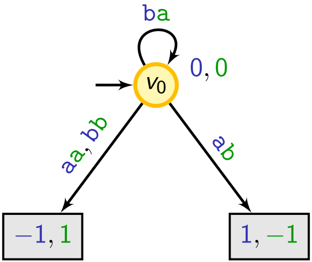

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Limit-Average Games *(CONCUR'11)*.

No universal existence in general!

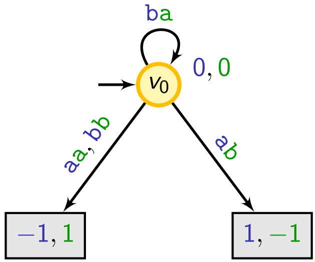

- By playing stationary strategy

$$\sigma_{A_2}(v_0) = (1-\epsilon)\cdot \mathtt{a} + \epsilon \cdot \mathtt{b},$$

A_2 ensures payoff $1 - 2\epsilon$

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Limit-Average Games *(CONCUR'11)*.

No universal existence in general!

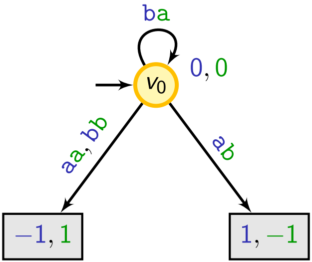

- By playing stationary strategy

$$\sigma_{A_2}(v_0) = (1-\epsilon)\cdot \mathtt{a} + \epsilon \cdot \mathtt{b},$$

  A_2 ensures payoff $1 - 2\epsilon$
- Hence any Nash equilibrium would have payoff $(-1, 1)$

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Limit-Average Games *(CONCUR'11)*.

No universal existence in general!

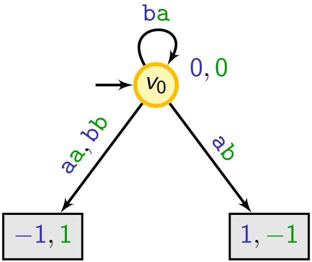

- By playing stationary strategy

$$\sigma_{A_2}(v_0) = (1-\epsilon)\cdot \mathtt{a} + \epsilon \cdot \mathtt{b},$$

  A_2 ensures payoff $1-2\epsilon$
- Hence any Nash equilibrium would have payoff $(-1, 1)$
- If A_2 plays a forever, then A_1 will play b forever, yielding payoff $(0, 0)$, which is not a Nash equilibrium

No universal existence in general!

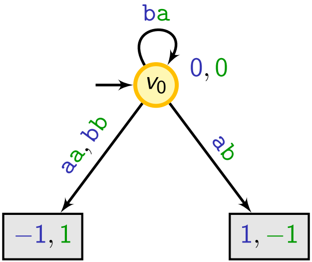

- By playing stationary strategy

$$\sigma_{A_2}(v_0) = (1-\epsilon)\cdot \mathtt{a} + \epsilon \cdot \mathtt{b},$$

  A_2 ensures payoff $1 - 2\epsilon$
- Hence any Nash equilibrium would have payoff $(-1, 1)$
- If A_2 plays a forever, then A_1 will play b forever, yielding payoff $(0, 0)$, which is not a Nash equilibrium
- If A_2 plays b with some positive probability p at some round (first time this occurs), then by playing b before and a at that precise round, A_1 can ensure payoff $p > 0$

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Limit-Average Games *(CONCUR'11)*.

No universal existence in general!

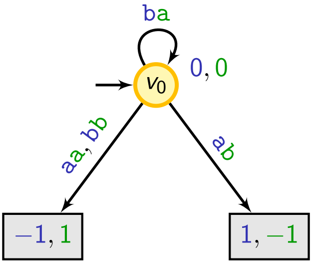

- By playing stationary strategy

$$\sigma_{A_2}(v_0) = (1-\epsilon)\cdot \mathtt{a} + \epsilon \cdot \mathtt{b},$$

  A_2 ensures payoff $1-2\epsilon$
- Hence any Nash equilibrium would have payoff $(-1, 1)$
- If A_2 plays a forever, then A_1 will play b forever, yielding payoff $(0, 0)$, which is not a Nash equilibrium
- If A_2 plays b with some positive probability p at some round (first time this occurs), then by playing b before and a at that precise round, A_1 can ensure payoff $p > 0$

$\rightsquigarrow$ There is no Nash equilibrium!

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Limit-Average Games *(CONCUR'11)*.

Outline

We focus on a simple scenario

Restrictions

- Turn-based games
- Payoffs given by ω-regular objectives: ϕ_A objective of player $A \in \mathrm{Agt}$
- Pure strategy profiles

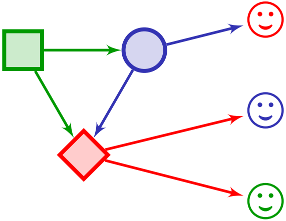

We focus on a simple scenario

Restrictions

- Turn-based games
- Payoffs given by ω-regular objectives: ϕ_A objective of player $A \in \mathrm{Agt}$
- Pure strategy profiles

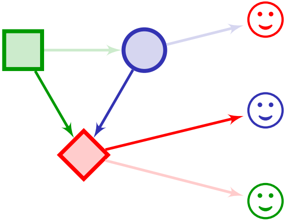

is a Nash equilibrium with payoff $(0, 1, 0)$

We focus on a simple scenario

Restrictions

- Turn-based games
- Payoffs given by ω-regular objectives: ϕ_A objective of player $A \in \mathrm{Agt}$
- Pure strategy profiles

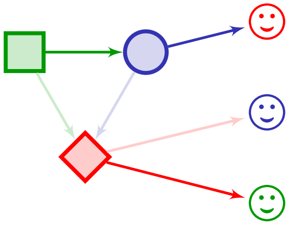

is not a Nash equilibrium

A simple characterization for ω-regular objectives

Player $\square$ loses along that play

ϕ_A: objective of player A

A simple characterization for ω-regular objectives

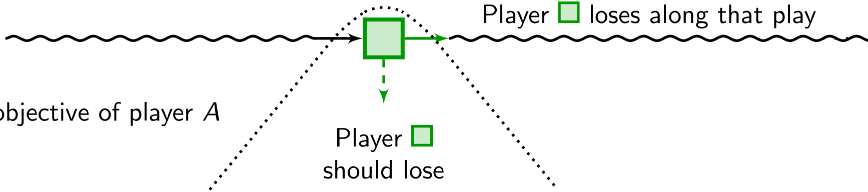

ϕ_A: objective of player A

A simple characterization for ω-regular objectives

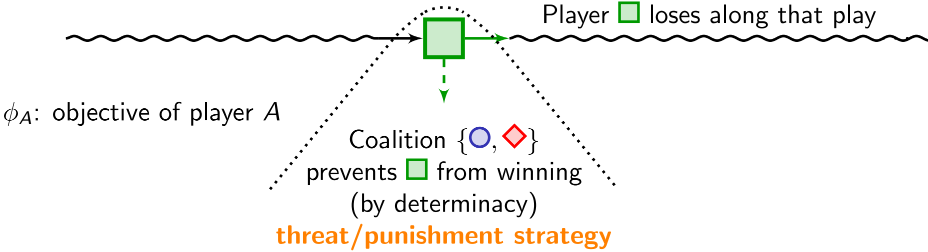

A simple characterization for ω-regular objectives

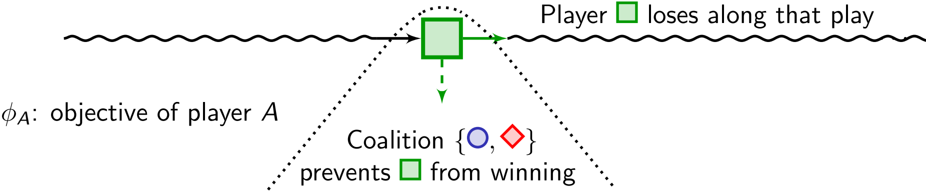

$$\neg\phi_{\square} \Rightarrow \mathbf{G}(p_{\square} \Rightarrow \mathbf{X}\, W_{\{\circ,\diamond\}})$$

where $p_{\square}$ labels $\square$-states and $W_{\{\circ,\diamond\}}$ is the set of winning states for the coalition $\{\circ, \diamond\}$ for winning objective $\neg\phi_{\square}$.

A simple characterization for ω-regular objectives

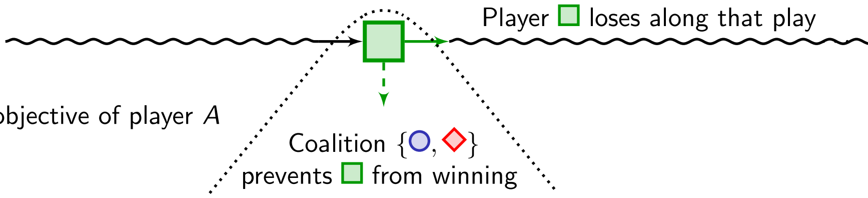

ϕ_A: objective of player A

Main outcomes of Boolean Nash equilibria in turn-based games can be characterized by an LTL formula:

$$\Phi_{\mathsf{NE}} = \bigwedge_{A \in \mathsf{Agt}} \Big(\neg \phi_A \Rightarrow \mathbf{G}(p_A \Rightarrow \mathbf{X}\, W_{\{-A\}}) \Big)$$

where p_A labels A-states and $W_{\{-A\}}$ is the set of winning states for the coalition $\{-A\} \stackrel{\text{def}}{=} \mathsf{Agt} \setminus \{A\}$ against A for the objective $\neg\phi_A$. These sets should be precomputed.

A simple characterization for ω-regular objectives

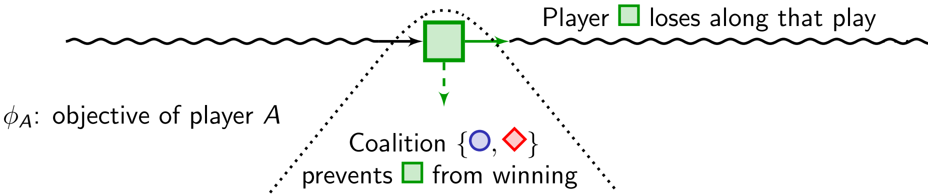

Main outcomes of Boolean Nash equilibria in turn-based games can be characterized by an LTL formula:

$$\Phi_{\mathsf{NE}} = \bigwedge_{A \in \mathsf{Agt}} \Big(\neg\phi_A \Rightarrow \mathbf{G}(p_A \Rightarrow \mathbf{X}\, W_{\{-A\}})\Big)$$

where p_A labels A-states and $W_{\{-A\}}$ is the set of winning states for the coalition $\{-A\} \stackrel{\text{def}}{=} \mathsf{Agt} \setminus \{A\}$ against A for the objective $\neg\phi_A$. These sets should be precomputed.

(valid for prefix-independent objectives)

Decidability of the constrained existence problem

Constrained existence problem

Given two thresholds $L, U \in \mathbb{Q}^+$, does there exists a Nash equilibrium σ such that for every $A \in \mathsf{Agt}$:

$$L_A \leq \mathbb{E}^{\sigma}_{v_0}(\mathsf{payoff}_A) \leq U_A?$$

[Umm08] Ummels. The Complexity of Nash Equilibria in Infinite Multiplayer Games *(FoSSaCS'08)*.

Decidability of the constrained existence problem

Constrained existence problem

Given two thresholds $L, U \in \mathbb{Q}^+$, does there exists a Nash equilibrium σ such that for every $A \in \mathsf{Agt}$:

$$L_A \leq \mathbb{E}^{\sigma}_{v_0}(\mathsf{payoff}_A) \leq U_A?$$

Theorem [Umm08]

One can decide the *pure* constrained existence problem in finite turn-based multiplayer games for ω-regular objectives.

Examples of complexity results for single objectives:

Objectives	Reach.	Safety	Büchi	co-Büchi	Parity
Complexity	NP-c.		P-c.	NP-c.	

Note: it extends to "ω-regular" preference relations with a finite image.

[Umm08] Ummels. The Complexity of Nash Equilibria in Infinite Multiplayer Games *(FoSSaCS'08)*.

An example of NP-hardness result

By reduction from a SAT instance:

$$\varphi = \bigwedge_{1 \leq i \leq n} C_i \qquad \text{with } C_i = \bigvee_{j=1}^{3} \ell_{i,j} \qquad \ell_{i,j} \in \{x_1, \neg x_1, x_2, \neg x_2, \ldots, x_k, \neg x_k\}$$

An example of NP-hardness result

By reduction from a SAT instance:

$$\varphi = \bigwedge_{1 \le i \le n} C_i \qquad \text{with } C_i = \bigvee_{j=1}^{3} \ell_{i,j} \qquad \ell_{i,j} \in \{x_1, \neg x_1, x_2, \neg x_2, \ldots, x_k, \neg x_k\}$$

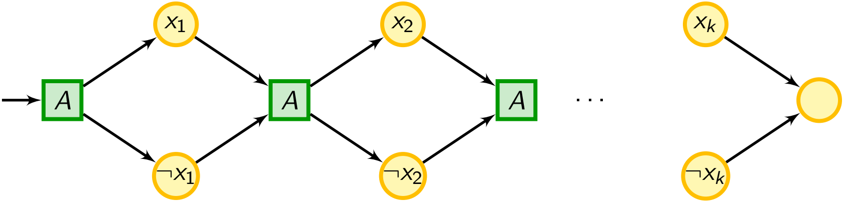

- Player A_i for clause C_i, with objective to reach $\{\ell_{i,j} \mid j = 1, 2, 3\}$
- Player A: reach the rightmost state

An example of NP-hardness result

By reduction from a SAT instance:

$$\varphi = \bigwedge_{1 \le i \le n} C_i \qquad \text{with } C_i = \bigvee_{j=1}^{3} \ell_{i,j} \qquad \ell_{i,j} \in \{x_1, \neg x_1, x_2, \neg x_2, \ldots, x_k, \neg x_k\}$$

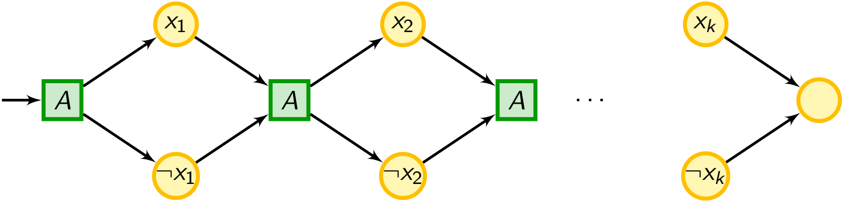

- Player A_i for clause C_i, with objective to reach $\{\ell_{i,j} \mid j = 1, 2, 3\}$
- Player A: reach the rightmost state

> φ is satisfiable iff there is a Nash equilibrium with payoff 1 for everyone in the game

The universal existence problem: ω-regular objectives

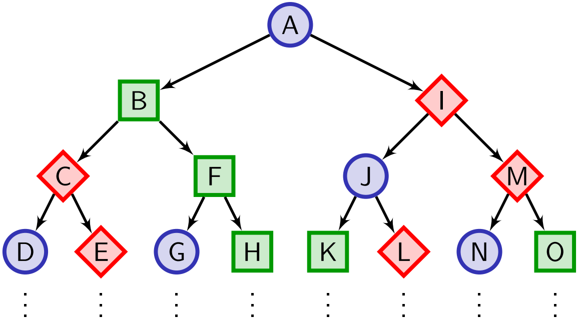

The universal existence problem: ω-regular objectives

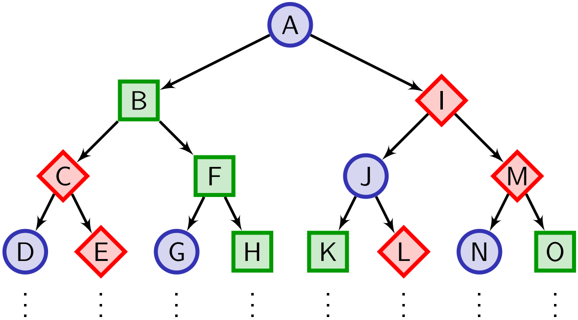

- If ◯ has a winning strategy from A, then ◯ should play it forever
- Otherwise ◯ plays any strategy, until (by chance) a new blue node, for instance J, is visited, from which ◯ has a winning strategy; ◯ then switches to such a winning strategy, forever

The universal existence problem: ω-regular objectives

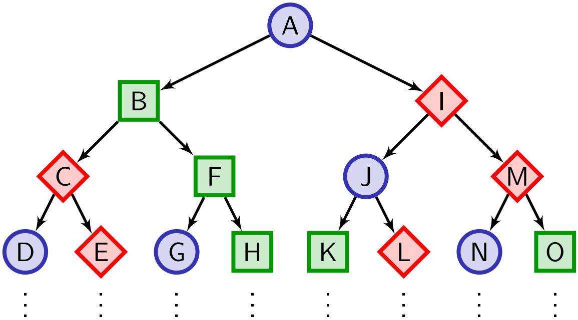

- If the game proceeds through B and □ has a winning strategy from B, then □ should play it forever
- If the game proceeds through B but □ has no winning strategy from B, then □ should play any strategy, until (by chance) a new green node, for instance H, is visited, from which □ has a winning strategy; □ then switches to such a winning strategy, forever

The universal existence problem: ω-regular objectives

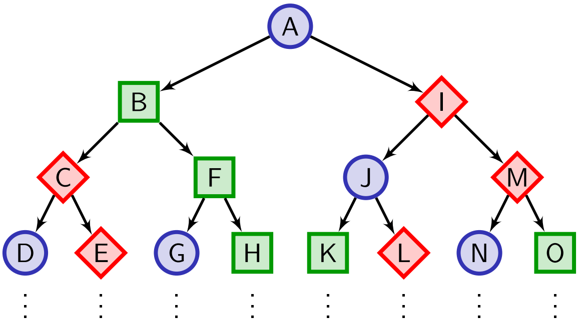

- If the game proceeds through C and ◇ has a winning strategy from C, then ◇ should play it forever
- If the game proceeds through C but ◇ has no winning strategy from C, then ◇ should play any strategy, until (by chance) a new red node, for instance E, is visited, from which ◇ has a winning strategy; ◇ then switches to such a winning strategy, forever

The universal existence problem: ω-regular objectives

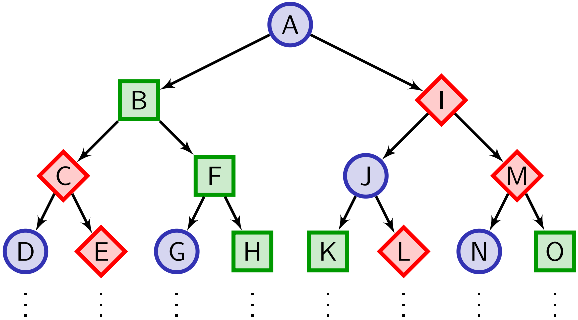

- Outside the main outcome, all players play the adequate threat or punishment strategy: this is the coalition strategy that makes the deviator lose (NB: determinacy required!)

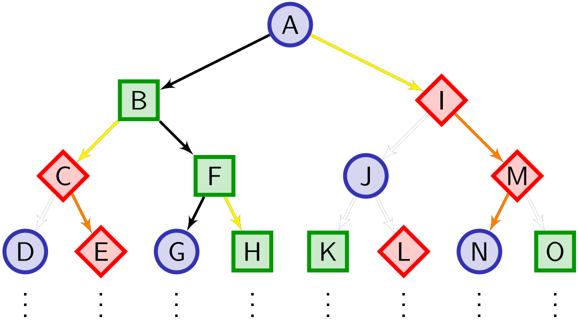
A
B
I
C
F
J
M
D
E
G
H
K
L
N
O

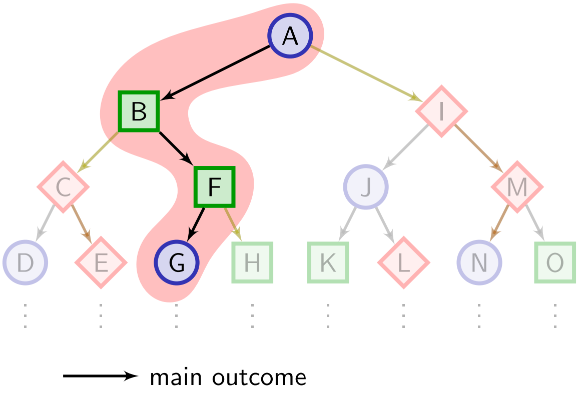
A
B
I
C
F
J
M
D
E
G
H
K
L
N
O
main outcome

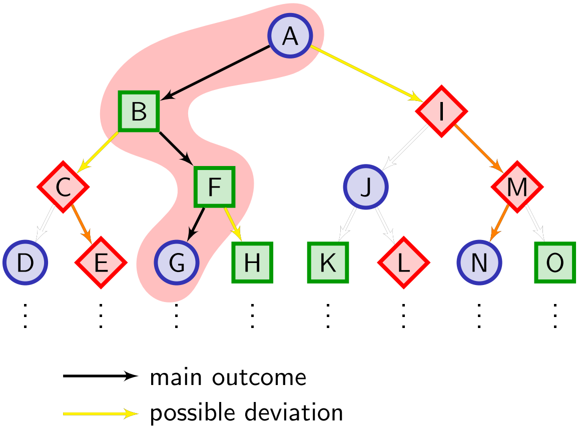
A
B
I
C
F
J
M
D
E
G
H
K
L
N
O
main outcome
possible deviation

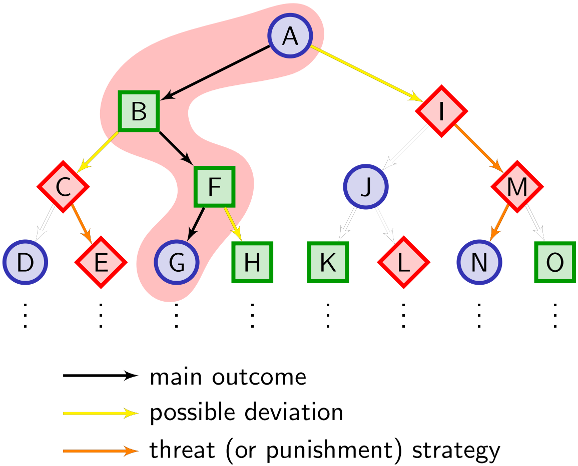
A
B
I
C
F
J
M
D
E
G
H
K
L
N
O
main outcome
possible deviation
threat (or punishment) strategy

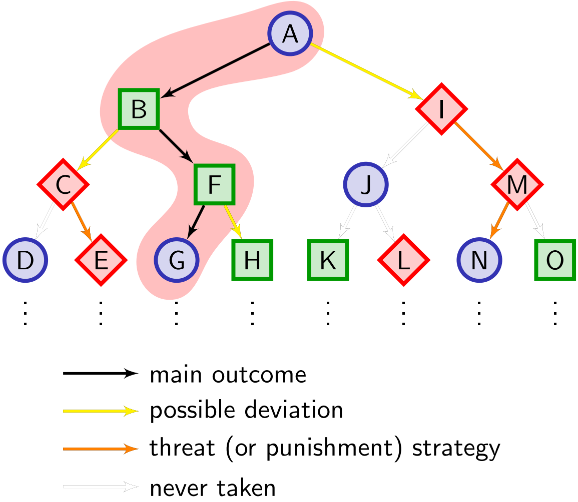
A
B
I
C
F
J
M
D
E
G
H
K
L
N
O
main outcome
possible deviation
threat (or punishment) strategy
never taken

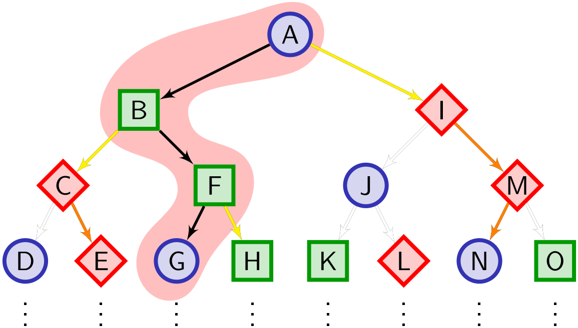

main outcome

possible deviation

threat (or punishment) strategy

never taken

Questions:

- why is it correct?

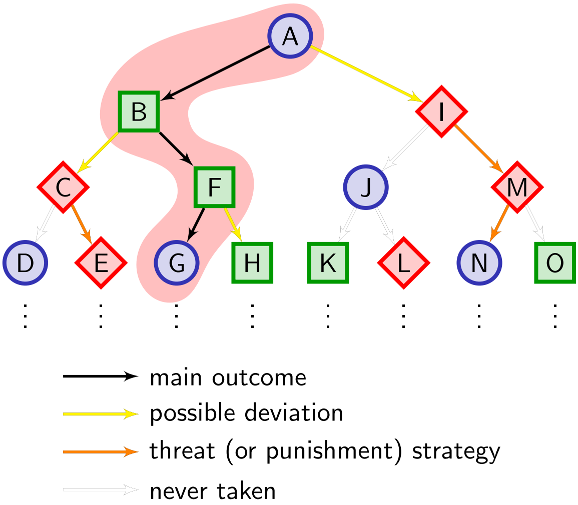

Questions:

- why is it correct?
- what immediate extension can be handled?

The universal existence problem: ω-regular objectives

Universal existence [Umm11]

In infinite-duration turn-based deterministic games on finite graphs with ω-regular objectives, there is always a pure Nash equilibrium. Moreover, one can compute a witness.

[Umm11] Ummels. Stochastic multiplayer games: theory and algorithms *(PhD thesis)*.
[LeR13] Le Roux. Infinite sequential Nash equilibrium *(LMCS)*.

The universal existence problem

Universal existence [Umm11]

In infinite-duration turn-based deterministic games on finite graphs with ω-regular objectives, there is always a pure Nash equilibrium. Moreover, one can compute a witness.

Universal existence [LeR13]

In infinite turn-based deterministic games with Borel measurable countable preferences, with no ascending infinite chains, there is always a pure Nash equilibrium.

[Umm11] Ummels. Stochastic multiplayer games: theory and algorithms *(PhD thesis)*.
[LeR13] Le Roux. Infinite sequential Nash equilibrium *(LMCS)*.

Outline

1. Verification and game theory
2. Games on graphs
   - The general model
   - Focus on a simple scenario
   - Adding probabilities to the setting?
   - Concurrent games
3. Conclusion

Stochastic turn-based games

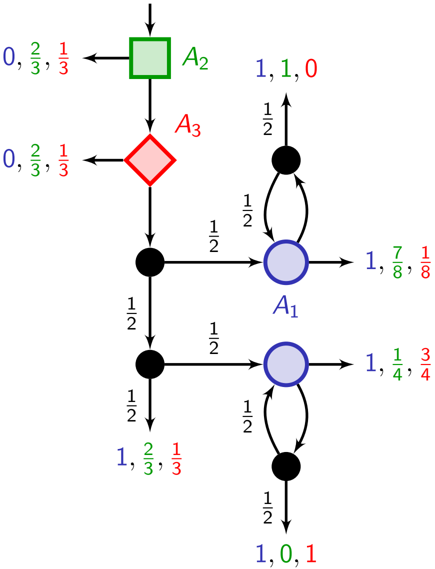

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Stochastic Multiplayer Games *(LMCS)*.

Stochastic turn-based games

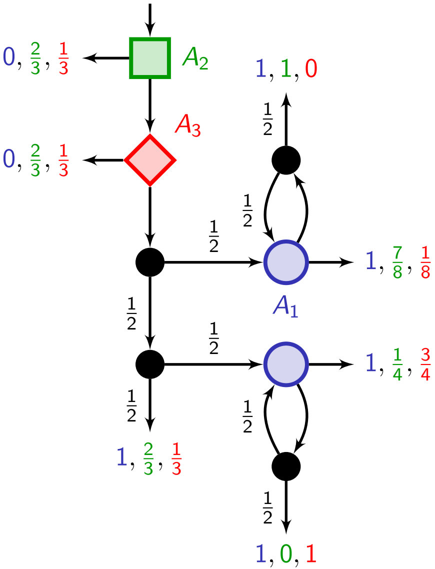

Along a Nash equilibrium where $p_{A_1} \geq 1$:

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Stochastic Multiplayer Games *(LMCS)*.

Stochastic turn-based games

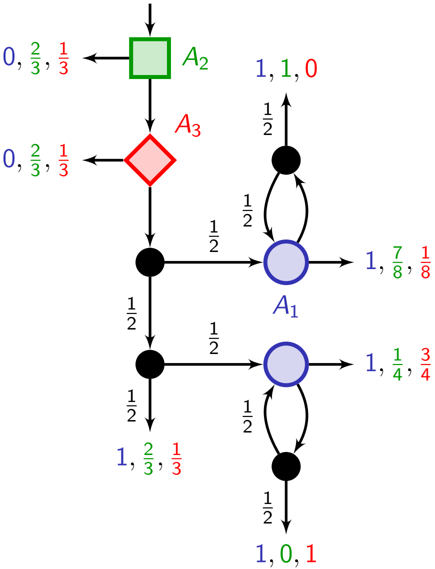

Along a Nash equilibrium where $p_{A_1} \geq 1$:

- $p_{A_2} + p_{A_3} = 1$

Stochastic turn-based games

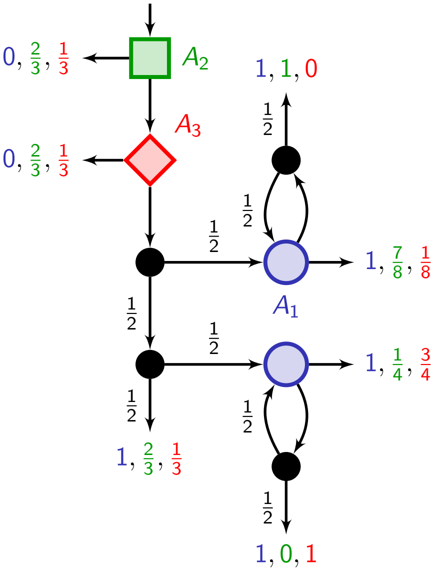

Along a Nash equilibrium where $p_{A_1} \geq 1$:

- $p_{A_2} + p_{A_3} = 1$
- $p_{A_2} \geq \frac{2}{3}$ and $p_{A_3} \geq \frac{1}{3}$

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Stochastic Multiplayer Games *(LMCS)*.

Stochastic turn-based games

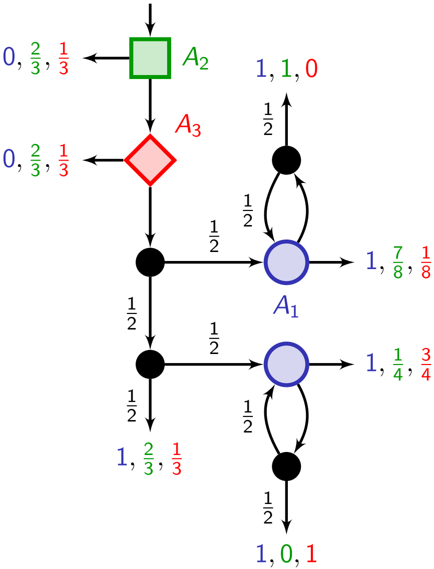

Along a Nash equilibrium where $p_{A_1} \geq 1$:

- $p_{A_2} + p_{A_3} = 1$
- $p_{A_2} = \frac{2}{3}$ and $p_{A_3} = \frac{1}{3}$

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Stochastic Multiplayer Games *(LMCS)*.

Stochastic turn-based games

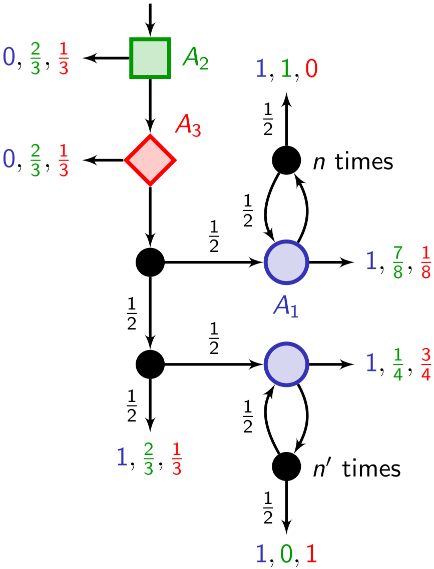

Along a Nash equilibrium where $p_{A_1} \geq 1$:

- $p_{A_2} + p_{A_3} = 1$
- $p_{A_2} = \frac{2}{3}$ and $p_{A_3} = \frac{1}{3}$
- $p_{A_2} = \frac{2}{3} + \frac{1}{16}\left(\frac{1}{2^{n'}} - \frac{1}{2^n}\right)$

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Stochastic Multiplayer Games *(LMCS)*.

Stochastic turn-based games

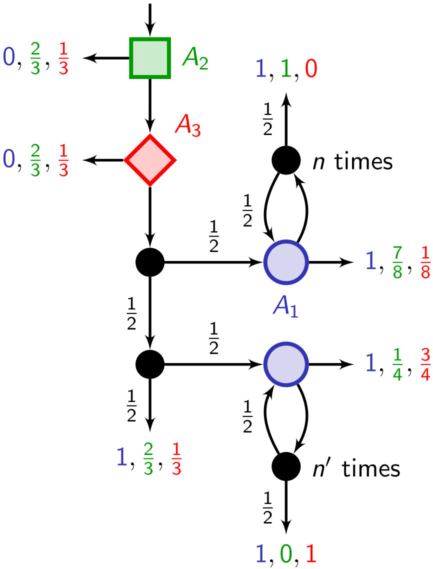

Along a Nash equilibrium where $p_{A_1} \geq 1$:

- $p_{A_2} + p_{A_3} = 1$
- $p_{A_2} = \frac{2}{3}$ and $p_{A_3} = \frac{1}{3}$
- $p_{A_2} = \frac{2}{3} + \frac{1}{16}\left(\frac{1}{2^{n'}} - \frac{1}{2^n}\right)$

$\rightsquigarrow$ $n = n'$

Stochastic turn-based games

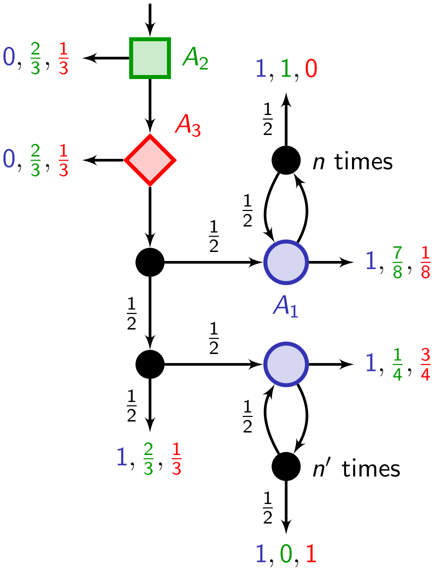

Along a Nash equilibrium where $p_{A_1} \geq 1$:

- $p_{A_2} + p_{A_3} = 1$
- $p_{A_2} = \frac{2}{3}$ and $p_{A_3} = \frac{1}{3}$
- $p_{A_2} = \frac{2}{3} + \frac{1}{16}\left(\frac{1}{2^{n'}} - \frac{1}{2^n}\right)$

$\rightsquigarrow$ $n = n'$

One can simulate a two-counter machine if we constrain $p_{A_1} \geq 1$!!

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Stochastic Multiplayer Games *(LMCS)*.

Stochastic turn-based games

Undecidability results [UW11]

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Stochastic Multiplayer Games *(LMCS)*.

Stochastic turn-based games

Undecidability results [UW11]

- The constrained existence problem for pure strategies in stochastic turn-based games is undecidable.

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Stochastic Multiplayer Games *(LMCS)*.

Stochastic turn-based games

Undecidability results [UW11]

- The constrained existence problem for pure strategies in stochastic turn-based games is undecidable.
- The constrained existence problem for mixed strategies in deterministic turn-based games is undecidable.

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Stochastic Multiplayer Games *(LMCS)*.

Short summary for turn-based ω-regular games

[UW11,Umm11,LeR13]

- There always exists a Nash equilibrium for Boolean ω-regular objectives
- One can decide the constrained existence of a Nash equilibrium (and compute one!)
- One cannot decide the existence of a mixed (i.e. stochastic) Nash equilibrium

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Stochastic Multiplayer Games *(LMCS)*
[Umm11] Ummels. Stochastic multiplayer games: theory and algorithms *(PhD thesis, RWTH Aachen University)*
[LeR13] Le Roux. Infinite sequential Nash equilibrium *(LMCS)*.

Short summary for turn-based ω-regular games

[UW11,Umm11,LeR13]

- There always exists a Nash equilibrium for Boolean ω-regular objectives
- One can decide the constrained existence of a Nash equilibrium (and compute one!)
- One cannot decide the existence of a mixed (i.e. stochastic) Nash equilibrium

$\rightsquigarrow$ this is why we will restrict to pure equilibria in det. games

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Stochastic Multiplayer Games *(LMCS)*
[Umm11] Ummels. Stochastic multiplayer games: theory and algorithms *(PhD thesis, RWTH Aachen University)*
[LeR13] Le Roux. Infinite sequential Nash equilibrium *(LMCS)*.

Outline

Can this theory be extended to concurrent games?

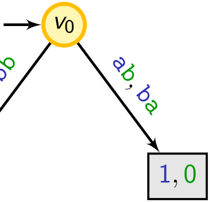

There is no universal existence, even for simple Boolean objectives.

Can this theory be extended to concurrent games?

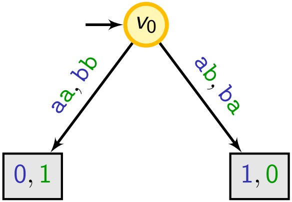

There is no universal existence, even for simple Boolean objectives.

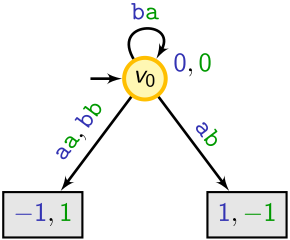

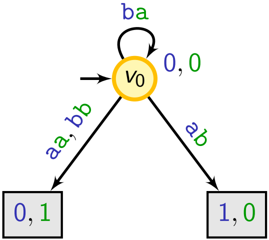

There is no pure Nash equilibrium

There is a pure Nash equilibrium

Can this theory be extended to concurrent games?

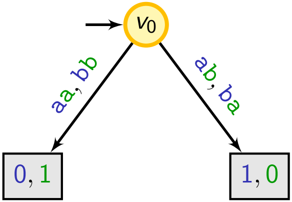

There is no universal existence, even for simple Boolean objectives.

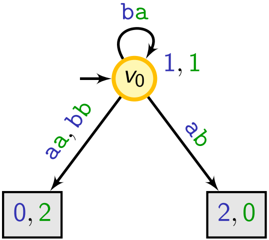

There is no pure Nash equilibrium

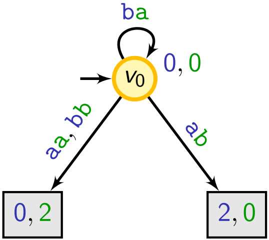

There is a pure Nash equilibrium

Existence becomes NP-hard

Hardness

The existence problem is NP-hard for reachability objectives.

By reduction from a SAT instance:

$$\varphi = \bigwedge_{1 \leq i \leq n} C_i \qquad \text{with } C_i = \bigvee_{j=1}^{3} \ell_{i,j} \qquad \ell_{i,j} \in \{x_1, \neg x_1, x_2, \neg x_2, \ldots, x_k, \neg x_k\}$$

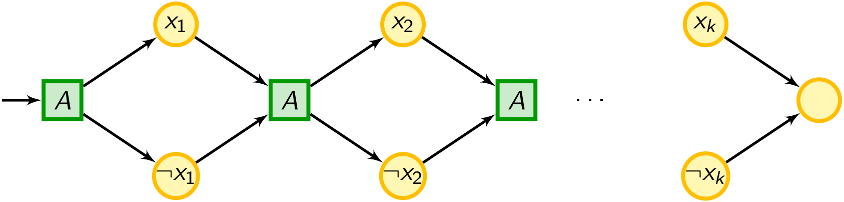

φ is satisfiable iff there is a Nash equilibrium with payoff 1 for everyone in the game

Existence becomes NP-hard

Hardness

The existence problem is NP-hard for reachability objectives.

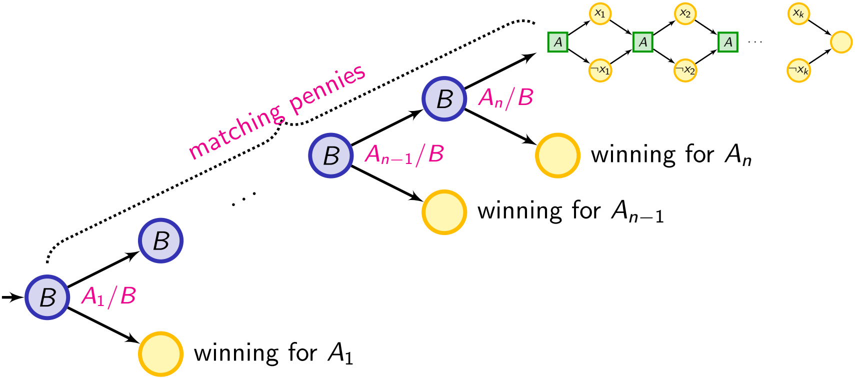

φ is satisfiable iff there is a Nash equilibrium in the game

Who is a suspect? Who knows what?

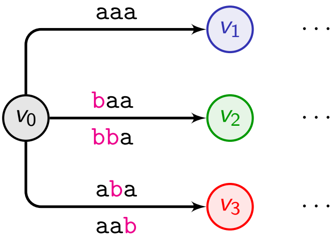

Who is a suspect? Who knows what?

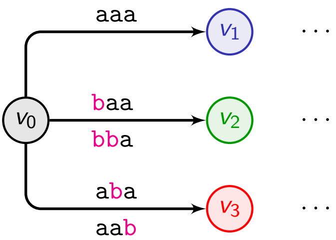

Assume that the normal move is $v_0 \xrightarrow{\texttt{aaa}} v_1$

- what does that mean if the game proceeds to v_2?
- what does that mean if the game proceeds to v_3?

Who is a suspect? Who knows what?

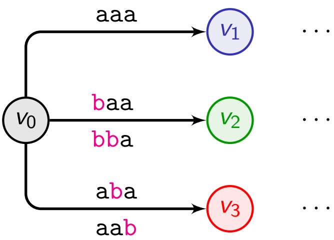

Assume that the normal move is $v_0 \xrightarrow{\texttt{aaa}} v_1$

- what does that mean if the game proceeds to v_2?
- what does that mean if the game proceeds to v_3?
  - either player A_2 deviated alone (playing b instead of a);

Who is a suspect? Who knows what?

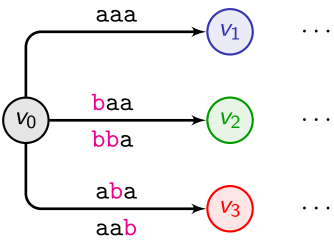

Assume that the normal move is $v_0 \xrightarrow{\texttt{aaa}} v_1$

- what does that mean if the game proceeds to v_2?
- what does that mean if the game proceeds to v_3?
  - either player A_2 deviated alone (playing b instead of a);
  - or A_3 deviated alone (playing b instead of a).

Who is a suspect? Who knows what?

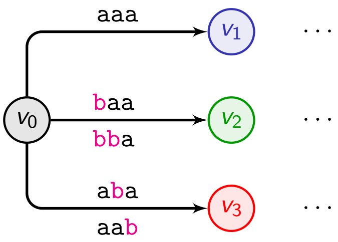

Assume that the normal move is $v_0 \xrightarrow{\texttt{aaa}} v_1$

- what does that mean if the game proceeds to v_2?
  - either player A_1 deviated alone (playing b instead of a);
- what does that mean if the game proceeds to v_3?
  - either player A_2 deviated alone (playing b instead of a);
  - or A_3 deviated alone (playing b instead of a).

Who is a suspect? Who knows what?

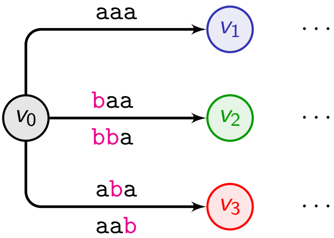

Assume that the normal move is $v_0 \xrightarrow{\texttt{aaa}} v_1$

- what does that mean if the game proceeds to v_2?
  - either player A_1 deviated alone (playing b instead of a);
  - or both players A_1 and A_2 played b instead of a.
- what does that mean if the game proceeds to v_3?
  - either player A_2 deviated alone (playing b instead of a);
  - or A_3 deviated alone (playing b instead of a).

Who is a suspect? Who knows what?

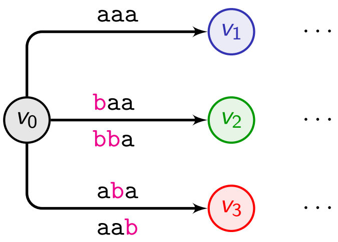

Assume that the normal move is $v_0 \xrightarrow{\texttt{aaa}} v_1$

- what does that mean if the game proceeds to v_2?
  - either player A_1 deviated alone (playing b instead of a);
  - ~~or both players A_1 and A_2 played b instead of a.~~
- what does that mean if the game proceeds to v_3?
  - either player A_2 deviated alone (playing b instead of a);
  - or A_3 deviated alone (playing b instead of a).

Who is a suspect? Who knows what?

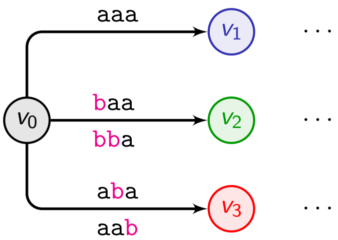

- $\mathrm{susp}\big((v_0, v_2), \mathtt{aaa}\big) = \{A_1\}$
- $\mathrm{susp}\big((v_0, v_3), \mathtt{aaa}\big) = \{A_2, A_3\}$

Assume that the normal move is $v_0 \xrightarrow{\mathtt{aaa}} v_1$

- what does that mean if the game proceeds to v_2?
  - either player A_1 deviated alone (playing b instead of a);
  - ~~or both players A_1 and A_2 played b instead of a.~~
- what does that mean if the game proceeds to v_3?
  - either player A_2 deviated alone (playing b instead of a);
  - or A_3 deviated alone (playing b instead of a).

Who is a suspect? Who knows what?

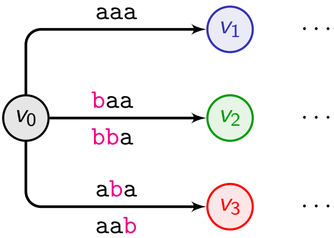

- $\mathrm{susp}\big((v_0, v_2), \mathtt{aaa}\big) = \{A_1\}$
  Everyone knows that A_1 is the deviator
- $\mathrm{susp}\big((v_0, v_3), \mathtt{aaa}\big) = \{A_2, A_3\}$

Assume that the normal move is $v_0 \xrightarrow{\mathtt{aaa}} v_1$

- what does that mean if the game proceeds to v_2?
  - either player A_1 deviated alone (playing b instead of a);
  - ~~or both players A_1 and A_2 played b instead of a.~~
- what does that mean if the game proceeds to v_3?
  - either player A_2 deviated alone (playing b instead of a);
  - or A_3 deviated alone (playing b instead of a).

Who is a suspect? Who knows what?

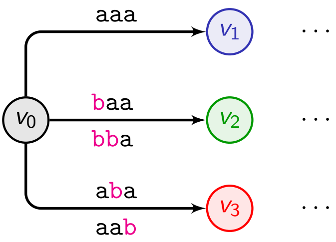

- $\mathrm{susp}\big((v_0, v_2), \mathtt{aaa}\big) = \{A_1\}$
  Everyone knows that A_1 is the deviator
- $\mathrm{susp}\big((v_0, v_3), \mathtt{aaa}\big) = \{A_2, A_3\}$
  A_1 knows that the deviator is either A_2 or A_3; A_2 knows the identity of the deviator; and so does A_3

Assume that the normal move is $v_0 \xrightarrow{\mathtt{aaa}} v_1$

- what does that mean if the game proceeds to v_2?
  - either player A_1 deviated alone (playing b instead of a);
  - ~~or both players A_1 and A_2 played b instead of a.~~
- what does that mean if the game proceeds to v_3?
  - either player A_2 deviated alone (playing b instead of a);
  - or A_3 deviated alone (playing b instead of a).

Who is a suspect? Who knows what?

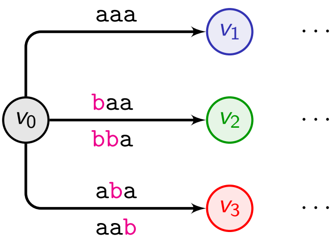

- $\mathrm{susp}\big((v_0, v_2), \mathtt{aaa}\big) = \{A_1\}$
  Everyone knows that A_1 is the deviator
- $\mathrm{susp}\big((v_0, v_3), \mathtt{aaa}\big) = \{A_2, A_3\}$
  A_1 knows that the deviator is either A_2 or A_3; A_2 knows the identity of the deviator; and so does A_3
- $\mathrm{susp}\big((v_0, v_1), \mathtt{aaa}\big) = \{A_1, A_2, A_3\}$

Assume that the normal move is $v_0 \xrightarrow{\mathtt{aaa}} v_1$

- what does that mean if the game proceeds to v_2?
  - either player A_1 deviated alone (playing b instead of a);
  - ~~or both players A_1 and A_2 played b instead of a.~~
- what does that mean if the game proceeds to v_3?
  - either player A_2 deviated alone (playing b instead of a);
  - or A_3 deviated alone (playing b instead of a).

Construction of the suspect game abstraction

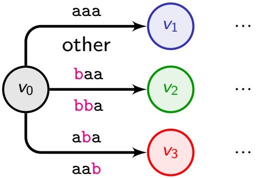

Two players: Eve (light)
Adam (dark)

Construction of the suspect game abstraction

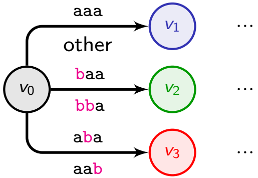

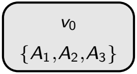

Two players: Eve (light)
Adam (dark)

Construction of the suspect game abstraction

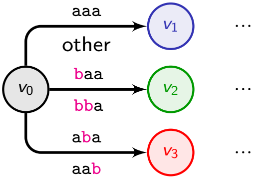

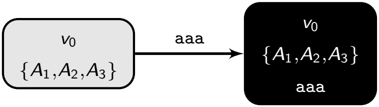

Two players: Eve (light)
Adam (dark)

Construction of the suspect game abstraction

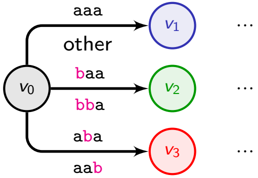

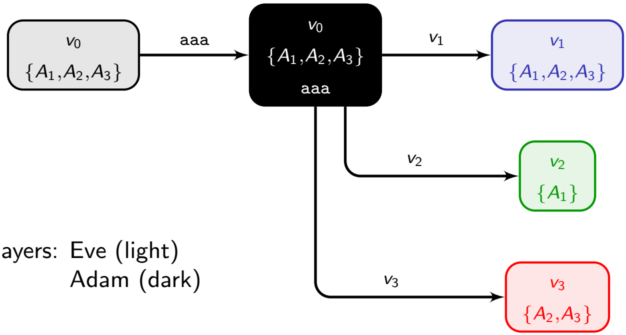

Two players: Eve (light)
Adam (dark)

Construction of the suspect game abstraction

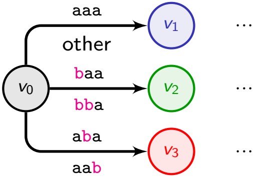

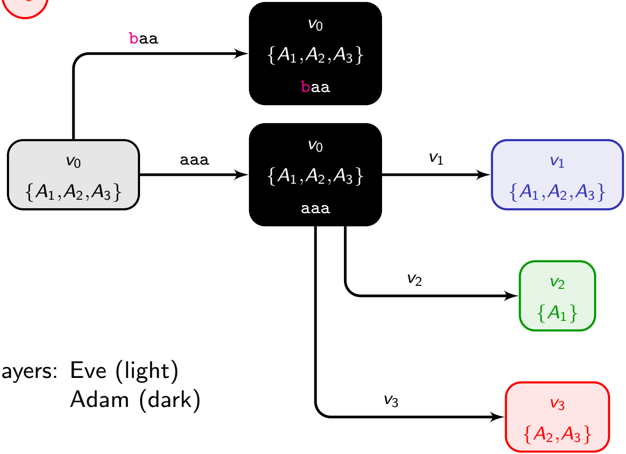

Two players: Eve (light)
Adam (dark)

Construction of the suspect game abstraction

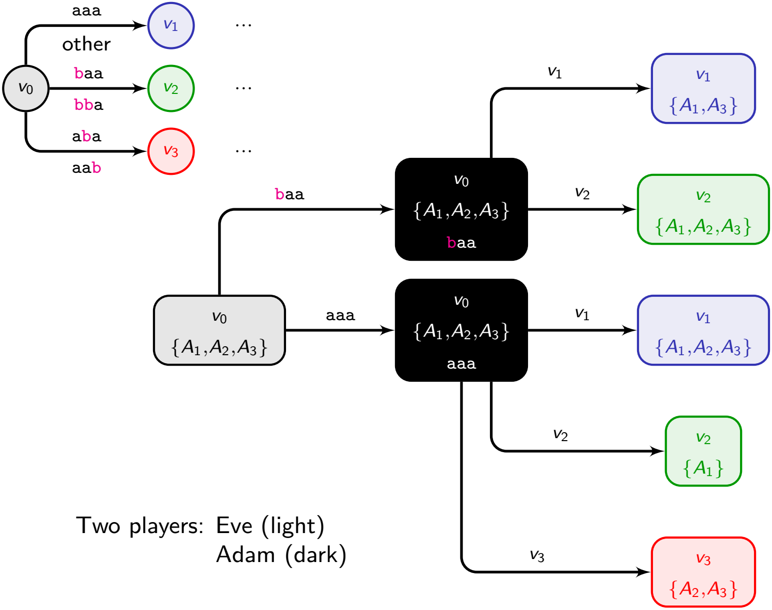

Two players: Eve (light)
Adam (dark)

Construction of the suspect game abstraction

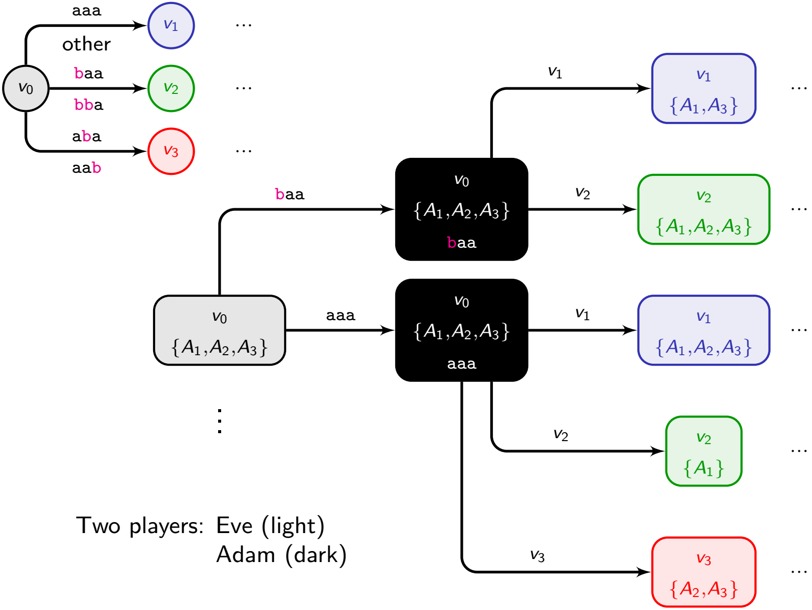

Two players: Eve (light)
Adam (dark)

Correctness of the suspect game construction

Winning condition

A strategy ζ for Eve in the suspect game is winning for some $\alpha \in \mathbb{R}^{\mathsf{Agt}}$ if the unique outcome of ζ where Adam complies to Eve has payoff α, and for every other outcome ρ of ζ, for every $A \in \mathsf{susp}(\rho)$, $\mathsf{payoff}_A(\rho) \leq \alpha_A$.

Correctness of the suspect game construction

Winning condition

A strategy ζ for Eve in the suspect game is winning for some $\alpha \in \mathbb{R}^{\mathsf{Agt}}$ if the unique outcome of ζ where Adam complies to Eve has payoff α, and for every other outcome ρ of ζ, for every $A \in \mathsf{susp}(\rho)$, $\mathsf{payoff}_A(\rho) \leq \alpha_A$.

Correctness

Let $\alpha \in \mathbb{R}^{\mathsf{Agt}}$. There is a Nash equilibrium in the original game with payoff α if and only if Eve has a winning strategy for α in the suspect game.

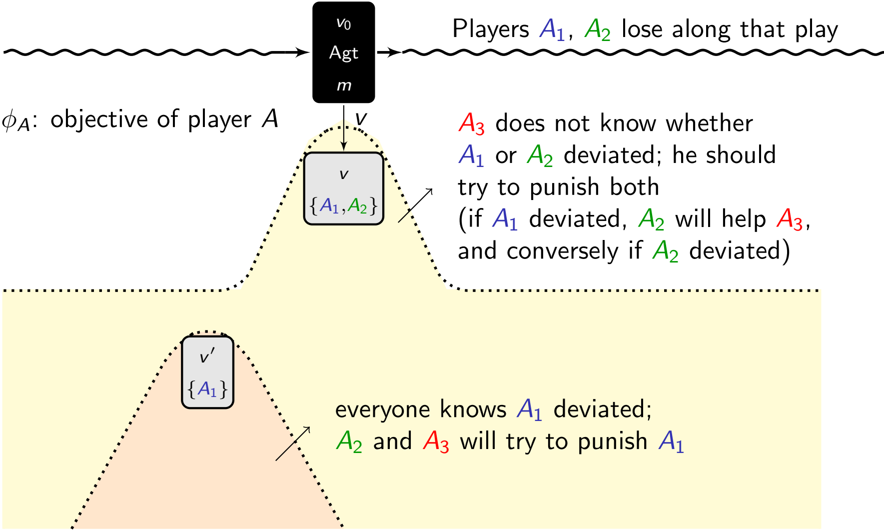
v_0
Agt
m
Players A_1, A_2 lose along that play
ϕ_A: objective of player A
v
v
$\{A_1, A_2\}$
A_3 does not know whether A_1 or A_2 deviated; he should try to punish both (if A_1 deviated, A_2 will help A_3, and conversely if A_2 deviated)
v'
$\{A_1\}$
everyone knows A_1 deviated; A_2 and A_3 will try to punish A_1

From an algorithmic point-of-view

From an algorithmic point-of-view

- In the orange part: compute the winner (Eve or Adam) of the zero-sum game, where Eve's objective is $\neg\phi_{A_1}$ (Eve wants to show that there is no profitable deviation for A_1)

From an algorithmic point-of-view

- In the orange part: compute the winner (Eve or Adam) of the zero-sum game, where Eve's objective is $\neg\phi_{A_1}$ (Eve wants to show that there is no profitable deviation for A_1)
- We remove the orange part, and replace the root vertex by a winning state for the previously computed winner

From an algorithmic point-of-view

- In the orange part: compute the winner (Eve or Adam) of the zero-sum game, where Eve's objective is $\neg\phi_{A_1}$ (Eve wants to show that there is no profitable deviation for A_1)
- We remove the orange part, and replace the root vertex by a winning state for the previously computed winner
- In the yellow part: compute the winner (Eve or Adam) of the zero-sum game, where Eve's objective is

  $$(\neg\phi_{A_1} \wedge \neg\phi_{A_2}) \vee \mathsf{Reach}(\mathsf{win}_{\mathsf{Eve}})$$

  where $\mathsf{win}_{\mathsf{Eve}}$ is an already computed winning state for Eve

From an algorithmic point-of-view

- In the orange part: compute the winner (Eve or Adam) of the zero-sum game, where Eve's objective is $\neg\phi_{A_1}$ (Eve wants to show that there is no profitable deviation for A_1)
- We remove the orange part, and replace the root vertex by a winning state for the previously computed winner
- In the yellow part: compute the winner (Eve or Adam) of the zero-sum game, where Eve's objective is

  $$(\neg\phi_{A_1} \wedge \neg\phi_{A_2}) \vee \mathsf{Reach}(\mathsf{win}_{\mathsf{Eve}})$$

  where $\mathsf{win}_{\mathsf{Eve}}$ is an already computed winning state for Eve
- It is then just a matter to find an infinite play satisfying the appropriate property

From an algorithmic point-of-view

- In the orange part: compute the winner (Eve or Adam) of the zero-sum game, where Eve's objective is $\neg\phi_{A_1}$ (Eve wants to show that there is no profitable deviation for A_1)
- We remove the orange part, and replace the root vertex by a winning state for the previously computed winner
- In the yellow part: compute the winner (Eve or Adam) of the zero-sum game, where Eve's objective is

$$(\neg\phi_{A_1} \wedge \neg\phi_{A_2}) \vee \mathsf{Reach}(\mathsf{win}_{\mathsf{Eve}})$$

  where $\mathsf{win}_{\mathsf{Eve}}$ is an already computed winning state for Eve
- It is then just a matter to find an infinite play satisfying the appropriate property

The approach can be extended to various settings!

Some results

Examples of complexity results

- For single objectives:

Objectives	Reach.	Safety	Büchi	co-Büchi	Parity
Complexity	NP-c.		P-c.	NP-c.	$\mathsf{P}^{\mathsf{NP}}_{\parallel}$-c.

[BBMU15] Bouyer, Brenguier, Markey, Ummels. Pure Nash equilibria in concurrent deterministic games *(LMCS)*

Some results

Examples of complexity results

- For single objectives:

Objectives	Reach.	Safety	Büchi	co-Büchi	Parity
Complexity	NP-c.		P-c.	NP-c.	$\mathsf{P}^{\mathsf{NP}}_{\parallel}$-c.

- For combinations of Büchi objectives:

Combinations	Subset	Lexico.	Count.	Bool. circuit
Complexity	P-c.	NP-.c		PSPACE-c.

[BBMU15] Bouyer, Brenguier, Markey, Ummels. Pure Nash equilibria in concurrent deterministic games *(LMCS)*

Some results

Examples of complexity results

- For single objectives:

Objectives	Reach.	Safety	Büchi	co-Büchi	Parity
Complexity	NP-c.		P-c.	NP-c.	$\mathsf{P}^{\mathsf{NP}}_{\parallel}$-c.

- For combinations of Büchi objectives:

Combinations	Subset	Lexico.	Count.	Bool. circuit
Complexity	P-c.	NP-.c		PSPACE-c.

- For combinations of reachability objectives:

Combinations	Subset	Lexico.	Count.	Bool. circuit
Complexity	NP-c.	PSPACE-.c		

Extensions of this approach

Partial information monitoring

- Public signal [Bou18]
- Communication graphs [BT19]

[Bou18] Bouyer. Games on graphs with a public signal monitoring *(FoSSaCS'18)*.
[BT19] Bouyer, Tomasset. Nash equilibria in games over graphs equipped with a communication mechanism *(MFCS'19)*.
[Bre16] Brenguier. Robust equilibria in mean-payoff games *(FoSSaCS'16)*.
[COT18] Condurache, Oualhadj, Troquard. The complexity of rational synthesis for concurrent games *(CONCUR'18)*.

Extensions of this approach

Partial information monitoring

- Public signal [Bou18]
- Communication graphs [BT19]

Other solution concepts

- Robust equilibria [Bre16]
- Rational synthesis [COT18]

[Bou18] Bouyer. Games on graphs with a public signal monitoring *(FoSSaCS'18)*.
[BT19] Bouyer, Tomasset. Nash equilibria in games over graphs equipped with a communication mechanism *(MFCS'19)*.
[Bre16] Brenguier. Robust equilibria in mean-payoff games *(FoSSaCS'16)*.
[COT18] Condurache, Oualhadj, Troquard. The complexity of rational synthesis for concurrent games *(CONCUR'18)*.

Outline

Wrap-up

General objective

- Import game theory solutions to the verification field, where interactivity plays also a role
  Ex: Distributed systems interacting in some environment

Applications?

- Smart grids: decentralized control of EV charging [GBLM19]
  - stochastic setting
  - ad-hoc approximated solutions
- Cassting project: smart houses that produce energy with solar panels [BDGHM16]
  - deterministic setting
  - setting with universal existence
  - exact computation
- PRISM-games: medium access control, Aloha protocol, robot coordination, power control [KNPS19]
  - stochastic setting
  - approximated value iteration for computing ϵ-SPE

[BDGHM16] Brihaye, Dhar, Geeraerts, Haddad, Monmege. Efficient energy distribution in a smart grid using multi-player games *(Cassting'16)*
[KNPS19] Kwiatkowska, Norman, Parker, Santos. Equilibria-based probabilistic model checking for concurrent stochastic games *(FM'19)*.
[GBLM19] González, Bouyer, Lasaulce, Markey. Optimisation en présence de contraintes en probabilités et processus markoviens contrôlés *(GRETSI'19)*

Wrap-up

General objective

- Import game theory solutions to the verification field, where interactivity plays also a role

  Ex: Distributed systems interacting in some environment
- Relevant questions:
  - assumptions made in the game theory field relevant?
  - solution concepts adapted to the context?

Wrap-up

General objective

- Import game theory solutions to the verification field, where interactivity plays also a role

  Ex: Distributed systems interacting in some environment
- Relevant questions:
  - assumptions made in the game theory field relevant?
  - solution concepts adapted to the context?

Nash equilibria in games on graphs

- The setting of pure Nash equilibria in turn-based det. games rather well-understood
- Probabilistic setting much more complicated
- Concurrent games: a rather generic approach based on the suspect game construction

Going further?

More relevant solution concepts?

- Temporal aspects weakens the concept of Nash equilibrium:
  Will a rational agent/process focus on punishing a deviator, instead of pursuing her own objective?
- Another solution concept: subgame-perfect equilibrium (SPE)

[Umm06] Ummels. Rational Behaviour and Strategy Construction in Infinite Multiplayer Games *(FSTTCS'06)*
[LP14] Le Roux, Pauly. Infinite sequential games with real-valued payoffs *(CSL-LICS'14)*

Going further?

More relevant solution concepts?

- Temporal aspects weakens the concept of Nash equilibrium:
  Will a rational agent/process focus on punishing a deviator, instead of pursuing her own objective?
- Another solution concept: subgame-perfect equilibrium (SPE)

Universal existence [Umm06]

For Boolean parity (even Borel) objectives, there is always a pure SPE.

A simple example with no SPE [LP14]

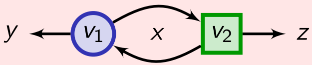

- A_1: $x >_{A_1} y >_{A_1} z$
- A_2: $y >_{A_2} z >_{A_2} x$

[Umm06] Ummels. Rational Behaviour and Strategy Construction in Infinite Multiplayer Games *(FSTTCS'06)*
[LP14] Le Roux, Pauly. Infinite sequential games with real-valued payoffs *(CSL-LICS'14)*

Going further?

A simple example with no SPE [LP14]

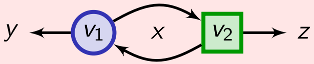

- A_1: $x >_{A_1} y >_{A_1} z$
- A_2: $y >_{A_2} z >_{A_2} x$

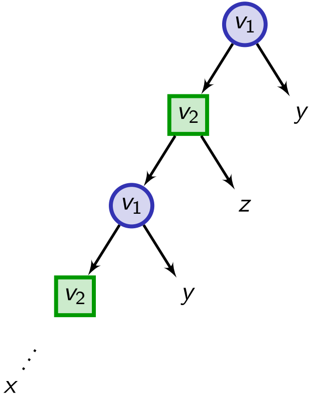

Going further?

More relevant solution concepts?

- Temporal aspects weakens the concept of Nash equilibrium:
  Will a rational agent/process focus on punishing a deviator, instead of pursuing her own objective?
- Another solution concept: subgame-perfect equilibrium (SPE)

[Umm06] Ummels. Rational Behaviour and Strategy Construction in Infinite Multiplayer Games *(FSTTCS'06)*
[LP14] Le Roux, Pauly. Infinite sequential games with real-valued payoffs *(CSL-LICS'14)*