Nash equilibria in games on graphs with public signal monitoring

Patricia Bouyer-Decitre

LSV, CNRS & ENS Paris-Saclay, France

What this talk is about

- pure Nash equilibria in game graphs
- imperfect information monitoring
- public signals
- computability of Nash equilibria

Two-player turn-based games

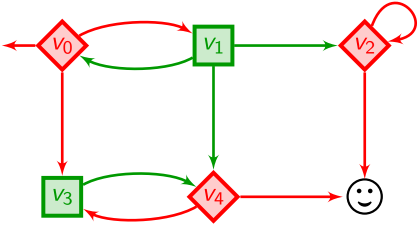

- Game graph $G = (V, E)$
- V partitioned into $V_\diamond$ and $V_\square$

Two-player turn-based games

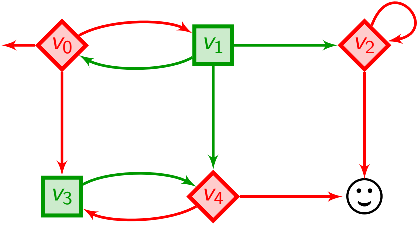

$\sigma_\Diamond(v_0) = v_3$, $\sigma_\Diamond(v_2) = \sigma_\Diamond(v_4) = ☺$
is a memoryless strategy for $\Diamond$

- Game graph $G = (V, E)$
- V partitioned into $V_\Diamond$ and $V_\Box$

- Strategy for player i:

$$\sigma_i \colon V^* V_i \to V$$

  s.t. $(\text{last}(h), \sigma_i(h)) \in E$ if $\text{last}(h) \in V_i$
- Memoryless if $\sigma_i(h) = \sigma_i(h')$ if $\text{last}(h) = \text{last}(h')$, that is:

$$\sigma_i \colon V_i \to V$$

- Given $(\sigma_\Diamond, \sigma_\Box)$, unique outcome

Two-player turn-based games

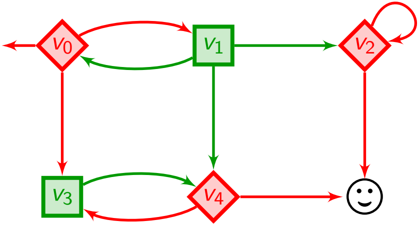

$\sigma_\Diamond(v_0) = v_3$, $\sigma_\Diamond(v_2) = \sigma_\Diamond(v_4) = ☺$ is a memoryless strategy for $\Diamond$

Goal of $\Diamond$: Reach ☺
Goal of $\Box$: Avoid ☺

- Game graph $G = (V, E)$
- V partitioned into $V_\Diamond$ and $V_\Box$

- Strategy for player i:

$$\sigma_i \colon V^* V_i \to V$$

  s.t. $(\mathrm{last}(h), \sigma_i(h)) \in E$ if $\mathrm{last}(h) \in V_i$
- Memoryless if $\sigma_i(h) = \sigma_i(h')$ if $\mathrm{last}(h) = \mathrm{last}(h')$, that is:

$$\sigma_i \colon V_i \to V$$

- Given $(\sigma_\Diamond, \sigma_\Box)$, unique outcome

- Zero-sum hyp.: $\mathrm{payoff}_\Box(\rho) = -\mathrm{payoff}_\Diamond(\rho)$

Two-player turn-based games

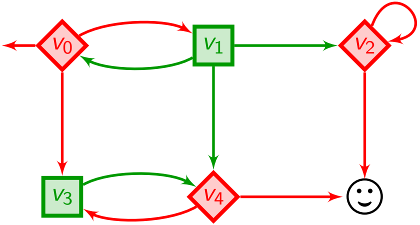

$\sigma_\diamond(v_0) = v_3$, $\sigma_\diamond(v_2) = \sigma_\diamond(v_4) = ☺$ is a memoryless strategy for $\diamond$

Goal of $\diamond$: Reach ☺
Goal of $\square$: Avoid ☺
$\mathrm{payoff}_\diamond(\rho) \in \{-1, 1\}$
$\mathrm{payoff}_\diamond(\rho) = 1$ iff ρ visits ☺

- Game graph $G = (V, E)$
- V partitioned into $V_\diamond$ and $V_\square$

- Strategy for player i:

$$\sigma_i\colon V^* V_i \to V$$

  s.t. $(\mathrm{last}(h), \sigma_i(h)) \in E$ if $\mathrm{last}(h) \in V_i$
- Memoryless if $\sigma_i(h) = \sigma_i(h')$ if $\mathrm{last}(h) = \mathrm{last}(h')$, that is:

$$\sigma_i\colon V_i \to V$$

- Given $(\sigma_\diamond, \sigma_\square)$, unique outcome

- Zero-sum hyp.: $\mathrm{payoff}_\square(\rho) = -\mathrm{payoff}_\diamond(\rho)$

Two-player turn-based games

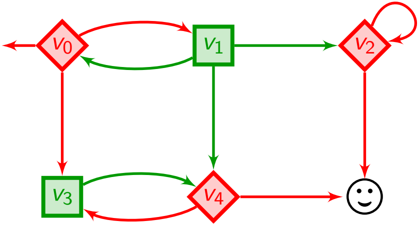

$\sigma_\diamond(v_0) = v_3$, $\sigma_\diamond(v_2) = \sigma_\diamond(v_4) = ☺$ is a memoryless strategy for $\diamond$

Goal of $\diamond$: Reach ☺
Goal of $\square$: Avoid ☺
$\text{payoff}_\diamond(\rho) \in \{-1, 1\}$
$\text{payoff}_\diamond(\rho) = 1$ iff ρ visits ☺

$\sigma_\diamond$ ensures payoff $+1$ for $\diamond$: it is a winning strategy

- Game graph $G = (V, E)$
- V partitioned into $V_\diamond$ and $V_\square$

- Strategy for player i:

$$\sigma_i \colon V^* V_i \to V$$

  s.t. $(\text{last}(h), \sigma_i(h)) \in E$ if $\text{last}(h) \in V_i$
- Memoryless if $\sigma_i(h) = \sigma_i(h')$ if $\text{last}(h) = \text{last}(h')$, that is:

$$\sigma_i \colon V_i \to V$$

- Given $(\sigma_\diamond, \sigma_\square)$, unique outcome

- Zero-sum hyp.: $\text{payoff}_\square(\rho) = -\text{payoff}_\diamond(\rho)$

Two-player turn-based games

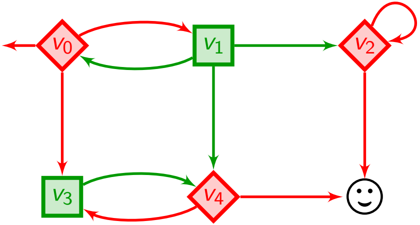

$\sigma_\Diamond(v_0) = v_3$, $\sigma_\Diamond(v_2) = \sigma_\Diamond(v_4) = ☺$ is a memoryless strategy for $\Diamond$

Goal of $\Diamond$: Reach ☺
Goal of $\Box$: Avoid ☺
$\mathrm{payoff}_\Diamond(\rho) \in \{0, 1\}$
$\mathrm{payoff}_\Diamond(\rho) = 1$ iff ρ visits ☺

$\sigma_\Diamond$ ensures payoff $+1$ for $\Diamond$: it is a winning strategy

- Game graph $G = (V, E)$
- V partitioned into $V_\Diamond$ and $V_\Box$

- Strategy for player i:

$$\sigma_i \colon V^* V_i \to V$$

  s.t. $(\mathrm{last}(h), \sigma_i(h)) \in E$ if $\mathrm{last}(h) \in V_i$
- Memoryless if $\sigma_i(h) = \sigma_i(h')$ if $\mathrm{last}(h) = \mathrm{last}(h')$, that is:

$$\sigma_i \colon V_i \to V$$

- Given $(\sigma_\Diamond, \sigma_\Box)$, unique outcome

- Zero-sum hyp.: $\mathrm{payoff}_\Box(\rho) = -\mathrm{payoff}_\Diamond(\rho)$

Non-zero-sum multiplayer games

- Several players $\mathrm{Agt} = \{A_1, \ldots, A_N\}$

Non-zero-sum multiplayer games

- Several players $\mathrm{Agt} = \{A_1, \ldots, A_N\}$
- Each player A plays according to a strategy σ_A

Non-zero-sum multiplayer games

- Several players $\mathsf{Agt} = \{A_1, \ldots, A_N\}$
- Each player A plays according to a strategy σ_A
- Each player A has a payoff function

$$\mathsf{payoff}_A : V^\omega \to \mathbb{R}$$

Non-zero-sum multiplayer games

- Several players $\mathsf{Agt} = \{A_1, \ldots, A_N\}$
- Each player A plays according to a strategy σ_A
- Each player A has a payoff function

$$\mathsf{payoff}_A : V^\omega \to \mathbb{R}$$

- Obviously non-zero-sum...

Non-zero-sum multiplayer games

- Several players $\mathsf{Agt} = \{A_1, \ldots, A_N\}$
- Each player A plays according to a strategy σ_A
- Each player A has a payoff function

$$\mathsf{payoff}_A : V^\omega \to \mathbb{R}$$

- Obviously non-zero-sum...
- Selfishness hypothesis: each player wants to maximize her own payoff!

Non-zero-sum multiplayer games

- Several players $\mathsf{Agt} = \{A_1, \ldots, A_N\}$
- Each player A plays according to a strategy σ_A
- Each player A has a payoff function

$$\mathsf{payoff}_A : V^\omega \to \mathbb{R}$$

- Obviously non-zero-sum...
- Selfishness hypothesis: each player wants to maximize her own payoff!
- Need of solution concepts to describe the kind of interactions between the players

Non-zero-sum multiplayer games

- Several players $\mathrm{Agt} = \{A_1, \ldots, A_N\}$
- Each player A plays according to a strategy σ_A
- Each player A has a payoff function

$$\mathsf{payoff}_A : V^\omega \to \mathbb{R}$$

- Obviously non-zero-sum...
- Selfishness hypothesis: each player wants to maximize her own payoff!
- Need of solution concepts to describe the kind of interactions between the players
- The simplest: Nash equilibria

Nash equilibria in turn-based games

Nash equilibrium

A strategy profile $(\sigma_A)_{A \in \mathrm{Agt}}$ is a Nash equilibrium if no player can improve her payoff by unilaterally changing her strategy.

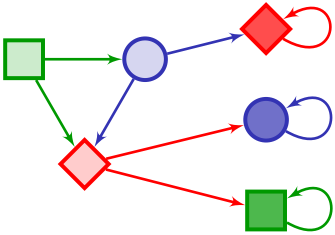

Nash equilibria in turn-based games

Nash equilibrium

A strategy profile $(\sigma_A)_{A \in \mathsf{Agt}}$ is a Nash equilibrium if no player can improve her payoff by unilaterally changing her strategy.

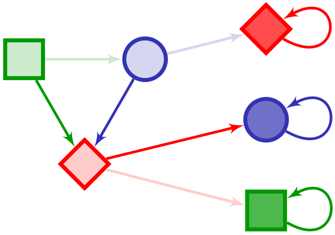

is a Nash equilibrium with payoff $(0, 1, 0)$

Nash equilibria in turn-based games

Nash equilibrium

A strategy profile $(\sigma_A)_{A \in \mathsf{Agt}}$ is a Nash equilibrium if no player can improve her payoff by unilaterally changing her strategy.

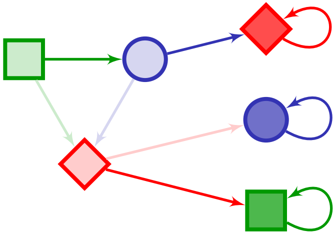

is not a Nash equilibrium

Characterization of Boolean Nash equilibria in turn-based games

ψ_A: objective of player A

Characterization of Boolean Nash equilibria in turn-based games

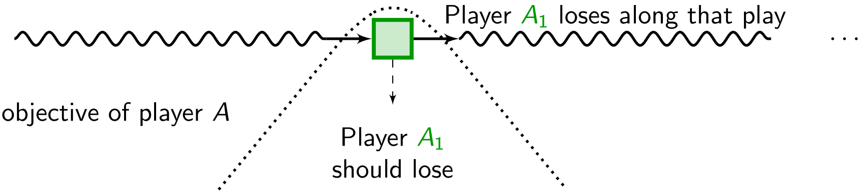

ψ_A: objective of player A

Characterization of Boolean Nash equilibria in turn-based games

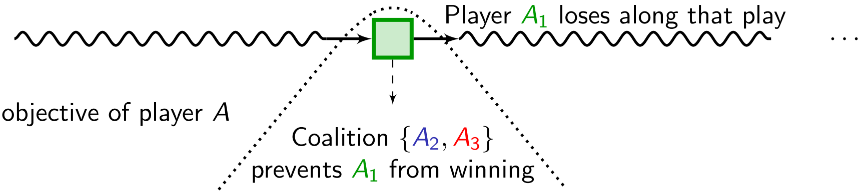

ψ_A: objective of player A

Characterization of Boolean Nash equilibria in turn-based games

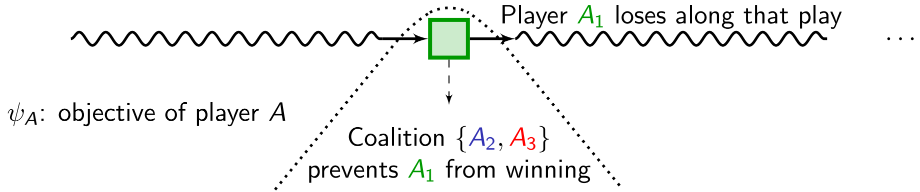

Main outcomes of Boolean Nash equilibria in turn-based games can be characterized by an LTL formula:

$$\Phi_{\mathsf{NE}} = \bigwedge_{A \in \mathsf{Agt}} \left(\neg \psi_A \Rightarrow \mathbf{G} \neg W_A \right)$$

where W_A is the set of winning states for A (should be precomputed).

Existing results in the framework of turn-based games

[UW11,Umm11]

- There always exists a Nash equilibrium for Boolean ω-regular objectives

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Stochastic Multiplayer Games *(LMCS)*
[Umm11] Ummels. Stochastic multiplayer games: theory and algorithms (RWTH Aachen University)

Existing results in the framework of turn-based games

[UW11,Umm11]

- There always exists a Nash equilibrium for Boolean ω-regular objectives
- One can decide the constrained existence of a Nash equilibrium (and compute one!)

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Stochastic Multiplayer Games *(LMCS)*
[Umm11] Ummels. Stochastic multiplayer games: theory and algorithms (RWTH Aachen University)

Existing results in the framework of turn-based games

[UW11,Umm11]

- 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 (RWTH Aachen University)

Existing results in the framework of turn-based games

[UW11,Umm11]

- 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 restrict to pure equilibria

[UW11] Ummels, Wojtczak. The Complexity of Nash Equilibria in Stochastic Multiplayer Games *(LMCS)*
[Umm11] Ummels. Stochastic multiplayer games: theory and algorithms (RWTH Aachen University)

Extension to partially informed players?

Invisible actions in concurrent games [BBMU15]

Extension to partially informed players?

Invisible actions in concurrent games [BBMU15]

The matching-penny game:

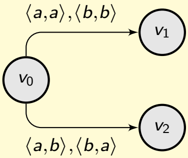

[BBMU15] Bouyer, Brenguier, Markey, Ummels. Pure Nash equilibria in concurrent games *(LMCS)*

Extension to partially informed players?

Invisible actions in concurrent games [BBMU15]

The matching-penny game:

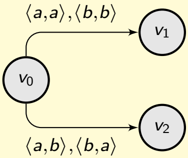

There is no Nash equilibrium.

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

Extension to partially informed players?

Invisible actions in concurrent games [BBMU15]

The matching-penny game:

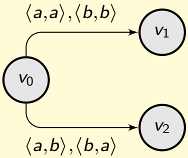

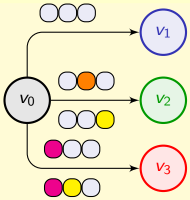

There is no Nash equilibrium.

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

Extension to partially informed players?

Invisible actions in concurrent games [BBMU15]

The matching-penny game:

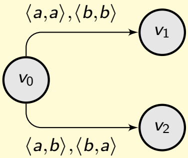

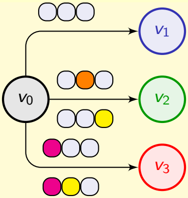

There is no Nash equilibrium.

- $\mathrm{susp}\Big((v_0, v_3), \bigcirc\bigcirc\bigcirc\Big) = \{A_1\}$

Extension to partially informed players?

Invisible actions in concurrent games [BBMU15]

The matching-penny game:

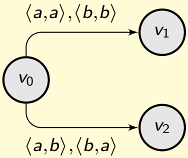

There is no Nash equilibrium.

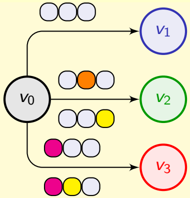

- $\mathrm{susp}\Big((v_0, {\color{red}v_3}), \bigcirc\bigcirc\bigcirc\Big) = \{A_1\}$
- $\mathrm{susp}\Big((v_0, {\color{green}v_2}), \bigcirc\bigcirc\bigcirc\Big) = \{A_2, A_3\}$

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

Extension to partially informed players?

Invisible actions in concurrent games [BBMU15]

The matching-penny game:

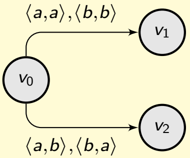

There is no Nash equilibrium.

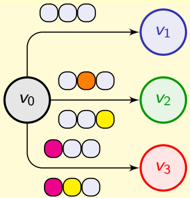

- $\mathrm{susp}\Big((v_0, v_3), \bigcirc\bigcirc\bigcirc\Big) = \{A_1\}$
- $\mathrm{susp}\Big((v_0, v_2), \bigcirc\bigcirc\bigcirc\Big) = \{A_2, A_3\}$

Solution via the suspect game abstraction,
a structure to track suspect players

Extension to partially informed players?

Invisible actions in concurrent games [BBMU15]

The matching-penny game:

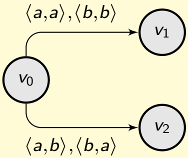

There is no Nash equilibrium.

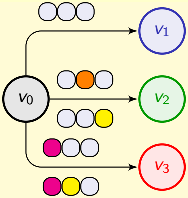

- $\mathrm{susp}\Big((v_0, v_3), \Big) = \{A_1\}$
- $\mathrm{susp}\Big((v_0, v_2), \Big) = \{A_2, A_3\}$

Solution via the suspect game abstraction,
a structure to track suspect players

Can we add more partial information to that framework?

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

Concurrent games with signals

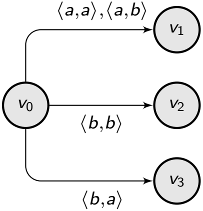

[Tom98] Tomala. Pure equilibria of repeated games with public observation *(International Journal of Game Theory)*

Concurrent games with signals

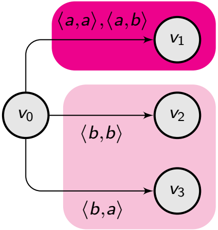

- Signal for player A_1: and

- On playing a, player A_1 will receive
- On playing b, player A_1 will receive

[Tom98] Tomala. Pure equilibria of repeated games with public observation *(International Journal of Game Theory)*

Concurrent games with signals

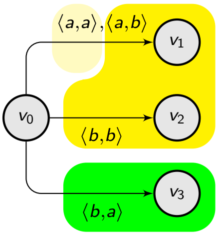

- Signal for player A_1: ● and ●
- Signal for player A_2: ●, ● and ●

- On playing a, player A_1 will receive ●
- On playing b, player A_1 will receive ●
- On playing a, player A_2 will receive either ● or ●
- On playing b, player A_2 will receive ●

[Tom98] Tomala. Pure equilibria of repeated games with public observation *(International Journal of Game Theory)*

Concurrent games with signals

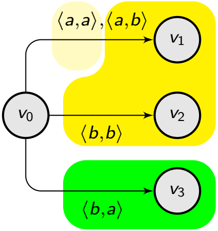

- Signal for player A_1: and
- Signal for player A_2: , and

- On playing a, player A_1 will receive
- On playing b, player A_1 will receive

- On playing a, player A_2 will receive either or
- On playing b, player A_2 will receive

Public signal

Same signal to every player!

[Tom98] Tomala. Pure equilibria of repeated games with public observation *(International Journal of Game Theory)*

A concurrent game with signals is a tuple

$$\mathcal{G} = \langle V, v_{\mathsf{init}}, \mathsf{Agt}, \mathsf{Act}, \Sigma, \mathsf{Allow}, \mathsf{Tab}, (\ell_A)_{A\in\mathsf{Agt}}, (\mathsf{payoff}_A)_{A\in\mathsf{Agt}}\rangle$$

where:

- V is a finite set of vertices,
- $v_{\mathsf{init}} \in V$ is the initial vertex,
- Agt is a finite set of players,
- Act is a finite set of actions,
- Σ is a finite alphabet,
- $\mathsf{Allow}\colon V \times \mathsf{Agt} \to 2^{\mathsf{Act}} \setminus \{\emptyset\}$ is a mapping indicating the actions available to a given player in a given state,
- $\mathsf{Tab}\colon V \times \mathsf{Act}^{\mathsf{Agt}} \to V$ associates, with a given state and a given move of the players (i.e., an element of $\mathsf{Act}^{\mathsf{Agt}}$), the state resulting from that move,
- for every $A \in \mathsf{Agt}$, $\ell_A\colon \left(\mathsf{Act}^{\mathsf{Agt}} \times V\right) \to \Sigma$ is a signal,
- for every $A \in \mathsf{Agt}$, $\mathsf{payoff}_A\colon V \times \left(\mathsf{Act}^{\mathsf{Agt}} \times V\right)^{\omega} \to \mathbb{D}$ is a payoff function for player A

Strategies

- What player A sees from history $h = v_0 \xrightarrow{m_0} v_1 \xrightarrow{m_1} \ldots \xrightarrow{m_{k-1}} v_k$:

$$\pi_A(h) = v_0 \cdot m_0(A) \cdot \ell_A(m_0, v_1) \cdot m_1(A) \ldots m_{k-1}(A) \cdot \ell_A(m_{k-1}, v_k)$$

$\leadsto$ perfect recall hypothesis

Strategies

- What player A sees from history $h = v_0 \xrightarrow{m_0} v_1 \xrightarrow{m_1} \ldots \xrightarrow{m_{k-1}} v_k$:

$$\pi_A(h) = v_0 \cdot m_0(A) \cdot \ell_A(m_0, v_1) \cdot m_1(A) \ldots m_{k-1}(A) \cdot \ell_A(m_{k-1}, v_k)$$

$\rightsquigarrow$ perfect recall hypothesis

- Undistinguishability relation for player A:

$$h \sim_A h' \qquad \text{iff} \qquad \pi_A(h) = \pi_A(h')$$

Strategies

- What player A sees from history $h = v_0 \xrightarrow{m_0} v_1 \xrightarrow{m_1} \dots \xrightarrow{m_{k-1}} v_k$:

$$\pi_A(h) = v_0 \cdot m_0(A) \cdot \ell_A(m_0, v_1) \cdot m_1(A) \dots m_{k-1}(A) \cdot \ell_A(m_{k-1}, v_k)$$

$\rightsquigarrow$ perfect recall hypothesis

- Undistinguishability relation for player A:

$$h \sim_A h' \qquad \text{iff} \qquad \pi_A(h) = \pi_A(h')$$

- A strategy for player A is a (partial) function:

$$\sigma_A \colon V \cdot \left(\mathrm{Act}^{\mathrm{Agt}} \cdot V\right)^* \to \mathrm{Act}$$

such that $h \sim_A h'$ implies $\sigma_A(h) = \sigma_A(h')$.

Strategies

- What player A sees from history $h = v_0 \xrightarrow{m_0} v_1 \xrightarrow{m_1} \ldots \xrightarrow{m_{k-1}} v_k$:

$$\pi_A(h) = v_0 \cdot m_0(A) \cdot \ell_A(m_0, v_1) \cdot m_1(A) \ldots m_{k-1}(A) \cdot \ell_A(m_{k-1}, v_k)$$

$\rightsquigarrow$ perfect recall hypothesis

- Undistinguishability relation for player A:

$$h \sim_A h' \qquad \text{iff} \qquad \pi_A(h) = \pi_A(h')$$

- A strategy for player A is a (partial) function:

$$\sigma_A \colon V \cdot \left(\mathrm{Act}^{\mathrm{Agt}} \cdot V\right)^* \to \mathrm{Act}$$

  such that $h \sim_A h'$ implies $\sigma_A(h) = \sigma_A(h')$.
- A strategy profile is a tuple $\sigma_{\mathrm{Agt}} = (\sigma_A)_{A \in \mathrm{Agt}}$ where σ_A is a strategy for player A.

Discussion on the perfect-recall assumption

In most existing frameworks, strategies are defined through observation maps

$$O_A\colon V \to \Sigma \qquad\qquad \sigma_A\colon \Sigma^* \to \mathsf{Act}$$

Discussion on the perfect-recall assumption

In most existing frameworks, strategies are defined through observation maps

$$O_A \colon V \to \Sigma \qquad\qquad \sigma_A \colon \Sigma^* \to \mathsf{Act}$$

- This choice is suitable for distributed synthesis and Nash equilibria (for instance)...

Discussion on the perfect-recall assumption

In most existing frameworks, strategies are defined through observation maps

$$O_A \colon V \to \Sigma \qquad\qquad \sigma_A \colon \Sigma^* \to \mathsf{Act}$$

- This choice is suitable for distributed synthesis and Nash equilibria (for instance)...
- but I think this choice is **not** suitable in general

Discussion on the perfect-recall assumption

In most existing frameworks, strategies are defined through observation maps

$$O_A\colon V \to \Sigma \qquad\qquad \sigma_A\colon \Sigma^* \to \mathsf{Act}$$

- This choice is suitable for distributed synthesis and Nash equilibria (for instance)...
- but I think this choice is **not** suitable in general

Example (Subgame-perfect equilibrium)

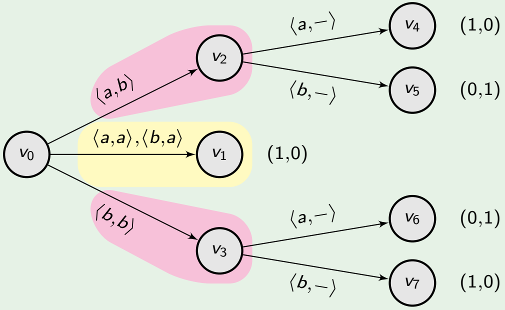

Discussion on the perfect-recall assumption

In most existing frameworks, strategies are defined through observation maps

$$O_A\colon V \to \Sigma \qquad\qquad \sigma_A\colon \Sigma^* \to \mathsf{Act}$$

- This choice is suitable for distributed synthesis and Nash equilibria (for instance)...
- but I think this choice is **not** suitable in general

Example (Subgame-perfect equilibrium)

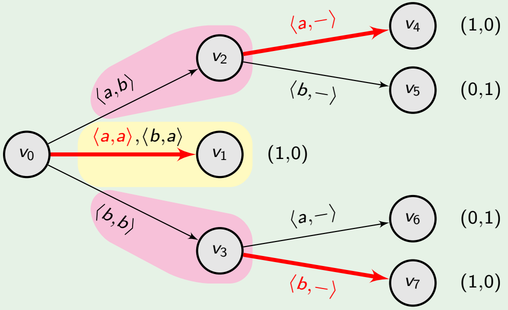

Digression on payoff functions

Payoff functions

- Payoff function for player A ($\mathbb{D} \subseteq \mathbb{R}$):

$$\mathsf{payoff}_A \colon V \cdot \left(\mathsf{Act}^{\mathsf{Agt}} \cdot V\right)^{\omega} \to \mathbb{D}$$

Digression on payoff functions

Payoff functions

- Payoff function for player A ($\mathbb{D} \subseteq \mathbb{R}$):

$$\mathsf{payoff}_A \colon V \cdot \left(\mathsf{Act}^{\mathsf{Agt}} \cdot V\right)^{\omega} \to \mathbb{D}$$

- payoff_A is privately visible whenever

$$\pi_A(\rho) = \pi_A(\rho') \text{ implies } \mathsf{payoff}_A(\rho) = \mathsf{payoff}_A(\rho')$$

Digression on payoff functions

Payoff functions

- Payoff function for player A ($\mathbb{D} \subseteq \mathbb{R}$):

$$\mathsf{payoff}_A \colon V \cdot \left(\mathsf{Act}^{\mathsf{Agt}} \cdot V\right)^{\omega} \to \mathbb{D}$$

- payoff_A is privately visible whenever

$$\pi_A(\rho) = \pi_A(\rho') \text{ implies } \mathsf{payoff}_A(\rho) = \mathsf{payoff}_A(\rho')$$

- If signal ℓ is public ($\ell_A = \ell$ for every A), payoff_A is publicly visible whenever

$$\ell(\rho) = \ell(\rho') \text{ implies } \mathsf{payoff}_A(\rho) = \mathsf{payoff}_A(\rho')$$

Digression on payoff functions (cont'd)

Some payoff functions

- Boolean ω-regular payoff function (for Ω):

$$\mathsf{payoff}(\rho) = \begin{cases} 1 & \text{if } \rho \in \Omega \\ 0 & \text{otherwise} \end{cases}$$

Digression on payoff functions (cont'd)

Some payoff functions

- Boolean ω-regular payoff function (for Ω):

$$\text{payoff}(\rho) = \begin{cases} 1 & \text{if } \rho \in \Omega \\ 0 & \text{otherwise} \end{cases}$$

- Mean-payoff (limsup or liminf) w.r.t. weight function w:

$$\begin{cases} \underline{\text{MP}}_w(\rho) = \liminf_{n\to\infty} \sum_{i=0}^{n} w\left(v_i \xrightarrow{m_i} v_{i+1}\right) \\ \overline{\text{MP}}_w(\rho) = \limsup_{n\to\infty} \sum_{i=0}^{n} w\left(v_i \xrightarrow{m_i} v_{i+1}\right) \end{cases}$$

Digression on payoff functions (cont'd)

Some payoff functions

- Boolean ω-regular payoff function (for Ω):

$$\mathsf{payoff}(\rho) = \begin{cases} 1 & \text{if } \rho \in \Omega \\ 0 & \text{otherwise} \end{cases}$$

- Mean-payoff (limsup or liminf) w.r.t. weight function w:

$$\begin{cases} \underline{\mathsf{MP}}_w(\rho) = \liminf_{n\to\infty} \sum_{i=0}^{n} w\left(v_i \xrightarrow{m_i} v_{i+1}\right) \\ \overline{\mathsf{MP}}_w(\rho) = \limsup_{n\to\infty} \sum_{i=0}^{n} w\left(v_i \xrightarrow{m_i} v_{i+1}\right) \end{cases}$$

For public visibility, we will assume that atomic propositions/atomic weights are defined w.r.t. the signal alphabet Σ.

An example

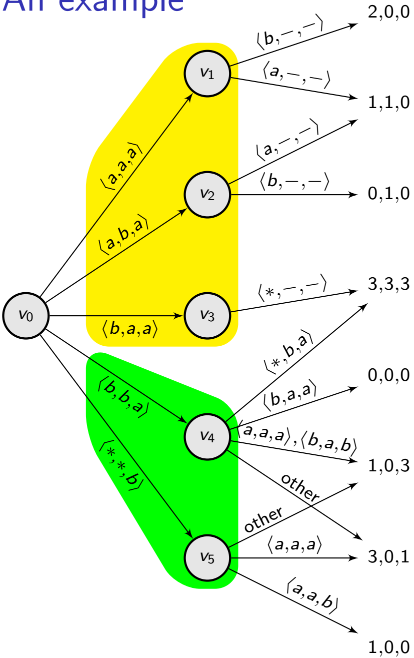

- Three players concurrent game with public signal

An example

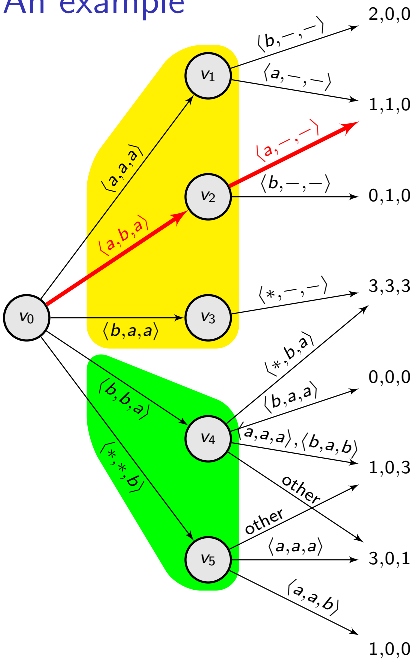

- Three players concurrent game with public signal
- Consider the (partial) strategy profile σ_{Agt}. Can we complete it into a Nash equilibrium?

An example

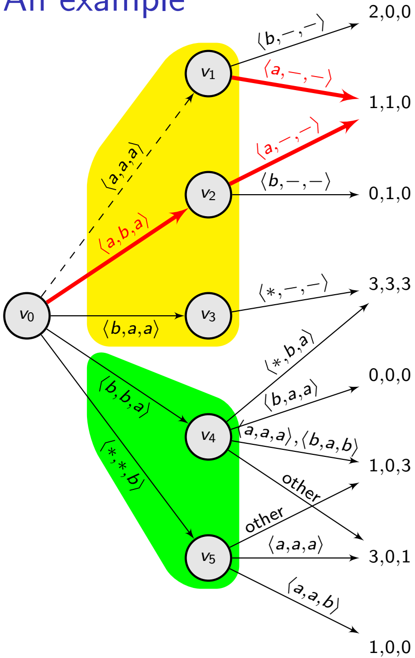

- Three players concurrent game with public signal
- Consider the (partial) strategy profile σ_{Agt}. Can we complete it into a Nash equilibrium?
- This is an A_2-deviation, which is invisible to both A_1 and A_3. A_1 has to play a and cannot deviate to $2,0,0$.

An example

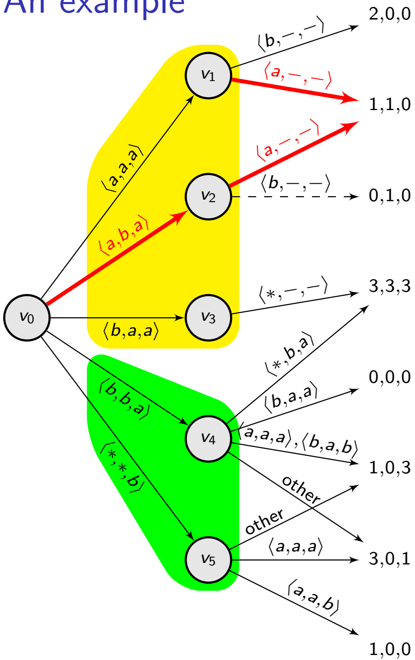

- Three players concurrent game with public signal
- Consider the (partial) strategy profile σ_{Agt}. Can we complete it into a Nash equilibrium?
- This is an A_2-deviation, which is invisible to both A_1 and A_3. A_1 has to play a and cannot deviate to $2, 0, 0$.
- This is a non-profitable A_1-deviation.

An example

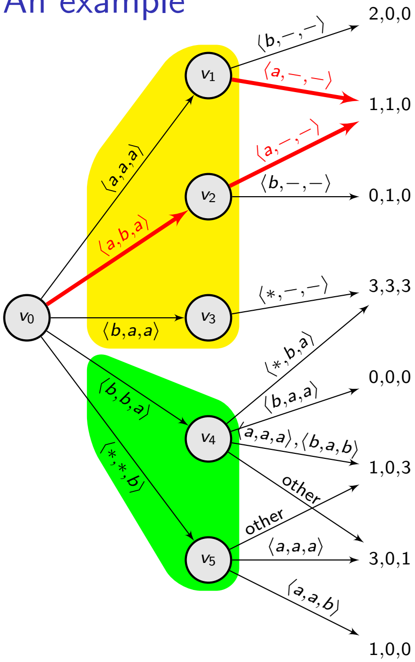

- Three players concurrent game with public signal
- Consider the (partial) strategy profile σ_{Agt}. Can we complete it into a Nash equilibrium?
- This is an A_2-deviation, which is invisible to both A_1 and A_3. A_1 has to play a and cannot deviate to $2, 0, 0$.
- This is a non-profitable A_1-deviation.
- No one (alone) can deviate to v_3.

An example

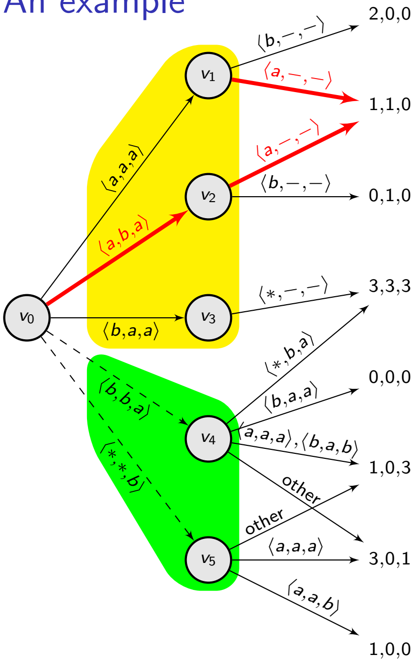

- Three players concurrent game with public signal
- Consider the (partial) strategy profile σ_{Agt}. Can we complete it into a Nash equilibrium?
- This is an A_2-deviation, which is invisible to both A_1 and A_3. A_1 has to play a and cannot deviate to $2, 0, 0$.
- This is a non-profitable A_1-deviation.
- No one (alone) can deviate to v_3.
- A_1 can deviate to v_4 and A_3 can deviate to v_5: A_2 knows there has been a deviation, but (s)he doesn't whether A_1 or A_3 did so, and whether the game proceeds to v_4 or v_5. On the other hand, both A_1 and A_3 know!

An example

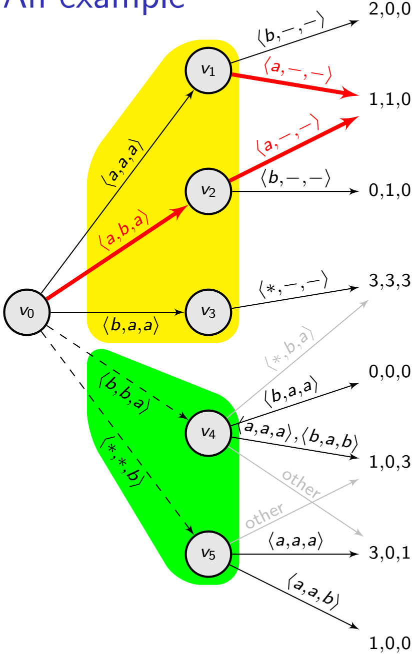

- Three players concurrent game with public signal
- Consider the (partial) strategy profile σ_{Agt}. Can we complete it into a Nash equilibrium?
- This is an A_2-deviation, which is invisible to both A_1 and A_3. A_1 has to play a and cannot deviate to $2, 0, 0$.
- This is a non-profitable A_1-deviation.
- No one (alone) can deviate to v_3.
- A_1 can deviate to v_4 and A_3 can deviate to v_5: A_2 knows there has been a deviation, but (s)he doesn't whether A_1 or A_3 did so, and whether the game proceeds to v_4 or v_5. On the other hand, both A_1 and A_3 know! But if the game proceeds to v_4, A_3 can help A_2 punishing A_1, and if the game proceeds to v_5, A_1 can help A_2 punishing A_3.

An example

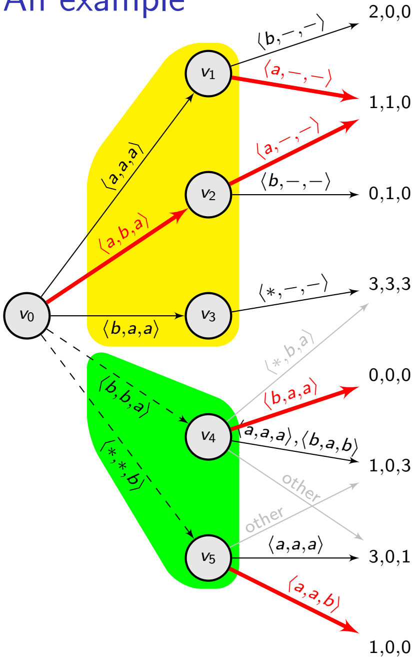

- Three players concurrent game with public signal
- Consider the (partial) strategy profile σ_{Agt}. Can we complete it into a Nash equilibrium?
- This is an A_2-deviation, which is invisible to both A_1 and A_3. A_1 has to play a and cannot deviate to $2,0,0$.
- This is a non-profitable A_1-deviation.
- No one (alone) can deviate to v_3.
- A_1 can deviate to v_4 and A_3 can deviate to v_5: A_2 knows there has been a deviation, but (s)he doesn't whether A_1 or A_3 did so, and whether the game proceeds to v_4 or v_5. On the other hand, both A_1 and A_3 know! But if the game proceeds to v_4, A_3 can help A_2 punishing A_1, and if the game proceeds to v_5, A_1 can help A_2 punishing A_3.

How to systematically track all undistinguishable behaviours and all individual deviations? Is that always possible?

First undecidability results

One cannot decide the existence problem in games with signals with three players and publicly visible qualitative ω-regular payoff functions.

$\rightsquigarrow$ by reduction from the distributed synthesis problem (construction for reachability properties taken in [BK10])

[BK10] Berwanger, Kaiser. Information Tracking in Games on Graphs *(Journal of Logic, Language and Information)*
[DDG+10] Degorre, Doyen, Gentilini, Raskin, Toruńczyk. Energy and Mean-Payoff Games with Imperfect Information *(CSL'10)*

First undecidability results

One cannot decide the existence problem in games with signals with three players and publicly visible qualitative ω-regular payoff functions.

$\rightsquigarrow$ by reduction from the distributed synthesis problem (construction for reachability properties taken in [BK10])

One cannot decide the constrained existence of a Nash equilibrium in a game with public signals, for a mixture of limsup and liminf mean-payoff functions which are privately visible. Even for two players.

$\rightsquigarrow$ by reduction from blind mean-payoff games (proven undecidable in [DDG+10])

[BK10] Berwanger, Kaiser. Information Tracking in Games on Graphs *(Journal of Logic, Language and Information)*
[DDG+10] Degorre, Doyen, Gentilini, Raskin, Toruńczyk. Energy and Mean-Payoff Games with Imperfect Information *(CSL'10)*

The epistemic game abstraction

Inspired by:

- the standard powerset construction [Rei84]
- the epistemic unfolding for coordination/distributed synthesis [BKP11]
- the suspect game [BBMU15]
- the deviator game [Bre16]

[Rei84] Reif. The complexity of two-player games of incomplete information *(J. Comp. and Syst. Sc.)*
[BKP11] Berwanger, Kaiser, Puchala. Perfect-information construction for coordination in games *(FSTTCS'11)*
[BBMU15] Pure Nash equilibria in concurrent games *(Log. Meth. in Comp. Sc.)*
[Bre16] Brenguier. Robust equilibria in mean-payoff games *(FoSSaCS'16)*

The epistemic game abstraction

Inspired by:

- the standard powerset construction [Rei84]
- the epistemic unfolding for coordination/distributed synthesis [BKP11]
- the suspect game [BBMU15]
- the deviator game [Bre16]

The idea is to track all possible undistinguishable behaviours, including the single-player deviations

[Rei84] Reif. The complexity of two-player games of incomplete information *(J. Comp. and Syst. Sc.)*
[BKP11] Berwanger, Kaiser, Puchala. Perfect-information construction for coordination in games *(FSTTCS'11)*
[BBMU15] Pure Nash equilibria in concurrent games *(Log. Meth. in Comp. Sc.)*
[Bre16] Brenguier. Robust equilibria in mean-payoff games *(FoSSaCS'16)*

The epistemic game abstraction (cont'd)

Epistemic states (type-1)

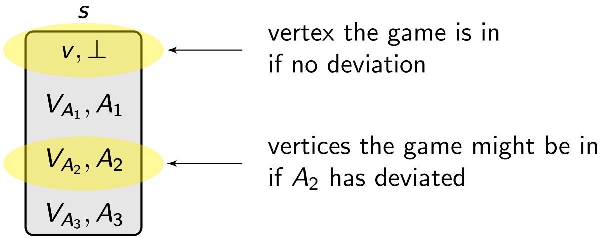

Captures set of histories that some of the players do not distinguish.
A_i cannot distinguish between the normal outcome (no deviation) and deviations of other players leading to some $v \in V_{A_j}$ with $j \neq i$

The epistemic game abstraction (cont'd)

Epistemic states (type-2)

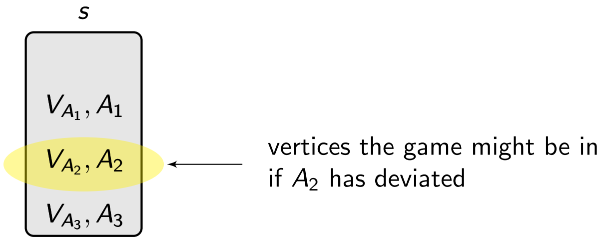

Captures set of histories that some of the players do not distinguish.
A_i cannot distinguish between the possible deviations of other players (but he knows there has been a deviation)

The example again

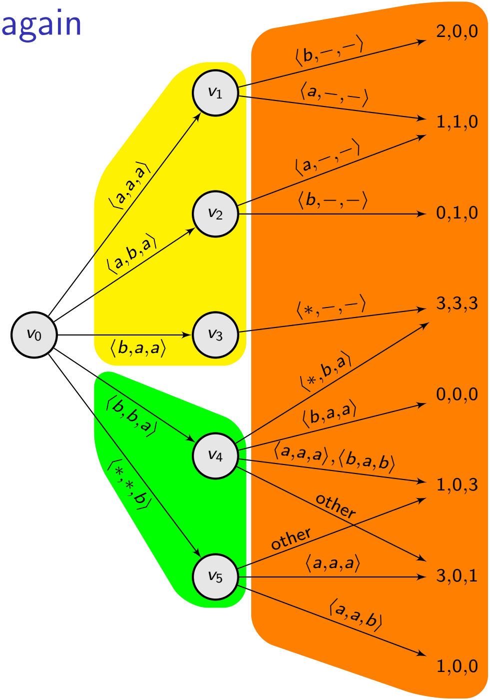

The example again

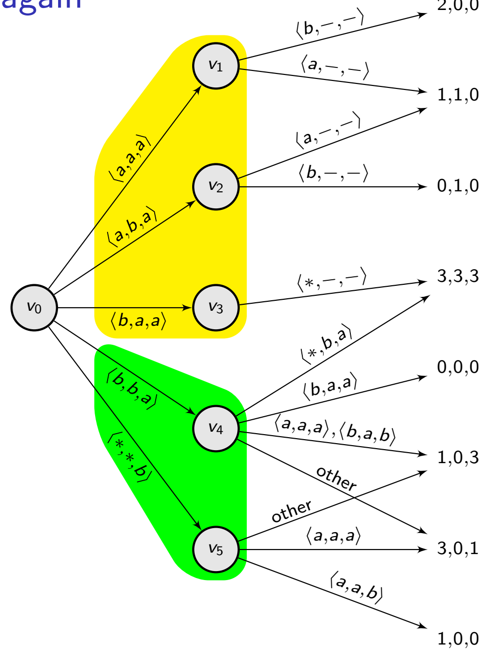

Example of construction

s_0

$v_0, \perp$
$\emptyset, A_1$
$\emptyset, A_2$
$\emptyset, A_3$

Example of construction

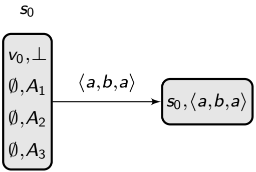

Example of construction

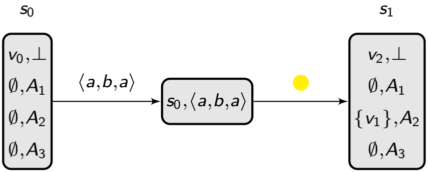

Example of construction

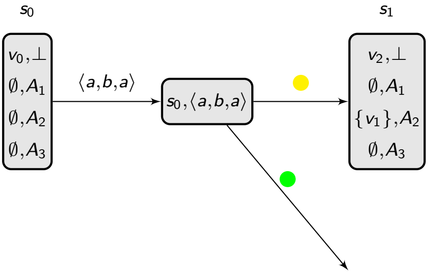

Example of construction

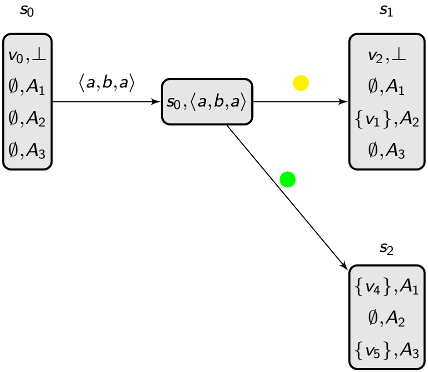

Example of construction

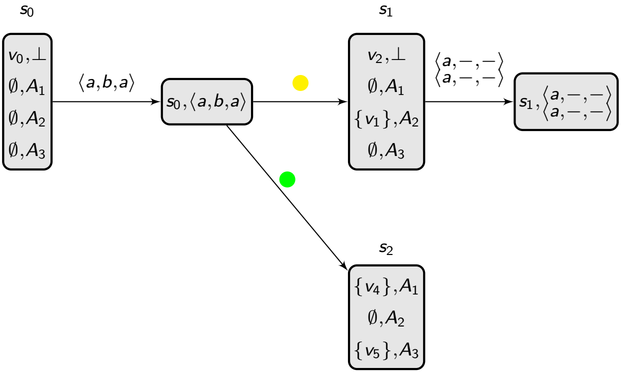

Example of construction

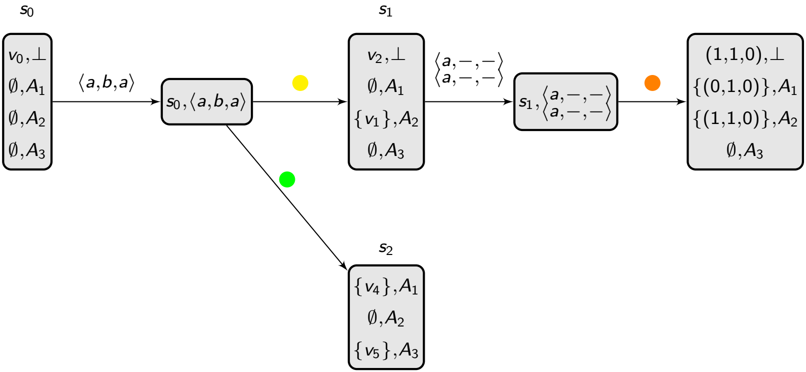

Example of construction

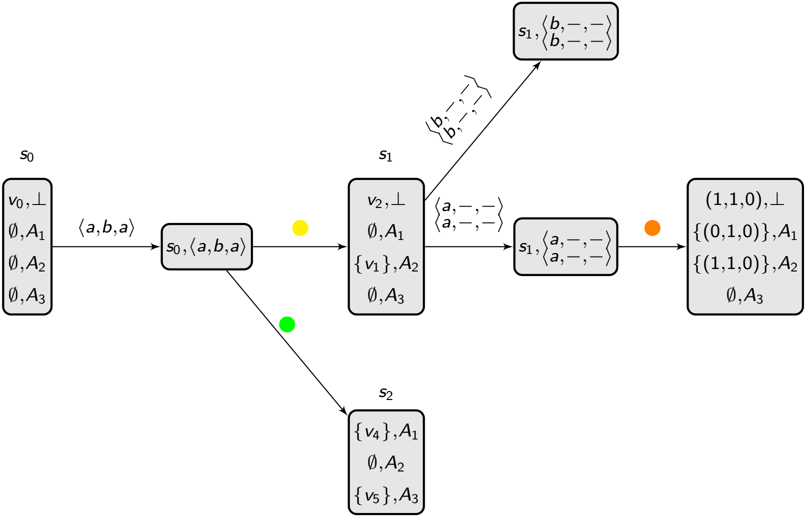

Example of construction

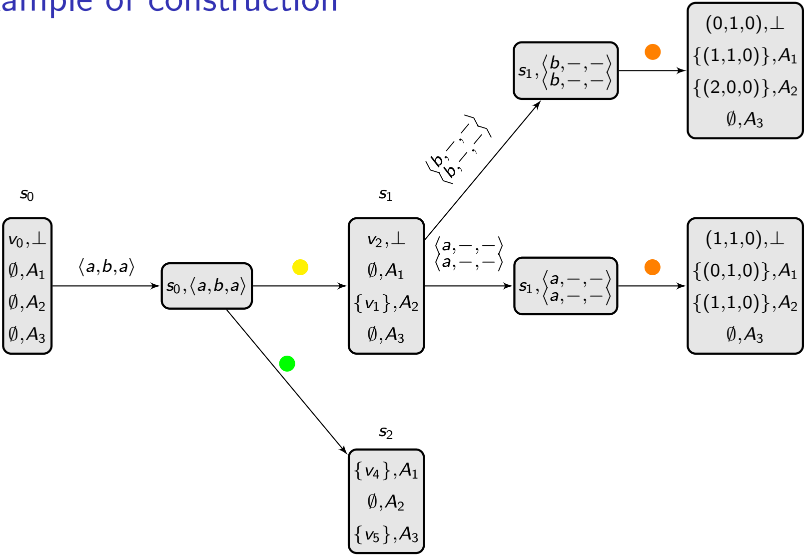

Example of construction

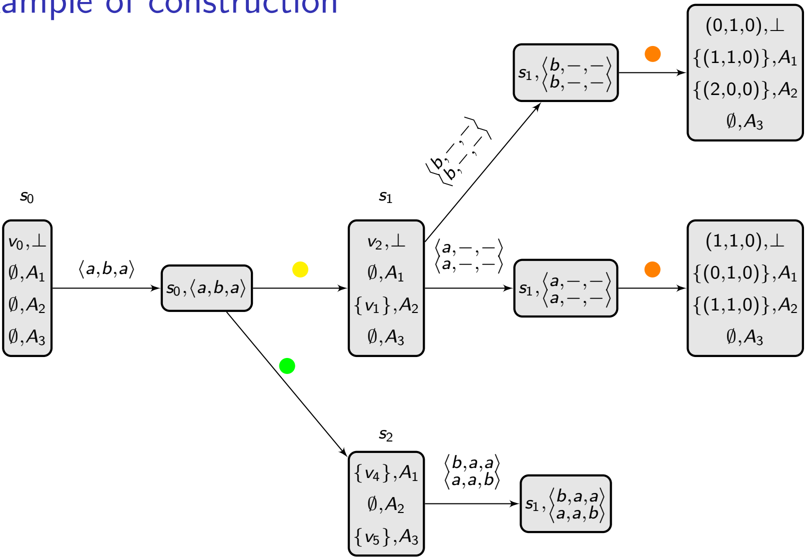

Example of construction

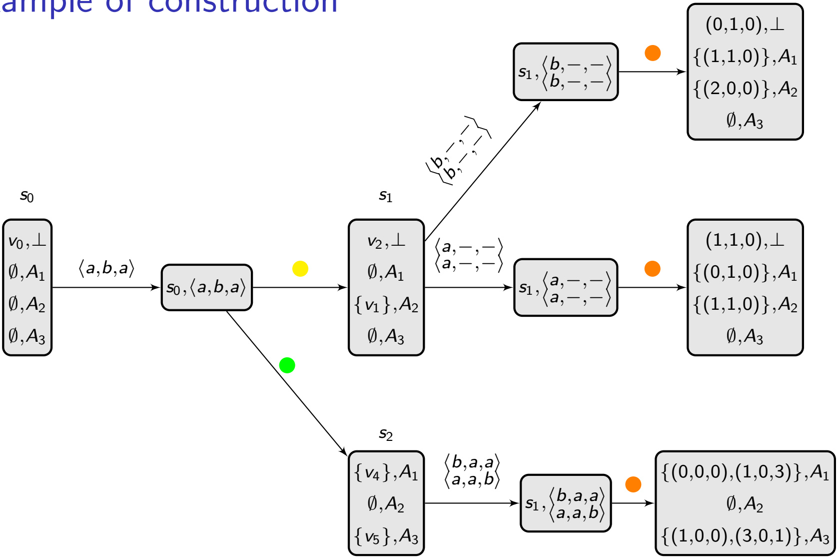

Example of construction

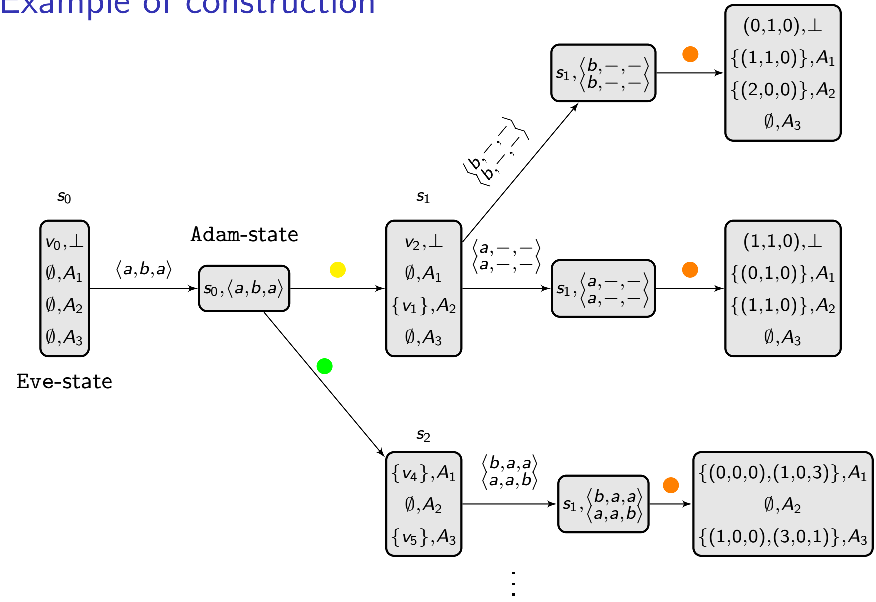

Properties of the epistemic game

- To every history H in the epistemic game, one can associate sets
  - $\mathit{concrete}_{\perp}(H)$: at most one concrete real history (if no deviation)
  - $\mathit{concrete}_{A}(H)$: all possible A-deviations
  - $\mathit{concrete}(H) = \bigcup_{A \in \mathsf{Agt} \cup \{\perp\}} \mathit{concrete}_{A}(H)$

Properties of the epistemic game

- To every history H in the epistemic game, one can associate sets
  - $\mathit{concrete}_{\perp}(H)$: at most one concrete real history (if no deviation)
  - $\mathit{concrete}_{A}(H)$: all possible A-deviations
  - $\mathit{concrete}(H) = \bigcup_{A \in \mathsf{Agt} \cup \{\perp\}} \mathit{concrete}_{A}(H)$

H history in the epistemic game. For every $h_1 \neq h_2 \in \mathit{concrete}(H)$,

$$h_1 \sim_A h_2 \qquad \text{iff} \qquad h_1, h_2 \notin \mathit{concrete}_A(H)$$

Properties of the epistemic game (cont'd)

Winning condition for Eve

A strategy $\sigma_{\texttt{Eve}}$ is said winning for payoff $p \in \mathbb{R}^{\mathrm{Agt}}$ from s_0 whenever $\mathrm{payoff}(\mathit{concrete}_{\perp}(\mathrm{out}_{\perp}(\sigma_{\texttt{Eve}}, s_0))) = p$, and for every $R \in \mathrm{out}(\sigma_{\texttt{Eve}}, s_0)$, for every $A \in \mathrm{Agt}$, for every $\rho \in \mathit{concrete}_A(R)$, $\mathrm{payoff}_A(\rho) \leq p_A$.

Properties of the epistemic game (cont'd)

Winning condition for $\mathtt{Eve}$ (publicly visible payoffs)

A strategy $\sigma_{\mathtt{Eve}}$ is said winning for p from s_0 whenever $\mathsf{payoff}'(\mathsf{out}_\perp(\sigma_{\mathtt{Eve}}, s_0)) = p$, and for every $R \in \mathsf{out}(\sigma_{\mathtt{Eve}}, s_0)$, for every $A \in \mathsf{susp}(R)$, $\mathsf{payoff}'_A(R) \leq p_A$.

Properties of the epistemic game (cont'd)

Winning condition for $\texttt{Eve}$ (publicly visible payoffs)

A strategy $\sigma_{\texttt{Eve}}$ is said winning for p from s_0 whenever $\text{payoff}'(\text{out}_{\perp}(\sigma_{\texttt{Eve}}, s_0)) = p$, and for every $R \in \text{out}(\sigma_{\texttt{Eve}}, s_0)$, for every $A \in \text{susp}(R)$, $\text{payoff}'_A(R) \leq p_A$.

Proposition

There is a Nash equilibrium in $\mathcal{G}$ with payoff p from v_0 if and only if $\texttt{Eve}$ has a winning strategy for p in $\mathcal{E}_{\mathcal{G}}$ from s_0.

Application to ω-regular objectives

Player A_1 loses along that play

ψ_A: objective of player A

Application to ω-regular objectives

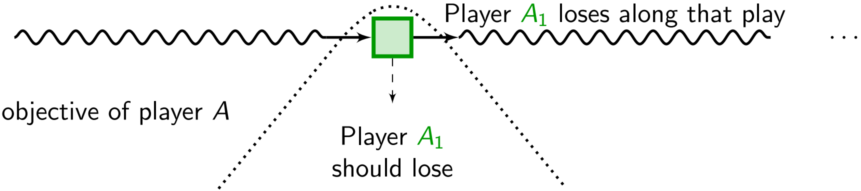

ψ_A: objective of player A

Application to ω-regular objectives

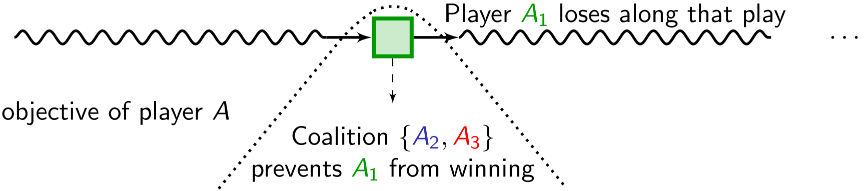

ψ_A: objective of player A

Application to ω-regular objectives

ψ_A: objective of player A

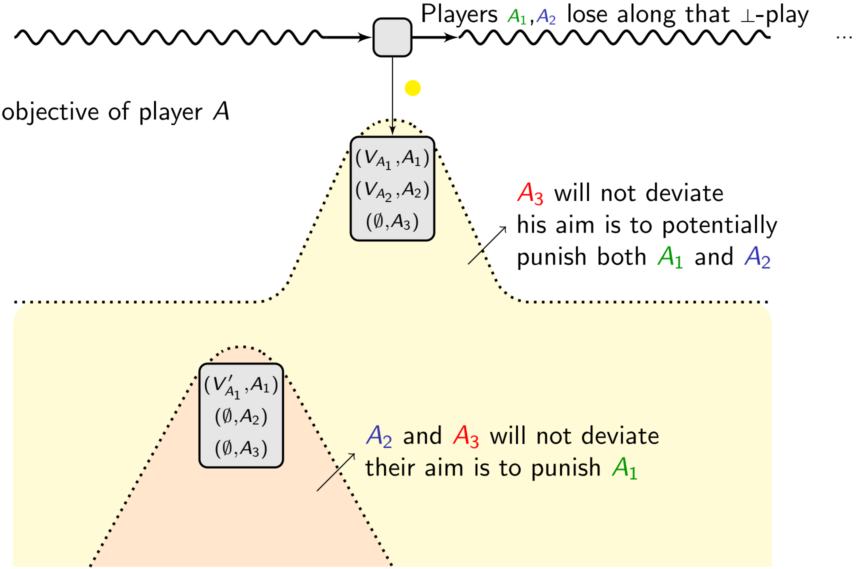

Application to ω-regular objectives (cont'd)

- This amounts to solving two-player turn-based games with generalized (i.e. conjunctions of) ω-regular objectives

[CDHR07] Chatterjee, Doyen, Henzinger, Raskin. Algorithms for ω-regular games with imperfect information *(LMCS)*

Application to ω-regular objectives (cont'd)

- This amounts to solving two-player turn-based games with generalized (i.e. conjunctions of) ω-regular objectives

Theorem

One can decide the (constrained) existence of a Nash equilibrium in a game with public signal and publicly visible payoff functions associated with parity conditions in EXPSPACE. It is EXPTIME-hard.

[CDHR07] Chatterjee, Doyen, Henzinger, Raskin. Algorithms for ω-regular games with imperfect information *(LMCS)*

Application to ω-regular objectives (cont'd)

- This amounts to solving two-player turn-based games with generalized (i.e. conjunctions of) ω-regular objectives

Theorem

One can decide the (constrained) existence of a Nash equilibrium in a game with public signal and publicly visible payoff functions associated with parity conditions in EXPSPACE. It is EXPTIME-hard.

- EXPTIME-hardness: same proof as for the distributed synthesis problem [CDHR07]

[CDHR07] Chatterjee, Doyen, Henzinger, Raskin. Algorithms for ω-regular games with imperfect information *(LMCS)*

Application to ω-regular objectives (cont'd)

- This amounts to solving two-player turn-based games with generalized (i.e. conjunctions of) ω-regular objectives

Theorem

One can decide the (constrained) existence of a Nash equilibrium in a game with public signal and publicly visible payoff functions associated with parity conditions in EXPSPACE. It is EXPTIME-hard.

- EXPTIME-hardness: same proof as for the distributed synthesis problem [CDHR07]
- Can be extended to (finite) preorders over such objectives

[CDHR07] Chatterjee, Doyen, Henzinger, Raskin. Algorithms for ω-regular games with imperfect information *(LMCS)*

Application to ω-regular objectives (cont'd)

- This amounts to solving two-player turn-based games with generalized (i.e. conjunctions of) ω-regular objectives

Theorem

One can decide the (constrained) existence of a Nash equilibrium in a game with public signal and publicly visible payoff functions associated with parity conditions in EXPSPACE. It is EXPTIME-hard.

- EXPTIME-hardness: same proof as for the distributed synthesis problem [CDHR07]
- Can be extended to (finite) preorders over such objectives
- May even probably be extended to privately visible or invisible payoff functions (needs to be checked)

[CDHR07] Chatterjee, Doyen, Henzinger, Raskin. Algorithms for ω-regular games with imperfect information *(LMCS)*

Application to mean-payoff payoff functions

The mean-payoff payoff publicly visible functions can be used in the epistemic game, and the winning condition for $\texttt{Eve}$ rewrites as:

A strategy for $\texttt{Eve}$ is said winning for payoff $p \in \mathbb{R}^{\mathsf{Agt}}$ from s_0 whenever $\mathsf{MP}(\mathsf{out}_\perp(\sigma_{\texttt{Eve}}, s_0)) = p$, and for every $\rho \in \mathsf{out}(\sigma_{\texttt{Eve}}, s_0)$, for every $A \in \mathsf{susp}(\rho)$, $\mathsf{MP}_A(\rho) \leq p_A$.

[Bre16] Brenguier. Robust equilibria in mean-payoff games *(FoSSaCS'16)*
[BR15] Brenguier, Raskin. Pareto curves of multidimensional mean-payoff games *(CAV'15)*

Application to mean-payoff payoff functions

The mean-payoff payoff publicly visible functions can be used in the epistemic game, and the winning condition for Eve rewrites as:

A strategy for Eve is said winning for payoff $p \in \mathbb{R}^{\mathsf{Agt}}$ from s_0 whenever $\mathsf{MP}(\mathsf{out}_\perp(\sigma_{\mathtt{Eve}}, s_0)) = p$, and for every $\rho \in \mathsf{out}(\sigma_{\mathtt{Eve}}, s_0)$, for every $A \in \mathsf{susp}(\rho)$, $\mathsf{MP}_A(\rho) \leq p_A$.

Inspired by [Bre16], we can reduce the constrained existence problem of a Nash equilibrium to the polyhedron problem [BR15].

[Bre16] Brenguier. Robust equilibria in mean-payoff games *(FoSSaCS'16)*
[BR15] Brenguier, Raskin. Pareto curves of multidimensional mean-payoff games *(CAV'15)*

Application to mean-payoff payoff functions (cont'd)

The polyhedron problem

In a multi-dimensional mean-payoff two-player turn-based game, the polyhedron problem aks, given a polyhedron π, whether there is a strategy for $\mathtt{Eve}$ which ensures a payoff vector which belongs to π.

Application to mean-payoff payoff functions (cont'd)

The polyhedron problem

In a multi-dimensional mean-payoff two-player turn-based game, the polyhedron problem aks, given a polyhedron π, whether there is a strategy for $\mathtt{Eve}$ which ensures a payoff vector which belongs to π.

- [BR15]: if there is a solution, there is one solution with a payoff of polynomial size.

Application to mean-payoff payoff functions (cont'd)

The polyhedron problem

In a multi-dimensional mean-payoff two-player turn-based game, the polyhedron problem aks, given a polyhedron π, whether there is a strategy for $\mathtt{Eve}$ which ensures a payoff vector which belongs to π.

- [BR15]: if there is a solution, there is one solution with a payoff of polynomial size.
- [BR15]: the polyhedron problem is Σ_2P-complete ($\Sigma_2\mathsf{P} = \mathsf{NP}^{\mathsf{NP}}$)

Application to mean-payoff payoff functions (cont'd)

The polyhedron problem

In a multi-dimensional mean-payoff two-player turn-based game, the polyhedron problem aks, given a polyhedron π, whether there is a strategy for $\mathtt{Eve}$ which ensures a payoff vector which belongs to π.

- [BR15]: if there is a solution, there is one solution with a payoff of polynomial size.
- [BR15]: the polyhedron problem is $\Sigma_2\mathsf{P}$-complete ($\Sigma_2\mathsf{P} = \mathsf{NP}^{\mathsf{NP}}$)

Theorem

One can decide the (constrained) existence of a Nash equilibrium in a game with public signal and publicly visible mean-payoff payoff functions, in NP, with a coNEXPTIME oracle. This in particular can be solved in EXPSPACE. It is EXPTIME-hard.

Conclusion

We have:

- proposed a framework for games over graphs with a public signal monitoring
  Note: framework inspired by [Tom98]
- proposed an abstraction called the epistemic game abstraction, which allows to characterize Nash equilibria in the original game
- used it to propose several decidability results.

[Tom98] Tomala. Pure equilibria of repeated games with public observation *(International Journal of Game Theory)*
[RT98] Renault, Tomala. Repeated proximity games *(International Journal of Game Theory)*

Conclusion

We have:

- proposed a framework for games over graphs with a public signal monitoring
  Note: framework inspired by [Tom98]
- proposed an abstraction called the epistemic game abstraction, which allows to characterize Nash equilibria in the original game
- used it to propose several decidability results.

We want:

- work out the precise complexities
- understand whether one can extend the approach to other communication architectures ([RT98]??)

[Tom98] Tomala. Pure equilibria of repeated games with public observation *(International Journal of Game Theory)*
[RT98] Renault, Tomala. Repeated proximity games *(International Journal of Game Theory)*