Permissive Equilibria in Multiplayer Reachability Games

Benjamin Monmege[1]
joint work with Aline Goeminne[2]

1. Aix-Marseille Univ, CNRS, LIS, Marseille, France.
2. F.R.S.-FNRS & UMONS – Université de Mons, Belgium.

CAALM 2025

Permissiveness?

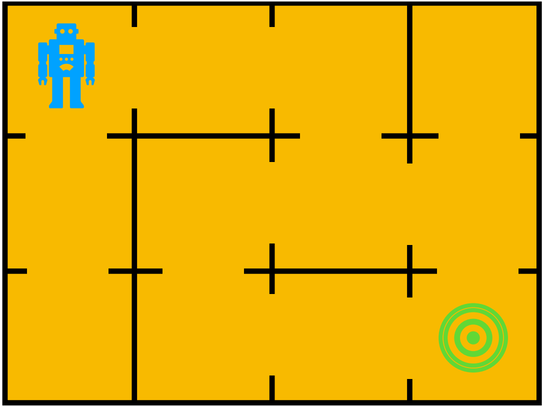

Permissiveness?

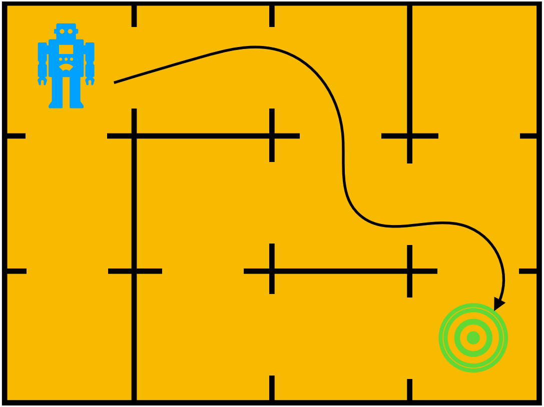

Permissiveness?

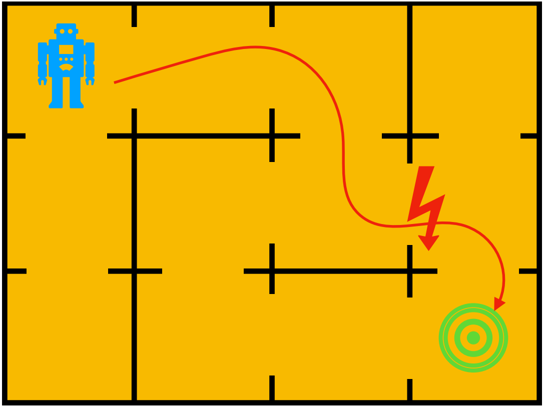

Permissiveness?

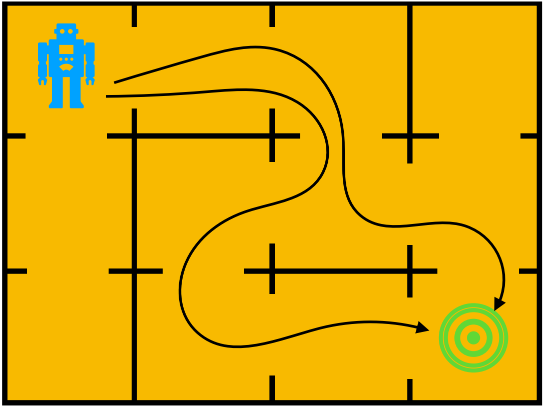

- From discrete controller synthesis, allow more behaviours to be able to react to uncertainty (unavailability of a transition, e.g.)

Permissiveness?

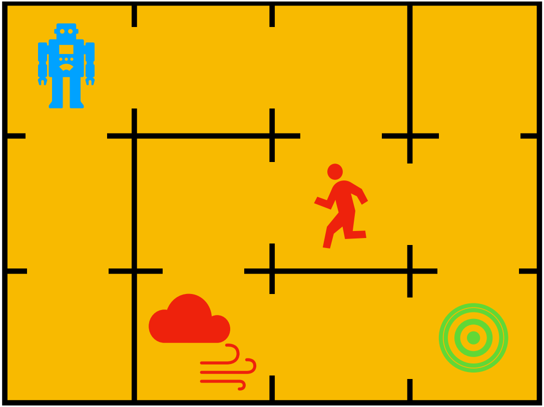

- From discrete controller synthesis, allow more behaviours to be able to react to uncertainty (unavailability of a transition, e.g.)
- In the presence of uncontrollable/antagonistic agents, even more critical!

Permissiveness in games on graphs

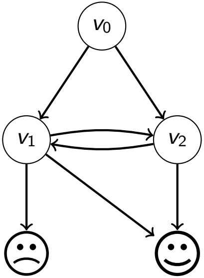

Objective: player ○ wants to reach ☺ from v_0.

[1] Bernet, Janin, Walukiewicz, 2002.

Permissiveness in games on graphs

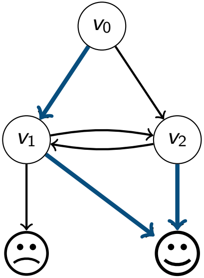

Objective: player ◯ wants to reach ☺ from v_0.

$\rightsquigarrow$ $\sigma_{\bigcirc}$ is a **winning strategy**

[1] Bernet, Janin, Walukiewicz, 2002.

Permissiveness in games on graphs

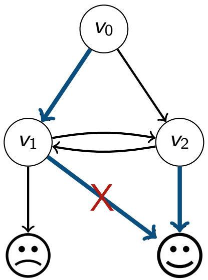

Objective: player ◯ wants to reach ☺ from v_0.

⇝ $\sigma_{\bigcirc}$ is a **winning strategy**

If the edge (v_1, ☺) becomes **unavailable**...

[1] Bernet, Janin, Walukiewicz, 2002.

Permissiveness in games on graphs

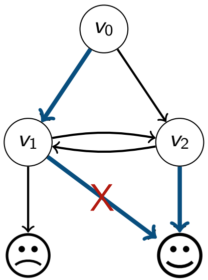

Objective: player ◯ wants to reach ☺ from v_0.

⇝ $\sigma_{\bigcirc}$ is a **winning strategy**

If the edge $(v_1, ☺)$ becomes **unavailable**...

⇝ choosing (v_1, v_2) is also winning

[1] Bernet, Janin, Walukiewicz, 2002.

Permissiveness in games on graphs

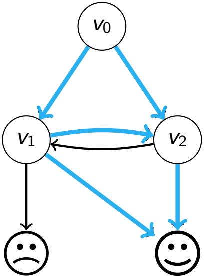

Objective: player ◯ wants to reach ☺ from v_0.

⇝ $\sigma_{\bigcirc}$ is a **winning strategy**

If the edge $(v_1, ☺)$ becomes **unavailable**...

⇝ choosing (v_1, v_2) is also winning
⇝ allowing permissiveness (**multi-strategies**)

[1] Bernet, Janin, Walukiewicz, 2002.

Permissiveness in games on graphs

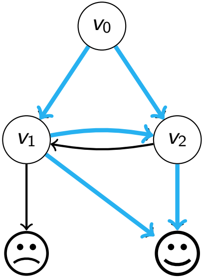

Objective: player ◯ wants to reach ☺ from v_0.

⇝ $\sigma_{\bigcirc}$ is a **winning strategy**

If the edge $(v_1, ☺)$ becomes **unavailable**...

⇝ choosing (v_1, v_2) is also winning
⇝ allowing permissiveness (**multi-strategies**)

- Is there a most permissive winning multi-strategy w.r.t. to set inclusion?

[1] Bernet, Janin, Walukiewicz, 2002.

Permissiveness in games on graphs

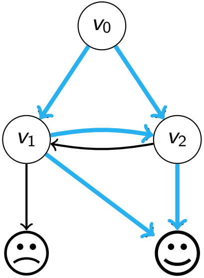

Objective: player ◯ wants to reach ☺ from v_0.

⇝ $\sigma_{\bigcirc}$ is a **winning strategy**

If the edge $(v_1, ☺)$ becomes **unavailable**...

⇝ choosing (v_1, v_2) is also winning
⇝ allowing permissiveness (**multi-strategies**)

- Is there a most permissive winning multi-strategy w.r.t. to set inclusion?
- Not always... but exists for safety games[1]

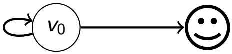

[1]Bernet, Janin, Walukiewicz, 2002.

Permissiveness in games on graphs

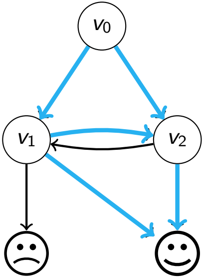

Objective: player ◯ wants to reach ☺ from v_0.

⇝ $\sigma_{\bigcirc}$ is a **winning strategy**

If the edge (v_1, ☺) becomes **unavailable**...

⇝ choosing (v_1, v_2) is also winning
⇝ allowing permissiveness (**multi-strategies**)

- Is there a most permissive winning multi-strategy w.r.t. to set inclusion?
- Not always... but exists for safety games[1]
- In parity games[1]: there exists a finite-memory permissive winning strategy that allows all behaviours of memoryless winning strategies, computable in time $O(n^{d/2+1})$

[1] Bernet, Janin, Walukiewicz, 2002.

Permissiveness in games on graphs

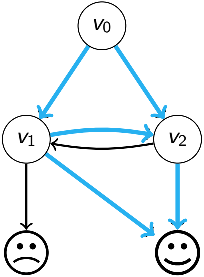

Objective: player ◯ wants to reach ☺ from v_0.

⇝ $\sigma_{\bigcirc}$ is a **winning strategy**

If the edge $(v_1, ☺)$ becomes **unavailable**...

⇝ choosing (v_1, v_2) is also winning
⇝ allowing permissiveness (**multi-strategies**)

- Is there a most permissive winning multi-strategy w.r.t. to set inclusion?
- Not always... but exists for safety games[1]
- In parity games[1]: there exists a finite-memory permissive winning strategy that allows all behaviours of memoryless winning strategies, computable in time $O(n^{d/2+1})$
- Open(?): find a quasi-polynomial algorithm instead?...

[1] Bernet, Janin, Walukiewicz, 2002.

Quantitative comparison of multi-strategies: penalty[2]

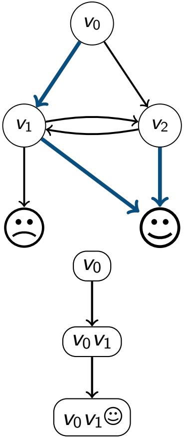

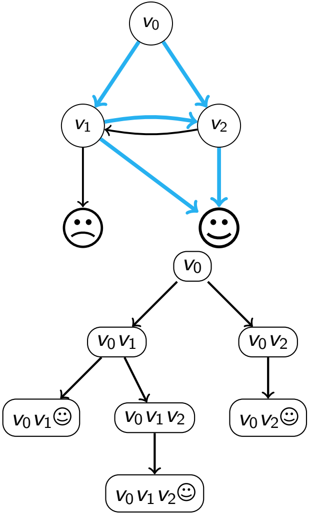

[2] Bouyer, Duflot, Markey, Renault, 2009.

Quantitative comparison of multi-strategies: penalty[2]

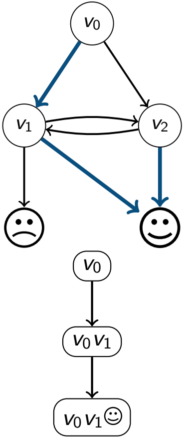

Penalty = 3

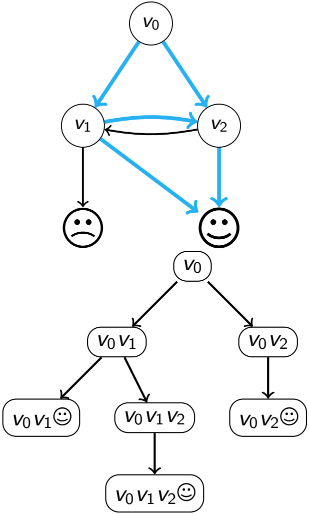

Penalty = 2

[2] Bouyer, Duflot, Markey, Renault, 2009.

Quantitative comparison of multi-strategies: penalty[2]

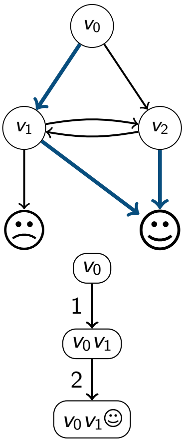

Penalty = 3

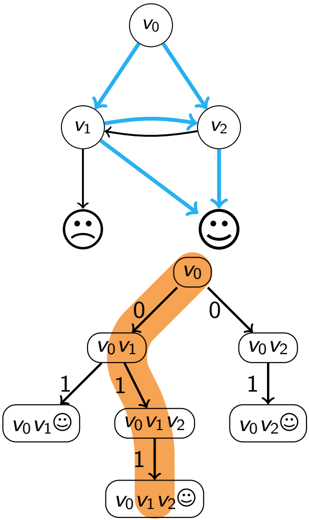

Penalty = 2

[2] Bouyer, Duflot, Markey, Renault, 2009.

Penalty-constrained existence of winning multi-strategies

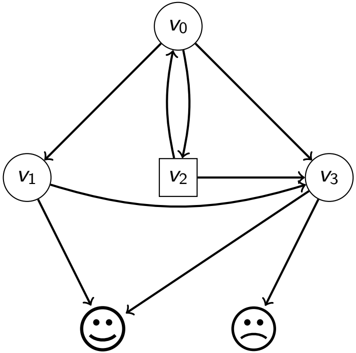

- Reachability objective for player ◯, antagonistic objective for player □

Penalty-constrained existence of winning multi-strategies

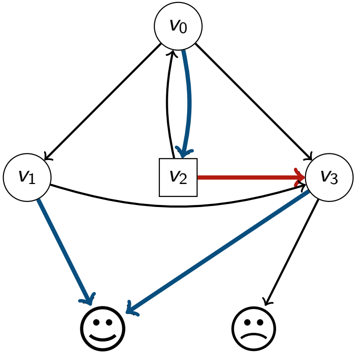

- Reachability objective for player $\bigcirc$, antagonistic objective for player $\square$
- **Strategy for player i:** $\sigma_i : V^* V_i \to V$
  Ex: $(\sigma_{\bigcirc}, \sigma_{\square})$

Penalty-constrained existence of winning multi-strategies

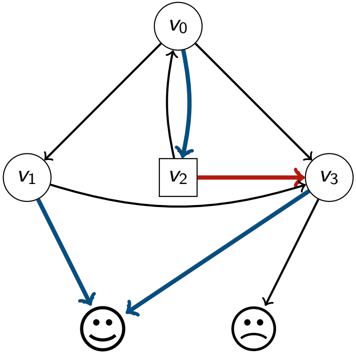

- Reachability objective for player $\bigcirc$, antagonistic objective for player $\square$
- **Strategy for player i:** $\sigma_i : V^* V_i \to V$
  Ex: $(\sigma_{\bigcirc}, \sigma_{\square})$
- $\rightsquigarrow \langle \sigma_{\bigcirc}, \sigma_{\square} \rangle_{v_0}$ the **outcome**.
  Ex: $\langle \sigma_{\bigcirc}, \sigma_{\square} \rangle_{v_0} = v_0 v_2 v_3$☺.

Penalty-constrained existence of winning multi-strategies

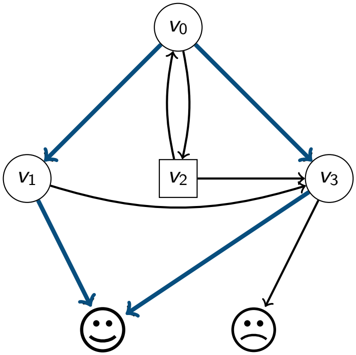

- Reachability objective for player $\bigcirc$, antagonistic objective for player $\square$
- **Strategy for player i:** $\sigma_i : V^* V_i \to V$
  Ex: $(\sigma_{\bigcirc}, \sigma_{\square})$
- $\rightsquigarrow \langle \sigma_{\bigcirc}, \sigma_{\square} \rangle_{v_0}$ the **outcome**.
  Ex: $\langle \sigma_{\bigcirc}, \sigma_{\square} \rangle_{v_0} = v_0 v_2 v_3$☺.
- **Multi-strategy for $\bigcirc$:**
  $\Theta_{\bigcirc} : V^* V_i \to \mathcal{P}(V) \setminus \{\emptyset\}$
- $\Theta_{\bigcirc}$ winning if all consistent plays are winning

Penalty-constrained existence of winning multi-strategies

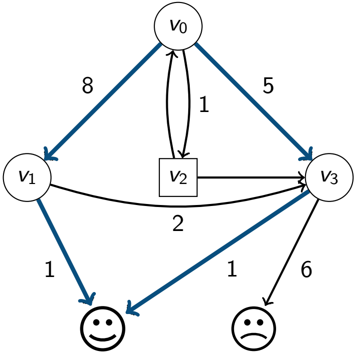

- Reachability objective for player $\bigcirc$, antagonistic objective for player $\square$
- **Strategy for player i:** $\sigma_i : V^* V_i \to V$
  Ex: $(\sigma_{\bigcirc}, \sigma_{\square})$
- $\rightsquigarrow \langle \sigma_{\bigcirc}, \sigma_{\square} \rangle_{v_0}$ the **outcome**.
  Ex: $\langle \sigma_{\bigcirc}, \sigma_{\square} \rangle_{v_0} = v_0 v_2 v_3$☺.
- **Multi-strategy for $\bigcirc$:**
  $\Theta_{\bigcirc} : V^* V_i \to \mathcal{P}(V) \setminus \{\emptyset\}$
- $\Theta_{\bigcirc}$ winning if all consistent plays are winning
- Penalty: maximum (over the consistent plays) of the penalty along that play...
  7 for $\Theta_{\bigcirc}$ because of the play $v_0 v_3$☺

Penalty-constrained existence of winning multi-strategies

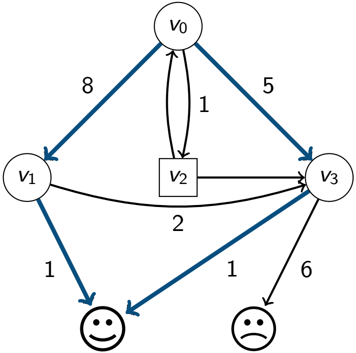

- Reachability objective for player $\bigcirc$, antagonistic objective for player $\square$
- **Strategy for player i:** $\sigma_i : V^* V_i \to V$
  Ex: $(\sigma_{\bigcirc}, \sigma_{\square})$
- $\rightsquigarrow \langle \sigma_{\bigcirc}, \sigma_{\square} \rangle_{v_0}$ the **outcome**.
  Ex: $\langle \sigma_{\bigcirc}, \sigma_{\square} \rangle_{v_0} = v_0 v_2 v_3$☺.
- **Multi-strategy for $\bigcirc$:**
  $\Theta_{\bigcirc} : V^* V_i \to \mathcal{P}(V) \setminus \{\emptyset\}$
- $\Theta_{\bigcirc}$ winning if all consistent plays are winning
- Penalty: maximum (over the consistent plays) of the penalty along that play...
  7 for $\Theta_{\bigcirc}$ because of the play $v_0 v_3$☺

Given a game, and $m \in \mathbb{N}$, $\exists$ a winning multi-strategy $\Theta_{\bigcirc}$ with a penalty at most m?

Penalty-constrained existence of winning multi-strategies

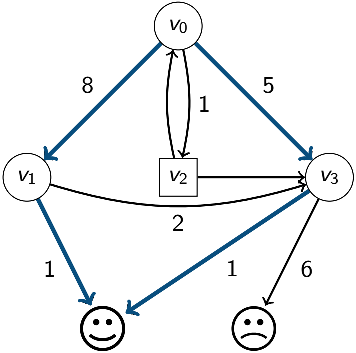

- Reachability objective for player $\bigcirc$, antagonistic objective for player $\square$
- **Strategy for player i:** $\sigma_i : V^* V_i \to V$
  Ex: $(\sigma_\bigcirc, \sigma_\square)$
- $\rightsquigarrow \langle \sigma_\bigcirc, \sigma_\square \rangle_{v_0}$ the **outcome**.
  Ex: $\langle \sigma_\bigcirc, \sigma_\square \rangle_{v_0} = v_0 v_2 v_3$☺.
- **Multi-strategy for $\bigcirc$:**
  $\Theta_\bigcirc : V^* V_i \to \mathcal{P}(V) \setminus \{\emptyset\}$
- $\Theta_\bigcirc$ winning if all consistent plays are winning
- Penalty: maximum (over the consistent plays) of the penalty along that play...
  7 for $\Theta_\bigcirc$ because of the play $v_0 v_3$☺

Given a game, and $m \in \mathbb{N}$, $\exists$ a winning multi-strategy $\Theta_\bigcirc$ with a penalty at most m?

- Can be solved in P [Bouyer, Duflot, Markey, Renault, 2009.]

Equilibria in multiplayer reachability games

Multiplayer reachability games

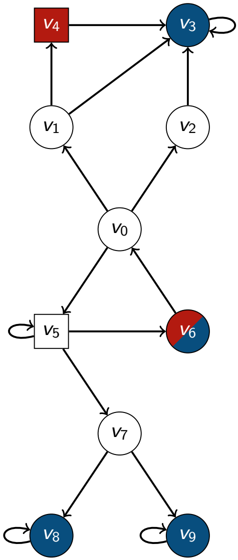

- n players, here only 2: $\bigcirc$ and $\square$
- An initial vertex: v_0
- Target set for each player: $F_{\bigcirc} = \{v_3, v_6, v_8, v_9\}$ and $F_{\square} = \{v_4, v_6\}$

Multiplayer reachability games

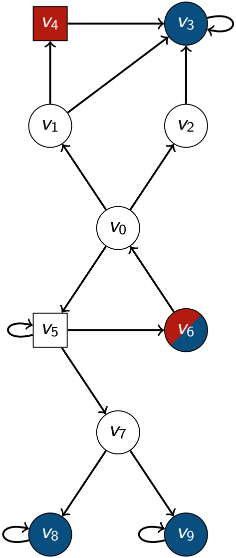

- n players, here only 2: $\bigcirc$ and $\square$
- An initial vertex: v_0
- Target set for each player: $F_\bigcirc = \{v_3, v_6, v_8, v_9\}$ and $F_\square = \{v_4, v_6\}$
- Infinite outcomes: the play continues even when one player has reached their objective
- In $(v_0 v_5 v_6)^\omega$, both players win

Strategies and Nash equilibria

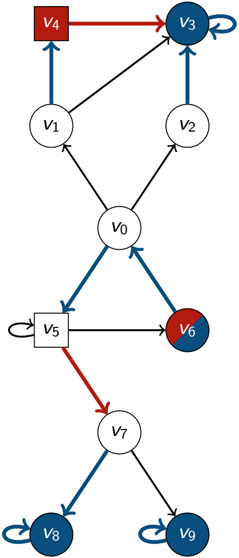

- **Strategy:** $\sigma_i : V^* V_i \to V$
  Ex: $(\sigma_{\bigcirc}, \sigma_{\square})$

Strategies and Nash equilibria

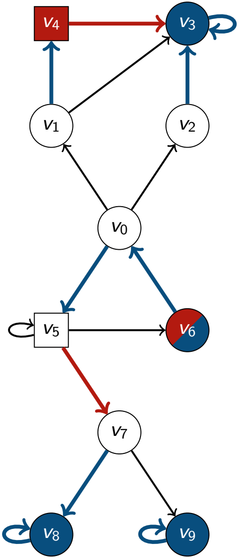

- **Strategy:** $\sigma_i : V^* V_i \to V$
  Ex: $(\sigma_{\bigcirc}, \sigma_{\Box})$
- **Strategy profile:** $\sigma = (\sigma_1, \dots, \sigma_n)$
  $\rightsquigarrow \langle \sigma \rangle_{v_0}$ the **outcome**.
  Ex: $\langle \sigma_{\bigcirc}, \sigma_{\Box} \rangle_{v_0} = v_0 v_5 v_7 v_8^{\omega}$.

Strategies and Nash equilibria

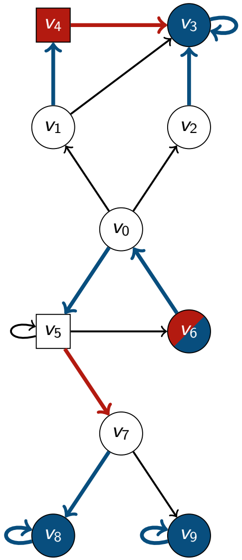

- **Strategy:** $\sigma_i : V^* V_i \to V$
  Ex: $(\sigma_{\bigcirc}, \sigma_{\square})$
- **Strategy profile:** $\sigma = (\sigma_1, \ldots, \sigma_n)$
  $\rightsquigarrow \langle \sigma \rangle_{v_0}$ the **outcome**.
  Ex: $\langle \sigma_{\bigcirc}, \sigma_{\square} \rangle_{v_0} = v_0 v_5 v_7 v_8^{\omega}$.

Nash equilibrium

A **strategy profile** σ is a Nash equilibrium (NE) if **no** player has an incentive to **deviate** unilaterally.

Strategies and Nash equilibria

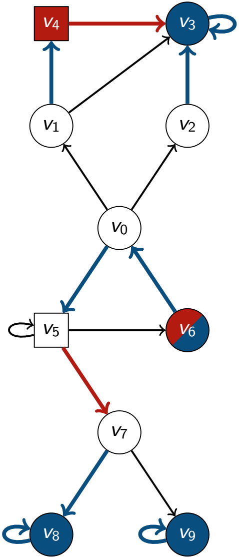

- **Strategy:** $\sigma_i : V^* V_i \to V$
  Ex: $(\sigma_{\bigcirc}, \sigma_{\square})$
- **Strategy profile:** $\sigma = (\sigma_1, \ldots, \sigma_n)$
  $\rightsquigarrow \langle \sigma \rangle_{v_0}$ the **outcome**.
  Ex: $\langle \sigma_{\bigcirc}, \sigma_{\square} \rangle_{v_0} = v_0 v_5 v_7 v_8^\omega$.

Nash equilibrium

A **strategy profile** σ is a Nash equilibrium (NE) if **no** player has an incentive to **deviate** unilaterally.

Counter-example:

- $(\sigma_{\bigcirc}, \sigma_{\square})$ is **not** an NE
- In $\langle \sigma_{\bigcirc}, \sigma_{\square} \rangle_{v_0}$, only player $\bigcirc$ wins

Strategies and Nash equilibria

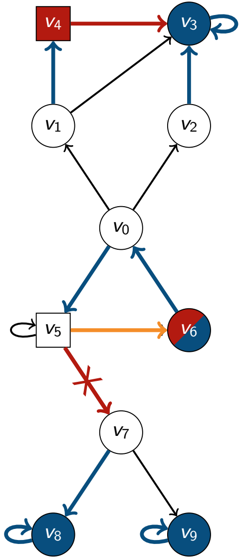

- **Strategy:** $\sigma_i : V^* V_i \to V$
  Ex: $(\sigma_{\bigcirc}, \sigma_{\square})$
- **Strategy profile:** $\sigma = (\sigma_1, \dots, \sigma_n)$
  $\rightsquigarrow \langle \sigma \rangle_{v_0}$ the **outcome**.
  Ex: $\langle \sigma_{\bigcirc}, \sigma_{\square} \rangle_{v_0} = v_0 v_5 v_7 v_8^{\omega}$.

Nash equilibrium

A **strategy profile** σ is a Nash equilibrium (NE) if **no** player has an incentive to **deviate** unilaterally.

Counter-example:

- $(\sigma_{\bigcirc}, \sigma_{\square})$ is **not** an NE
- In $\langle \sigma_{\bigcirc}, \sigma_{\square} \rangle_{v_0}$, only player $\bigcirc$ wins
- $\sigma_{\square}$ is a profitable deviation
- $\langle \sigma_{\bigcirc}, \sigma_{\square} \rangle_{v_0} = (v_0 v_5 v_6)^{\omega}$ winning for $\square$

What is known...

- Nash equilibria always exist in reachability games, and even for objectives that are prefix-independent...[3]

[3] Brihaye, De Pril, and Schewe, 2013
[4] Bruyère, 2017
[5] Ummels, 2006. Condurache, Filiot, Raskin, 2016.
[6] Brihaye, Bruyère, Goeminne, Thomasset, 2019

What is known...

- Nash equilibria always exist in reachability games, and even for objectives that are prefix-independent...[3] and finite memory suffices if the corresponding two-player zero-sum games have optimal finite-memory strategies

[3] Brihaye, De Pril, and Schewe, 2013

[4] Bruyère, 2017

[5] Ummels, 2006. Condurache, Filiot, Raskin, 2016.

[6] Brihaye, Bruyère, Goeminne, Thomasset, 2019

What is known...

- Nash equilibria always exist in reachability games, and even for objectives that are prefix-independent...[3] and finite memory suffices if the corresponding two-player zero-sum games have optimal finite-memory strategies

Does there exist a NE where a subset of players wins?

[3] Brihaye, De Pril, and Schewe, 2013
[4] Bruyère, 2017
[5] Ummels, 2006. Condurache, Filiot, Raskin, 2016.
[6] Brihaye, Bruyère, Goeminne, Thomasset, 2019

What is known...

- Nash equilibria always exist in reachability games, and even for objectives that are prefix-independent...[3] and finite memory suffices if the corresponding two-player zero-sum games have optimal finite-memory strategies

Does there exist a NE where a subset of players wins?

Characterization of outcomes ρ of Nash equilibria

There exists a NE σ whose outcome is ρ
if and only if
at every vertex v of the play, belonging to player i, if player i has a winning strategy from v, then they win in ρ.

[3] Brihaye, De Pril, and Schewe, 2013
[4] Bruyère, 2017
[5] Ummels, 2006. Condurache, Filiot, Raskin, 2016.
[6] Brihaye, Bruyère, Goeminne, Thomasset, 2019

What is known...

- Nash equilibria always exist in reachability games, and even for objectives that are prefix-independent...[3] and finite memory suffices if the corresponding two-player zero-sum games have optimal finite-memory strategies

Does there exist a NE where a subset of players wins?

Characterization of outcomes ρ of Nash equilibria

There exists a NE σ whose outcome is ρ
if and only if
at every vertex v of the play, belonging to player i, if player i has a winning strategy from v, then they win in ρ.

- P-complete for Büchi and Müller games[4]; NP-complete for reachability, safety, co-Büchi, parity and Streett games [5]

[3] Brihaye, De Pril, and Schewe, 2013
[4] Bruyère, 2017
[5] Ummels, 2006. Condurache, Filiot, Raskin, 2016.
[6] Brihaye, Bruyère, Goeminne, Thomasset, 2019

What is known...

- Nash equilibria always exist in reachability games, and even for objectives that are prefix-independent...[3] and finite memory suffices if the corresponding two-player zero-sum games have optimal finite-memory strategies

Does there exist a NE where a subset of players wins?

Characterization of outcomes ρ of Nash equilibria

There exists a NE σ whose outcome is ρ
if and only if
at every vertex v of the play, belonging to player i, if player i has a winning strategy from v, then they win in ρ.

- P-complete for Büchi and Müller games[4]; NP-complete for reachability, safety, co-Büchi, parity and Streett games [5]
- P-complete for limsup games
- NP-complete for quantitative reachability[6], supremum, infimum, liminf, mean-payoff

[3] Brihaye, De Pril, and Schewe, 2013
[4] Bruyère, 2017
[5] Ummels, 2006. Condurache, Filiot, Raskin, 2016.
[6] Brihaye, Bruyère, Goeminne, Thomasset, 2019

What is known...

- Nash equilibria always exist in reachability games, and even for objectives that are prefix-independent...[3] and finite memory suffices if the corresponding two-player zero-sum games have optimal finite-memory strategies

Does there exist a NE where a subset of players wins?

Characterization of outcomes ρ of Nash equilibria

There exists a NE σ whose outcome is ρ
if and only if
at every vertex v of the play, belonging to player i, if player i has a winning strategy from v, then they win in ρ.

- P-complete for Büchi and Müller games[4]; NP-complete for reachability, safety, co-Büchi, parity and Streett games [5]
- P-complete for limsup games
- NP-complete for quantitative reachability[6], supremum, infimum, liminf, mean-payoff
- open for discounted games

[3] Brihaye, De Pril, and Schewe, 2013
[4] Bruyère, 2017
[5] Ummels, 2006. Condurache, Filiot, Raskin, 2016.
[6] Brihaye, Bruyère, Goeminne, Thomasset, 2019

Multi-strategies and permissive Nash equilibria

Multi-strategies

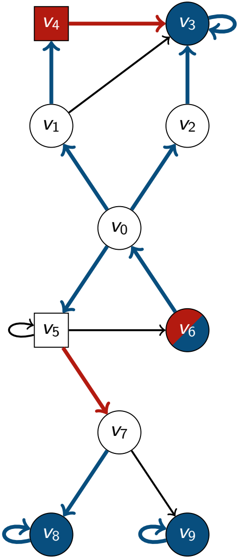

- **Multi-strategy (for both players):**
  $\Theta_i : V^* V_i \to \mathcal{P}(V) \setminus \{\emptyset\}$
  Ex: $(\Theta_{\bigcirc}, \Theta_{\square})$

Multi-strategies

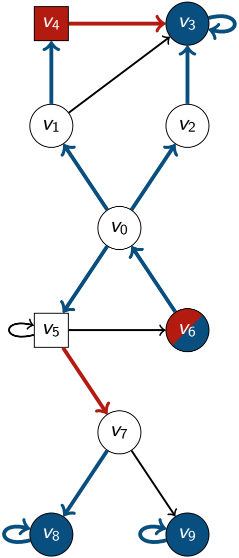

- **Multi-strategy (for both players):**
  $\Theta_i : V^* V_i \to \mathcal{P}(V) \setminus \{\emptyset\}$
  Ex: $(\Theta_{\bigcirc}, \Theta_{\Box})$
- **Multi-strategy profile:** $\Theta = (\Theta_1, \dots, \Theta_n)$
  $\rightsquigarrow \langle \Theta \rangle_{v_0}$ the **set** of outcomes
  Ex: $\langle \Theta_{\bigcirc}, \Theta_{\Box} \rangle_{v_0} = \{ v_0 v_1 v_4 v_3^\omega, v_0 v_2 v_3^\omega, v_0 v_5 v_7 v_8^\omega \}$

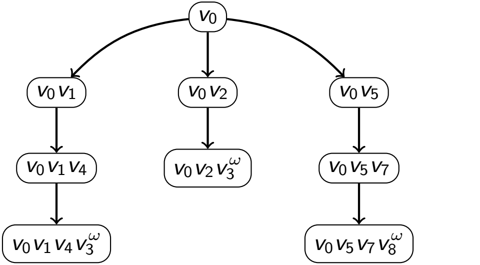

- can be seen as a **tree** $\mathcal{T}$

Permissive Nash equilibria

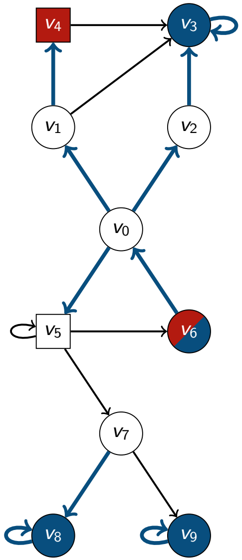

- a strategy σ_i is **consistent** with a multi-strategy Θ_i if for all $hv \in V^* V_i$:

$$\sigma_i(hv) \in \Theta_i(hv).$$

Permissive Nash equilibria

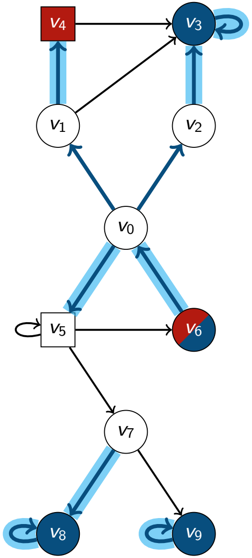

- a strategy σ_i is **consistent** with a multi-strategy Θ_i if for all $hv \in V^* V_i$:

$$\sigma_i(hv) \in \Theta_i(hv).$$

Permissive Nash equilibria

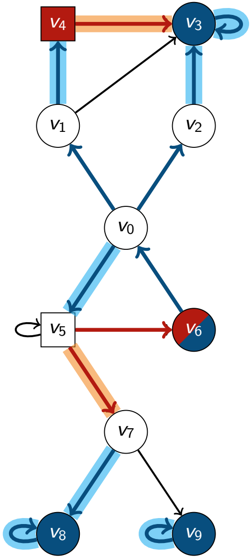

- a strategy σ_i is **consistent** with a multi-strategy Θ_i if for all $hv \in V^* V_i$:

$$\sigma_i(hv) \in \Theta_i(hv).$$

- a strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ is **consistent** with a multi-strategy profile $\Theta = (\Theta_1, \ldots, \Theta_n)$ if for each $1 \leq i \leq n$, σ_i is consistent with Θ_i.

Permissive Nash equilibria

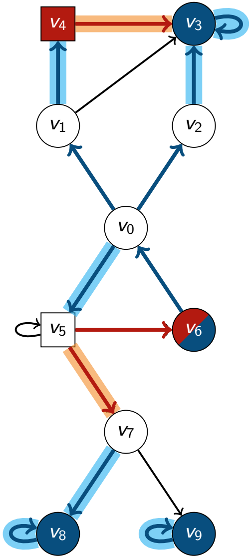

- a strategy σ_i is **consistent** with a multi-strategy Θ_i if for all $hv \in V^* V_i$:

$$\sigma_i(hv) \in \Theta_i(hv).$$

- a strategy profile $\sigma = (\sigma_1, \ldots, \sigma_n)$ is **consistent** with a multi-strategy profile $\Theta = (\Theta_1, \ldots, \Theta_n)$ if for each $1 \leq i \leq n$, σ_i is consistent with Θ_i.

Permissive Nash equilibrium

A **multi-strategy profile** Θ is a permissive NE if each strategy profile σ consistent with Θ is an NE.

Permissive Nash equilibria

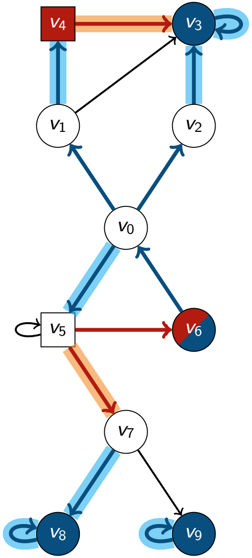

- a strategy σ_i is **consistent** with a multi-strategy Θ_i if for all $hv \in V^* V_i$:

$$\sigma_i(hv) \in \Theta_i(hv).$$

- a strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ is **consistent** with a multi-strategy profile $\Theta = (\Theta_1, \dots, \Theta_n)$ if for each $1 \leq i \leq n$, σ_i is consistent with Θ_i.

Permissive Nash equilibrium

A **multi-strategy profile** Θ is a permissive NE if each strategy profile σ consistent with Θ is an NE.

Counter-example:

- $(\Theta_{\bigcirc}, \Theta_{\Box})$ is **not** a permissive NE;
- because $(\sigma_{\bigcirc}, \sigma_{\Box})$ is **not** an NE.

Penalties

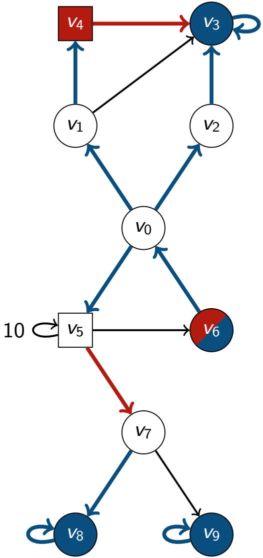

- $w : E \to \mathbb{N}$ a **weight function** (by default: 1)

Penalties

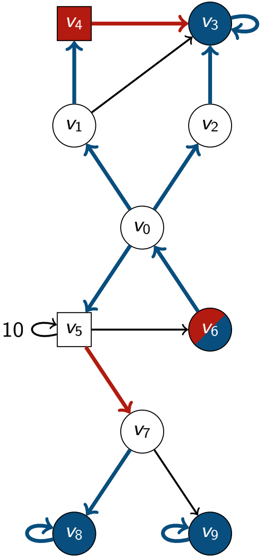

- $w : E \to \mathbb{N}$ a **weight function** (by default: 1)

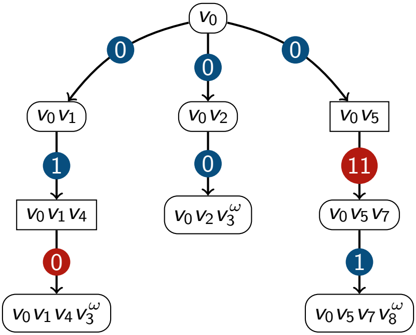

Penalties

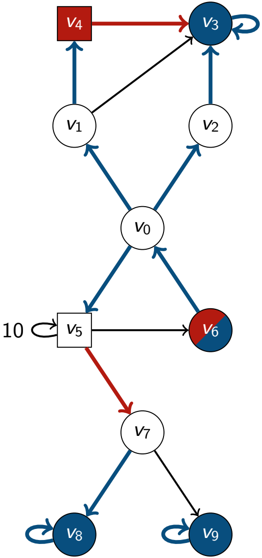

- $w : E \to \mathbb{N}$ a **weight function** (by default: 1)

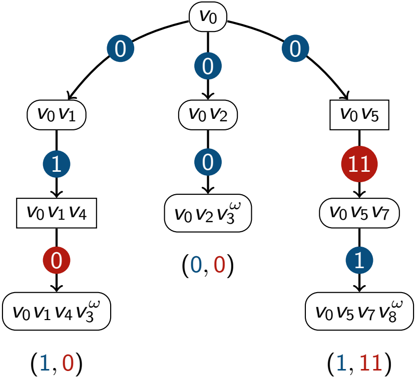

Penalties

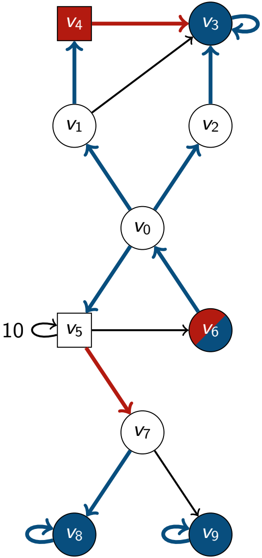

- $w : E \to \mathbb{N}$ a **weight function** (by default: 1)

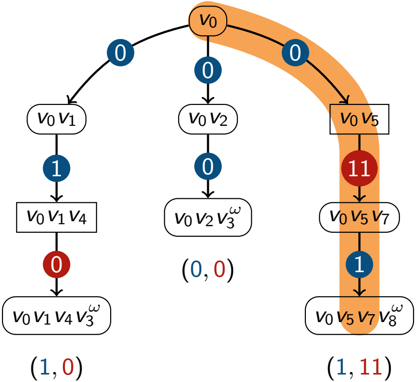

Penalties : $(1, 11)$

Studied problems

Constrained penalty problems

Constrained penalty problem

Given $(m_1, \ldots, m_n) \in (\mathbb{N} \cup \{+\infty\})^n$, does there exist a **permissive NE** Θ such that for each $1 \leq i \leq n$:

$$\mathsf{Penalty}_i(\Theta) \leq m_i.$$

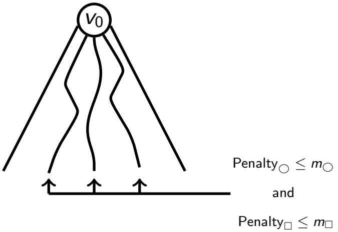

Constrained penalty problems

Constrained penalty problem

Given $(m_1, \ldots, m_n) \in (\mathbb{N} \cup \{+\infty\})^n$, does there exist a **permissive NE** Θ such that for each $1 \leq i \leq n$:

$$\text{Penalty}_i(\Theta) \leq m_i.$$

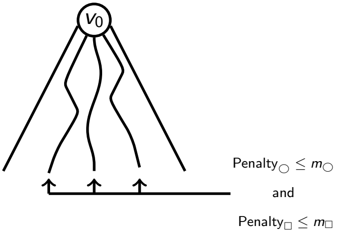

Strongly winning with constrained penalty problem

Given $(m_1, \ldots, m_n) \in (\mathbb{N} \cup \{+\infty\})^n$ and a coalition Win,
does there exist a **permissive NE** Θ such that for each player i:

$$\text{Penalty}_i(\Theta) \leq m_i$$

and Θ is **strongly winning w.r.t. Win.**

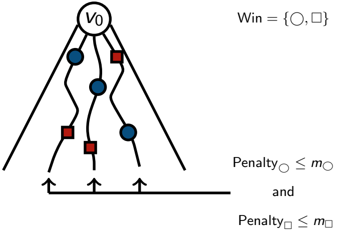

Constrained penalty problems

Constrained penalty problem

Given $(m_1, \ldots, m_n) \in (\mathbb{N} \cup \{+\infty\})^n$, does there exist a **permissive NE** Θ such that for each $1 \leq i \leq n$:

$$\mathrm{Penalty}_i(\Theta) \leq m_i.$$

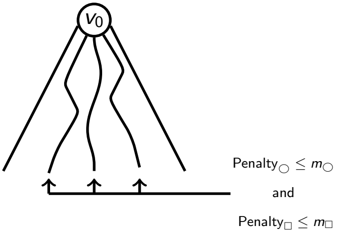

Strongly winning with constrained penalty problem

Given $(m_1, \ldots, m_n) \in (\mathbb{N} \cup \{+\infty\})^n$ and a coalition Win,
does there exist a **permissive NE** Θ such that for each player i:

$$\mathrm{Penalty}_i(\Theta) \leq m_i$$

and Θ is **strongly winning w.r.t. Win.**

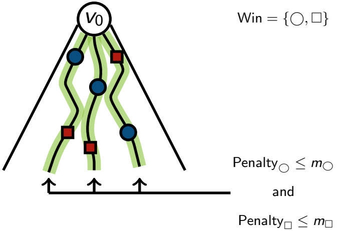

Constrained penalty problems

Weakly winning with constrained penalty problem

Given $(m_1, \ldots, m_n) \in (\mathbb{N} \cup \{+\infty\})^n$ and a coalition Win,
does there exist a **permissive NE** Θ such that for each player i:

$$\text{Penalty}_i(\Theta) \leq m_i$$

and Θ is **weakly winning w.r.t.** Win.

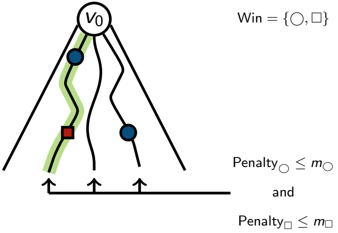

Constrained penalty problems

Weakly winning with constrained penalty problem

Given $(m_1, \ldots, m_n) \in (\mathbb{N} \cup \{+\infty\})^n$ and a coalition Win,
does there exist a **permissive NE** Θ such that for each player i:

$$\mathsf{Penalty}_i(\Theta) \leq m_i$$

and Θ is **weakly winning w.r.t. Win.**

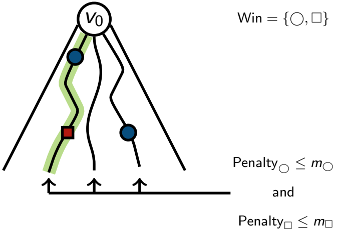

Theorem: If $m_1, \ldots, m_n$ are encoded in **unary**,
the constrained penalty problems belong to PSPACE.

How to solve these problems?

Key idea

Characterization of Outcomes of permissive Nash equilibria

Let $\mathcal{T}$ be an infinite tree,

$$\text{there exists a permissive NE } (\Theta_1, \ldots, \Theta_n) \text{ such that}$$
$$\langle \Theta_1, \ldots, \Theta_n \rangle_{v_0} = \mathcal{T}$$
$$\text{if and only if}$$
$$\mathcal{T} \text{ is a } \underline{\text{good}} \text{ tree.}$$

Key idea

Characterization of Outcomes of permissive Nash equilibria

Let $\mathcal{T}$ be an infinite tree,

there exists a permissive NE $(\Theta_1, \ldots, \Theta_n)$ such that
$$\langle \Theta_1, \ldots, \Theta_n \rangle_{v_0} = \mathcal{T}$$
if and only if
$\mathcal{T}$ is a good tree.

$\rightsquigarrow$ Does there exist a tree $\mathcal{T}$ such that

- each $\rho \in \mathcal{T}$ and each player i, $\mathrm{Penalty}_i(\rho) \leq m_i$;
- $\mathcal{T}$ satisfies the property of being strongly/weakly winning;
- $\mathcal{T}$ is a good tree.

Characterization of outcomes of permissive Nash equilibria

Good tree

No internal deviations

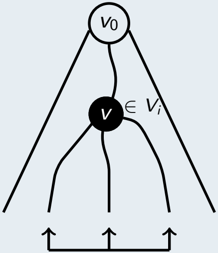

player i is winning everywhere/nowhere

Characterization of outcomes of permissive Nash equilibria

Good tree

No internal deviations

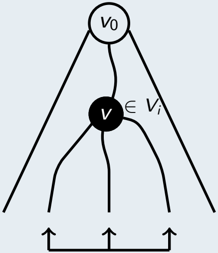

player i is winning everywhere/nowhere

Intuition

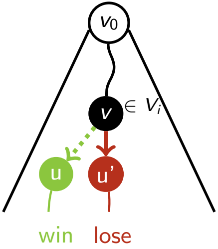

Characterization of outcomes of permissive Nash equilibria

Internal deviation

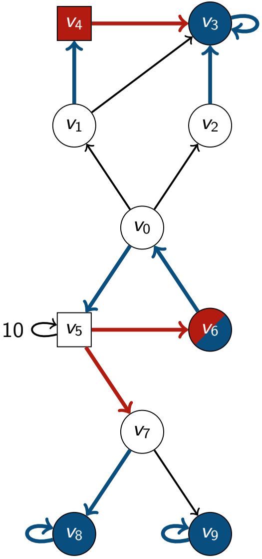

- $(\Theta_{\bigcirc}, \Theta_{\Box})$ is **not** a permissive NE.

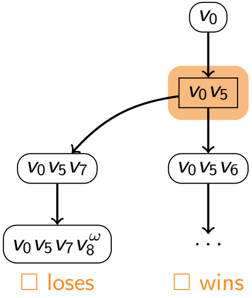

- $\mathcal{T}$ **cannot** be the outcome of a permissive NE.

Characterization of outcomes of permissive Nash equilibria

Good tree

No external deviations

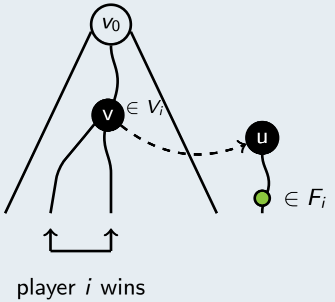

Characterization of outcomes of permissive Nash equilibria

Good tree

No external deviations

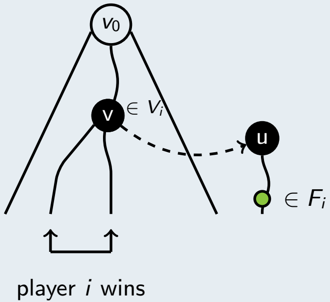

Intuition

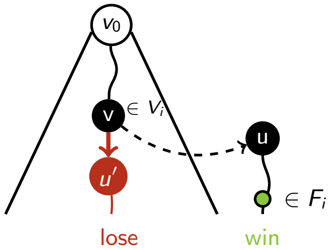

Characterization of outcomes of permissive Nash equilibria

External deviation

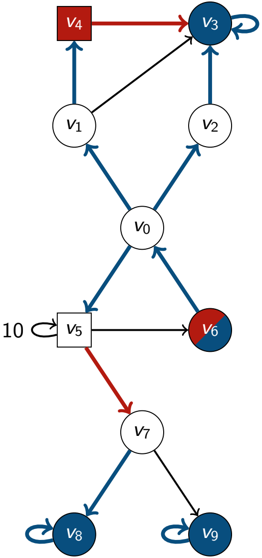

- $(\Theta_{\bigcirc}, \Theta_{\square})$ is **not** a permissive NE.

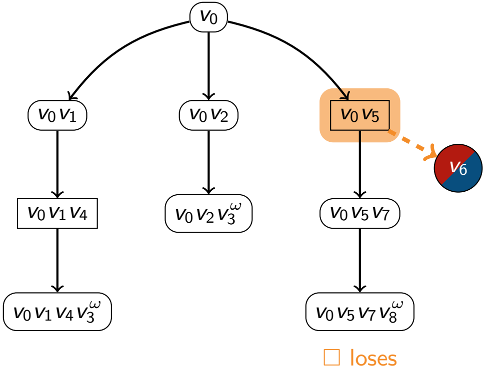

- the tree **cannot** be the outcome of a permissive NE

Deciding the constrained penalty problems

Finite symbolic tree

If there exists a tree $\mathcal{T}$ that

- satisfies the constraints given by the problem;
- is good;

then there exists a tree $\mathcal{T}'$ that

- also satisfies the constraints and is good;
- has a **finite representation** where leaves are sent back to ancesters, and the finite tree has polynomial height.

Deciding the constrained penalty problems

Finite symbolic tree

If there exists a tree $\mathcal{T}$ that

- satisfies the constraints given by the problem;
- is good;

then there exists a tree $\mathcal{T}'$ that

- also satisfies the constraints and is good;
- has a **finite representation** where leaves are sent back to ancesters, and the finite tree has polynomial height.

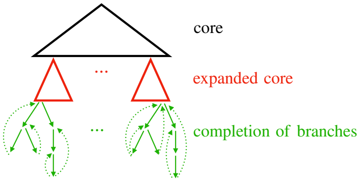

- core: finite subtree of $\mathcal{T}$ where all players of Win have won (finite by König's lemma)

Deciding the constrained penalty problems

Finite symbolic tree

If there exists a tree $\mathcal{T}$ that

- satisfies the constraints given by the problem;
- is good;

then there exists a tree $\mathcal{T}'$ that

- also satisfies the constraints and is good;
- has a **finite representation** where leaves are sent back to ancesters, and the finite tree has polynomial height.

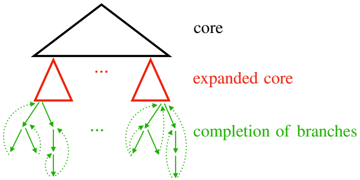

- core: finite subtree of $\mathcal{T}$ where all players of Win have won (finite by König's lemma)
- expanded core: continue to expand the tree so that we create no internal deviations (new players must win...)

Deciding the constrained penalty problems

Finite symbolic tree

If there exists a tree $\mathcal{T}$ that

- satisfies the constraints given by the problem;
- is good;

then there exists a tree $\mathcal{T}'$ that

- also satisfies the constraints and is good;
- has a **finite representation** where leaves are sent back to ancesters, and the finite tree has polynomial height.

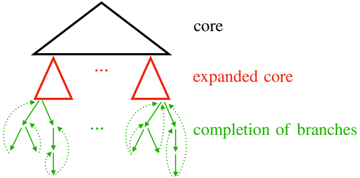

- core: finite subtree of $\mathcal{T}$ where all players of Win have won (finite by König's lemma)
- expanded core: continue to expand the tree so that we create no internal deviations (new players must win...)
- continue to complete the branches until all successor of a node are of a similar type (taken from a finite set) to some ancestors: these completions have polynomial length

Deciding the constrained penalty problems

Finite symbolic tree

If there exists a tree $\mathcal{T}$ that

- satisfies the constraints given by the problem;
- is good;

then there exists a tree $\mathcal{T}'$ that

- also satisfies the constraints and is good;
- has a **finite representation** where leaves are sent back to ancesters, and the finite tree has polynomial height.

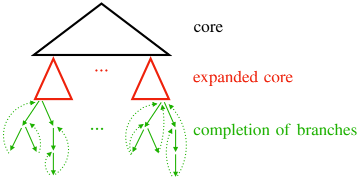

- core: finite subtree of $\mathcal{T}$ where all players of Win have won (finite by König's lemma)
- expanded core: continue to expand the tree so that we create no internal deviations (new players must win...)
- continue to complete the branches
- finally, compress the core and expanded core (by copy-pasting the subtrees) to make their height polynomial

Deciding the constrained penalty problems

Finite symbolic tree

If there exists a tree $\mathcal{T}$ that

- satisfies the constraints given by the problem;
- is good;

then there exists a tree $\mathcal{T}'$ that

- also satisfies the constraints and is good;
- has a **finite representation** where leaves are sent back to ancesters, and the finite tree has polynomial height.

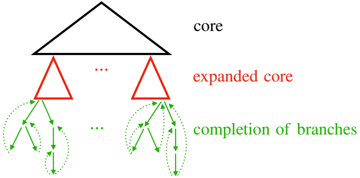

- core: finite subtree of $\mathcal{T}$ where all players of Win have won (finite by König's lemma)
- expanded core: continue to expand the tree so that we create no internal deviations (new players must win...)
- continue to complete the branches
- finally, compress the core and expanded core

- This finite symbolic tree and the characterization of the outcomes of permissive NEs $\rightsquigarrow$ APTIME algorithm if thresholds are encoded in unary.

Generalisation to subgame perfect equilibria

Subgame perfect equilibria: avoid non-credible threats

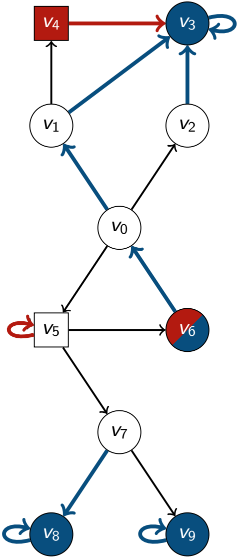

- Nash equilibrium, but player □ plays a non-credible threat in v_5: going to v_6 is profitable for them...

Subgame perfect equilibria: avoid non-credible threats

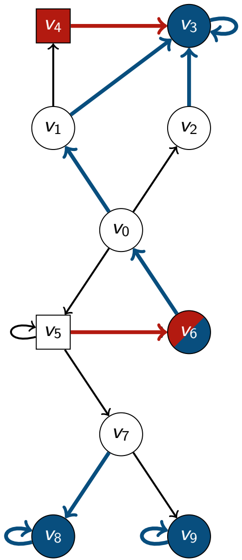

- Nash equilibrium, but player □ plays a non-credible threat in v_5: going to v_6 is profitable for them...

Subgame perfect equilibrium

A **strategy profile** σ is a subgame perfect equilibrium (SPE) if it is a Nash equilibrium in all subgames (from all possible histories from v_0).

Subgame perfect equilibria: avoid non-credible threats

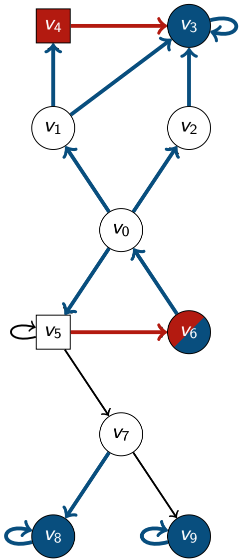

- Nash equilibrium, but player □ plays a non-credible threat in v_5: going to v_6 is profitable for them...

Subgame perfect equilibrium

A **strategy profile** σ is a subgame perfect equilibrium (SPE) if it is a Nash equilibrium in all subgames (from all possible histories from v_0).

Permissive SPE

A **multi-strategy profile** Θ is a permissive SPE if each strategy profile σ consistent with Θ is an SPE.

Subgame perfect equilibria: avoid non-credible threats

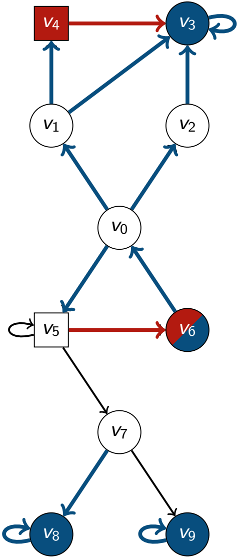

- Nash equilibrium, but player □ plays a non-credible threat in v_5: going to v_6 is profitable for them...

Subgame perfect equilibrium

A **strategy profile** σ is a subgame perfect equilibrium (SPE) if it is a Nash equilibrium in all subgames (from all possible histories from v_0).

Permissive SPE

A **multi-strategy profile** Θ is a permissive SPE if each strategy profile σ consistent with Θ is an SPE.

- The other edges from v_5 are disallowed in any permissive SPE: more robustness!
- Two notions of penalty: the main penalty as before, and the retaliation penalty considering all other subgames

What is known...

- SPE always exist in reachability games[7], and other qualitative objectives like Borel objectives[8] and quantitative reachability games[9].

[7] Brihaye, Bruyère, De Pril, Gimbert, 2012
[8] Grädel, Ummels, 2008.
[9] De Pril, 2013.
[10] Brice, Raskin, van den Bogaard, 2022.
[11] Brihaye, Bruyère, Goeminne, Raskin, 2018.
[12] Brihaye, Bruyère, Goeminne, Raskin, van den Bogaard, 2019

What is known...

- SPE always exist in reachability games[7], and other qualitative objectives like Borel objectives[8] and quantitative reachability games[9].

Does there exist an SPE where a subset of players wins?

[7] Brihaye, Bruyère, De Pril, Gimbert, 2012
[8] Grädel, Ummels, 2008.
[9] De Pril, 2013.
[10] Brice, Raskin, van den Bogaard, 2022.
[11] Brihaye, Bruyère, Goeminne, Raskin, 2018.
[12] Brihaye, Bruyère, Goeminne, Raskin, van den Bogaard, 2019

What is known...

- SPE always exist in reachability games[7], and other qualitative objectives like Borel objectives[8] and quantitative reachability games[9].

Does there exist an SPE where a subset of players wins?

Characterization of outcomes of SPEs
...

[7] Brihaye, Bruyère, De Pril, Gimbert, 2012
[8] Grädel, Ummels, 2008.
[9] De Pril, 2013.
[10] Brice, Raskin, van den Bogaard, 2022.
[11] Brihaye, Bruyère, Goeminne, Raskin, 2018.
[12] Brihaye, Bruyère, Goeminne, Raskin, van den Bogaard, 2019

What is known...

- SPE always exist in reachability games[7], and other qualitative objectives like Borel objectives[8] and quantitative reachability games[9].

Does there exist an SPE where a subset of players wins?

Characterization of outcomes of SPEs
...

- NP-complete for parity games and mean-payoff games[10]
- PSPACE-complete for reachability and safety[11], and for quantitative reachability games[12]

[7] Brihaye, Bruyère, De Pril, Gimbert, 2012
[8] Grädel, Ummels, 2008.
[9] De Pril, 2013.
[10] Brice, Raskin, van den Bogaard, 2022.
[11] Brihaye, Bruyère, Goeminne, Raskin, 2018.
[12] Brihaye, Bruyère, Goeminne, Raskin, van den Bogaard, 2019

Constrained penalty problems

Constrained penalty problem

Given $(m_1, \dots, m_n) \in (\mathbb{N} \cup \{+\infty\})^n$ and $(r_1, \dots, r_n) \in (\mathbb{N} \cup \{+\infty\})^n$, does there exist a **permissive SPE** Θ such that for each player i:

$$\mathrm{Penalty}_i(\Theta) \leq m_i \qquad \text{and} \qquad \text{Retaliate-penalty}_i(\Theta) \leq r_i$$

Constrained penalty problems

Constrained penalty problem

Given $(m_1, \dots, m_n) \in (\mathbb{N} \cup \{+\infty\})^n$ and $(r_1, \dots, r_n) \in (\mathbb{N} \cup \{+\infty\})^n$, does there exist a **permissive SPE** Θ such that for each player i:

$$\text{Penalty}_i(\Theta) \leq m_i \qquad \text{and} \qquad \text{Retaliate-penalty}_i(\Theta) \leq r_i$$

Strongly/weakly winning with constrained penalty problem

Given $(m_1, \dots, m_n) \in (\mathbb{N} \cup \{+\infty\})^n$ and a coalition Win, does there exist a **permissive SPE** Θ such that for each player i:

$$\text{Penalty}_i(\Theta) \leq m_i \qquad \text{and} \qquad \text{Retaliate-penalty}_i(\Theta) \leq r_i$$

and Θ is **strongly/weakly winning w.r.t.** Win.

Constrained penalty problems

Theorem: The constrained penalty problems for SPEs belong to PSPACE if penalty upperbounds are encoded in unary.

Constrained penalty problems

Theorem: The constrained penalty problems for SPEs belong to PSPACE if penalty upperbounds are encoded in unary.

Characterization of outcomes of permissive SPEs

Let $\mathcal{T}$ be an infinite tree,

there exists a permissive SPE $(\Theta_1, \ldots, \Theta_n)$ such that
$\langle \Theta_1, \ldots, \Theta_n \rangle_{v_0} = \mathcal{T}$
if and only if
there is a good forest (indexed by initial history) that has $\mathcal{T}$ has tree from v_0.

Constrained penalty problems

Theorem: The constrained penalty problems for SPEs belong to PSPACE if penalty upperbounds are encoded in unary.

Characterization of outcomes of permissive SPEs

Let $\mathcal{T}$ be an infinite tree,

there exists a permissive SPE $(\Theta_1, \ldots, \Theta_n)$ such that
$\langle \Theta_1, \ldots, \Theta_n \rangle_{v_0} = \mathcal{T}$
if and only if
there is a good forest (indexed by initial history) that has $\mathcal{T}$ has tree from v_0.

- The rule over external deviations in good trees is now replaced by the absence of deviations in-between the trees of the forest...

Constrained penalty problems

Theorem: The constrained penalty problems for SPEs belong to PSPACE if penalty upperbounds are encoded in unary.

Characterization of outcomes of permissive SPEs

Let $\mathcal{T}$ be an infinite tree,

there exists a permissive SPE $(\Theta_1, \ldots, \Theta_n)$ such that
$\langle \Theta_1, \ldots, \Theta_n \rangle_{v_0} = \mathcal{T}$
if and only if
there is a good forest (indexed by initial history) that has $\mathcal{T}$ has tree from v_0.

- The rule over external deviations in good trees is now replaced by the absence of deviations in-between the trees of the forest...
- Still able to find a compact representation of good forests, and thus a PSPACE algorithm

Conclusion

- permissiveness in multiplayer reachability games $\rightsquigarrow$ permissive equilibria (Nash equilibria, subgame perfect equilibria)
- penalties to compare multi-strategies (main penalties, retaliation penalties)
- decision problems related to the existence of permissive equilibria with constraints on the penalties
- relevant permissive equilibria
  $\rightsquigarrow$ strongly/weakly winning with constrained penalty problems
- those problems belong to PSPACE if the thresholds are encoded in unary

Perspectives

- decrease the space dependency to be only polynomial in the logarithm of the penalty thresholds? or matching lower bound on complexity?
- extension to more general ω-regular objectives and weighted games
- other definitions of penalties: discounted, mean-payoff...
- extensions in the time setting: already preliminary works[13] showing that the problem is difficult... but more tricky definition of penalties... what if we stick to simpler penalties, to extend the study to equilibria?

[13] Bouyer, Fang, Markey, 2015. Clement, Jéron, Markey, Mentré, 2020.