On Automata-Based Preferences in Graph Games

GT Verif 2026 – Rennes

Christophe Grandmont

September 11, 2026

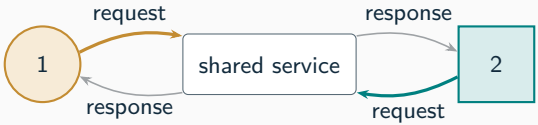




Joint work with Véronique Bruyère, Noémie Meunier, and Jean-François Raskin

Reactive systems

Continuous interactions between multiple **independent agents** with their own interests.

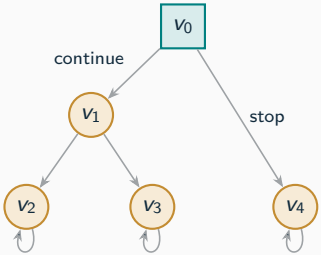


- Objectives are neither fully aligned nor entirely antagonistic.
- A service can be unavailable, available once, or available repeatedly.

$\rightsquigarrow$ Ensure a specification under **rational assumptions** on the agents.

Games played on graphs

$\rightsquigarrow$ Model interactions with **games played on graphs**.



A finite directed graph (V, E).

Each vertex belongs to one player, who chooses the next move.

Here: two players, $\circ$ and $\square$.

$$\pi = v_0 v_1 v_2 v_2 \cdots = v_0 v_1 v_2^\omega$$

A **play** is an infinite path.

A **strategy** chooses a (legal) move after every history:

$$\sigma_i : V^* V_i \longrightarrow V.$$

From Boolean to three payoffs

An agent wants to visit a target set T.

Boolean objective: payoff 1 if T is visited, 0 otherwise

What if repeated visit is **better**?

Three levels of satisfaction

$$\mathrm{pay}(\pi) = \begin{cases} 0 & \text{if } T \text{ is never visited,} \\ 1 & \text{if } T \text{ is visited finitely often, but at least once,} \\ 2 & \text{if } T \text{ is visited infinitely often.} \end{cases} \qquad 0 < 1 < 2.$$

$\rightsquigarrow$ Each payoff still describes an **ω-regular set** of plays.

Regular-payoff games

For each player $i \in \mathcal{P}$:

- a **finite** set of payoffs P_i,
- a payoff function $\mathrm{pay}_i : V^\omega \to P_i$,
- a preference relation $\ltimes_i \subseteq P_i \times P_i$.

One automaton for each payoff class

Each class $X_p^i = \mathrm{pay}_i^{-1}(p)$ is given by a **deterministic parity automaton** $\mathcal{A}_p^i$.

These sets **partition** V^ω: every play has exactly one payoff for player i.

$$i \text{ prefers } \rho \text{ to } \pi \iff \underbrace{\mathrm{pay}_i(\pi) \ltimes_i \mathrm{pay}_i(\rho)}_{\text{compare their payoffs}}.$$

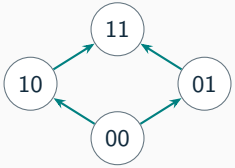
$\rightsquigarrow$ **Regular-payoff games** (RP games).

More general preferences

Payoffs need not be numbers, and preferences need not form an order.

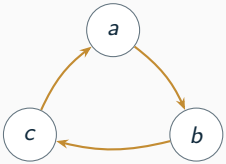
Two Boolean objectives

$$P_i = \{0,1\}^2$$



Subset order: 10 and 01 are incomparable.

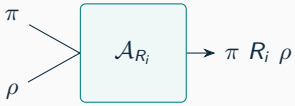
Or even cyclic preferences



No transitivity or totality assumption.

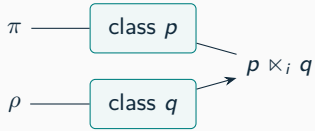
Link with ω-automatic preferences

Previous work



One automaton reads both plays **synchronously**.

Here



Classify each play **independently**, then compare.

ω-recognizable relations

$$R_i = \bigcup_{p \ltimes_i q} X_p^i \times X_q^i.$$

Exactly the ω-recognizable relations, a strict subclass of ω-automatic relations.

Bruyère, G., Raskin, MFCS 2025; Bruyère, Filiot, G., Raskin, 2026.

Nash equilibria

A strategy profile $\sigma = (\sigma_i)_{i \in \mathcal{P}}$ induces one outcome $\mathrm{Out}_{v_0}(\sigma)$.

Nash equilibrium (NE)

No player has an incentive to **unilaterally deviate**:

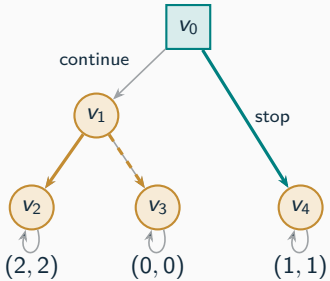
$$\forall i \in \mathcal{P},\ \forall \tau_i, \qquad \mathrm{pay}_i\big(\mathrm{Out}_{v_0}(\sigma)\big) \not< _i\ \mathrm{pay}_i\big(\mathrm{Out}_{v_0}(\tau_i, \sigma_{-i})\big).$$

- Player i may change their **whole strategy**.
- All other players keep their strategies, including their reactions after a deviation.

$\rightsquigarrow$ Are those reactions themselves rational?

A Nash equilibrium, but...

Payoffs $(p_\circ, p_\square)$, with $0 < 1 < 2$ for both players.



What happens if v_1 is reached?

The profile tells $\circ$ to choose v_3 (payoff 0).

But choosing v_2 gives $\circ$ payoff 2!

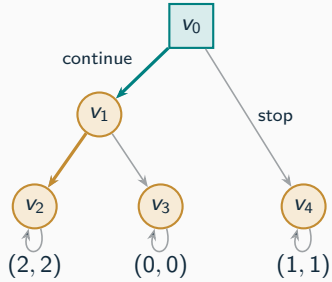
$\rightsquigarrow$ A **non-credible threat**.

The threat occurs outside the NE outcome, so the NE condition does not rule it out.

Subgame perfect equilibria

Subgame perfect equilibrium (SPE)

A strategy profile that is an **NE after every history**, including histories inconsistent with the profile.



At v_1, ○ chooses v_2.

At v_0, □ now prefers to continue: $2 > 1$.

$\rightsquigarrow$ This profile is an **SPE**, with payoff $(2, 2)$.

After a history hv, past events still count:

$$\text{pay}_{i \restriction h}(\rho) = \text{pay}_i(h\rho), \qquad \text{for a continuation } \rho \text{ from } v.$$

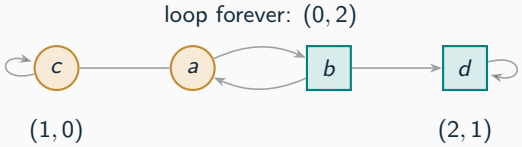
On existence of NEs and SPEs

Existence of NEs

If every $\ltimes_i$ is a strict partial order, there **exists an NE**.

Bruyère, G., Raskin, MFCS 2025.

Not true for SPE!



- ○: loop forever $<$ exit at c $<$ exit at d.
- □: exit at c $<$ exit at d $<$ loop forever.

There is an NE from a, but **no SPE**.

Example from Solan–Vieille (2003).

Constrained existence

Constrained NE / SPE existence problem

Given an initial vertex v_0 and payoffs $d_i \in P_i$, does there exist an NE / SPE with outcome π such that

$$\forall i \in \mathcal{P}, \qquad \mathrm{pay}_i(\pi) = d_i \quad \text{or} \quad d_i \ltimes_i \mathrm{pay}_i(\pi)?$$

With numerical payoffs: can every player meet a required level?

Theorem

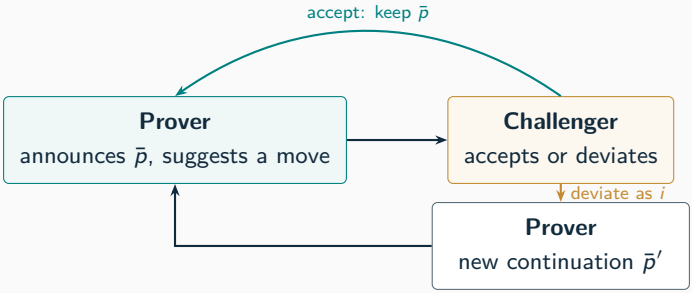
The constrained **NE** and **SPE** existence problems are EXPTIME-**complete** for RP games.

Hardness already holds for **two players**.

Input: one explicit deterministic parity automaton per payoff, and an explicit preference relation for each player.

Solve the SPE existence problem

A **two-player zero-sum game**, with payoff automata as observers.



Prover must justify the announced payoffs and prevent profitable unilateral deviations **after every history**.

SPE with payoff $\bar{p}$ $\iff$ Prover wins.

Construction from Noémie Meunier's thesis (2016) originally tailored for total orders, never published.

EXPTIME **algorithms**

Upper bound

Arena position + announced payoffs + automata states + deviation history $\rightsquigarrow$ **exponentially many states**.

Emerson–Lei objective: a Boolean combination of "(in)finitely often" conditions.

Polynomially many acceptance sets $\rightsquigarrow$ solving the game takes **exponential time**.

Matching lower bound

Reduce acceptance of a **linear-space alternating Turing machine**.

Players simulate a computation; finite automata check local errors when a player challenges a declaration.

For NEs: after the first deviation, keep testing that same deviator.

Conclusion

From **Boolean objectives** to **finite payoff classes**, with arbitrary preference relations.

From **NEs** to **SPEs**: rational behaviour after every history.

Constrained existence	NE	SPE
General setting	EXPTIME-c	EXPTIME-c
Fixed number of payoffs per player	PSPACE-c	EXPTIME-c

Thank you! Questions?