

Examining Games and a Way to Repair Them

Muhammad Najib

Department of Computer Science, University of Kaiserslautern

This is a summary of a series of works with Julian Gutierrez, Lewis Hammond, Anthony W. Lin, Giuseppe Perelli, and Mike Wooldridge.

The works published in IJCAI 2019, CONCUR 2019, AIJ, KR 2021.

1. Examining


- Correctness
- Tractability
- Cooperation and Probability

2. Repairing

Part I: Examining


How should we define correctness in MAS?

How Should We Define Multi-Agent Systems Correctness?



Classical notion of correctness ignores agents preferences

How Should We Define Multi-Agent Systems Correctness?



Correctness with respect to **rational choices** of agents

Not all behaviours are equal, but some are more unequal than others



- Autonomous cars crossing an intersection

Not all behaviours are equal, but some are more unequal than others



- Autonomous cars crossing an intersection
- Most of them (are expected to) cross without crashing with each other

Not all behaviours are equal, but some are more unequal than others



- Autonomous cars crossing an intersection
- Most of them (are expected to) cross without crashing with each other
- Cross and crash is also a *possible* behaviour of the system

Not all behaviours are equal, but some are more unequal than others



- Autonomous cars crossing an intersection
- Most of them (are expected to) cross without crashing with each other
- Cross and crash is also a *possible* behaviour of the system
- But cross and crash is not a *rational* behaviour

Not all behaviours are equal, but some are more unequal than others

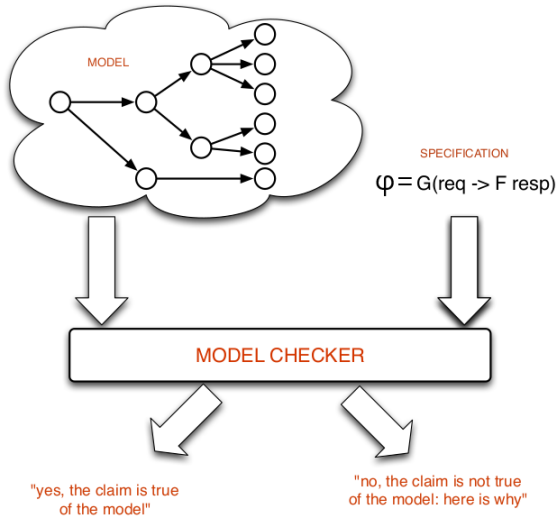


- Autonomous cars crossing an intersection
- Most of them (are expected to) cross without crashing with each other
- Cross and crash is also a *possible* behaviour of the system
- But cross and crash is not a *rational* behaviour
- They would rather do something else (not crash), thus it's not a stable behaviour

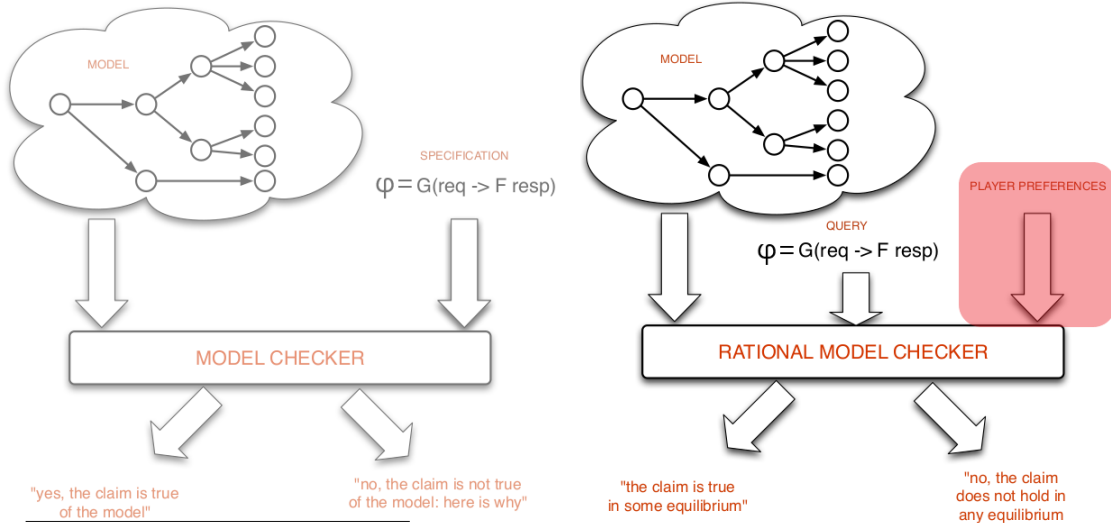
How do we define correctness in MAS?

- Is the system correct with respect to the set of **stable** behaviours?
- Stable behaviours in a group of *intelligent* agents \Rightarrow game theory
- Turn MAS into multi-player game

From Verification to Rational Verification



From Verification to Rational Verification¹



¹M. Wooldridge et al. "Rational Verification: From Model Checking to Equilibrium Checking". In: *AAAI*. 2016, pp. 4184–4191.

Rational Verification

- Game \mathcal{G} , each Player i is associated with a LTL goal γ_i
- Each player chooses a strategy; resolves non-determinism.
- A LTL property φ

LTL Game

A multi-player LTL game is a tuple $\mathcal{G}_{\text{LTL}} = (\mathcal{M}, \lambda, (\gamma_i)_{i \in N})$

- $\mathcal{M} = (N, (Ac_i)_{i \in N}, St, s_0, tr)$ is a concurrent game arena (CGA) ^a,
- γ_i is the LTL goal for player i .
- $\lambda : St \rightarrow 2^{\text{AP}}$ is a labelling function

^aAs usual: N agents; Ac_i actions of player i ; St states; s_0 initial state; tr transition function.

Useful Games

A (2-player) *parity* game is a tuple $H = (V_0, V_1, E, \alpha)$

- zero-sum turn-based
- $V = V_0 \cup V_1$
- $E \subseteq V \times V$
- $\alpha : V \rightarrow \mathbb{N}$ is a labelling priority function

Player 0 wins if the smallest priority that occurs infinitely often in the infinite play is even. Otherwise, player 1 wins. Can be solved in $\text{NP} \cap \text{coNP}^a$.

^aMarcin Jurdziński. "Deciding the winner in parity games is in $\text{UP} \cap \text{co-UP}$ ". In: *Information Processing Letters* (1998).

Useful Games

A multi-player *parity* game is a tuple $\mathcal{G}_{\text{PAR}} = (\mathcal{M}, (\alpha_i)_{i \in \mathbb{N}})$

- $\mathcal{M} = (\mathbb{N}, (\text{Ac}_i)_{i \in \mathbb{N}}, \text{St}, s_0, \text{tr})$ is a concurrent game arena (CGA) ²,
- $\alpha_i : \text{St} \rightarrow \mathbb{N}$ is the goal of player i , given as a priority function over St .

²As usual: \mathbb{N} agents; Ac_i actions of player i ; St states; s_0 initial state; tr transition function.

Strategies and Plays

Strategy

Finite state machine $\sigma_i = \langle S_i, s_i^0, \delta_i, \tau_i \rangle$

- S_i , internal state (s_i^0 initial state);
- $\delta_i : S_i \times Ac \rightarrow S_i$ internal transition function;
- $\tau_i : S_i \rightarrow Ac_i$ action function.

A strategy is a **recipe** for the agent prescribing the action to take at every time-step of the game execution. A **strategy profile** $\vec{\sigma} = \langle \sigma_1, \dots, \sigma_N \rangle$ assigns a strategy to each agent in the arena.

Play

Given a strategy assigned to every agent in A , denoted $\vec{\sigma}$, there is a unique possible execution $\pi(\vec{\sigma})$ called **play**.

Note that plays can only be **ultimately periodic**.

Nash Equilibria

Payoff Function

Let w_i be γ_i if \mathcal{G} is an LTL game, and be α_i if \mathcal{G} is a Parity game. For a strategy profiles $\vec{\sigma}$ in \mathcal{G} , we have

$$\text{pay}_i(\pi(\vec{\sigma})) = \begin{cases} 1, & \text{if } \pi(\vec{\sigma}) \models w_i \\ 0, & \text{otherwise} \end{cases}$$

Nash Equilibria

Payoff Function

Let w_i be γ_i if \mathcal{G} is an LTL game, and be α_i if \mathcal{G} is a Parity game. For a strategy profiles $\vec{\sigma}$ in \mathcal{G} , we have

$$\text{pay}_i(\pi(\vec{\sigma})) = \begin{cases} 1, & \text{if } \pi(\vec{\sigma}) \models w_i \\ 0, & \text{otherwise} \end{cases}$$

Nash Equilibrium

For a game \mathcal{G} , a strategy profile $\vec{\sigma}$ is a *Nash equilibrium* of \mathcal{G} if, for every player i and strategy $\sigma'_i \in \text{Str}_i$, we have

$$\text{pay}_i(\pi(\vec{\sigma})) \geq \text{pay}_i(\pi((\vec{\sigma}_{-i}, \sigma'_i))) .$$

i.e., **no player can benefit by changing its strategy unilaterally.**

Rational Verification: Decision Problems

E-Nash

Given: Game \mathcal{G} , temporal property φ .

Quest: Is there any Nash Equilibrium $\vec{\sigma}$ in \mathcal{G} such that $\pi(\vec{\sigma}) \models \varphi$?

Rational Verification: Decision Problems

E-Nash

Given: Game \mathcal{G} , temporal property φ .

Quest: Is there any Nash Equilibrium $\vec{\sigma}$ in \mathcal{G} such that $\pi(\vec{\sigma}) \models \varphi$?

A-Nash

Given: Game \mathcal{G} , temporal property φ .

Quest: Does $\pi(\vec{\sigma}) \models \varphi$ hold for **every** Nash Equilibrium $\vec{\sigma}$ in \mathcal{G} ?

Rational Verification: Decision Problems

E-Nash

Given: Game \mathcal{G} , temporal property φ .

Quest: Is there any Nash Equilibrium $\vec{\sigma}$ in \mathcal{G} such that $\pi(\vec{\sigma}) \models \varphi$?

A-Nash

Given: Game \mathcal{G} , temporal property φ .

Quest: Does $\pi(\vec{\sigma}) \models \varphi$ hold for **every** Nash Equilibrium $\vec{\sigma}$ in \mathcal{G} ?

Both decision problems above can be reduced to the following

Non-Emptiness

Given: Game \mathcal{G} .

Quest: Is there any Nash Equilibrium in \mathcal{G} ?

NE Characterisation

Theorem (NE characterisation)

Let $NE(\mathcal{G})$ be the set of Nash equilibria in \mathcal{G} . A strategy profile $\vec{\sigma} \in NE(\mathcal{G})$

if and only if

the path $\pi = \pi(\vec{\sigma})$ is such that, for every $k \in \mathbb{N}$, the pair (s_k, \vec{a}^k) of the k -th position of π is **punishing secure**³ for every $j \in Lose(\pi)$.⁴ Where $\vec{a}^k = \langle a_1, \dots, a_n \rangle$ is an action profile at k .

Along π , no player j can unilaterally get its goal γ_j achieved.

³**Punishing secure**: agent j does not have a strategy σ'_j that wins against $\vec{\sigma}_{-j}$, i.e. $\pi(\vec{\sigma}_{-j}, \sigma'_j) \models \gamma_j$.

⁴Here $Lose(\pi) = \{j \in \mathbb{N} : \pi \not\models \gamma_j\}$ are the agents that are **not satisfied** over π .

NE Characterisation via Local Reasoning

- Memory is needed to satisfy LTL goal
- Memory is NOT necessary for (2-player) parity games (memoryless/positional determinacy)
- Reason locally by converting each γ_i into deterministic parity word automaton (DPW)
 $\mathcal{A}_i = \langle 2^{AP}, Q, q^0, \rho, \alpha \rangle$.
- Then build $\mathcal{G}_{\text{LTL}} = (\mathcal{M}, \lambda, (\gamma_i)_{i \in \mathbb{N}})$ into $\mathcal{G}_{\text{PAR}} = (\mathcal{M}', (\alpha'_i)_{i \in \mathbb{N}})$, where $\mathcal{M}' = (\mathbb{N}, (A_{C_i})_{i \in \mathbb{N}}, \text{St}', s'_0, \text{tr}')$ and $(\alpha'_i)_{i \in \mathbb{N}}$:
 - $\text{St}' = \text{St} \times \prod_{i \in \mathbb{N}} Q_i$ and $s'_0 = (s_0, q_1^0, \dots, q_n^0)$;
 - for each state $(s, q_1, \dots, q_n) \in \text{St}'$ and action profile \vec{a} ,
 $\text{tr}'((s, q_1, \dots, q_n), \vec{a}) = (\text{tr}(s, \vec{a}), \rho_1(q_1, \lambda(s)), \dots, \rho_n(q_n, \lambda(s)))$;
 - $\alpha'_i(s, q_1, \dots, q_n) = \alpha_i(q_i)$.

Lemma (Goal Invariance)


Let \mathcal{G}_{LTL} be an LTL game and \mathcal{G}_{PAR} its associated Parity game. Then, for every strategy profile $\vec{\sigma}$ and player i , it is the case that $\pi(\vec{\sigma}) \models \gamma_i$ in \mathcal{G}_{LTL} if and only if $\pi(\vec{\sigma}) \models \alpha_i$ in \mathcal{G}_{PAR} .

Theorem (NE Invariance)

Let \mathcal{G}_{LTL} be an LTL game and \mathcal{G}_{PAR} its associated Parity game. Then, $NE(\mathcal{G}_{\text{LTL}}) = NE(\mathcal{G}_{\text{PAR}})$.

⁵Julian Gutierrez et al. “Automated temporal equilibrium analysis: Verification and synthesis of multi-player games”. In: *Artificial Intelligence* (2020).

Visualising NE Characterisation



$\bigcap_{j \in \text{Lose}} \text{Pun}_j(\mathcal{G}_{\text{PAR}})$ is the **punishing region** for Lose

Computing Punishing Region

For a \mathcal{G}_{PAR} and a (to-be-punished) player j . We turn \mathcal{G}_{PAR} into a 2-player zero-sum parity game $H_j = (V_0, V_1, E, \alpha)$ between player j (Player 1) and (coalition) player N_{-j} (Player 0). Circular states are in V_0 .



punishing region for $\text{Lose} = \bigcap_{j \in \text{Lose}} \text{Pun}_j(\mathcal{G}_{\text{PAR}})$

Computing Punishing Region

For a \mathcal{G}_{PAR} and a (to-be-punished) player j . We turn \mathcal{G}_{PAR} into a 2-player zero-sum parity game $H_j = (V_0, V_1, E, \alpha)$ between player j (Player 1) and (coalition) player N_{-j} (Player 0). Circular states are in V_0 .

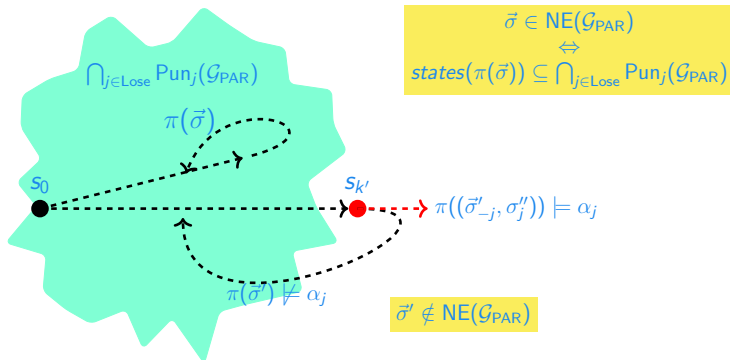


punishing region for $\text{Lose} = \bigcap_{j \in \text{Lose}} \text{Pun}_j(\mathcal{G}_{\text{PAR}})$

Corollary

Computing $\text{Pun}_i(\mathcal{G}_{\text{PAR}})$ can be done in polynomial time with respect to the size of the underlying graph of the game \mathcal{G}_{PAR} and exponential in the size of the priority function α_i , that is, to the size of the range of α_i . Moreover, there is a memoryless strategy $\vec{\sigma}_i$ that is a punishment against player i in every state $s \in \text{Pun}_i(\mathcal{G}_{\text{PAR}})$.

Finding NE Run



How do we compute $\pi(\vec{\sigma})$? Is there such run $\pi(\vec{\sigma})$ inside the punishing region?

Finding NE Run

- $\pi(\vec{\sigma})$ must be accepting for each $\alpha_i, i \in \text{Win} = N \setminus \text{Lose}$.
- Solve emptiness problem of DPWs intersection $\bigtimes_{i \in \text{Win}} \mathcal{A}^i$
- Intersection of DPWs might involve exponential blowup
- Each parity condition $\alpha = (F_1, \dots, F_n)$ is a Streett condition $((E_1, C_1), \dots, (E_m, C_m))$ with $m = \lceil \frac{n}{2} \rceil$ and $(E_i, C_i) = (F_{2i+1}, \bigcup_{j \leq i} F_{2j})$, for each $0 \leq i \leq m$
- Intersection of (Deterministic Streett Word Automata) DSWs $\bigtimes_{i \in \text{Win}} \mathcal{S}_i$ and nonemptiness check can be done in polynomial time

The Procedure

1. $\mathcal{G}_{\text{LTL}} \Rightarrow \mathcal{G}_{\text{PAR}}$
 2. For each $\text{Win} \subseteq N$ do:
 - 2.1 Compute punishing region
 $\bigcap_{j \in \text{Lose}} \text{Pun}_j(\mathcal{G}_{\text{PAR}})$
 - 2.2 Construct DSW $\times_{i \in \text{Win}} \mathcal{S}_i$
 - 2.3 If $\mathcal{L}(\times_{i \in \text{Win}} \mathcal{S}_i) \neq \emptyset$ then return “YES”
 3. Return “NO”
- Step 1 can be done in 2EXPTIME: the number of states is doubly exponential in the size of LTL goals, but priority functions $(\alpha_i)_{i \in N}$ is only singly exponential.
 - Step 2 at most executed exponential in the number of players
 - Step 2.1 is polynomial in the number of states and exponential in the number of priorities
 - Step 2.2 and 2.3 are both polynomial in the number of states
 - Overall we have 2EXPTIME procedure.

The Procedure

1. $\mathcal{G}_{\text{LTL}} \Rightarrow \mathcal{G}_{\text{PAR}}$
 2. For each $\text{Win} \subseteq N$ do:
 - 2.1 Compute punishing region
 $\bigcap_{j \in \text{Lose}} \text{Pun}_j(\mathcal{G}_{\text{PAR}})$
 - 2.2 Construct DSW $\times_{i \in \text{Win}} \mathcal{S}_i$
 - 2.3 If $\mathcal{L}(\times_{i \in \text{Win}} \mathcal{S}_i) \neq \emptyset$ then return “YES”
 3. Return “NO”
- Step 1 can be done in 2EXPTIME: the number of states is doubly exponential in the size of LTL goals, but priority functions $(\alpha_i)_{i \in N}$ is only singly exponential.
 - Step 2 at most executed exponential in the number of players
 - Step 2.1 is polynomial in the number of states and exponential in the number of priorities
 - Step 2.2 and 2.3 are both polynomial in the number of states
 - Overall we have 2EXPTIME procedure.

The Procedure

1. $\mathcal{G}_{\text{LTL}} \Rightarrow \mathcal{G}_{\text{PAR}}$
 2. For each $\text{Win} \subseteq N$ do:
 - 2.1 Compute punishing region
 $\bigcap_{j \in \text{Lose}} \text{Pun}_j(\mathcal{G}_{\text{PAR}})$
 - 2.2 Construct DSW $\times_{i \in \text{Win}} S_i$
 - 2.3 If $\mathcal{L}(\times_{i \in \text{Win}} S_i) \neq \emptyset$ then return "YES"
 3. Return "NO"
- Step 1 can be done in 2EXPTIME: the number of states is doubly exponential in the size of LTL goals, but priority functions $(\alpha_i)_{i \in N}$ is only singly exponential.
 - Step 2 at most executed exponential in the number of players
 - Step 2.1 is polynomial in the number of states and exponential in the number of priorities
 - Step 2.2 and 2.3 are both polynomial in the number of states
 - Overall we have 2EXPTIME procedure.

The Procedure

1. $\mathcal{G}_{\text{LTL}} \Rightarrow \mathcal{G}_{\text{PAR}}$
 2. For each $\text{Win} \subseteq \mathbb{N}$ do:
 - 2.1 Compute punishing region
 $\bigcap_{j \in \text{Lose}} \text{Pun}_j(\mathcal{G}_{\text{PAR}})$
 - 2.2 Construct DSW $\bigtimes_{i \in \text{Win}} S_i$
 - 2.3 If $\mathcal{L}(\bigtimes_{i \in \text{Win}} S_i) \neq \emptyset$ then return "YES"
 3. Return "NO"
- Step 1 can be done in 2EXPTIME: the number of states is doubly exponential in the size of LTL goals, but priority functions $(\alpha_i)_{i \in \mathbb{N}}$ is only singly exponential.
 - Step 2 at most executed exponential in the number of players
 - Step 2.1 is polynomial in the number of states and exponential in the number of priorities
 - Step 2.2 and 2.3 are both polynomial in the number of states
 - Overall we have 2EXPTIME procedure.

The Procedure

1. $\mathcal{G}_{\text{LTL}} \Rightarrow \mathcal{G}_{\text{PAR}}$
 2. For each $\text{Win} \subseteq N$ do:
 - 2.1 Compute punishing region
 $\bigcap_{j \in \text{Lose}} \text{Pun}_j(\mathcal{G}_{\text{PAR}})$
 - 2.2 Construct DSW $\times_{i \in \text{Win}} S_i$
 - 2.3 If $\mathcal{L}(\times_{i \in \text{Win}} S_i) \neq \emptyset$ then return "YES"
 3. Return "NO"
- Step 1 can be done in 2EXPTIME: the number of states is doubly exponential in the size of LTL goals, but priority functions $(\alpha_i)_{i \in N}$ is only singly exponential.
 - Step 2 at most executed exponential in the number of players
 - Step 2.1 is polynomial in the number of states and exponential in the number of priorities
 - Step 2.2 and 2.3 are both polynomial in the number of states
 - Overall we have 2EXPTIME procedure.

The Procedure

1. $\mathcal{G}_{\text{LTL}} \Rightarrow \mathcal{G}_{\text{PAR}}$
 2. For each $\text{Win} \subseteq \mathbb{N}$ do:
 - 2.1 Compute punishing region
 $\bigcap_{j \in \text{Lose}} \text{Pun}_j(\mathcal{G}_{\text{PAR}})$
 - 2.2 Construct DSW $\times_{i \in \text{Win}} \mathcal{S}_i$
 - 2.3 If $\mathcal{L}(\times_{i \in \text{Win}} \mathcal{S}_i) \neq \emptyset$ then return "YES"
 3. Return "NO"
- Step 1 can be done in 2EXPTIME: the number of states is doubly exponential in the size of LTL goals, but priority functions $(\alpha_i)_{i \in \mathbb{N}}$ is only singly exponential.
 - Step 2 at most executed exponential in the number of players
 - Step 2.1 is polynomial in the number of states and exponential in the number of priorities
 - Step 2.2 and 2.3 are both polynomial in the number of states
 - Overall we have 2EXPTIME procedure.

Rational Verification: Complexity

Theorem (Complexity)

For the case of both the specification φ and the agents goals γ_i expressed as LTL formulas, E-Nash and A-Nash are 2EXPTIME-Complete.⁶

⁶Wooldridge et al., “Rational Verification: From Model Checking to Equilibrium Checking”; Julian Gutierrez, Paul Harrenstein, and Michael J. Wooldridge. “From model checking to equilibrium checking: Reactive modules for rational verification”. In: *Artificial Intelligence* 248 (2017), pp. 123–157.

EVE (Equilibrium Verification Environment)

- Simple Reactive Modules Language (SRML)⁷ as modelling language
- Supports general-sum multi-player LTL games, bisimulation-invariant strategies, and perfect recall.
- Supports Non-emptiness, E-Nash, and A-Nash
- Synthesise strategies
- **Open-source:** <https://github.com/eve-mas/eve-parity>
- **EVE Online:** <http://eve.cs.ox.ac.uk/>

⁷Based on the Reactive Modules language used by PRISM and MOCHA.

Chapter 2: Tractability⁸

- 2EXPTIME is rather slow
- What can we do to improve?
- Use different goals and properties: GR(1) and mean-payoff value

⁸Julian Gutierrez et al. “On Computational Tractability for Rational Verification”. In: *IJCAI*. 2019, pp. 329–335.

The language of *General Reactivity of rank 1*, denoted GR(1), is the fragment of LTL of formulae written in the following form:

$$(\mathbf{GF}\psi_1 \wedge \dots \wedge \mathbf{GF}\psi_m) \rightarrow (\mathbf{GF}\varphi_1 \wedge \dots \wedge \mathbf{GF}\varphi_n),$$

each ψ_i and φ_i is a Boolean combination of atomic propositions.


$$(\mathbf{GFreq}_1 \wedge \mathbf{GFreq}_2) \rightarrow \mathbf{GFack}$$

⁹Roderick Bloem et al. "Synthesis of Reactive(1) designs". In: *J. Comput. Syst. Sci.* 78.3 (2012), pp. 911–938.

Mean-payoff value

For an infinite sequence $\beta \in \mathbb{R}^\omega$ of real numbers, let $\text{mp}(\beta)$ be the *mean-payoff* value of β , defined as follows:

$$\text{mp}(\beta) = \liminf_{n \rightarrow \infty} \frac{1}{n} \sum_{i=0}^{n-1} \beta[i]$$



$$\beta^1 = 0000000000 \dots \quad \text{mp}(\beta^1) = 0$$

$$\beta^2 = 0101010101 \dots \quad \text{mp}(\beta^2) = 0.5$$

$$\beta^3 = 0102010201 \dots \quad \text{mp}(\beta^3) = 3/4$$

A multi-player GR(1) game is a tuple $\mathcal{G}_{\text{GR}(1)} = \langle \mathcal{M}, (\gamma_i)_{i \in N}, \lambda \rangle$

- $\mathcal{M} = \langle N, Ac, St, s_0, tr \rangle$ is an arena,
- γ_i is the GR(1) goal for player i .

A multi-player mp game is a tuple $\mathcal{G}_{\text{mp}} = \langle \mathcal{M}, (w_i)_{i \in N}, \lambda \rangle$,

- $\mathcal{M} = \langle N, Ac, St, s_0, tr \rangle$ is an arena
- $w_i : St \rightarrow \mathbb{Z}$ maps states to integer numbers, for each player i

Cases

E-Nash


Given: Game \mathcal{G} , temporal property φ .

Quest: Is there any Nash Equilibrium $\vec{\sigma}$ in \mathcal{G} such that $\pi(\vec{\sigma}) \models \varphi$?

	γ_i	φ	E-NASH
	LTL	LTL	2EXPTIME-complete
GR(1) games {	GR(1)	LTL	?
	GR(1)	GR(1)	?
mp games {	mp	LTL	?
	mp	GR(1)	?

E-Nash in GR(1) games: NE characterisation

Theorem (NE characterisation)



E-Nash in GR(1) games: Computing punishment regions

Theorem (Computing $\text{Pun}_j(\mathcal{G})$)

For a given GR(1) game \mathcal{G} , computing $\text{Pun}_j(\mathcal{G})$ of player j can be done in polynomial time with respect to the size of both \mathcal{G} and γ_j .

E-Nash in GR(1) games: the procedure

1. Guess a set $\text{Win} \subseteq N$ of winners;
2. For each player $j \in \text{Lose} = N \setminus \text{Win}$, a loser in the game, compute its punishment region $\text{Pun}_j(\mathcal{G})$;
3. Find desired path $\pi(\vec{\sigma})$ consisting of states in $\bigcap_{j \in \text{Lose}} \text{Pun}_j(\mathcal{G})$. Any deviation by player j must remain inside $\text{Pun}_j(\mathcal{G})$, that is, a path $\pi(\vec{\sigma})$ satisfying the following three conditions:
 - $\text{states}(\pi(\vec{\sigma})) \subseteq \bigcap_{j \in \text{Lose}} \text{Pun}_j(\mathcal{G})$
 - $\text{states}(\pi(\vec{\sigma}_{-j}, \sigma'_j)) \subseteq \text{Pun}_j(\mathcal{G})$, for every $j \in \text{Lose}$ and σ'_j of j
 - $\pi(\vec{\sigma}) \models \varphi \wedge \bigwedge_{i \in \text{Win}} \gamma_i$

E-Nash in GR(1) games: the procedure

1. Guess a set $\text{Win} \subseteq N$ of winners;
2. For each player $j \in \text{Lose} = N \setminus \text{Win}$, a loser in the game, compute its punishment region $\text{Pun}_j(\mathcal{G})$;
3. Find desired path $\pi(\vec{\sigma})$ consisting of states in $\bigcap_{j \in \text{Lose}} \text{Pun}_j(\mathcal{G})$. Any deviation by player j must remain inside $\text{Pun}_j(\mathcal{G})$, that is, a path $\pi(\vec{\sigma})$ satisfying the following three conditions:
 - $\text{states}(\pi(\vec{\sigma})) \subseteq \bigcap_{j \in \text{Lose}} \text{Pun}_j(\mathcal{G})$
 - $\text{states}(\pi(\vec{\sigma}_{-j}, \sigma'_j)) \subseteq \text{Pun}_j(\mathcal{G})$, for every $j \in \text{Lose}$ and σ'_j of j
 - $\pi(\vec{\sigma}) \models \varphi \wedge \bigwedge_{i \in \text{Win}} \gamma_i$

Complexities for GR(1) and LTL specifications:

- If φ is a GR(1) specification: FPT
- If φ is an LTL specification: PSPACE

E-Nash in mp games: NE characterisation

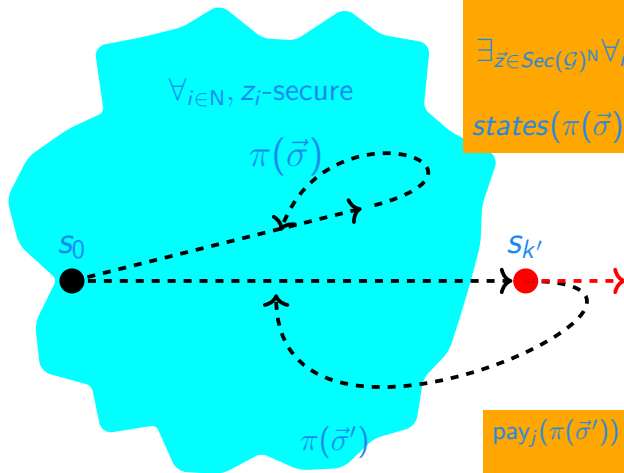
Theorem (NE characterisation)

For every mp game \mathcal{G} and ultimately periodic path $\pi = (s_0, \vec{a}^0), (s_1, \vec{a}^1), \dots$, the following are equivalent

1. There is $\vec{\sigma} \in NE(\mathcal{G})$ such that $\pi = \pi(\vec{\sigma})$;
2. There exists $\vec{z} \in \mathbb{R}^N$, where $z_i \in \{pun_i(s) : s \in St\}$ such that, for every $i \in N$
 - 2.1 $z_i \leq pay_i(\pi)$, and
 - 2.2 for all $k \in \mathbb{N}$, the pair (s_k, \vec{a}^k) is z_i -secure for i .

Along π , no player i can unilaterally get a payoff greater than z_i .

E-Nash in mp games: NE characterisation



$$\vec{\sigma} \in \text{NE}(\mathcal{G})$$

$$\Leftrightarrow$$

$$\exists z \in \text{Sec}(\mathcal{G}) \forall i \in N, \text{pay}_i(\pi(\vec{\sigma})) \geq z_i$$

$$\Leftrightarrow$$

$$\text{states}(\pi(\vec{\sigma})) \subseteq \bigcap_i \text{Pun}_i(\mathcal{G}, \leq z_i)$$

$$\text{pay}_j(\pi(\vec{\sigma}')) < \text{pay}_j(\pi((\vec{\sigma}'_{-j}, \sigma''_j)))$$

$$\Leftrightarrow$$

$$\vec{\sigma}' \notin \text{NE}(\mathcal{G})$$

E-Nash in mp games: the procedure

1. Guess a vector $\vec{z} \in \mathbb{R}^N$ of values, each being a punishment value for a player i
2. For each i , compute its z_i -punishment region $\text{Pun}_i(\mathcal{G}, \leq z_i)$;
3. Find (u.p.) path $\pi(\vec{\sigma})$ consisting of states in $\bigcap_i \text{Pun}_i(\mathcal{G}, \leq z_i)$. Any deviation by player i must remain inside $\text{Pun}_i(\mathcal{G}, \leq z_i)$, that is, an ultimately periodic path $\pi(\vec{\sigma})$ satisfying that:
 - $\text{states}(\pi(\vec{\sigma})) \subseteq \bigcap_i \text{Pun}_i(\mathcal{G}, \leq z_i)$
 - $\text{states}(\pi(\vec{\sigma}_{-i}, \sigma'_i)) \subseteq \text{Pun}_i(\mathcal{G}, \leq z_i)$, for every i and σ'_i of i
 - $\pi(\vec{\sigma}) \models \varphi$ and $\forall_i, \text{pay}_i(\pi(\vec{\sigma})) \geq z_i$

E-Nash in mp games: the procedure

1. Guess a vector $\vec{z} \in \mathbb{R}^N$ of values, each being a punishment value for a player i
2. For each i , compute its z_i -punishment region $\text{Pun}_i(\mathcal{G}, \leq z_i)$;
3. Find (u.p.) path $\pi(\vec{\sigma})$ consisting of states in $\bigcap_i \text{Pun}_i(\mathcal{G}, \leq z_i)$. Any deviation by player i must remain inside $\text{Pun}_i(\mathcal{G}, \leq z_i)$, that is, an ultimately periodic path $\pi(\vec{\sigma})$ satisfying that:
 - $\text{states}(\pi(\vec{\sigma})) \subseteq \bigcap_i \text{Pun}_i(\mathcal{G}, \leq z_i)$
 - $\text{states}(\pi(\vec{\sigma}_{-i}, \sigma'_i)) \subseteq \text{Pun}_i(\mathcal{G}, \leq z_i)$, for every i and σ'_i of i
 - $\pi(\vec{\sigma}) \models \varphi$ and $\forall_i, \text{pay}_i(\pi(\vec{\sigma})) \geq z_i$

Complexities for GR(1) and LTL specifications:

- If φ is a GR(1) specification: NP-complete
- If φ is an LTL specification: PSPACE-complete

Complexity Results

γ_i	φ	E-NASH
LTL	LTL	2EXPTIME-complete
GR(1)	LTL	PSPACE-complete
GR(1)	GR(1)	FPT
mp	LTL	PSPACE-complete
mp	GR(1)	NP-complete

- NON-EMPTINESS (E-NASH when $\varphi = \top$):
 - LTL games: 2EXPTIME-complete
 - GR(1) games: FPT
 - mp games: NP-complete
- A-NASH: 2EXPTIME, PSPACE, FPT, PSPACE, coNP.

Chapter 3: Cooperation and Probability¹⁰

- Players can make a binding agreements and form coalitions
- Coalitions can collectively achieve goals
- Cooperative games
- Solution concept: **Core**

¹⁰Julian Gutierrez et al. "Rational Verification for Probabilistic Systems". In: *KR*. to appear. 2021.

Rational Verification in Cooperative Games

- Game \mathcal{G} , each Player i is associated with a LTL goal γ_i
- Each player chooses a strategy; resolves non-determinism.
- A LTL property φ

Rational Verification in Cooperative Games

- Game \mathcal{G} , each Player i is associated with a LTL goal γ_i
- Each player chooses a strategy; resolves non-determinism.
- A LTL property φ

E-Core

Is there any core $\vec{\sigma}$ in \mathcal{G} such that $\pi(\vec{\sigma}) \models \varphi$?

Rational Verification in Cooperative Games

- Game \mathcal{G} , each Player i is associated with a LTL goal γ_i
- Each player chooses a strategy; resolves non-determinism.
- A LTL property φ

E-Core

Is there any core $\vec{\sigma}$ in \mathcal{G} such that $\pi(\vec{\sigma}) \models \varphi$?

A-Core

Does $\pi(\vec{\sigma}) \models \varphi$ hold for every core $\vec{\sigma}$ in \mathcal{G} ?

Rational Verification

- Game \mathcal{G} , each Player i is associated with a LTL goal γ_i
- A strategy profile $\vec{\sigma}$

Core-Membership

Is $\vec{\sigma}$ a core in the game \mathcal{G} ?

Theorem (Complexity)

For the case of both the specification φ and the agents goals γ_i expressed as LTL formulas, E-Core, A-Core, and Core-Membership are 2EXPTIME-Complete.¹¹

¹¹Julian Gutierrez, Sarit Kraus, and Michael Wooldridge. "Cooperative Concurrent Games". In: AAMAS. 2019.

Complexity Map

	Non-Cooperative	Cooperative
E-(Nash/Core)	2EXPTIME-Complete	2EXPTIME-Complete
A-(Nash/Core)	2EXPTIME-Complete	2EXPTIME-Complete
(NE/Core)-Membership	PSPACE-Complete	2EXPTIME-Complete

Complexity Map

	Non-Cooperative	Cooperative
E-(Nash/Core)	2EXPTIME-Complete	2EXPTIME-Complete
A-(Nash/Core)	2EXPTIME-Complete	2EXPTIME-Complete
(NE/Core)-Membership	PSPACE-Complete	2EXPTIME-Complete

Without probabilistic behaviours...

A Case for Probabilistic Systems

- Real life scenarios often involve probabilities
- Probabilities grant us power, e.g., the dining philosopher problem

This Work

- Rational verification for probabilistic systems
- Cooperative and non-cooperative games
- Goals and specifications are LTL formulae
- Concurrent actions, infinite horizon, infinite memory
- *Qualitative setting*: almost-surely satisfaction

Concurrent Stochastic Games (CSGs)

Definition (CSG Arena)

A *concurrent stochastic game arena* (CSGA) is a tuple $\mathcal{M} = (\mathcal{N}, \text{St}, s^0, (A_{c_i})_{i \in \mathcal{N}}, \text{tr})$, where

- $\text{tr} : \text{St} \times \vec{A}_{c_i} \rightarrow \mathcal{D}(\text{St})$ is probabilistic transition function

Definition (CSG)

A *concurrent stochastic game* (CSG) is a tuple $\mathcal{G} = (\mathcal{M}, (\gamma_i)_{i \in \mathcal{N}}, \lambda)$, where \mathcal{M} is a CSGA, γ_i is a LTL formula that represents the *goal* of player i , and $\lambda : \text{St} \rightarrow 2^{\text{AP}}$ a labelling function.

Concurrent Stochastic Parity Games (CSPGs)

Definition (CSG Arena)

A *concurrent stochastic game arena* (CSGA) is a tuple $\mathcal{M} = (\mathbb{N}, \text{St}, s^0, (\text{Ac}_i)_{i \in \mathbb{N}}, \text{tr})$, where

- $\text{tr} : \text{St} \times \vec{\text{Ac}} \rightarrow \text{D}(\text{St})$ is probabilistic transition function

Definition (CSPG)

A *concurrent stochastic parity game* (CSPG) is a tuple $\mathcal{G}_{\text{PAR}} = (\mathcal{M}, (\alpha_i)_{i \in \mathbb{N}})$, where $\alpha_i : \text{St} \rightarrow \mathbb{N}$ is the goal of player i , given as a priority function over the set of states St . A path π satisfies a priority function α , denoted by $\pi \models \alpha$, if the minimum number occurring infinitely often in the infinite sequence $\alpha(\pi_0)\alpha(\pi_1)\alpha(\pi_2)\dots$ is even.

Definition (Strategy)

A strategy for player i can be understood (abstractly) as a function $\sigma_i : St^+ \rightarrow D(Ac_i)$ that assigns to every non-empty finite sequence of states a probability distribution over player i 's set of actions.

Definition (Strategy as Transducer)

a strategy in \mathcal{G} for player i is a transducer $\sigma_i = (Q_i, q_i^0, \delta_i, \tau_i)$

A strategy is

- *memoryless* if there exists a transducer encoding the strategy with $|Q_i| = 1$
- *finite-memory* if $|Q_i| < \infty$
- *deterministic* if $\tau_i : Q_i \times \text{St} \rightarrow \text{Ac}_i$, such that for every $q_i \in Q_i$ and every $s \in \text{St}$, we have that $\tau_i(q_i, s) \in \text{Ac}_i(s)$

Satisfaction Conditions

LTL goals:

- For a given game \mathcal{G} and a strategy profile $\vec{\sigma}$, a formula φ is said to be *almost-surely* satisfied, denoted $\vec{\sigma} \models AS(\varphi)$, iff, $\Pr_{\mathcal{C}_{\vec{\sigma}}}(\{\pi \in \text{Paths}(\mathcal{C}_{\vec{\sigma}}, s^0) : \pi \models \varphi\}) = 1$.
- φ is satisfied with *non-zero* probability, denoted $\vec{\sigma} \models NZ(\varphi)$ iff $\Pr_{\mathcal{C}_{\vec{\sigma}}}(\{\pi \in \text{Paths}(\mathcal{C}_{\vec{\sigma}}, s^0) : \pi \models \varphi\}) > 0$.
- $NZ(\varphi) \equiv \neg AS(\neg\varphi)$

Parity goals:

- $\vec{\sigma} \models AS(\alpha)$ if and only if $\Pr_{\mathcal{C}_{\vec{\sigma}}}(\{\pi \in \text{Paths}(\mathcal{C}_{\vec{\sigma}}, s^0) : \pi \models \alpha\}) = 1$.
- $\vec{\sigma} \models NZ(\alpha)$ if and only if $\Pr_{\mathcal{C}_{\vec{\sigma}}}(\{\pi \in \text{Paths}(\mathcal{C}_{\vec{\sigma}}, s^0) : \pi \models \alpha\}) > 0$.

Definition (Deviation)

A deviation is a joint strategy $\vec{\sigma}_A$ for the coalition $A \subseteq N$, with $A \neq \emptyset$.

Definition (Beneficial Deviation)


For a strategy profile $\vec{\sigma}$, we say $\vec{\sigma}'_A$ is a beneficial deviation from $\vec{\sigma}$ if $A \subseteq \text{Lose}(\vec{\sigma})$ and for all $\vec{\sigma}'_{-A}$, we have $A \subseteq \text{Win}((\vec{\sigma}'_A, \vec{\sigma}'_{-A}))$.

Definition (Core)

The core of a game \mathcal{G} , denoted $\text{core}(\mathcal{G})$, is then defined to be the set of strategy profiles that admit no beneficial deviation.

Example

Consider a game with two players $N = \{1, 2\}$ and two variables $AP = \{p, q\}$, with player 1's action set being $Ac_1 = \{a, \bar{a}\}$ and player 2's being $Ac_2 = \{b, \bar{b}\}$. Let $\gamma_1 = \mathbf{F}p$ and $\gamma_2 = \mathbf{F}q$.



consider a strategy profile $\vec{\sigma}$ in which player 1/2 always chooses action a/b in s_0 (i.e., chooses a/b with probability 1)

consider a strategy profile $\vec{\sigma}'$ in which player 1/2 chooses action \bar{a}/\bar{b} with non-zero probability

Non-Coop: NE-Membership

NE-Membership

Given: Game \mathcal{G} , strategy profile $\vec{\sigma}$.

Question: Is $\vec{\sigma}$ a Nash equilibrium in the game \mathcal{G} ?

In general, infinite memory strategies are needed to play ω -regular games with almost-sure winning conditions. Here, we assume that $\vec{\sigma}$ can be represented by some finite state machine.

Non-Coop: NE-Membership

1. For $i \in N$:
 - 1.1 If $\pi(\vec{\sigma}) \not\models AS(\gamma_i)$ then
 - 1.1.1 If there is σ'_i s.t. $(\vec{\sigma}_{-i}, \sigma'_i) \models AS(\gamma_i)$ then return “NO”
 2. Return “YES”
- 1.1 amounts to qualitative model checking LTL formula γ_i over a Markov chain (PSPACE)
 - 1.1.1 amounts to qualitative model checking LTL formula γ_i over a MDP (2EXPTIME)
 - Lower-bound: reduce to qualitative model checking LTL formula γ_i over a MDP

Theorem

NE-Membership for probabilistic systems is 2EXPTIME-complete

Non-Coop: E/A-Nash

E-NASH

Given: Game \mathcal{G} , temporal property φ .

Quest: Is there any Nash equilibrium $\vec{\sigma}$ in \mathcal{G} such that $\pi(\vec{\sigma}) \models AS(\varphi)$?

- Use the similar construction and NE characterisation to deterministic games
- Use Qualitative Parity Logic (QPL) Realizability problem to find NE run inside punishing region
- Procedure is 2EXPTIME, lower-bound via LTL model checking over MDPs.

Theorem

E-Nash and A-Nash for probabilistic systems are 2EXPTIME-complete

Coop: E/A-Core

E-CORE

Given: Game \mathcal{G} , temporal property φ .

Quest: Is there any core $\vec{\sigma}$ in \mathcal{G} such that $\vec{\sigma} \models AS(\varphi)$?

- Turn the game into its corresponding parity game
- Check for each possible winning coalition $Win \subseteq N$ s.t. for each possible losing coalition $Lose \subseteq N \setminus Win$, there is no beneficial deviation
- Use QPL to solve some problems in the procedure
- Procedure is 2EXPTIME, lower-bound via LTL model checking over MDPs.

Theorem

E-Core and A-Core for probabilistic systems are 2EXPTIME-complete

Coop: Core-Membership

CORE-MEMBERSHIP

Given: Game \mathcal{G} , a strategy profile $\vec{\sigma}$.

Quest: Is $\vec{\sigma}$ a core in \mathcal{G} ?^a

^aAgain, we assume that $\vec{\sigma}$ can be represented by some FSM

- For each $\text{Lose} \subseteq N \setminus \text{Win}(\vec{\sigma})$, check if there is beneficial deviation
- This amounts to model checking LTL formula over a MDP (2EXPTIME)
- Lower-bound: reduce to qualitative model checking LTL formula γ_i over a MDP

Theorem

Core-Membership for probabilistic systems is 2EXPTIME-complete

Results

Deterministic	Non-Coop.	Coop.
E-(Nash/Core)	2EXPTIME-C	2EXPTIME-C
A-(Nash/Core)	2EXPTIME-C	2EXPTIME-C
(NE/Core)-Mbrshp	PSPACE-C	2EXPTIME-C

Table 1: Complexity results for deterministic systems.

Probabilistic	Non-Coop.	Coop.
E-(Nash/Core)	2EXPTIME-C	2EXPTIME-C
A-(Nash/Core)	2EXPTIME-C	2EXPTIME-C
(NE/Core)-Mbrshp	2EXPTIME-C	2EXPTIME-C

Table 2: Complexity results for probabilistic systems.

Part II: Repairing

Dealing with missing or bad equilibria

Problem

Individually rational choices can cause outcomes that are highly undesirable, e.g., there is **no equilibrium** or the temporal specification is **not satisfied**.

Question

The problem with this is intrinsic in the system. Can we **repair** it in order to gain (desirable) equilibria?

Solution

Equilibrium Design: **redesign** the game such that individually rational behaviour leads to **desired outcomes**.

Equilibrium Design via Subsidy Scheme

Subsidy scheme

Let $\mathcal{G} = (A, w_1, \dots, w_n)$ be a **Mean-payoff game**.

A **subsidy scheme** for \mathcal{G} is a function $\kappa : \mathbb{N} \times \text{St} \rightarrow \mathbb{N}$.

The **cost** of κ is $\text{cost}(\kappa) = \sum_{i \in \mathbb{N}} \sum_{s \in \text{St}} \kappa(i)(s)$.

Subsidised game

For a **Mean-payoff game** $\mathcal{G} = (A, w_1, \dots, w_n)$ and a subsidy scheme κ , the subsidised game $(\mathcal{G}, \kappa) = (A, w'_1, \dots, w'_n)$ is obtained by updating every player's objective with $w'_i(s) = w_i + \kappa(i)(s)$, for every $s \in \text{St}$.

Intuition

Designers can **incentivise players** to achieve outcomes that are desirable from the temporal specification point of view.

Equilibrium Design Implementation

Definition (Weak Implementation)

For a given game \mathcal{G} , a temporal specification φ and a budget $\beta \in \mathbb{N}$, find a subsidy scheme κ with $\text{cost}(\kappa) \leq \beta$ such that $(\mathcal{G}, \kappa, \varphi)$ solves E-NASH positively.

Definition (Strong Implementation)

For a given game \mathcal{G} , a temporal specification φ and a budget $\beta \in \mathbb{N}$, find a subsidy scheme κ with $\text{cost}(\kappa) \leq \beta$ such that $(\mathcal{G}, \kappa, \varphi)$ solves A-NASH positively.

Filling the toolbox

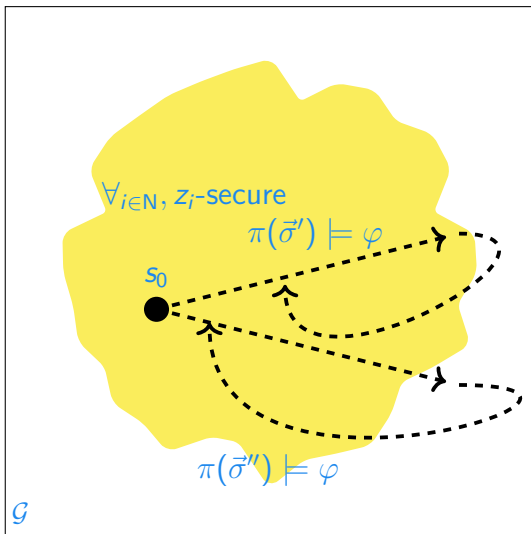
Theorem (Counting subsidy schemes)

The number of subsidy schemes with cost bounded by β is

$$|\mathcal{K}(\mathcal{G}, \beta)| = \frac{\beta + 1}{m} \cdot \binom{\beta + m}{\beta + 1}$$

where $m = |St| \cdot |N|$

Solving Weak Implementation: Intuition




E-NASH:

$\exists \vec{\sigma} \in NE(\mathcal{G}), \pi(\vec{\sigma}) \models \varphi?$

No :-(
The text "No :-(" is written in a grey box, indicating a negative answer to the E-NASH question.

Solving Weak Implementation: Intuition



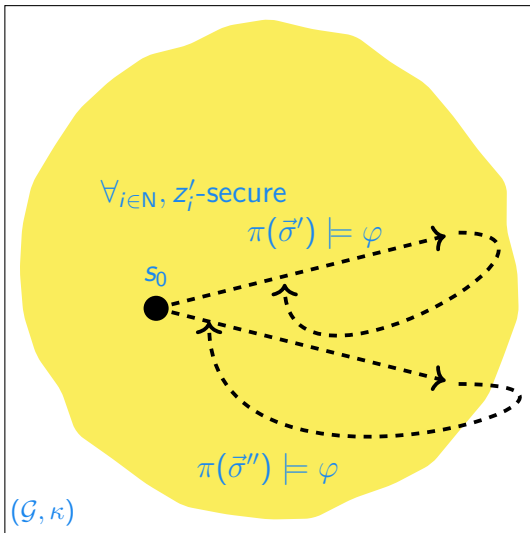
E-NASH:

$\exists \vec{\sigma} \in NE(\mathcal{G}), \pi(\vec{\sigma}) \models \varphi?$

No :-(
The text "No :-(" is written in a larger font size than the text above it.

If only we can modify
the perimeter...

Solving Weak Implementation: Intuition



Apply subsidy scheme

$$\kappa \in \mathcal{K}(\mathcal{G}, \beta)$$

E-NASH:

$$\exists \vec{\sigma} \in \text{NE}(\mathcal{G}), \pi(\vec{\sigma}) \models \varphi?$$

YES :-)

Solving Weak Implementation: Complexity

Complexity

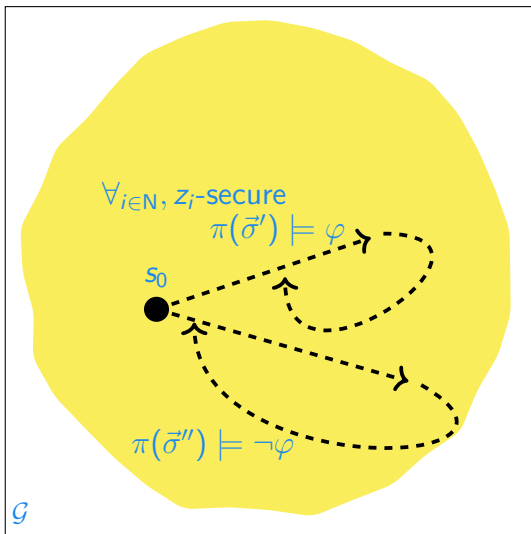
- For LTL specifications: PSPACE-complete (bottleneck is LTL model-checking)

Solving Weak Implementation: Complexity

Complexity

- For **LTL** specifications: **PSPACE-complete** (bottleneck is **LTL** model-checking)
- For **GR(1)** specifications: **NP-complete** (**GR(1)** model checking is poly, all guesses are made together)

Solving Strong Implementation: Intuition

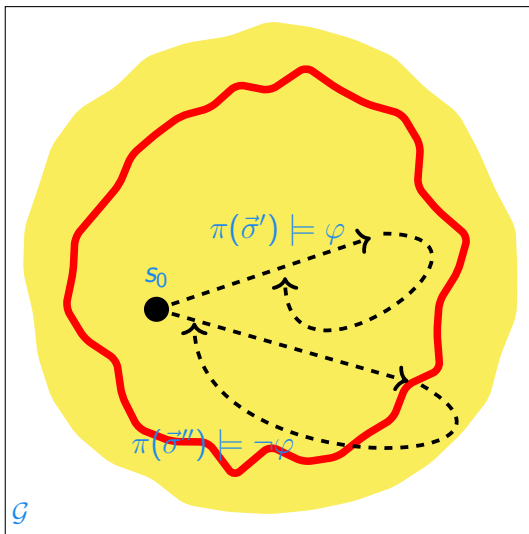


A-NASH:

$\forall \vec{\sigma} \in NE(\mathcal{G}), \pi(\vec{\sigma}) \models \varphi?$

No :-(
The diagram shows a yellow irregular shape representing the set of strategies \mathcal{G} . A black dot labeled s_0 is located on the left side of the shape. Two dashed paths with arrows originate from s_0 . The upper path leads to a point labeled $\pi(\vec{\sigma}') \models \varphi$. The lower path leads to a point labeled $\pi(\vec{\sigma}'') \models \neg\varphi$. Above the upper path, the text $\forall i \in N, z_i\text{-secure}$ is written. In the bottom-left corner of the diagram, the symbol \mathcal{G} is written.

Solving Weak Implementation: Intuition



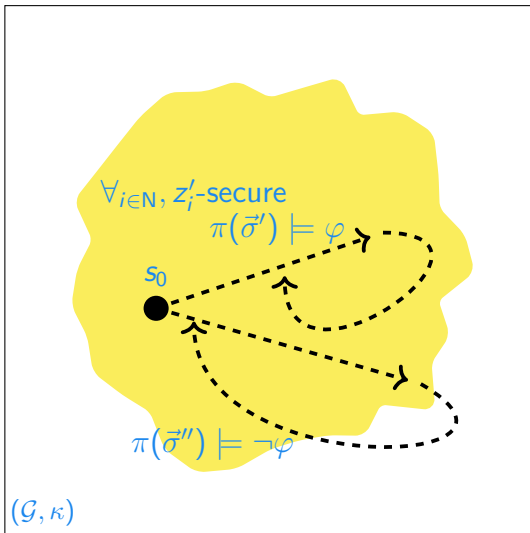
A-NASH:

$\forall \vec{\sigma} \in NE(\mathcal{G}), \pi(\vec{\sigma}) \models \varphi?$

No :-(
The text "No :-(" is written in a light gray box.

If only we can modify
the perimeter...
The text "If only we can modify the perimeter..." is written in a light gray box.

Solving Weak Implementation: Intuition



Apply subsidy scheme

$$\kappa \in \mathcal{K}(\mathcal{G}, \beta)$$

A-NASH:

$$\forall \vec{\sigma} \in NE(\mathcal{G}), \pi(\vec{\sigma}) \models \varphi?$$

YES :-)

Solving Strong Implementation

Complexity

- For LTL specifications: PSPACE-complete (alternating quantification is absorbed)

Solving Strong Implementation

Complexity

- For LTL specifications: PSPACE-complete (alternating quantification is absorbed)
- For GR(1) specifications: Σ_2^P -complete (extra alternation is unavoidable)

Optimizing the budget

For a given game \mathcal{G} , we say that β is the **optimal budget** if it is the **minimum** required to solve weak or strong implementation, respectively.

Definition (Optimality)

Opt-WI For a game \mathcal{G} , compute the optimal budget β for the Weak Implementation.

Opt-SI For a game \mathcal{G} , compute the optimal budget β for the Strong Implementation.

Solving Optimality

Weak Implementation Complexity

- For **LTL** specifications: **FPSPACE-complete** (binary search is absorbed)
- For **GR(1)** specifications: **FP^{NP}-complete**. Hardness via **TSP** problem.

Strong Implementation Complexity

- For **LTL** specifications: **PSPACE-complete** (binary search is absorbed)
- For **GR(1)** specifications: **FP ^{Σ_2^P} -complete**. Hardness via **WEIGHTED MINQSAT₂** problem.

Checking the budget

Definition (Exactness)

Exact-WI For a game \mathcal{G} , check whether b is optimal for the Weak Implementation.

Exact-SI For a game \mathcal{G} , check whether b is the optimal for the Strong Implementation.

Checking uniqueness of the scheme

Definition (Uniqueness)

- UOpt-WI** For a game \mathcal{G} , check whether there is a **unique** subsidy scheme κ for the optimal budget β that solves the Weak Implementation.
- UOpt-SI** For a game \mathcal{G} , check whether there is a **unique** subsidy scheme κ for the optimal budget β that solves the Strong Implementation.

Complexity table summary

	LTL Spec.	GR(1) Spec.
Weak Implementation	PSPACE-complete	NP-complete
Strong Implementation	PSPACE-complete	Σ_2^P -complete
OPT-WI	FPSPACE-complete	FP^{NP} -complete
OPT-SI	FPSPACE-complete	$FP^{\Sigma_2^P}$ -complete
EXACT-WI	PSPACE-complete	D^P -complete
EXACT-SI	PSPACE-complete	D_2^P -complete
UOPT-WI	PSPACE-complete	Δ_2^P -complete
UOPT-SI	PSPACE-complete	Δ_3^P -complete

Epilogue

- An approach to multi-agent systems correctness
- Decision problems and procedures to solve them
- A quest for tractable cases
- A different model with cooperative and probabilistic behaviour
- Future investigation: imperfect information, more quantitative flavour in probabilistic model, learning agents,...

Thank you!