On Computational Tractability for Rational Verification

Julian Gutierrez[1] Muhammad Najib[2]
Giuseppe Perelli[3] Michael Wooldridge[4]

[1]Faculty of Information Technology, Monash University
[2]Department of Computer Science, University of Kaiserslautern
[3]Department of Computer Science, La Sapienza University of Rome
[4]Department of Computer Science, University of Oxford

LAMAS, New Zealand, May 2020

This Paper

Motivation:
Rational verification (*i.e., verification for rational agents*) is intractable in general: 2EXPTIME with LTL goals[1].

Approach:
Using more tractable languages for goals and specifications: goals given by GR(1) formulae and mean-payoff utility functions.

Contribution:
Improved complexity bounds for rational verification.

[1]**GHW**, **Artificial Intelligence**, **2017; WGHMPT**, **AAAI**, **2016**.

Model Checking

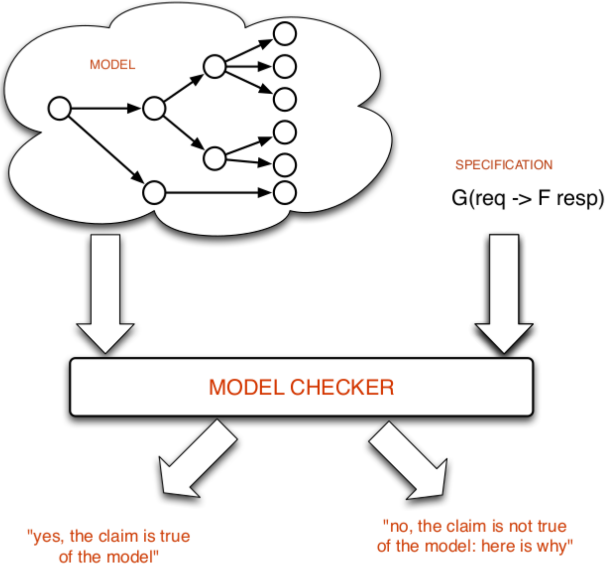

Equilibrium Checking

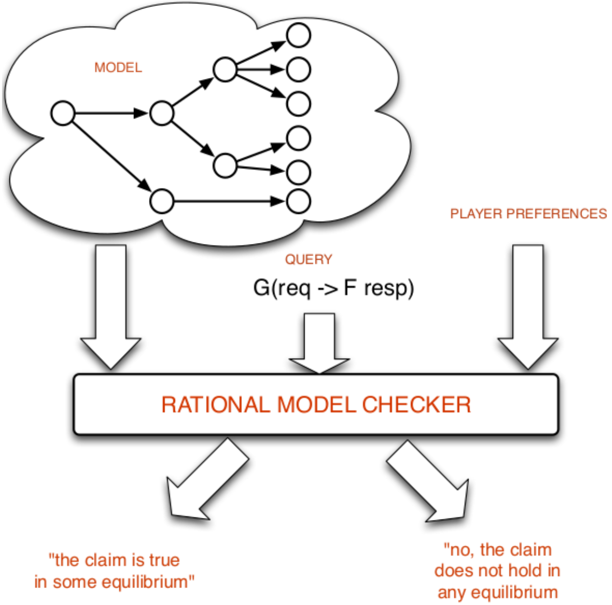

Rational Verification

E-NASH:
Given: a multiagent system $\mathcal{G}$ and a temporal logic formula φ.
Question: Is it the case that $\rho(\vec{\sigma}) \models \varphi$ in **some** $\vec{\sigma} \in NE(\mathcal{G})$?

Other rational verification problem:

- A-NASH: the dual of E-NASH (**all** $\vec{\sigma} \in NE(\mathcal{G})$)
- NON-EMPTINESS: special case of E-NASH ($\varphi = \top$)

GR(1)

The language of *General Reactivity of rank 1*, denoted GR(1), is the fragment of LTL of formulae written in the following form[2]:

$$(\mathbf{GF}\psi_1 \wedge \ldots \wedge \mathbf{GF}\psi_m) \rightarrow (\mathbf{GF}\phi_1 \wedge \ldots \wedge \mathbf{GF}\phi_n),$$

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

$$(\mathbf{GF}\mathit{req}_1 \wedge \mathbf{GF}\mathit{req}_2) \rightarrow \mathbf{GF}\mathit{ack}$$

[2] **BJPPS**, **JCSS**, **2012**.

Mean-payoff value

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

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

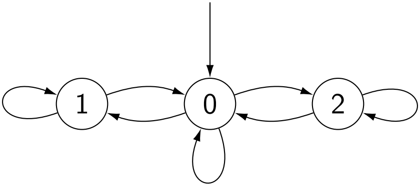

$$\beta^1 = 00000000000\ldots \qquad \mathrm{mp}(\beta^1) = 0$$

$$\beta^2 = 01010101010\ldots \qquad \mathrm{mp}(\beta^2) = 0.5$$

$$\beta^3 = 01020102010\ldots \qquad \mathrm{mp}(\beta^3) = 3/4$$

Games

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

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

A multi-player mp game is a tuple $\mathcal{G}_{\mathsf{mp}} = \langle A, (\mathsf{w}_i)_{i \in \mathsf{N}} \rangle$,

- $A = \langle \mathsf{N}, \mathsf{Ac}, \mathsf{St}, s_0, \mathsf{tr}, \lambda \rangle$ is an arena
- $\mathsf{w}_i : \mathsf{St} \to \mathbb{Z}$ maps states to integer numbers, for each player i

Nash Equilibria

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

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

i.e., **no player can improve its payoff unilaterally.**.

Cases

	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)

In a given GR(1) game $\mathcal{G}$, there exists a Nash Equilibrium if and only if there exists an ultimately periodic path π, such that, for every $k \in \mathbb{N}$, the pair $(s_k, \vec{a}^k)$ of the k-th iteration of π is punishing secure for every $j \in Lose(\pi)$.

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

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

Theorem (NE characterisation[3])

$\vec{\sigma} \in \mathsf{NE}(\mathcal{G})$
$\Leftrightarrow$
$states(\pi(\vec{\sigma})) \subseteq \bigcap_{j \in L} \mathsf{Pun}_j(\mathcal{G})$

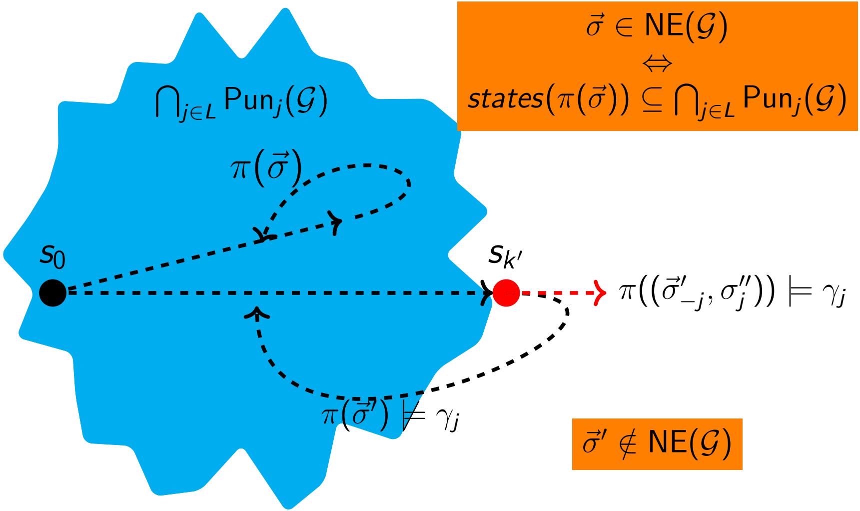

$\vec{\sigma}' \notin \mathsf{NE}(\mathcal{G})$

[3] If L is empty, then $states(\pi(\vec{\sigma})) \subseteq \mathsf{St}$, which is trivially true.

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

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

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

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

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

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* $\pi = (s_0, \vec{a}^0), (s_1, \vec{a}^1), \ldots$, *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

$$\vec{\sigma} \in \mathsf{NE}(\mathcal{G})$$
$$\Leftrightarrow$$
$$\exists_{\vec{z} \in Sec(\mathcal{G})^{\mathsf{N}}} \forall_{i \in \mathsf{N}}, \mathsf{pay}_i(\pi(\vec{\sigma})) \geq z_i$$
$$\Leftrightarrow$$
$$states(\pi(\vec{\sigma})) \subseteq \bigcap_i \mathsf{Pun}_i(\mathcal{G}, \leq_{z_i})$$

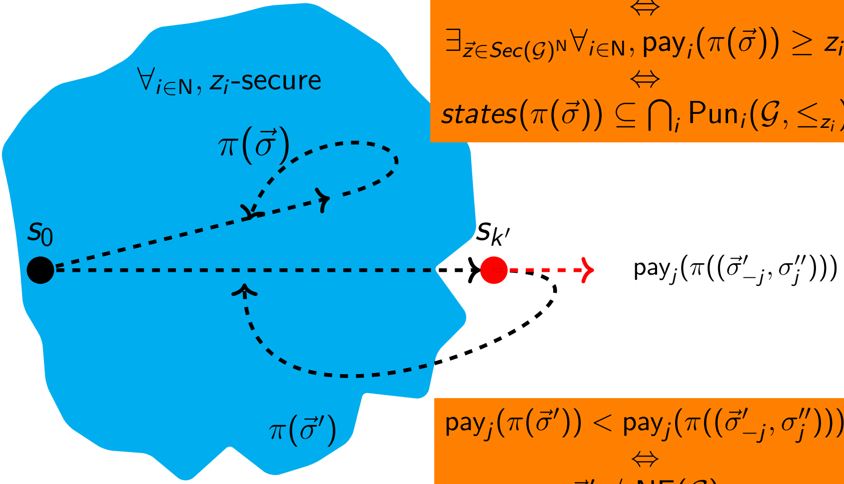

$$\mathsf{pay}_j(\pi(\vec{\sigma}')) < \mathsf{pay}_j(\pi((\vec{\sigma}'_{-j}, \sigma''_j)))$$
$$\Leftrightarrow$$
$$\vec{\sigma}' \notin \mathsf{NE}(\mathcal{G})$$

E-NASH in mp games: the procedure

1. Guess a vector $\vec{z} \in \mathbb{R}^{\mathsf{N}}$ of values, each being a punishment value for a player i, that is, a vector $\vec{z}$ in $Sec_1 \times \ldots \times Sec_{|\mathsf{N}|}$;
2. For each i, compute its z_i-punishment region $\mathrm{Pun}_i(\mathcal{G}, \leq_{z_i})$;
3. Find (u.p.) path $\pi(\vec{\sigma})$ consisting of states in $\bigcap_i \mathrm{Pun}_i(\mathcal{G}, \leq_{z_i})$. Any deviation by player i must remain inside $\mathrm{Pun}_i(\mathcal{G}, \leq_{z_i})$, that is, an ultimately periodic path $\pi(\vec{\sigma})$ satisfying that:
   - $states(\pi(\vec{\sigma})) \subseteq \bigcap_i \mathrm{Pun}_i(\mathcal{G}, \leq_{z_i})$
   - $states(\pi(\vec{\sigma}_{-i}, \sigma'_i)) \subseteq \mathrm{Pun}_i(\mathcal{G}, \leq_{z_i})$, for every i and σ'_i of i
   - $\pi(\vec{\sigma}) \models \varphi$ and $\forall_i, \mathrm{pay}_i(\pi(\vec{\sigma})) \geq z_i$

E-Nash in mp games: the procedure

1. Guess a vector $\vec{z} \in \mathbb{R}^{\mathrm{N}}$ of values, each being a punishment value for a player i, that is, a vector $\vec{z}$ in $Sec_1 \times \ldots \times Sec_{|\mathrm{N}|}$;
2. For each i, compute its z_i-punishment region $\mathrm{Pun}_i(\mathcal{G}, \leq_{z_i})$;
3. Find (u.p.) path $\pi(\vec{\sigma})$ consisting of states in $\bigcap_i \mathrm{Pun}_i(\mathcal{G}, \leq_{z_i})$. Any deviation by player i must remain inside $\mathrm{Pun}_i(\mathcal{G}, \leq_{z_i})$, that is, an ultimately periodic path $\pi(\vec{\sigma})$ satisfying that:
   - $states(\pi(\vec{\sigma})) \subseteq \bigcap_i \mathrm{Pun}_i(\mathcal{G}, \leq_{z_i})$
   - $states(\pi(\vec{\sigma}_{-i}, \sigma'_i)) \subseteq \mathrm{Pun}_i(\mathcal{G}, \leq_{z_i})$, for every i and σ'_i of i
   - $\pi(\vec{\sigma}) \models \varphi$ and $\forall_i, \mathrm{pay}_i(\pi(\vec{\sigma})) \geq z_i$

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:
  - LTL games: 2EXPTIME-complete
  - GR(1) games: FPT
  - mp games: NP-complete
- A-Nash: 2EXPTIME, PSPACE, FPT, PSPACE, coNP.

Proof techniques

In each case, we rely on:

- LTL specifications, LTL goals: LTL synthesis.
- LTL specificadtion, GR(1) goals: LTL model checking.
- GR(1) specification, GR(1) goals: Streett automata.
- LTL specifications, mp goals: $\mathsf{LTL}^{\mathsf{Lim}}$.
- GR(1) specifications, mp goals: LP.