Verification of Cooperative and Concurrent Multi-Player Mean-Payoff Games

Muhammad Najib

Heriot-Watt University

To appear in CSL'24

Joint work with: Julian Gutierrez (Monash), Anthony W. Lin (Kaiserslautern-Landau), Thomas Steeples and Mike Wooldridge (Oxford)

Games and AI

- Long and illustrious history: starting from Turing's 'imitation game'
- Concurrent multi-player games for modelling multi-agent AI systems (ATL, PRISM,...)
  - played in **infinite** sequence of **rounds**
  - **multiple** players/agents[1] chooses actions **simultaneously**
  - each player has a **preference/goal**

[1] we use these terms interchangeably

Games and AI

- Long and illustrious history: starting from Turing's 'imitation game'
- Concurrent multi-player games for modelling multi-agent AI systems (ATL, PRISM,...)
  - played in **infinite** sequence of **rounds**
  - **multiple** players/agents[1] chooses actions **simultaneously**
  - each player has a **preference/goal**
- Concurrent multi-player **mean-payoff** games:
  - Played over a **weighted** graph
  - A play generates an infinite sequence of numbers (weights): $r^0 r^1 r^2 \cdots \in \mathbb{R}^\omega$
  - Players want to maximise a **mean-payoff**: $\mathrm{mp}(r) = \liminf_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} r^i$

[1]we use these terms interchangeably

Games and AI

- Long and illustrious history: starting from Turing's 'imitation game'
- Concurrent multi-player games for modelling multi-agent AI systems (ATL, PRISM,...)
  - played in **infinite** sequence of **rounds**
  - **multiple** players/agents[1] chooses actions **simultaneously**
  - each player has a **preference/goal**
- Concurrent multi-player **mean-payoff** games:
  - Played over a **weighted** graph
  - A play generates an infinite sequence of numbers (weights): $r^0 r^1 r^2 \cdots \in \mathbb{R}^\omega$
  - Players want to maximise a **mean-payoff**: $\mathrm{mp}(r) = \liminf_{n\to\infty} \frac{1}{n} \sum_{i=0}^{n-1} r^i$
- Much research has been done on (2-player) zero-sum, (multi-player) general sum in non-cooperative settings (NE, SPNE)

[1] we use these terms interchangeably

Cooperative AI

- Recently emerged as a prominent topic[2,3,4]
- Agents can communicate (negotiate, reach agreements,...) and benefit from cooperation
- Use mean-payoff games to model *resource-sensitive* cooperative AI systems
- What **outcomes** can/cannot arise given the possibility of cooperation? (rational verification)
- To predict the outcomes, use **solution concept** from cooperative game theory

[2] Allan Dafoe et al. "Cooperative AI: machines must learn to find common ground". In: *Nature* (2021).
[3] Vincent Conitzer and Caspar Oesterheld. "Foundations of Cooperative AI". In: *AAAI*. 2023.
[4] Elisa Bertino et al. *Artificial Intelligence and Cooperation*. Tech. rep. Computing Community Consortium, 2020.

Concurrent multi-player mean-payoff games

Concurrent multi-player mean-payoff game $\mathcal{G} = (A, (\mathrm{w}_i)_{i \in \mathrm{N}})$

- **Arena** $A = \langle \mathrm{N}, \{\mathrm{Ac}_i\}_{i \in \mathrm{N}}, \mathrm{St}, s_{init}, \mathsf{tr}, \mathsf{lab} \rangle$
- **weight function** $\mathrm{w}_i : \mathrm{St} \to \mathbb{Z}$ is a mapping, for every player i, every state of the arena into an integer number.

Player i's Payoff

For an infinite sequence of weights, $w_i = w_i^0 w_i^1 w_i^2 \cdots \in \mathbb{Z}^\omega$, define the payoff $\mathsf{pay}_i(w_i) = \mathsf{mp}(w_i) = \liminf_{n \to \infty} \frac{1}{n} \sum_{t=0}^{n-1} w_i^t$.

Strategies

- A strategy for i can be understood abstractly as a function $\sigma_i : \mathrm{St}^+ \to \mathrm{Ac}_i$ which maps sequences (or histories) of states into a chosen action for player i.
- **memoryless** strategy $\sigma_i : \mathrm{St} \to \mathrm{Ac}_i$ chooses an action based only on the current state of the environment
- **finite-memory** strategy represented by a finite state machine $\sigma_i = (Q_i, q_i^0, \delta_i, \tau_i)$,
  - Q_i is a finite and non-empty set of *internal states*
  - q_i^0 is the *initial state*
  - $\delta_i : Q_i \times \mathrm{St} \to Q_i$ is a deterministic *internal transition function*
  - $\tau_i : Q_i \to \mathrm{Ac}_i$ an *action function*

Strategies

- A strategy for i can be understood abstractly as a function $\sigma_i : \mathrm{St}^+ \to \mathrm{Ac}_i$ which maps sequences (or histories) of states into a chosen action for player i.
- **memoryless** strategy $\sigma_i : \mathrm{St} \to \mathrm{Ac}_i$ chooses an action based only on the current state of the environment
- **finite-memory** strategy represented by a finite state machine $\sigma_i = (Q_i, q_i^0, \delta_i, \tau_i)$,
  - Q_i is a finite and non-empty set of *internal states*
  - q_i^0 is the *initial state*
  - $\delta_i : Q_i \times \mathrm{St} \to Q_i$ is a deterministic *internal transition function*
  - $\tau_i : Q_i \to \mathrm{Ac}_i$ an *action function*

In this work, we assume that players have **finite but unbounded memory**[a] strategies.

1. Practically realisable
2. sufficient to implement LTL specifications

[a]There is previous work in the **memoryless** setting.

Strategies in games

- When each player has chosen a strategy we have a strategy profile $\vec{\sigma} = (\sigma_1, \ldots, \sigma_n)$
- Given a game $\mathcal{G} = \langle A, (\mathsf{w}_i)_{i \in \mathrm{N}} \rangle$ and a strategy profile $\vec{\sigma}$, an outcome $\pi(\vec{\sigma})$ in A induces
  - a sequence $\mathsf{lab}(\pi(\vec{\sigma})) = \mathsf{lab}(s^0)\mathsf{lab}(s^1)\cdots$ of sets of atomic propositions
  - and for each player i, the sequence $\mathsf{w}_i(\pi(\vec{\sigma})) = \mathsf{w}_i(s^0)\mathsf{w}_i(s^1)\cdots$ of weights
- The **payoff** of player i is $\mathsf{pay}_i(\vec{\sigma}) = \mathsf{mp}(\mathsf{w}_i(\pi(\vec{\sigma})))$

Solution Concepts

- Non-cooperative: NE
  - a strategy profile from which no **individual** player has any incentive to **unilaterally deviate**
- Cooperative: the core (introduced by Aumann (2005 Nobel in Economics)[5])
  - the set of strategy profiles from which no **coalition** has any incentive to **deviate**

[5] Robert J Aumann. "The core of a cooperative game without side payments". In: *Trans. of the American Math. Soc.* (1961).

Solution Concepts

- Non-cooperative: NE
  - a strategy profile from which no **individual** player has any incentive to **unilaterally deviate**
- Cooperative: the core (introduced by Aumann (2005 Nobel in Economics)[5])
  - the set of strategy profiles from which no **coalition** has any incentive to **deviate**

strategy in the core

$\vec{\sigma} \in \mathrm{Core}(\mathcal{G})$ if for every coalition $C \subseteq \mathrm{N}$ and (partial) strategy profile $\vec{\sigma}'_C$, there is some (partial) counter-strategy profile $\vec{\sigma}'_{-C}$ such that $\mathsf{pay}_i(\vec{\sigma}) \geq \mathsf{pay}_i(\vec{\sigma}'_C, \vec{\sigma}'_{-C})$

[5] Aumann, "The core of a cooperative game without side payments".

Solution Concepts

- Non-cooperative: NE
  - a strategy profile from which no **individual** player has any incentive to **unilaterally deviate**
- Cooperative: the core (introduced by Aumann (2005 Nobel in Economics)[5])
  - the set of strategy profiles from which no **coalition** has any incentive to **deviate**

strategy in the core

$\vec{\sigma} \in \mathrm{Core}(\mathcal{G})$ if for every coalition $C \subseteq \mathrm{N}$ and (partial) strategy profile $\vec{\sigma}'_C$, there is some (partial) counter-strategy profile $\vec{\sigma}'_{-C}$ such that $\mathsf{pay}_i(\vec{\sigma}) \geq \mathsf{pay}_i(\vec{\sigma}'_C, \vec{\sigma}'_{-C})$

Main differences:

- Players can act in coalitions (as opposed to individuals)
- Counter-strategy can be different from the original strategy

[5] Aumann, "The core of a cooperative game without side payments".

An example: coordination game

- $\mathrm{N} = \{1, 2\}$
- Players are initially in m
- player 1 gets 1 when **both** chooses L
- player 2 gets 1 when **both** chooses R

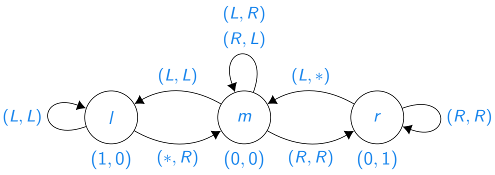

An example: coordination game

- σ_1 prescribes L^ω, σ_2 prescribes R^ω
- (σ_1, σ_2) is a NE, albeit a "bad" one as $\mathsf{pay}_1((\sigma_1, \sigma_2)) = \mathsf{pay}_2((\sigma_1, \sigma_2)) = 0$

An example: coordination game

- σ_1 prescribes L^ω, σ_2 prescribes R^ω
- (σ_1, σ_2) is a NE, albeit a "bad" one as $\text{pay}_1((\sigma_1, \sigma_2)) = \text{pay}_2((\sigma_1, \sigma_2)) = 0$
- But (σ_1, σ_2) is **not** in the core: $\{1, 2\}$ can agree to alternately go L and R
  σ_1', σ_2' prescribe $(LR)^\omega$
- $\text{pay}_1((\sigma_1', \sigma_2')) = \text{pay}_2((\sigma_1', \sigma_2')) = \frac{1}{4}$

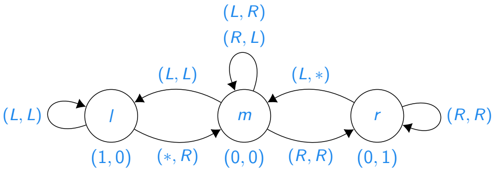

An example: coordination game

- σ_1 prescribes L^ω, σ_2 prescribes R^ω
- (σ_1, σ_2) is a NE, albeit a "bad" one as $\mathrm{pay}_1((\sigma_1, \sigma_2)) = \mathrm{pay}_2((\sigma_1, \sigma_2)) = 0$
- But (σ_1, σ_2) is **not** in the core: $\{1, 2\}$ can agree to alternately go L and R
  σ_1', σ_2' prescribe $(LR)^\omega$
- $\mathrm{pay}_1((\sigma_1', \sigma_2')) = \mathrm{pay}_2((\sigma_1', \sigma_2')) = \frac{1}{4}$

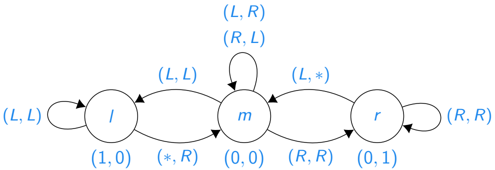

(σ_1', σ_2') corresponds to the liveness property $\varphi := \mathbf{GF}l \wedge \mathbf{GF}r$
i.e., $\pi((\sigma_1', \sigma_2')) \models \varphi$.

An example: coordination game

- σ_1 prescribes L^ω, σ_2 prescribes R^ω
- (σ_1, σ_2) is a NE, albeit a "bad" one as $\mathrm{pay}_1((\sigma_1, \sigma_2)) = \mathrm{pay}_2((\sigma_1, \sigma_2)) = 0$
- But (σ_1, σ_2) is **not** in the core: $\{1, 2\}$ can agree to alternately go L and R
  σ_1', σ_2' prescribe $(LR)^\omega$
- $\mathrm{pay}_1((\sigma_1', \sigma_2')) = \mathrm{pay}_2((\sigma_1', \sigma_2')) = \frac{1}{4}$

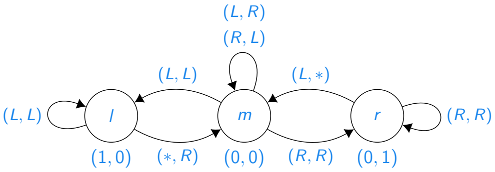

(σ_1', σ_2') corresponds to the liveness property $\varphi := \mathbf{GF}l \wedge \mathbf{GF}r$
i.e., $\pi((\sigma_1', \sigma_2')) \models \varphi$.

- **All** strategy profiles in the core satisfy φ
- **All** strategy profiles in the core require **memory**

How to characterise the core?

- Previous work[6] in the memoryless setting: guess a correct strategy profile (poly size)
- Strategies have arbitrarily large memories: **no bounds** on the search space

[6] Thomas Steeples, Julian Gutierrez, and Michael Wooldridge. "Mean-payoff games with ω-regular specifications". In: *AAMAS*. 2021.

[7] Romain Brenguier and Jean-François Raskin. "Pareto Curves of Multidimensional Mean-Payoff Games". In: *CAV*. 2015.

How to characterise the core?

- Previous work[6] in the memoryless setting: guess a correct strategy profile (poly size)
- Strategies have arbitrarily large memories: **no bounds** on the search space
- Can we characterise the core using Pareto optimality[7]?

[6] Steeples, Gutierrez, and Wooldridge, "Mean-payoff games with ω-regular specifications".
[7] Brenguier and Raskin, "Pareto Curves of Multidimensional Mean-Payoff Games".

How to characterise the core?

- Previous work[6] in the memoryless setting: guess a correct strategy profile (poly size)
- Strategies have arbitrarily large memories: **no bounds** on the search space
- Can we characterise the core using Pareto optimality[7]? No

Proposition

There exist games $\mathcal{G}$ such that $\vec{\sigma} \in \mathrm{Core}(\mathcal{G})$ and $\vec{\sigma}$ is not Pareto optimal.

Proposition

There exist games $\mathcal{G}$ such that $\vec{\sigma}$ is Pareto optimal and $\vec{\sigma} \notin \mathrm{Core}(\mathcal{G})$.

[6] Steeples, Gutierrez, and Wooldridge, "Mean-payoff games with ω-regular specifications".

[7] Brenguier and Raskin, "Pareto Curves of Multidimensional Mean-Payoff Games".

How to characterise the core?

- In general, the core does not coincide with Pareto optimality
- But PO is still useful!

[8] Brenguier and Raskin, "Pareto Curves of Multidimensional Mean-Payoff Games".

How to characterise the core?

- In general, the core does not coincide with Pareto optimality
- But PO is still useful!
- For a given game $\mathcal{G}$ and $C \subseteq \mathrm{N}$, we **sequentialise** into 2-player multi-mean-payoff game $G^C = (V_1, V_2, E, w)$, where
    - C acts as player 1 who owns V_1
    - $-C$ acts as player 2 who owns V_2
    - $w : V_1 \cup V_2 \to \mathbb{Z}^c$ corresponds to k-dimensional vectors representing the weight functions of C
- $\text{val}(G^C, s)$ is the set of values that can be **enforced** by C, and $\text{val}(G^C, s) =\downarrow \mathrm{PO}(G^C, s)$
- Brenguier and Raskin[8] showed that
    1. $\text{val}(G^C, s)$ can be represented as finite union of polyhedra
    2. For every polyhedron P, there is a vector $\vec{v} \in P$ whose representation is of poly size

[8] Brenguier and Raskin, "Pareto Curves of Multidimensional Mean-Payoff Games".

How to characterise the core?

Can we characterise the set of values corresponding to the core as a polyhedron?

How to characterise the core?

Can we characterise the set of values corresponding to the core as a polyhedron?

How to characterise the core?

Want: characterise the set of values corresponding to the core as a polyhedron

How to characterise the core?

Want: characterise the set of values corresponding to the core as a polyhedron

Lemma

The set $\mathrm{val}(G^C)$ *can be represented by a finite union of a set of polyhedra* $\mathrm{PS}(G^C)$*, and each polyhedron* $P_j^C \in \mathrm{PS}(G^C)$ *is polynomially representable.*

How to characterise the core?

Want: characterise the set of values corresponding to the core as a polyhedron

Lemma

The set $\mathsf{val}(G^C)$ *can be represented by a finite union of a set of polyhedra* $\mathrm{PS}(G^C)$*, and each polyhedron* $P_j^C \in \mathrm{PS}(G^C)$ *is polynomially representable.*

Lemma

If $\vec{\sigma} \in \mathrm{Core}(\mathcal{G})$ *then for each* $C \subseteq \mathrm{N}$ *and* $P_j^C \in \mathrm{PS}(G^C)$ *there is a half-space* H *of* P_j^C *such that vector* $(\mathsf{pay}_i(\vec{\sigma}))_{i \in C}$ *is in* $\overline{H}$ *(i.e., closed complement of* H*)*

$(\mathsf{pay}_i(\vec{\sigma}))_{i \in C}$ is NOT **strictly** contained in P_j^C

An intuitive example

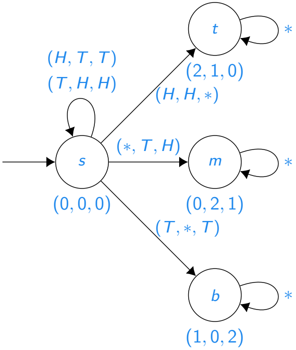

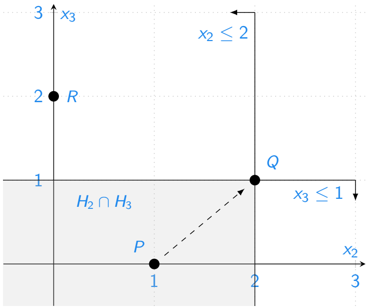

Figure 1: Left: Arena for the example. Right: Graphical representation of $\mathrm{val}(G^{\{2,3\}})$. Coordinates P, Q, R corresponds to the set $\mathrm{PO}(G^{\mathrm{N}}) = \{(2,1,0),(0,2,1),(1,0,2)\}$. There is a beneficial deviation by $\{2,3\}$ (dashed arrow) from P (the $\{1,2\}$-Pareto optimal value) to Q (the $\{2,3\}$-Pareto optimal value).

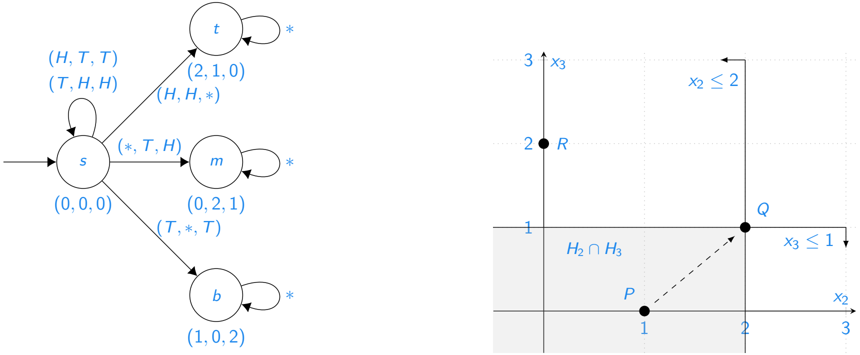

Figure 2: Left: Arena for the example. Right: Graphical representation of $\mathrm{val}(G^{\{2,3\}})$. Coordinates P, Q, R corresponds to the set $\mathrm{PO}(G^{\mathrm{N}}) = \{(2,1,0),(0,2,1),(1,0,2)\}$. There is a beneficial deviation by $\{2,3\}$ (dashed arrow) from P (the $\{1,2\}$-Pareto optimal value) to Q (the $\{2,3\}$-Pareto optimal value).

An intuitive example modified

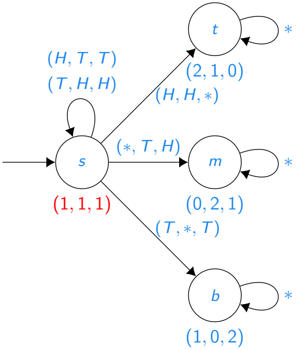

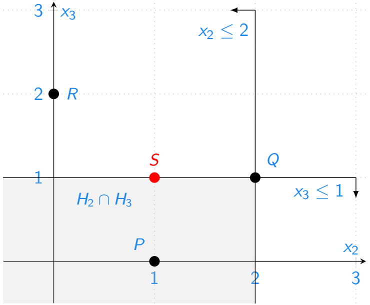

Figure 3: Left: Arena for the modified example. Right: Graphical representation of $\mathrm{val}(G'^{\{2,3\}})$. Coordinates P, Q, R, S corresponds to the set $\mathrm{PO}(G'^{\mathrm{N}}) = \{(2,1,0), (0,2,1), (1,0,2), (1,1,1)\}$. There is no beneficial deviation from S.

An intuitive example modified

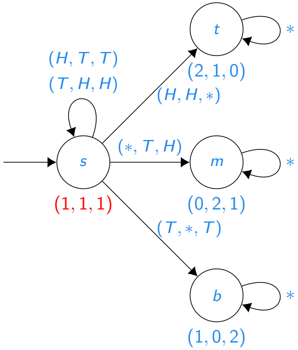

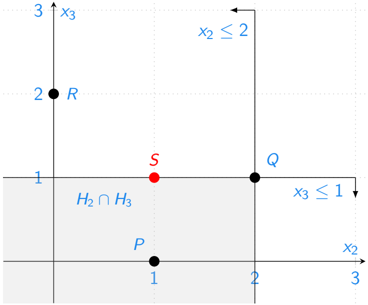

Figure 3: Left: Arena for the modified example. Right: Graphical representation of $\mathsf{val}(G'^{\{2,3\}})$. Coordinates P, Q, R, S corresponds to the set $\mathrm{PO}(G'^{\mathrm{N}}) = \{(2,1,0),(0,2,1),(1,0,2),(1,1,1)\}$. There is no beneficial deviation from S.

$S \in \overline{H}_3$. Indeed, for each $C \subseteq \mathrm{N}$ there is such a "**blocking**" half-space. If we take the intersection of such **blocking** half-spaces and $\mathsf{val}(G'^{\mathrm{N}})$ we obtain $\{(1,1,1)\}$

From intuition to characterisation

If the intersection of such **blocking** half-spaces and $\mathsf{val}(G'^{\mathrm{N}})$ is non-empty, then the core is non-empty.

Theorem

$\mathrm{Core}(\mathcal{G}) \neq \varnothing$ iff there exists a set of **blocking** *half-spaces I such that $R = \bigcap_{H \in I} \overline{H} \cap \mathsf{val}(G^{\mathrm{N}}) \neq \varnothing$*

[9] Brenguier and Raskin, "Pareto Curves of Multidimensional Mean-Payoff Games".

From intuition to characterisation

If the intersection of such **blocking** half-spaces and $\mathsf{val}(G'^{\mathrm{N}})$ is non-empty, then the core is non-empty.

Theorem

$\mathrm{Core}(\mathcal{G}) \neq \varnothing$ *iff there exists a set of* **blocking** *half-spaces* I *such that* $R = \bigcap_{H \in I} \overline{H} \cap \mathsf{val}(G^{\mathrm{N}}) \neq \varnothing$

R is a polyhedron representable polynomially wrt $\mathcal{G}$, as such, there exists a **polynomial witness**[9] $\vec{x} \in R$.

Theorem

Given a game $\mathcal{G}$, *if the core is non-empty, then there is* $\vec{\sigma} \in \mathrm{Core}(\mathcal{G})$ *such that* $(\mathsf{pay}_i(\vec{\sigma}))_{i \in \mathrm{N}}$ *can be represented polynomially in the size of* $\mathcal{G}$.

[9] Brenguier and Raskin, "Pareto Curves of Multidimensional Mean-Payoff Games".

Strategies in the core and how to find them

- Strategies have arbitrarily large memories: **no bounds** on the search space
- We do not have to guess the strategies, only need to guess a polynomial witness vector $\vec{x} \in R$

Strategies in the core and how to find them

- Strategies have arbitrarily large memories: **no bounds** on the search space
- We do not have to guess the strategies, only need to guess a polynomial witness vector $\vec{x} \in R$
- Finding strategies in the core can be reduced to finding $\vec{x} \in R$

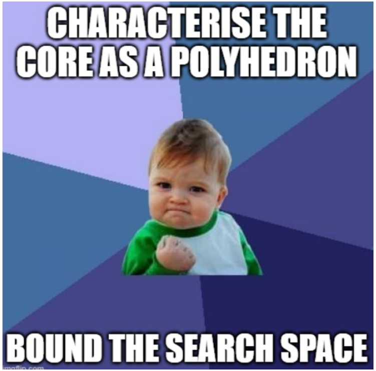

Rational Verification: Decision Problems

- **Universality: all** strategy profiles in the core **satisfy** a LTL property φ (A-CORE)
- **Existence:** there **exists** a strategy profile in the core satisfying a LTL property φ (E-CORE)
- **Stability:** Is the core non-empty? (NON-EMPTINESS)

Coordination game revisited

- Non-Emptiness returns YES

Coordination game revisited

- Non-Emptiness returns YES
- A-Core with $\varphi := \mathbf{GF}l \wedge \mathbf{GF}r$ returns YES

Coordination game revisited

- NON-EMPTINESS returns YES
- A-CORE with $\varphi := \mathbf{GF}l \wedge \mathbf{GF}r$ returns YES
- E-CORE with $\varphi := \mathbf{G}m$ returns NO

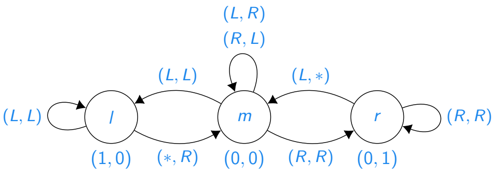

Solving Non-Emptiness

Given: game $\mathcal{G}$

Non-Emptiness: Is it the case that $\mathrm{Core}(\mathcal{G}) \neq \varnothing$?

Procedure:

1. Guess a vector $\vec{x} \in R$
2. Check if $\vec{x}$ admits beneficial deviations, then return **NO**; otherwise return **YES**

- Step 1 is in NP
- Step 2 involves calling Σ_2^P oracle
- Non-Emptiness can be solved in Σ_3^P

Solving E-Core and A-Core

Given: Game $\mathcal{G}$, formula φ.

E-CORE: Is it the case that there exists some $\vec{\sigma} \in \mathrm{Core}(\mathcal{G})$ such that $\vec{\sigma} \models \varphi$?

A-CORE: Is it the case that for all $\vec{\sigma} \in \mathrm{Core}(\mathcal{G})$, we have $\vec{\sigma} \models \varphi$?

Observations:

- A witness to E-CORE would be a path π such that
  1. $\mathsf{pay}_i(\pi))_{i \in \mathrm{N}} \geq (\mathsf{pay}_i(\vec{\sigma}))$ for some $\vec{\sigma} \in \mathrm{Core}(\mathcal{G})$
  2. $\pi \models \varphi$
- φ has an **ultimately periodic** model of size $2^{O(|\varphi|)}$,[10] thus the size of representation of $\mathsf{pay}_i(\pi)$ is polynomial wrt $\mathcal{G}$

[10] A. P. Sistla and E. M. Clarke. "The complexity of propositional linear temporal logics". In: *J. ACM* (1985).

Solving E-Core and A-Core

Procedure:

1. Guess a vector $\vec{x} \in R$ and set of states $S \subseteq \mathrm{St}$
2. If all $s \in S$ is "**safe**", then
   2.1 Produce a subgame $\mathcal{G}[S]$ by removing all states $s \notin S$
   2.2 If there is π in $\mathcal{G}[S]$ with $\mathrm{pay}_i(\pi) \geq x_i, \forall i \in \mathrm{N}$ and $\pi \models \varphi$, return **YES**
3. Return **NO**

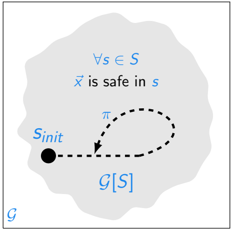

Solving E-Core and A-Core

Procedure:

1. Guess a vector $\vec{x} \in R$ and set of states $S \subseteq \mathrm{St}$
2. If all $s \in S$ is "**safe**", then
   2.1 Produce a subgame $\mathcal{G}[S]$ by removing all states $s \notin S$
   2.2 If there is π in $\mathcal{G}[S]$ with $\mathsf{pay}_i(\pi) \geq x_i, \forall i \in \mathrm{N}$ and $\pi \models \varphi$, return **YES**
3. Return **NO**

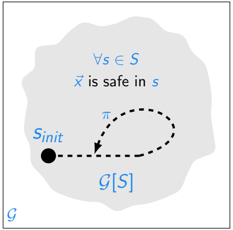

- Step 1 is in NP
- Step 2 is in Σ_2^P
- Step 2.1 is in PTIME
- Step 2.2:
  - If φ is in the full LTL: PSPACE
  - If φ is in an "easy" fragment of LTL (e.g., GR(1)): PTIME
- Solving E-Core is in PSPACE (full LTL) or Σ_3^P (easy LTL fragment)
- Solving A-Core is in PSPACE (full LTL) or Π_3^P (easy LTL fragment)

Complexity landscape

Problem	Finite Memory	Memoryless[11]	NE[12,13]
Non-Emptiness	Σ_3^P-c	Σ_2^P	NP-c
E-Core with LTL spec.	PSPACE-c		PSPACE-c
A-Core with LTL spec.	PSPACE-c		PSPACE-c
E-Core with GR(1) spec.	Σ_3^P-c	Σ_2^P	NP-c
A-Core with GR(1) spec.	Π_3^P-c	Π_2^P	coNP-c

Figure 4: Summary of complexity results. The NE column shows complexity results for the corresponding decision problems with NE.

[11] Steeples, Gutierrez, and Wooldridge, "Mean-payoff games with ω-regular specifications".
[12] M. Ummels and D. Wojtczak. "The Complexity of Nash Equilibria in Limit-Average Games". In: *CONCUR*. 2011.
[13] Julian Gutierrez et al. "On Computational Tractability for Rational Verification". In: *IJCAI*. 2019.

Concluding remarks

- Characterisation of the core using discrete geometry
- Showed that in our setting, the core admits polynomial witnesses
- Tight complexity bounds for rational verification problems

Future work:

- Can we establish sufficient and necessary conditions of non-emptiness of the core in a broader sense, e.g., quasi-concavity of utility functions[14]
- What can we do when the core is empty? Modify the games, e.g., utility functions, norms to limit actions,...
- Can we use our characterisation here to extend ATL* with mean-payoff semantics?

[14]Herbert E Scarf. "On the existence of a cooperative solution for a general class of N-person games". In: *J. of Economic Theory* (1971).