Towards the Verification of General Game Playing

Francesco Belardinelli

Imperial College London & Université d'Evry

joint work with V. Malvone

AGAPE kick-off meeting, 28 March 2019

Background

- **General Game Playing (GGP)**: paradigm in AI focusing on developing algorithms and techniques to play more than one game successfully.
- **Game Description Language (GDL)**: logic-based language to represent games in GGP.

Research Question

How can we **express** *and* **verify** *game-theoretical notions and solution concepts in GDL?*

Verification by Model Checking

1. Model system S as a graph/game M_S
2. Model property P as a formula ϕ_P
3. Check $M_S \models \phi_P$

In our case M_S is specified in GDL.

But what about ϕ_P?

Research Question

How can we **express** *and* **verify** *game-theoretical notions and solution concepts in GDL?*

Verification by Model Checking

1. Model system S as a graph/game M_S
2. Model property P as a formula ϕ_P
3. Check $M_S \models \phi_P$

In our case M_S is specified in GDL.

But what about ϕ_P?

State of the Art

[RvdHW09] extends GDL with operators for individual and joint strategies in Alternating-time temporal logic (ATL).

$(M, s) \models \langle\langle A \rangle\rangle \psi$ iff agents in A have a joint strategy such that,
no matter what the agents in $Ag \setminus A$ do,
$(M, s) \models \psi$

Normally, ψ is a formula in LTL:

$$\psi \quad = \quad p \mid \neg\phi \mid \phi \wedge \phi \mid X\phi \mid \phi U \phi$$

State of the Art

Expressivity:

- **Termination**: the game will eventually end.

$$\langle\langle \emptyset \rangle\rangle F terminal$$

- **Playability**: every player has at least one move in every non-terminal state.

$$\langle\langle \emptyset \rangle\rangle G(\neg terminal \rightarrow \bigwedge_{i \in Ag} \bigvee_{a \in Act_i, a \neq fin_i} legal(i, a))$$

- **Strong Winnability**: some player has a strategy to win.

$$\bigvee_{i \in Ag} \langle\langle i \rangle\rangle F win_i$$

- **Fairness**: all players might win.

$$\langle\langle Ag \rangle\rangle F win_i$$

- **Collusion**: coalition C has a strategy to enforce a win for agent i.

$$\langle\langle C \rangle\rangle F win_i$$

State of the Art

Expressivity:

- **Termination**: the game will eventually end.

$$\langle\langle\emptyset\rangle\rangle F terminal$$

- **Playability**: every player has at least one move in every non-terminal state.

$$\langle\langle\emptyset\rangle\rangle G(\neg terminal \rightarrow \bigwedge_{i \in Ag} \bigvee_{a \in Act_i, a \neq fin_i} legal(i,a))$$

- **Strong Winnability**: some player has a strategy to win.

$$\bigvee_{i \in Ag} \langle\langle i \rangle\rangle F win_i$$

- **Fairness**: all players might win.

$$\langle\langle Ag \rangle\rangle F win_i$$

- **Collusion**: coalition C has a strategy to enforce a win for agent i.

$$\langle\langle C \rangle\rangle F win_i$$

State of the Art

Expressivity:

- **Termination**: the game will eventually end.

$$\langle\langle \emptyset \rangle\rangle F terminal$$

- **Playability**: every player has at least one move in every non-terminal state.

$$\langle\langle \emptyset \rangle\rangle G(\neg terminal \rightarrow \bigwedge_{i \in Ag} \bigvee_{a \in Act_i, a \neq fin_i} legal(i, a))$$

- **Strong Winnability**: some player has a strategy to win.

$$\bigvee_{i \in Ag} \langle\langle i \rangle\rangle F win_i$$

- **Fairness**: all players might win.

$$\langle\langle Ag \rangle\rangle F win_i$$

- **Collusion**: coalition C has a strategy to enforce a win for agent i.

$$\langle\langle C \rangle\rangle F win_i$$

State of the Art

Expressivity:

- **Termination**: the game will eventually end.

$$\langle\langle\emptyset\rangle\rangle F terminal$$

- **Playability**: every player has at least one move in every non-terminal state.

$$\langle\langle\emptyset\rangle\rangle G(\neg terminal \rightarrow \bigwedge_{i \in Ag} \bigvee_{a \in Act_i, a \neq fin_i} legal(i, a))$$

- **Strong Winnability**: some player has a strategy to win.

$$\bigvee_{i \in Ag} \langle\langle i \rangle\rangle F win_i$$

- **Fairness**: all players might win.

$$\langle\langle Ag \rangle\rangle F win_i$$

- **Collusion**: coalition C has a strategy to enforce a win for agent i.

$$\langle\langle C \rangle\rangle F win_i$$

State of the Art

Expressivity:

- **Termination**: the game will eventually end.

$$\langle\langle \emptyset \rangle\rangle F terminal$$

- **Playability**: every player has at least one move in every non-terminal state.

$$\langle\langle \emptyset \rangle\rangle G(\neg terminal \rightarrow \bigwedge_{i \in Ag} \bigvee_{a \in Act_i, a \neq fin_i} legal(i, a))$$

- **Strong Winnability**: some player has a strategy to win.

$$\bigvee_{i \in Ag} \langle\langle i \rangle\rangle F win_i$$

- **Fairness**: all players might win.

$$\langle\langle Ag \rangle\rangle F win_i$$

- **Collusion**: coalition C has a strategy to enforce a win for agent i.

$$\langle\langle C \rangle\rangle F win_i$$

State of the Art

Main result: Model checking ATL on GDL models is EXPTIME-c, *when assuming perfect information*.

Bad, isn't it?
The complexity is given on *compact representations*.

State of the Art

Main result: Model checking ATL on GDL models is EXPTIME-c, *when assuming perfect information*.

Bad, isn't it?
The complexity is given on *compact representations*.

State of the Art

Main result: Model checking ATL on GDL models is EXPTIME-c, *when assuming perfect information*.

Bad, isn't it?

The complexity is given on *compact representations*.

State of the Art

Main result: Model checking ATL on GDL models is EXPTIME-c, *when assuming perfect information*.

Bad, isn't it?
The complexity is given on *compact representations*.

Imperfect Information

But what about imperfect information?

- Players often have imperfect/incomplete information about the state of the game (poker, battleships, etc.)
- In auctions bidders do not know other bidders' preferences.

[Thi10]: Game Description Language with Imperfect Information (GDL-II)

Verification of ATL under Imperfect Information

Unfortunately, imperfect information makes things hard(er)

	perfect inf.	**imperfect inf.**
memoryless	PTIME-c. [AHK02]	Δ_2^P-c. [JD06]
perfect recall		undec. [DT11]

UROP Project:

- Extend GDL to express and verify **efficiently** game-theoretical notions and solution concepts **under imperfect information**.

Research Challenge: how to tame imperfect information

[DT11]: the source of undecidability lies in the interplay between. . .

1. agents having incomparable observations
2. agents using private communication

What happens if we drop 1 or 2?

Hereafter we focus on dropping 2.

Research Challenge: how to tame imperfect information

[DT11]: the source of undecidability lies in the interplay between. . .

1. agents having incomparable observations
2. agents using private communication

What happens if we drop 1 or 2?

Hereafter we focus on dropping 2.

Idea: drop private communication

- Public Announcement Logic is decidable [GG97]
- Epistemic planning is easier [DMP18]
- LTLK synthesis is decidable [vdMW05]

Research question: *are there meaningful restrictions on imperfect information that preserves tractability/decidability?*

MAS with Public Actions Only

PA-MAS

A MAS S **has public actions only** if any action is visible to any agent.

That is, no private communication can take place.

PA-MAS capture many interesting scenarios in CS/Game Theory

- card/board games: poker, bridge, etc.
- *open-outcry auctions*
- communication by broadcast, including tweeting [LvdMR00]
- planning via public actions [KG15]
- recording contexts [FHMV95]

Concurrent Game Structures with Imperfect Information

iCGS

An iCGS $S = \langle Ag, AP, \{Act_a\}_{a \in Ag}, S, S_0, \delta, \lambda, \{\sim_a\}_{a \in Ag} \rangle$ includes

- agents in Ag
- atomic propositions in AP
- actions in Act_a and joint actions in $ACT = \prod_{a \in Ag} Act_a$
- states in S with initial states in $S_0 \subseteq S$
- transition function $\delta : S \times ACT \to S$
- labelling function $\lambda : AP \to 2^S$
- **indistinguishability relation** $\sim_a \subseteq S \times S$ for every agent $a \in Ag$.

Formal models for PA-MAS:

PA-iCGS

An iCGS S **has public actions only** if for every agent $a \in Ag$, states $s, s' \in S$, and joint actions $J, J' \in ACT$,

$$\text{if } J \neq J' \text{ then } \delta(s, J) \not\sim_a \delta(s', J')$$

Decidable Model Checking

[DT11]: Model checking ATL on iCGS **with perfect recall** is undecidable.

Theorem ([BLMR17a, BLMR17b])

Model checking ATL on iCGS with public actions only is decidable.

[BLMR18]: Model checking ATL remains decidable if we allow for a finite number of "private" actions.

$\Rightarrow$ useful to model collusion in auctions.

Take-home Message:

- Complex specifications can *in principle* be checked on MAS as long as evolution is *almost always* via public actions.

The UROP Project

How can we **express** *and* **verify** *game-theoretical notions and solution concepts in GDL* **under imperfect information***?*

We leverage on

- Extension of GDL with ATL operators [RvdHW09]
- Decidability results for iCGS with public actions only [BLMR17a, BLMR17b, BLMR18]

Project Outline:

1. State of the art on GDL and ATL.
2. Interpretation of ATL on CGSi-like models for GDL-II.
3. Expressivity, applications to auction.
4. Model Checking (particularly on PA-iCGS).
5. Implementation (?)

The UROP Project

How can we **express** *and* **verify** *game-theoretical notions and solution concepts in GDL* **under imperfect information***?*

We leverage on

- Extension of GDL with ATL operators [RvdHW09]
- Decidability results for iCGS with public actions only [BLMR17a, BLMR17b, BLMR18]

Project Outline:

1. State of the art on GDL and ATL.
2. Interpretation of ATL on CGSi-like models for GDL-II.
3. Expressivity, applications to auction.
4. Model Checking (particularly on PA-iCGS).
5. Implementation (?)

The UROP Project

How can we **express** *and* **verify** *game-theoretical notions and solution concepts in GDL* **under imperfect information***?*

We leverage on

- Extension of GDL with ATL operators [RvdHW09]
- Decidability results for iCGS with public actions only [BLMR17a, BLMR17b, BLMR18]

Project Outline:

1. State of the art on GDL and ATL.
2. Interpretation of ATL on CGSi-like models for GDL-II.
3. Expressivity, applications to auction.
4. Model Checking (particularly on PA-iCGS).
5. Implementation (?)

AGAPE

1. *How can we verify properties for any number of players in the system?*
   $\Rightarrow$ Parametric model checking (?)
2. *How can we manipulate efficiently numerical values in model checkers?*
3. *Strategically equivalent $\approx$ bisimilar?*

Cannons for a paper castle?

AGAPE

1. *How can we verify properties for any number of players in the system?*
   $\Rightarrow$ Parametric model checking (?)
2. *How can we manipulate efficiently numerical values in model checkers?*
3. *Strategically equivalent $\approx$ bisimilar?*

Cannons for a paper castle?

AGAPE

1. *How can we verify properties for any number of players in the system?*
   $\Rightarrow$ Parametric model checking (?)
2. *How can we manipulate efficiently numerical values in model checkers?*
3. *Strategically equivalent* $\approx$ *bisimilar?*

Cannons for a paper castle?

AGAPE

1. *How can we verify properties for any number of players in the system?*
   $\Rightarrow$ Parametric model checking (?)
2. *How can we manipulate efficiently numerical values in model checkers?*
3. *Strategically equivalent* $\approx$ *bisimilar?*

Cannons for a paper castle?

References

R. Alur, T. A. Henzinger, and O. Kupferman.
Alternating-time temporal logic.
J. ACM, 49(5):672–713, 2002.

F. Belardinelli, A. Lomuscio, A. Murano, and S. Rubin.
Verification of multi-agent systems with imperfect information and public actions.
In *17*, pages 1268–1276, 2017.

Francesco Belardinelli, Alessio Lomuscio, Aniello Murano, and Sasha Rubin.
Verification of broadcasting multi-agent systems against an epistemic strategy logic.
In Carles Sierra, editor, *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017, Melbourne, Australia, August 19-25, 2017*, pages 91–97. ijcai.org, 2017.

F. Belardinelli, A. Lomuscio, A. Murano, and S. Rubin.
Decidable verification of multi-agent systems with bounded private actions.
In *Proc. of the 17th International Conference on Autonomous Agents and Multiagent Systems (AAMAS18)*, 2018.

Catalin Dima, Bastien Maubert, and Sophie Pinchinat.
Relating paths in transition systems: The fall of the modal mu-calculus.
ACM Trans. Comput. Log., 19(3):23:1–23:33, 2018.

C. Dima and F.L. Tiplea.
Model-checking ATL under imperfect information and perfect recall semantics is undecidable.
CoRR, abs/1102.4225, 2011.

R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi.
Reasoning about Knowledge.
MIT, 1995.

References

J. Gerbrandy and W. Groeneveld.
Reasoning About Information Change.
6(2):147–169, 1997.

W. Jamroga and J. Dix.
Model checking abilities under incomplete information is indeed Δ_p^2-complete.
In *EUMAS'06*, pages 14–15, 2006.

Filippos Kominis and Hector Geffner.
Beliefs in multiagent planning: From one agent to many.
In *ICAPS*, pages 147–155. AAAI Press, 2015.

Alessio R. Lomuscio, Ron van der Meyden, and Mark Ryan.
Knowledge in multiagent systems: Initial configurations and broadcast.
ACM Trans. Comput. Logic, 1(2):247–284, October 2000.

Ji Ruan, Wiebe van der Hoek, and Michael Wooldridge.
Verification of games in the game description language.
J. Log. Comput., 19(6):1127–1156, 2009.

Michael Thielscher.
A general game description language for incomplete information games.
In *Proceedings of the Twenty-Fourth AAAI Conference on Artificial Intelligence*, AAAI'10, pages 994–999. AAAI Press, 2010.

R. van der Meyden and T. Wilke.
Synthesis of distributed systems from knowledge-based specifications.
In *05*, LNCS 3653, pages 562–576. Springer-Verlag, 2005.