Reasoning about interactions when all actions are public (abstract of talk at ACTIONS@KR'18)*

Sasha Rubin
University of Naples "Federico II"
rubin@unina.it

Many multi-agent planning tasks are undecidable or computationally hard (assuming perfect-recall and partial-observability): For instance:

1. Finding winning strategies in games with reachability objectives is undecidable for multi-player games [18] and ExpTime-complete with two players [21].
2. Synthesising reactive modules for LTL objectives is undecidable for distributed systems [20] and 2ExpTime-complete for a single module [19].
3. Model checking ATL with perfect recall is PTime-complete assuming perfect-information and undecidable assuming partial-observability [1; 7].
4. Deciding the existence of a Nash equilibrium in a game with LTL objectives is 2ExpTime-complete assuming perfect-information [16] and undecidable assuming partial-observability [10].

There are a number of ways to lower the complexity. These include:

1. limit the recall of the agents, e.g., to imperfect-recall [1],
2. limit the observations of the agents, e.g., to form a hierarchy of some kind [20; 5; 8; 22; 20; 14].
3. limit the actions of the agents, e.g., via broadcasting environments [15].

I will talk about a new class of multi-agent systems, called PA-MAS (public-action multi-agent systems), in which agents have perfect-recall, no restrictions on their observations, but all their actions are observable to all the agents. Such a class, similarly to broadcasting environments [15], can provide more-or-less accurate models of many scenarios relevant to AI:

1. In community-card games such as Texas hold'em, each player is privately dealt some cards, which are combined with "community cards" that are dealt face up on the table. Moreover, all bidding is also public. Single rounds, or a bounded number of rounds, of such games can be modeled as PA-MAS. Such single rounds appear, for instance, as endgames and other simplified forms of Poker [17; 13; 9; 23]. The main limitation of our framework for giving a useful model of Poker is that we don't model probabilities which are central for optimal play in games like Poker.
2. There are epistemic puzzles in which all communication is public, e.g., the muddy children puzzle, the Russian cards puzzle, the consecutive numbers puzzle, and the sum-and-products puzzle [24].
3. In distributed systems one of the basic communication primitives is to broadcast a message to all other components. Rational distributed computing is a refinement of distributed computing in which agents have preferences rather than simple objectives [11]. For instance, in the classic m-out-of-n secret-sharing problem, there are n agents and initially each agent $i \in Ag$ privately holds a "share" f_i of a secret f_0, and any of the m "good" agents can collaborate to learn the secret in spite of the remaining $n - m$ "bad" agents. In a rational version of this scenario the preference of each agent is to learn the secret, i.e., she prefers to learn the secret rather than not to learn it, and a secondary preference is that the least number of other agents learn the secret. We can model such a secret-sharing scenario with broadcasting components.

I will talk about a new logic called SLi for reasoning about MAS. The logic SLi extends strategy logic (SL), which can already talk about complex notions like Nash equilibria and Pareto optimality since it has first-order variables x that vary over strategies. The extension is in two directions.

First, strategies in SLi are required to be uniform for the agents that use them (cf. [6]) and we use the subjective interpretation of strategic quantifiers (cf. [12]). That is, $\exists x \varphi$ is satisfied in a history h roughly means that there is a strategy σ, for every agent i that uses strategy x in the subformula φ, we have that (i) if h' is indistinguishable to h by agent i then $\sigma(h) = \sigma(h')$ (this is the uniformity condition) and (2) if h' is indistinguishable to h by agent i then φ is satisfied in the history h' (this is the subjective interpretation).

Second, SLi includes a symbol for equality between variables. Thus the formula $x = y$ tests if the strategies x and y are equal. This distinction is inspired by the distinction between first-order logic with equality and without equality.

*Based on joint work with Francesco Belardinelli, Alessio Lomuscio, and Aniello Murano [3; 4].

Moreover, it allows us to express complex properties such as the existence of an evolutionary stable strategy and uniqueness of Nash equilibria (a different extension of SL was introduced for reasoning about uniqueness of equilibria in [2]).

Although model-checking MAS against SLi specifications is undecidable, we have the following:

Theorem 1. *Model-checking PA-MAS in which all actions are public against SLi specifications is decidable.*

This theorem illustrates that reasoning about MAS becomes easier under the not unreasonable assumption that all actions of all agents are public.

References

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

[2] B. Aminof, V. Malvone, A. Murano, and S. Rubin. Graded modalities in strategy logic. *Inf. Comput.*, 261(Part):634–649, 2018.

[3] F. Belardinelli, A. Lomuscio, A. Murano, and S. Rubin. Verification of broadcasting multi-agent systems against an epistemic strategy logic. In C. Sierra, editor, *IJCAI'17*, pages 91–97. ijcai.org, 2017.

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

[5] D. Berwanger, A. B. Mathew, and M. van den Bogaard. Hierarchical information patterns and distributed strategy synthesis. In *ATVA15*, volume 9364 of *LNCS*, pages 378–393. Springer, 2015.

[6] P. Cermák, A. Lomuscio, F. Mogavero, and A. Murano. Practical verification of multi-agent systems against slk specifications. *Inf. Comput.*, 261(Part):588–614, 2018.

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

[8] B. Finkbeiner and S. Schewe. Uniform distributed synthesis. In *LICS05*, pages 321–330. IEEE Computer Society, 2005.

[9] L. Friedman. Optimal bluffing strategies in poker. *Management Science*, 17(12):B–764, 1971.

[10] J. Gutierrez, G. Perelli, and M. Wooldridge. Imperfect information in reactive modules games. In *KR'15*, pages 390–400, 2016.

[11] J. Halpern and V. Teague. Rational secret sharing and multiparty computation. In *STOC'04*, pages 623–632. ACM Press, 2004.

[12] W. Jamroga and W. van der Hoek. Agents that know how to play. *Fund. Inf.*, 62:1–35, 2004.

[13] H. W. Kuhn. A simplified two-person poker. *Contributions to the Theory of Games*, 1:97–103, 1950.

[14] O. Kupferman and M. Y. Vardi. Synthesizing distributed systems. In *LICS'01*, pages 389–398. IEEE Computer Society, 2001.

[15] R. v. Meyden and T. Wilke. Synthesis of distributed systems from knowledge-based specifications. In *CONCUR'05*, pages 562–576, 2005.

[16] F. Mogavero, A. Murano, G. Perelli, and M. Vardi. Reasoning about strategies: On the model-checking problem. *ACM Trans. Comp. Log.*, 15(4):34:1–34:47, 2014.

[17] J. Nash and L. Shapley. A simple three-person poker game. *Essays on Game Theory*, 1950.

[18] G. L. Peterson and J. H. Reif. Multiple-person alternation. In *FOCS79*, pages 348–363. IEEE Computer Society, 1979.

[19] A. Pnueli and R. Rosner. On the synthesis of a reactive module. In *POPL'89*, pages 179–190, 1989.

[20] A. Pnueli and R. Rosner. Distributed reactive systems are hard to synthesize. In *FOCS'90*, pages 746–757. IEEE Computer Society, 1990.

[21] J. H. Reif. The complexity of two-player games of incomplete information. *J. Comput. Syst. Sci.*, 29(2):274–301, 1984.

[22] S. Schewe and B. Finkbeiner. Distributed synthesis for alternating-time logics. In K. S. Namjoshi, T. Yoneda, T. Higashino, and Y. Okamura, editors, *ATVA'07*, pages 268–283. Springer, 2007.

[23] D. Szafron, R. G. Gibson, and N. R. Sturtevant. A parameterized family of equilibrium profiles for three-player kuhn poker. In *AAMAS'13*, pages 247–254. IFAAMAS, 2013.

[24] H. van Ditmarsch and B. Kooi. *One Hundred Prisoners and a Light Bulb*. Copernicus, 2015.