

Verification of Broadcasting Multi-Agent Systems against an Epistemic Strategy Logic

Francesco Belardinelli
Laboratoire IBISC, UEVE
and IRIT Toulouse
belardinelli@ibisc.fr

Alessio Lomuscio
Department of Computing
Imperial College London
a.lomuscio@imperial.ac.uk

Aniello Murano and Sasha Rubin
DIETI
Università degli Studi di Napoli
murano@na.infn.it
rubin@unina.it

Abstract

We study a class of synchronous, perfect-recall multi-agent systems with imperfect information and broadcasting, i.e., fully observable actions. We define an epistemic extension of strategy logic with incomplete information and the assumption of uniform and coherent strategies. In this setting, we prove that the model checking problem, and thus rational synthesis, is non-elementary decidable. We exemplify the applicability of the framework on a rational secret-sharing scenario. This paper has been accepted for publication in the proceedings of the 26th International Joint Conference on Artificial Intelligence (IJCAI 2017).

1 Introduction

In this contribution we introduce ESL, an epistemic extension of Strategy Logic (SL) [20; 9], based on synchronous perfect-recall strategies (Section 2). The language introduced can express rational synthesis [12; 21; 17], but its model-checking problem is undecidable. However, we identify a significant class BA-iCGS of systems: those having broadcast (i.e., fully observable) actions (Section 2) and prove that model checking BA-iCGS against ESL is non-elementary decidable (Section 4). This is a tight result as a matching lower-bound already holds in the perfect-information case. We illustrate our formalism on a rational secret-sharing scenario with broadcast actions (Section 3).

Related Work. SL and knowledge have been combined before in the context of MAS. In [7; 8], an epistemic variant of SL [20] was introduced. However, this was limited to epistemic sentences, whereas we consider the full combined language, and the approach assumed observational semantics, whereas we here consider synchronous perfect recall. Although not studied in these papers these formalisms have an undecidable model checking problem if evaluated under synchronous perfect recall. Also, [5] defines a variant of SL with uniform strategies. They achieve decidability by a variation of the tradition of assuming a hierarchy on the observations. In this paper we do not make any hierarchical assumptions.

A key aspect of the work here presented is that it relies on broadcasting to achieve decidability in the context of a very expressive specification language. The notion of broadcast

has already been studied in the context of knowledge [11; 18]. A further important result in this area is that for broadcast systems the synthesis problem of specifications in LTL and knowledge is decidable [19]. However, ESL is strictly more expressive and synthesis, which in our case can be expressed via model checking, can also be shown to be decidable. An approach to reasoning about strategies and knowledge under broadcast was also recently presented in [4]. However, their logic is considerably less expressive than ours, as it is based on ATL and not SL. In particular, it cannot express Nash equilibria and rational synthesis, which are essential features of this contribution.

Rational synthesis has been studied before in the context of perfect information. In [17] the strong-rational synthesis problem with LTL objectives (and aggregation of finitely many objectives), is shown to be 2EXPTIME-complete. In [13], Equilibrium Logic is introduced to reason about Nash equilibria in games with LTL and CTL objectives. However, both cases assume perfect information of the agents. Synthesis under imperfect information has been first tackled in [14] albeit for a restricted class of CGS, viz. *reactive modules*.

2 Strategy Logic with Imperfect Information

We consider concurrent game structures enriched with indistinguishability relations. [15; 6].

Definition 1 (iCGS). *An imperfect information concurrent game structure (iCGS) is a tuple $S = \langle Ag, AP, \{Act_a\}_{a \in Ag}, S, S_0, tr, \{\sim_a\}_{a \in Ag}, \lambda \rangle$, where:*

1. Ag is the finite non-empty set of agent names.
2. AP is the finite non-empty set of atomic propositions.
3. Act_a is the finite non-empty set of actions.
4. S is the finite non-empty set of states and $S_0 \subseteq S$ is the non-empty set of initial states.
5. $tr : S \times ACT \rightarrow S$ is the transition function, where $ACT = \prod_{a \in Ag} Act_a$ is the set of all joint actions.
6. $\sim_a \subseteq S^2$ is the indistinguishability relation for agent a , which is an equivalence relation.
7. $\lambda : AP \rightarrow 2^S$ is the labelling function that assigns to each atom p the set of states $\lambda(p)$ in which p holds.

We focus on a particular class of iCGS, those having broadcast actions only. This definition is reported from [4].

Definition 2 (BA-iCGS). An iCGS S only has broadcast actions if for every agent $a \in \text{Ag}$, states $s, s' \in S$, and joint actions $J, J' \in \text{ACT}$, if $J \neq J'$ and $s \sim_a s'$ then $\text{tr}(s, J) \not\sim_a \text{tr}(s', J')$. In this case we call S a broadcast iCGS. We write BA-iCGS for the set of broadcast iCGS.

Broadcast iCGS arise naturally in several MAS scenarios, including epistemic puzzles (e.g., the muddy children puzzle) and games (e.g., battleship). In Section 3 we discuss an application to rational synthesis.

We introduce ESL, an epistemic extension of SL. We interpret ESL on iCGS with history-based semantics. Fix a finite set of *atomic propositions (atoms)* AP , a finite set of *agents* Ag , and an infinite set Var of strategy variables x_0, x_1, \dots . The *formulas over* AP , Ag , and Var are built according to the following grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \wedge \varphi \mid \times\varphi \mid \varphi \cup \varphi \mid \langle\langle x \rangle\rangle\varphi \mid (x, a)\varphi \mid \mathbb{K}_a\varphi \mid \mathbb{C}_A\varphi \mid \mathbb{D}_A\varphi$$

where $p \in AP$, $x \in \text{Var}$, $a \in \text{Ag}$, and $A \subseteq \text{Ag}$. The set of ESL *formulas* is the one generated by the grammar. For a detailed account of the interpretation of ESL on iCGS we refer to [3]

We now introduce the main decision problem of this work.

Definition 3 (Model Checking). Let \mathcal{C} be a class of iCGS and \mathcal{F} a sublanguage of ESL. Model checking \mathcal{C} against \mathcal{F} specifications is the following decision problem: given $S \in \mathcal{C}$ and $\varphi \in \mathcal{F}$ as input, decide whether $S \models \varphi$.

Model checking iCGS against ATL is undecidable [10]. Thus, we obtain:

Proposition 1. Model checking iCGS against ESL is undecidable.

3 Rational Synthesis under Imperfect Information

Several questions in computer science can be cast as the problem of deciding if there exists a joint winning strategy for a coalition of agents against a coalition of adversarial agents (and computing one if it exists). In the verification literature this problem is called *synthesis*.

However, as argued in [21; 17; 1], the partition of agents into “good” and “bad” is often insufficient, and it is more appropriate to view agents as rational. That is, agents have preferences over outcomes and act in a way that increases their own utility. Then, instead of reasoning about winning strategies, one should reason about *rational* strategy profiles, i.e., that satisfy some notion of equilibrium. Application domains include rational distributed computing and rational cryptography [1], and negotiating systems with self-interested agents [16]. Technically, suppose we are given an iCGS S representing the multi-agent system, and LTLK-formulas γ_a representing the objective of agent $a \in \text{Ag}$. Here, LTLK is the logic consisting of the set of path-formulas of ATL^*K . We can then talk about Nash equilibria $\bar{\sigma}$ in games of the form $G = \langle S, \{\gamma_a\}_{a \in \text{Ag}} \rangle$. Rational synthesis considers the following decision problem (sometimes called *E-NASH*):

Definition 4 (Rational Synthesis for LTLK objectives, cf. [17]). Given an iCGS S , LTLK-formulas γ_a for every $a \in \text{Ag}$,

and an LTLK-formula φ , decide whether there exists a Nash equilibrium $\bar{\sigma}$ in the game $G = \langle S, \{\gamma_a\}_{a \in \text{Ag}} \rangle$ such that the path induced by $\bar{\sigma}$ satisfies φ .

The dual problem, called *Strong Rational Synthesis* (sometimes called *A-NASH*), concerns deciding whether all Nash equilibria induce a path that satisfies φ [17].

We now show that rational synthesis for LTLK objectives reduces to model checking against ESL. Suppose $\text{Ag} = \{a_1, a_2, \dots, a_n\}$, and let \bar{x} be an n -tuple of variables. The following formula $\text{RatSyn}_\varphi(\bar{x})$ in ESL expresses that \bar{x} is a Nash equilibrium whose induced execution satisfies φ : $(x_1, a_1) \dots (x_n, a_n) \left[\varphi \wedge \bigwedge_{a \in \text{Ag}} (\langle\langle y \rangle\rangle(y, a)\gamma_a \rightarrow \gamma_a) \right]$.

Lemma 1. Rational synthesis for LTLK objectives is reducible to model checking against the ESL-formula $\langle\langle x_1 \rangle\rangle \dots \langle\langle x_n \rangle\rangle \text{RatSyn}_\varphi(\bar{x})$.

A universally quantified formula is used for Strong Rational Synthesis. It is important to observe that ESL can express other equilibrium concepts such as subgame-perfect equilibria, concepts that capture deviations by groups of players such as k -resilience and t -immunity, and the combination (k, t) -robustness that captures fault-tolerance [1]. Also, ESL is able to express the existence of Nash equilibria w.r.t. epistemic objectives, which, to the best of our knowledge, has not yet been considered in the literature.

Rational Secret-Sharing with Broadcast. We illustrate the model-checking problem for BA-iCGS against ESL with a simple scenario inspired by [2] that uses broadcast. In the classic m -out-of- n secret-sharing problem, for $\text{Ag} = \{1, 2, \dots, n\}$, initially each agent $i \in \text{Ag}$ privately holds a “share” f_i of a secret f_0 , and any m “good” agents can collaborate to learn the secret in spite of the remaining $n - m$ “bad” agents. In the rational version of this scenario, the objective of each agent is to learn the secret, i.e., she prefers to learn the secret rather than not to learn it. Richer, non-binary, preferences can also be handled, including the fact that an agent may prefer that the least number of other agents learn the secret. For simplicity we do not consider such extensions.

We can model this scenario as an iCGS as follows. The secret is the value of a variable s initially hidden from all agents (formally, a variable v with finite domain D is modelled as $|D|$ -many atomic propositions); agent i 's share is modelled as a private variable f_i ; each agent has a private variable s_i that represents what she thinks the secret is; at every step, every agent broadcasts a message (from some fixed finite set of M messages). Finally, the objective γ_i of each agent i can be formalised as the LTLK-formula $\text{FG}\mathbb{K}_i(s_i = s)$: *from some point on, agent i knows the secret*. Thus, the ESL-formula $\langle\langle x_1 \rangle\rangle \dots \langle\langle x_n \rangle\rangle \text{RatSyn}_\varphi(\bar{x})$ expresses that there is a Nash equilibrium satisfying φ in the rational secret-sharing scenario. For instance, one can use φ to express that agents make “true” statements, e.g., that if agent i broadcasts “my share is x ”, then indeed $f_i = x$. Observe that by using ESL specifications we can naturally express secrecy and strategic concepts.

4 Model Checking BA-iCGS against ESL

In this section we prove the main technical result of this paper.

Theorem 5. *Model checking BA-iCGS against ESL specifications is decidable and non-elementary complete.*

For the non-elementary lower-bound we use the observation that model-checking SL on CGS (i.e., with perfect-information) is non-elementary [20], together with the fact that by encoding the last joint action into the states, one can translate a CGS S into a BA-iCGS S' such that for all sentences φ in ESL, we have that $S \models \varphi$ iff $S' \models \varphi$ [4].

For the non-elementary upper-bound, we reduce the model-checking problem of BA-iCGS against ESL specifications to model checking regular-trees against Monadic Second-Order Logic (MSO). The naïve approach is to code every tuple (S, h, χ) by a tuple of functions $(\widehat{S}, \widehat{h}, \widehat{\chi})$ each of whose domain is the set ACT^* of finite sequences of joint actions, and whose ranges are finite (to be specified later). This encoding allows us to build, for every ESL-sentence φ , an MSO-formula Φ , such that $(S, h, \chi) \models \varphi$ iff $T \models \Phi(\widehat{S}, \widehat{h}, \widehat{\chi})$, where T is the infinite ACT-ary tree generated by \widehat{S}, \widehat{h} , and $\widehat{\chi}$. The latter problem is decidable if \widehat{S}, \widehat{h} , and $\widehat{\chi}$ are regular functions (a function $f : D^* \rightarrow L$ is *regular* if, for each $l \in L$, the set $f^{-1}(l) \subseteq D^*$ is accepted by a finite automaton). Since φ is a sentence we can choose χ arbitrarily, in particular so that it is regular (on the other hand, both \widehat{S} and \widehat{h} are always regular).

Application to Rational Synthesis. By the discussion in Section 4, we immediately get the first part of the following:

Corollary 1. *Rational synthesis for LTLK objectives on BA-iCGS is decidable. Moreover, if a given instance returns “yes”, then a finite-state Nash equilibrium can be computed.*

5 Conclusions

We defined ESL, a combination of Strategy Logic and Epistemic Logic. We observed that model checking and synthesis are undecidable under synchronous perfect-recall semantics. However, we showed that a noteworthy subclass of systems, those that admit only broadcast actions, admit decidable model checking and synthesis, and identified tight bounds for the model-checking problem.

We have illustrated the expressivity of the formalism by phrasing rational synthesis under incomplete information, a previously unexplored set-up, as an instance of model checking for ESL. This has the noteworthy consequence that rational synthesis is decidable in the framework. It follows that we can decide expressive strategic properties of rational secret-sharing scenarios like the one presented in Section 3 under the assumption of non-randomised strategies. We leave the exploration of other scenarios for future work.

Acknowledgements. This research was partly supported by EPSRC (grant EP/I00529X), INdAM (grant “Logica e Automi per il Model Checking”), the ANR JCJC (project SVeDaS) and an INdAM Marie Curie fellowship to S. Rubin. The authors thank Benjamin Aminof for fruitful discussions.

References

- [1] I. Abraham, L. Alvisi, and J.Y. Halpern. Distributed computing meets game theory: combining insights from two fields. *SIGACT News*, 42(2):69–76, 2011.
- [2] I. Abraham, D. Dolev, R. Gonen, and J. Y. Halpern. Distributed computing meets game theory. In *PODC'06*, pages 53–62, 2006.
- [3] F. Belardinelli, A. Lomuscio, A. Murano, and S. Rubin. Verification of broadcasting multi-agent systems against an epistemic strategy logic. In *IJCAI'17*, 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] R. Berthon, B. Maubert, A. Murano, S. Rubin, and M. Vardi. Hierarchical strategic reasoning. In *LICS'17*, 2017.
- [6] N. Bulling and W. Jamroga. Comparing variants of strategic ability: how uncertainty and memory influence general properties of games. *JAAMAS*, 28(3):474–518, 2014.
- [7] P. Cermák. A model checker for strategy logic. Master’s thesis, Dept. of Comp., Imperial, 2014.
- [8] P. Čermák, A. Lomuscio, F. Mogavero, and A. Murano. MCMAS-SLK: A model checker for the verification of strategy logic specifications. In *CAV'14*, pages 524–531, 2014.
- [9] K. Chatterjee, T.A. Henzinger, and N. Piterman. Strategy Logic. *IC*, 208(6):677–693, 2010.
- [10] C. Dima and F.L. Tiplea. Model-checking ATL under imperfect information and perfect recall semantics is undecidable. *CoRR*, abs/1102.4225, 2011.
- [11] R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi. *Reasoning about Knowledge*. MIT, 1995.
- [12] D. Fisman, O. Kupferman, and Y. Lustig. Rational synthesis. In *TACAS'10*, pages 190–204, 2010.
- [13] J. Gutierrez, P. Harrenstein, and M. Wooldridge. Reasoning about equilibria in game-like concurrent systems. *Ann. Pure Appl. Log.*, 168(2):373–403, 2017.
- [14] J. Gutierrez, G. Perelli, and M. Wooldridge. Imperfect information in reactive modules games. In *KR'15*, pages 390–400, 2016.
- [15] W. Jamroga and W. van der Hoek. Agents that know how to play. *Fund. Inf.*, 62:1–35, 2004.
- [16] N. Jennings et al. Automated negotiation: Prospects, methods, and challenges. *Int. J. of Gr. Dec and Neg.*, 10(199-215):199–215, 2001.
- [17] O. Kupferman, G. Perelli, and M.Y. Vardi. Synthesis with rational environments. *Ann. Math. Art. Int.*, 78(1):3–20, 2016.
- [18] A. Lomuscio, R. van der Meyden, and M. Ryan. Knowledge in multi-agent systems. *ACM Trans. Comp. Log.*, 1(2):246–282, 2000.
- [19] R. van der Meyden and T. Wilke. Synthesis of distributed systems from knowledge-based specifications. In *CONCUR'05*, pages 562–576, 2005.
- [20] F. Mogavero, A. Murano, G. Perelli, and M.Y. Vardi. Reasoning about strategies. *ACM Trans. Comp. Log.*, 15(4):34:1–34:47, 2014.
- [21] M. Wooldridge, J. Gutierrez, P. Harrenstein, E. Marchioni, G. Perelli, and A. Toumi. Rational verification. In *AAAI'16*, pages 4184–4191, 2016.