Recent Results on Strategy Logic with Imperfect Information (short paper)*

Bastien Maubert

Universit degli studi di Napoli “Federico II”, Napoli, Italy

Abstract. In this communication we present recent advances in the field of logics for strategic reasoning, and more precisely about the popular Strategy Logic (SL) in the context of systems with imperfect information. As a consequence of Reif's seminal result on multiplayer games with imperfect information, model checking SL is undecidable when agents observe the game imperfectly but remember everything they observe. However, in the case of multiplayer games or distributed synthesis la Pnueli and Rosner, some additional hypothesis are known to bring back decidability. The main ones consist in assuming that information is hierarchical among agents, or that actions are public. We report recent results that, under similar assumptions, establish the decidability of the model-checking problem for different variants and extensions of SL.

Keywords: Strategy Logic · Imperfect Information · Knowledge · Pushdown systems · Model checking.

1 Introduction

Since the introduction of Alternating-time Temporal Logic (ATL) by Alur, Henzinger and Kupferman [1] to reason about strategic abilities of agents and coalitions in multi-agent systems, research on the topic has thrived to develop more expressive formalisms and solve related algorithmic problems on different classes of systems. One important contribution has been the introduction of Strategy Logic (SL) [9, 18], a very expressive logic that extends ATL and gives the possibility to express a number of important game-theoretic concepts. In particular, equilibria such as Nash equilibria, subgame-perfect equilibria or core equilibria can all be expressed very naturally in SL.

Since multi-agent systems often display some kind of imperfect information due to agents having only partial view of their common environment and/or of other agents' state, variants of both ATL and SL have been introduced to model and reason about systems with imperfect information. An important challenge is that, while imperfect-information games are decidable for two players [25], they become undecidable when more players are involved [20], which is typically the case in multi-agent systems. However it is known that assuming a total order on the agents where "higher" agents know more than "lower" ones, a situation known as *hierarchical information*, can bring back decidability of a number of related strategic problems [22, 12, 19, 11]. Similarly, variants around the idea that all new information must be observed similarly by all agents also yield decidability results [17, 24, 7].

Recently these ideas have been applied to logics for strategic reasoning, and in particular Strategy Logic, leading to a number of decidability results, the first ones obtained for such logics in the context of agents with imperfect information and memory (alternative *memoryless* semantics exist in which agents have no memory of the past and are restricted to using only positional strategies, which usually also leads to decidable model-checking problems on finite systems). We briefly present several such results, obtained in collaboration with a number of co-authors over the last few years.

2 Strategy Logic

Before considering extensions with imperfect information, we first recall usual SL with perfect information.

Let us fix a finite set of *atomic propositions* AP, a finite set of *agents* or *players* Ag and a finite set of *variables* Var.

Definition 1. *The syntax of classic* SL *is defined by the following grammar:*

$$\varphi := p \mid \neg\varphi \mid \varphi \vee \varphi \mid \mathbf{X}\varphi \mid \varphi\mathbf{U}\varphi \mid \exists s.\varphi \mid (a,s)\varphi$$

where $p \in \mathrm{AP}$*,* $s \in \mathrm{Var}$*, and* $a \in \mathrm{Ag}$*.*

Boolean operators and temporal operators, $\mathbf{X}$ ("next") and $\mathbf{U}$ ("until"), have the usual meaning. The *strategy quantifier* $\exists s$ is a first-order-like quantification on strategies: $\exists s.\varphi$ reads as "there exists a strategy s such that φ holds", where s is a strategy variable. The *binding operator* (a,s) assigns a strategy to an agent: $(a,s)\varphi$ reads as "when agent a plays strategy s, φ holds".

The models of SL are classic concurrent game structures.

Definition 2 (CGS). *A* concurrent game structure *(or CGS for short) is a tuple* $\mathcal{G} = (\mathrm{Ac}, V, \Delta, \ell, v_\iota)$ *where* Ac *is a finite non-empty set of* actions*,* V *is a finite non-empty set of* positions*,* $\Delta : V \times \mathrm{Ac}^{\mathrm{Ag}} \to V$ *is a* transition function*,* $\ell : V \to 2^{\mathrm{AP}}$ *is a* labelling function*, and* $v_\iota \in V$ *is an* initial position*.*

In a position $v \in V$, each player a chooses an action $\alpha_a \in \mathrm{Ac}$ and the game proceeds to position $\Delta(v, \boldsymbol{\alpha})$, where $\boldsymbol{\alpha} \in \mathrm{Ac}^{\mathrm{Ag}}$ stands for the *joint action* $(\alpha_a)_{a\in\mathrm{Ag}}$. A *play* is an infinite sequence of positions and joint actions $\lambda = v_0\boldsymbol{\alpha}_1 v_1 \ldots$ that respect the transition relation, and a history is a finite prefix of a play. We let λ_i denote v_i, and Hist be the set of histories. A *strategy* is a function $\sigma : \mathrm{Hist} \to \mathrm{Ac}$, and we let *Str* be the set of all strategies.

An *assignment* $\chi : \mathrm{Ag} \cup \mathrm{Var} \to \mathit{Str}$ is a function assigning strategies to agents and variables. A history λ is *consistent* with an assignment χ if, at each step of λ, all agents follow the strategy assigned to them by χ.

The semantics of SL formulas is defined given a CGS $\mathcal{G}$, an assignment χ and a history λ. The semantics for Boolean and temporal connectives being standard, we just give the semantics of the strategy quantifier and the binding operator, and refer the reader to [18] for more details. A formula $\exists s.\varphi$ is true in $\mathcal{G}, \chi, \lambda$, written $\mathcal{G}, \chi, \lambda \models \exists s.\varphi$, if there exists a strategy σ such that $\mathcal{G}, \chi[s \mapsto \sigma], \lambda \models \varphi$, and $\mathcal{G}, \chi, \lambda \models (a, s)\varphi$ if $\mathcal{G}, \chi[a \mapsto \chi(s)], \lambda \models \varphi$.

For instance, if each agent $a_i \in \mathrm{Ag} = \{a_1, \ldots, a_n\}$ has an LTL objective ψ_i, the following SL formula expresses the existence of a Nash equilibrium:

$$\exists s_1 \ldots \exists s_n (a_1, s_1) \ldots (a_n, s_n) \bigwedge_i \forall s_i' \left[((a_i, s_i')\psi_i) \to \psi_i\right]$$

3 SL with imperfect information

Imperfect information extensions of SL have been considered for instance in [2, 8], considering purely semantical aspects or focusing on agents without memory. The first decidability result for an imperfect-information extension of SL for agents with perfect recall was published in [5] (see [6] for an extended version).

In games with imperfect information, the fact that the lack of knowledge reduces players' power is captured by the notion of *uniform strategies*: players can only use strategies that assign the same move to situations that they cannot distinguish. In SL, a problem when introducing imperfect information is to know for which player(s) a strategy should be uniform: indeed, quantification on strategies and assignment of a strategy to an agent are dissociated, and in addition a strategy can be assigned to different agents. The solution adopted in [5] was to dissociate players and observational power: each strategy quantifier $\exists^o s$ is quantified with an observation symbol o, interpreted in the model by an equivalence relation on positions: we specify directly with respect to which observation the quantified strategy should be uniform.

We thus consider a finite set of *observation symbols* Obs, and extend concurrent game structures with a component called *observation interpretation*: a CGS is now a tuple $\mathcal{G} = (\mathrm{Ac}, V, \Delta, \ell, v_\iota, \mathcal{O})$, where $(\mathrm{Ac}, V, \Delta, \ell, v_\iota)$ is as before, and $\mathcal{O} : \mathrm{Obs} \to 2^{V \times V}$ maps each observation symbol $o \in \mathrm{Obs}$ to an equivalence relation $\mathrm{Obs}(o) = \sim_o$ over the positions, that represents indistinguishability of positions. To capture agents with synchronous perfect recall, we then extend these relations to histories by letting $\lambda \sim_o \lambda'$ if $|\lambda| = |\lambda'|$ and $\lambda_i \sim_o \lambda_i'$ for every $i \in \{0, \ldots, |\lambda| - 1\}$. We then say that a strategy σ is *o-uniform* if for all histories λ, λ' such that $\lambda \sim_o \lambda'$, $\sigma(\lambda) = \sigma(\lambda')$.

In the resulting logic, called $\mathsf{SL}_{\mathrm{iR}}$, we can express for instance the problem of distributed synthesis for two players a_1 and a_2 against an opponent a_3. Assume that agent a_i observes the game through observation relation $\sim_i$, and that agents a_1, a_2 aim at reaching a set of positions $\mathrm{Reach} \subseteq V$. Define $\mathrm{Obs}(o_i) = \sim_i$, and assume that positions in Reach are labeled with proposition p_{win}. Then the existence of winning strategies for a_1, a_2 is expressed in $\mathsf{SL}_{\mathrm{iR}}$ as follows:

$$\varphi = \exists^{o_1} s_1 \, \exists^{o_2} s_2 \, \forall^{o_3} s_3 \, (a_1, s_1)(a_2, s_2)(a_3, s_3)\mathbf{F} p_{\mathrm{win}}$$

As a direct consequence of Peterson and Reif's seminal result on the undecidability of multiplayer games with imperfect information [21], model checking $\mathsf{SL}_{\mathrm{iR}}$ is undecidable already for the single formula φ. Our main result in [5] is that it is decidable on *hierarchical instances*, i.e., on the set of inputs where deeper quantifiers (in the syntactic tree) are parameterized with finer observation relations. For instance, if in CGS $\mathcal{G}$ we have $\mathrm{Obs}(o_3) \subseteq \mathrm{Obs}(o_2) \subseteq \mathrm{Obs}(o_1)$, meaning that agent a_3 observes better than agent a_2 who in turn observes better than agent a_1, then $(\mathcal{G}, \varphi)$ is a hierarchical instance.

Theorem 1. *Model checking* $\mathsf{SL}_{\mathrm{iR}}$ *on hierarchical instances is decidable.*

The proof goes through $\mathsf{QCTL}^*_{\mathrm{iR}}$, an imperfect-information extension of Quantified CTL^* (QCTL^*) that we introduced to serve as a "compilation" logic, an intermediary between $\mathsf{SL}_{\mathrm{iR}}$ and the tree automata machinery used to solve the model-checking problem for $\mathsf{QCTL}^*_{\mathrm{iR}}$. The complexity of the model-checking problem is nonelementary: essentially we gain one more exponential for every alternation between existential and universal strategy quantifiers, and at every change of observation relation.

While Theorem 1 considers finite concurrent game structures, we recently extended this result to a class of infinite systems obtained as unfoldings of pushdown systems. Informally a pushdown CGS, or PGS, is a CGS equipped with a stack. Transitions depend both on the state (position) and the top symbol of the stack, and each transition determines both the new state and an operation to execute on the stack (essentially, either "pop" the top symbol, or "push" one or several symbols). The set of possible configurations in a PGS is thus infinite. We proved in [16] that, as long as the stack is perfectly observed by all agents, then the result from Theorem 1 can be extended to pushdown concurrent game structures.

Theorem 2. *Model checking* $\mathsf{SL}_{\mathrm{iR}}$ *on pushdown CGS with visible stack is decidable for hierarchical instances.*

Actually we proved a stronger result: the problem remains decidable even if we consider the much more general class of *higher-order pushdown systems*, always as long as the (higher-order) stack is visible.

We now move to extensions of SL with imperfect information *and knowledge operators.*

4 SL with knowledge operators

In systems with imperfect information it becomes relevant and useful to be able to express and reason about what agents know or not about the state of the system, about what other agents know, and about each other's strategic abilities. Epistemic extensions of logics for strategic reasoning have thus been studied, for instance in ATL [14, 13] or SL [2]. Typically, such logics add to the language epistemic operators, the most important one being the knowledge operator $K_a\varphi$, which reads as "agent a knows that φ holds" [10]. However, since such logics usually also require agents to play according to their knowledge (uniform strategies), they are undecidable for agents with perfect recall. A first decidability result for SL with imperfect information and epistemic operators for agents with perfect recall was established in [4] in the case where all actions are public.

In [15] we established a second decidability result for such an epistemic extension of Strategy Logic with perfect recall, in the context of hierarchical information. This result extends the one in [5] by adding epistemic operators to $\mathsf{SL}_{\mathrm{iR}}$ and extending the automata construction to deal with these.

In this work we pointed out an important subtlety that arises when studying together knowledge and strategic ability. We noted that two different semantics are found in the literature, usually without mentioning that a choice is made and what this choice represents. Informally, one semantics captures agents that know everyone's strategy, while the other represents the situation where agents do not know anyone's strategy. We call them respectively the *informed semantics* and the *uninformed semantics*. These two semantics are different: as one would expect, knowing everyone's strategy makes it possible to infer more from one's observations, and thus under the informed semantics agents know more facts about the system. The choice of semantics also impacts the complexity of the model-checking problem: on the one hand it is proved in [17], which implicitly uses the informed semantics, that the distributed synthesis for LTL with knowledge is undecidable already for systems with hierarchical information. On the other hand, Puchala proves in [23] that the same problem with uninformed semantics is decidable when information is hierarchical.

Our main result in [15] generalizes the latter, by showing that model checking ESL (the extension of $\mathsf{SL}_{\mathrm{iR}}$ with knowledge operators) is decidable on hierarchical instances. Here an instance is said to be hierarchical if, as before, deeper quantifiers have finer observations, and in addition epistemic formulas do not refer to strategies that are quantified before the knowledge operator. This is satisfied if epistemic formulas do not talk about the future, or quantify on strategies for all agents before doing so.

Theorem 3. *Model checking* ESL *with uninformed semantics is decidable on hierarchical instances.*

Our last result in this line of work is a first decidability result for SL with knowledge and informed semantics. Defining formally this informed semantics for ESL is not easy, but in [3] we provide such a definition for ESL with Boolean Goals (ESL[BG]), an important fragment of ESL; we study in detail the difference between informed and uninformed semantics, and fix an inconsistency present in the semantics of previous epistemic extensions of SL. We also prove that model checking ESL[BG] is decidable on systems where all actions are public.

Theorem 4. *Model checking* ESL[BG] *with informed semantics is decidable on systems with public actions.*

Acknowledgment. We thank our co-authors on the publications mentioned in this communication: Francesco Belardinelli, Raphaël Berthon, Sophia Knight, Alessio Lomuscio, Aniello Murano, Olivier Serre, Sasha Rubin and Moshe Vardi.

References

1. Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time temporal logic. Journal of the ACM **49**(5), 672–713 (2002)
2. Belardinelli, F.: A logic of knowledge and strategies with imperfect information. In: LAMAS'15. pp. 1–15 (2015)
3. Belardinelli, F., Knight, S., Lomuscio, A., Maubert, B., Murano, A., Rubin, S.: Reasoning about agents that may know other agents' strategies. In: Proceedings of the 30th International Joint Conference on Artificial Intelligence. IJCAI/AAAI Press (2021)
4. Belardinelli, F., Lomuscio, A., Murano, A., Rubin, S.: Verification of broadcasting multi-agent systems against an epistemic strategy logic. In: Proceedings of the 26th International Joint Conference on Artificial Intelligence. vol. 17, pp. 91–97 (2017)
5. Berthon, R., Maubert, B., Murano, A., Rubin, S., Vardi, M.Y.: Strategy logic with imperfect information. In: Proceedings of the 32nd Annual ACM/IEEE Symposium on Logic in Computer Science. pp. 1–12. IEEE Computer Society (2017)
6. Berthon, R., Maubert, B., Murano, A., Rubin, S., Vardi, M.Y.: Strategy logic with imperfect information. ACM Transactions on Computational Logic (TOCL) **22**(1), 1–51 (2021)
7. Bouyer, P.: Games on graphs with a public signal monitoring. In: FOSSACS'18. pp. 530–547. Springer (2018), http://arxiv.org/abs/1710.07163
8. Cermák, P., Lomuscio, A., Mogavero, F., Murano, A.: MCMAS-SLK: A model checker for the verification of strategy logic specifications. In: CAV'14. pp. 525–532 (2014)
9. Chatterjee, K., Henzinger, T.A., Piterman, N.: Strategy logic. Information and Computation **208** (2010)
10. Fagin, R., Halpern, J.Y., Moses, Y., Vardi, M.Y.: Reasoning about knowledge, vol. 4. MIT press Cambridge (1995)
11. Finkbeiner, B., Schewe, S.: Uniform distributed synthesis. In: Proceedings of the 20th IEEE Symposium on Logic in Computer Science. pp. 321–330. IEEE Computer Society (2005)
12. Gastin, P., Sznajder, N., Zeitoun, M.: Distributed synthesis for well-connected architectures. Form. Method Syst. Des. **34**(3), 215–237 (2009). https://doi.org/10.1007/s10703-008-0064-7
13. Jamroga, W., Ågotnes, T.: What agents can achieve under incomplete information. In: Proceedings of the fifth international joint conference on Autonomous agents and multiagent systems. pp. 232–234. ACM (2006)
14. Jamroga, W., van der Hoek, W.: Agents that know how to play. Fundam. Inform. **63**(2-3), 185–219 (2004)
15. Maubert, B., Murano, A.: Reasoning about knowledge and strategies under hierarchical information. In: KR'18. pp. 530–540 (2018), https://aaai.org/ocs/index.php/KR/KR18/paper/view/17996
16. Maubert, B., Murano, A., Serre, O.: Reasoning about strategies on collapsible pushdown arenas with imperfect information (2020)

17. van der Meyden, R., Wilke, T.: Synthesis of distributed systems from knowledge-based specifications. In: CONCUR'05. pp. 562–576 (2005)
18. Mogavero, F., Murano, A., Perelli, G., Vardi, M.Y.: Reasoning about strategies: On the model-checking problem. ACM Transactions on Computational Logic **15**(4), 34:1–34:47 (2014)
19. Peterson, G., Reif, J., Azhar, S.: Decision algorithms for multiplayer noncooperative games of incomplete information. Computers & Mathematics with Applications **43**(1), 179–206 (2002)
20. Peterson, G., Reif, J.H., Azhar, S.: Lower bounds for multiplayer noncooperative games of incomplete information. Computers & Mathematics with Applications **41**(7), 957–992 (2001)
21. Peterson, G.L., Reif, J.H.: Multiple-person alternation. In: Proceedings of the 20th Annual Symposium on Foundations of Computer Science. pp. 348–363. IEEE Computer Society (1979)
22. Pnueli, A., Rosner, R.: Distributed reactive systems are hard to synthesize. In: Proceedings of the 31st Annual Symposium on Foundations of Computer Science. pp. 746–757. IEEE Computer Society (1990)
23. Puchala, B.: Asynchronous omega-regular games with partial information. In: MFCS'10. pp. 592–603 (2010)
24. Ramanujam, R., Simon, S.: A communication based model for games of imperfect information. In: CONCUR'10. pp. 509–523. Springer (2010)
25. Reif, J.H.: The complexity of two-player games of incomplete information. Journal of computer and system sciences **29**(2), 274–301 (1984). https://doi.org/10.1016/0022-0000(84)90034-5