Reasoning about agents that may know other agents' strategies

Francesco Belardinelli[1] , **Sophia Knight**[2] , **Alessio Lomuscio**[1] ,
Bastien Maubert[3] , **Aniello Murano**[3] and **Sasha Rubin**[4]
[1]Imperial College London, UK
[2]University of Minnesota Duluth, USA
[3]Università degli Studi di Napoli "Federico II", Italy
[4]University of Sydney, Australia

Abstract

We study the semantics of knowledge in strategic reasoning. Most existing works either implicitly assume that agents do not know one another's strategies, or that all strategies are known to all; and some works present inconsistent mixes of both features. We put forward a novel semantics for Strategy Logic with Knowledge that cleanly models whose strategies each agent knows. We study how adopting this semantics impacts agents' knowledge and strategic ability, as well as the complexity of the model-checking problem.

1 Introduction

Reasoning about strategic abilities of agents and groups of agents is a topic of great interest in the formal study of multi-agent systems. To that end many logics have been introduced, such as Game Logic [Pauly and Parikh, 2003], STIT [Herzig and Lorini, 2010], and ATL [Alur *et al.*, 2002].

Among these logics, Strategy Logic (SL) [Chatterjee *et al.*, 2010; Mogavero *et al.*, 2014] has proved very successful in recent years, as it combines a natural syntax with high expressive power that allow it to easily express complex game-theoretic notions such as existence of Nash equilibria. Its model checking problem is decidable, and model-checking algorithms can usually be used to synthesize strategies that satisfy a given specification. Because imperfect information is a crucial aspect of multi-agent systems, SL has recently been extended to take into account imperfect information [Berthon *et al.*, 2021] and allow for epistemic reasoning [Maubert and Murano, 2018; Belardinelli *et al.*, 2020].

As discussed in [Maubert and Murano, 2018], defining the semantics of knowledge in a strategic context involves subtleties with far-reaching consequences, that have often gone unheeded. Two semantics exist in the literature, and most works adopt one or the other without mentioning it. One corresponds to agents that do not know each other's strategies (called *uninformed semantics* in [Maubert and Murano, 2018]), while the other corresponds to agents who know everyone's strategies (*informed semantics*). The former is used in all existing epistemic extensions of ATL and SL (e.g., [van der Hoek and Wooldridge, 2003; Jamroga and van der Hoek, 2004; Guelev *et al.*, 2011; Belardinelli *et al.*, 2017b]), while the latter is adopted in works on distributed synthesis [van der Meyden and Wilke, 2005] and knowledge-based policies in POMDPs [Saffidine *et al.*, 2018].

Intuitively, the uninformed semantics corresponds to situations where agents do not collaborate, or cannot communicate information about their strategies, and have no a priori knowledge of how other agents behave. On the other hand, the informed semantics corresponds to situations where the strategies followed by each agent are common knowledge. Knowing others' strategies gives additional power to agents, as it allows them to infer more information than what they learn from their observations alone. One can think for instance of the Hanabi card game, where by knowing how other players choose their actions, one can interpret them and infer valuable information. However, this refined notion of knowledge comes at a price in terms of computational complexity.

Indeed some problems are decidable with the uninformed semantics and undecidable with the informed one. This is the case for distributed synthesis. This central problem in multi-agent systems is, in general, undecidable, already for simple reachability objectives without epistemic content [Peterson *et al.*, 2001]. But it becomes decidable if there is a hierarchy of information among agents, or when all actions are observed similarly by all agents (see [Berthon *et al.*, 2021] for a discussion). When objectives are no longer purely temporal but can also refer to knowledge, distributed synthesis for systems with hierarchical information remains decidable with the uninformed semantics of knowledge [Puchala, 2010], but it becomes undecidable for the informed semantics [van der Meyden and Wilke, 2005]. This gives even more reason to consider carefully, when studying a system, the assumptions made on the knowledge agents have of each other's strategies.

Part of these results have been generalized to logics for strategic reasoning that capture distributed synthesis, but only for the uninformed semantics. Decidability has been generalized to an epistemic extension of SL in the case of hierarchical information [Maubert and Murano, 2018], and decidability in the case of public actions has been generalized to epistemic extensions of ATL [Belardinelli *et al.*, 2017b] and SL [Belardinelli *et al.*, 2020]. However no result exists for the informed semantics which, as far as we know, has never been formalized in the context of logics for strategic reasoning.

Contribution. We address this issue and propose a novel semantics of SL with imperfect information and epistemic

operators (SLK) in which one can specify, for each agent, the set of agents whose strategy she is informed of. This allows us to capture the case where all agents know everyone's strategies (informed semantics) and the case where no one knows anyone else's strategy (uninformed semantics), as well as intermediate situations (for instance, some agents may be allowed to privately communicate their strategies). We illustrate in detail the impact that knowing others' strategies has on the agents' knowledge as well as on their strategizing power. Doing so, we also shed light on an inconsistency in existing semantics of SL with knowledge, which is solved by our semantics. We then show that for the new semantics, model checking SLK is already undecidable under hierarchical information, but is decidable when all actions are public.

2 SLK with informed semantics

Fix AP a finite set of *atomic propositions*, Ag a finite set of *agents*, and Var a finite set of *variables*.

2.1 Syntax

We consider the flat fragment of SL with Boolean Goals [Mogavero *et al.*, 2014] extended with knowledge operators. This simple syntax allows us to focus on the most important aspects of the new semantics, avoiding technical issues that would arise with the full syntax of SL and diminish the clarity of the exposition (see Remark 2 for more detail).

Definition 1. The syntax of $\mathsf{SLK}_{\text{inf}}[\text{BG}]$ is given by the following grammar:

$$\Phi = \exists s.\Phi \mid \neg\Phi \mid \Phi \vee \Phi \mid \theta$$
$$\theta = (\text{Ag}, \bar{s})\varphi \mid \neg\theta \mid \theta \vee \theta$$
$$\varphi = p \mid \neg\varphi \mid \varphi \vee \varphi \mid X\varphi \mid \varphi U \varphi \mid K_a\varphi$$

where $p \in \text{AP}$, $s \in \text{Var}$, $\bar{s} \in \text{Ag} \to \text{Var} \times \{\langle\rangle, []\}$ and $a \in \text{Ag}$.

X and U are the temporal operators "next" and "until", $\exists s$ is the existential quantifier on strategies, and K_a is the individual knowledge operator for agent a. The group binding $(\text{Ag}, \bar{s})$ assigns to each agent a strategy variable and a symbol indicating how the uniformity constraint should be treated. Intuitively, $\langle\rangle$ (resp. $[]$) indicates a strategy that was quantified existentially (resp. universally). We may write a group binding as a sequence of individual bindings $\langle a_1, s_1\rangle \dots [a_n, s_n]$, where symbols $\langle\rangle$ and $[]$ are split and used as delimiters.

We use abbreviations $\top = p \vee \neg p$, $\bot = \neg\top$ and $\varphi \to \varphi' = \neg\varphi \vee \varphi'$ for boolean connectives, $F\varphi = \top U \varphi$, $G\varphi = \neg F \neg\varphi$ for temporal operators, and finally $\forall s.\varphi = \neg\exists s.\neg\varphi$.

A variable s appears *free* in a formula φ if it appears out of the scope of a strategy quantifier, and an agent a appears free in φ if a temporal operator (either X or U) appears in φ out of the scope of any binding (in which case all agents are free, since group bindings treat all agents at once). We let *free* (φ) be the set of variables and players that appear free in φ. If *free* (φ) is empty, φ is a *sentence*.

2.2 Models

We present concurrent game structures with imperfect information on both positions and actions. Towards our definition of the informed semantics, we specify for each agent a a set of agents $\mathcal{A}_a$ whose strategies agent a is informed of.

Definition 2. A *concurrent game structure* (CGS) is a tuple $\mathcal{G} = (\text{Ac}, V, E, \ell, V_\iota, \{\sim_a\}_{a\in\text{Ag}}, \{\mathcal{A}_a\}_{a\in\text{Ag}})$ where Ac is a finite set of *actions*, V is a finite set of *positions*, $E : V \times \text{Ac}^{\text{Ag}} \to V$ is a *transition function*, $\ell : V \to 2^{\text{AP}}$ is a *valuation function*, $V_\iota \subseteq V$ is a set of *initial positions*, $\sim_a \subseteq (V \times V) \cup (\text{Ac} \times \text{Ac})$ is an equivalence relation called an *observation relation*, and $\mathcal{A}_a \subseteq \text{Ag}$ is the set of agents whose strategies agent a is *informed* of.

As in [Puchala, 2010] we consider equivalence relations on positions *and actions*. This, together with the inclusion of actions in histories (see below), allows us to model public actions in a very natural way.

Joint actions. In a position $v \in V$, each player a chooses an action $\alpha^a \in \text{Ac}$ and the game proceeds to position $E(v, \boldsymbol{\alpha})$, where $\boldsymbol{\alpha} \in \text{Ac}^{\text{Ag}}$ stands for the *joint action* $(\alpha^a)_{a\in\text{Ag}}$. Given a joint action $\boldsymbol{\alpha} = (\alpha^a)_{a\in\text{Ag}}$ and $a \in \text{Ag}$, we let $\boldsymbol{\alpha}^a = \alpha^a$.

Histories, plays and strategies. A *history* (resp. *play*) is a finite (resp. infinite) sequence of positions and joint actions $\rho = v_0\boldsymbol{\alpha}_1 v_1 \dots \boldsymbol{\alpha}_n v_n$ (resp. $\pi = v_0\boldsymbol{\alpha}_1 v_1 \dots$) such that $v_0 \in V_\iota$ and for all i with $0 \le i < n$ (resp. $i \ge 0$), $E(v_i, \boldsymbol{\alpha}_{i+1}) = v_{i+1}$. For a history $\rho = v_0\boldsymbol{\alpha}_1 v_1 \dots \boldsymbol{\alpha}_n v_n$ we let $\text{last}(\rho) = v_n$ and $|\rho| = n$. For a play $\pi = v_0\boldsymbol{\alpha}_1 v_1 \dots$ and $i \ge 0$, we write $\pi_{\le i} = v_0 \dots \boldsymbol{\alpha}_i v_i$ for its prefix history ending at position i. We define $\rho_{\le i}$ similarly for histories ρ when $i \le |\rho|$. We let Hist be the set of histories. A *strategy* is a function $\sigma : \text{Hist} \to \text{Ac}$, and we let *Str* be the set of all strategies.

Assignments. An *assignment* $\chi : \text{Ag} \cup \text{Var} \rightharpoonup \mathit{Str}$ is a partial function assigning strategies to agents and variables. Given a tuple of strategies $\bar{\sigma} \in \mathit{Str}^{\text{Ag}}$, $\chi[\text{Ag} \mapsto \bar{\sigma}]$ is the assignment that maps each agent a to $\bar{\sigma}^a$, and is equal to χ on variables; similarly, $\chi[s \mapsto \sigma]$ is the assignment that maps variable s to σ and is otherwise equal to χ. When χ is defined for all agents and variables we call it *total*. Given a total assignment χ and a set of agents $A \subseteq \text{Ag}$, we let $\chi\restriction_A$ be the restriction of χ to $A \cup \text{Var}$. A history $\rho = v_0\boldsymbol{\alpha}_1 \dots \boldsymbol{\alpha}_n v_n$ is *consistent* with χ if, for every $i < n$ and agent a in the domain of χ, $\boldsymbol{\alpha}^a_{i+1} = \chi(a)(\rho_{\le i})$.

Outcomes. For a total agent assignment χ and a history ρ, we let $\text{Out}(\chi, \rho)$ be the only play π that extends ρ consistently with χ: $\pi = \rho \cdot \boldsymbol{\alpha}_1 v_1 \boldsymbol{\alpha}_2 v_2 \dots$ where for all $i \ge 0$ and $a \in \text{Ag}$, $\boldsymbol{\alpha}^a_{i+1} = \chi(a)(\pi_{\le|\rho|+i})$ and $v_{i+1} = E(v_i, \boldsymbol{\alpha}_{i+1})$.

Indistinguishability. Two positions v, v' are *indistinguishable* to agent a if $v \sim_a v'$, and similarly for actions. Two joint actions $\boldsymbol{\alpha}$ and $\boldsymbol{\beta}$ are indistinguishable to agent a, written $\boldsymbol{\alpha} \sim_a \boldsymbol{\beta}$, if they are component-wise indistinguishable, i.e., if $\boldsymbol{\alpha}^b \sim_a \boldsymbol{\beta}^b$ for each $b \in \text{Ag}$.

Synchronous perfect recall. We consider agents that remember their observations of all past positions and actions. Each observation relation $\sim_a$ is thus extended to histories as follows: given two histories $\rho = v_0\boldsymbol{\alpha}_1 \dots \boldsymbol{\alpha}_n v_n$ and $\rho' = v'_0\boldsymbol{\alpha}'_1 \dots \boldsymbol{\alpha}'_m v'_m$ we write $\rho \sim_a \rho'$ if $m = n$, $v_i \sim_a v'_i$ for all $i \in \{0, \dots, n\}$, and $\boldsymbol{\alpha}_i \sim_a \boldsymbol{\alpha}'_i$ for all $i \in \{1, \dots, n\}$.

2.3 Informed semantics

We now introduce several notions needed to define the informed semantics.

Informed perfect recall. To model agents that know some of the agents' strategies, we refine the observation relation on histories as follows. Let χ be a total assignment, a an agent and $\mathcal{A}_a$ the subset of agents whose strategies a knows. Two histories ρ and ρ' are *indistinguishable* for agent a, written $\rho \sim_a^\chi \rho'$, if the following two conditions hold:

1. $\rho \sim_a \rho'$
2. ρ and ρ' are both consistent with $\chi\restriction_{\mathcal{A}_a}$.

Given an agent a and total agent assignment χ, we say that a strategy σ is $\sim_a^\chi$*-uniform* if for all histories ρ and ρ' such that $\rho \sim_a^\chi \rho'$ it holds that $\sigma(\rho) = \sigma(\rho')$.

Remark 1. Relation $\sim_a^\chi$ is not an equivalence relation (it is not reflexive): histories inconsistent with $\chi\restriction_{\mathcal{A}_a}$ are not related to any history; it is an equivalence relation plus isolated points. The only impact this has on the semantics concerns nested knowledge operators, on which it has interesting consequences that we do not discuss here for lack of space.

Indistinguishable assignments. When evaluating an epistemic formula $K_a\varphi$ in a history ρ and assignment χ, existing variants of SL with knowledge such as [Maubert and Murano, 2018] and [Belardinelli *et al.*, 2017a] evaluate φ in all equivalent histories $\rho' \sim_a \rho$ with *the same assignment* χ. While quantifying on equivalent histories for the usual (uninformed) relation $\sim_a$ models agents that do not know the strategies being used, fixing the assignment χ indicates on the other hand agents *that know all agents' strategies*. As a consequence of this clash, this semantics neither properly captures the knowledge of agents who are unaware of each other's strategies, nor does it properly model the knowledge of agents who do know each other's strategies, as we show in the next section.

To remedy this we introduce an equivalence relation on assignments. Two assignments χ and χ' are indistinguishable to agent a, written $\chi \sim_a \chi'$, if they agree on all variables, and on the agents' strategies that a is informed of. Formally, $\chi \sim_a \chi'$ if $\chi'\restriction_{\mathcal{A}_a} = \chi\restriction_{\mathcal{A}_a}$. The semantics of knowledge will quantify over all equivalent histories and assignments.

Interpreting strategy quantification. In introducing imperfect information to Strategy Logic, it is difficult to determine for whom a strategy should be uniform, at the level of strategy quantification. One solution is to specify it explicitly, as in [Berthon *et al.*, 2021; Maubert and Murano, 2018]. Another is to require that it be uniform for all agents that are bound to it after the quantification, as done in [Belardinelli *et al.*, 2020]. However neither of these solutions can be used to define the informed semantics. Indeed, the informed perfect-recall relation $\sim_a^\chi$ depends on the assignment χ, which is only known after all agents have been bound to their strategy.

We thus propose a new way of interpreting the uniformity of strategies: instead of putting the uniformity constraint at the level of strategy quantification, we postpone it until after the group binding. In addition to allowing for a natural definition of the informed semantics, it is also a very neat solution to defining the semantics of SL with imperfect information, independent from the matter of informed knowledge.

There are two dual ways of dealing with the uniformity constraint. Either we require that the strategy being assigned should be uniformed and the remaining subformula should hold, or we say that if the strategy being assigned is uniform, then the subformula should hold. The first case corresponds to existentially quantified strategies, the second to universally quantified ones. Given a group binding $(\mathrm{Ag}, \bar{s})$, we let $\mathrm{Ag}_{\bar{s}}^{\langle\rangle} = \{a \in \mathrm{Ag} \mid \bar{s}(a) = (s, \langle\rangle) \text{ for some } s\}$ and $\mathrm{Ag}_{\bar{s}}^{[]} = \{a \in \mathrm{Ag} \mid \bar{s}(a) = (s, []) \text{ for some } s\}$.

Definition 3. Let $\mathcal{G} = (\mathrm{Ac}, V, E, \ell, V_\iota, \{\sim\}_{a\in\mathrm{Ag}}, \{\mathcal{A}_a\}_{a\in\mathrm{Ag}})$ be a CGS, χ a total assignment, ρ a history and φ an $\mathsf{SLK}_{\text{inf}}[\mathrm{BG}]$-formula. We define $\mathcal{G}, \chi, \rho \models \varphi$ as follows:

$$
\begin{array}{lll}
\mathcal{G},\chi,\rho \models p & \text{if} & p \in \ell(\mathrm{last}(\rho)) \\
\mathcal{G},\chi,\rho \models \neg\varphi & \text{if} & \mathcal{G},\chi,\rho \not\models \varphi \\
\mathcal{G},\chi,\rho \models \varphi \vee \varphi' & \text{if} & \mathcal{G},\chi,\rho \models \varphi \text{ or } \mathcal{G},\chi,\rho \models \varphi' \\
\mathcal{G},\chi,\rho \models \exists s.\varphi & \text{if} & \exists\, \sigma \in \mathit{Str} \text{ s.t. } \mathcal{G},\chi[s \mapsto \sigma],\rho \models \varphi \\
\mathcal{G},\chi,\rho \models (\mathrm{Ag},\bar{s})\varphi & \text{if} & \\
\quad \text{if } \forall a \in \mathrm{Ag}_{\bar{s}}^{[]},\ \chi'(a) \text{ is } \sim_a^{\chi'}\text{-uniform, then} & & \\
\quad\quad \forall a \in \mathrm{Ag}_{\bar{s}}^{\langle\rangle},\ \chi'(a) \text{ is } \sim_a^{\chi'}\text{-uniform, and} & & \\
\quad\quad\quad \mathcal{G},\chi',\rho \models \varphi, \text{ where } \chi' = \chi[\mathrm{Ag} \mapsto \bar{s}] & & \\
\mathcal{G},\chi,\rho \models K_a\varphi & \text{if} & \forall \rho' \in \mathrm{Hist} \text{ s.t. } \rho \sim_a^\chi \rho' \text{ and} \\
 & & \forall \chi' \text{ s.t. } \chi \sim_a \chi',\ \nu,\chi',b,\rho' \models \varphi
\end{array}
$$

and, writing $\pi = \mathrm{Out}(\chi, \rho)$:

$$
\begin{array}{lll}
\mathcal{G},\chi,\rho \models X\varphi & \text{if} & \mathcal{G},\chi,\pi_{\leq|\rho|+1} \models \varphi \\
\mathcal{G},\chi,\rho \models \varphi U \varphi' & \text{if} & \exists\, i \geq 0 \text{ s.t. } \mathcal{G},\chi,\pi_{\leq|\rho|+i} \models \varphi', \\
 & & \text{and } \forall\, j \text{ s.t. } 0 \leq j < i, \\
 & & \mathcal{G},\chi,\pi_{\leq|\rho|+j} \models \varphi
\end{array}
$$

We explain the case for the binding operator, which is the novelty of this semantics. The idea is that the two dual bindings $\langle\rangle$ and $[]$ are used so that the first line restricts universal quantifications to uniform strategies by ignoring those that are not, while the second line forces existential quantifiers to pick uniform strategies. In both cases, uniformity of strategies is evaluated for the "informed" semantics with respect to each agent's strategy as assigned by the binding.

The semantics of a sentence φ does not depend on χ, and we write $\mathcal{G}, v_\iota \models \varphi$ if $\mathcal{G},\chi,v_\iota \models \varphi$ for some χ.

Public actions. A CGS has *public actions* if for all $\alpha \in \mathrm{Ac}$ and $a \in \mathrm{Ag}$, $\sim_a \cap\, \mathrm{Ac} \times \mathrm{Ac}$ is the identity relation, meaning that all agents observe all actions perfectly.

Hierarchical information. A CGS presents *hierarchical information* if there exists an enumeration $\mathrm{Ag} = \{a_1, \ldots, a_n\}$ of the agents such that $\sim_{a_{i+1}} \subseteq \sim_{a_i}$ for all $1 \leq i < n$.

Remark 2. The reason why we base our logic on the flat fragment of SL with Boolean Goals is twofold. First, to express the uniformity constraint on strategies for the informed semantics, we need to know the strategies that will be used by all agents. The group binding of the Boolean Goals fragment makes this possible. Second, if agents can change strategy after time 0, as is possible in the non-flat fragment, it becomes unclear how to define the informed semantics: should agents know only the current strategy of other agents, and believe that they have been using it from the start? Or should they know the sequence of strategies that other agents have used? Both interpretations could be defined and studied, but we leave this matter for future work. Besides, the flat fragment of SL with Boolean Goals is already very expressive and can express, for instance, distributed synthesis, rational synthesis, and existence of Nash equilibria.

3 Illustrating the informed semantics

We discuss some aspects of our semantics, that we illustrate in the classic scenario of matching pennies.

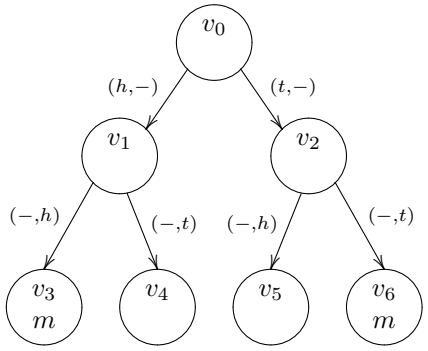

Figure 1: CGS $\mathcal{G}$ for the game of matching pennies

Consider the game of matching pennies $\mathcal{G}$ illustrated in Figure 1, where the players are called 1 and 2 (loops on final positions are omitted). Each player has a coin that only they can see. First, player one places her hidden coin heads or tails up, then player 2 does the same with her hidden coin. The proposition m stands for the coins matching. We suppose that the players only observe their own actions, so $v_3 \sim_1 v_4$, $v_5 \sim_1 v_6$, $v_1 \sim_2 v_2$, $v_3 \sim_2 v_5$, and $v_4 \sim_2 v_6$. In this example we omit the action names in the histories without ambiguity.

3.1 Impact on knowledge

We first illustrate how restricting the quantification on equivalent plays to those that are consistent with known strategies affects the knowledge of the agents. The usual semantics for the knowledge operator in epistemic extensions of SL, that we can find for instance in [Maubert and Murano, 2018] and [Belardinelli *et al.*, 2020], and that we shall refer to as "uninformed semantics" and denote $\models_u$, is the following:

$$\mathcal{G}, \chi, \rho \models_u K_a\varphi \quad \text{if for all } \rho' \text{ s.t. } \rho' \sim_a \rho,\ \mathcal{G}, \chi, \rho' \models_u \varphi$$

It differs from our semantics in two respects: (i) it uses the usual "uninformed" equivalence relation $\sim_a$ on histories instead of the informed one $\sim_a^\chi$, and (ii) it does not quantify over indistinguishable assignments.

To illustrate the first difference (i), consider the formula

$$\alpha = \forall s_1.\forall s_2.[1, s_1][2, s_2]XX(K_1 m \vee K_1 \neg m)$$

which means that no matter which strategies are used, in the end player 1 knows whether the coins match or not. With the uninformed semantics it clearly does not hold, as player 1 does not observe player 2's coin (so that $v_0v_1v_3 \sim_1 v_0v_1v_4$ and $v_0v_2v_5 \sim_1 v_0v_2v_6$), nor does she know her strategy, so $\mathcal{G}, v_0 \not\models_u \alpha$. Now with the informed semantics, it is enough to assume that player 1 knows player 2's strategy (i.e. $2 \in \mathcal{A}_1$) to make the formula true. Indeed, for any strategies s_1 and s_2 for players 1 and 2, the only history that player 1 considers possible with the uninformed relation $\sim_1$ besides the actual outcome, is not consistent with strategy s_2 and is thus discarded by $\sim_1^\chi$. For instance if the outcome is $v_0v_1v_3$, i.e. s_1 and s_2 are the strategies that play heads, then the only history equivalent for $\sim_1$ is $v_0v_1v_4$, in which player 2 plays tails, which is not consistent with s_2. So $\mathcal{G}, v_0 \models \alpha$.

We now illustrate the second difference (ii). Since the assignment does not impact the truth value of a formula without temporal operators, the only case in which considering equivalent assignments makes a difference is for formulas that involve knowledge about the future. So consider formula

$$\beta = \forall s_1.\forall s_2.[1, s_1][2, s_2](K_1 XXm \vee K_1 XX\neg m)$$

which means that for any pair of strategies s_1 and s_2, player 1 knows, before any move is made, whether in the end the coins will match. Intuitively this can only be the case if player 1 knows the strategies of both players. We check that indeed, with our semantics, the formula holds if we assume that $\mathcal{A}_1 = \{1, 2\}$, and it does not if $\mathcal{A}_1 \subsetneq \{1, 2\}$. Fix a pair of strategies (s_1, s_2), let χ assign s_1 to 1 and s_2 to 2, and assume for instance that these strategies lead to matching coins, i.e. $\mathcal{G}, \chi, v_0 \models XXm$. First, note that the only history indistinguishable from v_0 via $\sim_1^\chi$ is v_0. Now if $\mathcal{A}_1 \subsetneq \{1, 2\}$, meaning that player 1 does not know one of the strategies, then there exists at least one assignment $\chi' \sim_1 \chi$ that has a different outcome, i.e. such that $\mathcal{G}, \chi', v_0 \models XX\neg m$. The same reasoning can be made if (s_1, s_2) lead to coins that do not match. As a result, $\mathcal{G}, v_0 \not\models \beta$. Now assume that player 1 knows both strategies, i.e., $\mathcal{A}_1 = \{1, 2\}$. For any assignment χ, $\chi \sim_1 \chi'$ implies that $\chi' = \chi$. It follows that either $\mathcal{G}, \chi, v_0 \models K_1 XXm$, if $\mathcal{G}, \chi, v_0 \models XXm$, or $\mathcal{G}, \chi, v_0 \models K_1 XX\neg m$, if $\mathcal{G}, \chi, v_0 \models XX\neg m$. So $\mathcal{G}, v_0 \models \beta$.

One can see that also with the usual "uninformed" semantics, it holds that $\mathcal{G}, v_0 \models \beta$: indeed, since it only considers the actual assignment, it corresponds in the previous paragraph to the case where $\mathcal{A}_1 = \{1, 2\}$, i.e., the case where 1 knows both strategies. This is in contrast with what we illustrated with point (i), where the uninformed semantics corresponds to agents that do not know each other's strategies, and it demonstrates an inconsistency in the usual semantics. The semantics we propose solves this problem.

3.2 Impact on strategizing power

Our semantics also makes a difference for agents' strategic ability. Consider the following alternate "uninformed" semantics, which uses relation $\sim_a$ instead of $\sim_a^{\chi'}$:

$$\begin{array}{l} \mathcal{G}, \chi, \rho \models_u (\mathrm{Ag}, \bar{s})\varphi \quad \text{if} \\ \text{if } \forall a \in \mathrm{Ag}_{\bar{s}}^{[]},\ \chi'(a) \text{ is } \sim_a\text{-uniform, then} \\ \forall a \in \mathrm{Ag}_{\bar{s}}^{\langle\rangle},\ \chi'(a) \text{ is } \sim_a\text{-uniform, and} \\ \nu, \chi', \rho \models \varphi, \text{ where } \chi' = \chi_{\bar{s}}^{\nu} \end{array}$$

Consider again the example from Figure 1, and assume that player 2 knows player 1's strategy, i.e., $1 \in \mathcal{A}_2$. Formula

$$\gamma = \exists s_2.\forall s_1.[1, s_1]\langle 2, s_2\rangle XXm$$

means that player 2 has a strategy to ensure that the coins match. Knowing player 1's strategy, intuitively this should be true: if player 1 plans to play heads, then play heads, otherwise play tails. However with the semantics that uses the usual "uninformed" notion of uniformity, this formula does not hold: indeed, because v_1 and v_2 are indistinguishable to

player 2, any $\sim_2$-uniform strategy s_2 must play the same action in histories v_0v_1 and v_0v_2. If it is heads, then the coins will not match if 1 plays tails, and vice versa. So $\mathcal{G}, v_0 \not\models_u \gamma$.

However with our informed semantics, the formula holds. Indeed consider the strategy s_2 such that $s_2(v_0.v_1) = h$ and $s_2(v_0.v_2) = t$. Clearly it ensures that the coins match. We just need to see that it is also uniform for the informed notion of uniformity. Let s_1 be a strategy and $\chi = [1 \mapsto s_1, 2 \mapsto s_2]$. To see that s_2 is $\sim_2^\chi$-uniform, observe that either v_0v_1 or v_0v_2 is not consistent with s_1. Therefore it is not the case that $v_0v_1 \sim_2^\chi v_0v_2$, and s_1 does not have to assign the same action in both histories. As a result, $\mathcal{G}, v_0 \models \gamma$.

4 Model Checking $\mathsf{SLK}_{\text{inf}}[\text{BG}]$

We consider the *model-checking problem* for $\mathsf{SLK}_{\text{inf}}[\text{BG}]$ which consists in deciding, given a CGS $\mathcal{G}$, an initial position $v_\iota \in V_\iota$ and an $\mathsf{SLK}_{\text{inf}}[\text{BG}]$ sentence Φ, whether $\mathcal{G}, v_\iota \models \Phi$.

4.1 Undecidability for hierarchical information

It is proved in [van der Meyden and Wilke, 2005] that distributed synthesis for LTL with knowledge is undecidable, even for systems with hierarchical information. The setting there corresponds to agents that know the strategy of everyone but the environment, which we can model in our informed semantics. Since distributed synthesis for LTLK objectives is expressible in $\mathsf{SLK}_{\text{inf}}[\text{BG}]$, we obtain that:

Theorem 1. *Model checking* SLK *is undecidable, even when restricted to CGS with hierarchical information.*

4.2 Decidability for public actions

We now show that model checking $\mathsf{SLK}_{\text{inf}}[\text{BG}]$ is decidable when all actions are public. This is in line with the results in [van der Meyden and Wilke, 2005], where the informed semantics is used implicitly, and distributed synthesis for LTLK is shown to be decidable on broadcast systems, which are very close to our CGS with public actions.

As in [Laroussinie *et al.*, 2015; Berthon *et al.*, 2021; Belardinelli *et al.*, 2020], we solve the model-checking problem by reducing it to the model-checking problem for QCTL^*, for which we briefly recall the syntax and semantics.

Definition 4. The syntax of QCTL^* is defined as follows:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid E\psi \mid \exists p.\varphi$$
$$\psi ::= \varphi \mid \neg\psi \mid \psi \vee \psi \mid X\psi \mid \psi U \psi$$

where $p \in \text{AP}$.

Note that the set AP may not be the same as in the previous section (in the reduction from $\mathsf{SLK}_{\text{inf}}$ to QCTL^* we will use a bigger finite set of atomic propositions for QCTL^*).

We consider the *tree semantics* for QCTL^*, where formulas are evaluated on infinite trees, typically representing unfoldings of finite Kripke structures.

Trees. Let Σ be a finite alphabet. An infinite *tree* over Σ and AP is a mapping $t : \Sigma^* \to 2^{\text{AP}}$. An element $x \in \Sigma^*$ is a *node*, the empty word ϵ is the *root*, a node y is a *child* of a node x if it is of the form $y = x \cdot d$ with $d \in \Sigma$, and $t(x)$ is the *label* of node x. A *path* $\lambda = x_0x_1 \dots$ is an infinite sequence of nodes such that for all $i \geq 0$, x_{i+1} is a child of x_i. We let $\lambda_i = x_i$ denote the node at position i, and $\lambda_{\geq i} = x_ix_{i+1} \dots$ is the suffix of λ starting at position i. We say that λ *starts* in node x_0, and let $\text{Paths}(x)$ be the set of paths that start in x. The *subtree* of t rooted in node x is the tree $t_x : y \mapsto t(x \cdot y)$. A tree is *regular* if it has finitely many different subtrees. Finally, given a tree t and an interpretation $X_p \subseteq \Sigma^*$ for an atom $p \in \text{AP}$, we define the tree

$$t \otimes X_p : x \mapsto \begin{cases} t(x) \cup \{p\} & \text{if } x \in X_p \\ t(x) \setminus \{p\} & \text{otherwise.} \end{cases}$$

Definition 5. Let t be a tree, x a node and λ a path. The semantics of QCTL^* is defined as follows (we omit negation and disjunction):

$$\begin{array}{lll} t, x \models p & \text{if} & p \in t(x) \\ t, x \models E\psi & \text{if} & \exists \lambda \in \text{Paths}(x) \text{ s.t. } t, \lambda \models \psi \\ t, x \models \exists p.\varphi & \text{if} & \exists X_p \subseteq \Sigma^* \text{ s.t. } t \otimes X_p, x \models \varphi \\ t, \lambda \models \varphi & \text{if} & t, \lambda_0 \models \varphi \\ t, \lambda \models X\psi & \text{if} & t, \lambda_{\geq 1} \models \psi \\ t, \lambda \models \psi U \psi' & \text{if} & \exists i \geq 0 \text{ s.t. } t, \lambda_{\geq i} \models \psi' \text{ and} \\ & & \forall j < i,\ t, \lambda_{\geq j} \models \psi \end{array}$$

We write $t \models \varphi$ for $t, \epsilon \models \varphi$.

Reduction. For the rest of this section we fix a CGS $\mathcal{G} = (\text{Ac}, V, E, \ell, V_\iota, \{\sim_a\}_{a \in \text{Ag}}, \{\mathcal{A}_a\}_{a \in \text{Ag}})$ with public actions and an $\mathsf{SLK}_{\text{inf}}[\text{BG}]$ sentence Φ. The reduction is based on the one in [Belardinelli *et al.*, 2020], where the two main ideas are that when actions are public and transitions deterministic, **(1)** each history is entirely determined by its initial position and the sequence of joint actions, and **(2)** two indistinguishable histories share the same sequence of joint actions. Note that only the first point relies on determinism.

We thus work on the tree consisting of all sequences of joint actions. A node $x = \boldsymbol{\alpha}_1 \dots \boldsymbol{\alpha}_n$ in this tree represents the set of all histories ρ sharing this sequence of joint actions. By point 2 above, two histories that are indistinguishable for some agent are represented by the same node. This makes it possible to check uniformity of strategies locally in the tree.

Formally, let $\Sigma = \text{Ac}^{\text{Ag}}$, and for a node $x \in \Sigma^*$ and an initial position $v \in V_\iota$, we let ρ_v^x be the history in $\mathcal{G}$ that starts in v and has sequence of joint actions x. This is uniquely defined because the transition function is total and deterministic, and all histories can be written in this form. Point 2 above can be formalized as:

Lemma 2. *If* $\rho_v^x \sim_a \rho_u^y$ *then* $x = y$.

We introduce propositions $\text{AP}_{\text{act}} = \{p_{\boldsymbol{\alpha}} \mid \boldsymbol{\alpha} \in \text{Ac}^{\text{Ag}}\}$ that mark in each node the last joint action performed, and propositions $\text{AP}_{\text{pos}} = \{p_\nu \mid \nu : V_\iota \to V\}$ that indicate the current position $\nu(v_\iota)$ given that the initial position was v_ι. We then define the tree $t_{\mathcal{G}}$ over Σ and $\text{AP}_{\text{act}} \cup \text{AP}_{\text{pos}}$ as

$$t_{\mathcal{G}} : \boldsymbol{\alpha}_1 \dots \boldsymbol{\alpha}_n \mapsto \{p_{\boldsymbol{\alpha}_n}, p_\nu\}$$

where $\nu : v_\iota \mapsto \text{last}(\rho_{v_\iota}^{\boldsymbol{\alpha}_1 \cdots \boldsymbol{\alpha}_n})$. So each node is labeled by exactly one atom in AP_{act} and one in AP_{pos}. Also, since $\mathcal{G}$ is finite, the tree $t_{\mathcal{G}}$ is regular.

To keep track of strategies we introduce propositions $\text{AP}_{\text{str}} = \{p^{\tau}_s \mid \tau : V_\iota \to \text{Ac} \text{ and } s \in \text{Var}(\Phi)\}$, where τ is called a *local strategy* and $\text{Var}(\Phi)$ is the set of strategy variables in Φ (we assume without loss of generality that each variable is quantified at most once in Φ). Proposition p^{τ}_s in node x means that on history ρ^x_v, strategy s chooses action $\tau(v)$. To represent agents' knowledge we introduce propositions $\text{AP}_{\approx} = \{p^a_{\approx} \mid a \in \text{Ag} \text{ and } \approx \subseteq V_\iota \times V_\iota\}$, where $\approx$ is called a *local observation relation*. The intended meaning of $p^a_{\approx}$ holding in a node x is that histories ρ^x_v and ρ^x_u are indistinguishable to agent a if and only if $v \approx u$. By Lemma 2, indistinguishability relations can be completely characterized by such propositions. To capture the interpretation of "indistinguishable to agent a" corresponding to the informed semantics, these propositions will be placed on the tree *after* those encoding strategies, and consistently with them.

We define for each subformula φ of Φ a QCTL^* formula $[\varphi]^{\mathcal{B}}_{v_\iota}$. The translation is parameterized by a mapping $\mathcal{B} : \text{Ag} \rightharpoonup \text{Var}$ that represents the current binding of agents to strategy variables, and a position $v_\iota \in V_\iota$ representing the initial position of the current history. For atomic propositions:

$$[p]^{\mathcal{B}}_{v_\iota} \;=\; \bigvee_{\nu \mid p \in \ell(\nu(v_\iota))} p_\nu$$

This formula grabs the unique ν for which p_ν holds in the current node, and checks whether p holds in the current position as indicated by ν given that the initial position is v_ι. Boolean connectives simply distribute:

$$[\neg\varphi]^{\mathcal{B}}_{v_\iota} = \neg[\varphi]^{\mathcal{B}}_{v_\iota} \qquad [\varphi \vee \varphi']^{\mathcal{B}}_{v_\iota} = [\varphi]^{\mathcal{B}}_{v_\iota} \vee [\varphi']^{\mathcal{B}}_{v_\iota}$$

For strategic quantification, let $\{\tau_1, \dots, \tau_k\}$ be an enumeration of the possible local strategies, with $k = |\text{Ac}|^{|V_\iota|}$. Let

$$\begin{aligned} [\exists s.\varphi]^{\mathcal{B}}_{v_\iota} \;&=\; \exists p^s_{\tau_1} \dots \exists p^s_{\tau_k}.\, \text{strat}_s \wedge [\varphi]^{\mathcal{B}}_{v_\iota} \qquad \text{where} \\ \text{strat}_s \;&=\; AG \bigvee_i (p^s_{\tau_i} \wedge \bigwedge_{j \neq i} \neg p^s_{\tau_j}) \end{aligned}$$

checks that each node is labeled by exactly one local strategy.

Next comes the binding operator $(\text{Ag}, \bar{s})$, which we represent as a mapping $\mathcal{B}' : \text{Ag} \to \text{Var} \times \{\langle\rangle, []\}$. Let $\{a_1, \dots, a_m\}$ be an enumeration of the agents and $\{\approx_1, \dots, \approx_l\}$ an enumeration of the possible equivalence relations over V_ι. For each agent a we define the initial local observation relation $\approx^\iota_a = \sim_a \cap V_\iota \times V_\iota$. Also, given a local observation relation $\approx_a$ for agent a (in some node x), a tuple $\bar{\tau}$ of local strategies for each agent, as well as a next joint action $\boldsymbol{\alpha}$ and resulting current positions ν (in node $x \cdot \boldsymbol{\alpha}$), we define the updated local observation relation for agent a as $\approx'_a (\approx_a, \bar{\tau}, \boldsymbol{\alpha}, \nu) =$

$$\left\{ (v_\iota, u_\iota) \mid \begin{array}{c} v_\iota \approx_a u_\iota,\, \nu(v_\iota) \sim_a \nu(u_\iota) \text{ and} \\ \forall b \in \mathcal{A}_a, \tau_b(v_\iota) = \boldsymbol{\alpha}_b = \tau_b(u_\iota) \end{array} \right\}$$

This definition means that two histories $\rho^{x \cdot \boldsymbol{\alpha}}_{v_\iota}$ and $\rho^{x \cdot \boldsymbol{\alpha}}_{u_\iota}$ are indistinguishable for a if $\rho^x_{v_\iota}$ and $\rho^x_{u_\iota}$ are indistinguishable ($v_\iota \approx_a u_\iota$), the new positions $\text{last}(\rho^{x \cdot \boldsymbol{\alpha}}_{v_\iota})$ and $\text{last}(\rho^{x \cdot \boldsymbol{\alpha}}_{u_\iota})$ are indistinguishable ($\nu(v_\iota) \sim_a \nu(u_\iota)$), and for each agent b whose strategy is known to a, the last action $\boldsymbol{\alpha}_b$ played by b corresponds to strategy τ_b in both $\rho^x_{v_\iota}$ and $\rho^x_{u_\iota}$ (so $\rho^{x \cdot \boldsymbol{\alpha}}_{v_\iota}$ and $\rho^{x \cdot \boldsymbol{\alpha}}_{u_\iota}$ are consistent with τ_b). Given a binding $\mathcal{B}$, we let $\text{Ag}^{\langle\rangle}_{\mathcal{B}}$ (resp. $\text{Ag}^{[]}_{\mathcal{B}}$) be the set of agents a such that $\mathcal{B}(a) = (s, \langle\rangle)$ (resp. $(s, [])$) for some variable s. Define

$$\begin{aligned} [\mathcal{B}'\varphi]^{\mathcal{B}}_{v_\iota} \;=\; &\exists p^{a_1}_{\approx_1} \dots \exists p^{a_1}_{\approx_l} \dots \exists p^{a_m}_{\approx_1} \dots \exists p^{a_m}_{\approx_l}. \\ &\bigwedge_{a \in \text{Ag}} \text{uniq}_a \wedge \bigwedge_{a \in \text{Ag}} \text{inf}_a \wedge \\ &\bigwedge_{a \in \text{Ag}^{[]}_{\mathcal{B}'}} \text{unif}^{\mathcal{B}'(a)}_a \to \bigwedge_{a \in \text{Ag}^{\langle\rangle}_{\mathcal{B}'}} \text{unif}^{\mathcal{B}'(a)}_a \wedge [\varphi]^{\mathcal{B}'}_{v_\iota} \end{aligned}$$

where $\text{uniq}_a = AG \bigvee_{\approx} (p^a_{\approx} \wedge \bigwedge_{\approx' \neq \approx} \neg p^a_{\approx'})$ checks that each node is marked with one unique local observation relation,

$$\begin{aligned} \text{inf}_a \;=\; p^a_{\approx^\iota_a} \wedge AG \bigwedge_{\approx, \bar{\tau}} \Big[(p^a_{\approx} \wedge \bigwedge_{b \in \text{Ag}} p^{\mathcal{B}'(b)}_{\tau_b}) \to \\ AX \bigwedge_{\boldsymbol{\alpha}, \nu} \big((p_{\boldsymbol{\alpha}} \wedge p_\nu) \to p^a_{\approx'_a(\approx_a, \bar{\tau}, \boldsymbol{\alpha}, \nu)} \big) \Big] \end{aligned}$$

checks that the local observation relation is initially correct and that in each step it is updated correctly, and

$$\text{unif}^s_a \;=\; AG \bigwedge_{\approx, \tau} (p^a_{\approx} \wedge p^s_\tau) \to \text{loc-unif}(\approx, \tau)$$

checks that strategy s is uniform for agent a, with

$$\text{loc-unif}(\approx, \tau) = \begin{cases} \texttt{true} & \text{if for all } u_\iota, u'_\iota \in V_\iota \text{ s.t. } u_\iota \approx u'_\iota, \\ & \tau(u_\iota) = \tau(u'_\iota) \\ \texttt{false} & \text{otherwise.} \end{cases}$$

For the knowledge operator, let

$$[K_a\varphi]^{\mathcal{B}}_{v_\iota} \;=\; \bigwedge_{\approx} p^a_{\approx} \to \bigwedge_{u_\iota \mid v_\iota \approx u_\iota} [\varphi]^{\mathcal{B}}_{u_\iota}$$

We are left with the temporal operators:

$$[X\varphi]^{\mathcal{B}}_{v_\iota} \;=\; \bigwedge_{\bar{\tau}} (\bigwedge_a p^{\mathcal{B}(a)}_{\tau_a}) \to AX \left(p_{\boldsymbol{\alpha}(\bar{\tau}, v_\iota)} \to [\varphi]^{\mathcal{B}}_{v_\iota} \right)$$

where $\boldsymbol{\alpha}(\bar{\tau}, v_\iota)$ is the joint action where each agent a plays action $\tau_a(v_\iota)$. Similarly, we let

$$[\varphi U \varphi']^{\mathcal{B}}_{v_\iota} =$$
$$A\Big[G\Big(\bigwedge_{\bar{\tau}} (\bigwedge_a p^{\mathcal{B}(a)}_{\tau_a}) \to X(p_{\boldsymbol{\alpha}(\bar{\tau}, v_\iota)}) \Big) \to [\varphi]^{\mathcal{B}}_{v_\iota} U [\varphi']^{\mathcal{B}}_{v_\iota} \Big]$$

Lemma 3. *For every initial position $v_\iota \in V_\iota$ and every initial binding $\mathcal{B}$, it holds that $\mathcal{G}, v_\iota \models \Phi$ iff $t_{\mathcal{G}} \models [\Phi]^{\mathcal{B}}_{v_\iota}$.*

Because model checking QCTL^* on regular trees is decidable [Laroussinie and Markey, 2014], we obtain that

Theorem 4. *Model checking $\mathsf{SLK}_{\text{inf}}$ on CGS with public actions is decidable.*

5 Conclusion and future work

In this work we defined the first semantics for Strategy Logic that can model agents that know other agents' strategies. This semantics reflects the intuitive idea that knowing others' strategies refines agents' knowledge about the system, and increases their ability to strategize. It is also general enough to capture both situations found in the literature, which correspond to the two extreme cases. Two important questions left for future work are the implications of the informed semantics on higher-order knowledge, and trying to extend this new informed semantics to the full syntax of SL.

References

[Alur *et al.*, 2002] Rajeev Alur, Thomas A. Henzinger, and Orna Kupferman. Alternating-time temporal logic. *J. ACM*, 49(5):672–713, 2002.

[Belardinelli *et al.*, 2017a] Francesco Belardinelli, Alessio Lomuscio, Aniello Murano, and Sasha Rubin. Verification of broadcasting multi-agent systems against an epistemic strategy logic. In *IJCAI*, volume 17, pages 91–97, 2017.

[Belardinelli *et al.*, 2017b] Francesco Belardinelli, Alessio Lomuscio, Aniello Murano, and Sasha Rubin. Verification of multi-agent systems with imperfect information and public actions. In *AAMAS'17*, pages 1268–1276, 2017.

[Belardinelli *et al.*, 2020] Francesco Belardinelli, Alessio Lomuscio, Aniello Murano, and Sasha Rubin. Verification of multi-agent systems with public actions against strategy logic. *Artif. Intell.*, 285:103302, 2020.

[Berthon *et al.*, 2021] R. Berthon, B. Maubert, A. Murano, S. Rubin, and M. Vardi. Strategy logic with imperfect information. *ACM Trans. Comput. Logic*, 22(1), 2021.

[Chatterjee *et al.*, 2010] Krishnendu Chatterjee, Thomas A. Henzinger, and Nir Piterman. Strategy Logic. *Inf. Comput.*, 208(6):677–693, 2010.

[Gardy, 2017] Patrick Gardy. *Semantics of Strategy Logic*. PhD thesis, 2017.

[Guelev *et al.*, 2011] Dimitar P. Guelev, Catalin Dima, and Constantin Enea. An alternating-time temporal logic with knowledge, perfect recall and past: axiomatisation and model-checking. *J. Appl. Non-Class. Log.*, 21(1):93–131, 2011.

[Herzig and Lorini, 2010] Andreas Herzig and Emiliano Lorini. A dynamic logic of agency i: Stit, capabilities and powers. *Journal of Logic, Language and Information*, 19(1):89, 2010.

[Jamroga and van der Hoek, 2004] Wojciech Jamroga and Wiebe van der Hoek. Agents that know how to play. *Fundam. Inform.*, 63(2-3):185–219, 2004.

[Laroussinie and Markey, 2014] François Laroussinie and Nicolas Markey. Quantified CTL: expressiveness and complexity. *Log. Methods Comput. Sci.*, 10(4), 2014.

[Laroussinie *et al.*, 2015] François Laroussinie, Nicolas Markey, and Arnaud Sangnier. ATLsc with partial observation. In *GandALF'15*, pages 43–57, 2015.

[Maubert and Murano, 2018] Bastien Maubert and Aniello Murano. Reasoning about knowledge and strategies under hierarchical information. In *KR'18*, pages 530–540. AAAI Press, 2018.

[Mogavero *et al.*, 2014] Fabio Mogavero, Aniello Murano, Giuseppe Perelli, and Moshe Y. Vardi. Reasoning about strategies: On the model-checking problem. *ACM Trans. Comput. Log.*, 15(4):34:1–34:47, 2014.

[Pauly and Parikh, 2003] Marc Pauly and Rohit Parikh. Game logic-an overview. *Studia Logica*, 75(2):165–182, 2003.

[Peterson *et al.*, 2001] Gary Peterson, John Reif, and Salman Azhar. Lower bounds for multiplayer noncooperative games of incomplete information. *Comput. Math. with Appl.*, 41(7):957–992, 2001.

[Puchala, 2010] Bernd Puchala. Asynchronous omega-regular games with partial information. In *MFCS'10*, pages 592–603, 2010.

[Saffidine *et al.*, 2018] Abdallah Saffidine, François Schwarzentruber, and Bruno Zanuttini. Knowledge-based policies for qualitative decentralized pomdps. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 32, 2018.

[van der Hoek and Wooldridge, 2003] Wiebe van der Hoek and Michael Wooldridge. Cooperation, knowledge and time: Alternating-time Temporal Epistemic Logic and its applications. *Stud. Log.*, 75(1):125–157, 2003.

[van der Meyden and Wilke, 2005] Ron van der Meyden and Thomas Wilke. Synthesis of distributed systems from knowledge-based specifications. In Martín Abadi and Luca de Alfaro, editors, *CONCUR*, volume 3653 of *Lecture Notes in Computer Science*, pages 562–576. Springer, 2005.