Simulations for Multi-Agent Systems with Imperfect Information*

Patrick Gardy and Yuxin Deng

Shanghai Key Laboratory of Trustworthy Computing, East China Normal University
{gardy.patrick, yxdeng}@sei.ecnu.edu.cn

Abstract Equivalence-checking and simulations are well-known methods used to reduce the size of a system in order to verify it more efficiently. While Alur et al. proposed a notion of simulation sound and complete for ATL as early as 1998, there have been very few works on equivalence-checking performed on extensions of ATL* with probabilities, imperfect information, counters etc. In the case of multi-agent systems (MASs) with imperfect information, the lack of sound and complete algorithm mostly follows from the undecidability of ATL model-checking. However, while ATL is undecidable overall, there exist sub-classes of MASs for which ATL becomes decidable. In this paper, we propose a notion of simulation sound for ATL/ATL* on any MASs and complete on naive MASs. Using our simulations we design an equivalence-checking algorithm sound and complete for MASs with public actions.

1 Introduction

With the rise of multi-agent systems (MASs), the software verification community has tried to extend methods useful for the verification of closed systems to multi-agent systems. The usual model represents each agent's local control through a transition graph with the edges labeled by the actions of all agents involved in the system. This way the agents may influence the state of one another, but each has its own separate control-graph. The overall system is then built as the product of all the agents' local systems. In many practical cases, some agents have only a partial view of the overall system and may not know the control-graph or the exact state of other agents. This can either follow from a faulty communication or be a design choice, either for security or cost purposes. To model this imperfect information, some partial observation relations are attached to each agent.

Many formalisms have been proposed in order to specify expected behaviors of MASs. Among the most famous ones we cite ω-regular conditions [1] and ATL, ATL* [2,13,15], the go-to adaptation of CTL, CTL* to multi-agent systems. Initially defined on MASs with perfect information, these formalisms were quickly adapted and studied in the context of imperfect information (for example in [11,14] for ATL*).

* Supported by the National Natural Science Foundation of China (61672229, 61832015).

A need for equivalence-checking Simpler formalisms like Buchi conditions and ATL enjoy a polynomial model-checking for perfect information, making them target choices for practical applications. The situation is however drastically different in the presence of imperfect information. Thereby ATL goes from polynomial to exponential time model-checking (Δ_2^P to be precise) for positional strategies while it is outright undecidable for perfect recall strategies. The algorithm for positional strategies scales poorly and methods of minimizing the models are necessary to improve the practical uses. In this line of work, a proven concept consists in finding smaller and smaller models of the system and proving at each step that the new model despite its reduced size satisfies the same properties as the bigger one. Such method makes heavy use of an equivalence-check subroutine between two models. There are many ways to perform an equivalence-check: simulations [12,9], trace-equivalence [3], testing [17], etc. This idea was put in application in [5]. In their paper, Belardinelli et al. proposed a notion of simulation sound for ATL and discussed different modelizations of the three-ballot voting protocol (3BVP). ATL was shown to be a logic of choice to model security properties of voting protocols [23,4]. The authors of [5] proposed three models of the 3BVP and showed that each model can simulate the others. We can then check ATL security properties on the smallest model, gaining a considerable amount of time and space.

Contributions We propose a notion of simulation for games with imperfect information by extending the one of [5]. This simulation is sound for ATL/ATL*, works with both positional and perfect recall strategies, and (with a minor change in the definition) works for both the objective and subjective semantics. Our notion, unlike the one in [5], does not require perfect replication of the partial observation but instead focuses on similarity of results. To be more precise, for four states q, s, q', s' with q, q' similar, s, s' similar and $s' \sim_C q'$, we do not require the states q, s to have the same observation C. This makes our notion of simulation coarser than the only other existing one.

Due to the undecidability of ATL with perfect recall strategies and imperfect information, our notion is not proven to be complete[1]. We however prove completeness on *naive games*, a subclass of MASs with imperfect information. A naive game is one where by design the imperfect information is “state based” in the sense that no history can augment the information of an agent. The concept is illustrated later in Figure 4. Using our result on naive games, we develop an equivalence-checking algorithm for MASs with public actions, which is both sound and complete. The proof proceeds by restructuring public-actions MASs into naive MASs equivalent on all ATL formulas. To perform equivalence checking, both public-actions MASs are transformed into naive games which are then checked using our notion of simulation.

Related works ATL was proven undecidable in perfect recall strategies and Δ_2^P with positional strategies [11,14]. To regain decidability for perfect recall

[1] Continuing the tradition in multi-agent systems with the exception of the initial paper on alternating refinement relations [3].

strategies, there are two possibilities. The first option is to restrict the MASs to public actions [6]. A MAS has public actions whenever any agent can see the actions played by all other agents. In such case, ATL* model-checking is 2-EXPTIME. The second option is to use hierarchical observations (and other derivative options) for which ATL/ATL* model-checking is Non-Elementary. A MAS has hierarchical observation whenever there is an order on the agents such that an agent A dominated by another agent B has a strictly less complicated partial observation relation than B.

In a slightly more distant fashion, we mention the work of Berthon et al. [7] on strategy logic with imperfect information and also the work of Laroussinie et al. [16] on ATL with strategy contexts and partial observations (both logics extend ATL*). Each paper proposes small fragments on which the model-checking is decidable in the presence of partial observations.

There are two main related works on equivalence-checking. The first is by Alur et al. [3] on alternating refinement relations with two main contributions: alternating simulations (sound and complete for ATL/ATL*) and alternating trace containment (sound and complete for LTL). The second [5] proposes a simulation sound for ATL* in the presence of imperfect information with an application to model the 3BVP. The protocol is a voting process that does not rely on cryptographic methods for its security [25]. Interestingly, some practical problems and security failures were quickly detected in the 3BVP following its presentation [22]. In [5], the authors proposed different modelizations possible for the protocol as MASs with imperfect information. They discussed the size of each modelization before showing all the models to be equivalent. In a more distant fashion we also cite [26] which proposes a concept of simulation sound for ATL on probabilistic MASs.

Outline In Section 2, we introduce games with imperfect information (used to represent MASs) and ATL*. Section 3 covers the notion of simulation with its soundness relative to ATL* for games with imperfect information. Section 4 discusses the completeness of our notion for the subclass of naive games. In Section 5, we present an algorithm to perform equivalence checking on games with public actions based on the work done in previous sections. Finally, we conclude in Section 6.

2 Games, Imperfect information and ATL*

Games with imperfect information

For the rest of the paper, fix AP a finite set of atomic propositions. A multi-agent system is usually represented in the following way: each agent has its own control-graph whose edges are labeled by tuples of actions (one per agent), the overall system is then represented by a product of all local control-graphs of the agents. To model this product, we use the notion of concurrent game structures. This is the method used in the open-source tool MCMAS [18,24] and the ISPL language it uses.

Definition 1. *A* concurrent game structure with imperfect information *(CGS for short) is a tuple* $\mathcal{G} := (S, \mathtt{Agt}, \mathtt{Act}, \mathtt{Label}, \Delta, \{\sim_P\}_{P\in\mathtt{Agt}})$ *where* S *is a nonempty set of states;* $\mathtt{Agt} = \{P_1, ..., P_n\}$ *is a nonempty finite set of agents;* $\mathtt{Act}$ *is a nonempty finite set of actions;* $\mathtt{Label} : S \to 2^{AP}$ *is a labeling function;* $\Delta : S \times \mathtt{JAct} \to S$ *is a transition function with* $\mathtt{JAct} := \prod_{i\in\mathtt{Agt}} \mathtt{Act}$ *the set of joint actions (where the* i^{th} *component represents the choice of the agent* P_i*); and for each* $P \in \mathtt{Agt}$*,* $\sim_P \in S \times S$ *is an equivalence relation marking the partial observation of agent* P*.*

A CGS is said to have *perfect information* when $\sim_P = \{(s, s) \mid s \in S\}$ for each $P \in \mathtt{Agt}$. A *path* (or outcome) $\rho = s_0 s_1 \dots$ in a CGS $\mathcal{G}$ is a (finite or infinite) sequence of states such that for every $j \geq 0$, $s_{j+1} = \Delta(s_j, \overline{a}_j)$ for some joint action $\overline{a}_j \in \mathtt{JAct}$. We let $\mathtt{Path}_\mathcal{G}$ denote the set of paths in $\mathcal{G}$. When clear from context, we will drop the game from the notation. We write $|\rho| \in \mathbb{N} \cup \{\infty\}$ for the length of ρ, $\mathtt{last}(\rho)$ for the last state of ρ (when it is finite), and $\mathtt{Prefix}(\rho)$ for the set of all prefixes of ρ. Finally, we write $\rho_{<i+1}$ for the prefix of length i of ρ. Given two paths ρ and ρ', and an agent P we write $\rho \sim_P \rho'$ if for all index i, $\rho(i) \sim_P \rho'(i)$. We then call a set of agents with common knowledges the set of agents A such that $\rho \sim_P \rho'$ iff $P \in A$.

A function $\delta : S^+ \to \mathtt{Act}$ is called a *strategy* (with perfect recall and no randomness). We denote by $\mathtt{Strat}_\mathcal{G}$ the set of strategies. We say that a strategy δ conforms to the partial observation of a player P if for any two paths ρ and π of the same length such that $\rho(i) \sim_P \pi(i)$ for any i, we have $\delta(\pi) = \delta(\rho)$. Consider a state s, a coalition of agents $\mathbf{C} \subseteq \mathtt{Agt}$ and a set of strategies $\overline{\delta_\mathbf{C}} = (\delta_P)_{P\in\mathbf{C}}$ for players in $\mathbf{C}$. A path ρ is compatible with $\overline{\delta_\mathbf{C}}$ and s when $\rho(1) = s$ and for all $0 < i < |\rho|$ there exists a joint action $\overline{a}$ such that $\overline{a}(P) = \delta_P(\rho_{<i})$ for each agent P in $\mathbf{C}$ and $\rho(i+1) = \Delta(\mathtt{last}(\rho_{\leq i}), \overline{a})$. There are two ways to define outcomes in games with imperfect information: *objective* and *subjective*. The objective outcome $\mathtt{Out}_{obj}(\overline{\delta_\mathbf{C}}, s)$ is the set of all paths compatible with $\overline{\delta_\mathbf{C}}$ starting from s, thus it differentiates the initial state from similar states. The subjective semantics makes no such distinction, $\mathtt{Out}_{sub}(\overline{\delta_\mathbf{C}}, s) = \bigcup_{s' \sim_P s, P\in\mathbf{C}} \mathtt{Out}_{obj}(\overline{\delta_\mathbf{C}}, s')$. In order to analyze outcomes, we need the last concept: *traces*. A trace of a path is the projection of the path onto the set of atomic propositions $\mathtt{AP}$.

ATL^* on games with imperfect information

ATL^* is a well-known and widely used logic introduced in [2] for games with perfect information as an extension of the logic CTL^* for closed systems. It extends relatively simply to games with imperfect information, only using a little semantic change on the quantification operator. ATL^* is defined with respect to a set of agents $\mathtt{Agt}$ and a set of atomic propositions $\mathtt{AP}$ by the following grammar (note that as usual we do not allow the universal quantifier when dealing with simulations):

$$\begin{aligned} \mathsf{ATL}^* \ni \phi &:= \langle\!\langle \mathbf{C} \rangle\!\rangle \varphi \mid \phi \wedge \phi \mid \phi \vee \phi \\ \varphi &:= \mathbf{p} \mid \neg\mathbf{p} \mid \mathbf{X}\,\varphi \mid \varphi\,\mathbf{U}\,\varphi \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \phi \end{aligned}$$

where $\mathbf{p}$ is an atomic proposition and $\mathbf{C}$ is a subset of $\mathtt{Agt}$.

The ϕ-type formulas are *state formulas* and are evaluated on a state s of a *CGS* $\mathcal{G}$. The semantic interpretation of boolean operators is as usual. We recall that there are two semantics to define outcomes, subjective and objective. This gives rise to two semantics for the quantification, with the first being the objective definition and the second being the subjective definition:

$$\mathcal{G}, s \models_{obj} \langle\!\langle \mathbf{C} \rangle\!\rangle \varphi \quad \text{iff} \quad \begin{cases} \exists\overline{\delta} = \{\delta_P\}_{P\in\mathbf{C}} \in \mathtt{Strat} \text{ s.t. } \forall P \in \mathbf{C},\ \delta_P \text{ conforms to the} \\ \text{information of } P \text{ and } \forall\rho \in \mathtt{Out}_{obj}(\overline{\delta}, s) \text{ it holds } \mathcal{G}, \rho, 1 \models \varphi \end{cases}$$

$$\mathcal{G}, s \models_{sub} \langle\!\langle \mathbf{C} \rangle\!\rangle \varphi \quad \text{iff} \quad \begin{cases} \exists\overline{\delta} = \{\delta_P\}_{P\in\mathbf{C}} \in \mathtt{Strat} \text{ s.t. } \forall P \in \mathbf{C},\ \delta_P \text{ conforms to the} \\ \text{information of } P \text{ and } \forall\rho \in \mathtt{Out}_{sub}(\overline{\delta}, s) \text{ it holds } \mathcal{G}, \rho, 1 \models \varphi \end{cases}$$

The φ-type formulas are called *path-formulas* and are evaluated with respect to a path within the *CGS*. The semantics of the boolean operators and the atomic propositions is standard. The other operators follow the semantics below.

$$\begin{array}{lll} \mathcal{G}, \rho, i \models \mathbf{X}\,\varphi & \text{iff} & \mathcal{G}, \rho, i+1 \models \varphi \\ \mathcal{G}, \rho, i \models \varphi_1\,\mathbf{U}\,\varphi_2 & \text{iff} & \exists j > i.\ \mathcal{G}, \rho, j \models \varphi_2 \text{ and } \forall i < k < j.\ \mathcal{G}, \rho, k \models \varphi_1 \\ \mathcal{G}, \rho, i \models \phi & \text{iff} & \mathcal{G}, \rho(i) \models \phi \end{array}$$

We call ATL the fragment of ATL^* obeying the syntax

$$\begin{array}{rcl} \mathsf{ATL} \ni \phi & := & \langle\!\langle \mathbf{C} \rangle\!\rangle \varphi \mid \langle\!\langle \mathbf{C} \rangle\!\rangle \mathbf{X}\,\varphi \mid\ \mid \langle\!\langle \mathbf{C} \rangle\!\rangle \varphi\,\mathbf{U}\,\varphi \\ \varphi & := & \mathbf{p} \mid \neg\mathbf{p} \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \phi \end{array}$$

3 Simulation in games with imperfect information

In [5] the authors propose a notion of equivalence sound for ATL that works for both the subjective and the objective semantics. This notion is however rather restrictive. We develop our own notion, which shares some similarities with the one of [5], yet is more general. The simulation we propose is also sound for ATL^*, works on both subjective and objective semantics. Besides those properties already present in [5], our simulations do not require a perfect replication of the partial information. By "replication of partial information", we mean the following. Consider three states q, s', q' with q, q' similar and $s' \sim_P q'$, there is no need for the existence of a state s with $s \sim_P q$ and s, q similar. Finally our notion is complete on a small class of games: *naive games*, and from this completeness one can deduce an equivalence-checking algorithm for games with public actions.

For the rest of the paper, we consider two games $\mathcal{G}, \mathcal{G}'$ that build upon the same atomic propositions and upon the same set $\mathtt{Agt}$ of agents. Simulation − or equivalence-checking in general − in multi-agent systems is parameterized by a coalition of agents (made of all agents to be existentially quantified in the formulas we are interested in). Therefore we also fix a coalition $\mathbf{C} \subseteq \mathtt{Agt}$ as a parameter. We first describe the simulation and soundness for the objective semantics. The case for the subjective semantics is similar and will be discussed in the end. The main idea behind our algorithm is to keep track of all imperfect information scenarios possible through a tracker. We represent the tracker, written Λ, as a relation on $S \times S \times 2^{\texttt{Agt}} \times S' \times S' \times 2^{\texttt{Agt}}$.

Definition 2. *A simulation of $\mathcal{G}$ by $\mathcal{G}'$ for $\mathbf{C}$ is a relation $\mathcal{R} \subseteq S \times S'$ such that there is another relation $\Lambda \subseteq S \times S \times 2^{\texttt{Agt}} \times S' \times S' \times 2^{\texttt{Agt}}$ where*

1. *for each $(q, q') \in \mathcal{R}$, $\texttt{Label}(q) = \texttt{Label}(q')$.*
2. *for any (q, q') in $\mathcal{R}$, we have $(q, q, \mathbf{C}, q', q', \mathbf{C}) \in \Lambda$*
3. – *for each $(q, q') \in \mathcal{R}$, there is a function $\mathcal{T}_{q,q'} : \texttt{JAct}^{\mathbf{C}}_{\mathcal{G}} \mapsto \texttt{JAct}^{\mathbf{C}}_{\mathcal{G}'}$*
   – *for each $(q, q') \in \mathcal{R}$ and each $\overline{a} \in \texttt{JAct}^{\mathbf{C}}_{\mathcal{G}}$ there exists a function $\mathcal{U}^{\overline{a}}_{q,q'} : \texttt{JAct}^{\texttt{Agt}\setminus\mathbf{C}}_{\mathcal{G}'} \mapsto \texttt{JAct}^{\texttt{Agt}\setminus\mathbf{C}}_{\mathcal{G}}$*

   such that the following two properties hold:
   (a) *consider any $(q_1, q_2, A, q'_1, q'_2, B) \in \Lambda$, any two joint actions $\overline{a}, \overline{b} \in \texttt{JAct}^{\mathbf{C}}_{\mathcal{G}}$ such that $\overline{a}(A) = \overline{b}(A)$, and any two joint actions $\overline{c'}$ and $\overline{d}' \in \texttt{JAct}^{\texttt{Agt}\setminus\mathbf{C}}_{\mathcal{G}'}$. Write k_1 for the successor of q_1 by $\overline{a} \cdot \mathcal{U}^{\overline{a}}_{q_1,q'_1}(\overline{c'})$, k_2 for the successor of q_2 by $\overline{b} \cdot \mathcal{U}^{\overline{b}}_{q_2,q'_2}(\overline{d'})$, k'_1 for the successor of q'_1 by $\mathcal{T}_{q_1,q'_1}(\overline{a}) \cdot \overline{c'}$, k'_2 for the successor of q'_2 by $\mathcal{T}_{q_2,q'_2}(\overline{b}) \cdot \overline{d}'$, C the set of agent with information common to k_1, k_2; and D the set of agents with information common to k'_1, k'_2. Then $(k_1, k_2, E, k'_1, k'_2, F) \in \Lambda$ where $E = A \cap C$ and $F = B \cap D$.*
   (b) *for each $(q, q') \in \mathcal{R}$, each joint action $\overline{a} \in \texttt{JAct}^{\mathbf{C}}_{\mathcal{G}}$, there is a joint action $\overline{c}' \in \texttt{JAct}^{\texttt{Agt}\setminus\mathbf{C}}_{\mathcal{G}'}$ such that the pair consisting of a successor of q by $\overline{a} \cdot \mathcal{U}^{\overline{a}}_{q,q'}(\overline{c'})$ and a successor of q' by $\mathcal{T}_{q,q'}(\overline{a}) \cdot \overline{c'}$ is in $\mathcal{R}$.*
4. *for each $(q_1, q_2, A, q'_1, q'_2, B) \in \Lambda$*

$$\forall \overline{a}, \overline{b} \in \texttt{JAct}^{\mathbf{C}}_{\mathcal{G}}. \quad \left[\overline{a}(A) = \overline{b}(A)\right] \Rightarrow \left[\mathcal{T}_{q_1,q'_1}(\overline{a})(B) = \mathcal{T}_{q_2,q'_2}(\overline{b})(B)\right] \tag{1}$$

The above definition of simulations may look complicated but is in fact relatively similar to the one of ATL^* with the addition of the syntactic sugar to manage the tracker Λ. Indeed, Points 1 and 3.b are similar to the requirements of the simulations for ATL with perfect information [3]. Points 2 and 3.a are there to build the tracker properly. Intuitively, the tracker can be built based on $\mathcal{R}$ by a fix-point algorithm using Point 2 for initialization and Point 3.a as recurrence relation. Point 4 enforces the simulation to make coherent choices for the scenarios in the tracker. Note that if the tracker is larger than the one of the definition above, but the property in Point 4 still holds for the larger tracker, then the soundness for ATL will also hold. Note also that, while it may not look obvious, this kind of simulations is closed by union. The tracker for the union of two simulations is simply the union of the trackers from each simulation.

We provide a small example for the games on Figure 1. There exists a simulation of domain (where we omit the last states for clarity)

$$\mathcal{R} := \{(A, A'), (B, B'), (B, C'), (C, B'), (C, C')\}$$

and where the tracker is made of

$$\Lambda := \begin{cases} (A, A, \{P_i\}_{i\leq 3}, A', A', \{P_i\}_{i\leq 3}) \\ (B, C, \{P_1\}, B', C', \{P_2\}) \\ (C, B, \{P_1\}, B', C', \{P_2\}) \\ (C, B, \{P_1\}, C', B', \{P_2\}) \\ (B, C, \{P_1\}, C', B', \{P_2\}) \end{cases}$$

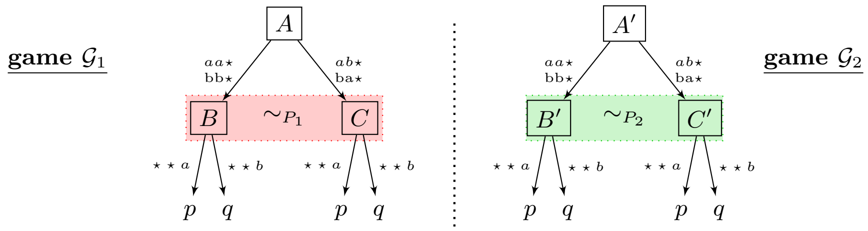

Figure 1. Two games bisimilar, each with 3 agents. The bisimilarity is relatively trivial as only the third player is active on the B, C, B' and C' states.

Remark 1. Using a naive approach, finding if there exists a simulation takes an exponential time.

Strategic characterization To establish the soundness of simulations for ATL^*, we restate simulations as relations between strategies. We need a few notations first. An *existential profile* $\overline{\delta}$ is a set of strategies $(\delta_{P_1}, \ldots, \delta_{P_n})$, one per agents in $\mathbf{C}$. Universal profiles are defined similarly as sets of strategies from the agents in $\mathtt{Agt}\backslash\mathbf{C}$. We write $\mathtt{Profile}^*_\star$ with $* \in \{\mathbf{C}, \mathtt{Agt}\backslash\mathbf{C}\}$ and $\star \in \{\mathcal{G}, \mathcal{G}'\}$ for the set of $*$-profiles in the $\star$-game. A *strategic characterization* is a set $\{\mathcal{S}^{\mathbf{C}}_{q,q'}, \mathcal{S}^{\mathtt{Agt}\backslash\mathbf{C}}_{q,q'}\}_{q,q'\in Z}$ of functions on some domain $Z \subseteq S \times S'$ where the functions are of the form $\mathcal{S}^{\mathbf{C}}_{q,q'} : \mathtt{Profile}^{\mathbf{C}}_{\mathcal{G}} \mapsto \mathtt{Profile}^{\mathbf{C}}_{\mathcal{G}'}$ and $\mathcal{S}^{\mathtt{Agt}\backslash\mathbf{C}}_{q,q'} : \mathtt{Profile}^{\mathtt{Agt}\backslash\mathbf{C}}_{\mathcal{G}'} \mapsto \mathtt{Profile}^{\mathtt{Agt}\backslash\mathbf{C}}_{\mathcal{G}}$ that obey two features:

Feat.1 for all q, q', any two profiles $\overline{\delta}, \overline{\gamma'}$, and any two states s, s' belonging to the objective outcomes of $\overline{\delta}$ and $\mathcal{S}^{\mathtt{Agt}\backslash\mathbf{C}}_{q,q'}(\overline{\gamma'})$ and of $\mathcal{S}^{\mathbf{C}}_{q,q'}(\overline{\delta})$ and $\overline{\gamma'}$, the pair (s, s') belongs to the domain Z of the strategic profile.

Feat.2 for any pair of states q, q', any two profiles $\overline{\delta}, \overline{\gamma'}$, the objective outcomes of $\overline{\delta}$ and $\mathcal{S}^{\mathtt{Agt}\backslash\mathbf{C}}_{q,q'}(\overline{\gamma'})$ and of $\mathcal{S}^{\mathbf{C}}_{q,q'}(\overline{\delta})$ and $\overline{\gamma'}$ have the same traces starting from q and q', respectively.

Simulations can be linked to strategic characterizations via Theorem 1 below.

Theorem 1. *If there exists a simulation $\mathcal{R}$ of $\mathcal{G}$ by $\mathcal{G}'$, then there is a strategic characterization defined on $\mathcal{R}$.*

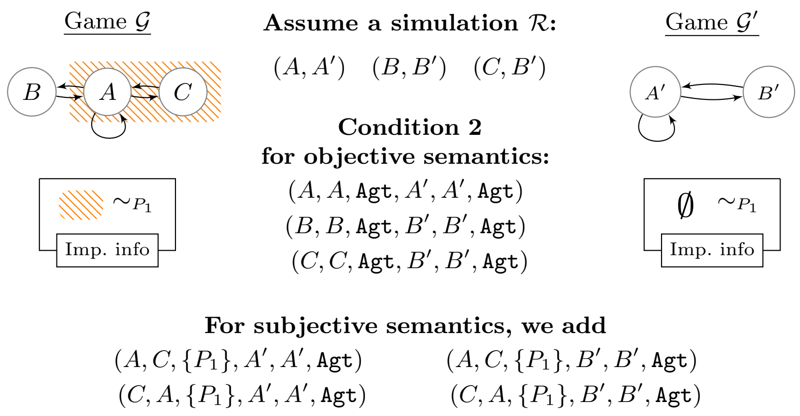

Figure 2. Illustration of Point 2 of simulation for the subjective semantic.

Simulation soundness for ATL

Theorem 2. *Let $\mathcal{R}$ be a simulation of $\mathcal{G}$ by $\mathcal{G}'$. For any $(q, q') \in \mathcal{R}$ and any $\Phi \in ATL^*$, if $q \models \Phi$ then $q' \models \Phi$ (for the objective semantics).*

Proof. Assume there is a simulation $\mathcal{R}$ of $\mathcal{G}$ by $\mathcal{G}'$. The proof is by induction on the nesting of quantifier operators. Consider the case where Φ has no nested quantification. If Φ holds on $\mathcal{G}$, then there is an existential winning strategy profile $\overline{\delta}$. Using Theorem 1, we obtain a strategy $\mathcal{S}^{\exists}(\overline{\delta})$. Then $\mathcal{S}^{\exists}(\overline{\delta})$ is a winning strategy in $\mathcal{G}'$ for the temporal property of Φ. Indeed, if there was a universal strategy $\overline{\gamma'}$ falsifying Φ against $\mathcal{S}^{\exists}(\overline{\delta})$, we could use $\mathcal{S}^{\forall}(\overline{\gamma'})$ to get a strategy falsifying the temporal property of Φ against $\overline{\delta}$, which would contradict the hypothesis that $\overline{\delta}$ is winning for Φ. The case where Φ has nested quantifications is similar, only using the induction hypothesis to check the sub-formulas. □

Simulation in the subjective semantics The notion of simulation in the subjective semantics is similar with the exception of the requirement on the tracker Λ (the second point of the definition). In the objective semantic, Point 2 provides an initialization of the tracker for the different possible starting states while Point 3.a provides a recurrence condition. Subjective semantics do not make a difference between a starting state q in $\mathcal{G}$ and a state h indistinguishable from q for some agent P. Thus a strategy δ for P must be conform to $q \sim_P h$. Something similar occurs in $\mathcal{G}'$. The tracker in a simulation between $\mathcal{G}$ and $\mathcal{G}'$ must handle this potential scenario, hence we adapt the tracker initialization (Point 2).

2. for any (q, q') in $\mathcal{R}$, any $h \in \bigcup_{P \in \mathbf{C}} \{h \mid q \sim_P h\}$, $h' \in \bigcup_{P \in \mathbf{C}} \{h' \mid q' \sim_P h'\}$, the following holds

$$(q, h, A, q', h', B) \in \Lambda \text{ where } \begin{cases} A = \{P \in \texttt{Agt} \mid q \sim_P h\} \\ B = \{P \in \texttt{Agt} \mid q' \sim_P h'\} \end{cases}$$

The proof of soundness is similar, using a definition of strategic characterization with subjective outcomes (in both features). The change in definition is illustrated in Figure 2. In the figure, we can see two games (on the left and on the right) with the imperfect information described just below (in $\mathcal{G}$, the information is for player P between A and C; in $\mathcal{G}'$ there is no imperfect information). For the relation $\mathcal{R}$, we describe the initialization of the tracker for both the objective and subjective semantics in the central part of the figure.

Remark 2. In the subjective semantics, it may be necessary to have some degree of imperfect information replication in order to establishing a simulation (some knowledge operators of epistemic logics can be expressed by subjective ATL). This is however covered through the definition : the tracker will enforce a minimum replication required.

Comparison to the existing notion of simulation Our notion is more general than the one of [5] as it needs not to reproduce similar observations. This way the game on the right of Figure 3, defined over a single (existential) agent P, is not similar for [5] to the game on the left since there is no state similar to h' in both the possibilities and the observation: r lacks the similar observation while s lacks the successor with similar label. Trivially, the games satisfy the same formulas with existential quantification over the single agent P. The two games are also similar for our notion. Indeed, we can build a relation $\mathcal{R}$ with (q, q'), (r, r'), (s, s') and (r, h'). The Λ relation follows trivially with $(*, *, P, *', *', P)$ for $(*, *') \in \mathcal{R}$ and $(r, s, \emptyset, h', s', P)$. Take $\mathcal{T}$ as the identity function plus $\mathcal{T}_{h',h}(2) \mapsto 0$. With this choice, the fourth condition is trivially satisfied.

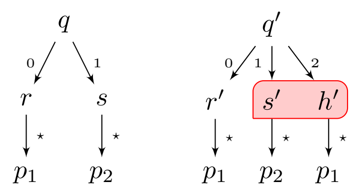

Figure 3. Two games similar with common observation in color.

4 Naive games and completeness

As ATL with perfect recall is undecidable [21], it is very unlikely that there exists a notion of simulation provably sound and complete for ATL. There exist some model restrictions which make the ATL model-checking decidable: hierarchical observations and the many derivatives (hierarchical information, dynamic hierarchies) [21,8], public actions [6]. The search for completeness relative to these fragments is not a vain quest, unlike the general case. In this section we identify a small subclass of games, *naive games*, for which our concept of simulation is complete. This concept will also prove itself crucial to develop an equivalence-checking algorithm in games with public actions in the next section. A game is naive when the imperfect information is state-based, meaning that two states can or cannot be distinguished by the same agents regardless of the histories; a

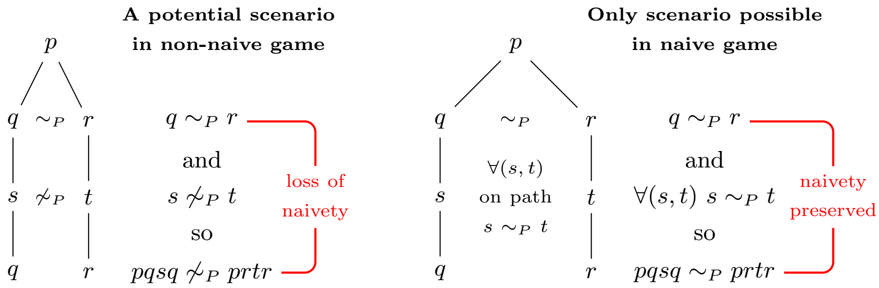

Figure 4. History influence on partial observation in both non-naive and naive games.

formal definition is given below and an illustration in Figure 4. From the definition, any game with a tree-shape structure is de-facto naive (see Figure 3 for example). This approach (restriction) on imperfect information is also used in the MCMAS tool [18].

Definition 3. *A* naive game *is a game in which for any two finite paths* ρ_A, ρ_B*,*

$$\{P \in \texttt{Agt} \mid \rho_A \sim_P \rho_B\} = \{P \in \texttt{Agt} \mid last(\rho_A) \sim_P last(\rho_B)\}$$

Note that the left-to-right inclusion is always true in *CGS*, naive games guarantee that the converse inclusion (right-to-left) also holds. Naive games are interesting for simulations because they have a very simplified tracker. The inputs are all of shape (h, k, A, h', k', B) where $A = \{P \in \texttt{Agt} \mid h \sim_P k\}$ and $B = \{P \in \texttt{Agt} \mid h' \sim_P k'\}$ whereas general inputs for non-naive games can also be of shape (h, k, C, h', k', D) with $C \subsetneq A$ and $D \subsetneq B$. They are incomparable with both games with public actions and games with hierarchical observations. On them, ATL model-checking is decidable.

Theorem 3. *ATL and ATL* model-checking are decidable on naive games with imperfect information.*

Proof (Sketch). The result is relatively trivial so we only provide a sketch of the proof. Transform the temporal objective into a parity automaton $\mathcal{A}$ and cross it with the *CGS*. Let $\mathcal{G}_{\mathcal{A}}$ be the result. We get a parity game with imperfect information for which the property of naive games still applies. On $\mathcal{G}_{\mathcal{A}}$, optimal strategies can be chosen positional even if we allow perfect recall strategies. This is because the imperfect information is fixed and will not evolve with the choices made previously by either player. We can then simply enumerate the positional strategies conform to imperfect information in $\mathcal{G}_{\mathcal{A}}$ and see if some works. □

Proving the completeness of our simulation on non-naive games seems an herculean task. It requires to build a formula which can fully encode all scenarios possible from an initial state. Such formula would require to not only handle the atomic propositions seen along the way but also the potential changes in imperfect information with other paths. With naive games, there are no changes in the imperfect information. This brings us back to a situation close to games with perfect information for which there exist sound and complete notions of alternating simulations [3]. Using similar ideas to the ones used to prove the completeness of alternating simulations for ATL, we prove that our simulations are complete for naive games.

Theorem 4. *Fix two naive games $\mathcal{G}$ and $\mathcal{G}'$. Let $\mathcal{R}$ be the set*

$$\mathcal{R} \ := \ \big\{(q, q') \mid q \in S,\ q' \in S' \ s.t\ \forall \phi \in ATL\ [q \models \phi \Rightarrow q' \models \phi]\big\}$$

then $\mathcal{R}$ is the domain of a simulation.

5 Equivalence checking in games with public actions

Games with public actions are games on which agents have perfect visibility of the other agents actions. On them, ATL enjoys a decidable model-checking [6]. Using the completeness of our simulations for naive games, we develop a sound and complete algorithm to check simulations on public action games.

Definition 4. *A game $\mathcal{G}$ has public actions when*

$$\left.\begin{array}{l}\forall P \in \texttt{Agt}\\ \forall q, q' \in \mathcal{G}\\ \forall \overline{a}, \overline{a'} \in \texttt{Act}^{\texttt{Agt}}\end{array}\right\} \quad \big[\overline{a} \neq \overline{a'} \text{ and } q \sim_P q' \ \Rightarrow\ \delta(q, \overline{a}) \not\sim_P \delta(q', \overline{a'})\big]$$

From the definition, any two histories of equal length are distinguishable as long as they start in the same initial state. So, in the objective semantics, games with public actions are equivalent to games with perfect information. Games with public actions are only interesting in that semantics if multiple starting states are considered. In the setting of this paper, it corresponds to using subjective semantics. In such cases, games with public actions are strictly more expressive than perfect information games. For the rest of this section we fix a game $\mathcal{G}$ with public actions and a coalition C of agents.

Lemma 1 (Consequence of Remark 2 in [6]). *Consider a strategy profile $\overline{\delta_C}$ for the coalition C, a starting state q, and a finite path ρ compatible with $\overline{\delta_C}$ starting in q. Then ρ has at most $|\{q' \mid q' \sim_P q, P \in C\}|$ outcomes indistinguishable from ρ in $\texttt{Out}_{sub}(\overline{\delta_C}, q)$.*

Intuitively, there is only a finite number of paths indistinguishable from the "objective" path. Each of theses paths can be identified by its starting state (within $\{q' \mid q' \sim_P q, P \in C\}$) and the sequence of actions played (common to all these paths).

So, as there are only a finite number of paths indistinguishable, we can track them easily within the state space. By doing so, we can go from public action games to naive games; this is what the lemma below does. In it we call an ATL formula principal when it has no closed sub-formula.

Theorem 5. *For each public action game $\mathcal{G}$, there exists a naive game $\mathcal{H}$ such that $\mathcal{G}$ and $\mathcal{H}$ satisfy exactly the same ATL principal formulas existentially quantifying over the coalition* $\mathbf{C}$ *of agents.*

Construction of the naive game

$\mathcal{H}$ is a version of $\mathcal{G}$ which records all possible paths indistinguishable from the current one for each agent. Each indistinguishable path will be summarized by the starting and finishing states. Each state q in $\mathcal{G}$ is augmented with a function $f : \mathtt{Agt} \mapsto 2^{G\times G}$, making the state space of $\mathcal{H}$ equal to $G\times(2^{G\times G})^{\mathtt{Agt}}$. Intuitively, if a path ends in a state q augmented by f_q with $(r, s) \in f_q(P)$, then it means there is a path indistinguishable from the current one starting in r and ending in s.

Remark 3. The construction can be seen as building an information set of a tree automaton for games with perfect information [10,19,20].

Formally, the state space of $\mathcal{H}$ is $G \times (2^{G\times G})^{\mathtt{Agt}}$. For each joint action $\overline{a}$ for $\mathtt{Agt}$, we create an edge from (q, f) to (q', f') when

- $q \xrightarrow{\overline{a}} q'$ in $\mathcal{G}$
- $f'(P) := \{(r, s') \mid \exists (r, s) \in f(P) \text{ and } s' \sim_P q' \text{ and } s \xrightarrow{\overline{a}} s'\}$ for every agent P.

The imperfect information is created inductively. Initially, it follows from $q \sim_P q'$ in $\mathcal{G}$ that

$$(q, f : P \mapsto \{(r, r) \mid r \sim_P q\})) \sim_P (q', f' : P \mapsto \{(r', r') \mid r' \sim_P q'\}) \text{ in } \mathcal{H}, \quad (2)$$

then inductively,

$$\left.\begin{array}{ll} (q, f) \sim_P (q', f') & (q, f) \xrightarrow{\overline{a}} (r, g) \\ (q', f') \xrightarrow{\overline{a}} (r', g') & r \sim_P r' \text{ in } \mathcal{G} \end{array}\right\} \Rightarrow \quad (r, g) \sim_P (r', g') \quad (3)$$

The induction trivially reaches a fixed point and terminates. The initial relation is reflexive (inherited from the relation on $\mathcal{G}$), symmetric (by definition) and transitive (inherited from the relation on $\mathcal{G}$ and the definition). At each step of the induction, these three properties are preserved. Indeed reflexivity is trivially preserved. The definition of (3) is symmetric, so the relation is also symmetric. Finally, the transitivity is preserved through the use of similar joint actions, as in lines 2 and 3 in (3). The relation thus defined is indeed an equivalence relation on states of $\mathcal{H}$ and therefore an imperfect information relation.

The set of initial states we consider in $\mathcal{H}$ is $\{(q, f) \mid f : P \mapsto \{(r, r) \mid r \sim_P q\}\}$. By definition of the imperfect information in $\mathcal{H}$:

$$\forall (q, f), (q', f') \in \mathcal{H}.\ \forall P \in \mathtt{Agt}. \qquad \big[(q, f) \sim_P (q', f') \text{ in } \mathcal{H} \Rightarrow q \sim_P q' \text{ in } \mathcal{G}\big] \quad (4)$$

The idea is partially illustrated in Figure 5, with the public-actions game on the left and the naive game on the right. The functions f and g are described at the top.

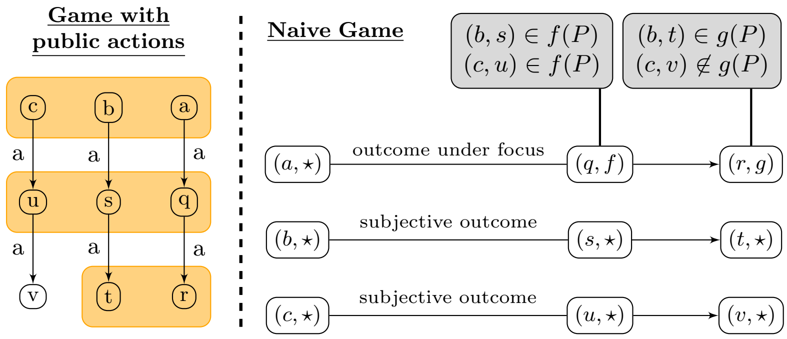

Figure 5. Construction (with a single agent P). The imperfect information on the public action game is represented by colored areas.

Correctness of the construction

Notations: In the following we write a state $(q, \star)$ of $\mathcal{H}$ for a pair of shape (q, f) for some function f, and write a state of $\mathcal{H}$ $(\star, f)$ for a pair (q, f) for some state q of $\mathcal{G}$. This allows us to ease the reading.

Lemma 2. *The following holds for any two paths ρ, ρ' and any agent P. Write $\rho'(1) := (q', \star)$, $\rho(|\rho|) := (\star, g)$ and $\rho'(|\rho|) := (s', \star)$. Then*

I *if $\rho \sim_P \rho'$ then $(q', s') \in g(P)$.*
II *if $\rho \not\sim_P \rho'$ then $(q', s') \notin g(P)$.*

Lemma 3. *Let P be any agent, (r, f) and (t, g) be any two states such that $(r, f) \sim_P (t, g)$. There are two paths ρ_C and ρ_D of shapes $\rho_C : (u, \star) \mapsto^* (r, f)$ and $\rho_D := (v, \star) \mapsto^* (t, g)$ such that $(v, t) \in f(P)$ and $(u, r) \in g(P)$.*

Lemma 4. *$\mathcal{H}$ is a naive game.*

Proof. Toward a contradiction, suppose the game is not naive. Then there must be two finite paths ρ_A, ρ_B and an agent $P \in \texttt{Agt}$ such that $\rho_A \not\sim_P \rho_B$ but $last(\rho_A) \sim_P last(\rho_B)$. Write $\texttt{last}(\rho_A) = (r, f)$ and $\texttt{last}(\rho_B) = (t, g)$. Since $(r, f) \sim_P (t, g)$, by Lemma 3, there are two paths ρ_C and ρ_D of the shapes $\rho_C : (u, \star) \mapsto^* (r, f)$ and $\rho_D := (v, \star) \mapsto^* (t, g)$ such that $(v, t) \in f(P)$ and $(u, r) \in g(P)$. By Lemma 2, since $(v, t) \in f(P)$, we have $\rho_A \sim_P \rho_D$. Then by Lemma 2 once again, since $last(\rho_D) = (t, g)$, we have $(first(\rho_A), r) \in g$. Applying one last time Lemma 2, since $(first(\rho_A), r) \in g$ and $last(\rho_B) = (t, g)$ we get $\rho_A \sim_P \rho_B$, which is a contradiction. □

Lemma 5. *A principal formula $\phi \in ATL$ existentially quantifying* **C** *holds from state q in $\mathcal{G}$ if and only if ϕ holds from (q, f) in $\mathcal{H}$ with $f : P \mapsto \{(r, r) \mid r \sim_P q\}$.*

Proof. For this we simply show an equivalence between paths in $\mathcal{G}$ and paths in $\mathcal{H}$ (from the starting states), in which a state q in $\mathcal{G}$ is always linked to a state of shape $(q, \star)$ in $\mathcal{H}$. We proceed by induction on the length of the paths. First note that for a state q in $\mathcal{G}$ there is a single initial state (q, f) in $\mathcal{H}$. We can therefore establish an equivalence between starting states. For the induction case, consider a path ρ in $\mathcal{H}$ and π in $\mathcal{H}$ and write $last(\rho) = (q, f)$ and $\mathtt{last}(\pi) = q$. For each joint-action $\overline{a}$ there is a single q' such that $q \xrightarrow{\overline{a}} q'$ in $\mathcal{G}$, and a single (q', f') such that $(q, f) \xrightarrow{\overline{a}} (q', f')$. We can therefore extend the correspondence one more step. And with the induction step sorted out, we can conclude the existence of a one-to-one correspondence between paths in both $\mathcal{G}$ and $\mathcal{H}$. Through a simple induction, we obtain that two paths (from the starting state) are indistinguishable in $\mathcal{H}$ if and only if their counterparts in $\mathcal{G}$ are indistinguishable. From the path correspondence, it is trivial to establish a correspondence between conform strategies, and to establish an equivalence between the formulas that can be satisfied (as long as we start from the appropriate starting state in $\mathcal{H}$). □

Theorem 5 then follows from the construction and Lemmas 4 and 5.

Sound and complete checking of public-actions games

By combining Theorems 4 and 5, we can obtain a sound and complete way for ATL principal formulas to check simulation on public-action games. The process is presented in Algorithm 1. The correctness of the algorithm is ensured by the following lemma whose proof is in annex:

Lemma 6. *In Algo 1, define $\mathcal{R}_{\mathcal{H}}$ as the largest simulation of $\mathcal{H}'$ by $\mathcal{H}$. Then*

$$\{(q, q') \mid \forall \phi \ \ \textit{principal in ATL} \ [q \models \phi \Rightarrow q' \models \phi]\}$$
$$= \left\{(q, q') \mid \exists \ f_{ini}, f'_{ini} \ s.t \begin{cases} (q, f_{ini}) \ \textit{is an initial state of} \ \mathcal{H} \\ (q', f'_{ini}) \ \textit{is an initial state of} \ \mathcal{H}' \\ ((q, f_{ini}), (q', f'_{ini})) \in \mathcal{R}_{\mathcal{H}} \end{cases} \right\}$$

Algorithm 1 does not work for non-principal ATL formulas. Indeed, two elements in the simulation relation $\mathcal{R}$ may not be starting states of $\mathcal{H}$, and therefore the correctness which only applies from starting states may not hold. The lemma below tells us precisely when our algorithm extends to non-principal formulas.

Lemma 7. *In Algorithm 1, if $\mathcal{R}$ satisfies*

$$\forall((r, f), (s, g)) \in \mathcal{R} \qquad \exists f_{ini}, g_{ini} \ \textit{such that} \begin{cases} ((r, f_{ini}), (s, g_{ini})) \in \mathcal{R} \\ (r, f_{ini}) \ \textit{is an initial state in} \ \mathcal{H} \\ (s, g_{ini}) \ \textit{is an initial state in} \ \mathcal{H} \end{cases}$$

Algorithm 1 Check for principal formulas in public-action games.

INPUT: Two games $\mathcal{G}$ and $\mathcal{G}'$ and two initial states q, q' respectively in $\mathcal{G}$ and $\mathcal{G}'$.
OUTPUT: Does $\mathcal{G}$ and $\mathcal{G}'$ satisfy the same principal formulas from q and q'.

1: $\mathcal{H} \rightarrow$ naive game satisfying the same ATL formulas as $\mathcal{G}$ through Theorem 5
2: $\mathcal{H}' \rightarrow$ naive game satisfying the same ATL formulas as $\mathcal{G}'$ through Theorem 5
3: Find the maximal simulation relation $\mathcal{R}$ of $\mathcal{H}$ by $\mathcal{H}'$
4: **if** $\exists f_{ini}, f'_{ini}$ such that $((q, f_{ini}), (q', f'_{ini})) \in \mathcal{R}$ **then**
5:     return $True$
6: **else**
7:     return $False$
8: **end if**

then

$$\{(q, q') \mid \forall\phi \quad \textit{(principal or not) in } \mathsf{ATL}\ [q \models \phi \Rightarrow q' \models \phi]\}$$

$$= \Big\{(q, q') \mid \exists\ f_{ini}, f'_{ini}\ s.t \begin{cases} (q, f_{ini}) \textit{ is an initial state} \\ (q', f'_{ini}) \textit{ is an initial state} \\ ((q, f_{ini}), (q', f'_{ini})) \in \mathcal{R}_{\mathcal{H}} \end{cases} \Big\}$$

The proof follows from the definition of the condition and Theorem 5. With the lemma above, we can develop an algorithm for non-principal formulas simply by requiring step 3 to find the maximal simulation relation $\mathcal{R}$ which satisfies the condition of the lemma above.

6 Concluding Remarks

We have proposed a notion of simulation sound for ATL on multi-agent systems in general and complete on naive systems where the information is state-based. Using the completeness of our concept of simulation for naive games, we have designed a simulation-checking algorithm for public-action games. A remaining interrogation is whether there is an equivalence-checking algorithm that is both sound and complete for ATL on hierarchical information systems.

References

1. de Alfaro, L., Henzinger, T.A.: Concurrent omega-regular games. In: Proc. LICS 2000 (2000)
2. Alur, R., Henzinger, T.A., Kupferman, O.: Alternating-time temporal logic. In: Proc. COMPOS 1997. Revised Lectures (1997)
3. Alur, R., Henzinger, T.A., Kupferman, O., Vardi, M.Y.: Alternating refinement relations. In: Proc. CONCUR 1998 (1998)
4. Baskar, A., Ramanujam, R., Suresh, S.P.: Knowledge-based modelling of voting protocols. In: Proc. TARK 2007 (2007)

5. Belardinelli, F., Condurache, R., Dima, C., Jamroga, W., Jones, A.V.: Bisimulations for verifying strategic abilities with an application to threeballot. In: Pro. AAMAS 2017 (2017)
6. Belardinelli, F., Lomuscio, A., Murano, A., Rubin, S.: Verification of multi-agent systems with imperfect information and public actions. In: Proc. AAMAS 2017. International Foundation for Autonomous Agents and Multiagent Systems (2017)
7. Berthon, R., Maubert, B., Murano, A., Rubin, S., Vardi, M.Y.: Strategy logic with imperfect information. In: Proc. LICS 2017. pp. 1–12 (2017)
8. Berwanger, D., Mathew, A.B., van den Bogaard, M.: Hierarchical information and the synthesis of distributed strategies. Acta Informatica (2018)
9. Blackburn, P., Rijke, M.d., Venema, Y.: Modal Logic. Cambridge Tracts in Theoretical Computer Science, Cambridge University Press (2001)
10. Chatterjee, K., Doyen, L.: The complexity of partial-observation parity games. vol. 6397, pp. 1–14 (10 2010)
11. Dima, C., Tiplea, F.L.: Model-checking ATL under imperfect information and perfect recall semantics is undecidable. CoRR (2011)
12. Goltz, U., Kuiper, R., Penczek, W.: Propositional temporal logics and equivalences. In: CONCUR '92 (1992)
13. Goranko, V., van Drimmelen, G.: Complete axiomatization and decidability of alternating-time temporal logic. Theor. Comput. Sci. (2006)
14. Jamroga, W., Dix, J.: Model checking abilities under incomplete information is indeed delta2-complete. In: Proc. EUMAS 2006 (2006)
15. Laroussinie, F., Markey, N., Oreiby, G.: On the expressiveness and complexity of atl. In: Seidl, H. (ed.) Proc. FOSSACS 2007 (2007)
16. Laroussinie, F., Markey, N., Sangnier, A.: Atlsc with partial observation. In: Proc. GandALF 2015 (2015)
17. Larsen, K.G., Skou, A.: Bisimulation through probabilistic testing. Information and Computation (1991)
18. Lomuscio, A., Qu, H., Raimondi, F.: Mcmas: an open-source model checker for the verification of multi-agent systems. International Journal on Software Tools for Technology Transfer (2017)
19. van der Meyden, R., Vardi, M.Y.: Synthesis from knowledge-based specifications. In: Proc. CONCUR 1998. pp. 34–49. Springer (1998)
20. van der Meyden, R., Wilke, T.: Synthesis of distributed systems from knowledge-based specifications. vol. 3653, pp. 562–576 (08 2005)
21. Peterson, G.L., Reif, J.H.: Multiple-person alternation. In: Proc. FOCS 1979 (1979)
22. Strauss, C.: A critical review of the triple ballot voting system, part2: Cracking the triple ballot encryption (2006)
23. Tabatabaei, M., Jamroga, W., Ryan, P.Y.: Expressing receipt freeness and coercion-resistance in logics of strategic ability preliminary attempt. In: Proc. PrAISe 2016 (2016)
24. VAS-Group: https://vas.doc.ic.ac.uk/software/mcmas/. In: Imperial college of London
25. Wikipedia: Three ballot voting system, https://en.wikipedia.org/wiki/threeballot
26. Zhang, C., Pang, J.: On probabilistic alternating simulations. In: Proc. TCS 2010. pp. 71–85 (2010)