

This is the peer reviewed version of the following article:

Towards the Combination of Model Checking and Runtime Verification on Multi-agent Systems / Ferrando, A., Malvone, V.. - 13616:(2022), pp. 140-152. (20th International Conference on Practical Applications of Agents and Multi-Agent Systems , PAAMS 2022 L'Aquila, ITALY JUL 13-15, 2022) [10.1007/978-3-031-18192-4_12].

Springer Science and Business Media Deutschland GmbH

Terms of use:

The terms and conditions for the reuse of this version of the manuscript are specified in the publishing policy. For all terms of use and more information see the publisher's website.

26/06/2026 22:39

(Article begins on next page)

Towards the Combination of Model Checking and Runtime Verification on Multi-Agent Systems

Angelo Ferrando¹ and Vadim Malvone²

¹University of Genova, Italy

²Télécom Paris, France

angelo.ferrando@unige.it, vadim.malvone@telecom-paris.fr

Abstract

Multi-Agent Systems (MAS) are notoriously complex and hard to verify. In fact, it is not trivial to model a MAS, and even when a model is built, it is not always possible to verify, in a formal way, that it is actually behaving as we expect. Usually, it is relevant to know whether an agent is capable of fulfilling its own goals. One possible way to check this is through Model Checking. Specifically, by verifying Alternating-time Temporal Logic (ATL) properties, where the notion of strategies for achieving goals can be described. Unfortunately, the resulting model checking problem is not decidable in general. In this paper, we present a verification procedure based on combining Model Checking and Runtime Verification, where sub-models of the MAS model belonging to decidable fragments are verified by a model checker, and runtime monitors are used to verify the rest. We present our technique and we show experimental results.

1 Introduction

Intelligent systems, such as Multi-Agent Systems (MAS), can be seen as a set of intelligent entities capable of proactively decide how to act to fulfill their own goals. These entities, called generally agents, are notoriously autonomous, *i.e.*, they do not expect input from an user to act, and social, *i.e.*, they usually communicate amongst each other to achieve common goals.

Software systems are not easy to trust in general. This is especially true in the case of complex and distributed systems, such as MAS. Because of this, we need verification techniques to verify that such systems behave as expected. More specifically, in the case of MAS, it is relevant to know whether the agents are capable of achieving their own goals, by themselves or by collaborating with other agents by forming a coalition. This is usually referred to as the process of finding a strategy for the agent(s).

A well-known formalism for reasoning about strategic behaviours in MAS is Alternating-time Temporal Logic (*ATL*) [1]. Before verifying *ATL* specifications, two questions need to be answered: (i) *does each agent know everything about the system?* (ii) *does the property require the agent to have memory of the system?* The first question concerns the model of the MAS. If each agent can distinguish each state of the model, then we have *perfect information*; otherwise, we have *imperfect information*.

The second question concerns the *ATL* property. If the property can be verified without the need for the agent to remember which states of the model have been visited before, then we have *imperfect recall*; otherwise, we have *perfect recall*.

The model checking problem for *ATL* giving a generic MAS is known to be undecidable. This is due to the fact that the model checking problem for *ATL* specifications under imperfect information and perfect recall has been proved to be undecidable [2]. Nonetheless, decidable fragments exist. Indeed, model checking *ATL* under perfect information is PTIME-complete [1], while under imperfect information and imperfect recall is PSPACE [3]. Unfortunately, MAS usually have imperfect information, and when memory is needed to achieve the goals, the resulting model checking problem becomes undecidable. Given the relevance of the imperfect information setting, even partial solutions to the problem are useful.

This is not the first time that a verification technique alone is not enough to complete the wanted task. Specifically, even if the verification of the entire model is not possible, there might still be sub-models of the model for which it is. Consequently, we could focus on these sub-models for which the model checking problem is still decidable; which are the sub-models with perfect information and perfect recall strategies. With more detail, given an *ATL* formula φ and a model of MAS M , our procedure extracts all the sub-models of M with perfect information that satisfy a sub-formula of φ . After this step, runtime monitors are used to check if the remaining part of φ can be satisfied at execution time. If this is the case, we can conclude at runtime the satisfaction of φ for the corresponding system execution. This is determined by the fact that the system has been observed behaving as expected, since it has verified at design time the sub-formula ψ of φ , and at runtime the remaining temporal part of φ (which consists in the part left to verify in φ , not covered by ψ). Note that, this does not imply that the system satisfies φ , indeed future executions may violate φ . The formal result over φ only concerns the current system execution, and how it has behaved in it. However, we will present preservation results on the initial model checking problem of φ on the model of the system M , as well. This will be obtained by linking the result obtained at runtime, with its static counterpart. Hence, we are going to show how the satisfaction (resp., violation) of φ at runtime in our approach can be propagated to the verification question over φ on model M . Before moving on with the related works in literature, it is important to linger on the main contribution of this work. As we mentioned previously, the problem of statically verify MAS with imperfect information and using perfect recall strategies is undecidable. Thus, the work presented in this paper cannot answer the same question (*i.e.*, we are not claiming decidability for a well-known undecidable problem). Instead, it is focused on gathering and extracting more information about the MAS under analysis at runtime, through runtime verification. This information can be used to better understand the system, and it is an improvement w.r.t. the undecidability of the original problem.

The intuition behind this work lies behind the relation amongst what can be observed at execution time (runtime), and what can be concluded at design time (statically). To the best of our knowledge, no such relation has ever been explored before in the strategic scenario. Usually, static verification of MAS mainly consists in verifying whether strategies for the agents exist to achieve some common goal (expressed as some sort of temporal property enriched with strategic flavour). Even though the two formal verification techniques may seem completely orthogonal, they are very close to each other. In fact, standard runtime verification of temporal properties (such as LTL) consists, in a certain way, in applying model checking at runtime over the all possible executions of a system (whose model may not be available). For the verification

of strategic properties as well such relation holds. However, because of the gap between the linearity of the properties verifiable by a runtime monitor, and the branching behaviour of strategic properties, the results that can be obtained through runtime verification are not so natural to propagate to the corresponding model checking problem. Which means, it is not obvious, given a result at runtime, to know what to conclude on the corresponding static verification problem. This is of paramount difference w.r.t. LTL, where a runtime violation can be propagated to a violation of the model checking problem as well. Nonetheless, as we are going to show in this paper, also for strategic properties it is possible to use runtime verification to propagate results on the initial model checking problem. In a nutshell, since it will be better clarified in the due course, static verification of strategic properties over a MAS consists in checking whether a strategy for a set of agents (coalition) can be used to achieve a common (temporal) goal. Now, this is done by analysing, through model checking, the possible executions inside the model in accordance with the strategies for the coalition. Even though at runtime such thorough analysis cannot be done, the observation of an execution of the system at runtime can bring much information. For instance, let us say that the current system execution satisfies the temporal property (the goal, without considering the strategic aspects). Then, this means that the agents at runtime were capable (at least once) to collaborate with each other to achieve a common goal (the temporal property). Note that, this does not imply that the agents will always behave (we are still not exhaustive at runtime), but gives us a vital information about the system: “if the agents want to achieve the goal, they can”. This runtime outcome can be propagated back to the initial model checking problem, and helps us to conclude the satisfaction of the strategic property when all the agents are assumed to collaborate (one single big coalition). Naturally, it might be possible that even with smaller coalitions the goal would still be achievable, but this is something that cannot be implicated with the only runtime information. On the other hand, if at runtime we observe a wrong behaviour, it means the agents were not capable of achieving the goal. Since we cannot claim which (if any) coalitions were actually formed to achieve the goal, we cannot assume that it is not possible with a greater coalition to achieve the goal. In fact, two scenarios are possible. 1) The agents did not form any coalition (each agent works alone). 2) The agents did form a coalition, but this was not enough to achieve the goal. In both cases, there is a common result that can be propagated back to the initial model checking problem, which is that without cooperating the agents cannot achieve the goal. This is true in case (1), since it is what has actually happened at runtime, and it is also true in (2), since by knowing that cooperating (at a certain level) is not enough to achieve the goal, it is also true that with less cooperation the same goal cannot be achieved neither. Note that, this does not imply that the agents will always wrongly behave, indeed with a greater coalition of agents it might still be possible to conclude the goal achievement. The vital information obtained in this way at runtime can be rephrased as: “if the agents do not cooperate, they cannot achieve the goal”.

2 Related Work

Model Checking on MAS. Several approaches for the verification of specifications in *ATL* and *ATL** under imperfect information and perfect recall have been recently put forward. In one line, restrictions are made on how information is shared amongst the agents, so as to retain decidability [4, 5]. In a related line, interactions amongst agents are limited to public actions only [6, 7]. These approaches are markedly dif-

ferent from ours as they seek to identify classes for which verification is decidable. Instead, we consider the whole class of iCGS and define a general verification procedure. In this sense, existing approaches to approximate *ATL* model checking under imperfect information and perfect recall have either focused on an approximation to perfect information [8, 9] or developed notions of bounded recall [10]. Related to bounded strategies, in [11] the notion of natural strategies is introduced and in [12] is provided a model checking solution for a variant of *ATL* under imperfect information.

Differently from these works, we introduce, for the first time, a technique that couples model checking and runtime verification to provide results. Furthermore, we always concludes with a result. Note that the problem is undecidable in general, thus the result might be inconclusive (but it is always returned). When the result is inconclusive for the whole formula, we present sub-results to give at least the maximum information about the satisfaction/violation of the formula under exam.

Runtime Verification. Runtime Verification (RV) has never been used before in a strategic context, where monitors check whether a coalition of agents satisfies a strategic property. This can be obtained by combining Model Checking on MAS with RV. The combination of Model Checking with RV is not new; in a position paper dating back to 2014, Hinrichs et al. suggested to “model check what you can, runtime verify the rest” [13]. Their work presented several realistic examples where such mixed approach would give advantages, but no technical aspects were addressed. Desai et al. [14] present a framework to combine model checking and runtime verification for robotic applications. They represent the discrete model of their system and extract the assumptions deriving from such abstraction. Kejstová et al. [15] extended an existing software model checker, DIVINE [16], with a runtime verification mode. The system under test consists of a user program in C or C++, along with the environment. Other blended approaches exist, such as a verification-centric software development process for Java making it possible to write, type check, and consistency check behavioural specifications for Java before writing any code [17]. Although it integrates a static checker for Java and a runtime assertion checker, it does not properly integrate model checking and RV. In all the previously mentioned works, both Model Checking and RV were used to verify temporal properties, such as LTL. Instead, we focus on strategic properties, we show how combining Model Checking of *ATL** properties with RV, and we can give results; even in scenarios where Model Checking alone would not suffice. Because of this, our work is closer in spirit to [13]; in fact, we use RV to support Model Checking in verifying at runtime what the model checker could not at static time. Finally, in [18], a demonstration paper presenting the tool deriving by this work may be found. Specifically, in this paper we present the theoretical foundations behind the tool.

3 Preliminaries

In this section we recall some preliminary notions. Given a set U , \bar{U} denotes its complement. We denote the length of a tuple v as $|v|$, and its i -th element as v_i . For $i \leq |v|$, let $v_{\geq i}$ be the suffix $v_i, \dots, v_{|v|}$ of v starting at v_i and $v_{\leq i}$ the prefix v_1, \dots, v_i of v . We denote with $v \cdot w$ the concatenation of the tuples v and w .

3.1 Models for Multi-agent systems

We start by giving a formal model for Multi-agent Systems by means of concurrent game structures with imperfect information [1, 19].

Definition 1. A concurrent game structure with imperfect information (iCGS) is a tuple $M = \langle Ag, AP, S, s_I, \{Act_i\}_{i \in Ag}, \{\sim_i\}_{i \in Ag}, d, \delta, V \rangle$ such that:

- $Ag = \{1, \dots, m\}$ is a nonempty finite set of agents (or players).
- AP is a nonempty finite set of atomic propositions (atoms).
- $S \neq \emptyset$ is a finite set of states, with initial state $s_I \in S$.
- For every $i \in Ag$, Act_i is a nonempty finite set of actions. Let $Act = \bigcup_{i \in Ag} Act_i$ be the set of all actions, and $ACT = \prod_{i \in Ag} Act_i$ the set of all joint actions.
- For every $i \in Ag$, \sim_i is a relation of indistinguishability between states. That is, given states $s, s' \in S$, $s \sim_i s'$ iff s and s' are observationally indistinguishable for agent i .
- The protocol function $d : Ag \times S \rightarrow (2^{Act} \setminus \emptyset)$ defines the availability of actions so that for every $i \in Ag$, $s \in S$, (i) $d(i, s) \subseteq Act_i$ and (ii) $s \sim_i s'$ implies $d(i, s) = d(i, s')$.
- The (deterministic) transition function $\delta : S \times ACT \rightarrow S$ assigns a successor state $s' = \delta(s, \vec{a})$ to each state $s \in S$, for every joint action $\vec{a} \in ACT$ such that $a_i \in d(i, s)$ for every $i \in Ag$, that is, \vec{a} is enabled at s .
- $V : S \rightarrow 2^{AP}$ is the labelling function.

By Def. 1 an iCGS describes the interactions of a group Ag of agents, starting from the initial state $s_I \in S$, according to the transition function δ . The latter is constrained by the availability of actions to agents, as specified by the protocol function d . Furthermore, we assume that every agent i has imperfect information of the exact state of the system; so in any state s , i considers epistemically possible all states s' that are i -indistinguishable from s [20]. When every \sim_i is the identity relation, i.e., $s \sim_i s'$ iff $s = s'$, we obtain a standard CGS with perfect information [1].

Given a set $\Gamma \subseteq Ag$ of agents and a joint action $\vec{a} \in ACT$, let \vec{a}_Γ and $\vec{a}_{\bar{\Gamma}}$ be two tuples comprising only of actions for the agents in Γ and $\bar{\Gamma}$, respectively.

A history $h \in S^+$ is a finite (non-empty) sequence of states. The indistinguishability relations are extended to histories in a synchronous, point-wise way, i.e., histories $h, h' \in S^+$ are *indistinguishable* for agent $i \in Ag$, or $h \sim_i h'$, iff (i) $|h| = |h'|$ and (ii) for all $j \leq |h|$, $h_j \sim_i h'_j$.

3.2 Syntax

To reason about the strategic abilities of agents in iCGS with imperfect information, we use Alternating-time Temporal Logic ATL^* [1].

Definition 2. State (φ) and path (ψ) formulas in ATL^* are defined as follows, where $q \in AP$ and $\Gamma \subseteq Ag$:

$$\begin{aligned} \varphi &::= q \mid \neg\varphi \mid \varphi \wedge \varphi \mid \langle\langle \Gamma \rangle\rangle\psi \\ \psi &::= \varphi \mid \neg\psi \mid \psi \wedge \psi \mid X\psi \mid (\psi U \psi) \end{aligned}$$

Formulas in ATL^* are all and only the state formulas.

As customary, a formula $\langle\langle\Gamma\rangle\rangle\Phi$ is read as “the agents in coalition Γ have a strategy to achieve Φ ”. The meaning of linear-time operators *next* X and *until* U is standard [21]. Operators $\llbracket\Gamma\rrbracket$, *release* R , *finally* F , and *globally* G can be introduced as usual. Formulas in the *ATL* fragment of *ATL** are obtained from Def. 2 by restricting path formulas ψ as follows, where φ is a state formula and R is the *release* operator:

$$\psi ::= X\varphi \mid (\varphi U\varphi) \mid (\varphi R\varphi)$$

In the rest of the paper, we will also consider the syntax of *ATL** in negative normal form (NNF):

$$\begin{aligned} \varphi &::= q \mid \neg q \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \langle\langle\Gamma\rangle\rangle\psi \mid \llbracket\Gamma\rrbracket\psi \\ \psi &::= \varphi \mid \psi \wedge \psi \mid \psi \vee \psi \mid X\psi \mid (\psi U\psi) \mid (\psi R\psi) \end{aligned}$$

where $q \in AP$ and $\Gamma \subseteq Ag$.

3.3 Semantics

When giving a semantics to *ATL** formulas we assume that agents are endowed with *uniform strategies* [19], *i.e.*, they perform the same action whenever they have the same information.

Definition 3. A uniform strategy for agent $i \in Ag$ is a function $\sigma_i : S^+ \rightarrow Act_i$ such that for all histories $h, h' \in S^+$, (i) $\sigma_i(h) \in d(i, last(h))$; and (ii) $h \sim_i h'$ implies $\sigma_i(h) = \sigma_i(h')$.

By Def. 3 any strategy for agent i has to return actions that are enabled for i . Also, whenever two histories are indistinguishable for i , then the same action is returned. Notice that, for the case of CGS (perfect information), condition (ii) is satisfied by any strategy σ . Furthermore, we obtain memoryless (or imperfect recall) strategies by considering the domain of σ_i in S , *i.e.*, $\sigma_i : S \rightarrow Act_i$.

Given an iCGS M , a path $p \in S^\omega$ is an infinite sequence $s_1 s_2 \dots$ of states. Given a joint strategy $\sigma_\Gamma = \{\sigma_i \mid i \in \Gamma\}$, comprising of one strategy for each agent in coalition Γ , a path p is σ_Γ -compatible iff for every $j \geq 1$, $p_{j+1} = \delta(p_j, \vec{a})$ for some joint action \vec{a} such that for every $i \in \Gamma$, $a_i = \sigma_i(p_{\leq j})$, and for every $i \in \bar{\Gamma}$, $a_i \in d(i, p_j)$. Let $out(s, \sigma_\Gamma)$ be the set of all σ_Γ -compatible paths from s .

We can now assign a meaning to *ATL** formulas on iCGS.

Definition 4. The satisfaction relation \models for an iCGS M , state $s \in S$, path $p \in S^\omega$, atom $q \in AP$, and *ATL** formula ϕ is defined as follows:

$$\begin{aligned} (M, s) \models q & \text{ iff } q \in V(s) \\ (M, s) \models \neg\varphi & \text{ iff } (M, s) \not\models \varphi \\ (M, s) \models \varphi \wedge \varphi' & \text{ iff } (M, s) \models \varphi \text{ and } (M, s) \models \varphi' \\ (M, s) \models \langle\langle\Gamma\rangle\rangle\psi & \text{ iff for some } \sigma_\Gamma, \text{ for all } p \in out(s, \sigma_\Gamma), (M, p) \models \psi \\ (M, p) \models \varphi & \text{ iff } (M, p_1) \models \varphi \\ (M, p) \models \neg\psi & \text{ iff } (M, p) \not\models \psi \\ (M, p) \models \psi \wedge \psi' & \text{ iff } (M, p) \models \psi \text{ and } (M, p) \models \psi' \\ (M, p) \models X\psi & \text{ iff } (M, p_{\geq 2}) \models \psi \\ (M, p) \models \psi U\psi' & \text{ iff for some } k \geq 1, (M, p_{\geq k}) \models \psi', \text{ and} \\ & \text{ for all } j, 1 \leq j < k \Rightarrow (M, p_{\geq j}) \models \psi \end{aligned}$$

We say that formula ϕ is *true* in an iCGS M , or $M \models \phi$, iff $(M, s_I) \models \phi$.
We now state the model checking problem.

Definition 5. *Given an iCGS M and a formula ϕ , the model checking problem concerns determining whether $M \models \phi$.*

Since the semantics provided in Def. 4 is the standard interpretation of ATL^* [1, 19], it is well known that model checking ATL , *a fortiori* ATL^* , against iCGS with imperfect information and perfect recall is undecidable [2]. In the rest of the paper we develop methods to obtain partial solutions to this by using Runtime Verification (RV).

3.4 Runtime verification and Monitors

Given a nonempty set of atomic propositions AP , we define a *trace* $\rho = ev_1 ev_2 \dots$, as a sequence of set of events in AP (i.e., for each i we have that $ev_i \in 2^{AP}$). For brevity, we name $\Sigma = 2^{AP}$ the powerset of atomic propositions. As usual, Σ^* is the *set of all possible finite traces* over Σ , and Σ^ω is the *set of all possible infinite traces* over Σ .

The standard formalism to specify formal properties in RV is Linear Temporal Logic (LTL) [22]. The syntax of LTL is as follows:

$$\psi ::= q \mid \neg\psi \mid \psi \wedge \psi \mid X\psi \mid (\psi U \psi)$$

where $q \in AP$ is an event (a proposition), ψ is a formula, U stands for *until*, and X stands for *next-time*.

Let $\rho \in \Sigma^\omega$ be an infinite sequence of events over Σ , the semantics of LTL is as follows:

$$\begin{aligned} \rho \models q & \quad \text{iff} \quad q \in \rho_1 \\ \rho \models \neg\psi & \quad \text{iff} \quad \rho \not\models \psi \\ \rho \models \psi \wedge \psi' & \quad \text{iff} \quad \rho \models \psi \text{ and } \rho \models \psi' \\ \rho \models X\psi & \quad \text{iff} \quad \rho_{\geq 2} \models \psi \\ \rho \models \psi U \psi' & \quad \text{iff} \quad \text{for some } k \geq 1, \rho_{\geq k} \models \psi', \text{ and for all } j, 1 \leq j < k \Rightarrow \rho_{\geq j} \models \psi \end{aligned}$$

Thus, given an LTL property ψ , we denote $\llbracket \psi \rrbracket$ the language of the property, i.e., the set of traces which satisfy ψ ; namely $\llbracket \psi \rrbracket = \{\rho \mid \rho \models \psi\}$.

Definition 6 (Monitor). *Let AP be the alphabet of atomic propositions, $\Sigma = 2^{AP}$ be its powerset, and ψ be an LTL property. Then, a monitor for ψ is a function $Mon_\psi : \Sigma^* \rightarrow \mathbb{B}_3$, where $\mathbb{B}_3 = \{\top, \perp, ?\}$:*

$$Mon_\psi(\rho) = \begin{cases} \top & \forall \rho' \in \Sigma^\omega \quad \rho \cdot \rho' \in \llbracket \psi \rrbracket \\ \perp & \forall \rho' \in \Sigma^\omega \quad \rho \cdot \rho' \notin \llbracket \psi \rrbracket \\ ? & \text{otherwise.} \end{cases}$$

Intuitively, a monitor returns \top if all continuations (ρ') of ρ satisfy ψ ; \perp if all possible continuations of ρ violate ψ ; $?$ otherwise. The first two outcomes are standard representations of satisfaction and violation, while the third is specific to RV. In more detail, it denotes when the monitor cannot conclude any verdict yet. This is closely related to the fact that RV is applied while the system is still running, and not all information about it are available. For instance, a property might be currently satisfied (resp., violated) by the system, but violated (resp., satisfied) in the (still unknown) future. The monitor can only safely conclude any of the two final verdicts (\top or \perp) if it is sure such verdict will never change. The addition of the third outcome symbol $?$ helps the monitor to represent its position of uncertainty w.r.t. the current system execution.

3.5 Negative and Positive Sub-models

Now, we recall two definitions of sub-models, defined in [23], that we will use in our verification procedure. We start with the definition of negative sub-models.

Definition 7 (Negative sub-model). *Given an iCGS $M = \langle Ag, AP, S, s_I, \{Act_i\}_{i \in Ag}, \{\sim_i\}_{i \in Ag}, d, \delta, V \rangle$, we denote with $M_n = \langle Ag, AP, S_n, s_I, \{Act_i\}_{i \in Ag}, \{\sim_i^n\}_{i \in Ag}, d_n, \delta_n, V_n \rangle$ a negative sub-model of M , formally $M_n \subseteq M$, such that:*

- *the set of states is defined as $S_n = S^* \cup \{s_\perp\}$, where $S^* \subseteq S$, and $s_I \in S^*$ is the initial state.*
- *\sim_i^n is defined as the corresponding \sim_i restricted to S^* .*
- *The protocol function is defined as $d_n : Ag \times S_n \rightarrow (2^{Act} \setminus \emptyset)$, where $d_n(i, s) = d(i, s)$, for every $s \in S^*$ and $d_n(i, s_\perp) = Act_i$, for all $i \in Ag$.*
- *The transition function is defined as $\delta_n : S_n \times ACT \rightarrow S_n$, where given a transition $\delta(s, \vec{a}) = s'$, if $s, s' \in S^*$ then $\delta_n(s, \vec{a}) = \delta(s, \vec{a}) = s'$ else if $s' \in S \setminus S^*$ and $s \in S_n$ then $\delta_n(s, \vec{a}) = s_\perp$.*
- *for all $s \in S^*$, $V_n(s) = V(s)$ and $V_n(s_\perp) = \emptyset$.*

Now, we present the definition of positive sub-models.

Definition 8 (Positive sub-model). *Given an iCGS $M = \langle Ag, AP, S, s_I, \{Act_i\}_{i \in Ag}, \{\sim_i\}_{i \in Ag}, d, \delta, V \rangle$, we denote with $M_p = \langle Ag, AP, S_p, s_I, \{Act_i\}_{i \in Ag}, \{\sim_i^p\}_{i \in Ag}, d_p, \delta_p, V_p \rangle$ a positive sub-model of M , formally $M_p \subseteq M$, such that:*

- *the set of states is defined as $S_p = S^* \cup \{s_\top\}$, where $S^* \subseteq S$, and $s_I \in S^*$ is the initial state.*
- *\sim_i^p is defined as the corresponding \sim_i restricted to S^* .*
- *The protocol function is defined as $d_p : Ag \times S_p \rightarrow (2^{Act} \setminus \emptyset)$, where $d_p(i, s) = d(i, s)$, for every $s \in S^*$ and $d_p(i, s_\top) = Act_i$, for all $i \in Ag$.*
- *The transition function is defined as $\delta_p : S_p \times ACT \rightarrow S_p$, where given a transition $\delta(s, \vec{a}) = s'$, if $s, s' \in S^*$ then $\delta_p(s, \vec{a}) = \delta(s, \vec{a}) = s'$ else if $s' \in S \setminus S^*$ and $s \in S_p$ then $\delta_p(s, \vec{a}) = s_\top$.*
- *for all $s \in S^*$, $V_p(s) = V(s)$ and $V_p(s_\top) = AP$.*

Note that, the above sub-models are still iCGSs.

We conclude this part by recalling two preservation results presented in [23].

We start with a preservation result from negative sub-models to the original model.

Lemma 1. *Given a model M , a negative sub-model with perfect information M_n of M , and a formula φ of the form $\varphi = \langle\langle A \rangle\rangle\psi$ (resp., $\llbracket A \rrbracket\psi$) for some $A \subseteq Ag$. For any $s \in S_n \setminus \{s_\perp\}$, we have that:*

$$M_n, s \models \varphi \Rightarrow M, s \models \varphi$$

We also consider the preservation result from positive sub-models to the original model.

Lemma 2. *Given a model M , a positive sub-model with perfect information M_p of M , and a formula φ of the form $\varphi = \langle\langle A \rangle\rangle\psi$ (resp., $\llbracket A \rrbracket\psi$) for some $A \subseteq Ag$. For any $s \in S_p \setminus \{s_\top\}$, we have that:*

$$M_p, s \not\models \varphi \Rightarrow M, s \not\models \varphi$$

4 Our procedure

In this section, we provide a procedure to handle games with imperfect information and perfect recall strategies, a problem in general undecidable. The overall model checking procedure is described in Algorithm 1. It takes in input a model M , a formula φ , and a trace h (denoting an execution of the system) and calls the function *Preprocessing*() to generate the negative normal form of φ and to replace all negated atoms with new positive atoms inside M and φ . After that, it calls the function *FindSub-models*() to generate all the positive and negative sub-models that represent all the possible sub-models with perfect information of M . Then, there is a while loop (lines 4-7) that for each candidate checks the sub-formulas true on the sub-models via *CheckSub-formulas*() and returns a result via *RuntimeVerification*(). For the algorithms and additional details regarding the procedures *Preprocessing*(), *FindSub-models*(), and *CheckSub-formulas*() see [23].

Algorithm 1 *ModelCheckingProcedure* (M, φ, h)

```

1: Preprocessing( $M, \varphi$ );
2:  $candidates = FindSub-models(M, \varphi)$ ;
3:  $finalresult = \emptyset$ ;
4: while  $candidates$  is not empty do
5:   extract  $\langle M_n, M_p \rangle$  from  $candidates$ ;
6:    $result = CheckSub-formulas(\langle M_n, M_p \rangle, \varphi)$ ;
7:    $finalresult = RuntimeVerification(M, \varphi, h, result) \cup finalresult$ ;
8: return  $finalresult$ ;
```

Now, we will focus on the last step, the procedure *RuntimeVerification*(). It is performed at runtime, directly on the actual system. In previous steps, the sub-models satisfying (resp., violating) sub-properties φ' of φ are generated, and listed into the set *result*. In Algorithm 2, we report the algorithm performing runtime verification on the actual system. Such algorithm gets in input the model M , an ATL property φ to verify, an execution trace h of events observed by executing the actual system, and the set *result* containing the sub-properties of φ that have been checked on sub-models of M . First, in lines 1-4, the algorithm updates the model M with the atoms corresponding to the sub-properties verified previously on sub-models of M . This step is necessary to keep track explicitly inside M of where the sub-properties are verified (resp., violated). This last aspect depends on which sub-model had been used to verify the sub-property (whether negative or positive). After that, the formula φ needs to be updated accordingly to the newly introduced atoms. This is obtained through updating the formula, by generating at the same time two new versions ψ_n and ψ_p for the corresponding negative and positive versions (lines 6-14). Once ψ_n and ψ_p have been generated, they need to be converted into their corresponding LTL representation to be verified at runtime. Note that, ψ_n and ψ_p are still ATL properties, which may contain strategic operators. Thus, this translation is obtained by removing the strategic operators, leaving only the temporal ones (and the atoms). The resulting two new LTL properties φ_n and φ_p are so obtained (lines 15-16). Finally, by having these two LTL properties, the algorithm proceeds generating (using the standard LTL monitor generation algorithm [24]) the corresponding monitors Mon_{φ_n} and Mon_{φ_p} . Such monitors are then used by Algorithm 2 to check φ_n and φ_p over an execution trace h given in input. The latter consists in a trace observed by executing the system modelled by M (so, the actual system). Analysing h the monitor can conclude the satisfaction (resp., violation) of the LTL property under analysis. However, only certain results can actually be considered

valid. Specifically, when $Mon_{\varphi_n}(h) = \top$, or when $Mon_{\varphi_p}(h) = \perp$. The other cases are considered undefined, since nothing can be concluded at runtime. The reason why line 17 and line 20's conditions are enough to conclude \top and \perp (resp.) directly follow from the following lemmas.

Algorithm 2 *RuntimeVerification* ($M, \varphi, h, result$)

```

1:  $k = ?$ ;
2: for  $s \in S$  do
3:   take set  $atoms$  from  $result(s)$ ;
4:    $UpdateModel(M, s, atoms)$ ;
5:  $\varphi_{mc} = \emptyset$ ;
6: for  $\langle s, \psi, atom \rangle \in result$  do
7:    $\varphi_{mc} = \varphi_{mc} \cup \psi$ ;
8:  $\varphi_{rv} = SubFormulas(\varphi) \setminus \varphi_{mc}$ ;
9:  $\psi_n = \varphi, \psi_p = \varphi$ ;
10: while  $result$  is not empty do
11:   extract  $\langle s, \psi, vatom_\psi \rangle$  from  $result$ ;
12:   if  $v = n$  then
13:      $\psi_n = UpdateFormula(\psi_n, \psi, vatom_\psi)$ ;
14:   else
15:      $\psi_p = UpdateFormula(\psi_p, \psi, vatom_\psi)$ ;
16:  $\varphi_n = FromATLtoLTL(\psi_n, n)$ ;
17:  $\varphi_p = FromATLtoLTL(\psi_p, p)$ ;
18:  $Mon_{\varphi_p} = GenerateMonitor(\varphi_p)$ ;
19:  $Mon_{\varphi_n} = GenerateMonitor(\varphi_n)$ ;
20: if  $Mon_{\varphi_n}(h) = \top$  then
21:    $k = \top$ ;
22: if  $Mon_{\varphi_p}(h) = \perp$  then
23:    $k = \perp$ ;
24:  $\varphi_{unchk} = \emptyset$ ;
25: for  $\varphi' \in \varphi_{rv}$  do
26:    $Mon_{\varphi'} = GenerateMonitor(\varphi')$ ;
27:   if  $Mon_{\varphi'}(h) = ?$  then
28:      $\varphi_{rv} = \varphi_{rv} \setminus \varphi'$ ;
29:      $\varphi_{unchk} = \varphi_{unchk} \cup \varphi'$ ;
30: return  $\langle k, \varphi_{mc}, \varphi_{rv}, \varphi_{unchk} \rangle$ ;

```

We start with a preservation result from the truth of the monitor output to ATL* model checking.

Lemma 3. *Given a model M and a formula φ , for any history h of M starting in s_I , we have that:*

$$Mon_{\varphi_{LTL}}(h) = \top \implies M, s_I \models \varphi_{Ag}$$

where φ_{LTL} is the variant of φ where all strategic operators are removed and φ_{Ag} is the variant of φ where all strategic operators are converted into $\langle\langle Ag \rangle\rangle$.

Proof. First, consider the formula $\varphi = \langle\langle \Gamma \rangle\rangle \psi$, in which $\Gamma \subseteq Ag$ and ψ is a temporal formula without quantifications. So, $\varphi_{LTL} = \psi$ and $\varphi_{Ag} = \langle\langle Ag \rangle\rangle \psi$. By Def.6 we know that $Mon_{\varphi_{LTL}}(h) = \top$ if and only if for all path p in S^ω we have that $h \cdot p$ is in $\llbracket \varphi_{LTL} \rrbracket$. Note that, the latter is the set of paths that satisfy ψ , i.e., $\llbracket \varphi_{LTL} \rrbracket = \{p \mid M, p \models \psi\}$. By Def.2 we know that $M, s_I \models \varphi_{Ag}$ if and only if there exist a strategy profile σ_{Ag} such that for all paths p in $out(s_I, \sigma_{Ag})$ we have that $M, p \models \psi$. Notice that, since the strategic operator involves the whole set of agents, $out(s_I, \sigma_{Ag})$ is composed by a single path. Thus, to guarantee that φ_{Ag} holds in M , our objective is

to construct from s_I the history h as prefix of the unique path in $out(s_I, \sigma_{Ag})$. Since we have $\langle\langle Ag \rangle\rangle$ as strategic operator, this means that there is a way for the set of agents to construct h starting from s_I and the set $out(s_I, \sigma_{Ag})$ becomes equal to $\{p\}$, where $p = h \cdot p'$, for any $p' \in S^\omega$. From the above reasoning, the result follows.

To conclude the proof, note that if we have a formula with more strategic operators then we can use a classic bottom-up approach. \square

Now, we present a preservation result from the falsity of the monitor output to ATL* model checking.

Lemma 4. *Given a model M and a formula φ , for any history h of M starting in s_I , we have that:*

$$Mon_{\varphi_{LTL}}(h) = \perp \implies M, s_I \not\models \varphi_\emptyset$$

where φ_{LTL} is the variant of φ where all strategic operators are removed and φ_\emptyset is the variant of φ where all strategic operators are converted into $\langle\langle \emptyset \rangle\rangle$.

Proof. First, consider the formula $\varphi = \langle\langle \Gamma \rangle\rangle \psi$, in which $\Gamma \subseteq Ag$ and ψ is a temporal formula without quantifications. So, $\varphi_{LTL} = \psi$ and $\varphi_\emptyset = \langle\langle \emptyset \rangle\rangle \psi$. By Def.6 we know that $Mon_{\varphi_{LTL}}(h) = \perp$ if and only if for all path p in S^ω we have that $h \cdot p$ is not in $\llbracket \varphi_{LTL} \rrbracket$. Note that, the latter is the set of paths that satisfy ψ , i.e., $\llbracket \varphi_{LTL} \rrbracket = \{p \mid M, p \models \psi\}$. By Def.2 we know that $M, s_I \not\models \varphi_\emptyset$ if and only if for all strategy profiles σ_\emptyset , there exists a path p in $out(s_I, \sigma_\emptyset)$ such that $M, p \not\models \psi$. Notice that, since the strategic operator is empty then $out(s_I, \sigma_\emptyset)$ is composed by all the paths in M . Thus, to guarantee that φ_\emptyset does not hold in M , our objective is to select a path p in $out(s_I, \sigma_\emptyset)$ starting from s_I , where $p = h \cdot p'$, for any $p' \in S^\omega$. Given the assumption that $h \cdot p$ is not in $\llbracket \varphi_{LTL} \rrbracket$ then the result follows.

To conclude the proof, note that if we have a formula with more strategic operators then we can use a classic bottom-up approach. \square

It is important to evaluate in depth the meaning of the two lemmas presented above, we do this in the following remark.

Remark 1. *Lemma 3 and 4 show a preservation result from runtime verification to ATL* model checking that needs to be discussed. If our monitor returns true we have two possibilities:*

1. *the procedure found a negative sub-model in which the original formula φ is satisfied then it can conclude the verification procedure by using RV only by checking that the atom representing φ holds in the initial state of the history h given in input;*
2. *a sub-formula φ' is satisfied in a negative sub-model and at runtime the formula φ_{Ag} holds on the history h given in input.*

While case 1. gives a preservation result for the formula φ given in input, case 2. checks formula φ_{Ag} instead of φ . That is, it substitutes Ag as coalition for all the strategic operators of φ but the ones in φ' . So, our procedure approximates the truth value by considering the case in which all the agents in the game collaborate to achieve the objectives not satisfied in the model checking phase. That is, while in [8, 9] the approximation is given in terms of information, in [10] is given in terms of recall of the strategies, and in [23] the approximation is given by generalizing the logic, here we

give results by approximating the coalitions. Furthermore, we recall that our procedure produces always results, even partial. This aspect is strongly relevant in concrete scenario in which there is the necessity to have some sort of verification results. For example, in the context of swarm robots [25], with our procedure we can verify macro properties such as "the system works properly" since we are able to guarantee fully collaboration between agents because this property is relevant and desirable for each agent in the game. The same reasoning described above, can be applied in a complementary way for the case of positive sub-models and the falsity.

To conclude this section we show and prove the complexity of our procedure.

Theorem 1. *Algorithm 1 terminates in $2EXPTIME$. Moreover, Algorithm 1 is sound: if the value returned is different from \perp , then $M \models \varphi_{Ag}$ iff $k = \top$.*

Proof. The preprocessing phase is polynomial in the size of the model and the formula. As described in [23], $FindSub-models()$ terminates in $EXPTIME$. The while loop in lines 3-7 needs to check all the candidates and in the worst case the size of the list of candidates is equal to the size of the set of states of M (i.e., polynomial in the size of M). About $CheckSub-formulas()$, as described in [23], the complexity is $2EXPTIME$ due to the ATL^* model checking that is called in it. Finally, Algorithm 2 terminates in $2EXPTIME$. In particular, loops in lines 2, 6, and 10 terminate in polynomial time with respect to the size of the model and the size of the formula. As described in [24], to generate a monitor requires $2EXPTIME$ in the size of the formula and the execution of a monitor is linear in the size of the formula. So, the total complexity is still determined by the subroutines and directly follows.

About the soundness, suppose that the value returned is different from \perp . In particular, either $k = \top$ or $k = \perp$. If $M \models \varphi_{Ag}$ and $k = \perp$, then by Algorithm 1 and 2, we have that $Mon_{\varphi_p}(h) = \perp$. Now, there are two cases: (1) h is an history of M (2) there exists an history h' of M that differs from h for some atomic propositions added to h in lines 2-4 of Algorithm 2. For (1), we know that h is in M and thus $Mon_{\varphi_p}(h) = \perp$ implies $M \not\models \varphi_{\emptyset}$ by Lemma 4 that implies $M \not\models \varphi_{Ag}$ by the semantics in Def. 4, a contradiction. Hence, $k = \top$ as required. For (2), suppose that h has only one additional atomic proposition $atom_{\psi}$. The latter means that $CheckSub-formulas()$ found a positive sub-model M_p in which $M_p, s \models \psi$, for some $s \in S_p$. By Lemma 2, for all $s \in S_p$, we know that if $M_p, s \not\models \psi$ then $M, s \not\models \psi$. So, h over-approximates h' , i.e. there could be some states that in h are labeled with $atom_{\psi}$ but they don't satisfy ψ in h . Thus, if $Mon_{\varphi_p}(h) = \perp$ then $M \not\models \varphi_{\emptyset}$ by Lemma 4 that implies $M \not\models \varphi_{Ag}$, a contradiction. Hence, $k = \top$ as required. Obviously, we can generalize the above reasoning in case h and h' differ for multiple atomic propositions. On the other hand, if $k = \top$ then by Algorithm 1 and 2, we have that $Mon_{\varphi_n}(h) = \top$. Again, there are two cases: (1) h is an history of M (2) there exists an history h' of M that differs from h for some atomic propositions added to h in lines 2-4 of Algorithm 2. For (1), we know that h is in M and thus $Mon_{\varphi_n}(h) = \top$ implies $M \models \varphi_{Ag}$ by Lemma 3 as required. For (2), suppose that h has only one additional atomic proposition $atom_{\psi}$. The latter means that $CheckSub-formulas()$ found a negative sub-model M_n in which $M_n, s \models \psi$, for some $s \in S_n$. By Lemma 1, for all $s \in S_n$, we know that if $M_n, s \models \psi$ then $M, s \models \psi$. So, h under-approximates h' , i.e. there could be some states that in h are not labeled with $atom_{\psi}$ but they satisfy ψ in M . Thus, if $Mon_{\varphi_n}(h) = \top$ then $M \models \varphi_{Ag}$ by Lemma 3, as required. \square

5 Our tool

The algorithms presented previously have been implemented in Java¹. The resulting tool implementing Algorithm 1 allows to extract all sub-models with perfect information (*CheckSub-formulas()*) that satisfy a strategic objective from a model given in input. The extracted sub-models, along with the corresponding sub-formulas, are then used by the tool to generate and execute the corresponding monitors over a system execution (Algorithm 2).

In more detail, as shown in Figure 1, the tool expects a model in input formatted as a Json file. This file is then parsed, and an internal representation of the model is generated. After that, the verification of a sub-model against a sub-formula is achieved by translating the sub-model into its equivalent ISPL (Interpreted Systems Programming Language) program, which then is verified by using the model checker MCMAS²[26]. This corresponds to the verification steps performed in *CheckSub-formulas()* (i.e., where static verification through MCMAS is used). For each sub-model that satisfies this verification step, the tool produces a corresponding tuple; which contains the information needed by Algorithm 2 to complete the verification at runtime.

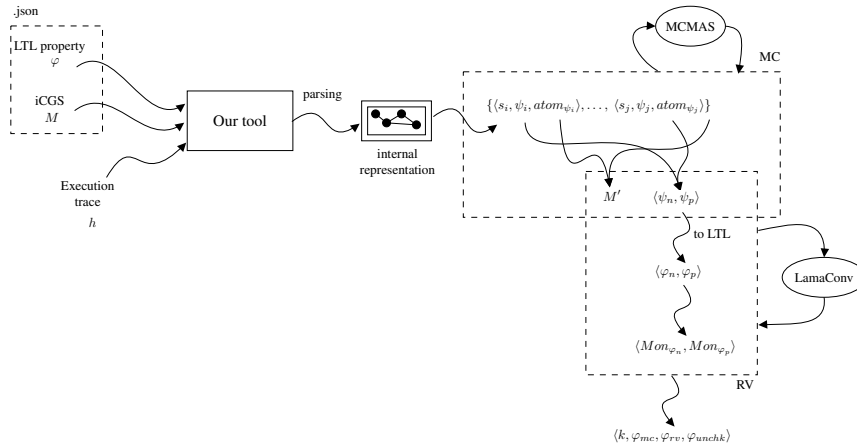


Figure 1: Overview of the implemented tool.

The entire manipulation, from parsing the model formatted in Json, to translating the latter to its equivalent ISPL program, has been performed by extending an existent Java library [27]; the rest of the tool derives directly from the algorithms presented in this paper. The monitors generated by Algorithm 2 at lines 18 and 19 are obtained using LamaConv [28], which is a Java library capable of translating expressions in temporal logic into equivalent automata and generating monitors out of these automata. For generating monitors, LamaConv uses the algorithm presented in [24].

5.1 Experiments

We tested our tool on a large set of automatically and randomly generated iCGSs; on a machine with the following specifications: Intel(R) Core(TM) i7-7700HQ CPU @ 2.80GHz, 4 cores 8 threads, 16 GB RAM DDR4. The objective of these experiments

¹The tool can be found at <https://github.com/AngeloFerrando/StrategyRV>

²<https://vas.doc.ic.ac.uk/software/mcmass/>




Figure 2: Success rate of our tool when applied to a set of randomly generated iCGSs.

was to show how many times our algorithm returned a conclusive verdict. For each model, we ran our procedure and counted the number of times a solution was returned. Note that, our approach concludes in any case, but since the general problem is undecidable, the result might be inconclusive (*i.e.*, ?). In Figure 2, we report our results by varying the percentage of imperfect information (*x* axis) inside the iCGSs, from 0% (perfect information, *i.e.*, all states are distinguishable for all agents), to 100% (no information, *i.e.*, no state is distinguishable for any agent). For each percentage selected, we generated 10000 random iCGSs and counted the number of times our algorithm returned with a conclusive result (*i.e.*, \top or \perp). As it can be seen in Figure 2, our tool concludes with a conclusive result more than 80% of times. We do not observe any relevant difference amongst the different percentage of information used in the experiments. This is mainly due to the completely random nature of the iCGSs used. In more detail, the results we obtained completely depend on the topology of the iCGSs, so it is very hard to precisely quantify the success rate. However, the results obtained by our experiments using our procedure are encouraging. Unfortunately, no benchmark of existing iCGSs – to test our tool on – exists, thus these results may vary on more realistic scenarios. Nonetheless, considering the large set of iCGSs we experimented on, we do not expect substantial differences.

Other than testing our tool w.r.t. the success rate over a random set of iCGSs, we evaluated the execution time as well. Specifically, we were much interested in analysing how such execution time is divided between *CheckSub-formulas()* and Algorithm 2. *I.e.*, how much time is spent on verifying the models statically (through model checking), and how much is spent on verifying the temporal properties (through runtime verification). Figure 3 reports the results we obtained on the same set of randomly generated used in Figure 2. The results we obtained are intriguing, indeed we can note a variation in the percentage of time spent on the two phases (*y*-axis) moving from low percentages to high percentages of imperfect information in the iCGSs (*x*-axis). When the iCGS is close to have perfect information (low percentages on *x*-axis),




Figure 3: How the execution time of our tool when applied to a set of randomly generated iCGSs is divided.

we may observe that most of the execution time is spent on performing static verification ($\sim 70\%$), which corresponds to *CheckSub-formulas()*. On the other hand, when imperfect information grows inside the iCGS (high percentage on x-axis), we may observe that most of the execution time is spent on performing runtime verification ($\sim 90\%$ in occurrence of absence of information). The reason for this change in the execution behaviour is determined by the number of candidates extracted by the *FindSub-models()* function. When the iCGS has perfect information, such function only extracts a single candidate (*i.e.*, the entire model), since *FindSub-models()* generates only one tuple. Such single candidate can be of non-negligible size, and the resulting static verification, time consuming; while the subsequent runtime verification is only performed once on the remaining temporal parts of the property to verify. On the other hand, when the iCGS has imperfect information, *FindSub-models()* returns a set of candidates that can grow exponentially w.r.t. the number of states of the iCGS. Nonetheless, such candidates are small in size, since *FindSub-models()* splits the iCGS into multiple smaller iCGSs with perfect information. Because of this, the static verification step is applied on small iCGSs and require less execution time; while the runtime verification step is called for each candidate (so an exponential number of times) and is only influenced by the size of the temporal property to verify.

In conclusion, it is important to emphasise that, even though the monitor synthesis is computationally hard (*i.e.*, $2EXPTIME$), the resulting runtime verification process is polynomial in the size of the history analysed. Naturally, the actual running complexity of a monitor depends on the formalism used to describe the formal property. In this work, monitors are synthesised from LTL properties. Since LTL properties are translated into Moore machines [24]; because of this, the time complexity w.r.t. the length of the analysed trace is linear. This can be understood intuitively by noticing that the Moore machine so generated has finite size, and it does not change at runtime. Thus, the number of execution steps for each event in the trace is constant.

6 Conclusions and Future work

The work presented in this paper follows a standard combined approach of formal verification techniques, where the objective is to get the best of both. We considered the model checking problem of MAS using strategic properties that is undecidable in general, and showed how runtime verification can help by verifying part of the properties at execution time. The resulting procedure has been presented both on a theoretical (theorems and algorithms) and a practical level (prototype implementation). It is important to note that this is the first attempt of combining model checking and runtime verification to verify strategic properties on a MAS. Thus, even though our solution might not be optimal, it is a milestone for the corresponding lines of research. Additional works will be done to improve the technique and, above all, its implementation. For instance, we are planning to extend this work considering a more predictive flavour. This can be done by recognising the fact that by verifying at static time part of the system, we can use this information at runtime to predict future events and conclude the runtime verification in advance.

References

- [1] R. Alur, T.A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *J. ACM*, 49(5):672–713, 2002.
- [2] C. Dima and F.L. Tiplea. Model-checking ATL under imperfect information and perfect recall semantics is undecidable. *CoRR*, abs/1102.4225, 2011.
- [3] P.Y. Schobbens. Alternating-Time Logic with Imperfect Recall. *ENTCS*, 85(2):82–93, 2004.
- [4] Raphaël Berthon, Bastien Maubert, and Aniello Murano. Decidability results for atl^* with imperfect information and perfect recall. In Kate Larson, Michael Winikoff, Sanmay Das, and Edmund H. Durfee, editors, *Proceedings of the 16th Conference on Autonomous Agents and MultiAgent Systems, AAMAS 2017, São Paulo, Brazil, May 8-12, 2017*, pages 1250–1258. ACM, 2017.
- [5] Raphaël Berthon, Bastien Maubert, Aniello Murano, Sasha Rubin, and Moshe Y. Vardi. Strategy logic with imperfect information. *ACM Trans. Comput. Log.*, 22(1):5:1–5:51, 2021.
- [6] F. Belardinelli, A. Lomuscio, A. Murano, and S. Rubin. Verification of multi-agent systems with imperfect information and public actions. In *AAMAS 2017*, pages 1268–1276, 2017.
- [7] 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.
- [8] F. Belardinelli, A. Lomuscio, and V. Malvone. An abstraction-based method for verifying strategic properties in multi-agent systems with imperfect information. In *Proceedings of AAI*, 2019.
- [9] Francesco Belardinelli and Vadim Malvone. A three-valued approach to strategic abilities under imperfect information. In *Proceedings of the 17th International Conference on Knowledge Representation and Reasoning*, pages 89–98, 2020.

- [10] F. Belardinelli, A. Lomuscio, and V. Malvone. Approximating perfect recall when model checking strategic abilities. In *KR2018*, pages 435–444, 2018.
- [11] Wojciech Jamroga, Vadim Malvone, and Aniello Murano. Natural strategic ability. *Artif. Intell.*, 277, 2019.
- [12] Wojciech Jamroga, Vadim Malvone, and Aniello Murano. Natural strategic ability under imperfect information. In Edith Elkind, Manuela Veloso, Noa Agmon, and Matthew E. Taylor, editors, *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19, Montreal, QC, Canada, May 13-17, 2019*, pages 962–970. International Foundation for Autonomous Agents and Multiagent Systems, 2019.
- [13] Timothy L. Hinrichs, A. Prasad Sistla, and Lenore D. Zuck. Model check what you can, runtime verify the rest. In Andrei Voronkov and Margarita V. Korovina, editors, *HOWARD-60: A Festschrift on the Occasion of Howard Barringer's 60th Birthday*, volume 42 of *EPiC Series in Computing*, pages 234–244. EasyChair, 2014.
- [14] Ankush Desai, Tommaso Dreossi, and Sanjit A. Seshia. Combining model checking and runtime verification for safe robotics. In Shuvendu K. Lahiri and Giles Reger, editors, *Runtime Verification - 17th International Conference, RV 2017, Seattle, WA, USA, September 13-16, 2017, Proceedings*, volume 10548 of *Lecture Notes in Computer Science*, pages 172–189. Springer, 2017.
- [15] Katarína Kejstová, Petr Rockai, and Jiri Barnat. From model checking to runtime verification and back. In Shuvendu K. Lahiri and Giles Reger, editors, *Runtime Verification - 17th International Conference, RV 2017, Seattle, WA, USA, September 13-16, 2017, Proceedings*, volume 10548 of *Lecture Notes in Computer Science*, pages 225–240. Springer, 2017.
- [16] Jiří Barnat, Luboš Brim, Vojtěch Havel, Jan Havlíček, Jan Kriho, Milan Lenčo, Petr Ročkai, Vladimír Štill, and Jiří Weiser. DiVinE 3.0—an explicit-state model checker for multithreaded C & C++ programs. In *International Conference on Computer Aided Verification*, pages 863–868. Springer, 2013.
- [17] Daniel M. Zimmerman and Joseph R. Kiniry. A verification-centric software development process for java. In Byoungju Choi, editor, *Proceedings of the Ninth International Conference on Quality Software, QSIC 2009, Jeju, Korea, August 24-25, 2009*, pages 76–85. IEEE Computer Society, 2009.
- [18] Angelo Ferrando and Vadim Malvone. Strategy rv: A tool to approximate atl model checking under imperfect information and perfect recall. In *Proceedings of the 20th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '21*, page 1764–1766, Richland, SC, 2021. International Foundation for Autonomous Agents and Multiagent Systems.
- [19] W. Jamroga and W. van der Hoek. Agents that know how to play. *Fund. Inf.*, 62:1–35, 2004.
- [20] R. Fagin, J.Y. Halpern, Y. Moses, and M.Y. Vardi. *Reasoning about Knowledge*. MIT, 1995.

- [21] C. Baier and J. P. Katoen. *Principles of Model Checking (Representation and Mind Series)*. 2008.
- [22] Amir Pnueli. The temporal logic of programs. In *18th Annual Symposium on Foundations of Computer Science, Providence, Rhode Island, USA, 31 October - 1 November 1977*, pages 46–57. IEEE Computer Society, 1977.
- [23] Angelo Ferrando and Vadim Malvone. Towards the verification of strategic properties in multi-agent systems with imperfect information. *CoRR*, abs/2112.13621, 2021.
- [24] Andreas Bauer, Martin Leucker, and Christian Schallhart. Runtime verification for LTL and TLTL. *ACM Trans. Softw. Eng. Methodol.*, 20(4):14:1–14:64, 2011.
- [25] Panagiotis Kouvaros and Alessio Lomuscio. Parameterised verification for multi-agent systems. *Artif. Intell.*, 234:152–189, 2016.
- [26] A. Lomuscio and F. Raimondi. Model checking knowledge, strategies, and games in multi-agent systems. In *Proceedings of the 5th International Joint Conference on Autonomous agents and Multi-Agent Systems (AAMAS06)*, pages 161–168. ACM Press, 2006.
- [27] Francesco Belardinelli, Vadim Malvone, and Abbas Slimani. A tool for verifying strategic properties in mas with imperfect information, 2020.
- [28] Torben Scheffel and Malte Schmitz et al. LamaConv- logics and automata converter library, 2016.