Decidability Results for ATL* with Imperfect Information and Perfect Recall

Raphaël Berthon
ENS Rennes
Rennes, France
raphael.berthon@ens-rennes.fr

Bastien Maubert
University of Naples Federico II
Naples, Italy
bastien.maubert@gmail.com

Aniello Murano
University of Naples Federico II
Naples, Italy
murano@na.infn.it

ABSTRACT

Alternating-time Temporal Logic (ATL*) is a central logic for multiagent systems. Its extension to the imperfect information setting (ATL$_i^*$) is well known to have an undecidable model-checking problem when agents have perfect recall. Studies have thus mostly focused either on agents without memory, or on alternative semantics to retrieve decidability. In this work we establish new decidability results for agents with perfect recall: We first prove a meta-theorem that allows the transfer of decidability results for classes of multiplayer games with imperfect information, such as games with hierarchical observation, to the model-checking problem for ATL$_i^*$. We then establish that model checking ATL* with strategy context and imperfect information is decidable when restricted to *hierarchical instances*.

1. INTRODUCTION

In formal verification, *model checking* is a well-established method to automatically check systems' correctness [7, 33, 8]. It consists in modelling the system as a mathematical structure, expressing a desired property as a formula from a suitable logic, and checking whether the model satisfies the formula. In the nineties, interest has arisen in the verification of *multiagent systems* (MAS), in which various entities (the *agents*) interact and can form coalitions to attain their objectives. This has led to the development of logics to reason about strategic abilities in MAS [1, 2, 6, 26, 27, 28, 37].

Alternating-time Temporal Logic (ATL*) [2] plays a central role in this line of work. Interpreted on *concurrent game structures (CGS)*, it extends CTL* with *strategic modalities*, which express the existence of strategies for coalitions of agents to force the system's behaviour to satisfy certain temporal properties. ATL* has been extended in many ways, and notably with *strategy contexts* [5, 24]: In ATL*, strategies of all agents are forgotten at each new strategic modality. In ATL* *with strategy context* (ATL$_{sc}^*$), instead, they are stored in a strategy context, and are forgotten only when replaced by a new strategy or when the formula explicitly unbinds the agent from her strategy. This makes ATL$_{sc}^*$ expressive enough to capture important game theoretic concepts, such as the existence of Nash Equilibria [24].

In many real-life scenarios, such as when some information

is private or hidden for security reasons, agents do not know precisely what is the current state of the system, but have a partial view, or observation, of the state. This fundamental feature of MAS is called *imperfect information*, and it is known to quickly bring about undecidability when involved in strategic problems, especially when agents have *perfect recall* of the past, which is a usual and important assumption in games with imperfect information and epistemic temporal logics [12]. For instance solving multiplayer games with imperfect information and perfect recall, *i.e.*, deciding the existence of a distributed winning strategy in such games, is already undecidable for reachability objective [31]. Since such games are easily captured by ATL* with imperfect information (ATL$_i^*$), model checking ATL$_i^*$ with perfect recall is also undecidable [2].

However, restricting attention to cases where some sort of hierarchy exists on the different agents' information yields decidability for several problems related to the existence of strategies: Synthesis of distributed systems, which implicitly uses perfect recall and is undecidable in general [32], is decidable for hierarchical architectures [22]. Actually, for branching-time specifications, distributed synthesis is decidable exactly on architectures free from *information forks*, for which the problem can be reduced to the hierarchical case [13]. For richer specifications from alternating-time logics, being free of information forks is no longer sufficient, but distributed synthesis is decidable precisely on hierarchical architectures [34]. Similarly, solving multiplayer games with imperfect information and perfect recall, *i.e.*, checking for the existence of winning distributed strategies, is decidable for ω-regular winning conditions when there is a hierarchy among players, each one observing more than those below [30, 22]. Recently, it has been proven that this assumption can be relaxed: the problem remains decidable if the hierarchy can change along a play, or even if transient phases without such a hierarchy are allowed [4]. Note that hierarchical information occurs naturally, for instance when agents are assigned different levels of security clearance.

Our contribution. In this work we establish several decidability results for model checking ATL$_i^*$ with perfect recall, with and without strategy context, all related to notions of hierarchy. Our first result is a theorem that allows the transfer of decidability results for classes of multiplayer games with imperfect information, such as those mentioned above, to the model-checking problem for ATL$_i^*$. This theorem essentially states that if solving multiplayer games with imperfect information, perfect recall and omega-regular objectives is decidable on some class of concurrent game struc-

tures, then model checking ATL_i^* with perfect recall is also decidable on this class of models (a simple bottom-up algorithm that evaluates innermost strategic modalities in every state of the model suffices). As a direct consequence we easily obtain new decidability results for the model checking of ATL_i^* on several classes of concurrent game structures.

Our second contribution concerns ATL^* with imperfect information and strategy context ($\mathsf{ATL}_{sc,i}^*$). Because there are in general infinitely many possible strategy contexts, the bottom-up approach used for ATL_i^* fails here. Instead we build upon the proof presented in [24] that establishes the decidability of model checking ATL_{sc}^* by reduction to the model-checking problem for Quantified CTL^* (QCTL^*). The latter extends CTL^* with second-order quantification on atomic propositions, and it has been well studied [36, 20, 21, 14, 23]. QCTL_i^*, an imperfect-information extension of QCTL^*, has recently been introduced, and its model-checking problem was proven decidable for the class of *hierarchical formulas* [3]. In this paper we define a notion of *hierarchical instances* for the $\mathsf{ATL}_{sc,i}^*$ model-checking problem: an $\mathsf{ATL}_{sc,i}^*$ formula φ together with a concurrent game structure $\mathcal{G}$ is a hierarchical instance if the observations of agents appearing in strategy quantifiers get more refined as one goes down φ's syntactic tree. We adapt the proof from [24] and prove the model-checking problem for $\mathsf{ATL}_{sc,i}^*$ on hierarchical instances decidable by reduction to the model-checking problem for hierarchical QCTL_i^* formulas.

Related work. The model-checking problem for ATL_i^* is known to be decidable when agents have no memory [35], and the case of agents with bounded memory reduces to that of no memory. Another way to retrieve decidability is to assume that all agents in a coalition have the same information, either because their observations of the system are the same, or because they can communicate and share their observations [10, 15, 16, 18, 19]. This idea was also used recently to establish a decidability result for $\mathsf{ATL}_{sc,i}^*$ [25] when all agents have the same observation of the game.

The results we establish here thus strictly extend previously known results on the decidability of model checking ATL_i^* and $\mathsf{ATL}_{sc,i}^*$ with perfect recall and standard semantics, and they hold for vast, natural classes of instances, that all rely on notions of hierarchy, which seems to be inherent to all decidable cases of strategic problems for multiple entities with imperfect information and perfect recall.

Outline. After setting some basic definitions in Section 2, we present our meta-theorem on the model checking problem for ATL_i^* in Section 3. In Section 4 we prove that when restricted to hierarchical instances, model checking $\mathsf{ATL}_{sc,i}^*$ is decidable, and we conclude in Section 5.

2. PRELEMINARIES

Let Σ be an alphabet. A *finite* (resp. *infinite*) *word* over Σ is an element of Σ^* (resp. Σ^ω). The empty word is noted ϵ, and $\Sigma^+ = \Sigma^* \setminus \{\epsilon\}$. The *length* of a word is $|w| := 0$ if w is the empty word ϵ, if $w = w_0 w_1 \ldots w_n$ is a finite nonempty word then $|w| := n + 1$, and for an infinite word w we let $|w| := \omega$. Given a word w and $0 \leq i, j \leq |w| - 1$, we let w_i be the letter at position i in w and $w[i, j]$ be the subword of w that starts at position i and ends at position j. For $n \in \mathbb{N}$ we let $[n] := \{1, \ldots, n\}$. Finally, for the rest of the paper, let us fix a countably infinite set of *atomic propositions* $\mathcal{AP}$ and let $AP \subset \mathcal{AP}$ be some finite subset of atomic propositions.

2.1 Kripke structures

A *Kripke structure* over AP is a tuple $\mathcal{S} = (S, R, \ell)$ where S is a set of *states*, $R \subseteq S \times S$ is a left-total[1] *transition relation* and $\ell : S \to 2^{AP}$ is a *labelling function*.

A *pointed Kripke structure* is a pair $(\mathcal{S}, s)$ where $s \in \mathcal{S}$. A *path* in a structure $\mathcal{S} = (S, R, \ell)$ is an infinite word λ over S such that for all $i \in \mathbb{N}$, $(\lambda_i, \lambda_{i+1}) \in R$. For $s \in S$, $\mathrm{Paths}(s)$ is the set of all paths that start in s.

2.2 Infinite trees

Let X be a finite set. An X-*tree* τ is a nonempty set of words $\tau \subseteq X^+$ such that

- there exists $r \in X$, called the *root* of τ, such that each $u \in \tau$ starts with r;
- if $u \cdot x \in \tau$ with $x \in X$ and $u \neq \epsilon$, then $u \in \tau$, and
- if $u \in \tau$ then there exists $x \in X$ such that $u \cdot x \in \tau$.

The elements of a tree τ are called *nodes*. If $u \cdot x \in \tau$, we say that $u \cdot x$ is a *child* of u. Similarly to Kripke structures, a *path* is an infinite sequence of nodes $\lambda = u_0 u_1 \ldots$ such that for all i, u_{i+1} is a child of u_i, and $Paths(u)$ is the set of paths that start in node u. An AP-*labelled* X-*tree*, or (AP, X)-*tree* for short, is a pair $t = (\tau, \ell)$, where τ is an X-tree called the *domain* of t and $\ell : \tau \to 2^{AP}$ is a *labelling*.

DEFINITION 1 (TREE UNFOLDINGS). *Let* $\mathcal{S} = (S, R, \ell)$ *be a Kripke structure over* AP, *and let* $s \in S$. *The tree-unfolding of* $\mathcal{S}$ *from* s *is the* (AP, S)-*tree* $t_{\mathcal{S}}(s) = (\tau, \ell')$, *where* τ *is the set of all finite paths that start in* s, *and for every* $u \in \tau$, $\ell'(u) = \ell(s)$, *where* s *is the last letter of* u.

3. ATL* WITH IMPERFECT INFORMATION

In this section we recall the syntax and semantics of ATL^* with imperfect information and synchronous perfect-recall semantics, or ATL_i^* for short, and establish a meta-theorem on the decidability of its model-checking problem.

3.1 Definitions

We first introduce the models of the logics we study. For the rest of the paper, let us fix a non-empty finite set of *agents* Ag and a non-empty finite set of *moves* M.

DEFINITION 2. *A concurrent game structure with imperfect information (or* CGS_i *for short) over* AP *is a tuple* $\mathcal{G} = (V, E, \ell, \{\sim_a\}_{a \in \mathrm{Ag}})$ *where* V *is a non-empty finite set of* positions, $E : V \times \mathrm{M}^{\mathrm{Ag}} \to V$ *is a* transition function, $\ell : V \to 2^{AP}$ *is a* labelling function *and for each agent* $a \in \mathrm{Ag}$, $\sim_a \subseteq V \times V$ *is an equivalence relation.*

In a position $v \in V$, each agent a chooses a move $m_a \in \mathrm{M}$, and the game proceeds to position $E(v, \boldsymbol{m})$, where $\boldsymbol{m} \in \mathrm{M}^{\mathrm{Ag}}$ stands for the *joint move* $(m_a)_{a \in \mathrm{Ag}}$ (note that we assume $E(v, \boldsymbol{m})$ to be defined for all v and $\boldsymbol{m}$[2]). For each position $v \in V$, $\ell(v)$ is the finite set of atomic propositions that hold in v, and for $a \in \mathrm{Ag}$, equivalence relation $\sim_a$ represents the observation of agent a: for two positions $v, v' \in V$, $v \sim_a v'$ means that agent a cannot tell the difference between v and v'. We may write $v \in \mathcal{G}$ for $v \in V$. A *pointed* CGS_i $(\mathcal{G}, v)$ is a CGS_i $\mathcal{G}$ together with a position $v \in \mathcal{G}$.

---

[1]*i.e.*, for all $s \in S$, there exists s' such that $(s, s') \in R$.
[2]This assumption, as well as the choice of a unique set of moves for all agents, is made to ease presentation. All the results presented here also hold when the set of available moves depends on the agent and the position.

In Section 3.2 we also use *nondeterministic CGS_i*, which are as in Definition 2 except that they have a *transition relation* $E \subseteq V \times M^{\mathrm{Ag}} \times V$ instead of a transition function. In a position v, after every agent has chosen a move, forming a joint move $\boldsymbol{m} \in M^{\mathrm{Ag}}$, a special agent called Nature (not in Ag) chooses a next position v' such that $(v, \boldsymbol{m}, v') \in E$ (see [4] for detail). In the following, unless explicitly specified, CGS_i always refers to deterministic CGS_i. The following definitions also concern deterministic CGS_i, but they can be adapted to nondeterministic ones in an obvious way.

A *finite* (resp. *infinite*) *play* is a finite (resp. infinite) word $\rho = v_0 \ldots v_n$ (resp. $\pi = v_0 v_1 \ldots$) such that for all i with $0 \leq i < |\rho| - 1$ (resp. $i \geq 0$), there exists a joint move $\boldsymbol{m}$ such that $E(v_i, \boldsymbol{m}) = v_{i+1}$. A finite (resp. infinite) play ρ (resp. π) *starts* in a position v if $\rho_0 = v$ (resp. $\pi_0 = v$). We let $\mathrm{Plays}(\mathcal{G}, v)$ be the set of plays, either finite or infinite, that start in v.

In this work we consider agents with synchronous perfect recall, meaning that the observational equivalence relation for each agent a is extended to finite plays the following way: $\rho \sim_a \rho'$ if $|\rho| = |\rho|'$ and $\rho_i \sim_a \rho'_i$ for every $i \in \{0, \ldots, |\rho| - 1\}$. A *strategy for agent a* is a function $\sigma : V^+ \to M$ such that $\sigma(\rho) = \sigma(\rho')$ whenever $\rho \sim_a \rho'$. The latter constraint captures the essence of imperfect information, which is that agents can base their strategic choices only on the information available to them, and removing this constraint yields the semantics of classic ATL with perfect information.

A *strategy profile* for a coalition $A \subseteq \mathrm{Ag}$ is a mapping σ_A that assigns a strategy to each agent $a \in A$; for $a \in A$, we may write σ_a instead of $\sigma_A(a)$. An infinite play π *follows* a strategy profile σ_A for a coalition A if for all $i \geq 0$, there exists a joint move $\boldsymbol{m}$ such that $E(\pi_i, \boldsymbol{m}) = \pi_{i+1}$ and for each $a \in A$, $m_a = \sigma_a(\pi[0, i])$. For a strategy profile σ_A and a position $v \in V$, we define the outcome $\mathrm{Out}(v, \sigma_A)$ of σ_A in v as the set of infinite plays that start in v and follow σ_A.

The syntax of ATL_i^* is the same as that of ATL^*, and is given by the following grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \langle A \rangle \varphi \mid \mathbf{X}\varphi \mid \varphi \mathbf{U}\varphi,$$

where $p \in \mathcal{AP}$ and $A \subseteq \mathrm{Ag}$.

$\mathbf{X}$ and $\mathbf{U}$ are the classic *next* and *until* operators, respectively, while the *strategic* operator $\langle A \rangle$ quantifies over strategy profiles for coalition A.

The semantics of ATL_i^* is defined with regards to a CGS_i $\mathcal{G} = (V, E, \ell, \{\sim_a\}_{a \in \mathrm{Ag}})$, an infinite play π and a position $i \geq 0$ along this play, by induction on formulas:

$$
\begin{aligned}
&\mathcal{G}, \pi, i \models p && \text{if } p \in \ell(\pi_i) \\
&\mathcal{G}, \pi, i \models \neg\varphi && \text{if } \mathcal{G}, \pi, i \not\models \varphi \\
&\mathcal{G}, \pi, i \models \varphi \vee \varphi' && \text{if } \mathcal{G}, \pi, i \models \varphi \text{ or } \mathcal{G}, \pi, i \models \varphi' \\
&\mathcal{G}, \pi, i \models \langle A \rangle \varphi && \text{if there exists a strategy profile } \sigma_A \text{ s.t.} \\
& && \quad \text{for all } \pi' \in \mathrm{Out}(\pi_i, \sigma_A),\ \mathcal{G}, \pi', 0 \models \varphi \\
&\mathcal{G}, \pi, i \models \mathbf{X}\varphi && \text{if } \mathcal{G}, \pi, i+1 \models \varphi \\
&\mathcal{G}, \pi, i \models \varphi \mathbf{U}\varphi' && \text{if there exists } j \geq i \text{ s.t. } \mathcal{G}, \pi, j \models \varphi' \text{ and,} \\
& && \quad \text{for all } k \text{ s.t. } i \leq k < j,\ \mathcal{G}, \pi, k \models \varphi.
\end{aligned}
$$

An ATL_i^* formula φ is *closed* if every temporal operator ($\mathbf{X}$ or $\mathbf{U}$) in φ is in the scope of a strategic operator $\langle A \rangle$. Since the semantics of a closed formula φ does not depend on the future, we may write $\mathcal{G}, v \models \varphi$, meaning that $\mathcal{G}, \pi, 0 \models \varphi$ for any infinite play π that starts in v.

The *model-checking problem for ATL_i^** consists in deciding, given a closed ATL_i^* formula φ and a finite pointed CGS_i $(\mathcal{G}, v)$, whether $\mathcal{G}, v \models \varphi$.

3.2   Model checking ATL_i^*

It is well known that the model-checking problem for ATL_i^* is undecidable for agents with perfect recall [2], as it can easily express the existence of distributed winning strategies for multiplayer reachability games with imperfect information and perfect recall, which was proved undecidable by Peterson, Reif and Azhar [29]. A direct proof of this undecidability result for ATL_i^* is also presented in [11]. However, there are classes of multiplayer games with imperfect information that are decidable. For many years, the only known decidable case was that of hierarchical games, in which there is a total preorder among players, each player observing at least as much as those below her in this preorder [30, 22]. Recently, this result has been extended by relaxing the assumption of hierarchical observation. In particular, it has been shown that the problem remains decidable if the hierarchy can change along a play, or if transient phases without such a hierarchy are allowed [4]. We establish that these results transfer to the model-checking problem for ATL_i^*.

We remind that a concurrent game with imperfect information is a pair $((\mathcal{G}, v), W)$ where $(\mathcal{G}, v)$ is a pointed *nondeterministic* CGS_i and W is a property of infinite plays called the *winning condition*. The *strategy problem* is, given such a game, to decide whether there exists a strategy profile for the grand coalition Ag to enforce the winning condition against Nature (for more details see, *e.g.*, [4]).

Before stating our meta-theorem we need to introduce a couple of notions. First we introduce a notion of abstraction over a group of agents. Informally, abstracting a CGS_i $\mathcal{G}$ over an agent consists in erasing her from the group of agents and letting Nature play for her in $\mathcal{G}$.

DEFINITION 3   (ABSTRACTION). *Let $A \subseteq \mathrm{Ag}$ be a group of agents and let $\mathcal{G} = (V, E, \ell, \{\sim_a\}_{a \in \mathrm{Ag}})$ be a CGS_i. The abstraction of $\mathcal{G}$ from A is the nondeterministic CGS_i over set of agents $\mathrm{Ag} \setminus A$ defined as $\mathcal{G}{\uparrow}^A := (V, E', \ell, \{\sim_a\}_{a \in \mathrm{Ag} \setminus A})$, where for every $v \in V$ and $\boldsymbol{m} \in M^{\mathrm{Ag} \setminus A}$,*

$$(v, \boldsymbol{m}, v') \in E' \text{ if } \exists \boldsymbol{m}' \in M^A \text{ s.t. } E(v, (\boldsymbol{m}, \boldsymbol{m}')) = v'.$$

Thanks to this notion we can define the following problem:

DEFINITION 4   (*A*-STRATEGY PROBLEM). *The A-strategy problem takes as input a pointed CGS_i $(\mathcal{G}, v)$, a set $A \subseteq \mathrm{Ag}$ of agents and a winning condition W, and returns the answer to the strategy problem for the game $((\mathcal{G}{\uparrow}^{\mathrm{Ag} \setminus A}, v), W)$.*

The A-strategy problem for $(\mathcal{G}, v)$ with winning condition W thus consists in deciding whether there is a strategy profile for agents in A to enforce W against everybody else.

Finally we introduce the following notion, which simply captures the change of initial position in a game from a position v to another position v' reachable from v:

DEFINITION 5   (INITIAL SHIFTING). *Let $\mathcal{G}$ be a CGS_i and let $v, v' \in \mathcal{G}$. The pointed CGS_i $(\mathcal{G}, v')$ is an* initial shifting *of $(\mathcal{G}, v)$ if v' is reachable from v in $\mathcal{G}$.*

We are now ready to state our first result.

THEOREM 1. *If $\mathcal{C}$ is a class of pointed CGS_i closed under initial shifting and such that the A-strategy problem with ω-regular objective is decidable on $\mathcal{C}$, then model checking ATL_i^* is decidable on $\mathcal{C}$.*

PROOF. Let $\mathcal{C}$ be such a class of pointed CGS_i, and let $(\varphi, (\mathcal{G}, v))$ be an instance of the model-checking problem for ATL_i^* on $\mathcal{C}$. A bottom-up algorithm consists in evaluating each innermost subformula of φ of the form $\langle A \rangle \varphi'$, where φ' is thus an LTL formula, on each position v' of $\mathcal{G}$ reachable from v. Evaluating $\langle A \rangle \varphi'$ on v' amounts to solving an instance of the A-strategy problem[3] with ω-regular objective (recall that LTL properties are ω-regular). By assumption $(\mathcal{G}, v) \in \mathcal{C}$, and because $\mathcal{C}$ is closed by initial shifting and v' is reachable from v, we have that $(\mathcal{G}, v') \in \mathcal{C}$. Also by assumption, the A-strategy problem for ω-regular winning conditions is decidable on $\mathcal{C}$. We thus have an algorithm to evaluate each $\langle A \rangle \varphi'$ on each v'. One can then mark positions of the game with fresh atomic propositions indicating where these formulas hold, and repeat the procedure until all strategic operators have been eliminated. It then remains to evaluate a boolean formula in the initial position v. $\square$

Let us recall for which classes of nondeterministic CGS_i the strategy problem is known to be decidable. A (nondeterministic or deterministic) CGS_i $\mathcal{G}$ has *hierarchical observation* if there exists a total preorder $\preccurlyeq$ over Ag such that if $a \preccurlyeq b$ and $v \sim_a v'$, then $v \sim_b v'$. This notion was refined in [4] to take into account the agents' memory, using the notion of *information set*: for a finite play $\rho \in \mathrm{Plays}(\mathcal{G}, v)$ and an agent a, the *information set* of agent a after ρ is $I^a(\rho) := \{\rho' \in \mathrm{Plays}(\mathcal{G}, v) \mid \rho \sim_a \rho'\}$. A finite play ρ yields *hierarchical information* if there is a total preorder $\preccurlyeq$ over Ag such that if $a \preccurlyeq b$, then $I^a(\rho) \subseteq I^b(\rho)$. If all finite plays in $\mathrm{Plays}(\mathcal{G}, v)$ yield hierarchical information for the same preorder over agents, $(\mathcal{G}, v)$ yields *static hierarchical information*. If this preorder can vary depending on the play, $(\mathcal{G}, v)$ yields *dynamic hierarchical information*. The last generalisation consists in allowing for transient phases without hierarchical information: if every infinite play in $\mathrm{Plays}(\mathcal{G}, v)$ has infinitely many prefixes that yield hierarchical information, $(\mathcal{G}, v)$ yields *recurring hierarchical information*.

PROPOSITION 1. *Hierarchical observation as well as static, dynamic and recurring hierarchical information are preserved by abstraction.*

PROPOSITION 2. *Hierarchical observation as well as static, dynamic and recurring hierarchical information are preserved by initial shifting.*

This is obvious for hierarchical observation. For the other cases we establish Lemma 1 below. It is then easy to check that Proposition 2 holds.

LEMMA 1. *If a finite play $v \cdot \rho \cdot v' \cdot \rho'$ yields hierarchical information in $(\mathcal{G}, v)$, so does $v' \cdot \rho'$ in $(\mathcal{G}, v')$, with the same preorder among agents.*

Let $\mathcal{C}_{\mathrm{obs}}$ (resp. $\mathcal{C}_{\mathrm{stat}}$, $\mathcal{C}_{\mathrm{dyn}}$, $\mathcal{C}_{\mathrm{rec}}$) be the class of pointed CGS_i with hierarchical observation (resp. static, dynamic, recurring hierarchical information). We instantiate Theorem 1 to obtain three decidability results for ATL_i^*.

THEOREM 2. *Model checking ATL_i^* is decidable on the class of CGS_i with hierarchical observation.*

[3]Observe that if $A = \mathrm{Ag}$ then $\mathcal{G}\!\uparrow^{\mathrm{Ag}\setminus A} = \mathcal{G}$, and Nature thus does not do anything. This is coherent with the fact that for agents with perfect recall $\langle \mathrm{Ag} \rangle \varphi \equiv \mathbf{E}\varphi$, where $\mathbf{E}$ is the CTL path quantifier, even for imperfect information.

PROOF. By Proposition 2, $\mathcal{C}_{\mathrm{obs}}$ is closed under initial shifting. It is proven in [22] that the strategy problem is decidable for games with hierarchical observation and ω-regular objectives. Since, by Proposition 1, all pointed nondeterministic CGS_i obtained by abstracting agents from CGS_i in $\mathcal{C}_{\mathrm{obs}}$ also yield hierarchical observation, we get that the A-strategy problem with ω-regular objectives is decidable on $\mathcal{C}_{\mathrm{obs}}$. We can therefore apply Theorem 1 on $\mathcal{C}_{\mathrm{obs}}$. $\square$

It is proven in [4] that the strategy problem with ω-regular objectives is also decidable for games with static hierarchical information and for games with dynamic hierarchical information. Since Proposition 1 and Proposition 2 also hold for $\mathcal{C}_{\mathrm{stat}}$ and $\mathcal{C}_{\mathrm{dyn}}$, with the same argument as in the proof of Theorem 2, we obtain the following results as consequences of Theorem 1:

THEOREM 3. *Model checking ATL_i^* is decidable on the class of CGS_i with static hierarchical information.*

THEOREM 4. *Model checking ATL_i^* is decidable on the class of CGS_i with dynamic hierarchical information.*

Note that in fact, since $\mathcal{C}_{\mathrm{obs}} \subset \mathcal{C}_{\mathrm{stat}} \subset \mathcal{C}_{\mathrm{dyn}}$, Theorem 2 and Theorem 3 are also obtained as corollaries of Theorem 4, but we wanted to illustrate how Theorem 1 can be applied to obtain decidability results for different classes of CGS_i.

REMARK 1. *The last result in [4] establishes that the strategy problem is decidable for games with recurring hierarchical information, but only for* observable ω-regular winning conditions, *i.e., when all agents can tell whether a play is winning or not. Now considering ATL_i^* on $\mathcal{C}_{dyn}$ we could require atomic propositions to be observable for all agents; in that case we could evaluate the inner-most strategy quantifiers using the above-mentioned result. But then the fresh atomic propositions that mark positions where these subformulas hold (see the proof of Theorem 1) would not, in general, be observable by all agents. So on $\mathcal{C}_{rec}$ we could obtain a decision procedure for the fragment of ATL_i^* without nested non-trivial strategy quantifiers, where "non-trivial" means for coalitions other than the empty coalition or the one made of all agents (which, we recall, are simply the CTL path quantifiers). We do not state it explicitly due to lack of space and because it does not seem of much interest.*

Concerning complexity, the strategy problem for games with imperfect information and hierarchical observation is already nonelementary [32, 29], hence the following result:

COROLLARY 1. *Model checking ATL_i^* is nonelementary on games with hierarchical observation, hence also for games with static or dynamic hierarchical information.*

EXAMPLE 1. *Our decidability results typically apply to systems with different security levels, where higher levels have access to more data (i.e., can observe more). In such systems, by Theorem 4, we can model check all ATL_i^* formulas, even if the distribution of clearance levels between agents can vary in different scenarios/plays and also along time (an agent may get access to a higher security clearance).*

We now turn to ATL with imperfect information and strategy context, and study its model-checking problem.

4. ATL$_i$ WITH STRATEGY CONTEXT

While in ATL strategies for all agents are forgotten each time a new strategy quantifier is met, in ATL with strategy context (ATL$_{sc}$) [5, 9, 24] agents keep using the same strategy as long as the formula does not say otherwise. In this section we consider ATL$_{sc}$ with imperfect information (ATL$_{sc,i}$). As far as we know, the only existing work on this logic is [25], which proved its model-checking problem to be decidable in the case where all agents have the same observation of the game. We extend significantly this result by establishing that the model-checking problem is decidable as long as strategy quantification is *hierarchical*, in the sense that if there is a strategy quantification for agent a nested in a strategy quantification for agent b, then b should observe no more than a. In other terms, innermost strategic quantifications should concern agents who observe more.

4.1 Syntax and semantics

The models are still CGS$_i$. To remember which agents are currently bound to a strategy, and what these strategies are, the semantics uses *strategy contexts*. Formally, a strategy context for a set of agents $B \subseteq \mathrm{Ag}$ is a strategy profile σ_B. We define the composition of strategy contexts as follows. If σ_B is a strategy context for B and σ_A is a new strategy profile for coalition A, we let $\sigma_A \circ \sigma_B$ be the strategy context for $A \cup B$ defined as $\sigma_{A \cup B} : a \mapsto \begin{cases} \sigma_A(a) & \text{if } a \in A, \\ \sigma_B(a) & \text{otherwise} \end{cases}$.

So if a is assigned a strategy by σ_A, her strategy in $\sigma_A \circ \sigma_B$ is $\sigma_A(a)$. If she is not assigned a strategy by σ_A her strategy remains the one given by σ_B, if any.

Also, given a strategy context σ_B and a set of agents $A \subseteq \mathrm{Ag}$, we let $(\sigma_B)_{\setminus A}$ be the strategy context obtained by restricting σ_B to the domain $B \setminus A$.

Finally, because agents who do not change their strategy keep playing the one they were assigned, if any, we cannot forget the past at each strategy quantifier, as in the semantics of ATL$_i^*$ (see Section 3.1). We thus define the outcome of a strategy profile σ_A after a finite play ρ, written $\mathrm{Out}(\rho, \sigma_A)$, as the set of infinite plays π that start with ρ and then follow σ_A: $\pi \in \mathrm{Out}(\rho, \sigma_A)$ if $\pi = \rho \cdot \pi'$ for some π', and for all $i \geq |\rho| - 1$, there exists a joint move $\boldsymbol{m} \in \mathrm{M}^{\mathrm{Ag}}$ such that $E(\pi_i, \boldsymbol{m}) = \pi_{i+1}$ and for each $a \in A$, $m_a = \sigma_a(\pi[0, i])$.

To differentiate from ATL*, in ATL$_{sc}^*$ the strategy quantifier for a coalition A is written $\langle A \rangle$ instead of $\langle A \rangle$. ATL$_{sc}^*$ also has an additional operator, $(\!|A|\!)$, that releases agents in A from their current strategy, if they have one. The syntax of ATL$_{sc,i}^*$ is the same as that of ATL$_{sc}^*$ and is thus given by the following grammar:

$$\varphi ::= p \mid \neg\varphi \mid \varphi \vee \varphi \mid \langle A \rangle \varphi \mid (\!|A|\!)\varphi \mid \mathbf{X}\varphi \mid \varphi\mathbf{U}\varphi,$$

where $p \in \mathcal{AP}$ and $A \subseteq \mathrm{Ag}$. We use standard abbreviations: $\top := p \vee \neg p$, $\bot := \neg\top$, $\mathbf{F}\varphi := \top\mathbf{U}\varphi$, and $\mathbf{G}\varphi := \neg\mathbf{F}\neg\varphi$.

REMARK 2. *In [24] the syntax of ATL$_{sc}^*$ contains in addition operators $\langle \overline{A} \rangle$ and $(\!|\overline{A}|\!)$ for complement coalitions. While they add expressivity when the set of agents is not fixed, and are thus of interest when considering expressivity or satisfiability, they are redundant if we consider model checking, which is our case in this work. To simplify presentation we thus choose not to consider them here.*

The semantics of ATL$_{sc,i}^*$ is defined with regards to a CGS$_i$ $\mathcal{G} = (V, E, \ell, \{\sim_a\}_{a \in \mathrm{Ag}})$, an infinite play π, a position $i \in \mathbb{N}$

along this play, and a strategy context σ_B. The semantics is defined by induction on formulas:

$\mathcal{G}, \pi, i \models_{\sigma_B} p$ if $p \in \ell(\pi_i)$
$\mathcal{G}, \pi, i \models_{\sigma_B} \neg\varphi$ if $\mathcal{G}, \pi, i \not\models_{\sigma_B} \varphi$
$\mathcal{G}, \pi, i \models_{\sigma_B} \varphi \vee \varphi'$ if $\mathcal{G}, \pi, i \models_{\sigma_B} \varphi$ or $\mathcal{G}, \pi, i \models_{\sigma_B} \varphi'$
$\mathcal{G}, \pi, i \models_{\sigma_B} \langle A \rangle \varphi$ if there exists a strategy profile σ_A s.t. for all $\pi' \in \mathrm{Out}(\pi[0, i], \sigma_A \circ \sigma_B)$, $\mathcal{G}, \pi', i \models_{\sigma_A \circ \sigma_B} \varphi$
$\mathcal{G}, \pi, i \models_{\sigma_B} (\!|A|\!)\varphi$ if $\mathcal{G}, \pi, i \models_{(\sigma_B)_{\setminus A}} \varphi$
$\mathcal{G}, \pi, i \models_{\sigma_B} \mathbf{X}\varphi$ if $\mathcal{G}, \pi, i+1 \models_{\sigma_B} \varphi$
$\mathcal{G}, \pi, i \models_{\sigma_B} \varphi\mathbf{U}\varphi'$ if there exists $j \geq i$ s.t. $\mathcal{G}, \pi, j \models_{\sigma_B} \varphi'$ and, for all k such that $i \leq k < j$, $\mathcal{G}, \pi, k \models_{\sigma_B} \varphi$.

The notion of closed formula is as defined in Section 3.1 and once more, the semantics of a closed formula φ being independent from the future, we may write $\mathcal{G}, v \models_{\sigma_B} \varphi$ instead of $\mathcal{G}, \pi, 0 \models_{\sigma_B} \varphi$ for any infinite play π that starts in position v. We also write $\mathcal{G}, v \models \varphi$ if $\mathcal{G}, v \models_{\sigma_\emptyset} \varphi$, that is if φ holds in v with the empty strategy context.

The *model-checking problem for* ATL$_{sc,i}^*$ consists in deciding, given a closed ATL$_{sc,i}^*$ formula φ and a finite pointed CGS$_i$ ($\mathcal{G}, v$), whether $\mathcal{G}, v \models \varphi$.

We now present QCTL* with imperfect information, or QCTL$_i^*$ for short, before proving our main result on the model-checking problem for ATL$_{sc,i}^*$ by reducing it to the model-checking problem for a decidable fragment of QCTL$_i^*$.

4.2 QCTL* with imperfect information

Quantified CTL*, or QCTL* for short, is an extension of CTL* with second-order quantifiers on atomic propositions that has been well studied [36, 20, 21, 23]. It has recently been further extended to take into account imperfect information, resulting in the logic called QCTL* with imperfect information, or QCTL$_i^*$ [3]. We briefly present this logic, as well as a decidability result on its model-checking problem proved in [3] and that we rely on to establish our result on the model checking of ATL$_{sc,i}^*$.

Imperfect information is incorporated into QCTL* by considering Kripke models with internal structure in the form of local states, like in distributed systems (see for instance [17]), and then parameterising quantifiers on atomic propositions with observations that define what portions of the states a quantifier can "observe". The semantics is then adapted to capture the idea of quantification on atomic propositions being made with partial observation.

Let us fix a collection $\{L_i\}_{i \in [n]}$ of n disjoint finite sets of *local states*. We also let $X_n = L_1 \times \ldots \times L_n$.

DEFINITION 6. *A compound Kripke structure (CKS) over AP is a Kripke structure $\mathcal{S} = (S, R, \ell)$ such that $S \subseteq X_n$.*

The syntax of QCTL$_i^*$ is that of QCTL*, except that quantifiers over atomic propositions are parameterised by a set of indices that defines what local states the quantifier can "observe". It is thus defined by the following grammar:

$$\varphi := p \mid \neg\varphi \mid \varphi \vee \varphi \mid \mathbf{E}\varphi \mid \exists^o p.\ \varphi \mid \mathbf{X}\varphi \mid \varphi\mathbf{U}\varphi$$

where $p \in \mathcal{AP}$ and $o \subset \mathbb{N}$ is a finite set of indices. As usual, we let $\mathbf{A}\varphi := \neg\mathbf{E}\neg\varphi$.

A finite set $o \subset \mathbb{N}$ is called an *observation*, and two states $s = (l_1, \ldots, l_n)$ and $s' = (l'_1, \ldots, l'_n)$ are *o-indistinguishable*, written $s \approx_o s'$, if for all $i \in [n] \cap o$, it holds that $l_i = l'_i$.

The intuition is that a quantifier with observation o must choose the valuation of atomic propositions *uniformly* with respect to o. Note that in [3], two semantics are considered for QCTL_i^*, just like in [23] for QCTL^*: the structure semantics and the tree semantics. In the former, formulas are evaluated directly on the structure, while in the latter the structure is first unfolded into an infinite tree. Here we only present the tree semantics, as it is this one that allows us to capture agents with perfect recall. But we first need a few more definitions.

For $p \in \mathcal{AP}$, two labelled trees $t = (\tau, \ell)$ and $t' = (\tau', \ell')$ are *equivalent modulo p*, written $t \equiv_p t'$, if $\tau = \tau'$ and for each node $u \in \tau$, $\ell(u) \setminus \{p\} = \ell'(u) \setminus \{p\}$. So $t \equiv_p t'$ if they are the same trees, except for the labelling of proposition p.

This notion of equivalence modulo p is the one used to define quantification on atomic propositions in QCTL^*: intuitively, an existential quantification over p chooses a new labelling for valuation p, all else remaining the same, and the evaluation of the formula continues from the current node with the new labelling. For imperfect information we need to express the fact that this new labelling for a proposition is done uniformly with regards to the quantifier's observation.

First, we define the notion of indistinguishability between two nodes in the unfolding of a CKS. Let o be an observation, let τ be an X_n-tree (which may be obtained by unfolding some pointed CKS), and let $u = s_0 \ldots s_i$ and $u' = s_0' \ldots s_j'$ be two nodes in τ. The nodes u and u' are *o-indistinguishable*, written $u \approx_o u'$, if $i = j$ and for all $k \in \{0, \ldots, i\}$, we have $s_k \approx_o s_k'$. Observe that this definition corresponds to the notion of synchronous perfect recall in CGS_i (see Section 3.1). We now define what it means for the labelling of an atomic proposition to be uniform with regards to an observation.

DEFINITION 7. *Let $t = (\tau, \ell)$ be a labelled X_n-tree, let $p \in \mathcal{AP}$ be an atomic proposition and $o \subset \mathbb{N}$ an observation. Tree t is o-uniform in p if for every pair of nodes $u, u' \in \tau$ such that $u \approx_o u'$, we have $p \in \ell(u)$ iff $p \in \ell(u')$.*

The satisfaction relation $\models_t$ (t is for *tree semantics*) is now defined as follows, where $t = (\tau, \ell)$ is a labelled X_n-tree, λ is a path in τ and $i \in \mathbb{N}$ a position along that branch:

$$
\begin{array}{lll}
t, \lambda, i \models_t p & \text{if } p \in \ell(\lambda_i) \\
t, \lambda, i \models_t \neg\varphi & \text{if } t, \lambda, i \not\models_t \varphi \\
t, \lambda, i \models_t \varphi \vee \varphi' & \text{if } t, \lambda, i \models_t \varphi \text{ or } t, \lambda, i \models_t \varphi' \\
t, \lambda, i \models_t \mathbf{E}\varphi & \text{if there exists } \lambda' \in Paths(\lambda_i) \\
& \quad \text{such that } t, \lambda', 0 \models_t \varphi \\
t, \lambda, i \models_t \exists^o p.\, \varphi & \text{if there exists } t' \equiv_p t \text{ such that} \\
& \quad t' \text{ is } o\text{-uniform in } p \text{ and } t', \lambda, i \models_t \varphi \\
t, \lambda, i \models_t \mathbf{X}\varphi & \text{if } t, \lambda, i+1 \models_t \varphi \\
t, \lambda, i \models_t \varphi\mathbf{U}\varphi' & \text{if there exists } j \geq i \text{ such that } t, \lambda, j \models_t \varphi' \\
& \quad \text{and for } i \leq k < j,\ t, \lambda, j \models_t \varphi
\end{array}
$$

Similarly to ATL_i^* and $\mathsf{ATL}_{sc,i}^*$, we say that a QCTL_i^* formula is *closed* if all temporal operators are in the scope of a path quantifier. The semantics of such formulas depending only on the current node, for a closed formula φ we may write $t \models_t \varphi$ for $t, r \models_t \varphi$, where r is the root of t, and given a CGS_i $\mathcal{G}$, a state s and a QCTL_i^* formula φ, we write $\mathcal{S}, s \models_t \varphi$ if $t_{\mathcal{S}}(s) \models_t \varphi$.

REMARK 3. *In [3] the syntax is presented with path formulas distinguished from state formulas, and the semantics is defined accordingly. To make the presentation more uniform with that of $\mathsf{ATL}_{sc,i}$ we chose here a different, but equivalent, presentation.*

REMARK 4. *Note that when n is fixed, the propositional quantifier with perfect information from QCTL^* is equivalent to the QCTL_i^* quantifier that observes all the components, i.e., the quantifier parameterised with observation $[n]$.*

The model-checking problem for QCTL_i^* is the following: given a closed QCTL_i^* formula φ and a finite pointed CKS $(\mathcal{S}, s)$, decide whether $\mathcal{S}, s \models_t \varphi$.

We now define the class of QCTL_i^* formulas for which the model-checking problem is known to be decidable with the tree semantics.

DEFINITION 8. *A QCTL_i^* formula φ is hierarchical if for all subformulas φ_1, φ_2 of the form $\varphi_1 = \exists^{o_1} p_1.\, \varphi_1'$ and $\varphi_2 = \exists^{o_2} p_2.\, \varphi_2'$ where φ_2 is a subformula of φ_1', we have $o_1 \subseteq o_2$.*

The following result is proved in [3], where $\mathsf{QCTL}_{i,\subset}^*$ is the set of hierarchical QCTL_i^* formulas:

THEOREM 5 ([3]). *Model checking $\mathsf{QCTL}_{i,\subset}^*$ with tree semantics is decidable.*

4.3 Model checking $\mathsf{ATL}_{sc,i}^*$

We establish that model checking $\mathsf{ATL}_{sc,i}^*$ is decidable on a class of instances whose definition relies on the notion of *hierarchical observation*.

DEFINITION 9. *Let $\mathcal{G} = (V, E, \ell, \{\sim_a\}_{a \in \mathrm{Ag}})$ be a CGS_i, and let $a, b \in \mathrm{Ag}$ be two agents. Agent a observes no more than agent b in $\mathcal{G}$, written $a \preccurlyeq_{\mathcal{G}} b$, if for every pair of positions $v, v' \in V$, $v \sim_b v'$ implies $v \sim_a v'$. We say that $A \subseteq \mathrm{Ag}$ is hierarchical in $\mathcal{G}$ if $\preccurlyeq_{\mathcal{G}}$ is a total preorder on A.*

If a set of agents A is hierarchical in a CGS_i $\mathcal{G}$, we thus may talk about maximal and minimal agents in A, referring to maximal and minimal elements of A for the relation $\preccurlyeq_{\mathcal{G}}$.

The essence of the requirement that makes the problem decidable is the same as for the decidability result on QCTL_i^* (Theorem 5): nesting of quantifiers (here, strategy quantifiers) should be hierarchical, with those observing more inside those observing less. However, unlike in QCTL_i^*, in $\mathsf{ATL}_{sc,i}^*$ observations are not part of formulas, but rather they are given by the models. We thus define the notion of hierarchical $\mathsf{ATL}_{sc,i}^*$ formula with respect to a given CGS_i:

DEFINITION 10. *Let Φ be an $\mathsf{ATL}_{sc,i}^*$ formula and $\mathcal{G}$ a CGS_i. We say that Φ is hierarchical in $\mathcal{G}$ if:*

- *for every subformula φ of the form $\varphi = \langle A \rangle \varphi'$, A is hierarchical in $\mathcal{G}$, and*
- *for all subformulas φ_1, φ_2 of the form $\varphi_1 = \langle A_1 \rangle \varphi_1'$ and $\varphi_2 = \langle A_2 \rangle \varphi_2'$ where φ_2 is a subformula of φ_1', maximal agents of A_1 observe no more than minimal agents of A_2.*

An instance $(\Phi, (\mathcal{G}, v))$ of the model-checking problem for $\mathsf{ATL}_{sc,i}^*$ is hierarchical if Φ is hierarchical in $\mathcal{G}$.

In the rest of the section we establish the following:

THEOREM 6. *Model checking $\mathsf{ATL}_{sc,i}^*$ is decidable on the class of hierarchical instances.*

EXAMPLE 2. *Consider the security levels scenario of Example 1, and assume that $a \preccurlyeq_{\mathcal{G}} b \preccurlyeq_{\mathcal{G}} c$. Then $\langle a \rangle [b \cdot] \langle c \rangle \mathbf{G} p$, which says that a and c can collaborate against an unreliable agent b to ensure some safety property, as long as agent c*

can adapt her strategy to that of agent b, forms a hierarchical instance with $\mathcal{G}$. On the other hand, $\langle c \rangle [\cdot b \cdot] \langle a \rangle \mathbf{G} p$ does not form a hierarchical instance with $\mathcal{G}$.

Further, the decidable fragment of $\mathsf{ATL}^*_{sc,i}$ is not restricted to models where there is a total order on agents' observations: assume a fourth agent d that observes more than a and b, but whose security level is incomparable to that of c. On such models, the following formulas form hierarchical instances that we can model check: $\langle a, b, c \rangle \mathbf{F} p \vee \langle a, b, d \rangle \mathbf{F} p$, which means that a and b can achieve p by collaborating with c or with d, and $[\cdot a, b \cdot](\langle c \rangle \mathbf{F} p \wedge \langle d \rangle \mathbf{G} q)$, which means that for all strategies of a and b, c can enforce that p is reached, and d can enforce that q always holds.

To establish Theorem 6 we build upon the proof in [24] that establishes the decidability of the model-checking problem for ATL^*_{sc} by reduction to the model-checking problem for QCTL^*. The main difference is that we reduce to the model-checking problem for QCTL^*_i instead, using quantifiers parameterised with observations corresponding to agents' observations. We also need a couple of adjustments to obtain formulas in the decidable fragment $\mathsf{QCTL}^*_{i,\subset}$.

Let $(\Phi, (\mathcal{G}, v_\iota))$ be a hierarchical instance of the $\mathsf{ATL}^*_{sc,i}$ model-checking problem, where $\mathcal{G} = (V, E, \ell, \{\sim_a\}_{a \in \mathrm{Ag}})$ is a CGS_i over AP. In the reduction we will transform Φ into an equivalent QCTL^*_i formula Φ' in which we need to refer to the current position in the model $\mathcal{G}$, and also to talk about moves taken by agents. To do so, we consider the additional sets of atomic propositions $AP_v := \{p_v \mid v \in V\}$ and $AP_m := \{p^a_m \mid a \in \mathrm{Ag} \text{ and } m \in \mathrm{M}\}$, that we take disjoint from AP.

First we define the CKS $\mathcal{S}_{\mathcal{G}}$ on which Φ' will be evaluated. Since the models of the two logics use different ways to represent imperfect information (equivalence relations on positions for CGS_i and local states for CKS) this requires a bit of work. First, for each $v \in V$ and $a \in \mathrm{Ag}$, let us define $[v]_a$ as the equivalence class of v for relation $\sim_a$. Now, noting $\mathrm{Ag} = \{a_1, \ldots, a_n\}$, we define for each $i \in [n]$ the set $L_i := \{[v]_{a_i} \mid v \in V\}$ of local states for agent a_i. Since we need to know the actual position of the CGS_i to define the dynamics, we also let $L_{n+1} := V$. States of $\mathcal{S}_{\mathcal{G}}$ will thus be tuples in $L_1 \times \ldots \times L_n \times L_{n+1}$. For each $v \in \mathcal{G}$, let $s_v := ([v]_{a_1}, \ldots, [v]_{a_n}, v)$ be its corresponding state in $\mathcal{S}_{\mathcal{G}}$.

We can now define $\mathcal{S}_{\mathcal{G}} := (S, R, \ell')$, where

- $S := \{s_v \mid v \in V\}$,
- $R := \{(s_v, s_{v'}) \mid \exists \boldsymbol{m} \in \mathrm{M}^{\mathrm{Ag}} \text{ s.t. } E(v, \boldsymbol{m}) = v'\}$, and
- $\ell'(s_v) := \ell(v) \cup \{p_v\}$.

To make the connection between finite plays in $\mathcal{G}$ and nodes in tree unfoldings of $\mathcal{S}_{\mathcal{G}}$, let us define, for every finite play $\rho = v_0 \ldots v_k$, the node $u_\rho := s_{v_0} \ldots s_{v_k}$ in $t_{\mathcal{S}_{\mathcal{G}}}(s_{v_0})$ (which exists, by definition of $\mathcal{S}_{\mathcal{G}}$ and of tree unfoldings). Observe that the mapping $\rho \mapsto u_\rho$ is in fact a bijection between the set of finite plays starting in a given position v and the set of nodes in $t_{\mathcal{S}_{\mathcal{G}}}(s_v)$.

Now it should be clear that giving to a propositional quantifier in QCTL^*_i observation $o_i := \{i\}$, for $i \in [n]$, amounts to giving him the same observation as agent a_i. Formally, one can prove the following lemma, simply by applying the definitions of observational equivalence in the two frameworks:

LEMMA 2. *For all finite plays ρ, ρ' starting in position v, $\rho \sim_{a_i} \rho'$ iff $u_\rho \approx_{o_i} u_{\rho'}$ in $t_{\mathcal{S}_{\mathcal{G}}}(s_v)$.*

We now describe the translation[4] from $\mathsf{ATL}_{sc,i}$ formulas to QCTL^*_i formulas. First we recall the translation from [24] for the perfect-information case.

The translation from ATL_{sc} to QCTL^* is parameterised by a coalition $B \subset \mathrm{Ag}$, that conveys the set of agents who are currently bound to a strategy. It is defined by induction on Φ as follows:

$$\overline{p}^B := p \qquad\qquad \overline{\neg\varphi}^B := \neg\overline{\varphi}^B$$
$$\overline{\varphi \vee \varphi'}^B := \overline{\varphi}^B \vee \overline{\varphi'}^B \qquad \overline{(\!|A|\!)\varphi}^B := \overline{\varphi}^{B \setminus A}$$
$$\overline{\mathbf{X}\varphi}^B := \mathbf{X}\overline{\varphi}^B \qquad \overline{\varphi \mathbf{U} \varphi'}^B := \overline{\varphi}^B \mathbf{U} \overline{\varphi'}^B$$

The only non-trivial case is for formulas of the form $\langle A \rangle \varphi$. For the rest of the section, we let $\mathrm{M} = \{m_1, \ldots, m_l\}$. Now, if $A = \{a_{i_1}, \ldots, a_{i_k}\}$, we define

$$\overline{\langle A \rangle \varphi}^B := \exists m_1^{a_{i_1}} \ldots m_l^{a_{i_1}} \ldots m_1^{a_{i_k}} \ldots m_l^{a_{i_k}} p_{\mathrm{out}}.$$
$$\left(\Phi_{\mathrm{strat}}(A) \wedge \Phi_{\mathrm{out}}(A \cup B) \wedge \mathbf{A}(\mathbf{G} p_{\mathrm{out}} \to \overline{\varphi}^{A \cup B}) \right),$$

where

$$\Phi_{\mathrm{strat}}(A) := \bigwedge_{a \in A} \mathbf{AG} \bigvee_{m \in \mathrm{M}} (m^a \wedge \bigwedge_{m' \neq m} \neg m'^a)$$

and

$$\Phi_{\mathrm{out}}(A) := p_{\mathrm{out}} \wedge \mathbf{AG}\left[\neg p_{\mathrm{out}} \to \mathbf{AX} \neg p_{\mathrm{out}}\right] \wedge \mathbf{AG}\left[p_{\mathrm{out}} \to \right.$$
$$\left. \bigvee_{v \in V} \bigvee_{\boldsymbol{m} \in \mathrm{M}^A} \left(p_v \wedge p_{\boldsymbol{m}} \wedge \mathbf{AX}\left(\bigvee_{v' \in E(v,\boldsymbol{m})} p_{v'} \leftrightarrow p_{\mathrm{out}} \right) \right) \right].$$

In $\Phi_{\mathrm{out}}(A)$, for $\boldsymbol{m} = (m_a)_{a \in A} \in \mathrm{M}^A$, notation $p_{\boldsymbol{m}}$ stands for the propositional formula $\bigwedge_{a \in A} m_a^a$ which characterises the joint move $\boldsymbol{m}$ that agents in A play in v. Also, $E(v, \boldsymbol{m})$ is the set of possible next positions when the current one is v and agents in A play $\boldsymbol{m}$, and it is defined as $E(v, \boldsymbol{m}) := \{E(v, (\boldsymbol{m}, \boldsymbol{m}')) \mid \boldsymbol{m}' \in \mathrm{M}^{\mathrm{Ag} \setminus A}\}$.

The idea of this translation is the following: first, for each agent $a \in A$ and each possible move $m \in \mathrm{M}$, an existential quantification on the atomic proposition m^a "chooses" for each finite play ρ of $(\mathcal{G}, v_\iota)$ (or, equivalently, for each node u_ρ of $t_{\mathcal{S}_{\mathcal{G}}}(s_{v_\iota})$) whether agent a plays move m in ρ or not, coded by m^a being chosen to be true a in ρ or not. Formula $\Phi_{\mathrm{strat}}(A)$ ensures that each agent a chooses exactly one move in each finite play, and thus that atomic propositions m^a characterise a strategy for her. An atomic proposition p_{out} is then used to mark the paths that follow the currently fixed strategies: formula $\Phi_{\mathrm{out}}(A \cup B)$ states that p_{out} marks exactly the outcome of strategies just chosen for agents in A, as well as those of agents in B, that were chosen previously by a strategy quantifier "higher" in Φ.

Note that we simplified slightly $\Phi_{\mathrm{strat}}(A)$ and $\Phi_{\mathrm{out}}(A)$, using the fact that unlike in [24], we have assumed in our definition of CGS_i that the set of available moves is the same for all agents in all positions (see Footnote 2).

It is proven in [24] that this translation is correct, in the sense that for every ATL_{sc} closed formula φ and pointed perfect-information concurrent game structure $(\mathcal{G}, v)$, letting

---

[4]Here we abuse language: the construction depends on the model $\mathcal{G}$ and is therefore not a translation in the usual sense.

$\mathcal{S}_{\mathcal{G}}$ be as described above but removing the local states for all agents and keeping only the L_{n+1} component, we have:

$$\mathcal{G}, v \models \varphi \text{ iff } t_{\mathcal{S}_{\mathcal{G}}}(s_v) \models_t \overline{\varphi}^{\emptyset}.$$

We now explain how to adapt this translation to the case of imperfect information. Observe that the only difference between ATL^*_{sc} and $\mathsf{ATL}^*_{sc,i}$ is that in the latter, strategies must be defined uniformly over indistinguishable finite plays, *i.e.*, a strategy σ for an agent a must be such that if $\rho \sim_a \rho'$, then $\sigma(\rho) = \sigma(\rho')$. To enforce that the strategies coded by atomic propositions m^a in $\overline{\langle\!\langle A \rangle\!\rangle \varphi}^B$ are uniform, we use the propositional quantifiers with partial observation of QCTL^*_i. Formally, we define a translation $\smile^B$ from $\mathsf{ATL}^*_{sc,i}$ to QCTL^*_i. It is defined exactly as the one from ATL^*_{sc} to QCTL^*, except for the following inductive case.

If $A = \{a_{i_1}, \ldots, a_{i_k}\}$ we let

$$\widetilde{\langle\!\langle A \rangle\!\rangle \varphi}^B := \exists^{o_{i_1}} m_1^{a_{i_1}} \ldots m_l^{a_{i_1}} \ldots \exists^{o_{i_k}} m_1^{a_{i_k}} \ldots m_l^{a_{i_k}} \exists p_{\text{out}}.$$
$$\left(\Phi_{\text{strat}}(A) \wedge \Phi_{\text{out}}(A \cup B) \wedge \mathbf{A}(\mathbf{G} p_{\text{out}} \rightarrow \widetilde{\varphi}^{A \cup B}) \right),$$

where $\Phi_{\text{strat}}(A)$ and $\Phi_{\text{out}}(A)$ are defined as before, and $\exists p_{\text{out}}$ is a macro for $\exists^{\{1,\ldots,n+1\}} p_{\text{out}}$ (see Remark 4).

So the only difference from the previous translation is that now, the labelling of each atomic proposition m^{a_i} must be o_i-uniform. This means that if two nodes u and u' in $t_{\mathcal{S}_{\mathcal{G}}}(s_{v_\iota})$ are o_i-indistinguishable, then u is labelled with m^{a_i} if and only if u' also is. In other words, in the strategy coded by atomic propositions m^{a_i}, agent a_i plays m in u if and only if she also plays it in u', and thus this strategy is uniform (recall that, by Lemma 2, observation o_i correctly reflects agent a_i's observation in $t_{\mathcal{S}_{\mathcal{G}}}(s_{v_\iota})$). It is then clear that this translation is correct:

$$\mathcal{G}, v_\iota \models \Phi \text{ iff } t_{\mathcal{S}_{\mathcal{G}}}(s_{v_\iota}) \models_t \widetilde{\Phi}^{\emptyset}. \tag{1}$$

However, even if we have taken $(\Phi, (\mathcal{G}, v_\iota))$ to be a hierarchical instance, $\widetilde{\Phi}^{\emptyset}$ is not in the decidable fragment $\mathsf{QCTL}^*_{i,\subset}$. Indeed, with the current definition of observations $\{o_i\}_{i \in [n]}$, hierarchical observation in $\mathcal{G}$ does not imply hierarchical observation in $\mathcal{S}_{\mathcal{G}}$: since $o_i = \{i\}$, for $i \neq j$ it is never the case that $o_i \subseteq o_j$. Still, we note that if agent a_j observes no more than agent a_i, then letting a_i see also what agent a_j sees does not increase her knowledge of the situation:

LEMMA 3. *If $a_j \preccurlyeq_{\mathcal{G}} a_i$, then for all finite plays ρ, ρ' that start in the same position, $u_\rho \approx_{o_i} u_{\rho'}$ iff $u_\rho \approx_{o_i \cup o_j} u_{\rho'}$.*

In the light of this Lemma 3, we can safely redefine observations as follows: for each $i \in [n]$, we let

$$o'_i := \bigcup_{j \mid a_j \preccurlyeq_{\mathcal{G}} a_i} o_j.$$

Observe that in fact $o'_i = \{j \mid a_j \preccurlyeq_{\mathcal{G}} a_i\}$. Informally, a quantifier with observation o'_i sees what agent a_i observes (note that $\preccurlyeq_{\mathcal{G}}$ is reflexive), as well as what agents that see no more than a_i observe.

Let us define a new version of the translation $\smile^B$. First, Φ being hierarchical in $\mathcal{G}$, for each subformula of Φ of the form $\langle\!\langle A \rangle\!\rangle \varphi$ we have that A is hierarchical in $\mathcal{G}$. It is thus possible to choose for agents in A an indexing $A = \{a_{i_1}, \ldots, a_{i_k}\}$ such that for all $1 \leq c < d \leq k$, we have $a_{i_c} \preccurlyeq_{\mathcal{G}} a_{i_d}$.

Now the translation remains the same as before except for the following inductive case:

If $A = \{a_{i_1}, \ldots, a_{i_k}\}$, where for all $1 \leq c < d \leq k$, we have $a_{i_c} \preccurlyeq_{\mathcal{G}} a_{i_d}$, we let

$$\widetilde{\langle\!\langle A \rangle\!\rangle \varphi}^B := \exists^{o'_{i_1}} m_1^{a_{i_1}} \ldots m_l^{a_{i_1}} \ldots \exists^{o'_{i_k}} m_1^{a_{i_k}} \ldots m_l^{a_{i_k}} \exists p_{\text{out}}.$$
$$\left(\Phi_{\text{strat}}(A) \wedge \Phi_{\text{out}}(A \cup B) \wedge \mathbf{A}(\mathbf{G} p_{\text{out}} \rightarrow \widetilde{\varphi}^{A \cup B}) \right),$$

where $\Phi_{\text{strat}}(A)$ and $\Phi_{\text{out}}(A)$ are defined as before.

From Lemma 3 we have that this new translation is still correct in the sense of Equation (1). In addition, for all $1 \leq c < d \leq k$ we have $o'_{i_c} \subseteq o'_{i_d}$.

Now consider formula $\widetilde{\Phi}^{\emptyset}$. Because Φ is hierarchical in $\mathcal{G}$, for every pair of subformulas φ_1, φ_2 of the form $\varphi_1 = \langle\!\langle A_1 \rangle\!\rangle \varphi'_1$ and $\varphi_2 = \langle\!\langle A_2 \rangle\!\rangle \varphi'_2$ where φ_2 is a subformula of φ'_1, maximal agents of A_1 observe no more than minimal agents of A_2. It is then easy to see that $\widetilde{\Phi}^{\emptyset}$ would be hierarchical if there were not the perfect-information quantifications on atomic proposition p_{out} that break the monotony of observations along subformulas when there are nested strategic quantifiers. We explain how to remedy this last problem.

We remove altogether proposition p_{out}, and we use instead the formula $\psi_{\text{out}}(A)$ defined below to characterise which paths are in the outcome of the currently-fixed strategies:

$$\psi_{\text{out}}(A) := \mathbf{G} \left(\bigwedge_{v \in V} \bigwedge_{\boldsymbol{m} \in \mathbf{M}^A} p_v \wedge p_{\boldsymbol{m}} \rightarrow \mathbf{X} \bigvee_{v' \in E(v, \boldsymbol{m})} p_{v'} \right).$$

Clearly, this formula holds in a path λ of $t_{\mathcal{S}_{\mathcal{G}}}(s_{v_\iota})$ marked with propositions m^a characterising strategies for agents in A, if at each point along λ corresponding to some position v, the next point in λ corresponds to a position v' that can be attained from v when agents in A each play the move prescribed by their current strategy. The last modification to $\smile^B$ is thus the following:

If $A = \{a_{i_1}, \ldots, a_{i_k}\}$, where for all $1 \leq c < d \leq k$, we have $a_{i_c} \preccurlyeq_{\mathcal{G}} a_{i_d}$, we let

$$\widetilde{\langle\!\langle A \rangle\!\rangle \varphi}^B := \exists^{o'_{i_1}} m_1^{a_{i_1}} \ldots m_l^{a_{i_1}} \ldots \exists^{o'_{i_k}} m_1^{a_{i_k}} \ldots m_l^{a_{i_k}}.$$
$$\Phi_{\text{strat}}(A) \wedge \mathbf{A} \left(\psi_{\text{out}}(A \cup B) \rightarrow \widetilde{\varphi}^{A \cup B} \right),$$

where $\Phi_{\text{strat}}(A)$ is defined as before.

It follows from the above considerations that this translation is still correct in the sense of Equation (1), and one can check that $\widetilde{\Phi}^{\emptyset}$ is a hierarchical QCTL^*_i formula. We conclude the proof by recalling that by Theorem 5, model checking $\mathsf{QCTL}^*_{i,\subset}$ is decidable.

Concerning complexity, model checking ATL_{sc} being already nonelementary [24], so is it for $\mathsf{ATL}_{sc,i}$.

5. CONCLUSION

In this work we established new decidability results for the model-checking problem of ATL^* with imperfect information and perfect recall as well as its extension with strategy context. Should new decidable classes of multiplayer games with imperfect information be discovered, and assuming the reasonable property of closure under initial shifting, our transfer theorem (Theorem 1) would entail new decidability results also for ATL^*_i. As for $\mathsf{ATL}^*_{sc,i}$, it would be interesting to investigate whether a meaningful notion of hierarchical instances based on, *e.g.*, dynamic or recurring hierarchical information instead of hierarchical observation as here, could lead to stronger decidability results.

6. ACKNOWLEDGEMENTS

REFERENCES

[1] T. Agotnes, V. Goranko, and W. Jamroga. Alternating-Time Temporal Logics with Irrevocable Strategies. In *TARK*, pages 15–24, 2007.

[2] R. Alur, T. Henzinger, and O. Kupferman. Alternating-Time Temporal Logic. *J. ACM*, 49(5):672–713, 2002.

[3] R. Berthon, B. Maubert, and A. Murano. Quantified CTL with imperfect information. *CoRR*, abs/1611.03524, 2016.

[4] D. Berwanger, A. B. Mathew, and M. van den Bogaard. Hierarchical information patterns and distributed strategy synthesis. In *ATVA 2015*, pages 378–393, 2015.

[5] T. Brihaye, A. D. C. Lopes, F. Laroussinie, and N. Markey. ATL with strategy contexts and bounded memory. In *LFCS 2009*, pages 92–106, 2009.

[6] K. Chatterjee, T. Henzinger, and N. Piterman. Strategy Logic. 208(6):677–693, 2010.

[7] E. Clarke and E. Emerson. Design and Synthesis of Synchronization Skeletons Using Branching-Time Temporal Logic. In *LP'81*, LNCS 131, pages 52–71, 1981.

[8] E. Clarke, O. Grumberg, and D. Peled. *Model Checking*. 2002.

[9] A. Da Costa, F. Laroussinie, and N. Markey. ATL with Strategy Contexts: Expressiveness and Model Checking. In *FSTTCS'10*, LIPIcs 8, pages 120–132, 2010.

[10] C. Dima, C. Enea, and D. P. Guelev. Model-checking an alternating-time temporal logic with knowledge, imperfect information, perfect recall and communicating coalitions. In *GANDALF 2010*, pages 103–117, 2010.

[11] C. Dima and F. L. Tiplea. Model-checking ATL under imperfect information and perfect recall semantics is undecidable. *CoRR*, abs/1102.4225, 2011.

[12] R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about knowledge*, volume 4. MIT press Cambridge, 1995.

[13] B. Finkbeiner and S. Schewe. Uniform distributed synthesis. In *LICS 2005*, pages 321–330, 2005.

[14] T. French. Decidability of quantifed propositional branching time logics. In *Australian Joint Conference on Artificial Intelligence*, pages 165–176. Springer, 2001.

[15] D. P. Guelev and C. Dima. Model-checking strategic ability and knowledge of the past of communicating coalitions. In *DALT 2008*, pages 75–90, 2008.

[16] D. P. Guelev, C. Dima, and C. Enea. An alternating-time temporal logic with knowledge, perfect recall and past: axiomatisation and model-checking. *Journal of Applied Non-Classical Logics*, 21(1):93–131, 2011.

[17] J. Y. Halpern and M. Y. Vardi. The complexity of reasoning about knowledge and time. i. lower bounds. *Journal of Computer and System Sciences*, 38(1):195–237, 1989.

[18] W. Jamroga and A. Murano. Module checking of strategic ability. In *AAMAS'15*, pages 227–235. ACM, 2015.

[19] P. Kazmierczak, T. Ågotnes, and W. Jamroga. Multi-agency is coordination and (limited) communication. In *PRIMA 2014*, pages 91–106, 2014.

[20] O. Kupferman. Augmenting branching temporal logics with existential quantification over atomic propositions. In *CAV'95*, LNCS 939, pages 325–338. Springer, 1995.

[21] O. Kupferman, P. Madhusudan, P. S. Thiagarajan, and M. Y. Vardi. Open systems in reactive environments: Control and synthesis. In *CONCUR'00*, LNCS 1877, pages 92–107. Springer, 2000.

[22] O. Kupferman and M. Y. Vardi. Synthesizing distributed systems. In *LICS'01*, pages 389–398, 2001.

[23] F. Laroussinie and N. Markey. Quantified CTL: expressiveness and complexity. *Logical Methods in Computer Science*, 10(4), 2014.

[24] F. Laroussinie and N. Markey. Augmenting ATL with strategy contexts. *Inf. Comput.*, 245:98–123, 2015.

[25] F. Laroussinie, N. Markey, and A. Sangnier. ATLsc with partial observation. In *GandALF'15*, pages 43–57, 2015.

[26] F. Mogavero, A. Murano, G. Perelli, and M. Vardi. Reasoning About Strategies: On the Model-Checking Problem. 15(4):34:1–42, 2014.

[27] F. Mogavero, A. Murano, and M. Y. Vardi. Relentful strategic reasoning in alternating-time temporal logic. *J. Log. Comput.*, 26(5):1663–1695, 2016.

[28] M. Pauly. A Modal Logic for Coalitional Power in Games. 12(1):149–166, 2002.

[29] G. Peterson, J. Reif, and S. Azhar. Lower bounds for multiplayer noncooperative games of incomplete information. *Computers & Mathematics with Applications*, 41(7):957–992, 2001.

[30] G. Peterson, J. Reif, and S. Azhar. Decision algorithms for multiplayer noncooperative games of incomplete information. *Computers & Mathematics with Applications*, 43(1):179–206, 2002.

[31] G. L. Peterson and J. H. Reif. Multiple-person alternation. In *FOCS'79*, pages 348–363, 1979.

[32] A. Pnueli and R. Rosner. Distributed reactive systems are hard to synthesize. In *FOCS'90*, pages 746–757, 1990.

[33] J. Queille and J. Sifakis. Specification and Verification of Concurrent Programs in Cesar. In *81*, LNCS 137, pages 337–351, 1981.

[34] S. Schewe and B. Finkbeiner. Distributed synthesis for alternating-time logics. In *ATVA'07*, pages 268–283, 2007.

[35] P. Schobbens. Alternating-Time Logic with Imperfect Recall. *ENTCS*, 85(2):82–93, 2004.

[36] A. Sistla. *Theoretical Issues in the Design and Cerification of Distributed Systems*. PhD thesis, Harvard University, Cambridge, MA, USA, 1983.

[37] D. Walther, W. van der Hoek, and M. Wooldridge. Alternating-Time Temporal Logic with Explicit Strategies. In *TARK*, pages 269–278, 2007.