

ATL_{sc} with partial observation*

François Laroussinie
LIAFA
Univ. Paris Diderot & CNRS, France

Nicolas Markey
LSV
ENS Cachan & CNRS, France

Arnaud Sangnier
LIAFA
Univ. Paris Diderot & CNRS, France

Alternating-time temporal logic with strategy contexts (ATL_{sc}) is a powerful formalism for expressing properties of multi-agent systems: it extends CTL with *strategy quantifiers*, offering a convenient way of expressing both collaboration and antagonism between several agents. Incomplete observation of the state space is a desirable feature in such a framework, but it quickly leads to undecidable verification problems. In this paper, we prove that *uniform* incomplete observation (where all players have the same observation) preserves decidability of the model checking problem, even for very expressive logics such as ATL_{sc}.

1 Introduction

Model checking is a powerful technique for automatically checking properties of computerized systems [Pnu77, CE82, QS82]. Model-checking algorithms classically take as input a model of the system under analysis (e.g. a finite-state automaton), and a formal property (expressed e.g. in some temporal logic, such as LTL or CTL) to be checked; they then automatically and exhaustively verify whether the set of behaviors of the model satisfies the property.

During the last 15 years, model checking has been extended to handle complex systems, whose behaviors are the result of the interactions of several components. Games played on graphs are a convenient model for representing such interactions, and temporal logics have been hence proposed in order to express relevant properties in such a setting. One of the specification language to navigate in the execution trees of multi-agents systems is the temporal logic ATL [AHK02] (Alternating-Time Temporal Logic); it is an extension of the branching temporal logic CTL which allows to express properties such as the fact that a component can enforce a certain behavior independently of the actions performed by the other components. ATL has then be enriched in different ways to obtain more expressive logics for multi-agent systems. In particular, ATL_{sc} (ATL with strategy contexts) [BDLM09, LM15] and Strategy Logic [CHP07, MMV10] are two powerful extensions with similar properties in terms of expressive power and algorithmic properties. It was furthermore proved that those two logics have decidable, but Tower-complete model-checking algorithms.

In the approaches cited above, it is always assumed that all the players in the games have perfect observation of the state of the game, and that they also have perfect recall of the sequence of states that have been visited. In other words, they can choose an action to perform based on the entire sequence of states visited before. However, in many applications, components only have bounded memory, and most often they do not have the ability to fully observe all the other components of the system. While considering imperfect recall—the hypothesis that each player can only store into a finite memory the history of the seen states seen—can greatly simplify verification algorithms (since the number of strategies in the systems becomes finite), partial observation is known to make ATL model checking undecidable [AHK02, DT11]. Such results obviously carry over to more expressive logics like ATL_{sc}. Decidability can be regained by

*This work was partly supported by ERC Starting grant EQualIS (FP7-308087) and by FET project Cassting (FP7-601148).

restricting to imperfect-recall strategies [Sch04], or by considering hierarchical information [BMV15] or special communication architectures in distributed synthesis [KV01, Sch08].

In this paper, we consider a restricted case of partial observation, where all the players have the same information about the state space. We call such a case *uniform* partial observation. We prove that under this hypothesis, model checking concurrent game structure is decidable, even for the powerful logic ATL_{sc}. In particular, it is decidable whether there exists a strategy, based only on a subset of atomic propositions (assuming that the precise states and the other propositions are not visible), to enforce a given property. Note also that the restriction to uniform observation is not significant when one looks for a strategy of a single agent a against all other players, since the semantics we use requires that the strategy be winning for any outcome (hence the exact observation of the opponent players is irrelevant). The decidability proof for model-checking under uniform partial observation is obtained by adapting a previous approach developed in [DLM12, LM15], which consists in transforming the model-checking problem for ATL_{sc}^{*} into a model-checking problem for QCTL, an extension of CTL with propositional quantification. A similar technique also allows us to prove that when restricting the strategy quantifiers to range over memoryless strategies, then the model-checking problem for ATL_{sc}^{*} with partial observation is again decidable. We finally prove that satisfiability checking for ATL_{sc} with partial observation (i.e., deciding whether there exists a game structure with partial observation satisfying a given formula of ATL_{sc}^{*}) is undecidable, even in the case of turn-based games (where satisfiability is decidable under full observation [LM13]).

2 Definitions

2.1 Game structures with partial observation

In this paper, we consider concurrent games with partial observation. They correspond to classical concurrent game structures [AHK02] where, for each agent, an equivalence relation over the states of the structure defines sets of states that are observationally equivalent for this player. Observation equivalence extends to sequences of states in the obvious way. The strategies of the agents then have to be compatible with their observation, in the sense that after two observationally equivalent plays, a strategy has to return the same action. In this preliminary section, we formalize this setting.

All along this paper, we consider a set AP of atomic propositions. We recall that a *Kripke structure* over AP is a tuple $\langle Q, R, \ell \rangle$ where Q is a countable set of states, $R \subseteq Q \times Q$ is the transition relation and $\ell: Q \rightarrow 2^{\text{AP}}$ is a state-labeling function.

Definition 1. A concurrent game structure with partial observation (CGSO) is a tuple $\mathcal{C} = \langle Q, R, \ell, \text{Agt}, \mathcal{M}, \text{Mov}, \text{Edge}, (\sim_a)_{a \in \text{Agt}} \rangle$ where: $\langle Q, R, \ell \rangle$ is a finite-state Kripke structure; $\text{Agt} = \{a_1, \dots, a_p\}$ is a finite set of agents (or players); $\mathcal{M} = \{m_1, \dots, m_r\}$ is a finite set of moves (or actions); $\text{Mov}: Q \times \text{Agt} \rightarrow 2^{\mathcal{M}} \setminus \{\emptyset\}$ defines the set of available moves of each agent in each state; $\text{Edge}: Q \times \mathcal{M}^{\text{Agt}} \rightarrow Q$ is a transition table associating, with each state q and each set of moves of the agents, the resulting state q' , with the requirement that $(q, q') \in R$; finally, $(\sim_a)_{a \in \text{Agt}}$ assigns to each player an equivalence relation over Q .

In the following, we assume w.l.o.g. (w.r.t existence of specific strategies) that $\text{Mov}(q, a) = \mathcal{M}$ for every $q \in Q$ and $a \in \text{Agt}$: all the actions are available to all the players at any time. A move vector is a vector $\mathbf{m} \in \mathcal{M}^{\text{Agt}}$; for such a vector and for an agent $a \in \text{Agt}$, we denote by $\mathbf{m}[a]$ the move of agent a in \mathbf{m} . A *turn-based game structure with partial observation (TBGSO)* is a CGSO for which there exists a mapping $\text{Own}: Q \rightarrow \text{Agt}$ such that for any $q \in Q$ and for any two move vectors \mathbf{m} and \mathbf{m}' if $\mathbf{m}[\text{Own}(q)] = \mathbf{m}'[\text{Own}(q)]$, then $\text{Edge}(q, \mathbf{m}) = \text{Edge}(q, \mathbf{m}')$.

As we mentioned above, each relation \sim_a with $a \in \text{Agt}$ characterizes the observation of the agent a in the game structure: if $q \sim_a q'$, then agent a is not able to make a distinction between the states q and q' ; in such a case, we say that q and q' are a -equivalent. As an important special case, an observation relation $(\sim_a)_{a \in \text{Agt}}$ is said to be *uniform* when $\sim_a = \sim_{a'}$ for all $a, a' \in \text{Agt}$. In that case we might substitute the set $(\sim_a)_{a \in \text{Agt}}$ by a unique equivalence relation \sim .

A *finite path* in the CGSO is a finite non-empty sequence of states $\rho = q_0 q_1 q_2 \dots q_k$ such that $(q_i, q_{i+1}) \in R$ for all $i \in \{0, \dots, k-1\}$. An *infinite path* (or run) is an infinite sequence of states such that each finite prefix is a finite path. We denote by Path (resp. InfPath) the set of finite (resp. infinite) paths. Let $|\rho|$ the length of the path ρ (with $|\rho| = \infty$ if ρ is infinite). For $0 \leq i < |\rho|$, we write $\rho(i)$ to represent the $i+1$ -th element of the path ρ . For a path ρ , we write $\text{first}(\rho)$ for its first element $\rho(0)$ and, when $|\rho| < \infty$, we write $\text{last}(\rho)$ for its last element $\rho(|\rho|)$. For $0 \leq i < |\rho|$, we denote by $\rho_{\leq i}$ the prefix of ρ until position i , i.e. the finite path $\rho(0)\rho(1)\dots\rho(i)$. We extend the equivalence relation \sim_a for $a \in \text{Agt}$ to paths as follows : two paths ρ and ρ' are a -equivalent (written $\rho \sim_a \rho'$) if, and only if, $|\rho| = |\rho'|$ and $\rho(i) \sim_a \rho'(i)$ for every $0 \leq i < |\rho|$.

Given a CGSO \mathcal{C} and one of its states q_0 , we write $\mathcal{T}_{\mathcal{C}}(q_0)$ for the *execution tree* of \mathcal{C} from q_0 : formally, $\mathcal{T}_{\mathcal{C}}(q_0)$ is the pair $\langle T, \ell \rangle$ where T is the set of all finite paths (called *nodes* in the context of trees) in \mathcal{C} with first state q_0 , and ℓ labels each node ρ with the labeling of $\text{last}(\rho)$ in \mathcal{C} . It will be convenient in the sequel to see execution trees as infinite-state Kripke structures. To alleviate notations, we still write $\mathcal{T}_{\mathcal{C}}(q_0)$ for the Kripke structure $\langle T, R, \ell \rangle$ where $\langle T, \ell \rangle$ is the tree defined above, and $R \subseteq T \times T$ is the transition relation such that $(\rho, \rho') \in R$ whenever ρ is the prefix of ρ' of length $|\rho'| - 1$.

A *strategy* for agent $a \in \text{Agt}$ is a function $f_a: \text{Path} \rightarrow \mathcal{M}$; it associates with any finite path a move to be played by agent a after this path. A strategy f_a for agent $a \in \text{Agt}$ is said to be *memoryless* whenever for any two finite paths ρ and ρ' such that $\rho(|\rho| - 1) = \rho'(|\rho'| - 1)$, it holds $f_a(\rho) = f_a(\rho')$. Hence the decision of a memoryless strategy depends only on the current control state; for this reason, we may simply give such a strategy as a fonction $f_a: Q \rightarrow \mathcal{M}$. A strategy f_A for a coalition of agents $A \subseteq \text{Agt}$ is a set of strategy $\{f_a\}_{a \in A}$ assigning a strategy f_a to each agent $a \in A$ (note that a strategy for agent a is equivalent to a strategy for coalition $\{a\}$). Given a strategy $f_A = \{f_a\}_{a \in A}$ for coalition A , we say that a path ρ respects f_A from a finite path π if, and only if, for all $0 \leq i < |\pi|$, we have $\rho(i) = \pi(i)$ and for all $|\pi| \leq i < |\rho| - 1$, we have that $\rho(i+1) = \text{Edge}(\rho(i), \mathbf{m})$ where \mathbf{m} is a move vector satisfying $\mathbf{m}(a) = f_a(\rho_{\leq i})$ for all $a \in A$. Given a finite path π , we denote by $\text{Out}(\pi, f_A)$ the set of infinite paths ρ' such that ρ respects the strategy f_A from π . Given a strategy $g_A = \{g_a\}_{a \in A}$ for a coalition A and a strategy $f_B = \{f_b\}_{b \in B}$ for a coalition B , we denote by $g_A \circ f_B$, the strategy $\{h_c\}_{c \in A \cup B}$ for coalition $A \cup B$ such that $h_c = g_c$ for all $c \in A$ and $h_c = f_c$ for all $c \in B \setminus A$. Finally given a $f_B = \{f_b\}_{b \in B}$ for a coalition B and a set of agents $A \subseteq \text{Agt}$, we denote by $(f_B)_{|A}$ (resp. $(f_B)_{\setminus A}$) the strategy $\{f_b\}_{b \in B \cap A}$ (resp. $\{f_b\}_{b \in B \setminus A}$) for coalition $B \cap A$ (resp. $B \setminus A$).

Partial observation comes into the play by restricting the space of allowed strategies: in our setting, we only consider strategies that are *compatible* with the observation in the game, which means that after any two a -equivalent finite paths ρ and ρ' , the strategies for agent a have to take the same decisions (i.e. $f_a(\rho) = f_a(\rho')$). We could equivalently define a compatible strategy for a as a function from the quotient set Path/\sim_a to \mathcal{M} , such that if $[\rho]$ is the equivalence class of ρ with respect to \sim_a , then $f_a([\rho])$ gives the move to play for a from any history equivalent to ρ . A strategy $f_A = \{f_a\}_{a \in A}$ for coalition A is *compatible* if f_a is compatible for all $a \in A$. We write $\text{Strat}(A)$ to denote the unrestricted strategies for coalition A , $\text{Strat}^{\sim}(A)$ for the set of compatible strategies, and $\text{Strat}_0^{\sim}(A)$ is the set of compatible memoryless strategies for A .

2.2 ATL with strategy contexts

We will be interested in the logic ATL_{sc}, which extends the alternating-time temporal logic of [AHK02] with strategy contexts. We assume a fixed set of atomic propositions AP and a fixed set of agents Agt.

Definition 2. *The formulas of ATL_{sc}^{*} are defined by the following grammar:*

$$\begin{aligned} \text{ATL}_{sc}^* \ni \varphi_s, \psi_s ::= & P \mid \neg \varphi_s \mid \varphi_s \vee \psi_s \mid \langle A \cdot \rangle \varphi_p \mid \langle \bar{A} \cdot \rangle \varphi_p \mid \langle A \rangle \varphi_p \mid \langle \bar{A} \rangle \varphi_p \\ \varphi_p, \psi_p ::= & \varphi_s \mid \neg \varphi_p \mid \varphi_p \vee \psi_p \mid \mathbf{X} \varphi_p \mid \varphi_p \mathbf{U} \psi_p \end{aligned}$$

where P ranges over AP and A ranges over 2^{Agt} .

We interpret ATL_{sc}^{*} formulas over CGSOs, within a context (*i.e.*, a preselected strategy for a coalition): *state formulas* of the form φ_s in the grammar above are evaluated over states, while *path formulas* of the form φ_p are evaluated along infinite paths. In order to have a uniform definition, we evaluate all formulas at a given position along a path.

In ATL_{sc}^{*}, contrary to the case of classical ATL [AHK02], when strategy quantifiers assign new strategies to some players, the other players keep on playing their previously-assigned strategies. This is what “strategy contexts” refer to. Informally, formula $\langle A \cdot \rangle \varphi_p$ holds at position n along ρ under the context F if it is possible to extend F with a *strategy* for the coalition A such that the outcomes of the resulting strategy after $\rho_{\leq n}$ all satisfy φ_p . In Section 4, we also consider the strategy quantifiers $\langle A \cdot \rangle_0$, which quantifies only on *memoryless* strategies. Finally strategies can be dropped from the context using the operator $\langle - \rangle$. Notice that in this paper, all strategy quantifiers are restricted to compatible strategies.

We now define the semantics formally. Let \mathcal{C} be a CGSO with agent set Agt, ρ be an infinite path of \mathcal{C} , and $n \in \mathbb{N}$. Let $B \subseteq \text{Agt}$ be a coalition, and $f_B \in \text{Strat}^{\sim}(B)$. That a (state or path) formula φ holds at a position n along ρ in \mathcal{C} under strategy context f_B , denoted $\mathcal{C}, \rho, n \models_{f_B} \varphi$, is defined inductively as follows (omitting atomic propositions and Boolean operators):

$$\begin{aligned} \mathcal{C}, \rho, n \models_{f_B} \langle A \rangle \varphi_s & \quad \text{iff} \quad \mathcal{C}, \rho, n \models_{(f_B) \setminus A} \varphi_s \\ \mathcal{C}, \rho, n \models_{f_B} \langle \bar{A} \rangle \varphi_s & \quad \text{iff} \quad \mathcal{C}, \rho, n \models_{(f_B) \setminus A} \varphi_s \\ \mathcal{C}, \rho, n \models_{f_B} \langle A \cdot \rangle \varphi_p & \quad \text{iff} \quad \exists g_A \in \text{Strat}^{\sim}(A). \forall \rho' \in \text{Out}(\rho_{\leq n}, g_A \circ f_B). \mathcal{C}, \rho', n \models_{g_A \circ f_B} \varphi_p \\ \mathcal{C}, \rho, n \models_{f_B} \langle \bar{A} \cdot \rangle \varphi_p & \quad \text{iff} \quad \mathcal{C}, \rho, n \models_{f_B} \langle \text{Agt} \setminus A \cdot \rangle \varphi_p \\ \mathcal{C}, \rho, n \models_{f_B} \mathbf{X} \varphi_p & \quad \text{iff} \quad \mathcal{C}, \rho, n+1 \models_{f_B} \varphi_p \\ \mathcal{C}, \rho, n \models_{f_B} \varphi_p \mathbf{U} \psi_p & \quad \text{iff} \quad \exists l \geq 0. \mathcal{C}, \rho, n+l \models_{f_B} \psi_p \text{ and } \forall 0 \leq m < l. \mathcal{C}, \rho, n+m \models_{f_B} \varphi_p \end{aligned}$$

Finally, we write $\mathcal{C}, q_0 \models \varphi_s$ when $\mathcal{C}, \rho, 0 \models_{f_\emptyset} \varphi_s$ (with empty context) for all path ρ such that $\rho(0) = q_0$. Notice that this does not depend on the choice of ρ . The usual shorthands such as \mathbf{F} and \mathbf{G} are defined as for CTL*. It will also be convenient to use the constructs $\llbracket \cdot \rrbracket \varphi_p$ as a shorthand for $\neg \langle A \cdot \rangle \neg \varphi_p$, and $\langle \cdot \rangle \varphi_s$ as a shorthand for $\langle A \cdot \rangle \perp \mathbf{U} \varphi_s$. The CTL* universal path quantifiers \mathbf{E} and \mathbf{A} can be expressed in ATL_{sc}^{*}. For instance, $\mathbf{A} \varphi_p$ is equivalent to $\langle \emptyset \rangle \langle \emptyset \cdot \rangle \varphi_p$. Finally the fragment ATL_{sc} of ATL_{sc}^{*} is defined as usual, by restricting the set of path formulas to

$$\varphi_p, \psi_p ::= \neg \varphi_p \mid \mathbf{X} \varphi_s \mid \varphi_s \mathbf{U} \psi_s.$$

Remark 3. *The strategy quantifiers for complement coalitions (namely $\langle \bar{A} \cdot \rangle$ and $\langle \bar{A} \rangle$) are only useful in case the set Agt is not known in advance, as it is the case when dealing with satisfiability or with expressiveness questions.*

Remark 4. As opposed to Strategy Logic [CHP07, MMV10], the (existential) strategy quantifiers in ATL_{sc} include an implicit (universal) quantification over the outcomes of the strategies of the unassigned players. This is indeed a quantification over all the outcomes, and not over the outcomes that would result only from compatible strategies.

Example 5. Consider the CGSO \mathcal{C} in Figure 1. It is a turn-based CGS with two players (Player a_1 plays in circle node, and Player a_2 plays in box node). The set \mathcal{M} is $\{1, 2, 3\}$. The partial observation for Player a_1 is then given by \sim_{a_1} is \equiv_P (i.e. two states are equivalent if, and only if, the truth value of P is the same in both states). Now consider the formula $\varphi = \langle a_1 \cdot \rangle_0 \mathbf{F} f$. To see that φ holds for true in q_0 with the standard semantics (where the strategy quantifiers are not restricted to compatible strategies), it is sufficient to consider the (memoryless) strategy for a_1 consisting in choosing the move m_1 from q_2 and the move m_2 (or m_3) from q_3 . Note that this strategy is not compatible (after q_0q_2 and q_0q_3 , agent a_1 should choose the same move), thus we need to consider another one to show that φ is satisfied with the partial-observation semantics; for example, the strategy consisting in choosing m_1 after q_0q_2 and q_0q_3 , m_2 after $q_0q_1q_3$ and m_1 after $q_0q_3q_2$. If furthermore we look for a memoryless compatible strategy for a_1 , we can use the strategy assigning the move m_1 to q_2 and q_3 , hence the formula $\langle a_1 \cdot \rangle_0 \mathbf{F} f$ also holds true in q_0 . Therefore, formula φ holds true at q_0 . On the other hand, formula $\langle a_1 \cdot \rangle_0 \mathbf{XXX} f$ does not hold true in q_0 , but formula $\langle a_1 \cdot \rangle \mathbf{XXX} f$, where we drop the memory constraint, does hold true.




Fig. 1: The turn-based CGS \mathcal{C}

In the sequel we will consider the *model-checking problem* of ATL_{sc}^* over CGSOs which takes as input a CGSO \mathcal{C} , a control state q_0 and an ATL_{sc}^* formula φ and which answers yes if $\mathcal{C}, q_0 \models \varphi$ and no otherwise. We will also consider the *satisfiability problem* which given a formula of ATL_{sc}^* determines whether there exists a CGSO \mathcal{C} and a control state q_0 such that $\mathcal{C}, q_0 \models \varphi$. We recall that when considering concurrent game structures (with full observation), ATL_{sc}^* model-checking is decidable, but its satisfiability problem is undecidable (except when restricting to turn-based games) [LM15].

2.3 From concurrent games to turn-based games

Following our definitions, turn-based game structures can be seen as special cases of concurrent game structures, where in each location only one player may have several non-equivalent moves. In this section, we prove that any partial-observation CGSO can be turned into an *equivalent* partial-observation TBGSO (where *equivalent* will be made precise later). While all players play at the same time in a CGSO, they play one after the other (in any predefined order) in the correspondig TBGSO, but the intermediary states are made undistinguishable to all players, so that no player can gain information from playing after another one. Figure 2 schematically represents this transformation in the case of two players. Obviously, since we add intermediary states, we also have to modify the ATL_{sc}^* formula to be checked, by making the intermediary states “invisible”; this is a classical construction in temporal logics. In the end, we prove the following result:




Fig. 2: Turning a CGSO into an equivalent TBGSO

Theorem 6. For any CGSO $\mathcal{C} = \langle Q, R, \ell, \text{Agt}, \mathcal{M}, \text{Mov}, \text{Edge}, (\sim_a)_{a \in \text{Agt}} \rangle$, we can compute a TBGSO $\mathcal{T} = \langle Q', R', \ell', \text{Agt}, \mathcal{M}, \text{Mov}', \text{Edge}', (\sim_a)_{a \in \text{Agt}} \rangle$ with $Q \subseteq Q'$, for which there is a logspace translation $\varphi \in \text{ATL}_{sc}^* \mapsto \tilde{\varphi} \in \text{ATL}_{sc}^*$ such that for any $\varphi \in \text{ATL}_{sc}^*$ and any state q of \mathcal{C} , it holds:

$$\mathcal{C}, q \models_{\emptyset} \varphi \quad \Leftrightarrow \quad \mathcal{T}, q \models_{\emptyset} \tilde{\varphi}.$$

Furthermore, if \mathcal{C} is uniform, then so is \mathcal{T} .

2.4 QCTL* in a nutshell

As we explain in the sequel, under some restrictions (uniformity or restriction of the considered strategies), the model-checking problem of ATL_{sc}* over CGSOs can be reduced to the model-checking problem of the temporal logic Quantified CTL* over Kripke Structures. QCTL* extends the classical branching time temporal logic CTL* with *atomic propositions quantifiers* $\exists P. \varphi_s$ (and its dual $\forall P. \varphi_s$), where P is an atomic proposition in AP:

$$\begin{aligned} \text{QCTL}^* \ni \varphi_s, \psi_s ::= & P \mid \neg \varphi_s \mid \varphi_s \vee \psi_s \mid \mathbf{E} \varphi_p \mid \exists P. \varphi_s \\ \varphi_p, \psi_p ::= & \varphi_s \mid \neg \varphi_p \mid \varphi_p \vee \psi_p \mid \mathbf{X} \varphi_p \mid \varphi_p \mathbf{U} \psi_p \end{aligned}$$

where P ranges over AP. We briefly review QCTL* here, and refer to [LM14] for more details and examples.

QCTL* formulas are then interpreted over a (classical) Kripke structure $\mathcal{S} = \langle Q, R, \ell \rangle$. Informally $\exists P. \varphi_s$ is used to specify the existence of a valuation for P over \mathcal{S} such that φ_s is satisfied. We consider two different semantics of $\exists P. \varphi$: the *structure semantics* where one looks for a valuation of the structure \mathcal{S} , and the *tree semantics* where one considers the valuation of its execution tree.

For $X \subseteq \text{AP}$, two Kripke structures $\mathcal{S} = \langle Q, R, \ell \rangle$ and $\mathcal{S}' = \langle Q', R', \ell' \rangle$ are *X-equivalent* (denoted by $\mathcal{S} \equiv_X \mathcal{S}'$) whenever $Q = Q'$, $R = R'$, and $\ell \equiv_X \ell'$ (i.e., $\ell(q) \cap X = \ell'(q) \cap X$ for any $q \in Q$). The *structure semantics* of QCTL* (whose satisfaction relation is denoted by \models_{st}) is derived from the semantics of CTL* by adding the following rule:

$$\mathcal{S}, \rho, n \models_{\text{st}} \exists P. \varphi_s \quad \text{iff} \quad \exists \mathcal{S}'. \mathcal{S}' \equiv_{\text{AP} \setminus \{P\}} \mathcal{S} \text{ and } \mathcal{S}', \rho, n \models_{\text{st}} \varphi_s.$$

In other terms, $\exists P. \varphi_s$ means that it is possible to (re)label the Kripke structure with P in order to make φ_s hold. As for CGSO, we write $\mathcal{S}, q_0 \models_{\text{st}} \varphi_s$ whenever $\mathcal{S}, \rho, 0 \models_{\text{st}} \varphi_s$ for all path ρ such that $\rho(0) = q_0$. The *tree semantics* (whose satisfaction relation is denoted by \models_{tr}) is defined as: $\mathcal{S}, q_0 \models_{\text{tr}} \varphi_s$ if, and only if, $\mathcal{T}_{\mathcal{S}}(q_0), q_0 \models_{\text{st}} \varphi_s$ (where $\mathcal{T}_{\mathcal{S}}(q_0)$ is the execution tree of \mathcal{S} from q_0 , seen as an infinite-state Kripke structure).

Example 7. Consider formula $\mathbf{EX} \varphi \wedge \forall P. (\mathbf{EX}(\varphi \wedge P) \Rightarrow \mathbf{AX}(P \Rightarrow \varphi))$, which we write $EX_1 \varphi$ in the sequel. This formula states the existence of a unique immediate successor satisfying φ . We will reuse it later in our construction.

Theorem 8 ([LM14]). 1. For the structure semantics, the model-checking problem of QCTL* is PSPACE-complete and the satisfiability problem is undecidable.

2. For the tree semantics, the model-checking and satisfiability problems of QCTL* are decidable, and Tower-complete.

3 Model checking uniform CGSOs

As we have already mentioned, it is well known that the model-checking of ATL over CGSO is undecidable [AHK02, DT11]. Since ATL is a fragment of ATL_{sc}^* , this undecidability result holds also for ATL_{sc}^* . We prove in this section that decidability can be regained by restricting to uniform partial observation, i.e., by assuming that the observation is the same for all the agents. The idea consists in reducing the model-checking problem of ATL_{sc}^* over uniform CGSOs to the model-checking problem for QCTL* with tree semantics over a complete Kripke structure representing the possible observations.

In the sequel, we consider a uniform CGSO $\mathcal{C} = \langle Q, R, \ell, \text{Agt}, \mathcal{M}, \text{Mov}, \text{Edge}, (\sim_a)_{a \in \text{Agt}} \rangle$ with the finite set of states $Q = \{q_0, \dots, q_s\}$, the finite set of moves $\mathcal{M} = \{m_1, \dots, m_r\}$ and $\sim_a = \sim$ for all $a \in \text{Agt}$. For each $q \in Q$, we denote by $[q]$ the equivalence class of q with respect to \sim , i.e. $[q] \subseteq 2^Q$ is such that $q' \in [q]$ if, and only if, $q \sim q'$. We also consider an ATL_{sc}^* formula φ_s whose $\langle \cdot \cdot \cdot \rangle$ -depth (i.e. the maximal number of nested modalities $\langle \cdot \cdot \cdot \rangle$) is λ (we assume here that modalities $\langle \cdot \cdot \cdot \rangle$ and $\langle \bar{\cdot} \rangle$ have been previously removed from φ_s , as the set of agents is known). The idea behind our proof is to consider a Kripke structure $\mathcal{S}_{\mathcal{C}}$ as a complete graph whose set of nodes is the set of equivalence classes $[q]$. We then use the power of QCTL* (with the tree semantics) to build paths and strategies in this Kripke Structure. The fact that, in the Kripke structure, the paths go through equivalent classes of states (and not through states), guarantees that the corresponding strategies are compatible with the partial observation.

Before we build the Kripke Structure, we need to introduce some sets of fresh atomic propositions (not appearing in \mathcal{C} or in φ_s) which we will use to encode paths and strategies in the Kripke structure. Let $AP_Q = \{q_i^\ell \mid 1 \leq i \leq s \text{ and } 0 \leq \ell \leq \lambda\}$, $AP_{\mathcal{M}}^a = \{m_1^a, \dots, m_r^a\}$ for every $a \in \text{Agt}$, and $AP_{\mathcal{M}} = \bigcup_{a \in \text{Agt}} AP_{\mathcal{M}}^a$ and $AP_S = \{s_i \mid 1 \leq i \leq s\}$ be those sets of fresh atomic propositions. We assume that AP_Q , $AP_{\mathcal{M}}$ and AP_S are subsets of AP. Intuitively, propositions in AP_Q will be used to fix a path of \mathcal{C} at each level κ of quantification (a path at level κ being a sequence of states q_i^κ); propositions in $AP_{\mathcal{M}}^a$ will be used to label the choices corresponding to a strategy of agent a .

The Kripke structure $\mathcal{S}_{\mathcal{C}}$ associated with \mathcal{C} is then defined as $\langle Q_{\mathcal{C}}, R_{\mathcal{C}}, \ell_{\mathcal{C}} \rangle$ where $Q_{\mathcal{C}} = \{[q] \mid q \in Q\}$, $R_{\mathcal{C}} = Q_{\mathcal{C}} \times Q_{\mathcal{C}}$ and $\ell_{\mathcal{C}}([q]) = \{s_i \mid q_i \in [q]\}$. Compatible strategies can then be encoded as functions labeling the execution trees of this Kripke structure (from the state where we evaluate the formula): a strategy f_a for some agent a is represented as a function $\tilde{f}_a: Q_{\mathcal{C}}^+ \rightarrow AP_{\mathcal{M}}^a$ that labels the execution trees $\mathcal{T}_{\mathcal{S}_{\mathcal{C}}}(q)$ of $\mathcal{S}_{\mathcal{C}}$ from the current state q with proposition in $AP_{\mathcal{M}}^a$. However note that $\mathcal{S}_{\mathcal{C}}$ being a complete graph, paths in this structure might not correspond to paths in the associated CGSO \mathcal{C} ; we will use atomic propositions q_i to identify concrete paths of \mathcal{C} .

Given a coalition B in Agt, an integer $0 \leq \kappa \leq \lambda$, and a subformula of φ_s at $\langle \cdot \cdot \cdot \rangle$ -depth κ , we define a QCTL* formula $\tilde{\varphi}^{B, \kappa}$ inductively as follows:

$$\tilde{P}^{B, \kappa} = \bigvee_{\{i \mid P \in \ell(q_i)\}} q_i^\kappa \qquad \widetilde{\langle A \rangle} \varphi_s^{B, \kappa} = \tilde{\varphi}_s^{B \setminus A, \kappa}$$

$$\begin{aligned} \widetilde{\varphi \wedge \psi}^{B,\kappa} &= \widetilde{\varphi}^{B,\kappa} \wedge \widetilde{\psi}^{B,\kappa} & \widetilde{\neg \varphi}^{B,\kappa} &= \neg \widetilde{\varphi}^{B,\kappa} \\ \widetilde{\varphi_p \mathbf{U} \psi_p}^{B,\kappa} &= \widetilde{\varphi_p}^{B,\kappa} \mathbf{U} \widetilde{\psi_p}^{B,\kappa} & \widetilde{\mathbf{X} \varphi_p}^{B,\kappa} &= \mathbf{X} \widetilde{\varphi_p}^{B,\kappa} \end{aligned}$$

For a formula of the shape¹ $\langle a \cdot \rangle \varphi_p$, the construction is as follows:

$$\begin{aligned} \widetilde{\langle a \cdot \rangle \varphi_p}^{B,\kappa} &= \exists m_1^a \dots m_\kappa^a. \left[\Phi_{\text{strat}}(\{a\}) \wedge \forall q_1^{\kappa+1} \dots q_s^{\kappa+1}. \left(\Phi_{\text{path}}(\kappa+1) \wedge \right. \right. \\ &\quad \left. \left. \bigvee_{1 \leq i \leq s} (q_i^\kappa \wedge q_i^{\kappa+1}) \wedge \Phi_{\text{out}}(\kappa+1, B \cup \{a\}) \right) \Rightarrow \mathbf{A} \left(\mathbf{G} \bigvee_{1 \leq i \leq s} q_i^{\kappa+1} \Rightarrow \widetilde{\varphi_p}^{B \cup \{a\}, \kappa+1} \right) \right] \end{aligned}$$

with

$$\begin{aligned} \Phi_{\text{strat}}(A) &= \bigwedge_{a \in A} \mathbf{A} \mathbf{G} \left(\bigvee_{1 \leq j \leq r} (m_j^a \wedge \bigwedge_{j' \neq j} \neg m_{j'}^a) \right) \\ \Phi_{\text{path}}(\kappa+1) &= \mathbf{E} \mathbf{G} \left(\bigvee_{1 \leq i \leq s} (q_i^{\kappa+1} \wedge s_i \wedge \bigwedge_{i' \neq i} \neg q_{i'}^{\kappa+1}) \wedge \mathbf{E} \mathbf{X}_1 \left(\bigvee_{1 \leq i \leq s} q_i^{\kappa+1} \right) \right) \\ \Phi_{\text{out}}(\kappa+1, B) &= \mathbf{E} \mathbf{G} \left(\bigvee_{1 \leq i \leq s} \bigvee_{\left\{ (m, q_j) \mid \begin{array}{l} m \in \mathcal{M}^{\text{Agt}} \wedge \\ q_j = \text{Edge}(q_i, m) \end{array} \right\}} q_i^{\kappa+1} \wedge \widetilde{\mathbf{m}}^B \wedge \mathbf{E} \mathbf{X} q_j^{\kappa+1} \right) \end{aligned}$$

where $\widetilde{\mathbf{m}}^B$ is a shorthand for the formula $\bigwedge_{\{(b, m_k) \mid b \in B \wedge m[b] = m_k\}} m_k^b$. We point out the fact that $\Phi_{\text{out}}(\kappa+1, B)$ is based on the transition table Edge of \mathcal{C} ; consequently, its size is in $O(|Q|^2 \cdot |\mathcal{M}|^{|\text{Agt}|})$ (i.e., in $O(|Q| \cdot |\text{Edge}|)$).

The intuition behind formula $\widetilde{\langle a \cdot \rangle \varphi_p}^{B,\kappa}$ is as follows: the formula first “selects” a strategy for agent a under the form of a labeling of the execution tree of $\mathcal{S}_\mathcal{C}$ with m_i^a ; it then uses subformula $\Phi_{\text{strat}}(\{a\})$ to ensure that this labeling correctly encodes a strategy for a ; finally, it checks that all the outcomes of the selected strategy satisfy φ_p . The latter task is achieved by considering all the labelings of the structure with $q_i^{\kappa+1}$, and, for the labelings that correspond to one outcome of the selected strategy, by checking the formula φ_p . Ensuring that a labeling corresponds to an outcome is done as follows:

1. it corresponds to an infinite branch in the execution tree of $\mathcal{S}_\mathcal{C}$, and each node labeled by $q_i^{\kappa+1}$ should correspond to a node $[q_i]$ (labeled by s_i) in $\mathcal{S}_\mathcal{C}$ (both points are ensured by formula² $\Phi_{\text{path}}(\kappa+1)$);
2. at the present position, one of the propositions $q_i^{\kappa+1}$ has to match with one of the state q_i^κ of the previous level (this ensures that the path labeled by $q_i^{\kappa+1}$ starts from the “current state” considered in the game);
3. the branch obtained by this labeling effectively follows the choices dictated by the labels m^b encoding the strategies for $b \in B \cup \{a\}$; this is checked by the formula $\Phi_{\text{out}}(\kappa+1, B \cup \{a\})$.

Finally, the formula checks that the corresponding path satisfies the formula $\widetilde{\varphi_p}^{B \cup \{a\}, \kappa+1}$.

The correctness of the reduction is stated in the following theorem:

¹For the sake of readability, we restrict to one-player coalitions here; the construction easily extends to the the general case with a coalition (including the empty coalition).

²See Example 7 for the definition of $\mathbf{E} \mathbf{X}_1$

Theorem 9. Let φ_s be an ATL_{-sc}^* state-formula, \mathcal{C} be a uniform CGSO, and q_α be a state of \mathcal{C} . Then:

$$\mathcal{C}, q_\alpha \models \varphi_s \quad \text{if, and only if,} \quad \mathcal{S}_{\mathcal{C}}, [q_\alpha] \models_{\text{tr}} \exists q_\alpha^0. (\tilde{\varphi}_s^{\emptyset,0} \wedge q_\alpha^0)$$

Proof. First we point out the fact that none of the formula used in the reduction checks that the considered strategies are compatible, but in fact this is guaranteed because we evaluate the formula over the Kripke structure $\mathcal{S}_{\mathcal{C}}$, and two equivalent paths in \mathcal{C} with respect to \sim are necessarily matched to a unique path in the execution tree of $\mathcal{S}_{\mathcal{C}}$.

We now prove that our reduction is correct. Let $\mathcal{T}_{\mathcal{S}_{\mathcal{C}}}([q_\alpha]) = \langle T_{\mathcal{C}}, R'_{\mathcal{C}}, \ell'_{\mathcal{C}} \rangle$ be the execution tree of the Kripke structure $\mathcal{S}_{\mathcal{C}}$ from the state $[q_\alpha]$. Note that nodes in $\mathcal{T}_{\mathcal{S}_{\mathcal{C}}}([q_\alpha])$ are thus finite paths of the form $[q_0][q_1] \dots [q_k]$ (with $q_0 = q_\alpha$). For a finite path π in \mathcal{C} such that $\pi(0) = q_\alpha$, we denote by $[\pi]$ the path in $\mathcal{S}_{\mathcal{C}}$ such that $|\pi| = |[\pi]|$ and for all $0 \leq i < |\pi|$, we have $[\pi](i) = [\pi(i)]$. We also denote by $\text{path}(\pi)$ the path in $\mathcal{T}_{\mathcal{S}_{\mathcal{C}}}([q_\alpha])$ such that $|\text{path}(\pi)| = |\pi|$ and for all $0 \leq i < |\pi|$, $\text{path}(\pi)(i) = [\pi_{\leq i}]$. We now consider the tree $\mathcal{T}' = \langle T_{\mathcal{C}}, R'_{\mathcal{C}}, \ell' \rangle$ where ℓ' extends the function $\ell'_{\mathcal{S}}$ with propositions of type m_j^a and q_i^k . This extension is used to encode a strategy context and to describe a path of \mathcal{C} . Given a compatible strategy $f_B = \{f_b\}_{b \in B}$ for the coalition $B = \{b_1, \dots, b_h\}$, an infinite path ρ in \mathcal{C} such $\rho(0) = q_\alpha$ and a position $n \geq 0$, we say that ℓ' is:

- an f_B -labeling whenever, for every node $\gamma \in T_{\mathcal{C}}$ with $\gamma = [\pi]$ for some finite path π in \mathcal{C} , for any $b \in B$, for any $1 \leq j \leq r$, we have $m_j^b \in \ell'(\gamma)$ if, and only if, $f_b(\pi) = m_j$;
- a (κ, ρ) -labeling if the following two conditions are verified:
 1. for all $j \geq 0$, it holds $q_i^k \in \ell'([\rho_{\leq j}])$ if, and only if, $\rho(j) = q_i$
 2. if $q_i \in \ell'(\pi)$ for a node $\pi \in T_{\mathcal{C}}$, then there exists $j \geq 0$ such that $\pi = [\rho_{\leq j}]$.

In other words, the propositions q_1^k, \dots, q_s^k label a unique branch in the tree, that can be matched with the path ρ .

We say that for $0 \leq \kappa \leq \lambda$, a subformula φ of φ_s occurs at $\langle \cdot \cdot \rangle$ -depth κ , if the number of $\langle \cdot \cdot \rangle$ -quantifier above φ in tree representing formula φ_s is κ .

Proposition 10. Let φ be some subformula of φ_s which occurs at $\langle \cdot \cdot \rangle$ -depth κ with $0 \leq \kappa \leq \lambda$, ρ be a run of \mathcal{C} such that $\rho(0) = q_\alpha$, n be a natural number, f_B be a compatible strategy for coalition B . Let $\mathcal{T}' = \langle T_{\mathcal{C}}, R'_{\mathcal{C}}, \ell' \rangle$ be the tree defined as above. If ℓ' is an f_B -labeling and a (κ, ρ) -labeling, then we have:

$$\mathcal{C}, \rho, n \models_{f_B} \varphi \quad \text{if, and only if,} \quad \mathcal{T}', \text{path}(\rho), n \models_{\text{st}} \tilde{\varphi}^{B, \kappa}.$$

Proof. The proof is done by structural induction over φ . In the following, we let $q_{\alpha'} = \rho(n)$.

- case $\varphi = P$: we have $\rho(n) \models_{f_B} P$ if, and only if, $P \in \ell(q_{\alpha'})$. As ℓ' is a (κ, ρ) -labeling, we know that $\mathcal{T}', \text{path}(\rho), n \models q_\alpha^k$. By definition of $\tilde{P}^{B, \kappa}$, the implication follows. Conversely assume $\mathcal{T}', \text{path}(\rho), n \models \tilde{P}^{B, \kappa}$, we know that $\rho(n)$ has to be labeled by P , because ℓ' is a (κ, B) -labeling.
- case $\varphi = \varphi_p \mathbf{U} \psi_p$: if $\mathcal{C}, \rho, n \models_{f_B} \varphi_p \mathbf{U} \psi_p$, then there exists $i \geq n$ s.t. $\mathcal{C}, \rho, i \models_{f_B} \psi_p$ and for any $n \leq j < i$, we have $\mathcal{C}, \rho, j \models_{f_B} \varphi_p$. By i.h., we get $\mathcal{T}', \text{path}(\rho), i \models \tilde{\psi}_p^{B, \kappa}$ and, for any j , $\mathcal{T}', \text{path}(\rho), j \models \tilde{\varphi}_p^{B, \kappa}$; from this we obtains $\mathcal{T}', \text{path}(\rho), n \models \tilde{\varphi}^{B, \kappa}$. The converse is similar.
- case $\varphi = \langle \cdot \cdot \rangle \varphi_p$: Assume $\mathcal{C}, \rho, n \models_{f_B} \langle \cdot \cdot \rangle \varphi_p$. Then there exists a \sim -compatible strategy f_a s.t. for any $\rho' \in \text{Out}(\rho_{\leq n}, f_a \circ f_C)$, we have $\rho' \models_{f_a \circ f_B} \varphi_p$. From this strategy f_a , we deduce a valuation for propositions m_1^a, \dots, m_r^a extending ℓ' over T (because f_a is \sim -compatible), and satisfying $\Phi_{\text{strat}}(\{a\})$. Now extend ℓ' with a valuation for $q_1^{k+1}, \dots, q_s^{k+1}$ following the run ρ' (i.e. for every state $\rho'(i) = q_\beta$, the corresponding node $[\rho'](i)$ is labeled by q_β^{k+1} , and only $[\rho']$ -nodes are labeled by these propositions). Let ℓ'' be this new valuation for \mathcal{T}' . Then clearly we have:

- $\mathcal{T}', \text{path}(\rho), n \models \Phi_{\text{path}}(\kappa + 1)$, since q propositions label a path;
- $\mathcal{T}', \text{path}(\rho), n \models q_{\alpha'}^{\kappa} \wedge q_{\alpha'}^{\kappa+1}$, because the current position belongs to the runs ρ and ρ' , and the new run ρ' is issued from the current position;
- $\mathcal{T}', \text{path}(\rho), n \models \Phi_{\text{out}}(\kappa + 1, C \cup \{a\})$, meaning that the labeling of $q^{\kappa+1}$ propositions follows the "correct" path ρ' from $\text{Out}(\rho_{\leq n}, f_a \circ f_B)$;
- finally, $\mathcal{T}', \text{path}(\rho'), n \models \widetilde{\varphi}_p^{B \cup \{a\}, \kappa+1}$ by i.h.

Therefore we have: $\mathcal{T}', \text{path}(\rho), n \models \widetilde{\varphi}^{B, \kappa}$.

We now prove the converse implication. Assume $\mathcal{T}', \text{path}(\rho), n \models \widetilde{\varphi}^{B, \kappa}$. From the existence of a labeling for m_1^a, \dots, m_r^a satisfying $\Phi_{\text{strat}}(\{a\})$, we deduce a \sim -compatible strategy f_a in \mathcal{C} for every finite runs issued from $\rho(\leq n)$. Now consider a valuation for $q_1^{\kappa+1}, \dots, q_s^{\kappa+1}$. Either it makes the left-hand side subformula of the implication to be false, and there is no consequence, or this subformula is true and in this case, the valuation describes a run ρ' in \mathcal{C} issued from $\rho(n)$ and belonging in $\text{Out}(\rho_{\leq n}, f_a \circ f_B)$; this run has to satisfy $\widetilde{\varphi}_p^{B \cup \{a\}, \kappa+1}$, and by i.h. we get that $\mathcal{C}, \rho', n \models_{f_a \circ f_B} \varphi_p$. □

Corollary 11. *The model-checking problem of ATL_{sc}^{*} over uniform CGSOs is decidable (and Tower-complete).*

Remark 12. *Our algorithm can be used to decide whether one player (with partial observation of the system) has a compatible strategy to win against all the other players, whatever is the observation of the other players (since the implicit quantification in strategy quantifiers ranges over all the outcomes). As a consequence, when considering the fragment of ATL in which strategy quantifiers always involve at most one-player coalitions, the model-checking problem is decidable for possibly non-uniform partial observation. It can be checked that indeed the undecidability proof of [DT11] involves a strategy quantification over a coalition of two players with different observation.*

4 Restriction to memoryless strategies

In this section, we show that another way to obtain decidability for the model-checking problem of ATL_{sc}^{*} over game structure with (possibly non-uniform) partial observation is by restricting the set of considered strategies to memoryless strategies. We denote by ATL_{sc,0}^{*}, the temporal logic obtained from ATL_{sc}^{*} by replacing respectively the quantifiers over strategies $\langle A \cdot \rangle$ by $\langle A \cdot \rangle_0$ and $\langle \bar{A} \cdot \rangle$ by $\langle \bar{A} \cdot \rangle_0$. Given a CGSO \mathcal{C} a path ρ of \mathcal{C} , and $n \in \mathbb{N}$ such that $n < |\rho|$ and $f_B \in \text{Strat}_0^{\sim}(B)$, the semantics of these two operators is given by:

$$\begin{aligned} \mathcal{C}, \rho, n \models_{f_B} \langle A \cdot \rangle_0 \varphi_p & \quad \text{iff} \quad \exists g_A \in \text{Strat}_0^{\sim}(A). \forall \rho' \in \text{Out}(\rho_{\leq n}, g_A \circ f_B). \mathcal{C}, \rho', n \models_{g_A \circ f_B} \varphi_p \\ \mathcal{C}, \rho, n \models_{f_B} \langle \bar{A} \cdot \rangle_0 \varphi_p & \quad \text{iff} \quad \exists g_{\bar{A}} \in \text{Strat}_0^{\sim}(\text{Agt} \setminus A). \forall \rho' \in \text{Out}(\rho_{\leq n}, g_{\bar{A}} \circ f_B). \mathcal{C}, \rho', n \models_{g_{\bar{A}} \circ f_B} \varphi_p \end{aligned}$$

In [LM15], a reduction from model-checking ATL_{sc,0}^{*} and ATL_{sc,0}^{*} to QCTL^{*} model-checking problem over concurrent game structure with perfect observation is given, the main idea is to use the structure semantics instead of the tree semantics. In this framework the complexity is simpler, since both problems are PSPACE-complete: the number of memoryless strategies for one player being bounded (by $|\mathcal{M}|^{|\mathcal{Q}|}$), we can easily enumerate all of them, and store each strategy within polynomial space. When considering partial observation, we can use the same approach, we need to also ensure that the chosen strategies are compatible, but this can easily be achieved when considering memoryless strategies, since one only needs to check that a strategy f_a proposes for two equivalent states (w.r.t. \sim_a) the same move.

Let $\mathcal{C} = \langle Q, R, \ell, \text{Agt}, \mathcal{M}, \text{Mov}, \text{Edge}, (\sim_a)_{a \in \text{Agt}} \rangle$ be a CGSO with finite set of states $Q = \{q_0, \dots, q_s\}$, finite set of moves $\mathcal{M} = \{m_1, \dots, m_r\}$ and the finite set of agents $\text{Agt} = \{a_1, \dots, a_p\}$. For each $a \in \text{Agt}$, we suppose that the number of equivalence classes for \sim_a is n_a and we use the notation E_i^a with $1 \leq i \leq n_a$ to represent the i -th equivalence class of \sim_a . We denote by $\mathcal{K}_{\mathcal{C}}$ the Kripke structure $\langle Q, R, \ell_{\mathcal{K}} \rangle$ underlying \mathcal{C} based on the same states and the same transitions, and where the labeling function extends ℓ by adding labels from the sets $\{P_q \mid q \in Q\} \cup \{P_i^{\sim_a} \mid a \in \text{Agt} \wedge 1 \leq i \leq n_a\}$ in such a way that the following two conditions hold:

- $P_q \in \ell_{\mathcal{K}}(q')$ if, and only if, $q' = q$,
- for all $a \in \text{Agt}$, $P_i^{\sim_a} \in \ell_{\mathcal{K}}(q)$ if, and only if, $q \in E_i^a$.

Below we show how to translate a formula φ_s of $\text{ATL}_{sc,0}^*$ into a formula of $\widehat{\text{QCTL}}^*$ such that we will have $\mathcal{C}, q_0 \models \varphi_s$ for a state q_0 of \mathcal{C} iff $\mathcal{K}_{\mathcal{C}}, q_0 \models_{\text{st}} \widehat{\varphi}_s$. Note that at the opposite of the previous section, we consider here the *structure semantics* [LM14] when performing the model-checking of $\mathcal{K}_{\mathcal{C}}$: this means that instead of ranging over labelings of the execution tree, propositional quantification ranges over labelings of the Kripke structure. This is due to the fact that the compatible memoryless strategies we take here into account are functions mapping equivalent states (and not anymore path of equivalent states) to the same move. Given a coalition B in Agt and φ a subformula of φ_s , we define a $\widehat{\text{QCTL}}^*$ formula $\widehat{\varphi}^B$ inductively as follows:

$$\begin{aligned} \widehat{P}^B &= P & \widehat{(A)} \varphi_s &= \widehat{\varphi}_s^{B \setminus A} \\ \widehat{\varphi \wedge \psi}^B &= \widehat{\varphi}^B \wedge \widehat{\psi}^B & \widehat{\neg \varphi}^B &= \neg \widehat{\varphi}^B \\ \widehat{\varphi_p \mathbf{U} \psi_p}^B &= \widehat{\varphi}_p^B \mathbf{U} \widehat{\psi}_p^B & \widehat{\mathbf{X} \varphi_p}^B &= \mathbf{X} \widehat{\varphi}_p^B \end{aligned}$$

For a formula of the shape $\langle \cdot \rangle_0 \varphi_p$, we let³:

$$\widehat{\langle \cdot \rangle_0 \varphi_p}^B = \exists m_1^r \dots m_k^a. \left(\Phi_{\text{strat}_0}(\{a\}) \wedge \mathbf{A} \left[(\Phi'_{\text{out}}(\{a\} \cup B)) \Rightarrow \widehat{\varphi}_p^{B \cup \{a\}} \right] \right)$$

where

$$\Phi_{\text{strat}_0}(A) = \bigwedge_{a \in A} \left[\mathbf{AG} \left(\bigvee_{1 \leq j \leq r} (m_j^a \wedge \bigwedge_{j' \neq j} \neg m_{j'}^a) \right) \wedge \bigwedge_{1 \leq c \leq n_a} \bigwedge_{1 \leq j \leq r} \mathbf{EF} (P_c^{\sim_a} \wedge m_j^a) \Rightarrow \mathbf{AG} (P_c^{\sim_a} \Rightarrow m_j^a) \right]$$

and

$$\Phi'_{\text{out}}(B) = \mathbf{G} \left(\bigwedge_{q \in Q} \bigwedge_{\left\{ (m, q') \mid \begin{array}{l} m \in \mathcal{M}^{\text{Agt} \setminus B} \\ q' = \text{Edge}(q, m) \end{array} \right\}} \left[(P_q \wedge \widehat{\mathbf{m}}^B) \Rightarrow \mathbf{X} P_{q'} \right] \right).$$

where $\widehat{\mathbf{m}}^B$ is a shorthand for the formula $\bigwedge_{\{(b, m_k) \mid b \in B \wedge \mathbf{m}[b] = m_k\}} m_k^b$. The last Formula $\Phi'_{\text{out}}(B)$ characterizes the outcomes of the strategies in use for some coalition B . We point out that the $\widehat{\text{QCTL}}^*$ formula $\widehat{\varphi}_s^{\emptyset}$ has size $O(|\Phi| \cdot |Q| \cdot (|\text{Agt}| \cdot |\mathcal{M}|^2 + |Q| \cdot |\text{Edge}|))$.

A memoryless and compatible strategy f_B can easily be encoded by labeling states of $\mathcal{K}_{\mathcal{C}}$ with the moves of the coalition B thanks to the propositions m_1^a, \dots, m_r^a for $a \in B$. The $\Phi_{\text{strat}_0}(B)$ is here to ensure both that the labels m^a correspond to a strategy for coalition B and to check that such a strategy is memoryless and compatible: for all *reachable* states labeled with some move, the formula verifies that all

³As in the previous case, the extension to $\langle \cdot \rangle_0$ is straightforward.

the equivalent states are labeled by the same move. Note that this property ranges over reachable states: we could have unreachable states that are labeled in an improper way but this is not a problem as such state will never be reached from the current position (and consequently they cannot impact on the truth value of the formula).

We will now consider the Kripke structure $\mathcal{K}'_{\mathcal{C}} = \langle Q, R, \ell' \rangle$ where ℓ' extends the function $\ell_{\mathcal{K}}$ with propositions of type m_j^a . Given a compatible memoryless strategy $f_B = \{f_b\}_{b \in B}$ for the coalition B , we say that ℓ' is an f_B -labeling if for every node $q \in Q$, for any $b \in B$ and for any $1 \leq j \leq r$, we have $m_j^b \in \ell'(q)$ iff $f_b(q) = m_j$ (we recall that a memoryless strategy f_b can be seen as a function from Q to \mathcal{M}).

Proposition 13. *Let ρ be a run of \mathcal{C} , n be a natural number, f_B be a compatible strategy in $\text{Strat}_0^{\sim}(B)$ for coalition B . If ℓ' is an f_B -labeling, then we have:*

$$\mathcal{C}, \rho, n \models_{f_B} \varphi_s \quad \text{iff} \quad \mathcal{K}'_{\mathcal{C}}, \rho, n \models \widehat{\varphi}_s^B$$

Proof. The proof is done by structural induction over Φ .

- $\Phi = \varphi_p \mathbf{U} \psi_p$: Assume $\mathcal{C}, \rho, n \models_{f_B} \varphi_p \mathbf{U} \psi_p$. Therefore there exists $i \geq n$ s.t. $\mathcal{C}, \rho, i \models_{f_B} \psi_p$ and for any $n \leq j < i$, we have $\mathcal{C}, \rho, j \models_{f_B} \varphi_p$. For every position between n and i , the induction hypothesis can be applied and we deduce $\mathcal{K}'_{\mathcal{C}}, \rho, n \models \widehat{\varphi}_p^B \mathbf{U} \widehat{\psi}_p^B$. The converse is done similarly.
- $\Phi = \langle a \rangle \varphi_p$: (1) \Rightarrow (2). Assume $\mathcal{C}, \rho, n \models_{f_B} \langle a \rangle \varphi_p$. There exists a memoryless compatible strategy f_a s.t. for any $\rho' \in \text{Out}(\rho_{\leq n}, f_a \circ f_B)$, we have $\mathcal{C}, \rho', n \models_{f_a \circ f_B} \varphi_p$. Thus we can find a labeling of propositions m_1^a, \dots, m_r^a for $\mathcal{K}'_{\mathcal{C}}$ to represent f_a . This labeling completes the existing one for strategy context F_B and the formula $\Phi_{\text{strat}_0}(\{a\} \cup B)$ is then satisfied on $\mathcal{K}'_{\mathcal{C}}, \rho, n$. And any $\mathcal{K}'_{\mathcal{C}}$ -run satisfying $\Phi'_{\text{out}}(\{a\} \cup B)$ belongs to the set of outcomes generated by the strategy context $f_a \circ f_B$, and then satisfies $\widehat{\varphi}_p^{\{a\} \cup B}$ by induction hypothesis.
(2) \Rightarrow (1). Now assume $\mathcal{K}'_{\mathcal{C}}, \rho, n \models \widehat{\varphi}_p^B$. Therefore there exists a labeling for m_1^a, \dots, m_r^a such that $\Phi_{\text{strat}_0}(\{a\})$ holds true. Such a labeling defines a memoryless and compatible strategy for the reachable states (from ρ, n). And finally every run satisfying $\Phi'_{\text{out}}(\{a\} \cup B)$ has to satisfy $\widehat{\varphi}_p^{\{a\} \cup B}$. By induction hypothesis we get the stated result. \square

Thus we have:

Theorem 14. *The model-checking problem of $\text{ATL}_{sc,0}^*$ is PSPACE-complete.*

Proof. QCTL* model checking can be solved in polynomial space for the structure semantics. This immediately gives a PSPACE algorithm for our problem. Conversely, $\text{ATL}_{sc,0}$ model checking (with perfect observation) is PSPACE-hard [BDLM09]. This extends immediately to the partial-observation setting. \square

Remark 15. *Quantification over memoryless strategies could also be achieved using the tree semantics, following the presentation of Section 3. To do so, it suffices to label each state with its name (hence adding a few extra atomic propositions) and to require that the labeling of the execution tree with strategies satisfies that whenever some state s is labeled with some move m_i , then all the occurrences of s must be labeled with the same mode m_i .*

While this has little interest for model checking $\text{ATL}_{sc,0}^$ in terms of complexity, it shows that model checking remains decidable for the logic involving quantification over both memoryful, bounded-memory and memoryless strategies.*

5 Satisfiability and partial observation

In this section, we consider satisfiability checking: given a formula Φ , we look for a game structure \mathcal{C} , an equivalence \sim over states, and a control state q_0 , such that $\mathcal{C}, q_0 \models \Phi$. This problem is undecidable for ATL_{sc} (and ATL_{sc}^*) in the classical setting of perfect-observation games.

First note that considering partial observation makes the problem different: there exists formulas that are satisfiable under partial observation, and not satisfiable for full observation. Consider formula $\Phi = \mathbf{AX}(\langle a_1 \cdot \rangle \mathbf{X}f) \wedge \neg \langle a_1 \cdot \rangle \mathbf{XX}f$. Clearly Φ is satisfiable when considering games with partial observation: for example, one can consider the turn-based structure in Figure 3, where a_1 plays in circle nodes and a_2 plays in box nodes, and where \sim is \equiv_P : from q_0 , there is no \sim -strategy for a_1 ensuring the property $\mathbf{XX}f$ (because from the strategy should play the same move after q_0q_1 and q_0q_2), but the subformula $\langle a_1 \cdot \rangle \mathbf{X}f$ holds for true in q_1 and in q_2 .




Fig. 3: The turn-based CGS \mathcal{C} s.t. $q_0 \models \Phi$

But formula Φ is not satisfiable in the classical case: assume that $\mathbf{AX}(\langle a_1 \cdot \rangle \mathbf{X}f)$ is satisfied in some state q , then from every q -successor q' , there is a strategy $f_{q'}$ for a_1 to ensure $\mathbf{X}f$. Therefore the strategy for a_1 consisting, from q , to choose an arbitrary move, and then to choose the strategy $f_{q'}$ for every possible successor q' ensures the property $\mathbf{XX}f$.

From a decidability point of view, considering partial observation does not make satisfiability problems to be simpler: in fact this problem remains undecidable for ATL_{sc} . Furthermore, it is even worse than in the classical setting: while the turn-based satisfiability is decidable when perfect information is assumed, it is undecidable when one considers the partial observation case.

Theorem 16. *Satisfiability problems for ATL_{sc}^* , ATL_{sc} , $\text{ATL}_{sc,0}^*$ and $\text{ATL}_{sc,0}$ (with partial observation) are undecidable even restricted to turn-based structures.*

Proof. We can mostly reuse the proof of Troquard and Walther [TW12] (we have slightly modified in [LM15]). The key idea of their proof is to reduce $S5^n$ satisfiability to ATL_{sc} satisfiability. Given an $S5^n$ formula Φ , one can build an ATL_{sc} formula $\langle \bar{\theta} \cdot \rangle \widehat{\Phi}$ such that Φ is satisfiable iff $\langle \bar{\theta} \cdot \rangle \widehat{\Phi}$ is satisfiable. Without considering the details of the proof, one can just note that $\widehat{\Phi}$ uses Boolean operators and strategies quantifiers $\langle a_i \cdot \rangle$ (for $1 \leq i \leq n$) and formulas of the form $\mathbf{X}P$: there is no \mathbf{U} modality and there is no nesting of \mathbf{X} . Therefore with such a formula, every strategies quantifier is interpreted in a unique state w and we only consider the moves done in this state w : adding an equivalence \sim over states or considering memoryless strategies do not change the semantics of $\widehat{\Phi}$.

Assume that Φ is satisfiable. From the proof of [TW12], one can build a game structure \mathcal{C} satisfying $\langle \bar{\theta} \cdot \rangle \widehat{\Phi}$. Moreover one can use the reduction of Theorem 6 to get a turn-based games with partial information.

Now assume that $\langle \bar{\theta} \cdot \rangle \widehat{\Phi}$ is satisfiable. Therefore there exists a game structure \mathcal{C} with an equivalence \sim such that $\mathcal{C}, q \models \langle \bar{\theta} \cdot \rangle \widehat{\Phi}$. This structure is based on a set of agents $\{1, \dots, k\}$ with $k \leq n$. And there is a strategy F for Agt such that $\mathcal{C}, q \models_F \widehat{\Phi}$ and $\widehat{\Phi}$ may only modify the choices for players a_1, \dots, a_n . If $k > n$,

we can replace the players $n + 1, \dots, k$ by their first move selected by F from q . It gives a game structure based on $\text{Agt} = \{a_1, \dots, a_n\}$. And we can use the same construction of a corresponding $S5^n$ model for Φ as it is done in [LM15], and as explained before, considering partial information does not change the construction since every strategy is applied from the same unique initial state of the game. \square

6 Conclusion

In this paper, we have proved that the model-checking of ATL_{sc} over games with *uniform* partial observation is decidable. It would be interesting to study whether our uniformity requirement on the observation of each player in the game could be relaxed in order to be able to analyze richer models. One possible direction is to look at games with hierarchical information [BMV15], but we currently could not find a way of extending our algorithm to non-uniform observation.

References

- [AHK02] R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.
- [BDLM09] Th. Brihaye, A. Da Costa, F. Laroussinie, and N. Markey. ATL with strategy contexts and bounded memory. In LFCS’09, LNCS 5407, p. 92–106. Springer, 2009.
- [BMV15] D. Berwanger, A. B. Mathew, and M. Van den Bogaard. Hierarchical information patterns and distributed strategy synthesis. In ATVA’15, LNCS. Springer, 2015.
- [CE82] E. M. Clarke and E. A. Emerson. Design and synthesis of synchronization skeletons using branching-time temporal logic. In LOP’81, LNCS 131, p. 52–71. Springer, 1982.
- [CHP07] K. Chatterjee, T. A. Henzinger, and N. Piterman. Strategy logic. In CONCUR’07, LNCS 4703, p. 59–73. Springer, 2007.
- [DLM12] A. Da Costa, F. Laroussinie, and N. Markey. Quantified CTL: Expressiveness and model checking. In CONCUR’12, LNCS 7454, p. 177–192. Springer, 2012.
- [DT11] C. Dima and F. L. Țiplea. Model-checking ATL under imperfect information and perfect recall semantics is undecidable. Research Report 1102.4225, arXiv, 2011.
- [KV01] O. Kupferman and M. Y. Vardi. Synthesizing distributed systems. In LICS’01, p. 389–398. IEEE Comp. Soc. Press, 2001.
- [LM13] F. Laroussinie and N. Markey. Satisfiability of ATL with strategy contexts. In GandALF’13, Electronic Proceedings in Theoretical Computer Science 119, p. 208–223, 2013.
- [LM14] F. Laroussinie and N. Markey. Quantified CTL: expressiveness and complexity. *Logical Methods in Computer Science*, 10(4), 2014.
- [LM15] F. Laroussinie and N. Markey. Augmenting ATL with strategy contexts. *Information and Computation*, 2015. To appear.
- [MMV10] F. Mogavero, A. Murano, and M. Y. Vardi. Reasoning about strategies. In FSTTCS’10, Leibniz International Proceedings in Informatics 8, p. 133–144. Leibniz-Zentrum für Informatik, 2010.
- [Pnu77] A. Pnueli. The temporal logic of programs. In FOCS’77, p. 46–57. IEEE Comp. Soc. Press, 1977.
- [QS82] J.-P. Queille and J. Sifakis. Specification and verification of concurrent systems in CESAR. In SOP’82, LNCS 137, p. 337–351. Springer, 1982.

- [Sch04] P.-Y. Schobbens. Alternating-time logic with imperfect recall. In LCMAS'03, ENTCS 85(2), p. 82–93. Elsevier Science, 2004.
- [Sch08] S. Schewe. *Synthesis of Distributed Systems*. PhD thesis, Saarland University, Germany, 2008.
- [TW12] N. Troquard and D. Walther. On satisfiability in ATL with strategy contexts. In JELIA'12, LNCS 7519, p. 398–410. Springer, 2012.