

A Hennessy-Milner Theorem for ATL with Imperfect Information

Francesco Belardinelli^{1,2}, Catalin Dima³, Vadim Malvone², and Ferucio Tiplea⁴

¹ Imperial College London, UK

² Université d'Évry, France

³ Université Paris-Est Créteil, France

⁴ Universitatea Al. I. Cuza, Iași, Romania

Abstract. We show that a history-based variant of alternating bisimulation with imperfect information allows it to be related to a variant of Alternating-time Temporal Logic (ATL) with imperfect information by a full Hennessy-Milner theorem. The variant of ATL we consider has a *common knowledge* semantics, which requires that the uniform strategy available for a coalition to accomplish some goal must be common knowledge inside the coalition, while other semantic variants of ATL with imperfect information do not accommodate a Hennessy-Milner theorem. We also show that the existence of a history-based alternating bisimulation between two finite Concurrent Game Structures with imperfect information (iCGS) is undecidable.

1 Introduction

Alternating-time Temporal Logic (ATL) [3] is a powerful logic for specifying strategic abilities of individual agents and coalitions in multi-agent game structures. Crucially, ATL has been extended to games of imperfect information [17] with various flavors related to the agents' knowledge of the existence of strategies for accomplishing the coalition's goals [2, 8, 9]. In this contribution, we focus on the *common knowledge* (*ck*) interpretation of ATL under imperfect information, which was first put forward in [17], along with its *objective* and *subjective* interpretations. However, differently from the latter, to the best of our knowledge, the *ck* interpretation has nowhere else been considered in the literature. Nonetheless, the *ck* interpretation allows us to prove a Hennessy-Milner theorem for ATL under imperfect information for the memoryful notion of bisimulation we introduce in this paper. This result is in marked contrast with the situation for the other interpretations, which do not enjoy the Hennessy-Milner property [16].

The literature on bisimulations for modal logics is extensive, an in-depth survey of model equivalences for various temporal logics appears in [15]. The landscape for logics of strategic abilities, including ATL, is comparatively more sparse. A proof of the Hennessy-Milner property for ATL* with perfect information was already given in the paper introducing alternating bisimulations [4]. Since then, there have been numerous attempts to extend bisimulations to

more expressive languages (including Strategy Logic recently [7]), as well as to contexts of imperfect information [1, 10, 5]. In [10, 18] non-local model equivalences for ATL with imperfect information have been put forward. However, these works do not deal with the imperfect information/perfect recall setting here considered, nor do they provide a local account of bisimulations. Further, in [5] the authors consider a *memoryless* notion of bisimulation for ATL, under imperfect information. Unfortunately, their definition does not allow for the Hennessy-Milner property. We also note the results from [11] which show that ATL with imperfect information is incompatible in expressive power when compared with the modal-epistemic μ -calculus, contrary to what is known for the perfect information case. The present contribution extends the notion of alternating bisimulation to the setting of imperfect information and perfect recall so that it satisfies the Hennessy-Milner property: two game structures are bisimilar iff they satisfy the same formulas in ATL.

The classic proof for Hennessy-Milner type properties typically uses bisimulation games played between SPOILER and DUPLICATOR. These bisimulation games are turn-based, perfect information, safety games (in regards of DUPLICATOR's objective) played on a tree whose nodes are labeled with pairs of states (or histories, in case of a memoryful semantics) of the two game structures checked for bisimulation. Hence, such games are determined, and determinacy plays a crucial role since, when there is no bisimulation between the two structures, the bisimulation game cannot be won by DUPLICATOR, and hence SPOILER has a winning strategy, which is then used for exhibiting a formula that is satisfied in one structure but not in the other.

The extension of this proof technique to ATL with imperfect information has to cope with the fact that any notion of bisimulation has to account for the fact that coalitions have to choose action profiles in indistinguishable states in a “uniform” way: agents that do not distinguish between two states must choose the same actions in both. Uniformity entails a slightly more involved notion of bisimulation which utilizes *strategy simulators* [5]. Then, any bisimulation game has to encode these strategy simulators, in the sense that DUPLICATOR is given the choice of a uniform strategy in some common-knowledge neighbourhood in one of the game structures and the SPOILER has to reply with a uniform strategy in the related common-knowledge neighbourhood of the other game structure.

The problem raised by this generalization is that positions in a bisimulation game are normally labeled with histories, not common-knowledge neighbourhoods, as bisimulations relate the former, not the latter. So, we need both a SPOILER and a DUPLICATOR who have imperfect information at each position of the bisimulation game. On the other hand, as it is the case with bisimulations for the perfect information case, for each choice of strategies in the two structures, the outcomes of one strategy have to be related with the outcomes of the other strategy. But this requires both SPOILER and DUPLICATOR to be *perfectly-informed!*

The solution we propose is a 4-player bisimulation game played between the SPOILER coalition {I-SPOIL, P-SPOIL} and the DUPLICATOR coalition {I-

DUPL, P-DUPL}, where both I-players have imperfect information, while both P-players have perfect information. We show that such a game can be won by the DUPLICATOR coalition if and only if there exists a bisimulation between the two game structures.

Further, we provide a Gale-Stewart type determinacy theorem [14] for the bisimulation game, showing that exactly one of the two coalitions has a winning joint strategy. The key point is that, when DUPLICATOR does not have a winning strategy, the strategic choices for I-SPOIL can be defined in a uniform way that is only based on I-SPOIL’s observations. To the best of our knowledge, this is the first example of a class of multi-player, imperfect information, zero-sum (reachability) games played over infinite trees that are determined. Note that, for technical reasons, our Hennessy-Milner theorem only holds for ATL with the “yesterday” modality **Y**.

Moreover, we analyse the problem of checking the existence of a bisimulation between two given game structures. We show that this problem is undecidable in general by building on the undecidability of the model-checking problem for ATL with imperfect information and perfect recall. More specifically, given a Turing machine M , we build a game structure in which a two-agent coalition has a strategy for avoiding an error state if and only if M halts when starting with an empty tape. We then build a second, unrelated, simple game structure in which the same coalition always has an avoiding strategy. Finally, we prove that the two structures are bisimilar if and only if M halts.

Scheme of the paper. In Sec. 2 we recall the syntax and semantics of ATL according to various flavors of imperfect information (and perfect recall). Sec. 3 extends the bisimulation in [5] to the case of perfect recall, and shows that bisimilar game structures satisfy the same formulas in ATL. Then, in Sec. 4 we introduce our variant of bisimulation games, for which we prove that the DUPLICATOR coalition has a winning strategy if and only if there exists a bisimulation between the two given game structures. In Sec. 5 we prove the Gale-Stewart determinacy theorem for our bisimulation games, which allows us to prove the Hennessy-Milner theorem. Finally, in Sec. 6 we show that checking the existence of a bisimulation between two given game structures is undecidable in general.

2 ATL with Imperfect Information

In this section we present the syntax and semantics of the Alternating-time Temporal Logic ATL^* [3]. In the rest of the paper we assume a set AP of atomic propositions (or atoms) and a set Ag of agents.

Definition 1 (ATL^*). *History formulas ϕ and path formulas ψ in ATL^* are defined by the following BNF, where $p \in AP$ and $A \subseteq Ag$:*

$$\begin{aligned} \phi &::= p \mid \neg\phi \mid \phi \rightarrow \phi \mid \langle\langle A \rangle\rangle\psi \\ \psi &::= \phi \mid \neg\psi \mid \psi \rightarrow \psi \mid X\psi \mid \mathbf{Y}\psi \mid \psi U \psi \end{aligned}$$

The formulas in ATL^ are all and only the history formulas.*

The ATL* operator $\langle\langle A \rangle\rangle$ intuitively means that ‘the agents in coalition A have a (collective) strategy to achieve ...’, where the goals are LTL formulas built by using operators ‘next’ X and ‘until’ U . We define A -formulas as the formulas in ATL* in which A is the only coalition appearing in ATL* modalities.

Notice that we talk about *history* formulas, rather than *state* formulas as customary, as such formulas will be interpreted on histories rather than states as per perfect recall.

We provide ATL* with both the objective and subjective variants [17] of the history-based semantics with imperfect information and perfect recall, as well as a novel interpretation based on *common knowledge* [13].

Definition 2 (iCGS). *Given sets AP of atoms and Ag of agents, a concurrent game structure with imperfect information, or iCGS, is a tuple $\mathcal{G} = \langle Ag, S, s_0, Act, \{\sim_i\}_{i \in Ag}, d, \rightarrow, \pi \rangle$ where*

- S is a non-empty set of states and $s_0 \in S$ is the initial state of \mathcal{G} .
- Act is a finite non-empty set of actions. A tuple $\vec{a} = (a_i)_{i \in Ag} \in Act^{Ag}$ is called a joint action.
- For every agent $i \in Ag$, \sim_i is an equivalence relation on S , called the indistinguishability relation for agent i .
- $d : Ag \times S \rightarrow (2^{Act} \setminus \{\emptyset\})$ is the protocol function, satisfying the property that, for all states $s, s' \in S$ and any agent i , $s \sim_i s'$ implies $d(i, s) = d(i, s')$. That is, the same (non-empty) set of actions is available to agent i in indistinguishable states.
- $\rightarrow \subseteq S \times Act^{Ag} \times S$ is the transition relation such that, for every state $s \in S$ and joint action $\vec{a} \in Act^{Ag}$, $(s, \vec{a}, s') \in \rightarrow$ for some state $s' \in S$ iff $a_i \in d(i, s)$ for every agent $i \in Ag$. We normally write $s \xrightarrow{\vec{a}} r$ for $(s, \vec{a}, r) \in \rightarrow$.
- $\pi : S \rightarrow 2^{AP}$ is the state-labeling function.

Runs. Given an iCGS \mathcal{G} , a *run* is a finite or infinite sequence $\rho = s_0 \vec{a}_0 s_1 \dots$ in $((S \cdot Act^{Ag})^* \cdot S) \cup (S \cdot Act^{Ag})^\omega$ such that for every $j \geq 0$, $s_j \xrightarrow{\vec{a}_j} s_{j+1}$. Given a run $\rho = s_0 \vec{a}_0 s_1 \dots$ and $j \geq 0$, $\rho[j]$ denotes the $j + 1$ -th state s_j in the sequence and $\rho[j, k]$ denotes the sequence of states from the $j + 1$ -th state to the $k + 1$ -th state; while $\rho_{\geq j}$ (or $\rho[\geq j]$) denotes run $s_j \vec{a}_j s_{j+1} \dots$ starting from $\rho[j]$, and $\rho_{\leq j}$ (or $\rho[\leq j]$) denotes run $s_0 \vec{a}_0 s_1 \dots \vec{a}_{j-1} s_j$. Further, with $act_i(h, m)$ we denote the m -th action of agent i in history h .

We call finite runs *histories*, denote them as $h \in H$, their length as $|h| \in \mathbb{N}$, and their last element $h_{|h|-1}$ as *last*(h); whereas infinite runs are called *paths* and denoted as $\lambda, \lambda' \in P$. We denote the set of all histories (resp. paths) in an iCGS \mathcal{G} as $Hist(\mathcal{G})$ (resp. $Path(\mathcal{G})$). Notice that states are instances of histories of length 1. Accordingly, several notions defined below for histories can also be applied to states. Finally, we write $h \leq \rho$ to say that h is the prefix of ρ , that is $h = \rho[\leq |h|]$.

For a coalition $A \subseteq Ag$ of agents, a *joint A -action* denotes a tuple $\vec{a}_A = (a_i)_{i \in A} \in Act^A$ of actions, one for each agent in A . For coalitions $A \subseteq B \subseteq Ag$ of agents, a joint A -action \vec{a}_A is *extended* by a joint B -action \vec{b}_B , denoted $\vec{a}_A \sqsubseteq \vec{b}_B$,

if for every $i \in A$, $a_i = b_i$. Also, a joint A -action \vec{a}_A is *enabled* at state $s \in S$ if for every agent $i \in A$, $a_i \in d(i, s)$.

Epistemic neighbourhoods. We extend the indistinguishability relations \sim_i , for $i \in Ag$, to histories in a synchronous, point-wise manner: $h \sim_i h'$ iff $|h| = |h'|$ and for all $m \leq |h|$, $h_m \sim_i h'_m$ and $act_i(h, m) = act_i(h', m)$.

Given a coalition $A \subseteq Ag$ of agents, the *collective knowledge relation* \sim_A^E is defined as $\bigcup_{i \in A} \sim_i$, while the *common knowledge relation* \sim_A^C is the transitive closure $(\bigcup_{i \in A} \sim_i)^+$ of \sim_A^E . Then, $C_A^G(h) = \{h' \in H \mid h' \sim_A^C h\}$ is the *common knowledge neighbourhood* (CKN) of history h for coalition A in the iCGS \mathcal{G} . We will omit the superscript \mathcal{G} whenever it is clear from the context.

Uniform strategies. We introduce a notion of strategy for the interpretation of $\langle\langle A \rangle\rangle$ modalities.

Definition 3 (Strategy). A (uniform, memoryfull) strategy for an agent $i \in Ag$ is a function $\sigma : H \rightarrow Act$ that is compatible with d and \sim_i , that is, for all histories $h, h' \in H$, $\sigma(h) \in d(i, last(h))$ and $h \sim_i h'$ implies $\sigma(h) = \sigma(h')$.

We denote by Σ_R the set of all memoryfull uniform strategies.

A strategy for a coalition A of agents is a set $\sigma_A = \{\sigma_a \mid a \in A\}$ of strategies, one for each agent in A . Given coalitions $A \subseteq B \subseteq Ag$, a strategy σ_A for coalition A , a state $s \in S$, and a joint B -action $\vec{b}_B \in Act^B$ that is enabled at s , we say that \vec{b}_B is *compatible with σ_A (in s)* whenever $\sigma_A(s) \subseteq \vec{b}_B$. For states $s, s' \in S$ and strategy σ_A , we write $s \xrightarrow{\sigma_A(s)} r$ if $s \xrightarrow{\vec{a}} r$ for some joint action $\vec{a} \in Act^{Ag}$ that is compatible with σ_A .

We define three notions of *outcome* of strategy σ_A at history h , corresponding to the *objective*, *subjective*, and *common knowledge* interpretation of alternating-time operators. Fix a history h and a strategy σ_A for coalition A .

1. The set of *objective outcomes* of σ_A at h is defined as $out_{obj}(h, \sigma_A) = \{\lambda \in P \mid \lambda_{\leq |h|} = h \text{ and for all } j \geq |h|, \lambda[j] \xrightarrow{\sigma_A(\lambda_{\leq j})} \lambda[j+1]\}$.
2. The set of *subjective outcomes* of σ_A at h is defined as $out_{subj}(h, \sigma_A) = \bigcup_{i \in A, h' \sim_i h} out_{obj}(h', \sigma_A)$.
3. The set of *common knowledge (ck) outcomes* of σ_A at h is defined as $out_{ck}(h, \sigma_A) = \bigcup_{h' \in C_A(h)} out_{obj}(h', \sigma_A)$.

Intuitively, objective outcomes are paths beginning with the current history h and consistent with the current joint strategy σ_A ; whereas subjective (resp. common knowledge) outcomes are paths beginning with some history h' *indistinguishable* from h according to collective (resp. common) knowledge (as well as consistent with σ_A). Again, notions of outcomes from states can be obtained from the definitions above, as states are a particular type of histories.

Definition 4. Given an iCGS \mathcal{G} , a history formula ϕ , path formula ψ , and $m \in \mathbb{N}$, the *subjective* (resp. *objective*, *common knowledge*) *satisfaction* of ϕ at history h and of ψ in path λ , denoted $(\mathcal{G}, h) \models_x \phi$ and $(\mathcal{G}, \lambda, m) \models_x \psi$ for

$x \in \{\text{subj}, \text{obj}, \text{ck}\}$, is defined recursively as follows (clauses for Boolean operators are immediate and thus omitted):

$$\begin{array}{ll}
(\mathcal{G}, h) \models_x p & \text{iff } p \in \pi(\text{last}(h)) \\
(\mathcal{G}, h) \models_x \langle\langle A \rangle\rangle \psi & \text{iff for some } \sigma_A \in \Sigma_R, \\
& \text{for all } \lambda \in \text{out}_x(h, \sigma_A), (\mathcal{G}, \lambda, |h|) \models_x \psi \\
(\mathcal{G}, \lambda, m) \models_x \phi & \text{iff } (\mathcal{G}, \lambda_{\leq m}) \models_x \phi \\
(\mathcal{G}, \lambda, m) \models_x X\psi & \text{iff } (\mathcal{G}, \lambda, m+1) \models_x \psi \\
(\mathcal{G}, \lambda, m) \models_x Y\psi & \text{iff } m \geq 1 \text{ and } (\mathcal{G}, \lambda, m-1) \models_x \psi \\
(\mathcal{G}, \lambda, m) \models_x \psi U \psi' & \text{iff for some } j \geq m, (\mathcal{G}, \lambda, j) \models_x \psi', \text{ and} \\
& \text{for all } k, m \leq k < j \text{ implies } (\mathcal{G}, \lambda, k) \models_x \psi
\end{array}$$

Remark 1. The individual and common knowledge operators K_i and C_A of epistemic logic [13] can be added to the syntax of ATL^* with the following (memoryful) interpretation:

$$\begin{array}{l}
(\mathcal{G}, h) \models_x K_i \phi \text{ iff for all } h' \sim_i h, (\mathcal{G}, h') \models_x \phi \\
(\mathcal{G}, h) \models_x C_A \phi \text{ iff for all } h' \in C_A(h), (\mathcal{G}, h') \models_x \phi
\end{array}$$

Within the subjective or the common knowledge interpretation of ATL^* , the individual knowledge operator becomes a derived operator, as we have $(\mathcal{G}, h) \models_x K_i \phi$ iff $(\mathcal{G}, h) \models_x \langle\langle i \rangle\rangle \phi U \phi$ for both $x \in \{\text{subj}, \text{ck}\}$. It is known that there exists no such definition for the knowledge operators in ATL^* within the objective interpretation. Furthermore, and only for the case of the common knowledge interpretation, we may similarly derive the common knowledge operator as well: $(\mathcal{G}, h) \models_{\text{ck}} C_A \phi$ iff $(\mathcal{G}, h) \models_{\text{ck}} \langle\langle A \rangle\rangle \phi U \phi$.

Example 1. We describe a coordination scenario comprising of two agents, 1 and 2, who have to agree on a meeting. But 1 does not know where she is, in Paris or London, and therefore which is the time zone, while 2 does not know if it is winter time or summer time. Agent 1 can choose either go to the meeting (g) or wait one hour (w) whereas 2 can choose either to go at 3pm (3) or at 4pm (4), local time. Now suppose it is 3pm GMT. In London, in the winter (s_2) 1 and 2 coordinate if 1 goes to the meeting and 2 goes at 3pm local time. They also meet if 1 waits one hour and 2 goes at 4pm. All other combined actions are unsuccessful. Analogously for Paris in the winter (s_1), and London in the summer (s_3). The iCGS \mathcal{G} depicted in Fig. 1 shows the described scenario. Since 1 and 2 have partial observability, 1 (resp. 2) cannot distinguish between states s_2 and s_1 (resp. s_3). After the initial choice, 1 and 2 stay indefinitely in either s_4 or s_5 . Finally, we use two atoms, to denote success (s) and failure (f), respectively.

As an example of specification in ATL^* , consider formula $\varphi = \langle\langle \{1, 2\} \rangle\rangle Xs$. This formula can be read as: 1 and 2 have a joint strategy to meet. Note that φ is true in both s_1 and s_3 when considering the subjective interpretation. However, is the truth of φ in s_1 and s_3 justified from point of view of the rational behaviour of 1 and 2? Specifically, since φ is true in s_1 according to the subjective interpretation, both 1 and 2 know that they have a successful strategy, which




Fig. 1. The meeting scenario in Example 1.

consists in playing action g for 1 and action 4 for 2. But for this strategy to be successful (i.e., satisfying Xs for all outcomes) it assumes that 2 is playing action 3 in s_2 : is such an assumption rationally justified? Notice that in s_2 , 2 considers state s_3 epistemically possible, and in s_3 the joint action $(g, 3)$ leads to failure. Hence, it does not appear to be rational for 2 to play 3 in s_2 . Ever more so that, by playing 4 in s_2 and s_3 , 2 can coordinate with 1 to achieve success. This example shows a scenario where, even though both agents know that their strategies are successful in principle, they do not necessarily coordinate, as they do not know that the other agent knows her strategy to be successful. Indeed, we have that in both s_1 and s_2 formula φ is false according to the common knowledge interpretation. So, it is not the case that they have common knowledge of their strategies being successful.

3 Alternating Bisimulations

In this section we introduce a notion of bisimulation suitable for concurrent game structures with imperfect information. In particular, we show that it preserves the satisfaction of formulas in ATL^* , when interpreted under imperfect information and perfect recall. Firstly, we present several auxiliary notions. Hereafter $\mathcal{G} = \langle Ag, S, s_0, Act, \{\sim_i\}_{i \in Ag}, d, \rightarrow, \pi \rangle$ and $\mathcal{G}' = \langle Ag, S', s'_0, Act', \{\sim'_i\}_{i \in Ag}, d', \rightarrow', \pi' \rangle$ are two iCGS defined on the same set Ag of agents, with histories $h \in Hist(\mathcal{G})$ and $h' \in Hist(\mathcal{G}')$.

Partial strategies. A *partial (uniform, memoryful) strategy* for agent $i \in Ag$ is a partial function $\sigma : H \rightarrow Act$ such that for each $h, h' \in H$, $\sigma(h) \in d(i, last(h))$, and $h \sim_i h'$ implies $\sigma(h) = \sigma(h')$. We denote the domain of partial strategy σ as $dom(\sigma)$. Given a coalition $A \subseteq Ag$, a *partial strategy for A* is a tuple $(\sigma_i)_{i \in A}$ of partial strategies, one for each agent $i \in A$. The set of partial strategies for A is denoted as $PStr_A$. Given a set $Q \subseteq H$ of histories and coalition $A \subseteq Ag$, we denote by $PStr_A(Q)$ the set of partial strategies whose domain is Q :

$$PStr_A(Q) = \{(\sigma_i)_{i \in A} \in PStr_A \mid dom(\sigma_i) = Q \text{ for all } i \in A\}$$

Additionally, given a (total or partial) strategy σ_A and a history $h \in \text{dom}(\sigma_A)$, define the set of *successors of h by σ* as

$$\text{succ}(h, \sigma_A) = \{h\alpha s \mid \alpha \in \text{Act}^{Ag} \text{ with } \sigma_A(h) \sqsubseteq \alpha \text{ and } h \xrightarrow{\alpha} s\}$$

Further, we set $\text{succ}(\sigma_A) = \bigcup_{h \in \text{dom}(\sigma_A)} \text{succ}(h, \sigma_A)$.

Definition 5 (Strategy simulators). *Given a coalition $A \subseteq Ag$, an A -strategy simulator (or simply strategy simulator, when A is understood from the context) is a family $ST = (ST_{C_A(h), C_A(h')})_{h \in \text{Hist}(\mathcal{G}), h' \in \text{Hist}(\mathcal{G}'})$ of mappings $ST_{C_A(h), C_A(h')} : PStr_A(C_A(h)) \rightarrow PStr_A(C_A(h'))$ such that for all histories $h, k \in \text{Hist}(\mathcal{G})$ and $h', k' \in \text{Hist}(\mathcal{G}')$,*

$$\begin{aligned} \text{if } C_A(h) = C_A(k) \text{ and } C'_A(h') = C'_A(k') \\ \text{then } ST_{C_A(h), C'_A(h')} = ST_{C_A(k), C'_A(k')} \quad (1) \end{aligned}$$

Hereafter, we simplify the notation by writing $ST(\sigma)$ instead of the cumbersome $ST_{C_A(h), C_A(h')}(\sigma)$, whenever h and h' are clear from the context and $\sigma \in PStr(C_A(h))$.

We can now introduce the notion of (bi)simulation for iCGS.

Definition 6 (Memoryful Simulation). *Let $A \subseteq Ag$ be a coalition of agents. A relation $\Rightarrow_A \subseteq \text{Hist}(\mathcal{G}) \times \text{Hist}(\mathcal{G}')$ is a simulation for A iff there exists a strategy simulator ST such that for any two histories $h \in \text{Hist}(\mathcal{G})$, $h' \in \text{Hist}(\mathcal{G}')$, $h \Rightarrow_A h'$ implies the following:*

1. $\pi(\text{last}(h)) = \pi'(\text{last}(h'))$;
2. For every $i \in A$ and $k' \in \text{Hist}(\mathcal{G}')$, if $h' \sim'_i k'$ then for some $k \in \text{Hist}(\mathcal{G})$, $h \sim_i k$ and $k \Rightarrow_A k'$.
3. For every pair of histories $k \in C_A(h)$ and $k' \in C'_A(h')$ with $k \Rightarrow_A k'$, for every partial strategy $\sigma_A \in PStr_A(C_A(h))$ and every history $l' \in \text{succ}(k', ST(\sigma_A))$, there exist a history $l \in \text{succ}(k, \sigma_A)$ such that $l \Rightarrow_A l'$.

A relation \Leftrightarrow_A is a bisimulation iff both \Rightarrow_A and its converse $\Leftarrow_A = \{(h', h) \mid h \Leftrightarrow_A h'\}$ are simulations.

We also extend (bi)-simulation to paths $\lambda \in \text{Path}(\mathcal{G})$, $\lambda' \in \text{Path}(\mathcal{G}')$, by denoting $\lambda \Rightarrow_A \lambda'$ iff for all $j \geq 0$, $\lambda_{\leq j} \Rightarrow_A \lambda'_{\leq j}$.

The main result of this section, Theorem 1, shows that bisimilar iCGS satisfy the same formulas in ATL* under imperfect information and perfect recall. To prove this result, we need the following auxiliary lemma:

Lemma 1. *If $h \Rightarrow_A h'$ then for every strategy σ_A , there exists a strategy σ'_A such that*

- (*) *for every path $\lambda' \in \text{out}_x(h', \sigma'_A)$, for $x \in \{\text{subj}, \text{obj}, \text{ck}\}$, there exists a path $\lambda \in \text{out}_x(h, \sigma_A)$ such that $\lambda \Rightarrow_A \lambda'$.*

Proof. First, notice that point 3 in Def. 6 can be rewritten as:

3. For all histories $k \in C_A(h)$ and $k' \in C'_A(h')$ such that $k \Rightarrow_A k'$, for all partial strategies $\sigma_A \in PStr_A(C_A(h))$, there exists a mapping $\rho_{\sigma_A, k, k'} : succ(k', ST(\sigma_A)) \rightarrow succ(k, \sigma_A)$ such that for all histories $l' \in succ(k', ST(\sigma_A))$, $\rho_{\sigma_A, k, k'}(l') \Rightarrow_A l'$.

in which the mapping $\rho_{\sigma_A, k, k'}$ represents the *skolemization*, in the original point 3, of the existential quantifier over $l \in succ(k, \sigma_A)$, seen as a unary function on $l' \in succ(k', ST(\sigma_A))$ indexed by $\sigma_A \in PStr_A(C_A(h))$, $k \in C_A(h)$ and $k' \in C'_A(h')$.

We now define the sequence $(dom^n(\sigma_A))_{n \in \mathbb{N}}$, of sets of histories in \mathcal{G} such that $k \in dom^n(\sigma_A)$ iff k can be reached in at most n steps from $C_A(h)$ by applying actions compatible with strategy σ_A . Formally, $dom^0(\sigma_A) = C_A(h)$ and

$$dom^{n+1}(\sigma_A) = dom^n(\sigma_A) \cup \bigcup_{k \in dom^n(\sigma_A)} \{C_A(l) \mid l \in succ(k, \sigma_A)\}$$

Also, we denote by σ_A^n the partial strategy resulting from restricting σ_A to $dom^n(\sigma_A)$.

We then define inductively a sequence $(\bar{\sigma}_A^n)_{n \in \mathbb{N}}$ of partial strategies in \mathcal{G}' such that $dom(\bar{\sigma}_A^n) \subseteq dom(\bar{\sigma}_A^{n+1})$ for every $n \in \mathbb{N}$, and, at the same time, a sequence of mappings $\theta_A^n : dom(\bar{\sigma}_A^n) \rightarrow dom^n(\sigma_A)$, satisfying the following property:

$$\theta_A^{n+1}(k') \in succ(\theta_A^n(k'_{\leq |k'| - 1}), \sigma_A) \quad (2)$$

The sequences $(\bar{\sigma}_A^n)_{n \in \mathbb{N}}$ and $(\theta_A^n)_{n \in \mathbb{N}}$ are defined as follows:

1. $dom(\bar{\sigma}_A^0) = C'_A(h')$;
 $dom(\bar{\sigma}_A^{n+1}) = dom(\bar{\sigma}_A^n) \cup \bigcup \{succ(k', \bar{\sigma}_A^n) \mid k' \in dom(\bar{\sigma}_A^n)\}$.
2. For all $k' \in dom(\bar{\sigma}_A^0)$, $\bar{\sigma}_A^0(k') = ST(\sigma_A^0)(k')$.
3. For all $k' \in C'_A(h')$, we fix a unique $k \in C_A(h)$ such that $k \Rightarrow_A k'$ (which exists by point 2 in Def. 6), and define $\theta_A^0(k') = k$.
4. For all $k' \in dom(\bar{\sigma}_A^{n+1})$, let $l' = k'_{\leq |k'| - 1}$. Then, we set $\theta_A^{n+1}(k') = \rho_{\sigma_A, \theta_A^n(l'), l'}(k')$.
5. For all $k' \in dom(\bar{\sigma}_A^{n+1})$,
 $\bar{\sigma}_A^{n+1}(k') = (ST_{C_A(\theta_A^n(k')), C'_A(k')})(\sigma_A|_{C_A(\theta_A^n(k'))})(k')$.

We prove property (2) above, as well as the following: for every $k' \in dom(\bar{\sigma}_A^n)$,

$$\theta_A^n(k') \Rightarrow_A k' \quad (*) \quad \theta_A^n(k') \in dom^n(\sigma_A) \quad (**)$$

Property (*) holds by definition, since $\rho_{\sigma_A, \theta_A^n(l'), l'}(k') \Rightarrow_A k'$. Property (2) and (**) can be proved by induction on $n = |k'| - |h'|$, by observing that (**) holds for $n = 0$; if (2) holds for $n + 1$ then (**) holds for $n + 1$ too; and finally, (2) is an immediate consequence of the definition of θ_A^n , property (**) Def. 6. Note that property (**) ensures that indeed $\theta_A^n : dom(\bar{\sigma}_A^n) \rightarrow dom^n(\sigma_A)$ as desired.

The "limit" of the sequence of strategies $(\bar{\sigma}_A^n)_{n \in \mathbb{N}}$ is still a partial strategy and the domain of each $\bar{\sigma}_A^n$ might not be closed under the common knowledge indistinguishability relation \sim_A^C . So, we extend first the domain of each $\bar{\sigma}_A^n$ to one which is closed under \sim_A^C in \mathcal{G}' , by constructing the sequence of partial strategies $(\hat{\sigma}_A^n)_{n \in \mathbb{N}}$ and the sequence of mappings $\hat{\theta}_A^n : dom(\hat{\sigma}_A^n) \rightarrow dom^n(\sigma_A)$, as follows:

1. $dom(\hat{\sigma}_A^0) = dom(\bar{\sigma}_A^0) = C'_A(h')$;
 $dom(\hat{\sigma}_A^{n+1}) = dom(\bar{\sigma}_A^n) \cup \bigcup \{C'_A(l') \mid \exists k' \in dom(\hat{\sigma}_A^n)$
with $l' \in succ(k', \bar{\sigma}_A^n) \cap C'_A(k')\}$.
2. For all $k' \in dom(\hat{\sigma}_A^0)$, $\hat{\sigma}_A^0(k') = ST(\sigma_A^0)(k')$.
3. For all $k' \in C'_A(h')$, $\hat{\theta}_A^0(k') = \theta_A^0(k')$.
4. For all $k' \in dom(\hat{\sigma}_A^{n+1})$, let $l' = k'_{\leq |k'|-1}$. Then, we set $\hat{\theta}_A^{n+1}(k') = \rho_{\sigma_A, \theta_A^n}(l', l')(k')$.
5. For all $k' \in dom(\hat{\sigma}_A^{n+1})$,
 $\hat{\sigma}_A^{n+1}(k') = (ST_{C_A(\theta_A^n(k')), C'_A(k')})(\sigma_A|_{C_A(\theta_A^n(k'))})(k')$.

We observe that properties (*) and (**) still hold for $\hat{\sigma}_A^n$ and $\hat{\theta}_A^n$, though property (2) does not in general. In this way we get that $dom(\hat{\sigma}_A^n) \supseteq dom(\bar{\sigma}_A^n)$ and for every $k' \in dom(\bar{\sigma}_A^n)$, $n \in \mathbb{N}$, $\hat{\sigma}_A^n(k') = \bar{\sigma}_A^n(k')$. As a result, the “limit” partial strategy $\hat{\sigma}_A = \bigcup_{n \in \mathbb{N}} \hat{\sigma}_A^n$ defined as $\hat{\sigma}_A(k') = \hat{\sigma}_A^{|k'|-|h'|}(k')$ is also uniform and

its domain $dom(\hat{\sigma}_A)$ is closed under \sim_A^C . We then transform it into a (total) uniform strategy σ'_A by imposing a fixed action $a_0 \in Act$ wherever $\hat{\sigma}_A^n$ is undefined, that is, $\sigma'_A(k') = \hat{\sigma}_A(k')$ for $k' \in dom(\hat{\sigma}_A)$ and $\sigma'_A(k') = a_0$ otherwise.

Finally, to prove property (*) for the common knowledge semantics, consider a path $\lambda' \in out_{ck}^{\mathcal{G}'}(h', \sigma'_A)$ and the sequence $(\theta_A^n(\lambda'_{\leq |h'|+n}))_{n \in \mathbb{N}}$ of histories in \mathcal{G} . By construction, $\theta_A^{n+1}(\lambda'_{\leq |h'|+n+1}) \in succ(\theta_A^n(\lambda'_{\leq |h'|+n}), \sigma_A)$ and $\theta_A^n(\lambda'_{\leq |h'|+n}) \Rightarrow_A \lambda'_{\leq |h'|+n}$, which means that this sequence of histories is in fact a path λ in \mathcal{G} which is compatible with σ_A and satisfies $\lambda \Rightarrow_A \lambda'$, which ends the proof.

By using Lemma 1 we are finally able to prove the main preservation result of this paper.

Theorem 1. *Let $h \in Hist(\mathcal{G})$ and $h' \in Hist(\mathcal{G}')$ be histories such that $h \Leftrightarrow_A h'$, and $\lambda \in Path(\mathcal{G})$ and $\lambda' \in Path(\mathcal{G}')$ be paths such that $\lambda \Leftrightarrow_A \lambda'$. Then, for every history A -formula ϕ , path A -formula ψ , $m \in \mathbb{N}$, and $x \in \{subj, obj, ck\}$,*

$$\begin{aligned} (\mathcal{G}, h) \models_x \phi \text{ iff } (\mathcal{G}', h') \models_x \phi \\ (\mathcal{G}, \lambda, m) \models_x \psi \text{ iff } (\mathcal{G}', \lambda', m) \models_x \psi \end{aligned}$$

Proof. The proof is by mutual induction on the structure of ϕ and ψ .

The case for propositional atoms is immediate as for $x \in \{subj, obj, ck\}$, $(\mathcal{G}, h) \models_x p$ iff $p \in \pi(last(h))$, iff $p \in \pi'(last(h'))$ by item 1 in Def. 6, iff $(\mathcal{G}', h') \models_x p$. The inductive cases for propositional connectives are also immediate.

For $\psi = \phi$, suppose that $(\mathcal{G}, \lambda, m) \models_x \psi$, that is, $(\mathcal{G}, \lambda_{\leq m}) \models_x \phi$. By assumption, $\lambda_{\leq m} \Leftrightarrow_A \lambda'_{\leq m}$ as well, and by induction hypothesis $(\mathcal{G}', \lambda'_{\leq m}) \models_x \phi$. Thus, $(\mathcal{G}', \lambda', m) \models_x \psi$.

For $\psi = X\psi'$, suppose that $(\mathcal{G}, \lambda, m+1) \models_x \psi'$. By the induction hypothesis, $(\mathcal{G}', \lambda', m+1) \models_x \psi'$. Thus, $(\mathcal{G}', \lambda', m) \models_x \psi$. The inductive cases for $\psi = Y\psi'$ and $\psi = \psi'U\psi''$ is similar.

Finally, for $\phi = \langle\langle A \rangle\rangle\psi$, $(\mathcal{G}, h) \models_x \phi$ iff for some strategy σ_A , for all $\lambda \in out_x^{\mathcal{G}}(h, \sigma_A)$, $(\mathcal{G}, \lambda, |h|) \models_x \psi$. By Lemma. 1, there exists strategy σ'_A s.t. for all

$\lambda' \in \text{out}_x^{\mathcal{G}'}(h', \sigma'_A)$, there exists $\lambda \in \text{out}_x^{\mathcal{G}}(h, \sigma_A)$ s.t. $\lambda \Leftarrow_A \lambda'$. Since $|h| = |h'|$, by the induction hypothesis $(\mathcal{G}, \lambda, |h|) \models_x \psi$ iff $(\mathcal{G}', \lambda', |h'|) \models_x \psi$. Hence, $(\mathcal{G}', h') \models_x \phi$.

4 Bisimulations Games

In this section we introduce bisimulation games played on two iCGS and we prove that the existence of a winning strategy for the DUPLICATOR coalition is equivalent to the existence of a bisimulation between the iCGS.

Definition 7 (Bisimulation Game). *Given iCGSs \mathcal{G} and \mathcal{G}' , defined on the same sets Ag of agents and AP of atoms, a relation $R \subseteq \text{Hist}(\mathcal{G}) \times \text{Hist}(\mathcal{G}')$, and a pair $(h_0, h'_0) \in \text{Hist}(\mathcal{G}) \times \text{Hist}(\mathcal{G}')$ of histories, we define the bisimulation game $\mathcal{B}(\mathcal{G}, \mathcal{G}', R, h_0, h'_0)$ as a turn-based game of imperfect information between four players: P-DUPL, P-SPOIL, called *P-players*, and I-DUPL, I-SPOIL, called *I-players*, organized in two coalitions: the DUPLICATOR coalition $\{\text{P-DUPL}, \text{I-DUPL}\}$ and the SPOILER coalition $\{\text{P-SPOIL}, \text{I-SPOIL}\}$, with both *P-players* having perfect information while both *I-players* have **the same** imperfect information.*

At a higher-level, the bisimulation game is a turn-based game in which the I-players are in charge of defining the strategy simulators, in the sense that I-SPOIL chooses a partial strategy for A over some common knowledge neighbourhood in one of the game structures, and I-DUPL responds with an appropriate partial strategy for A in the other game structure. Then the perfectly-informed players come into play, by appropriately defining mappings between histories compatible with the chosen strategies, which represent "skolemizations" of conditions (2) and (3) in Def. 6.

The necessity for I-SPOIL and I-DUPL to only have imperfect information comes from the fact that the same strategy profile has to be chosen by both players at positions which belong to the same common knowledge neighborhood in both game structures, since perfect information might be used by each player to trick the other player by choosing a strategy which is not uniform for some agent in coalition A.

More formally, the game proceeds as follows:

0. *The positions of the game form a **labeled tree**, denoted $T(\mathcal{B})$, with the root position labeled (h_0, h'_0) . The rest of positions and their labels are given below.*
1. *Each position (h, h') where $\pi(h) \neq \pi(h')$ or $(h, h') \notin R$ is **winning** for the SPOILER coalition.*
2. *Each position labeled $(h, h') \in \text{Hist}(\mathcal{G}) \times \text{Hist}(\mathcal{G}')$ belongs to I-SPOIL, and both I-players receive observation $C_A(h) \times C'_A(h')$. In each such position I-SPOIL may choose between two types of transitions:*
 - (a) *For each $\sigma_A \in P\text{Str}(C_A(h))$, a transition to a successor (of the current position in the tree) labeled (h, h', σ_A, L) .*
 - (b) *For each $\sigma'_A \in P\text{Str}(C'_A(h'))$ a transition to a successor labeled (h, h', σ'_A, R) .*

3. Each position (h, h', σ_A, L) belongs to I-DUPL and both I-players observe σ_A . I-DUPL may choose, for each $\sigma'_A \in PStr(C'_A(h'))$, a transition to a successor labeled $(h, h', \sigma_A, \sigma'_A, L)$.
4. Each position (h, h', σ'_A, R) belongs to I-DUPL and both I-players observe σ'_A . I-DUPL may choose, for each $\sigma_A \in PStr(C_A(h))$, a transition to a successor labeled $(h, h', \sigma_A, \sigma'_A, R)$.
5. Each position $(h, h', \sigma_A, \sigma'_A, L)$ belongs to P-SPOIL and P-SPOIL may choose, for each $k' \in C'_A(h')$, a transition to a successor labeled $(h, h', \sigma_A, \sigma'_A, k', L)$. In all positions at points 5-12, both I-players observe $C_A(h) \times C'_A(h')$.
6. Each position $(h, h', \sigma_A, \sigma'_A, R)$ belongs to P-SPOIL, and P-SPOIL may choose, for each $k \in C_A(h)$, a transition to a successor labeled $(h, h', \sigma_A, \sigma'_A, k, R)$.
7. Each position $(h, h', \sigma_A, \sigma'_A, k', L)$ belongs to P-DUPL, and P-DUPL may choose, for each $k \in C_A(h)$, a transition to a successor labeled $(h, h', \sigma_A, \sigma'_A, k, k', L)$.
8. Each position $(h, h', \sigma_A, \sigma'_A, k, R)$ belongs to P-DUPL, and P-DUPL may choose, for each $k' \in C'_A(h')$, a transition to a successor labeled $(h, h', \sigma_A, \sigma'_A, k, k', R)$.
9. Each position $(h, h', \sigma_A, \sigma'_A, k, k', L)$ belongs to P-SPOIL, This position is **winning** for the SPOILER coalition if $\pi(k) \neq \pi'(k')$ or $(k, k') \notin R$. In this position P-SPOIL may choose, for each $l' \in succ(k', \sigma'_A)$, a transition to a successor labeled $(h, h', \sigma_A, \sigma'_A, k, k', l', L)$.
10. Each position $(h, h', \sigma_A, \sigma'_A, k, k', R)$ belongs to P-SPOIL, This position is **winning** for the SPOILER coalition if $\pi(k) \neq \pi'(k')$ or $(k, k') \notin R$. In this position P-SPOIL may choose, for each $l \in succ(k, \sigma_A)$, a transition to a successor labeled $(h, h', \sigma_A, \sigma'_A, k, k', l, R)$.
11. Each position $(h, h', \sigma_A, \sigma'_A, k, k', l', L)$ belongs to P-DUPL, and P-DUPL may choose, for each $l \in succ(k, \sigma_A)$, a transition to a successor labeled (l, l') from where Rule 1 above applies.
12. Each position $(h, h', \sigma_A, \sigma'_A, k, k', l, R)$ belongs to P-DUPL, and P-DUPL may choose, for each $l' \in succ(k', \sigma'_A)$, a transition to a successor labeled (l, l') from where Rule 1 above applies.

In the sequel, given a position $p \in T(\mathcal{B})$, we denote $Obs(p)$ the set of positions which give the same observation as p to any of the I-players. Also, the set of strategies for the DUPLICATOR (resp. SPOILER) coalition is denoted Σ_{Dupl} (resp. Σ_{Spoil}). Further, the set of positions which are compatible with a strategy $\sigma \in \Sigma_{Dupl} \cup \Sigma_{Spoil}$ is denoted $Comp(\sigma)$. Finally, for each position p we denote $lab(p)$ its label, as per Def. 7 of bisimulation game.

Next, we prove that bisimulation relations and bisimulation games are equivalent characterisations of iCGS. To this end, given a history $h_0 \in Hist(\mathcal{G})$, we define the *pointed* iCGS $\mathcal{G}(h_0)$ in which the initial state is h_0 and the transitions are modified accordingly.

Theorem 2. *For any A-bisimulation relation R between $\mathcal{G}(h_0)$ and $\mathcal{G}'(h_0)$ the DUPLICATOR coalition has a strategy to win the bisimulation game $\mathcal{B}(\mathcal{G}, \mathcal{G}', R, h_0, h'_0)$.*

Conversely, if the DUPLICATOR coalition has a joint strategy σ_D to win the game $\mathcal{B}(\mathcal{G}, \mathcal{G}', R, h_0, h'_0)$, then there exists an A -bisimulation \Leftrightarrow_A with $\Leftrightarrow_A \subseteq R \cap \{(h, h') \mid (h, h') \in \text{out}(p_{h_0, h'_0}, \sigma_D)\}$, where p_{h_0, h'_0} is the initial position of the bisimulation game $\mathcal{B}(\mathcal{G}, \mathcal{G}', R, h_0, h'_0)$.

Proof. We prove this theorem by double inclusion.

\Rightarrow Suppose that \Leftrightarrow_A is an A -bisimulation. For convenience, we utilize, as in the proof of Lemma 1, the restated variant ($\hat{3}$) of point (3) in Def. 6 of A -simulations, which assumes a mapping $\rho_{\sigma_A, k, k'} : \text{succ}(k', ST(\sigma_A)) \rightarrow \text{succ}(k, \sigma_A)$ that ensures that for any $l' \in \text{succ}(k', ST(\sigma_A))$, we have $\rho_{\sigma_A, k, k'}(l') \Leftrightarrow_A l'$. Since \Leftrightarrow_A is also a reverse simulation, we symmetrically consider $\rho'_{\sigma'_A, k, k'} : \text{succ}(k, ST'(\sigma'_A)) \rightarrow \text{succ}(k', \sigma'_A)$ s.t. $\rho'_{\sigma'_A, k, k'}(l') \Leftrightarrow_A l$ for any $l \in \text{succ}(k, ST'(\sigma'_A))$.

Similarly, we restate point (2) in Def. 6 in functional terms:

$\hat{2}$ For for each $\sigma_A \in PStr(C_A(h))$ there exists a mapping $\theta_{\sigma_A}^{\leftarrow} : C'_A(h') \rightarrow C_A(h)$ such that for any $i \in A$, if $k'_1 \sim_i k'_2$ then $\theta_{\sigma_A}^{\leftarrow}(k'_1) \sim_i \theta_{\sigma_A}^{\leftarrow}(k'_2)$.

To see that this formulation is equivalent to item (2) in Def. 6, note first that this point restates as the first-order formula $\varphi = \forall k' \in \text{Hist}(\mathcal{G}') \exists k \in \text{Hist}(\mathcal{G}) (h' \sim_i k' \rightarrow h \sim_i k \wedge k \Rightarrow_A k')$. This formula is equivalent to $\forall \sigma_A. (\varphi \wedge \sigma_A \in PStr(C_A(h)))$ by the Universal Generalization Rule since σ_A is not free in φ . Then $\theta_{\sigma_A}^{\leftarrow} : C'_A(h') \rightarrow C_A(h)$ corresponds to the *skolemization* of $\exists k \in \text{Hist}(\mathcal{G})$ (seen as a unary function indexed by σ_A).

By symmetry, for each $\sigma'_A \in PStr(C'_A(h'))$ we denote $\theta_{\sigma'_A}^{\rightarrow} : C_A(h) \rightarrow C'_A(h')$ the reverse mapping, which exists since \Leftrightarrow_A is also a (reverse) simulation between \mathcal{G}' and \mathcal{G} .

Then, we define the strategy profile $(\sigma_{ID}, \sigma_{PD})$ for the DUPLICATOR coalition as follows: for any position p ,

1. If $\text{lab}(p) = (h, h', \sigma_A, L)$ then $\sigma_{ID}(p) = ST(\sigma_A)$, and if $\text{lab}(p) = (h, h', \sigma'_A, R)$ then $\sigma_{ID}(p) = ST'(\sigma'_A)$.
2. If $\text{lab}(p) = (h, h', \sigma_A, \sigma'_A, k', L)$ then $\sigma_{PD}(p) = \theta_{\sigma_A}^{\leftarrow}(k')$, and if $\text{lab}(p) = (h, h', \sigma_A, \sigma'_A, k, R)$ then $\sigma_{PD}(p) = \theta_{\sigma'_A}^{\rightarrow}(k)$.
3. If $\text{lab}(p) = (h, h', \sigma_A, \sigma'_A, k, k', l', L)$ then $\sigma_{PD}(p) = \rho_{\sigma_A, k, k'}(l')$ and if $\text{lab}(p) = (h, h', \sigma_A, \sigma'_A, k, k', l, R)$ then $\sigma_{PD}(p) = \rho'_{\sigma'_A, k, k'}(l)$.

Since \Leftrightarrow_A is an A -bisimulation and ST, ST' are strategy simulators that do not depend on h or h' , strategy σ_{ID} is uniform, that is, for all positions p, p' belonging to I-DUPL and this player receives the same sequence of observations along the history that leads to p and the history that leads to p' , we must have $\sigma_{ID}(p) = \sigma_{ID}(p')$. Then, all the runs that are compatible with the strategy profile σ_D never reach a position (h, h') where SPOILER wins:

- a. For $\text{lab}(p) = (h, h', \sigma_A, \sigma'_A, k', L)$, $\text{lab}(\text{succ}(p, \sigma_D)) = (h, h', \sigma_A, \sigma'_A, \theta_{\sigma_A}^{\leftarrow}(k'), k', L)$. But $\theta_{\sigma_A}^{\leftarrow}(k') \Rightarrow_A k$ by point ($\hat{2}$) for Def. 6, which implies that $\text{succ}(p, \sigma_D)$

is not winning for SPOILER. A similar argument holds for $lab(p) = (h, h', \sigma_A, \sigma'_A, k, R)$.

- b. For $lab(p) = (h, h', \sigma_A, \sigma'_A, k, k', l', L)$, $lab(succ(p, \sigma_D)) = (h, h', \sigma_A, \sigma'_A, k, k', \rho_{\sigma_A, k, k'}(l'), l', L)$. But $\rho_{\sigma_A, k, k'}(l')$ $\Rightarrow_A l'$ by point (3) for Def. 6, which implies that $succ(p, \sigma_D)$ is not winning for SPOILER. A similar argument holds for $lab(p) = (h, h', \sigma_A, \sigma'_A, k, k', l, R)$.

\Leftarrow Suppose now that we have a winning joint strategy $\sigma_D = (\sigma_{ID}, \sigma_{PD})$ for the DUPLICATOR coalition. Then, for each position p that is consistent with σ_D , with label $lab(p) = (h, h') \in Hist(\mathcal{G}) \times Hist(\mathcal{G}')$, we set $h \Leftrightarrow_A^{\sigma_D} h'$.

The strategy simulators are then defined as follows: for each $h \Leftrightarrow_A^{\sigma_D} h'$ with $(h, h') = lab(p)$ for some position p in the bisimulation game, and each $\sigma_A \in PStr(C_A(h))$, note first that we have a I-SPOIL transition to a position $p_{\sigma_A}^1$ labeled (h, h', σ_A, L) and then a I-DUPL transition to a position $p_{\sigma_A}^2$ labeled $(h, h', \sigma_A, \sigma_{ID}(p_{\sigma_A}^1), L)$. Then, we set $ST_{C_A(h), C'_A(h')}(\sigma_A) = \sigma_{ID}(p_{\sigma_A}^1)$. Note that this definition is independent of the choice of p since, by construction, all positions \bar{p} with $lab(\bar{p}) = lab(p)$ are indistinguishable for I-DUPL, as he observes only $C_A(h) \times C'_A(h')$ and σ_A . Hence, $\sigma_{ID}(p_{\sigma_A}^1) = \sigma_{ID}(\bar{p}_{\sigma_A}^1)$, where $\bar{p}_{\sigma_A}^1$ is the position resulting by I-SPOIL choosing σ_A in position \bar{p} . This ensures that ST is indeed a strategy simulator according to Equation (1).

Furthermore, the mappings $\theta_{\sigma_A}^{\leftarrow}$ are defined as follows: given position p with $lab(p) = (h, h')$ as above, then for each $\sigma_A \in PStr(C_A(h))$, denote first $p_{\sigma_A}^1$ the position resulting from I-SPOIL executing transition σ_A ; further denote $p_{\sigma_A}^2$ the position which belongs to P-SPOIL after P-DUPL executes action $\sigma_D(p_{\sigma_A}^1)$. Note then that, in position $p_{\sigma_A}^2$, for each $k' \in C'_A(h')$, P-SPOIL has a move to a position $p_{\sigma_A, k'}^3$ which is labeled $(h, h', \sigma_A, \sigma_{ID}(p_{\sigma_A}^1), k', L)$ which belongs to P-DUPL. Then we define $\theta_{\sigma_A}^{\leftarrow}(k') = \sigma_{PD}(p_{\sigma_A, k'}^3)$.

This definition is dependent on the choice of the starting position p , but this is not an issue for our definition of $\Leftrightarrow_A^{\sigma_D}$ since there is no requirement for building the *maximal* bisimulation associated with a bisimulation game. Note further that this definition, together with the fact that σ_D is winning and hence position $p_{\sigma_A, k'}^4 = succ(p_{\sigma_A, k'}^3, \sigma_{PD}(p_{\sigma_A, k'}^3))$ is not winning for the SPOILER coalition, implies that $\pi(\theta_{\sigma_A}^{\leftarrow}(k')) = \pi'(k')$ and further ensures that $\theta_{\sigma_A}^{\leftarrow}(k')$ satisfies the restated point (2) for Def. 6.

Finally, by proceeding from position $p_{\sigma_A, k'}^4$, which again belongs to P-SPOIL, for each P-SPOIL's choice of some history $l' \in succ(k', \sigma_{ID}(p_{\sigma_A}^1))$, the game proceeds to a position $p_{\sigma_A, k', l'}^5$ that belongs to P-DUPL and is labeled with $lab(p_{\sigma_A, k', l'}^5) = (h, h', \sigma_A, \sigma_{ID}(p_{\sigma_A}^1), \sigma_{PD}(p_{\sigma_A, k'}^3), k', l', L)$. We then define $\rho_{\sigma_A, \theta_{\sigma_A}^{\leftarrow}(k'), k'}(l') = \sigma_{PD}(p_{\sigma_A, k', l'}^5)$. Also note that this definition is dependent on the choice of the initial position p with no loss of generality, and the fact that σ_D is winning ensures that the position resulting from $p_{\sigma_A, k', l'}^5$ by P-DUPL's choice and labeled $(\sigma_{PD}(p_{\sigma_A, k', l'}^5), l')$, is not winning for SPOILER. In particular, $\pi(p_{\sigma_A, k', l'}^5) = \pi'(l')$ and $\rho_{\sigma_A, \theta_{\sigma_A}^{\leftarrow}(k'), k'}$ satisfies the restated point (3) for Def. 6.

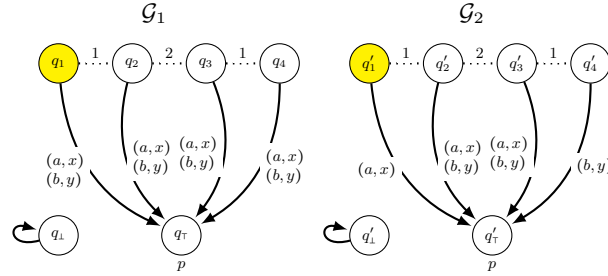


Fig. 2. Counterexample for the Hennessy-Milner property for the subjective and objective semantics.

Similar considerations give us the definitions for $\theta_{\sigma'_A}^{\rightarrow}$ and $\rho'_{\sigma'_A, k, k'}$ for each $\sigma'_A \in PStr(C'_A(h'))$, $k' \in C'_A(h')$ and $k \in C_A(h)$ with $k \iff_A k'$. This completes the proof of Theorem 2.

We conclude this section with some immediate properties about our bisimulation games and bisimulation relations.

- Proposition 1.**
1. *The set of bisimulations associated with the same strategy simulator forms a complete lattice w.r.t. set inclusion.*
 2. *If two iCGS \mathcal{G} and \mathcal{G}' are bisimilar, then the DUPLICATOR coalition has a winning strategy in the bisimulation game $\mathcal{B}(\mathcal{G}, \mathcal{G}', Tot)$ where Tot is the total relation $Hist(\mathcal{G}) \times Hist(\mathcal{G}')$.*

The second claim follows by observing that if the DUPLICATOR coalition has a strategy to win a bisimulation game $\mathcal{B}(R) = \mathcal{B}(\mathcal{G}, \mathcal{G}', h_0, h'_0, R)$ for some R , then they also have a strategy to win the bisimulation game $\mathcal{B}(Tot) = \mathcal{B}(\mathcal{G}, \mathcal{G}', h_0, h'_0, Tot)$, and the construction in Theorem 2 can be used to show that the bisimulation associated with $\mathcal{B}(R)$ is included in the bisimulation associated with $\mathcal{B}(Tot)$. Hence, this latter is maximal w.r.t. all bisimulations that share the strategy simulator constructed as in Theorem 2.

5 The Hennessy-Milner Property

We now show that the notion of bisimulation introduced in Sec. 3 enjoys the Hennessy-Milner property. To this end, we need to define A -equivalence between iCGS. Specifically, given iCGS \mathcal{G} and \mathcal{G}' with histories $h_0 \in Hist(\mathcal{G})$ and $h'_0 \in Hist(\mathcal{G}')$, we say that the pointed iCGS $\mathcal{G}(h_0)$ and $\mathcal{G}'(h'_0)$, having h_0 and h'_0 as respective initial histories, are A -equivalent iff for every A -formula ϕ , $(\mathcal{G}, h_0) \vDash_x \phi$ iff $(\mathcal{G}', h'_0) \vDash_x \phi$.

Theorem 3. *The notion of bisimulation in Def. 6 enjoys the Hennessy-Milner property, that is, the pointed iCGS $\mathcal{G}(h_0)$ and $\mathcal{G}'(h'_0)$ are A -equivalent for the common knowledge semantics if and only if they are A -bisimilar.*

Before proving Theorem 3, as a counterexample for the subjective and objective semantics, we recall the example used in [6] depicted in Figure 2. In each

state agent 1 can execute actions $\{a, b, c\}$ while agent 2 can execute $\{x, y, z\}$. The transitions shown lead to q_\top and q'_\top , while the omitted transitions lead to q_\perp and q'_\perp , respectively. We can check that states q_i and q'_j , with $i, j \in \{1, 2, 3, 4\}$, are $\{1, 2\}$ -equivalent, and it holds the same also for states q_\perp (resp., q_\top) and q'_\perp (resp., q'_\top). Therefore, \mathcal{G}_1 and \mathcal{G}_2 are $\{1, 2\}$ -equivalent, i.e., they satisfy the same $\{1, 2\}$ -formulas in ATL^* . However, there is no $\{1, 2\}$ -bisimulation between the two iCGS. In particular, for any $i, j \in \{1, 2, 3, 4\}$, state q_i cannot be $\{1, 2\}$ -bisimilar with any state q'_j .

The lack of a bisimulation between the two structures in Fig. 2 follows by observing that the SPOILER coalition wins the appropriate bisimulation game, since in the initial position (q_1, q'_1) , I-SPOIL may choose strategy σ which produces action tuple (a, x) in each state in $C_{\{1, 2\}}(q_1) = \{q_1, q_2, q_3, q_4\}$. If I-DUPL responds with strategy $\sigma'(q'_1) = \sigma'(q'_2) = \sigma'(q'_3) = \sigma'(q_4) = (a, x)$, then P-SPOIL will choose state q'_4 , and P-DUPL has no good choice of some state in \mathcal{G}_1 whose successor by σ is labeled $\neg p$, which is the label of the successor of q'_4 by σ' . Similar situations occur for all the other choices by I-DUPL.

To prove Theorem 3, we actually prove the following Gale-Stewart-type theorem for the bisimulation games introduced in Def. 7:

Theorem 4 (Gale-Stewart theorem for bisimulation games). *Each bisimulation game is **determined**: either the DUPLICATOR coalition or the SPOILER coalition wins the game.*

Proof. We follow the pattern of Gale-Stewart games by proving first that, at positions where one coalition does not have a winning strategy, the other one has a “defensive” strategy, and then showing that any defensive strategy is winning.

Formally, a **defensive** strategy for the SPOILER coalition in the bisimulation game is a joint strategy $\sigma_S = (\sigma_{IS}, \sigma_{PS})$ such that, for any position p of the game which is compatible with σ_D , DUPLICATOR does not have a winning strategy starting from the set of positions that have the same observability as p and are compatible with σ_S . Defensive strategies for DUPLICATOR are defined similarly.

We then have:

Lemma 2. *If DUPLICATOR coalition does not have a winning strategy, then SPOILER coalition has a defensive strategy.*

The proof of this claim works similarly to the classical case [14], by building the defensive strategy by induction on the level of the position in the tree of positions of the bisimulation game, such that, at each position belonging to P-SPOIL or I-SPOIL, we identify one “defensive” action for these agent. The difficulty is to build a uniform strategy for I-SPOIL, i.e., at two positions with identical observations for I-SPOIL, her actions are identical.

The argument is similar for the base case and the inductive cases, and starts by assuming the following: for some position p with $\text{lab}(p) = (\bar{h}, \bar{h}')$, the SPOILER coalition does not have a defensive strategy from $\text{Obs}(p)$, but neither DUPLICATOR has a winning strategy from $\text{Obs}(p)$. Therefore the following property,

which formalizes the lack of a defensive strategy w.r.t. L-transitions (i.e. steps 2.a, 3, 5, 7, 9 and 11) in the bisimulation game, holds:

$$\begin{aligned}
& \forall p_1 \in Obs(p) \quad \forall \sigma \in PStr(C_A(lab(p_1))) \\
& \quad \exists \sigma'_{\sigma, p_1} \in PStr(C'_A(lab_2(p_1))) \quad \forall k' \in C'_A(lab_2(p_1)) \\
& \quad \quad \exists k_{\sigma, p_1, k'} \in C_A(lab_1(p_1)) \quad \forall l' \in succ(k', \sigma'_{\sigma, p_1}) \\
& \quad \exists l_{\sigma, p_1, k', l'} \in succ(k_{\sigma, p_1, k'}, \sigma) \quad \exists \sigma_D = (\sigma_{ID}, \sigma_{PD}) \in \Sigma_{Dupl} \\
& \quad \quad \text{with } \sigma_D \text{ winning from position } (l_{\sigma, p_1, k', l'}, l') \quad (3)
\end{aligned}$$

Note that a similar property holds w.r.t. R-transitions.

Notice that, by default, nothing excludes having some $p_1, p_2 \in Obs(p)$ such that, for some, $\sigma \in PStr(C_A(lab_1(p)))$, $\sigma'_{\sigma, p_1} \neq \sigma'_{\sigma, p_2}$. But I-DUPL can choose the same $\sigma'_{\sigma, p}$ and P-DUPL can then choose $l_{\sigma, p, k', l'}$ at **all positions** $p_1 \in Obs(p)$, because choices of k' for P-DUPL in Formula 3 are quantified over the whole $C'_A(lab_2(p_1)) = C'_A(lab_2(p))$. In other words, because I-DUPL's choice of $\sigma'_{\sigma, p}$ combined with P-DUPL choice of $l_{\sigma, p, k', l'}$ are "defensive" at position p , they are "defensive" at any other position $p_1 \in Obs(p)$. This way, I-DUPL's choice can be made **uniform** w.r.t. her observations, which gives a winning strategy for the DUPLICATOR coalition at p , fact which contradicts the initial assumption.

Formally, from any position $p_1 \in Obs(p)$, the following strategy for DUPLICATOR coalition is winning:

- Denote p_σ^1 the successor of p_1 after I-SPOIL chooses σ . Then I-DUPL chooses $\sigma'_{\sigma, p}$ at p_σ^1 .
- Denote the resulting position $p_{\sigma, \sigma'_{\sigma, p}}^1$. Denote further by $p_{\sigma, \sigma'_{\sigma, p}, k'}^1$ the successor of $p_{\sigma, \sigma'_{\sigma, p}}^1$ after P-SPOIL has chosen $k' \in C'_A(lab_2(p))$. Then P-DUPL chooses $k_{\sigma, p, k'} \in C_A(lab_1(p))$ at $p_{\sigma, \sigma'_{\sigma, p}, k'}^1$.
- Denote the resulting position $p_{\sigma, \sigma'_{\sigma, p}, k_{\sigma, p, k'}}^1$. Also denote $p_{\sigma, \sigma'_{\sigma, p}, k_{\sigma, p, k'}, l'}^1$ the successor of $p_{\sigma, \sigma'_{\sigma, p}, k_{\sigma, p, k'}}^1$ after P-SPOIL chooses $l' \in succ(k', \sigma'_{\sigma, p})$. Then P-DUPL chooses $l_{\sigma, p, k', l'} \in succ(k_{\sigma, p, k'}, \sigma)$ at $p_{\sigma, \sigma'_{\sigma, p}, k_{\sigma, p, k'}, l'}^1$.

Formula 3 implies that from the resulting position, which is labeled $(l_{\sigma, p, k', l'}, l')$, the DUPLICATOR coalition has a winning strategy. Hence, we have a winning strategy for the DUPLICATOR coalition from $Obs(p)$, which contradicts the initial assumption.

As a result, the SPOILER coalition must have a defensive strategy from $Obs(p)$, which can be built by negating Formula 3, after which the construction continues by induction on the observation class of the resulting positions labeled $(l_{\sigma, p, k', l'}, l')$. A similar argument shows that, when SPOILER coalition does not have a winning strategy, DUPLICATOR coalition has a defensive strategy.

It remains to show that a defensive strategy for DUPLICATOR is winning. This follows by observing that any infinite path in $T(\mathcal{B})$ which is compatible with a defensive strategy for DUPLICATOR must not pass through a position which is winning for SPOILER, hence is an infinite path which is winning for DUPLICATOR, which ends the proof.

We can now proceed with the proof of Theorem 3.

Proof. Theorem 3 Assume that there exists no bisimulation between \mathcal{G} and \mathcal{G}' which, by Proposition 1, means that in the bisimulation game $\mathcal{B}(\mathcal{G}, \mathcal{G}', s_0, s'_0, Tot)$ the DUPLICATOR coalition has no winning strategy. By the determinacy theorem, the SPOILER coalition has a winning strategy $\sigma_S = (\sigma_{IS}, \sigma_{PS})$. Since each position in the bisimulation game has a finite number of successors, as a consequence of König's Lemma, there exists a finite set P_{σ_S} of winning positions for SPOILER such that all runs compatible with σ_S pass through one position of P_{σ_S} .

Pick then a position p labeled (h, h') such that, on all runs starting from p and compatible with σ_S , the first position labeled with some (l, l') occurring on the run after p is a winning position for SPOILER. Note that the following property, formalizing the fact that σ_S is winning, holds:

$$\begin{aligned} \exists \sigma \in PStr(C_A(h)) \forall \sigma' \in PStr(C'_A(h')) \exists k'_{\sigma'} \in C'_A(h') \forall k \in C_A(h) \\ (\pi(k) = \pi'(k'_{\sigma'}) \rightarrow \exists l'_{\sigma', k'_{\sigma'}, k} \in succ(k', \sigma') \\ \forall l \in succ(k, \sigma) \pi(l) \neq \pi(l'_{\sigma', k'_{\sigma'}, k})) \end{aligned} \quad (4)$$

where $\sigma = \sigma_{IS}(p)$, $k'_{\sigma'} = \sigma_{PS}(p_1)$, p_1 is the successor of p after I-SPOIL chooses σ and P-DUPL answers with σ' , and $l'_{\sigma', k'_{\sigma'}, k} = \sigma_{PS}(p_2)$ where p_2 is the successor of p_1 after P-SPOIL chooses k' and P-DUPL answers with k . Note that the implication with premise $\pi(k) = \pi'(k'_{\sigma'})$ is needed since P-DUPL's choices with $\pi(k) \neq \pi'(k'_{\sigma'})$ are immediately winning for SPOILER, and then there is no need to proceed with steps 11-12 corresponding with the successors of k and k' .

So, if we define the formula

$$\varphi(P_{\sigma_S}) = \langle\langle A \rangle\rangle X \left(\bigwedge_{\sigma' \in PStr(C'_A(h'))} (\mathbf{Y} \pi'(k'_{\sigma'}) \rightarrow \bigvee_{k \in C_A(h)} \neg \pi'(l'_{\sigma', k'_{\sigma'}, k})) \right) \quad (5)$$

then Formula 4 implies that $(\mathcal{G}, h) \models_{ck} \varphi(P_{\sigma_S})$ but, on the other hand, $(\mathcal{G}', h') \not\models_{ck} \varphi(P_{\sigma_S})$.

To see this, note that, in $\varphi(P_{\sigma_S})$ the coalition operator $\langle\langle A \rangle\rangle$ encodes $\exists \sigma \in PStr(C_A(h))$ in 4. Further, the conjunction indexed by σ' in 5 corresponds to the universal quantifier on σ' in 4. The $k'_{\sigma'}$ in 5 represents the skolemization of the existential quantification over $k'_{\sigma'}$ in 4. The last disjunction in 5 corresponds with the existential quantification over k in 4. The existential quantification over $l'_{\sigma', k'_{\sigma'}, k}$ in 4 is encoded in 5 by its skolemization, denoted $l'_{\sigma', k'_{\sigma'}, k}$ too. Finally, the universal quantifier over $l \in succ(k, \sigma)$ in 4 and the last property connecting l to $l'_{\sigma', k'_{\sigma'}, k}$ is encoded in 5 by $\neg \pi'(l'_{\sigma', k'_{\sigma'}, k})$.

The yesterday operator Y is needed because we must encode the part of 4 referring to $\pi(k)$, which refers to the current position. Unfortunately, a formula like $\langle\langle A \rangle\rangle (\pi(k) \rightarrow X \psi)$ which would simulate more easily the implication $\pi(k) = \pi'(k'_{\sigma'}) \rightarrow \exists l'_{\sigma', k'_{\sigma'}, k} \in succ(k', \sigma') \dots$ from 4 would not be ATL but rather ATL*. But, in order to correctly simulate quantifier order from 4, in 5 $\pi(k)$ must lie

within the scope of $\langle\langle A \rangle\rangle X$, which refers to the positions one time step after the current position. Hence, in the scope of $\langle\langle A \rangle\rangle X$ we need to recover the value of $\pi(k)$ at the previous position, hence we utilize Y . We believe Y might not be needed for the full ATL^* , a topic for further research.

The proof can then be completed by induction as follows: we modify the bisimulation game by appending a new winning condition for SPOILER: all positions in $Obs(p)$ are labeled as winning, with the formula φ witnessing this. The set of atomic propositions for both iCGS is augmented with p_φ and, for each (h, h') labeling a position $p_1 \in Obs(p)$, we augment $\pi(h)$ with p_φ while $\pi'(h')$ is left unchanged. This provides us with a new bisimulation game in which (the appropriately updated) strategy profile σ_S is still a winning strategy for SPOILER, there is a strictly smaller set of positions P'_{σ_S} which are winning for SPOILER, and all runs compatible with σ_S pass through one position of P'_{σ_S} .

The argument ends when we obtain some $P_{\sigma_S}^m$ for which $Obs(P_{\sigma_S}^m)$ is a singleton, which means that $(h_0, h'_0) \in P_{\sigma_S}^m$. Then the formula $\varphi(P_{\sigma_S}^m)$ built as in Equation 5 is the witness that (\mathcal{G}, h_0) is not A -equivalent with (\mathcal{G}', h'_0) .

6 Undecidability Result

In this section we show that deciding the existence of bisimulations between iCGS is undecidable in general. We state immediately the main result of this section.

Theorem 5. *The problem of checking whether two CGS \mathcal{G}_1 and \mathcal{G}_2 defined on the same set Ag of agents are A -bisimilar, for some set $A \subseteq Ag$ of agents, is undecidable.*

The proof of Theorem 5 can be outlined as follows: given any deterministic Turing Machine, by building on [12] we construct a 3-agent iCGS which has the property that two agents (call them 1 and 2) have a winning strategy to avoid an error state if and only if the TM never stops when starting with an empty tape. We note that this strategy, when it exists, is unique. Then, we construct a second 3-agent “simple” iCGS in which there exists a unique strategy for agents 1 and 2 (without memory) for avoiding an error state. Finally, we prove that these two iCGS are bisimilar if and only if the TM never stops when starting with an empty tape, which is sufficient to derive the undecidability of the former problem.

We start with the construction of the second iCGS, which is depicted in Figure 3. Note that the transitions only represent the actions of agents 1 and 2, agent 3’s role is to “solve nondeterminism” in states s_{init} , s_{gen} and s_{tr} . First of all, we prove the following lemma.

Lemma 3. *In the iCGS depicted in Figure 3 there exists a unique strategy for agents 1 and 2 to avoid state s_{err} .*

Proof. Note that, in all the states except s_{amb}^1 and s_{amb}^2 , coalition $\{1, 2\}$ must play (ok, ok) to avoid s_{err} . To further understand why 1 and 2 need to play

the same action in the remaining two states, consider history $h = s_{init} \xrightarrow{ok,ok} s_{gen} \xrightarrow{ok,ok} s_{tr} \xrightarrow{ok,ok} s_{amb}^2$. Note that if we have $h \sim_1 h'$ (and $h \neq h'$), then $h' = s_{init} \xrightarrow{ok,ok} s_{gen} \xrightarrow{ok,ok} s_{amb}^1 \xrightarrow{a_1,a_2} s_{namb}^1$. So, for any joint strategy $\sigma = (\sigma_1, \sigma_2)$ for 1 and 2, ensuring $s_{err} \notin succ(h', \sigma_1(h'))$ requires that $\sigma_1(h') = ok$, which also implies that $\sigma_1(h) = ok$ by 1-uniformity of σ . A similar argument shows that $h \sim_2 h''$ (and $h \neq h''$) implies that $h'' = s_{init} \xrightarrow{ok,ok} s_{gen} \xrightarrow{ok,ok} s_{tr} \xrightarrow{ok,ok} s_{gen}$, and since the only way to ensure $s_{err} \notin succ(h'', \sigma_2(h''))$ is by choosing $\sigma_2(h'') = ok$, we must also have $\sigma_2(h) = ok$.

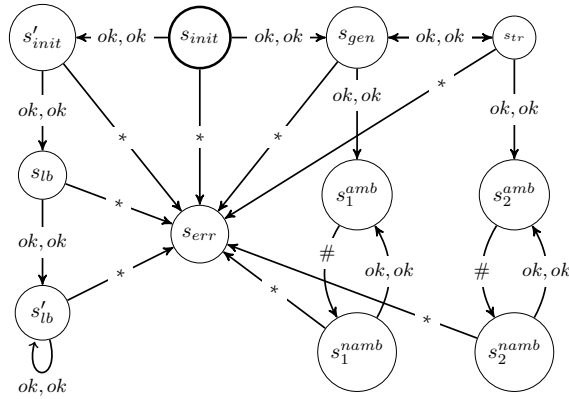


Fig. 3. The *iCGS* \mathcal{G}_S , where $*$ = $\neg ok, ok \mid ok, \neg ok \mid \neg ok, \neg ok$ and $\#$ = $\neg ok, ok \mid ok, \neg ok \mid \neg ok, \neg ok \mid ok, ok$. The indistinguishability relation for player 1 has three classes: s_{gen} , s_{err} , and s_{set} , where $s_{set} = S \setminus \{s_{gen}, s_{err}\}$. The indistinguishability relation for player 2 has three classes: s_{tr} , s_{err} , and s'_{set} , where $s'_{set} = S \setminus \{s_{tr}, s_{err}\}$.

By generalizing these observations, a strategy that avoids the error state for all outcomes can only be constructed if in every history h ending in s_{amb}^1 the joint action is (ok, ok) . This is because:

1. if $h \sim_1 h'$ and $last(h') = s_{namb}^2$, then in s_{namb}^2 the only “good” transition for agent 1 is ok , and
2. if $h \sim_2 h''$ and $last(h'') \in \{s_{gen}, s_{namb}^1\}$, then in both these states the only “good” transition for agent 2 is ok too.

A similar remark holds for histories ending in the other “ambiguous” state, s_{amb}^2 . So there is only one joint uniform strategy for $\{1, 2\}$ that avoids s_{err} , which is choosing (ok, ok) at every history.

We now turn to the construction of the first iCGS, which is adapted from [12]. We give the construction for a simple deterministic *TM* $M = \langle \mathcal{Q}, \Gamma, \delta \rangle$ with states $\mathcal{Q} = \{q_0, q_1, q_2\}$ and tape symbols $\Gamma = \{B, a\}$, whose transition function δ is in Table 1.

$\rho^{(1,n+1)}[2k+2, 2k+4] = s_z^1 \xrightarrow{i,i} s_z^0 \xrightarrow{i,(q,r,R)} s_{r,z}$, for some $z \in \Gamma$ representing the contents of tape cell $(n+1)$ in configuration k .

(b) On the other hand, if the transition table gives $\delta(q, x) = (r, y, L)$ (i.e. an L-move of the head), then

$$\rho^{(1,n)}[2k+2, 2k+4] = s_{q,x} \xrightarrow{i,(q,r,L)} s_y^0 \xrightarrow{i,i} s_y^1$$

$$\rho^{(1,n-1)}[2k+2, 2k+4] = s_z^1 \xrightarrow{i,i} s_z^0 \xrightarrow{(q,r,L),i} s_{r,z}$$
, for some $z \in \Gamma$.

(c) Otherwise (i.e., the R/W head is not pointing cells $n-1$ to $n+1$),

$$\rho^{(1,n)}[2k, 2k+2] = s_z^1 \xrightarrow{i,i} s_z^0 \xrightarrow{i,i} s_z^1$$
 where z is the contents of cell n in configuration k .

Note that two steps are needed along each run to encode the transition of the R/W head from cell n to cell $(n+1)$ for an R-move, or to cell $(n-1)$ for a L-move.

2. Every run ρ starting with $s_{init} \xrightarrow{i,i} s_{gen} \xrightarrow{i,i} s_{tr} \xrightarrow{i,i} s_{tr}^1$, simulates a move of the R/W head between the $(n-1)$ -th and the n -th cell, (which we call in the sequel the n -th frontier), for $n \geq 1$. We call such runs as $(2, n)$ -runs and denote them $\rho^{(2,n)}$. Formally, for every $k \geq n$, depending on the transition between the k -th and the $(k+1)$ -th configuration of the TM:

(a) If the R/W head moves from the n -th cell to the $(n-1)$ -th by executing $\delta(q, x) = (r, y, L)$, then

$$\rho^{(2,n)}[2k+3, 2k+5] = s_{tr}^1 \xrightarrow{i,(q,r,L)} s_{q,r,L} \xrightarrow{(q,r,L),i} s_{tr}^1.$$

(b) If the R/W head moves from the $(n-1)$ -th cell to the n -th by executing transition $\delta(q, x) = (r, y, R)$, then

$$\rho^{(2,n)}[2k+3, 2k+5] = s_{tr}^1 \xrightarrow{(q,r,R),i} s_{q,r,R} \xrightarrow{i,(q,r,R)} s_{tr}^1.$$

(c) Otherwise, $\rho^{(2,n)}[2k+3, 2k+5] = s_{tr}^1 \xrightarrow{i,i} s_{tr}^0 \xrightarrow{i,i} s_{tr}^1$.

Finally, run $\rho_0 = s_{init} \xrightarrow{i,i} s'_{init} \xrightarrow{i,i} s_{ib} \xrightarrow{i,(q_0)} s'_{lb} \left(\xrightarrow{i,i} s'_{lb} \right)^\omega$ simulates the “left bound” of the tape and the start of the TM with the R/W head on the initial state q_0 . Let GR denote the set $\{\rho_0\} \cup \{\rho^{(1,n)}, \rho^{(2,n)} \mid n \in \mathbb{N}\}$ of “good” runs, and σ^{win} the unique strategy for agents 1 and 2 that simulates the infinite run of the TM, when it exists.

In what follows, we group states of \mathcal{G}_M in five sets:

1. $S_{namb}^0 = \{s_{init}, s'_{init}, s_{lb}, s'_{lb}\}$.
2. $S_{namb}^1 = \{s_{q,z} \mid q \in \mathcal{Q} \text{ and } z \in \Gamma\} \cup \{s_\perp^1, s_B^1, s_a^1\}$.
3. $S_{amb}^1 = \{s_a^0, s_B^0, s_\perp^2\}$.
4. $S_{namb}^2 = \{s_{tr}^1, s_\perp^3\}$.
5. $S_{amb}^2 = \{s_{q,r,x} \mid q, r \in \mathcal{Q}, x \in \{R, L\}\} \cup \{s_\perp^4, s_{tr}^0\}$.

We also call states in $S_{amb} = S_{amb}^1 \cup S_{amb}^2$ as “ambiguous” and in $S_{namb} = S \setminus S_{amb}$ “nonambiguous”. Note that, in each nonambiguous state, all outgoing transitions which avoid s_{err} are labeled with a unique tuple of actions, while all ambiguous states do not have transitions leading to s_{err} .

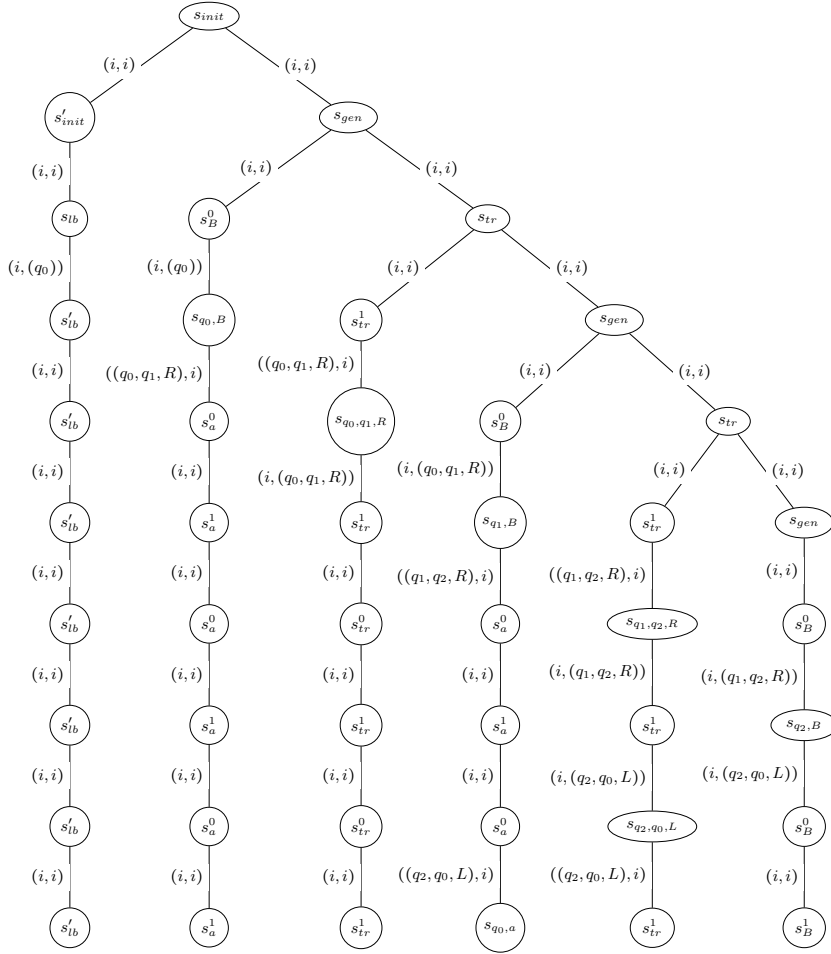


Fig. 5. Simulating three computation steps of the Turing machine in Table 1.

Before defining the actual bisimulation between \mathcal{G}_M and \mathcal{G}_S , note that runs that “deviate” from runs in GR (and therefore associated with a “wrong” strategy that cannot avoid s_{err}) contain either a transition from a nonambiguous state to s_{err} , or a transition from an ambiguous state which is not consistent with the TM computation, as explain below.

The line of reasoning guaranteeing that every strategy σ consistent with a run which “deviates” from a run in GR will not be able to avoid s_{err} is similar to the proof of Lemma 3 above. Assume that there exists a single partial strategy, defined on $Hist^{\leq i}(\mathcal{G}_M)$, which avoids s_{err} and all histories compatible with this strategy are prefixes of length $\leq n$ of runs in GR . Note first that each such history ending in states in S_{namb} can only be completed with a transition that simulates correctly the unique run up to level $i+1$.

To see what happens with the other type of histories, consider some history $h = \rho^{(1,n)}[\leq i] \xrightarrow{a_1, a_2} s$ with $h[i] \notin S_{namb}$. Note then that if $h \sim_1 h'$ then $h'[\leq i] \leq \rho^{(2,n)}$ and $h'[i] \in S_{namb}$ is a nonambiguous state. Therefore, there exists a unique $h'[i] \xrightarrow{a_1, b_2} s'$, and, moreover, this transition has to correctly simulate the n -th frontier at level $i+1$. It then follows that a_1 is the “good” decision agent 1 has to make to correctly simulate the TM run on h . A similar argument holds for $h'' \sim_2 h$ by noting that $h''[\leq i] \leq \rho^{(2,n-1)}$. Also similar arguments hold if we start with $h \leq \rho^{(2,n)}$.

Finally, the bisimulation relation between \mathcal{G}_M and \mathcal{G}_S is guided by the intuition that histories ending in nonambiguous states S_{namb}^1 resp. S_{namb}^2 , ”behave similarly” with histories ending in s_{namb}^1 , resp. s_{namb}^2 , while histories ending in ambiguous states S_{amb}^1 , resp. S_{namb}^2 behave similarly with histories ending in s_{amb}^1 , resp. s_{amb}^2 .

Specifically, for every $h \in Hist(\mathcal{G}_M)$ with $h < \rho^{(1,n)}$ and $h \not\leq \rho^{(2,n)}$, we set $h \iff_{1,2} \chi$ for χ defined as follows:

1. $\chi[i] = h[i]$ if $h[i] \in S_{namb}^0 \cup \{s_{err}\}$.
2. $\chi[i] = s_{amb}^1$ if $h[i] \in S_{amb}^1$.
3. $\chi[i] = s_{namb}^1$ if $h[i] \in S_{namb}^1$.
4. For $y \in \{1, 2\}$, $act_y(\chi, i) = ok$ if and only if $act_y(h, i)$ is the “correct” action executed by agent y for simulating the contents of the n -th cell at level i along $\rho^{(1,n)}$.

Similarly, for every $h \in Hist(\mathcal{G}_M)$ with $h < \rho^{(2,n)}$ and $h \not\leq \rho^{(1,n+1)}$, we set $h \iff_{1,2} \chi$ for χ defined as follows:

1. $\chi[i] = s_{amb}^2$ if $h[i] \in S_{amb}^2$.
2. $\chi[i] = s_{namb}^2$ if $h[i] \in S_{namb}^2$.
3. For $y \in \{1, 2\}$, $act_y(\chi, i) = ok$ if and only if $act_y(h, i)$ is the “correct” action executed by agent y for simulating the n -th frontier at level i along $\rho^{(2,n)}$.

Note that $\iff_{1,2}$ is in fact functional.

The strategy simulator ST can be constructed using the functional relation $\iff_{1,2}$ as follows: for any joint strategy σ and any history $h \in Hist(\mathcal{G}_M)$ compatible with σ , take the unique $\chi_h \in Hist(\mathcal{G}_S)$ with $h \iff_{1,2} \chi_h$ and define the partial strategy τ_σ with $\tau_\sigma(\chi_h[< |\chi_h|]) = act(\chi_h, |< \chi_h|)$. Then note that, whenever we have two different joint strategies σ^1 and σ^2 which share some compatible histories, then for any h compatible with both we have that $\tau^{\sigma^1}(\chi_h[< |\chi_h|]) = \tau^{\sigma^2}(\chi_h[< |\chi_h|])$. This means that the following definition correctly constructs a strategy simulator: for each $h \in Hist(\mathcal{G}_M)$, each partial strategy $\bar{\sigma} \in PStr(C_A(h))$, $h' \in C_A(h)$ and for each joint strategy σ with $\sigma|_{C_A(h)} = \bar{\sigma}$,

$$ST(\bar{\sigma})(h) = \tau^\sigma(\chi_h[< |\chi_h|])$$

because, as noted above, different τ^σ agree on the same χ_h , so the choice of σ is not important as long as it is compatible with h .

The inverse strategy simulator can be chosen as any inverse function of ST , i.e. any function ST' with $ST \circ ST'$ being the identity function.

Hence $\Leftrightarrow_{1,2}$ is indeed an $\{1,2\}$ -bisimulation with ST and ST' strategy simulators, if and only if M never stops when starting with an empty tape, which ends the proof of the undecidability theorem.

7 Conclusions and Future Work

In this paper we advanced the state of the art in the model theory of logics for strategic reasoning in multi-agent systems. Specifically, in Sec. 2 we considered the *common knowledge* interpretation of the Alternating-time Temporal Logic ATL under the assumption of imperfect information (and perfect recall), which has so far received little attention in the literature. For this context of imperfect information, we introduced a novel notion of alternating bisimulation in Sec. 3 and were able to prove the preservation of ATL formulas in bisimilar iCGS (Theorem 1). Further, in order to show that the common knowledge interpretation enjoys the Hennessy-Milner property, in Sec. 4 we introduced an imperfect information variant bisimulation games and showed that the DUPLICATOR coalition has a winning strategy if and only if there exists a bisimulation between the two given iCGS (Theorem 2). Finally, in Sec. 5 we proved the Gale-Stewart determinacy Theorem 4, which allows us to prove the Hennessy-Milner Theorem 3. We also provided counterexamples to the Hennessy-Milner property for the objective and subjective interpretation of ATL. To conclude, in Sec. 6 we showed that checking the existence of an alternating bisimulation between two iCGS is undecidable in general (Theorem 5).

We note that our Hennessy-Milner theorem utilizes the "yesterday" modality for technical reasons. As noted in the proof of Theorem 3, Formula 4 might be encoded with an ATL* formula which does not utilize \mathbf{Y} . The translation of this theorem to the full ATL* is left for future research.

As another direction for future research, we plan to investigate under which conditions our Gale-Stewart-type theorem can be generalized to a full determinacy theorem for multi-agent games.

References

1. T. Ågotnes, V. Goranko, and W. Jamroga. Alternating-time temporal logics with irrevocable strategies. In *Proceedings of TARK XI*, pages 15–24, 2007.
2. T. Ågotnes, V. Goranko, W. Jamroga, and M. Wooldridge. Knowledge and ability. In *Handbook of Logics for Knowledge and Belief*. College Publications, 2015.
3. R. Alur, T. A. Henzinger, and O. Kupferman. Alternating-time temporal logic. *Journal of the ACM*, 49(5):672–713, 2002.
4. R. Alur, T. A. Henzinger, O. Kupferman, and M. Y. Vardi. Alternating refinement relations. In *In Proceedings of the Ninth International Conference on Concurrency Theory (CONCUR'98), volume 1466 of LNCS*, pages 163–178. Springer-Verlag, 1998.
5. F. Belardinelli, R. Condurache, C. Dima, W. Jamroga, and A. V. Jones. Bisimulations for verifying strategic abilities with an application to threeballot. In *Proc. of the 16th International Conference on Autonomous Agents and Multiagent Systems (AAMAS17)*, 2017.

6. F. Belardinelli, R. Condurache, C. Dima, W. Jamroga, and M. Knapik. Bisimulations for verifying strategic abilities with an application to the threeballot voting protocol. *Information & Computation*, 2020 (to appear).
7. F. Belardinelli, C. Dima, and A. Murano. Bisimulations for logics of strategies: A study in expressiveness and verification. In M. Thielscher, F. Toni, and F. Wolter, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the Sixteenth International Conference, KR 2018, Tempe, Arizona, 30 October - 2 November 2018.*, pages 425–434. AAAI Press, 2018.
8. N. Bulling, J. Dix, and W. Jamroga. Model checking logics of strategic ability: Complexity. In *Specification and Verification of Multi-agent Systems*, pages 125–159. Springer, 2010.
9. N. Bulling and W. Jamroga. Comparing variants of strategic ability: how uncertainty and memory influence general properties of games. *Autonomous Agents and Multi-Agent Systems*, 28(3):474–518, 2014.
10. M. Dastani and W. Jamroga. Reasoning about strategies of multi-agent programs. In *Proceedings of AAMAS2010*, pages 625–632, 2010.
11. C. Dima, B. Maubert, and S. Pinchinat. Relating paths in transition systems: The fall of the modal mu-calculus. *ACM Trans. Comput. Log.*, 19(3):23:1–23:33, 2018.
12. C. Dima and F. Tiplea. Model-checking ATL under imperfect information and perfect recall semantics is undecidable. *CoRR*, abs/1102.4225, 2011.
13. R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning about Knowledge*. MIT Press, Cambridge, 1995.
14. D. Gale and F. M. Stewart. Infinite games with perfect information. In *Contributions to the theory of games, vol. 2*, Annals of Mathematics Studies, no. 28, pages 245–266. Princeton University Press, Princeton, N. J., 1953.
15. U. Goltz, R. Kuiper, and W. Penczek. Propositional temporal logics and equivalences. In *Proceedings of CONCUR '92*, pages 222–236, 1992.
16. M. Hennessy and R. Milner. On observing nondeterminism and concurrency. In J. de Bakker and J. van Leeuwen, editors, *Automata, Languages and Programming*, pages 299–309, Berlin, Heidelberg, 1980. Springer Berlin Heidelberg.
17. W. Jamroga and W. van der Hoek. Agents that know how to play. *Fundamenta Informaticae*, 62:1–35, 2004.
18. M. Melissen. *Game Theory and Logic for Non-repudiation Protocols and Attack Analysis*. PhD thesis, University of Luxembourg, 2013.