

Satisfiability of Alternating-Time Temporal Epistemic Logic Through Tableaux

Francesco Belardinelli
Laboratoire IBISC – Université d’Evry

Abstract

In this paper we present a tableau-based method to decide the satisfiability of formulas in ATEL, an extension of the alternating-time temporal logic ATL including epistemic modalities for individual knowledge. Specifically, we analyse satisfiability of ATEL formulas under a number of conditions. We evaluate the assumptions of synchronicity and of a unique initial state, which have been proposed in the context of Interpreted Systems. Also, we consider satisfiability at an initial state as opposed to any state in the system. We introduce a tableau-based decision procedure for each of these combinations. Moreover, we adopt an agent-based approach to satisfiability, namely, the decision procedure returns a set of agents inducing a concurrent game structure that satisfies the relevant specification.

1 Introduction

Formal languages for knowledge reasoning and representation find increasingly application to the analysis of concurrent and reactive systems (van Harmelen, Lifschitz, and Porter 2007). In particular, multi-modal epistemic logics (Fagin et al. 1995; Meyer and Hoek 1995) have proved to be an invaluable tool for the specification and verification of systems of autonomous agents, especially in combination with temporal operators (Gammie and van der Meyden 2004; Lomuscio, Qu, and Raimondi 2009; Kacprzak et al. 2008). Among these multi-agent logics, the alternating-time temporal epistemic logic ATEL has been proposed as a formal language to represent and reason about individual and group strategies, as well as the knowledge thereof (van der Hoek and Wooldridge 2003; Jamroga and van der Hoek 2004). The theoretical properties of ATEL have been thoroughly investigated, also w.r.t. the assumptions of imperfect information and commitment in strategies, thus producing a wealth of results (Agotnes, Goranko, and Jamroga 2007; Bulling, Dix, and Jamroga 2010).

This paper aims to make a novel contribution to the formal analysis of logics of knowledge and strategies by introducing a tableau-based decision procedure for ATEL formulas. Tableaux are a well-established technique to decide satisfiability, actively researched both in the realm of logic and in

theoretical computer science (D’Agostino et al. 1999). Recently, a stream of papers has appeared on tableau methods for various flavours of temporal logics as well as multi-agent epistemic logics (Goranko and Shkatov 2009a; 2009b; 2009c; Ajspur, Goranko, and Shkatov 2013). While we explicitly acknowledge the influence of these works and make use of part of their formal machinery, we substantially extend the object of investigation. Specifically, our work differs from the contributions above in three ways. Firstly, we adopt an agent-based perspective and consider agents as the basic components of our epistemic concurrent game models. This means that the decision procedure returns not just a model satisfying the formula if successful, instead a system of agents is provided. Secondly, we follow the paradigm of Interpreted Systems (Fagin et al. 1995) and consider forms of interaction between the temporal and epistemic dimensions. In particular, we introduce epistemic concurrent game models which are synchronous or have a unique initial state. Thirdly, we analyse two different notions of satisfiability, namely satisfiability in some initial state, as opposed to satisfiability in any state. We maintain that the former notion is typical in the modelling and verification of concurrent systems (Baier and Katoen 2008), while the latter has traditionally been studied in mathematical logic (Blackburn, de Rijke, and Venema 2001). We will see that all these choices do have an impact on tableaux construction.

The motivation for the present work comes also from the fact that, besides the theoretical interest of algorithmic decision techniques, tableaux for ATEL can in principle be used to synthesize agent systems capable of enforcing behaviours specified as ATEL formulas. Thus, the investigations carried out hereafter can be seen as a preliminary contribution to bridge the gap between knowledge representation and model synthesis.

Related Work. This contribution builds on a series of papers on tableaux for multi-agent modal logics. Specifically, (Goranko and Shkatov 2009a) puts forward incremental tableaux for (non-epistemic) ATL; while in (Ajspur, Goranko, and Shkatov 2013) an epistemic logic with group knowledge is considered. In (Goranko and Shkatov 2009b; 2009c) the linear- and branching-time temporal epistemic logics LTLK and CTLK are given tableau-based decision procedures. However, we extend the object of investigation as detailed above. In (Walther 2005) tableaux for ATEL are

presented, but their construction is top-down, while we follow the references above and provide incremental tableaux. This point is crucial because, even though in the worst case the computational complexity remains the same, the incremental procedure needs not to build all formula types. Hence, for the purpose of implementation its performance is, on average, better. On a related subject, we remark that in (Goranko and Jamroga 2004) a translation from ATEL to ATL is provided, but at the cost of an exponential blow-up.

Scheme of the paper. In Section 2 we present the epistemic concurrent game models and the logic ATEL; while Section 3 introduces the Hintikka structures for ATEL. In Section 4 we define the tableaux for satisfiability in any state of the model; while Section 5 is dedicated to satisfiability at initial states. We conclude in Section 6 and point to future work. For reasons of space we provide sketches of proofs only in selected cases. An extended version with some more details is available for the reader's perusal (Beldardinelli 2014).

2 Epistemic Concurrent Game Models

In this section we present the formal machinery that will be used in the rest of the paper. We first introduce the epistemic concurrent game models, inspired to (van der Hoek and Wooldridge 2003; Jamroga and van der Hoek 2004), by adopting an agent-based perspective. Then, we consider the alternating-time epistemic logic ATEL. We start with the notion of *agent*.

Definition 1 (Agent) An agent is a tuple $i = \langle L_i, Act_i, Pr_i \rangle$ where

- L_i is a set of local states l_i, l'_i, \dots ;
- Act_i is the finite set of actions $\sigma_i, \sigma'_i, \dots$;
- $Pr_i : L_i \mapsto 2^{Act_i}$ is the protocol function.

Intuitively, at each moment the agent i is situated in some local state $l_i \in L_i$ representing all the information she has. In this respect we follow the typical approach to Multi-agent Systems (Fagin et al. 1995; Wooldridge 2000). Also, we assume that agents perform the actions in Act_i according to the protocol function Pr_i .

As we are interested in agent interactions, we introduce their composition. Given a set $Ag = \{i_0, \dots, i_n\}$ of agents, we define the set L of *global states* s, s', \dots (resp. the set Act of *joint actions* σ, σ', \dots) as the cartesian product $L_0 \times \dots \times L_n$ (resp. $Act_0 \times \dots \times Act_n$). In the rest of the paper we fix a set AP of atomic propositions and represent the j -th element of a tuple $t = \langle t_1, \dots, t_n \rangle$ as t_j or $t(j)$.

Definition 2 (ECGM) Given a set $Ag = \{i_0, \dots, i_n\}$ of agents $i = \langle L_i, Act_i, Pr_i \rangle$, an epistemic concurrent game model is a tuple $\mathcal{P} = \langle Ag, I, \tau, \pi \rangle$ where

- $I \subseteq L$ is the set of initial global states;
- $\tau : L \times Act \mapsto L$ is the global transition function, where $\tau(s, \sigma)$ is defined iff $\sigma_i \in Pr_i(l_i)$ for every $i \in Ag$;
- $\pi : AP \rightarrow 2^L$ is an interpretation of atomic propositions.

Intuitively, an ECGM evolves from an initial state in I as specified by the global transition function τ , which returns a successor state for each enabled joint action σ . In

line with the tradition of Interpreted Systems (Fagin et al. 1995), the agent i_0 will be reserved to model the environment the agents act in. That is, i_0 will encode all the relevant information about the system that does not appear in the private local state of any agent. In this capacity we denote i_0 as Env . Further, we observe that standard notions of concurrent game models can be defined in the present framework. Specifically, given a global state $s = \langle l_0, \dots, l_n \rangle$ and an agent i , we introduce the functions $D_i(s) = Pr_i(l_i)$ for the set of enabled actions, $d_i(s) = |Pr_i(l_i)|$ for the number of enabled actions, and $D(s) = \{\sigma \in Act \mid \text{for every } i \in Ag, \sigma_i \in Pr_i(l_i)\} = D_0(s) \times \dots \times D_n(s)$ for the set of joint enabled actions (Alur, Henzinger, and Kupferman 2002; Goranko and Shkatov 2009a).

We now fix some more notation. The *transition relation* \rightarrow on global states is defined so that $s \rightarrow s'$ if there exists $\sigma \in Act$ and $s \xrightarrow{\sigma} s'$, i.e., $\tau(s, \sigma) = s'$. We will also consider the transitive closure \rightarrow^* . A *run* λ from a state s , or *s-run*, is an infinite sequence s^0, s^1, \dots s.t. $s_i \rightarrow s_{i+1}$ and $s^0 = s$. For $n, m \in \mathbb{N}$, with $n \leq m$, we define $\lambda(n) = s^n$ and $\lambda[n, m] = s^n, s^{n+1}, \dots, s^m$. A state s' is *reachable from* s iff $s \rightarrow^* s'$. Now, \mathcal{R} is the set of all s_0 -runs, for some initial state $s_0 \in I$, and $S = \{\lambda(n) \mid \lambda \in \mathcal{R}, n \in \mathbb{N}\}$ is the set of states reachable from any initial state. Further, let \sharp be a placeholder for arbitrary actions. Given a group $A \subseteq Ag$ of agents, an *A-action* σ_A is an $|Ag|$ -tuple s.t. (i) $\sigma_A(i) \in Act_i$ for $i \in A$, and (ii) $\sigma_A(j) = \sharp$ for $j \notin A$. We denote by Act_A the set of all such *A-actions* and $D_A(s) = \{\sigma_A \in Act_A \mid \text{for every } i \in A, \sigma_A(i) \in Pr_i(l_i)\}$. A joint action σ extends an *A-action* σ_A , or $\sigma_A \sqsubseteq \sigma$, iff $\sigma_A(i) = \sigma(i)$ for all $i \in A$. The *outcome of an A-action* σ_A at a state s , denoted by $out(s, \sigma_A)$, is the set of all states s' for which there exists a joint action $\sigma \in Act$ s.t. $\sigma_A \sqsubseteq \sigma$ and $\tau(s, \sigma) = s'$. As for the epistemic component of ECGM, two global states $s = \langle l_0, \dots, l_n \rangle$ and $s' = \langle l'_0, \dots, l'_n \rangle$ are *indistinguishable* for agent i , written $s \sim_i s'$, iff $l_i = l'_i$ (Fagin et al. 1995).

In this paper we consider two conditions on ECGM.

Definition 3 An ECGM \mathcal{P}

- is *synchronous* iff for every $\lambda, \lambda' \in \mathcal{R}$, $i \in Ag$, $n \in \mathbb{N}$, if $\lambda(n) \sim_i \lambda'(n')$ then $n = n'$;
- has a unique initial state iff $|I| = 1$.

Synchronicity and uniqueness of initial state are typically considered in the context of Interpreted Systems, where their impact on axiomatisations and the satisfiability problem has been thoroughly assessed (Halpern and Vardi 1986; 1989; Halpern, van der Meyden, and Vardi 2004). The meaning of a unique initial state is intuitively clear; while in synchronous ECGM the agents are assumed to have access to a global clock. Hereafter, we denote the class of all ECGM satisfying either uniqueness of initial state or synchronicity with the superscripts *uis* and *sync*. Hence, for instance, $ECGM^{uis, sync}$ is the class of all synchronous ECGM with a unique initial state.

We now present the alternating-time epistemic logic ATEL as a specification language for ECGM.

Definition 4 (ATEL) The ATEL formulas φ are defined in BNF as follows:

$\varphi ::= p \mid \neg\varphi \mid \varphi \rightarrow \varphi \mid \langle\langle A \rangle\rangle X\varphi \mid \langle\langle A \rangle\rangle G\varphi \mid \langle\langle A \rangle\rangle \varphi U \varphi \mid K_i\varphi$
where $A \subseteq Ag$, $i \neq Env$, and $Env \notin A$.

The language ATEL is an extension of the alternating-time logic ATL enriched with an epistemic operator K_i for each agent different from the environment (van der Hoek and Wooldridge 2003). Also, the environment Env does not appear in coalition modalities. This is once more in line with the semantics of Interpreted Systems, where the environment is seen as adversarial w.r.t. the other agents (Fagin et al. 1995). The ATEL formulas (i) $\langle\langle A \rangle\rangle X\varphi$, (ii) $\langle\langle A \rangle\rangle G\varphi$ and (iii) $\langle\langle A \rangle\rangle \varphi U \varphi'$ are read as “the agents in A have a strategy to...” (i) “...enforce φ at the next state”, (ii) “...always enforce φ ”, and (iii) “...enforce φ until φ' ”; while $\langle\langle A \rangle\rangle F\varphi$ is defined as $\langle\langle A \rangle\rangle \top U \varphi$. The epistemic formula $K_i\varphi$ intuitively means that “agent i knows φ ”. For ease of presentation we do not consider operators for group knowledge, even if we envisage such an extension. In what follows we define $(\langle\langle A \rangle\rangle X)^k \phi$ by induction on k as (i) $(\langle\langle A \rangle\rangle X)^0 \phi = \phi$, and (ii) $(\langle\langle A \rangle\rangle X)^{k+1} \phi = (\langle\langle A \rangle\rangle X)^k \langle\langle A \rangle\rangle X \phi$. Also, the positive (resp. negative) X -formulas have the form $\langle\langle A \rangle\rangle X \phi$ (resp. $\neg \langle\langle A \rangle\rangle X \phi$).

To interpret ATEL formulas on ECGM we need to introduce the notion of a *strategy* for a set A of agents. Let γ be an ordinal with $1 \leq \gamma \leq \omega$, a γ -recall A -strategy is a mapping $F_A[\gamma] : \bigcup_{1 \leq n < 1+\gamma} S^n \mapsto \bigcup_{s \in S} D_A(s)$ s.t. $F_A[\gamma](\kappa) \in D_A(\text{lst}(\kappa))$ for every $\kappa \in \bigcup_{1 \leq n < 1+\gamma} S^n$, where $\text{lst}(\kappa)$ is the last element of κ and $1 + \gamma = \gamma$ for $\gamma = \omega$. Intuitively, a strategy returns an enabled A -action for every finite sequence of states. We remark that, according to standard terminology in concurrent game models (Bulling, Dix, and Jamroga 2010), the agents in ECGM have *perfect information*, that is, their strategies are determined by all information available at each global state. Further, the *outcome* of strategy $F_A[\gamma]$ at state s , or $\text{out}(s, F_A[\gamma])$, is the set of all s -runs λ s.t. $\lambda(i+1) \in \text{out}(\lambda(i), F_A[\gamma](\lambda[1..i]))$ for all $i \geq 0$ and $j = \max(i - \gamma + 1, 0)$. Depending on γ , we can define positional strategies, perfect-recall strategies, etc. Since these distinctions are not relevant for the present discussion, as they all generate the same class of models, hereafter we assume that γ is fixed and omit it.

Definition 5 (Semantics of ATEL) We define whether an ECGM \mathcal{P} satisfies a formula φ at state s , or $(\mathcal{P}, s) \models \varphi$, as follows:

$$\begin{aligned} (\mathcal{P}, s) \models p & \quad \text{iff } s \in \pi(p) \\ (\mathcal{P}, s) \models \neg\varphi & \quad \text{iff } (\mathcal{P}, s) \not\models \varphi \\ (\mathcal{P}, s) \models \varphi \rightarrow \varphi' & \quad \text{iff } (\mathcal{P}, s) \not\models \varphi \text{ or } (\mathcal{P}, s) \models \varphi' \\ (\mathcal{P}, s) \models \langle\langle A \rangle\rangle X\varphi & \quad \text{iff there is an } A\text{-strategy } F_A \text{ s.t. for all} \\ & \quad \lambda \in \text{out}(s, F_A), (\mathcal{P}, \lambda(1)) \models \varphi \\ (\mathcal{P}, s) \models \langle\langle A \rangle\rangle G\varphi & \quad \text{iff there is an } A\text{-strategy } F_A \text{ s.t. for all} \\ & \quad \lambda \in \text{out}(s, F_A), i \geq 0, (\mathcal{P}, \lambda(i)) \models \varphi \\ (\mathcal{P}, s) \models \langle\langle A \rangle\rangle \varphi U \varphi' & \quad \text{iff there is an } A\text{-strategy } F_A \text{ s.t. for all} \\ & \quad \lambda \in \text{out}(s, F_A) \text{ there is } k \geq 0 \\ & \quad \text{s.t. } (\mathcal{P}, \lambda(k)) \models \varphi', \text{ and for all } j, \\ & \quad 0 \leq j < k \text{ implies } (\mathcal{P}, \lambda(j)) \models \varphi \\ (\mathcal{P}, s) \models K_i\varphi & \quad \text{iff for all } s' \in S, s \sim_i s' \text{ implies } (\mathcal{P}, s') \models \varphi \end{aligned}$$

As to the notion of satisfaction, multiple choices are possible. In mathematical logic satisfaction typically takes into account all the states in a system (Blackburn, de Rijke, and

Venema 2001); while in the modelling of concurrent systems satisfaction w.r.t. initial states only is also considered (Baier and Katoen 2008). This remark motivates the following definition.

Definition 6 (Satisfaction) An ATEL-formula θ is S -satisfied (resp. I -satisfied) in a ECGM \mathcal{P} iff there exists $s \in S$ (resp. $s \in I$) s.t. $(\mathcal{P}, s) \models \theta$.

While for the class of all ECGM the two notions are equivalent, under the assumptions of a unique initial state or synchronicity this is not longer the case. For instance, the ATEL-formula $\langle\langle \emptyset \rangle\rangle G\phi \wedge \langle\langle A \rangle\rangle F\neg K_i\phi$ is not I -satisfiable in $ECGM^{uis}$, but it is S -satisfiable in the same class. Similarly, the ATEL-formula $(\langle\langle \emptyset \rangle\rangle X)^k \phi \wedge (\langle\langle A \rangle\rangle X)^k \neg K_i\phi$ is not I -satisfiable in $ECGM^{uis, sync}$, but it is S -satisfiable. Indeed, for ECGM with a unique initial state, the operator $\langle\langle \emptyset \rangle\rangle G$ acts as a universal modality, traversing the whole state space. We maintain that both notions of satisfaction are of interest and will analyse both. In the following lemma we report the equivalences between classes of formulas, where, for instance, $ATEL_I^{uis}$ is the set of ATEL-formulas I -satisfiable in the class $ECGM^{uis}$.

Lemma 1 The classes of ATEL formulas in different boxes are different.

$ATEL_S = ATEL_S^{uis} = ATEL_S^{sync} = ATEL_S^{sync, uis}$ $= ATEL_I = ATEL_I^{sync}$
$ATEL_I^{uis}$
$ATEL_I^{sync, uis}$

By Lemma 1 we have to consider three distinct cases: $ATEL_S$, $ATEL_I^{uis}$, and $ATEL_I^{sync, uis}$, which will be the main focus of the paper. Hereafter we simply refer to *satisfaction* to indicate any of the two notions above. Moreover, satisfiability can be also considered w.r.t. different sets of agents (Walther et al. 2006). Let Ag_θ be the set of agents mentioned in an ATEL-formula θ , while $Ag_\theta^+ = Ag_\theta \cup \{Env\}$ (remember that Env does appear in θ).

Definition 7 An ATEL-formula θ is

1. Ag -satisfiable, for some $Ag \supseteq Ag_\theta^+$, iff θ is satisfiable in an ECGM $\mathcal{P} = \langle Ag, I, \tau, \pi \rangle$;
2. generally satisfiable iff θ is satisfiable in an ECGM $\mathcal{P} = \langle Ag, I, \tau, \pi \rangle$ for some $Ag \supseteq Ag_\theta^+$.

By “absorbing” the agents in $Ag \setminus Ag_\theta$ in the environment and by adding dummy agents, we can now prove the following result that extends Corollary 2.25 in (Goranko and Shkatov 2009a) to ATEL.

Theorem 2 Every ATEL-formula θ is generally satisfiable iff it is Ag_θ^+ -satisfiable.

As a result, it is sufficient to consider Ag_θ^+ -satisfiability when checking for models of ATEL formulas.

3 Hintikka Structures for ECGM

Tableau-based methods do not provide directly the formula in question with a relevant model (ECGM in the present case), but they first go through the construction of *Hintikka*

	α	α_1	α_2
	$\neg\neg\phi$	ϕ	ϕ
	$\neg(\phi \rightarrow \psi)$	ϕ	$\neg\psi$
	$\langle\langle A \rangle\rangle G\phi$	ϕ	$\langle\langle A \rangle\rangle X \langle\langle A \rangle\rangle G\phi$
	$K_i\phi$	ϕ	$K_i\phi$

	β	β_1	β_2
	$\phi \rightarrow \psi$	$\neg\phi$	ψ
	$\langle\langle A \rangle\rangle(\phi U \psi)$	ψ	$\phi \wedge \langle\langle A \rangle\rangle X \langle\langle A \rangle\rangle(\phi U \psi)$
	$\neg\langle\langle A \rangle\rangle(\phi U \psi)$	$\neg\phi \wedge \neg\psi$	$\neg\psi \wedge \neg\langle\langle A \rangle\rangle X \langle\langle A \rangle\rangle(\phi U \psi)$
	$\neg\langle\langle A \rangle\rangle G\phi$	$\neg\phi$	$\neg\langle\langle A \rangle\rangle X \langle\langle A \rangle\rangle G\phi$

Figure 1: α - and β -formulas in ATEL.

structures. The main difference between ECGM and Hintikka structures is that, while the former define truth conditions for every ATEL formula on the set AP of atomic propositions, the latter is only concerned with formulas ‘related’ to the satisfaction of the formula to be checked. To specify this relation between ATEL formulas, we introduce the distinction between α - and β -formulas of ATEL in Fig. 1. Intuitively, the states in a Hintikka structures are sets of formulas that contain α_1 and α_2 (resp. β_1 or β_2) whenever they contain α (resp. β), according to the semantics of the main logical symbol. This idea is made precise by the following definition.

Definition 8 A set Δ of ATEL formulas is

- consistent if it does not contain both ψ and $\neg\psi$.
- downward saturated if
 1. $\alpha \in \Delta$ implies $\alpha_1 \in \Delta$ and $\alpha_2 \in \Delta$;
 2. $\beta \in \Delta$ implies $\beta_1 \in \Delta$ or $\beta_2 \in \Delta$.

The states of a Hintikka structure will be consistent and downward saturated sets of ATEL formulas. To introduce the epistemic concurrent game Hintikka structures, we need to present the notions of *co-action* and *co-strategy* from (Goranko and van Drimmelen 2006). First, a *co-A-move* at state s is a function $\sigma_A^c : D_A(s) \rightarrow D(s)$ s.t. $\sigma_A \sqsubseteq \sigma_A^c(\sigma_A)$; $D_A^c(s)$ denotes the set of all co-A-moves at s . Further, $out(s, \sigma_A^c) = \{\tau(s, \sigma_A^c(\sigma_A)) \mid \sigma_A \in D_A(s)\}$ is the *outcome* of σ_A^c at s . Then, a γ -recall *co-A-strategy* is a mapping $F_A^c[\gamma] : \bigcup_{1 \leq n < 1+\gamma} S^n \mapsto \bigcup_{s \in S} D_A^c(s)$ s.t. $F_A^c[\gamma](\kappa) \in D_A^c(lst(\kappa))$ for every $\kappa \in \bigcup_{1 \leq n < 1+\gamma} S^n$. Finally, the *outcome* of $F_A^c[\gamma]$ at s is the set $out(s, F_A^c[\gamma])$ of all s -runs λ s.t. $\lambda(i+1) \in out(\lambda(i), F_A^c[\gamma](\lambda[j, i]))$ for all $i \geq 0$ and $j = \max(i - \gamma + 1, 0)$. Co-strategies are used to define the labelling function of Hintikka structures in the cases of ATEL-formulas $\neg\langle\langle A \rangle\rangle X\varphi$ and $\neg\langle\langle A \rangle\rangle G\varphi$. In what follows we assume γ fixed and omit it.

Definition 9 (ECGHS) Let \mathcal{A} be a finite set of agent indexes i_0, \dots, i_n . An epistemic concurrent game Hintikka structure, or ECGHS, is a tuple $\mathcal{H} = \langle \S, I, \{Act_i\}_{i \in \mathcal{A}}, \{\sim_i\}_{i \in \mathcal{A} \setminus \{Env\}}, \tau, H \rangle$ where

- \S is a set of states;
- $I \subseteq \S$ is a set of initial states;
- for each $i \in \mathcal{A}$, Act_i is a finite set of actions;
- $\tau : \S \times Act \rightarrow \S$ is the transitions function;

- for each $i \in \mathcal{A} \setminus \{Env\}$, \sim_i is a symmetric, reflexive and transitive relation on \S ;
- H is a labelling function from \S to sets of ATEL formulas s.t.

- H1 if $\neg\phi \in H(s)$ then $\phi \notin H(s)$;
- H2 every $H(s)$ is downward saturated;
- H3 if $\langle\langle A \rangle\rangle X\varphi \in H(s)$ then there is an A -strategy F_A s.t. for all $\lambda \in out(s, F_A)$, $\varphi \in H(\lambda(1))$;
- H4 if $\neg\langle\langle A \rangle\rangle X\varphi \in H(s)$ then there is a co- A -strategy F_A^c s.t. for all $\lambda \in out(s, F_A^c)$, $\neg\varphi \in H(\lambda(1))$;
- H5 if $\neg\langle\langle A \rangle\rangle G\varphi \in H(s)$ there is a co- A -strategy F_A^c s.t. for all $\lambda \in out(s, F_A^c)$ there is $i \geq 0$ s.t. $\neg\varphi \in H(\lambda(i))$;
- H6 if $\langle\langle A \rangle\rangle \varphi U \varphi' \in H(s)$ there is an A -strategy F_A s.t. for all $\lambda \in out(s, F_A)$ there is $k \geq 0$ s.t. $\varphi' \in H(\lambda(k))$, and for all j , $0 \leq j < k$ implies $\varphi \in H(\lambda(j))$;
- H7 if $s \sim_i s'$ then $K_i\varphi \in H(s)$ iff $K_i\varphi \in H(s')$;
- H8 if $\neg K_i\varphi \in H(s)$ then there is $s' \in \S$ s.t. $s \sim_i s'$ and $\neg\varphi \in H(s')$.

Notice that conditions H3-8 mimic the satisfaction clauses for ATEL formulas in Def. 5. Also, ECGHS are syntactic structures, where each state is labelled by a finite set of ATEL formulas. In addition, the main difference w.r.t. ECGM consists in that global states are primitive. This implies that the indistinguishability relations \sim_i are also primitive. Further, as for ECGM, we can define a notion of run, as well as the set \mathcal{R} of all runs. Hence, we can introduce the following classes of ECGHS.

Definition 10 An ECGHS \mathcal{H}

- is synchronous iff for every $\lambda, \lambda' \in \mathcal{R}$, $i \in \mathcal{A}$, if $\lambda(n) \sim_i \lambda'(n')$ then $n = n'$;
- has a unique initial state iff $|I| = 1$.

Further, similarly to the case for ECGM, we introduce two notions of satisfaction.

Definition 11 (Satisfaction) An ECGHS \mathcal{H} \S -satisfies (resp. I -satisfies) an ATEL-formula θ iff $\theta \in H(s)$ for some $s \in \S$ (resp. $s \in I$).

The tableau procedure that we present in the following sections ultimately generates a Hintikka structure. From this ECGHS we are always able to extract an ECGM, as stated in the following result.

Theorem 3 Given a finite set \mathcal{A} of agent indexes and an ECGHS \mathcal{H} on \mathcal{A} that \S -satisfies (resp. I -satisfies) an ATEL-formula θ , we can derive a finite set \mathcal{A} of agents and an ECGM \mathcal{P} on \mathcal{A} that also S -satisfies (resp. I -satisfies) θ . Further, this can be done so as to preserve synchronicity and uniqueness of initial state.

As a consequence of Theorem 3, satisfiability w.r.t. Hintikka structures suffices to obtain satisfiability w.r.t. ECGM. Also, we underline that Theorem 3 is constructive as it provides a set \mathcal{A} of agents generating the ECGM \mathcal{P} . In this respect our approach is agent-based. In the following sections we present the tableau-based procedures for deciding satisfiability of ATEL formulas.

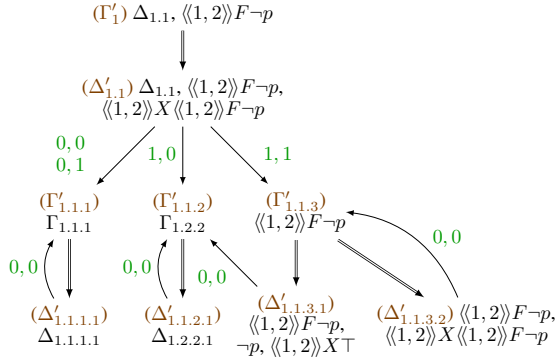


Figure 2: fragment of the pre-tableau for $\Gamma'_1 = \{\theta_1, \langle(1,2)\rangle F(\neg p \wedge \neg K_2 p)\}$.

4 Tableaux for \mathcal{S} -Satisfiability

In this section we introduce a tableau-based decision procedure for \mathcal{S} -satisfiability of ATEL formulas. Specifically, we build on the methods developed for (non-epistemic) ATL and multi-agent epistemic CMAEL(CD) in (Goranko and Shkatov 2009a; Ajspur, Goranko, and Shkatov 2013). As stated in Lemma 1, for \mathcal{S} -satisfiability the assumptions of a unique initial state or synchronicity do not affect satisfiability of ATEL formulas. Nonetheless, we deem the tableaux for \mathcal{S} -satisfiability interesting in themselves, as well as useful for a comparison with the case of \mathcal{I} -satisfiability. Indeed in this section we introduce notions and present results that will be used also in Section 5.

Similarly to (Goranko and Shkatov 2009a; Ajspur, Goranko, and Shkatov 2013), the tableau procedure for ATEL consists of two phases: a *construction phase*, where a *pre-tableau* is populated with states and pre-states (i.e., sets of formulas), and an *elimination phase*, where the pre-tableau is pruned into a proper *tableau* by removing pre-states together with “unsatisfiable” states. The aim of the whole construction is to obtain a tableau from which it is possible to extract a Hintikka structure iff the relevant formula is satisfiable. We start with the construction phase.

4.1 Construction Phase

The construction of the pre-tableau $\text{pre-}\mathcal{T}^\theta$ for an ATEL-formula θ proceeds in stages. At each stage rules are applied to create either states from pre-states or pre-states from states. The process starts with the pre-state $\Gamma = \{\theta\}$ containing the ATEL formula to be checked. Then at each stage we apply one of the rules \mathbf{SR}_S , \mathbf{K}_S and \mathbf{Next}_S that adapt the corresponding rules in (Goranko and Shkatov 2009a; 2009c). In what follows we say that a set Δ of formulas is a *minimal downward saturated* extension of a set Γ if (i) Δ is downward saturated, (ii) $\Gamma \subseteq \Delta$, and (iii) there is no downward saturated Δ' s.t. $\Gamma \subseteq \Delta' \subset \Delta$. The first rule we present allows us to build states from pre-states.

Definition 12 (Saturation Rule \mathbf{SR}_S) Given a pre-state Γ s.t. \mathbf{SR}_S has not been applied to Γ , for every consistent and minimal downward saturated extension Δ of Γ ,

- if the pre-tableau contains a state Δ' that coincides with Δ , set $\Gamma \Rightarrow \Delta'$;
- else,
 - add Δ as a new state to the pre-tableau;
 - if Δ contains no positive nor negative X -formula, then add $\langle\langle Ag_\theta \rangle\rangle X\top$ to Δ ;
 - set $\Gamma \Rightarrow \Delta$.

The (finite) set of states $\{\Delta \mid \Gamma \Rightarrow \Delta\}$ is denoted by $\mathbf{states}(\Gamma)$. To introduce the next rule we need some more notation. Hereafter, $\Delta|_{K_i}$ denotes the set $\{K_i\psi \mid K_i\psi \in \Delta\}$. Also, we do not distinguish between a finite set of formulas and their conjunction, the context will disambiguate. The next rule is designed to ensure the satisfaction of epistemic possibilities of the form $\neg K_i\phi$.

Definition 13 (Rule \mathbf{K}_S) Given a state Δ s.t. \mathbf{K}_S has not been applied to Δ , for every $\neg K_i\phi \in \Delta$ and consistent $\Gamma = \{\neg\phi\} \cup \Delta|_{K_i}$,

- if the pre-tableau contains a pre-state $\Gamma' = \Gamma$, set $\Delta \xrightarrow{\neg K_i\phi} \Gamma'$;
- else, add Γ as a new pre-state and set $\Delta \xrightarrow{\neg K_i\phi} \Gamma$.

Finally, we consider the rule \mathbf{Next}_S , which builds the next level of pre-states.

Definition 14 (Rule \mathbf{Next}_S) Given a state Δ s.t. \mathbf{Next}_S has not been applied to Δ ,

1. order linearly all positive X -formulas of Δ followed by all negative ones:

$$L = \langle\langle A_0 \rangle\rangle X\phi_0, \dots, \langle\langle A_{m-1} \rangle\rangle X\phi_{m-1}, \\ \neg\langle\langle B_0 \rangle\rangle X\psi_0, \dots, \neg\langle\langle B_{l-1} \rangle\rangle X\psi_{l-1}$$

(because of \mathbf{SR}_S , L is always non-empty.) Let $r_\Delta = m+l$ and $D(\Delta) = \{0, \dots, r_\Delta - 1\}^{|Ag_\theta^+|}$. Further, for every $\sigma \in D(\Delta)$, we set $N(\sigma) = \{i \mid \sigma_i \geq m\}$ and $\text{neg}(\sigma) = \sum_{i \in N(\sigma)} (\sigma_i - m) \pmod{l}$.

2. For each $\sigma \in D(\Delta)$ consider every consistent

$$\Gamma_\sigma = \{\neg\psi_q \mid \neg\langle\langle B_q \rangle\rangle X\psi_q \in \Delta, \text{neg}(\sigma) = q, Ag_\theta^+ \setminus B_q \subseteq N(\sigma)\} \\ \cup \{\phi_p \mid \langle\langle A_p \rangle\rangle X\phi_p \in \Delta \text{ and } \sigma_i = p \text{ for all } i \in A_p\} \quad (1)$$

where $\Gamma_\sigma = \{\top\}$ if the set on the RHS of (1) is empty.

- If the pre-tableau contains a pre-state $\Gamma' = \Gamma_\sigma$, set $\Delta \xrightarrow{\sigma} \Gamma'$;
- else, add Γ_σ as a new pre-state and set $\Delta \xrightarrow{\sigma} \Gamma_\sigma$.

We remark that by definition of \mathbf{SR}_S , \mathbf{K}_S and \mathbf{Next}_S , no state nor pre-state in $\text{pre-}\mathcal{T}^\theta$ is inconsistent, i.e., it does not contain both ϕ and $\neg\phi$. Also, we denote the (finite) set of pre-states $\{\Gamma \mid \Delta \xrightarrow{\sigma} \Gamma \text{ for some } \sigma\}$ as $\mathbf{pre-states}(\Delta)$. We now briefly describe the construction of the pre-tableau for an ATEL-formula θ . We begin at stage $k = 0$ by creating a pre-state $\Gamma = \{\theta\}$ for the input formula θ . Then, at each stage $k > 0$, we alternately apply the rule \mathbf{SR}_S to obtain the states in $\mathbf{states}(\Gamma)$, followed by \mathbf{Next}_S and \mathbf{K}_S to create the next level of pre-states and the epistemically related pre-states. The construction phase terminates, after a finite number of stages, when \mathbf{K}_S and \mathbf{Next}_S do not produce any new pre-state, nor \mathbf{SR}_S new states. We conclude this section with an example.




Figure 3: fragment of the pre-tableau for $\theta_1 = \langle\langle 1 \rangle\rangle F \neg K_2 p \wedge K_1 \langle\langle 2 \rangle\rangle G p$ obtained by applying SR_S and Next_S only.

Example 1 Fig. 3 shows a fragment of the pre-tableau for the ATEL-formula $\theta_1 = \langle\langle 1 \rangle\rangle F \neg K_2 p \wedge K_1 \langle\langle 2 \rangle\rangle G p$ obtained by applying rules SR_S and Next_S only. We present this restricted structure as it will be used later on. Also, the environment's actions are omitted. Indeed, since no negative X -formula appears in the pre-tableau, it suffices to consider each action σ_1, σ_2 in Fig. 3 as a shorthand for the set $\{\sigma_{Env}, \sigma_1, \sigma_2 \mid 0 \leq \sigma_{Env} < r_\Delta\}$. ■

4.2 Elimination phase

After completion of the pre-tableau $\text{pre-}\mathcal{T}^\theta$ we step into the elimination phase, where we prune $\text{pre-}\mathcal{T}^\theta$ first by removing all pre-states, and then by deleting “unsatisfiable” states. We have already noticed that by construction no state in the pre-tableau is inconsistent. Moreover, if a state either (i) has no successor for an enabled joint action, or (ii) it does not fulfil an epistemic possibility, or (iii) it does not fulfil a temporal eventuality, then it is removed. We make this intuition precise in the following elimination rules, adapted from (Goranko and Shkatov 2009a).

Definition 15 (Rule PR) For every state Δ and pre-state Γ in $\text{pre-}\mathcal{T}^\theta$, and $\Delta' \in \text{states}(\Gamma)$,

- if $\Delta \xrightarrow{C} \Gamma$ then $\Delta \xrightarrow{C} \Delta'$, where C can be either an action σ or an epistemic possibility $\neg K_i \phi$;
- remove Γ from $\text{pre-}\mathcal{T}^\theta$.

We now describe the process of state elimination also in stages. We start with the tableau \mathcal{T}_0^θ , obtained by applying **PR** to $\text{pre-}\mathcal{T}^\theta$. At each stage n we apply one of the rules **E1-3** below to obtain the tableau \mathcal{T}_{n+1}^θ .

E1 If for some $\sigma \in D(\Delta)$ there is no $\Delta' \in \mathcal{T}_n^\theta$ s.t. $\Delta \xrightarrow{\sigma} \Delta'$, then $\mathcal{T}_{n+1}^\theta = \mathcal{T}_n^\theta \setminus \{\Delta\}$.

E2_S If for some $\neg K_i \phi \in \Delta$ there is no $\Delta' \in \mathcal{T}_n^\theta$ s.t. $\Delta \xrightarrow{\neg K_i \phi} \Delta'$, then $\mathcal{T}_{n+1}^\theta = \mathcal{T}_n^\theta \setminus \{\Delta\}$.

Rule **E1** removes states with no successor for an enabled action; while **E2_S** deletes states with unrealized epistemic possibilities. Rule **E3** to be introduced deals with the realisation of temporal eventualities, namely the satisfaction of formulas of the form $\langle\langle A \rangle\rangle \phi U \psi$ and $\neg \langle\langle A \rangle\rangle G \phi$. To present

E3 we consider the following notation. Let Δ be a state in \mathcal{T}_0^θ and $\langle\langle A \rangle\rangle X \phi$ (resp. $\neg \langle\langle B \rangle\rangle X \psi$) the p -th (resp. q -th) formula in the linear ordering L of X -formulas of Δ in the application of Next_S . Then,

$$D(\Delta, \langle\langle A \rangle\rangle X \phi) = \{\sigma \in D(\Delta) \mid \sigma_i = p \text{ for every } i \in A\}$$

$$D(\Delta, \neg \langle\langle B \rangle\rangle X \psi) = \{\sigma \in D(\Delta) \mid \text{neg}(\sigma) = q \text{ and } \text{Ag}_\theta^+ \setminus B \subseteq N(\sigma)\}$$

We can now define when a temporal eventuality is realized.

Definition 16 (Realization of temporal eventualities) A temporal eventuality $\beta = \langle\langle A \rangle\rangle \phi U \psi$ (resp. $\neg \langle\langle A \rangle\rangle G \phi$) in $\Delta \in \mathcal{T}_n^\theta$ is realized at Δ' iff

1. $\beta_1 \in \Delta$, or
2. $\beta_2 \in \Delta$ and for every $\sigma \in D(\Delta, \langle\langle A \rangle\rangle X \langle\langle A \rangle\rangle \phi U \psi)$ (resp. $D(\Delta, \neg \langle\langle A \rangle\rangle X \langle\langle A \rangle\rangle G \phi)$) there exists $\Delta' \in \mathcal{T}_n^\theta$ s.t. (i) $\Delta \xrightarrow{\sigma} \Delta'$; and (ii) β is realized at $\Delta' \in \mathcal{T}_n^\theta$.

We can now state the final elimination rule.

E3 If $\Delta \in \mathcal{T}_n^\theta$ contains an unrealized temporal eventuality, then $\mathcal{T}_{n+1}^\theta = \mathcal{T}_n^\theta \setminus \{\Delta\}$.

The elimination phase terminates when none of the rules **E1-3** is any longer applicable (obviously, this happens after a finite number of stages.) The tableau thus obtained is denoted as \mathcal{T}^θ . In particular, we say that \mathcal{T}_θ is *open* if $\theta \in \Delta$ for some $\Delta \in \mathcal{T}_\theta$; otherwise, it is *closed*. In the next sections we show the soundness and completeness of the tableau procedure, that is, the tableau \mathcal{T}^θ is open iff θ is S -satisfiable.

4.3 Soundness and Completeness

We first outline the soundness proof for the tableau procedure presented above, which amounts to the following result:

Theorem 4 (Soundness) If θ is S -satisfiable, then \mathcal{T}^θ is open.

The proof of Theorem 4 consists of two parts. First, we use the following lemma to show that if θ is S -satisfiable, then at least one of the states in $\text{states}(\{\theta\})$ is S -satisfiable.

Lemma 5 Let Γ be an S -satisfiable pre-state in $\text{pre-}\mathcal{T}^\theta$. Then, at least one $\Delta \in \text{states}(\Gamma)$ is S -satisfiable.

Then, we show that no satisfiable state is removed during the elimination phase. The soundness result then follows, as θ belongs to $\Delta \in \mathcal{T}^\theta$, for some $\Delta \in \mathbf{states}(\{\theta\})$. To show that no S -satisfiable state is removed, we prove by induction on the number n of stages that, for every $\Delta \in \mathcal{T}_n^\theta$, if Δ is S -satisfiable, then Δ is not removed at stage n . The base case for $n = 0$ is trivial as no elimination rule has been applied. As regards the inductive step, we consider the following lemma:

Lemma 6 *If $\Delta \in \mathcal{T}_n^\theta$ is S -satisfiable, then*

1. *for every $\sigma \in D(\Delta)$ there is $\Delta' \in \mathcal{T}_n^\theta$ s.t. $\Delta \xrightarrow{\sigma} \Delta'$;*
2. *for every $\neg K_i \phi \in \Delta$ there is $\Delta' \in \mathcal{T}_n^\theta$ s.t. $\Delta \xrightarrow{\neg K_i \phi} \Delta'$;*
3. *every temporal eventuality is realized at Δ .*

We conclude that if Δ is S -satisfiable, then it will not be eliminated by **E1-3**. Hence, some $\Delta \supseteq \{\theta\}$ belongs to \mathcal{T}^θ , i.e., \mathcal{T}^θ is open.

The completeness proof consists in showing that if the tableau \mathcal{T}^θ is open, then the ATEL-formula θ is S -satisfiable.

Theorem 7 (Completeness) *If \mathcal{T}^θ is open, then θ is S -satisfiable.*

The proof amounts to building an ECGHS \mathcal{H}_θ starting from the open tableau \mathcal{T}^θ . By Theorem 3 the desired result immediately follows. Specifically, when building the Hintikka structure for θ we take care of satisfying conditions H5, H6 and H8 for temporal eventualities and epistemic possibilities. In this task the key concept is that of a *temporal component* $F(\xi, \Delta)$ for an eventuality ξ and state Δ (Goranko and Shkatov 2009a; 2009c). If $\xi \in \Delta$ then $F(\xi, \Delta)$ is a tree rooted at Δ that realizes ξ . If $\xi \notin \Delta$ then $F(\xi, \Delta)$ is a tree with all the successor of Δ . We now outline the construction of the ECGHS \mathcal{H}_θ satisfying an ATEL-formula θ .

The construction proceeds in stages. First, we list all states in \mathcal{T}^θ as $\Delta_0, \dots, \Delta_{n-1}$ and all temporal eventualities as ξ_0, \dots, ξ_{m-1} . Further, we assume that the temporal components are arranged in an $m \times n$ matrix, whose rows (resp. columns) are marked with the corresponding temporal eventualities (resp. states). Hence, $F_{p,q} = F(\xi_p, \Delta_q)$. Now, if θ is the temporal eventuality ξ_p , for $p < m$, we set $\mathfrak{S}_0^0 = F_{p,q}$ for some q s.t. $\theta \in \Delta_q$. Otherwise, we set $\mathfrak{S}_0^0 = F_{0,q}$, where q is as above. For the induction step, suppose that we have defined \mathfrak{S}_k^i . Then, we construct \mathfrak{S}_k^{i+1} by realizing all temporal eventualities $\xi \in \mathfrak{S}_k^i$. That is, we keep track of the last realized eventuality, say ξ_i , and replace every deadlocked state w such that $l(w) = \Delta_j$ with the temporal component $F_{(i+1) \bmod m, j}$. We repeat the procedure until all temporal eventualities have been realized. Further, given \mathfrak{S}_k^i we construct \mathfrak{S}_{k+1}^i by adding for every $\Delta \in \mathfrak{S}_k^i$ and every epistemic possibility $\neg K_i \phi \in \Delta$ a state $\Delta' \in \mathcal{T}^\theta$ s.t. $\neg \phi \in \Delta'$. Notice that such Δ' always exists as $\Delta \in \mathcal{T}^\theta$. We proceed in this way to realize all epistemic possibilities contained in \mathfrak{S}_k^i . To guarantee that the ECGHS we are building remains finite, whenever the component we are about to add, say $F_{p,q}$, is already contained in \mathfrak{S}_k^i , instead of replacing the state w (s.t. $l(w) = \Delta_q$) with $F_{p,q}$, we connect every

“predecessor” v of w to the root of $F_{p,q}$. The procedure is repeated until no more components are added, thus obtaining an ECGHS \mathcal{H}_θ .

Lemma 8 *If \mathcal{T}^θ is open, then there exists an ECGHS \mathcal{H}_θ that ξ -satisfies θ .*

Complexity. We conclude this section with a discussion of the complexity of the decision procedure outlined above. It can be shown that the construction and elimination phases can be performed in time exponential in the size of the formula. Hence, the procedure here provided meets the lower-bound of the satisfiability problem for ATEL (Walther 2005). In this respect, it is optimal. Further, as discussed in (Goranko and Shkatov 2009a), since the tableau construction is incremental, in many cases it allows to decide satisfiability in less time than top-down tableaux, which need to build all formula types.

5 Tableaux for I -Satisfiability

In this section we introduce the tableau-based decision procedure for I -satisfiability. Specifically, in Lemma 1 we remarked that the notions of I - and S -satisfiability coincide over the class of all (resp., synchronous) ECGM. On the other hand, for ECGM with a unique initial state we need to consider a different construction for I -satisfiability.

5.1 Construction Phase

We start with the rules to build the pre-tableau $\text{pre-}\mathcal{T}^\theta$, but first we need to introduce some new terminology. In what follows we index each state and pre-state with a finite set $\{\vec{k}\}$ of time-stamps (we assume the set is enumerated). Intuitively, these represent the times at which a state (resp. pre-state) occurs in a fragment of a run. Operations on a time-stamp $\{\vec{k}\}$ (such as addition) are defined on each member of $\{\vec{k}\}$. Further, in the construction phase we define a sequence $\Gamma_0^{\{0\}}, \dots, \Gamma_n^{\{0\}}$ of “initial” pre-states, starting from $\Gamma_0^{\{0\}} = \{\theta\}$. Also, we introduce a mark mk that intuitively keeps track of the initial state a state (resp. pre-state) stems from. We set $mk(\Gamma_0^{\{0\}}) = \text{undef}$ at the beginning. Finally, let \mapsto be the transitive closure of $\Rightarrow \cup \rightarrow$. We now provide the saturation rule $\mathbf{SR}_I^{\text{vis, sync}}$ for I -satisfiability w.r.t. synchronous ECGM with a unique initial state.

Definition 17 (Saturation Rule $\mathbf{SR}_I^{\text{vis, sync}}$) *Given a pre-state $\Gamma^{\{\vec{k}\}}$ s.t. $\mathbf{SR}_I^{\text{vis, sync}}$ has not been applied to $\Gamma^{\{\vec{k}\}}$, consider every consistent and minimal downward saturated extensions $\Delta^{\{\vec{k}\}}$ of $\Gamma^{\{\vec{k}\}}$,*

- *if the pre-tableau contains a state $\Delta^{\{\vec{m}\}} = \Delta^{\{\vec{k}\}}$ and $mk(\Delta^{\{\vec{m}\}}) = mk(\Gamma^{\{\vec{k}\}})$, then*
 - *rename $\Delta^{\{\vec{m}\}}$ and all pre-states and states $\Theta_i^{\{\vec{m}+i\}}$ s.t. $\Delta^{\{\vec{m}\}} \mapsto \Theta_i^{\{\vec{m}+i\}}$ as $\Delta^{\{\vec{m}\} \cup \{\vec{k}\}}$ and $\Theta_i^{\{\vec{m}+i\} \cup \{\vec{k}+i\}}$. Then, set $\Gamma^{\{\vec{k}\}} \Rightarrow \Delta^{\{\vec{m}\} \cup \{\vec{k}\}}$.*
- *else,*
 - *add $\Delta^{\{\vec{k}\}}$ as a new state to the pre-tableau;*

- if $\Delta^{\{\vec{k}\}}$ contains no positive nor negative X -formulas, then add $\langle\langle Ag_\theta \rangle\rangle X \top$ to $\Delta^{\{\vec{k}\}}$;
- set $\Gamma^{\{\vec{k}\}} \Rightarrow \Delta^{\{\vec{k}\}}$ and $mk(\Delta^{\{\vec{k}\}}) = mk(\Gamma^{\{\vec{k}\}})$ unless $\{\vec{k}\} = \{0\}$ and $mk(\Gamma^{\{0\}}) = \text{undef}$. In this case, set $mk(\Delta^{\{0\}}) = \Delta^{\{0\}}$.

Hence, $\mathbf{SR}_I^{uis, sync}$ modifies \mathbf{SR}_S by keeping track of time-stamps as well as of mark mk : the former are needed for synchronicity, while the latter registers the initial state. In particular, for all $\Delta, \Delta' \in \mathbf{states}(\Gamma)$, $mk(\Delta) = mk(\Delta')$, and for every state $\Delta', \Delta^{\{0\}} \mapsto \Delta'$ iff $mk(\Delta') = \Delta^{\{0\}}$. The rule \mathbf{SR}_I^{uis} for I -satisfiability w.r.t. $ECGM^{uis}$ can be obtained from $\mathbf{SR}_I^{uis, sync}$ by omitting time-stamps. Further, we need to modify the rule \mathbf{Next}_S as follows.

Definition 18 (Rule $\mathbf{Next}_I^{uis, sync}$) Given a state $\Delta^{\{\vec{k}\}}$ s.t. $\mathbf{Next}_I^{uis, sync}$ has not been applied to $\Delta^{\{\vec{k}\}}$,

1. (as in rule \mathbf{Next}_S)
2. For each $\sigma \in D(\Delta^{\{\vec{k}\}})$ consider every consistent

$$\Gamma_\sigma^{\{\vec{k}+1\}} = \{\phi_p \mid \langle\langle A_p \rangle\rangle X \phi_p \in \Delta^{\{\vec{k}\}}, \sigma_i = p \text{ for all } i \in A_p\} \cup \{\neg\psi_q \mid \neg\langle\langle B_q \rangle\rangle X \psi_q \in \Delta^{\{\vec{k}\}}, \text{neg}(\sigma) = q, Ag_\theta^+ \setminus B_q \subseteq N(\sigma)\} \quad (2)$$

where $\Gamma_\sigma^{\{\vec{k}+1\}} = \{\top\}$ if the set on the RHS of (2) is empty.

- If the pre-tableau contains a pre-state $\Gamma^{\{\vec{m}\}} = \Gamma_\sigma^{\{\vec{k}+1\}}$ s.t. $mk(\Gamma^{\{\vec{m}\}}) = mk(\Gamma_\sigma^{\{\vec{k}+1\}})$, then
 - rename $\Gamma^{\{\vec{m}\}}$ and all pre-states and states $\Theta_i^{\{\vec{m}+i\}}$ s.t. $\Gamma^{\{\vec{m}\}} \mapsto \Theta_i^{\{\vec{m}+i\}}$ as $\Gamma^{\{\vec{m}\} \cup \{\vec{k}+1\}}$ and $\Theta_i^{\{\vec{m}+i\} \cup \{\vec{k}+i+1\}}$. Then, set $\Delta^{\{\vec{k}\}} \xrightarrow{\sigma} \Gamma^{\{\vec{m}\} \cup \{\vec{k}+1\}}$;
- else,
 - add $\Gamma_\sigma^{\{\vec{k}+1\}}$ as a new pre-state to the pre-tableau;
 - set $\Delta^{\{\vec{k}\}} \xrightarrow{\sigma} \Gamma_\sigma^{\{\vec{k}+1\}}$ and $mk(\Gamma_\sigma^{\{\vec{k}+1\}}) = mk(\Delta^{\{\vec{k}\}})$.

Also in this case, the rule \mathbf{Next}_I^{uis} is obtained by omitting indexes from states and pre-states in $\mathbf{Next}_I^{uis, sync}$. Now we consider the epistemic operator. Similarly to the above, we have to introduce a distinct case for each of $ECGM^{uis}$ and $ECGM^{uis, sync}$. Moreover, for each case we have to introduce two rules that apply at different stages of the construction phase.

Definition 19 (Rule \mathbf{K}_1^{uis}) Given a state Δ s.t. \mathbf{K}_1^{uis} has not been applied to Δ , for every $\neg K_i \phi \in \Delta$,

- if $\langle\langle Ag_\theta \rangle\rangle F(\neg\phi \wedge \Delta|_{K_i}) \notin mk(\Delta)$, then add $\Gamma_{i+1} = mk(\Delta) \cup \{\langle\langle Ag_\theta \rangle\rangle F(\neg\phi \wedge \Delta|_{K_i})\}$ if it is consistent, where i is the greatest index s.t. Γ_i appears in the pre-tableau;
- set $mk(\Gamma_{i+1}) = \text{undef}$.

Intuitively, when Δ contains an epistemic possibility $\neg K_i \phi$ and the root of Δ does not anticipate to fulfil this possibility (i.e., $\langle\langle Ag_\theta \rangle\rangle F(\neg\phi \wedge \Delta|_{K_i}) \notin mk(\Delta)$), then by \mathbf{K}_1^{uis} we attempt to construct a branch of the pre-tableau that satisfies $\neg K_i \phi$ by adding a pre-state Γ_{i+1} as above. To define the corresponding rule for $ECGM^{uis, sync}$ we

introduce the notion of a *loop* on a state Δ , that is, a sequence $\Delta_0, \Gamma_1, \Delta_1 \dots, \Gamma_n, \Delta_n$ s.t. (i) $\Delta = \Delta_0 = \Delta_n$; (ii) $\Gamma_{i+1} \in \mathbf{pre-states}(\Delta_i)$; (iii) $\Delta_i \in \mathbf{states}(\Gamma_i)$; and (iv) $\Delta_i \neq \Delta$ for $0 < i < n$. Let n be the *length* of the loop. We denote the set of lengths of all loops for a state Δ as $loop(\Delta)$. Also, for $\neg K_i \phi \in \Delta^{\{\vec{k}\}}$, we use ζ as a shorthand for $\neg\phi \wedge \Delta^{\{\vec{k}\}}|_{K_i}$.

Definition 20 (Rule $\mathbf{K}_1^{uis, sync}$) Given a state $\Delta^{\{\vec{k}\}}$ s.t. $\mathbf{K}_1^{uis, sync}$ has not been applied to $\Delta^{\{\vec{k}\}}$, for every $\neg K_i \phi \in \Delta^{\{\vec{k}\}}$, $Y \subseteq \{\vec{k}\}$, and $W \subseteq loop(\Delta^{\{\vec{k}\}})$,

- if $\eta = \bigwedge_{y \in Y} (\langle\langle Ag_\theta \rangle\rangle X)^y (\zeta \wedge \langle\langle \emptyset \rangle\rangle G(\zeta \rightarrow \bigwedge_{w \in W} (\langle\langle Ag_\theta \rangle\rangle X)^w \zeta)) \notin \Delta^{\{\vec{k}\}}$, then add $\Gamma_{i+1}^{\{0\}} = mk(\Delta^{\{\vec{k}\}}) \cup \{\eta\}$ if it is consistent, where i is the smallest index s.t. $\Gamma_i^{\{0\}}$ appears in the pre-tableau;
- set $mk(\Gamma_{i+1}^{\{0\}}) = \text{undef}$.

The *rationale* for $\mathbf{K}_1^{uis, sync}$ comes from the fact that satisfiability w.r.t. $ECGM^{sync, uis}$ needs stronger conditions. Indeed, we have to ensure that for every $\neg K_i \phi$ in $\Delta^{\{\vec{k}\}}$ there exists a $\Delta^{\{\vec{k}\}}$ s.t. $\neg\phi \in \Delta^{\{\vec{k}\}}$. In particular, $\Delta^{\{\vec{k}\}}$ and $\Delta^{\prime\{\vec{k}\}}$ have to share the same time-stamp $\{\vec{k}\}$.

The construction phase for satisfiability w.r.t. $ECGM^{uis}$ (resp. $ECGM^{uis, sync}$) goes then as follows. As customary, we start with the initial pre-state $\Gamma_0^{\{0\}} = \{\theta\}$, which contains only the input formula. Then, we alternately apply \mathbf{SR}_I and \mathbf{Next}_I (we omit superscripts to refer to both cases) until these rules are no longer applicable. This procedure builds a purely temporal pre-tableau, i.e., a pre-tableau whose \rightarrow -edges are marked by joint actions only. At this point, we apply \mathbf{K}_1^{uis} (resp. $\mathbf{K}_1^{uis, sync}$). If a new initial pre-states $\Gamma_i^{\{0\}}$ is created, then we repeat the procedure above. We proceed in this way until none of the rules is any longer applicable. Finally, we apply the rules \mathbf{K}_2^{uis} (resp. $\mathbf{K}_2^{uis, sync}$).

Definition 21 (Rule $\mathbf{K}_2^{uis, sync}$) Given a state $\Delta^{\{\vec{k}\}}$ s.t. $\mathbf{K}_2^{uis, sync}$ has not been applied to $\Delta^{\{\vec{k}\}}$, for every $\neg K_i \phi \in \Delta^{\{\vec{k}\}}$ and state $\Delta^{\prime\{\vec{k}\}}$,

- if $mk(\Delta^{\prime\{\vec{k}\}}) = mk(\Delta^{\{\vec{k}\}})$ and $\{\neg\phi\} \cup \Delta^{\{\vec{k}\}}|_{K_i} \subseteq \Delta^{\prime\{\vec{k}\}}$, then set $\Delta^{\{\vec{k}\}} \xrightarrow{\neg K_i \phi} \Delta^{\prime\{\vec{k}\}}$.

The rule \mathbf{K}_2^{uis} is obtained from $\mathbf{K}_2^{uis, sync}$ by omitting time-stamps, so we can relate states with different time-stamps. Notice that neither \mathbf{K}_2^{uis} nor $\mathbf{K}_2^{uis, sync}$ creates new pre-states. Instead, they link a state containing an epistemic possibility with a state that fulfils that possibility. Also in the present case we can prove that the construction phase terminates after a finite number of steps. We conclude this section with an example.

Example 2 In Fig. 3 we reported a fragment of the pre-tableau for the ATEL-formula $\theta_1 = \langle\langle 1 \rangle\rangle F \neg K_2 p \wedge K_1 \langle\langle 2 \rangle\rangle G p$. We now observe that it also corresponds to the fragment of the pre-tableau obtainable by applying rules \mathbf{SR}_I^{uis} and \mathbf{Next}_I^{uis} only, modulo the mark mk . Since state

$\Delta_{1,1}$ in Fig. 3 contains an epistemic possibility, we apply rule \mathbf{K}_1^{uis} and create a new initial pre-state $\Gamma'_1 = \{\theta_1, \langle\langle 1, 2 \rangle\rangle F(\neg p \wedge \neg K_2 p)\}$. The fragment of the pre-tableau for Γ'_1 is shown in Fig. 2. ■

5.2 Elimination Phase

For I -satisfiability the elimination phase makes use of most of the rules for S -satisfiability, with some exceptions. First, by **PR** we remove all pre-states and re-direct incoming and outgoing \rightarrow -edges. Further, we proceed in stages and at each stage we apply one of the elimination rules **E1**, **E3**, and the following new rule:

E2^{uis, sync} If for some $\neg K_i \phi \in \Delta^{\vec{k}}$ there is no $\Delta'^{\vec{k}} \in \mathcal{T}_n^\theta$ s.t. $mk(\Delta'^{\vec{k}}) = mk(\Delta^{\vec{k}})$, $mk(\Delta'^{\vec{k}}) \Rightarrow^* \Delta'^{\vec{k}}$, and $\Delta^{\vec{k}} \xrightarrow{\neg K_i \phi} \Delta'^{\vec{k}}$, then $\mathcal{T}_{n+1}^\theta = \mathcal{T}_n^\theta \setminus \{\Delta^{\vec{k}}\}$.

This means that epistemic possibilities have to be satisfied by states that have the same time-stamp and are reachable from an initial state. The corresponding rule **E2^{uis}** can be obtained by omitting time-stamps. Finally, for synchronous ECGM only, we have also to consider the following rule to manipulate time-stamps:

Et_I If $\mathcal{T}_{n+1}^\theta = \mathcal{T}_n^\theta \setminus \{\Delta^{\vec{k}}\}$ and in \mathcal{T}_n^θ it was the case that $\Delta^{\vec{k}} \Rightarrow^* \Delta'^{\vec{m}}$, then in \mathcal{T}_{n+1}^θ rename $\Delta'^{\vec{m}}$ as $\Delta'^{\vec{m} \setminus \vec{k}}$ (remove $\Delta'^{\vec{m}}$ if $\vec{m} \setminus \vec{k} = \emptyset$).

The elimination phase terminates when it is no longer possible to apply any of the elimination rules above. As in the case of S -satisfiability, the structure obtained at the end of the elimination phase is the tableau \mathcal{T}^θ , which is said to be open if $\theta \in \Delta$, for some $\Delta \in \mathcal{T}^\theta$ s.t. $\Delta \in \mathbf{states}(\Gamma_i)$, for some initial Γ_i . Hereafter we outline the soundness and completeness proof of the tableau procedure for I -satisfiability.

5.3 Soundness and Completeness

The soundness proof for the tableau-based decision procedure amounts to the following result, analogous to Theorem 4.

Theorem 9 *If θ is I -satisfiable, then \mathcal{T}^θ is open.*

The structure of the proof is similar to Theorem 4, i.e., we show that the elimination rules do not remove satisfiable state. However, we need to consider different elimination rules. Hereafter we briefly consider the case of **E2^{uis}**. Suppose that $\Delta' \in \mathcal{T}_n^\theta$ is S -satisfiable and reachable from Δ , also $\neg K_i \phi \in \Delta'$, and there exists a sequence $\Theta_0 \rightarrow \dots \rightarrow \Theta_n$ of states s.t. $\Theta_0 = \Delta$, $\Theta_n = \Delta'$, and each Θ_i is S -satisfiable (the existence of such a sequence follows from Lemma 5.2 in (Goranko and Shkatov 2009a)). In particular, we can deduce that there exists an ECGM \mathcal{P} with a unique initial state, and $s_0 \in \mathcal{P}$ s.t. $(\mathcal{P}, s_0) \models \Theta_0 = \Delta$, $s_i \rightarrow s_{i+1}$ and $(\mathcal{P}, s_i) \models \Theta_i$. Since $\neg K_i \phi \in \Delta' = \Theta_n$, there exists $s' \in \mathcal{P}$ s.t. $s' \sim_i s$ and $(\mathcal{P}, s') \models \neg \phi \wedge \Delta|_{K_i}$. By the uniqueness of the initial state, we have that s' is reachable from s_0 . Hence $(\mathcal{P}, s_0) \models \langle\langle Ag_\theta \rangle\rangle F(\neg \phi \wedge \Delta|_{K_i})$. As a result we obtain that $\Gamma_{i+1} = \Delta \cup \{\langle\langle Ag_\theta \rangle\rangle F(\neg \phi \wedge \Delta|_{K_i})\}$ is also I -satisfiable, and by Lemma 5 some Δ_{i+1} saturating Γ_{i+1} is also I -satisfiable. Now, we remark that if Θ is

satisfiable and $\Delta \rightarrow^* \Theta$ then $\Delta_{i+1} \rightarrow^* \Theta$. In particular, Δ' is reachable from Δ_{i+1} . Moreover, since $\langle\langle Ag_\theta \rangle\rangle F(\neg \phi \wedge \Delta|_{K_i}) \in \Delta_{i+1}$ and Δ_{i+1} is I -satisfiable, the eventuality $\langle\langle Ag_\theta \rangle\rangle F(\neg \phi \wedge \Delta|_{K_i})$ is realized in some Δ'' reachable from Δ_{i+1} . As a consequence, when we apply rule **K₂^{uis}** we obtain that $\Delta' \xrightarrow{\neg K_i \phi} \Delta''$ and Δ' cannot be eliminated by an application of **E2^{uis}**. As a final result, we have that the I -satisfiable set Δ containing θ is never removed from \mathcal{T}^θ , thus it remains open.

To prove the completeness of the tableaux for I -satisfiability, we have to prove the following result, corresponding to Theorem 7.

Theorem 10 (Completeness) *If \mathcal{T}^θ is open, then θ is I -satisfiable.*

Also in this case, the proof amounts to build an ECGHS \mathcal{H}_θ starting from the open tableau \mathcal{T}^θ . Moreover, \mathcal{H}_θ has to have a unique initial state, or being synchronous. The key remark here is that we can restrict the construction of temporal components to states having the same mark mk . Thus, we state the following result, from which Theorem 10 immediately follows.

Lemma 11 *If \mathcal{T}^θ is open, then there exists an ECGHS \mathcal{H}_θ that I -satisfies θ .*

Complexity. We briefly discuss the complexity of the procedure above and notice that it is non-elementary, as in the worst case we have to reinitiate the tableau construction for each epistemic possibility to be fulfilled. We are not aware of any contribution to the complexity of I -satisfiability for ATEL formulas in either $ECGM^{sync}$ or $ECGM^{uis, sync}$. So, we here provide an upper bound for this problem.

6 Conclusions

In this paper we presented a tableau-based method for deciding the satisfiability of ATEL formulas in a number of settings. We considered satisfiability in some initial state as opposed to any state of the model. We argued that the former problem is of interest for verification purposes, while the latter notion is typical of logical investigations. Further, we considered the assumptions of synchronicity and of a unique initial state. These are standard conditions in the literature on Interpreted Systems and their impact on the satisfiability problem has been assessed in relation with various temporal epistemic logics, based on both linear and branching time (Halpern and Vardi 1986; 1989). This paper aimed at extending this type of investigations to alternating time epistemic logic.

There is a number of possible extensions for the present results. In one direction, other assumptions from Interpreted Systems, such as perfect knowledge and no learning, can be imposed on ECGM and corresponding decision procedures developed. It is well-known that the decision problem under these conditions becomes particularly hard (Halpern and Vardi 1986; 1989). So, incremental tableau techniques can contribute to alleviate the decision task for such cases, at least on average, since we do not have to build all formula types as in top-down tableaux. On another direction, notions such as imperfect information and commitment in strategies

have been considered in relation to ATL. It would be of interest to check their impact also w.r.t. satisfiability of ATEL formulas. In yet another direction, modalities for group knowledge can be added to ATEL. This extension has not been considered here, so as not to make the formal language too cumbersome. We leave all these issues for future research.

References

- Agotnes, T.; Goranko, V.; and Jamroga, W. 2007. Alternating-time temporal logics with irrevocable strategies. In *Proceedings of the 11th Conference on Theoretical Aspects of Rationality and Knowledge, TARK '07*, 15–24. New York, NY, USA: ACM.
- Ajspur, M.; Goranko, V.; and Shkatov, D. 2013. Tableau-based decision procedure for the multiagent epistemic logic with all coalitional operators for common and distributed knowledge. *Logic Journal of the IGPL* 21(3):407–437.
- Alur, R.; Henzinger, T. A.; and Kupferman, O. 2002. Alternating-time temporal logic. *J. ACM* 49(5):672–713.
- Baier, C., and Katoen, J.-P. 2008. *Principles of Model Checking (Representation and Mind Series)*. The MIT Press.
- Belardinelli, F. 2014. Satisfiability of alternating-time temporal epistemic logic through tableaux. In *Proceedings of the 2nd International Workshop on Strategic Reasoning*.
- Blackburn, P.; de Rijke, M.; and Venema, Y. 2001. *Modal Logic*, volume 53 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press.
- Bulling, N.; Dix, J.; and Jamroga, W. 2010. Model checking logics of strategic ability: Complexity*. In Dastani, M.; Hindriks, K. V.; and Meyer, J.-J. C., eds., *Specification and Verification of Multi-agent Systems*. Springer US. 125–159.
- D’Agostino, M.; Gabbay, D. M.; Hhnle, R.; and Posegga, J., eds. 1999. *Handbook of Tableau Methods*. Springer.
- Fagin, R.; Halpern, J.; Moses, Y.; and Vardi, M. 1995. *Reasoning About Knowledge*. The MIT Press.
- Gammie, P., and van der Meyden, R. 2004. MCK: Model checking the logic of knowledge. In *Proceedings of 16th International Conference on Computer Aided Verification (CAV’04)*, volume 3114 of *LNCS*, 479–483. Springer-Verlag.
- Goranko, V., and Jamroga, W. 2004. Comparing semantics of logics for multi-agent systems. *Synthese* 139(2):241–280.
- Goranko, V., and Shkatov, D. 2009a. Tableau-based decision procedures for logics of strategic ability in multiagent systems. *ACM Trans. Comput. Log.* 11(1).
- Goranko, V., and Shkatov, D. 2009b. Tableau-based decision procedure for full coalitional multiagent temporal-epistemic logic of linear time. In Sierra, C.; Castelfranchi, C.; Decker, K. S.; and Sichman, J. S., eds., *AAMAS (2)*, 969–976. IFAA-MAS.
- Goranko, V., and Shkatov, D. 2009c. Tableau-based decision procedure for the full coalitional multiagent logic of branching time. In Baldoni, M.; Baroglio, C.; Bentahar, J.; Boella, G.; Cossentino, M.; Dastani, M.; Dunin-Keplicz, B.; Fortino, G.; Gleizes, M. P.; Leite, J.; Mascardi, V.; Padget, J. A.; Pavón, J.; Polleres, A.; Fallah-Seghrouchni, A. E.; Torroni, P.; and Verbrugge, R., eds., *MALLOW*, volume 494 of *CEUR Workshop Proceedings*. CEUR-WS.org.
- Goranko, V., and van Drimmelen, G. 2006. Complete axiomatization and decidability of alternating-time temporal logic. *Theor. Comput. Sci.* 353(1-3):93–117.
- Halpern, J., and Vardi, M. 1986. The complexity of reasoning about knowledge and time. In *ACM Symposium on Theory of Computing (STOC ’86)*, 304–315. Baltimore, USA: ACM Press.
- Halpern, J., and Vardi, M. 1989. The complexity of reasoning about knowledge and time 1: lower bounds. *Journal of Computer and System Sciences* 38(1):195–237.
- Halpern, J.; van der Meyden, R.; and Vardi, M. 2004. Complete axiomatizations for reasoning about knowledge and time. *SIAM Journal on Computing* 33(3):674–703.
- Jamroga, W., and van der Hoek, W. 2004. Agents that know how to play. *Fundam. Inform.* 63(2-3):185–219.
- Kacprzak, M.; Nabialek, W.; Niewiadomski, A.; Penczek, W.; Pórola, A.; Szreter, M.; Wozna, B.; and Zbrzezny, A. 2008. Verics 2007 - a model checker for knowledge and real-time. *Fundamenta Informaticae* 85(1-4):313–328.
- Lomuscio, A.; Qu, H.; and Raimondi, F. 2009. MCMAS: A model checker for the verification of multi-agent systems. In Bouajjani, A., and Maler, O., eds., *CAV*, volume 5643 of *Lecture Notes in Computer Science*, 682–688. Springer.
- Meyer, J.-J. C., and Hoek, W. 1995. *Epistemic Logic for AI and Computer Science*, volume 41 of *Cambridge Tracts in Theoretical Computer Science*. Cambridge University Press.
- van der Hoek, W., and Wooldridge, M. 2003. Cooperation, knowledge, and time: Alternating-time temporal epistemic logic and its applications. *Studia Logica* 75(1):125–157.
- van Harmelen, F.; Lifschitz, V.; and Porter, B. 2007. *Handbook of Knowledge Representation*. San Diego, USA: Elsevier Science.
- Walther, D.; Lutz, C.; Wolter, F.; and Wooldridge, M. 2006. Atl satisfiability is indeed exptime-complete. *J. Log. Comput.* 16(6):765–787.
- Walther, D. 2005. Atel with common and distributed knowledge is exptime-complete. In *Proceedings of Methods for Modalities 4*.
- Wooldridge, M. 2000. Computationally Grounded Theories of Agency. In *Proc. of ICMAS*. IEEE Press. 13–22.