



<http://www.diva-portal.org>

This is the published version of a paper published in *ACM Transactions of Computational Logic*.

Citation for the original published paper (version of record):

Goranko, V., Cerrito, S., David, A. (2015)

Optimal Tableaux Method for Constructive Satisfiability Testing and Model Synthesis in the Alternating-time Temporal Logic ATL+.

ACM Transactions of Computational Logic, 17(1): 4

<http://dx.doi.org/10.1145/2811261>

Access to the published version may require subscription.

N.B. When citing this work, cite the original published paper.

Permanent link to this version:

<http://urn.kb.se/resolve?urn=urn:nbn:se:su:diva-123439>

Optimal Tableau Method for Constructive Satisfiability Testing and Model Synthesis in the Alternating-Time Temporal Logic ATL^+

SERENELLA CERRITO and AMÉLIE DAVID, Laboratoire IBISC - Université

Evry Val-d'Essonne, France

VALENTIN GORANKO, Department of Philosophy, Stockholm University and Department of Mathematics, University of Johannesburg (visiting professor)

We develop a sound, complete, and practically implementable tableau-based decision method for constructive satisfiability testing and model synthesis for the fragment ATL^+ of the full alternating-time temporal logic ATL^* . The method extends in an essential way a previously developed tableau-based decision method for ATL and works in $2EXPTIME$, which is the optimal worst-case complexity of the satisfiability problem for ATL^+ . We also discuss how suitable parameterizations and syntactic restrictions on the class of input ATL^+ formulas can reduce the complexity of the satisfiability problem.

Categories and Subject Descriptors: F.4.1 [Mathematical Logic and Formal Languages]: Mathematical Logic—Computational logic; Mechanical theorem proving; F.4.3 [Mathematical Logic and Formal Languages]: Formal Languages—Decision problems

General Terms: Theory

Additional Key Words and Phrases: Alternating-time temporal logics, ATL^+ , decision procedure, model synthesis, satisfiability, tableaux

ACM Reference Format:

Serenella Cerrito, Amélie David, and Valentin Goranko. 2015. Optimal tableau method for constructive satisfiability testing and model synthesis in the alternating-time temporal logic ATL^+ . *ACM Trans. Comput. Logic* 17, 1, Article 4 (October 2015), 34 pages.

DOI: <http://dx.doi.org/10.1145/2811261>

1. INTRODUCTION

The alternating-time temporal logic ATL^* was introduced and studied in Alur et al. [2002] as a multiagent extension of the branching-time temporal logic CTL^* , where the path quantifiers are generalized to “strategic quantifiers,” indexed with coalitions of agents A , and range existentially over collective strategies of A and then universally over all paths (computations) enabled by the selected collective strategy of A . ATL^* was proposed as logical framework for specification and verification of properties of open systems modeled as concurrent game models, in which all agents effect state transitions collectively by taking simultaneous actions at each state. The language of ATL^* allows expressing statements of the type “Coalition A has a collective strategy to guarantee the satisfaction of the objective Φ on every play enabled by that strategy.” The syntactic fragment ATL of ATL^* allows only state formulae, where all occurrences of temporal

This article is an essentially extended version of Cerrito et al. [2014].

Authors' addresses: S. Cerrito and A. David, Laboratoire IBISC-IBGBI, 23 boulevard de France, 91037 Evry Cedex; emails: {serena, adavid}@ibisc.univ-evry.fr; V. Goranko, Department of Philosophy, SE - 10691 Stockholm, Sweden; email: valentin.goranko@philosophy.su.se.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies show this notice on the first page or initial screen of a display along with the full citation. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, to republish, to post on servers, to redistribute to lists, or to use any component of this work in other works requires prior specific permission and/or a fee. Permissions may be requested from Publications Dept., ACM, Inc., 2 Penn Plaza, Suite 701, New York, NY 10121-0701 USA, fax +1 (212) 869-0481, or permissions@acm.org.

© 2015 ACM 1529-3785/2015/10-ART4 \$15.00

DOI: <http://dx.doi.org/10.1145/2811261>

operators must be immediately preceded by strategic quantifiers. The fragment ATL^+ of ATL^* extends ATL by allowing any Boolean combinations of ATL objectives in the scope of a strategic quantifier. It is considerably more expressive than ATL, which is reflected in the high— 2EXPTIME —worst-case complexity lower bound of the satisfiability problem for ATL^+ (inherited from the lower bound for CTL^+ , see [Johannsen and Lange 2003]) as opposed to the EXPTIME -completeness of the satisfiability problem for ATL [van Drimmelen 2003; Walther et al. 2006]. The matching 2EXPTIME upper bound is provided by the automata-based method for deciding satisfiability in the full ATL^* developed in Schewe [2008].

The contribution of this article is the development of a sound, complete, and terminating tableau-based decision method for constructive satisfiability testing of ATL^+ formulae. We also claim that our approach is intuitive and conceptually simple, as well as practically implementable and even manually usable, despite the inherently high worst-case complexity of the problem. The tableau method presented here is based on the general methodology going back to Pratt [1980] (for PDL), Wolper [1985] (for LTL), and [Ben-Ari et al. 1983; Emerson and Halpern 1985] (for CTL), and further adapted for ATL in Goranko and Shkatov [2009] to which the reader is referred for more details. A recent implementation of such a method is reported in David [2013]. The tableau method for ATL^+ is an essential extension of the one for ATL because it has to deal with much more complex (and computationally expensive) path objectives that can be assigned to the agents. It is also rather different from the above-mentioned automata-based method in Schewe [2008].

The article is structured as follows. In Section 2, we offer brief technical preliminaries on concurrent game models, syntax, and semantics of ATL^* and ATL^+ . Section 3 develops the technical machinery needed for the presentation of the tableau method itself, in Section 4. In Section 5, we prove the soundness of the tableau method, whereas in Section 6 we prove its completeness and demonstrate with examples how satisfying models can be extracted from the final open tableau. We also estimate the worst-case complexity of the procedure. In Section 7, we offer a brief comparison with the automata-based method in Schewe [2008].

2. PRELIMINARIES

We assume that the reader has basic familiarity with the branching-time logic CTL^* (see, e.g., Emerson [1990]). Also, basic knowledge on ATL^* [Alur et al. 2002] and the tableaux-based decision procedure for ATL in Goranko and Shkatov [2009], on which this article builds, would be beneficial.

2.1. Concurrent Game Models, Strategies, and Co-strategies

For technical reasons that will become clearer later in the soundness and the completeness proofs, we define a more general, *nondeterministic* version of the concurrent game structure with respect to Alur et al. [2002]. For the moment, we can say that the basic idea is to avoid several definitions of the notion of Realization Witness Trees for very similar structures (models, tableaux, and Hintikka structures). Note that the very notion of tableau will be defined as a nondeterministic labeled CGS (see the beginning of Section 4).

Notation: Given a set X , we denote the power set of X by $\mathcal{P}(X)$.

Definition 2.1. A (*nondeterministic*) *Concurrent Game Structure* (CGS) is a tuple

$$S = (\mathbb{A}, \text{St}, \{\text{Act}_a\}_{a \in \mathbb{A}}, \{\text{act}_a\}_{a \in \mathbb{A}}, \text{out})$$

comprising:

- a finite, non-empty set of *players (agents)* $\mathbb{A} = \{1, \dots, k\}$
- a non-empty set of *states* St ,

—a set of actions $\text{Act}_a \neq \emptyset$ for each $a \in \mathbb{A}$.

For any $A \subseteq \mathbb{A}$ we denote $\text{Act}_A := \prod_{a \in A} \text{Act}_a$ and use σ_A to denote a tuple from Act_A .

In particular, $\text{Act}_{\mathbb{A}}$ is the set of all possible *action profiles* in \mathcal{S} .

—for each $a \in \mathbb{A}$, a map $\text{act}_a : \text{St} \rightarrow \mathcal{P}(\text{Act}_a) \setminus \{\emptyset\}$ defining for each state s the actions available to a at s ,

—a *transition relation* $\text{out} \subseteq \text{St} \times \text{Act}_{\mathbb{A}} \times \text{St}$.

Whenever $\langle s, \sigma_{\mathbb{A}}, s' \rangle \in \text{out}$, for $\sigma_{\mathbb{A}} = \langle \sigma_1, \dots, \sigma_k \rangle$, then $\sigma_a \in \text{act}_a(s)$ for every $a \in \mathbb{A}$. Given a pair $\langle s, \sigma_{\mathbb{A}} \rangle$, the set of states $s' \in \text{St}$ such that $\langle s, \sigma_{\mathbb{A}}, s' \rangle \in \text{out}$ is denoted $\text{out}(s, \sigma_{\mathbb{A}})$ and called the set of *successor (outcome) states of $\sigma_{\mathbb{A}}$ at s* .

When $\text{out}(s, \sigma_{\mathbb{A}})$ is a singleton, the CGS is said to be *deterministic*. In such cases, by a slight abuse of notation, we use $\text{out}(s, \sigma_{\mathbb{A}})$ to denote a state s' rather than the singleton $\{s'\}$.

Definition 2.2.

- (1) Given a set of formulae (of some language) Θ , a CGS \mathcal{S} with a state space St is *state-labeled by Θ* if there is a mapping $l : \text{St} \rightarrow \mathcal{P}(\Theta)$ assigning to every state in \mathcal{S} a set of formulae from Θ , called the *label* of that state.
- (2) A *Concurrent Game Model (CGM)* is a *deterministic* CGS state-labeled by a fixed set of atomic propositions Prop (i.e., a tuple $(\mathcal{M} = (\mathbb{A}, \text{St}, \{\text{Act}_a\}_{a \in \mathbb{A}}, \{\text{act}_a\}_{a \in \mathbb{A}}, \text{out}, \text{Prop}, L)$ where
 - $(\mathbb{A}, \text{St}, \{\text{Act}_a\}_{a \in \mathbb{A}}, \{\text{act}_a\}_{a \in \mathbb{A}}, \text{out})$ is a deterministic CGS,
 - Prop is a set of *atomic propositions*, and
 - $L : \text{St} \rightarrow \mathcal{P}(\text{Prop})$ is a *(state-)labeling function*.

CGMs represent multiagent discrete transition systems that function as follows. At any moment, the system is in a given state in which each agent selects an action from those available to her at that state. All agents execute their actions synchronously, and the combination of these actions together with the current state determines a transition to a unique successor state in the model. A *play* in a CGM is an infinite sequence of subsequent successor states; that is, an infinite sequence $s_0 s_1 \dots \in \text{St}^\omega$ of states such that for each $i \geq 0$ there exists an action profile $\sigma_{\mathbb{A}} = \langle \sigma_1, \dots, \sigma_k \rangle$ such that $\text{out}(s_i, \sigma_{\mathbb{A}}) = s_{i+1}$. A *history* is a finite prefix of a play. We denote by $\text{Plays}_{\mathcal{M}}$ and $\text{Hist}_{\mathcal{M}}$, respectively, the set of plays and set of histories in a CGM \mathcal{M} . For a state $s \in \text{St}$, we define $\text{Plays}_{\mathcal{M}}(s)$ and $\text{Hist}_{\mathcal{M}}(s)$ as the set of plays and set of histories with initial state s . Given a sequence of states λ , we denote by λ_0 its initial state, by λ_i its $(i + 1)$ th state, by $\lambda_{\leq i}$ the prefix $\lambda_0 \dots \lambda_i$ of λ , and by $\lambda_{\geq i}$ the suffix $\lambda_i \lambda_{i+1} \dots$ of λ . When $\lambda = \lambda_0 \dots \lambda_\ell$ is finite, we say that it has length ℓ and write $|\lambda| = \ell$. Furthermore, we put $\text{last}(\lambda) = \lambda_\ell$.

For any coalition $A \subseteq \mathbb{A}$, a given CGM \mathcal{M} and state $s \in \text{St}$, an *A-co-action* at s in \mathcal{M} is a mapping $\text{Act}_A^c : \text{Act}_A \rightarrow \text{Act}_{\mathbb{A} \setminus A}$ that assigns to every collective action of A at the state s a collective action at s for the complementary coalition $\mathbb{A} \setminus A$.

We use $\text{act}_A(s)$ to denote the set of all A -actions that can be played by the coalition A at state s (i.e., $\text{act}_A(s) = \prod_{a \in A} \text{act}_a(s)$). We also use $\text{act}_A^c(s)$ to denote the set of all A -co-actions available at state s and σ_A^c for an element of this set.

A (*perfect recall*) *strategy* for an agent a in \mathcal{M} is a mapping $F_a : \text{Hist}_{\mathcal{M}} \rightarrow \text{Act}_a$ such that for all $h \in \text{Hist}_{\mathcal{M}}$ we have $F_a(h) \in \text{act}_a(\text{last}(h))$. Intuitively, it assigns an admissible action for agent a after any history h of the game. We denote by $\text{Strat}_{\mathcal{M}}(a)$ the set of all strategies of agent a . A (collective) strategy of a set (*coalition*) of agents $A \subseteq \mathbb{A}$ is a tuple $(F_a)_{a \in A}$ of strategies, one for each agent in A . When $A = \mathbb{A}$, this is called a *strategy profile*. We denote by $\text{Strat}_{\mathcal{M}}(A)$ the set of collective strategies of coalition A . A play $\lambda \in \text{Plays}_{\mathcal{M}}$ is *consistent with a collective strategy $F_A \in \text{Strat}_{\mathcal{M}}(A)$* if for every $i \geq 0$ there exists an action profile $\sigma_{\mathbb{A}} = \langle \sigma_1, \dots, \sigma_k \rangle$ such that $\text{out}(\lambda_i, \sigma_{\mathbb{A}}) = \lambda_{i+1}$ and $\sigma_a = F_a(\lambda_{\leq i})$

for all $a \in A$. The set of plays with initial state s that are consistent with F_A is denoted $\text{Plays}_{\mathcal{M}}(s, F_A)$.

Likewise, a (*perfect-recall*) A -co-strategy in \mathcal{M} for a coalition of agents A (possibly reduced to just one agent a) is a mapping $F_{A \setminus A} : \text{Hist}_{\mathcal{M}} \times \text{Strat}_{\mathcal{M}}(A) \rightarrow \text{Act}_{A \setminus A}$ that assigns to each $h \in \text{Hist}_{\mathcal{M}}$ and every collective strategy $F_A \in \text{Strat}_{\mathcal{M}}(A)$ an A -co-action $F_{A \setminus A}(h, F_A) \in \text{act}_A^c(\text{last}(h))$.

2.2. The Logic ATL* and Fragments

The logic ATL* is a multiagent extension of CTL* with *strategic quantifiers* $\langle\langle A \rangle\rangle$ indexed with coalitions A of agents. There are two types of formulae in ATL*: *state formulae*, that are evaluated at states, and *path formulae*, that are evaluated on plays. To simplify the presentation, we work with formulae in negation normal form over a fixed set Prop of atomic propositions and primitive temporal operators *Always* \square and *Until* \mathcal{U} . The syntax of the full language ATL* and its fragments ATL⁺ and ATL can then be defined as follows, where $l \in \text{Prop} \cup \{\neg p \mid p \in \text{Prop}\}$ is a literal, \mathbb{A} is a fixed finite set of agents, and $A \subseteq \mathbb{A}$:

$$\text{State formulae: } \varphi := l \mid (\varphi \vee \varphi) \mid (\varphi \wedge \varphi) \mid \langle\langle A \rangle\rangle \Phi \mid \llbracket A \rrbracket \Phi \quad (1)$$

$$\text{ATL* -path formulae: } \Phi := \varphi \mid \bigcirc \Phi \mid \square \Phi \mid (\Phi \mathcal{U} \Phi) \mid (\Phi \vee \Phi) \mid (\Phi \wedge \Phi) \quad (2)$$

$$\text{ATL+ -path formulae: } \Phi := \varphi \mid \bigcirc \varphi \mid \square \varphi \mid (\varphi \mathcal{U} \varphi) \mid (\Phi \vee \Phi) \mid (\Phi \wedge \Phi) \quad (3)$$

$$\text{ATL-path formulae: } \Phi := \bigcirc \varphi \mid \square \varphi \mid (\varphi \mathcal{U} \varphi) \quad (4)$$

Note that the state formulae have the same definition but define different sets in all three cases. To keep the notation lighter, we list the members of the set A in $\langle\langle A \rangle\rangle$ without using $\{\}$. When the length of a formula is measured, A will be assumed given by a bit vector. Parentheses will be omitted whenever safe, but they will be important when conjunctions and disjunctions are composed.

Hereafter, we use φ, ψ, η to denote arbitrary state formulae and Φ, Ψ to denote path formulae. By an ATL⁺ formula, we mean by default a *state* formula of ATL⁺; likewise for ATL. We define $\top := p \vee \neg p$, $\perp := \neg \top$ and the temporal operators *Sometime* \diamond by $\diamond \varphi := \top \mathcal{U} \varphi$ and *Release* \mathcal{R} by $\psi \mathcal{R} \varphi := \square \varphi \vee \varphi \mathcal{U} (\varphi \wedge \psi)$. Note that $\langle\langle A \rangle\rangle \psi \mathcal{R} \varphi$ and $\llbracket A \rrbracket \psi \mathcal{R} \varphi$ are ATL⁺ state formulae.

CTL* can be regarded as the fragment of ATL* where $\langle\langle \emptyset \rangle\rangle$ represents the path quantifier \forall and $\langle\langle \mathbb{A} \rangle\rangle$ represents \exists . The semantics of ATL* (inherited by ATL⁺) is defined in a given CGM \mathcal{M} , state $s \in \mathcal{M}$, and a path λ in \mathcal{M} just like the semantics of CTL*, with the added clauses for the strategic quantifiers:

- $\mathcal{M}, s \models p$, for any proposition $p \in \text{Prop}$, if and only if $p \in L(s)$.
- $\mathcal{M}, s \models \neg p$ if and only if $\mathcal{M}, s \not\models p$.
- $\mathcal{M}, s \models \varphi \wedge \psi$ if and only if $\mathcal{M}, s \models \varphi$ and $\mathcal{M}, s \models \psi$.
- $\mathcal{M}, s \models \varphi \vee \psi$ if and only if $\mathcal{M}, s \models \varphi$ or $\mathcal{M}, s \models \psi$.
- $\mathcal{M}, s \models \langle\langle A \rangle\rangle \Phi$ if and only if there exists an A -strategy F_A such that, for all computations λ consistent with F_A , $\mathcal{M}, \lambda \models \Phi$.
- $\mathcal{M}, s \models \llbracket A \rrbracket \Phi$ if and only if there exists an A -co-strategy F_A^c such that, for all computations λ consistent with F_A^c , $\mathcal{M}, \lambda \models \Phi$.
- $\mathcal{M}, \lambda \models \varphi$ if and only if $\mathcal{M}, \lambda_0 \models \varphi$.
- $\mathcal{M}, \lambda \models \bigcirc \varphi$ if and only if $\mathcal{M}, \lambda_{\geq 1} \models \varphi$.
- $\mathcal{M}, \lambda \models \square \varphi$ if and only if for all positions $i \geq 0$, $\mathcal{M}, \lambda_{\geq i} \models \varphi$.
- $\mathcal{M}, \lambda \models \varphi \mathcal{U} \psi$ if and only if there exists a position $i \geq 0$ where $\mathcal{M}, \lambda_{\geq i} \models \psi$ and for all positions $0 \leq j < i$, $\mathcal{M}, \lambda_{\geq j} \models \varphi$.




Fig. 1. A CGM.

- $\mathcal{M}, \lambda \models \Phi \wedge \Psi$ if and only if $\mathcal{M}, \lambda \models \Phi$ and $\mathcal{M}, \lambda \models \Psi$.
- $\mathcal{M}, \lambda \models \Phi \vee \Psi$ if and only if $\mathcal{M}, \lambda \models \Phi$ or $\mathcal{M}, \lambda \models \Psi$.

Valid, satisfiable, and equivalent formulae in ATL* are defined as usual. Here are some important equivalences in LTL [Emerson 1990] and in ATL* [Alur et al. 2002; Goranko and van Drimmelen 2006], used further:

- $\Box\Psi \equiv \Psi \wedge \Box\Psi$; $\Phi\mathcal{U}\Psi \equiv \Psi \vee (\Phi \wedge \Box(\Phi\mathcal{U}\Psi))$;
- $\langle\langle A \rangle\rangle\Box\Psi \equiv \Psi \wedge \langle\langle A \rangle\rangle\Box\Psi$; $\langle\langle A \rangle\rangle\Phi\mathcal{U}\Psi \equiv \Psi \vee (\Phi \wedge \langle\langle A \rangle\rangle\Box(\Phi\mathcal{U}\Psi))$;
- $\llbracket A \rrbracket\Box\Psi \equiv \Psi \wedge \llbracket A \rrbracket\Box\Psi$; $\llbracket A \rrbracket\Phi\mathcal{U}\Psi \equiv \Psi \vee (\Phi \wedge \llbracket A \rrbracket\Box(\llbracket A \rrbracket\Phi\mathcal{U}\Psi))$;
- $\llbracket A \rrbracket\Box\varphi \equiv \neg\langle\langle A \rangle\rangle\Box\neg\varphi \equiv \langle\langle \emptyset \rangle\rangle\Box\varphi$; $\langle\langle A \rangle\rangle\langle\langle B \rangle\rangle\Phi \equiv \langle\langle B \rangle\rangle\Phi$;
- For every state formula φ : $\langle\langle A \rangle\rangle(\varphi \wedge \Psi) \equiv \varphi \wedge \langle\langle A \rangle\rangle\Psi$, $\langle\langle A \rangle\rangle(\varphi \vee \Psi) \equiv \varphi \vee \langle\langle A \rangle\rangle\Psi$.

Remark 2.3. It is known [Alur et al. 2002] that, when restricted to ATL formulae, the semantics just presented (based on perfect-recall strategies) is equivalent to the semantics based on positional (or memoryless) strategies, where the prescribed actions only depend on the current state, not on the whole history. This is no longer the case for ATL⁺. For example, the formula $\langle\langle 1 \rangle\rangle\Diamond(p \wedge \langle\langle 1 \rangle\rangle\Diamond q) \rightarrow \langle\langle 1 \rangle\rangle(\Diamond p \wedge \Diamond q)$ in a two-agent language is valid in the semantics with perfect-recall strategies (which can be freely composed) but not in the semantics with positional strategies (which cannot be freely composed). Indeed, in the concurrent game model of Figure 1, the antecedent of the above implication, namely $\langle\langle 1 \rangle\rangle\Diamond(p \wedge \langle\langle 1 \rangle\rangle\Diamond q)$, is true at state S_0 no matter what strategy—perfect-recall or positional—is considered, whereas the consequent, namely $\langle\langle 1 \rangle\rangle(\Diamond p \wedge \Diamond q)$, is true at S_0 only with respect to perfect-recall strategies. To be more precise, with respect to the state S_0 only two cases of memoryless strategy F for player 1 are possible: $F(S_0) = a$ and $F(S_0) = b$. Since the strategy is positional, these actions would be applied every time the play reaches S_0 , and neither of them guarantees that the play will eventually visit both a state satisfying p and a state satisfying q . On the other hand, a perfect-recall strategy F such that $F(S_0) = a$ and $F(S_0S_1S_0) = b$ guarantees the satisfaction of the objective $\Diamond p \wedge \Diamond q$.

Here, we assume that the semantics is based on perfect-recall strategies.

The (*constructive*) *satisfiability decision problem* for ATL⁺ is defined as follows:

Given a state formula φ in ATL⁺, does there exist a CGM \mathcal{M} and a state s in \mathcal{M} such that $\mathcal{M}, s \models \varphi$? If so, construct such a satisfying pair (\mathcal{M}, s) .

Remark 2.4. There are three variants of the satisfiability problem: *tight*, where it is assumed that all agents in the model are mentioned in the formula; *loose* where just one additional agent, not mentioned in the formula is allowed in the model; and *general*, where any number of additional agents, not mentioned in the formula, are allowed in the model. These variants are really different, but the general satisfiability

is immediately reducible to the loose satisfiability by adding just one extra agent a to the language. Furthermore, this extra agent can be easily added superfluously to the formula, for example by adding a conjunct $\langle\langle a \rangle\rangle \circ \top$, thus reducing loose to tight satisfiability. So, hereafter, we only consider the tight satisfiability version. For further details and discussion on this issue, see Walther et al. [2006] and Goranko and Shkatov [2009].

3. DECOMPOSITION AND CLOSURE OF ATL⁺ FORMULAE

We partition the set of ATL⁺ formulae into *primitive* and *nonprimitive* formulae. The primitive formulae are \top , \perp , the literals, and all ATL⁺ *successor formulae* of the form $\langle\langle A \rangle\rangle \circ \psi$ or $\llbracket A \rrbracket \circ \psi$, where $A \subseteq \mathbb{A}$ and $A' \subset \mathbb{A}$, each with *successor component* ψ . The nonprimitive formulae are classified as α -, β -, and γ -formulae. An α -formula in our syntax is a conjunction $\varphi \wedge \psi$ with (conjunctive) α -components φ and ψ , plus the formulae of the form $\llbracket A \rrbracket \circ \psi$ whose α -components are both $\langle\langle \emptyset \rangle\rangle \circ \psi$; a β -formula is a disjunction $\varphi \vee \psi$ with (disjunctive) β -components φ and ψ . The rest of the nonprimitive formulae are classified as γ -formulae. That is, a γ -formula is one of the form $\llbracket A \rrbracket \Phi$ or $\langle\langle A \rangle\rangle \Phi$, where Φ is an ATL⁺ path formula whose main operator is not \circ and $A \neq \mathbb{A}$.

The need to introduce the new category of γ -formulae with respect to the partition of nonprimitive formulae into α - and β -classes done in Goranko and Shkatov [2009] is the following. In ATL, each strategic quantifier is necessarily followed by a temporal operator, and, for instance $\langle\langle A \rangle\rangle \square \varphi$ can be seen as an α -formula, whereas $\langle\langle A \rangle\rangle \varphi \mathcal{U} \psi$ can be seen as a β -formula. However, typical state formulae in ATL⁺ have the form $\langle\langle A \rangle\rangle (\Phi_1 \vee \Phi_2)$, $\langle\langle A \rangle\rangle (\Phi_1 \wedge \Phi_2)$, $\llbracket A \rrbracket (\Phi_1 \wedge \Phi_2)$, $\llbracket A \rrbracket (\Phi_1 \vee \Phi_2)$. Now, these four types of formulae cannot reasonably be classified as α - or β -formulae. Note, in particular, that the strategic quantifier $\langle\langle A \rangle\rangle$ in general distributes neither on \vee nor on \wedge , and the same applies to $\llbracket A \rrbracket$. Thus, a new category of γ -formulae is created containing also $\langle\langle A \rangle\rangle \square \varphi$ and $\langle\langle A \rangle\rangle \varphi \mathcal{U} \psi$ as special cases, and needing a special analysis.

Thus, α - and β -formulae will be decomposed in the tableau as usual, whereas the case of γ -formulae $\langle\langle A \rangle\rangle \Phi$ and $\llbracket A \rrbracket \Phi$ is special and needs extra work because their tableau decomposition will depend on the structure of Φ .

3.1. γ -Decomposition and γ -Components of γ -Formulae

We denote the set of ATL⁺ state formulae by ATL_s^+ and the set of ATL⁺ path formulae by ATL_p^+ . We define a γ -decomposition function $\text{dec} : \text{ATL}_p^+ \rightarrow \mathcal{P}(\text{ATL}_s^+ \times \text{ATL}_p^+)$ with the following intuitive meaning: For any $\Phi \in \text{ATL}_p^+$ and pair $\langle \psi, \Psi \rangle \in \text{dec}(\Phi)$, ψ is a state formula true at the current state and Ψ is a path formula expressing what must be true at the next state of a possible play starting at the current state. Thus, the set $\text{dec}(\Phi)$ is interpreted as a disjunction describing all possible types of paths starting from the current state and satisfying Φ .

We emphasize that, although the domain of dec is the whole set ATL_p^+ , dec will only be used to analyze Φ in the contexts $\langle\langle A \rangle\rangle \Phi$ and $\llbracket A \rrbracket \Phi$ (where Φ does not have \circ as main connective), and, as we will see, its role is just auxiliary to the rewriting of the (always quantified) γ -formulae in a special form useful to obtain a key property of our tableau calculus (see Lemma 3.1).

Base Cases:

★ $\text{dec}(\varphi) = \{ \langle \varphi, \top \rangle \}$, $\text{dec}(\circ \varphi) = \{ \langle \top, \varphi \rangle \}$ for any ATL⁺ state formula φ .

The other base cases derive from the well-known LTL equivalences listed in Section 2.2:

★ $\text{dec}(\square \varphi) = \{ \langle \varphi, \square \varphi \rangle \}$

★ $\text{dec}(\varphi \mathcal{U} \psi) = \{ \langle \varphi, \varphi \mathcal{U} \psi \rangle, \langle \psi, \top \rangle \}$.




Fig. 2. The three cases for disjunctive path objectives in a γ -formula.

Recursive Steps:

★ $\text{dec}(\Phi_1 \wedge \Phi_2) = \text{dec}(\Phi_1) \otimes \text{dec}(\Phi_2)$, where

$$\text{dec}(\Phi_1) \otimes \text{dec}(\Phi_2) := \{ \langle \psi_i \wedge \psi_j, \Psi_i \wedge \Psi_j \rangle \mid \langle \psi_i, \Psi_i \rangle \in \text{dec}(\Phi_1), \langle \psi_j, \Psi_j \rangle \in \text{dec}(\Phi_2) \}.$$

★ $\text{dec}(\Phi_1 \vee \Phi_2) = \text{dec}(\Phi_1) \cup \text{dec}(\Phi_2) \cup (\text{dec}(\Phi_1) \oplus \text{dec}(\Phi_2))$, where

$$\text{dec}(\Phi_1) \oplus \text{dec}(\Phi_2) :=$$

$$\{ \langle \psi_i \wedge \psi_j, \Psi_i \vee \Psi_j \rangle \mid \langle \psi_i, \Psi_i \rangle \in \text{dec}(\Phi_1), \langle \psi_j, \Psi_j \rangle \in \text{dec}(\Phi_2), \Psi_i \neq \top, \Psi_j \neq \top \}.$$

Note that the operations \otimes and \oplus are associative, up to logical equivalence.

The conjunctive case should be clear: Every path satisfying $\Phi_1 \wedge \Phi_2$ combines a type of path satisfying Φ_1 with a type of path satisfying Φ_2 . To understand the disjunctive case, note that, as will be seen in Section 4, the construction of the tableau is step-by-step. Therefore, for a given prestate under construction, when we have a formula of the form $\langle\langle A \rangle\rangle(\Phi_1 \vee \Phi_2)$, where, for instance $\Phi_1 = \square\varphi_1$ and $\Phi_2 = \square\varphi_2$, we do not know in advance which of $\square\varphi_1$ or $\square\varphi_2$ would be completed; so it is important to keep both possibilities at the current state, if possible. This idea is expressed by the use of $\text{dec}(\Phi_1) \oplus \text{dec}(\Phi_2)$ in the above union, where we keep both disjuncts true at the present state and delay the choice. This is why the state formulae ψ_i and ψ_j are connected by \wedge but the path formulae Ψ_i and Ψ_j are connected by \vee . Moreover, the \oplus operation avoids the construction of a pair $\langle \psi_i \wedge \psi_j, \Psi_i \vee \Psi_j \rangle$ where either Ψ_i or Ψ_j is \top , because that case would already be included in $\text{dec}(\Phi_1)$ or in $\text{dec}(\Phi_2)$. The three cases for paths satisfying the disjunction $\Phi_1 \vee \Phi_2$ can be illustrated by the picture in Figure 2.

Now, let $\zeta = \langle\langle A \rangle\rangle\Phi$ or $\zeta = \llbracket A \rrbracket\Phi$ be a γ -formula to be decomposed. Each pair $\langle \psi, \Psi \rangle \in \text{dec}(\Phi)$ is then converted to a γ -component $\gamma(\psi, \Psi)$ as follows:

$$\gamma(\psi, \Psi) = \psi \quad \text{if } \Psi = \top \tag{5}$$

$$\gamma(\psi, \Psi) = \psi \wedge \langle\langle A \rangle\rangle \circ \langle\langle A \rangle\rangle \Psi \quad \text{if } \zeta \text{ is of the form } \langle\langle A \rangle\rangle\Phi, \tag{6}$$

$$\gamma(\psi, \Psi) = \psi \wedge \llbracket A \rrbracket \circ \llbracket A \rrbracket \Psi \quad \text{if } \zeta \text{ is of the form } \llbracket A \rrbracket\Phi \tag{7}$$

Thus, the role of dec is to associate with any γ -formula ζ a set of formulae that are simpler in some precise sense, viz. its γ -components, so that ζ is equivalent to the disjunction of its γ -components. This key property is Item 3 of Lemma 3.1 (the first two items being just auxiliary claims), and it is the core distinction between the proposed calculus for ATL⁺ in this work and the tableau calculus for ATL in Goranko and Shkatov [2009].

LEMMA 3.1. *For any γ -formula $\Theta = \langle\langle A \rangle\rangle\Phi$ or $\Theta = \llbracket A \rrbracket\Phi$ of ATL⁺, the following properties hold:*

- (1) $\Phi \equiv \bigvee \{ \psi \wedge \circ \Psi \mid \langle \psi, \Psi \rangle \in \text{dec}(\Phi) \}$.
- (2) $\langle\langle A \rangle\rangle\Phi \equiv \bigvee \{ \langle\langle A \rangle\rangle(\psi \wedge \circ \Psi) \mid \langle \psi, \Psi \rangle \in \text{dec}(\Phi) \}$, and respectively, $\llbracket A \rrbracket\Phi \equiv \bigvee \{ \llbracket A \rrbracket(\psi \wedge \circ \Psi) \mid \langle \psi, \Psi \rangle \in \text{dec}(\Phi) \}$.
- (3) $\Theta \equiv \bigvee \{ \gamma(\psi, \Psi) \mid \langle \psi, \Psi \rangle \in \text{dec}(\Phi) \}$.

PROOF. Claim 1. We prove the claim by induction on the path formula Φ . It is equivalent to the following property $P(\Phi)$:

For every CGM \mathcal{M} and a play λ in it, $\mathcal{M}, \lambda \models \Phi$ if and only if there exists $\langle \psi, \Psi \rangle \in \text{dec}(\Phi)$ such that $\mathcal{M}, \lambda_0 \models \psi$ and $\mathcal{M}, \lambda_{\geq 1} \models \Psi$.

The base cases are $\Phi = \varphi$, $\Phi = \bigcirc\varphi$, $\Phi = \square\varphi$, and $\Phi = \varphi\mathcal{U}\psi$. For each of these, the property $P(\Phi)$ follows immediately from the definitions of dec and γ -components and—for the latter two cases—the well-known fixed point LTL equivalences for the temporal operators, listed at the end of Section 2.2.

For the inductive steps there are two cases to consider:

Case 1: $\Phi = \Phi_1 \wedge \Phi_2$. We have that:

$\mathcal{M}, \lambda \models \Phi$ if and only if

$\mathcal{M}, \lambda \models \Phi_1$ and $\mathcal{M}, \lambda \models \Phi_2$, if and only if (by the induction hypothesis):

- (i) there is $\langle \psi_1, \Psi_1 \rangle \in \text{dec}(\Phi_1)$, such that $\mathcal{M}, \lambda_0 \models \psi_1$ and $\mathcal{M}, \lambda_{\geq 1} \models \Psi_1$,
and
- (ii) there is $\langle \psi_2, \Psi_2 \rangle \in \text{dec}(\Phi_2)$ such that $\mathcal{M}, \lambda_0 \models \psi_2$ and $\mathcal{M}, \lambda_{\geq 1} \models \Psi_2$.

These two are the case if and only if

$\mathcal{M}, \lambda_0 \models \psi_1 \wedge \psi_2$ and $\mathcal{M}, \lambda_{\geq 1} \models \Psi_1 \wedge \Psi_2$, if and only if

$\mathcal{M}, \lambda_0 \models \psi$ and $\mathcal{M}, \lambda_{\geq 1} \models \Psi$ where $\psi = \psi_1 \wedge \psi_2$, $\Psi = \Psi_1 \wedge \Psi_2$ and $\langle \psi, \Psi \rangle \in \text{dec}(\Phi)$.

This completes the proof of $P(\Phi)$ for $\Phi = \Phi_1 \wedge \Phi_2$.

Case 2: $\Phi = \Phi_1 \vee \Phi_2$. We have that $\mathcal{M}, \lambda \models \Phi$ if and only if $\mathcal{M}, \lambda \models \Phi_1$ or $\mathcal{M}, \lambda \models \Phi_2$. By inductive hypotheses for Φ_1 and Φ_2 and from the fact that $\text{dec}(\Phi_1) \cup \text{dec}(\Phi_2) \subseteq \text{dec}(\Phi)$, we obtain the direction from left to right in property $P(\Phi)$. For the converse direction, we only need to consider the case that does not follow directly from the inductive hypotheses for Φ_1 and Φ_2 , viz. when there exists $\langle \psi, \Psi \rangle \in (\text{dec}(\Phi_1) \oplus \text{dec}(\Phi_2))$ such that $\mathcal{M}, \lambda_0 \models \psi$ and $\mathcal{M}, \lambda_{\geq 1} \models \Psi$. In this case, $\psi = \psi_1 \wedge \psi_2$ and $\Psi = \Psi_1 \vee \Psi_2$ for some $\langle \psi_1, \Psi_1 \rangle \in \text{dec}(\Phi_1)$ and $\langle \psi_2, \Psi_2 \rangle \in \text{dec}(\Phi_2)$ such that $\Psi_1 \neq \top$, $\Psi_2 \neq \top$. Suppose $\mathcal{M}, \lambda_{\geq 1} \models \Psi_1$. Since we also have $\mathcal{M}, \lambda_0 \models \psi_1$, by the inductive hypothesis for Φ_1 , it follows that $\mathcal{M}, \lambda \models \Phi_1$, hence $\mathcal{M}, \lambda \models \Phi$. Likewise, when $\mathcal{M}, \lambda_{\geq 1} \models \Psi_2$.

Claim 2. We consider the case of $\Theta = \langle\langle A \rangle\rangle\Phi$; the case of $\llbracket A \rrbracket\Phi$ is analogous. The implication from right to left of the claimed equivalence follows from Claim 1 and the monotonicity of $\langle\langle A \rangle\rangle$ (in the sense that if $\Psi \models \Phi$ then $\langle\langle A \rangle\rangle\Psi \models \langle\langle A \rangle\rangle\Phi$). For the converse direction, first recall that every ATL^+ path formula Ξ is a positive Boolean combination of subformulae of the types φ , $\bigcirc\varphi$, $\square\varphi$, $\varphi\mathcal{U}\psi$ where φ, ψ are ATL^+ state formulae. Let the set of these subformulae of Ξ be $S(\Xi)$. Now, we introduce some ad hoc notation for special sets of formulae in $S(\Xi)$ and their subformulae:

— $L(\Xi)$ is the set of all state formulae in $S(\Xi)$;

— $N(\Xi) := \{\varphi \mid \bigcirc\varphi \in S(\Xi)\}$;

— $B(\Xi) := \{\varphi \mid \square\varphi \in S(\Xi)\}$;

— $U(\Xi) := \{\varphi\mathcal{U}\psi \mid \varphi\mathcal{U}\psi \in S(\Xi)\}$;

— $U_1(\Xi) := \{\varphi \mid \varphi\mathcal{U}\psi \in S(\Xi)\}$;

— $U_2(\Xi) := \{\psi \mid \varphi\mathcal{U}\psi \in S(\Xi)\}$;

Without loss of generality, we can assume that Φ is in a DNF over the set of formulae in $S(\Phi)$ (i.e., $\Phi = \Phi_1 \vee \dots \vee \Phi_m$) where each Φ_i is a conjunction of formulae from $S(\Phi)$.

Now, to prove the implication from left to right, take any CGM \mathcal{M} and state s in it, such that $\mathcal{M}, s \models \langle\langle A \rangle\rangle\Phi$. Take and fix any collective strategy F_A of A such that $\mathcal{M}, \lambda \models \Phi$ for every play λ starting at s and consistent with F_A . We denote that set of plays by $\text{OutPath}(s, F_A)$. Then, for every play $\lambda \in \text{out}_{\mathcal{M}}(s, \sigma_A)$, we have that $\mathcal{M}, \lambda \models \Phi_i$ for some $i = 1, \dots, m$. Without restriction of generality, we can assume that the set of

Φ_i 's for which there is a $\lambda \in \text{out}_{\mathcal{M}}(s, \sigma_A)$ such that $\mathcal{M}, \lambda \models \Phi_i$ is $\{\Phi_1, \dots, \Phi_n\}$ for some $n \leq m$.

Let Φ_i be any of these. We associate with it a pair $\langle \psi_i, \Psi_i \rangle \in \text{dec}(\Phi_i)$ as follows. First, note that all formulae from $L(\Phi_i)$ and $B(\Phi_i)$ are true at s . Furthermore, let $E_i(s)$ be the subset of those formulae from $U_2(\Phi_i)$ that are true at s in \mathcal{M} . Thus, for every play $\lambda \in \text{OutPath}(s, F_A)$ satisfying Φ_i the following hold:

- (i) $\mathcal{M}, \lambda \models \varphi$ for each $\varphi \in L(\Phi_i)$.
- (ii) $\mathcal{M}, \lambda \models \bigcirc \varphi$ for each $\bigcirc \varphi \in S(\Phi_i)$.
- (iii) $\mathcal{M}, \lambda \models \varphi \wedge \bigcirc \square \varphi$ for each $\square \varphi \in S(\Phi)$.
- (iv) $\mathcal{M}, \lambda \models \psi$ for each $\psi \in E_i(s)$.
- (v) $\mathcal{M}, \lambda \models \varphi \wedge \bigcirc \varphi \mathcal{U} \psi$ for each $\psi \in U_2(\Phi_i) - E_i(s)$.

Now, suppose $\Phi_i = \Psi_{i1} \wedge \dots \wedge \Psi_{ik}$ for some $\Psi_{i1} \wedge \dots \wedge \Psi_{ik} \in S(\Phi)$. Then $\text{dec}(\Phi_i) = \text{dec}(\Psi_{i1}) \otimes \dots \otimes \text{dec}(\Psi_{ik})$. (Recall that the operations \otimes and \oplus are associative, up to logical equivalence, so there is no need to put parentheses.) Thus, for every $\langle \psi, \Psi \rangle \in \text{dec}(\Phi_i)$, ψ is a conjunction of all formulae from $L(\Phi_i) \cup B(\Phi_i)$ and, for every conjunct of Φ_i of the type $\varphi \mathcal{U} \psi$, at least one of the respective formulae coming from $U_1(\Phi_i)$ and $U_2(\Phi_i)$. We now select $\langle \psi_i, \Psi_i \rangle \in \text{dec}(\Phi_i)$ to be the one where the conjuncts taken from $U_2(\Phi_i)$ are exactly those in $E_i(s)$. Then we claim that, for every play $\lambda \in \text{OutPath}(s, F_A)$ satisfying Φ_i , it is the case that $\mathcal{M}, \lambda \models \psi_i \wedge \bigcirc \Psi_i$. Indeed, this follows from the list of properties (i–v) just given and from the definition of $\text{dec}(\Psi_{i1}) \otimes \dots \otimes \text{dec}(\Psi_{ik})$. Note further that if Ψ_i above is \top , then $\mathcal{M}, \lambda \models \psi_i \wedge \bigcirc \Psi_i$ for all paths λ starting at s , so we can assume without affecting what follows that no Ψ_i above is \top .

After having selected such a pair $\langle \psi_i, \Psi_i \rangle \in \text{dec}(\Phi_i)$ for each $\Phi_i \in \{\Phi_1, \dots, \Phi_n\}$, we use these n pairs (or, those of them for which $\Psi_i \neq \top$) to construct the pair $\langle \psi, \Psi \rangle \in \text{dec}(\Phi_1) \oplus \dots \oplus \text{dec}(\Phi_n)$ such that $\psi = \psi_1 \wedge \dots \wedge \psi_n$ and $\Psi = \Psi_1 \vee \dots \vee \Psi_n$.

Finally, we claim that, by virtue of the construction, $\mathcal{M}, \lambda \models \psi \wedge \bigcirc \Psi$ for every play $\lambda \in \text{OutPath}(s, F_A)$ satisfying Φ . Therefore, the strategy σ_A is a witness of the truth of $\mathcal{M}, s \models \langle\langle A \rangle\rangle(\psi \wedge \bigcirc \Psi)$; hence, $\mathcal{M}, s \models \bigvee \{ \langle\langle A \rangle\rangle(\psi \wedge \bigcirc \Psi) \mid \langle \psi, \Psi \rangle \in \text{dec}(\Phi) \}$. This completes the proof of the implication left-to-right of Claim 2.

Claim 3. This claim follows easily from Claim 2 by noting that:

- $\langle\langle A \rangle\rangle(\psi \wedge \bigcirc \square \Psi) \equiv \psi \wedge \langle\langle A \rangle\rangle \bigcirc \square \Psi \equiv \psi \wedge \langle\langle A \rangle\rangle \bigcirc \langle\langle A \rangle\rangle \square \Psi$, because ψ is a state formula. Note that the second equivalence is due to the fact that the semantics of $\langle\langle \rangle\rangle$ is based on perfect recall strategies, that can be composed. More precisely, it essentially assumes that any strategy at s ensuring that every successor satisfies $\langle\langle A \rangle\rangle \square \Psi$ can be composed with the family of strategies, one for every such successor s' witnessing the truth of $\langle\langle A \rangle\rangle \square \Psi$ on all plays starting at s' , into one perfect recall strategy that guarantees the truth of $\bigcirc \square \Psi$ on all plays starting at s . (This, in general, cannot be done if only positional strategies are considered because those applied at the different successors of s may interfere with each other.)
- Likewise, $\llbracket A \rrbracket(\psi \wedge \bigcirc \square \Psi) \equiv \psi \wedge \llbracket A \rrbracket \bigcirc \square \Psi \equiv \psi \wedge \llbracket A \rrbracket \bigcirc \llbracket A \rrbracket \square \Psi$.

Therefore, for each $\langle \psi, \Psi \rangle \in \text{dec}(\Phi)$ the γ -component $\gamma(\psi, \Psi)$ is equivalent to its respective disjunct on the right-hand side of Claim 2. \square

Example 3.2. We use two syntactically similar yet different running examples:

$$\theta = \langle\langle 1 \rangle\rangle(p \mathcal{U} q \vee \square q) \wedge \langle\langle 2 \rangle\rangle(\diamond p \wedge \square \neg q)$$

and

$$\vartheta = \langle\langle 1 \rangle\rangle(p \mathcal{U} q \vee \square q) \wedge \llbracket 2 \rrbracket(\diamond p \wedge \square \neg q).$$

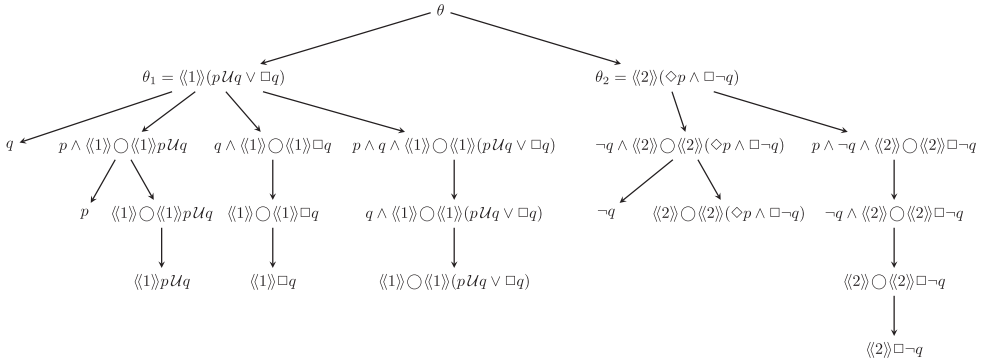


Fig. 3. Closure of the formula $\theta = \langle\langle 1 \rangle\rangle(pUq \vee \Box q) \wedge \langle\langle 2 \rangle\rangle(\Diamond p \wedge \Box \neg q)$.

First, we consider θ . It is an α -formula with conjunctive components

$\theta_1 = \langle\langle 1 \rangle\rangle(pUq \vee \Box q)$ and $\theta_2 = \langle\langle 2 \rangle\rangle(\Diamond p \wedge \Box \neg q)$.

Furthermore, θ_1 is a γ -formula of the form $\langle\langle A \rangle\rangle\Phi$ where the main connective of Φ is \vee . So, $\text{dec}(\theta_1) = \text{dec}(pUq) \cup \text{dec}(\Box q) \cup (\text{dec}(pUq) \oplus \text{dec}(\Box q))$, where $\text{dec}(pUq) = \{\langle p, pUq \rangle, \langle q, \top \rangle\}$ and $\text{dec}(\Box q) = \{\langle q, \Box q \rangle\}$.

Thus, $\text{dec}(\theta_1) = \{\langle p, pUq \rangle, \langle q, \top \rangle, \langle q, \Box q \rangle, \langle p \wedge q, pUq \vee \Box q \rangle\}$, hence $\theta_1 \equiv (p \wedge \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle pUq) \vee q \vee (q \wedge \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle \Box q) \vee (p \wedge q \wedge \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle (pUq \vee \Box q))$.

Likewise, θ_2 is a γ -formula of the form $\langle\langle A \rangle\rangle\Phi$, and the main connective of Φ is \wedge . So $\text{dec}(\theta_2) = \text{dec}(\Diamond p) \otimes \text{dec}(\Box \neg q)$, with $\text{dec}(\Diamond p) = \{\langle T, \Diamond p \rangle, \langle p, T \rangle\}$ and $\text{dec}(\Box \neg q) = \{\langle \neg q, \Box \neg q \rangle\}$.

Thus, $\text{dec}(\theta_2) = \{\langle \top \wedge \neg q, \Diamond p \wedge \Box \neg q \rangle, \langle p \wedge \neg q, \top \wedge \Box \neg q \rangle\}$
 $= \{\langle \neg q, \Diamond p \wedge \Box \neg q \rangle, \langle p \wedge \neg q, \Box \neg q \rangle\}$ and

$\theta_2 \equiv (\neg q \wedge \langle\langle 2 \rangle\rangle \circ \langle\langle 2 \rangle\rangle (\Diamond p \wedge \Box \neg q)) \vee (p \wedge \neg q \wedge \langle\langle 2 \rangle\rangle \circ \langle\langle 2 \rangle\rangle \Box \neg q)$.

For ϑ , the γ -decomposition is similar, we only replace $\langle\langle 2 \rangle\rangle$ by $\llbracket 2 \rrbracket$. Thus, we obtain
 $\vartheta_1 \equiv (p \wedge \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle pUq) \vee (q) \vee (q \wedge \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle \Box q) \vee (p \wedge q \wedge \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle (pUq \vee \Box q))$
and $\vartheta_2 \equiv (\neg q \wedge \llbracket 2 \rrbracket \circ \llbracket 2 \rrbracket (\Diamond p \wedge \Box \neg q)) \vee (p \wedge \neg q \wedge \llbracket 2 \rrbracket \circ \llbracket 2 \rrbracket \Box \neg q)$.

The *closure* $cl(\psi)$ of an ATL^+ state formula ψ is the least set of ATL^+ formulae such that $\psi, \top, \perp \in cl(\psi)$, and $cl(\psi)$ is closed under taking of successor-, α -, β -, and γ -components. For any set of state formulae Γ , we define

$$cl(\Gamma) := \bigcup \{cl(\psi) \mid \psi \in \Gamma\}.$$

We denote by $|\psi|$ the length of ψ and by $\|\Gamma\|$ the cardinality of Γ .

Example 3.3. The construction of the closure of the formula θ from Example 3.2 is given in Figure 3. Each node of the tree represents an element of the closure. Children of an interior node are respective components of the parent formula, according to the definition of closure.

The closure of ϑ is similar to the one of θ except that every $\llbracket 2 \rrbracket$ is replaced by $\langle\langle 2 \rangle\rangle$.

LEMMA 3.4. For any ATL^+ state formula φ , $\|cl(\varphi)\| < 2^{|\varphi|^2}$.

PROOF. Every formula in $cl(\varphi)$ has length less than $2|\varphi|$ and is built from symbols in φ , so there can be at most $|\varphi|^{2|\varphi|} = 2^{2|\varphi| \log_2 |\varphi|} < 2^{|\varphi|^2}$ such formulae. \square

This estimate is rather crude, but $\|cl(\varphi)\|$ can reach size exponential in $|\varphi|$. Indeed, consider the formulae $\phi_k = \langle\langle 1 \rangle\rangle(p_1 U q_1 \wedge (p_2 U q_2 \wedge (\dots \wedge p_k U q_k) \dots))$ for $k = 1, 2, \dots$ and distinct $p_1, q_1, \dots, p_k, q_k, \dots \in \text{Prop}$. Then $|\phi_k| = O(k)$, whereas the number of different γ -components of ϕ_k is 2^k ; hence, $\|cl(\phi_k)\| > 2^k$.

3.2. Full Expansions of Sets of ATL⁺ Formulae

As part of the tableau construction, we need a procedure that, for any given finite set of ATL⁺ state formulae Γ , produces all “full expansions” (called in Goranko and Shkatov [2009] “downward saturated extensions”; see Remark 3.7) defined here:

Definition 3.5. Let Γ, Δ be sets of ATL⁺ state formulae and $\Gamma \subseteq \Delta \subseteq cl(\Gamma)$.

- (1) Δ is *patently inconsistent* if it contains \perp or a pair of formulae φ and $\neg\varphi$.
- (2) Δ is a *full expansion* of Γ if it is not patently inconsistent and satisfies the following closure conditions:
 - if $\varphi \wedge \psi \in \Delta$ then $\varphi \in \Delta$ and $\psi \in \Delta$;
 - if $\varphi \vee \psi \in \Delta$ then $\varphi \in \Delta$ or $\psi \in \Delta$;
 - if $\varphi \in \Delta$ is a γ -formula, then at least one γ -component of φ is in Δ and exactly one of these γ -components in Δ , denoted $\gamma(\varphi, \Delta)$, is designated as *the γ -component in Δ linked to the γ -formula φ* , as explained below.

The family of all full expansions of Γ will be denoted by $FE(\Gamma)$. It can be constructed by a simple iterative procedure that starts with $\{\Gamma\}$ and repeatedly, until saturation, takes a set X from the currently constructed family, selects a formula $\varphi \in X$, and applies the closure rule just given corresponding to its type. Clearly, this procedure terminates on every finite input set of formulae Γ and produces a family of at most $2^{\|cl(\Gamma)\|}$ sets. Furthermore, due to Lemma 3.1, we have the following:

PROPOSITION 3.6. *For any finite set of ATL⁺ state formulae Γ :*

$$\bigwedge \Gamma \equiv \bigvee \left\{ \bigwedge \Delta \mid \Delta \in FE(\Gamma) \right\}.$$

PROOF. Lemma 3.1 implies that every set extension step (as just described) applied to a family of sets \mathcal{F} preserves the formula $\bigvee \{ \bigwedge \Delta \mid \Delta \in \mathcal{F} \}$ up to logical equivalence. At the beginning, that formula is $\bigwedge \Gamma$. \square

Remark 3.7. Instead of full expansions, the tableau construction in Goranko and Shkatov [2009] uses “minimal downward saturated extensions,” where “downward saturated extension” corresponds to “full expansion.” The minimality condition means that if one full expansion is contained in another one, then it is omitted. This could be problematic because sometimes nonminimal full expansions may be needed. For instance, if $\Gamma = \{ \langle\langle 1 \rangle\rangle(p\mathcal{U}q), p \wedge \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle(p\mathcal{U}q) \}$, then

$$FE(\Gamma) = \{ \{ \langle\langle 1 \rangle\rangle(p\mathcal{U}q), p \wedge \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle(p\mathcal{U}q), p, \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle(p\mathcal{U}q) \}, \{ q, \langle\langle 1 \rangle\rangle(p\mathcal{U}q), p \wedge \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle(p\mathcal{U}q), p, \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle(p\mathcal{U}q) \} \}.$$

Although the second full expansion contains the first one, we might have to consider both alternatives in the tableau where Γ is only part of the label of a state for the sake of satisfying an eventuality of the type $\varphi\mathcal{U}\psi$. However, we have no concrete example showing that such situations may occur indeed.

4. TABLEAU-BASED DECISION PROCEDURE FOR ATL⁺

The tableau procedure consists of three major phases: *pretableau construction*, *prestate elimination*, and *state elimination*. Given an input formula η , it essentially constructs a (nondeterministic) CGS that is state-labeled by the closure set of the input formula $cl(\eta)$; that is, a directed graph T^η (called a *tableau*) where each node is labeled by a subset of $cl(\eta)$ (see Definition 2.2(1)) and directed edges between nodes relating them to successor nodes.

The pretableau construction phase produces the so-called *pretableau* \mathcal{P}^η for the input formula η , with two kinds of nodes: *states* and *prestates*. States are fully expanded




Fig. 4. Application of the rule **(SR)** on the prestate $\Gamma_0^\theta = \{\langle\langle 1 \rangle\rangle(p\mathcal{U}q \vee \Box q) \wedge \langle\langle 2 \rangle\rangle(\Diamond p \wedge \Box \neg q)\}$.

sets meant to represent states of a CGM, whereas prestates can be any finite sets of formula from $cl(\eta)$ and only play a temporary role in the construction of \mathcal{P}^η . States and prestates are labeled uniquely so they can be identified with their labels. The prestate elimination phase creates a smaller graph T_0^η out of \mathcal{P}^η , called the *initial tableau for η* , by eliminating all the prestates from \mathcal{P}^η and accordingly redirecting its edges. Finally, the state elimination phase removes, step-by-step, all the states (if any) that cannot be satisfied in a CGM because they lack necessary successors or because they contain unrealized eventualities. Eventually, the elimination procedure produces a (possibly empty) subgraph T^η of T_0^η , called the *final tableau for η* . If some state Δ of T^η contains η , the tableau procedure declares η satisfiable, and a partly defined CGM (called *Hintikka game structure*) satisfying η can be extracted from it by another procedure (described in Section 6.2); otherwise, it declares η unsatisfiable.

4.1. Pretableau Construction Phase

The pretableau construction phase for an input formula η starts with an initial prestate (with label $\{\eta\}$) and consists of alternating application of two construction rules until saturation: **(SR)**, expanding prestates into states, and **(Next)**, creating successor prestates from states. This phase closely resembles the corresponding one for the ATL tableaux in Goranko and Shkatov [2009], with the only essential difference being the γ -decomposition of γ -formulae used here by the rule **(SR)**, which causes, as we will see, a possibly exponential blow-up of the size of the tableaux and, eventually, of the entire worst-case time complexity, as compared to the ATL tableaux. Another (minor) difference with respect to Goranko and Shkatov [2009] is in the formulation of both rules, because here we work with formulae in negation normal form.

Rule (SR) Given a prestate Γ , do the following:

- (1) For each full expansion Δ of Γ add to the pretableau a state with label Δ .
- (2) For each of the added states Δ , if Δ does not contain any formulae of the form $\langle\langle A \rangle\rangle \circ \varphi$ or $\llbracket A \rrbracket \circ \varphi$, add the formula $\langle\langle A \rangle\rangle \top$ to it;
- (3) For each state Δ obtained at steps 1 and 2, link Γ to Δ via a \implies edge;
- (4) If, however, the pretableau already contains a state Δ' with label Δ , do not create another copy of it but only link Γ to Δ' via a \implies edge.

Example 4.1. For the formula θ from Example 3.2 the initial prestate is

$$\Gamma_0^\theta = \{\langle\langle 1 \rangle\rangle(p\mathcal{U}q \vee \Box q) \wedge \langle\langle 2 \rangle\rangle(\Diamond p \wedge \Box \neg q)\}.$$

It has two full expansions:

$$\Delta_1 = \{\theta, \theta_1, \theta_2, p, \neg q, \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle p\mathcal{U}q, \langle\langle 2 \rangle\rangle \circ \langle\langle 2 \rangle\rangle (\Diamond p \wedge \Box \neg q)\}, \text{ and}$$

$$\Delta_2 = \{\theta, \theta_1, \theta_2, p, p \wedge \neg q, \neg q, \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle p\mathcal{U}q, \langle\langle 2 \rangle\rangle \circ \langle\langle 2 \rangle\rangle \Box \neg q\}.$$

Likewise, for the formula ϑ the initial prestate is

$$\Gamma_0^\vartheta = \{\langle\langle 1 \rangle\rangle(p\mathcal{U}q \vee \Box q) \wedge \llbracket 2 \rrbracket (\Diamond p \wedge \Box \neg q)\}$$

and it has two full expansions:

$$\Delta_1 = \{\vartheta, \vartheta_1, \vartheta_2, p, \neg q, \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle p\mathcal{U}q, \llbracket 2 \rrbracket \circ \llbracket 2 \rrbracket (\Diamond p \wedge \Box \neg q)\}, \text{ and}$$

$$\Delta_2 = \{\vartheta, \vartheta_1, \vartheta_2, p, p \wedge \neg q, \neg q, \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle p\mathcal{U}q, \llbracket 2 \rrbracket \circ \llbracket 2 \rrbracket \Box \neg q\}.$$

In the following, by *enforceable successor formula* we mean a formula of the form $\langle\langle A \rangle\rangle \psi$, and, by *unavoidable successor formula*, one of the form $\llbracket A' \rrbracket \psi$.

Rule (Next) Given a state Δ , do the following, where σ is a shorthand for $\sigma_{\mathbb{A}}$:

- (1) List all primitive successor formulae of Δ in such a way that all enforceable successor formulae precede all unavoidable ones where $A \neq \mathbb{A}$; let the result be the list

$$\mathbb{L} = \langle\langle A_0 \rangle\rangle \varphi_0, \dots, \langle\langle A_{m-1} \rangle\rangle \varphi_{m-1}, \llbracket A'_0 \rrbracket \psi_0, \dots, \llbracket A'_{l-1} \rrbracket \psi_{l-1}$$

Let $r_{\Delta} = m + l$; denote by $D(\Delta)$ the set $\{0, \dots, r_{\Delta} - 1\}^{|\mathbb{A}|}$. Then, for every $\sigma \in D(\Delta)$, denote $N(\sigma) := \{i \mid \sigma_i \geq m\}$, where σ_i is the i th component of the tuple σ , and let $\text{co}(\sigma) := [\sum_{i \in N(\sigma)} (\sigma_i - m)] \bmod l$.

- (2) For each $\sigma \in D(\Delta)$ create a prestate:

$$\begin{aligned} \Gamma_{\sigma} = & \{\varphi_p \mid \langle\langle A_p \rangle\rangle \varphi_p \in \Delta \text{ and } \sigma_a = p \text{ for all } a \in A_p\} \\ & \cup \{\psi_q \mid \llbracket A'_q \rrbracket \psi_q \in \Delta, \text{co}(\sigma) = q, \text{ and } \mathbb{A} - A'_q \subseteq N(\sigma)\} \end{aligned}$$

If Γ_{σ} is empty, add \top to it. Then connect Δ to Γ_{σ} with $\xrightarrow{\sigma}$.

If, however, $\Gamma_{\sigma} = \Gamma$ for some prestate Γ that has already been added to the pretableau, only connect Δ to Γ with $\xrightarrow{\sigma}$.

Remark 4.2. Rule **(Next)** ensures that every prestate Γ of Δ (i.e., every element of the finite set of prestates that are targets of \longrightarrow edges outgoing from Δ) satisfies the following:

- if $\{\langle\langle A_i \rangle\rangle \varphi_i, \langle\langle A_j \rangle\rangle \varphi_j\} \subseteq \Delta$ and $\{\varphi_i, \varphi_j\} \subseteq \Gamma$, then $A_i \cap A_j = \emptyset$;
- Γ contains at most one formula of the form ψ such that $\llbracket A \rrbracket \psi \in \Delta$, since the number $\text{co}(\sigma)$ is uniquely determined for every $\sigma \in D(\Delta)$;
- if $\{\langle\langle A_i \rangle\rangle \varphi_i, \llbracket A' \rrbracket \psi\} \subseteq \Delta$ and $\{\varphi_i, \psi\} \subseteq \Gamma$, then $A_i \subseteq A'$.

Here is some intuition on the rule **(Next)** (see also [Goranko and Shkatov 2009]). This rule must ensure that for each $\langle\langle A \rangle\rangle \varphi$ from \mathbb{L} there is a respective A -action at Δ that guarantees φ in the label of every successor and that for every $\llbracket A' \rrbracket \psi$ from \mathbb{L} there is a A' -co-action at Δ that ensures ψ in the label of the respective successors.

Now, the actions at Δ are defined so that every agent's action represents a choice of that agent of a formula from \mathbb{L} for the satisfaction of which the agent chooses to act. When all agents in some A_p choose action p , then they act together for satisfying $\langle\langle A_p \rangle\rangle \varphi_p$, so this is the required A_p -action. As for the co-actions, the idea is that for any fixed $\llbracket A'_q \rrbracket \psi_q$ in \mathbb{L} , all agents in the complement of A'_q may choose to act in favor of some $\llbracket \cdot \rrbracket$ -formula by simply selecting an action of the type $\sigma_i \geq m$. Then, for every A'_q -action the agents in $\mathbb{A} - A'_q$ can synchronize their actions to ensure that the resulting action profile σ satisfies $\text{co}(\sigma) = q$, thereby ensuring ψ_q in the successor state. In fact, any agent who chooses to act *co-strategically* (i.e., in favor of a $\llbracket \cdot \rrbracket$ -formula) can always synchronize her action with all other agents acting co-strategically to ensure that the resulting action profile σ satisfies $\text{co}(\sigma) = j$, for any value $j = 0, \dots, l - 1$. So, every such agent i is able, once all other agents have chosen their actions, to unilaterally enforce in the successor state any ψ_q such that $\llbracket A'_q \rrbracket \psi_q$ in \mathbb{L} and $i \notin A'_q$.

The rules **(SR)** and **(Next)** are applied alternatively until saturation, which is bound to occur because every label is a subset of $cl(\eta)$. Then the construction phase is over. The graph built in this phase is called *pretableau* for the input formula η and denoted by \mathcal{P}^{η} . Given a pretableau, if Γ is a prestate, we denote by $\mathbf{states}(\Gamma)$ the finite set of states that are targets of \Longrightarrow edges outgoing from Γ , and, if Δ is a state, we denote by $\mathbf{prestates}(\Delta)$ the finite set of prestates that are targets of \longrightarrow edges outgoing from Δ .

Before providing an example of how rule **(Next)** works, we give an example for the computation of the function co .

Example 4.3. Let the input formula, containing two agents 1 and 2, be such that at some step of the pretableau construction there is a state containing the next four primitive formulae: $\{\langle\langle 1 \rangle\rangle \circ \varphi_1, \langle\langle 1, 2 \rangle\rangle \circ \varphi_2, \llbracket 2 \rrbracket \circ \varphi_3, \llbracket 1 \rrbracket \circ \varphi_4\}$. The computation of the functions N and co and the successor prestate Γ_σ for each action profile gives:

σ	$N(\sigma)$	$\text{co}(\sigma)$	Γ_σ	σ	$N(\sigma)$	$\text{co}(\sigma)$	Γ_σ
0, 0	\emptyset	0	$\{\varphi_1\}$	2, 0	$\{1\}$	0	$\{\varphi_3\}$
0, 1	\emptyset	0	$\{\varphi_1\}$	2, 1	$\{1\}$	0	$\{\top\}$
0, 2	$\{2\}$	0	$\{\varphi_1\}$	2, 2	$\{1, 2\}$	0	$\{\varphi_3\}$
0, 3	$\{2\}$	1	$\{\varphi_1, \varphi_4\}$	2, 3	$\{1, 2\}$	1	$\{\varphi_4\}$
1, 0	\emptyset	0	$\{\top\}$	3, 0	$\{1\}$	1	$\{\top\}$
1, 1	\emptyset	0	$\{\varphi_2\}$	3, 1	$\{1\}$	1	$\{\top\}$
1, 2	$\{2\}$	0	$\{\top\}$	3, 2	$\{1, 2\}$	1	$\{\varphi_4\}$
1, 3	$\{2\}$	1	$\{\varphi_4\}$	3, 3	$\{1, 2\}$	0	$\{\varphi_3\}$

Example 4.4 (Continuation of Example 4.1 for θ). For Δ_1 , the list of successor formulae is

$$\mathbb{L} = \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle p \mathcal{U} q, \langle\langle 2 \rangle\rangle \circ \langle\langle 2 \rangle\rangle (\diamond p \wedge \square \neg q)$$

So, $m = 2$, $l = 0$ and $r_{\Delta_1} = 2$.

Because there are no unavoidable successor formulae, we do not need to compute $N(\sigma)$ and $\text{co}(\sigma)$. Then,

$$\begin{aligned} \Gamma_{(0,0)} &= \{\langle\langle 1 \rangle\rangle p \mathcal{U} q\} = \Gamma_1 & \Gamma_{(1,0)} &= \{\top\} = \Gamma_3 \\ \Gamma_{(0,1)} &= \{\langle\langle 1 \rangle\rangle p \mathcal{U} q, \langle\langle 2 \rangle\rangle (\diamond p \wedge \square \neg q)\} = \Gamma_2 & \Gamma_{(1,1)} &= \{\langle\langle 2 \rangle\rangle (\diamond p \wedge \square \neg q)\} = \Gamma_4. \end{aligned}$$

For Δ_2 , the list of successor formulae is

$$\mathbb{L} = \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle p \mathcal{U} q, \langle\langle 2 \rangle\rangle \circ \langle\langle 2 \rangle\rangle \square \neg q$$

So, $m = 2$, $l = 0$ and $r_{\Delta_2} = 2$.

Here again, we do not compute $N(\sigma)$ and $\text{co}(\sigma)$. Then,

$$\begin{aligned} \Gamma_{(0,0)} &= \{\langle\langle 1 \rangle\rangle p \mathcal{U} q\} = \Gamma_1 & \Gamma_{(1,0)} &= \{\top\} = \Gamma_3 \\ \Gamma_{(0,1)} &= \{\langle\langle 1 \rangle\rangle p \mathcal{U} q, \langle\langle 2 \rangle\rangle \square \neg q\} = \Gamma_5 & \Gamma_{(1,1)} &= \{\langle\langle 2 \rangle\rangle \square \neg q\} = \Gamma_6. \end{aligned}$$

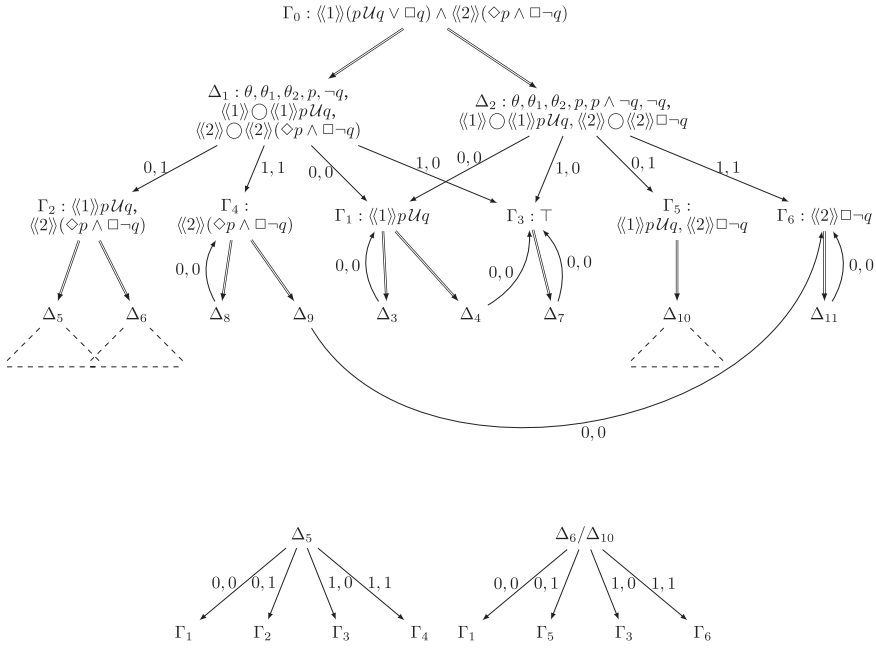
Applying rule **(SR)** to the so-obtained prestates, we have:

$$\begin{aligned} \mathbf{states}(\Gamma_1) &= \{\Delta_3 : \{\langle\langle 1 \rangle\rangle p \mathcal{U} q, p, \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle p \mathcal{U} q\}, \Delta_4 : \{\langle\langle 1 \rangle\rangle p \mathcal{U} q, q, \langle\langle 1, 2 \rangle\rangle \circ \top\}\}, \\ \mathbf{states}(\Gamma_2) &= \{\Delta_5 : \{\langle\langle 1 \rangle\rangle p \mathcal{U} q, \langle\langle 2 \rangle\rangle (\diamond p \wedge \square \neg q), p, \neg q, \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle p \mathcal{U} q, \langle\langle 2 \rangle\rangle \circ \langle\langle 2 \rangle\rangle (\diamond p \wedge \square \neg q)\}, \Delta_6 : \{\langle\langle 1 \rangle\rangle p \mathcal{U} q, \langle\langle 2 \rangle\rangle (\diamond p \wedge \square \neg q), p, p \wedge \neg q, \neg q, \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle p \mathcal{U} q, \langle\langle 2 \rangle\rangle \circ \langle\langle 2 \rangle\rangle \square \neg q\}\}, \\ \mathbf{states}(\Gamma_3) &= \{\Delta_7 : \{\top, \langle\langle 1, 2 \rangle\rangle \circ \top\}\}, \\ \mathbf{states}(\Gamma_4) &= \{\Delta_8 : \{\langle\langle 2 \rangle\rangle (\diamond p \wedge \square \neg q), \neg q, \langle\langle 2 \rangle\rangle \circ \langle\langle 2 \rangle\rangle (\diamond p \wedge \square \neg q)\}, \Delta_9 : \{\langle\langle 2 \rangle\rangle (\diamond p \wedge \square \neg q), p \wedge \neg q, \neg q, \langle\langle 2 \rangle\rangle \circ \langle\langle 2 \rangle\rangle \square \neg q\}\}, \\ \mathbf{states}(\Gamma_5) &= \{\Delta_{10} : \{\langle\langle 1 \rangle\rangle p \mathcal{U} q, \langle\langle 2 \rangle\rangle \square \neg q, p, \neg q, \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle p \mathcal{U} q, \langle\langle 2 \rangle\rangle \circ \langle\langle 2 \rangle\rangle \square \neg q\}\}, \\ \mathbf{states}(\Gamma_6) &= \{\Delta_{11} : \{\langle\langle 2 \rangle\rangle \square \neg q, \neg q, \langle\langle 2 \rangle\rangle \circ \langle\langle 2 \rangle\rangle \square \neg q\}\}. \end{aligned}$$

The pretableau for θ is given in Figure 5.

Example 4.5 (Continuation of Example 4.1 for ϑ). For Δ_1 , the list of successor formulae is

$$\mathbb{L} = \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle p \mathcal{U} q, \llbracket 2 \rrbracket \circ \llbracket 2 \rrbracket (\diamond p \wedge \square \neg q)$$

Fig. 5. The pretableau for θ .

So, $m = 1, l = 1$ and $r_{\Delta_1} = 2$. Therefore,

$$\begin{aligned} N(0, 0) &= \emptyset & N(1, 0) &= \{1\} \\ N(0, 1) &= \{2\} & N(1, 1) &= \{1, 2\} \end{aligned}$$

and also $\text{co}(0, 0) = \text{co}(0, 1) = \text{co}(1, 0) = \text{co}(1, 1) = 0$. Then,

$$\Gamma_{(0,0)} = \Gamma_{(0,1)} = \{\langle\langle 1 \rangle\rangle pUq\} = \Gamma_1, \text{ and } \Gamma_{(1,0)} = \Gamma_{(1,1)} = \{\langle\langle 2 \rangle\rangle (\Diamond p \wedge \Box \neg q)\} = \Gamma_2.$$

For Δ_2 , the list of successor formulae is

$$\mathbb{L} = \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle pUq, \langle\langle 2 \rangle\rangle \circ \langle\langle 2 \rangle\rangle \Box \neg q$$

So, $m = 1, l = 1$ and $r_{\Delta_2} = 2$. Here also

$$\begin{aligned} N(0, 0) &= \emptyset & N(1, 0) &= \{1\} \\ N(0, 1) &= \{2\} & N(1, 1) &= \{1, 2\} \end{aligned}$$




and $\text{co}(0, 0) = \text{co}(0, 1) = \text{co}(1, 0) = \text{co}(1, 1) = 0$. Then,

$$\Gamma_{(0,0)} = \Gamma_{(0,1)} = \{\langle\langle 1 \rangle\rangle pUq\} = \Gamma_1, \text{ and } \Gamma_{(1,0)} = \Gamma_{(1,1)} = \{\langle\langle 2 \rangle\rangle \Box \neg q\} = \Gamma_3.$$

In the same way, we obtain:

$$\begin{aligned} \mathbf{states}(\Gamma_1) &= \{\Delta_3 : \{\langle\langle 1 \rangle\rangle pUq, p, \langle\langle 1 \rangle\rangle \circ \langle\langle 1 \rangle\rangle pUq\}, \Delta_4 : \{\langle\langle 1 \rangle\rangle pUq, q, \langle\langle 1, 2 \rangle\rangle \circ \top\}\}; \\ \mathbf{states}(\Gamma_2) &= \{\Delta_5 : \{\langle\langle 2 \rangle\rangle (\Diamond p \wedge \Box \neg q), \neg q, \langle\langle 2 \rangle\rangle \circ \langle\langle 2 \rangle\rangle (\Diamond p \wedge \Box \neg q)\}, \Delta_6 : \{\langle\langle 2 \rangle\rangle (\Diamond p \wedge \Box \neg q), p \wedge \neg q, p, \neg q, \langle\langle 2 \rangle\rangle \circ \langle\langle 2 \rangle\rangle \Box \neg q\}\}; \\ \mathbf{states}(\Gamma_3) &= \{\Delta_7 : \{\langle\langle 2 \rangle\rangle \Box \neg q, \neg q, \langle\langle 2 \rangle\rangle \circ \langle\langle 2 \rangle\rangle \Box \neg q\}\}; \\ \mathbf{states}(\Gamma_4) &= \{\Delta_8 : \{\top, \langle\langle 1, 2 \rangle\rangle \circ \top\}\}. \end{aligned}$$

The pretableau for ϑ is given in Figure 6.

Fig. 6. The pretableau for ϑ .Fig. 7. The initial tableau for θ .Fig. 8. The initial and final tableau for ϑ .

4.2. The Prestate and State Elimination Phases: Eventualities

First, we remove from \mathcal{P}^η all the prestates and the \implies edges, as follows. For every prestate Γ in \mathcal{P}^η , put $\Delta \xrightarrow{\sigma} \Delta'$ for all states Δ in \mathcal{P}^η with $\Delta \xrightarrow{\sigma} \Gamma$ and all $\Delta' \in \mathbf{states}(\Gamma)$; then, remove Γ from \mathcal{P}^η . The graph obtained after eliminating all prestates is called the *initial tableau*, denoted by T_0^η . The initial tableau for the formula θ in our running example is given in Figure 7, and the initial tableau for ϑ is given in Figure 8.

The elimination phase starts with T_0^η and goes through stages. At stage $n+1$, we remove exactly one state from the tableau T_n^η obtained at the previous stage by applying one of the elimination rules described later, thus obtaining the tableau T_{n+1}^η . The set of states of T_m^η is noted S_m^η .

The first elimination rule (**ER1**), defined later, is used to eliminate all states with missing successors for some action vectors determined by the rule (**Next**). If, due to a previous state elimination, any state has an outgoing action vector for which the corresponding successor state is missing, we delete the state. The reason is clear: If Δ is to be satisfiable, then for each $\sigma \in D(\Delta)$ there should exist a satisfiable Δ' that Δ reaches via σ . Formally, the rule is stated as follows, where $D(\Delta)$ is defined in the rule (**Next**):

Rule (ER1): If, for some $\sigma \in D(\Delta)$, all states Δ' with $\Delta \xrightarrow{\sigma} \Delta'$ have been eliminated at earlier stages, then obtain T_{n+1}^η by eliminating Δ (together with its adjacent edges) from T_n^η .

The aim of the next elimination rule is to make sure that there are no *unrealized eventualities*. In ATL, there are only two kinds of eventualities: $\langle\langle A \rangle\rangle\varphi \mathcal{U}\psi$ and $\llbracket A \rrbracket\varphi \mathcal{U}\psi$. The situation is more complex in ATL⁺. For instance, should the formula $\langle\langle A \rangle\rangle(\Box\varphi \vee \psi_1 \mathcal{U}\psi_2)$ be considered an eventuality? Our solution for ATL⁺ is to consider all γ -formulae as *potential eventualities*. In order to properly define the notion of *realization of a potential eventuality*, we first define a Boolean-valued function *Real* that takes as arguments two elements: an ATL⁺ path-formula Φ and a set Θ of ATL⁺ state-formulae. This function allows us to check the realization of a potential eventuality of the form $\langle\langle A \rangle\rangle\Phi$ and $\llbracket A \rrbracket\Phi$ (where Φ is the first argument of *Real*) at a given state labeled by Θ (where Θ is the second argument of *Real*).

- $Real(\Phi \wedge \Psi, \Theta) = Real(\Phi, \Theta) \wedge Real(\Psi, \Theta)$
- $Real(\Phi \vee \Psi, \Theta) = Real(\Phi, \Theta) \vee Real(\Psi, \Theta)$
- $Real(\varphi, \Theta) = true$ if and only if $\varphi \in \Theta$
- $Real(\Box\varphi, \Theta) = false$
- $Real(\Box\varphi, \Theta) = true$ if and only if $\varphi \in \Theta$
- $Real(\varphi \mathcal{U}\psi, \Theta) = true$ if and only if $\psi \in \Theta$

Definition 4.6 (Descendant Potential Eventualities). Let $\xi \in \Delta$ be a potential eventuality of the form $\langle\langle A \rangle\rangle\Phi$ or $\llbracket A \rrbracket\Phi$. Suppose the γ -component $\gamma(\xi, \Delta)$ in Δ linked to ξ is, respectively, of the form $\psi \wedge \langle\langle A \rangle\rangle\Box\langle\langle A \rangle\rangle\Psi$ or $\psi \wedge \llbracket A \rrbracket\Box\llbracket A \rrbracket\Psi$. Then, the *successor potential eventuality* of ξ with respect to $\gamma(\xi, \Delta)$ is the γ -formula $\langle\langle A \rangle\rangle\Psi$ (respectively, $\llbracket A \rrbracket\Psi$) and it will be denoted by ξ_Δ^1 . The notion of *descendant potential eventuality of ξ of degree d* , for $d > 1$, is defined inductively as follows:

- any successor eventuality of ξ (with respect to some γ -component of ξ) is a descendant eventuality of ξ of degree 1;
- any successor eventuality of a descendant eventuality ξ^n of ξ of degree n is a descendant eventuality of ξ of degree $n + 1$.

We also consider ξ to be a descendant eventuality of itself of degree 0.

Example 4.7. Let $\xi = \langle\langle A \rangle\rangle(\Box p \vee q \mathcal{U}r)$ be a potential eventuality such that $\xi \in \Theta$, where Θ is the labeling of a state Δ . Let us see the different cases that can occur and the corresponding result of the function *Real*.

- (1) $p \in \Theta$ and $r \in \Theta$. In that case, $Real(\Box p, \Theta) = true$ and $Real(q \mathcal{U}r, \Theta) = true$, so $Real(\Box p \vee q \mathcal{U}r, \Theta) = true$. This is indeed correct since $q \mathcal{U}r$ is immediately realized.
- (2) $p \notin \Theta$ and $r \in \Theta$. This case is similar to the previous one even if $Real(\Box p, \Theta) = false$, indeed $Real(\Box p \vee q \mathcal{U}r, \Theta) = true$.
- (3) $p \notin \Theta$ and $r \notin \Theta$. Here $Real(\Box p \vee q \mathcal{U}r, \Theta) = false$ and the potential eventuality ξ is not immediately realized. The rule (**SR**) guarantees that $q \in \Theta$. This case means that the part $\Box p$ of ξ is skipped and the part $q \mathcal{U}r$ will be continued. Therefore the

next potential eventuality to be realized is $\langle\langle A \rangle\rangle q \mathcal{U} r$. The immediate realization of this new potential eventuality will be checked again at next states.

- (4) $p \in \Theta$ and $r \notin \Theta$. The potential eventuality ξ is immediately realized since $Real(\Box p, \Theta) = true$, but two subcases can be distinguished to explain why this is correct:
- (a) $q \in \Theta$. Here, both possibilities to do either $\Box p$ or $q \mathcal{U} r$ are kept. So the successor potential eventuality is $\langle\langle A \rangle\rangle(\Box p \vee q \mathcal{U} r)$ and its immediate realization will be checked again at next states.
 - (b) $q \notin \Theta$. This means that only the part $\Box p$ will be kept and the successor potential eventuality is $\langle\langle A \rangle\rangle \Box p$. This case can be declared immediately realized since the construction rules of the tableau guarantee that $\langle\langle A \rangle\rangle \Box p$ is correctly treated.

Example 4.8 (Continuation of Example 4.4). In Δ_5 , we have $\xi = \langle\langle 1 \rangle\rangle(p \mathcal{U} q \vee \Box q)$ with $Real(p \mathcal{U} q \vee \Box q, \Delta_1) = Real(p \mathcal{U} q, \Delta_1) \vee Real(\Box q, \Delta_1) = false \vee false = false$, since $q \notin \Delta_1$, and $\xi' = \langle\langle 2 \rangle\rangle(\Diamond p \wedge \Box \neg q)$ with $Real(\Diamond p \wedge \Box \neg q, \Delta_1) = Real(\Diamond p, \Delta_1) \wedge Real(\Box \neg q, \Delta_1) = true \wedge true = true$ since $p, \neg q \in \Delta_1$.

The successor eventuality of $\xi = \langle\langle 1 \rangle\rangle(p \mathcal{U} q \vee \Box q)$ with respect to $\gamma(\xi, \Delta_1)$ is $\xi_{\Delta_1}^1 = \langle\langle 1 \rangle\rangle p \mathcal{U} q$ in $\Delta_3, \Delta_4, \Delta_5, \Delta_6$. For each $n > 1$, the descendant eventuality of degree n of ξ with respect to $\gamma(\xi, \Delta_1)$ is $\xi_{\Delta_1}^n = \xi_{\Delta_1}^1$ in $\Delta_3, \Delta_4, \Delta_5, \Delta_6, \Delta_{10}$. The successor eventuality of $\xi' = \langle\langle 2 \rangle\rangle(\Diamond p \wedge \Box \neg q)$ with respect to $\gamma(\xi', \Delta_5)$ is $\xi_{\Delta_5}^1 = \langle\langle 2 \rangle\rangle(\Diamond p \wedge \Box \neg q)$ in $\Delta_5, \Delta_6, \Delta_8$. For each $n > 1$, the descendant eventualities of degree n of ξ' with respect to $\gamma(\xi', \Delta_5)$ are $\xi_{\Delta_5}^n = \xi_{\Delta_5}^1$ in $\Delta_5, \Delta_6, \Delta_8$ and Δ_9 ; and $\xi_{\Delta_5}^n = \langle\langle 2 \rangle\rangle \Box \neg q$ in Δ_{10} and Δ_{11} .

Now, let $\mathbb{L} = \langle\langle A_0 \rangle\rangle \circ \varphi_0, \dots, \langle\langle A_{m-1} \rangle\rangle \circ \varphi_{m-1}, \llbracket A'_0 \rrbracket \circ \psi_0, \dots, \llbracket A'_{l-1} \rrbracket \circ \psi_{l-1}$ be the list of all primitive successor formulae of $\Delta \in S_0^n$, induced as part of an application of **(Next)**. We use the following notation:

$$D(\Delta, \langle\langle A_p \rangle\rangle \circ \varphi_p) := \{\sigma \in D(\Delta) \mid \sigma_a = p \text{ for every } a \in A_p\}$$

$$D(\Delta, \llbracket A'_q \rrbracket \circ \psi_q) := \{\sigma \in D(\Delta) \mid \text{co}(\sigma) = q \text{ and } \mathbb{A} - A'_q \subseteq N(\sigma)\}$$

Next, we define recursively what it means for an eventuality ξ to be realized at a state Δ of a tableau T_n^n , followed by our second elimination rule.

Definition 4.9 (Realization of Potential Eventualities). Let $\Delta \in S_n^n$ and $\xi \in \Delta$ be a potential eventuality of the form $\langle\langle A \rangle\rangle \Phi$ or $\llbracket A \rrbracket \Phi$. Then:


- (1) If $Real(\Phi, \Delta) = true$ then ξ is realized at Δ in T_n^n .
- (2) Else, let ξ_{Δ}^1 be the successor potential eventuality of ξ with respect to $\gamma(\xi, \Delta)$. If for every $\sigma \in D(\Delta, \langle\langle A \rangle\rangle \circ \xi_{\Delta}^1)$ (respectively, $\sigma \in D(\Delta, \llbracket A \rrbracket \circ \xi_{\Delta}^1)$), there exists $\Delta' \in T_n^n$ with $\Delta \xrightarrow{\sigma} \Delta'$ and ξ_{Δ}^1 is realized at Δ' in T_n^n , then ξ is realized at Δ in T_n^n .

Rule (ER2): If $\Delta \in S_n^n$ contains a potential eventuality that is not realized at $\Delta \in T_n^n$, then obtain T_{n+1}^n by removing Δ (together with its adjacent edges) from S_n^n .

Example 4.10 (Continuation of Example 4.8). The potential eventuality $\xi'' = \langle\langle 1 \rangle\rangle(p \mathcal{U} q)$ is not realized at Δ_5 , so, by Rule **(ER2)**, we remove the state Δ_5 from T_0^θ and obtain the tableau T_1^θ . The same applies to Δ_6 for ξ'' , so we also remove Δ_6 from T_1^θ and obtain T_2^θ with Rule **(ER2)**. In T_2^θ , there is no more move vector $(0, 1)$ for the state Δ_1 , so, by Rule **(ER1)**, we remove Δ_1 from T_2^θ and obtain T_3^θ . In the same way, Δ_{10} is removed by Rule **(ER2)** and Δ_2 by Rule **(ER1)**.

For the case of ϑ , it is easy to see that no state gets eliminated, so the final tableau is the same as the initial one.

The elimination phase is completed when no more applications of elimination rules are possible. Then we obtain the *final tableau* for η , denoted by T^η . It is declared *open* if

Fig. 9. The final tableau for θ .

η belongs to some state in it, otherwise *closed*. The procedure for deciding satisfiability of η returns “No” if \mathcal{T}^η is closed, “Yes” otherwise.

Example 4.11 (Continuation of Example 4.10). At the end of the elimination phase, Δ_1 and Δ_2 are no longer in \mathcal{T}^θ . Thus, \mathcal{T}^θ is closed, and the formula $\theta = \langle\langle 1 \rangle\rangle(p \cup q \vee \square q) \wedge \langle\langle 2 \rangle\rangle(\diamond p \wedge \square \neg q)$ is declared unsatisfiable. The final tableau for θ is given in Figure 9.

Respectively, the final tableau for ϑ is open; hence, ϑ is declared satisfiable. Indeed, a CGM can be extracted from the final tableau. We explain in Section 6.2 how this can be done in a systematic way.

5. TERMINATION AND SOUNDNESS

The termination of the tableau procedure is straightforward because there are only finitely many states and prestates that can be added in the construction phase.

THEOREM 5.1. *The tableau method for ATL⁺ is sound.*

Soundness of the tableau procedure with respect to unsatisfiability means that if a formula is satisfiable then its final tableau is open. To prove that, we essentially follow the same procedure as in the soundness proof for the tableau-based decision procedure for ATL in Goranko and Shkatov [2009].

The soundness proof consists of three main claims. First, we show that when a prestate Γ is satisfiable then at least one of the states in **states**(Γ) is satisfiable. Then, we prove that when a state Δ is satisfiable then all the prestates in **prestates**(Δ) are satisfiable. Finally, we show that no satisfiable states are eliminated in the elimination phase. Here, we take the input formula of the tableau procedure to be η .

The first step of the proof consists in showing that rule **(SR)** is sound:

LEMMA 5.2. *Let Γ be a prestate of \mathcal{P}^η and let $\mathcal{M}, s \models \Gamma$ for some CGM \mathcal{M} and some $s \in \mathcal{M}$. Then, $\mathcal{M}, s \models \Delta$ holds for at least one $\Delta \in \mathbf{states}(\Gamma)$.*

PROOF. Straightforward from Proposition 3.6. \square

The aim of the next two lemmas is to show that the rule **(Next)** creates only satisfiable prestates from satisfiable states.

We recall that we use $\mathbf{act}_A(s)$ to denote the set of all A -actions that can be played by the coalition A at state s (i.e., $\mathbf{act}_A(s) = \Pi_{a \in A} \mathbf{act}_a(s)$). We also use $\mathbf{act}_A^c(s)$ to denote the set of all A -co-actions available at state s and σ_A^c for an element of this set. Let $\sigma_A \in \mathbf{act}_A(s)$. We say that an action profile $\sigma_{\bar{A}}$ extends an A -action σ_A , denoted by $\sigma_{\bar{A}} \supseteq \sigma_A$, if $\sigma_{\bar{A}}(a) = \sigma_A(a)$ for every $a \in A$. We also use $\mathbf{Out}(s, \sigma_A)$ to denote the set of all states s' for which there exists an action profile $\sigma_{\bar{A}} \in \mathbf{act}_{\bar{A}}(s)$ that extends σ_A and such that $\mathbf{out}(s, \sigma_{\bar{A}}) = s'$. We define in a same way \supseteq and $\mathbf{Out}(s, \sigma_A^c)$ for an A -co-action $\sigma_A^c \in \mathbf{act}_A^c(s)$.

Lemma 5.3 states a semantic property, independent of the tableau construction.

LEMMA 5.3. *Let $\Theta = \{\langle\langle A_1 \rangle\rangle \circ \varphi_1, \dots, \langle\langle A_m \rangle\rangle \circ \varphi_m, \llbracket A' \rrbracket \circ \psi\}$ be a set of formulae such that $A_i \cap A_j = \emptyset$ for every $1 \leq i, j \leq m$, $i \neq j$ and $A_i \subseteq A'$ for every $1 \leq i \leq m$. Let $\mathcal{M}, s \models \Theta$ for some GCM \mathcal{M} and $s \in \mathcal{M}$. Let $\sigma_{A_i} \in \mathbf{act}_{A_i}(s)$ be an A_i -action witnessing the truth of $\langle\langle A_i \rangle\rangle \circ \varphi_i$ at s , for each $1 \leq i \leq m$, and let, finally, $\sigma_{A'}^c \in \mathbf{act}_{A'}^c(s)$ be an A' -co-action witnessing the truth of $\llbracket A' \rrbracket \circ \psi$ at s . Then there exists $s' \in \mathbf{Out}(s, \sigma_{A_1}) \cap \dots \cap \mathbf{Out}(s, \sigma_{A_m}) \cap \mathbf{Out}(s, \sigma_{A'}^c)$ such that $\mathcal{M}, s' \models \{\varphi_1, \dots, \varphi_m, \psi\}$.*

PROOF. Let $A = A_1 \cup \dots \cup A_m$. Since $A_i \cap A_j = \emptyset$ for every $1 \leq i, j \leq m$, the actions $\sigma_{A_1}, \dots, \sigma_{A_m}$ can be combined to get an A -action σ_A . This last can be arbitrarily extended to an A' -action $\sigma_{A'}$ because $A_i \subseteq A'$ for every $1 \leq i \leq m$. Finally, the so obtained $\sigma_{A'}$ can be completed by the A' -co-action $\sigma_{A'}^c$. The resulting action σ_A leads from s to the desired s' . \square

Lemma 5.4 states that satisfiability propagates from states to their successor prestates created via rule **(Next)**.

LEMMA 5.4. *If $\Delta \in \mathcal{P}^\eta$ is a satisfiable state then all the prestates Γ obtained by applying the rule **(Next)** are satisfiable.*

PROOF. Follows by induction on the number of steps in the construction of the tableau and from Lemma 5.3 and Remark 4.2. \square

Thus, the rule **(SR)** generates at least one satisfiable state from a satisfiable prestate and the rule **(Next)** generates only satisfiable prestates from a satisfiable state. Hence, we can conclude that the construction phase of the tableau procedure is sound.

We now move to the elimination phase.

LEMMA 5.5. *Let Δ be a state in \mathcal{T}_n^η . If Δ is satisfiable, then Rule **(ER1)** cannot eliminate Δ from \mathcal{T}_n^η .*

PROOF. By Lemma 5.3, a satisfiable state Δ generates only satisfiable successor prestates, and, by Lemma 5.4, each of these prestates in turn generates at least one satisfiable state. Therefore, if Δ is satisfiable, then for every action profile $\sigma_A \in \text{act}_A(s)$ there is a state Δ' such that $\Delta \xrightarrow{\sigma_A} \Delta'$. Therefore the Rule **(ER1)** cannot eliminate a satisfiable state. \square

It remains to be proved that a satisfiable state cannot be eliminated by Rule **(ER2)** either. We recall that Rule **(ER2)** eliminates each state containing an eventuality that is not realized at that state. So, we need to prove that if a state Δ is satisfiable, then every eventuality $\xi \in \Delta$ is realized at Δ throughout the elimination phase.

Note that the structure underlying a tableau can be seen as a nondeterministic CGS, where edges outgoing from a tableau state can lead to different successors even if they are labeled by the same action vector. The following two definitions will be used to extract deterministic transitions from nondeterministic ones.

Definition 5.6 (Outcome Set of σ_A at s). Let S be a nondeterministic concurrent game structure, let s be a state, and let $\sigma_A \in \text{act}_A(s)$. An *outcome set* of σ_A at s is a set of states X such that for every $\sigma_A \sqsupseteq \sigma_A$ there exists exactly one $s' \in X$ such that $s \xrightarrow{\sigma_A} s'$.

Definition 5.7 (Outcome Set of σ_A^c at s). Let S be a nondeterministic concurrent game structure, let s be a state, and let $\sigma_A^c \in \text{act}_A^c(s)$. An *outcome set* of σ_A^c at s is a set of states X such that for every $\sigma_A \in \text{act}_A(s)$ there exists exactly one $s' \in X$ such that $s \xrightarrow{\sigma_A^c(\sigma_A)} s'$.

In particular, both definitions can be applied to a tableau where the states s and s' are taken to be tableau states Δ and Δ' .

Some notation. Consider a concurrent game structure S that is state-labeled by a set Θ of state formulae of ATL^+ and suppose that the elements of Θ are listed by any enumeration E where enforceable next-time formulae appear before unavoidable next-time formulae, in particular, in the list \mathbb{L} given in the definition of Rule **(Next)**. Then:

- (1) Whenever we write $\langle\langle A_p \rangle\rangle \circ \varphi_p \in \Theta$, we mean that $\langle\langle A_p \rangle\rangle \circ \varphi_p$ is the p th enforceable next-time formula according to E . In particular, when S is a tableau, E is usually

assumed to be the listing of the successor formulae of Θ induced by the application of the rule **(Next)** to Θ .

We use the notation $\llbracket A'_q \rrbracket \circ \psi_q \in \Theta$ likewise.

- (2) Given $\langle\langle A_p \rangle\rangle \circ \varphi \in \Theta$, we denote by $\sigma_{A_p}[\langle\langle A_p \rangle\rangle \circ \varphi]$ a (somehow selected) A -action enforcing φ in any associated successor state.

In particular, when S is a tableau, we denote by $\sigma_{A_p}[\langle\langle A_p \rangle\rangle \circ \varphi_p]$ the unique A_p -action $\sigma_{A_p} \in \text{act}_{A_p}(\Theta)$ in the tableau such that $\sigma_{A_p}(a) = p$ for every $a \in A_p$.

- (3) Likewise, given a formula $\llbracket A' \rrbracket \circ \psi \in \Theta$, where $A'_q \neq \mathbb{A}$, we denote by $\sigma_{A'_q}^c[\llbracket A'_q \rrbracket \circ \psi]$ a (somehow selected) A'_q -co-action $\sigma_{A'_q}^c$ enforcing ψ in any associated successor state.

In particular, when S is a tableau, we denote by $\sigma_{A'_q}^c[\llbracket A'_q \rrbracket \circ \psi_q]$ the unique A'_q -co-action $\sigma_{A'_q}^c \in \text{act}_{A'_q}^c(\Delta)$ in the tableau satisfying the following condition (with notation referring to the definition of rule **(Next)**): $\text{co}(\sigma_{A'_q}^c(\sigma_{A'_q}^c)) = q$ and $\mathbb{A} - A'_q \subseteq N(\sigma_{A'_q}^c(\sigma_{A'_q}^c))$ for every $\sigma_{A'_q} \in \text{act}_{A'_q}(\Delta)$.

In order to prove that the rule **(ER2)** does not eliminate any satisfiable states, we need to show that if a tableau T_n^η contains a state Δ that is satisfiable and contains an eventuality ξ , then ξ is realized at Δ . Thus, we prove that T_n^η “contains” a structure (more precisely, a tree) that “witnesses” the realization of ξ at Δ in T_n^η . This tree will emulate a tree of runs effected by a strategy or co-strategy that “realizes” an eventuality in a model. This simulation is done step-by-step, and each step (i.e., A -action (in the case of $\langle\langle A \rangle\rangle \Phi$) or A -co-action (in the case of $\llbracket A \rrbracket \Phi$)) corresponds to a tableau action or co-action associated with a respective eventuality. The fact that this step-by-step simulation can be done is proved in the next two lemmas (together with their corollaries).

LEMMA 5.8. *Let $\langle\langle A_p \rangle\rangle \circ \varphi_p \in \Delta \in S_n^\eta$ and let $\mathcal{M}, s \models \Delta$ for some CGM \mathcal{M} and state $s \in \mathcal{M}$. Let, furthermore, $\sigma_{A_p} \in \text{act}_{A_p}(s)$ be an A_p -action witnessing the truth of $\langle\langle A_p \rangle\rangle \circ \varphi_p$ at s . Then, there exists in T_n^η an outcome set X of $\sigma_{A_p}[\langle\langle A_p \rangle\rangle \circ \varphi_p]$ such that for each $\Delta' \in X$ there exists $s' \in \text{Out}(s, \sigma_{A_p})$ such that $\mathcal{M}, s' \models \Delta'$.*

PROOF. We consider the following set of prestates (from the pretableau construction):

$$Y = \{\Gamma \in \text{prestates}(\Delta) \mid \Delta \xrightarrow{\sigma_{\mathbb{A}}} \Gamma \text{ for some } \sigma_{\mathbb{A}} \sqsupseteq \sigma_{A_p}[\langle\langle A_p \rangle\rangle \circ \varphi_p]\}$$

For every $\Gamma \in Y$, it follows immediately from the rule **(Next)** that Γ (which must contain φ_p) is either of the form

$$\{\varphi_1, \dots, \varphi_m, \psi\}, \text{ where } \{\langle\langle A_1 \rangle\rangle \circ \varphi_1, \dots, \langle\langle A_m \rangle\rangle \circ \varphi_m, \llbracket A' \rrbracket \circ \psi\} \subseteq \Delta$$

or of the form

$$\{\varphi_1, \dots, \varphi_m\} \text{ where } \{\langle\langle A_1 \rangle\rangle \circ \varphi_1, \dots, \langle\langle A_m \rangle\rangle \circ \varphi_m\} \subseteq \Delta$$

We can reduce the latter case to the former by adding the valid formula $\llbracket \mathbb{A} \rrbracket \circ \top$ (equivalent to $\langle\langle \emptyset \rangle\rangle \circ \top$).

Since $\mathcal{M}, s \models \Delta$, by Lemma 5.3, there exists $s' \in \text{Out}(s, \sigma_{A_p})$ with $\mathcal{M}, s' \models \Gamma$. Then Γ can be extended to a fully expanded set Δ' containing at least one successor formula ($\langle\langle \mathbb{A} \rangle\rangle \circ \top$, if nothing else) such that $\mathcal{M}, s' \models \Delta'$. This is done by choosing, for every β - or γ -formula to be processed in the procedure computing the family of full expansions, a disjunct (respectively, a γ -component) that is actually true in \mathcal{M} at s' (if there are several such options, the choice is arbitrary) and adding it to the current set. \square

COROLLARY 1. *Let $\langle\langle A_p \rangle\rangle \circ \varphi_p \in \Delta$ for $\Delta \in S_n^\eta$ and let $\mathcal{M}, s \models \Delta$ for some CGM \mathcal{M} and state $s \in \mathcal{M}$. Let, furthermore, $\sigma_{A_p} \in \text{act}_{A_p}(s)$ be an A_p -action witnessing the truth of*

$\langle\langle A_p \rangle\rangle \circ \varphi_p$ at s and let $\chi \in cl(\eta)$ be a β -formula (respectively, a γ -formula) and ψ be one of its β -components (respectively, γ -components). Then there exists in T_n^n an outcome set X_ψ of $\sigma_{A_p}[\langle\langle A_p \rangle\rangle \circ \varphi_p]$ such that for every $\Delta' \in X_\psi$ there exists $s' \in \text{Out}(s, \sigma_{A_p})$ such that $\mathcal{M}, s' \models \Delta'$, and moreover, if $\mathcal{M}, s' \models \psi$, then $\psi \in \Delta'$.

PROOF. Construct X_ψ just as X was constructed in the proof of the preceding lemma, but with a single modification: When dealing with the formula χ , instead of choosing arbitrarily between the different options for ψ , choose ψ , which is true at s' . \square

Likewise, we obtain the following for unavoidable formulae:

LEMMA 5.9. Let $\llbracket A'_q \rrbracket \circ \psi_q \in \Delta \in S_n^\eta$ and let $\mathcal{M}, s \models \Delta$ for some CGM \mathcal{M} and state $s \in \mathcal{M}$. Let, furthermore, $\sigma_{A'_q}^c \in \text{act}_{A'_q}^c(s)$ be an A'_q -co-action witnessing the truth of $\llbracket A'_q \rrbracket \circ \psi_q$ at s . Then, there exists in T_n^n an outcome set X of $\sigma_{A'_q}^c[\llbracket A'_q \rrbracket \circ \psi_q]$ such that for each $\Delta' \in X$ there exists $s' \in \text{Out}(s, \sigma_{A'_q}^c)$ such that $\mathcal{M}, s' \models \Delta'$.

The proof is analogous to the proof of Lemma 5.8.

COROLLARY 2. Let $\llbracket A'_q \rrbracket \circ \psi_q \in \Delta \in S_n^\eta$ and let $\mathcal{M}, s \models \Delta$ for some CGM \mathcal{M} and state $s \in \mathcal{M}$. Let, furthermore, $\sigma_{A'_q}^c \in \text{act}_{A'_q}^c(s)$ be an A'_q -co-action witnessing the truth of $\llbracket A'_q \rrbracket \circ \psi_q$ at s and let $\chi \in cl(\eta)$ be a β -formula (respectively, a γ -formula), whose associated β_i -component ($i \in \{1, 2\}$) (respectively, i th γ -component ($i \geq 1$)) is χ_i . Then there exists in T_n^n an outcome set X_{χ_i} of $\sigma_{A'_q}^c[\llbracket A'_q \rrbracket \circ \psi_q]$ such that for every $\Delta' \in X_{\chi_i}$ there exists $s' \in \text{Out}(s, \sigma_{A'_q}^c)$ such that $\mathcal{M}, s' \models \Delta'$, and moreover, if $\mathcal{M}, s' \models \chi_i$, then $\chi_i \in \Delta'$.

In what follows, we make use of the notion of *tree*. In our context, we use such a term as a synonym for a “directed, connected, and acyclic graph, every node of which, except the root, has exactly one incoming edge.” We denote a tree as a pair (R, \rightarrow) , where R is the set of nodes and \rightarrow is the parent-child relation (the edges).

The first kind of tree that we define is the so-called *realization witness tree*. Intuitively, such tree witnesses the satisfaction of a given potential eventuality ξ at a state and simulates a tree of runs effected in a model by (co-)strategies. Our definition is more general than the one in Goranko and Shkatov [2009] because we want this notion to be applicable in a broader context, including tableaux, CGMs, and concurrent game Hintikka structures (to be defined later).

The two definitions given here implicitly use the notion of descendant potential eventuality of degree d and its associate notation (see Definition 4.6). That notion was defined in the context of tableaux; however, it is applicable to any CGS that is state-labeled by a set of state formulae. We recall that, given a potential eventuality $\xi = \langle\langle A \rangle\rangle \Phi (\llbracket A \rrbracket \Phi)$, by convention ξ itself is taken to be its (unique) descendant potential eventuality of degree 0 and that if ξ^i is a descendant eventuality of degree i of ξ , then a γ -component of ξ^i will have the form $\psi \wedge \langle\langle A \rangle\rangle \Phi^{i+1}$ (respectively, $\psi \wedge \llbracket A \rrbracket \Phi^{i+1}$) and $\langle\langle A \rangle\rangle \Phi^{i+1}$ (respectively, $\llbracket A \rrbracket \Phi^{i+1}$) will be a descendant potential eventuality of ξ having degree $d = i + 1$.

A piece of terminology that will be used often further: Given sets X, Y and a mapping $c : X \rightarrow Y$, we sometimes say that *the set X is Y -colored by c* and that for any $x \in X$, the value $c(x)$ is *the Y -color of x under the coloring c* .

Definition 5.10 (Realization Witness Tree for Enforceable Potential eventualities). Let \mathcal{S} be any (nondeterministic) CGS with a state space St that is state-labeled by some set of ATL^+ formulae Γ , with a labeling function c . Let $s \in \text{St}$ and let $\xi \in s$ be a

potential eventuality of the form $\langle\langle A \rangle\rangle\Phi$. A *realization witness tree* for ξ at s is a finite tree $\mathcal{R} = (R, \rightarrow)$, where the set of nodes R is St-colored so that:

- (1) the root of \mathcal{R} is colored with s and is of depth 0;
- (2) if an interior node w of depth i of \mathcal{R} is colored with s' where $c(s') = \Theta$, then there exists a successor ξ^{i+1} of ξ^i such that $\langle\langle A \rangle\rangle\bigcirc_{\xi^{i+1}} \in \Theta$;
- (3) for every interior node $w \in \mathcal{R}$ of depth i colored with s' where $c(s') = \Theta$, the children of w are colored bijectively with vertices from an outcome set of $\sigma_A[\langle\langle A \rangle\rangle\bigcirc_{\xi^i}]$ at s' ;
- (4) if a leaf of depth i of \mathcal{R} is colored with s' where $c(s') = \Theta$, then $\xi^i = \langle\langle A \rangle\rangle\Phi \in \Theta$ is such that $Real(\Phi, \Theta) = true$.

Definition 5.11 (Realization Witness Tree for Unavoidable Potential eventualities). Let S be any (nondeterministic) CGS with a state space St that is state-labeled by some set of ATL⁺ formulae Γ , with a labeling function c . Let $s \in \text{St}$ and let $\xi \in s$ be a potential eventuality of the form $\llbracket A \rrbracket\Phi$. A *realization witness tree* for ξ at s is a finite tree $\mathcal{R} = (R, \rightarrow)$, where the set of nodes R is St-colored so that:

- (1) the root of \mathcal{R} is colored with s and is of depth 0;
- (2) if an interior node w of depth i of \mathcal{R} is colored with s' where $c(s') = \Theta$, then there exists a successor ξ^{i+1} of ξ^i such that $\llbracket A \rrbracket\bigcirc_{\xi^{i+1}} \in \Theta$;
- (3) for every interior node $w \in \mathcal{R}$ of depth i colored with s' where $c(s') = \Theta$, the children of w are colored bijectively with vertices from an outcome set of $\sigma_A^c[\llbracket A \rrbracket\bigcirc_{\xi^i}]$ at s' ;
- (4) if a leaf of depth i of \mathcal{R} is colored with s' where $c(s') = \Theta$, then $\xi^i = \llbracket A \rrbracket\Phi \in \Theta$ is such that $Real(\Phi, \Delta') = true$.

We are going to apply these definitions for the case when the CGS S is a tableau \mathcal{T}_n^η , with states being (identified with) the sets of formulae in their labels.

LEMMA 5.12. *Let $\mathcal{R} = (R, \rightarrow)$ be a realization witness tree for a potential eventuality ξ at $\Delta \in S_n^\eta$. For every Δ' of depth i , coloring a node of R , ξ^i is realized at Δ' in \mathcal{T}_n^η . In particular for $i = 0$; thus, ξ is realized at Δ in \mathcal{T}_n^η .*

PROOF. We prove this lemma by induction on the length of the longest path from a node colored by Δ to a leaf of \mathcal{R} .

Base Case: The length of the longest path from a node w colored by Δ to a leaf of \mathcal{R} is 0. Then, w is a leaf and $Real(\xi, \Delta) = true$. Thus, by Item 1 of Definition 4.9, ξ is realized at $\Delta \in S_n^\eta$.

Induction Step: The length of the longest path from a node w colored by Δ to a leaf of \mathcal{R} is $l > 0$. Then w is an interior node of depth i , so $\langle\langle A \rangle\rangle\xi^i \in \Delta$ (respectively, $\llbracket A \rrbracket\xi^i \in \Delta$) and there exists an action (respectively, a co-action) such that for all children w' of w , where each w' is colored by Δ' , $\xi_\Delta^{i+1} \in \Delta'$. Let \mathcal{R}' be a subtree of \mathcal{R} whose root is w' . The length of the longest path from a node w' colored by Δ' to a leaf of \mathcal{R}' is at most $l - 1$. Thus, by induction hypothesis, ξ_Δ^{i+1} is realized at $\Delta' \in \mathcal{T}_n^\eta$ and ξ^j is realized at Δ'' in \mathcal{T}_n^η . Therefore ξ_Δ^{i+1} is realized at Δ' in \mathcal{T}_n^η , and w respects Item 2 of Definition 4.9. We conclude that ξ^i is realized at $\Delta \in S_n^\eta$. \square

We now prove the existence of a realization witness tree for any satisfiable state of a tableau containing a potential eventuality.

LEMMA 5.13. *Let $\xi \in \Delta$ be a potential eventuality and $\Delta \in S_n^\eta$ be satisfiable. Then, there exists a realization witness tree $\mathcal{R} = (R, \rightarrow)$ for ξ at $\Delta \in S_n^\eta$. Moreover, every Δ' , coloring a node of R , is satisfiable.*

PROOF. We only give the proof for potential eventualities of the type $\langle\langle A \rangle\rangle\Phi$. The case of potential eventualities of type $\llbracket A \rrbracket\Phi$ is similar.

When dealing with realization of potential eventualities, we have two cases:

- (1) $Real(\Phi, \Delta) = true$. This case is straightforward: The realization witness tree consists of only the root, colored with Δ .
- (2) $Real(\Phi, \Delta) = false$. This case means that there is a successor potential eventuality ξ_Δ^1 such that $\langle\langle A \rangle\rangle \circ \xi_\Delta^1 \in \Delta$.

As Δ is satisfiable, there exists a CGM \mathcal{M} and a state $s \in \mathcal{M}$ such that $\mathcal{M}, s \models \Delta$, and in particular, $\mathcal{M}, s \models \langle\langle A \rangle\rangle \circ \xi_\Delta^1$. Thus, there exists an A -action $\sigma_A \in \text{act}_A^c(s)$ such that $\mathcal{M}, s' \models \xi_\Delta^1$ for all $s' \in \text{Out}(s, \sigma_A)$, that is an A -action witnessing the truth of $\langle\langle A \rangle\rangle \circ \xi_\Delta^1$ at s .

We know that Δ is satisfiable and that $\langle\langle A \rangle\rangle \circ \xi_\Delta^1$ is an enforceable successor formula. Let p be the position of $\langle\langle A \rangle\rangle \circ \xi_\Delta^1$ in the list made at the application of the rule **(Next)** on Δ . Note that ξ is a γ -formula $\in cl(\eta)$, where at least one of its γ -components, obtained from a pair $\langle\psi, \Psi\rangle$, is such that $Real(\Psi, FE(\psi)) = true$. Let χ be such a γ -component. So, Corollary 1 is applicable to Δ , and, according to that corollary, there exists an outcome set X_χ of $\sigma_A[\langle\langle A \rangle\rangle \circ \xi_\Delta^1]$ at Δ such that, for every $\Delta' \in X_\chi$, there exists $s' \in \text{Out}(s, \sigma)$ such that $\mathcal{M}, s' \models \Delta'$; and, moreover, if $\mathcal{M}, s' \models \chi$, then $\chi \in \Delta'$. We start building the realization witness tree \mathcal{R} with a simple tree whose root r is colored with Δ and whose leaves are colored bijectively with a node from X . This first tree respects Items 1 to 3 of Definition 5.10; some of the leaves respect Item 4 of this definition, but not all of them. The next part treats these leaves.

Since $\mathcal{M}, s \models \xi_\Delta^1$ with $\xi_\Delta^1 = \langle\langle A \rangle\rangle \Phi'$ for every $s' \in \text{Out}(s, \sigma_A)$, it follows that for every such s' there exists a perfect-recall A -strategy $F_A^{s'}$ such that for every $\lambda \in \text{Plays}(s', F_A^{s'})$, $\mathcal{M}, \lambda \models \Phi'$. Then, playing σ_A followed by playing $F_A^{s'}$ constitutes a perfect recall strategy F_A witnessing the truth of ξ at s .

Then we continue the construction of \mathcal{R} as follows. Let S' be the set of all states s'' appearing as part of a play consistent with $F_A^{s'}$, containing a descendant eventuality ξ^i of ξ and satisfying the requirement that $\mathcal{M}, s'' \not\models \chi$, for all γ -components χ obtained from a pair $\langle\psi, \Psi\rangle$ such that $Real(\Psi, FE(\psi)) = true$. For every $s' \in \text{Out}(s, \sigma_A)$, we follow the perfect recall strategy $F_A^{s'}$, matching every state $s'' \in S'$ with a node w'' of \mathcal{R} and matching every A' -action of $F_A^{s'}$ at s'' with the tableau $\sigma_A[\langle\langle A \rangle\rangle \circ \xi^i] \in \text{act}_A(\Delta'')$ where ξ^i is the descendant eventuality of ξ in w'' and Δ'' is the state coloring the node w'' . We follow in this way each $F_A^{s'}$ along each run until we reach a state t where $\mathcal{M}, t \models \chi$ (χ is as described earlier). This means that we have reached a leaf of \mathcal{R} ; this leaf respects Item 4 of Definition 5.10. As $\mathcal{M}, s \models \xi$, such a state can be reached for each run, and we thus obtain a finite tree \mathcal{R} .

Thus, the so constructed realization witness tree conforms to Definition 5.10. \square

LEMMA 5.14. *Let Δ be a state in \mathcal{T}_n^η . If Δ is satisfiable then Rule **(ER2)** cannot eliminate Δ from \mathcal{T}_n^η .*

PROOF. Let $\Delta \in \mathcal{T}_n^\eta$ be a satisfiable state.

If Δ contains no eventuality, then Rule **(ER2)** is not applicable.

If Δ contains an eventuality ξ , then Lemma 5.13 ensures that there exists a realization witness tree for Δ and, by Lemma 5.12, we know that ξ is realized at Δ in \mathcal{T}_n^η . Therefore, Rule **(ER2)** cannot eliminate Δ for \mathcal{T}_n^η . \square

THEOREM 5.15 (SOUNDNESS). *If η is satisfiable, then \mathcal{T}^η is open.*

PROOF. Lemmas 5.4–5.14 ensure that if Δ is satisfiable, then Δ cannot be eliminated from \mathcal{T}_n^η due to Rule **(ER1)** or Rule **(ER2)**. Moreover, Lemma 5.2 ensures that if the input formula η is satisfiable, then at least one state containing η (created from the

initial prestate) is satisfiable. Thus, this state cannot be eliminated and therefore the final tableau T^n is open. \square

6. COMPLETENESS, MODEL SYNTHESIS AND COMPLEXITY

6.1. Hintikka Structures

The tableau procedure actually attempts to build not a CGM of the input formula but a state-labeled nondeterministic CGS from which structures of a special kind can be extracted, which essentially are partly defined CGMs. Following Pratt [1980], Ben-Ari et al. [1983], and Goranko and Shkatov [2009], we will call them *Hintikka structures*. Here, we give the definition of a Hintikka structure for a given ATL^+ formula η and show how to obtain a CGM for η from a Hintikka structure for η . Later, we explain how to extract a Hintikka structure “satisfying” the input formula from its open final tableau.

Definition 6.1. A *Concurrent Game Hintikka Structure* (CGHS) is a deterministic CGS $\mathcal{H} = (\mathbb{A}, \text{St}, \{\text{Act}_a\}_{a \in \mathbb{A}}, \{\text{act}_a\}_{a \in \mathbb{A}}, \text{out}, H)$ that is state-labeled by a given set Γ of ATL^+ -formulae with a state-labeling function H . Let $s \in \text{St}$ be a state of \mathcal{H} . A Hintikka structure \mathcal{H} satisfies the following constraints:

- H1. If $\varphi \in H(s)$ then $\neg\varphi \notin H(s)$;
- H2. If an α -formula belongs to $H(s)$, then both its α -components do;
- H3. If a β -formula belongs to $H(s)$, then one of its β -components does;
- H4. If a γ -formula belongs to $H(s)$, then one of its γ -components does;
- H5. If $\langle\langle A \rangle\rangle \psi \in H(s)$, then there exists an A -action $\sigma_A \in \text{act}_{\mathbb{A}}(s)$ such that $\psi \in H(s')$ for all $s' \in \text{Out}(s, \sigma_A)$. Likewise, if $\llbracket A \rrbracket \psi \in H(s)$, then there exists an A -co-action $\sigma_A^c \in \text{act}_A^c(s)$ such that $\psi \in H(s')$ for all $s' \in \text{Out}(s, \sigma_A^c)$.
- H6. If a potential eventuality $\xi = \langle\langle A \rangle\rangle \Phi$ (respectively, $\xi = \llbracket A \rrbracket \Phi$) belongs to $H(s)$, then there exists a realization witness tree, rooted at s in \mathcal{H} for $\xi = \langle\langle A \rangle\rangle \Phi$ (respectively, $\xi = \llbracket A \rrbracket \Phi$) at s .

Remark 6.2. The condition H6 is well defined because a Hintikka structure is obtained by coloring via H from a deterministic CGS for which the notion of realization witness tree is defined.

Definition 6.3. Let $\mathcal{H} = (\mathbb{A}, \text{St}, \{\text{Act}_a\}_{a \in \mathbb{A}}, \{\text{act}_a\}_{a \in \mathbb{A}}, \text{out}, H)$ be a CGHS and η be an ATL^+ -formula. We say that \mathcal{H} is a CGHS for η , if $\eta \in H(s)$ for some $s \in \text{St}$.

We now show that from any CGHS for a given formula η a CGM satisfying η can be obtained.

THEOREM 6.4. *Let $\mathcal{H} = (\mathbb{A}, \text{St}, \{\text{Act}_a\}_{a \in \mathbb{A}}, \{\text{act}_a\}_{a \in \mathbb{A}}, \text{out}, H)$ be a CGHS for a given ATL^+ -formula η . Let further $\mathcal{M} = (\mathbb{A}, \text{St}, \{\text{Act}_a\}_{a \in \mathbb{A}}, \{\text{act}_a\}_{a \in \mathbb{A}}, \text{out}, \text{Prop}, L)$ be the CGM obtained from \mathcal{H} by setting, for every $s \in \text{St}$, $L(s) = H(s) \cap \text{Prop}$. Then, for every $s \in \text{St}$ and every ATL^+ formula φ , $\varphi \in H(s)$ implies $\mathcal{M}, s \models \varphi$. In particular, \mathcal{M} satisfies η .*

PROOF. Suppose $\varphi \in H(s)$. We will prove that $\mathcal{M}, s \models \varphi$ by induction on the structure of the state formula φ .

Base. If $\varphi \in \text{Prop} \cup \{\top\}$ belongs to $H(s)$, it is immediate that $\mathcal{M}, s \models \varphi$, by definition of L and H1.

Inductive Step.

- $\neg\varphi$ is $\psi_1 \wedge \psi_2$. By H2, we get that $\psi_1 \in H(s)$ and $\psi_2 \in H(s)$. By inductive hypothesis $\mathcal{M}, s \models \psi_1$ and $\mathcal{M}, s \models \psi_2$. Therefore $\mathcal{M}, s \models \varphi$.
- $\neg\varphi$ is $\psi_1 \vee \psi_2$. By H3 we get that either $\psi_1 \in H(s)$ or $\psi_2 \in H(s)$. By inductive hypothesis, either $\mathcal{M}, s \models \psi_1$, or $\mathcal{M}, s \models \psi_2$. Therefore $\mathcal{M}, s \models \varphi$.

- φ is $\langle\langle A \rangle\rangle \circ \psi$ or $\llbracket A \rrbracket \circ \psi$. An application of H5 and the inductive hypothesis to ψ imply that $\mathcal{M}, s \models \varphi$.
- φ is $\langle\langle A \rangle\rangle \Phi$ or φ is $\llbracket A \rrbracket \Phi$, where Φ is a path formula whose main operator is different from \circ ; that is, φ is a γ -formula. Here, we only present in detail the first case, the second one being quite similar. We need to prove the existence of a (perfect recall) strategy F_A such that, for each branch λ in \mathcal{M} stemming from s and consistent with that strategy, $\mathcal{M}, \lambda \models \Phi$. This will imply that $\mathcal{M}, s \models \varphi$. Since $\varphi = \langle\langle A \rangle\rangle \Phi \in H(s)$ by hypothesis, then H6 guarantees the existence of a realization witness tree T on \mathcal{H} for φ . By construction, T provides a partial finite strategy Fp_A , defined only for the finite set of histories occurring in T and having length strictly less than the height of T . We want to show that Fp_A can be extended to a strategy F_A , defined for all the histories in \mathcal{H} , and that T can be extended to a possibly infinite tree T' such that:
 - Each node of T' is also a node of \mathcal{H} , and each labeled edge of T' is also a labeled edge of \mathcal{H} .
 - All paths in T' are consistent with F_A ; hence, T' witnesses the truth of $\langle\langle A \rangle\rangle \Phi$ at state s of \mathcal{M} by instantiating the quantifier $\langle\langle A \rangle\rangle$ to F_A .

Next, we show how to construct F_A and T' . Let us consider any finite path in T of the form $\lambda_{\leq n}$, where $\lambda_0 = s$ and λ_n is a leaf. By construction of T , each node λ_i , for $1 \leq i \leq n$, is a node of \mathcal{H} , and each labeled edge of T is a labeled edge of \mathcal{H} . The descendant potential eventuality φ^n of φ belongs to the color of λ_n by construction of T . Since λ_n is a node of \mathcal{H} and $\varphi^n \in H(\lambda_n)$, by H4 some γ -component χ of φ^n belongs to $H(\lambda_n)$. This formula χ is either of the form ψ or of the form $\psi \wedge \langle\langle A \rangle\rangle \circ \varphi^{n+1}$ (the second case occurs, for instance, when φ has the form $\langle\langle A \rangle\rangle \square \theta$).

In the first case, any extension of the partial strategy Fp_A and any extension of $\lambda_{\leq n}$ to an infinite path will do.

In the second case, we apply H2 to get $\psi \in H(\lambda_n)$ and $\langle\langle A \rangle\rangle \circ \varphi^{n+1} \in H(\lambda_n)$. By H5, there exists an A-action $\sigma_A \in \text{act}_A(H(\lambda_n))$ such that $\varphi^{n+1} \in H(s')$ for all $s' \in \text{Out}(\lambda_n, \sigma_A)$. Playing this A-action σ_A after the partial strategy Fp_A gives us a new partial strategy Fp'_A defined for histories whose length is less than or equal to n . The set of successors of λ_n for T' is the set $\text{Out}(\lambda_n, \sigma_A)$. For each $s' \in \text{Out}(\lambda_n, \sigma_A)$, we can again apply H2, H4, and H5 to get a new partial strategy Fp''_A defined for histories whose length is inferior or equal to $n+1$. For any $s' \in \text{Out}(\lambda_n, \sigma_A)$, its successors are obtained by an application of Fp''_A . An infinite iteration of this procedure will give us the complete strategy F_A and the way to extend the finite tree so as to get T' . \square

6.2. Completeness and Model Synthesis

THEOREM 6.5. *The tableau method for ATL^+ is complete.*

Completeness of the procedure means that an open tableau for η implies existence of a CGM model for η . So, we start with an open tableau T^η for η , and we want to prove that η is indeed satisfiable. The proof is constructive, as we will build from T^η a Hintikka structure \mathcal{H}_η that can be turned into a model for η . In order to construct that Hintikka structure, first we extract special T^η -trees associated with potential eventualities, which can be seen as building modules to be used to construct the entire structure. Eventually, we show that the so constructed structure is a Hintikka structure for η .

First, we need to define *edge-labeling* of a tree.

Definition 6.6. Let $\mathcal{W} = (W, \rightsquigarrow)$ be a tree and Y be a non-empty set. An *edge-labeling* of \mathcal{W} by Y is a mapping l from the set of edges of \mathcal{W} to the set of non-empty subsets of Y .

Definition 6.7. Given a tableau T^η , a tree $\mathcal{W} = (W, \rightsquigarrow)$ is a T^η -tree if the following conditions hold:

- \mathcal{W} is S^η -colored by some coloring mapping c .
- \mathcal{W} is edge-labeled by $\bigcup_{(\Delta \in S^\eta)} \text{act}_\mathbb{A}(\Delta)$, by some edge-labeling mapping l ;
- $l(w \rightsquigarrow w') \subseteq \text{act}_\mathbb{A}(\Delta)$ for every $w \in W$ with $c(w) = \Delta$;
- For every interior node $w \in W$ with $c(w) = \Delta$ and every successor $\Delta' \in T^\eta$ of Δ , there exists exactly one $w' \in W$ such that $l(w \rightsquigarrow w') = \{\sigma \mid \Delta \xrightarrow{\sigma} \Delta'\}$.

Definition 6.8. Let $\Delta \in S^\eta$. A T^η -tree \mathcal{W} is *rooted at* Δ if the root r of \mathcal{W} is colored with Δ .

For the purpose of our construction, we distinguish two kinds of T^η -trees: *simple* or *realizing*. Their definitions are given next. Realizing T^η -trees will deal especially with potential eventualities.

Definition 6.9. A tree $\mathcal{W} = (W, \rightsquigarrow)$ is *simple* if it has no interior nodes except the root.

Simple T^η -trees can be seen as one-step modules.

Definition 6.10. Let $\mathcal{W} = (W, \rightsquigarrow)$ be a T^η -tree rooted at Δ and $\xi \in \Delta$ a potential eventuality. The tree \mathcal{W} is a *realizing T^η -tree for* ξ , denoted \mathcal{W}_ξ , if there exists a subtree \mathcal{R}_ξ of \mathcal{W} rooted at Δ such that \mathcal{R}_ξ is a realization witness tree for ξ rooted at $\Delta \in T^\eta$.

LEMMA 6.11. *Let $\Delta \in S^\eta$. Then, there exists a simple T^η -tree rooted at Δ .*

PROOF. We construct a simple T^η -tree \mathcal{W} rooted at Δ as follows. The root of \mathcal{W} is a node r such that $c(r) = \Delta$. For every successor state Δ' of $\Delta \in T^\eta$, let $Moves(\Delta, \Delta') = \{\sigma \mid \Delta \xrightarrow{\sigma} \Delta'\}$. Note that, by construction of the tableau, the family $\{Moves(\Delta, \Delta') \mid \Delta' \text{ is a successor of } \Delta\}$ is a partition of the set $\text{act}_\mathbb{A}(\Delta)$ of all action profiles applied at Δ . Now, for each set X of that family, we select one successor Δ' of Δ such that $X = \text{act}(\Delta, \Delta')$ and add a successor t to \mathcal{W} such that $c(t) = \Delta'$ and $l(r \rightsquigarrow t) = \{\sigma \mid \Delta \xrightarrow{\sigma} \Delta'\}$. \square

Example 6.12 (Continuation of Example 4.10). Consider the final tableau T^ϑ for the formula $\vartheta = \langle\langle 1 \rangle\rangle(p \sqcup q \vee \square q) \wedge \llbracket 2 \rrbracket(\diamond p \wedge \square \neg q)$.

We have seen in Example 4.10 that $S^\vartheta = \{\Delta_1, \dots, \Delta_8\}$.


We have listed possible simple T^ϑ -trees rooted at each Δ in Figure 10.

To show the existence of a realizing T^η -tree for ξ at Δ , we first prove the existence of a realization witness tree \mathcal{R}_ξ for ξ at Δ .

LEMMA 6.13. *Let T^η be a tableau for η and ξ be a potential eventuality realized at $\Delta \in T^\eta$. Then, there exists a realization witness tree \mathcal{R}_ξ for ξ at Δ in T^η .*

PROOF. We give detailed proof only for the case where $\xi = \langle\langle A \rangle\rangle \Phi$; the other case is similar, just replace $\langle\langle A \rangle\rangle$ by $\llbracket A \rrbracket$ in the proof. Suppose that ξ is realized at Δ in T^η . We define the *rank of ξ at Δ in T^η* , denoted $\text{rank}(\xi, \Delta, T^\eta)$, to be the minimal length of a chain of descendant potential eventualities $\xi = \xi^0, \dots, \xi^n = \langle\langle A \rangle\rangle \Phi_n$ ensuring the realization of ξ ; that is, $\text{Real}(\Phi_n, \Delta_j) = \text{true}$ for some state Δ_j descendant of Δ in T^η . We prove the existence of a realization witness tree \mathcal{R} for ξ at Δ in T^η by induction on $\text{rank}(\xi, \Delta, T^\eta)$.

Base: $\text{rank}(\xi, \Delta, T^\eta) = 0$. Here, ξ is immediately realized and \mathcal{R}_ξ contains only the root colored with Δ .

Fig. 10. Simple T^θ -trees.

Inductive Step: $\text{rank}(\xi, \Delta, T^\eta) = k$ where $k > 0$. Since ξ is realized at $\Delta \in T^\eta$ and $\text{rank}(\xi, \Delta, T^\eta) > 0$, by Definition 4.9, we have that for every $\sigma \in D(\Delta, \langle\langle A \rangle\rangle \circ \xi^1)$ there exists $\Delta' \in T^\eta$ such that $\Delta \xrightarrow{\sigma} \Delta'$ and ξ^1 is realized at $\Delta' \in T^\eta$. We build a tree T rooted at a node r colored with Δ where the children v of r are colored bijectively with the set of Δ' obtained earlier. Then, $\text{rank}(\xi^1, \Delta', T^\eta) = k - 1$ and we can apply the inductive hypothesis to obtain a realization witness tree \mathcal{R}_{ξ^1} for ξ^1 at Δ' in T^η for each Δ' . Thus, replacing each node v of T by the corresponding \mathcal{R}_{ξ^1} gives us \mathcal{R}_ξ . \square

LEMMA 6.14. *Let $\xi \in \Delta \in S^\eta$ be a potential eventuality. Then, there exists a finite realizing T^η -tree for ξ rooted at Δ .*

PROOF. Since T^η is open, ξ is realized at Δ in T^η . To construct the realizing T^η -tree \mathcal{W}_ξ for ξ rooted at Δ , we start from the realization witness tree \mathcal{R}_ξ , whose existence is given by Lemma 6.13, and, provisionally, we take \mathcal{W}_ξ to be \mathcal{R}_ξ . The problem with \mathcal{R}_ξ is that for some $\sigma \in \text{act}_\mathbb{A}(\Delta)$ at some node w of \mathcal{R}_ξ , there is no edge $w \rightsquigarrow w'$ such that $l(w \rightsquigarrow w') \ni \sigma$. Therefore, to extend \mathcal{W}_ξ into a realizing T^η -tree, for every such node w , we pick one of the successor states of $c(w)$ via σ , say Δ' and add a node w' to \mathcal{W}_ξ such that $c(w') = \Delta'$ and $l(w \rightsquigarrow w') \ni \sigma$. \square

Example 6.15 (Continuation of Example 6.12). We now extract in Figure 11 a possible realizing tree from the open tableau for every state Δ_i using realization witness trees.

We now construct a final structure, denoted by \mathfrak{F} , from simple and realizing T^η -trees. This construction is made step by step. At the end of the construction, we prove that \mathfrak{F} is indeed a Hintikka structure.

Step 1. We define a grid \mathcal{F} of size $m \times n$, where m is the number of eventualities occurring in T^η and n the number of states of T^η . Each row of that grid is labeled by one of the potential eventualities and each column by a state of T^η previously ordered by name ($\Delta_i < \Delta_j$ if $i < j$). We denote by ξ_i the eventuality associated with row $0 \leq i \leq m$, we denote by Δ_j the state associated with the column $0 \leq j \leq n$. The content $\mathcal{F}(i, j)$ of each intersection between a row i and a column j of \mathcal{F} is as follows: If $\xi_i \in \Delta_j$, then $\mathcal{F}(i, j)$ is the realizing T^η -tree for ξ_i rooted at Δ_j , whose existence is ensured by Lemma 6.14; otherwise, $\mathcal{F}(i, j)$ is the simple T^η -tree rooted at Δ_j , whose existence is ensured by Lemma 6.11.

Example 6.16 (Continuation of Example 6.15). The grid \mathcal{F} for our example has a size 4×8 and is represented in Figure 12.

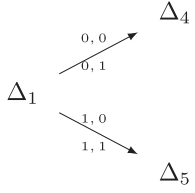
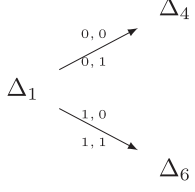
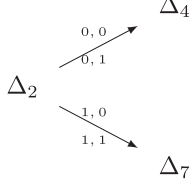
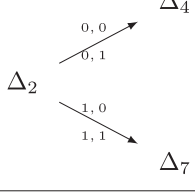
State	Eventualities	Rank	Realization Witness Tree	Realizing Tree
Δ_1	$\xi_1 = \langle\langle 1 \rangle\rangle(p\mathcal{U}q \vee \Box q)$	1	$\Delta_1 \xrightarrow[0,1]{0,0} \Delta_4$	
Δ_1	$\xi_2 = \llbracket 2 \rrbracket(\Diamond p \wedge \Box \neg q)$	1	$\Delta_1 \xrightarrow[1,1]{1,0} \Delta_6$	
Δ_2	$\xi_1 = \langle\langle 1 \rangle\rangle(p\mathcal{U}q \vee \Box q)$	1	$\Delta_2 \xrightarrow[0,1]{0,0} \Delta_4$	
Δ_2	$\xi_2 = \llbracket 2 \rrbracket(\Diamond p \wedge \Box \neg q)$	1	$\Delta_2 \xrightarrow[1,1]{1,0} \Delta_7$	
Δ_3	$\xi_3 = \xi_1^1 = \langle\langle 1 \rangle\rangle p\mathcal{U}q$	1	$\Delta_3 \xrightarrow{0,0} \Delta_4$	$\Delta_3 \xrightarrow{0,0} \Delta_4$
Δ_4	$\xi_3 = \xi_1^1 = \langle\langle 1 \rangle\rangle p\mathcal{U}q$	0	Δ_4	$\Delta_4 \xrightarrow{0,0} \Delta_8$
Δ_5	$\xi_2 = \llbracket 2 \rrbracket(\Diamond p \wedge \Box \neg q)$	1	$\Delta_5 \xrightarrow{0,0} \Delta_6$	$\Delta_5 \xrightarrow{0,0} \Delta_6$
Δ_6	$\xi_2 = \llbracket 2 \rrbracket(\Diamond p \wedge \Box \neg q)$	0	Δ_6	$\Delta_6 \xrightarrow{0,0} \Delta_7$
Δ_7	$\xi_4 = \xi_2^1 = \llbracket 2 \rrbracket \Box \neg q$	0	Δ_7	$\Delta_7 \xrightarrow{0,0} \Delta_7$

Fig. 11. Eventualities and realizing T^ϑ -trees.

Step 2. We make a queue \mathcal{Q} that will contain potential eventualities occurring in \mathcal{T}^η . The first element of \mathcal{Q} is either η , if η is a potential eventuality, or the potential eventuality associated to the first column of the grid just defined. Let ξ_i be the first element of the queue, so that $\mathcal{Q}(0) = \xi_i$. Then we add to \mathcal{Q} all the potential eventualities following the order of grid's rows and cycling if necessary; that is, $\mathcal{Q}(k) = \xi_{((i+k) \bmod m)}$ for $k \in [1, m-1]$.

Step 3. Let Δ be one of the states containing η . Next, we take the element $\mathcal{F}(\mathcal{Q}(0), \Delta)$ of the grid. The root of $\mathcal{F}(\mathcal{Q}(0), \Delta)$ is then the root of \mathfrak{F} . Then we take one-by-one in order all the elements of the rest of the queue and do the following:

\mathcal{F}		0	1	2	3
		ξ_1	ξ_2	ξ_3	ξ_4
0	Δ_1				
1	Δ_2				
2	Δ_3	$\Delta_3 \xrightarrow{0,0} \Delta_4$	$\Delta_3 \xrightarrow{0,0} \Delta_4$	$\Delta_3 \xrightarrow{0,0} \Delta_4$	$\Delta_3 \xrightarrow{0,0} \Delta_4$
3	Δ_4	$\Delta_4 \xrightarrow{0,0} \Delta_8$	$\Delta_4 \xrightarrow{0,0} \Delta_8$	$\Delta_4 \xrightarrow{0,0} \Delta_8$	$\Delta_4 \xrightarrow{0,0} \Delta_8$
4	Δ_5	$\Delta_5 \xrightarrow{0,0} \Delta_5$	$\Delta_5 \xrightarrow{0,0} \Delta_6$	$\Delta_5 \xrightarrow{0,0} \Delta_5$	$\Delta_5 \xrightarrow{0,0} \Delta_5$
5	Δ_6	$\Delta_6 \xrightarrow{0,0} \Delta_7$	$\Delta_6 \xrightarrow{0,0} \Delta_7$	$\Delta_6 \xrightarrow{0,0} \Delta_7$	$\Delta_6 \xrightarrow{0,0} \Delta_7$
6	Δ_7	$\Delta_7 \xrightarrow{0,0} \Delta_7$	$\Delta_7 \xrightarrow{0,0} \Delta_7$	$\Delta_7 \xrightarrow{0,0} \Delta_7$	$\Delta_7 \xrightarrow{0,0} \Delta_7$
7	Δ_8	$\Delta_8 \xrightarrow{0,0} \Delta_8$	$\Delta_8 \xrightarrow{0,0} \Delta_8$	$\Delta_8 \xrightarrow{0,0} \Delta_8$	$\Delta_8 \xrightarrow{0,0} \Delta_8$

Fig. 12. The grid \mathcal{F} .

Let $\mathcal{Q}(i)$ be the current element of the queue to be treated. For every dead-end state $w \in \mathfrak{F}$, that is a state without successors, such that $c(w) = \Delta_j$, we add the tree $\mathcal{F}(\mathcal{Q}(i), \Delta_j)$ by identifying the dead-end state w with the root of $\mathcal{F}(\mathcal{Q}(i), \Delta_j)$;

Example 6.17. From the grid \mathcal{F} , we can extract in four steps a partial structure \mathfrak{F} realizing all the eventualities (see Figure 13).

Step 4. Finally, we ensure that \mathfrak{F} finite. While there is a dead-end state in \mathfrak{F} , say w with $c(w) = \Delta_j$, we choose a component from the row $\mathcal{F}(\Delta_j)$ as follows:

- With priority, we choose a component $\mathcal{F}(i, \Delta_j)$, $0 \leq i \leq m$ already occurring in \mathfrak{F} . Let r be the root of the component $\mathcal{F}(i, \Delta_j)$, $0 \leq i \leq m$ inside \mathfrak{F} . Then, we add an arrow \rightsquigarrow between every predecessor v of w and the root r and label this arrow with $l(v \rightsquigarrow w)$. Then we delete the node $w \in \mathfrak{F}$.
- Otherwise, if the chosen component $\mathcal{F}(i, \Delta_j)$ is not already occurring in \mathfrak{F} , then we add the new component to \mathfrak{F} as usual by merging the root of the component with the dead-end state w .

When there are no longer dead-ends in \mathfrak{F} , the structure is completed, and we have obtained our final structure.

Example 6.18. The final structure \mathfrak{F} for the formula $\vartheta = \langle\langle 1 \rangle\rangle(p\mathcal{U}q \vee \Box q) \wedge \langle\langle 2 \rangle\rangle(\Diamond p \wedge \Box \neg q)$ is given in Figure 14.

By keeping only the propositional variables in the state labels, we obtain the following CGM satisfying ϑ in Figure 15.

LEMMA 6.19. *Let \mathcal{T} be a T^n -tree rooted at $\Delta = c(w)$. Then, the following holds:*

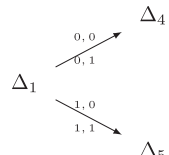
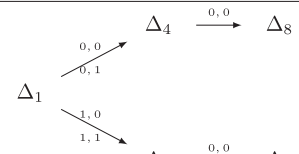
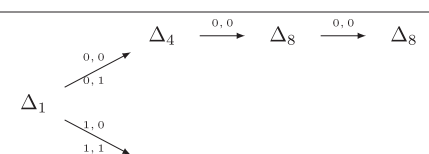
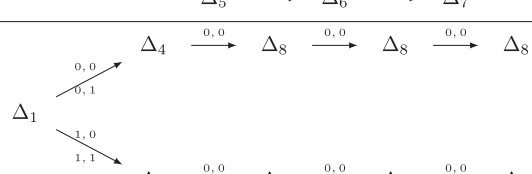
Eventualities	Added grid elements	Partial structure \mathfrak{F}
$\{\xi_1, \xi_2, \xi_3, \xi_4\}$	$\mathcal{F}(\xi_1, \Delta_1)$	
$\{\xi_2, \xi_3, \xi_4\}$	$\mathcal{F}(\xi_2, \Delta_4); \mathcal{F}(\xi_2, \Delta_5)$	
$\{\xi_3, \xi_4\}$	$\mathcal{F}(\xi_3, \Delta_8); \mathcal{F}(\xi_3, \Delta_6)$	
$\{\xi_4\}$	$\mathcal{F}(\xi_4, \Delta_8); \mathcal{F}(\xi_4, \Delta_8)$	

Fig. 13. Eventualities and partial structures.




Fig. 14. The final structure \mathfrak{F} for the formula $\vartheta = \langle\langle 1 \rangle\rangle(p \sqcup q \vee \square q) \wedge \langle\langle 2 \rangle\rangle(\diamond p \wedge \square \neg q)$.




Fig. 15. Concurrent game model satisfying ϑ .

- (1) If $\langle\langle A \rangle\rangle \circ \varphi \in \Delta$, then there exists an A -action $\sigma_A \in \text{act}_A(\Delta)$ such that $\varphi \in c(w') = \Delta'$ where $l(w \rightsquigarrow w') \ni \sigma$ for every $\sigma \sqsupseteq \sigma_A$.
- (2) If $\llbracket A \rrbracket \circ \varphi \in \Delta$ then there exists a co- A -action $\sigma_A^c \in \text{act}_A^c(\Delta)$ such that $\varphi \in c(w') = \Delta'$ where $l(w \rightsquigarrow w') \ni \sigma$ for every $\sigma \sqsupseteq \sigma_A^c$.

PROOF. We recall that all successor formulae of $\Delta \in S^\eta$ are ordered at the application of the rule **(Next)** to Δ .

(1) Suppose that $\langle\langle A \rangle\rangle \circ \varphi \in \Delta$. Then, the required A -action is $\sigma_A[\langle\langle A \rangle\rangle \circ \varphi]$. Indeed, it immediately follows from the rule **(Next)** that for every $\sigma \sqsupseteq \sigma_A$ in the pretableau \mathcal{P}^η , if $\Delta \xrightarrow{\sigma} \Gamma$, then $\varphi \in \Gamma$ and $\varphi \in \Delta'$ since Δ' is a full expansion of Γ . The statement (1) of the lemma follows.

(2) Suppose that $\llbracket A \rrbracket \circ \varphi \in \Delta$. There are two cases to consider:

Case 1: $A \neq \mathbb{A}$. We consider an arbitrary $\sigma_A \in \text{act}_A(\Delta)$. Then, σ_A can be extended to an action vector $\sigma' \sqsupseteq \sigma$. Let $N(\sigma_A)$ be the set $\{i \mid \sigma_A(i) \geq m\}$, where m is the number of enforceable successor formulae in Δ , and let $\text{co}(\sigma_A) = (\sum_{i \in N(\sigma_A)} (\sigma_A(i) - m)) \bmod l$, where l is the number of unavoidable successor formulae in Δ . Now, we consider $\sigma' \sqsupseteq \sigma_A$ defined as follows: $\sigma'_b = ((q - \text{co}(\sigma_A)) \bmod l + m)$ and $\sigma'_{a'} = m$ for any $a' \in \mathbb{A} - (A \cup \{b\})$, where $b \in \mathbb{A} - A$. Thus, we have $\mathbb{A} - A \subseteq N(\sigma)$ and also $\text{co}(\sigma') = (\text{co}(\sigma_A) + (q - \text{co}(\sigma_A))) \bmod l + m = q$. Therefore, for this arbitrarily chosen σ_A there exists at least one state, say Δ' , such that $\Delta \xrightarrow{\sigma} \Delta'$ and $\varphi \in \Delta'$.

Case 2: $A = \mathbb{A}$. Then, by virtue of (H2), $\langle\langle \emptyset \rangle\rangle \circ \neg\varphi \in \Delta$ and thus, by the rule **(Next)**, $\neg\varphi \in \Gamma$ for every successor Γ of Δ . Then, $\neg\varphi \in \Delta'$ for every Δ' that is a successor of Δ in \mathcal{T}^η and hence the coloring set of every leaf of \mathcal{T} . Then, the (unique) co- \mathbb{A} -action, which is an identity function, has the required properties.

The statement (2) of the lemma follows. \square

THEOREM 6.20. *If \mathcal{T}^η is open, then η is satisfiable.*

PROOF. The structure \mathfrak{F} constructed from \mathcal{T}^η is a Hintikka structure. Indeed, H1–H4 of Definition 6.1 are satisfied since the nodes of \mathfrak{F} are nodes of \mathcal{T}^η . H5 of the same definition essentially follows from Lemma 6.19. Whenever a node w of \mathfrak{F} contains a potential eventuality ξ , this means that this eventuality will stay in the queue (see the earlier construction of \mathfrak{F}) until realized. Moreover, if the \mathcal{T}^η -tree \mathcal{W} chosen to complete \mathfrak{F} from w does not realize ξ , either ξ or one of its descendants is present in each newly generated dead-end of \mathfrak{F} . So, when it is the turn to realize ξ , we add to each dead-end state the realizing \mathcal{T}^η -tree for ξ . This, together with Lemma 6.19, guarantees that there exists a realization witness tree for ξ on \mathfrak{F} at w . Thus, H6 of Definition 6.1 is satisfied, too.

By construction, the structure \mathfrak{F} is a CGHS for η ; thus, Theorem 6.4 can be applied to obtain from it a model for η . Thus, η is satisfiable. \square

6.3. Complexity

THEOREM 6.21. *The tableau procedure for ATL^+ runs in 2EXPTIME.*

PROOF. The argument generally follows the calculations computing the complexity of the tableau method for ATL in section 4.7 of Goranko and Shkatov [2009], with one essential difference: $\|cl(\eta)\|$ for any ATL formula η is linear in its length $|\eta|$, whereas $\|cl(\eta)\|$ for an ATL^+ formula η can be exponentially large in $|\eta|$, as shown after Lemma 3.4. This exponential blow-up, combined with the worst-case exponential in $\|cl(\eta)\|$ number of states in the tableau, accounts for the 2EXPTIME worst-case complexity of the tableau method for ATL^+ , which is the expected optimal lower bound. It is also an upper bound

for the tableau method because no further exponential blow-ups occur in the prestate- and state-elimination phases. \square

There are various ways to restrict or parametrize the set of ATL^+ formulae in order to avoid the exponential blow-up of their closure sets. As suggested by the example after Lemma 3.4, the main cause for that blow-up of the number of γ -components of γ -formulae $\varphi = \langle\langle A \rangle\rangle\Phi$ or $\varphi = \llbracket A \rrbracket\Phi$ in ATL^+ is the nesting of conjunctions and disjunctions in the path formula Φ , which are not separated by temporal operators. Let us call the number of such nestings the *superficial Boolean depth* of Φ and denote it by $\delta_0(\Phi)$. Then, let the *nested Boolean depth* of any ATL^+ formula Ψ , denoted $\delta(\Psi)$, be the maximal superficial Boolean depth $\delta_0(\Phi)$ of a path subformula Φ of Ψ . For instance, $\delta(\langle\langle 1 \rangle\rangle \bigcirc \langle\langle 1 \rangle\rangle (p \mathcal{U} \neg q)) = 0$, $\delta(\langle\langle 1 \rangle\rangle (\Box p \vee ((q \wedge p) \mathcal{U} \neg q))) = 1$, $\delta(\langle\langle 1 \rangle\rangle (\Diamond q \wedge (\Box p \wedge (q \mathcal{U} \neg q)))) = 2$. Now, if this number for a formula η is bounded above, the size of the closure $\|\eta\|$ becomes polynomially bounded in $|\eta|$ because the nesting of \wedge and \vee when they are separated by a temporal operator does not have a multiplicative effect on the number of γ -components. Consequently, the complexity of the tableau method is reduced to single exponential time, caused only by the maximal possible number of states in the tableau, just like in ATL . Thus, we have the following.

PROPOSITION 6.22. *The tableau procedure for ATL^+ applied to a class of ATL^+ formulae of bounded nested Boolean depth runs in EXPTIME.*

COROLLARY 6.23. *The tableau procedure for ATL^+ applied to ATL formulae runs in EXPTIME.*

7. CONCLUSION

We developed a sound, complete, and terminating tableau-based decision method for constructive satisfiability testing of ATL^+ formulae and argued for its practical usability and implementability. The method is amenable to further extension to the full ATL^* , but this is left to future work.

Some comparison with the automata-based method for satisfiability testing in ATL^* presented in Schewe [2008] are in order. The two methods appear to be quite different and, although eventually working in the same worst-case complexity, the double exponential blow-ups seem to occur in different ways; namely, in the automata-based method, one exponential blow-up occurs in converting the formula into an automaton, whereas the other is in the time complexity of checking non-emptiness of the resulting automaton. It would be instructive to compare the practical implications and efficiency of both methods, and we leave such systematic comparison to the future, when (hopefully) both methods are implemented. For now, we only mention that the formula θ from our running example, the tableau for which is worked out explicitly and in detail in this article, is translated with the method from Schewe [2008] into an automaton with 2^{12} alphabet symbols and more than 100 states. Of course, this comparison cannot serve as an argument for general practical superiority in efficiency of the tableau-based method. Still, the technical details of both methods illustrated in that example indicate that, although the worst-case exponential blow-ups are bound to occur in both methods, they seem to be more controllable and avoidable in the tableau-based method, at the expense of its lesser automaticity and higher degree of user control. Thus, we would argue that both methods have generally incomparable pros and cons and, consequently, are of independent interest, both theoretically and practically.

ACKNOWLEDGMENTS

We thank the anonymous reviewers of Cerrito et al. [2014] and of this article for their helpful remarks and suggestions and for several corrections.

REFERENCES

- R. Alur, T. A. Henzinger, and O. Kupferman. 2002. Alternating-time temporal logic. *Journal of the ACM* 49, 5, 672–713.
- Mordechai Ben-Ari, Amir Pnueli, and Zohar Manna. 1983. The temporal logic of branching time. *Acta Informatica* 20, 207–226.
- Serenella Cerrito, Amélie David, and Valentin Goranko. 2014. Optimal tableaux-based decision procedure for testing satisfiability in the alternating-time temporal logic ATL^+ . In *Proceedings of IJCAR 2014 (LNCS)*, Vol. 8562. Springer, 277–291.
- Amélie David. 2013. TATL: Implementation of ATL tableau-based decision procedure. In *Proceedings of TABLEAUX 2013 (LNCS)*, Vol. 8123. Springer, 97–103.
- E. Emerson and J. Halpern. 1985. Decision procedures and expressiveness in the temporal logic of branching time. *Journal of Computation and System Sciences* 30, 1, 1–24.
- E. Allen Emerson. 1990. Temporal and modal logics. In *Handbook of Theoretical Computer Science*, J. van Leeuwen (Ed.). Vol. B. MIT Press, 995–1072.
- Valentin Goranko and Dmitry Shkatov. 2009. Tableau-based decision procedures for logics of strategic ability in multiagent systems. *ACM Transactions in Computer Logic* 11, 1, 1–49.
- V. Goranko and G. van Drimmelen. 2006. Complete axiomatization and decidability of alternating-time temporal logic. *Theoretic and Computational Science* 353, 93–117.
- J. Johannsen and M. Lange. 2003. CTL+ is complete for double exponential time. In *Proceedings of ICALP'03 (LNCS)*, Vol. 2719. Springer, 767–775.
- Vaughan R. Pratt. 1980. A near optimal method for reasoning about action. *Journal of Computer System Science* 20, 231–254.
- Sven Schewe. 2008. ATL^* satisfiability is 2EXPTIME-complete. In *Proceedings of ICALP (Part 2) (LNCS)*, Vol. 5126. Springer, 373–385.
- Govert van Drimmelen. 2003. Satisfiability in alternating-time temporal logic. In *Proceedings of the 18th Annual IEEE Symposium on Logic in Computer Science (LICS'03)*. IEEE Computer Society Press, 208–217.
- Dirk Walther, Carsten Lutz, Frank Wolter, and Michael Wooldridge. 2006. ATL satisfiability is indeed ExpTime-complete. *Journal of Logic and Computation* 16, 6, 765–787.
- Pierre Wolper. 1985. The tableau method for temporal logic: An overview. *Logique et Analyse* 28, 110–111, 119–136.

Received July 2014; revised May 2015; accepted July 2015