A tableau-based decision procedure for a branching-time interval temporal logic

Davide Bresolin and Angelo Montanari

Department of Mathematics and Computer Science,
University of Udine, Italy
E-mail: {bresolin|montana}@dimi.uniud.it

Abstract

Propositional interval temporal logics are very expressive temporal logics, with simple syntax and semantics, which allow one to naturally express statements that refer to time intervals and continuous processes. Most of them feature temporal operators that only allow one to express properties of a single timeline. In this paper we develop a branching-time interval neighborhood logic that interleaves operators that quantify over possible timelines with operators that quantify over intervals belonging to a given timeline. We define its syntax and semantics, and we provide it with a doubly-exponential tableau-based decision procedure.

1 Introduction

Propositional interval temporal logics are very expressive temporal logics, with simple syntax and semantics, which allow one to naturally express statements that refer to time intervals and continuous processes. Among them, we mention Halpern and Shoham's Modal Logic of Time Intervals (HS) [9], Venema's CDT logic [13], Moszkowski's Propositional Interval Temporal Logic (PITL) [11], and Goranko, Montanari, and Sciavicco's Propositional Neighborhood Logic (PNL), which is the propositional counterpart of Chaochen and Hansen's first-order Neighborhood Logic [2,6] (an up-to-date survey of the field can be found in [7]). Most of them feature temporal operators that only allow one to express properties of a single timeline, with the exception of Paech's Branching Regular Logic (BRL) [12]. BRL is a *local* branching-time interval logic, whose operators quantify over different timelines[1]. In [12] the author provides a Gentzen-style system for BRL and she states some expressiveness and complexity results.

[1] Locality forces a propositional variable to be true over an interval if and only if it is true at its starting point [11].

In this paper we develop a branching-time interval neighborhood logic that interleaves operators that quantify over possible timelines with operators that quantify over intervals belonging to a given timeline. We define its syntax and semantics, and we provide it with a doubly-exponential tableau-based decision procedure. Unlike the case of BRL, we do not impose any semantic restriction, such as locality, to get decidability.

When defining a temporal logic, there are basically two possible choices for the underlying temporal structure. Either time is *linear* (at any time there is only one possible future) or it has a *branching*, tree-like structure (any time may have many different futures). In the case of point-based temporal logics, both these alternatives have been successfully explored, and several meaningful logics have been developed (we only mention the linear temporal logic LTL and its many variants, and the branching-time temporal logics CTL and CTL^* [4]). On the contrary, interval-based temporal logics are usually interpreted over linear temporal structures. Even those interval logics which are interpreted over branching-time temporal structures, such as Halpern and Shoham's HS (in its original formulation) and Goranko, Montanari, and Sciavicco's branching CDT ($BCDT^+$) [5], only feature temporal operators that express properties of single timelines.

The main difference between interval-based and point-based temporal logics is that the former can express properties of *pairs* of time points (think of intervals as constructed out of points), rather than *single* time points. In most cases, this feature leads to undecidability since it prevents one from the possibility of reducing interval-based temporal logics to point-based ones. However, by imposing suitable *syntactic and/or semantic restrictions*, such a reduction can be defined, thus allowing one to benefit from the good computational properties of point-based logics [10]. A possible syntactic restriction is the choice of a suitable subset of interval modalities. As for semantic restrictions, one possibility is to constrain the truth value of propositional variables over intervals. We already mention the locality assumption. Another simplifying assumption often made is that of homogeneity, which states that a propositional variable is true over an interval if and only it it is true over all of its subintervals. Another possibility is to achieve decidability by constraining the classes of interval structures over which formulas are interpreted. The problem of identifying expressive enough, yet decidable, *genuinely* interval-based logics, that is, logics which are not directly translated into point-based logics and not invoking semantic restrictions, is still largely unexplored.

While various tableau methods have been developed for linear and branching time point-based temporal logics, not much work has been done on tableau methods for interval-based temporal logics. One reason for this disparity is that operators of interval temporal logics are in many respects more difficult to deal with [8]. In [5,8], Goranko et al. outline a tableau method for Venema's CDT logic interpreted over partial orders ($BCDT^+$). The method can be easily adapted to variations and subsystems of $BCDT^+$, thus providing a general tableau method for propositional interval logics. However, while most existing tableau methods for temporal logics are terminating methods for decidable logics, and thus they yield decision procedures, that for BCDT^+ only provides a semi-decision procedure for unsatisfiability. In [1], we propose a tableau-based decision procedure for the future fragment of (strict) Propositional Neighborhood Logic, that we call Right Propositional Neighborhood Logic (RPNL^- for short), interpreted over natural numbers. By combining syntactic restrictions (future temporal operators) and semantic ones (the domain of natural numbers), we succeeded in developing a tableau-based decision procedure for RPNL^-. To the best of our knowledge, RPNL^- is the first non-trivial case of a *genuine* propositional interval logic for which a tableau-based decision procedure has been given.

In this paper, we generalize the method to a suitable branching-time propositional interval temporal logic, that we call Branching-Time Right (Propositional) Neighborhood Logic (BTNL[R]^- for short). Such a logic combines the interval neighborhood operators of RPNL^- with the path quantifiers of CTL [4]. Emerson and Halpern's tableau-based decision procedure for CTL [3] consists of an initial construction phase, followed by an elimination phase. The elimination phase encompasses both a local pruning, that removes local inconsistencies, and a global pruning, that removes nodes including eventualities which are not fulfilled by the current graph. It turns out that a CTL formula is satisfiable if and only if the final graph is not empty. By combining such a tableau method for CTL and the one we developed for RPNL^- [1], we have been able to devise a doubly-exponential tableau-based decision procedure for BTNL[R]^-.

The rest of the paper is organized as follows. In Section 2, we introduce the syntax and semantics of BTNL[R]^-. In Section 3, we present our decision procedure, we prove its soundness and completeness, and we address complexity issues. In Section 4, we show our procedure at work on a simple example. Conclusions provide an assessment of the work and outline future research directions.

2 The Logic BTNL[R]^-

2.1 Tree-like structures

According to a commonly accepted perspective [4], the underlying temporal structure of branching-time temporal logics has a branching-like nature where each time point may have many successors points. The structure of time thus corresponds to an infinite tree. We shall further assume that the timeline defined by every (infinite) path in the tree is isomorphic to $\langle \mathbb{N}, < \rangle$. We allow a node in the tree to have infinitely many (possibly, uncountably many) successors, while we require each node to have at least one successor. It will turn out that, as far as our logic is concerned, such trees are undistinguishable

from trees with finite branching.
Given a directed graph $\mathbb{D} = \langle D, R \rangle$, a *finite R-sequence* over $\mathbb{D}$ is a sequence of nodes $d_1 d_2 \ldots d_n$, with $n \geq 2$ and $d_i \in D$ for $i = 1, \ldots, n$, such that $R(d_i, d_{i+1})$ for $i = 1, \ldots, n-1$. *Infinite R-sequences* can be defined analogously. We define a *path* π in $\mathbb{D}$ as a finite or infinite R-sequence. In the following, we shall take advantage of a relation $R^+ \subseteq D \times D$ such that $R^+(d_i, d_j)$ if and only if d_i and d_j are respectively the first and the last element of a finite R-sequence.
Temporal structures for branching time logics are infinite tree defined as follows.

Definition 2.1 An infinite *tree* is an infinite directed graph $\mathbb{D} = \langle D, R \rangle$, with a distinguished element $d_0 \in D$, called the *root* of the tree, where D is the set of nodes, called *time points*, and the set of edges $R \subseteq D \times D$ is a relation such that:

- for every $d(\neq d_0) \in D$, $R^+(d_0, d)$, that is, every $d(\neq d_0) \in D$ is R-reachable from d_0;
- for every $d(\neq d_0) \in D$, there exists at most one $d' \in D$ such that $R(d', d)$ (together with the previous one, this condition guarantees that every $d(\neq d_0) \in D$ has exactly one R-predecessor);
- there exists no d' such that $R(d', d_0)$, that is, d_0 has no R-predecessors;
- for every $d(\neq d_0) \in D$, there exists at least one $d' \in D$ such that $R(d, d')$, that is, every $d \in D$ has at least one R-successor.

It is not difficult to show that infinite trees are *acyclic* graphs, that is, there exist no finite paths which start from and end at the same node.
Given an infinite tree $\mathbb{D} = \langle D, R \rangle$, we can define a partial order $<$ over D such that, for every $d, d' \in D$, $d < d'$ if and only if $R^+(d, d')$. It is immediate to see that, for every infinite path π in $\mathbb{D}$, $\langle \pi, < \rangle$ is isomorphic to $\langle \mathbb{N}, < \rangle$.
Intervals over infinite trees are defined as follows. A *(strict) interval* over $\mathbb{D} = \langle D, R \rangle$ is a pair $[d_i, d_j]$ such that $d_i, d_j \in D$ and $d_i < d_j$. We denote the set of all strict intervals over a tree $\mathbb{D}$ as $\mathbb{I}(\mathbb{D})^-$. For any pair of intervals $[d_i, d_j]$ and $[d_j, d_k]$, we say that $[d_j, d_k]$ is a *right neighbor* of $[d_i, d_j]$. A point $d \in D$ belongs to the interval $[d_i, d_j]$ if $d_i \leq d \leq d_j$. In the following, we shall interpret our logic over *interval structures* $\langle \mathbb{D}, \mathbb{I}(\mathbb{D})^- \rangle$, where $\mathbb{D}$ is an infinite tree and $\mathbb{I}(\mathbb{D})^-$ is the set of all strict intervals over it.

2.2 Syntax and semantics of BTNL[R]$^-$

In this section, we give syntax and semantics of BTNL[R]$^-$ interpreted over infinite trees. BTNL[R]$^-$ is a propositional interval temporal logic based on the neighborhood relation between intervals. Its formulas are built from a set AP of propositional letters $p, q, \ldots$, by using the Boolean connectives $\neg$ and $\vee$ and the future temporal operators $E\langle A \rangle$ and $E[A]$. The other classical propositional connectives, as well as the logical constants $\top$ (true) and $\bot$

(false), are defined in the usual way. Furthermore, we introduce the temporal operator $A[A]$ as a shorthand for $\neg E\langle A\rangle\neg$ and the temporal operator $A\langle A\rangle$ as a shorthand for $\neg E[A]\neg$, and, from now on, we identify $\neg E\langle A\rangle\psi$ with $A[A]\neg\psi$ and $\neg E[A]\psi$ with $A\langle A\rangle\neg\psi$.

Formulas of BTNL[R]$^-$, denoted by $\varphi, \psi, \ldots$, are recursively defined by the following grammar:

$$\varphi = p \mid \neg\varphi \mid \varphi \vee \varphi \mid E\langle A\rangle\varphi \mid E[A]\varphi.$$

We denote by $|\varphi|$ the size of φ, that is, the number of symbols in φ, where the quantifier $E\langle A\rangle$ (resp., $E[A]$) is counted as one symbol. Whenever there are no ambiguities, we call a BTNL[R]$^-$ formula just a formula. A formula of the form $E\langle A\rangle\psi$, $E[A]\psi$, $A\langle A\rangle\psi$ or $A[A]\psi$, is called a *temporal formula*. A temporal formula whose main temporal operator is either $E\langle A\rangle$ or $E[A]$ is called an *existential formula*, while a temporal formula whose main temporal operator is either $A\langle A\rangle$ or $A[A]$ is called a *universal formula*.

A *model* for a formula is a tuple $\mathbf{M}^- = \langle\langle\mathbb{D}, \mathbb{I}^-(\mathbb{D})\rangle, \mathcal{V}\rangle$, where the pair $\langle\mathbb{D}, \mathbb{I}(\mathbb{D})^-\rangle$ is an interval structure, and $\mathcal{V} : \mathbb{I}(\mathbb{D})^- \longrightarrow 2^{AP}$ is a *valuation function* that assigns to every interval the set of propositional letters true on it. We place ourselves in the most general setting, and we do not impose any locality or monotonicity constraint on the valuation function. As an example, it may happen that $\mathcal{V}$ maps an interval $[d_i, d_j]$ to some set of atomic propositions and that it maps an interval contained in/that contains $[d_i, d_j]$ to a different set.

The semantics of BTNL[R]$^-$ is defined recursively by the *satisfiability relation* $\Vdash$ as follows. Let $\mathbf{M}^- = \langle\langle\mathbb{D}, \mathbb{I}^-(\mathbb{D})\rangle, \mathcal{V}\rangle$ be some given model, and let $[d_i, d_j] \in \mathbb{I}(\mathbb{D})^-$:

- for every propositional letter $p \in AP$, $\mathbf{M}^-, [d_i, d_j] \Vdash p$ iff $p \in \mathcal{V}([d_i, d_j])$;
- $\mathbf{M}^-, [d_i, d_j] \Vdash \neg\psi$ iff $\mathbf{M}^-, [d_i, d_j] \nVdash \psi$;
- $\mathbf{M}^-, [d_i, d_j] \Vdash \psi_1 \vee \psi_2$ iff $\mathbf{M}^-, [d_i, d_j] \Vdash \psi_1$, or $\mathbf{M}^-, [d_i, d_j] \Vdash \psi_2$;
- $\mathbf{M}^-, [d_i, d_j] \Vdash E\langle A\rangle\psi$ iff there exists $d_k \in D$, $d_j < d_k$, such that $\mathbf{M}^-, [d_j, d_k] \Vdash \psi$;
- $\mathbf{M}^-, [d_i, d_j] \Vdash E[A]\psi$ iff there exists an infinite path $\pi = d_j d_{j+1} \ldots$ rooted at d_j such that, for every d_k in π, with $d_j < d_k$, $\mathbf{M}^-, [d_j, d_k] \Vdash \psi$.

Since our logic has only future time operators, without loss of generality, we can restrict our attention to *standard trees*, namely, trees such that the root d_0 has only one R-successor d_1. In such a case, the corresponding interval structure has an *initial interval* $[d_0, d_1]$, and we say that a formula φ is satisfied by the interval structure if and only if φ holds at the initial interval.

3 A tableau-based decision procedure for BTNL[R]$^-$

To check the satisfiability of a formula φ, we build a graph, that we call a *tableau* for φ, whose nodes represent points of the temporal domain $\mathbb{D}$ and whose edges represent the relation R connecting a point to its successors in the tree. We shall take advantage of such a construction to reduce the problem of finding a model for φ to the problem of testing whether the tableau satisfies some suitable properties or not.

Let $\mathbf{M}^- = \langle\langle\mathbb{D}, \mathbb{I}^-(\mathbb{D})\rangle, \mathcal{V}\rangle$ be a model for φ and let d_j be a point in $\mathbb{D}$. We have that, given an interval $[d_i, d_j]$ ending in d_j, every right neighbor of it is also a right neighbor of every other interval $[d_k, d_j]$ ending in d_j. Hence, every *temporal formula* φ holds over $[d_i, d_j]$ if and only if it holds over every other interval $[d_k, d_j]$ ending in d_j. We denote by $\mathrm{REQ}(d_j)$ the set of temporal formulas which hold over all intervals ending in d_j.

The building blocks for the tableau construction are φ-atoms. For every interval $[d_i, d_j] \in \mathbb{I}^-(\mathbb{D})$, we introduce a pair of sets of formulas $(\mathcal{A}_{[d_i,d_j]}, \mathcal{C}_{[d_i,d_j]})$, that we call an *atom* for φ (φ-atom for short). The set $\mathcal{A}_{[d_i,d_j]}$ is a subset of $\mathrm{REQ}(d_i)$, which collects the set of requests in $\mathrm{REQ}(d_i)$ relevant to the interval $[d_i, d_j]$. In general, $\mathcal{A}_{[d_i,d_j]}$ may differ from $\mathcal{A}_{[d_i,d_k]}$ for $j \neq k$. The set $\mathcal{C}_{[d_i,d_j]}$ contains all and only the formulas that (should) hold over $[d_i, d_j]$. We can associate with every point $d_j \in D$ the set of φ-atoms $\{(\mathcal{A}_{[d_i,d_j]}, \mathcal{C}_{[d_i,d_j]}) : d_i < d_j\}$, which includes all φ-atoms paired with intervals ending in d_j. These sets of atoms are the (macro)nodes of the tableau for φ.

3.1 Basic notions

Let φ be a BTNL[R]$^-$-formula to be checked for satisfiability and let AP be the set of its propositional variables.

Definition 3.1 The *closure* $\mathrm{CL}(\varphi)$ of φ is the set of all subformulas of φ and their negations (we identify $\neg\neg\psi$ with ψ).

Definition 3.2 The set of *temporal requests* of φ is the set $\mathrm{TF}(\varphi)$ of all temporal formulas in $\mathrm{CL}(\varphi)$.

By induction on the structure of φ, it can be easily proved that $|\,\mathrm{CL}(\varphi)| \leq 2\cdot|\varphi|$ and $|\,\mathrm{TF}(\varphi)| \leq 2 \cdot |\varphi|$.

We are now ready to introduce the key notion of φ-atom.

Definition 3.3 Let φ be a BTNL[R]$^-$-formula. A *φ-atom* is a pair $(\mathcal{A}, \mathcal{C})$, with $\mathcal{A} \subseteq \mathrm{TF}(\varphi)$ and $\mathcal{C} \subseteq \mathrm{CL}(\varphi)$, such that:
- for every $\psi \in \mathrm{CL}(\varphi)$, $\psi \in \mathcal{C}$ iff $\neg\psi \notin \mathcal{C}$;
- for every $\psi_1 \vee \psi_2 \in \mathrm{CL}(\varphi)$, $\psi_1 \vee \psi_2 \in \mathcal{C}$ iff $\psi_1 \in \mathcal{C}$ or $\psi_2 \in \mathcal{C}$;
- for every $\psi \in \mathrm{TF}(\varphi)$, if $\psi \in \mathcal{A}$ then $\neg\psi \notin \mathcal{A}$;
- for every $A[A]\psi \in \mathcal{A}$, $\psi \in \mathcal{C}$;
- for every $E[A]\psi \in \mathcal{A}$, $\psi \in \mathcal{C}$.

Formulas in $\mathcal{C}$ are called *current formulas*, while (temporal) formulas in $\mathcal{A}$ are called *active requests*.

As for the set $\mathcal{A}$, it is worth pointing out that there may exist a φ-atom $(\mathcal{A}, \mathcal{C})$

and a temporal formula ψ such that neither ψ nor $\neg\psi$ belongs to $\mathcal{A}$.
We denote the set of all φ-atoms by $\mathcal{A}_\varphi$. It is not difficult to show that $|\mathcal{A}_\varphi| \leq 2^{2|\varphi|}$.
Atoms come into play in the proposed tableau method as follows. The method associates an atom $(\mathcal{A}, \mathcal{C})$ with any interval $[d_i, d_j]$. The set $\mathcal{A}$ includes all formulas of the form $A[A]\psi$ that belong to $\mathrm{REQ}(d_i)$ as well as some formulas of the forms $E[A]\psi$, $A\langle A\rangle\psi$, and $E\langle A\rangle\psi$ in $\mathrm{REQ}(d_i)$; the set $\mathcal{C}$ includes all formulas $\psi \in \mathrm{CL}(\varphi)$ which (should) hold over $[d_i, d_j]$. Moreover, for all formulas of the forms $A[A]\psi$ and $E[A]\psi$ in $\mathcal{A}$, we put ψ into $\mathcal{C}$, while for any formula of the forms $A\langle A\rangle\psi$ and $E\langle A\rangle\psi$ in $\mathcal{A}$, it may happen that $\psi \in \mathcal{C}$, but this is not necessarily the case.

Atoms are connected by the following binary relation.

Definition 3.4 Let X_φ be a binary relation over $\mathcal{A}_\varphi$ such that, for every pair of atoms $(\mathcal{A}, \mathcal{C}), (\mathcal{A}', \mathcal{C}') \in \mathcal{A}_\varphi$, $(\mathcal{A}, \mathcal{C}) X_\varphi (\mathcal{A}', \mathcal{C}')$ if (and only if):

- $\mathcal{A}' \subseteq \mathcal{A}$;
- for every $A[A]\psi \in \mathcal{A}$, $A[A]\psi \in \mathcal{A}'$;
- for every $A\langle A\rangle\psi \in \mathcal{A}$, $A\langle A\rangle\psi \in \mathcal{A}'$ iff $\neg\psi \in \mathcal{C}$.

In the next section we shall show that for any pair of points $d_i < d_j$, the relation X_φ connects the atom associated with the interval $[d_i, d_j]$ to the atom associated with the interval $[d_i, d_{j+1}]$, where d_{j+1} is an R-successor of d_j. In particular, it will turn out that, if $(\mathcal{A}, \mathcal{C})$ is associated with the interval $[d_i, d_j]$, then for every formula $A[A]\psi \in \mathrm{REQ}(d_i)$ and every atom $(\mathcal{A}', \mathcal{C}')$ such that $(\mathcal{A}, \mathcal{C}) X_\varphi (\mathcal{A}', \mathcal{C}')$, we have that $A[A]\psi \in \mathcal{A}'$ (and thus $\psi \in \mathcal{C}'$), while for every formula $A\langle A\rangle\psi \in \mathcal{A}$, if $[d_i, d_j]$ satisfies ψ then $A\langle A\rangle\psi \notin \mathcal{A}'$. This guarantees that temporal requests of the form $A[A]\psi$ are propagated through X_φ-successors, while temporal requests of the form $A\langle A\rangle\psi$ are discarded once that ψ has been satisfied by the set of current formulas of some atom.

3.2 Tableau construction

Definition 3.5 A *node* is a set N of φ-atoms such that, for any pair $(\mathcal{A}, \mathcal{C})$, $(\mathcal{A}', \mathcal{C}') \in N$ and any $\psi \in \mathrm{TF}(\varphi)$, $\psi \in \mathcal{C} \Leftrightarrow \psi \in \mathcal{C}'$. We denote by $\mathcal{N}_\varphi$ the set of all nodes that can be built from $\mathcal{A}_\varphi$ and by $Init(\mathcal{N}_\varphi)$ the subset of all *initial nodes*, that is, the set $\{\{(\emptyset, \mathcal{C})\} \in \mathcal{N}_\varphi : \varphi \in \mathcal{C}\}$. Furthermore, for any node N, we denote by $\mathrm{REQ}(N)$ the set $\{\psi \in \mathrm{TF}(\varphi) : \exists(\mathcal{A}, \mathcal{C}) \in N(\psi \in \mathcal{C})\}$ (or, equivalently, the set $\{\psi \in \mathrm{TF}(\varphi) : \forall(\mathcal{A}, \mathcal{C}) \in N(\psi \in \mathcal{C})\}$).

From Definition 3.5, it follows that $|\mathcal{N}_\varphi| \leq 2^{2^{2\cdot|\varphi|}}$.

Nodes can be viewed as (maximal) collections of intervals ending at the same point of the temporal domain, that is, we can associate every node N with a point $d_j \in D$ and every atom $(\mathcal{A}, \mathcal{C}) \in N$ with some interval $[d_i, d_j]$ ending in d_j. Accordingly, we have that $\mathrm{REQ}(N) = \mathrm{REQ}(d_j)$. The relation between a node N, associated with point d_j, and a node M, associated with an R-

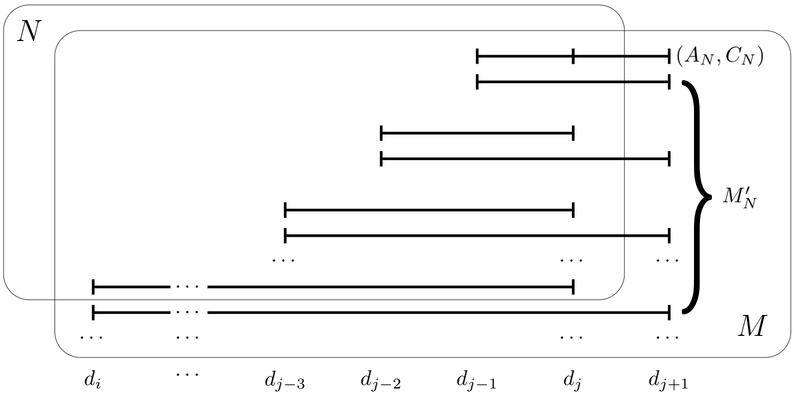

Figure 1. Connecting two nodes.

successor d_{j+1} of d_j, as well as the relations between intervals ending in d_j and intervals ending in d_{j+1} (and, thus, between atoms in N and atoms in M), are graphically depicted in Figure 1. We have that *(i)* the interval $[d_j, d_{j+1}]$ is a right neighbor of all intervals $[d_i, d_j]$ ending in d_j, and *(ii)* for every interval $[d_i, d_j]$ there exists an interval $[d_i, d_{j+1}]$. Thus, M should contain:

- an atom $(\mathcal{A}_N, \mathcal{C}_N)$ associated to the interval $[d_j, d_{j+1}]$;
- for every atom $(\mathcal{A}, \mathcal{C}) \in N$, an atom $(\mathcal{A}', \mathcal{C}')$ such that $(\mathcal{A}, \mathcal{C}) X_\varphi (\mathcal{A}', \mathcal{C}')$.

Definition 3.6 The *tableau* for a BTNL[R]$^-$-formula φ is a (finite) directed graph $\mathcal{T}_\varphi = \langle \mathcal{N}_\varphi, \mathcal{R}_\varphi \rangle$, where for any pair $N, M \in \mathcal{N}_\varphi$, $(N, M) \in \mathcal{R}_\varphi$ if and only if $M = \{(\mathcal{A}_N, \mathcal{C}_N)\} \cup M'_N$, where

1. $(\mathcal{A}_N, \mathcal{C}_N)$ is an atom such that $\mathcal{A}_N \subseteq \mathrm{REQ}(N)$ and for all universal formulas $\psi \in \mathrm{REQ}(N)$, $\psi \in \mathcal{A}_N$;
2. for every $(\mathcal{A}, \mathcal{C}) \in N$, there exists $(\mathcal{A}', \mathcal{C}') \in M'_N$ such that $(\mathcal{A}, \mathcal{C}) X_\varphi (\mathcal{A}', \mathcal{C}')$;
3. for every $(\mathcal{A}', \mathcal{C}') \in M'_N$, there exists $(\mathcal{A}, \mathcal{C}) \in N$ such that $(\mathcal{A}, \mathcal{C}) X_\varphi (\mathcal{A}', \mathcal{C}')$.

Let $N, M \in \mathcal{N}_\varphi$. If $(N, M) \in \mathcal{R}_\varphi$, we say that M is an $\mathcal{R}_\varphi$-*successor* of N. We say that M is an $\mathcal{R}_\varphi$-*descendant* of N if there exists a (finite) path from N to M in $\mathcal{T}_\varphi$.

Definition 3.7 Given a (finite or infinite) path $\pi = N_1 N_2 \ldots$ in $\mathcal{T}_\varphi$, an *atom path* in π is a sequence of atoms $(\mathcal{A}_1, \mathcal{C}_1)(\mathcal{A}_2, \mathcal{C}_2) \ldots$ such that:

- for every $i \geq 1$, $(\mathcal{A}_i, \mathcal{C}_i) \in N_i$;
- for every $1 \leq i (< n)$, $(\mathcal{A}_i, \mathcal{C}_i) X_\varphi (\mathcal{A}_{i+1}, \mathcal{C}_{i+1})$.

Given a node N and an atom $(\mathcal{A}, \mathcal{C}) \in N$, we say that the atom $(\mathcal{A}', \mathcal{C}')$ is an X_φ-*descendant* of $(\mathcal{A}, \mathcal{C})$ if and only if there exists a node M such that $(\mathcal{A}', \mathcal{C}') \in M$, and there exists a path π from N to M such that there is an atom path from $(\mathcal{A}, \mathcal{C})$ to $(\mathcal{A}', \mathcal{C}')$ in π.

Definition 3.8 An infinite path $\pi = N_1 N_2 \ldots$ in $\mathcal{T}_\varphi$ is a *fulfilling path* if and only if, for every $i \geq 1$, every atom $(\mathcal{A}, \mathcal{C}) \in N_i$, and every formula $A\langle A \rangle \psi \in \mathcal{A}$, either $\psi \in \mathcal{C}$ or there exist N_j, with $j > i$, and $(\mathcal{A}', \mathcal{C}') \in N_j$ such that $(\mathcal{A}', \mathcal{C}')$

is an X_φ-descendant of $(\mathcal{A}, \mathcal{C})$ in π and $\psi \in \mathcal{C}'$.

Definition 3.9 A *substructure* is a subgraph $\langle \mathcal{N}, \mathcal{R} \rangle \subseteq \mathcal{T}_\varphi$ such that:

- there exists a node $N_0 \in \mathcal{N} \cap Init(\mathcal{N}_\varphi)$ (*initial node*) such that all other nodes in $\mathcal{N}$ are $\mathcal{R}$-reachable from it;
- for every node $N \in \mathcal{N}$, there exists a fulfilling path (in $\langle \mathcal{N}, \mathcal{R} \rangle$) starting from N.

Substructures represent *candidate models* for φ. The truth of formulas devoid of temporal operators and of formulas of the form $A[A]\psi$, indeed, follows from Definition 3.3. Moreover, the truth of formulas of the form $A\langle A \rangle \psi$ follows from Definition 3.8. However, to obtain a model for φ we must also guarantee the truth of formulas of the forms $E[A]\psi$ and $E\langle A \rangle \psi$. To this end, we introduce the notion of *fulfilling substructure*.

Definition 3.10 A substructure $\langle \mathcal{N}, \mathcal{R} \rangle \subseteq \mathcal{T}_\varphi$ is *fulfilling* if and only if, for every node $N \in \mathcal{N}$ and every atom $(\mathcal{A}, \mathcal{C}) \in N$, the following conditions hold:

(F1) for every existential formula $\psi \in \mathcal{C}$, there exist an $\mathcal{R}$-successor M of N and an atom $(\mathcal{A}', \mathcal{C}') \in M$ such that $\psi \in \mathcal{A}'$ and $\mathcal{A}' \subseteq \mathrm{REQ}(N)$, and for all universal formulas $\xi \in \mathrm{REQ}(N)$, $\xi \in \mathcal{A}'$;

(F2) for every formula $E\langle A \rangle \psi \in \mathcal{A}$, either $\psi \in \mathcal{C}$ or there exist an $\mathcal{R}$-descendant M of N and an X_φ-descendant $(\mathcal{A}', \mathcal{C}')$ of $(\mathcal{A}, \mathcal{C})$ in M such that $\psi \in \mathcal{C}'$;

(F3) for every formula $E[A]\psi \in \mathcal{A}$, there exist a fulfilling path $\pi = N_0 N_1 N_2 \ldots$ and an atom path $(\mathcal{A}_0, \mathcal{C}_0)(\mathcal{A}_1, \mathcal{C}_1)(\mathcal{A}_2, \mathcal{C}_2) \ldots$ in π such that: *(i)* $(\mathcal{A}_0, \mathcal{C}_0) = (\mathcal{A}, \mathcal{C})$; *(ii)* $N_0 = N$; *(iii)* for every $i \geq 0$, $E[A]\psi \in \mathcal{A}_i$; *(iv)* for every formula $A\langle A \rangle \theta \in \mathcal{A}_0$, there exists $j \geq 0$ such that $\theta \in \mathcal{C}_j$.

Theorem 3.11 *If the formula φ is satisfiable (in a standard tree), then there exists a fulfilling substructure $\langle \mathcal{N}, \mathcal{R} \rangle \subseteq \mathcal{T}_\varphi$.*

Proof Let $\mathbf{M}^- = \langle \langle \mathbb{D}, \mathbb{I}(\mathbb{D})^- \rangle, \mathcal{V} \rangle$ be a model for φ such that $\mathbb{D} = \langle D, R \rangle$, with distinguished element d_0, is a standard tree. For every interval $[d_i, d_j] \in \mathbb{I}(\mathbb{D})^-$, we define an atom $(\mathcal{A}_{[d_i,d_j]}, \mathcal{C}_{[d_i,d_j]})$ as follows.

- $\mathcal{A}_{[d_i,d_j]}$ contains exactly:
  - · all formulas $A[A]\psi \in \mathrm{REQ}(d_i)$;
  - · all formulas $A\langle A \rangle \psi \in \mathrm{REQ}(d_i)$ such that, for every $d_i < d_l < d_j$, $\mathbf{M}^-, [d_i, d_l] \Vdash \neg\psi$;
  - · all formulas $E[A]\psi \in \mathrm{REQ}(d_i)$ such that there exists an infinite path $\pi = d_i d_{i+1} \ldots d_j d_{j+1} \ldots$, starting from d_i and containing d_j, such that $\mathbf{M}^-, [d_i, d_k] \Vdash \psi$ for every $d_k \in \pi$;
  - · all formulas $E\langle A \rangle \psi \in \mathrm{REQ}(d_i)$ such that there exists $d_k \geq d_j$ such that $\mathbf{M}^-, [d_i, d_k] \Vdash \psi$ and, for every $d_i < d_l < d_k$, $\mathbf{M}^-, [d_i, d_l] \Vdash \neg\psi$;
- $\mathcal{C}_{[d_i,d_j]}$ contains exactly all formulas $\psi \in \mathrm{CL}(\varphi)$ such that $\mathbf{M}^-, [d_i, d_j] \Vdash \psi$.

It is easy to check that, for every $[d_i, d_j] \in \mathbb{I}(\mathbb{D})^-$ and for every R-successor

d_{j+1} of d_j, $(\mathcal{A}_{[d_i,d_j]}, \mathcal{C}_{[d_i,d_j]})$ is an atom such that $(\mathcal{A}_{[d_i,d_j]}, \mathcal{C}_{[d_i,d_j]}) X_\varphi (\mathcal{A}_{[d_i,d_{j+1}]}, \mathcal{C}_{[d_i,d_{j+1}]})$.

For every $d_j \in D$, with $d_j \neq d_0$, let $N_j = \{(\mathcal{A}_{[d_i,d_j]}, \mathcal{C}_{[d_i,d_j]}) : d_i < d_j\}$, $\mathcal{N} = \{N_j : d_j \in D, d_j \neq d_0\}$, and $\mathcal{R} = \mathcal{R}_\varphi \cap (\mathcal{N} \times \mathcal{N})$. It is easy to check that for all $d_j \neq d_0$, N_j is a node and that $\langle \mathcal{N}, \mathcal{R} \rangle$ is a fulfilling substructure. $\square$

The next theorem shows that a model for φ can be obtained by unfolding a fulfilling substructure $\langle \mathcal{N}, \mathcal{R} \rangle$, starting from its initial node N_0.

Theorem 3.12 *If there exists a fulfilling substructure $\langle \mathcal{N}, \mathcal{R} \rangle \subseteq \mathcal{T}_\varphi$, then the formula φ is satisfiable.*

Proof Let $\langle \mathcal{N}, \mathcal{R} \rangle$ be a fulfilling substructure. To define a model for φ, we first build an infinite tree $\mathbb{D} = \langle D, R \rangle$ by *unfolding* the fulfilling substructure $\langle \mathcal{N}, \mathcal{R} \rangle$ from its initial node N_1 as follows:

- D is the (infinite) set of all finite $\mathcal{R}$-sequences $N_1 N_2 \dots N_k$ starting from the initial node N_1, including the empty sequence ε (the root of the tree);
- R is such that, for any pair of points $d, d' \in D$, $(d, d') \in R$ if and only if $d = N_1 \dots N_k$ and $d' = N_1 \dots N_k N_{k+1}$ and N_{k+1} is an $\mathcal{R}$-successor of N_k.

Notice that $\mathbb{D}$ is a standard tree with distinguished element $d_0 = \varepsilon$, since the only successor of d_0 is $d_1 = N_1$.

Given an arbitrary order of the nodes of $\mathcal{N}$, we define a *total* order $\lhd$ over finite $\mathcal{R}$-sequences, that is, on points of D, as follows:

- given two $\mathcal{R}$-sequences d and d' such that the length of d is less than the length of d', we have that $d \lhd d'$;
- given two $\mathcal{R}$-sequences d and d' of the same length, $d \lhd d'$ if and only if d precedes d' on the lexicographical order based on the given (arbitrary) order of nodes of $\mathcal{N}$.

In order to build a model for φ, we define a suitable (partial) labelling function $L : \mathbb{I}(\mathbb{D})^- \to A_\varphi$. For any $d_j = N_1 \dots N_k$, L associates an atom $(\mathcal{A}, \mathcal{C}) \in N_k$ with any interval $[d_i, d_j]$. We define L by (infinite) induction on the total order $\lhd$.

Base case. We start by defining the labelling of the initial interval $[d_0, d_1]$ (where $d_0 = \varepsilon$ and $d_1 = N_1$). Since N_1 is the initial node of $\langle \mathcal{N}, \mathcal{R} \rangle$, we have that $N_1 = \{(\emptyset, \mathcal{C})\}$, with $\varphi \in \mathcal{C}$. We put $L([d_0, d_1]) = (\emptyset, \mathcal{C})$. Two cases may arise.

$\mathcal{C}$ contains no existential formulas. In such a case, we define the labelling of an infinite branch starting from d_1 (remind that we admit only infinite models). Since $\langle \mathcal{N}, \mathcal{R} \rangle$ is a substructure, there exists a fulfilling path $N_1 N_2 \dots$ starting from N_1. Let $\pi = d_1 d_2 \dots$ be the corresponding infinite branch in $\mathbb{D}$ (where, for every $i \geq 1$, $d_i = N_1 \dots N_i$). We define the labelling $L([d_i, d_j])$ of every interval $[d_i, d_j]$, with $d_i, d_j \in \pi$, in such a way that:

- for all $j > 1$, $L([d_{j-1}, d_j]) = (\mathcal{A}', \mathcal{C}')$ is such that $\mathcal{A}' \subseteq \mathrm{REQ}(N_{j-1})$ and, for all universal formulas $\theta \in \text{REQ}(N_{j-1})$, $\theta \in \mathcal{A}'$ (see condition 1 of Definition 3.6);

- for all $j > 1$ and all intervals $[d_i, d_j]$, with $d_j > d_{i+1}$, $L([d_i, d_{j-1}])X_\varphi$ $L([d_i, d_j])$ (see conditions 2 and 3 of Definition 3.6);
- for every interval $[d_i, d_j]$, with $L([d_i, d_j]) = (\mathcal{A}', \mathcal{C}')$, if there exists a formula $A\langle A\rangle\theta \in \mathcal{A}'$, then there exists a point $d_k \geq d_j$ such that $L([d_i, d_k]) = (\mathcal{A}'', \mathcal{C}'')$ and $\theta \in \mathcal{C}''$ (see Definition 3.8).

Definitions 3.6 and 3.8 guarantee that there exists a labelling with such properties.

$\mathcal{C}$ *contains at least one existential formula*. In such a case, for every existential formula $\psi \in \mathcal{C}$, we guarantee that ψ gets satisfied by properly labelling an infinite branch starting from d_1. Let $\psi \in \mathcal{C}$ be an existential formula. Two cases may arise (depending on the structure of ψ).

- $\psi = E\langle A\rangle\theta$. By condition F1, there exists an $\mathcal{R}$-successor N_2 of N_1 and an atom $(\mathcal{A}', \mathcal{C}') \in N_2$ such that $\psi \in \mathcal{A}'$. Let $d_2 = N_1N_2$. We put $L([d_1, d_2]) = (\mathcal{A}', \mathcal{C}')$. By condition F2, either $\theta \in \mathcal{C}'$ (and thus θ is satisfied over the interval $[d_1, d_2]$) or there exists a path $N_2N_3 \ldots N_k$ and a corresponding atom path $(\mathcal{A}', \mathcal{C}') = (\mathcal{A}_2, \mathcal{C}_2)(\mathcal{A}_3, \mathcal{C}_3) \ldots (\mathcal{A}_k, \mathcal{C}_k)$, such that $\theta \in \mathcal{C}_k$. In the latter case, let $d_2d_3 \ldots d_k$ be the branch in $\mathbb{D}$ corresponding to $N_2N_3 \ldots N_k$. We put, for every $3 \leq i \leq k$, $L([d_1, d_i]) = (\mathcal{A}_i, \mathcal{C}_i)$, in order to satisfy θ over $[d_1, d_k]$. In both cases, we extend the finite branch to an infinite one (starting from N_2 or from N_k, respectively) and we define the labelling of all other intervals on the branch as in the case in which $\mathcal{C}$ contains no existential formulas.
- $\psi = E[A]\theta$. By condition F1, there exists an $\mathcal{R}$-successor N_2 of N_1 and an atom $(\mathcal{A}', \mathcal{C}') \in N_2$ such that $\psi \in \mathcal{A}'$. Let $d_2 = N_1N_2$. We put $L([d_1, d_2]) = (\mathcal{A}', \mathcal{C}')$. By condition F3, there exist a fulfilling infinite path $N_2N_3 \ldots$ and a corresponding infinite atom path $(\mathcal{A}', \mathcal{C}') = (\mathcal{A}_2, \mathcal{C}_2)(\mathcal{A}_3, \mathcal{C}_3) \ldots$, such that $E[A]\theta \in \mathcal{A}_i$ for every $i \geq 2$. Let $d_2d_3 \ldots$ be the infinite branch in $\mathbb{D}$ corresponding to $N_2N_3 \ldots$. We put, for every $i \geq 3$, $L([d_1, d_i]) = (\mathcal{A}_i, \mathcal{C}_i)$, in order to satisfy θ over $[d_1, d_i]$. As before, we define the labelling of all other intervals on the branch as in the case in which $\mathcal{C}$ contains no existential formulas.

We repeat such a procedure for every existential formula in $\mathcal{C}$.

Inductive step. Let $d \in D$ such that *(i)* L is defined over all intervals $[d', d]$; *(ii)* for all $d' \lhd d$, either d' has been already taken into consideration or L is not defined over any interval $[d'', d']$; *(iii)* d has not been taken into consideration yet. Consider the set $\text{REQ}(N)$, with N such that $d = N_1 \ldots N$, and suppose there exists an existential formula $\psi \in \text{REQ}(N)$. Two cases may arise. Either ψ is satisfied by the current labelling L, and we are done, or there are no branches starting from d that satisfy ψ. In the latter case, we satisfy ψ by defining a suitable labelling of an unlabeled infinite branch starting from d, as we have done in the base case of the induction. By repeating such a procedure for every existential formula $\psi \in \mathrm{REQ}(N)$, we guarantee that all existential formulas in $\mathrm{REQ}(N)$ are satisfied.

The model satisfying φ contains all and only the labelled (infinite) branches of $\mathbb{D}$. Let $\mathbb{D}'$ be the infinite tree obtained from $\mathbb{D}$ by removing all unlabeled branches and let $\mathcal{V}$ be a valuation function $\mathcal{V}$ such that, for every $p \in AP$ and $[d_i, d_j] \in \mathbb{I}(\mathbb{D}')^-$, $p \in \mathcal{V}([d_i, d_j])$ if and only if $L([d_i, d_j]) = (\mathcal{A}_{[d_i,d_j]}, \mathcal{C}_{[d_i,d_j]})$ and $p \in \mathcal{C}_{[d_i,d_j]}$. We prove by induction on the structure of $\psi \in \mathrm{CL}(\varphi)$ that, for every $[d_i, d_j] \in \mathbb{I}(\mathbb{D}')^-$, we have that $\mathbf{M}^-, [d_i, d_j] \Vdash \psi$ if and only if $L([d_i, d_j]) = (\mathcal{A}_{[d_i,d_j]}, \mathcal{C}_{[d_i,d_j]})$ and $\psi \in \mathcal{C}_{[d_i,d_j]}$.

- The base case, as well as the case of the propositional connectives $\neg$ and $\vee$, are straightforward.
- Let ψ be the formula $E\langle A\rangle\chi$. Suppose that $E\langle A\rangle\chi \in \mathcal{C}_{[d_i,d_j]}$. By the definition of L, there exists an infinite path $d_j d_{j+1} \dots$ and an interval $[d_j, d_k]$ such that $\chi \in \mathcal{C}_{[d_j,d_k]}$. By inductive hypothesis, we have that $\mathbf{M}^-, [d_j, d_k] \Vdash \chi$, and thus $\mathbf{M}^-, [d_i, d_j] \Vdash E\langle A\rangle\chi$.

  As for the opposite implication, assume by contradiction that $\mathbf{M}^-, [d_i, d_j] \Vdash E\langle A\rangle\chi$ and $E\langle A\rangle\chi \notin \mathcal{C}_{[d_i,d_j]}$. By atom definition, this implies that $\neg E\langle A\rangle\chi = A[A]\neg\chi \in \mathcal{C}_{[d_i,d_j]}$. By the definition of L, we have that $A[A]\neg\chi \in \mathcal{A}_{[d_j,d_k]}$ for every $d_k > d_j$, and thus $\neg\chi \in \mathcal{C}_{[d_j,d_k]}$. By inductive hypothesis, this implies that $\mathbf{M}^-, [d_j, d_k] \Vdash \neg\chi$ for every $d_k > d_j$, and thus $\mathbf{M}^-, [d_i, d_j] \Vdash A[A]\neg\chi$, which contradicts the hypothesis that $\mathbf{M}^-, [d_i, d_j] \Vdash E\langle A\rangle\chi$.
- Let ψ be the formula $E[A]\chi$. Suppose that $E[A]\chi \in \mathcal{C}_{[d_i,d_j]}$. By the definition of L, there exists an infinite path $\pi = d_j d_{j+1} \dots$ such that, for every $d_k \in \pi$, $d_k > d_j$, $E[A]\chi \in \mathcal{A}_{[d_j,d_k]}$. By atom definition, this implies that $\chi \in \mathcal{C}_{[d_j,d_k]}$ and, by inductive hypothesis, we have that $\mathbf{M}^-, [d_j, d_k] \Vdash \chi$, for every $d_k \in \pi$, $d_k > d_j$, and thus $\mathbf{M}^-, [d_i, d_j] \Vdash E[A]\chi$.

  As for the opposite implication, assume by contradiction that $\mathbf{M}^-, [d_i, d_j] \Vdash E[A]\chi$ and $E[A]\chi \notin \mathcal{C}_{[d_i,d_j]}$. By atom definition, this implies that $\neg E[A]\chi = A\langle A\rangle\neg\chi \in \mathcal{C}_{[d_i,d_j]}$. By the definition of L, we have that, for every infinite path $d_j d_{j+1} \dots$ starting from d_j, there exists $d_k \in \pi$, $d_k > d_j$ such that $\neg\chi \in \mathcal{C}_{[d_j,d_k]}$. By inductive hypothesis, this implies that for every infinite path $\pi = d_j d_{j+1} \dots$ there exists a point $d_k \in \pi$, $d_k > d_j$, such that $\mathbf{M}^-, [d_j, d_k] \Vdash \neg\chi$, and thus $\mathbf{M}^-, [d_i, d_j] \Vdash A\langle A\rangle\neg\chi$, which contradicts the hypothesis that $\mathbf{M}^-, [d_i, d_j] \Vdash E[A]\chi$.

Since $\langle \mathcal{N}, \mathcal{R}\rangle$ is a substructure, $\varphi \in \mathcal{C}_{[d_0,d_1]}$, and thus $\mathbf{M}^-, [d_0, d_1] \Vdash \varphi$. □

3.3 The decision procedure

In this section, we present a decision procedure for BTNL[R]$^-$, that progressively removes from $\mathcal{T}_\varphi$ nodes that cannot contribute to fulfilling substructures.

Algorithm 1 *Let φ be the formula we want to test for satisfiability. The* *decision procedure works as follows.*

(i) *Build the (unique) initial tableau* $\mathcal{T}_\varphi = \langle \mathcal{N}_\varphi, \mathcal{R}_\varphi \rangle$.

(ii) *Look for a fulfilling substructure by repeatedly applying the following deletion rules, until no more nodes in the tableau can be deleted:*
  - *delete any node which is not* $\mathcal{R}_\varphi$*-reachable from an initial node;*
  - *delete any node such that there are no fulfilling paths starting from it;*
  - *delete any node which does not satisfy the conditions of Definition 3.10.*

(iii) *Let* $\mathcal{T}^* = \langle \mathcal{N}^*, \mathcal{R}^* \rangle$ *be the final tableau. If* $\mathcal{T}^*$ *is not empty, return* true*, otherwise return* false.

The check for the existence of fulfilling paths can be performed as follows. Given a formula $A\langle A\rangle\psi \in \mathrm{CL}(\varphi)$, we execute the following marking procedure. First, for all nodes N, mark all atoms $(\mathcal{A}, \mathcal{C}) \in N$ such that $A\langle A\rangle\psi \in \mathcal{A}$ and $\psi \in \mathcal{C}$. Then, for all nodes N, mark all unmarked atoms $(\mathcal{A}, \mathcal{C}) \in N$ such that there exists an $\mathcal{R}_\varphi$-successor M of N that contains a marked atom $(\mathcal{A}', \mathcal{C}')$ such that $(\mathcal{A}, \mathcal{C}) X_\varphi (\mathcal{A}', \mathcal{C}')$. Repeat this last step until no more atoms can be marked. Then, delete all nodes that either contain an unmarked atom $(\mathcal{A}, \mathcal{C})$ with $A\langle A\rangle\psi \in \mathcal{A}$ or have no $\mathcal{R}_\varphi$-successors.

The other non trivial step of the algorithm is the removal of nodes that do not satisfy the conditions of Definition 3.10. Given a node N, condition F1 can be easily checked by visiting the $\mathcal{R}_\varphi$-successors of N, while condition F2 can be checked by visiting the $\mathcal{R}_\varphi$-descendants of N. Finally, given a node N_0, an atom $(\mathcal{A}_0, \mathcal{C}_0) \in N_0$, and a formula $E[A]\psi \in \mathcal{A}_0$, condition F3 is satisfied if we can find a *finite* path of nodes $N_0 N_1 \ldots N_j N_{j+1} \ldots N_k$ and a corresponding path of atoms $(\mathcal{A}_0, \mathcal{C}_0)(\mathcal{A}_1, \mathcal{C}_1) \ldots (\mathcal{A}_j, \mathcal{C}_j)(\mathcal{A}_{j+1}, \mathcal{C}_{j+1}) \ldots (\mathcal{A}_k, \mathcal{C}_k)$ such that *(i)* $N_j = N_k$, *(ii)* $(\mathcal{A}_j, \mathcal{C}_j) = (\mathcal{A}_k, \mathcal{C}_k)$, *(iii)* for all $0 \leq i \leq k$, $E[A]\psi \in \mathcal{A}_i$, *(iv)* for every formula $A\langle A\rangle\theta \in \mathcal{A}_0$, there exists $i \geq 0$ such that $\theta \in \mathcal{C}_i$. Furthermore, to guarantee that the infinite path

$$N_0 N_1 \ldots N_j N_{j+1} \ldots N_{k-1} N_j N_{j+1} \ldots N_{k-1} \ldots$$

is a fulfilling one, it suffices to check that, for every atom $(\mathcal{A}', \mathcal{C}') \in N_j$ and every formula $A\langle A\rangle\xi \in \mathcal{C}'$, either $\xi \in \mathcal{C}'$ or there exists a node N_l, with $j < l < k$, and an atom $(\mathcal{A}'', \mathcal{C}'') \in N_l$ such that $\xi \in \mathcal{C}''$ and $(\mathcal{A}'', \mathcal{C}'')$ is a X_φ-descendant of $(\mathcal{A}', \mathcal{C}')$.

As for complexity issues, we have that:

- $|\mathcal{T}_\varphi| = 2^{2^{O(|\varphi|)}}$;
- all checkings of step (ii) of the algorithm can be done in time polynomial in the size of $|\mathcal{T}_\varphi|$;
- after deleting at most $|\mathcal{N}_\varphi|$ nodes, the algorithm terminates.

Hence, checking the satisfiability for a BTNL[R]$^-$ formula has an overall time bound of $2^{2^{O(|\varphi|)}}$, that is, doubly exponential in the size of φ.

4 The decision procedure at work

In this section we apply the proposed decision procedure to the (satisfiable) formula $\varphi = E[A]p$. We show only a portion of the whole tableau, which is sufficiently large to include a fulfilling substructure for φ, and thus to prove that φ is satisfiable.

When searching for a fulfilling substructure for φ, we must take into consideration atoms which have been obtained by suitably combining one set of active requests with one set of current formulas. The sets of active requests and current formulas are the following ones (the left column reports the sets of active requests, while the middle and the right columns report the sets of current formulas):

$\emptyset$;	$\mathcal{C}_0 = \{E[A]p, p\}$;	$\mathcal{C}_3 = \{A\langle A\rangle\neg p, \neg p\}$.
$\mathcal{A}_0 = \{E[A]p\}$;	$\mathcal{C}_1 = \{E[A]p, \neg p\}$;	
$\mathcal{A}_1 = \{A\langle A\rangle\neg p\}$;	$\mathcal{C}_2 = \{A\langle A\rangle\neg p, p\}$;	

As an example, consider the initial node $N_0 = \{(\emptyset, \mathcal{C}_0)\}$. Figure 2 depicts a portion of $\mathcal{T}_\varphi$ that is $\mathcal{R}_\varphi$-reachable from N_0.

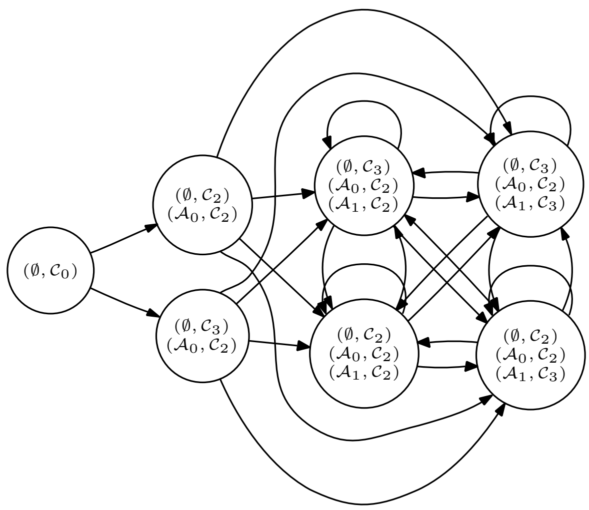

Figure 2. A portion of the tableau for $E[A]p$.

The only atoms with $A\langle A\rangle$-formulas in their set of active requests are $(\mathcal{A}_1, \mathcal{C}_2)$ and $(\mathcal{A}_1, \mathcal{C}_3)$, since $A\langle A\rangle\neg p \in \mathcal{A}_1$. The atom $(\mathcal{A}_1, \mathcal{C}_3)$ immediately fulfills $A\langle A\rangle\neg p$, since $\neg p \in \mathcal{C}_3$. The atom $(\mathcal{A}_1, \mathcal{C}_2)$ does not fulfill $A\langle A\rangle\neg p$, but we have that $(\mathcal{A}_1, \mathcal{C}_2) X_\varphi (\mathcal{A}_1, \mathcal{C}_3)$. Thus, since every node that contains $(\mathcal{A}_1, \mathcal{C}_2)$ can reach a node containing $(\mathcal{A}_1, \mathcal{C}_3)$, for every node of the substructure of Figure 2 there exists a fulfilling path.

As for condition F1, the only sets of current formulas which contains an existential formula are $\mathcal{C}_1$ and $\mathcal{C}_0$. $\mathcal{C}_1$ does not belong to any node in Figure 2, and thus we can ignore it. $\mathcal{C}_0$ only belongs to the initial node, whose two successors include the atom $(\mathcal{A}_0, \mathcal{C}_2)$ with $E[A]p \in \mathcal{A}_0$. Hence, condition F1 is satisfied. Since there are no atoms containing formulas of the form $E\langle A\rangle\psi$, condition F2 is trivially satisfied. As for condition F3, consider the atom $(\mathcal{A}_0, \mathcal{C}_2)$. It is easy to see that, for every node N containing $(\mathcal{A}_0, \mathcal{C}_2)$, there exists a fulfilling infinite path starting from N such that every node contains $(\mathcal{A}_0, \mathcal{C}_2)$. Thus, condition F3 is satisfied. This allows us to conclude that the substructure depicted in Figure 2 is fulfilling, and thus our decision procedure correctly concludes that the formula $E[A]p$ is satisfiable.

5 Conclusions and further work

In this paper, we proposed a new propositional interval temporal logic, interpreted over infinite trees, which combines the interval neighborhood operators $\langle A \rangle$ and $[A]$ of RPNL$^-$ with the path quantifiers A and E of branching time temporal logics, and we provided it with a doubly-exponential tableau-based decision procedure. We do not know yet if the satisfiability problem for BTNL[R]$^-$ is doubly EXPTIME-complete or not (we conjecture it is not). As for possible extensions of the work, we are studying the decision problem for combinations of path quantifiers operators with other sets of interval logic operators, e.g., those of PNL [6].

Acknowledgements

We would like to thank the anonymous referees for their helpful comments and suggestions. This work has been partially funded by the Italy/South Africa joint project "Temporal logics in computer and information sciences" and by the INTAS project on "Algebraic and deduction methods in non-classical logics and their applications to computer science".

References

[1] D. Bresolin and A. Montanari, *A tableau-based decision procedure for right propositional neighborhood logic*, in: *Proc. of the International Conference TABLEAUX 2005*, vol. 3702 of *LNAI* (2005), 63–77.

[2] Z. Chaochen and M. R. Hansen, *An adequate first order interval logic*, in: W. de Roever, H. Langmark and A. Pnueli, editors, *Compositionality: the Significant Difference*, vol. 1536 of *LNCS* (1998), 584–608.

[3] E. Emerson and J. Halpern, *Decision procedures and expressiveness in the temporal logic of branching time*, Journal of Computer and System Sciences **30** (1985), 1–24.

[4] E. A. Emerson, *Temporal and modal logic*, in: J. van Leeuwen, editor, *Handbook of Theoretical Computer Science*, vol. B, MIT Press, 1990 995–1072.

[5] V. Goranko, A. Montanari, and G. Sciavicco, *A general tableau method for propositional interval temporal logic*, in: *Proceedings of the International Conference TABLEAUX 2003*, vol. 2796 of *LNAI* (2003), 102–116.

[6] V. Goranko, A. Montanari, and G. Sciavicco, *Propositional interval neighborhood temporal logics*, Journal of Universal Computer Science **9** (2003), 1137–1167.

[7] V. Goranko, A. Montanari, and G. Sciavicco, *A road map of interval temporal logics and duration calculi*, Journal of Applied Non-Classical Logics **14** (2004), 9–54.

[8] V. Goranko, A. Montanari, G. Sciavicco, and P. Sala, *A general tableau method for propositional interval temporal logics: theory and implementation*, Journal of Applied Logic (2005), to appear.

[9] J. Halpern and Y. Shoham, *A propositional modal logic of time intervals*, Journal of the ACM **38** (1991), 935–962.

[10] A. Montanari, *Propositional interval temporal logics: some promising paths*, in: *Proc. of the 12th International Symposium on Temporal Representation and Reasoning (TIME)* (2005), 201–203.

[11] B. Moszkowski, "Reasoning about digital circuits," Tech. rep. stan-cs-83-970, Dept. of Computer Science, Stanford University, Stanford, CA (1983).

[12] B. Paech, *Gentzen-systems for propositional temporal logics*, in: *Proceedings of CSL '88, 2nd Workshop on Computer Science Logic*, vol. 385 of *LNCS* (1989), 240–253.

[13] Y. Venema, *A modal logic for chopping intervals*, Journal of Logic and Computation **1** (1991), 453–476.