

An On-the-fly Tableau-based Decision Procedure for *PDL*-Satisfiability

Pietro Abate^{a,1}, Rajeev Goré^{a,2} and Florian Widmann^{b,3,4}

^a *Computer Sciences Laboratory
The Australian National University
Canberra, Australia*

^b *Computer Sciences Laboratory and Logic & Computation Programme
The Australian National University and NICTA
Canberra, Australia*

Abstract

We present a tableau-based algorithm for deciding satisfiability for propositional dynamic logic (*PDL*) which builds a finite rooted tree with ancestor loops and passes extra information from children to parents to separate good loops from bad loops during backtracking. It is easy to implement, with potential for parallelisation, because it constructs a pseudo-model “on the fly” by exploring each tableau branch independently. But its worst-case behaviour is 2EXPTIME rather than EXPTIME. A prototype implementation in the TWB (<http://twb.rsise.anu.edu.au>) is available.

Keywords: propositional dynamic logic, automated reasoning, tableau calculus, decision procedure

1 Introduction

Propositional dynamic logic (*PDL*) is a logic for reasoning about programs [14,8]. Its formulae consist of traditional Boolean formulae plus “action modalities” built from a finite set of atomic programs using sequential composition (;), non-deterministic choice (\cup), repetition (*), and test (?). The satisfiability problem for *PDL* is EXPTIME-complete [15]. Unlike EXPTIME-complete description logics with algorithms exhibiting good average-case behaviour, no decision procedures for *PDL*-satisfiability are satisfactory from both a theoretical (soundness and completeness) and practical (average case behaviour) viewpoint as we explain below.

The earliest decision procedures for *PDL* are due to Fischer and Ladner [8] and Pratt [15]. Fischer and Ladner’s method is impractical because it first constructs

¹ Email: Pietro.Abate@pps.jussieu.fr

² Email: Rajeev.Gore@rsise.anu.edu.au

³ Email: Florian.Widmann@rsise.anu.edu.au

⁴ National ICT Australia is funded by the Australian Government’s Dept of Communications, Information Technology and the Arts and the Australian Research Council through Backing Australia’s Ability and the ICT Centre of Excellence program.

the set of all consistent subsets of the set of all subformulae of the given formula, which always requires exponential time in all cases. On the other hand, Pratt [15] essentially builds a multi-pass (explained shortly) tableau method. Most subsequent decision procedures for other fix-point logics like propositional linear temporal logic (PLTL) [18], computation tree logic (CTL) [4,7] and the modal μ -calculus [13] trace back to Pratt [15], and they all share one main disadvantage as explained next.

In these multi-pass procedures, a “state” is a node which contains only diamond-like-formulae (“eventualities”), box-like-formulae, atoms and negated atoms. The first pass constructs a rooted tableau of nodes containing formula-sets, but allows cross-branch arcs from a state n on one branch to a (previously constructed) state m on a different branch if applying the tableau construction to n would duplicate m . Thus the first pass constructs a “pseudo-model” which is a potentially exponential-sized cyclic graph (rather than a cyclic tree where m would have to be an ancestor of n). The subsequent passes check that the “pseudo-model” is a real model by pruning inconsistent nodes and pruning nodes containing “unfulfilled eventualities”.

Although efficient model-checking techniques can check the “pseudo-model” in time which is linear in its size, these multi-pass methods can construct an exponential-sized cyclic graph needlessly. One solution is to check for fulfilled eventualities “on the fly”, as the graph is built, and although such methods exist for model-checking [6,5], we know of no such decision procedures for *PDL*. The only implementation of a multiple-pass method for *PDL* that we know of is in LoTRec (www.irit.fr/Lotrec) but it is not optimal as it treats disjunctions naively.

Baader [3] gave a single-pass tableau-based decision procedure for a description logic with role definitions involving union, composition and transitive closure of roles: essentially *PDL* without test. His method constructs a (cyclic tree) tableau using the semantics of the *PDL* operators. To separate “good loops” from “bad loops”, Baader must decide equality of regular languages, a PSPACE-complete problem which in practice may require exponential time. Instead of solving these problems “on the fly”, they can be reduced to a simple check on the identity of states in a deterministic minimal automaton created from the positive regular expressions appearing in the initial formula during a pre-processing stage [3, page 27]. But since the pre-computed automaton can be of exponential size, this alternative may require exponential time needlessly. Baader’s method is double-exponential in the worst-case. The “test” construct is essential to express “while” loops but creates a mutual recursion between the Boolean language and the regular language. It is not obvious to us how to extend Baader’s method to “test”. DLP (<http://www.cs.bell-labs.com/cm/cs/who/pfps/dlp>) implements this method restricted to test-free formulae where $*$ applies only to atomic programs.

De Giacomo and Massacci [9] gave an optimal *PDL*-satisfiability test using labelled formulae like $\sigma : \varphi$ to capture that “possible world σ makes formula φ true”. They first give a NEXPTIME algorithm for deciding *PDL*-satisfiability and then discuss ways to obtain an EXPTIME version using various known results. But an actual EXPTIME algorithm, and its soundness and completeness proofs, are not given. A deterministic implementation of their NEXPTIME algorithm by Schmidt and Tishkovsky struck problems with nested stars, but a solution is forthcoming [16].

Other decision procedures for fix-point logics use resolution calculi, translation

methods, automata-theoretic methods, and game theoretic methods: see [1] for references. We know of no implementations for *PDL* based on these methods.

Here, we give a sound, complete and terminating decision procedure for *PDL* with the following advantages and disadvantages:

One-pass nature: our method constructs a single-rooted finite tree (with loops from leaves to ancestors). As there are no cross-branch edges, we can use depth-first, left-to-right search, reclaiming the space used for each branch via backtracking.

Proofs: Full elementary proofs of soundness and completeness are available.

Ease of implementation: our rules are easy to implement since our tableau nodes contain sets of formulae and some easily defined extra information whose manipulation requires only set intersection, set membership, and min/max on integers. But these low-level details make the rules cumbersome to describe.

Potential for optimisation: there is potential to optimise our (tree) tableaux using successful techniques from (one-pass) tableaux for description logics [11].

Ease of generating counter-models: the soundness proof immediately gives an effective procedure for turning an “open” tableau into a *PDL*-model.

Ease of generating proofs: unlike existing Gentzen calculi for fix-point logics [2,12], our tableau calculus gives a cut-free Gentzen-style calculus with “cyclic proofs” with an optimal rather than worst-case bound for the finitised omega rule.

Potential for parallelisation: our rules build the branches independently but combine their results during backtracking, enabling a parallel implementation.

Prototype: a (sequential) prototype implementation in the Tableau Work Bench (twb.rsise.anu.edu.au) allows to test arbitrary *PDL* formulae over the web.

Complexity: our method has worst-case double-exponential time complexity.

Generality: Our method for *PDL* fits into a class of similar “one pass” methods for other fix-point logics like PLTL [17] and CTL [1]. Further experimental work is required to determine if our methods can be optimised to exhibit good average-case behaviour using techniques like sound global caching [10].

2 Syntax, Semantics and Hintikka Structures

Definition 2.1 Let AFml and APrg be two disjoint and countably infinite sets of propositional atoms and *atomic programs*, respectively. The set Fml of all formulae and the set Prg of all *programs* are defined inductively as follows:

- (i) $\text{AFml} \subseteq \text{Fml}$ and $\text{APrg} \subseteq \text{Prg}$
- (ii) if $\varphi, \psi \in \text{Fml}$ then $\neg\varphi \in \text{Fml}$ and $\varphi \wedge \psi \in \text{Fml}$ and $\varphi \vee \psi \in \text{Fml}$ and $\varphi? \in \text{Prg}$
- (iii) if $\varphi \in \text{Fml}$ and $\alpha \in \text{Prg}$ then $\langle\alpha\rangle\varphi \in \text{Fml}$ and $[\alpha]\varphi \in \text{Fml}$
- (iv) if $\alpha \in \text{Prg}$ and $\beta \in \text{Prg}$ then $(\alpha;\beta) \in \text{Prg}$ and $\alpha \cup \beta \in \text{Prg}$ and $\alpha^* \in \text{Prg}$.

Let p, q range over members of AFml and a, b range over members of APrg . A $\langle\rangle$ -formula is any formula $\langle\alpha\rangle\varphi$, a $\langle q \rangle$ -formula is a $\langle\rangle$ -formula $\langle\alpha\rangle\varphi$ with $\alpha \notin \text{APrg}$, and a $\langle*\rangle$ -formula is any formula $\langle\alpha^*\rangle\varphi$. $\text{Fml}\langle\rangle$ is the set of all $\langle\rangle$ -formulae, $\text{Fml}\langle q \rangle$ is the set of all $\langle q \rangle$ -formulae, and $\text{Fml}\langle*\rangle$ is the set of all $\langle*\rangle$ -formulae.

Table 1
 Smullyan's α - and β -notation to classify formulae

α	$\varphi \wedge \psi$	$[\alpha \cup \beta]\varphi$	$[\alpha*]\varphi$	$\langle \psi? \rangle \varphi$	$\langle \alpha; \beta \rangle \varphi$	$[\alpha; \beta]\varphi$
α_1	φ	$[\alpha]\varphi$	φ	φ	$\langle \alpha \rangle \langle \beta \rangle \varphi$	$[\alpha][\beta]\varphi$
α_2	ψ	$[\beta]\varphi$	$[\alpha][\alpha*]\varphi$	ψ		
	β	$\varphi \vee \psi$	$\langle \alpha \cup \beta \rangle \varphi$	$\langle \alpha* \rangle \varphi$	$[\psi?]\varphi$	
	β_1	φ	$\langle \alpha \rangle \varphi$	φ	φ	
	β_2	ψ	$\langle \beta \rangle \varphi$	$\langle \alpha \rangle \langle \alpha* \rangle \varphi$	$\sim \psi$	

Definition 2.2 A *transition frame* is a pair (W, R) where W is a non-empty set of worlds and R a function that maps each atomic program a to a binary relation R_a over W . A *model* (W, R, V) is a transition frame (W, R) and a valuation function $V : \text{AFml} \rightarrow 2^W$ mapping each atomic proposition p to a set $V(p)$ of worlds.

Definition 2.3 Let $M = (W, R, V)$ be a model. The functions $\tau_M : \text{Fml} \rightarrow 2^W$ and $\rho_M : \text{Prg} \rightarrow 2^{W \times W}$ are defined inductively as follows:

$$\begin{aligned}
 \tau_M(p) &:= V(p) & \rho_M(a) &:= R_a & \tau_M(\neg\varphi) &:= W \setminus \tau_M(\varphi) \\
 \tau_M(\varphi \wedge \psi) &:= \tau_M(\varphi) \cap \tau_M(\psi) & \tau_M(\varphi \vee \psi) &:= \tau_M(\varphi) \cup \tau_M(\psi) \\
 \tau_M([\alpha]\varphi) &:= \{w \mid \forall v \in W. (w, v) \in \rho_M(\alpha) \Rightarrow v \in \tau_M(\varphi)\} \\
 \tau_M(\langle \alpha \rangle \varphi) &:= \{w \mid \exists v \in W. (w, v) \in \rho_M(\alpha) \ \& \ v \in \tau_M(\varphi)\} \\
 \rho_M(\alpha \cup \beta) &:= \rho_M(\alpha) \cup \rho_M(\beta) & \rho_M(\varphi?) &:= \{(w, w) \mid w \in \tau_M(\varphi)\} \\
 \rho_M(\alpha; \beta) &:= \{(w, v) \mid \exists u \in W. (w, u) \in \rho_M(\alpha) \ \& \ (u, v) \in \rho_M(\beta)\} \\
 \rho_M(\alpha*) &:= \{(w, v) \mid \exists k \in \mathbb{N}. \exists w_0, \dots, w_k \in W. (w_0 = w \ \& \ w_k = v \ \& \\
 & \quad \forall i \in \{0, \dots, k-1\}. (w_i, w_{i+1}) \in \rho_M(\alpha))\}
 \end{aligned}$$

For $w \in W$ and $\varphi \in \text{Fml}$, we write $M, w \Vdash \varphi$ iff $w \in \tau_M(\varphi)$.

Definition 2.4 Formula $\varphi \in \text{Fml}$ is *satisfiable* iff there is a model $M = (W, R, V)$ and a $w \in W$ such that $M, w \Vdash \varphi$. Formula $\varphi \in \text{Fml}$ is *valid* iff $\neg\varphi$ is not satisfiable.

Definition 2.5 Formula $\varphi \in \text{Fml}$ is in *negation normal form* if \neg appears only immediately before propositional atoms. For every $\varphi \in \text{Fml}$, we obtain a formula $\text{nnf}(\varphi)$ in negation normal form by pushing negations inward repeatedly (*e.g.* using de Morgan's laws) so $\varphi \leftrightarrow \text{nnf}(\varphi)$ is valid. We define $\sim\varphi := \text{nnf}(\neg\varphi)$.

We use Smullyan's α/β -notation to categorise formulae via Table 1 and use bolding to differentiate it from the use of α and β as members of Prg . So if α (respectively β) is any formula pattern in the first row then α_1 and α_2 (respectively β_1 and β_2) are its corresponding patterns in the second and third row.

Proposition 2.6 All formulae $\alpha \leftrightarrow \alpha_1 \wedge \alpha_2$ and $\beta \leftrightarrow \beta_1 \vee \beta_2$ in Table 1 are valid.

Definition 2.7 A *structure* (W, R, L) [for $\varphi \in \text{Fml}$] is a transition frame (W, R) and a labelling function $L : W \rightarrow 2^{\text{Fml}}$ which associates with each world $w \in W$ a set $L(w)$ of formulae [and has $\varphi \in L(v)$ for some world $v \in W$].

Definition 2.8 For a given $\varphi \in \text{Fml}$ the (infinite) set $\text{pre}(\varphi)$ is defined as:

$$\text{pre}(\varphi) := \{\psi \in \text{Fml} \mid \exists k \in \mathbb{N}. \exists \alpha_1, \dots, \alpha_k \in \text{Prg}. \psi = \langle \alpha_1 \rangle \dots \langle \alpha_k \rangle \varphi\} .$$

For all formulae φ and ψ , the binary relation \rightsquigarrow on formulae is defined as: $\varphi \rightsquigarrow \psi$ iff (exactly) one of the following conditions is true:

- $\exists \chi \in \text{Fml}. \exists \alpha, \beta \in \text{Prg}. \varphi = \langle \alpha; \beta \rangle \chi$ & $\psi = \langle \alpha \rangle \langle \beta \rangle \chi$
- $\exists \chi \in \text{Fml}. \exists \alpha, \beta \in \text{Prg}. \varphi = \langle \alpha \cup \beta \rangle \chi$ & ($\psi = \langle \alpha \rangle \chi$ or $\psi = \langle \beta \rangle \chi$)
- $\exists \chi \in \text{Fml}. \exists \alpha \in \text{Prg}. \varphi = \langle \alpha^* \rangle \chi$ & ($\psi = \chi$ or $\psi = \langle \alpha \rangle \langle \alpha^* \rangle \chi$)
- $\exists \chi, \phi \in \text{Fml}. \varphi = \langle \phi^? \rangle \chi$ & $\psi = \chi$.

Intuitively, using Table 1, the “ \rightsquigarrow ” relates a $\langle \phi \rangle$ -formulae α (respectively β), to α_1 (respectively β_1 and β_2) while $\text{pre}(\varphi)$ captures that $\langle \alpha^* \rangle \varphi$ can be “reduced” to $\langle \alpha \rangle \langle \alpha^* \rangle \varphi$, which can be reduced to $\langle \alpha_1 \rangle \dots \langle \alpha_k \rangle \langle \alpha^* \rangle \varphi$. Note that $\varphi \in \text{pre}(\varphi)$.

Definition 2.9 Let $H = (W, R, L)$ be a structure, $\varphi \in \text{Fml}$ a formula, $\beta \in \text{Prg}$ a program, and $w \in W$ a state. A *fulfilling chain* for (φ, β, w) in H is a finite sequence $(w_0, \psi_0), \dots, (w_n, \psi_n)$ of world-formula pairs with $n \geq 0$ such that:

- $w_i \in W$, $\psi_i \in \text{pre}(\varphi)$, and $\psi_i \in L(w_i)$ for all $0 \leq i \leq n$
- $w_0 = w$, $\psi_0 = \langle \beta \rangle \varphi$, $\psi_n = \varphi$, and $\psi_i \neq \varphi$ for all $0 \leq i \leq n-1$
- for all $0 \leq i \leq n-1$, if $\psi_i = \langle a \rangle \chi$ for some $a \in \text{APrg}$ and $\chi \in \text{Fml}$ then $\psi_{i+1} = \chi$ and $w_i R_a w_{i+1}$; otherwise $\psi_i \rightsquigarrow \psi_{i+1}$ and $w_i = w_{i+1}$.

Each ψ_i is in $L(w_i)$, the chain starts at $(w_0, \langle \beta \rangle \varphi)$, ends at (w_n, φ) , and no other w_i is paired with φ . Formulae ψ_i, ψ_{i+1} are \rightsquigarrow -related and corresponding worlds w_i, w_{i+1} are equal unless $\psi_i = \langle a \rangle \chi$, in which case $\psi_{i+1} = \chi$ and $w_i R_a w_{i+1}$. Thus eventuality $\langle \beta \rangle \varphi \in w_0$ is fulfilled by $\varphi \in w_n$ and w_n is β -reachable from w_0 .

Definition 2.10 A *pre-Hintikka structure* $H = (W, R, L)$ [for $\varphi \in \text{Fml}$] is a structure [for φ] that satisfies H1-H5 (below) for every $w \in W$ where α and β are formulae as defined in Table 1. A *Hintikka structure* $H = (W, R, L)$ [for $\varphi \in \text{Fml}$] is a pre-Hintikka structure [for φ] that additionally satisfies H6 below:

$$\text{H1} : \neg p \in L(w) \Rightarrow p \notin L(w)$$

$$\text{H2} : \alpha \in L(w) \Rightarrow \alpha_1 \in L(w) \ \& \ \alpha_2 \in L(w)$$

$$\text{H3} : \beta \in L(w) \Rightarrow \beta_1 \in L(w) \ \text{or} \ \beta_2 \in L(w)$$

$$\text{H4} : \langle a \rangle \varphi \in L(w) \Rightarrow \exists v \in W. w R_a v \ \& \ \varphi \in L(v)$$

$$\text{H5} : [a] \varphi \in L(w) \Rightarrow \forall v \in W. w R_a v \Rightarrow \varphi \in L(v)$$

$$\text{H6} : \langle \alpha^* \rangle \varphi \in L(w) \Rightarrow \text{there exists a fulfilling chain for } (\varphi, \alpha^*, w) \text{ in } H .$$

H3 “locally unwinds” the fix-point semantics of $\langle \alpha * \rangle \varphi$, but does not guarantee a *least* fix-point which requires φ be true eventually. H6 “globally” ensures all $\langle * \rangle$ -formulae are fulfilled. H2 captures the *greatest* fix-point semantics of $[\alpha *] \varphi$.

Theorem 2.11 *A formula $\varphi \in \text{Fml}$ in negation normal form is satisfiable iff there exists a Hintikka structure for φ .*

3 An Overview of the Algorithm

To track unfulfilled eventualities and to avoid “at a world” cycles, our algorithm stores additional information in each tableau node using *histories* and *variables* [17]. Histories are passed from parents to children and variables from children to parents.

Our algorithm starts at a root containing a given formula ϕ and some default history values. It builds a tree by repeatedly applying α -/ β -rules to decompose formulae via the semantics of *PDL*. The β -rule for $\langle \alpha * \rangle \varphi$ has a left child that fulfils this eventuality by reducing it to φ , and a right child that procrastinates fulfilment by “reducing” it to $\langle \alpha \rangle \langle \alpha * \rangle \varphi$. The rules modify the histories and variables as appropriate for their intended purpose.

But naive application of the α -/ β -rules to formulae like $\langle a * * \rangle \varphi$ with nested stars can lead to “at a world” cycles: e.g. $\langle a * * \rangle \varphi, \dots, \langle a * \rangle \langle a * * \rangle \varphi, \dots, \langle a * * \rangle \varphi$. A solution is to use the histories to reduce one particular $\langle \alpha \rangle$ -formula until α becomes atomic by forcing the rules to concentrate on this task, and to block previously reduced diamonds and boxes if they lead to “at a world” cycles. The application of α / β -rules stops when all non-blocked leaves contain only atoms, negated atoms, and all $\langle \rangle$ -formulae and all $[]$ -formulae begin with outermost atomic programs only.

For each such leaf node l , and for each $\langle a \rangle \xi$ -formula in l , the $\langle \rangle$ -rule creates a successor node containing $\{\xi\} \cup \Delta$, where $\Delta = \{\psi \mid [a]\psi \in l\}$. These successors are then saturated to produce new leaves using the α - and β -rules, and the $\langle \rangle$ -rule creates the successors of these new leaves, and so on.

If left unchecked, this procedure can produce infinite branches since the same successors can be created again and again on the same branch. To obtain termination, the $\langle \rangle$ -rule creates a successor containing $\{\xi\} \cup \Delta$ for l only if this successor has not already been created previously higher up on the current branch.

So if the successor $\{\xi\} \cup \Delta$ exists already, the current branch is “blocked” from re-creating it. The resulting loop may be “bad” since every β -node on this branch for an eventuality $\langle \alpha * \rangle \varphi$ may procrastinate, so $\langle \alpha * \rangle \varphi$ is never fulfilled. To track this potentially unfulfilled eventuality, we assign the height of the blocking node to the pair $(\xi, \langle \alpha * \rangle \varphi)$ via a variable *uev* as long as ξ is a decomposition of $\langle \alpha * \rangle \varphi$.

During backtracking, our rules “merge” the *uev* entries of the children and also modify the resulting *uev* to reverse-track the decomposition of $\langle \alpha * \rangle \varphi$. In particular, a *uev* entry becomes undefined at a node if the eventuality it tracks can be fulfilled in the sub-tableau rooted at this node. Conversely, if a node at height h receives a *uev* entry with value at least h then the eventuality tracked by this *uev* entry definitely cannot be fulfilled, so the parent of this (blocking) node is then unsatisfiable.

Whether or not the initial formula ϕ is satisfiable is determined by the status of the root node. Due to technicalities caused by “at a world” cycles, the status

can be one of the values “unsatisfiable”, “open” or “barred” (to be explained later). The initial formula ϕ is *PDL*-satisfiable iff the status of the root node is “open”.

4 A One-pass Tableau Algorithm for *PDL*

Definition 4.1 A *tableau node* x is of the form $(\Gamma :: \text{HCr}, \text{Nx}, \text{BD}, \text{BB} :: \text{stat}, \text{uev})$ where: Γ is a set of formulae; HCr is a list of pairs (φ, Δ) where Δ is a set of formulae and $\varphi \in \Delta$; Nx is either \perp or a formula designated to be the principal formula of the rule applied to x ; BD is the set of “Blocked Diamonds”; BB is the set of “Blocked Boxes”; stat has one of the values **unsat**, **open**, or **barred**; and uev is a partial function from $\text{Fml}\langle \rangle \times \text{Fml}\langle * \rangle$ to $\mathbb{N}_{>0}$ (the positive natural numbers).

Definition 4.2 A *tableau* for a formula set $\Gamma \subseteq \text{Fml}$ and histories HCr , Nx , BD , and BB is a tree of tableau nodes with root $(\Gamma :: \text{HCr}, \text{Nx}, \text{BD}, \text{BB} :: \text{stat}, \text{uev})$ where the children of a node x are obtained by a single application of a rule to x (*i.e.* only one rule can be applied to a node) but where the parent can inherit some information from the children. A tableau is *expanded* if no rules can be applied to any of its leaves. On any branch of a tableau, a node t is an *ancestor* of a node s iff t lies above s on the unique path from the root down to s .

The list HCr is a history for detecting ancestor-loops and guarantees termination. The choice of principal formula is free if $\text{Nx} = \perp$, but is pre-determined as the formula in Nx otherwise. When a diamond formula in the parent is decomposed to give a formula $\varphi \in \text{Fml}\langle \varphi \rangle$ in the current node, we set the Nx -value of the child to φ to ensure that φ is decomposed next. Together with the histories BD and BB , this allows us to block $\langle \alpha * \rangle$ -formulae and $[\alpha *]$ -formulae from creating “at a world” cycles. The variables stat and uev have their values determined by the children of a node. Formally, $\text{stat} = \mathbf{unsat}$ at node x if x is definitely unsatisfiable. Informally, $\text{stat} = \mathbf{barred}$ if all descendants of node x are unsatisfiable or lead to an “at a world” cycle. Finally, $\text{stat} = \mathbf{open}$ indicates that the node is potentially satisfiable, but as it may be on a loop, this is something which we can determine only later as we backtrack towards the root.

Definition 4.3 The partial function $\text{uev}_\perp : \text{Fml}\langle \rangle \times \text{Fml}\langle * \rangle \rightarrow \mathbb{N}_{>0}$ is the constant function that is undefined for all pairs of formulae: *i.e.* $\forall \psi_1, \psi_2. \text{uev}_\perp(\psi_1, \psi_2) = \perp$. The partial functions $\text{tst} : \text{Fml} \rightarrow \text{Fml}$ and $\text{bl} : \text{Fml} \times 2^{\text{Fml}} \rightarrow 2^{\text{Fml}}$ are defined as:

$$\text{tst}(\chi) := \begin{cases} \chi & \text{if } \chi \in \text{Fml}\langle \varphi \rangle \\ \perp & \text{otherwise} \end{cases} \quad \text{bl}(\chi, \Gamma) := \begin{cases} \Gamma & \text{if } \chi \in \text{Fml}\langle \varphi \rangle \\ \emptyset & \text{otherwise.} \end{cases}$$

The function tst returns \perp when the formula being tested is not a $\langle \rangle$ -formula, or is a $\langle \rangle$ -formula but its program is atomic. The function uev tracks unfulfilled eventualities, so uev_\perp flags that all eventualities are fulfilled, and $\text{uev}(\chi_1, \chi_2)$ defined flags a potentially unfulfilled eventuality. If a node has $\text{stat} = \mathbf{unsat}$ or $\text{stat} = \mathbf{barred}$ then its uev is irrelevant so it is arbitrarily set to uev_\perp .

4.1 The Rules

We use Γ and Δ for sets of formulae and write $\varphi_1, \dots, \varphi_n, \Delta_1, \dots, \Delta_m$ for the partition $\{\varphi_1\} \uplus \dots \uplus \{\varphi_n\} \uplus \Delta_1 \uplus \dots \uplus \Delta_m$ of formulae in a node. To save space, we often omit histories/variables which are passed unchanged from parents/children to children/parents. Most rules are applicable only if some side-conditions hold, and most involve actions that change histories downwards or variables upwards.

Terminal Rules.

$$(id) \frac{(\Gamma :: \dots :: \text{stat}, \text{uev})}{\{p, \neg p\} \subseteq \Gamma \text{ for some } p \in \text{AFml}}$$

Action for (id) : $\text{stat} := \mathbf{unsat}$ and $\text{uev} := \text{uev}_\perp$.

$$\langle (*)_2 \rangle \frac{(\langle \alpha * \rangle \varphi, \Gamma :: \text{Nx}, \text{BD} :: \text{stat}, \text{uev})}{\text{Nx} \in \{\perp, \langle \alpha * \rangle \varphi\} \ \& \ \langle \alpha * \rangle \varphi \in \text{BD}}$$

Action for $\langle (*)_2 \rangle$: $\text{stat} := \mathbf{barred}$ and $\text{uev} := \text{uev}_\perp$.

An id -node is clearly unsatisfiable. The principal formula of the $\langle (*)_2$ -rule is unfulfillable because it causes an “at a world” cycle, so this rule terminates the current branch. Note both rules may be applicable to a node.

Linear (α) Rules.

$$\begin{array}{ll} (\wedge) \frac{(\varphi \wedge \psi, \Gamma :: \text{Nx} :: \text{uev})}{(\varphi, \psi, \Gamma :: \text{Nx} :: \text{uev}_1)} & ([\cup]) \frac{([\alpha \cup \beta] \varphi, \Gamma :: \text{Nx} :: \text{uev})}{([\alpha] \varphi, [\beta] \varphi, \Gamma :: \text{Nx} :: \text{uev}_1)} \\ ([;]) \frac{([\alpha; \beta] \varphi, \Gamma :: \text{Nx} :: \text{uev})}{([\alpha][\beta] \varphi, \Gamma :: \text{Nx} :: \text{uev}_1)} & ([*]) \frac{([\alpha *] \varphi, \Gamma :: \text{Nx}, \text{BB} :: \text{uev})}{(\Gamma_1 :: \text{Nx}, \text{BB}_1 :: \text{uev}_1)} \end{array}$$

Common Side Condition: $\text{Nx} = \perp$.

Common Action: $\text{uev}(\chi_1, \chi_2) := \text{uev}_1(\chi_1, \chi_2)$ if $\chi_1 \in \Gamma$ else $\text{uev}(\chi_1, \chi_2) := \perp$.

Extra Action for $([*])$: $\Gamma_1 := \Gamma$ if $[\alpha *] \varphi \in \text{BB}$ else $\Gamma_1 := \{\varphi\} \cup \{[\alpha][\alpha *] \varphi\} \cup \Gamma$,

$$\text{BB}_1 := \{[\alpha *] \varphi\} \cup \text{BB}.$$

Most rules are standard but for the histories since they just capture the transformations in Table 1. The $[*]$ -rule just deletes $[\alpha *] \varphi$ if $[\alpha *] \varphi \in \text{BB}$ since this indicates that it has already been expanded once “at this world”. Otherwise it captures the fix-point nature of $[\alpha *] \varphi$ via Prop. 2.6 and then puts $[\alpha *] \varphi$ into BB_1 .

The next two rules have individual side-conditions and actions as shown.

$$\langle \langle ; \rangle \rangle \frac{(\langle \alpha; \beta \rangle \varphi, \Gamma :: \text{Nx}, \text{BD} :: \text{uev})}{(\langle \alpha \rangle \langle \beta \rangle \varphi, \Gamma :: \text{Nx}_1, \text{BD}_1 :: \text{uev}_1)} \quad \text{Nx} \in \{\perp, \langle \alpha; \beta \rangle \varphi\}$$

Actions for $\langle \langle ; \rangle \rangle$:

$$\begin{array}{l} \text{Nx}_1 := \text{tst}(\langle \alpha \rangle \langle \beta \rangle \varphi) \\ \text{BD}_1 := \text{bl}(\langle \alpha \rangle \langle \beta \rangle \varphi, \text{BD}) \end{array} \quad \text{uev}(\chi_1, \chi_2) := \begin{cases} \text{uev}_1(\langle \alpha \rangle \langle \beta \rangle \varphi, \chi_2) & \text{if } \chi_1 = \langle \alpha; \beta \rangle \varphi \\ \text{uev}_1(\chi_1, \chi_2) & \text{if } \chi_1 \in \Gamma \\ \perp & \text{otherwise} \end{cases}$$

$$\langle\langle ? \rangle\rangle \frac{\langle\langle \psi? \rangle\rangle\varphi, \Gamma :: N_x, BD_1 :: \text{uev}}{\langle\psi, \varphi, \Gamma :: N_{x_1}, BD_1 :: \text{uev}_1\rangle} \quad N_x \in \{\perp, \langle\psi? \rangle\varphi\}$$

Actions for $\langle\langle ? \rangle\rangle$:

$$\begin{aligned} N_{x_1} &:= \text{tst}(\varphi) \\ BD_1 &:= \text{bl}(\varphi, BD) \end{aligned} \quad \text{uev}(\chi_1, \chi_2) := \begin{cases} \text{uev}_1(\varphi, \chi_2) & \text{if } \chi_1 = \langle\psi? \rangle\varphi \\ \text{uev}_1(\chi_1, \chi_2) & \text{if } \chi_1 \in \Gamma \\ \perp & \text{otherwise} \end{cases}$$

These rules just capture the transformations in Table 1 except for the histories. Their choice of principal formula is free if $N_x = \perp$, but is restricted to the formula in N_x otherwise. If the decomposition χ of the principal $\langle\langle \cdot \rangle\rangle$ -formula is a $\langle\langle \varphi \rangle\rangle$ -formula, we put N_{x_1} of the child to be χ to enforce that χ is the principal formula of the child. The actions for uev ensure that $\text{uev}(\chi_1, \chi_2)$, where χ_1 is the principal $\langle\langle \cdot \rangle\rangle$ -formula, inherits its value from the corresponding $\langle\langle \cdot \rangle\rangle$ -formulae in the child: *e.g.* $\text{uev}(\langle\langle \alpha; \beta \rangle\rangle\varphi, \chi_2) = \text{uev}_1(\langle\langle \alpha \rangle\rangle\langle\langle \beta \rangle\rangle\varphi, \chi_2)$ reverse-tracks the decomposition of $\langle\langle \alpha; \beta \rangle\rangle\varphi$ into $\langle\langle \alpha \rangle\rangle\langle\langle \beta \rangle\rangle\varphi$. Also, $\text{uev}(\chi_1, \chi_2)$ is only defined if χ_1 is in the parent.

Universal Branching (β) Rules.

$$\langle\langle \vee \rangle\rangle \frac{(\varphi_1 \vee \varphi_2, \Gamma :: N_x :: \text{stat}, \text{uev})}{(\varphi_1, \Gamma :: N_x :: \text{stat}_1, \text{uev}_1) \mid (\varphi_2, \Gamma :: N_x :: \text{stat}_2, \text{uev}_2)} \quad N_x = \perp$$

$$\langle\langle [?] \rangle\rangle \frac{([\psi?] \varphi, \Gamma :: N_x :: \text{stat}, \text{uev})}{(\sim\psi, \Gamma :: N_x :: \text{stat}_1, \text{uev}_1) \mid (\varphi, \Gamma :: N_x :: \text{stat}_2, \text{uev}_2)} \quad N_x = \perp$$

$$\text{Action for } \langle\langle \vee \rangle\rangle \text{ and } \langle\langle [?] \rangle\rangle \text{ for } i = 1, 2: \text{uev}'_i(\chi_1, \chi_2) := \begin{cases} \text{uev}_i(\chi_1, \chi_2) & \text{if } \chi_1 \in \Gamma \\ \perp & \text{otherwise} \end{cases}$$

$$\langle\langle \cup \rangle\rangle \frac{\langle\langle \alpha_1 \cup \alpha_2 \rangle\rangle\varphi, \Gamma :: N_x, BD :: \text{stat}, \text{uev}}{\langle\langle \alpha_1 \rangle\rangle\varphi, \Gamma :: N_{x_1}, BD_1 :: \text{stat}_1, \text{uev}_1 \mid \langle\langle \alpha_2 \rangle\rangle\varphi, \Gamma :: N_{x_2}, BD_2 :: \text{stat}_2, \text{uev}_2}$$

Side-condition for $\langle\langle \cup \rangle\rangle$: $N_x \in \{\perp, \langle\alpha_1 \cup \alpha_2 \rangle\varphi\}$

Action for $\langle\langle \cup \rangle\rangle$ for $i = 1, 2$:

$$\begin{aligned} N_{x_i} &:= \text{tst}(\langle\langle \alpha_i \rangle\rangle\varphi) \\ BD_i &:= \text{bl}(\langle\langle \alpha_i \rangle\rangle\varphi, BD) \end{aligned} \quad \text{uev}'_i(\chi_1, \chi_2) := \begin{cases} \text{uev}_i(\langle\langle \alpha_i \rangle\rangle\varphi, \chi_2) & \text{if } \chi_1 = \langle\alpha_1 \cup \alpha_2 \rangle\varphi \\ \text{uev}_i(\chi_1, \chi_2) & \text{if } \chi_1 \in \Gamma \\ \perp & \text{otherwise} \end{cases}$$

$$\langle\langle * \rangle\rangle_1 \frac{\langle\langle \alpha * \rangle\rangle\varphi, \Gamma :: N_x, BD :: \text{stat}, \text{uev}}{(\varphi, \Gamma :: N_{x_1}, BD_1 :: \text{stat}_1, \text{uev}_1) \mid \langle\langle \alpha \rangle\rangle\langle\langle \alpha * \rangle\rangle\varphi, \Gamma :: N_{x_2}, BD_2 :: \text{stat}_2, \text{uev}_2}$$

Side-condition for $\langle\langle * \rangle\rangle_1$: $N_x \in \{\perp, \langle\alpha * \rangle\varphi\}$ & $\langle\alpha * \rangle\varphi \notin BD$

Action for $\langle\langle * \rangle\rangle_1$:

$$\begin{aligned}
 \text{Nx}_1 &:= \text{tst}(\varphi) \\
 \text{BD}_1 &:= \text{bl}(\varphi, \{\langle \alpha^* \rangle \varphi\} \cup \text{BD}) \\
 \text{uev}'_1(\chi_1, \chi_2) &:= \begin{cases} \perp & \text{if } \chi_1 = \chi_2 = \langle \alpha^* \rangle \varphi \\ \text{uev}_1(\varphi, \chi_2) & \text{if } \chi_1 = \langle \alpha^* \rangle \varphi \neq \chi_2 \\ \text{uev}_1(\chi_1, \chi_2) & \text{if } \chi_1 \in \Gamma \\ \perp & \text{otherwise} \end{cases} \\
 \\
 \text{Nx}_2 &:= \text{tst}(\langle \alpha \rangle \langle \alpha^* \rangle \varphi) \\
 \text{BD}_2 &:= \text{bl}(\langle \alpha \rangle \langle \alpha^* \rangle \varphi, \{\langle \alpha^* \rangle \varphi\} \cup \text{BD}) \\
 \text{uev}'_2(\chi_1, \chi_2) &:= \begin{cases} \text{uev}_2(\langle \alpha \rangle \langle \alpha^* \rangle \varphi, \chi_2) & \text{if } \chi_1 = \langle \alpha^* \rangle \varphi \\ \text{uev}_2(\chi_1, \chi_2) & \text{if } \chi_1 \in \Gamma \\ \perp & \text{otherwise} \end{cases}
 \end{aligned}$$

The $\langle * \rangle_1$ -rule captures the fix-point nature of the $\langle * \rangle$ -formulae according to Prop. 2.6 as long as the principal formula is not blocked via BD. The choice of the principal formulae in the first child is either free if φ is not a $\langle g \rangle$ -formula or is φ if φ is a $\langle g \rangle$ -formula. In the latter case we also block the regeneration of $\langle \alpha^* \rangle \varphi$ and thus avoid an “at a world” cycle by putting $\langle \alpha^* \rangle \varphi$ into BD_1 . The right child is treated similarly but uses $\langle \alpha \rangle \langle \alpha^* \rangle \varphi$ instead of φ .

Actions for all β -rules:

$$\begin{aligned}
 \text{stat} &:= \begin{cases} \mathbf{unsat} & \text{if } \text{stat}_1 = \mathbf{unsat} \ \& \ \text{stat}_2 = \mathbf{unsat} \\ \mathbf{open} & \text{if } \text{stat}_1 = \mathbf{open} \ \text{or} \ \text{stat}_2 = \mathbf{open} \\ \mathbf{barred} & \text{otherwise} \end{cases} \\
 \text{min}_\perp(f, g)(\chi_1, \chi_2) &:= \begin{cases} \perp & \text{if } f(\chi_1, \chi_2) = \perp \ \text{or} \ g(\chi_1, \chi_2) = \perp \\ \min(f(\chi_1, \chi_2), g(\chi_1, \chi_2)) & \text{otherwise} \end{cases} \\
 \text{uev} &:= \begin{cases} \text{uev}_\perp & \text{if } \text{stat} \neq \mathbf{open} \\ \text{uev}'_1 & \text{if } \text{stat}_1 = \mathbf{open} \neq \text{stat}_2 \\ \text{uev}'_2 & \text{if } \text{stat}_1 \neq \mathbf{open} = \text{stat}_2 \\ \min_\perp(\text{uev}'_1, \text{uev}'_2) & \text{if } \text{stat}_1 = \mathbf{open} = \text{stat}_2 \end{cases}
 \end{aligned}$$

The intuitions are:

uev'_i : the definitions of uev'_i ensure that the pairs (χ_1, χ_2) , where χ_1 is the principal $\langle \rangle$ -formula, get the values from their corresponding $\langle \rangle$ -formulae in the children.

In the $\langle * \rangle_1$ -rule, a special case sets the value of $\text{uev}'_1(\chi_1, \chi_2)$ to \perp if χ_1 and χ_2 are equal to the principal formula $\langle \alpha^* \rangle \varphi$ of this rule since the eventuality $\langle \alpha^* \rangle \varphi$ is no longer unfulfilled as the left child fulfils it. Note that $\text{uev}'(\chi_1, \chi_2)$ is only defined if χ_1 is in the parent.

min_\perp : the definition of min_\perp ensures that we take the minimum of $f(\chi_1, \chi_2)$ and $g(\chi_1, \chi_2)$ only when both functions are defined for (χ_1, χ_2) .

uev : if $\text{stat} \neq \mathbf{open}$, the uev is irrelevant, so we arbitrarily set it as undefined. If only one child has $\text{stat} = \mathbf{open}$, we take its uev' . If both children have $\text{stat} = \mathbf{open}$, we take the minimum value of entries that are defined in uev'_1 and uev'_2 .

All previous rules modify existing uev-entries, but never create new ones. The next rule is the only rule that creates uev-entries (by identifying loops).

Existential Branching Rule.

$$\begin{array}{c}
 \langle a_1 \rangle \varphi_1, \dots, \langle a_n \rangle \varphi_n, \langle a_{n+1} \rangle \varphi_{n+1}, \dots, \langle a_{n+m} \rangle \varphi_{n+m}, [-]\Delta, \Gamma \\
 :: \text{HCr}, \text{Nx}, \text{BD}, \text{BB} :: \text{stat}, \text{uev} \\
 (\langle \rangle) \frac{}{\varphi_1, \Delta_1 :: \text{HCr}_1, \text{Nx}_1, \text{BD}_1, \text{BB}_1 \quad | \dots | \quad \varphi_n, \Delta_n :: \text{HCr}_n, \text{Nx}_n, \text{BD}_n, \text{BB}_n} \\
 :: \text{stat}_1, \text{uev}_1 \qquad \qquad \qquad :: \text{stat}_n, \text{uev}_n
 \end{array}$$

where:

- (1) $n + m \geq 0$
- (2) $\Gamma \subseteq (\text{AFml} \cup \{-q \mid q \in \text{AFml}\})$
- (3) $[-]\Delta \subseteq \{[a]\psi \mid a \in \text{APrg} \ \& \ \psi \in \text{Fml}\}$
- (4) $\Delta_i := \{\psi \mid [a_i]\psi \in [-]\Delta\}$ for $i = 1, \dots, n$
- (5) $\forall p \in \text{AFml}. \{p, \neg p\} \not\subseteq \Gamma$
- (6) $\forall i \in \{1, \dots, n\}. \forall j \in \{1, \dots, \text{len}(\text{HCr})\}. (\varphi_i, \{\varphi_i\} \cup \Delta_i) \neq \text{HCr}[j]$
- (7) $\forall k \in \{n+1, \dots, n+m\}. \exists j \in \{1, \dots, \text{len}(\text{HCr})\}. (\varphi_k, \{\varphi_k\} \cup \Delta_k) = \text{HCr}[j]$

Actions for $(\langle \rangle)$: for $i = 1, \dots, n$: $\text{HCr}_i := \text{HCr} @ [(\varphi_i, \{\varphi_i\} \cup \Delta_i)]$,

$$\text{Nx}_i := \text{tst}(\varphi_i), \quad \text{BD}_i := \emptyset, \quad \text{BB}_i := \emptyset$$

$$\text{stat} := \begin{cases} \mathbf{unsat} & \text{if } \exists i \in \{1, \dots, n\}. \text{stat}_i \neq \mathbf{open} \text{ or} \\ & (\exists \psi \in \text{Fml}\langle * \rangle. \varphi_i \in \text{pre}(\psi) \ \& \\ & \perp \neq \text{uev}_i(\varphi_i, \psi) > \text{len}(\text{HCr})) \\ \mathbf{open} & \text{otherwise} \end{cases}$$

$$\begin{array}{l}
 \text{uev}_k(\cdot, \cdot) := j \in \{1, \dots, \text{len}(\text{HCr})\} \text{ such that } (\varphi_k, \{\varphi_k\} \cup \Delta_k) = \text{HCr}[j] \\
 \text{for } k = n+1, \dots, n+m
 \end{array}$$

$$\text{uev}(\chi_1, \chi_2) := \begin{cases} \text{uev}_i(\varphi_i, \chi_2) & \text{if } \text{stat} = \mathbf{open} \ \& \ \chi_2 \in \text{Fml}\langle * \rangle \ \& \ \chi_1 \in \text{pre}(\chi_2) \\ & \ \& \ \chi_1 = \langle a_i \rangle \varphi_i \text{ for an } i \in \{1, \dots, n+m\} \\ \perp & \text{otherwise} \end{cases}$$

Some intuitions are in order:

- (1) If $n = 0$, the application of the rule generates no new nodes and stat vacuously evaluates to **open**. If $m = n = 0$, we additionally have $\text{uev} := \text{uev}_\perp$.
- (2) The set Γ contains only propositional atoms or their negations.
- (3) The set $[-]\Delta$ contains only formulae of the type $[a]\varphi$. Thus (2) and (3) imply that the $\langle \rangle$ -rule is applicable only if the node contains no α - or β -formulae.

- (4) The set Δ_i contains all formulae that must belong to the i^{th} child, which fulfils $\langle a_i \rangle \varphi_i$, so that we can build a Hintikka structure later on.
- (5) The node must not contain a contradiction.
- (6) If $n > 0$, then each $\langle a_i \rangle \varphi_i$ for $1 \leq i \leq n$ is not “blocked” by an ancestor and has a child containing the formula set $\varphi_i \cup \Delta_i$ thereby generating the required successor for $\langle a_i \rangle \varphi_i$. Note that $\text{len}(\text{HCr})$ denotes the length of HCr.
- (7) If $m > 0$, then each $\langle a_k \rangle \varphi_k$ for $n + 1 \leq k \leq n + m$ is “blocked” from creating its required child $\{\varphi_k\} \cup \Delta_k$ because some ancestor does the job. This ancestor must not only consist of the formulae $\{\varphi_k\} \cup \Delta_k$ but it must also have been created to fulfil $\langle a \rangle \varphi_k$ for some $a \in \text{APrg}$. Note that the values a_k and a are ignored when looking for loops since we are interested only in the contents of the required child.

HCr_i : is the HCr of the parent extended with an extra entry to record the “history” of worlds created on the path from the root down to the i^{th} child using “@” as list concatenation. Note that we store a *pair* $(\varphi_k, \varphi_k \cup \Delta_k)$, not just $\varphi_k \cup \Delta_k$. That is, we remember that the node $\varphi_k \cup \Delta_k$ was created to fulfil $\langle a \rangle \varphi_k$ for some $a \in \text{APrg}$.

stat : the parent is unsatisfiable if some child has $\text{stat} \neq \mathbf{open}$. But it is also unsatisfiable if some child, say the i^{th} , and some eventuality $\langle \alpha^* \rangle \chi$ in it “loops lower” because $\varphi_i \in \text{pre}(\langle \alpha^* \rangle \chi)$ and $\text{uev}_i(\varphi_i, \langle \alpha^* \rangle \chi)$ is defined and greater than the length of the current HCr. Intuitively, the latter tells us that the eventuality $\langle \alpha^* \rangle \chi$ occurs in the sub-tableau rooted at the parent but cannot be fulfilled.

uev_k : for $n + 1 \leq k \leq n + m$, the k^{th} child is blocked by a higher (proxy) child. For every such k we set uev_k to be the *constant* function which maps every formula-pair to the *level* j of its proxy child. This is just a temporary function used to define uev as explained next. The blocking child itself must have been created to fulfil a $\langle \cdot \rangle$ -formula $\langle a' \rangle \varphi_k$, as indicated by the first component of $\text{HCr}[j]$.

$\text{uev}(\chi_1, \chi_2)$: If $\text{stat} = \mathbf{unsat}$ then uev is undefined everywhere. Else, for each $\chi_1 = \langle a_i \rangle \varphi_i$ with $i \in \{1, \dots, n + m\}$, and each χ_2 with $\langle a_i \rangle \varphi_i \in \text{pre}(\chi_2)$, we take $\text{uev}(\langle a_i \rangle \varphi_i, \chi_2)$ from the formulae-pair (φ_i, χ_2) of the corresponding (real) child if $\langle a_i \rangle \varphi_i$ is “unblocked”, or set it to the level of the proxy child higher in the branch if it is “blocked”. For all other formulae-pairs, uev is undefined. The intuition is that a defined $\text{uev}(\chi_1, \chi_2)$ flags a “loop” which starts at the parent and eventually “loops” up to some blocking proxy. The value of $\text{uev}(\chi_1, \chi_2)$ tells us the level of the proxy because we cannot classify this “loop” as “good” or “bad” until we backtrack to that level. The uev of each $\langle a_i \rangle \varphi_i$ is taken from *the* child created specifically to contain φ_i , a fact which is vital in the proofs.

$\text{BD}_i, \text{BB}_i, \text{Nx}_i$: each child has no blocked diamond- or box-formulae, and its principal formula is determined by the form of φ_i .

The $\langle \cdot \rangle$ - and *id*-rules are mutually exclusive via their side-conditions. Our rules are designed so that at least one rule is applicable to any node. As shown in the next section, we need to build only one fully expanded tableau, hence if multiple rules are applicable to a node, the choice of rule is immaterial. Of course, in our implementation, we give priority to the *id*-rule since it may close a branch sooner. Other heuristics, like preferring linear rules over branching rules, are also useful.

4.2 Termination, Soundness, and Completeness

Definition 4.4 Let $x = (\Gamma :: \text{HCr}, \text{Nx}, \text{BD}, \text{BB} :: \text{stat}, \text{uev})$ be a tableau node, φ a formula, and Δ a set of formulae. We write $\varphi \in x [\Delta \subseteq x]$ to mean $\varphi \in \Gamma [\Delta \subseteq \Gamma]$. The parts of x are written as HCr_x , Nx_x , BD_x , BB_x , stat_x , and uev_x . Node x is *closed* iff $\text{stat}_x = \mathbf{unsat}$, *open* iff $\text{stat}_x = \mathbf{open}$, and *barred* iff $\text{stat}_x = \mathbf{barred}$.

Definition 4.5 Let x be a $\langle \rangle$ -node in a tableau T (i.e. a $\langle \rangle$ -rule was applied to x). Then x is also called a *state* and the children of x are called *core-nodes*. Using the notation of the $\langle \rangle$ -rule, a formula $\langle a_i \rangle \varphi_i \in x$ is *blocked* iff $n + 1 \leq i \leq n + m$. For every not blocked $\langle a_i \rangle \varphi_i \in x$, the *successor* of $\langle a_i \rangle \varphi_i$ is the i^{th} child of the $\langle \rangle$ -rule. For every blocked $\langle a_i \rangle \varphi_i \in x$ there exists a unique core-node y on the path from the root of T to x such that $\{\varphi_i\} \cup \Delta_i$ is the set of formulae of y , and y is the successor of a formula $\langle a' \rangle \varphi_i$ in the parent of y . We call y the *virtual successor* of $\langle a_i \rangle \varphi_i$, and also call the formula φ_i in the (possibly virtual) successor of $\langle a_i \rangle \varphi_i$ a *core-formula*.

A state is another term for a $\langle \rangle$ -node but a core-node can be any type of node (even a state). A state arises from a core-node by α - and β -rules. Note that the core-formula in a core-node y is well-defined and unique: if x_1 and x_2 are states and y is the (possibly virtual) successor of $\langle a_1 \rangle \varphi_1 \in x_1$ and $\langle a_2 \rangle \varphi_2 \in x_2$, then $\varphi_1 = \varphi_2$.

Let ϕ be a formula in negation normal form, and T an expanded tableau with root $r = (\{\phi\} :: [], \perp, \emptyset, \emptyset :: \text{stat}, \text{uev})$ with stat and uev determined by r 's children.

Theorem 4.6 T is a finite tree.

Theorem 4.7 If the root $r \in T$ is open, there is a Hintikka structure for ϕ .

Theorem 4.8 If the root $r \in T$ is not open then ϕ is not satisfiable.

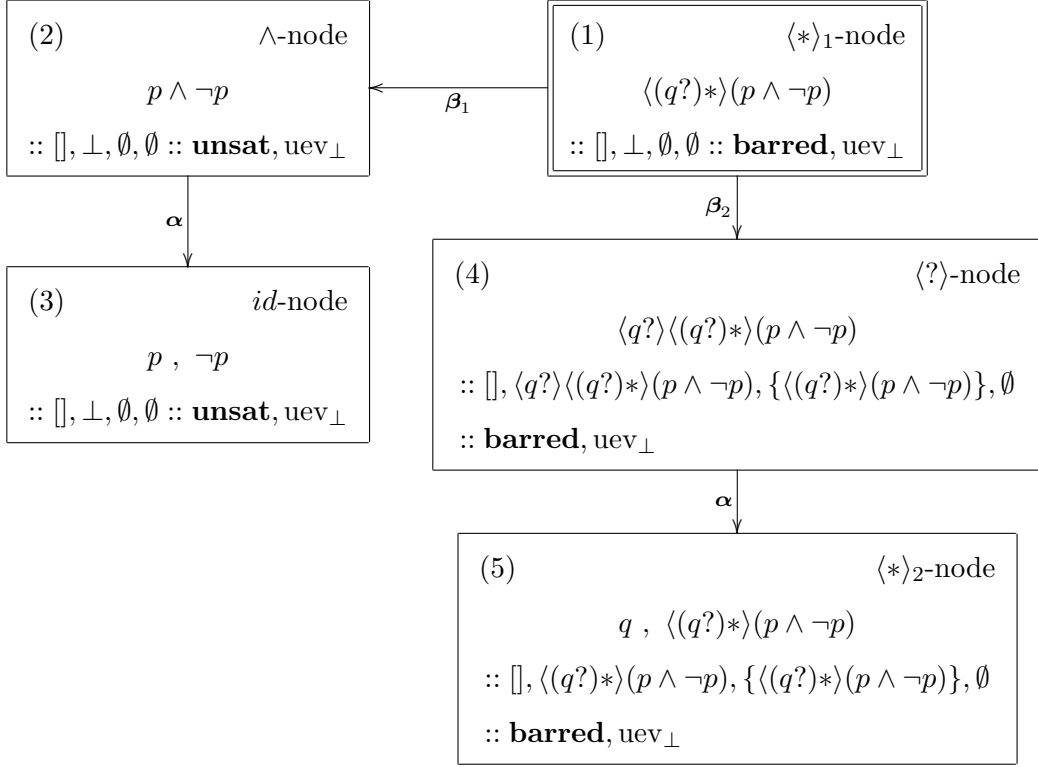
Theorem 4.9 If $|\phi| = n$, our procedure has worst-case time complexity in $O(2^{2^n})$.

The length of a branch in a tableau is bounded, essentially by the number of core-nodes on that branch. The number of core-nodes itself is bounded, essentially by the cardinality of the power set of the set $\text{cl}(\phi)$ of all formula that can appear in the tableau. The size of $\text{cl}(\phi)$ is polynomial in n , hence the length of a branch is in $O(2^n)$. Thus the overall (worst case) number of nodes in a tableau is in $O(2^{2^n})$.

4.3 Fully Worked Examples

The first simple example illustrates how the procedure avoids infinite loops due to “at a world” cycles by blocking $\langle \alpha^* \rangle \varphi$ - and $[\alpha^*] \varphi$ -formulae from regenerating. The formula $\langle (q^?)^* \rangle (p \wedge \neg p)$ is obviously not satisfiable. Hence, any expanded tableau with root $\langle (q^?)^* \rangle (p \wedge \neg p)$ should not be open. Figure 1 shows such a tableau where each node is classified as a ρ -node if rule ρ is applied to that node in the tableau.

The initial formula $\langle (q^?)^* \rangle (p \wedge \neg p)$ in node (1) is decomposed into a β_1 -child $p \wedge \neg p$ and a β_2 -child $\langle q^? \rangle \langle (q^?)^* \rangle (p \wedge \neg p)$ according to the $\langle * \rangle_1$ -rule. The formula $p \wedge \neg p$ in node (2) is then decomposed according to the \wedge -rule and node (3) is marked as closed because it contains a contradiction. Node (2) inherits the status from node (3) unchanged according to the α -rules and, thus, is closed too.


 Fig. 1. A first example: a closed tableau for $\langle (q?)^* \rangle (p \wedge \neg p)$

Because the β_2 -formula $\langle q? \rangle \langle (q?)^* \rangle (p \wedge \neg p)$ is a $\langle q \rangle$ -formula, the $\langle * \rangle_1$ -rule puts this formula into its Nx_2 , the Nx -value of node (4), and thus forces node (4) to have $\langle q? \rangle \langle (q?)^* \rangle (p \wedge \neg p)$ as its principal formula. For the same reason, the $\langle * \rangle_1$ -rule puts its own principal formula $\langle (q?)^* \rangle (p \wedge \neg p)$ into its BD_2 , the BD -value of node (4). Hence node (4) decomposes $\langle q? \rangle \langle (q?)^* \rangle (p \wedge \neg p)$ according to the $\langle ? \rangle$ -rule. Again, the resulting node (5) is forced to have $\langle (q?)^* \rangle (p \wedge \neg p)$ as its principal formula via its Nx -value, and gets its BD -value unchanged from node (4).

Node (5) has the same principal formula as node (1), so applying the $\langle * \rangle_1$ -rule to node (5) would cause the procedure to enter an “at a world” (infinite) cycle. Because the history BD of node (5) contains $\langle (q?)^* \rangle (p \wedge \neg p)$, the $\langle * \rangle_1$ -rule is blocked on node (5), but the $\langle * \rangle_2$ -rule is not. Hence the branch is terminated and the status of node (5) is set to **barred** (thereby avoiding the “at a world” cycle).

Node (4) inherits the status from node (5) unchanged and node (1) is marked **barred** also according to the definition of $stat$ in the β -rules. Therefore the tableau is not open. Note that the variable uev does not play a role in this example as it is irrelevant for nodes that are closed or barred.

The second example demonstrates the role of uev . The formula $[a^*]p \rightarrow [(a; a)^*]p$ is valid. Hence, its negation $\phi := [a^*]p \wedge \langle (a; a)^* \rangle \neg p$, which is already in negation normal form, is unsatisfiable and the root of any expanded tableau for ϕ should not be open. Figure 2 shows such a tableau. The unlabelled edges in Fig. 2 link states to core-nodes. We omit the histories BD and BB as they do not play an important role in this example. Each partial function UEV_i maps the formula-pair (ψ_i, χ_i) in Table 2 to 1 and is undefined otherwise as explained be-

Table 2
 Definitions for the example in Fig. 2

UEV_i	$i = 1$	$i = 2$	$i = 3$	$i = 4$
ψ_i	$\langle a \rangle \langle a \rangle \langle (a; a)^* \rangle \neg p$	$\langle a; a \rangle \langle (a; a)^* \rangle \neg p$	$\langle (a; a)^* \rangle \neg p$	$\langle a \rangle \langle (a; a)^* \rangle \neg p$
χ_i	$\langle (a; a)^* \rangle \neg p$	$\langle (a; a)^* \rangle \neg p$	$\langle (a; a)^* \rangle \neg p$	$\langle (a; a)^* \rangle \neg p$

low. The histories are $HCR_1 := [(\varphi_1, \Delta_1)]$ where $\varphi_1 := \langle a \rangle \langle (a; a)^* \rangle \neg p$ and $\Delta_1 := \{[a^*]p, \langle a \rangle \langle (a; a)^* \rangle \neg p\}$ and $HCR_2 := HCR_1 @ [(\varphi_2, \Delta_2)]$ where $\varphi_2 := \langle (a; a)^* \rangle \neg p$ and $\Delta_2 := \{[a^*]p, \langle (a; a)^* \rangle \neg p\}$.

The dotted frame at (7a) indicates that its child, an *id*-node, is not shown due to space restrictions. Thus the marking of the nodes (3a) and (7a) in Fig. 2 with **unsat** is straightforward. The leaf (9) is a $\langle \rangle$ -node, but it is “blocked” from creating its successor containing $\Delta := \{[a^*]p, \langle a \rangle \langle (a; a)^* \rangle \neg p\}$ because there is a $j \in \mathbb{N}$ such that $HCr_9[j] = HCr_2[j] = (\langle a \rangle \langle (a; a)^* \rangle \neg p, \Delta)$: namely $j = 1$. Thus the $\langle \rangle$ -rule computes $UEV_1(\langle a \rangle \varphi_1, \langle (a; a)^* \rangle \neg p) = 1$ as stated above and also puts $\text{stat}_9 := \mathbf{open}$. As node (7a) is closed, nodes (8), (7b), (7), (6), and (5) inherit their functions UEV_i from their open children via the corresponding α - and β -rules.

The crux of our method occurs at node (4), a $\langle \rangle$ -node with $HCr_4 = []$ and hence $\text{len}(HCr_4) = 0$. The $\langle \rangle$ -rule thus finds a child node (5) and a pair of formulae $(\psi, \chi) := (\langle a \rangle \langle (a; a)^* \rangle \neg p, \langle (a; a)^* \rangle \neg p)$ where ψ is a core-formula, $\psi \in \text{pre}(\chi)$, and $1 = UEV_4(\psi, \chi) = \text{uev}_5(\psi, \chi) > \text{len}(HCr_4) = 0$. Thus node (4) “sees” a child (5) that “loops lower”, meaning that node (5) is the root of an “isolated” subtree which fails to fulfil its eventuality $\langle (a; a)^* \rangle \neg p$. The $\langle \rangle$ -rule marks (4) as closed via $\text{stat}_4 = \mathbf{unsat}$. The propagation of **unsat** to the root is simple.

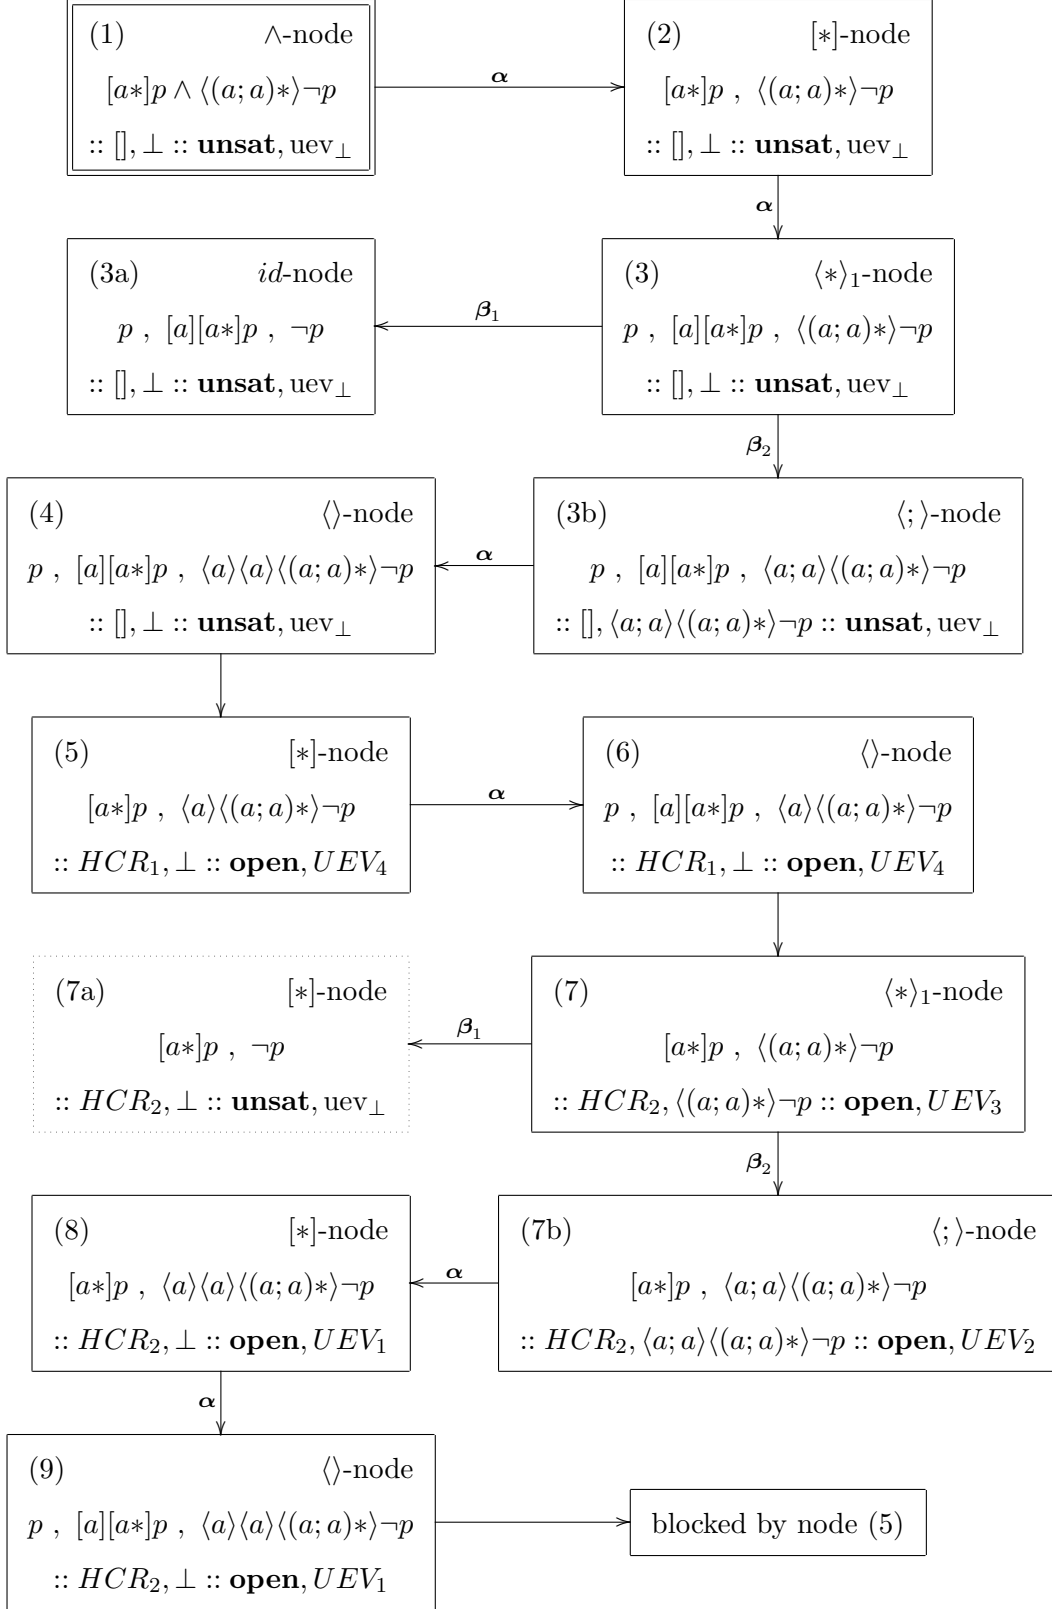
What if the omitted child of (7a), and hence (7a) itself, had been open? Then UEV_3 in (7) would be undefined everywhere via the $\langle * \rangle_1$ -rule, regardless of uev_{7b} . Thus $\langle (a; a)^* \rangle \neg p$ in (7) would be fulfilled via the β_1 -child (7a). Hence UEV_4 would be undefined everywhere, and node (4) would not be closed.

5 Conclusion and Further Work

We have given a sound, complete and terminating procedure for checking *PDL*-satisfiability. Unfortunately, its worst-case time-complexity is in 2EXPTIME rather than in EXPTIME , thus our procedure is sub-optimal. We now outline some further practical and theoretical work which may eliminate this disadvantage.

First, we believe that a small refinement of our histories will allow our calculus to classify a loop as “bad” or “good” at the looping leaf, as is done by Baader’s procedure [3], but with no pre-computation of automata. Thus it should be possible to extend DLP to handle our method. Further experimental work is required to determine if such an extension will remain practical.

Second, recent work has shown that global caching can indeed deliver optimality of tableau procedures soundly [10]. The histories used in our calculus make it harder to extend sound global caching to it since nodes are now sensitive to their context in the tree under construction. Further theoretical work is required to extend sound global caching to handle such context sensitivity.


 Fig. 2. A second example: a closed tableau for $[a^*]p \wedge \langle (a; a)^* \neg p$

References

- [1] Abate, P., R. Goré, and F. Widmann. One-pass tableaux for computation tree logic. In N. Dershowitz and A. Voronkov, editors, *Proc. LPAR 2007*, LNCS 4790:32-46, Springer, 2007.
- [2] Alberucci L., and G. Jäger. About cut elimination for logics of common knowledge. *Annal of Pure and Applied Logic*, 133(1-3):73–99, 2005.
- [3] Baader, F. Augmenting concept languages by transitive closure of roles: an alternative to terminological cycles. Technical Report RR-90-13, DFKI, Kaiserslautern, Germany, 1990.
- [4] Ben-Ari, M., Z. Manna, and A. Pnueli. The temporal logic of branching time. In *Proc. Eighth ACM Symposium on Principles of Programming Languages*, pages 164–176, 1981.
- [5] Bhat, G., and R. Cleaveland. Efficient on-the-fly model checking for CTL*. In *Proc. Symposium on Logics in Computer Science*, pages 388–397, 1995.
- [6] Cleaveland, R. Tableau-based model checking in the propositional mu-calculus. *Acta Informatica*, 27:725–747, 1990.
- [7] Emerson, E. A., and J. Y. Halpern. Decision procedures and expressiveness in the temporal logic of branching time. *Journal of Computer and System Science*, 30:1–24, 1985.
- [8] Fischer, M. J., and R. E. Ladner. Propositional dynamic logic of regular programs. *Journal of Computer Systems and Science*, 18:194-211, 1979.
- [9] Giacomo, G. D., and F. Massacci. Combining deduction and model checking into tableaux and algorithms for Converse-PDL. *Information and Computation*, 160:109–169, 2000.
- [10] Goré, R., and L. A. Nguyen. Exptime tableaux for ALC using sound global caching. In *DL'07: Proc. International Workshop on Description Logics, 2007*, 8 pages.
- [11] Horrocks, I., and P. F. Patel-Schneider. Optimising description logic subsumption. *Journal of Logic and Computation*, 9(3):267–293, OUP, 1999.
- [12] Jäger, G., M. Kretz, and T. Studer. Cut-free common knowledge. *Journal of Applied Logic*, to appear.
- [13] Kozen, D., and R. Parikh. An elementary proof of the completeness of PDL. *Theoretical Computer Science*, 14:113–118, 1981.
- [14] Pratt, V. Semantical considerations on Floyd-Hoare logic. In *Proc. 17th IEEE Symposium on Foundations Of Computer Science*, pages 109–121, 1976.
- [15] Pratt, V. A near-optimal method for reasoning about action. *Journal of Computer Systems and Science*, 20:231–254, 1980.
- [16] Schmidt, R., and D. Tishkovsky. Personal communication, September 2007. <http://www.cs.man.ac.uk/~schmidt/pdl-tableau>
- [17] Schwendimann, S. A new one-pass tableau calculus for PLTL. In H. de Swart, editor, *Proc. TABLEAUX'98*, LNAI 1397:277-291. Springer, 1998.
- [18] Wolper, P. Temporal logic can be more expressive. *Information and Computation*, 56:72–99, 1983.

Appendix: Termination, Soundness and Completeness

Definition 5.1 Let $G = (W, R)$ be a directed graph (*e.g.* a tableau where R is just the child-of relation between nodes). A *path* π in G is a finite or infinite sequence x_0, x_1, x_2, \dots of nodes in W such that $x_i R x_{i+1}$ for all x_i except the last node if π is finite.

Termination

Theorem 4.6 T is a finite tree.

Proof Sketch It is obvious that T is a tree and that every node in T can contain only formulae from the negation normal form analogue $\text{cl}(\phi)$ of the Fisher-Ladner closure [8]. The definition of $\text{cl}(\phi)$ has been omitted to save space, but $\text{cl}(\phi)$ is finite. Hence there are only a finite number of different sets that can be assigned to nodes, in particular core-nodes, and the number of pairs (φ, Δ) with $\varphi \in \Delta \subseteq \text{cl}(\phi)$ is finite. As each core-node is assigned such a pair and the $\langle \rangle$ -rule ensures core-nodes on a branch possess different pairs, the number of core-nodes on a branch is finite.

It is not obvious that the number of nodes between consecutive core-nodes on a branch is finite since $\langle \alpha * \rangle$ - and $[\alpha *]$ -formulae like $\langle a * * \rangle \varphi$ can “regenerate” on a branch without passing a core-node (*e.g.* $\langle a * * \rangle \varphi \rightsquigarrow \langle a * \rangle \langle a * * \rangle \varphi \rightsquigarrow \langle a * * \rangle \varphi$). However, it is relatively easy to see that formulae of the form $\langle \alpha * \rangle \varphi$ or $[\alpha *] \varphi$ are the only potential “troublemakers” between two states. For formulae of the form $[\alpha *] \varphi$ regeneration between two core-nodes is clearly ruled out by the history BB and the $[*]$ -rule. For formulae of the form $\langle \alpha * \rangle \varphi$, the job is done by the history BD and the $\langle * \rangle_1$ and $\langle * \rangle_2$ -rules. In the latter case, it is crucial that the procedure chooses the decomposition of a principal $\langle \alpha \rangle$ -formula as the principal formula of the child, provided that the decomposition is also a $\langle \alpha \rangle$ -formula.

As the number of nodes between two core-nodes is finite, and there are only finitely many core-nodes on any branch, all branches in T are finite. Every node has finite degree so König’s lemma completes the proof. \square

Soundness

Theorem 4.7 If the root $r \in T$ is open, there is a Hintikka structure for ϕ .

Proof. By construction, T is a finite tree. Let T_p (“p” for pruned) be the subgraph that consists of all nodes x having the following property: there is a path of open nodes from r to x inclusive. The edges of T_p are exactly the edges of T that connect two nodes in T_p . Clearly, T_p is also a finite tree with root r . Intuitively, T_p is the result of pruning all subtrees of T that have a closed or barred root.

Next, we extend T_p to a finite cyclic tree T_l (“l” for looping) by doing the following for every state x : for every formula $\langle a \rangle \varphi \in x$ having a *virtual* successor y , which must lie on the path from r to x , we add the edge (x, y) to T_l . These new edges are called *backward edges*. Note that as *id*-nodes are closed by construction of T , all leaves of T_p must be states where all $\langle \rangle$ -formulae (if any) are blocked. Hence every formula $\langle a \rangle \varphi$ of every leaf has a virtual successor.

Finally, following Ben-Ari et al. [4], the cyclic tree T_l is used to generate a structure $H = (W, R, L)$ as described next. Let W be the set of all states of T_l .

For every $a \in \text{APrg}$ and every $s, t \in W$, let $s R_a t$ iff s contains a formula $\langle a \rangle \psi$ and there exists a path $x_0 = s, x_1, \dots, x_{k+1} = t$ in T_1 such that x_1 is the (possibly virtual) successor of $\langle a \rangle \psi$ and each $x_i, 1 \leq i \leq k$ is an α - or a β -node. Thus state t is a “saturation” of x_1 using only α - and β -rules. Note that $s R_a t$ and $s R_b t$ is possible for $a \neq b$, because two formulae $\langle a \rangle \psi \in s$ and $\langle b \rangle \psi \in s$ might have the same virtual successor: see point (7) of the $\langle \rangle$ -rule. It is also possible that $s R_a t$ and $s R_a u$ for $t \neq u$.

If we consider the root r of T_1 as a core-node for a moment, it is not hard to see that for every state $s \in T_1$ there exists a unique core-node $x \in T_1$ and a unique path π of the form $x_0 = x, x_1, \dots, x_k = s$ in T_1 such that either $k = 0$ (and thus $s = x$) or $k > 0$ and each $x_i, 0 \leq i \leq k - 1$ is not a state. We set $L(s)$ to be the union of all formulae of all nodes on π . Intuitively, we form $L(s)$ by adding back all the principal formulae of the α - and β -rules which were applied to obtain s from x .

It is almost straightforward to check that H is a pre-Hintikka structure for ϕ . There are only two things that deserve extra comments: Firstly, it is not possible that T_1 contains a $\langle * \rangle_2$ -node as it would be barred. Secondly, assume that $y \in T_1$ is a $[*]$ -node with principal formula $[\alpha*]\varphi$ and s is a state such that y lies on the path π to s that defines the set $L(s)$, which contains $[\alpha*]\varphi$, as described above. Then either φ and $[\alpha][\alpha*]\varphi$ are contained in the child of y in T_1 , or – as the first node x on π is a core-node with $\text{BB}_z = \emptyset$ – there exists another $[*]$ -node on π that also has $[\alpha*]\varphi$ as principal formula and its child in T_1 contains φ and $[\alpha][\alpha*]\varphi$. As the child of an α -node that lies on π must lie on π too, in both cases, there is a node on π containing φ and $[\alpha][\alpha*]\varphi$. Thus φ and $[\alpha][\alpha*]\varphi$ are also contained in $L(s)$.

To show that H is even a Hintikka structure we use Lemma 5.2 to conclude $H6$ as is shown next.

Suppose $\langle \alpha* \rangle \varphi \in L(s)$. If we also have $\varphi \in L(s)$ then $(s, \langle \alpha* \rangle \varphi), (s, \varphi)$ is a fulfilling chain for $(\varphi, \alpha*, s)$ and we are done. Otherwise, the finiteness of the tableau and the fact that H is a pre-Hintikka structure give us a sequence $\sigma = (s, \varphi_0), \dots, (s, \varphi_m)$ such that:

- $\varphi_i \in \text{pre}(\langle \alpha* \rangle \varphi)$ and $\varphi_i \in L(s)$ for all $0 \leq i \leq m$
- $\varphi_0 = \langle \alpha* \rangle \varphi$ and $\varphi_m = \langle a \rangle \varphi'$ for some $a \in \text{APrg}$ and $\varphi' \in \text{Fml}$
- $\varphi_i \rightsquigarrow \varphi_{i+1}$ for all $0 \leq i \leq m - 1$.

Applying Lemma 5.2 for the state s and the formula $\varphi_m = \langle a \rangle \varphi'$ gives us a sequence $\sigma' := (y_0, \psi_0), \dots, (y_n, \psi_n)$ with the properties stated in Lemma 5.2. Let y_n, \dots, y_{n+m} be an arbitrary path in T_1 such that y_{n+m} is a state. Next, we replace each $y_i, 1 \leq i \leq n$ in σ' with the first state s_i that appears on the path $y_i, \dots, y_n, \dots, y_{n+m}$.

It is easy to check that the combined sequence σ, σ' is a fulfilling chain for $(\varphi, \alpha*, s)$ in H if we contract all consecutive repetitions of pairs. This concludes the proof. \square

Lemma 5.2 *Let $y \in T_1$ be a node and $\psi \in y$ a formula such that $\psi \in \text{pre}(\langle \alpha* \rangle \varphi)$. There exists a finite sequence $\sigma' = (y_0, \psi_0), \dots, (y_n, \psi_n)$ of pairs with $n \geq 0$ such that:*

- y_0, \dots, y_n is a path in T_1

- $y_i \in T_1$, $\psi_i \in \text{pre}(\varphi)$, and $\psi_i \in y_i$ for all $0 \leq i \leq n$
- $y_0 = y$, $\psi_0 = \psi$, $\psi_n = \varphi$, and $\psi_i \neq \varphi$ for all $0 \leq i \leq n-1$
- for all $0 \leq i \leq n-1$, either $\psi_i = \psi_{i+1}$ or: if $\psi_i = \langle a \rangle \chi$ for some $a \in \text{APrg}$ and $\chi \in \text{Fml}$ then y_i is a state else $\psi_i \rightsquigarrow \psi_{i+1}$.

Proof. We inductively construct σ' starting with $(y_0, \psi_0) := (y, \psi)$. Most of the required properties of σ' follow directly from its construction and we leave it to the reader to check that they hold.

Step 1 Let (y_i, ψ_i) be the last pair of σ' . We distinguish three cases: either y_i is an α - or β -node and ψ_i is not the principal formula in y_i ; or y_i is an α - or β -node and ψ_i is the principal formula in y_i ; or y_i is a state.

If y_i is an α - or β -node and ψ_i is not the principal formula in y_i , we set $\psi_{i+1} := \psi_i$ and we choose y_{i+1} to be a successor of y_i in T_1 such that $\text{uev}_{y_i}(\psi_i, \langle \alpha^* \rangle \varphi) = \text{uev}_{y_{i+1}}(\psi_{i+1}, \langle \alpha^* \rangle \varphi)$. Note that such a y_{i+1} always exists since the value of $\text{uev}_{y_i}(\psi_i, \langle \alpha^* \rangle \varphi)$ is determined by one of its open children during the construction of T and hence T_1 . But it does not have to be unique. We then repeat Step 1.

If y_i is an α - or β -node and ψ_i is the principal formula in y_i , we look at all pairs (x, χ) such that x is a child of y_i in T_1 and ψ_i is decomposed into $\chi \in x$ and $\psi_i \rightsquigarrow \chi$ holds. By construction of T and hence T_1 there is at least one open child such that the corresponding pair (x, χ) obeys $\text{uev}_{y_i}(\psi_i, \langle \alpha^* \rangle \varphi) = \text{uev}_x(\chi, \langle \alpha^* \rangle \varphi)$. Let (y_{i+1}, ψ_{i+1}) be such a pair. If $\psi_{i+1} = \varphi$ we stop and return σ' ; otherwise we repeat Step 1.

If y_i is a state, it is not too hard to see that ψ_i must be of the form $\langle a \rangle \chi$ for some $a \in \text{APrg}$ and $\chi \in \text{Fml}$. We set $(y_{i+1}, \psi_{i+1}) := (x, \chi)$ where x is the (possibly virtual) successor of $\psi_i = \langle a \rangle \chi$ and repeat Step 1. Note that if x is a non-virtual successor of ψ_i , we have $\text{uev}_{y_i}(\psi_i, \langle \alpha^* \rangle \varphi) = \text{uev}_{y_{i+1}}(\psi_{i+1}, \langle \alpha^* \rangle \varphi)$ by construction of T and hence T_1 . Also note that if x is a virtual successor of ψ_i then $\psi_{i+1} = \chi$ is the core-formula of y_{i+1} by construction of T and hence T_1 .

The only way for Step 1 to terminate is by finding $\psi_{i+1} = \varphi$. It is not difficult to see that the resulting (finite) sequence σ' fulfils all requirements and the proof is completed. Hence the rest of the proof shows that σ' as constructed by Step 1 is finite. Step 1 maintains the following invariant:

(†) For all appropriate $i \in \mathbb{N}$ we have $\text{uev}_{y_i}(\psi_i, \langle \alpha^* \rangle \varphi) = \text{uev}_{y_{i+1}}(\psi_{i+1}, \langle \alpha^* \rangle \varphi)$ unless y_{i+1} is the *virtual* successor of $\psi_i \in y_i$.

In other words, the values of $\text{uev}_{y_i}(\psi_i, \langle \alpha^* \rangle \varphi)$ and $\text{uev}_{y_{i+1}}(\psi_{i+1}, \langle \alpha^* \rangle \varphi)$ can differ only if (y_i, y_{i+1}) is a backward edge in T_1 . We distinguish two cases: either $\text{uev}_{y_0}(\psi_0, \langle \alpha^* \rangle \varphi)$ is undefined or it is defined. In both cases we show that the path y_0, y_1, \dots can only have a finite number of backward edges. As every infinite path in T_1 must use an infinite number of backward edges since T and T_p are finite trees, this proves that Step 1 terminates.

Case 1. If $\text{uev}_{y_0}(\psi_0, \langle \alpha^* \rangle \varphi)$ is undefined, the path y_0, y_1, \dots cannot contain a backward edge as shown next. Assume for a contradiction that y_i with $i \geq 0$ is the first node such that (y_i, y_{i+1}) is a backward edge. Since the initial $\text{uev}_{y_0}(\psi_0, \langle \alpha^* \rangle \varphi)$ was

undefined, by (†) we know that $\text{uev}_{y_i}(\psi_i, \langle \alpha^* \rangle \varphi)$ is undefined. But y_i is a state and as $\psi_i \in y_i$, which must be of the form $\langle a \rangle \chi$ for some $a \in \text{APrg}$ and $\chi \in \text{Fml}$, has a virtual successor z , $\text{uev}_{y_i}(\psi_i, \langle \alpha^* \rangle \varphi)$ is defined to be the height of z by the application of the $\langle \rangle$ -rule to y_i during the construction of the tableau. Thus $\text{uev}_{y_i}(\psi_i, \langle \alpha^* \rangle \varphi)$ is both defined and undefined, which is a contradiction.

Case 2. If $h := \text{uev}_{y_0}(\psi_0, \langle \alpha^* \rangle \varphi)$ is defined, the path y_0, y_1, \dots can only contain a finite number of backward edges as shown next. Let y_i with $i \geq 0$ be the first node such that (y_i, y_{i+1}) is a backward edge. If no such node exists, we are obviously done. Otherwise, we have $\text{uev}_{y_i}(\psi_i, \langle \alpha^* \rangle \varphi) = h$ by (†). This means by construction of the tableau that there exists a set $\Delta \subseteq \text{Fml}$ such that $(\psi_{i+1}, \{\psi_{i+1}\} \cup \Delta) = \text{HCr}_{y_i}[h]$. Thus y_{i+1} is the h^{th} core-node (child of a $\langle \rangle$ -node) on the path from the root r to y_i in T_1 and we have $\text{len}(\text{HCr}_{y_{i+1}}) = h$ by construction of HCr .

If $\text{uev}_{y_{i+1}}(\psi_{i+1}, \langle \alpha^* \rangle \varphi)$ had a value equal to or greater than h then the $\langle \rangle$ -rule would cause the parent of y_{i+1} in T_1 to be marked as closed since ψ_{i+1} is the core-formula of y_{i+1} ; but we know this is not the case. Hence $\text{uev}_{y_{i+1}}(\psi_{i+1}, \langle \alpha^* \rangle \varphi)$ is either undefined or has a value h' that is strictly smaller than h .

If $\text{uev}_{y_{i+1}}(\psi_{i+1}, \langle \alpha^* \rangle \varphi)$ is undefined, we can prove exactly as in Case 1 that the path y_{i+1}, y_{i+2}, \dots cannot contain a backward edge. On the other hand, if $h' := \text{uev}_{y_{i+1}}(\psi_{i+1}, \langle \alpha^* \rangle \varphi)$ is defined, we can inductively repeat the arguments in Case 2 for the sequence $(y_{i+1}, \psi_{i+1}), (y_{i+2}, \psi_{i+2}), \dots$. The induction is well-defined because of $h' < h$, meaning that eventually this inductive argument must terminate because all such h -values must be in $\mathbb{N}_{>0}$. \square

Completeness

Definition 5.3 Let $M = (W, R, V)$ be a model, $w \in W$ a state and $\varphi \in \text{Fml}$ a formula of the form $\varphi = \langle \alpha_1 \rangle \dots \langle \alpha_k \rangle \psi$ for some $k > 0$ and $\alpha_1, \dots, \alpha_k \in \text{Prg}$ and $\psi \in \text{Fml}$. A *witness chain* for (φ, ψ, M, w) is a finite sequence $(w_0, \psi_0), \dots, (w_n, \psi_n)$ of world-formula pairs with $n > 0$ such that:

- (1) $w_i \in W$, $\psi_i \in \text{pre}(\psi)$, and $M, w_i \Vdash \psi_i$ for all $0 \leq i \leq n$
- (2) $w_0 = w$, $\psi_0 = \varphi$, $\psi_n = \psi$, and $\psi_i \neq \psi$ for all $0 \leq i \leq n-1$
- (3) $\forall i, j \in \{0, \dots, n\}. i \neq j \Rightarrow (w_i, \psi_i) \neq (w_{i+1}, \psi_{i+1})$
- (4) for all $0 \leq i \leq n-1$, if $\psi_i = \langle a \rangle \chi$ for some $a \in \text{APrg}$ and $\chi \in \text{Fml}$ then $\psi_{i+1} = \chi$ and $w_i R_a w_{i+1}$; otherwise $\psi_i \rightsquigarrow \psi_{i+1}$ and $w_i = w_{i+1}$.

Proposition 5.4 *In the setting of Def. 5.3, we have:*

- (1) for every $1 \leq i \leq k$ there exists an $m < n$ such that $(w_0, \psi_0), \dots, (w_m, \psi_m)$ is a witness chain for $(\varphi, \langle \alpha_i \rangle \dots \langle \alpha_k \rangle \psi, M, w)$
- (2) if $\alpha_k = \beta^*$ for some $\beta \in \text{Prg}$ then $\psi_{n-1} = \langle \beta^* \rangle \psi$.

Proposition 5.5 *Let $M = (W, R, V)$ be a model, $w \in W$ a state and $\varphi \in \text{Fml}$ a formula of the form $\varphi = \langle \alpha_1 \rangle \dots \langle \alpha_k \rangle \psi$ for some $k > 0$ and $\alpha_1, \dots, \alpha_k \in \text{Prg}$ and $\psi \in \text{Fml}$. If $M, w \Vdash \varphi$ then there exists a witness chain for (φ, ψ, M, w) .*

From now on, let Γ_y denote the set of formulae of a node $y \in T$. We say that a finite set of formulae Γ is satisfiable iff $\bigwedge_{\varphi \in \Gamma} \varphi$ is satisfiable.

Lemma 5.6 *Let $x \in T$ with $\text{BD}_x = \emptyset$ and principal formula $\varphi \in \text{Fml}\langle \mathcal{A} \rangle$ of the form $\varphi = \langle \alpha_1 \rangle \dots \langle \alpha_k \rangle \psi$ for some $k > 0$ and $\alpha_1, \dots, \alpha_k \in \text{Prg}$ and $\psi \in \text{Fml} \setminus \text{Fml}\langle \rangle$. Let $M = (W, R, V)$ be a model and $w \in W$ a world such that (M, w) satisfies Γ_x . Furthermore let $\sigma = (w_0, \psi_0), \dots, (w_n, \psi_n)$ be a witness chain for (φ, ψ, M, w) . Then there exists a finite path $\pi = z_0, z_1, \dots, z_m$ in T with the following properties:*

- (i) $m \leq n$, $z_0 = x$, $\text{BD}_{z_m} = \emptyset$, and the only state (if any) is z_m
- (ii) $w_i = w$, $\psi_i \in z_i$, and (M, w) satisfies Γ_{z_i} for all $0 \leq i \leq m$
- (iii) $\psi_i \in \text{Fml}\langle \mathcal{A} \rangle$ is the principal formula of z_i for all $0 \leq i \leq m - 1$
- (iv) $\psi_m = \psi$ or $\psi_m = \langle a \rangle \chi$ for some $a \in \text{APrg}$ and $\chi \in \text{Fml}$.

Proof. We inductively construct π starting with $z_0 = x$, such that the following invariant holds:

- (\sharp) $m < n$ and for all $0 \leq i \leq m$: $w_i = w$ and (M, w) satisfies Γ_{z_i} and $\psi_i \in \text{Fml}\langle \mathcal{A} \rangle$ is the principal formula of z_i .

Note that (\sharp) holds for the initial path $\pi = z_0$. Also note that if π fulfils (\sharp) then no node on π can be a state and $\psi_i \in z_i$ for all $0 \leq i \leq m$.

Step 2 Let z_m be the last node of π . It cannot be an *id*-node because it is satisfiable, nor a $\langle * \rangle_2$ -node for the following reason: Assume that z_m were a $\langle * \rangle_2$ -node. Then $\psi_m \in \text{BD}_{z_m}$ due to the $\langle * \rangle_2$ -rule and there must be an ancestor node z of z_m in T which inserted ψ_m into the BD of its child such that ψ_m is contained in the BD of all nodes between z (exclusive) and z_m (inclusive). As $\text{BD}_{z_0} = \emptyset$ by assumption, the node z must lie on π , i.e. $z = z_{m'}$ for some $m' < m$. Due to the tableau rules and the fact that z inserted ψ_m , the node z must be a $\langle * \rangle_1$ -node with principal formula ψ_m ; but that – together with (\sharp) – entails $(w_{m'}, \psi_{m'}) = (w, \psi_m) = (w_m, \psi_m)$ which is not possible because σ is a witness chain. Hence z_m is not a $\langle * \rangle_2$ -node.

Let z_{m+1} be the child of z_m where ψ_m is decomposed into ψ_{m+1} . Such a child must exist because we have $m < n$ and $\psi_m \rightsquigarrow \psi_{m+1}$ due to the definition of the witness chain σ and the fact that $\psi_m \in \text{Fml}\langle \mathcal{A} \rangle$. The same reasoning also gives us $w = w_m = w_{m+1}$ and $M, w \Vdash \psi_{m+1}$. Moreover, the set Γ_{z_m} is satisfied by (M, w) by (\sharp) and $\Gamma_{z_{m+1}} = (\Gamma_{z_m} \setminus \{\psi_m\}) \cup \{\psi_{m+1}\}$ by construction of the tableau T . Hence the set $\Gamma_{z_{m+1}}$ is satisfied by (M, w) .

Now we distinguish whether or not ψ_{m+1} is a $\langle \mathcal{A} \rangle$ -formula.

If ψ_{m+1} is a $\langle \mathcal{A} \rangle$ -formula, it must be the principal formula of z_{m+1} due to the tableau rules and the fact that we have $\psi_m \in \text{Fml}\langle \mathcal{A} \rangle$. Moreover, we have $m+1 < n$ because $\psi_{m+1} \neq \psi = \psi_n$ and $\psi \notin \text{Fml}\langle \rangle$. Thus our invariant (\sharp) for π extended by ψ_{m+1} still holds and we repeat Step 2.

If ψ_{m+1} is not a $\langle \mathcal{A} \rangle$ -formula, we have $\text{BD}_{z_{m+1}} = \emptyset$ due to the tableau rules and the fact that $\psi_m \in \text{Fml}\langle \mathcal{A} \rangle$. Furthermore, we have $\psi_m = \psi$ or $\psi_m = \langle a \rangle \chi$ for some $a \in \text{APrg}$ and $\chi \in \text{Fml}$ because σ is a witness chain. Thus π extended by ψ_{m+1} fulfils all the required properties of the lemma which concludes the proof in this case.

As σ is finite, Step 2 must terminate after a finite number of repetitions which means that we have found a path π that proves this lemma. \square

Lemma 5.7 *Let $x \in T$ with $\text{BD}_x = \emptyset$ and $M = (W, R, V)$ be a model and $w \in W$ a world such that (M, w) satisfies Γ_x . Then there exists a finite path $\pi = z_0, z_1, \dots, z_n$ in T with the following properties: $z_0 = x$, z_n is the only state on π , and (M, w) satisfies Γ_{z_i} for all $0 \leq i \leq n$.*

Proof. We inductively construct π starting with $z_0 = x$ such that the following invariant holds:

(\dagger) (M, w) satisfies Γ_y for every node y on π and the last node z_i of π has $\text{BD}_{z_i} = \emptyset$.

Note that the initial $\pi = z_0$ fulfils the invariant by assumption.

Step 3 Let z_i be the last node of π . If z_i is a state, we stop and return π . Otherwise, we distinguish two cases: either the principal formula of z_i is not a $\langle \rangle$ -formula or it is a $\langle \rangle$ -formula.

If the principal formula of z_i is not a $\langle \rangle$ -formula, we choose z_{i+1} to be a successor of z_i in T such that (M, w) satisfies $\Gamma_{z_{i+1}}$. The existence of z_{i+1} is guaranteed by Prop. 2.6, the fact that (M, w) satisfies Γ_{z_i} by (\dagger), and the fact that z_i cannot be an *id*-node because z_i is satisfiable nor a $\langle * \rangle_2$ -node because z_i 's principal formula is not a $\langle \rangle$ -formula. As z_i 's principal formula is not a $\langle \rangle$ -formula and $\text{BD}_{z_i} = \emptyset$ by (\dagger), we also have $\text{BD}_{z_{i+1}} = \emptyset$ by a simple inspection of the tableau rules. We then repeat Step 3.

If the principal formula φ of z_i is a $\langle \rangle$ -formula, it is also a $\langle \mu \rangle$ -formula because z_i is not a state. Hence it must be of the form $\varphi = \langle \alpha_1 \rangle \dots \langle \alpha_k \rangle \psi$ for some $k > 0$ and $\alpha_1, \dots, \alpha_k \in \text{Prg}$ and $\psi \in \text{Fml} \setminus \text{Fml}\langle \rangle$. As (M, w) satisfies Γ_{z_i} by (\dagger) and $\varphi \in \Gamma_{z_i}$, we have $M, w \Vdash \varphi$. Thus Prop. 5.5 gives us a sequence $\sigma := (w_0, \psi_0), \dots, (w_n, \psi_n)$ with the properties stated in Prop. 5.5.

Next we apply Lemma 5.6 to z_i and obtain a path τ with the properties of Lemma 5.6. Finally, the new π is obtained from the old π by appending τ – minus the first node z_i which is already the last node of π – to the old π . As (M, w) satisfies Γ_y for all y on τ and the last node y' on τ has $\text{BD}_{y'} = \emptyset$, the new π fulfils (\dagger). We then repeat Step 3.

As T is finite, it is easy to see that Step 3 terminates, meaning that the last node z_n of the finite path π is the only state on π . \square

Lemma 5.8 *For every closed node $x = (\Gamma :: \dots :: \dots)$ in T , the set Γ_x is not satisfiable. In particular, if r is closed then ϕ is not satisfiable.*

Proof. We use well-founded induction on the (strict) descendant relation of T . As T is a finite tree, the descendant relation is clearly well-founded. Thus we can use the following induction hypothesis for every node $x \in T$:

IH: for every descendant y of x , if y is closed then the set Γ_y is not satisfiable.

If a leaf $x \in T$ is closed, it must be an *id*-node as a state with no children is always open. Hence, our theorem follows from the fact that $\{p, \neg p\} \subseteq x$ for some $p \in \text{AFml}$. Note that this can be seen as the base case of the induction as leaves do not have descendants.

If x is a closed α -node then its child must be closed as well so we can apply the induction hypothesis and the claim follows from the fact that – in the sense of

Table 1 – the formulae of the form $\alpha \leftrightarrow \alpha_1 \wedge \alpha_2$ are valid (Prop. 2.6).

If x is a closed β -node then both children are closed as well so we can apply the induction hypothesis and the claim follows from the fact that – in the sense of Table 1 – the formulae of the form $\beta \leftrightarrow \beta_1 \vee \beta_2$ are valid (Prop. 2.6). Note that x cannot be a $\langle * \rangle_2$ -node as it would not be closed in this case.

If x is a closed $\langle \rangle$ -node (*i.e.* a closed state) then it has at least one child and there are three possibilities for why it was marked as closed by the $\langle \rangle$ -rule:

- (1) Some child x_0 of x is closed.
- (2) Some child x_0 of x is barred.
- (3) Some open child x_0 of x with core-formula φ has $\text{uev}_{x_0}(\varphi, \langle \alpha * \rangle \psi) > h := \text{len}(\text{HCr}_x)$ for some $\alpha \in \text{Prg}$ and $\psi \in \text{Fml}$ with $\varphi \in \text{pre}(\langle \alpha * \rangle \psi)$.

Case 1. In the first case, it is not too hard to see that the satisfiability of Γ_x implies the satisfiability of Γ_{x_0} since the $\langle \rangle$ -rule preserves satisfiability from parent to child. By the induction hypothesis, we know that Γ_{x_0} is not satisfiable, therefore Γ_x cannot be satisfiable either.

Case 2. In the second case, we assume that Γ_{x_0} is satisfiable and derive a contradiction. We can then prove the claim as in the first case.

So, for a contradiction, let $M = (W, R, V)$ be a model and $w \in W$ a world such that (M, w) satisfies Γ_{x_0} . As $\text{BD}_{x_0} = \emptyset$ by the $\langle \rangle$ -rule, we can apply Lemma 5.7 which gives us a path π in T with the properties stated in Lemma 5.7. Let y be the last node of π , hence y is a state. It is a descendant of x_0 , therefore the induction hypothesis applies to it. By Lemma 5.7, (M, w) satisfies Γ_y , hence y cannot be closed; but this means that y must be open as states can only be closed or open by the $\langle \rangle$ -rule. It is now easy to see that all nodes on π must also be open due to the construction of the variable stat in the α - and β -rules. But this is a contradiction to the assumption that x_0 , which is the first node on π , is barred.

Case 3. In the third case, we assume that Γ_{x_0} is satisfiable and derive a contradiction. We can then prove the claim as in the first case.

So, for a contradiction, let $M = (W, R, V)$ be a model and $w \in W$ a world such that (M, w) satisfies Γ_{x_0} . In particular, we have $M, w \Vdash \varphi$ by assumption since $\varphi \in x_0$. As $\varphi \in \text{pre}(\langle \alpha * \rangle \psi)$, it is of the form $\varphi = \langle \alpha_1 \rangle \dots \langle \alpha_{k-1} \rangle \langle \alpha * \rangle \psi$ for some $\alpha_1, \dots, \alpha_{k-1} \in \text{Prg}$. Furthermore, let ψ be of the form $\psi = \langle \alpha_{k+1} \rangle \dots \langle \alpha_{k+l} \rangle \psi'$ for some $\alpha_{k+1}, \dots, \alpha_{k+l} \in \text{Prg}$ and $\psi' \in \text{Fml} \setminus \text{Fml}\langle \rangle$. Note that $l = 0$ is possible: in this case we already have $\psi \in \text{Fml} \setminus \text{Fml}\langle \rangle$.

Applying Prop. 5.5 to M and $\varphi = \langle \alpha_1 \rangle \dots \langle \alpha_{k+l} \rangle \psi'$ with $\alpha_k := \alpha *$ gives us a witness chain $\sigma = (w_0, \psi_0), \dots, (w_n, \psi_n)$ for (φ, ψ', M, w) . According to Prop 5.4, there exists an $n' \leq n$ with $\psi_{n'} = \psi = \langle \alpha_{k+1} \rangle \dots \langle \alpha_{k+l} \rangle \psi'$ and $\psi_{n'-1} = \langle \alpha * \rangle \psi$. Our plan is to “walk down” the tableau T – starting from x_0 – in a way that is “consistent” with σ which will lead to a contradiction when we “reach” $\psi_{n'}$.

As $\text{BD}_{x_0} = \emptyset$ by the $\langle \rangle$ -rule, we can apply Lemma 5.6 which gives as a path $\pi_1 = z_0, z_1, \dots, z_m$ in T with the properties stated in Lemma 5.6. We can then apply Lemma 5.7 to z_m which gives us a path π_2 with the properties stated in Lemma 5.7. Let s be the last node of π_2 , hence s is a state. It is a descendant of x_0 , therefore the induction hypothesis applies to it. Thus s cannot be closed because (M, w)

satisfies Γ_s by Lemma 5.7; but this means that s must be open as states can only be closed or open by the $\langle \rangle$ -rule. If we join π_1 and π_2 to obtain π , it is now easy to see that all nodes on π must also be open due to the construction of the variable stat in the α - and β -rules.

By assumption we have $\text{uev}_{x_0}(\varphi, \langle \alpha^* \rangle \psi) > h$. As all nodes on π_1 are open and $\psi_i \in \text{Fml}\langle \emptyset \rangle$ is the principal formula of z_i for all $0 \leq i \leq m-1$, we also have $\text{uev}_{z_i}(\psi_i, \langle \alpha^* \rangle \psi) > h$ for all $0 \leq i \leq m-1$ by definition of the α - and β -rules. We now distinguish whether or not $n' \leq m$.

If $n' \leq m$ then we have $\text{uev}_{z_{n'-1}}(\langle \alpha^* \rangle \psi, \langle \alpha^* \rangle \psi) > h$ as $\psi_{n'-1} = \langle \alpha^* \rangle \psi$; but as $\langle \alpha^* \rangle \psi$ is the principal formula of $z_{n'-1}$, this is only possible if the first child of $z_{n'-1}$, which is $z_{n'}$ as $\psi_{n'} = \psi$ by definition of $\psi_{n'}$, is not open according to the construction of uev in the $\langle * \rangle_1$ -rule. This, however, is a contradiction to the fact that all nodes on π_1 , in particular $z_{n'}$, are open.

If $n' > m$, we must have $\psi_m = \langle a \rangle \chi$ for some $a \in \text{APrg}$ and $\chi \in \text{Fml}$ as $\psi_m = \psi'$ is clearly not possible. Furthermore, we have $\text{uev}_{z_m}(\langle a \rangle \chi, \langle \alpha^* \rangle \psi) > h$ by definition of the α - and β -rules. As ψ_m is the first node on π_2 and all nodes on π_2 are open, we also have $\text{uev}_y(\langle a \rangle \chi, \langle \alpha^* \rangle \psi) > h$ for all nodes y on π_2 by definition of the α - and β -rules. In particular, we have $\text{uev}_s(\langle a \rangle \chi, \langle \alpha^* \rangle \psi) > h$. Let x_1 be the (possibly virtual) successor of $\langle a \rangle \chi \in s$, that contains $\psi_{m+1} = \chi$. Then $\sigma' := (w_{m+1}, \psi_{m+1}), \dots, (w_n, \psi_n)$ is clearly a witness chain for $(\psi_{m+1}, \psi', M, w_{m+1})$ which is strictly shorter than σ and still contains $\psi_{n'}$ and $\psi_{n'-1}$. Note that $n' > m+1$ as $\psi_m = \langle a \rangle \chi \neq \langle \alpha^* \rangle \psi$. Additionally, we make the following two claims:

- (1) $\text{uev}_{x_1}(\psi_{m+1}, \langle \alpha^* \rangle \psi) > h$ and x_1 is a descendant of x_0 (*i.e.* the induction hypothesis holds in the subtree rooted at x_1).
- (2) (M, w_{m+1}) satisfies Γ_{x_1} .

Before we prove the two claims, we show their consequences: Basically, the two claims and the properties of σ' allow us to inductively repeat the proof for x_1 , w_{m+1} , ψ_{m+1} , and σ' instead of x_0 , w , φ , and σ , respectively. As σ' is strictly shorter than σ , this is possible only a finite number of times. Hence we must eventually end up in the case “ $n' \leq m$ ” of the proof that yields a contradiction. Therefore the only thing left is to show that the two claims hold.

Claim 1. We distinguish whether x_1 is a virtual successor of $\langle a \rangle \chi \in s$ or not.

If x_1 is not virtual, that is it is a child of s in T , it is obviously a descendant of x_0 as every node – in particular s – on π is a descendant of x_0 . Furthermore, it follows directly from $\text{uev}_s(\langle a \rangle \chi, \langle \alpha^* \rangle \psi) > h$ and $\psi_{m+1} = \chi$ and the definition of the $\langle \rangle$ -rule that $\text{uev}_{x_1}(\psi_{m+1}, \langle \alpha^* \rangle \psi) > h$.

If x_1 is a virtual successor, a glance at the definition of uev_s in the $\langle \rangle$ -rule reveals that x_1 must lie on the path from x_0 to s (it could be x_0) as we have $\text{uev}_s(\langle a \rangle \chi, \langle \alpha^* \rangle \psi) > h$ and $h = \text{len}(\text{HCr}_x)$. Thus x_1 is a descendant of x and has $\text{uev}_{x_0}(\chi, \langle \alpha^* \rangle \psi) > h$ as we have already established this on our way from x_0 down to s .

Claim 2. By definition of the $\langle \rangle$ -rule, Γ_{x_1} is of the form $\psi_{m+1} \cup \Delta$ where $[a]\Delta \subseteq \Gamma_s$. We know $M, w_{m+1} \Vdash \psi_{m+1}$ because of the properties of σ . We also know that (M, w_m) in particular satisfies $[a]\Delta$ since we have established that $\Gamma_s \supseteq [a]\Delta$ is satisfied by (M, w) and $w = w_m$. As w_{m+1} is a successor world of w (*i.e.* $w R_a w_{m+1}$),

this implies that (M, w_{m+1}) satisfies Δ , and hence Γ_{x_1} . \square

Theorem 4.8 *If the root $r \in T$ is not open then ϕ is not satisfiable.*

Proof. If r is closed, the claim follows directly from Lemma 5.8. If r is barred, we assume that Γ_{x_0} is satisfiable and derive a contradiction.

So, for a contradiction, let $M = (W, R, V)$ be a model and $w \in W$ a world such that (M, w) satisfies $\Gamma_r = \phi$. As $\text{BD}_r = \emptyset$ by construction of T , we can apply Lemma 5.7 which gives us a path π with the properties stated in Lemma 5.7. Let y be the last node of π , hence y is a state. It cannot be closed because of Lemma 5.8 and the fact that (M, w) satisfies Γ_y ; but this means that y must be open as states can only be closed or open by construction. It is easy to see that all nodes on π must also be open due to the construction of the variable stat in the α - and β -rules. But this is a contradiction to the assumption that r , which is the first node on π , is barred. \square