A Proof System for Finite Trees

Patrick Blackburn,[1] Wilfried Meyer-Viol[2] and Maarten de Rijke[3]

[1] Computerlinguistik, Universität des Saarlandes, Postfach 1150,
D-66041 Saarbrücken, Germany
patrick@coli.uni-sb.de

[2] Dept. of Computing, Imperial College, 180 Queen's Gate,
London SW7 2BZ, England
wm3@doc.ic.ac.uk

[3] Dept. of Computer Science, University of Warwick,
Coventry CV4 7AL, England
mdr@dcs.warwick.ac.uk

Abstract. In this paper we introduce a description language for finite trees. Although we briefly note some of its intended applications, the main goal of the paper is to provide it with a sound and complete proof system. We do so using standard axioms from *modal provability logic* and *modal logics of programs*, and prove completeness by extending techniques due to Van Benthem and Meyer-Viol (1994) and Blackburn and Meyer-Viol (1994). We conclude with a proof of the EXPTIME-completeness of the satisfiability problem, and a discussion of issues related to complexity and theorem proving.

1 Introduction

In this paper we introduce a modal language for describing the internal structure of trees, provide it with an axiom system which we prove to be complete with respect to the class of all finite trees, and prove the decidability and EXPTIME-completeness of its satisfiability problem. But before getting down to the technicalities, some motivation.

In many applications, finite trees are the fundamental data structure. Moreover, in many of these applications one wishes to specify how the nodes *within a single tree* relate to each other; that is, it is often the *internal* perspective that is fundamental. By way of contrast, most work on logics of trees in the computer science literature takes an *external* perspective on tree structure. For example, in the work of Courcelle (1985) and Maher (1988), variables range over entire trees. This is a natural choice for work on the semantics of programming languages, but unsuitable for the applications mentioned below. And although the internal perspective on trees has been explored in the logical literature (the classic example is Rabin's (1969) monadic second order theory SnS), such explorations have usually been for extremely powerful languages. It is interesting to explore (modal) fragments of these systems, and that is the purpose of the present paper.

Although the work that follows is concerned solely with technical issues, the reader may find it helpful to consider the sort of applications we have in mind.

One has already arisen in theoretical and computational linguistics. In contemporary linguistics, grammars are often considered to be a set of *constraints* (i.e. axioms) which grammatical structures must satisfy. To specify such grammars, it is crucial to have the ability to specify how tree nodes are related to each other and what properties they must possess. Moreover, it is desirable that such specifications be given in a simple, machine implementable system. A substantial body of work already exists which models the most commonly encountered grammatical formalisms using internal logics of trees: we draw the reader's attention to Backofen *et al.* (1995), Blackburn *et al.* (1993, 1994, 1995), Kracht (1993, 1995), and Rogers (1996). Modal logics of the type considered here have been shown to provide an appropriate level of expressivity for this application.

Another possible application is the formal treatment of corrections in graphical user interfaces. Many competing 'undo mechanisms' have been proposed, differing mainly in the way they allow users to jump through the histories of their actions, and in the way they perceive these histories. In multi-user applications where several agents submit commands concurrently such histories are finite trees, and the complexities of the possible action sequences call for simple, yet expressive description languages (see Berlage (1994)). Examination of the literature suggests that modal languages may be an appropriate modeling tool here as well.

2 The Language $\mathcal{L}$

$\mathcal{L}$ is a propositional modal language with eight modalities: $\langle l \rangle$, $\langle r \rangle$, $\langle u \rangle$ and $\langle d \rangle$ explore the left-sister, right-sister, mother-of and daughter-of relations, while $\langle l+ \rangle$, $\langle r+ \rangle$, $\langle u+ \rangle$ and $\langle d+ \rangle$ explore their transitive closures. The formal definition of $\mathcal{L}$'s syntax is as follows. We suppose we have fixed a non-empty, finite or countably infinite, set of atomic symbols A whose elements are typically denoted by p.

$$\phi ::= p \mid \bot \mid \top \mid \neg\phi \mid \phi \wedge \phi \mid \langle x \rangle \phi \mid \langle x+ \rangle \phi$$
$$x ::= l \mid r \mid u \mid d.$$

We sometimes write $\mathcal{L}(\mathrm{A})$ to emphasize the dependence on A. We employ the usual boolean abbreviations.

We interpret $\mathcal{L}(\mathrm{A})$ on *finite ordered trees* whose nodes are *labeled* with symbols drawn from A. We assume that the reader is familiar with finite trees and such concepts as 'daughter-of', 'mother-of', 'sister-of', 'root-node', 'terminal-node', and so on. If a node has no sister to the immediate right we call it a last node, and if it has no sister to the immediate left we call it a first node. Note that the root node is both first and last. A labeling of a finite tree associates a subset of A with each tree node.

Formally, we present finite ordered trees as tuples $\mathbf{T} = (T, R_l, R_r, R_u, R_d)$. Here T is the set of tree nodes and R_l, R_r, R_u and R_d are the left-sister, right-sister, mother-of and daughter-of relations respectively. A pair $(\mathbf{T}, V)$, where $\mathbf{T}$

is a finite tree and $V : A \longrightarrow Pow(T)$, is called a *model*, and we say that V is a *labeling function* or a *valuation*. Let $(R_x)^+$ denote the transitive closure of R_x. Then we interpret $\mathcal{L}(A)$ on models as follows:

Definition 1 (Truth). For any model $\mathbf{M}$ $(= (T, R_l, R_r, R_u, R_d, V))$ define:

$$
\begin{array}{rcl}
\mathbf{M}, t \models p & \text{iff} & p \in V(t) \text{ for all } p \in \mathrm{A} \\
\mathbf{M}, t \models \neg\phi & \text{iff} & \mathbf{M}, t \not\models \phi \\
\mathbf{M}, t \models \phi \wedge \psi & \text{iff} & \mathbf{M}, t \models \phi \text{ and } \mathbf{M}, t \models \psi \\
\mathbf{M}, t \models \langle x \rangle \phi & \text{iff} & \exists t' \, (t R_x t' \text{ and } \mathbf{M}, t' \models \phi), \text{ where } x \in \{l, r, u, d\} \\
\mathbf{M}, t \models \langle x{+} \rangle \phi & \text{iff} & \exists t' \, (t (R_x)^+ t' \text{ and } \mathbf{M}, t' \models \phi), \text{ where } x \in \{l, r, u, d\}.
\end{array}
$$

If $\mathbf{M}, t \models \phi$, then we say ϕ is *satisfied* in $\mathbf{M}$ at t. For any formula ϕ, if there is a model $\mathbf{M}$ and a node t in $\mathbf{M}$ such that $\mathbf{M}, t \models \phi$, then we say that ϕ is *satisfiable*. If ϕ is true at all nodes in a model $\mathbf{M}$ then we say it is *valid in the model* $\mathbf{M}$. If a formula ϕ is valid in all models then we say it is *valid* and write $\models \phi$.

The following defined operators will prove useful. First we define *duals* of the basic operators: $[x]\phi := \neg\langle x \rangle \neg\phi$ and $[x{+}]\phi := \neg\langle x{+} \rangle \neg\phi$, for all $x \in \{l, r, u, d\}$. We also define operators for talking about the *reflexive* transitive closure of four basic relations: $\langle x{*} \rangle \phi := \phi \vee \langle x{+} \rangle \phi$ and $[x{*}]\phi := \neg\langle x{*} \rangle \neg\phi$, for all $x \in \{l, r, u, d\}$. Next we define the following constants: *first* $:= [l]\bot$, *last* $:= [r]\bot$, *start* $:= [u]\bot$ and *term* $:= [d]\bot$. Note that *first*, *last*, *start* and *term* are constants true only at left nodes, right nodes, the root node, and terminal nodes, respectively.

3 A Proof System for $\mathcal{L}$

We now introduce a logic called LOFT (Logic Of Finite Trees). LOFT is the smallest set of $\mathcal{L}$ formulas that (a) contains all tautologies, (b) contains all instances of the axiom schemas given below, (c) is closed under modus ponens (if ϕ and $\phi \rightarrow \psi$ belong to LOFT then so does ψ), and (d) is closed under generalisation (if ϕ belongs to LOFT then so do $[l]\phi$, $[r]\phi$, $[u]\phi$, $[d]\phi$, $[l{+}]\phi$, $[r{+}]\phi$, $[u{+}]\phi$, and $[d{+}]\phi$). Note that this is a purely *syntactical* description of LOFT. The completeness theorem proved below shows that LOFT really does deserve its name: LOFT consists of precisely the formulas of $\mathcal{L}$ valid on finite trees.

It remains to specify the axiom schemas. These fall naturally into four groups. The first group is the simplest. Schema 1 is the fundamental schema of normal modal logic. Schemas 2l, 2r, 2u and 2d reflect the fact that both R_l and R_r, and R_u and R_d, are converse pairs of relations (these schemas are basic axioms of temporal logic), while schema 3 (familiar from modal logic) reflects the fact that R_l, R_r and R_u are partial functions.

1. $[x](\phi \rightarrow \psi) \rightarrow ([x]\phi \rightarrow [x]\psi)$ $\qquad (x \in \{l, r, u, d\})$

2l. $\phi \rightarrow [l]\langle r \rangle \phi$

2r. $\phi \rightarrow [r]\langle l\rangle\phi$
2u. $\phi \rightarrow [u]\langle d\rangle\phi$
2d. $\phi \rightarrow [d]\langle u\rangle\phi$
3. $\langle x\rangle\phi \rightarrow [x]\phi$ $\quad (x \in \{l, r, u\})$

The second group are (irreflexive analogs of) the Segerberg schemas used in modal logics of programs; they reflect the fact that the operators $[l+]$, $[r+]$, $[u+]$ and $[d+]$ make use of the transitive closure of the relations for $[l]$, $[r]$, $[u]$ and $[d]$ respectively.

4. $[x+]\phi \leftrightarrow [x][x*]\phi$ $\quad (x \in \{l, r, u, d\})$
5. $[x+](\phi \rightarrow [x]\phi) \rightarrow ([x]\phi \rightarrow [x+]\phi)$ $\quad (x \in \{l, r, u, d\})$.

The third group reflects the fact that we are working only with *finite trees.* Schema 7 (Löb's schema) is the crucial one. It is the key schema of *modal provability logic* and expresses a second-order fact about finite trees: the transitive closure of the 'daughter-of' relation, and of the 'to-the-right-of' relation, are both converse well founded.

6. $\langle u*\rangle \mathit{start} \wedge \langle d*\rangle \mathit{term} \wedge \langle l*\rangle \mathit{first} \wedge \langle r*\rangle \mathit{last}$
7. $[x+]([x+]\phi \rightarrow \phi) \rightarrow [x+]\phi$ $\quad (x \in \{r, d\})$.

The fourth group reflects the links between the vertically and horizontally scanning modalities.

8. $\langle d\rangle\phi \rightarrow [d](\mathit{first} \rightarrow \langle r*\rangle\phi)$
9. $\langle d\rangle\phi \rightarrow \langle d\rangle \mathit{first} \wedge \langle d\rangle \mathit{last}$
10. $\mathit{start} \rightarrow \mathit{first} \wedge \mathit{last}$.

4 Proving Completeness

In this section we prove the completeness of LOFT. (Proving that LOFT is sound with respect to finite trees is straightforward, though readers new to modal logic may find it helpful to refer to Goldblatt (1992) or Smoryński (1985) for further discussion of the Segerberg and Löb schemas.) Our proof uses ideas from provability logic and dynamic logic, and extends techniques used by Van Benthem and Meyer-Viol (1994) and Blackburn and Meyer-Viol (1994). The work falls into three phases. First, we show that LOFT is complete with respect to a certain class of finite pseudo-models. Although pseudo-models are not trees, they embody a great deal of useful information about LOFT, and in the second phase we show how to make use of this: we prove a sufficient condition (the *truth lemma for induced models*) under which pseudo-models induce genuine models on finite trees. In the third stage, the heart of the proof, we show that there is a (finite) inductive method for building induced models.

4.1 Preliminaries

The first notion we need is that of a *closure* of sentences. Recall that a set of formulas Σ is *closed under subformulas* iff for all $\phi \in \Sigma$, every subformula of ϕ is in Σ. Following Fischer and Ladner (1979) we define sets of formulas that are closed under a little more structure than simply subformulahood.

Definition 2 (Closures). Let Σ be a set of formulas. $Cl(\Sigma)$ is the smallest set of sentences containing Σ that is closed under subformulas and satisfies the following additional constraints.

1. If $\langle x+\rangle\phi \in Cl(\Sigma)$, then $\langle x\rangle\phi \in Cl(\Sigma)$, where $x \in \{l, r, u, d\}$.
2. If $\langle x+\rangle\phi \in Cl(\Sigma)$, then $\langle x\rangle\langle x+\rangle\phi \in Cl(\Sigma)$, where $x \in \{l, r, u, d\}$.
3. If $\langle d\rangle\phi \in Cl(\Sigma)$, then $\langle r*\rangle\phi \in Cl(\Sigma)$.
4. $\langle x\rangle\top \in Cl(\Sigma)$ for $x \in \{l, r, u, d\}$.
5. $\langle d\rangle$*first*, $\langle d\rangle$*last* $\in Cl(\Sigma)$.
6. $\langle l*\rangle$*first*, $\langle r*\rangle$*last*, $\langle u*\rangle$*start*, $\langle d*\rangle$*term* $\in Cl(\Sigma)$.
7. If $\phi \in Cl(\Sigma)$ and ϕ is not of the form $\neg\psi$, then $\neg\phi \in Cl(\Sigma)$.

$Cl(\Sigma)$ is called the *closure* of Σ. Observe that for every $\phi \in Cl(\Sigma)$ there is a $\psi \in Cl(\Sigma)$ such that ψ is equivalent to $\neg\phi$; we will often pretend that for every $\phi \in Cl(\Sigma)$, $\neg\phi$ is also in $Cl(\Sigma)$.

Lemma 3. *Let Σ be a finite set of formulas. Then $Cl(\Sigma)$ is finite too.*

Definition 4 (Atoms). If Σ is a set of formulas, then $At(\Sigma)$ consists of all the maximal consistent subsets of $Cl(\Sigma)$. In other words, $At(\Sigma)$ consists of all sets $\mathcal{A} \subseteq Cl(\Sigma)$ such that $\mathcal{A}$ is consistent, and if $\mathcal{B}$ is consistent and $\mathcal{A} \subseteq \mathcal{B} \subseteq Cl(\Sigma)$, then $\mathcal{A} = \mathcal{B}$. The elements of $At(\Sigma)$ are called *atoms* (*over* Σ).

Lemma 5 (Atoms exist). *If $\phi \in Cl(\Sigma)$ and ϕ is consistent, then there exists an atom $\mathcal{A} \in At(\Sigma)$ such that $\phi \in \mathcal{A}$.*

Proof. Use the usual Lindenbaum technique together with the observation that $At(\Sigma) = \{\mathcal{M} \cap Cl(\Sigma) \mid \mathcal{M}$ is a maximal consistent set in the usual sense $\}$. $\dashv$

Lemma 6 (Properties of atoms). *Let Σ be a set of formulas and $\mathcal{A} \in Cl(\Sigma)$.*

1. *If $\phi \in Cl(\Sigma)$, then $\phi \in \mathcal{A}$ iff $\neg\phi \notin \mathcal{A}$.*
2. *If $\phi \wedge \psi \in Cl(\Sigma)$ then $\phi \wedge \psi \in \mathcal{A}$ iff $\phi \in \mathcal{A}$ and $\psi \in \mathcal{A}$*
3. *If $\phi \rightarrow \psi \in Cl(\Sigma)$, then $\phi \rightarrow \psi$ and $\phi \in \mathcal{A}$ implies $\psi \in \mathcal{A}$.*
4. *If $\langle x+\rangle\phi \in Cl(\Sigma)$, then $\langle x+\rangle\phi \in \mathcal{A}$ iff $\langle x\rangle\phi \in \mathcal{A}$ or $\langle x\rangle\langle x+\rangle\phi \in \mathcal{A}$, where $x \in \{l, r, u, d\}$.*
5. *$\langle u*\rangle$start, $\langle d*\rangle$term, $\langle l*\rangle$first, $\langle r*\rangle$last, $\top \in \mathcal{A}$.*

Lemma 7. *Suppose $At(\Sigma) = \{\mathcal{A}_1, \ldots, \mathcal{A}_n\}$. Then $\vdash \bigwedge \mathcal{A}_1 \vee \cdots \vee \bigwedge \mathcal{A}_n$.*

Proof. Use the propositional tautology $\phi \leftrightarrow ((\phi \wedge \psi) \vee (\phi \wedge \neg\psi))$. $\dashv$

4.2 Pseudo-models

In this subsection we define a collection of finite pseudo-models with the following property: if ϕ is a consistent formula, then there is a (finite) pseudo-model that satisfies ϕ. Although this result is of interest in its own right (as we shall see at the end of the paper), of equal importance are the definitions and results we encounter along the way, for these will be used throughout.

Definition 8 (Canonical relations). Let $\mathcal{A}, \mathcal{B} \in At(\Sigma)$. For each $x \in \{l, r, u, d, l+, r+, u+, d+\}$ we define the *canonical relations* S_x on $At(\Sigma)$ as follows:

$$\mathcal{A} S_x \mathcal{B} \text{ iff } \bigwedge \mathcal{A} \wedge \langle x \rangle \bigwedge \mathcal{B} \text{ is consistent.}$$

Lemma 9. *Let $\mathcal{A}$ be an atom, $x \in \{l, r, u, d, l+, r+, u+, d+\}$, and $\psi \in Cl(\Sigma)$. If $\bigwedge \mathcal{A} \wedge \langle x \rangle \psi$ is consistent, then there is an atom $\mathcal{B}$ over Σ such that $\psi \in \mathcal{B}$ and $\mathcal{A} S_x \mathcal{B}$.*

Proof. Suppose that $\bigwedge \mathcal{A} \wedge \langle x \rangle \psi$ is consistent. We show how to construct the required atom $\mathcal{B}$ by 'forcing a choice' between the formulas in $At(\Sigma)$. For all formulas χ, $\psi \leftrightarrow (\psi \wedge \chi) \vee (\psi \wedge \neg\chi)$ is a propositional tautology, hence by simple modal reasoning: $\vdash \langle x \rangle \psi \leftrightarrow \langle x \rangle (\psi \wedge \chi) \vee \langle x \rangle (\psi \wedge \neg \chi)$. Hence by propositional logic, either $\bigwedge \mathcal{A} \wedge \langle x \rangle (\psi \wedge \chi)$ or $\bigwedge \mathcal{A} \wedge \langle x \rangle (\psi \wedge \neg\chi)$ is consistent. This observation enables us to construct the desired $\mathcal{B}$ 'behind the modality' $\langle x \rangle$ by working through all the formulas in $Cl(\Sigma)$. ⊣

Lemma 10. *Let $\mathcal{A}$ be an atom, $x \in \{l, r, u, d, l+, r+, u+, d+\}$, and $\langle x \rangle \psi \in Cl(\Sigma)$. Then $\langle x \rangle \psi \in \mathcal{A}$ iff there is an atom $\mathcal{B}$ such that $\psi \in \mathcal{A}$ and $\mathcal{A} S_x \mathcal{B}$.*

Proof. For the left to right direction, note that if $\langle x \rangle \psi \in \mathcal{A}$ then $\bigwedge \mathcal{A} \wedge \langle x \rangle \psi$ is consistent, and the result follows by the previous lemma. For the right to left direction note that if such a $\mathcal{B}$ exists, then $\bigwedge \mathcal{A} \wedge \langle x \rangle \bigwedge \mathcal{B}$ is consistent, thus so is $\bigwedge \mathcal{A} \wedge \langle x \rangle \psi$. As $\langle x \rangle \psi \in Cl(\Sigma)$, by maximality it belongs to $\mathcal{A}$. ⊣

Lemma 11. *Let $\mathcal{A}$ and $\mathcal{B}$ be atoms in $Cl(\Sigma)$. Then for all $x \in \{l, r, u, d\}$, if $\mathcal{A} S_{x+} \mathcal{B}$ then $\mathcal{A}(S_x)^+ \mathcal{B}$.*

Proof. Assume that $\mathcal{A} S_{x+} \mathcal{B}$ where x is either l, r u or d. That is, $\bigwedge \mathcal{A} \wedge \langle x+ \rangle \bigwedge \mathcal{B}$ is consistent. Let

$$\sigma := \bigvee \left\{ \bigwedge \mathcal{C} \mid \mathcal{A}(S_x)^+ \mathcal{C} \right\},$$

where $(S_x)^+$ is the transitive closure of S_x. Then $\sigma \wedge \langle x \rangle \neg \sigma$ is inconsistent, for otherwise $\sigma \wedge \langle x \rangle \bigwedge \mathcal{C}'$ would be consistent for at least one $\mathcal{C}'$ *not* reachable from $\mathcal{A}$ in finitely many S_x steps; but then $\bigwedge \mathcal{C} \wedge \langle x \rangle \bigwedge \mathcal{C}'$ would be consistent for at least one $\mathcal{C} \in At(\Sigma)$ with $\mathcal{A}(S_x)^+ \mathcal{C}$. Hence $\mathcal{A}(S_x)^+ \mathcal{C}'$ — a contradiction. Therefore

$$\begin{aligned} \vdash \sigma \wedge \langle x \rangle \neg \sigma \to \bot &\Rightarrow \vdash \sigma \to [x] \sigma \\ &\Rightarrow \vdash [x+](\sigma \to [x]\sigma), \text{ by generalization} \\ &\Rightarrow \vdash [x]\sigma \to [x+]\sigma, \text{ by axiom 5.} \end{aligned}$$

By simple modal reasoning, we have $\vdash \bigwedge \mathcal{A} \to [x]\sigma$, so $\vdash \bigwedge \mathcal{A} \to [x+]\sigma$. Then, as $\bigwedge \mathcal{A} \wedge \langle x+\rangle \bigwedge \mathcal{B}$ was assumed consistent, $\langle x+\rangle(\bigwedge \mathcal{B} \wedge \sigma)$ is consistent as well, and so $\bigwedge \mathcal{B} \wedge \sigma$ must be consistent. By the definition of σ this means that $\bigwedge \mathcal{B} \wedge \bigwedge \mathcal{C}$ is consistent for at least one atom $\mathcal{C}$ with $\mathcal{A}(S_x)^+\mathcal{C}$. By maximality $\mathcal{B} = \mathcal{C}$, and so $\mathcal{A}(S_x)^+\mathcal{B}$, as required. $\dashv$

Lemma 12. *Let $\mathcal{A} \in At(\Sigma)$, and let $x \in \{l, r, u, d\}$. Assume that $\langle x+\rangle\psi \in Cl(\Sigma)$. Then $\langle x+\rangle\psi \in \mathcal{A}$ iff for some $\mathcal{B} \in At(\Sigma)$ we have $\mathcal{A}(S_x)^+\mathcal{B}$ and $\psi \in \mathcal{B}$.*

Proof. Suppose $\langle x+\rangle\psi \in \mathcal{A}$. By Lemma 10 there is an atom $\mathcal{B}$ such that $\mathcal{A}S_{x+}\mathcal{B}$, hence by Lemma 11, $\mathcal{A}(S_x)^+\mathcal{B}$.

Conversely, suppose that $\mathcal{A} = \mathcal{A}_1 S_x \cdots S_x \mathcal{A}_k = \mathcal{B}$ and $\psi \in \mathcal{B}$. We show the desired result by induction on k. If $k = 1$, then $\mathcal{A}S_x\mathcal{B}$. We need to show $\langle x+\rangle\psi \in \mathcal{A}$. As $\mathcal{A}S_x\mathcal{B}$, $\langle x\rangle\psi \in \mathcal{A}$. By axiom 4, $\langle x\rangle\psi \to \langle x+\rangle\psi$, hence by maximality $\langle x+\rangle\psi \in \mathcal{A}$. For the induction step, assume that $k > 1$ and $\mathcal{A} = \mathcal{A}_1 S_x \mathcal{A}_2 \cdots S_x \mathcal{A}_k = \mathcal{B}$. By the induction hypothesis, $\langle x+\rangle\psi \in \mathcal{A}_2$. It follows that $\langle x\rangle\langle x+\rangle\psi \in \mathcal{A}_1 = \mathcal{A}$ (this uses the second closure condition on $Cl(\Sigma)$), and hence $\langle x+\rangle\psi \in \mathcal{A}$ by axiom 4. $\dashv$

Definition 13. Let a finite set of formulas Σ be given. Define the *canonical pseudo-model over* $At(\Sigma)$ to be the structure

$$\mathbf{P} = (At(\Sigma), S_l, S_r, S_u, S_d, (S_l)^+, (S_r)^+, (S_u)^+, (S_d)^+, V),$$

where $V(p) = \{\mathcal{A} \mid p \in \mathcal{A}\}$. We interpret $\mathcal{L}$ in the obvious way on pseudo-models.

Lemma 14 (Truth lemma for pseudo-models). *Let $\mathbf{P}$ be the pseudo-model over $At(\Sigma)$. For all $\mathcal{A} \in At(\Sigma)$ and all $\psi \in Cl(\Sigma)$, $\psi \in \mathcal{A}$ iff $\mathbf{P}, \mathcal{A} \models \psi$.*

Proof. By induction on the structure of ψ. The base case is clear and the boolean cases are trivial. It remains to examine the argument for the modalities.

First, let $x \in \{l, r, u, d\}$ and suppose that $\mathbf{P}, \mathcal{A} \models \langle x\rangle\psi$. This happens iff there is an atom $\mathcal{B}$ such that $\mathcal{A}S_x\mathcal{B}$ and $\mathbf{P}, \mathcal{B} \models \psi$. By the inductive hypothesis, this happens iff there is an atom $\mathcal{B}$ such that $\mathcal{A}S_x\mathcal{B}$ and $\psi \in \mathcal{B}$. By Lemma 10, this happens iff $\langle x\rangle\psi \in \mathcal{A}$, the desired result.

Next, let $x \in \{l+, r+, u+, d+\}$, and suppose that $\mathbf{P}, \mathcal{A} \models \langle x\rangle\psi$. This happens iff there is an atom $\mathcal{B}$ such that $\mathcal{A}(S_x)^+\mathcal{B}$ and $\mathbf{P}, \mathcal{B} \models \psi$. By the inductive hypothesis, this happens iff there is an atom $\mathcal{B}$ such that $\mathcal{A}(S_x)^+\mathcal{B}$ and $\psi \in \mathcal{B}$. By Lemma 12, this happens iff $\langle x\rangle\psi \in \mathcal{A}$, the desired result. $\dashv$

Theorem 15. LOFT *is complete with respect to the class of finite pseudo-models.*

Proof. Given a LOFT-consistent formula ψ, form the (finite) pseudo-model $\mathbf{P}$ over $At(\{\psi\})$. As ψ is consistent it belongs to some atom $\mathcal{A}$, hence by the above truth lemma $\mathbf{P}, \mathcal{A} \models \psi$. Thus every consistent sentence has a model, and completeness follows. $\dashv$

This gives us a completeness theorem for LOFT. Unfortunately it's not the one we want, since pseudo-models need not be based on finite trees. (The easiest way to see this is to observe that S_l, S_r, and S_u need not be partial functions.) However, as we shall now see, pseudo-models contain all the information needed to induce genuine models on finite trees.

4.3 Induced Models

In this subsection we prove the following result: if the nodes of a finite tree $\mathbf{T}$ are *sensibly decorated* with the atoms from some pseudo-model, then the pseudo-model induces a genuine model on $\mathbf{T}$.

Definition 16. Let $\mathbf{T} = (T, R_l, R_r, R_u, R_d)$ be a finite tree and Σ any finite set of sentences. A *decoration* of $\mathbf{T}$ by $At(\Sigma)$ is a function $h : T \longrightarrow At(\Sigma)$, and the *model induced by the decoration* on $\mathbf{T}$ is the pair $(\mathbf{T}, V)$, where V is the valuation on $\mathbf{T}$ defined by $t \in V(p)$ iff $p \in h(t)$. Suppose that h is a decoration with the following properties:

1. For all t, $t' \in T$, if tR_dt' then $h(t)S_dh(t')$.
2. For all t, $t' \in T$, if tR_rt' then $h(t)S_rh(t')$.
3. For all $t \in T$, if $\langle d\rangle\psi \in h(t)$ then there is a $t' \in T$ such that tR_dt' and $\psi \in h(t')$.
4. *start* $\in h(t)$ iff $t =$ *root*; *term* $\in h(t)$ (respectively: *first* $\in h(t)$, *last* $\in h(t)$) iff t is a terminal node (respectively: iff t is a first node, iff t is a last node).

Then h is called a *sensible decoration* of $\mathbf{T}$. (In short, a sensible decoration is simply a certain kind of order preserving morphism between a finite tree and the pseudo-model over $At(\Sigma)$.)

To prove a truth lemma for induced models, we need some additional facts.

Lemma 17. *Let $\mathcal{A}$, $\mathcal{B} \in At(\Sigma)$. Then $\mathcal{A}S_l\mathcal{B}$ iff $\mathcal{B}S_r\mathcal{A}$, and $\mathcal{A}S_u\mathcal{B}$ iff $\mathcal{B}S_d\mathcal{A}$.*

Proof. This is proved using the temporal logic axioms. We show that $\mathcal{A}S_d\mathcal{B}$ iff $\mathcal{B}S_u\mathcal{A}$; the other case is similar. Let $\mathcal{A}S_d\mathcal{B}$ and suppose for the sake of a contradiction that $\bigwedge\mathcal{B}\wedge\langle u\rangle\bigwedge\mathcal{A}$ is inconsistent. Thus $\vdash \bigwedge\mathcal{B} \rightarrow \neg\langle u\rangle\bigwedge\mathcal{A}$. Hence by generalisation $\vdash [d]\bigwedge\mathcal{B} \rightarrow [d]\neg\langle u\rangle\bigwedge\mathcal{A}$. As $\mathcal{A}S_d\mathcal{B}$, $\bigwedge\mathcal{A}\wedge\langle d\rangle\bigwedge\mathcal{B}$ is consistent, thus by simple modal reasoning, so is $\bigwedge\mathcal{A} \wedge \langle d\rangle\neg\langle u\rangle\bigwedge\mathcal{A}$. But by axiom 2d, $\vdash \bigwedge\mathcal{A} \rightarrow [d]\langle u\rangle\mathcal{A}$, therefore $\langle d\rangle(\langle u\rangle\mathcal{A}\wedge\neg\langle u\rangle\bigwedge\mathcal{A})$ is consistent — a contradiction. We conclude that $\mathcal{B}S_u\mathcal{A}$. A symmetric argument (using axiom 2u) establishes the converse, as required. ⊣

Corollary 18. *Let h be a sensible decoration of $\mathbf{T}$ and $x \in \{l, r, u, d\}$. Then for all nodes t, t' in $\mathbf{T}$, tR_xt' implies $h(t)S_xh(t')$.*

Proof. For r and d this is immediate from the definition of sensible decorations. For l and r it follows from the previous lemma. ⊣

Lemma 19. *For all atoms $\mathcal{A}$ and $\mathcal{B}$ in $Cl(\Sigma)$, and all $x \in \{l, r, u\}$, if $\mathcal{A}S_x\mathcal{B}$ and $\langle x\rangle\psi \in \mathcal{A}$ then $\psi \in \mathcal{B}$.*

Proof. As $\mathcal{A}S_x\mathcal{B}$, $\bigwedge \mathcal{A} \wedge \langle x\rangle \bigwedge \mathcal{B}$ is consistent, hence as $\langle x\rangle\psi \in \mathcal{A}$, $\langle x\rangle\psi \wedge \langle x\rangle \bigwedge \mathcal{B}$ is consistent. It is an easy consequence of axiom 3, the partial functionality axiom, that $\vdash \langle x\rangle\theta \wedge \langle x\rangle\chi \rightarrow \langle x\rangle(\theta \wedge \chi)$; thus it follows that $\langle x\rangle(\psi \wedge \bigwedge \mathcal{B})$ is consistent, and thus so is $\psi \wedge \bigwedge \mathcal{B}$. As $\psi \in Cl(\Sigma)$, by maximality we get $\psi \in \mathcal{B}$. $\dashv$

Lemma 20 (Effects of the constants). *Let $\mathcal{A} \in At(\Sigma)$. Then*

1. *start $\in \mathcal{A}$ iff no formula of the form $\langle u\rangle\phi$ or $\langle u+\rangle\phi$ is in $\mathcal{A}$;*
2. *term $\in \mathcal{A}$ iff no formula of the form $\langle d\rangle\phi$ or $\langle d+\rangle\phi$ is in $\mathcal{A}$;*
3. *first $\in \mathcal{A}$ iff no formula of the form $\langle l\rangle\phi$ or $\langle l+\rangle\phi$ is in $\mathcal{A}$;*
4. *last $\in \mathcal{A}$ iff no formula of the form $\langle r\rangle\phi$ or $\langle r+\rangle\phi$ is in $\mathcal{A}$.*

Proof. For the one step modalities the result is immediate. For the transitive closure modalities, note that by axiom 4, $\vdash \langle x+\rangle\phi \leftrightarrow (\langle x\rangle\phi \vee \langle x\rangle\langle x+\rangle\phi)$. So, assuming that $\langle x+\rangle\phi \in Cl(\Sigma)$, by the first and second closure conditions, we find that $\langle x+\rangle\phi$ is in $\mathcal{A}$ iff either $\langle x\rangle\phi$ or $\langle x\rangle\langle x+\rangle\phi)$ is in $\mathcal{A}$. This observation reduces the transitive closure case to the case for the one step modalities. $\dashv$

Lemma 21. *Let $\mathcal{A}, \mathcal{B} \in At(\Sigma)$, let $x \in \{l, r, u, d\}$, and $\mathcal{A}S_x\mathcal{B}$. If $\langle x+\rangle\psi \in \mathcal{A}$ then either $\psi \in \mathcal{B}$ or $\langle x+\rangle\psi \in \mathcal{B}$.*

Proof. Follows from axiom 4 and the first and second closure conditions. $\dashv$

Lemma 22 (Truth lemma for induced models). *Let h be a sensible decoration of $\mathbf{T}$ and $\mathbf{M} = (\mathbf{T}, V)$ be the model induced by h on $\mathbf{T}$. Then for all nodes t in $\mathbf{T}$, and all $\psi \in Cl(\Sigma)$, $\mathbf{M}, t \models \psi$ iff $\psi \in h(t)$.*

Proof. By induction on the structure of ψ. The base case is clear by definition, and the boolean cases are trivial. It remains to consider the modalities.

First we treat the case for the one step modalities. Suppose $\mathbf{M}, t \models \langle x\rangle\psi$, where $x \in \{l, r, u, d\}$. Then there is a node t' such that tR_xt' and $\mathbf{M}, t' \models \psi$. As h is a sensible decoration, by Corollary 18 $h(t)S_xh(t')$, and by the inductive hypothesis, $\psi \in h(t')$. By Lemma 10, $\langle x\rangle\psi \in h(t)$ as required.

For the converse, suppose $\mathbf{M}, t \not\models \langle x\rangle\psi$. Then *either* $x = u$ and t is the root node (respectively: $x = d$ and t is a terminal node, $x = l$ and t is a first node, $x = r$ and t is a last node) *or* there is at least one node t' such that tR_xt' but for all such nodes $\mathbf{M}, t' \not\models \psi$. Suppose the former. Then by Lemma 20, $\langle x\rangle\psi \notin h(t)$ for any ψ, the required result. So suppose that there is a t' such that tR_xt' but for all such nodes $\mathbf{M}, t' \not\models \psi$. As h is sensible, $h(t)S_xh(t')$ and by the inductive hypothesis $\psi \notin h(t')$. Now, if $x \in \{l, r, u\}$ then by Lemma 19, $\langle x\rangle\psi \notin h(t)$, the required result. On the other hand, if $x = d$ then we also have that $\langle x\rangle\psi \notin h(t)$, as otherwise we would contradict item 3 in the definition of sensible decorations. Either way, we have the required result.

It remains to treat the transitive closure operators. Suppose $\mathbf{M}, t \models \langle x+\rangle\psi$, where $x \in \{l, r, u, d\}$. Then there is a node t' such that $t(R_x)^+t'$; that is, there is a finite sequence of nodes $t = t_1 R_x \cdots R_x t_k = t'$ and $\mathbf{M}, t' \models \psi$. As h is sensible, $h(t) = h(t_1) S_x \cdots S_x h(t_k) = h(t')$, and by the induction hypothesis, $\psi \in h(t')$. Thus by Lemma 12, $\langle x+\rangle\psi \in h(t)$, as required.

Conversely, suppose $\mathbf{M}, t \not\models \langle x+\rangle\psi$. Then for all t' such that tR_x^+t' we have $\mathbf{M}, t' \not\models \psi$, and hence by the inductive hypothesis, $\psi \notin h(t')$. Suppose for the sake of a contradiction that $\langle x+\rangle\psi \in h(t)$. Then by Lemma 20, the constant corresponding to x (that is, *first*, *last*, *start* and *term* for l, r, u and d respectively) does not belong to $h(t)$. As h is a sensible decoration, this means that t has an R_x successor t_1. By Lemma 21, either ψ or $\langle x+\rangle\psi$ belongs to $h(t_1)$, so as $\psi \notin h(t_1)$, $\langle x+\rangle\psi \in h(t_1)$. We are now in the same position with respect to t_1 that we were in with respect to t, and can repeat the argument as many times as we wish, generating a sequence of nodes $tR_x t_1 R_x t_2 \ldots$ such that $\langle x+\rangle\psi \in h(t_i)$ for all i. But as t lives in a finite tree, it only has *finitely* many successors; hence, for some j, $h(t_j)$ must also contain the constant corresponding to x — but then by Lemma 20, it must also contain $\neg\langle x+\rangle\psi \in h(t_i)$. As atoms are consistent this is impossible. We conclude that $\langle x+\rangle\psi \notin h(t)$, the desired result. $\dashv$

4.4 Levels and Ranks

The truth lemma for induced models suggests the following strategy for proving completeness: given a consistent sentence ϕ, simultaneously build by induction a suitable finite tree and sensible decoration, and then use the induced model. This is essentially what we shall do, but there is a problem. We need to build a *finite* tree, so we must guarantee that the inductive construction halts after finitely many steps. It is here that the Löb axioms come into play. Roughly speaking, they enable us to assign to each atom two natural numbers: a vertical 'layer', and a horizontal 'rank'. These have the following property: when generating vertically we can always work with atoms of lower level, and when generating horizontally we can always work with atoms of lower rank. This will enable us to devise a terminating construction method. (The reader is warned, however, that these remarks are only intended to give the basic intuition; as we shall see, the real situation is more complex.)

The basic observation on which these ideas rest is the following:

Lemma 23. *Let $x \in \{r, d\}$. If $\langle x+\rangle\phi$ is consistent then so is $\phi \wedge [x][x*]\neg\phi$.*

Proof. The contrapositive of the Löb axiom is $\langle x+\rangle\phi \to \langle x+\rangle(\phi \wedge [x+]\neg\phi)$. Using the first Segerberg axiom, this can be rewritten as

$$\langle x+\rangle\phi \to \langle x+\rangle(\phi \wedge [x][x*]\neg\phi).$$

Hence, if $\langle x+\rangle\phi$ is consistent, so are $\langle x+\rangle(\phi \wedge [x][x*]\neg\phi)$ and $\phi \wedge [x][x*]\neg\phi$. $\dashv$

The following group of definitions and lemmas build on this to show that the set of atoms is 'vertically well behaved'.

Definition 24. S_Σ is the set of all atoms in $At(\Sigma)$ that contain *start*.

Note that S_Σ is non-empty for any choice of $At(\Sigma)$. To see this, note that by axiom 6, $\langle u{+}\rangle \mathit{start}$ is consistent, hence so is start. By our closure conditions, $\mathit{start} \in Cl(\Sigma)$, hence there is some atom in $At(\Sigma)$ containing start.

Lemma 25. *Suppose $At(\Sigma) \setminus \mathsf{S}_\Sigma$ is non-empty, and let $\mathsf{A} = \{\mathcal{A}_1, \ldots, \mathcal{A}_n\}$ and $\mathsf{B} = \{\mathcal{B}_1, \ldots, \mathcal{B}_m\}$ be disjoint non-empty sets of atoms with $\mathsf{A} \cup \mathsf{B} = At(\Sigma) \setminus \mathsf{S}_\Sigma$. Then for some $\mathcal{A} \in \mathsf{A}$,*

$$\bigwedge \mathcal{A} \wedge [d][d{+}](\bigwedge \mathcal{B}_1 \vee \cdots \vee \bigwedge \mathcal{B}_m)$$

is consistent.

Proof. Let $\mathbf{A}$ be $\bigwedge \mathcal{A}_1 \vee \cdots \vee \bigwedge \mathcal{A}_n$. As any atom is consistent, $\mathbf{A}$ is consistent, and as $\neg \mathit{start}$ belongs to every atom in A, $\langle d{+}\rangle \mathbf{A}$ is consistent. Let S_Σ be enumerated as $\mathcal{S}_1, \ldots, \mathcal{S}_l$ (this is possible, for S_Σ is finite) and let $\mathbf{S}$ be $\bigwedge \mathcal{S}_1 \vee \cdots \vee \bigwedge \mathcal{S}_l$. As $\langle d{+}\rangle \mathbf{A}$ is consistent, so is $\langle d{+}\rangle(\mathbf{S} \vee \mathbf{A})$, hence by the previous lemma

$$(\mathbf{S} \vee \mathbf{A}) \wedge [d][d*]\neg(\mathbf{S} \vee \mathbf{A})$$

is consistent too.

Let $\mathbf{B}$ be $\bigwedge \mathcal{B}_1 \vee \cdots \vee \bigwedge \mathcal{B}_m$. As $\mathsf{S}_\Sigma \cup \mathsf{A} \cup \mathsf{B} = At(\Sigma)$, by Lemma 7, $\vdash \mathbf{S} \vee \mathbf{A} \vee \mathbf{B}$, hence $\vdash \neg(\mathbf{S} \vee \mathbf{A}) \rightarrow \mathbf{B}$. Thus $(\mathbf{S} \vee \mathbf{A}) \wedge [d][d*]\mathbf{B}$ is consistent, hence $\mathbf{A} \wedge [d][d*]\mathbf{B}$ is consistent, hence for some $\mathcal{A} \in \mathsf{A}$, $\mathcal{A} \wedge [d][d*]\mathbf{B}$ is consistent, which yields the desired result. ⊣

Definition 26 (Levels 1). Let $Cl(\Sigma)$ be a closed set such that $At(\Sigma) \setminus \mathsf{S}_\Sigma$ is non-empty. Then the *levels* on $At(\Sigma) \setminus \mathsf{S}_\Sigma$ are defined as follows. L_0 is defined to be $\{\mathcal{A} \in (At(\Sigma) \setminus \mathsf{S}_\Sigma) \mid \mathit{term} \in \mathcal{A}\}$. For $i \geq 0$, V_i is $\bigcup_{0 \leq j \leq i} L_j$, and if $At(\Sigma) \setminus V_i$ is non-empty, then L_{i+1} exists and is defined to be

$$\left\{ \mathcal{A} \in (At(\Sigma) \setminus \mathsf{S}_\Sigma) \mid \mathcal{A} \notin V_i \text{ and } \bigwedge \mathcal{A} \wedge [d][d*] \bigvee_{\mathcal{B} \in V_i} \bigwedge \mathcal{B} \text{ is consistent} \right\}.$$

On the other hand, if $At(\Sigma) \setminus V_i$ is empty then there is no $i+1$-th level on $At(\Sigma) \setminus \mathsf{S}_\Sigma$.

Lemma 27. *Suppose $At(\Sigma) \setminus \mathsf{S}_\Sigma$ is non-empty. Then every atom in $At(\Sigma) \setminus \mathsf{S}_\Sigma$ belongs to exactly one level. Furthermore, there is a maximal level L_{max}.*

Proof. It is clear that each atom in $At(\Sigma) \setminus \mathsf{S}_\Sigma$ belongs to at most one level. Further, it follows by induction that no level is non-empty. For the base case let $\mathcal{A} \in At(\Sigma) \setminus \mathsf{S}_\Sigma$. By Lemma 6 item 5, $\langle d*\rangle \mathit{term} \in \mathcal{A}$. If $\mathit{term} \in \mathcal{A}$, then $\mathcal{A}$ belongs to L_0. On the other hand, if $\mathit{term} \notin \mathcal{A}$, then $\langle d{+}\rangle \mathit{term} \in \mathcal{A}$, and by Lemma 12, there is an atom $\mathcal{B}$ such that $\mathit{term} \in \mathcal{B}$ and $\mathcal{A}(S_d)^+\mathcal{B}$. Either way, some atom contains term and the base case of the induction is established. To drive through the inductive step of this argument, use Lemma 25. It follows by induction that no level is empty.

As there are only finitely many atoms, there is a maximum level L_{max}. Suppose for the sake of a contradiction that some atom $\mathcal{A}$ belongs to no level. Then $\mathcal{A} \notin V_{max}$, hence $At(\Sigma) \setminus V_{max}$ is non-empty, hence by Lemma 25 L_{max+1} exists and is non-empty; a contradiction. We conclude that every atom belongs to at least one level. $\dashv$

Definition 28 (Levels 2). For an arbitrary closed set $Cl(\Sigma)$, the *levels* on $At(\Sigma)$ are defined as follows. If $\mathsf{S}_\Sigma = At(\Sigma)$, then all atoms have level 0. On the other hand, if $At(\Sigma) \setminus \mathsf{S}_\Sigma$ is non-empty, then all atoms in $At(\Sigma) \setminus \mathsf{S}_\Sigma$ receive the level assigned by Definition 26, and all atoms in S_Σ are assigned the level L_{max+1}, where L_{max} is the maximum level assigned to an atom in $At(\Sigma) \setminus \mathsf{S}_\Sigma$.

The following lemma tells us that $At(\Sigma)$ really is 'vertically well behaved'.

Lemma 29. *Let $\mathcal{A}, \mathcal{B} \in At(\Sigma)$. Suppose $\mathcal{A} \in L_{i+1}$ where $i \geq 0$ and $\langle d \rangle \phi \in \mathcal{A}$. Then there is an atom $\mathcal{B} \in L_m$, where $m < i+1$, such that $\phi \in \mathcal{B}$ and $\mathcal{A} S_d \mathcal{B}$.*

Proof. Case 1: $\mathcal{A} \in L_{max+1}$. Let $\langle d \rangle \phi \in \mathcal{A}$, and suppose for the sake of a contradiction that there is no atom $\mathcal{B}$ in a lower level such that $\phi \in \mathcal{B}$ and $\mathcal{A} S_d \mathcal{B}$. Now, by Lemma 10 there is at least one atom $\mathcal{C}$ such that $\phi \in \mathcal{C}$ and $\mathcal{A} S_d \mathcal{C}$, hence by our initial supposition $\mathcal{C}$ must be in L_{max+1}, and hence *start* $\in \mathcal{C}$. As $\mathcal{A} S_d \mathcal{C}$, by Lemma 17, $\mathcal{C} S_u \mathcal{A}$. As $\top \in \mathcal{A}$, $\langle u \rangle \top \in \mathcal{C}$. But by Lemma 20 this contradicts the fact that *start* $\in \mathcal{C}$. We conclude that an appropriate atom $\mathcal{B}$ in a lower level exists.

Case 2: $\mathcal{A} \in L_{i+1}$ where $i + 1 \leq max$. Let $\langle d \rangle \phi \in \mathcal{A}$. Suppose for the sake of a contradiction that for all atoms $\mathcal{B} \in V_i$, $\bigwedge \mathcal{A} \wedge \langle d \rangle \bigwedge \mathcal{B}$ is inconsistent. This means that for all $\mathcal{B} \in V_i$, $\vdash \bigwedge \mathcal{A} \to [d] \neg \bigwedge \mathcal{B}$. Enumerate all the atoms in V_i as $\{\mathcal{B}_1, \ldots, \mathcal{B}_n\}$, and let $\mathbf{B}$ be $\bigwedge \mathcal{B}_1 \vee \cdots \vee \bigwedge \mathcal{B}_n$. It follows by simple modal reasoning that $\vdash \bigwedge \mathcal{A} \to [d] \neg \mathbf{B}$. Now, by our definition of levels, $\bigwedge \mathcal{A} \wedge [d] \mathbf{B}$ is consistent, therefore $\bigwedge \mathcal{A} \wedge [d](\mathbf{B} \wedge \neg \mathbf{B})$ is consistent also. But as $\langle d \rangle \phi$ belongs to $\mathcal{A}$, this implies that $\langle d \rangle (\mathbf{B} \wedge \neg \mathbf{B})$ is consistent, which is impossible. We conclude that the required atom $\mathcal{B}$ exists. $\dashv$

We now turn to a trickier task: ensuring that $At(\Sigma)$ is also 'horizontally well behaved'. We need the auxiliary notion of a downset.

Definition 30 (Downsets). Let $\mathcal{A}_0 \in L_{i+1}$, where $i \geq 0$. Then the *downset* of $\mathcal{A}_0$ is $\{\mathcal{D} \in V_i \mid \mathcal{A}_0 S_d \mathcal{D}\}$, and the *initial segment* of the downset is simply $\{\mathcal{D} \in V_i \mid \mathcal{A}_0 S_d \mathcal{D}$ and *first* $\in \mathcal{D}\}$.

Lemma 31. *Let $\mathcal{A}_0 \in L_{i+1}$, where $i \geq 0$. Then the downset of $\mathcal{A}_0$, and the initial segment of this downset, are both non-empty.*

Proof. As $\mathcal{A}_0 \in L_{i+1}$, where $i \geq 0$, *term* $\notin \mathcal{A}_0$. Hence by Lemma 20, there is some formula of the form $\langle d \rangle \psi \in \mathcal{A}_0$. By axiom 9 and the fifth closure condition, $\langle d \rangle$*first* $\in \mathcal{A}_0$. By the previous lemma, there is a $\mathcal{D} \in V_i$ such that *first* $\in \mathcal{D}$ and $\mathcal{A}_0 S_d \mathcal{D}$, thus the initial segment of $\mathcal{A}_0$'s downset is non-empty, and so is $\mathcal{A}_0$'s downset. $\dashv$

In order to proceed further, we must define a notion of *rank* on downsets. The basic ideas are similar to those underlying our notion of level; in particular, our initial observation concerning the Löb axiom does the real work. As a first step, we prove a horizontal analog of Lemma 25.

Lemma 32. *Let* D *be a downset of some atom* $\mathcal{A}_0$ *belonging to* L_{i+1}*, where* $i \geq 0$*, and let* I *be its initial segment. Suppose* $\mathsf{D} \setminus \mathsf{I}$ *is non-empty, and let* $\mathsf{A} = \{\mathcal{A}_1, \ldots, \mathcal{A}_n\}$ *and* $\mathsf{B} = \{\mathcal{B}_1, \ldots, \mathcal{B}_m\}$ *be disjoint non-empty sets of atoms such that* $\mathsf{A} \cup \mathsf{B} = \mathsf{D} \setminus \mathsf{I}$*. Then for some* $\mathcal{A} \in \mathsf{A}$*,*

$$\bigwedge \mathcal{A} \wedge [r][r{+}](\bigwedge \mathcal{B}_1 \vee \cdots \vee \bigwedge \mathcal{B}_m)$$

is consistent.

Proof. Let $\mathbf{A}$ be $\bigwedge \mathcal{A}_1 \vee \cdots \vee \bigwedge \mathcal{A}_n$. As any atom is consistent, so is $\mathbf{A}$. By the previous lemma, I is non-empty. Let $\mathbf{I}$ be $\bigwedge \mathcal{I}_1 \vee \cdots \vee \bigwedge \mathcal{I}_l$, where the $\mathcal{I}_j$ $(1 \leq j \leq l)$ are all and only the elements of I. As *first* does not belong to any atom in A, $\langle r{+}\rangle \mathbf{A}$ is consistent.

Let H (short for 'High') be the set of all atoms in $At(\Sigma) \setminus V_i$. Note that H is non-empty, for by our initial assumption there is at least one atom in L_{i+1}. Define $\mathbf{H}$ to be $\bigwedge \mathcal{H}_1 \vee \cdots \vee \bigwedge \mathcal{H}_p$, where the $\mathcal{H}_j$ $(1 \leq j \leq p)$ are all and only the elements of H. Let L (short for 'Low') be the set of all atoms in $V_i \setminus \mathsf{D}$. Note that it is possible that L is empty. If this is the case, we define $\mathbf{L}$ to be $\perp$, otherwise we define it to be $\bigwedge \mathcal{L}_1 \vee \cdots \vee \bigwedge \mathcal{L}_q$, where the $\mathcal{L}_j$ $(1 \leq j \leq q)$ are all and only the elements of L. Let Ψ be $\mathbf{I} \vee \mathbf{H} \vee \mathbf{L} \vee \mathbf{A}$. As $\langle r{+}\rangle \mathbf{A}$ is consistent, so is $\langle r{+}\rangle \Psi$, hence by Lemma 23, $\Psi \wedge [r][r{+}]\neg\Psi$ is consistent too.

Let $\mathbf{B}$ be $\bigwedge \mathcal{B}_1 \vee \cdots \vee \bigwedge \mathcal{B}_m$. As $At(\Sigma) = \mathsf{I} \cup \mathsf{H} \cup \mathsf{L} \cup \mathsf{A} \cup \mathsf{B}$, it follows from Lemma 7 that

$$\begin{aligned}
\vdash \mathbf{I} \vee \mathbf{H} \vee \mathbf{L} \vee \mathbf{A} \vee \mathbf{B} &\Rightarrow \vdash \neg(\mathbf{I} \vee \mathbf{H} \vee \mathbf{L} \vee \mathbf{A}) \rightarrow \mathbf{B} \\
&\Rightarrow \vdash \neg\Psi \rightarrow \mathbf{B} \\
&\Rightarrow \Psi \wedge [r][r{+}]\mathbf{B} \text{ is consistent} \\
&\Rightarrow \mathbf{A} \wedge [r][r{+}]\mathbf{B} \text{ is consistent.}
\end{aligned}$$

Hence for some $\mathcal{A} \in \mathsf{A}$, $\mathcal{A} \wedge [r][r{+}]\mathbf{B}$ is consistent, the required result. ⊣

Definition 33 (Ranks 1). Let D be a downset of some atom $\mathcal{A}_0$ belonging to L_{i+1} $(i \geq 0)$ with a non-empty initial segment I. Then the $\mathcal{A}_0$-*ranks* on $\mathsf{D} \setminus \mathsf{I}$ are defined as follows. R_0 is $\{\mathcal{D} \in (\mathsf{D} \setminus \mathsf{I}) \mid \mathit{last} \in \mathcal{D}\}$. For $i \geq 0$, H_i is defined to be $\bigcup_{0 \leq j \leq i} R_j$, and if $\mathsf{D} \setminus H_i$ is non-empty then R_{i+1} exists and is defined to be

$$\left\{ \mathcal{D} \in (\mathsf{D} \setminus \mathsf{I}) \mid \mathcal{D} \notin H_i \text{ and } \bigwedge \mathcal{D} \wedge [r][r*] \bigvee_{\mathcal{E} \in H_i} \bigwedge \mathcal{E} \text{ is consistent} \right\}.$$

On the other hand, if $\mathsf{D} \setminus H_i$ is empty then there is no $i+1$-th rank on $\mathsf{D} \setminus \mathsf{S}_\Sigma$.

Although the point should be clear, it's probably worth emphasizing that ranks are defined relative to some atom $\mathcal{A}_0$. Levels, on the other hand, were defined in absolute terms.

Next we prove a horizontal analog of Lemma 27.

Lemma 34. *Let* D *be a downset of some atom* $\mathcal{A}_0$ *belonging to* L_{i+1}, *where* $i \geq 0$, *with a non-empty initial segment* I. *Then every atom in* $\mathsf{D} \setminus \mathsf{I}$ *belongs to exactly one* $\mathcal{A}_0$*-rank. Furthermore, there is a maximal* $\mathcal{A}_0$*-rank* R_{max} *on* D.

Proof. It is clear that each atom in $\mathsf{D} \setminus \mathsf{I}$ belongs at most one $\mathcal{A}_0$-rank. Further, it follows by induction that no $\mathcal{A}_0$-rank is non-empty. For the base case, note that D is the downset of the atom $\mathcal{A}_0$ in L_{i+1}. At $\mathcal{A}_0$ has a non-empty downset, it contains a formula of the form $\langle d \rangle \psi$. By axiom 9 and the fifth condition on closures, it must also contain $\langle d \rangle last$, hence by Lemma 29 there is some atom $\mathcal{D} \in \mathsf{D}$ containing *last* and the base case is established. The inductive step follows from Lemma 32, and hence every $\mathcal{A}_0$-rank is non-empty. The remainder of the argument is essentially the same as that for Lemma 27. ⊣

Definition 35 (Ranks 2). For an arbitrary downset D of an atom $\mathcal{A}_0$, the $\mathcal{A}_0$-*ranks* on D are defined as follows. If $\mathsf{D} \setminus \mathsf{I}$ is empty, then all atoms in D have $\mathcal{A}_0$-rank 0. On the other hand, if $\mathsf{D} \setminus \mathsf{I}$ is non-empty, then all atoms in $\mathsf{D} \setminus \mathsf{I}$ receive the $\mathcal{A}_0$-rank assigned by Definition 33, and all atoms in I are assigned the $\mathcal{A}_0$-rank R_{max+1}, where R_{max} is the maximum $\mathcal{A}_0$-rank assigned to an atom in $\mathsf{D} \setminus \mathsf{I}$.

Lemma 36. *Suppose* $\mathcal{A}$ *has* $\mathcal{A}_0$*-rank* R_{i+1}, *where* $i \geq 0$ *and* $\langle r \rangle \phi \in \mathcal{A}$. *Then there is a* $\mathcal{B}$ *with* $\mathcal{A}_0$*-rank* R_m, *where* $m < i+1$, *such that* $\phi \in \mathcal{B}$ *and* $\mathcal{A} S_r \mathcal{B}$.

Proof. Essentially identical to the proof of Lemma 29. ⊣

The previous lemma is our first clue that downsets are horizontally well behaved, but we have more work to do. The next definition isolates the key concept required.

Definition 37 (Witnessing paths). Let $\mathcal{A} \in L_{i+1}$ and let D be its downset. A *witnessing path* for $\mathcal{A}$ is non-empty subset $\{\mathcal{D}_1, \ldots, \mathcal{D}_n\}$ of D such that:

1. $\mathcal{D}_i S_r \mathcal{D}_{i+1}$, for for all $i < n$.
2. *first* $\in \mathcal{D}_1$; and for all $i > 1$, *first* $\notin \mathcal{D}_i$.
3. *last* $\in \mathcal{D}_n$; and for all $i < n$, *last* $\notin \mathcal{D}_i$.
4. If $\langle d \rangle \psi \in \mathcal{A}$, then $\psi \in \mathcal{D}_i$ for some $0 \leq i \leq n$.

The reader should compare this definition with the definition of sensible decorations. Witnessing paths are designed to provide the structure demanded by the truth lemma for induced models, and to do so using atoms of lower level. Thus our goal is to prove that enough witnessing paths exist. First, a preliminary lemma.

Lemma 38. *Let $\mathcal{A} \in L_{i+1}$, and let $\mathcal{F}$ be any element of the initial segment of $\mathcal{A}$'s downset. If $\langle d\rangle\psi \in \mathcal{A}$, then $\langle r*\rangle\psi \in \mathcal{F}$.*

Proof. By Lemma 31, the initial segment of $\mathcal{A}$'s downset is non-empty, so such an $\mathcal{F}$ exists. Suppose for the sake of a contradiction that for some $\langle d\rangle\psi \in \mathcal{A}$, $\langle r*\rangle\psi \notin \mathcal{F}$. By the third closure condition, $\neg\langle r*\rangle\psi \in \mathcal{F}$, hence as *first* $\in \mathcal{F}$, *first* $\wedge \neg\langle r*\rangle\psi$ is consistent. Now, as $\mathcal{A}S_d\mathcal{F}$ holds, $\bigwedge \mathcal{A} \wedge \langle d\rangle \bigwedge \mathcal{F}$ is consistent, hence $\bigwedge \mathcal{A} \wedge \langle d\rangle(\mathit{first} \wedge \neg\langle r*\rangle\psi)$ is consistent. As $\langle d\rangle\psi \in \mathcal{A}$,

$$\bigwedge \mathcal{A} \wedge \langle d\rangle\psi \wedge \langle d\rangle(\mathit{first} \wedge \neg\langle r*\rangle\psi)$$

is also consistent. Hence by axiom 8,

$$\bigwedge \mathcal{A} \wedge [d](\mathit{first} \rightarrow \langle r*\rangle\psi) \wedge [d](\mathit{first} \rightarrow \neg\langle r*\rangle\psi)$$

is consistent, thus by simple modal reasoning,

$$\bigwedge \mathcal{A} \wedge [d]((\mathit{first} \rightarrow \langle r*\rangle\psi) \wedge (\mathit{first} \rightarrow \neg\langle r*\rangle\psi))$$

is consistent as well. As $\mathcal{A}$'s downset is non-empty, $\langle d\rangle\mathit{first} \in \mathcal{A}$, hence

$$\langle d\rangle\mathit{first} \wedge [d]((\mathit{first} \rightarrow \langle r*\rangle\psi) \wedge (\mathit{first} \rightarrow \neg\langle r*\rangle\psi))$$

is consistent, hence $\langle d\rangle(\langle r*\rangle\psi\wedge\neg\langle r*\rangle\psi)$ is consistent too — but this is impossible. We conclude that $\langle r*\rangle\psi \in \mathcal{F}$. ⊣

Lemma 39. *Let $\mathcal{A} \in L_{i+1}$ and let D be its downset. Then $\mathcal{A}$ has a witnessing path.*

Proof. Choose any element $\mathcal{F}$ of $\mathcal{A}$'s initial segment. We now construct a witnessing path for $\mathcal{A}$ whose first item is $\mathcal{F}$.

Case 1: $\mathcal{F}$ contains no formula of the form $\langle r\rangle\phi$. Suppose $\langle d\rangle\psi \in \mathcal{A}$. By the previous lemma, $\langle r*\rangle\psi \in \mathcal{F}$. As no formula of the form $\langle r\rangle\phi$ is in $\mathcal{F}$, no formula of the form $\langle r+\rangle\phi$ is in $\mathcal{F}$ either, and hence $\psi \in \mathcal{F}$. As a special case of this, note that by axiom 9 and the fifth closure condition, *last* $\in \mathcal{F}$. Hence $\{\mathcal{F}\}$ is a witnessing path for $\mathcal{A}$.

Case 2: $\mathcal{F}$ contains a formula of the form $\langle r\rangle\phi$. By Lemma 36, it is possible to construct a sequence $\mathcal{F} = \mathcal{D}_1 S_r \mathcal{D}_2 \cdots$, where all items in the sequence belong to $\mathcal{A}$'s downset, and such that $\mathcal{D}_i S_r \mathcal{D}_{i+1}$ implies that $\mathcal{D}_{i+1}$ has a strictly lower $\mathcal{A}$-rank than $\mathcal{D}_i$. Construct such a sequence that is closed under S_r successors. As $\langle x\rangle\psi \in \mathcal{F}$, this sequence has length at least 2. Moreover, the sequence must be finite: as each item in the sequence has a strictly lower $\mathcal{A}$-rank than all its predecessors, each item in the sequence is unique. As there are only finitely many atoms in $\mathcal{A}$'s downset, the sequence has length n, for some natural number n.

Clearly *first* $\in \mathcal{F}$. Moreover, for any atom $\mathcal{D}_i$ in the sequence $(2 \leq i \leq n)$, $\mathcal{D}_{i-1} S_r \mathcal{D}_i$, hence by Lemma 17 $\mathcal{D}_i S_l \mathcal{D}_{i-1}$. As $\top \in \mathcal{D}_{i-1}$, $\langle l\rangle\top \in \mathcal{D}_i$, hence by Lemma 20, *first* $\notin \mathcal{D}_i$. Next, note that there can be no formula of the form $\langle r\rangle\psi \in \mathcal{D}_n$. If there were, we could apply Lemma 36 to find an atom $\mathcal{D}_{n+1}$ such that $\mathcal{D}S_r\mathcal{D}_{n+1}$. But the sequence is closed under S_r successors, and $\mathcal{D}_n$ is the final item in the sequence, so this is impossible. Thus no formula of the form $\langle r\rangle\psi$ belongs to $\mathcal{D}_n$, hence by Lemma 20, *last* $\in \mathcal{D}_n$. We leave it to the reader to verify (again using Lemma 20) that *last* cannot belong to $\mathcal{D}_i$ for $i < n$.

It remains to verify that $\langle d\rangle\psi \in \mathcal{A}$ implies $\mathcal{A} \in \mathcal{D}_i$ for some $1 \leq i \leq n$. Suppose for the sake of a contradiction that this is not the case; that is, for some $\langle d\rangle\psi \in \mathcal{A}$, $\psi \notin \mathcal{D}_i$ for all $1 \leq i \leq n$. Now, by the previous lemma, $\langle d\rangle\psi \in \mathcal{A}$ implies that $\langle r*\rangle\psi \in \mathcal{F}$. As $\langle r*\rangle\psi \in \mathcal{F}$ and $\psi \notin \mathcal{F}$, we have that $\langle r+\rangle\psi \in \mathcal{F}$. By Lemma 21, for all $1 \leq i < n$, if $\langle r+\rangle\psi \in \mathcal{D}_i$ and $\mathcal{D}_iS_r\mathcal{D}_{i+1}$, then $\psi \in \mathcal{D}_{i+1}$ or $\langle r+\rangle\psi \in \mathcal{D}_{i+1}$. As by assumption ψ belongs to no item in the sequence, $\langle r+\rangle\psi$ belongs to them all, and in particular, $\langle r+\rangle\psi \in \mathcal{D}_n$. By Lemma 20, this means that $\neg$*last* $\in \mathcal{D}_n$, contradicting the fact that *last* $\in \mathcal{D}_n$. We conclude that our original supposition was false, and have the desired result. $\dashv$

4.5 Constructing and Decorating a Finite Tree

We can now simultaneously construct a finite tree $\mathbf{T}$ and a decoration h of $\mathbf{T}$ by induction. The construction will terminate after finitely many steps, and, as we shall see, results in a *sensible* decoration of $\mathbf{T}$.

So, suppose ψ is LOFT-consistent. Let T be a denumerably infinite set. We will use (finitely many) of its elements as the tree nodes, and decorate them with atoms taken from $At(\{start \wedge \langle d*\rangle\psi\})$.

Stage 1. Choose some $t_1 \in \mathrm{T}$ and an atom $\mathcal{A} \in At(\{start \wedge \langle d*\rangle\psi\})$ that contains $start \wedge \langle d*\rangle\psi$. (As ψ is consistent, so is $start \wedge \langle d*\rangle\psi$, so this is possible.) Define

$$\begin{aligned} T^1 &:= \{t_1\} \\ R^1_r &:= \emptyset \\ R^1_d &:= \emptyset \\ h^1 &:= \{(t_1, \mathcal{A})\}. \end{aligned}$$

Stage n+1. Suppose n stages of the construction have been performed. Call $t \in T^n$ *unsatisfied* if for some $\langle d\rangle\psi \in h^n(t)$ there is no $t' \in T^n$ such that tR_dt' and $\psi \in h^n(t')$.

if there are no unsatisfied nodes
then halt
else choose an unsatisfied node t. As $\langle d\rangle\psi \in h^n(t)$, by Lemma 39 $h^n(t)$ has a witnessing path $\{\mathcal{D}_1, \ldots, \mathcal{D}_k\}$. Let $t_1, \ldots, t_k \in \mathrm{T}\backslash T^n$. Define:

$$\begin{aligned} T^{n+1} &= T^n \cup \{t_1, \ldots, t_k\} \\ R^{n+1}_d &= R^n_d \cup \{(t, t_1), \ldots, (t, t_k)\} \\ R^{n+1}_r &= R^n_r \cup \{(t_i, t_{i+1}) \mid 1 \leq i < k\} \\ h^{n+1} &= h^n \cup \{(t_i, \mathcal{D}_i) \mid 1 \leq i \leq k\}. \end{aligned}$$

Lemma 40. *The above construction halts after finitely many steps. Moreover the number of nodes adjoined in the course of the construction is finite.*

Proof. Whenever we adjoin new R_d successors to a node t, we adjoin one new node for each element of the chosen witnessing path for $h(t)$. But witnessing paths are finite, thus at each stage we adjoin only finitely many nodes. Moreover, as witnessing paths are subsets of downsets, each element of $h(t)$'s witnessing path belongs to a strictly lower level that $h(t)$. As there are only finitely many levels, we can only adjoin new nodes finitely many times in the course of the construction. $\dashv$

Let max be the stage at which the construction halts. Define T to be T^{max}, R_d to be R_d^{max}, R_r to be R_r^{max}, and h to be h^{max}. Let R_l and R_u be the converses of R_r and R_d respectively. Let $\mathbf{T}$ be (T, R_l, R_r, R_u, R_d).

Lemma 41. $\mathbf{T}$ *is a finite tree and* h *is a sensible decoration of* $\mathbf{T}$.

Proof. That $\mathbf{T}$ is a finite tree follows straightforwardly from the nature of the inductive construction. To see that h is a sensible decoration of $\mathbf{T}$, argue as follows.

First, suppose the construction halts immediately after stage 1. By construction, $start \in h(t)$. As the construction halted after one step, t_1 was not unsatisfied, so there is no formula of the form $\langle d \rangle \psi \in h(t_1)$ and hence by Lemma 20, $term \in h(t)$. By axiom 10, $first, last \in h(t)$ also. It follows that h is a sensible decoration of $\mathbf{T}$.

So suppose the construction closed after max steps, where $max > 1$. The important point to observe is that because we used witnessing paths to satisfy tree nodes t, h fulfills the first three clauses in the definition of a sensible decoration. The fourth clause in the definition of sensible decorations insists that the constants be 'sensibly distributed'. Now, *first* and *last* are sensibly distributed in all witnessing paths. Further, *start* is sensibly distributed because $start \in h(t)$, where t is the root node in the tree, and thereafter the construction assigns atoms of lower level to tree nodes, and such nodes do not contain *start*. We leave it the reader to verify that *term* is also sensibly distributed. $\dashv$

Theorem 42. LOFT *is complete with respect to finite trees.*

Proof. Given a consistent formula ϕ, use the inductive construction to build a finite tree $\mathbf{T}$ and a decoration $h : \mathbf{T} \longrightarrow At(\{start \wedge \langle d* \rangle \phi\})$. Let $\mathbf{M}$ be the model induced by h on $\mathbf{T}$. By the previous lemma, h is a sensible decoration of $\mathbf{T}$, hence by the truth lemma for induced models (Lemma 22), $\mathbf{M}$ satisfies $start \wedge \langle d* \rangle \phi$ at the root node, and thus ϕ is true somewhere in this model. $\dashv$

5 Discussion

To conclude the paper we note some issues concerning complexity and theorem proving raised by this work. As a first step, note that LOFT is decidable. This could be proved by appealing to the results of Rabin (1969), but the completeness result yields it immediately.

Theorem 43. LOFT *is decidable.*

Proof. Because we are only working with finite trees, the set of satisfiable formulas is clearly RE. But the set of non-satisfiable formulas is also RE: by completeness, our axiomatisation recursively enumerates all the valid $\mathcal{L}$ formulas. So if a formula ϕ is not satisfiable on a finite tree, then its negation will eventually be generated.

What is the complexity of LOFT's satisfiability problem? The easiest way to answer this question is to think in terms of pseudo-models. We proved the following completeness theorem: if ϕ is consistent, then, by Lemma 22, it is satisfiable in a pseudo-model, namely, the pseudo-model over $At(\{\phi\})$. (The corresponding soundness theorem is clear: if ϕ is not consistent, it cannot belong to any atom in any closure, hence it cannot be satisfied in any pseudo-model at all.) As we now know that LOFT is the logic of finite trees, the completeness result for pseudo-models takes on a new significance. For a start, as $|At(\{\phi\})|$ is $O(2^{|\phi|})$, it gives an exponential upper bound on the size of pseudo-models needed to establish LOFT-satisfiability. This immediately yields:

Theorem 44. LOFT-*satisfiability is in NEXPTIME.*

But with a little more effort, one can do better.

Theorem 45. LOFT-*satisfiability is EXPTIME-complete.*

Proof. The lower bound is an immediate corollary of Spaan's (1993) analysis of the lower bound result for PDL. She notes that the following fragment of PDL is EXPTIME-hard: formulas of the form $\psi \wedge [a*]\theta$, (where ψ and θ contain only the atomic program a and no embedded modalities) that are satisfiable at the root of a finite binary tree. Trivially, this PDL fragment can be identified with an $\mathcal{L}$ fragment in the modalities $[d*]$ and $[d]$, hence LOFT-satisfiability is also EXPTIME-hard.

The upper bound can be proved by using the methods of Pratt (1979). We sketch what is involved. The reader who consults Pratt's paper will have no difficulty in filling in the details.

Following Pratt, we define $H(\Sigma)$, the set of *Hintikka sets* over Σ, to be subsets of $Cl(\Sigma)$ that have all the properties of atoms listed in lemma 6, but that may not be consistent. That is, $At(\Sigma) \subseteq H(\Sigma) \subseteq Cl(\Sigma)$. For $\mathcal{H}, \mathcal{H}' \in H(\Sigma)$ and $x \in \{l, r, u, d\}$, define $\mathcal{H}S'_x\mathcal{H}'$ to hold iff for some atomic formula p, $\langle x\rangle p \in \mathcal{H}$ and $p \in \mathcal{H}'$, and moreover, for all atomic formulas q, $[x]q \in \mathcal{H}$ implies $q \in \mathcal{H}'$. Define

$$\mathbf{D}^0 := (H(\Sigma), \{S'_x, (S'_x)^+\}_{x\in\{l,r,u,d\}}).$$

Given $\mathbf{D}^n$, one forms $\mathbf{D}^{n+1}$ by eliminating all Hintikka sets $\mathcal{H} \in \mathbf{D}^n$ such that $\langle x\rangle\phi \in \mathcal{H}$, but there is no $\mathcal{H}'$ such that $\phi \in \mathcal{H}'$ and $\mathcal{H}S'_x\mathcal{H}'$. The relations for $\mathbf{D}^{n+1}$ are defined by restricting the relations on $\mathbf{D}^n$ to this (smaller) set of Hintikka sets. This process terminates (there are only finitely many Hintikka sets) and yields a model. It is standard work to show that if ϕ is consistent, then this model satisfies ϕ.

From the point of view of complexity, two points are important. First, the process terminates after at most exponentially many steps, as there are only exponentially many Hintikka sets. Second, at each stage it is possible to calculate in polynomial time which Hintikka sets to eliminate. Thus 'elimination of Hintikka sets' is a deterministic EXPTIME-algorithm for LOFT-satisfiability. $\dashv$

However, while interesting in its own right, the above EXPTIME-completeness result for LOFT raises another question. For many applications we are not merely interested in whether or not ϕ is satisfiable: if ϕ is satisfiable, we would like to see a concrete finite tree that satisfies it. (This would be useful for applications in computational linguistics.) By the previous result, this problem is EXPTIME-hard, but at present we do not have tight upper and lower bounds.

Similar considerations apply to theorem proving for LOFT. It is clearly possible to devise tableaux systems for LOFT: working with pseudo-models is essentially the same as working with tableaux, and indeed the completeness result for pseudo-models gives us all that is required to define such systems. But a more interesting question is the following. Is it possible to develop a tableaux system that produces finite trees directly and is reasonably efficient on the formulas typically encountered in applications? Such issues are the focus of our ongoing work.

Acknowledgement. Wilfried Meyer-Viol was partially supported by the Engineering and Physical Sciences Research Counsil (EPSRC) under Grant Reference Number GR/K68776.

Part of the research was carried out while Maarten de Rijke was with CWI, Amsterdam; during this period he was supported by the Netherlands Organization for Scientific Research (NWO), project NF 102/62-356.

References

Backofen, R., Rogers, J., Vijay-Shanker, K.: A first-order axiomatisation of the theory of finite trees. *Journal of Logic, Language and Information*, 4:5–39, 1995.

van Benthem, J., Meyer-Viol, W.: *Logical Semantics of Programming.* Manuscript, 1994.

Berlage, T.: A selective undo mechanism for graphical user interfaces on command objects. *ACM Transactions on Computer-Human Interaction* 1(3):269–294, 1994.

Blackburn, P., Gardent, C., Meyer-Viol, W.: Talking about trees. In *Proceedings of the 6th Conference of the European Chapter of the Association for Computational Linguistics*, pages 21–29, Utrecht, The Netherlands, 1993.

Blackburn, P., Meyer-Viol, W.: Linguistics, logic and finite trees. *Bulletin of the IGPL* 2:2–29, 1994.

Blackburn, P., Gardent, C.: A specification language for lexical functional grammars. In *Proceedings of the 7th Conference of the European Chapter of the Association for Computational Linguistics*, pages 39–44, Dublin, Ireland, 1995.

Courcelle, B.: Fundamental properties of infinite trees. *Theoretical Computer Science* 25:95–169, 1985.

Fischer, M.J., Ladner, R.F.: Propositional dynamic logic of regular programs. *J. Comp. Syst. Sci.* 18:194–211, 1979.

Kracht, M.: Mathematical Aspects of command relations. In *Proceedings of the 6th Conference of the European Chapter of the Association for Computational Linguistics*, pages 240–249, Utrecht, The Netherlands, 1993.

Kracht, M.: Syntactic codes and grammar refinement. *Journal of Logic, Language and Information* 4: 41–60, 1995.

Maher, M.: Complete axiomatisations of the algebras of finite, rational and infinite trees. In *Proceedings of the 3rd International Symposium on Logic in Computer Science*, pages 348–357, Edinburgh, Scotland, 1988.

Goldblatt, R.: *Logics of Time and Computation.* 2nd edition, CSLI Lecture notes **7**, Center for the Study of Language and Information, Stanford, 1992.

Pratt, V.: Models of program logics. In *Proceedings of the 20th IEEE Symposium on Foundations of Computer Science*, pages 115-122, 1979.

Rabin, M.: Decidability of second-order theories and automata on infinite trees. *Transactions of the American Mathematical Society* 141:1–35, 1969.

Rogers, J.: *A Descriptive Approach to Language-Theoretic Complexity.* Studies in Logic, Language and Information, CSLI Publications, Stanford, 1996, to appear.

Smoryński, C.: *Self Reference and Modal Logic.* Springer Verlag, 1985.

Spaan, E.: *Complexity of Modal Logics.* PhD thesis, University of Amsterdam, 1993.