Logics with an existential modality

Natasha Alechina and Dmitry Shkatov

ABSTRACT. We consider multi-modal logics interpreted over edge-labelled graphs with a modality $\langle\#\rangle$, where $\langle\#\rangle\varphi$ means 'φ is accessible by an edge with *some* label'. In a logic with finitely many edge labels, $\langle\#\rangle$ is definable, but if the set of labels is infinite, it is an independent modality. We axiomatise multi-modal K, deterministic multi-modal K, and PDL with converse and a single nominal, extended with $\#$. The latter gives an axiomatisation of the logic PDL^{path} introduced in [3].

Keywords: Existential modality, multi-modal K, PDL, determinism

1 Introduction

In this paper, we consider multi-modal logics interpreted over edge-labelled graphs with a modality $\langle\#\rangle$, where $\langle\#\rangle\varphi$ means 'φ is accessible by an edge with *some* label'. We start by explaining what motivated our interest in the existential modality: namely, logics for modelling semi-structured data, such as data on the Web [1]. A collection of web pages can be represented as a graph with labelled edges. Edge labels come from some set I, which is either finite but very large, or even countably infinite. For example, I could be the set of all URLs, or all possible phrases in English (link names). Suppose we want to reason about constraints on possible paths in a graph, expressed as inclusions of regular expressions (inclusion constraints were introduced by Abiteboul and Vianu in [2]):

$$a; (b + c); \#; d^* \subseteq e; f$$

(if a data item is reachable by a path defined by $a; (b+c); \#; d^*$, that is: an a link followed by either a b or a c link, followed by an arbitrary link, followed by finitely many d links, then it is also reachable by a path $e; f$). We can study the implication problem for inclusion constraints (whether a set of constraints implies a constraint) by expressing it in a logical language; in [3], a logic called PDL^{path} was introduced for this purpose. The only unusual feature of PDL^{path} compared to other flavours of PDL, see e.g. [6, 5], is the wild card, or existential modality $\langle\#\rangle$, standing for 'any label'. In this paper, we consider axiomatisation and decidability problems for $\langle\#\rangle$, since, as far as we know, it has not been studied extensively before. The only reference we could find is in [4], where $\langle\#\rangle$ is used to make DPDL with intersection of programs badly undecidable. Clearly, in a language with finitely many edge labels $\langle\#\rangle$ is definable. It is easy to show (we do it below) that if the set of labels is infinite, then $\langle \# \rangle$ is not definable. But even if the set of labels I is finite, we may not want to write the very large formulas required to define $\langle \# \rangle$ (disjunction over all possible labels in I). For example, if I is the set of all URLs (which is finite but contains billions of elements), consider the difference between saying 'this web page does not have any outgoing links' as $\neg\langle \# \rangle\top$ and the alternative expression which involves a disjunction over all possible URLs.

The paper is organised as follows. First we study the logic $\mathsf{K}_\#$, obtained by adding $\langle \# \rangle$ to multi-modal K. We give a complete axiomatisation of $\mathsf{K}_\#$ and show that its complexity is the same as that of multi-modal K. However, an essential technical device we use in those proofs would not work with deterministic graphs. The device is as follows: given a formula φ and a model M for that formula, replace all edge labels in M which do not occur in φ by a single fresh edge label, and show that the resulting model still satisfies φ. This construction would not work if M were a deterministic model; we therefore believe that the completeness and decidability proof for deterministic multi-modal K, $\mathsf{DK}_\#$ is of an independent interest. This proof is given in section 4. In section 5 we briefly describe PDL^{path} introduced in [3] and in section 6 give a sound and weakly complete axiomatisation for it; decidability was proved in [3].

2 Logics $\mathsf{K}_\#$ and $\mathsf{DK}_\#$

Consider the propositional modal language $\mathcal{L}^I_\#$ containing (1) a countable set of propositional parameters **Par**; (2) propositional connectives $\neg$ ("not") and $\vee$ ("or"); (3) for every element i of the countable set I of modal indices, a modal operator $\langle i \rangle$; and (4) a modal operator $\langle \# \rangle$. All the other connectives, including the dual modalities $[i]$ and $[\#]$, can be defined in the usual way. The formulas of $\mathcal{L}^I_\#$ are defined by

$$\varphi := p \mid \neg\varphi \mid \varphi_1 \vee \varphi_2 \mid \langle i \rangle\, \varphi \mid \langle \# \rangle\, \varphi,$$

where $p \in \mathbf{Par}$ and $i \in I$. These formulas are evaluated on $\mathcal{L}^I_\#$-models.

DEFINITION 1. An $\mathcal{L}^I_\#$-model is a tuple $\mathcal{M} = (W, \{\mathcal{R}_i\}_{i \in I}, \mathcal{R}_\#, V)$, where $W \neq \emptyset$, $\mathcal{R}_i \subseteq W \times W$, $\mathcal{R}_\# = \bigcup_{i \in I} \mathcal{R}_i$, and V is a function from **Par** into 2^W.

$\mathcal{M}$ is deterministic if, for every $w \in W$ and every $i \in I$, there is no more than one v such that $w\mathcal{R}_i v$.

The truth definitions for formulas of $\mathcal{L}^I_\#$ are standard; in particular,

- $\mathcal{M}, w \Vdash \langle i \rangle\, \varphi$ iff $\exists v \in W(w\mathcal{R}_i v$ and $\mathcal{M}, v \Vdash \varphi)$
- $\mathcal{M}, w \Vdash \langle \# \rangle\, \varphi$ iff $\exists v \in W(w\mathcal{R}_\# v$ and $\mathcal{M}, v \Vdash \varphi)$.

It is easy to see that $\langle \# \rangle$ increases the expressive power of only those languages that contain at least a countable set I of modal indices; otherwise, $\langle \# \rangle\, \varphi$ can be defined as a finite disjunction of formulas of the form $\langle i \rangle\, \varphi$.

LEMMA 2. *Let φ be a formula not containing $\langle\#\rangle$, and let all labels occurring in φ be in the set L. Then φ is preserved with respect to the operation of removing R_i edges (where $i \notin L$).*

Proof. Let $\mathcal{M}, w \Vdash \varphi$, and let $\mathcal{M}'$ be obtained from $\mathcal{M}$ by removing all edges with labels not in L. Then $\mathcal{M}$ and $\mathcal{M}'$ are bisimilar with respect to $\{R_i : i \in L\}$, hence $\mathcal{M}', w \Vdash \varphi$. ■

LEMMA 3. *Let φ be a formula which does contain $\langle\#\rangle$, and let all labels occurring in φ be in the set L. Then,*

1. *φ is not guaranteed to be preserved with respect to removing non-L edges;*
2. *φ is preserved with respect to* renaming *non-L edges, provided that the new names are also not in L.*

COROLLARY 4. *$\langle\#\rangle$ is not definable in a language with an infinite set of labels I.*

Let us denote the logic of all $\mathcal{L}^I_\#$-models as $\mathsf{K}_\#$ and the logic of all deterministic $\mathcal{L}^I_\#$-models as $\mathsf{DK}_\#$. We are now going to formulate Hilbert-style axiomatisations of $\mathsf{K}_\#$ and $\mathsf{DK}_\#$. Since $\langle\#\rangle$ resembles the existential quantifier of first-order logic, it's not hard to see that the axiomatisation of $\mathsf{K}_\#$ should look as follows (π stands for either an arbitrary $i \in I$ or $\#$):

Axiom schemata:

(A0) All classical tautologies;

(K) $[\pi](\varphi \to \psi) \to ([\pi]\varphi \to [\pi]\psi)$;

(ER) $\langle i\rangle\varphi \to \langle\#\rangle\varphi$.

Inference rules:

(MP) $\dfrac{\vdash \varphi \to \psi, \vdash \varphi}{\vdash \psi}$; **(N)** $\dfrac{\vdash \varphi}{\vdash [\pi]\varphi}$;

(EL) $\dfrac{\vdash \langle i\rangle\varphi \to \psi}{\vdash \langle\#\rangle\varphi \to \psi}$, provided i does not occur in ψ.

Also, it is not difficult to guess that the axiomatisation of $\mathsf{DK}_\#$ can be obtained by adding to the axiom schemata above the 'axiom of determinism':

D $\langle i\rangle\, \varphi \to [i]\varphi$.

It is easy to show the following.

THEOREM 5. *$\mathsf{K}_\#$ is sound with respect to the class of all $\mathcal{L}_\#$-models. $\mathsf{DK}_\#$ is sound with respect to the class of all deterministic $\mathcal{L}_\#$-models.*

It is also easy to see that both $\mathsf{K}_\#$ and $\mathsf{DK}_\#$ are non-compact (consider the set $\Gamma = \{\, \langle\#\rangle\, p, \neg\langle i\rangle p : i \in I \,\}$), and hence don't have strongly-complete axiomatisations. In the next two sections, we prove weak completeness of

$\mathsf{K}_{\#}$ and $\mathsf{DK}_{\#}$, showing that every $\mathsf{K}_{\#}$-consistent and every $\mathsf{DK}_{\#}$-consistent formula has a model.

3 Completeness for $\mathsf{K}_{\#}$

We use the completeness-via-finite-models technique described in detail in [4].

Let us define $\sim \varphi$, the pseudo-negation of φ, as follows: if φ is $\neg\psi$, $\sim \varphi$ is ψ; otherwise, it is $\neg\varphi$. By the closure of the set of formulas Σ, we mean the smallest set $\mathsf{CL}(\Sigma)$ containing all the subformulas of formulas of Σ and their pseudo-negations. It is easy to see that $\mathsf{CL}(\Sigma)$ is finite whenever Σ is. A finite canonical model for a $\mathsf{K}_{\#}$-consistent formula φ is built out of $\{\varphi\}$-atoms, maximally consistent subsets of $\mathsf{CL}(\{\varphi\})$. (In general, for an arbitrary set of formulas Σ, a Σ-atom is a maximally consistent subset of $\mathsf{CL}(\Sigma)$.) The following is straightforward:

LEMMA 6. *If $\varphi \in \mathsf{CL}(\Sigma)$ is $\mathsf{K}_{\#}$-consistent, then there exists an atom A over Σ such that $\varphi \in A$.*

It is easy to see that every Σ-atom A has the following properties:

1. For every $\varphi \in \mathsf{CL}(\Sigma)$, exactly one of φ and $\sim \varphi$ belongs to A.

2. For every $\varphi \vee \psi \in \mathsf{CL}(\Sigma)$, $\varphi \in \mathsf{CL}(\Sigma)$ or $\psi \in \mathsf{CL}(\Sigma)$.

Now we define finite canonical models for $\mathsf{K}_{\#}$.

DEFINITION 7. Let Σ be a finite set of $\mathcal{L}^I_{\#}$-formulas and let $a \in I$ be such that a does not occur in Σ. The finite canonical model over Σ, $\mathcal{M}^{\Sigma}$, is the tuple $(At(\Sigma), \{\mathcal{R}^{\Sigma}_i\}_{i\in I}, \mathcal{R}^{\Sigma}_{\#}, V^{\Sigma})$, where

1. $At(\Sigma)$ is the set of all atoms over Σ;

2. $A\mathcal{R}^{\Sigma}_i A'$ iff $i \in \Sigma$ or $i = a$ and $\widehat{A} \wedge \langle i \rangle\, \widehat{A'}$ is consistent ($\widehat{X}$ stands for $\bigwedge_{\varphi \in X} \varphi$);

3. $A\mathcal{R}^{\Sigma}_{\#} A'$ iff $\widehat{A} \wedge \langle \# \rangle\, \widehat{A'}$ is consistent;

4. For every $p \in \mathbf{Par}$, $V^{\Sigma}(p) = \{\, A \in At(\Sigma) : p \in A \,\}$.

In a standard way (for details, see [4]), we can prove the following lemma.

LEMMA 8. *Let Σ be a set of $\mathcal{L}^I_{\#}$-formulas, A be an atom over Σ, and π be either an index occurring in Σ or $\pi = \#$. Then, for all $\langle\pi\rangle\varphi \in \mathsf{CL}(\Sigma)$, $\langle\pi\rangle\varphi \in A$ iff there is an atom A' such that $A\mathcal{R}_{\pi}A'$ and $\varphi \in A'$.*

From lemma 8 and the properties of atoms, we immediately obtain the following lemma.

LEMMA 9. *Let Σ be a set of $\mathcal{L}_{\#}$ formulas, $\mathcal{M}^{\Sigma}$ be the finite canonical model over Σ, and $\psi \in \mathsf{CL}(\Sigma)$. Then, for every $A \in At(\Sigma)$, $\mathcal{M}^{\Sigma}, A \Vdash \psi$ iff $\psi \in A$.*

All that remains to be done is to show that finite canonical models are $\mathcal{L}^I_{\#}$-models.

LEMMA 10. *Every finite canonical model* $\mathcal{M}^\Sigma = (At(\Sigma), \{\mathcal{R}^\Sigma_i\}_{i\in I}, \mathcal{R}^\Sigma_{\#}, V^\Sigma)$ *is an* $\mathcal{L}_{\#}$*-model.*

Proof. All we have to prove is that $\mathcal{R}^\Sigma_{\#} = \bigcup_{i\in I} \mathcal{R}^\Sigma_i$. First, right-to-left. Suppose that, for some $i \in I$, $A\mathcal{R}^\Sigma_i A'$ but $A\mathcal{R}^\Sigma_{\#} A'$ does not hold. Then, by definition 7, $\widehat{A} \wedge \langle i\rangle \widehat{A'}$ is consistent, but $\widehat{A} \wedge \langle \#\rangle \widehat{A'}$ is inconsistent. Then, $\vdash \langle \#\rangle \widehat{A'} \to \neg\widehat{A}$ and hence, by (ER), $\vdash \langle i\rangle \widehat{A'} \to \neg\widehat{A}$, which is impossible since $\widehat{A} \wedge \langle i\rangle \widehat{A'}$ is consistent.

Now, left-to-right. Suppose that $A\mathcal{R}^\Sigma_{\#} A'$. If, for some $i \in I$, $A\mathcal{R}^\Sigma_i A'$, then we are done. So, let us assume that for no $i \in I$ does $A\mathcal{R}^\Sigma_i A'$ hold. Then, $A\mathcal{R}^\Sigma_a A'$ holds. Indeed, if we suppose otherwise, then $\widehat{A}\wedge\langle \#\rangle \widehat{A'}$ is consistent and $\widehat{A} \wedge \langle a\rangle \widehat{A'}$ inconsistent. But then $\vdash \langle a\rangle \widehat{A'} \to \neg\widehat{A}$, and hence, by (EL) (note that $a \notin \Sigma$), $\vdash \langle \#\rangle \widehat{A'} \to \neg\widehat{A}$, which is impossible since $\widehat{A} \wedge \langle \#\rangle \widehat{A'}$ is consistent. ■

THEOREM 11. $\mathsf{K}_{\#}$ *is complete with respect to the class of all* $\mathcal{L}_{\#}$ *models.*

Proof. Immediately follows from lemmas 6, 9, and 10. ■

REMARK 12. If we had not added to the indices occurring in Σ a "new" index a, we would not have been able to prove that finite canonical models are $\mathcal{L}_{\#}$-models. As a counterexample, consider the set $\Sigma = \{\langle \#\rangle\, p \wedge \neg\langle b\rangle\, p\}$. Since $\langle \#\rangle\, p \wedge \neg\langle b\rangle\, p$ is consistent, in $\mathcal{M}^\Sigma$ there is an atom A such that $\langle \#\rangle\, p \wedge \neg\langle b\rangle\, p \in A$. Then, for some $B \in \mathcal{M}^\Sigma$ such that $p \in B$, we have $A\mathcal{R}^\Sigma_{\#} B$, but for no index $c \in \Sigma$ do we have $A\mathcal{R}^\Sigma_c B$.

In the proof of Theorem 11, we have constructed a model for a consistent formula φ which is of size $2^{|\varphi|}$. This gives us decidability of $\mathsf{K}_{\#}$. However, the complexity of checking whether φ is satisfiable by examining all models of size $2^{|\varphi|}$ is NEXPTIME. We can do better than that and show that the complexity of satisfiability problem for $\mathsf{K}_{\#}$ is no worse than that of multi-modal **K**.

THEOREM 13. *The satisfiability problem for* $\mathsf{K}_{\#}$ *is PSPACE-complete.*

Proof. There is a polynomial reduction from $\mathsf{K}_{\#}$ satisfiability to multi-modal **K** satisfiability (which is PSPACE-complete). A $\mathsf{K}_{\#}$ formula φ which contains labels $i_1, \ldots, i_n$ is satisfiable if and only if a formula φ' is satisfiable, where φ' is obtained from φ by replacing each subformula $\langle \#\rangle\psi$ with $\langle i_1\rangle\psi \vee \ldots \vee \langle i_{n+1}\rangle\psi$. From Lemma 3: rename all labels which do not occur in φ to be i_{n+1}. The resulting model satisfies φ'. The other direction: let $\mathcal{M}, w \Vdash \varphi'$. φ' does not contain $\langle \#\rangle$, so it is satisfied in model $\mathcal{M}'$ where all labels not in the set $\{i_1, \ldots, i_{n+1}\}$ are removed (by Lemma 2). But on $\mathcal{M}'$, φ and φ' are equivalent, so $\mathcal{M}', w \Vdash \varphi$.

PSPACE-hardness follows from PSPACE-completeness of multi-modal **K** and the fact that $\mathsf{K}_{\#}$ includes multi-modal **K**. ■

4 Completeness for $\mathsf{DK}_\#$

The above proof cannot be turned into a completeness proof for $\mathsf{DK}_\#$ in a straightforward way. If we simply replace in the definition of finite canonical models $\mathsf{K}_\#$-consistency by $\mathsf{DK}_\#$-consistency, we will not be able to prove that models so defined are deterministic, as the following example shows.

EXAMPLE 14. Consider the formula $\varphi = p \wedge \langle i \rangle\, q$ and the finite canonical model $\mathcal{M}^\varphi$ over φ. Then, among the points of $\mathcal{M}^\varphi$ (that is, among $\mathsf{DK}_\#$-atoms over $p \wedge \langle i \rangle\, q$) are $A = \{p, q, \langle i \rangle\, q, p \wedge \langle i \rangle\, q\}$ and $A' = \{\neg p, q, \langle i \rangle\, q, \neg(p \wedge \langle i \rangle\, q)\}$. Then, both $\widehat{A} \wedge \langle i \rangle\, \widehat{A}$ and $\widehat{A} \wedge \langle i \rangle\, \widehat{A'}$ are $\mathsf{DK}_\#$-consistent, which means that $A \mathcal{R}^\varphi_i A$ and $A \mathcal{R}^\varphi_i A'$.

Nevertheless, we will be able to reshape finite canonical models for $\mathsf{DK}_\#$ into deterministic models. First, we need to slightly adjust the definition of closure from the previous section. Given a subformula ψ of φ, we call the number of modal operators whose scope contains ψ the modal depth of ψ within φ, symbolically $\mathsf{md}_\varphi(\psi)$. Now, by the deterministic closure of the set of formulas Σ, we mean the smallest set $\mathsf{DCL}(\Sigma)$ containing (1) all the subformulas of formulas of Σ, (2) their pseudo-negations, and (3) for every $\varphi \in \Sigma$ and ψ such that $\mathsf{md}_\varphi(\psi) > 0$, if index i occurs in Σ, then $\langle i \rangle\, \psi \in \mathsf{DCL}(\Sigma)$ and $\langle i \rangle\, \sim \psi \in \mathsf{DCL}(\Sigma)$. It is easy to see that $\mathsf{DCL}(\Sigma)$ is finite whenever Σ is. In this section, Σ-atoms are taken to be maximally $\mathsf{DK}_\#$-consistent subsets of $\mathsf{DCL}(\Sigma)$.

DEFINITION 15. Let Σ be a finite set of $\mathcal{L}^I_\#$-formulas. The finite canonical model over Σ, $\mathcal{M}^\Sigma$, is the tuple $(At(\Sigma), \{\mathcal{R}^\Sigma_i\}_{i \in I}, \mathcal{R}^\Sigma_\#, V^\Sigma)$, where

1. $At(\Sigma)$ is the set of all atoms over Σ;
2. $A \mathcal{R}^\Sigma_i A'$ iff $\widehat{A} \wedge \langle i \rangle\, \widehat{A'}$ is consistent;
3. $A \mathcal{R}^\Sigma_\# A'$ iff $\widehat{A} \wedge \langle \# \rangle\, \widehat{A'}$ is consistent;
4. For every $p \in \mathbf{Par}$, $V^\Sigma(p) = \{\, A \in At(\Sigma) : p \in A \,\}$.

Proceeding exactly as in the completeness proof for $\mathsf{K}_\#$, we get the following two lemmas.

LEMMA 16. *Let Σ be a set of $\mathcal{L}^I_\#$ formulas, $\mathcal{M}^\Sigma$ be the finite canonical model for $\mathsf{DK}_\#$ over Σ, and $\psi \in \mathsf{DCL}(\Sigma)$. Then, for every $A \in At(\Sigma)$, $\mathcal{M}^\Sigma, A \Vdash \psi$ iff $\psi \in A$.*

LEMMA 17. *Every finite canonical model $\mathcal{M}^\Sigma = (At(\Sigma), \{\mathcal{R}^\Sigma_i\}_{i \in I}, \mathcal{R}^\Sigma_\#, V^\Sigma)$ is an $\mathcal{L}_\#$-model.*

Now, we have to reshape $\mathcal{M}^\Sigma$ into a deterministic model. We do so in two stages: first, we get rid of non-determinism with respect to indices not occurring in Σ, and then with respect to indices occurring in Σ. The first stage is easy.

LEMMA 18. *Let* $\mathcal{M}^\Sigma = (At(\Sigma), \{\mathcal{R}^\Sigma_i\}_{i\in I}, \mathcal{R}^\Sigma_\#, V^\Sigma)$ *be a finite canonical model for* $\mathsf{DK}_\#$ *. Then, there exists a model* ${\mathcal{M}'}^\Sigma = (At(\Sigma), \{{\mathcal{R}'}^\Sigma_i\}_{i\in I}, \mathcal{R}^\Sigma_\#, V^\Sigma)$ *such that (1) for every* $i \notin \Sigma$ *and every* $A, B, B' \in At(\Sigma)$*, if* $A{\mathcal{R}'}^\Sigma_i B$ *and* $A{\mathcal{R}'}^\Sigma_i B'$ *then* $B = B'$*, (2) the number of* R_i *edges with* $i \notin \Sigma$*, is at most* $|At(\Sigma)|^2$*, and (3) for every* $\psi \in \mathsf{DCL}(\Sigma)$ *and every* $X \in At(\Sigma)$*,* ${\mathcal{M}'}^\Sigma, X \Vdash \psi$ *iff* $\mathcal{M}^\Sigma, X \Vdash \psi$*.*

Proof. First, note that, by definition 15, if $A\mathcal{R}^\Sigma_\# B$ then $A\mathcal{R}^\Sigma_i B$ holds for every i not occurring in Σ. For every pair of atoms A and B connected by $\mathcal{R}^\Sigma_\#$, choose a single fresh link i_{AB} between them not occurring in Σ, and rename all other non-Σ links between A and B to be i_{AB}. The renaming of the links cannot affect the truth values of formulas in $\mathsf{DCL}(\Sigma)$, as the renamed links are indexed by indices not occurring in Σ and hence $\mathsf{DCL}(\Sigma)$. ■

At the second stage, we proceed as follows. We will take a submodel of ${\mathcal{M}'}^\varphi$ generated by the atom A_φ containing φ, unravel this submodel into a tree-like model with root A_φ and then prune the resultant tree, leaving only one $\mathcal{R}_i$ branch for every $i \in I$. We show that such a model still satisfies φ, since φ cannot tell apart the points on the branch we leave in the tree from the pruned ones. We need the versions of tree-likeness and unravelling that are slightly different from the standard ones.

DEFINITION 19. Let $\mathcal{M} = (W, \{\mathcal{R}_i\}_{i\in I}, \mathcal{R}_\#, V)$ be a $\mathcal{L}^I_\#$-model. $\mathcal{M}$ is tree-like if the structure $(W, \mathcal{R}_\#)$ is an irreflexive tree. $\mathcal{M}$ is strongly tree-like if $\mathcal{M}$ is tree-like and, for every $(w, v) \in \mathcal{R}_\#$, there exists exactly one $i \in I$ such that $(w, v) \in \mathcal{R}_i$.

THEOREM 20. *Let* $\mathcal{M} = (W, \{\mathcal{R}_i\}_{i\in I}, \mathcal{R}_\#, V)$ *be a rooted* $\mathcal{L}^I_\#$*-model with root* w*. Then, there exists a strongly tree-like* $\mathcal{L}^I_\#$*-model* $\mathcal{M}^\mathcal{T} = (W^T, \{\mathcal{R}^T_i\}_{i\in I}, \mathcal{R}^T_\#, V^T)$ *with root* w *such that, for every* $\mathcal{L}^I_\#$*-formula* φ*,* $\mathcal{M}, w \Vdash \varphi$ *iff* $\mathcal{M}^\mathcal{T}, w \Vdash \varphi$*.*

Proof. First, consider model $\mathcal{M}' = (W', \{\mathcal{R}'_i\}_{i\in I}, \mathcal{R}'_\#, V')$, where

1. W is the set of all possible sequences of the form $(w, w_1^{i_1}, \ldots, w_n^{i_n})$, where $w_1, \ldots, w_{n\geq 0} \in W$ and $i_1, \ldots, i_n \in I$;
2. $(w, w_1^{i_1}, \ldots, w_n^{i_n})\mathcal{R}'_j(w, w_1^{i_1}, \ldots, w_n^{i_n}, w_{n+1}^{i_{n+1}})$ if $w_n\mathcal{R}_j w_{n+1}$ and $j = i_{n+1}$;
3. $\mathcal{R}'_\# = \bigcup_{i\in I} \mathcal{R}'_i$;
4. $V'(p) = \{\, (w, w_1^{i_1}, \ldots, w_n^{i_n}) : w_n \in V(p)\,\}$, for every $p \in \mathbf{Par}$.

Next, take the submodel of $\mathcal{M}'$ generated by w. Let us call it $\mathcal{M}^\mathcal{T}$. It is clear that $\mathcal{M}^\mathcal{T}$ is strongly tree-like (the last member of the sequence serving as the second argument of each $\mathcal{R}'_i$ bears exactly one superscript). The truth-preservation is guaranteed by the existence of a bisimulation $Z \subseteq W \times W^T$ defined by $vZ(w, w_1^{i_1}, \ldots, w_n^{i_n})$ iff $w_n = v$, which connects the roots of the two models. ■

Now we show that, in tree-like models, for every formula φ, the value of φ at the root does not change if we replace a point v accessible from the root in k steps with another point v' such that v and v' agree on all the subformulas of φ of modal depth k. (In the statement of the following lemma, we use $w\mathcal{R}_{\#}^{k}v$ to mean that there are $u_1, \ldots, u_{k-1}$ such that $w\mathcal{R}_{\#}u_1\mathcal{R}_{\#}\ldots\mathcal{R}_{\#}u_{k-1}\mathcal{R}_{\#}v$; in particular, $w\mathcal{R}_{\#}^{0}v$ means that $w = v$. $\mathsf{Sub}(\varphi)$ stands for the set of all subformulas of φ.)

LEMMA 21. *Let φ be a $\mathcal{L}_{\#}^{I}$-formula, $\mathcal{M} = (W, \{\mathcal{R}_i\}_{i\in I}, \mathcal{R}_{\#}, V)$ a tree-like $\mathcal{L}_{\#}^{I}$-model, $w \in W$, and $v \in W$ such that $w\mathcal{R}_{\#}^{k}v$. Let $\mathcal{M}'$ be obtained from $\mathcal{M}$ by replacing the subtree generated by v by another subtree, with root v', such that, for every $\psi \in \mathsf{Sub}(\varphi)$ with $\mathsf{md}_{\varphi}(\psi) = k$, $\mathcal{M}, v \Vdash \psi$ iff $\mathcal{M}', v' \Vdash \psi$. Then, $\mathcal{M}, w \Vdash \varphi$ iff $\mathcal{M}', w' \Vdash \varphi$.*

Proof. By induction on k

Let $k = 0$. Then, $w = v$. Moreover, v and v' agree on all $\psi \in \mathsf{Sub}(\varphi)$ with $\mathsf{md}_{\varphi}(\psi) = 0$. As $\mathsf{md}_{\varphi}(\varphi) = 0$, w and v' agree on φ.

Assume that the statement of the lemma is true for $k = n$. We show that it is true for $k = n + 1$. Suppose that it is not. Then, v and v' agree on all $\psi \in \mathsf{Sub}(\varphi)$ with $\mathsf{md}_{\varphi}(\psi) = n + 1$ and $\mathcal{M}, w \Vdash \varphi$, but $\mathcal{M}', w' \not\Vdash \varphi$ (the other case is symmetrical). Since no changes have been made to w itself, φ should have a subformula $\langle i \rangle\, \chi$ with $\mathsf{md}_{\varphi}(\langle i \rangle\, \chi) = 0$ such that, for some u such that $w\mathcal{R}_i u$ and $u \in path(w, v)$, $\mathcal{M}, u \Vdash \chi$ but $\mathcal{M}', u \not\Vdash \chi$ (the other case is symmetrical). Now, $\mathsf{md}_{\varphi}(\chi) = \mathsf{md}_{\varphi}(\langle i \rangle\, \chi) + 1$ and $\mathsf{Sub}(\chi) \subseteq \mathsf{Sub}(\varphi)$; therefore, v and v' agree on all $\psi \in \mathsf{Sub}(\chi)$ with $\mathsf{md}_{\chi}(\psi) = n$. As $u\mathcal{R}_{\#}^{n}v$, applying the inductive hypothesis to the tree generated by u, we get $\mathcal{M}, u \Vdash \chi$ iff $\mathcal{M}', u \Vdash \chi$, a contradiction. ■

LEMMA 22. *Let $\mathcal{M}^{\varphi T}$ be a strongly tree-like model obtained from the canonical model over φ, $\mathcal{M}^{\varphi}$, by unravelling the submodel of $\mathcal{M}^{\varphi}$ generated by an atom A_{φ} containing φ. Then, for every $B, B' \in \mathcal{M}^{\varphi T}$ such that, for some C, $C\mathcal{R}_i B$ and $C\mathcal{R}_i B'$, and every ψ such that $\mathsf{md}_{\varphi}(\psi) > 0$, we have $\mathcal{M}^{\varphi T}, B \Vdash \psi$ iff $\mathcal{M}^{\varphi T}, B' \Vdash \psi$.*

Proof. Assume that there exist B and B' such that $C\mathcal{R}_i B$ and $C\mathcal{R}_i B'$, $\mathcal{M}^{\varphi T}, B \Vdash \psi$, and $\mathcal{M}^{\varphi T}, B' \not\Vdash \psi$. Then, $\mathcal{M}^{\varphi T}, C \Vdash \langle i \rangle\, \psi$ and $\mathcal{M}^{\varphi T}, C \Vdash \langle i \rangle \sim\!\psi$. Therefore, since $\langle i \rangle\, \psi, \langle i \rangle \sim\!\psi \in \mathsf{DCL}(\varphi)$, by lemma 16, $\langle i \rangle\, \psi, \langle i \rangle \sim\!\psi \in C$. This, however, is impossible, since by axiom (F), $\vdash_{\mathsf{DK}_{\#}} \langle i \rangle\, \psi, \langle i \rangle \sim\!\psi \to \perp$. ■

Now we can prove the completeness theorem.

THEOREM 23. *$\mathsf{DK}_{\#}$ is weakly complete with respect to the class of deterministic $\mathcal{L}_{\#}$-models.*

Proof. Let φ be a $\mathsf{DK}_{\#}$-consistent formula. Build the finite canonical model $\mathcal{M}^{\varphi}$ over φ. There is in $\mathcal{M}^{\varphi}$ an atom A_{φ} such that $\varphi \in A_{\varphi}$. By lemma 16, $\mathcal{M}^{\varphi}, A_{\varphi} \Vdash \varphi$. Remove, using the construction of lemma 18, all the "redundant" atomic links in $\mathcal{M}^{\varphi}$ indexed by i not occurring in Σ. Next, unravel $\mathcal{M}'^{\varphi}$ into a strongly tree-like model $\mathcal{M}'^{\varphi T}$ using the construction of theorem 20. Now, level by level, for every point C and label i at level n such that C can reach several points $B_1, \ldots, B_m$ by an edge labelled i, replace all B_js by B_1. Denote the resultant model by $\mathcal{M}'^{\varphi T'}$. By the lemmas above, $\mathcal{M}'^{\varphi T'}, A_\varphi \Vdash \varphi$. Lastly, construct $\mathcal{M}'^{\varphi T''}$ by replacing all identical copies of B_1 produced in construction of $\mathcal{M}'^{\varphi T'}$ by a single point B_1. $\mathcal{M}'^{\varphi T'}$ and $\mathcal{M}'^{\varphi T''}$ are obviously bisimilar, so $\mathcal{M}'^{\varphi T''}, A_\varphi \Vdash \varphi$. It is clear that $\mathcal{M}'^{\varphi T''}$ is deterministic. ■

The model for φ we have constructed in the proof of Theorem 23 is possibly infinite (it is an unravelling of a possibly cyclic model $\mathcal{M}^{\varphi}$ of size $2^{|\varphi|}$). We can however make it finite by pruning the tree at depth k, where k is the maximal depth of nesting of modal operators in φ (note that k is bound by $|\varphi|$). The branching factor of the tree is $|\varphi| + 2^{|\varphi|}$ (from every node A, there are at most $|\varphi|$ edges using labels from φ, and at most $2^{|\varphi|}$ fresh links - at most one for every other node B in the original canonical model). Given the branching factor and the depth of the finite tree model, the maximal number of nodes there is $(|\varphi| + 2^{|\varphi|})^{|\varphi|}$. This effective bounded model property for $\mathsf{DK}_{\#}$ gives us the following theorem:

THEOREM 24. *$\mathsf{DK}_{\#}$ is decidable.*

The best upper bound we have is NEXPTIME (guess a model for φ of size at most $(|\varphi| + 2^{|\varphi|})^{|\varphi|}$, which is $O(2^{|\varphi|^2})$, and verify that it satisfies φ).

5 Logic PDL^{path}

The language of PDL^{path} is an extension of the language of PDL, propositional dynamic logic. The language of PDL has two kinds of primitive symbols: propositional parameters and atomic transitions (or, modality indices). Indices are used to label edges in the transition system. Compound path expressions are built out of indices using binary operators $\circ$ (composition), $\cup$ (union) and a unary operator $*$ (finite iteration). In addition to these, the language of PDL^{path}, introduced in [3], has the modal identity constant id, the unary converse operator $\cdot^{-}$ and the wild card modality $\#$. Moreover, the language of PDL^{path} has a single nominal (a propositional letter that is true at exactly one point of a model) r, which is meant to mark the root of the graph. In the literature, PDL with the converse operator is referred to as converse PDL or CPDL. Thus, PDL^{path} is a fragment (since we have only one nominal) of CPDL with nominals augmented with the existential modality $\#$.

In this paper, we give a complete Hilbert-style axiomatisation for PDL^{path}. To that end, we need to extend the language of PDL^{path} as introduced in [3] with the "at" modality @ of hybrid logics, which we will need to axiomatically describe the behaviour of the nominal r.

DEFINITION 25. Given a countable set of indices $I = \{i_1, i_2, \ldots, i_n, \ldots\}$, path expressions over I are defined by the following BNF expression:

$$\Lambda_I := I \mid id \mid \# \mid \Lambda_I \circ \Lambda_I \mid \Lambda_I \cup \Lambda_I \mid \Lambda_I^* \mid \Lambda_I{}^-.$$

PDL^{path}-formulas over the set of path expressions Λ_I are defined as follows:

$$\varphi := \top \mid \bot \mid r \mid \neg\varphi \mid \varphi \vee \varphi \mid \langle \Lambda_I \rangle \varphi \mid @_r\varphi.$$

PDL^{path}-formulas are evaluated on path models.

DEFINITION 26. A path model $\mathcal{M}$ over the set of labels Λ_I is a tuple $(W, \{\mathcal{R}_\pi\}_{\pi\in\Lambda_I}, V)$, where

1. $W \neq \emptyset$;
2. V is a function assigning some $\{w\} \subseteq W$ to r.
3. $\{\mathcal{R}_\pi\}_{\pi\in\Lambda_I}$ is a collection of binary relations over W satisfying the following conditions:
   (a) $\mathcal{R}_\# = \bigcup_{i\in I} \mathcal{R}_i$;
   (b) $\mathcal{R}_{id} = \{\, (w, w) : w \in W \,\}$ (identity relation);
   (c) $\mathcal{R}_{\pi^-} = \mathcal{R}_\pi^-$ (converse);
   (d) $\mathcal{R}_{\pi_1\circ\pi_2} = \mathcal{R}_{\pi_1} \circ \mathcal{R}_{\pi_2}$ (composition);
   (e) $\mathcal{R}_{\pi_1\cup\pi_2} = \mathcal{R}_{\pi_1} \cup \mathcal{R}_{\pi_2}$ (union);
   (f) $\mathcal{R}_{\pi^*} = \mathcal{R}_\pi^*$ (reflexive-transitive closure);
   (g) For every $w, v \in W$, there is a sequence of points $u_1, \ldots, u_n$ such that (1) $w = u_1$, (2) $v = u_n$, and (3) for every $1 \leq i \leq n-1$, either, for some $i \in I$, $u_i\mathcal{R}_iu_{i+1}$, or, for some $i \in I$, $u_{i+1}\mathcal{R}_iu_i$ (connectedness).

The truth of PDL^{path}-formulas at a point in a path model is defined as follows.

DEFINITION 27. Let $\mathcal{M} = (W, \{R_\pi\}_{\pi\in\Lambda_I}, V)$ be a path model, $w, v \in W$. Then,

$\mathcal{M}, w \Vdash \top$		always;
$\mathcal{M}, w \Vdash \bot$		never;
$\mathcal{M}, w \Vdash r$	iff	$V(r) = \{w\}$;
$\mathcal{M}, w \Vdash \neg\varphi$	iff	$\mathcal{M}, w \nVdash \varphi$;
$\mathcal{M}, w \Vdash \varphi \vee \psi$	iff	$\mathcal{M}, w \Vdash \varphi$ or $\mathcal{M}, w \Vdash \psi$;
$\mathcal{M}, w \Vdash \langle\pi\rangle\varphi$	iff	for some $v \in W, w\mathcal{R}_\pi v$ and $\mathcal{M}, v \Vdash \varphi$;
$\mathcal{M}, w \Vdash @_r\varphi$	iff	$\mathcal{M}, v \Vdash \varphi$ and $V(r) = \{v\}$.

Here are some examples of properties definable in PDL^{path}: r defines the root; $\neg\langle\#\rangle\top$ defines leaf nodes; $\langle(\# \cup \#^-)^*\rangle r$ defines nodes connected to the root. To express a path constraint $\pi_1 \subseteq \pi_2$ (everything reachable from the root by a path π_1 is reachable by a path π_2), we can say $@_r[\pi_1]\langle\pi_2{}^-\rangle r$. Note that on connected graphs, $@_r\varphi$ is definable as $\langle(\# \cup \#^-)^*\rangle(r \wedge \varphi)$.

In [3], it was proved that PDL^{path} is decidable (the proof is similar to the proof of Theorem 13, reducing the satisfiability problem for PDL^{path} to the satisfiability problem for CPDL with nominals. CPDL with nominals is decidable in EXPTIME [5].

Now, we describe a Hilbert-style axiomatisation of PDL^{path}. Axiom schemata of PDL^{path} can be logically divided into four parts.

The first part describes the behaviour of propositional connectives and conventional modal operators $\langle\pi\rangle$ and $[\,\pi\,]$:

(A0) all classical tautologies;

(K) $[\,\pi\,]\,(\varphi \to \psi) \to ([\,\pi\,]\,\varphi \to [\,\pi\,]\,\psi)$;

(A1) $\langle\pi\rangle\,\varphi \leftrightarrow \neg[\,\pi\,]\,\neg\varphi$.

The second part describes the properties of path expression operators:

(A2) $\langle\pi_1 \circ \pi_2\rangle\,\varphi \leftrightarrow \langle\pi_1\rangle\,\langle\pi_2\rangle\,\varphi$;

(A3) $\langle\pi_1 \cup \pi_2\rangle\,\varphi \leftrightarrow \langle\pi_1\rangle\,\varphi \vee \langle\pi_2\rangle\,\varphi$;

(A4) $\langle\pi^*\rangle\,\varphi \leftrightarrow \varphi \vee \langle\pi\rangle\,\langle\pi^*\rangle\,\varphi$;

(A5) $[\pi^*](\varphi \to [\,\pi\,]\,\varphi) \to (\varphi \to [\pi^*]\varphi)$;

(A6) $\varphi \to [\,\pi^-\,]\,\langle\pi\rangle\,\varphi$;

(A7) $\varphi \to [\,\pi\,]\,\langle\pi^-\rangle\,\varphi$;

(A8) $\varphi \leftrightarrow \langle id\rangle\,\varphi$;

(ER) $\langle i\rangle\varphi \to \langle\#\rangle\varphi$.

The third part describes properties of $@_r$ operator:

(A9) $@_r(\varphi \to \psi) \to (@_r\varphi \to @_r\psi)$;

(A10) $@_r\varphi \leftrightarrow \neg @_r\neg\varphi$;

(A11) $r \wedge \varphi \to @_r\varphi$;

(A12) $@_r r$;

(A13) $\langle\pi\rangle\,@_r\varphi \to @_r\varphi$.

Finally, the following axiom pertains to connectedness:

(A14) $\langle(\# \cup \#^-)^*\rangle\, r$.

The inference rules are:

(MP) $\dfrac{\vdash \varphi \to \psi, \vdash \varphi}{\vdash \psi}$; **(N)** $\dfrac{\vdash \varphi}{\vdash [\pi]\varphi}$; **(NN)** $\dfrac{\vdash \varphi}{\vdash @_r\varphi}$;

(EL) $\dfrac{\vdash \langle i\rangle\varphi \to \psi}{\vdash \langle \#\rangle\varphi \to \psi}$, provided i does not occur in ψ.

In addition to the above axiom schemata and rules of inference, in the course of the following completeness proof, we will appeal to two additional rules of inference pertaining to the converse operator, whose derivability we establish in the following lemma.

LEMMA 28. *The following rules of inference are derivable in* PDL^{path} *:*

$$\frac{\vdash \varphi \to [\,\pi\,]\,\neg\psi}{\vdash \psi \to [\pi^-]\neg\varphi}; \quad \frac{\vdash \varphi \to [\pi^-]\neg\psi}{\vdash \psi \to [\,\pi\,]\,\neg\varphi}.$$

Proof. The first rule can be derived as follows.

1. $\varphi \to [\,\pi\,]\,\neg\psi$ – premise
2. ψ – assumption
3. $[\,\pi^-\,]\,(\varphi \to [\,\pi\,]\,\neg\psi)$ – by (N) from 1
4. $\psi \to [\,\pi^-\,]\,\langle\pi\rangle\,\psi$ – (A6)
5. $[\,\pi^-\,]\,\langle\pi\rangle\,\psi$ – by (MP) from 2, 4
6. $[\,\pi^-\,]\,(\langle\pi\rangle\,\psi \wedge (\varphi \to [\,\pi\,]\,\neg\psi))$ – from 3, 5 by (K)
7. $[\,\pi^-\,]\,(\neg\varphi \vee (\langle\pi\rangle\,\psi \wedge [\,\pi\,]\,\neg\psi))$ – by (K) and propositional reasoning from 6
8. $[\,\pi^-\,]\,\neg\varphi$ – by (A1) and propositional reasoning from 7
9. $\psi \to [\,\pi^-\,]\,\neg\varphi$ – from 2, 8.

The second rule can be derived analogously, relying on axiom (A7). ■

6 Completeness for PDL^{path}

In this section, we prove completeness of the above axiomatisation of PDL^{path} (its soundness is straightforward). As the language of PDL^{path} contains $\langle\#\rangle$ and $\langle\pi^*\rangle$, both of which give rise to non-compact logics, we can only prove weak completeness for PDL^{path} . As in the completeness proofs for $\mathsf{K}_\#$ and $\mathsf{DK}_\#$, we use the completeness-via-finite-models technique.

DEFINITION 29. Let Σ be a set of PDL^{path} -formulas over Λ_I. The closure of Σ, $\mathsf{CL}(\Sigma)$, is the smallest set such that

- if $\varphi \in \Sigma$ then $\mathsf{Sub}(\varphi) \subseteq \mathsf{CL}(\Sigma)$;
- if $\langle\pi^-\rangle\varphi \in \Sigma$ then $[\,\pi\,]\,\langle\pi^-\rangle\varphi \in \mathsf{CL}(\Sigma)$ (here and below, π ranges over all path labels);

- if $\langle \pi_1 \circ \pi_2 \rangle \varphi \in \mathsf{CL}(\Sigma)$ then $\langle \pi_1 \rangle \langle \pi_2 \rangle \varphi \in \mathsf{CL}(\Sigma)$;
- if $\langle \pi_1 \cup \pi_2 \rangle \varphi \in \mathsf{CL}(\Sigma)$ then $\langle \pi_1 \rangle \varphi \vee \langle \pi_2 \rangle \varphi \in \mathsf{CL}(\Sigma)$;
- if $\langle \pi^* \rangle \varphi \in \mathsf{CL}(\Sigma)$ then $\langle \pi \rangle \langle \pi^* \rangle \varphi \in \mathsf{CL}(\Sigma)$;
- if $\psi \in \mathsf{CL}(\Sigma)$ and $\psi \neq @_r\chi$ and $\psi \neq \neg @_r\chi$, then $@_r\psi \in \mathsf{CL}(\Sigma)$;
- $@_r r \in \mathsf{CL}(\Sigma)$;
- $\langle (\# \cup \#^-)^* \rangle r \in \mathsf{CL}(\Sigma)$;
- if $\varphi \in \mathsf{CL}(\Sigma)$, then $\sim \varphi \in \mathsf{CL}(\Sigma)$.

LEMMA 30. *Let Σ be a set of PDL^{path}-formulas. If Σ is finite, then $\mathsf{CL}(\Sigma)$ is finite, too.*

PDL^{path}-atoms are defined exactly as $\mathsf{K}_\#$-atoms. It is easy to show that, in addition to the properties satisfied by $\mathsf{K}_\#$-atoms, PDL^{path}-atoms have the following ones:

- for all $\langle \pi^- \rangle \varphi \in \mathsf{CL}(\Sigma)$, if $\varphi \in A$ then $[\,\pi\,] \langle \pi^- \rangle \varphi \in A$;
- for all $\langle \pi_1 \circ \pi_2 \rangle \varphi \in \mathsf{CL}(\Sigma)$, $\langle \pi_1 \circ \pi_2 \rangle \varphi \in A$ iff $\langle \pi_1 \rangle \langle \pi_2 \rangle \varphi \in A$;
- for all $\langle \pi_1 \cup \pi_2 \rangle \varphi \in \mathsf{CL}(\Sigma)$, $\langle \pi_1 \cup \pi_2 \rangle \varphi \in A$ iff $\langle \pi_1 \rangle \varphi \vee \langle \pi_2 \rangle \varphi \in A$;
- for all $\langle \pi^* \rangle \varphi \in \mathsf{CL}(\Sigma)$, $\langle \pi^* \rangle \varphi \in A$ iff $\langle \pi \rangle \langle \pi^* \rangle \varphi \in A$;
- for all $\langle id \rangle \varphi \in \mathsf{CL}(\Sigma)$, $\langle id \rangle \varphi \in A$ iff $\varphi \in A$.

LEMMA 31. *If $\varphi \in \mathsf{CL}(\Sigma)$ is PDL^{path}-consistent, then there exists an atom A over Σ such that $\varphi \in A$.*

Now we define the finite canonical PDL^{path}-model over Σ.

DEFINITION 32. Let Σ be a finite set of PDL^{path}-formulas over the set of path expressions Λ_I and let a be an index such that $a \in I$ but a does not occur in $\mathsf{CL}(\Sigma)$. First, define a family of binary relations $\{S_\pi\}$ on the set $At(\Sigma)$ of atoms over Σ, as follows:

- For all atoms $A, A' \in At(\Sigma)$, $AS_\pi A'$ iff $\pi \in \mathsf{CL}(\Sigma)$ or $\pi = a$, and $\widehat{A} \wedge \langle \pi \rangle \widehat{A'}$ is consistent.

Now, the finite canonical model $\mathcal{M}^\Sigma$ over Λ_I is a tuple $(W^\Sigma, \{\mathcal{R}^\Sigma_\pi\}_{\pi \in \Lambda_I}, V^\Sigma)$ such that

1. $W = At(\Sigma)$;
2. $V(r) = \{\, A \in At(\Sigma) : r \in A \,\}$;
3.
   - for every atomic c such that $c \in \mathsf{CL}(\Sigma)$ or $c = a$, $\mathcal{R}^\Sigma_c = S_c$;
   - $\mathcal{R}^\Sigma_\# = S_\#$;

- $\mathcal{R}^{\Sigma}_{id} = \{ (A, A) : A \in At(\Sigma) \}$;
- $\mathcal{R}^{\Sigma}_{\rho^-} = (\mathcal{R}^{\Sigma}_{\rho})^{-}$;
- $\mathcal{R}^{\Sigma}_{\pi_1 \circ \pi_2} = \mathcal{R}^{\Sigma}_{\pi_1} \circ \mathcal{R}^{\Sigma}_{\pi_2}$;
- $\mathcal{R}^{\Sigma}_{\pi_1 \cup \pi_2} = \mathcal{R}^{\Sigma}_{\pi_1} \cup \mathcal{R}^{\Sigma}_{\pi_2}$;
- $\mathcal{R}^{\Sigma}_{\pi^*} = (\mathcal{R}^{\Sigma})^{*}_{\pi}$.

It is easy to see that finite canonical models for PDL^{path} satisfy conditions (3a)–(3f) required by definition 26 of path models (indeed, conditions (3b)–(3f) are satisfied because of definition 32, and condition (3a) can be shown to be satisfied in the same way as in the proof of lemma 10); thus, finite canonical models are regular. This is enough to prove the existence lemma and the truth lemma for finite canonical models.

To prove the existence lemma for finite canonical models, we first need to show that, for every $\pi \in \Lambda_I$, $S_\pi \subseteq \mathcal{R}^{\Sigma}_{\pi}$.

LEMMA 33. *For every* $\pi \in \Lambda_I$, $S_\pi \subseteq \mathcal{R}^{\Sigma}_{\pi}$.

Proof. By induction on the complexity of π.

(0) The cases $\pi \in I$ and $\pi = \#$ are obvious, since for $\pi \in I \cup \{\#\}$, $\mathcal{R}^{\Sigma}_{\pi} = S_\pi$.

(1) Let π be id. Suppose that $AS_{id}B$, that is $\widehat{A} \wedge \langle id \rangle \widehat{B}$ is consistent. By (A8), $\widehat{A} \wedge \widehat{B}$ is consistent. Since both A and B are atoms, this is only possible if $A = B$. Therefore, $A\mathcal{R}^{\Sigma}_{id}B$.

(2) Let π be ρ^-. Suppose that $AS_{\rho^-}B$, that is $\widehat{A} \wedge \langle \rho^- \rangle \widehat{B}$ is consistent. This implies consistency of $\widehat{B} \wedge \langle \rho \rangle \widehat{A}$. Indeed, if we suppose otherwise, then $\vdash \widehat{B} \rightarrow \neg \langle \rho \rangle \widehat{A}$ and hence $\vdash \widehat{B} \rightarrow [\rho] \neg \widehat{A}$. Then, by lemma 28, $\vdash \widehat{A} \rightarrow [\rho^-] \neg \widehat{B}$, which means that, contrary to the assumption, $\widehat{A} \wedge \langle \rho^- \rangle \widehat{B}$ is inconsistent. Thus, $\widehat{B} \wedge \langle \rho \rangle \widehat{A}$ is consistent and hence $BS_\rho A$. By the inductive hypothesis, $B\mathcal{R}^{\Sigma}_{\rho}A$ and therefore $A\mathcal{R}^{\Sigma}_{\rho^-}B$, as required.

(3)-(5) The other cases are proved exactly as for PDL. ■

LEMMA 34 (Existence lemma). *Let* Σ *be a set of* PDL^{path}*-formulas over* Λ_I*,* A *be an atom over* Σ*, and* $\pi \in \Lambda_I$*. Then, for all* $\langle \pi \rangle \psi \in \mathsf{CL}(\Sigma)$*,* $\langle \pi \rangle \psi \in A$ *iff there is an atom* A' *such that* $A\mathcal{R}^{\Sigma}_{\pi}A'$ *and* $\psi \in A'$.

Proof. The left-to-right direction can be proved using the standard "forcing choices" technique: picking, for every $\psi \in \mathsf{CL}(\Sigma)$, either ψ itself of its pseudo-negation (in a consistency-preserving way), build an atom A' such that $\widehat{A} \wedge \langle \pi \rangle \widehat{A'}$ is consistent. Then, by lemma 33, $A\mathcal{R}^{\Sigma}_{\pi}A'$. The right-to-left direction is proved by induction on the complexity of π.

(0) $\pi \in I$. Suppose that there is an atom A' such that $\varphi \in A'$ and $A\mathcal{R}_\pi A'$. Then, by definition 32, $AS_\pi A'$, which means that $\widehat{A} \wedge \langle \pi \rangle \widehat{A'}$ is consistent. Then, as $\varphi \in A'$ and, thus, φ is one of the conjuncts of $\widehat{A'}$, $\widehat{A} \wedge \langle \pi \rangle \varphi$ is consistent, too. Then, as $\langle \pi \rangle \varphi \in \mathsf{CL}(\Sigma)$ and A is an atom, $\langle \pi \rangle \varphi \in A$.

(1) $\pi = \#$. Analogously to (0).

(2) $\pi = \rho^-$. Suppose that $A\mathcal{R}^\Sigma_{\rho^-}A'$ and $\psi \in A'$. Then, $A'\mathcal{R}^\Sigma_\rho A$. As we know, $[\,\rho\,]\,\langle\rho^-\rangle\psi \in A'$. But this implies $\langle\rho^-\rangle\psi \in A$; indeed, if we suppose otherwise then $\neg\langle\rho^-\rangle\psi \in A$ and so, by inductive hypothesis, $\langle\rho\rangle\,\neg\langle\rho^-\rangle\psi \in A'$, which is impossible since then A' would be inconsistent.

(3)-(5) As for PDL. ■

It is now easy to prove the truth lemma.

LEMMA 35. *Let Σ be a set of PDL^{path}-formulas, $\mathcal{M}^\Sigma$ be the finite canonical model over Σ, and $\psi \in \mathsf{CL}(\Sigma)$. Then, for every $A \in At(\Sigma)$, $\mathcal{M}^\Sigma, A \Vdash \psi$ iff $\psi \in A$.*

What remains to be done is ensure that we can reshape $\mathcal{M}^\Sigma$ into a model with exactly one root in a truth-preserving way. To that end, we will show that, given an atom $A \in \mathcal{M}^\Sigma$, if we take a submodel $\mathcal{M}^\Sigma_A$ of $\mathcal{M}^\Sigma$ generated by A, then $\mathcal{M}^\Sigma_A$ contains at most one root. This will be enough to prove weak completeness for PDL^{path}, since axiom (A14) ensures that $\mathcal{M}^\Sigma_A$ contains *at least one* root.

First, note the following simple fact.

LEMMA 36. *Let $\mathcal{M}$ be a regular model and $w \in \mathcal{M}$. Then, the submodel of $\mathcal{M}$ generated by w is also regular.*

Next, we prove that all the atoms of the submodel of $\mathcal{M}^\Sigma$ generated by A agree on formulas beginning with $@_r$.

LEMMA 37. *Let A be an atom, $\mathcal{M}^\Sigma_A$ be a submodel of $\mathcal{M}^\Sigma$ generated by A, and B and B' be atoms such that $B, B' \in \mathcal{M}^\Sigma_A$. Then, for every $@_r\psi \in \mathsf{CL}(\Sigma)$, $@_r\psi \in B$ iff $@_r\psi \in B'$.*

Proof. Assume that $@_r\psi \in B$ and $@_r\psi \notin B'$ (the other case is symmetrical) and, hence $\neg@_r\psi \in B'$.

Notice that, for any two atoms $X, X' \in \mathcal{M}^\Sigma_A$, if $X\mathcal{R}^\Sigma_{i\in I}X'$ and $@_r\psi \in X'$, then $@_r\psi \in X$. Indeed, otherwise, $\neg@_r\psi \in X$, which is impossible since, on the one hand, by (A13), $\neg@_r\psi \wedge \langle i\rangle\, @_r\psi$ is inconsistent and hence $\widehat{X} \wedge \langle i\rangle\, \widehat{X'}$ is inconsistent, and on the other, by definition 32, $X\mathcal{R}^\Sigma_i X'$ holds only if $\widehat{X} \wedge \langle i\rangle\, \widehat{X'}$ is consistent. Analogously, for any two atoms $X, X' \in \mathcal{M}^\Sigma_A$, if $X\mathcal{R}^\Sigma_i X'$ and $\neg@_r\psi \in X'$, then $\neg@_r\psi \in X$. For otherwise, $@_r\psi \in X$, which is impossible since, on the one hand, by (A10) and (A13), $@_r\psi \wedge \langle i\rangle\,\neg@_r\psi$ is inconsistent and hence $\widehat{X} \wedge \langle i\rangle\, \widehat{X'}$ is inconsistent, and on the other, by definition 32, $X\mathcal{R}^\Sigma_a X'$ holds only if $\widehat{X} \wedge \langle i\rangle\, \widehat{X'}$ is consistent. From the foregoing, it also follows that, for any $X, X' \in \mathcal{M}^\Sigma_A$ such that $X'\mathcal{R}^\Sigma_i X$, if $@_r\psi \in X'$ then $@_r\psi \in X$ and $\neg@_r\psi \in X'$ then $\neg@_r\psi \in X$.

Since $\mathcal{M}^\Sigma$, and hence, by lemma 36, also $\mathcal{M}^\Sigma_A$ are regular, $B \in \mathcal{M}^\Sigma_A$ implies that there is a chain of atomic transitions $\mathcal{R}^\Sigma_i$ connecting A and B (so that, to reach B from A, we can move forward as well as backward along $\mathcal{R}^\Sigma_i$'s in the chain). It follows that from $@_r\psi \in B$ we can infer $@_r\psi \in A$ (using the argument of the preceding paragraphs, "pull back" $@_r\psi$ along the chain connecting A and B). Analogously, from $\neg@_r\psi \in B'$ we can infer $\neg@_r\psi \in A$. This is impossible since A is an atom. ■

Next, we can show that $\mathcal{M}^{\Sigma}_{A}$ has at most one root.

LEMMA 38. *Let A be an atom, $\mathcal{M}^{\Sigma}_{A}$ be a submodel of $\mathcal{M}^{\Sigma}$ generated by A, and B and B' be atoms such that (1) $B, B' \in \mathcal{M}^{\Sigma}_{A}$ and (2) $B \neq B'$. Then, at most one of B and B' contains r.*

Proof. Assume that $r \in B$ and $r \in B'$. Since $B \neq B'$, there is $\psi \in \mathsf{CL}(\Sigma)$ such that $\psi \in B$ and $\sim \psi \in B'$. There are two cases to consider: (1) $@_r\psi \in \mathsf{CL}(\Sigma)$ and (2) $@_r\psi \notin \mathsf{CL}(\Sigma)$ and, hence, either $\psi = @_r\chi$ or $\psi = \neg @_r\chi$.

(1) Suppose that $@_r\psi \in \mathsf{CL}(\Sigma)$, and hence, $\neg @_r\psi \in \mathsf{CL}(\Sigma)$. As $\psi \in B$ and $r \in B$, we also have $@_r\psi \in B$ (due to (A11), otherwise B would be inconsistent). Analogously, $\sim \psi \in B$ and $r \in B$ imply $\neg @_r\psi \in B'$. However, since $@_r\psi \in B$ (by lemma 37), we also have $@_r\psi \in B'$, which is impossible.

(2) Suppose that $@_r\psi \notin \mathsf{CL}(\Sigma)$ and, hence, either (2a) $\psi = @_r\chi$ or (2b) $\psi = \neg @_r\chi$. The case (2a) is analogous to the case (1), and the case (2b) is symmetrical. ■

Now we can show that $\mathcal{M}^{\Sigma}_{A}$ is a path model.

LEMMA 39. *Let A be an atom and $\mathcal{M}^{\Sigma}_{A}$ be a submodel of $\mathcal{M}^{\Sigma}$ generated by A. Then, $\mathcal{M}^{\Sigma}_{A}$ is a path model.*

Proof. By lemma 36, $\mathcal{M}^{\Sigma}_{A}$ is regular and, by lemma 38, it has no more than one root. Moreover, (A14) guarantees that it has at least one root. ■

The foregoing gives us the following theorem.

THEOREM 40. *PDL^{path} is complete with respect to the class of all path models.*

7 PDL^{path} without connectedness

Now, we consider what happens if we want to drop from the semantic definition of PDL^{path} the requirement that path models should be connected. It is easy to see that all we have to do to axiomatise PDL^{path} without connectedness is to drop from the above axiomatisation of PDL^{path} axiom (A14). Then, we can still show that every consistent formula has a model with *exactly one* root.

The only difference between the completeness proof for PDL^{path} and the completeness proof for PDL^{path} without connectedness is that, in the latter case, we cannot prove the analogue of lemma 39, as the following example shows.

EXAMPLE 41. Consider the formula $\varphi = \neg\langle(\# \cup \#^{-})^{*}\rangle\, r$. Since now path models are allowed to be unconnected, it is consistent, and hence, there is, in the finite canonical model $\mathcal{M}^{\{\varphi\}}$ over $\{\varphi\}$, an atom A_{φ} such that $\varphi \in A_{\varphi}$. It is easy to see that the submodel $\mathcal{M}^{\{\varphi\}_{A_{\varphi}}}$ of $\mathcal{M}^{\{\varphi\}}$ generated by A_{φ}, does not contain an atom B such that $r \in B$.

Nevertheless, as the following lemma shows, given a finite canonical model for PDL^{path} without connectedness $\mathcal{M}^\Sigma$ and an atom A, we can always reshape $\mathcal{M}^\Sigma_A$ into a path model.

LEMMA 42. *Let A be an atom and $\mathcal{M}^\Sigma_A$ be a submodel of $\mathcal{M}^\Sigma$ generated by A such that no $X \in \mathcal{M}^\Sigma_A$ contains r. Then, there exists ${\mathcal{M}'}^\Sigma_A$ such that (1) ${\mathcal{M}'}^\Sigma_A$ is a path model, and (2) for every $X \in \mathcal{M}^\Sigma_A$ and every $\psi \in \mathsf{CL}(\Sigma)$, ${\mathcal{M}'}^\Sigma_A, X \Vdash \psi$ iff $\mathcal{M}^\Sigma_A, X \Vdash \psi$.*

Proof. Let us take an arbitrary atom $B \in \mathcal{M}^\Sigma_A$ and form the set $B_r = \{\, \chi : @_r\chi \in B \,\}$ (because of lemma 37, it does not matter which B we take).

First, note that B_r is consistent. Indeed, suppose that $\chi_1 \wedge \ldots \wedge \chi_n$ is inconsistent, where $\{\chi_1, \ldots, \chi_n\} = B_r$. Then, $\vdash \neg(\chi_1 \wedge \ldots \wedge \chi_n)$ and hence, by (NN), $\vdash @_r\neg(\chi_1 \wedge \ldots \wedge \chi_n)$. Therefore, due to (A10), $\vdash \neg @_r(\chi_1 \wedge \ldots \wedge \chi_n)$ and, due to (K) and **PL**, $\vdash \neg(@_r\chi_1 \wedge \ldots \wedge @_r\chi_n)$, which is impossible since then B would be inconsistent. Secondly, note that, as every $X \in \mathcal{M}^\Sigma_A$ contains $@_r r$ (due to (A12)), $r \in B_r$. Since B_r is consistent, by lemma 31, there exists an atom C such that $B_r \subseteq C$.

Next, obtain ${\mathcal{M}'}^\Sigma_A$ by adding to $\mathcal{M}^\Sigma_A$ the submodel $\mathcal{M}^\Sigma_C$ of $\mathcal{M}^\Sigma$ generated by C. It is easy to see that ${\mathcal{M}'}^\Sigma_A$ is a disjoint union of $\mathcal{M}^\Sigma_A$ and $\mathcal{M}^\Sigma_C$. Indeed, if for some $X \in \mathcal{M}^\Sigma_A$, some $X' \in \mathcal{M}^\Sigma_C$, and some $i \in I$ we would have either $X\mathcal{R}^\Sigma_i X'$ or $X'\mathcal{R}^\Sigma_i X$, then C would be in $\mathcal{M}^\Sigma_A$, which contradicts our assumption that no atom in $\mathcal{M}^\Sigma_A$ contains r. Now, first, by lemma 38, ${\mathcal{M}'}^\Sigma_A$ contains exactly one atom containing r (namely, C). Moreover, as both $\mathcal{M}^\Sigma_A$ and $\mathcal{M}^\Sigma_C$ are, by lemma 36, regular (since they are generated submodels of a regular model $\mathcal{M}^\Sigma$), ${\mathcal{M}'}^\Sigma_A$, being their disjoint union, is also regular. Therefore, ${\mathcal{M}'}^\Sigma_A$ is a path model. Secondly, as ${\mathcal{M}'}^\Sigma_A$ is a disjoint union of $\mathcal{M}^\Sigma_A$ and $\mathcal{M}^\Sigma_C$, for every $X \in \mathcal{M}^\Sigma_A$ and every $\psi \in \mathsf{CL}(\Sigma)$, ${\mathcal{M}'}^\Sigma_A, X \Vdash \psi$ iff $\mathcal{M}^\Sigma_A, X \Vdash \psi$. ■

Using the preceding lemma, we can prove the following theorem.

THEOREM 43. *PDL^{path} without axiom (A14) is complete with respect to the class of all (not necessarily connected) path models.*

8 Conclusions and future work

We have proved completeness and decidability of extensions of multi modal K and DK with the existential modality, and axiomatised PDL^{path}, which also contains this modality. In future work, we plan to investigate the decidability of deterministic PDL^{path}. This involves investigating an open problem of decidability of deterministic CPDL with nominals.

Acknowledgements We thank Brian Logan and the anonymous AiML referees for useful comments and suggestions. This research was supported by the EPSRC grant GR/M98050/01. Natasha Alechina would like to thank the Isaac Newton Institute for Mathematical Sciences for hospitality during work on this paper.

BIBLIOGRAPHY

[1] Serge Abiteboul, Peter Buneman, and Dan Suciu. *Data on the Web: From Relations to Semistructured Data and XML*. Morgan Kaufmann, 1999.

[2] Serge Abiteboul and Victor Vianu. Regular path queries with constraints. In *Proceedings of the Sixteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems (PODS'97)*, pages 122–133, 1997.

[3] Natasha Alechina, Maarten de Rijke, and Stéphane Demri. A modal perspective on path constraints. *Journal of Logic and Computation*, 13(6):939–956, 2003.

[4] Patrick Blackburn, Maarten de Rijke, and Yde Venema. *Modal Logic*. Cambridge University Press, 2001.

[5] Giuseppe De Giacomo. *Decidability of Class-Based Knowledge Representation Formalisms*. PhD thesis, Universitá degli Studi di Roma "La Sapienza", 1995.

[6] David Harel, Dexter Kozen, and Jerzy Tiuryn. *Dynamic Logic*. MIT Press, 2000.

Natasha Alechina

School of Computer Science
University of Nottingham
NG8 1BB Nottingham, UK.

nza@cs.nott.ac.uk

Dmitry Shkatov

Institute of Philosophy
Volkhonka 14
Moscow 119992, Russia

d_shkatov@hotmail.com