PDL with Intersection and Converse is Decidable

Carsten Lutz

Institute for Theoretical Computer Science
TU Dresden
Germany
lutz@tcs.inf.tu-dresden.de

Abstract. In its many guises and variations, propositional dynamic logic (PDL) plays an important role in various areas of computer science such as databases, artificial intelligence, and computer linguistics. One relevant and powerful variation is ICPDL, the extension of PDL with intersection and converse. Although ICPDL has several interesting applications, its computational properties have never been investigated. In this paper, we prove that ICPDL is decidable by developing a translation to the monadic second order logic of infinite trees. Our result has applications in information logic, description logic, and epistemic logic. In particular, we solve a long-standing open problem in information logic. Another virtue of our approach is that it provides a decidability proof that is more transparent than existing ones for PDL with intersection (but without converse).

1 Introduction

Propositional Dynamic Logic (PDL) has originally been proposed as a modal logic for reasoning about the behaviour of programs [22, 12, 13]. Since then, the adaptation of PDL to a growing number of applications has led to many modifications and extensions. Nowadays, these additional applications have become the main driving force behind the continuing interest in the PDL family of logics, see e.g. [14, 8, 2, 5, 1]. An important family of variations of PDL is obtained by adding an intersection operator on programs, and possibly additional program operators. Alas, the extension of PDL with intersection (IPDL) is notorious for being "theoretically difficult". This is mostly due to an intricate model theory: in contrast to most other extensions of PDL, the addition of intersection destroys the tree model property in a rather dramatic way. In particular, original PDL and many of its extensions can be decided by using automata on infinite trees [24] or embedding into the alternation-free fragment of Kozen's μ-calculus [16]. By adding intersection to PDL and destroying the tree model property, we leave this framework and thus lose the toolkit of results and techniques that have been established over the last twenty years. Consequently, the results obtained for IPDL are quickly summarized: the first result about the computational properties of PDL with intersection is due to Harel, who proved that satisfiability in IPDL

with deterministic programs is undecidable [15]. In 1984, Danecki showed that dropping determinism regains decidability [7]. He also establishes a 2-ExpTime upper bound. It was long unknown whether this upper bound is tight: only in 2004, the ExpTime lower bound for IPDL stemming from original PDL was improved to an ExpSpace one and then even to a tight 2-ExpTime one [17, 18]. An axiomatization for IPDL is long sought, but until now only the axiomatization of relatively weak fragments has been successfully accomplished [4].

It appears that virtually nothing is known about extensions of IPDL. Most strikingly, the natural extension of IPDL with converse programs (*ICPDL*) has never been investigated. The aim of this paper is to perform a first investigation of the computational properties of ICPDL: we show that satisfiability in ICPDL is decidable by developing a (satisfiability preserving) translation into the monadic second order logic of infinite trees (from now on simply called MSO). This result has several interesting consequences:

First, decidability of ICPDL implies decidability of the information logic DAL (*Data Analysis Logic*), a problem that has been open since DAL was proposed in 1985 [11]. The purpose of DAL is to aggregate data into sets that can be characterized using given properties, and, dually, to determine properties that best characterize a given set of data. Technically, DAL may be viewed as the variant of IPDL obtained by requiring all relations to be equivalence relations and admitting only the program operators $\cap$ and $\cup^*$, where the latter is a combination of PDL's operators $\cup$ and $\cdot^*$. In ICPDL, equivalence relations can be simulated using $(a \cup a^-)^*$ for some atomic program a. Thus, DAL can be viewed as a fragment of ICPDL.

Second, there is a close correspondence between variants of PDL and description logics (DLs). In particular, the description logic $\mathcal{ALC}_{\mathsf{reg}}$ [3, 14] is a syntactic variant of PDL without the test operator [23], and the intersection operator of IPDL corresponds to the intersection role constructor in description logics. The latter is a traditional constructor that is present in many DL formalisms, see e.g. [9, 6, 20, 21]. Decidability and complexity results play a central role in the area of description logic, but have never been obtained for the natural extension $\mathcal{ALC}^{\cap}_{\mathsf{reg}}$ of $\mathcal{ALC}_{\mathsf{reg}}$ with role intersection. Clearly, $\mathcal{ALC}^{\cap}_{\mathsf{reg}}$ is a syntactic variant of test-free ICPDL, and thus our decidability result carries over.

Third, ICPDL can be applied to obtain results in epistemic logic [10]. The basic observation is as in the case of DAL: ICPDL can simulate equivalence relations by writing $(a \cup a^-)^*$. Since union and transitive closure of programs can be combined to express the common knowledge operator of epistemic logic, and intersection of programs corresponds to the distributed knowledge operator, decidability of ICPDL can be used to obtain decidability for epistemic logic with both common knowledge and distributed knowledge. We should admit, however, that this approach is rather brute force: since the common knowledge and distributed knowledge operators of epistemic logic cannot be nested to build up more complex operations on relations, epistemic logic lacks much of the complexity of ICPDL. Therefore and as noted in [10], decidability can also be obtained using more standard techniques.

Apart from the applications just mentioned, we believe that there is an additional virtue of the MSO translation exhibited in this paper: without intending to derogate the admirable work of Danecki that provided the basic ideas for the tree encoding of ICPDL models developed in this paper [7], it seems fair to claim that Danecki's decidability proof for IPDL is rather intricate and difficult to understand. Moreover, the correctness is hard to verify since the only available presentation (a conference paper) lacks many non-trivial details. Although the MSO translation presented in the current paper also involves some non-trivial encodings, in our opinion it is the easiest proof of the decidability of IPDL that has been obtained so far. Together with the technical report accompanying this paper [19], the proofs are fully rigorous and readily checked in detail.

This paper is organized as follows. In Section 2, we introduce ICPDL. Section 3 prepares for the MSO translation by discussing, on an intuitive level, how ICPDL models can be abstracted into trees. The translation itself is exhibited in Section 4 which also contains a correctness proof. We discuss future work and conclude in Section 5.

2 The Language

Let Var and Prog be countably infinite sets of propositional variables and atomic programs, respectively. The sets of *ICPDL programs* and *ICPDL formulas* are defined by simultaneous induction as follows:

- each atomic program is a program;
- each propositional variable is a formula;
- if α and β are programs and φ is a formula, then the following are also programs:
$$\alpha^-,\ \alpha \cap \beta,\ \alpha \cup \beta,\ \alpha;\beta,\ \alpha^*,\ \varphi?$$
- if φ and ψ are formulas and α is a program, then the following are also formulas:
$$\neg\varphi,\ \langle\alpha\rangle\varphi$$

We use $\varphi_1 \wedge \varphi_2$ as an abbreviation for $\langle\varphi_1?\rangle\varphi_2$, $\varphi_1 \vee \varphi_2$ for $\neg(\neg\varphi_1 \wedge \neg\varphi_2)$, and $[\alpha]\,\varphi$ for $\neg\langle\alpha\rangle\neg\varphi$. Moreover, we use $\top$ to abbreviate an arbitrary (but fixed) propositional tautology, and $\bot$ for $\neg\top$.

The semantics of ICPDL is defined in the usual way through Kripke structures. A *Kripke structure* is a triple $K = (W, R, L)$, where

- W is a set of points,
- R assigns to each atomic program $a \in \mathsf{Prog}$ a binary relation $R(a)$ on W,
- L assigns to each atomic proposition $p \in \mathsf{Var}$ the set of points $L(p)$ in which it holds.

The extension of R to complex programs and the definition of the consequence relation $\models$ for ICPDL are, again, by simultaneous induction:

$$
\begin{array}{lll}
R(\alpha^-) & \text{is the converse of } R(\alpha) & \\
R(\alpha_1 \cap \alpha_2) & = & R(\alpha_1) \cap R(\alpha_2), \\
R(\alpha_1 \cup \alpha_2) & = & R(\alpha_1) \cup R(\alpha_2), \\
R(\alpha_1; \alpha_2) & = & R(\alpha_1) \circ R(\alpha_2). \\
R(\alpha^*) & \text{is the reflexive-transitive closure of } R(\alpha) & \\
R(\varphi?) & = & \{(w,w) \in W^2 \mid K, w \models \varphi\}
\end{array}
$$

$$
\begin{array}{lll}
K, w \models p & \text{iff} & w \in L(p) \text{ for } p \in \mathsf{Var} \\
K, w \models \neg\varphi & \text{iff} & K, w \not\models \varphi \\
K, w \models \langle\alpha\rangle\varphi & \text{iff} & \text{there is } w' : (w, w') \in R(\alpha) \text{ and } K, w' \models \varphi
\end{array}
$$

Let φ be a formula and $K = (W, R, L)$ a Kripke structure. Then K is a *model* of φ if there is a $w \in W$ with $K, w \models \varphi$. The formula φ is called *satisfiable* if it has a model.

3 ICPDL Models

Our aim is to devise a satisfiability preserving translation from ICPDL to MSO over infinite trees. The main difficulty is posed by the fact that ICPDL does not have the tree model property. This is witnessed e.g. by the formulas

$$\neg p \wedge \langle a \cap a^- \rangle p \text{ and } \neg p \wedge [b]\bot \wedge \langle (a; p?; a) \cap b^* \rangle \top$$

which both enforce a cycle of length 2.[1] To carry out the translation to MSO, it is important to develop a tree-shaped abstraction of ICPDL models. Such an abstraction is described in the current section. Although it provides the guiding intuitions for developing the translation to MSO, there is no need to formally establish the correctness of the abstraction beforehand. Therefore, our discussion will remain on an intuitive level.

Intersection

ICPDL's lack of the tree model property is clearly due to the intersection operator on relations. Even the simple formula $\langle a \cap b \rangle \top$ does not have a tree model: it enforces a Kripke structure K as shown on the left-hand side of Figure 1. For the MSO translation, we represent K using the tree displayed on the right-hand side of the same figure. In this tree, the left son represents the substructure of K that is obtained by dropping the b edge, and the right son describes the substructure obtained by dropping the a edge. The symbol "$\cap$" labelling the root node indicates that a parallelization operation is required to construct K from

[1] It is easy to modify these formulas such that they enforce a cycle whose length is exponential in the length of the formula.

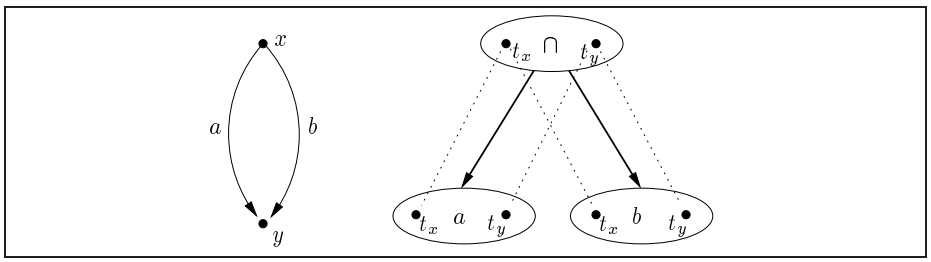

Fig. 1. Tree for intersection.

these two substructures: simply identify their roots and sinks. Intuitively, the root node represents the whole structure K.

The tree representation does not only encode the relational structure of K, but also records satisfaction of relevant formulas by states of K. The following definition fixes the set of formulas relevant for deciding satisfiability of an ICPDL formula φ: the (Fischer-Ladner) closure of φ.

Definition 1 (Closure). *The set of subprograms* $\mathsf{subp}(\alpha)$ *of ICPDL programs* α *and the set of subformulas* $\mathsf{subf}(\varphi)$ *of ICPDL formulas* φ *is defined simultaneously as follows:*

- $\mathsf{subp}(a) = \{a\}$ *if* a *is atomic;*
- $\mathsf{subp}(\alpha) = \{\alpha\} \cup \mathsf{subp}(\beta) \cup \mathsf{subp}(\gamma)$ *if* $\alpha = \beta \cap \gamma$ *or* $\alpha = \beta; \gamma$*;*
- $\mathsf{subp}(\alpha) = \{\alpha\} \cup \mathsf{subp}(\beta)$ *if* $\alpha = \beta^*$ *or* $\alpha = \beta^-$*;*
- $\mathsf{subp}(\varphi?) = \{\varphi?\} \cup \bigcup_{\langle\beta\rangle\psi \in \mathsf{subf}(\varphi)} \mathsf{subp}(\beta)$*;*
- $\mathsf{subf}(p) = \{p\}$ *if* $p \in \mathsf{Var}$*;*
- $\mathsf{subf}(\neg\varphi) = \{\neg\varphi\} \cup \mathsf{subf}(\varphi)$*;*
- $\mathsf{subf}(\langle\alpha\rangle\varphi) = \{\langle\alpha\rangle\varphi\} \cup \mathsf{subf}(\varphi) \cup \bigcup_{\psi? \in \mathsf{subp}(\alpha)} \mathsf{subf}(\psi)$.

Finally, we define the closure *of an ICPDL formula* φ *as*

$$\mathsf{cl}(\varphi) := \{\psi, \neg\psi \mid \psi \in \mathsf{cl}(\varphi)\}.$$

For x a state in a Kripke structure, the *type* of x is the set of formulas $\{\varphi \in \mathsf{cl}(\varphi_0) \mid K, x \models \varphi\}$, where φ_0 is the formula whose satisfiability is to be decided. In the tree representation of a model, each node stores the type of the root state and of the sink state of the substructure that this node represents. In the case of Figure 1, all three tree nodes store the type t_x of x and t_y of y since they all describe a substructure of K with root x and sink y. We say that t_x is stored *in the first place* of each node, and t_y is stored *in the second place*. Observe that distinct places in tree nodes may represent identical states in the model. This induces an equivalence relation on places, whose skeleton is given as dotted lines in Figure 1. This relation will play a central role in the translation to MSO.

Composition

Now consider a formula $\langle a; b\rangle\top$. It enforces the model on the left-hand side of Figure 2. Again, the right-hand side displays the corresponding tree abstraction

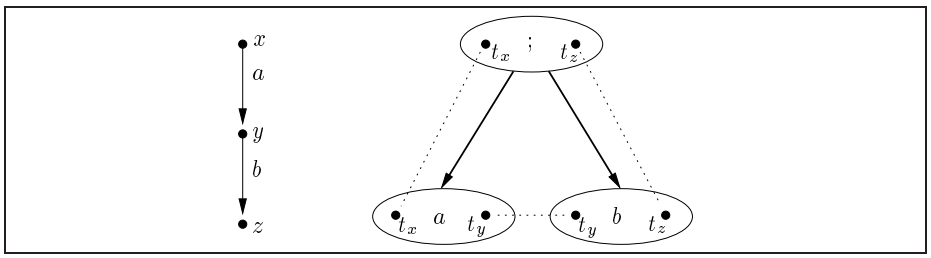

Fig. 2. Tree for composition.

with the dotted edges providing a skeleton for the equivalence relation on places. The symbol ";" of the root nodes indicates that the structure represented by the root node is obtained from the structures represented by the leaves through a composition operation: identify the sink of the left son with the root of the right son.

Kleene Star

Formulas $\langle a^*\rangle\top$ enforce an a-path of arbitrary length. To represent a path of length zero (i.e., a single state), we use a tree consisting of a single node labelled "=". The two places of this node are equivalent, i.e., represent the same state. To represent longer paths, we may repeatedly apply the composition operation to nodes labelled "a" and "=". A tree representation of a path of length two can be found in Figure 3.

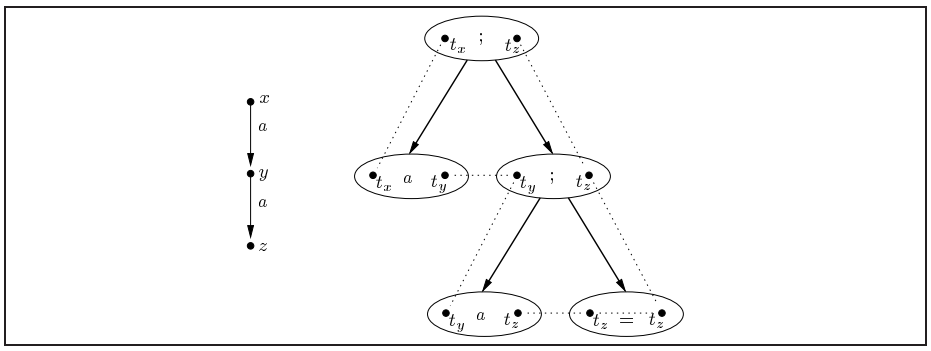

Fig. 3. Tree for Kleene star.

Observe the dotted edge connecting the two places of the "=" node. It should be clear that, by combining the representation schemata given in Figures 1 and 2 and by using "=" nodes, we can construct a tree representation of models enforced by any formula $\langle\alpha\rangle\top$, with α composed from the operators $\{\cup, \cap, \varphi?, ;, \cdot^*\}$ in an arbitrary way: the operator "$\cup$" requires no explicit represention in the tree structure and the operator "φ?" can be treated via a node labelled "=".

Converse

To deal with the converse operator, we take an approach that may not be what one would expect on first sight. As discussed later, the seemingly complicated treatment of converse allows to simplify other parts of the MSO translation. Consider a formula $\langle a^-\rangle\top$ and the enforced model given on the right-hand side of Figure 4.

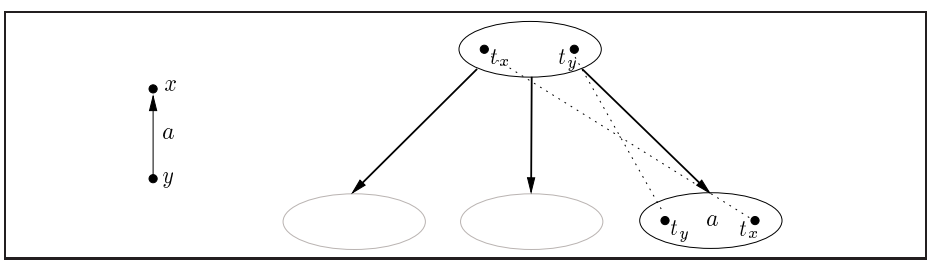

Fig. 4. Tree for converse programs

Until now, all considered models have been abstracted into *binary* trees. For dealing with converse, we switch to *ternary* trees. The Kripke structure from Figure 4 is represented by the tree given on the right-hand side of the same figure. The third son represents the structure in which there is an a-edge from root y to sink x, i.e., the horizontal mirror image of the Kripke structure on the left. In contrast, the root represents the original structure, where there is an a-edge from *sink* y to *root* x. Observe that the equivalence relation induced by the pointed edges swaps the places of the root and the third son as expected. Also observe that the root node does not have a particular type such as "$\cap$" or ";". We need not introduce a dedicated type for converse since, for technical reasons discussed below, *every* node in the tree has a third son whose places are obtained by swapping the places of the original node. Finally, note that the first and second son of the root are simply dummies. Although they will be required to exist for technical reasons, intuitively they carry no meaningful information.

Multiple Diamonds

So far, we have mostly concentrated on tree abstractions of models for simple formulas of the form $\langle\alpha\rangle\varphi$. Tree abstractions of models for arbitrarily shaped

formulas can be obtained by joining, in a suitable way, the tree abstractions of models for such simple formulas. Consider the formula $\langle a; b\rangle\top \wedge \langle c\rangle\top$, which enforces the structure shown on the left-hand side of Figure 5. As usual, the

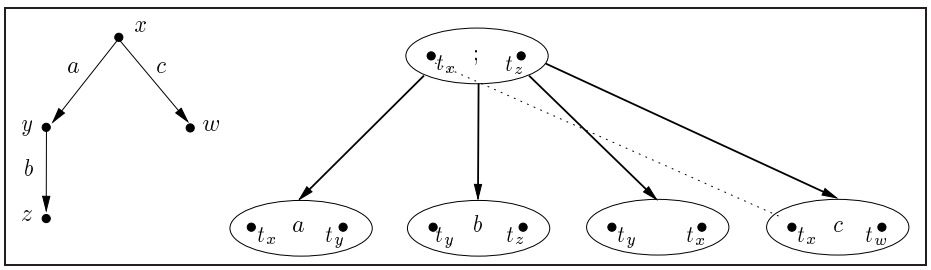

Fig. 5. Tree for multiple diamonds

tree abstraction is shown on the right-hand side. The root together with the first two sons are the tree abstraction of the substructure witnessing $\langle a; b\rangle\top$, where the dotted edges are as in Figure 2 but omitted for simplicity. The third son exists because every node is required to have a third son. The dotted edges connecting the root and the third son are as in Figure 4, but again omitted. Finally, the fourth son by itself (i.e., without the root) is the tree abstraction of the substructure witnessing $\langle c\rangle\top$.

The ratio of this representation is as follows: suppose that a state x in a Kripke structure satifies multiple diamonds $\langle\alpha_1\rangle\varphi_1, \ldots, \langle\alpha_k\rangle\varphi_k$. For $1 \leq i \leq k$, we take the representation of the model enforced by $\langle\alpha_i\rangle\varphi_i$ as a ternary tree as described above. Let these trees be $T_1, \ldots, T_k$. To join them into a single tree, we attach the roots of $T_2, \ldots, T_k$ as sons number 4 to $k+3$ to the root of T_1. Observe that, in the resulting tree, the first place of the root node is equivalent to the first place of sons number 4 to $k+3$. This is indicated by the dotted edge in Figure 5.

Using this method, we can deal with the problem that a state represented by the *left*-hand place of a tree node may have to satisfy more than a single diamond. What will we do if a state x represented by a right-hand place of a tree node has to satisfy diamonds $\langle\alpha_1\rangle\varphi_1, \ldots, \langle\alpha_k\rangle\varphi_k$? We simply exploit the fact that every node has a third son swapping the places: we attach the trees $T_1, \ldots, T_k$ representing the models enforced by the diamonds $\langle\alpha_1\rangle\varphi_1, \ldots, \langle\alpha_k\rangle\varphi_k$ as sons number 4 to $k+4$ to the third son of the node whose right-hand place represents x. By composing the dotted edges displayed in Figures 4 and 5, it is easily verified that, then, the second place of the root of T_1 is equivalent to the first place of the root of T_2 as required.

4 Translation to MSO

We now put the ideas developed in the previous section to work. The goal is to prove the main result of this paper:

Theorem 1. *Satisfiability in PDL with intersection and converse is decidable.*

Let φ_0 be an ICPDL formula whose satisfiability is to be decided. Moreover, let k be the number of diamond formulas $\langle\alpha\rangle\varphi$ in $\mathsf{cl}(\varphi_0)$. We translate φ_0 into an eqi-satisfiable formula φ_0^* of monadic second-order logic of the infinite $k+3$-ary tree. More precisely, we assume MSO models to have domain $\{1,\ldots,k+3\}^*$, which from now on we abbreviate with $[k+3]^*$. There are $k+3$ unary functions s_i mapping each node to it's i-th son.

Intuitively, the formula φ_0^* is constructed such that the models of φ_0^* are precisely the tree abstractions of models of φ_0. In particular, the intuition behind the $k+3$ successors is as explained in the previous section. The assembly of φ_0^* involves several steps. First, we fix the MSO signature used:

- unary predicates F_φ^1 and F_φ^2 for every $\varphi \in \mathsf{cl}(\varphi_0)$;
- unary predicates $T_=$, $T_\cap$, $T_;$, and $T_\bot$;
- a unary predicate T_a for each atomic program a.

The predicates F_φ^i are used to store types in the first and second place of tree nodes (c.f. previous section): if $\mathfrak{M}$ is an MSO model and $x \in [k+3]^*$, then $\{\varphi \mid \mathfrak{M} \models F_\varphi^1(x)\}$ is the type stored in the first place of x and $\{\varphi \mid \mathfrak{M} \models F_\varphi^2(x)\}$ is the type stored in the second place of x.

The predicates T_a, $T_=$, $T_\cap$, $T_;$, and $T_\bot$ are markers for the different kinds of nodes in trees. The only kind of node that was not discussed in the previous section is $T_\bot$. This kind of node is used when the i-th son is not needed, for some i with $3 < i \leq k+3$. For example, assume that $\mathfrak{M} \not\models F_\varphi^1(x)$ for some node $x \in [k+3]^*$ and all formulas $\varphi \in \mathsf{cl}(\varphi_0)$ of the form $\langle\alpha\rangle\varphi$. Then the sons $x4,\ldots,x(k+3)$ of x are not needed. Since our MSO models should be full $k+3$-ary trees, we simply mark such sons with $T_\bot$.

To ensure that the sets $\{\varphi \mid \mathfrak{M} \models F_\varphi^1(x)\}$ describe valid types, we have to describe the semantics of negation and of diamonds—recall that all other operators are merely abbreviations. Dealing with negation is easy:

$$\psi_1^* := \bigwedge_{\neg\varphi \in \mathsf{cl}(\varphi_0)} \forall x \,.\, F_{\neg\varphi}^1(x) \leftrightarrow \neg F_\varphi^1(x) \wedge F_{\neg\varphi}^2(x) \leftrightarrow \neg F_\varphi^2(x)$$

To treat diamonds, we need some preliminaries. First, we define a formula with two free variables that characterizes the identitiy of places as discussed in the previous section. More precisely, it is convenient to define four such formulas $\chi_{i,j}$, $i,j \in \{1,2\}$, as shown in Figure 6. Intuitively, we have $\mathfrak{M} \models \chi_{i,j}[x,y]$ iff the i'th place of x is equivalent to the j'th place of y. According to the idea of

$$\vartheta(P_1, P_2) := \forall z.(T_{=}(z) \to (P_1(z) \leftrightarrow P_2(z))) \wedge \tag{1}$$
$$\forall z.(T_{\cap}(z) \to (P_1(s) \leftrightarrow P_1(s_1(z)))) \wedge \tag{2}$$
$$\forall z.(T_{\cap}(z) \to (P_1(s) \leftrightarrow P_1(s_2(z)))) \wedge \tag{3}$$
$$\forall z.(T_{\cap}(z) \to (P_2(s) \leftrightarrow P_2(s_1(z)))) \wedge \tag{4}$$
$$\forall z.(T_{\cap}(z) \to (P_2(s) \leftrightarrow P_2(s_2(z)))) \wedge \tag{5}$$
$$\forall z.(T_{;}(z) \to (P_1(z) \leftrightarrow P_1(s_1(z)))) \wedge \tag{6}$$
$$\forall z.(T_{;}(z) \to (P_2(z) \leftrightarrow P_2(s_2(z)))) \wedge \tag{7}$$
$$\forall z.(T_{;}(z) \to (P_2(s_1(z)) \leftrightarrow P_1(s_2(z)))) \wedge \tag{8}$$
$$\forall z.(P_1(z) \leftrightarrow P_2(s_3(z))) \wedge \tag{9}$$
$$\forall z.(P_2(z) \leftrightarrow P_1(s_3(z))) \wedge \tag{10}$$
$$\bigwedge_{3<\ell\leq k+3} \forall z.(\neg T_{\perp}(s_\ell(z)) \to (P_1(z) \leftrightarrow P_1(s_\ell(z)))) \tag{11}$$

$$\chi_{i,j}(x, y) := \forall P_1, P_2.(P_i(x) \wedge \vartheta(P_1, P_2)) \to P_j(y)$$

Fig. 6. The formulas $\chi_{i,j}(x, y)$.

place equivalence, all equivalent places should have the same type:

$$\psi_2^* := \bigwedge_{i,j\in\{1,2\}} \forall x, y \,.\, \chi_{i,j}(x, y) \to (\bigwedge_{\varphi\in\mathsf{cl}(\varphi_0)} F^i_\varphi(x) \leftrightarrow F^j_\varphi(y))$$

We now define, for each program $\alpha \in \mathsf{subp}(\varphi_0)$, a formula σ_α that relates the *first* place of a node x to the *second* place of a node y iff the states represented by these two places are related via the program α: for each $\alpha \in \mathsf{subp}(\varphi_0)$, set:

- $\sigma_a(x, y) := \exists z.\chi_{1,1}(x, z) \wedge T_a(z) \wedge \chi_{2,2}(y, z)$;
- $\sigma_{\varphi?}(x, y) := \chi_{1,2}(x, y) \wedge F^1_\varphi(x)$;
- $\sigma_{\alpha\cup\beta}(x, y) := \sigma_\alpha(x, y) \vee \sigma_\beta(x, y)$;
- $\sigma_{\alpha\cap\beta}(x, y) := \sigma_\alpha(x, y) \wedge \sigma_\beta(x, y)$;
- $\sigma_{\alpha;\beta}(x, y) := \exists z, z'.\sigma_\alpha(x, z) \wedge \chi_{2,1}(z, z') \wedge \sigma_\beta(z', y)$;
- $\sigma_{\alpha^-}(x, y) := \sigma_\alpha(s_3(y), s_3(x))$;
- $\sigma_{\alpha^*}(x, y) := \chi_{1,2}[x, x] \vee \forall P.\big(\ (P(s_3(x)) \wedge \vartheta'_\alpha(P)) \to P(y)\ \big)$

  with

  $\vartheta'_\alpha(P) := \forall x, y, z.\big(\ (P(x) \wedge \chi_{2,1}(x, y) \wedge \sigma_\alpha(y, z)) \to P(z)\ \big)$

Some remarks are in order. To see why σ_a does not simply read $x = y \wedge T_a(x)$, consider Figure 1: the left place of the root node is clearly related to the right

place of the root node via the program a although the root is not labelled "a". In $\sigma_{\alpha;\beta}$, the middle conjunct is necessary since we only relate first places to second places. The formula σ_{α^-} is easily understood by considering the equivalence of places indicated in Figure 4. Finally, consider σ_{α^*}. The first disjunct reflects the fact that, in Kripke structures, α^* relates every state to itself. The formula $\vartheta'_\alpha(P)$ states that the set of nodes P is closed under making α-steps from second places of nodes in P: if $x \in P$, the second place of x is equivalent to the first place of some y, and y is related to some z via σ_α, then the second place of z can be reached from the second place of x by making an α transition and we add z to P. Note that, in the definition of σ_{α^*}, we put $s_3(x)$ into P as the initial element rather than x. This is necessary since σ_{α^*} relates first places to second places, but $\vartheta'_\alpha(P)$ closes off under making α-steps from *second* places of nodes in P. Moreover, the second place of $s_3(x)$ is clearly equivalent to the first place of x.

Using the formulas σ_α, we can now describe the semantics of diamonds:

$$\psi_3^* := \bigwedge_{\langle\alpha\rangle\varphi \in \mathsf{cl}(\varphi_0)} \forall x \,.\, F^1_{\langle\alpha\rangle\varphi}(x) \leftrightarrow \exists y.\sigma_\alpha(x,y) \wedge F^2_\varphi(y))$$

It pays off here that we require every node to have a third son with swapped places: due to this son, there is no need to explicitly describe the semantics of diamonds satisfied by second places, i.e., recorded via formulas $F^i_{\langle\alpha\rangle\varphi}(x)$ with $i = 2$. We thus save the definition of counterparts of the formulas σ_α that relate second places to first places. Also, there is no need to define counterparts of the formulas σ_α that relate first places to first places, or second places to second places: via the third son, such relationships can always be understood as a relationship from a first place to a second place.

Finally, we assemble φ_0^*:

$$\varphi_0^* := \psi_1^* \wedge \psi_2^* \wedge \psi_3^* \wedge \exists x.F^1_{\varphi_0}(x)$$

In [19], we prove correctness of the translation:

Lemma 1. *φ_0 is satisfiable in ICPDL iff φ_0^* is satisfiable in MSO.*

For the "if" direction, assume that φ_0^* is satisfiable in MSO, i.e. there is an MSO structure $\mathfrak{M}$ based on a tree of out-degree $k+3$ such that φ_0^* is satisfied in $\mathfrak{M}$. Let $P := [k+3]^* \times \{1,2\}$ be the set of *places*. We define the relation $\sim$ on P by setting $(x,i) \sim (y,j)$ iff $\mathfrak{M} \models \chi_{i,j}[x,y]$. It is not hard to show that $\sim$ is an equivalence relation. Let $[x,i]$ denote the equivalence class of $(x,i) \in P$ w.r.t. $\sim$. We define a Kripke structure $K = (W,R,L)$ as follows:

- $W = \{[x,i] \mid (x,i) \in P\}$;
- $R(a) = \{([x,1],[y,2]) \mid \mathfrak{M} \models \sigma_a[x,y]\}$ for all atomic programs a;
- $L(p) = \{[x,1] \mid x \in (F^1_p)^{\mathfrak{M}}\} \cup \{[x,2] \mid x \in (F^2_p)^{\mathfrak{M}}\}$ for all $p \in \mathsf{Var}$.

Note that K is well-defined: due to φ_2^*, $(x,1) \sim (y,1)$ implies that $x \in (F_p^1)^{\mathfrak{M}}$ iff $y \in (F_p^1)^{\mathfrak{M}}$ for all $p \in \mathsf{Var}$, and likewise for F_p^2. Additionally, by definition of σ_a, $(x,1) \sim (x',1)$ and $(y,2) \sim (y',2)$ implies that $\mathfrak{M} \models \sigma_a[x,y]$ iff $\mathfrak{M} \models \sigma_a[x',y']$, for all atomic programs a.

In [19], we then prove the following, central claim.

Claim. For all $x, y \in [k+3]^*$, $\varphi \in \mathsf{cl}(\varphi_0)$, and $\alpha \in \mathsf{subp}(\varphi_0)$, we have

1. $([x,1],[y,2]) \in R(\alpha)$ iff $\mathfrak{M} \models \sigma_\alpha[x,y]$;
2. $\mathfrak{M} \models F_\varphi^i[x]$ iff $K, [x,i] \models \varphi$

Since φ_0^* is satisfied in $\mathfrak{M}$, there is an $x \in [k+3]^*$ such that $\mathfrak{M} \models F_{\varphi_0}^1[x]$. By Point 2 of the claim, this implies that K is a model of φ_0.

For the "only if" direction, let $K = (W, R, L)$ be a model of φ_0, and let $w_0 \in W$ such that $K, w_0 \models \varphi_0$. To construct an MSO model with domain $[k+3]^*$ satisfying φ_0^* at the root, we inductively define three mappings

$$\begin{aligned} \tau_1 &: [k+3]^* \to W \\ p &: [k+3]^* \to \mathsf{subp}(\varphi_0) \cup \{\varepsilon, \bot\} \\ \tau_2 &: [k+3]^* \to W \end{aligned}$$

such that the following condition is satisfied:

$$\text{for all } x \in [k+3]^*, p(x) \neq \bot \text{ implies } (\tau_1(x), \tau_2(x)) \in R(p(x)), \qquad (\dagger)$$

where $R(\varepsilon)$ is defined as the identitiy relation on W. Intuitively, $\tau_1(x)$ identifies the state described by the first place of x, $\tau_2(x)$ identifies the state described by the second place of x, and $p(x)$ is the program that we want to hold between these two places. The case $p(x) = \bot$ means that the mapping $p(\cdot)$ carries no relevant information for the node x. Before we can start the definition, we need some preliminaries. First, we assume that the diamond formulas in $\mathsf{cl}(\varphi_0)$ are linearly ordered, and that $\mathcal{E}_i$ yields the i-th such formula (the numbering starts with 0). Second, we call a program α *determined* if the top-level operator is not "$\cup$". We inductively fix a choice function ch that maps every triple $(w, \alpha, w') \subseteq W \times \mathsf{subp}(\varphi_0) \times W$ with $(w,w') \in R(\alpha)$ to a determined program $\mathsf{ch}(w,\alpha,w') \in \mathsf{subp}(\alpha)$ such that $R(\mathsf{ch}(w,\alpha,w')) \subseteq R(\alpha)$ and $(w,w') \in R(\mathsf{ch}(w,\alpha,w'))$: let $(w,w') \in R(\alpha)$.

- if α is determined, set $\mathsf{ch}(w,\alpha,w') := \alpha$.
- if α is not determined, then $\alpha = \beta \cup \gamma$. By the semantics, $(w,w') \in R(\alpha)$ implies $(w,w') \in R(\beta)$ or $(w,w') \in R(\gamma)$. In the first case, set $\mathsf{ch}(w,\alpha,w') := \beta$ if β is determined, and $\mathsf{ch}(w,\alpha,w') := \mathsf{ch}(w,\beta,w')$ otherwise. In the second case, set $\mathsf{ch}(w,\alpha,w') := \gamma$ if γ is determined, and $\mathsf{ch}(w,\alpha,w') := \mathsf{ch}(w,\gamma,w')$ otherwise.

Now, the three mappings are defined simultaneously by making a case distinction as follows. To understand this definition, it may help to recall the intuitions laid out in Section 3.

1. To start, set $\tau_1(\varepsilon) := w_0$, $p(\varepsilon) := \varepsilon$, and $\tau_2(\varepsilon) := w_0$. (The choice of $p(\varepsilon)$ and $\tau_2(\varepsilon)$ is not crucial).
2. Let $\tau_1(x)$ be defined, $\tau_1(s_1(x))$ undefined, and $p = \alpha_1 \cap \alpha_2$. Then set, for $i \in \{1,2\}$: $\tau_1(s_i(x)) := \tau_1(x)$, $p(s_i(x)) := \mathsf{ch}(\tau_1(x), \alpha_i, \tau_2(x))$, and $\tau_2(s_i(x)) := \tau_2(x)$.
3. Let $\tau_1(x)$ be defined, $\tau_1(s_1(x))$ undefined, and $p = \alpha; \beta$. By (†) and the semantics, there is a $w \in W$ with $(\tau_1(x), w) \in R(\alpha)$ and $(w, \tau_2(x)) \in R(\beta)$. Set

$$\begin{array}{lll} \tau_1(s_1(x)) := \tau_1(x) & p(s_1(x)) := \mathsf{ch}(\tau_1(x), \alpha, w) & \tau_2(s_1(x)) := w \\ \tau_1(s_2(x)) := w & p(s_2(x)) := \mathsf{ch}(w, \beta, \tau_2(x)) & \tau_2(s_2(x)) := \tau_2(x) \end{array}$$

4. Let $\tau_1(x)$ be defined, $\tau_1(s_1(x))$ undefined, $p = \alpha^*$, and $\tau_1(x) = \tau_2(x)$. Set, for $i \in \{1,2\}$, $\tau_1(s_i(x)) := w_0$, $p(s_i(x)) := \varepsilon$, and $\tau_2(s_i(x)) := w_0$. Intuitively, the first and second successor of x are not needed. To nevertheless obtain a full $k+3$-ary tree, we "restart" at w_0.
5. Let $\tau_1(x)$ be defined, $\tau_1(s_1(x))$ undefined, $p = \alpha^*$, and $\tau_1(x) \neq \tau_2(x)$. By (†) and the semantics, there is a sequence $w_0, \ldots, w_n \in W$ such that $\tau_1(x) = w_0$, $\tau_2(x) = w_n$, $(w_i, w_{i+1}) \in R(\alpha)$ for $i < n$, and $w_i \neq w_j$ for $i < j \leq n$. Let $w_0, \ldots, w_n \in W$ be the shortest such sequence. Set

$$\begin{array}{lll} \tau_1(s_1(x)) := \tau_1(x) & p(s_1(x)) := \mathsf{ch}(\tau_1(x), \alpha, w_1) & \tau_2(s_1(x)) := w_1 \\ \tau_1(s_2(x)) := w_1 & p(s_2(x)) := \alpha^* & \tau_2(s_2(x)) := \tau_2(x) \end{array}$$

6. Let $\tau_1(x)$ be defined, $\tau_1(s_1(x))$ undefined, and $p \in \mathsf{Prog}$ or p of the form α^-. Set, for $i \in \{1,2\}$, $\tau_1(s_i(x)) := w_0$, $p(s_i(x)) := \varepsilon$, and $\tau_2(s_i(x)) := w_0$. Similar to Case 4, the first and second successor of x are not needed.
7. Let $\tau_1(x)$ be defined and $\tau_1(s_3(x))$ be undefined. Set $\tau_1(s_3(x)) := \tau_2(x)$, $\tau_2(s_3(x)) := \tau_1(x)$, and

$$p(s_3(x)) := \begin{cases} \mathsf{ch}(\tau_2(x), \alpha, \tau_1(x)) & \text{if } p(x) = \alpha^- \\ \bot & \text{if } p(x) \text{ is not of the form } \alpha^- \end{cases}$$

8. Let $\tau_1(x)$ be defined and $\tau_1(s_n(x))$ undefined for some n with $3 < n \leq k+3$, and $K, \tau_1(x) \models \mathcal{E}_{n-3} = \langle \alpha \rangle \varphi$. Then by the semantics there is a $w \in W$ with $(\tau_1(x), w) \in R(\alpha)$ and $K, w \models \varphi$. Set $\tau_1(s_n(x)) := \tau_1(x)$, $p(s_n(x)) := \mathsf{ch}(\tau_1(x), \alpha, w)$, and $\tau_2(s_n(x)) := w$.
9. Let $\tau_1(x)$ be defined and $\tau_1(s_n(x))$ undefined for some n with $3 < n \leq k+3$, and $K, \tau_1(x) \not\models \mathcal{E}_{n-3} = \langle \alpha \rangle \varphi$. Then set $\tau_1(s_n(x)) := w_0$, $p(s_n(x)) := \varepsilon$, and $\tau_2(s_n(x)) := w_0$. As in Cases 4 and 6, we restart at w_0 since the n-th successor of x is not needed.

Now we construct an MSO model $\mathfrak{M}$ as follows:

– for all $\varphi \in \mathsf{cl}(\varphi_0)$ and $i \in \{1,2\}$, set $(F^i_\varphi)^{\mathfrak{M}} := \{x \in [k+3]^* \mid K, \tau_i(x) \models \varphi\}$

- $T_{=}^{\mathfrak{M}} := \{x \in [k+3]^* \mid p(x) = \varepsilon\}$
  $\cup \{x \in [k+3]^* \mid p(x) = \varphi?$ for some formula $\varphi\}$
  $\cup \{x \in [k+3]^* \mid p(x) = \alpha^*$ for some $\alpha \in \mathsf{subp}(\varphi_0)$ and $\tau_1(x) = \tau_2(x)\}$
  $T_{\cap}^{\mathfrak{M}} := \{x \in [k+3]^* \mid p(x) = \alpha \cap \beta$ for some $\alpha, \beta \in \mathsf{subp}(\varphi_0)\}$
  $T_{;}^{\mathfrak{M}} := \{x \in [k+3]^* \mid p(x) = \alpha; \beta$ for some $\alpha, \beta \in \mathsf{subp}(\varphi_0)\}$
  $\cup \{x \in [k+3]^* \mid p(x) = \alpha^*$ for some $\alpha \in \mathsf{subp}(\varphi_0)$ and $\tau_1(x) \neq \tau_2(x)\}$
  $T_{\perp}^{\mathfrak{M}} := \{s_n(x) \mid K, \tau_1(x) \not\models \mathcal{E}_{n-3}\}$

- for $a \in \mathsf{prog}$, set $T_a^{\mathfrak{M}} := \{x \in [k+3]^* \mid p(x) = a\}$.

In [19], we show that $\mathfrak{M} \models \varphi_0^*[\varepsilon]$.

5 Conclusion

In this paper, we have proved decidability of ICPDL, i.e. PDL extended with intersection and converse. As laid out in the introduction, this result that has several interesting applications. One additional virtue of the presented decidability proof is that, compared to existing proofs for PDL with intersection (but without converse), it is relatively simple and fully rigorous. There is, however, a price to be paid for this simplicity: our translation to MSO only yields a non-elementary upper bound. Indeed, when translating the following sequence $(\varphi_i)_{i \in \mathbb{N}}$ of ICPDL formulas, we obtain a sequence of MSO formulas with a strictly increasing quantifier alternation depth:

$$\varphi_i := [(\cdots((a_0^*; a_1)^*; a_2)^*; \cdots; a_i)^*]p.$$

We believe that this upper bound is not tight. Indeed, it seems likely that satisfiability in ICPDL is 2-ExpTime-complete, just as satisfiability in IPDL. For proving this, however, it seems inevoidable to use the complex techniques of Danecki [7], in particular his "$\vdash$" relation. Therefore, we believe that it is useful and illustrative to first prove only decidability in a more transparent way. Pinpointing the exact computational complexity of ICPDL is left for future work. Another interesting question is whether or not there are useful fragments of ICPDL that involve both intersection and Kleene star and for which reasoning is in ExpTime—thus not harder than in PDL. We suspect that the set of program operators $\{\cup, \cap, \cdot^*, \cdot^-, \varphi?\}$ induces such a fragment. Note that the mentioned fragment of ICPDL is still strong enough to capture the information logic DAL.

Acknowledgements I am indebted to Ulrike Sattler, Lidia Tendera, and Martin Lange for many intense and fruitfull discussions about PDL with intersection.

References

1. L. Afanasiev, P. Blackburn, I. Dimitriou, B. Gaiffe, E. Goris, M. Marx, and M. de Rijke. PDL for ordered trees. *Journal of Applied Non-Classical Logic*, 2005. To appear.
2. N. Alechina, S. Demri, and M. de Rijke. A modal perspective on path constraints. *Journal of Logic and Computation*, 13(6):939–956, 2003.
3. F. Baader. Augmenting concept languages by transitive closure of roles: An alternative to terminological cycles. In *Proc. of IJCAI-91*, pages 446–451, Sydney, Australia, 1991.
4. P. Balbiani and D. Vakarelov. Iteration-free PDL with intersection: a complete axiomatization. In *Fundamenta Informaticae*, volume 45, pages 1–22. 2001.
5. D. Berardi, D. Calvanese, G. De Giacomo, M. Lenzerini, and M. Mecella. Automatic composition of e-services that export their behavior. In *Proc. of the 1st Int. Conf. on Service Oriented Computing (ICSOC 2003)*, volume 2910 of *LNCS*, pages 43–58. Springer, 2003.
6. D. Calvanese, G. De Giacomo, and M. Lenzerini. On the decidability of query containment under constraints. In *Proc. of PODS'98*, pages 149–158, 1998.
7. R. Danecki. Nondeterministic propositional dynamic logic with intersection is decidable. In *Proc. of the Fifth Symposium on Computation Theory*, volume 208 of *LNCS*, pages 34–53, Zaborów, Poland, Dec. 1984. Springer.
8. G. De Giacomo and M. Lenzerini. PDL-based framework for reasoning about actions. In *Proc. of AI*IA'95*, volume 992, pages 103–114. Springer, 1995.
9. F. M. Donini, M. Lenzerini, D. Nardi, and W. Nutt. The complexity of concept languages. *Information and Computation*, 134(1):1–58, 1997.
10. R. Fagin, J. Y. Halpern, Y. Moses, and M. Y. Vardi. *Reasoning About Knowledge*. MIT Press, 1995.
11. L. Farinas Del Cerro and E. Orlowska. DAL-a logic for data analysis. *Theoretical Computer Science*, 36(2-3):251–264, 1985.
12. M. J. Fischer and R. E. Ladner. Propositional modal logic of programs. In *Conference record of the ninth annual ACM Symposium on Theory of Computing*, pages 286–294. ACM Press, 1977.
13. M. J. Fischer and R. E. Ladner. Propositional dynamic logic of regular programs. *J. Comput. Syst. Sci.*, 18:194–211, 1979.
14. G. D. Giacomo and M. Lenzerini. Boosting the correspondence between description logics and propositional dynamic logics. In *Proc. of AAAI'94. Volume 1*, pages 205–212. AAAI Press, 1994.
15. D. Harel. Dynamic logic. In *Handbook of Philosophical Logic, Volume II*, pages 496–604. D. Reidel Publishers, 1984.
16. D. Kozen. Results on the propositional μ-calculus. In *Automata, Languages and Programming, 9th Colloquium*, volume 140 of *Lecture Notes in Computer Science*, pages 348–359. Springer-Verlag, 1982.
17. M. Lange. A lower complexity bound for propositional dynamic logic with intersection. In *Advances in Modal Logic Volume 5*. King's College Publications, 2005.
18. M. Lange and C. Lutz. 2-ExpTime lower bounds for propositional dynamic logics with intersection. 2005. Submitted.
19. C. Lutz. PDL with intersection and converse is decidable. LTCS-Report 05-05, Technical University Dresden, 2005. Available from http://lat.inf.tu-dresden.de/research/reports.html.

20. C. Lutz and U. Sattler. Mary likes all cats. In *Proc. of DL2000*, number 33 in CEUR-WS (http://ceur-ws.org/), pages 213–226, 2000.
21. F. Massacci. Decision procedures for expressive description logics with role intersection, composition and converse. In *Proc. of IJCAI-01*, pages 193–198, San Francisco, CA, Aug. 4–10 2001. Morgan Kaufmann Publishers, Inc.
22. V. Pratt. Considerations on floyd-hoare logic. In *FOCS: IEEE Symposium on Foundations of Computer Science (FOCS)*, 1976.
23. K. D. Schild. A correspondence theory for terminological logics: Preliminary report. In *Proc. of IJCAI-91*, pages 466–471. Morgan Kaufmann, 1991.
24. M.Y. Vardi and P. Wolper. Automata-theoretic techniques for modal logic of programs. Journal of Computer and System Sciences **32** (1986) 183–221