A trace semantics for Positive Core XPath

Pieter Hartel, Univ. of Twente, http://www.cs.utwente.nl/~pieter

***Abstract*— We provide a novel trace semantics for positive core XPath that exposes all intermediate nodes visited by the query engine. This enables a detailed analysis of all information relevant to the query. We give two examples of such analyses in the form of access control policies. We translate positive core XPath into Linear Temporal Logic, showing that branching structures can be linearised effectively. The translation is proved correct. We use the SPIN model checker in a proof of concept implementation to resolve the queries, and to perform the access control. The performance of the implementation is shown to be competitive.**

I. Introduction

Many approaches towards Access control on XML data use XPath (directly or indirectly) both for the queries and for access control (e.g. [5]). We are interested in combining a more flexible, logical approach [1] to access control, with the standard XPath based querying. An XPath (version 1.0 [10]) query is normally resolved by giving the answer set. This hides intermediate nodes visited by the query engine, which might contain sensitive information. We intend to expose this information so that it can be analysed, for example from the point of view of access control.

a) Example 1: Consider the family tree of Fig. 1 with query_1 asking for all family members with following siblings:

$$\mathsf{query}_1 = \mathsf{descendant}\ ::\ \texttt{"*"}[\mathsf{following\_sibling}\ ::\ \texttt{"*"}]$$

The answer set (i.e. Cain and Abel) does not reveal (1) the name of some of the following siblings (i.e. Seth), (2) that one of the members of the answer set is in fact a following sibling himself (i.e. Abel), and (3) the multiplicity of the answers (Cain is included for two reasons). So the answer set hides information that is available to the query engine. This information may be sensitive, and we are interested in making this information available for analysis. We achieve this by resolving a query not to the answer set but to the entire trace from the root produced by the query engine. For the example above there are three traces:

$$\begin{aligned}\mathsf{results}_1 = \{&[\mathsf{Root},\ \mathsf{Adam},\ \mathsf{Cain},\ (\mathsf{Abel},\ \mathsf{Seth}),\ \mathsf{Cain}],\\ &[\mathsf{Root},\ \mathsf{Adam},\ \mathsf{Cain},\ (\mathsf{Abel}),\ \mathsf{Cain}],\\ &[\mathsf{Root},\ \mathsf{Adam},\ \mathsf{Abel},\ (\mathsf{Seth}),\ \mathsf{Abel}]\}\end{aligned}$$

Some tags, like Abel and Seth in the first trace, are shown in parentheses to indicate that they are the result of exploring the predicate $[\mathsf{following\_sibling}\ ::\ \texttt{"*"}]$ of query_1. Other tags, such as Cain, are shown twice in the first trace because they have been visited twice: the first time while moving right from Adam to Cain and the second time returning from Seth to Cain.

We can now use the information contained in a trace for analysis purposes, such as access control. We give two examples.

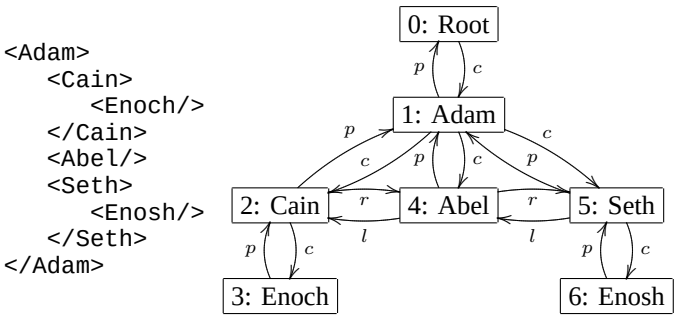

Fig. 1. Sample family tree in XML format (left) and in navigational format (right). The edge labels are: c for children, p for the parent, and l and r for the immediate sibling to the left and right respectively.

Firstly, suppose that (if only for historical reasons) the node tagged Cain should not be included in a trace that contains Abel also. Furthermore, we should like to be free to choose whether to access Cain first, or whether to access Abel first. This corresponds to (the object specification of) a Chinese wall policy [6], where Cain and Abel are in the same conflict of interest class. XPath is not powerful enough to formulate such a *general* policy because we do not know a-priori which axes to navigate to travel between members of a conflict of interest class. All we could hope to do is to formulate a *specific* policy for each query. To solve this problem we use Linear Temporal Logic (LTL, see Appendix A) to express the policy as follows (for generality extending the conflict of interest class to all children of Adam):

$$\mathsf{Chinese\_wall} = \Box(\ \mathsf{Cain}\ \rightarrow\ \neg\ \Diamond(\ \mathsf{Abel}\ \vee\ \ \mathsf{Seth}))$$

The formula Chinese_wall states that we should always (operator $\Box$) have that as soon as we encounter Cain, then we must not eventually (operator $\Diamond$) encounter either Abel or Seth. This corresponds to the mandatory aspect of the Chinese wall policy. The formula Chinese_wall does not insist that Cain is ever encountered, which corresponds to the discretionary aspect of the Chinese wall policy.

Secondly, using an idea of de Alfaro [12], suppose that every trace to a confidential node Cain must pass through an access control node Adam, thus blocking access via Abel. This can be formalised intuitively in LTL with past operators (an equivalent LTL expression with only future operators exists but it is less intuitive [32]):

$$\mathsf{Access\_control} = \Box(\ \mathsf{Cain}\ \rightarrow\ \Diamond^{-1}\ \ \mathsf{Adam})$$

The formula Access_control states that any access of Cain is due to some earlier access of Adam.

The two examples above only mention the object specification of an access control policy; we have tacitly assumed that

the subject making the query is identified and authenticated, that the authorisation is positive only, and that the privilege is assumed to be ”read”. Extension to more aspects of access control policies is future work.

Having motivated using LTL to express access control policies, a natural target for expressing a query is also LTL, so that we can combine them simply with a logical $\wedge$ operator, using the same formalism and implementation for both querying and access control. Therefore, the focus of the paper is on the semantics of positive core XPath because this can be translated efficiently into LTL. The main contributions are (1) a novel trace semantics for positive core XPath, (2) the translation of positive core XPath into LTL, (3) the correctness proof of the translation with respect to the trace semantics, and (4) a proof of concept implementation of the system.

The next section discusses related work. Sec. III motivates the positive core XPath subset. Sec. IV formalises undecorated XML trees. Sec. V defines the Kripke structure that forms the link between the formalised XML tree representation and the semantics of LTL. Sec. VI defines the embedding of positive core XPath into LTL via a translation algorithm. Sec. VII provides a natural semantics for positive core XPath. Sec. VIII presents the implementation of the positive core XPath engine using the SPIN model checker [20], and compares the performance of the implementation to that of state-of-the-art XPath query engines. The last Sec. concludes and gives ideas for future work. Appendix B gives a correctness proof of the positive core XPath translation with respect to the natural semantics.

II. Related work

Our work has similarities with the work of Afanasiev et al [2], who translate the downwards fragment of XPath into the existential fragment of Computation Tree Logic (CTL) [11], using the nuSMV model checker as the query engine. The differences include: (1) we are interested in trace semantics, whereas Afanasiev et al work with the standard semantics for answer sets; (2) our method of model building is orders of magnitude more efficient, and (3) we support all navigational axes, not just the downward axes. The efficiency of our method is mainly due to the judicious use of Embedded C code support provided by the SPIN model checker. Since SPIN supports LTL (and not CTL), we represent XPath queries using LTL, rather than CTL. While using the latter is more intuitive, we believe that our LTL rendering of XPath is still relatively simple.

Benedikt et al [3], and Marx [25], [24] study the expressive power of various fragments of XPath, including positive core XPath by embedding in various logics.

XML based access control offers three fundamental choices [23]. Should the XML data be filtered according to the access control policy: (1) before a query is applied, (2) after the query is applied, or (3) should the query be rewritten? Security views are an example of case (1). However, security views are expensive to compute and to maintain, which is why Fan et al [8] propose a method of avoiding to build security views using efficient query optimization techniques. Our approach to combining a query with an access control object specification is an example of case (3): the model checker ensures that only relevant parts of the state space are explored. Luo et al [23] also perform query rewriting, but consider forward axes only. Murata et al [28] use static analysis techniques to optimise query processing.

Bertino and Ferrari [4] present a versatile system for authoring XML based access control policies. Both the subject and the object are represented by XPath expressions. The policies themselves are again XML documents. Milau and Suciu [27] use XQuery (and thus also XPath) to state access control policies.

Fundulaki and Marx [14] use XPath to represent the object specification of an access control policy, which as we have shown is less powerful than using LTL for the same purpose.

Fu et al [13] use SPIN to model check XPath queries but their approach is radically different from ours in the sense that both the XML data and the query are part of the model. Fu et al use LTL formulae to specify liveness properties of the model, where we use LTL for the queries. Fu et al do not present performance data.

III. Positive core XPath

Full XPath is impractical to use as a tool for investigating the fundamental relation between query and access control. Several subsets have been defined, such as Core XPath [16], Simple XPath [2], and Navigational XPath [26]. We adopt a similar approach in that we omit expressions and focus on location paths and predicates. Contrary to some of the work cited earlier, we do support most (11 of the 13) axes, omitting attribute and namespace only. We omit negations for reasons to be explained later. Our subset is essentially positive core XPath, which is core XPath [16] without negations. The abstract syntax of positive core XPath is:

$$
\begin{aligned}
&\mathbb{X} \equiv \mathbb{X} \parallel \mathbb{X} \mid / \mathbb{X} \mid \mathbb{X} / \mathbb{X} \mid \mathbb{X}[\mathbb{Q}] \mid \mathbb{A} :: \mathbb{L} \\
&\mathbb{Q} \equiv \mathbb{X} \\
&\mathbb{A} \equiv \mathsf{self} \mid \\
&\qquad \mathsf{child} \mid \mathsf{descendant} \mid \mathsf{descendant\_or\_self} \mid \\
&\qquad \mathsf{parent} \mid \mathsf{ancestor} \mid \mathsf{ancestor\_or\_self} \mid \\
&\qquad \mathsf{preceding\_sibling} \mid \mathsf{following\_sibling} \mid \\
&\qquad \mathsf{preceding} \mid \mathsf{following}
\end{aligned}
$$

The node test in a step is restricted to a name test (i.e. kind tests are not supported, which is consistent with the use of undecorated XML data). We use location paths $\mathbb{X}$ by way of predicates $\mathbb{Q}$.

A. Disjunction

A typical answer contains several results. Hence we should expect the trace semantics of a query to be a set of traces. The semantics of the $\parallel$ operator applied to two queries is therefore the union of the traces returned for each query separately.

b) Example 2: Consider query_2 below, which in the standard semantics yields an answer set consisting of Cain and Seth:

$$\mathsf{query}_2 = \mathsf{descendant} :: \texttt{"*"} \\ [\mathsf{child} :: \mathsf{Enoch} \vee \mathsf{child} :: \mathsf{Enosh}]$$

The standard semantics for XPath prescribes that the result of the predicate should be a Boolean. In our interpretation we take an empty trace to mean false and a non-empty trace to represent true [33]. However, we should also like to preserve the traces resulting from the predicate, because all visited nodes must be kept for further analysis. This leads to the idea that the result should consist of two traces, both with an initial segment corresponding to $\mathsf{descendant} :: \texttt{"*"}$. Then the traces differ: one contains the trace corresponding to the left hand side of the $\vee$ operator, and the other takes care of the right hand side. In both cases a common trailing segment follows. The initial and trailing segment are effectively copied and concatenated to each intermediate segment, yielding the following result:

$$\mathsf{results}_2 = \{[\mathsf{Root}, \mathsf{Adam}, \mathsf{Cain}, (\mathsf{Enoch}), \mathsf{Cain}], \\ [\mathsf{Root}, \mathsf{Adam}, \mathsf{Seth}, (\mathsf{Enosh}), \mathsf{Seth}]\}$$

With this "copying" semantics in mind the $\vee$ and $\|$ operators are identified, thus obviating the need for a separate $\vee$ operator [3, Proposition 2] and so that query_2 is interpreted as:

$$\mathsf{query}_{2'} = \mathsf{descendant} :: \texttt{"*"} \\ [\mathsf{child} :: \mathsf{Enoch} \parallel \mathsf{child} :: \mathsf{Enosh}]$$

B. Conjunction

The trace semantics of a location path with a predicate is the concatenation of the trace of the location path and the trace of the predicate.

c) Example 3: Consider query_3, which in the standard semantics returns the singleton answer set $\{\mathsf{Adam}\}$:

$$\mathsf{query}_3 = \mathsf{descendant} :: \texttt{"*"} \\ [\mathsf{child} :: \mathsf{Cain} \wedge \mathsf{child} :: \mathsf{Abel}]$$

The resulting trace should contain an initial segment corresponding to the location path $\mathsf{descendant} :: \texttt{"*"}$. However for the predicate to succeed we must be sure that there is at least one non-empty trace corresponding to the left hand side of the $\wedge$ operator as well as a non-empty trace corresponding to the right hand side. Both non-empty traces must be returned as part of the full trace, which we achieve by concatenating the results. We have arbitrarily chosen to concatenate the right hand side trace onto the left hand side trace; interleaving or reordering would also be possible but this is subject to further work. In all cases a trailing segment will follow. The result then becomes:

$$\mathsf{result}_3 = [\mathsf{Root}, \mathsf{Adam}, (\mathsf{Cain}), (\mathsf{Abel}), \mathsf{Adam}]$$

This, however, is exactly the trace that would be returned by the following query:

$$\mathsf{query}_{3'} = \mathsf{descendant} :: \texttt{"*"} \\ [\mathsf{child} :: \mathsf{Cain}][\mathsf{child} :: \mathsf{Abel}]$$

Therefore we can dispense with the $\wedge$ operator also, as repeated use of predicates can achieve the desired effect [3, Proposition 2]. Using the fact that propositions with $\vee$, $\wedge$ and $\neg$ can always be written in conjunctive normal form [22], we can also remove all nested conjunctions and disjunctions.

C. Negation

Negation is a problem because it is unclear what trace to return for a negated predicate. Assume first that predicates are in conjunctive normal form, and that all occurrences of $\wedge$ and $\vee$ have been removed as described above. Then there are only atomic propositions and negated atomic propositions left.

d) Example 4: Consider query_4, which in the standard semantics returns the singleton answer set $\{\mathsf{Root}\}$:

$$\mathsf{query}_4 = \mathsf{descendant\_or\_self} :: \texttt{"*"}[\neg\ \mathsf{parent} :: \texttt{"*"}]$$

The question now is: which traces(s) to return for the predicate? (a) Should it be all possible traces that do not satisfy the predicate? This would be infinitely many with the 11 axes of XPath! (b) Or should the trace be empty? This would jeopardise our ability to analyse the trace properly: Consider our family tree again with a query asking for brothers not involved in fratricide. Should the query return $\{\mathsf{Seth}\}$? Or should it return an empty answer set because it violates the Chinese wall policy? (c) The most likely possibility is to label segments of the trace that correspond to negated steps, so that these can be distinguished from positive steps in the analysis. This, however, we leave as future work and for now omit negation. Note that often in policy specifications the same approach is taken: what is not explicitly allowed is forbidden, hence negative steps are not always necessary [19]. However, see Sec. VIII for an experiment with alternative (a) in SPIN. This concludes the motivation of the positive core XPath subset.

IV. XML DATA REPRESENTATION

XPath queries operate on an appropriate representation of the data that we assume to be bulk loaded; dealing with updates and inserts is beyond the scope of the paper. To provide efficient support for the 11 axes in queries, a representation of the XML data is needed that is slightly more sophisticated than a tree. Fig. 1 shows the navigational representation that we adopt. The four types of edges shown are p for parent, c for child, r for *immediate* following_sibling and l for *immediate* preceding_sibling. All other axes (except self) are supported by traversing more than one edge.

The nodes of the graph are represented by a given set $\mathbb{N}$, which in the case of our running example is:

$$\mathbb{N} \equiv \mathsf{Root} \mid \mathsf{Adam} \mid \mathsf{Cain} \mid \mathsf{Enoch} \mid \mathsf{Abel} \mid \mathsf{Seth} \mid \mathsf{Enosh}$$

The edges are represented by four functions, one for each type of edge (i.e. up_d, down_d, left_d, and $\mathsf{right}_\mathsf{d}$). In addition we need a function to return to the root (root_d), as well as a function to stay put (here_d). We show only the definition of down_d, the remaining functions are similar.

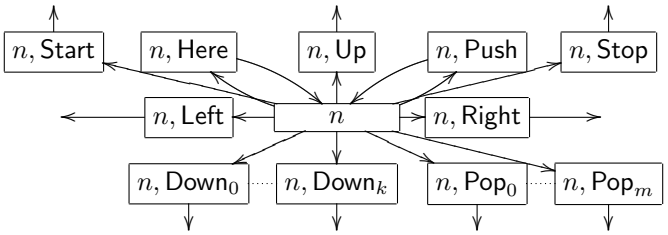

Fig. 2. State n showing the nine possible directions for reaching a successor state.

$$
\begin{array}{ll}
\mathsf{down_d} & :: \mathbb{N}\to\{\mathbb{N}\} \\
\mathsf{down_d(Root)} & = \{\mathsf{Adam}\} \\
\mathsf{down_d(Adam)} & = \{\mathsf{Cain},\ \mathsf{Abel},\ \mathsf{Seth}\} \\
\mathsf{down_d(Cain)} & = \{\mathsf{Enoch}\} \\
\mathsf{down_d(Seth)} & = \{\mathsf{Enosh}\} \\
\mathsf{down_d\ \_} & = \{\}
\end{array}
$$

We follow the approach of the work cited at the beginning of Sec. III to focus purely on the tags of XML data, omitting all other information, so that this concludes the presentation of our representation of an undecorated XML tree.

V. Kripke structure

Before we can give the translation of XPath into LTL we must develop a Kripke structure for the resulting logic. The structure is based on the definition of two sets, $\mathbb{N}$, given earlier to represent the nodes, and $\mathbb{D}$ to represent the directions corresponding to the axes (Here for self, Up for parent, Down for child, Left for immediate preceding_sibling, and Right for immediate following_sibling) as well as a further four directions (Start, Stop, Push, and Pop) to be discussed shortly.

$$
\begin{array}{l}
\mathbb{D} \equiv \mathsf{Start} \mid \mathsf{Here} \mid \mathsf{Up} \mid \mathsf{Down} \mid \mathsf{Left} \mid \mathsf{Right} \mid \\
\qquad \mathsf{Push} \mid \mathsf{Pop} \mid \mathsf{Stop}
\end{array}
$$

Fig. 2 shows all nodes in the Kripke structure that correspond to a single node of an XML tree. This representation is quadratic in the number of nodes of the original XML tree, which is clearly inefficient. We will come back to this issue in Sec. VIII, but we need to make the situation worse first by considering how to deal with predicates. Referring back to the introduction, we saw that predicate yields a trace segment that returns to the starting node of the segment, to linearise a finite branching structure. To support this we need a stack of nodes in the Kripke structure. The states of the Kripke structure are defined by the triple $\mathbb{S}$ below, where $\overline{\mathbb{N}}$ represents the stack (i.e. a list of nodes):

$$
\begin{array}{l}
\mathbb{S} \equiv (\mathbb{N},\ \mathbb{D},\ \overline{\mathbb{N}}) \\
\overline{\mathbb{N}} \equiv [\mathbb{N}]
\end{array}
$$

Given an XML tree with n nodes, and a query with predicates nested to a depth of d, the state space in the worst case grows as $n^{(d+1)}$. In practice the state space remains small as we shall see later (Sec. VIII).

We now have all ingredients to show the Kripke structure of our running example ϕ below. Here the function $\sigma(\mathsf{n},\mathsf{d},\mathsf{s})$ defines the set of all possible successor states. For example the successor state of $(\mathsf{n},\mathsf{Push},\mathsf{s})$ consists of the set of states $(\mathsf{n},\mathsf{d}',\mathsf{n}:\mathsf{s})$, where d' ranges over all possible 9 directions in $\mathbb{D}$, and where $\mathsf{n}:\mathsf{s}$ represents the current stack extended with the current node. Summarising, the interpretation of state $(\mathsf{n},\mathsf{d},\mathsf{s})$ is: we are now at node n going in the direction d, with current stack s.

$$
\begin{array}{ll}
\mathbb{M}\alpha \equiv (\{\alpha\},\ \alpha\to\{\alpha\},\ \alpha\to\{\mathbb{L}\}) & \\
\phi & :: \mathbb{M}\ \mathbb{S} \\
\phi & = (\{(\mathsf{n},\ \mathsf{d},\ \mathsf{s}) \mid \mathsf{n}\in\mathbb{N} \wedge\ \mathsf{d}\in\mathbb{D} \wedge\ \mathsf{s}\in\overline{\mathbb{N}}\},\ \sigma,\ \lambda) \\
\quad \textbf{where} & \\
\quad \sigma(\mathsf{n},\ \mathsf{Start},\ \mathsf{s}) & = \{(\mathsf{root_d}\ \mathsf{n},\ \mathsf{d}',\ \mathsf{s}) \mid \mathsf{d}'\in\mathbb{D}\} \\
\quad \sigma(\mathsf{n},\ \mathsf{Here},\ \mathsf{s}) & = \{(\mathsf{n},\ \mathsf{d}',\ \mathsf{s}) \mid \mathsf{d}'\in\mathbb{D}\} \\
\quad \sigma(\mathsf{n},\ \mathsf{Up},\ \mathsf{s}) & = \{(\mathsf{n}',\ \mathsf{d}',\ \mathsf{s}) \mid \mathsf{n}'\in\mathsf{up_d}\ \mathsf{n} \wedge\ \mathsf{d}'\in\mathbb{D}\} \\
\quad \sigma(\mathsf{n},\ \mathsf{Down},\ \mathsf{s}) & = \{(\mathsf{n}',\ \mathsf{d}',\ \mathsf{s}) \mid \mathsf{n}'\in\mathsf{down_d}\ \mathsf{n} \wedge\ \mathsf{d}'\in\mathbb{D}\} \\
\quad \sigma(\mathsf{n},\ \mathsf{Left},\ \mathsf{s}) & = \{(\mathsf{n}',\ \mathsf{d}',\ \mathsf{s}) \mid \mathsf{n}'\in\mathsf{left_d}\ \mathsf{n} \wedge\ \mathsf{d}'\in\mathbb{D}\} \\
\quad \sigma(\mathsf{n},\ \mathsf{Right},\ \mathsf{s}) & = \{(\mathsf{n}',\ \mathsf{d}',\ \mathsf{s}) \mid \mathsf{n}'\in\mathsf{right_d}\ \mathsf{n} \wedge\ \mathsf{d}'\in\mathbb{D}\} \\
\quad \sigma(\mathsf{n},\ \mathsf{Push},\ \mathsf{s}) & = \{(\mathsf{n},\ \mathsf{d}',\ \mathsf{n}:\mathsf{s}) \mid \mathsf{d}'\in\mathbb{D}\} \\
\quad \sigma(\mathsf{n},\ \mathsf{Pop},\ \mathsf{n}':\mathsf{s}) & = \{(\mathsf{n}',\ \mathsf{d}',\ \mathsf{s}) \mid \mathsf{d}'\in\mathbb{D}\} \\
\quad \sigma(\mathsf{n},\ \mathsf{Pop},\ []) & = \{\} \\
\quad \sigma(\mathsf{n},\ \mathsf{Stop},\ \mathsf{s}) & = \{\} \\
\quad \lambda(\mathsf{n},\ \mathsf{d},\ \mathsf{s}) & = \{\mathsf{n},\ \mathsf{d}\}
\end{array}
$$

We have tacitly assumed here that all nodes in the tree have a unique tag. If this is not the case, the Kripke structure must be extended with a unique identifier for each node. We will ensure that this is the case in the high performance SPIN models.

This concludes the presentation of the Kripke structure so that we can turn our attention to the translation of positive core XPath into LTL.

VI. Translation of positive core XPath into LTL

The function $\mathcal{T}_\mathsf{x}$ below translates an XPath query into an LTL formula. The function takes a query as its first argument, and an LTL formula ϕ which represents what should happen after we have dealt with the query. Consider for example the first clause of $\mathcal{T}_\mathsf{x}$. Since ϕ represents what happens after $\mathsf{xp_1}[\![\mathsf{xp_2}$, ϕ must happen after $\mathsf{xp_1}$ as well as $\mathsf{xp_2}$. This corresponds to the "copying" semantics alluded to in the introduction.

Consider also the second clause, which states that for an absolute query /xp we go from the current node in the Start direction, leading to the node Root in the next (X) step. Then we continue with xp, ultimately followed by ϕ. The remaining clauses are intended to be self explanatory.

$$
\begin{array}{ll}
\mathcal{T}_\mathsf{x} & :: \mathbb{X}\to\mathbb{T}\to\mathbb{T} \\
\mathcal{T}_\mathsf{x}[\![\mathsf{xp_1} \parallel \mathsf{xp_2}]\!]\phi & = \mathcal{T}_\mathsf{x}[\![\mathsf{xp_1}]\!]\phi\ \vee\ \mathcal{T}_\mathsf{x}[\![\mathsf{xp_2}]\!]\phi \\
\mathcal{T}_\mathsf{x}[\![/\ \mathsf{xp}]\!]\phi & = \mathsf{Start}\ \wedge\ \mathsf{X}(\mathsf{Root}\ \wedge\ \mathcal{T}_\mathsf{x}[\![\mathsf{xp}]\!]\phi) \\
\mathcal{T}_\mathsf{x}[\![\mathsf{xp_1}\ /\ \mathsf{xp_2}]\!]\phi & = \mathcal{T}_\mathsf{x}[\![\mathsf{xp_1}]\!](\mathcal{T}_\mathsf{x}[\![\mathsf{xp_2}]\!]\phi) \\
\mathcal{T}_\mathsf{x}[\![\mathsf{xp_1}[\mathsf{xp_2}]]\!]\phi & = \mathcal{T}_\mathsf{x}[\![\mathsf{xp_1}]\!](\mathsf{Push}\ \wedge \\
 & \qquad \mathsf{X}(\mathcal{T}_\mathsf{x}[\![\mathsf{xp_2}]\!](\mathsf{Pop}\ \wedge\ \mathsf{X}\ \phi))) \\
\mathcal{T}_\mathsf{x}[\![\mathsf{a}\ ::\ \mathsf{l}]\!]\phi & = \mathcal{T}_\mathsf{a}[\![\mathsf{a}]\!](\ \mathsf{l}\ \wedge\ \phi)
\end{array}
$$

The function $\mathcal{T}_a$ below follows the same pattern as $\mathcal{T}_x$. The first argument is an axis and the second argument ϕ corresponds to the query that must be matched after the current axis has been matched. For example the first clause states that the proposition Here must be true in the current state, and that ϕ must hold in the next state.

Also note the difference between descendant and descendant_or_self. In the former we check first that a move in the direction Down can be made, optionally followed by a further sequence of moves in the Down direction until finally a state is found in which ϕ is true. In the latter case we accept either a move to the current node (direction Here) or the moves implied by the axis descendant. The cases for the remaining axes are expected to be self explanatory.

$$
\begin{array}{ll}
\mathcal{T}_a & :: \mathbb{A}\rightarrow\mathbb{T}\rightarrow\mathbb{T} \\
\mathcal{T}_a[\![\mathsf{self}]\!]\phi & = \mathsf{Here} \wedge \mathsf{X}\ \phi \\
\mathcal{T}_a[\![\mathsf{child}]\!]\phi & = \mathsf{Down} \wedge \mathsf{X}\ \phi \\
\mathcal{T}_a[\![\mathsf{parent}]\!]\phi & = \mathsf{Up} \wedge \mathsf{X}\ \phi \\
\mathcal{T}_a[\![\mathsf{descendant}]\!]\phi & = \mathsf{Down} \wedge \mathsf{X}(\mathsf{Down\ U}\ \phi) \\
\mathcal{T}_a[\![\mathsf{ancestor}]\!]\phi & = \mathsf{Up} \wedge \mathsf{X}(\mathsf{Up\ U}\ \phi) \\
\mathcal{T}_a[\![\mathsf{descendant\_or\_self}]\!]\phi & = \mathcal{T}_a[\![\mathsf{self}]\!]\phi \vee \mathcal{T}_a[\![\mathsf{descendant}]\!]\phi \\
\mathcal{T}_a[\![\mathsf{ancestor\_or\_self}]\!]\phi & = \mathcal{T}_a[\![\mathsf{self}]\!]\phi \vee \mathcal{T}_a[\![\mathsf{ancestor}]\!]\phi \\
\mathcal{T}_a[\![\mathsf{following\_sibling}]\!]\phi & = \mathsf{Right} \wedge \mathsf{X}(\mathsf{Right\ U}\ \phi) \\
\mathcal{T}_a[\![\mathsf{preceding\_sibling}]\!]\phi & = \mathsf{Left} \wedge \mathsf{X}(\mathsf{Left\ U}\ \phi) \\
\mathcal{T}_a[\![\mathsf{following}]\!]\phi & = \mathsf{Up\ U}(\mathsf{Right} \wedge \\
 & \quad \mathsf{X}(\mathsf{Right\ U}(\mathsf{Down\ U}\ \phi))) \\
\mathcal{T}_a[\![\mathsf{preceding}]\!]\phi & = \mathsf{Up\ U}(\mathsf{Left} \wedge \\
 & \quad \mathsf{X}(\mathsf{Left\ U}(\mathsf{Down\ U}\ \phi)))
\end{array}
$$

A. Examples of the translation

We present some examples of the translation.

e) Example 5: Query_5 delivers the traces from the current context node to a child with tag Adam. There is one such trace from the Root.

$$
\begin{array}{l}
\mathsf{query}_5 = \mathsf{child} :: \mathsf{Adam} \\
\mathsf{ltl}_5 \quad = \mathcal{T}_x[\![\mathsf{query}_5]\!]\ \mathsf{Stop} \\
\qquad = \mathsf{Down} \wedge \mathsf{X}(\ \mathsf{Adam} \wedge \ \mathsf{Stop}) \\
\mathsf{result}_5 = [(\mathsf{Root},\ \mathsf{Down},\ []),\ (\mathsf{Adam},\ \mathsf{Stop},\ [])]
\end{array}
$$

The LTL translation ltl_5 and the trace result_5 satisfy: $\mathsf{result}_5 \models \mathsf{ltl}_5$ □

f) Example 6: The longer query_6 delivers the traces from the current context node to a descendant with tag Adam, then to a child Seth, then to a preceding sibling Abel. There is one such trace from the Root.

$$
\begin{array}{l}
\mathsf{query}_6 = \mathsf{descendant} :: \mathsf{Adam}\ / \\
\qquad \mathsf{child} :: \mathsf{Seth}\ / \\
\qquad \mathsf{preceding\_sibling} :: \mathsf{Abel}\ / \\
\qquad \mathsf{preceding\_sibling} :: \mathsf{Cain} \\
\mathsf{ltl}_6 \quad = \mathcal{T}_x[\![\mathsf{query}_6]\!]\ \mathsf{Stop} \\
\qquad = \mathsf{Down} \wedge \mathsf{X}(\ \mathsf{Down\ U}(\ \mathsf{Adam} \wedge \ \mathsf{Down} \wedge \\
\qquad\quad \mathsf{X}(\ \mathsf{Seth} \wedge \ \mathsf{Left} \wedge \\
\qquad\quad \mathsf{X}(\ \mathsf{Left\ U}(\ \mathsf{Abel} \wedge \ \mathsf{Left} \wedge \\
\qquad\quad \mathsf{X}(\ \mathsf{Left\ U}(\ \mathsf{Cain} \wedge \ \mathsf{Stop})))))))) \\
\mathsf{result}_6 = [(\mathsf{Root},\ \mathsf{Down},\ []),\ (\mathsf{Adam},\ \mathsf{Down},\ []), \\
\qquad (\mathsf{Seth},\ \mathsf{Left},\ []),\ (\mathsf{Abel},\ \mathsf{Left},\ []), \\
\qquad (\mathsf{Cain},\ \mathsf{Stop},\ [])]
\end{array}
$$

The trace result_6 is a model for the LTL formula ltl_6 with respect to the given Kripke structure: $\mathsf{result}_6 \models \mathsf{ltl}_6$ □

g) Example 7: Query_7 cannot be matched because the Root is not a proper descendant of itself.

$$
\begin{array}{l}
\mathsf{query}_7 = \mathsf{descendant} :: \mathsf{Root} \\
\mathsf{ltl}_7 \quad = \mathcal{T}_x[\![\mathsf{query}_7]\!]\ \mathsf{Stop} \\
\qquad = \mathsf{Down} \wedge \mathsf{X}(\ \mathsf{Down\ U}(\ \mathsf{Root} \wedge \ \mathsf{Stop})) \\
\mathsf{results}_7 = \{\}
\end{array}
$$

h) Example 1 revisited: We now revisit query_1 to demonstrate how predicates are translated.

$$
\begin{array}{l}
\mathsf{ltl}_1 \quad = \mathcal{T}_x[\![\mathsf{query}_1]\!]\ \mathsf{Stop} \\
\qquad = \mathsf{Down} \wedge \mathsf{X}(\ \mathsf{Down\ U}(\ \mathsf{Push} \wedge \\
\qquad\quad \mathsf{X}(\ \mathsf{Right} \wedge \mathsf{X}(\ \mathsf{Right\ U}(\ \mathsf{Pop} \wedge \mathsf{X}\ \ \mathsf{Stop}))))) \\
\mathsf{results}_1 = \{[(\mathsf{Root},\ \mathsf{Down},\ []),\ (\mathsf{Adam},\ \mathsf{Down},\ []), \\
\qquad (\mathsf{Cain},\ \mathsf{Push},\ []),\ (\mathsf{Cain},\ \mathsf{Right},\ [\mathsf{Cain}]), \\
\qquad (\mathsf{Abel},\ \mathsf{Right},\ [\mathsf{Cain}]),\ (\mathsf{Seth},\ \mathsf{Pop},\ [\mathsf{Cain}]), \\
\qquad (\mathsf{Cain},\ \mathsf{Stop},\ [])], \\
\qquad [(\mathsf{Root},\ \mathsf{Down},\ []),\ (\mathsf{Adam},\ \mathsf{Down},\ []), \\
\qquad (\mathsf{Cain},\ \mathsf{Push},\ []),\ (\mathsf{Cain},\ \mathsf{Right},\ [\mathsf{Cain}]), \\
\qquad (\mathsf{Abel},\ \mathsf{Pop},\ [\mathsf{Cain}]),\ (\mathsf{Cain},\ \mathsf{Stop},\ [])], \\
\qquad [(\mathsf{Root},\ \mathsf{Down},\ []),\ (\mathsf{Adam},\ \mathsf{Down},\ []), \\
\qquad (\mathsf{Abel},\ \mathsf{Push},\ []),\ (\mathsf{Abel},\ \mathsf{Right},\ [\mathsf{Abel}]), \\
\qquad (\mathsf{Seth},\ \mathsf{Pop},\ [\mathsf{Abel}]),\ (\mathsf{Abel},\ \mathsf{Stop},\ [])]\}
\end{array}
$$

As expected, all traces of the set $\mathsf{results}_1$ are models for the LTL formula ltl_1 with respect to the given Kripke structure, i.e.: $\bigwedge\{\mathsf{r} \models \mathsf{ltl}_1 \mid \mathsf{r} \in \mathsf{results}_1\}$ □

This concludes the translation of positive core XPath into LTL.

VII. Natural semantics for XPath

Borrowing ideas from Wadler's work [34], the semantics of positive core XPath below defines a relation between a trace and a query on the left hand side and a trace on the right hand side. The trace on the left hand side is the end point of the current trace, from which the current (context) node can be found. Consider for example the rule [abs] for absolute queries. The endpoint of the current trace is $(\mathsf{x}, \bot)$, where x is the current context node, and the direction in which to go is yet unknown ($\bot$). The premise of the rule asserts that the relative query xp started at the Root yields a trace xs', where the direction taken from the Root will be known (i.e. $\neq \bot$). The right hand side of the conclusion prepends the state $(\mathsf{x}, \mathsf{Start})$ to xs', to account for the fact that now we know in which direction to proceed from the original, initial node x. We hope that the remaining clauses are self explanatory. (As usual we omit explicit coercions, for example using the : operator for the concatenation of traces and traces, traces and elements etc.).

$$
\begin{array}{lll}
\mathbb{P} & \equiv & [(\mathbb{N},\ \mathbb{D})] \\
\rightarrow & :: & \langle \mathbb{P},\ \mathbb{X}\rangle \leftrightarrow \mathbb{P} \\
 & & \dfrac{\langle (x,\ \bot),\ xp_1\rangle \rightarrow xs'}{\langle (x,\ \bot),\ xp_1 \parallel xp_2\rangle \rightarrow xs'} \\
[\mathsf{bar}^1] & & \\
 & & \dfrac{\langle (x,\ \bot),\ xp_2\rangle \rightarrow xs'}{\langle (x,\ \bot),\ xp_1 \parallel xp_2\rangle \rightarrow xs'} \\
[\mathsf{bar}^2] & & \\
 & & \dfrac{\langle (\mathsf{Root},\ \bot),\ xp\rangle \rightarrow xs'}{\langle (x,\ \bot),\ /\ xp\rangle \rightarrow (x,\ \mathsf{Start}) : xs'} \\
[\mathsf{abs}] & & \\
 & & \dfrac{\langle (x,\ \bot),\ xp_1\rangle \rightarrow xs' : (x',\ \bot),\ \langle (x',\ \bot),\ xp_2\rangle \rightarrow xs''}{\langle (x,\ \bot),\ xp_1\ /\ xp_2\rangle \rightarrow xs' : xs''} \\
[\mathsf{slash}] & & \\
 & & \dfrac{\langle (x,\ \bot),\ xp_1\rangle \rightarrow xs' : (x',\ \bot),\ \langle (x',\ \bot),\ xp_2\rangle \rightarrow xs'' : (x'',\ \bot)}{\langle (x,\ \bot),\ xp_1[xp_2]\rangle \rightarrow (xs' : (x',\ \mathsf{Push}) : xs'' : (x'',\ \mathsf{Pop}) : (x',\ \bot)} \\
[\mathsf{pred}] & & \\
 & & \dfrac{xs' : (x',\ \bot) \in \mathcal{P}_a[\![a]\!](x,\ \bot)}{\langle (x,\ \bot),\ a :: l\rangle \rightarrow xs' : (x',\ \bot),\ \textbf{if } l = \text{"*"} \vee l = x'} \\
[\mathsf{step}] & &
\end{array}
$$

The semantic function $\mathcal{P}_x$ below provides a convenient interface to the natural semantics.

$$
\begin{array}{l}
\mathcal{P}_x :: \mathbb{X} \rightarrow (\mathbb{P} \rightarrow \{\mathbb{P}\}) \\
\mathcal{P}_x[\![xp]\!][(x,\ \bot)] = \{xs' : (x',\ \mathsf{Stop}) \mid xs' : (x',\ \bot) \in \langle (x,\ \bot),\ xp\rangle \rightarrow\}
\end{array}
$$

The rule [step] relies on the function $\mathcal{P}_a$ below to deal with the 11 axes of XPath.

$$
\begin{array}{ll}
\mathcal{P}_a :: \mathbb{A} \rightarrow (\mathbb{P} \rightarrow \{\mathbb{P}\}) \\
\mathcal{P}_a[\![\mathsf{self}]\!] = \mathsf{here}_p \\
\mathcal{P}_a[\![\mathsf{child}]\!] = \mathsf{down}_p \\
\mathcal{P}_a[\![\mathsf{parent}]\!] = \mathsf{up}_p \\
\mathcal{P}_a[\![\mathsf{descendant}]\!] = \mathsf{down}_p +_p \\
\mathcal{P}_a[\![\mathsf{ancestor}]\!] = \mathsf{up}_p +_p \\
\mathcal{P}_a[\![\mathsf{descendant\_or\_self}]\!] = \mathcal{P}_a[\![\mathsf{self}]\!] \vee_p \mathcal{P}_a[\![\mathsf{descendant}]\!] \\
\mathcal{P}_a[\![\mathsf{ancestor\_or\_self}]\!] = \mathcal{P}_a[\![\mathsf{self}]\!] \vee_p \mathcal{P}_a[\![\mathsf{ancestor}]\!] \\
\mathcal{P}_a[\![\mathsf{following\_sibling}]\!] = \mathsf{right}_p +_p \\
\mathcal{P}_a[\![\mathsf{preceding\_sibling}]\!] = \mathsf{left}_p +_p \\
\mathcal{P}_a[\![\mathsf{following}]\!] = \mathsf{horizontal}_p\ \mathsf{right}_p \\
\mathcal{P}_a[\![\mathsf{preceding}]\!] = \mathsf{horizontal}_p\ \mathsf{left}_p
\end{array}
$$

The function $\mathcal{P}_a$ in turn relies on a number of functions below to calculate the possible traces from the current node (again found in the endpoint of the current trace) in the direction indicated by the axis. For example down_p with a current node x yields a set of segments $[(x, \mathsf{Down}), (y, \bot)]$ where y ranges over all children of node x, as defined by the function down_d of Sec. IV. The result of down_d is empty if node x has no children.

$$
\begin{array}{ll}
\mathsf{go} :: \mathbb{D} \rightarrow (\mathbb{N} \rightarrow \{\mathbb{N}\}) \rightarrow (\mathbb{P} \rightarrow \{\mathbb{P}\}) \\
\mathsf{go\ d\ f}(x,\ \bot) = \{(x,\ d) : (y,\ \bot) \mid y \in f\ x\} \\
\mathsf{here}_p :: \mathbb{P} \rightarrow \{\mathbb{P}\} \\
\mathsf{here}_p = \mathsf{go\ Here\ here}_d \\
\mathsf{down}_p,\ \mathsf{up}_p :: \mathbb{P} \rightarrow \{\mathbb{P}\} \\
\mathsf{down}_p = \mathsf{go\ Down\ down}_d \\
\mathsf{up}_p = \mathsf{go\ Up\ up}_d \\
\mathsf{left}_p,\ \mathsf{right}_p :: \mathbb{P} \rightarrow \{\mathbb{P}\} \\
\mathsf{left}_p = \mathsf{go\ Left\ left}_d \\
\mathsf{right}_p = \mathsf{go\ Right\ right}_d
\end{array}
$$

The function $\mathsf{horizontal}_p$ is used by the axes preceding and following to discover trace segments corresponding to the nodes that precede the current node in XML document order.

$$
\begin{array}{l}
\mathsf{horizontal}_p :: (\mathbb{P} \rightarrow \{\mathbb{P}\}) \rightarrow (\mathbb{P} \rightarrow \{\mathbb{P}\}) \\
\mathsf{horizontal}_p\ fp = \mathsf{hhc}\ \vee_p ((\mathsf{up}_p +_p) \wedge_p\ \mathsf{hhc}) \\
\qquad \textbf{where} \\
\qquad h\ \ = fp +_p \\
\qquad \mathsf{hhc} = h\ \vee_p (h\ \wedge_p (\mathsf{down}_p +_p))
\end{array}
$$

Finally we need three operators ($+_p$, $\wedge_p$, and $\vee_p$) to glue trace segments together.

$$
\begin{array}{ll}
+_p :: (\mathbb{P} \rightarrow \{\mathbb{P}\}) \rightarrow (\mathbb{P} \rightarrow \{\mathbb{P}\}) \\
r +_p = r\ \vee_p (r\ \wedge_p\ r +_p) \\
\wedge_p,\ \vee_p :: (\mathbb{P} \rightarrow \{\mathbb{P}\}) \rightarrow (\mathbb{P} \rightarrow \{\mathbb{P}\}) \rightarrow (\mathbb{P} \rightarrow \{\mathbb{P}\}) \\
(r\ \wedge_p\ q)(x,\ \bot) = \{ys : zs \mid ys : (y,\ \bot) \in r(x,\ \bot) \wedge zs \in q(y,\ \bot)\} \\
(r\ \vee_p\ q)(x,\ \bot) = r(x,\ \bot) \cup\ q(x,\ \bot)
\end{array}
$$

This concludes the presentation of the Natural semantics of positive core XPath.

VIII. SPIN ENGINE

We now present two ways of representing the Kripke structure as an explicit state model for SPIN to show that in practical cases, the state space does not grow as in the worst case.

A. Pure Promela Model

The Promela model below is an optimised representation of the Kripke structure of Sec. V. The state consists of an `mtype` declaration introducing the nodes and directions, and three variables `tag`, `dir`, and `stack` representing the current tag, direction of travel, and stack.

```
mtype={ Root, Adam, Cain, Enoch, Abel,
        Seth, Enosh, Start, Here, Up, Down,
        Left, Right, Push, Pop, Stop };
mtype tag=Root;
mtype dir=Down;
byte  stack=0;
```

The XML tree is built using a series of macros `node`.... The first parameter is the node number as shown in Fig. 1, the second the tag and the remaining parameters are the node numbers of the parent, children, and the nodes immediately to the left and the right.

```
init{
nodeR(0,Root,1);
    node3(1,Adam,0,2,4,5);
        node1r(2,Cain,1,3,4);
            node0(3,Enoch,2);
        node1l0r(4,Abel,1,2,5);
        node1l1(5,Seth,1,4,6);
            node0(6,Enosh,5);
end: skip
}
```

We do not give the definitions of the macros as these are largely repetitive. Instead we show the expansion of the node with tag `Adam`. Starting at label `s1`, where 1 is the node number of `Adam`, there is a non-deterministic choice leading to all possible successor states of `s1`. Promela does not offer a "computed goto", so this has to be simulated for popping the stack. Promela models must be finite. Therefore, we limit the stack depth to 2, supporting a nesting level of 2 for predicates. (Using qualifier flattening [29] a nesting depth of 1 would be sufficient).

```
s1: if
 :: d_step{ tag=Adam; dir=Start }; goto s0
 :: d_step{ tag=Adam; dir=Here  }; goto s1
 :: d_step{ tag=Adam; dir=Up    }; goto s0
 :: d_step{ tag=Adam; dir=Down  }; goto s2
 :: d_step{ tag=Adam; dir=Down  }; goto s4
 :: d_step{ tag=Adam; dir=Down  }; goto s5
 :: d_step{ tag=Adam; dir=Push;
    stack=(stack<<4)|1 }; goto s1
 :: d_step{ (stack&15)==0 -> tag=Adam;
    dir=Pop; stack=(stack>>4) }; goto s0
 :: ...
 :: d_step{ (stack&15)==6 -> tag=Adam;
    dir=Pop; stack=(stack>>4) }; goto s6
 :: d_step{ tag=Adam; dir=Stop  }; goto end
 fi ;
```

B. Promela model with Embedded C code

Promela provides facilities to embed C code in the model [21]. We use this facility to separate parsing an XML file, and building an in-memory data structure in C on the one hand from the query processing with SPIN on the other hand. We use the eXpat library to parse the XML data [9]. The in-memory data structure follows the navigational format as shown in Fig. 1 and in the Kripke structure. For each node in the tree we `malloc()` a node with the appropriate number of children using the following C type definition:

```
typedef struct node* Nodeptr ;
typedef struct node {
    int tag ;
    int sz ; /* Number of children */
    Nodeptr parent, left, right ;
    Nodeptr child[sz] ;
} Node ;
```

The state of the Embedded C Promela model consists of five variables, where `tag`, `dir`, and `stack` are as in the pure Promela model. The added variable `ptr` points at the `Node` to which we are moving, and `nr` is used to index the appropriate `child`. The Promela model with Embedded C Code has fewer control states (106 versus 493 of the Pure Promela model) but it has more data states.

```
short tag ;
byte dir ;
int stack ;
c_state "Nodeptr ptr" "Global"
short nr ;
```

The `init` process below consists of the initialization where the C code which parses the XML tree is called. This is followed by a `do` statement with a non-deterministic choice for each of the nine directions, except for the `DOWN` direction, which has more cases to support nodes with many children efficiently. We show the cases for the direction `Up` and one of the cases for `Down`, the remaining cases are similar.

```
init {
    ... Initialisation calling XML parser ...
    do
    :: d_step{
        c_expr{ now.ptr->parent != NULL } ->
        c_code{
            now.tag = now.ptr->tag ;
            now.dir = Up ;
            now.ptr = now.ptr->parent ;
        }
    }
    :: d_step{ c_expr{ now.ptr->sz > 0 } ->
        c_code {
            now.tag = now.ptr->sym ;
            now.dir = Down ;
            now.ptr = now.ptr->child[0] ;
        }
    }
    ... Other cases ...
}
```

With the Kripke structure in place all that remains is to add the never claim generated by SPIN for the LTL formula that represents the query. The never claim specifies undesirable behaviour and SPIN will try to find a counter example. Therefore every counter example represents a match of the query, showing the details of the trace as required.

C. Performance

We discuss the performance of the pure Promela model first, and then compare the performance of the Promela model to that of state of the art XPath query engines. All our performance figures apply to a Sun SPARC Ultra-Enterprise Server running SunOS 5.8.

i) Pure Promela: The number of control states defined by our running example ϕ from Sec. V is 493. The number of data states defined by `tag`, `dir` and `stack` is at least $7 \times 9 \times 7^2 = 3087$. Multiplied by the number of control states, this yields over 1.5 M states. However a small percentage of these states is explored, as is shown in the second row of Table I. The columns correspond to the seven example queries discussed earlier in the paper. The first row shows the query number, the second shows the number of states stored to find at least one trace. (SPIN does not guarantee to find all traces.

SPIN version	$\mathsf{query}_{1\ldots7}$						
	1	2	3	4	5	6	7
pure	114	189	186	36	18	102	52
embedded C	36	91	46	7	5	58	41

TABLE I

SPIN PERFORMANCE FOR THE FAMILY TREE EXAMPLE IN TERMS OF THE NUMBER OF STATES STORED.

For query_2 SPIN returns all (2) traces, but in the case of query_1 only one of the three traces is found.)

The presence of data for query_4 in Table I is due to the fact that we translate the negated predicate into a negated LTL formula thus:

$$\mathcal{T}_x[\![\mathsf{xp}_1[\neg\ \mathsf{xp}_2]]\!]\phi = \mathcal{T}_x[\![\mathsf{xp}_1]\!](\neg(\mathsf{Push}\ \wedge\ \mathsf{X}(\mathcal{T}_x[\![\mathsf{xp}_2]\!](\mathsf{Pop}\ \wedge\ \mathsf{X}\ \phi))))$$

This means that SPIN will try to discover infinitely many counter examples, which for the purpose of this experiment has been capped at 100.

j) Promela with Embedded C Code: The Promela model with embedded C code performs better than the pure Promela model, because only the work relevant for the query processing is exposed to the model checker, the rest is hidden in the C code. The number of states explored is shown in the second row of Table I. The runtime of the query processing is not interesting since the XML tree corresponding to ϕ (Sec. V) is tiny.

k) Comparison with XML Task Force and MacMill: The third experiment repeats and extends the experiments of Afanasiev et al [2], using the standard XMark XML benchmark as the data base [31], with MacMill [7] and the XML taskforce query engine [17]. The six queries of Afanasiev et al are as follows:

$$\begin{aligned}
\mathsf{xmark}_1 &= /\ \mathsf{child}\ ::\ \mathsf{site}\ /\ \mathsf{child}\ ::\ \mathsf{regions}\ /\\
&\quad\ \mathsf{child}\ ::\ \mathsf{africa}\ /\ \mathsf{child}\ ::\ \mathsf{item}\ /\\
&\quad\ \mathsf{child}\ ::\ \mathsf{description}\ /\ \mathsf{child}\ ::\ \mathsf{parlist}\ /\\
&\quad\ \mathsf{child}\ ::\ \mathsf{listitem}\ /\ \mathsf{child}\ ::\ \mathsf{text}\\
\mathsf{xmark}_2 &= /\ \mathsf{descendant}\ ::\ \mathsf{item}\ /\ \mathsf{child}\ ::\ \mathsf{description}\ /\\
&\quad\ \mathsf{child}\ ::\ \mathsf{parlist}\ /\ \mathsf{child}\ ::\ \mathsf{listitem}\ /\\
&\quad\ \mathsf{child}\ ::\ \mathsf{text}\\
\mathsf{xmark}_3 &= /\ \mathsf{descendant}\ ::\ \mathsf{item}\ /\ \mathsf{descendant}\ ::\ \mathsf{text}\\
\mathsf{xmark}_4 &= \mathsf{descendant}\ ::\ \mathsf{open_auction}[\mathsf{child}\ ::\ \mathsf{bidder}]\\
\mathsf{xmark}_5 &= \mathsf{descendant}\ ::\ \mathsf{item}[\mathsf{child}\ ::\ \mathsf{payment}]\\
&\quad\ [\mathsf{child}\ ::\ \mathsf{location}]\\
\mathsf{xmark}_6 &= \mathsf{descendant}\ ::\ \mathsf{item}[\mathsf{descendant}\ ::\ \mathsf{payment}]
\end{aligned}$$

The extension consists of the three queries below which focus on antagonist axes. These examples originate from Grust et al [18].

$$\begin{aligned}
\mathsf{xmark}_7 &= \mathsf{descendant}\ ::\ \mathsf{open_auction}\ /\\
&\quad\ \mathsf{descendant}\ ::\ \mathsf{description}\\
\mathsf{xmark}_8 &= \mathsf{descendant}\ ::\ \mathsf{age}\ /\ \mathsf{ancestor}\ ::\ \mathsf{person}\\
\mathsf{xmark}_9 &= \mathsf{descendant}\ ::\ \mathsf{open_auction}\ /\ \mathsf{child}\ ::\ \mathsf{privacy}\ /\\
&\quad\ \mathsf{preceding_sibling}\ ::\ \mathsf{bidder}
\end{aligned}$$

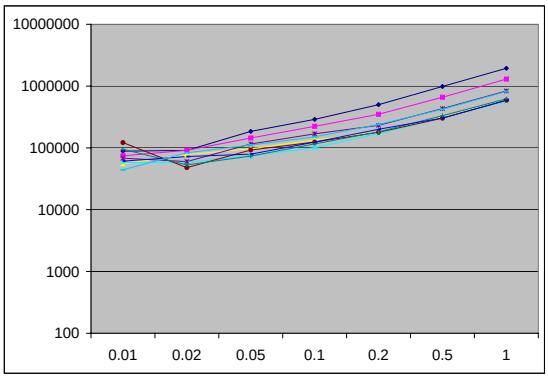

Fig. 3. MacMill query times (μSec) as a function of the XMark f parameter.

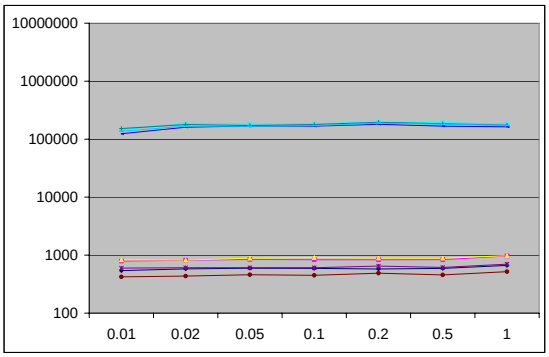

Fig. 4. SPIN query times (μSec) as a function of the XMark f parameter.

Using the scaling parameter settings $\mathsf{f} = 0.01, 0.02, 0.05, 0.1, 0.2, 0.5, 1.0$, the XML files generated by XMark range in size from 1.11MB to 111MB.

For each of the 9 XPath queries and 7 XML files, we report (a) the total time taken to read and parse the XML file and to execute the query, and (b) the time to process just the query itself. Each measurement is an average of ten experiments, with a standard deviation of 8% or less.

The query processing speed (defined as the size of the XML file divided by the time necessary to parse the XML input and to execute the query) is independent of the 9 XMark queries and the 7 data base sizes. The processing speed of MacMill is best with an average and standard deviation of $4885 \pm 2\%$ KB/s, for the TaskForce engine we found $2302 \pm 6\%$ KB/s, and for SPIN $2173 \pm 3\%$ KB/s. Overall we conclude that the SPIN implementation is competitive, which, going by the processing speed is MacMill : TaskForce : SPIN = 2.3 : 1.1 : 1.0

Fig. 3 reports the pure query processing times for MacMill, and Fig. 4 for the SPIN implementation. Each graph shows

the query time as a function of the XMark f parameter. With MacMill, the time increases with the f parameter because MacMill returns all answers, whereas SPIN stops after reporting the first answer. For queries 4, 7, 8 and 9 SPIN performs worse than for the other queries because these queries look for tags such as open_auction and age, which occur at the end of the XML tree (document order), whereas the fast queries look for tags such as item, which occur at the beginning of the tree. Overall SPIN is faster.

The SPIN implementation could be improved by fine tuning. On the other hand the range of XPath queries supported by the SPIN implementation is more limited than that of the MacMill and the XML task force engines. For example, the MacMill and the TaskForce engines report the entire answer set, including decorations, whereas the SPIN implementation is not guaranteed to report the entire answer set, and then only undecorated.

The fact that SPIN does not report all answers is not a problem with the main results of the paper, i.e. the trace semantics and its embedding in LTL. It should also be pointed out that because we supply the translated query *and* the access control policy together, the answer that the SPIN implementation does produce is guaranteed to satisfy the access control policy.

Being unable to provide all answers is a problem of the implementation with SPIN. To solve this problem with the current implementation of SPIN might require changes to the model, the translated query, and possibly the SPIN engine itself. We leave this a future work. On the other hand how often have you followed up all the hits that Google offers for a particular query? Google even provides an "I'm feeling lucky" mode, which gives just one hit. Therefore, there may be cases where not providing the entire answer set is acceptable.

l) Exponential time query complexity: Gottlob et al [17] report how naïve query processors suffer from exponential runtimes for relatively simple queries. The XML data base of their example is:

```
<A> <B/> <B/> </A>
```

The queries are $//a/b/\overbrace{\mathsf{parent} :: a/b}^{n \text{ times}}$. We have repeated the experiment to ensure that the SPIN implementation shows indeed linear behaviour, which is the case. However, to our surprise the SPIN compiler from LTL formula to never claims shows exponential runtimes, and so does the alternative compiler `ltl2ba` [15]. It is future work to investigate how to generate never claims directly from XPath queries. This would avoid exponential compilation times because we could exploit the regular structure of our LTL formulae.

IX. Conclusions and Future work

We define a novel trace semantics for positive core XPath that supports location paths, predicates, and 11 out of 13 axes. Expressions and negation are not currently supported. We show that positive core XPath can be translated into LTL. The translation is based on the idea that a branching structure as induced by location paths with predicates can be linearised effectively with the use of a stack. The translation is proved correct with respect to the trace semantics. The trace semantics provides opportunities for analysis. We give two examples showing that enforcing access control policies amounts to model checking the conjunction of the policy and the (LTL translation of) the query. Finally the SPIN model checker has been used as an efficient query engine, by providing it with a representation of an XML file and a never claim corresponding to the query translated into LTL. The performance of the SPIN implementation is comparable to that of the W3C XPath Taskforce Query engine.

Our SPIN implementation represents a successful experiment in creative laziness in the sense that we use existing tools (SPIN and the eXpat parser) for a new purpose (query processing) [30]. The necessary glue consists of a small Promela model and some C code (400 lines) that enable the model checker to traverse the XML tree, and a small compiler from XPath expressions into LTL (17 lines of Haskell). By comparison MacMill is 7Kloc. Our SPIN implementation has some undesirable features. In particular it stops after reporting one trace, and each new query must be compiled. This makes the current implementation unsuitable for practical use. A way forward would be to build, a query engine based on based on state of the art model checking technology. Instead of developing a tool from scratch one would use building blocks from a modular model checker and build an efficient special purpose tool with relative ease. This opens up a spectrum of possibilities ranging from a complete implementation from scratch (MacMill), via a partial implementation using existing model checker modules (future work) to our implementation with minimal glue.

Future work includes:

- Study the interaction of a wider class of security policies with our embedding of XPath queries into LTL, and extend the approach to embrace not only the object part of access control policies. It would also be of interest to investigate how XPath symmetries can continue to be exploited for query optimisation without undesirable interactions with the policy.
- Incorporate negation into the framework, as well as expressions and the two axes that we have omitted (namespace and attribute). Particularly the support of link edges for `id` and `idref` should not pose technical problems and should give rise to more interesting access control applications.
- Investigate the use of LTL formulae to cut down the search space by (1) compiling the query into more coarse grained filters, which can then be combined efficiently with the actual query, or (2) short circuiting the LTL compilation of the query expression using projection techniques, (3) using path equivalences to simplify the XPath expressions, or (4) adding further edges to shortcut recursive searches.

ACKNOWLEDGEMENTS

Loredana Afanasiev and Massimo Franceschet provided their CTL benchmark and commented on the approach. Sandro Etalle suggested using multi-lateral security as a motivating example. Gerard Holzmann and Theo Ruys answered many SPIN questions. Theo Ruys and Maurice van Keulen commented on a draft of the paper. Christoph Koch provided MacMill.

REFERENCES

[1] M. Abadi. Logic in access control. In *18th Annual IEEE Symp. on Logic in Computer ScienceC (LICS)*, pages 228–233, Ottawa, Canada, Jun 2003. IEEE Computer Society Press, Los Alamitos, California.

[2] L. Afanasiev, M. Franceschet, M. Marx, and M. de Rijke. CTL model checking for processing simple XPath queries. In *11th Int. Symp. on Temporal Representation and Reasoning (TIME)*, pages 117–124, Tatihou, France, Jul 2004. IEEE Computer Society Press, Los Alamitos, California.

[3] M. Benedikt, W. Fan, and G. M. Kuper. Structural properties of XPath fragments. In D. Calvanese, M. Lenzerini, and R. Motwani, editors, *9th International Conference on Database Theory (ICDT)*, volume LNCS 2572, pages 79–95, Siena, Italy, Jan 2003. Springer-Verlag, Berlin.

[4] E. Bertino and S. Castano. Securing XML documents with Author-X. *IEEE Internet Computing*, 5(3):21–31, May 2001.

[5] E. Bertino, S. Castano, E. Ferrari, and M. Mesiti. Specifying and enforcing access control policies for XML document sources. *World Wide Web*, 3(3):139–151, 2000.

[6] D. F. C. Brewer and M. J. Nash. The chinese wall security policy. In *10th IEEE Symposium on Security and Privacy (S&P)*, pages 1–3, Oakland, California, May 1989. IEEE Computer Society, Washington, DC.

[7] P. Buneman, M. Grohe, and Ch. Koch. Path queries on compressed XML. In J. C. Freytag, P. C. Lockemann, S. Abiteboul, M. J. Carey, P. G. Selinger, and A. Heuer, editors, *29th Int. Conf. on Very Large Data Bases (VLDB)*, pages 141–152, Berlin, Germany, Sep 2003. Morgan Kaufmann.

[8] C.-Y. Chan and M. Garofalakis. Secure XML querying with security views. In G. Weikum, A. Christian König, and S. Deßloch, editors, *SIGMOD Int. Conf. on Management of Data*, pages 587–598, Paris, France, Jun 2004. ACM Press, New York.

[9] J. Clark. *Expat XML Parser*. Open Software Technology Group, Fremont, California, Jul 2004.

[10] J. Clark and S. DeRose (eds.). *XML Path Language (XPath Version 1.0)*. W3C, Nov 1999.

[11] E. M. Clarke, O. Grumberg, and D. A. Peled. *Model Checking*. The MIT Press, Cmabridge, Massachusetts, 1999.

[12] L. de Alfaro. Model checking the world wide web. In G. Berry, H. Comon, and A. Finkel, editors, *13th Int. Conf. on Computer Aided Verification (CAV)*, volume LNCS 2102, pages 337–349, Paris, France, Jul 2001. Springer-Verlag, Berlin.

[13] X. Fu, T. Bultan, and J. Su. Analysis of interacting BPEL web services. In *13th conf. on World Wide Web*, pages 621–630, New York, NY, USA, 2004. ACM Press, New York.

[14] I. Fundulaki and M. Marx. Specifying access control policies for XML documents. In *9th ACM Symp. on access control models and technologies*, pages 61–69, IBM, Yorktown Heights, USA, Jun 2004. ACM Press, New York.

[15] P. Gastin and D. Oddoux. Fast LTL to Büchi automata translation. In G. Berry, H. Comon, and A. Finkel, editors, *13th Int. Conf. on Computer Aided Verification (CAV)*, volume LNCS 2102, pages 53–65, Paris, France, Jul 2001. Springer-verlag, Berlin.

[16] G. Gottlob, C. Koch, and R. Pichler. XPath processing in a nutshell. *SIGMOD Rec.*, 32(2):21–27, Jun 2003.

[17] G. Gottlob, Ch. Koch, and R. Pichler. Efficient algorithms for processing XPath queries. In *28th Int. Conf. on Very Large Data Bases (VLDB)*, pages 95–106, Hong Kong, China, Aug 2002. VLDB Endowment Inc.

[18] T. Grust, M. Van Keulen, and J. Teubner. Accelerating XPath evaluation in any RDBMS. *ACM Trans. Database Syst.*, 29(1):91–131, Mar 2004.

[19] J. Y. Halpern and V. Weissman. Using First-Order logic to reason about policies. In *16th IEEE Computer Security Foundations Workshop (CSFW)*, pages 187–201, Pacific Grove, California, Jun 2003. IEEE Computer Society Press, Los Alamitos, California.

[20] G. J. Holzmann. *The SPIN Model Checker: Primer and Reference manual*. Pearson Education Inc, Boston Massachusetts, 2004.

[21] G. J. Holzmann and R. Joshi. Model-Driven software verification. In S. Graf and L. Mounier, editors, *11th Int. SPIN Workshop: Model Checking Software*, volume LNCS 2989, pages 76–91, Barcelona, Spain, Apr 2004. Springer-Verlag Heidelberg.

[22] M. Huth and M. Ryan. *Logic in Computer Science*. Cambridge University Press, UK, 2004.

[23] B. Luo, D. Lee, W.-C. Lee, and P. Liu. QFilter: Fine-Grained Run-Time XML access control via NFA-based query rewriting. In *13th Conf. on Information and Knowledge Management (CIKM)*, pages 543–552, Washington D. C., Nov 2004. ACM Press, New York.

[24] M. Marx. Conditional XPath, the first order complete XPath dialect. In *23rd Principles of Database Systems (PODS)*, pages 13–22, Paris, France, Jun 2004. ACM Press, New York.

[25] M. Marx. XPath with conditional axis relations. In E. Bertino, S. Christodoulakis, D. Plexousakis, V. Christophides, M. Koubarakis, K. Böhm, and E. Ferrari, editors, *9th Int. Conf. on Extending Database Technology (EDBT)*, volume LNCS 2992, pages 477–494, Heraklion, Crete, Greece, Mar 2004. Springer-Verlag, Berlin.

[26] M. Marx and M. de Rijke. Semantic characterizations of navigational XPath. Technical report, Univ. of Amsterdam, 2004.

[27] G. Miklau and D. Suciu. Controlling access to published data using cryptography. In J. C. Freytag, P. C. Lockemann, S. Abiteboul, M. J. Carey, P. G. Selinger, and A. Heuer, editors, *29th Int. Conf. on Very Large Data Bases (VLDB)*, pages 898–909, Berlin, Germany, Sep 2003. Morgan Kaufmann.

[28] M. Murata, A. Tozawa, M. Kudo, and S. Hada. XML access control using static analysis. In *10th ACM conference on Computer and communication security*, pages 73–84, Washington D. C., 2003. ACM Press, new York.

[29] D. Olteanu, H. Meuss, T. Furche, and F. Bry. XPath: Looking forward. In A. B. Chaudhri, R. Unland, C. Djeraba, and W. Lindner, editors, *XML-Based Data Management and Multimedia Engineering (EDBT)*, volume LNCS 2490, pages 109–127, Prague, Czech Republic, Mar 2002. Springer-Verlag, Heidelberg.

[30] T. C. Ruys. Optimal scheduling using branch and bound with SPIN 4.0. In T. Ball and S. K. Rajamani, editors, *10th Int. SPIN Workshop on Model Checking Software*, volume LNCS 2648, pages 1–17, Portland, Oregon, May 2003. Springer-Verlag, Berlin.

[31] A. R. Schmidt, F. Waas, M. L. Kersten, M. J. Carey, I. Manolescu, and R. Busse. XMark: A benchmark for XML data management. In *28th Int. Conf. on Very Large Data Bases (VLDB)*, pages 974–985, Hong Kong, Aug 2002. VLDB Endowment Inc.

[32] Ph. Schnoebelen. The complexity of temporal logic model checking. In Ph. Balbiani, N.-Y. Suzuki, F. Wolter, and M. Zakharyaschev, editors, *Selected Papers from the 4th Workshop on Advances in Modal Logics (AiML'02)*, pages 393–436, Toulouse, France, Sep 2002. King's College Publication, London.

[33] P. L. Wadler. How to replace failure by a list of successes, a method for exception handling, backtracking and pattern matching in lazy functional languages. In J.-P. Jouannaud, editor, *2nd Functional programming languages and computer architecture (FPCA)*, volume LNCS 201, pages 113–128, Nancy, France, Sep 1985. Springer-Verlag, Berlin.

[34] P. L. Wadler. Two semantics for XPath. Technical note, Dept. of Comp. Sci, Univ. of Edinburgh, Jan 2000.

The appendices are included for the convenience of the reviewers, they will not be part of the final paper.

Appendix A – Linear Temporal Logic

We summarise the syntax and semantics of LTL here to make the paper self contained.

m) Syntax.: We use the fragment of temporal logic below, with proposition symbols $\mathbb{L}$ drawn from the sets $\mathbb{N}$ and $\mathbb{D}$.

$$\begin{aligned}\mathbb{T} \equiv\ & \mathsf{T} \mid \mathsf{F} \mid \mathbb{L} \mid \neg\, \mathbb{T} \mid \mathbb{T} \wedge \mathbb{T} \mid \mathbb{T} \vee \mathbb{T} \mid \mathbb{T} \rightarrow \mathbb{T} \mid \\ & \mathsf{X}\, \mathbb{T} \mid \Box\, \mathbb{T} \mid \Diamond\, \mathbb{T} \mid \mathbb{T}\ \mathsf{U}\ \mathbb{T}\end{aligned}$$

n) Semantics.: The semantics for finite traces is:

$$\begin{aligned}
&\models && :: [\alpha] \rightarrow \mathbb{T} \rightarrow \mathbb{B} \\
&[] \models _ && = \mathsf{False} \\
&\mathsf{xs} \models \mathsf{T} && = \mathsf{True} \\
&\mathsf{xs} \models \mathsf{F} && = \mathsf{False} \\
&(\mathsf{x} : \mathsf{xs}) \models \mathsf{l} && = \mathsf{l} \in \lambda\, \mathsf{x} \\
&\mathsf{xs} \models \neg\, \phi && = \neg(\mathsf{xs} \models \phi) \\
&\mathsf{xs} \models \phi \wedge \psi && = \mathsf{xs} \models \phi \wedge \mathsf{xs} \models \psi \\
&\mathsf{xs} \models \phi \vee \psi && = \mathsf{xs} \models \phi \vee \mathsf{xs} \models \psi \\
&\mathsf{xs} \models \phi \rightarrow \psi && = \mathsf{xs} \models (\neg\, \phi) \vee \psi \\
&(\mathsf{x} : \mathsf{xs}) \models \mathsf{X}\, \phi && = \mathsf{xs} \models \phi \\
&\mathsf{xs} \models \Box\, \phi && = \mathsf{xs} \models \neg\, \Diamond (\neg\, \phi) \\
&\mathsf{xs} \models \Diamond\, \phi && = \mathsf{xs} \models \mathsf{T}\ \mathsf{U}\ \phi \\
&\mathsf{xs} \models \phi\ \mathsf{U}\ \psi && = \mathsf{True}, \quad \textbf{if } \mathsf{xs} \models \psi \\
& && = \mathsf{xs} \models \mathsf{X}(\phi\ \mathsf{U}\ \psi), \quad \textbf{if } \mathsf{xs} \models \phi \\
& && = \mathsf{False}, \quad \textbf{otherwise}
\end{aligned}$$

Here the λ function is defined in the Kripke structure (See Sec. V).

Appendix B – Correctness

To prove the correctness of the translation $\mathcal{T}_\mathsf{x}$ with respect to the semantics $\mathcal{P}_\mathsf{x}$ we must have that for every query xp, trace ys and context node x holds:

$$\mathsf{ys} \in \mathcal{P}_\mathsf{x}[\![\mathsf{xp}]\!][(\mathsf{x},\ \mathsf{Start})] \text{ implies } \mathsf{ys} \models \mathcal{T}_\mathsf{x}[\![\mathsf{xp}]\!]\ \mathsf{Stop}$$

The proof follows directly from lemma I, the definition of $\mathcal{P}_\mathsf{x}$ and Lemma IV.

o) Lemma I: For every node $\mathsf{x} \in \mathbb{N}$, query $\mathsf{xp} \in \mathbb{X}$, and trace $\mathsf{xs} \in \mathbb{P}$ such that:

$$\langle (\mathsf{x},\ \bot),\ \mathsf{xp} \rangle \rightarrow \mathsf{xs}$$

there is a direction $\mathsf{d} \in \mathbb{D}$, a (possibly empty) trace $\mathsf{xs}' \in \mathbb{P}$, and a node $\mathsf{x}' \in \mathbb{N}$ such that:

$$\mathsf{xs} = (\mathsf{x},\ \mathsf{d}) :\ \mathsf{xs}'\ : (\mathsf{x}',\ \bot)$$

Proof: by induction on the shape of the derivation tree for $\ldots \rightarrow \ldots$. □

p) Lemma II: Given a function $\mathsf{f_p} \in \{\mathsf{here_p}, \mathsf{up_p}, \mathsf{down_p}, \mathsf{left_p}, \mathsf{right_p}\}$, then for every function $\mathsf{g_p} \in \{\mathsf{f_p}, \mathsf{f_p}\ +_\mathsf{p}, \mathsf{horizontal_p}\ \mathsf{f_p}\}$, every node $\mathsf{x} \in \mathbb{N}$, and trace $\mathsf{xs} \in \mathbb{P}$ such that:

$$\mathsf{xs} \in \mathsf{g_p}(\mathsf{x},\ \bot)$$

there is a direction $\mathsf{d} \in \mathbb{D}$, a (possibly empty) trace $\mathsf{xs}' \in \mathbb{P}$, and a node $\mathsf{x}' \in \mathbb{N}$ such that:

$$\mathsf{xs} = (\mathsf{x},\ \mathsf{d}) :\ \mathsf{xs}'\ : (\mathsf{x}',\ \bot)$$

Proof: by induction on the length of the trace xs. □

q) Lemma III: For every trace $\mathsf{ys}, \mathsf{zs} \in \mathbb{P}$, node $\mathsf{y}, \mathsf{y}' \in \mathbb{N}$, direction $\mathsf{d}, \mathsf{d}' \in \mathbb{D}$, axis $\mathsf{a} \in \mathbb{A}$, and LTL formula ϕ we have that if:

$$(\mathsf{y},\ \mathsf{d}) :\ \mathsf{ys}\ : (\mathsf{y}',\ \bot) \in \mathcal{P}_\mathsf{a}[\![\mathsf{a}]\!](\mathsf{y},\ \bot)$$

and:

$$(\mathsf{y}',\ \mathsf{d}') :\ \mathsf{zs} \models \phi$$

then:

$$(\mathsf{y},\ \mathsf{d}) :\ \mathsf{ys}\ : (\mathsf{y}',\ \mathsf{d}') :\ \mathsf{zs} \models \mathcal{T}_\mathsf{a}[\![\mathsf{a}]\!]\phi$$

Proof: by case analysis on the structure of a. □

r) Lemma IV: For every node $\mathsf{y}, \mathsf{y}' \in \mathbb{N}$, query $\mathsf{xp} \in \mathbb{X}$, direction $\mathsf{d}, \mathsf{d}' \in \mathbb{D}$, trace $\mathsf{ys}, \mathsf{zs} \in \mathbb{P}$, and LTL formula ϕ we have that if:

$$\langle (\mathsf{y},\ \bot),\ \mathsf{xp} \rangle \rightarrow (\mathsf{y},\ \mathsf{d}) :\ \mathsf{ys}\ : (\mathsf{y}',\ \bot)$$

and:

$$(\mathsf{y}',\ \mathsf{d}') :\ \mathsf{zs} \models \phi$$

then:

$$(\mathsf{y},\ \mathsf{d}) :\ \mathsf{ys}\ : (\mathsf{y}',\ \mathsf{d}') :\ \mathsf{zs} \models \mathcal{T}_\mathsf{x}[\![\mathsf{xp}]\!]\phi$$

Proof: by induction on the shape of the derivation tree for $\ldots \rightarrow \ldots$.

s) Case bar^1*::* Assume that the conclusion holds because the premise holds. Then by Lemma I we have:

$$\langle (\mathsf{x},\ \bot),\ \mathsf{xp_1} \rangle \rightarrow (\mathsf{x},\ \mathsf{d}) :\ \mathsf{xs}'\ : (\mathsf{x}',\ \bot)$$

$\Rightarrow$ (Assume that $(\mathsf{x}',\ \mathsf{d}') :\ \mathsf{zs} \models \phi$ and apply the induction hypothesis)

$$(\mathsf{x},\ \mathsf{d}) :\ \mathsf{xs}'\ : (\mathsf{x}',\ \mathsf{d}') :\ \mathsf{zs} \models \mathcal{T}_\mathsf{x}[\![\mathsf{xp_1}]\!]\phi$$

$\Leftrightarrow$ (Equation for $\vee$ of $\models$)

$$(\mathsf{x},\ \mathsf{d}) :\ \mathsf{xs}'\ : (\mathsf{x}',\ \mathsf{d}') :\ \mathsf{zs} \models \mathcal{T}_\mathsf{x}[\![\mathsf{xp_1}]\!]\phi \vee \mathcal{T}_\mathsf{x}[\![\mathsf{xp_2}]\!]\phi$$

$\Leftrightarrow$ (Equation for $\|$ of $\mathcal{T}_\mathsf{x}$)

$$(\mathsf{x},\ \mathsf{d}) :\ \mathsf{xs}'\ : (\mathsf{x}',\ \mathsf{d}') :\ \mathsf{zs} \models \mathcal{T}_\mathsf{x}[\![\mathsf{xp_1} \parallel \mathsf{xp_2}]\!]\phi$$

□

t) Case bar^2*::* The proof of this case is analogous to that of case bar^1.

u) Case abs*::* Assume that the conclusion holds because the premise holds. Then by Lemma I we have that:

$$\langle(\mathsf{Root},\ \bot),\ \mathsf{xp}\rangle \rightarrow (\mathsf{Root},\ \mathsf{d}) :\ \mathsf{xs}'\ : (\mathsf{x}',\ \bot)$$

$\Rightarrow$ (Assume that $(\mathsf{x}',\ \mathsf{d}') :\ \mathsf{zs}\ \models\ \phi$ and apply the induction hypothesis)

$$(\mathsf{Root},\ \mathsf{d}) :\ \mathsf{xs}'\ : (\mathsf{x}',\ \mathsf{d}') :\ \mathsf{zs}\ \models\ \mathcal{T}_{\mathsf{x}}[\![\mathsf{xp}]\!]\phi$$

$\Leftrightarrow$ (Equations for $\wedge$ and atom of $\models$)

$$(\mathsf{Root},\ \mathsf{d}) :\ \mathsf{xs}'\ : (\mathsf{x}',\ \mathsf{d}') :\ \mathsf{zs}\ \models\ \ \mathsf{Root}\ \wedge\ \mathcal{T}_{\mathsf{x}}[\![\mathsf{xp}]\!]\phi$$

$\Leftrightarrow$ (Equations for X and atom of $\models$)

$$(\mathsf{x},\ \mathsf{Start}) : (\mathsf{Root},\ \mathsf{d}) :\ \mathsf{xs}'\ : (\mathsf{x}',\ \mathsf{d}') :\ \mathsf{zs}$$
$$\models\ \ \mathsf{Start}\ \wedge\ \mathsf{X}(\ \mathsf{Root}\ \wedge\ \mathcal{T}_{\mathsf{x}}[\![\mathsf{xp}]\!]\phi)$$

$\Leftrightarrow$ (Equation for (unary) / of $\mathcal{T}_{\mathsf{x}}$)

$$(\mathsf{x},\ \mathsf{Start}) : (\mathsf{Root},\ \mathsf{d}) :\ \mathsf{xs}'\ : (\mathsf{x}',\ \mathsf{d}') :\ \mathsf{zs}\ \models\ \mathcal{T}_{\mathsf{x}}[\![/\ \mathsf{xp}]\!]\phi$$

□

v) Case slash*::* We omit the proof for this case, as it is a simplified version of the proof for case pred below.

w) Case pred*::* Firstly, assume that:

$$(\mathsf{x}',\ \mathsf{d}'') :\ \mathsf{zs}\ \models\ \phi$$

$\Leftrightarrow$ (Equations for X and atom of $\models$)

$$(\mathsf{x}'',\ \mathsf{Pop}) : (\mathsf{x}',\ \mathsf{d}'') :\ \mathsf{zs}\ \models\ \ \mathsf{Pop}\ \wedge\ \mathsf{X}\ \phi$$

Secondly, assume that the conclusion holds because the two premises hold. Then applying Lemma I to the second premise yields:

$$\langle(\mathsf{x}',\ \bot),\ \mathsf{xp}_2\rangle \rightarrow (\mathsf{x}',\ \mathsf{d}') :\ \mathsf{xs}''\ : (\mathsf{x}'',\ \bot)$$

$\Rightarrow$ (First assumption and induction hypothesis applied to second premise)

$$(\mathsf{x}',\ \mathsf{d}') :\ \mathsf{xs}''\ : (\mathsf{x}'',\ \mathsf{Pop}) : (\mathsf{x}',\ \mathsf{d}'') :\ \mathsf{zs}$$
$$\models\ \mathcal{T}_{\mathsf{x}}[\![\mathsf{xp}_2]\!](\ \mathsf{Pop}\ \wedge\ \mathsf{X}\ \phi)$$

$\Leftrightarrow$ (Equations for X and atom of $\models$)

$$(\mathsf{x}',\ \mathsf{Push}) : (\mathsf{x}',\ \mathsf{d}') :\ \mathsf{xs}''\ : (\mathsf{x}'',\ \mathsf{Pop}) : (\mathsf{x}',\ \mathsf{d}'') :\ \mathsf{zs}$$
$$\models\ \mathsf{Push}\ \wedge\ \mathsf{X}(\mathcal{T}_{\mathsf{x}}[\![\mathsf{xp}_2]\!](\ \mathsf{Pop}\ \wedge\ \mathsf{X}\ \phi))$$

$\Rightarrow$ (Induction hypothesis applied to the first premise)

$$(\mathsf{x},\ \mathsf{d}) :\ \mathsf{xs}'\ : (\mathsf{x}',\ \mathsf{Push}) : (\mathsf{x}',\ \mathsf{d}') :\ \mathsf{xs}''\ : (\mathsf{x}'',\ \mathsf{Pop}) : (\mathsf{x}',\ \mathsf{d}'') :\ \mathsf{zs}$$
$$\models\ \mathcal{T}_{\mathsf{x}}[\![\mathsf{xp}_1]\!](\mathsf{Push}\ \wedge\ \mathsf{X}(\mathcal{T}_{\mathsf{x}}[\![\mathsf{xp}_2]\!](\ \mathsf{Pop}\ \wedge\ \mathsf{X}\ \phi)))$$

$\Leftrightarrow$ (Equation for $\ldots[\ldots]$ of $\mathcal{T}_{\mathsf{x}}$)

$$(\mathsf{x},\ \mathsf{d}) :\ \mathsf{xs}'\ : (\mathsf{x}',\ \mathsf{Push}) : (\mathsf{x}',\ \mathsf{d}') :\ \mathsf{xs}''\ : (\mathsf{x}'',\ \mathsf{Pop}) : (\mathsf{x}',\ \mathsf{d}'') :\ \mathsf{zs}$$
$$\models\ \mathcal{T}_{\mathsf{x}}[\![\mathsf{xp}_1[\mathsf{xp}_2]]\!]\phi$$

□

x) Case step*::* The proof follows immediately from Lemmas II and III.

y) Lemma V: For every node $\mathsf{y} \in \mathbb{N}$, direction $\mathsf{d} \in \mathbb{D}$, trace $\mathsf{xs} \in \mathbb{P}$, query $\mathsf{xp}_1, \mathsf{xp}_2 \in \mathbb{X}$, and LTL formula ϕ we have that if:

$$(\mathsf{y},\ \mathsf{d}) :\ \mathsf{xs}\ \models\ \mathcal{T}_{\mathsf{x}}[\![\mathsf{xp}_1]\!](\mathcal{T}_{\mathsf{x}}[\![\mathsf{xp}_2]\!]\phi)$$

then there exists a node $\mathsf{y}' \in \mathbb{N}$, a direction $\mathsf{d}' \in \mathbb{D}$, and traces $\mathsf{ys}, \mathsf{zs} \in \mathbb{P}$ such that:

$$\mathsf{xs} = \mathsf{ys}\ : (\mathsf{y}',\ \mathsf{d}') :\ \mathsf{zs}$$

and:

$$(\mathsf{y}',\ \mathsf{d}') :\ \mathsf{zs}\ \models\ \mathcal{T}_{\mathsf{x}}[\![\mathsf{xp}_2]\!]\phi$$

Proof: by induction on the length of xs. □