
CONTAINMENT FOR CONDITIONAL TREE PATTERNS

ALESSANDRO FACCHINI^a, YOICHI HIRAI^b, MAARTEN MARX^c, AND EVGENY SHERKHONOV^d

^a IDSIA - Dalle Molle Institute for Artificial Intelligence
e-mail address: alessandro.facchini@idsia.ch

^b FireEye
e-mail address: yoichi.hirai@fireeye.com

^{c,d} ISLA, University of Amsterdam
e-mail address: {maartenmarx,e.sherkhonov}@uva.nl

ABSTRACT. A Conditional Tree Pattern (CTP) expands an XML tree pattern with labels attached to the descendant edges. These labels can be XML element names or Boolean CTP's. The meaning of a descendant edge labelled by A and ending in a node labelled by B is a path of child steps ending in a B node such that all intermediate nodes are A nodes. In effect this expresses the *until* B , A holds construction from temporal logic.

This paper studies the containment problem for CTP. For tree patterns (TP), this problem is known to be CONP-complete. We show that it is PSPACE-complete for CTP. This increase in complexity is due to the fact that CTP is expressive enough to encode an unrestricted form of label negation: $* \setminus a$, meaning "any node except an a -node". Containment of TP expanded with this type of negation is already PSPACE-hard.

CTP is a positive, forward, first order fragment of Regular XPath. Unlike TP, CTP expanded with disjunction is not equivalent to unions of CTP's. Like TP, CTP is a natural fragment to consider: CTP is closed under intersections and CTP with disjunction is equally expressive as positive existential first order logic expanded with the until operator.

1. INTRODUCTION

Tree Patterns, abbreviated as TP, are one of the most studied languages for XML documents and used in almost all aspects of XML data management. Tree patterns are a natural language: over trees, unions of tree patterns are equally expressive as positive first order logic [3]. Also, like relational conjunctive queries, the semantics of TP's can be given by embeddings from patterns to tree models [1]. Equivalence and containment of TP's is decidable in PTIME for several fragments [1], and CONP complete in general [15]. Tree patterns can be represented as trees, as in Figure 1, or be given a natural XPath-like syntax.

In this paper, we study the expansion of tree patterns with the conditional descendant axis. We call this expansion *Conditional Tree Patterns*, abbreviated as CTP. Where the descendant axis in TP can be written as the transitive reflexive closure of the XPath step

2012 ACM CCS: [Theory of computation]: Theory and algorithms for application domains—Database theory—Database query languages (principles) / Database query processing and optimization (theory).

Key words and phrases: XPath, XML, Conditional Tree Pattern, Containment, Complexity.

^{c,d}This research was supported by the Netherlands Organization for Scientific Research (NWO) under project number 612.001.012 (DEX)..

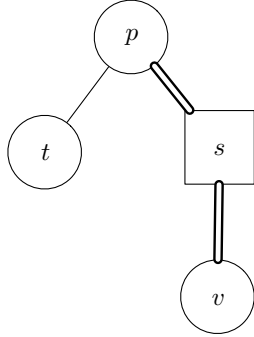


FIGURE 1. The tree pattern corresponding to the XPath expression $/p[t]//s[./v]$. The node in the square box denotes the output node.

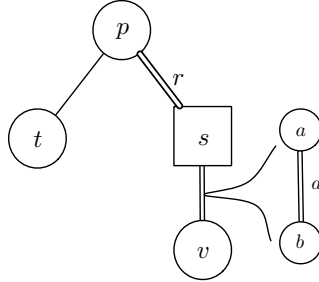


FIGURE 2. Conditional tree pattern corresponding to (1.1).

$\mathbf{child}::*$, the conditional descendant axis is the transitive closure of $\mathbf{child} :: p[F_1][F_2] \dots [F_n]$, where $n \geq 0$ and each F_i is an XPath expression which might contain the conditional descendant axis itself, followed by a child step. Syntactically, the expansion is straightforward: in the tree representation add labels representing conditional tree patterns themselves to the edges. Figure 2 contains an example which is equivalent to the following XPath-like formula in which we use $(\cdot)^*$ to denote the reflexive transitive closure of a path formula:

$$\begin{aligned} \mathbf{self} :: p[\mathbf{child} :: t]/(\mathbf{child} :: r)^*/\mathbf{child} :: s \\ [(\mathbf{child} :: a[(\mathbf{child} :: a)^*/\mathbf{child} :: b])^*/\mathbf{child} :: v] \end{aligned} \quad (1.1)$$

The edge in Figure 2 labeled by r corresponds to the expression $(\mathbf{child} :: r)^*/\mathbf{child} :: *$, and merely states that all intermediate nodes are labeled by r . The other labeled edge shows a nesting of patterns, corresponding to a nested transitive closure statement: all intermediate nodes have to be labeled by a and moreover have to start an a -labeled path ending in a b -node.

Conditional tree patterns are the forward fragment of Conditional XPath [13] without disjunction and negation. The conditional descendant axis is closely related to the strict until operator from temporal logic [13, 12].

Main results. Our main results concern the expressivity of CTP and the complexity of the containment problem. We consider two types of models: the standard XML trees in which each node has exactly one label, and trees in which nodes can have an arbitrary number of labels. These latter, called *multi-labeled trees*, are the models considered in temporal logic. All our results hold for both semantics. Models with multiple labels are a convenient logical abstraction for reasoning about tree patterns expanded with attribute value equalities. These are expressions of the form $@_a = c$, where a is an attribute, c a constant, meaning that it holds at a node if and only if the value of the a -attribute of the node equals c . (With that we can express XPath formulas like `//table[@border = '1']`).

Containment of tree patterns has been studied in [1, 15, 16]. The most relevant results for this paper are that containment of TP's is CONP-complete in general [15] and PSPACE-complete when the domain of labels is finite [16]. We show that containment of CTP's is PSPACE-complete (with both finite and infinite domain of labels). Interestingly, this increase in complexity is due to the fact that CTP is expressive enough to encode an unrestricted form of label negation: $* \setminus a$, meaning "any node except an a -node". We show that containment of TP's expanded with this form of negation is already PSPACE-hard. The matching upper bound for CTP containment is easily obtained by a translation into Existential CTL [11]. As a contrast we consider the expansion of TP with a *safe* form of propositional negation $n \setminus a$, which selects nodes with the label containing n and not a , instead of $* \setminus a$ [2]. Note that this construct only makes sense on models with multiple labels. With respect to expressivity, we show that most results for TP can be generalized to CTP. CTP's can be interpreted in trees by generalizing the TP-embeddings to simulations known from temporal logic. Similarly to the characterization for TP in [3], we show that CTP's with disjunction and union are equally expressive as positive first order logic expanded with an *until* operator. From this we obtain that like TP's, CTP's are closed under taking intersections.

Organization. The paper is organized as follows. This section is continued with a few more comparisons between TP and CTP, related work and a motivating example. Section 2 contains preliminaries. Section 3 contains the expressivity results and Section 4 the results on the complexity of the containment problem. We end with conclusions and open questions.

Comparing logical properties of TP and CTP. A characteristic difference between TP and Relational Conjunctive Queries (CQ) is the disjunction property: if $A \models B \vee C$, then $A \models B$ or $A \models C$. This holds for CQ, but not for TP. A counterexample is `//p` \models `/p` \vee `*/p`. The languages TP and CTP differ on the following:

Unions: TP expanded with disjunction is equally expressive as unions of TP [3]. However, CTP with disjunction is more expressive than unions of CTP.

Countermodels: If containment between two TP's fails, there is a countermodel for it of polynomial size [15]. Countermodels for CTP containment may be exponential.

Complexity: TP containment is CONP-complete [15] and CTP containment is PSPACE-complete. For both languages this remains true if we add disjunction to the language (for TP see [16]).

However, there are a few useful technical results that TP and CTP share.

Monotonicity: TP and CTP formulas are preserved under extensions of models at the leaves (i.e. when the original model is a *subtree* of the extension). This means that if a TP (CTP) formula holds in a tree, then it holds in the extensions of the tree.

Multiple output nodes: The containment problem for TP and CTP formulas with multiple output nodes can be reduced to containment of Boolean TP and CTP formulas (i.e. when the formula has a single output node which is the root).

Containment for unions: Containment for unions of TPs can be reduced in PTIME to checking a set of containments between TP formulas [15]. This reduction can be seen as a weak disjunction property. A similar property holds for CTP (see Proposition 4.2).

Multi-labeled models: There is a PTIME reduction from the containment problem over trees in which each node has exactly one label to the containment problem over multi-labeled trees. This holds for both TP and CTP.

For TP, most of these results are in [15]. Here we show how their proofs can be generalized to CTP.

Motivation. Tree patterns exhibit a nice tradeoff between expressive power and the complexity of static analysis. However, there are natural scenarios where tree patterns are not powerful enough, e.g. in Example 1.1. The conditional axis gives us more querying capabilities while preserving some of the nice properties of tree patterns (see the comparison above).

Example 1.1. Conditional tree patterns are used to query tree shaped structures. As an example, take the tree structure of the UNIX file system. In this file system, every file and directory has different access permissions (read, write or execute) for different type of users (the user, the group and others). Thus, the file system can be modeled as a tree where each node corresponds to a directory or a file (labeled by "dir" and "file" respectively) and has a required attribute for each pair (*user, access right*) which takes values from $\{0, 1\}$. A file can only be a leaf in the tree.

Assume we want to ask for all the files that are readable by the user. This means that we are looking for precisely those files for which the following permissions hold:

- the file is readable for the user,
- the directory in which the file resides is both readable and executable for the user,
- the same holds recursively for all the directories from the root to the file.

This query can be neatly expressed as a CTP path formula:

$$/(\text{child} :: \text{dir}[@_{(\text{user}, \text{read})} = 1][@_{(\text{user}, \text{execute})} = 1])^*/\text{child} :: \text{file}[@_{(\text{user}, \text{read})} = 1].$$

Additionally, CTP can express non-trivial constraints over the tree representing the file system. For instance, the formula

$$//\text{self} :: \text{file}[@_{(\text{other}, \text{read})} = 1] \rightarrow /(\text{child} :: \text{dir}[@_{(\text{other}, \text{read})} = 1][@_{(\text{other}, \text{execute})} = 1])^*/$$

imposes the constraint that for all files which are readable it holds that every directory on the path from the root to this file must be both readable and executable by others.

It is known that the conditional axis cannot be expressed by tree patterns or in Core XPath [13]. On the other hand, the queries from Example 1.1 can be expressed in the positive forward fragment of Regular XPath, which is more expressive than CTP. However, we believe this additional expressive power leads to increase in the complexity (of containment) than for conditional tree patterns.

1.1. Related Work. The complexity of tree patterns is studied in a number of papers and since [16] a virtually complete picture exists for the complexity of the containment problem for positive fragments of XPath. Miklau and Suciu [15] show that the complexity of the containment problem for TP with filters, wildcard and descendant is CONP -complete. Containment and equivalence for fragments of TP were studied before. The most interesting result is that containment for TP without the wildcard is in PTIME [1].

Our conditional tree patterns are a first order fragment of conjunctive regular path queries. Calvanese et al. [6] show that the complexity of containment of these queries is EXPSpace -complete, but these are interpreted on general graph models.

Conditional XPath [13] and conditional tree patterns are closely related to branching time temporal logic CTL [8]. The conditional child axis and the strict until connective are interdefinable. Kupferman and Vardi [11] show that the containment problem for $\exists\text{CTL}$, which is the restriction of CTL to formulas having only the \exists path quantifier in front of them, is a PSPACE -complete problem. The positive $\exists\text{CTL}$ -fragment without until was also studied by Miklau and Suciu [15]. They show that the containment problem for this fragment is equivalent to the TP-containment problem and thus also CONP -complete.

This paper studies the containment problem without presence of schema information. A number of complexity results for containment in TP w.r.t DTDs are given in [16]. In particular, containment for TP with filters, the wildcard and descendent w.r.t DTD is EXPTIME -complete. This hardness result together with an EXPTIME upper bound for Conditional XPath [13], which contains CTP, gives us EXPTIME -completeness of containment for CTP in the presence of DTDs.

2. PRELIMINARIES

We review the basic definitions of XML trees, Regular XPath and its semantics. Then we present Tree Patterns and Conditional Tree Patterns as sublanguages of Regular XPath. Tree patterns have an alternative semantics in terms of embeddings [15]. We give such an “embedding semantics” for Conditional Tree Patterns using “until-simulations” in Section 3.

2.1. Trees. We work with node-labeled unranked finite trees, where the node labels are elements of an infinite set Σ . Formally, a tree over Σ is a tuple (N, E, r, ρ) , where N , the set of nodes of the tree, is a prefix closed set of finite sequences of natural numbers, $E = \{(\langle n_1, \dots, n_k \rangle, \langle n_1, \dots, n_k, n_{k+1} \rangle) \mid \langle n_1, \dots, n_{k+1} \rangle \in N\}$ is the edge or child relation, r is the root of the tree, that is the empty sequence, and ρ is the function assigning to each node in N a finite subset of Σ . We refer a tree over Σ just as a tree if Σ is clear from the context.

Trees in which $\rho(\cdot)$ is always a singleton are called *single-labeled* or *XML trees*. Trees without this restriction are called *multi-labeled trees*.

We denote by E^+ the descendant relation, which is the transitive closure of the edge relation E , and by E^* the reflexive and transitive closure of E , and by $E(x)$ the set of all children of the node x . A node n is a *leaf* if $E(n)$ is empty. A *path* from a node n to a node m is a sequence of nodes $n = n_0, \dots, n_k = m$, with $k > 0$, such that for each $i \leq k$, $n_{i+1} \in E(n_i)$. We call a *branch* any maximal path starting from the root. If mE^+n , m is called an *ancestor* of n , and if mEn , m is called the *parent* of n .

If n is in N , by $T.n$ we denote the subtree of T rooted at n . A *pointed tree* is a pair T, n where T is a tree and n is a node of T . The *height* of a pointed tree T, n is the length of the longest path in $T.n$.

Given two trees $T_1 = (N_1, E_1, r_1, \rho_1)$ and $T_2 = (N_2, E_2, r_2, \rho_2)$ such that N_1 and N_2 are disjoint, we define the result of *fusion* of T_1 and T_2 , denoted as $T_1 \oplus T_2$, as the tree obtained by joining the trees T_1 and T_2 without the roots under a new common root labeled by the union of the labels of the roots of T_1 and T_2 . Formally, $T_1 \oplus T_2$ is the tree $T = (N, E, r, \rho)$, where $N = (N_1 \setminus \{r_1\}) \cup (N_2 \setminus \{r_2\}) \cup \{r\}$, $E = (E_1 \setminus \{\langle r_1, n \rangle \mid n \in N_1\}) \cup (E_2 \setminus \{\langle r_2, n \rangle \mid n \in N_2\}) \cup \{\langle r, n \rangle \mid \langle r_1, n \rangle \in E_1 \text{ or } \langle r_2, n \rangle \in E_2\}$ and

$$\rho(n) = \begin{cases} \rho_1(r_1) \cup \rho_2(r_2) & \text{if } n = r, \\ \rho_1(n) & \text{if } n \in N_1 \setminus \{r_1\}, \\ \rho_2(n) & \text{if } n \in N_2 \setminus \{r_2\}. \end{cases}$$

2.2. XPath and Tree Patterns. We define Tree Patterns and Conditional Tree Patterns as sublanguages of Regular XPath [17].

Definition 2.1 (Forward Regular XPath). Let Σ be an infinite domain of labels. Forward Regular XPath, RXPath for short, consists of node formulas φ and path formulas α which are defined by the following grammar

$$\begin{aligned} \varphi &::= p \mid \top \mid \neg\varphi \mid \varphi \vee \psi \mid \varphi \wedge \psi \mid \langle \alpha \rangle \varphi \\ \alpha &::= \downarrow \mid ?\varphi \mid \alpha; \alpha \mid \alpha \cup \alpha \mid \alpha^*, \end{aligned}$$

where $p \in \Sigma$.

We will use α^+ as an abbreviation of $\alpha; \alpha^*$.

For the semantics of RXPath, given a tree $T = (N, E, r, \rho)$ over Σ , the relation $\llbracket \alpha \rrbracket_T \subseteq N \times N$ for a path expression α and the satisfaction relation \models between pointed trees and node formulas are inductively defined as follows:

- $\llbracket \downarrow \rrbracket_T = E$,
- $\llbracket ?\varphi \rrbracket_T = \{(n, n) \in N \times N \mid T, n \models \varphi\}$,
- $\llbracket \alpha; \beta \rrbracket_T = \llbracket \alpha \rrbracket_T \circ \llbracket \beta \rrbracket_T$,
- $\llbracket \alpha \cup \beta \rrbracket_T = \llbracket \alpha \rrbracket_T \cup \llbracket \beta \rrbracket_T$,
- $\llbracket \alpha^* \rrbracket_T = (\llbracket \alpha \rrbracket_T)^*$,

and

- $T, n \models \top$,
- $T, n \models p$ iff $p \in \rho(n)$,
- $T, n \models \neg\varphi$ iff $T, n \not\models \varphi$,
- $T, n \models \varphi \wedge \psi$ iff $T, n \models \varphi$ and $T, n \models \psi$,
- $T, n \models \varphi \vee \psi$ iff $T, n \models \varphi$ or $T, n \models \psi$,
- $T, n \models \langle \alpha \rangle \varphi$ iff there is a node m with $(n, m) \in \llbracket \alpha \rrbracket_T$ and $T, m \models \varphi$.

Sometimes we will write $T \models \varphi$ to denote $T, r \models \varphi$. If the latter holds, we say that T is a *model* of φ .

(Conditional) Tree Patterns. Tree patterns are the conjunctive fragment of RXPath without unions of paths and with a strongly restricted Kleene star operation. Node formulas correspond to Boolean tree patterns, and path formulas to tree patterns with one output node.

In the following definitions we again let Σ be an infinite domain of labels. We define Tree Patterns by restricting the syntax of RXPath as follows:

Definition 2.2 (Tree Pattern). Tree Patterns (TP) consist of node formulas φ and path formulas α defined by the following grammar:

$$\begin{aligned}\varphi &::= p \mid \top \mid \varphi \wedge \varphi \mid \langle \alpha \rangle \varphi \\ \alpha &::= \downarrow \mid ?\varphi \mid \alpha; \alpha \mid \downarrow^+, \end{aligned}$$

where $p \in \Sigma$.

For example, the tree pattern from Figure 1 can be written as the path formula $\alpha = ?(p \wedge \langle \downarrow \rangle t); \downarrow^+; ?(s \wedge \langle \downarrow \rangle v)$. Conditional Tree Patterns are also defined by restricting the syntax of RXPath, where we allow conditional descendant paths.

Definition 2.3 (Conditional Tree Pattern). Conditional Tree Patterns (CTP) consist of node formulas φ and path formulas defined by the following grammar:

$$\begin{aligned}\varphi &::= p \mid \top \mid \varphi \wedge \varphi \mid \langle \alpha \rangle \varphi \\ \alpha &::= \downarrow \mid ?\varphi \mid \alpha; \alpha \mid (\downarrow; ?\varphi)^*; \downarrow, \end{aligned}$$

where $p \in \Sigma$.

The tree from Figure 2 can be written as the path formula

$$?(p \wedge \langle \downarrow \rangle t); (\downarrow; ?r)^*; \downarrow; ?(s \wedge \langle (\downarrow; ?(a \wedge \langle (\downarrow; ?a)^*; \downarrow \rangle b))^*; \downarrow \rangle v).$$

Note that the node formula $\langle (\downarrow; ?\varphi)^*; \downarrow \rangle \psi$ is exactly the strict until $\exists U(\psi, \varphi)$ from branching time temporal logic CTL [11]. We will abbreviate this formula simply as $\langle \varphi^+ \rangle \psi$. Because of the equivalences $\langle \alpha; \beta \rangle \varphi \equiv \langle \alpha \rangle \langle \beta \rangle \varphi$ and $\langle ?\varphi \rangle \psi \equiv \varphi \wedge \psi$, CTP node formulas can be given the following equivalent definition:

$$\varphi ::= p \mid \top \mid \varphi \wedge \varphi \mid \langle \downarrow \rangle \varphi \mid \langle \downarrow^+ \rangle \varphi \mid \langle \varphi^+ \rangle \varphi.$$

Even though the syntax is slightly different, CTP is the conjunctive forward only fragment of Conditional XPath [13].

Expansions. We consider two expansions of TP and CTP, with negated labels and with disjunction in node formulas together with union in paths. Negated labels is a restricted type of negation where only the construct $\neg p$ is allowed in the node formulas. We denote expansions of the language L with one or two of these features by L^S for $S \subseteq \{\neg, \vee\}$.

Query evaluation. In [9] it was shown that the query evaluation problem for Core XPath is PTIME-complete in the combined complexity. As noted in [13], using results on model checking for Propositional Dynamic Logic, this can be extended to Regular XPath, and thus to all our defined fragments.

Fact 2.4 ([13]). Let T be a tree, n_1, n_2 nodes in T , and α a Regular XPath path formula. The problem whether $(n_1, n_2) \in \llbracket \alpha \rrbracket_T$ is decidable in time $O(|T| \times |\alpha|)$ with $|T|$ the size of the tree and $|\alpha|$ the size of the formula.

2.3. Containment. As we are considering two kinds of expressions, path and node expressions, we have different notions of containment.

Definition 2.5. Let φ and ψ be two RXPath-node formulas. We say that φ is *contained* in ψ , notation $\varphi \subseteq \psi$, if for every T, n , $T, n \models \varphi$ implies $T, n \models \psi$. Let α and β be two RXPath-path formulas. We say that

- α is *contained as a binary query* in β , denoted $\alpha \subseteq^2 \beta$, if for any tree T , $\llbracket \alpha \rrbracket_T \subseteq \llbracket \beta \rrbracket_T$,
- α is *contained as a unary query* in β , denoted $\alpha \subseteq^1 \beta$, if for any tree T with root r , and any node n , $(r, n) \in \llbracket \alpha \rrbracket_T$ implies $(r, n) \in \llbracket \beta \rrbracket_T$.

Containment over single-labeled trees is denoted by \subseteq , and containment over multi-labeled models by \subseteq_{ML} .

Luckily, these three notions are closely related, and containment of path formulas can be reduced to containment of node formulas and vice versa (cf. also [16, 15]).

Proposition 2.6. *Let α and β be two RXPath-path formulas, φ and ψ RXPath-node formulas, in which negation is restricted to labels only.*

- (i) $\alpha \subseteq^2 \beta$ iff $\alpha \subseteq^1 \beta$,
- (ii) Let p be a label not occurring in α or in β . Then, $\alpha \subseteq^2 \beta$ iff $\langle \alpha \rangle \langle \downarrow \rangle p \subseteq \langle \beta \rangle \langle \downarrow \rangle p$,
- (iii) $\varphi \subseteq \psi$ iff $? \varphi \subseteq^2 ? \psi$.

All the above items also hold for the case of multi-labeled trees.

Proof. (i) (\Rightarrow) Let $T = (N, E, r, \rho)$ be a tree such that $(r, n) \in \llbracket \alpha \rrbracket_T$. By the assumption, we have $\llbracket \alpha \rrbracket_T \subseteq \llbracket \beta \rrbracket_T$. This implies that $(r, n) \in \llbracket \beta \rrbracket_T$, which was required to show.

(\Leftarrow) Let $T = (N, E, r, \rho)$ be a tree, n_1 and n_2 in N such that $(n_1, n_2) \in \llbracket \alpha \rrbracket_T$. Since α is from the forward fragment of Regular XPath, we have that $(n_1, n_2) \in \llbracket \alpha \rrbracket_{T.n_1}$. Using the assumption, we have $(n_1, n_2) \in \llbracket \beta \rrbracket_{T.n_1}$. Then by monotonicity, we obtain $(n_1, n_2) \in \llbracket \beta \rrbracket_T$, which was desired.

(ii) (\Rightarrow) Suppose $\alpha \subseteq^2 \beta$, and let $T, n \models \langle \alpha \rangle \langle \downarrow \rangle p$. Thus there is a descendant m of n in T such that $(n, m) \in \llbracket \alpha \rrbracket_T$ and $T, m \models \langle \downarrow \rangle p$. By the hypothesis we therefore have $(n, m) \in \llbracket \beta \rrbracket_T$ and thus $T, n \models \langle \beta \rangle \langle \downarrow \rangle p$.

(\Leftarrow) Assume $\langle \alpha \rangle \langle \downarrow \rangle p \subseteq \langle \beta \rangle \langle \downarrow \rangle p$. Consider a tree $T = (N, E, r, \rho)$ and a pair $(n, m) \in \llbracket \alpha \rrbracket_T$. Since p does not occur in α , we may assume that $T, n \not\models p$ for every node $n \in N$. We then define the tree $T' = (N', E', r, \rho')$, where

- $N' := N \cup \{x\}$,
- $E' := E' \cup \{(m, x)\}$,
- $\rho'(y) := \begin{cases} p & \text{if } y = x, \\ \rho(y) & \text{otherwise.} \end{cases}$

By definition of T' , we have $T', m \models \langle \downarrow \rangle p$ and $(n, m) \in \llbracket \alpha \rrbracket_{T'}$. Hence, $T', n \models \langle \alpha \rangle \langle \downarrow \rangle p$. By the assumption, it implies that $T', n \models \langle \beta \rangle \langle \downarrow \rangle p$. Since p holds at x only, and x is a child of m , we have that $(n, m) \in \llbracket \beta \rrbracket_{T'}$. By definition of T' , the path between n and m is also in T . Thus, $(n, m) \in \llbracket \beta \rrbracket_T$, as desired.

Item (iii) easily follows from the definitions. □

In light of Proposition 2.6, from now on we consider the containment problem of node formulas only. We now give an interesting example of CTP containment.

Example 2.7. Let us consider the CTP node formulas

$$\begin{aligned}\varphi &= \langle\langle\langle b^+ \rangle d^+\rangle\rangle\langle a^+ \rangle b \text{ and} \\ \psi &= \langle\langle\langle c^+ \rangle d^+\rangle\rangle\langle a^+ \rangle b.\end{aligned}$$

Although it is hard to see from the first glance, $\varphi \subseteq \psi$ holds. Indeed, in every model T of φ either there is a direct child of the root where $\langle a^+ \rangle b$ holds, or there is a path $r = v_1, \dots, v_n$ in T to the node v_n where $\langle a^+ \rangle b$ holds and $\langle b^+ \rangle d$ holds in every node v_i ($1 < i < n$) on the path. In the first case, ψ holds at the root since $\langle a^+ \rangle b$ holds at the direct child.

Now let's consider the second case. Let $j \in \{2, \dots, n-1\}$ be the least number with the property that v_j has a non-empty b -path to the d -descendant. If there is no such number, then $\langle c^+ \rangle d$ holds at every v_i , $1 < i < n$ (as each of them has a direct d -child) and, thus, $T, r \models \psi$.

Assume there is such a j . Since v_j has a b -node as a child, we have that $T, v_j \models \langle a^+ \rangle b$. Moreover, $\langle c^+ \rangle d$ holds at each v_i for $i < j$ since v_i has a d -node as a child. Thus in this case we obtain that $T, r \models \psi$ holds too.

3. EXPRESSIVITY

We extend the semantics of TP given by embeddings of queries into trees to CTP. Instead of embeddings we need a simulation known from temporal logic.

3.1. Interpreting Conditional Tree Patterns by simulations. The semantics of conditional tree patterns can be defined using simulations developed for LTL [5]. These simulations generalize the embeddings for tree patterns from [1, 15] with an additional clause for checking the labels on the edges.

We start with defining the tree pattern analogues of CTP node and path expressions.

Definition 3.1. A *conditional tree pattern* is a finite tree $(N, E, r, \bar{o}, \rho_N, \rho_E)$ with labeled nodes and edges, where N is the set of nodes of the tree, $E \subseteq N \times N$ is the set of edges, r is the root of the tree, \bar{o} is a k -tuple ($k > 0$) of output nodes, ρ_N is the function assigning to each node in N a finite set of labels from Σ and ρ_E is the function assigning to each pair in E either \downarrow or a Boolean conditional tree pattern.

A *Boolean conditional tree pattern* is a conditional tree pattern with a single output node which equals to the root.

A conditional tree patterns is said to have *multiple output nodes* if the number $k = |\bar{o}|$ is greater than 1.

A *tree pattern* is a conditional tree pattern whose edges are only labeled by \downarrow or \top .

To be consistent with the pictorial representation of TP, in CTP patterns an edge labeled with \downarrow is drawn as a single line, while an edge labeled with a CTP node formula is drawn as a double line with a CTP pattern as the label (e.g. as in Figure 2). The output nodes have the square shape.

CTP node and path expressions can be translated into (Boolean) conditional tree patterns with one output node and vice-versa. The translations are given in Appendix A. We denote the equivalent (Boolean) conditional tree pattern of a CTP path or node expression α or φ by $c(\alpha)$ and $c(\varphi)$, respectively.

Next we generalize the notion of TP-embeddings to CTP-simulations.

Definition 3.2. Let $T = (N, E, r, \bar{o}, \rho_N, \rho_E)$ be a conditional tree pattern as in the previous definition and $T' = (N', E', r', \rho')$ a tree. A total function $f : N \rightarrow N'$ is called a *simulation* from the pattern T into the pointed tree T', r' if it satisfies the following properties:

root preserving: $f(r) = r'$;

label preserving: if $p \in \rho_N(n)$, then $p \in \rho'(f(n))$;

child edge preserving: if nEn' and $\rho_E(n, n') = \downarrow'$, then $f(n)Ef(n')$;

conditional edge simulation: if nEn' and $\rho_E(n, n')$ is not equal to \downarrow , then $f(n)E'^+f(n')$ and for every x such that $f(n)E'^+xE'^+f(n')$ there is a simulation from the Boolean conditional tree pattern $\rho_E(n, n')$ into $T'.x$ (the subtree of T' rooted at x).

When the pattern is a tree pattern, \downarrow and \top are the only labels on edges. For the label \top , the 'conditional edge simulation' clause trivializes to checking that f is an embedding for the descendant edges. Simulations for tree patterns are thus equivalent to tree pattern embeddings [1, 15].

The next theorem states that simulations can be used to evaluate conditional tree patterns.

Theorem 3.3. *Let φ and α be a CTP node and path expression, respectively. Let T be a tree and n a node in T .*

- (i) $T, n \models \varphi$ if and only if there is a simulation from $c(\varphi)$ into $T.n$.
- (ii) $(n, n') \in \llbracket \alpha \rrbracket_T$ if and only if there is a simulation from $c(\alpha)$ into $T.n$ which relates the output node of $c(\alpha)$ to node n' .
- (iii) Items (i) and (ii) also hold when T is an infinite tree.

The proof is by mutual induction on the node and path expressions.

Remark 3.4. The notions of (conditional) tree pattern and embedding for tree patterns and simulation for conditional tree patterns are easily extended to work for expansions of these languages with negated labels (denoted by TP^\neg and CTP^\neg , respectively).

For (conditional) tree patterns, add a second node labelling function $\rho_N^-(\cdot)$ which also assigns each node a finite set of labels. These are interpreted as the labels which are false at the node.

For the embeddings and simulations, in Definition 3.2 we add a clause stating that also negated labels are preserved:

negated label preserving: If $p \in \rho_N^-(n)$, then $p \notin \rho'(f(n))$.

With these modifications, Theorem 3.3 also holds for TP^\neg and CTP^\neg .

3.2. Expressivity characterization. In this section we give expressivity characterization for CTP similar to the one for various fragments of XPath in [3]. The exact logical characterization allows one to compare different fragments of (Regular) XPath and derive non-trivial closure properties such as closure under intersection. We show that CTP path formulas correspond to a natural fragment of first order logic (**FO**) and are closed under intersections.

For φ a formula, let $\text{desc}_\varphi(x, y)$ be an abbreviation of the "until" formula $\text{desc}(x, y) \wedge \forall z(\text{desc}(x, z) \wedge \text{desc}(z, y) \rightarrow \varphi(z))$. Let $\exists^+(\text{child}, \text{desc}_\varphi)$ be the fragment of first order logic built up from the binary relations *child* and *desc*, label predicates $p(x)$ for each label $p \in \Sigma$ and equality $'='$, by closing under \wedge , \vee and \exists as well as under the rule:

$$\text{if } \varphi(x) \in \exists^+(\text{child}, \text{desc}_\varphi), \text{ then } \text{desc}_\varphi(x, y) \in \exists^+(\text{child}, \text{desc}_\varphi). \quad (3.1)$$

We use $\exists^+(child, desc_\varphi)(c, s)$ to denote $\exists^+(child, desc_\varphi)$ formulas with exactly two variables c and s free. Following [3] we restrict the **FO** fragment to its downward fragment $\exists^+(child, desc_\varphi)[down](c, s)$ by requiring that every bound variable as well as s is syntactically restricted to be a descendant of c or equal to c .

Note that $\exists^+(child, desc_\varphi)[down](c, s)$ without the rule (3.1) is the logic $\exists^+(child, desc)[down](c, s)$ from [3], and shown there to be equivalent to unions of tree patterns. Even though $\exists^+(child, desc_\varphi)[down](c, s)$ does not contain negation, not every formula is satisfiable (e.g. $desc(x, x)$).

We can now characterize conditional tree patterns in terms of $\exists^+(child, desc_\varphi)[down](c, s)$.

Theorem 3.5. *The following languages are equivalent in expressive power:*

- unions of the false symbol and CTP paths which can have disjunctions in the node formulas
- $\exists^+(child, desc_\varphi)[down](c, s)$.

Proof. A standard translation $TR_{xy}(\cdot)$ turns any CTP path formula α into an equivalent formula in $\exists^+(child, desc_\varphi)[down](c, s)$. The translation is essentially the semantics of path and node formulas written in the first order language.

$$\begin{aligned}
TR_{xy}(\emptyset) &= child(x, y) \wedge x = y \\
TR_{xy}(\downarrow) &= child(x, y) \\
TR_{xy}(?\varphi) &= x = y \wedge TR_x(\varphi) \\
TR_{xy}(\alpha_1; \alpha_2) &= \exists z.(TR_{xz}(\alpha_1) \wedge TR_{zy}(\alpha_2)), \\
&\quad \text{where } z \text{ is a fresh variable} \\
TR_{xy}((\downarrow; ?\varphi)^*; \downarrow) &= desc_{TR_z(\varphi)}(x, y) \\
TR_{xy}(\alpha_1 \cup \alpha_2) &= TR_{xy}(\alpha_1) \vee TR_{xy}(\alpha_2) \\
\\
TR_x(p) &= p(x) \\
TR_x(\top) &= x = x \\
TR_x(\varphi_1 \wedge \varphi_2) &= TR_x(\varphi_1) \wedge TR_x(\varphi_2) \\
TR_x(\varphi_1 \vee \varphi_2) &= TR_x(\varphi_1) \vee TR_x(\varphi_2) \\
TR_x(\langle \alpha \rangle \varphi) &= \exists y.(TR_{xy}(\alpha) \wedge TR_y(\varphi)), \\
&\quad \text{where } y \text{ is a fresh variable.}
\end{aligned}$$

By definition, $TR_{cs}(\alpha)$ is a formula in $\exists^+(child, desc_\varphi)[down](c, s)$ and equivalent to α .

For the other direction, let $\theta \in \exists^+(child, desc_\varphi)[down](c, s)$. We use our earlier notation $desc_\psi(x, y)$ for the "until" formulas. First we introduce two special formulas $\top(x)$ and $\perp(x)$ which stand for $x = x$ and $desc(x, x)$, respectively.

We will modify θ in several steps. First we replace $child(x, y)$ and $desc(x, y)$ by the equivalent $desc_\perp(x, y)$ and $desc_\top(x, y)$, respectively. Then eliminate all equalities by renaming variables. Bring all existential quantifiers inside θ to the front, bring the body into disjunctive normal form and distribute the disjunctions over the quantifiers. We then end up with a disjunction of formulas of the form $\exists \bar{x}\varphi(c, s)$, with φ a conjunction of formulas $desc_\psi(x, y)$, $p(x)$, $\top(x)$ and $\perp(x)$. If φ contains \perp replace $\exists \bar{x}\varphi(c, s)$ with \perp .

As our target language is closed under unions and \perp , we only need to translate the "conjunctive queries" $\exists \bar{x}\varphi(c, s)$. Consider the graph of the variables in φ in which two variables x, y are related if φ contains an atom $desc_\psi(x, y)$. If the graph is cyclic, it cannot be satisfied on a tree and $\exists \bar{x}\varphi(c, s)$ is equivalent to \perp . If it is a tree, c will be the root. Assuming that we can rewrite all ψ in the atoms $desc_\psi(x, y)$, we can rewrite it as a conditional tree pattern. If ψ is a boolean combination of $p(z)$, $\top(z)$ and $\perp(z)$ atoms, this is not hard:

Simply bring it into disjunctive normal form, and remove each conjunct containing \perp . If the result is not empty, then translate to a union of (trivial) conditional tree patterns. If it is empty, translate $desc_\psi(x, y)$ to a child step. If ψ contains subformulas of the form $desc_\varphi(x, y)$, we apply the current procedure to it.

Thus assume that the variable graph is a directed acyclic graph. We eliminate undirected cycles step by step. The length of the cycle equals the number of variables in it. Consider that the graph contains two paths π_1 and π_2 both going from x to y and without other common variables. We show that this subgraph is equivalent to a union of subgraphs with no or smaller cycles.

In the simplest case, both paths are of length 1 and thus consist of a $desc_\psi(x, y)$ atom. But then we can use the following equivalence to remove the cycle:

$$desc_\varphi(x, y) \wedge desc_\psi(x, y) \equiv desc_{\psi \wedge \psi}(x, y) \quad (3.2)$$

So assume the cycle looks like Figure 3.

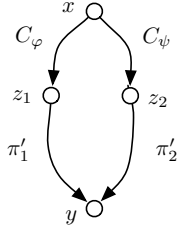


FIGURE 3. undirected cycle in φ

If this is satisfied on a tree (N, E, r, ρ) with assignment g , there are three possibilities: $g(z_1) = g(z_2)$, $g(z_1)E^+g(z_2)$ or $g(z_2)E^+g(z_1)$. In each case our original formula is equivalent to one with a smaller cycle. The three possibilities are

when $g(z_1) = g(z_2)$: $desc_{\varphi \wedge \psi}(x, z_1) \wedge z_1\pi'_1y \wedge z_1\pi'_2y$

when $g(z_1)E^+g(z_2)$: $desc_{\varphi \wedge \psi}(x, z_1) \wedge \psi(z_1) \wedge z_1\pi'_1y \wedge desc_\psi(z_1, z_2) \wedge z_2\pi'_2y$

when $g(z_2)E^+g(z_1)$: $desc_{\varphi \wedge \psi}(x, z_2) \wedge \varphi(z_2) \wedge z_2\pi'_2y \wedge desc_\varphi(z_2, z_1) \wedge z_1\pi'_1y$.

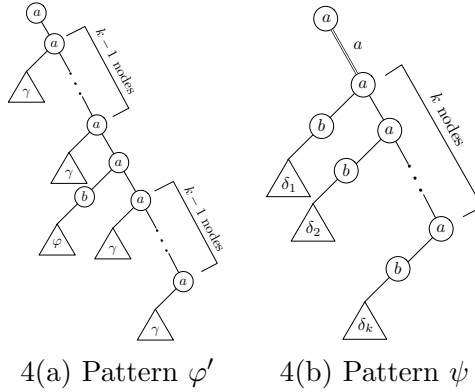
In the first case the length of the cycle decreased by two, in the two other cases by one.

Thus their disjunction is equivalent to the formula of Figure 3. We replace that formula by this disjunction, bring the result in disjunctive normal form and distribute the disjuncts out. We again have a disjunction of "conjunctive queries". As the new cycles are smaller, this procedure will terminate, and results in a (big) disjunction of trees. \square

An important consequence of this result is that CTP patterns are closed under intersection:

Theorem 3.6. *The intersection of two CTP paths is equivalent to \perp or a union of CTP paths.*

Theorem 3.5 together with the translation of Regular XPath into $\mathbf{FO}^*(c, s)$ from [17] implies that every union of CTP patterns with disjunctions is in the intersection of first order logic and positive downward $\mathbf{FO}^*(c, s)$. It is an intriguing open problem whether the converse also holds.

FIGURE 4. Patterns φ' and ψ from Proposition 4.2.

4. CONTAINMENT

Before we determine the exact complexity of the containment problem for CTP and expansions we prove a number of reductions. Most are generalizations from TP to CTP. The key new result is the encoding of negated labels in CTP.

4.1. Containment Preliminaries. The following reductions will be used later in our upper and lower bound proofs.

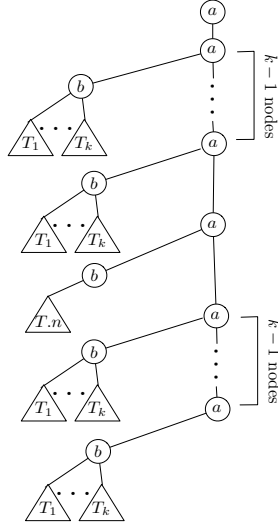
Multiple output nodes. The main difference between tree patterns and their XPath formulation is that tree patterns can have multiple output nodes. Kimelfeld and Sagiv [10] (Proposition 5.2) show that for the TP containment problem the number of output nodes is not important: the problem can be PTIME reduced to a containment problem of Boolean TPs. This is achieved, given $\varphi, \psi \in \text{TP}$, by adding a child labeled with a new label a_i to every output node X_i in both φ and ψ . Second, to every leaf that is not an output node or a newly added node, we add a child labeled with \top . The same result, using the same argument, holds for CTP.

Proposition 4.1. *Let $S \subseteq \{\vee, \neg\}$. For CTP^S patterns with multiple output nodes φ, ψ there are PTIME computable Boolean CTP^S , φ', ψ' such that $\varphi \subseteq \psi$ iff $\varphi' \subseteq \psi'$. The same holds for multi-labeled trees.*

Disjunctions in the consequent. We now show that the containment problem for both unions of CTP^\neg and unions of TP^\neg can be reduced to containments without unions. This is useful in lower bound proofs, as we can use a union in the consequent to express multiple constraints.

The proof of the following proposition is a slight modification of the proof for TP in [15].

Proposition 4.2. *Let φ be a CTP^\neg (TP^\neg) formula and Δ a finite set of CTP^\neg (resp. TP^\neg) formulas. Then there are PTIME computable CTP^\neg (TP^\neg) formulas φ' and ψ such that $\varphi \subseteq \bigvee \Delta$ iff $\varphi' \subseteq \psi$.*

FIGURE 5. Model for pattern φ' from Proposition 4.2.

If φ and Δ are in CTP (TP), φ' and ψ are in CTP (TP) as well.
The same holds for containment over multi-labeled trees.

Proof. First, for the case of TP^\neg , the proof from [15] can be readily applied here (using embeddings which also preserve negated labels, cf Remark 3.4).

Now we prove the proposition for CTP^\neg . For simplicity, we use the same letter to denote a CTP^\neg formula and its corresponding tree pattern representation. Assume Δ is $\{\delta_1, \dots, \delta_k\}$. If ι is a CTP^\neg pattern, by ι^a we denote the CTP^\neg pattern defined as having the root labeled by a , whose unique child is the root of ι . Consider new labels a and b occurring neither in φ nor Δ . Let $\gamma = \bigwedge_{0 < \ell \leq k} \delta_\ell^b$. Note that γ is consistent if every δ_ℓ is consistent. Now define φ' as the CTP^\neg tree pattern in Figure 4(a), and ψ as the CTP^\neg tree pattern in Figure 4(b).

We verify that $\varphi \subseteq \bigvee \Delta$ iff $\varphi' \subseteq \psi$. Without loss of generality, we assume that each of $\varphi, \delta_1, \dots, \delta_k$ is consistent.

For the direction from left to right, assume $\varphi \subseteq \bigvee \Delta$ and $T, n \models \varphi'$ for an arbitrary T, n . This means that there is a path $n = m_0, \dots, m_{2k-1}$ of nodes in T which are labeled with a and such that:

- for every $0 < \ell \leq k$: $T, m_i \models (\delta_\ell^b)^a$, with $i \in \{1, \dots, k-1\} \cup \{k+1, \dots, 2k-1\}$,
- $T, m_k \models (\varphi^b)^a$.

By the assumption, it follows that there is an index j such that $T, m_k \models (\delta_j^b)^a$. Then, by Theorem 3.3, there is a simulation f' from $(\delta_j^b)^a$ into T, m_k . The simulation f' can be extended to a simulation f from ψ into T, n in the obvious way such that f is a simulation from $(\delta_j^b)^a$ into T, m_{k-j+i} and the a -descendant edge in ψ is simulated on the path m_0, \dots, m_{k-j+1} . Thus, $T, n \models \psi$.

For the other direction, assume $\varphi' \subseteq \psi$ and $T, n \models \varphi$ for an arbitrary T rooted at n . As all δ_ℓ are consistent, there are trees $T_\ell, n_\ell \models \delta_\ell$, for each $\ell \in \{1, \dots, k\}$. Let T' be the tree created from T, n and the T_ℓ 's depicted in Figure 5 with root r . Then $T', r \models \varphi'$. Because we assumed $\varphi' \subseteq \psi$, then also $T', r \models \psi$, which by Theorem 3.3 implies the existence of a simulation g from ψ into T', r . In particular, the span of k a -nodes in ψ has to be simulated

on an a -path of length k in T' . This means that for some $j \in \{1, \dots, k\}$, g is a simulation from δ_j into $T.n$. By Theorem 3.3 again, $T, n \models \delta_j$, and thus $T, n \models \bigvee \Delta$.

Note that if φ and Δ are in CTP, then the constructed φ' and ψ are in CTP too. \square

From single-labeled to multi-labeled trees. We now show that the containment problem over single-labeled trees can be reduced to a containment problem over multi-labeled trees.

Proposition 4.3. *Let $S \subseteq \{\neg, \vee\}$. Given CTP^S (TP^S) patterns φ, ψ , there are PTIME computable CTP^S (TP^S) patterns φ', ψ' such that $\varphi \subseteq \psi$ iff $\varphi' \subseteq_{\text{ML}} \psi'$.*

Proof. Let p_1, \dots, p_n be the labels from Σ occurring in φ or ψ . We take φ'' as φ , and $\psi'' := \psi \vee \bigvee_{1 \leq i < j \leq n} (p_i \wedge p_j) \vee \bigvee_{1 \leq i < j \leq n} \langle \downarrow^+ \rangle (p_i \wedge p_j)$. The disjuncts added to ψ ensure that every node of a counterexample is labeled with at most one label. Now we show that $\varphi \subseteq \psi$ iff $\varphi'' \subseteq_{\text{ML}} \psi''$

(\Rightarrow) We show the contraposition. Let T be a multi-labeled tree such that $T \models \varphi''$ and $T \not\models \psi''$. Then let T' be the tree obtained from T by restricting the labels $\rho'(v) := \rho(v) \cap \{p_1, \dots, p_n\}$ for every node v . Since T does not satisfy ψ'' , the label of every node in T' contains at most one symbol. Because we did not change the valuation of the labels p_1, \dots, p_n , for each formula θ constructed from these tags, it holds that $T, v \models \theta$ iff $T', v \models \theta$. From this and the fact that the label of each node of T' is of size at most one, it follows that $T' \models \varphi''$ and $T' \not\models \psi''$. In order to make a single-labeled tree out of T' , we add a dummy symbol q in the label of a node if its label in T' is empty. For the resulting tree T'' it holds $T'' \models \varphi$ and $T'' \not\models \psi$.

(\Leftarrow) Let T be a single-labeled tree such that $T \models \varphi$ and $T \not\models \psi$. Then this tree can be considered as a multi-labeled tree. Moreover, $T \models \varphi''$ and $T \not\models \psi''$ since T does not satisfy any of the disjuncts in ψ'' .

Our transformation does not introduce negations, but only disjunctions in ψ'' . Thus the proposition is proven for the cases when S contains disjunction, where we take $\varphi' := \varphi''$ and $\psi' := \psi''$. For the remaining cases ($S \subseteq \{\neg\}$), although ψ'' contains disjunctions, by Proposition 4.2, there exist two PTIME computable CTP^S (TP^S) patterns φ' and ψ' such that $\varphi'' \subseteq_{\text{ML}} \psi''$ iff $\varphi' \subseteq_{\text{ML}} \psi'$. This concludes the proof. \square

Negated labels. We now show that CTP is expressive enough to encode label negation, as far as the containment problem is concerned.

The trick of the encoding lies in the fact that two nodes connected by a descendent edge are in either child or descendant of a child relation. Additionally, if we require that neither of these relations occur at the same time, we can faithfully encode label negation.

Proposition 4.4. *Let φ and ψ be CTP[¬] patterns. There are PTIME computable CTP patterns φ' and ψ' such that*

$$\varphi \subseteq \psi \text{ iff } \varphi' \subseteq \psi'.$$

This also holds for containment over multi-labeled trees.

Proof. Given a containment problem $\varphi \subseteq \psi$ in CTP[¬], we construct an equivalent containment problem $\varphi^\bullet \subseteq \psi^\circ$ with φ^\bullet in CTP and ψ° a union of CTP patterns. Applying Proposition 4.2 then yields the desired result.

Let $p_1, \neg p_1, \dots, p_n, \neg p_n$ be the labels appearing in φ and ψ and their negations. Let s be a new label and ζ be the formula $\langle \downarrow \rangle (s \wedge \bigwedge_i \langle \downarrow^+ \rangle p_i)$. We define the translation $(\cdot)^\bullet$ inductively:

$$\begin{aligned} \top^\bullet &= \zeta \\ (p_i)^\bullet &= \zeta \wedge \langle \downarrow \rangle (s \wedge \langle \downarrow \rangle p_i) \\ (\neg p_i)^\bullet &= \zeta \wedge \langle \downarrow \rangle (s \wedge \langle \downarrow \rangle \langle \downarrow^+ \rangle p_i) \\ (\theta \wedge \sigma)^\bullet &= \theta^\bullet \wedge \sigma^\bullet \\ (\langle \downarrow \rangle \varphi)^\bullet &= \zeta \wedge \langle \downarrow \rangle \varphi^\bullet \\ (\langle \varphi^+ \rangle \psi)^\bullet &= \zeta \wedge \langle \varphi^{\bullet+} \rangle \psi^\bullet. \end{aligned}$$

The translation is defined in such a way that for any φ , φ^\bullet implies ζ . Obviously $(\cdot)^\bullet$ is a PTIME translation. Let $\langle \downarrow^* \rangle \varphi$ denote $\varphi \vee \langle \downarrow^+ \rangle \varphi$. Let AX be the disjunction of the following formulas:

$$\langle \downarrow^* \rangle (s \wedge \langle \downarrow^* \rangle \zeta) \tag{4.1}$$

$$\bigvee_i \langle \downarrow^* \rangle (p_i^\bullet \wedge (\neg p_i)^\bullet) \tag{4.2}$$

$$\bigvee_{i \neq j} \langle \downarrow^* \rangle (p_i^\bullet \wedge p_j^\bullet). \tag{4.3}$$

In our intended models, none of these disjuncts are true at the root. The disjunct (4.1) is technical, (4.2) states that no node makes the encodings of both p and $\neg p$ true, and (4.3) ensures that no node makes the encoding of two different labels true. We define $\psi^\circ = \psi^\bullet \vee AX$. In case of multi-labeled trees the translation is the same, except that we do not include the disjunction (4.3) in the definition of ψ° . Note that although in ψ° we have disjunctions within the scope of a modality, we can (in PTIME) rewrite the formula into an equivalent union of disjunction free formulas.

We now show that $\varphi \subseteq \psi$ iff $\varphi^\bullet \subseteq \psi^\circ$.

(\Leftarrow) We show the contrapositive. Given a single-labeled tree $T = (N, E, r, l)$ such that $T, r \models \varphi$ and $T, r \not\models \psi$, we build a single-labeled tree $T' = (N', E', r', l')$ as follows. The set of nodes N' is the maximal E' -connected subset of $N \times \{0, 1, 2, 3\} \times \{p_1, \dots, p_n, \#\}$ which contains the root $r' = (r, 0, \#)$. The relation E' is defined as follows: $(n', 0, \#)$ is the parent of $(n, 0, \#)$ when n' is the parent of n in T ; $(n, 0, \#)$ is the parent of $(n, 1, \#)$; $(n, 1, \#)$ is the parent of $(n, 2, p_i), i = 1, \dots, n$, and $(n, 2, p_i)$ is the parent of $(n, 3, p_i), i = 1, \dots, n$. The labeling l' is defined as follows: s only labels nodes of the form $(n, 1, \#)$. Then, p_i labels $(n, 2, p_i)$ if $T, n \models p_i$ and p_i labels $(n, 3, p_i)$ if $T, n \not\models p_i$. No other nodes are labeled by p_i . In all other cases, we label a node by a fresh label z .

It is clear from the definition that T' is a single-labeled tree. Note that $T' \not\models AX$ and that ζ is true at all nodes of type $(n, 0, \#)$, and only at these nodes. Let θ be a formula in variables $\{p_1, \dots, p_n\}$, and $n \in N$. By induction on θ , we show $T, n \models \theta$ iff $T', (n, 0, \#) \models \theta^\bullet$.

- $\theta = \top$. This holds because by construction ζ is true at all nodes of type $(n, 0, \#)$.
- $\theta = p$. We have $T, n \models p \Leftrightarrow l(n) = p \Leftrightarrow l'(n, 2, p) = p \Leftrightarrow$ (since $(n, 2, p)$ is the only child of $(n, 1, \#)$ labeled with p and $(n, 1, \#)$ is labeled by s)
 $T', (n, 1, \#) \models s \wedge \langle \downarrow \rangle p \Leftrightarrow$ (since $(n, 1, \#)$ is the only child of $(n, 0, \#)$ and ζ is always true at $(n, 0, \#)$) $T', (n, 0, \#) \models \zeta \wedge \langle \downarrow \rangle (s \wedge \langle \downarrow \rangle p) \Leftrightarrow T', (n, 0, \#) \models p^\bullet$.

- $\theta = \neg p$. We have $T, n \models \neg p \Leftrightarrow l(n) \neq p \Leftrightarrow l'(n, 3, p) = p \Leftrightarrow$ (since $(n, 3, p)$ is the only descendent of $(n, 1, \#)$ labeled with p and $(n, 1, \#)$ is labeled by s)
 $T', (n, 1, \#) \models s \wedge \langle \downarrow \rangle \langle \downarrow^+ \rangle p \Leftrightarrow$ (since $(n, 1, \#)$ is the only child of $(n, 0, \#)$ and ζ is always true at $(n, 0, \#)$) $T', (n, 0, \#) \models \zeta \wedge \langle \downarrow \rangle (s \wedge \langle \downarrow \rangle \langle \downarrow^+ \rangle p) \Leftrightarrow T', (n, 0, \#) \models (\neg p)^\bullet$.
- $\theta = \varphi_1 \wedge \varphi_2$. We have $T, n \models \varphi_1 \wedge \varphi_2 \Leftrightarrow T, n \models \varphi_1$ and $T, n \models \varphi_2 \Leftrightarrow$ (by the induction hypothesis) $T', (n, 0, \#) \models \varphi_1^\bullet$ and $T', (n, 0, \#) \models \varphi_2^\bullet \Leftrightarrow T', (n, 0, \#) \models \varphi_1^\bullet \wedge \varphi_2^\bullet \Leftrightarrow T', (n, 0, \#) \models (\varphi_1 \wedge \varphi_2)^\bullet$.
- $\theta = \langle \downarrow \rangle \varphi$. We only show the right to left direction. $T', (n, 0, \#) \models (\langle \downarrow \rangle \varphi)^\bullet \Leftrightarrow T', (n, 0, \#) \models \zeta \wedge \langle \downarrow \rangle \varphi^\bullet \Leftrightarrow \exists m \in N'$ such that $(n, 0, \#)E'm$ and $T', m \models \varphi^\bullet$. But then $T', m \models \zeta$ because φ^\bullet implies ζ . Thus m must be of the form $(n', 0, \#)$ as only these make ζ true. But then $n' \in N$ and nEn' and we may apply the inductive hypothesis to get $T, n' \models \varphi$, and thus $T, n \models \langle \downarrow \rangle \varphi$.
- $\theta = \langle \varphi^+ \rangle \psi$. We only show the right to left direction. $T', (n, 0, \#) \models (\langle \varphi^+ \rangle \psi)^\bullet \Leftrightarrow T', (n, 0, \#) \models \zeta \wedge \langle \varphi^+ \rangle \psi^\bullet \Leftrightarrow \exists m \in N'$ such that $(n, 0, \#)E'^+m$ and $T', m \models \psi^\bullet$ and $\forall m'. (n, 0, \#)E'^+m'E'^+m$ it holds $T', m' \models \varphi^\bullet$. Now, because φ^\bullet implies ζ for all φ, m and all nodes m' are of the form $(n', 0, \#)$ and thus all in the original model N . Moreover they stand in the same way in the E relation. Thus we can apply the inductive hypothesis and obtain $T, n \models \langle \varphi^+ \rangle \psi$.

As a special case, $T', (r, 0, \#)$ satisfies φ^\bullet but not ψ^\bullet . Recall that $T', (r, 0, \#)$ does not satisfy the other disjuncts AX of ψ° either and thus, $T', (r, 0, \#) \models \varphi^\bullet$ and $T', (r, 0, \#) \not\models \psi^\circ$, as desired.

(\Rightarrow) Again we show the contrapositive. Suppose there is a model $T = (N, E, r, l)$ satisfying φ^\bullet but not ψ° at the root r . Then in particular $T, r \not\models AX$. Without loss of generality, we can assume that the simulation from (the conditional tree pattern corresponding to) φ^\bullet into T is surjective. (Otherwise the image of φ^\bullet is a subtree of T and thus by monotonicity ψ° cannot be satisfied at the root of this subtree.) In this model, as a consequence of the fact that φ^\bullet implies ζ , every branch has an initial segment satisfying ζ , immediately followed by a node labeled by s , which is followed by a segment where ζ is never satisfied because of the first disjunct (4.1) of AX .

We define a tree $T' = (N', E', r', l')$ whose set of nodes N' consists of the nodes of T where ζ is satisfied, i.e. $N' = \{n \in N \mid T, n \models \zeta\}$; E' is simply the restriction of E to N' and $r' = r$. We define $l'(n) = p_i$ iff $T, n \models p_i^\bullet$, and if $T, n \not\models p_i^\bullet$ for any i , then we label n with a fresh variable z . This definition of l' is well-defined because the falsity of the disjunction (4.3) in AX ensures that each node makes at most one p_i^\bullet true in T .

Let θ be a formula in variables $\{p_1, \dots, p_n\}$, and $n \in N$. By induction on θ , we show that

$$T, n \models \theta^\bullet \text{ iff } T', n \models \theta. \quad (4.4)$$

- $\theta = \top$. Then $T, n \models \top^\bullet \Leftrightarrow$ (by definition of $(\cdot)^\bullet$) $T, n \models \zeta \Leftrightarrow$ (by definition of N') $\Leftrightarrow T', n \models \top$.
- $\theta = p$. Then $T, n \models p^\bullet \Leftrightarrow$ (by definition of the labeling l') $l'(n) = p \Leftrightarrow T', n \models p$.
- $\theta = \neg p$. If $T, n \models (\neg p)^\bullet$, then by the falsity of AX , $T, n \not\models p^\bullet$, and thus $l'(n) \neq p$ and $T', n \not\models p$ and thus $T', n \models \neg p$. Conversely, $T', n \models \neg p \Leftrightarrow l'(n) \neq p \Leftrightarrow T, n \not\models p^\bullet$. But $T, n \models \zeta$ and thus either p^\bullet or $(\neg p)^\bullet$ must hold at T, n . Thus $T, n \models (\neg p)^\bullet$.
- $\theta = \varphi_1 \wedge \varphi_2$. Then $T, n \models (\varphi_1 \wedge \varphi_2)^\bullet \Leftrightarrow T, n \models \varphi_1^\bullet$ and $T, n \models \varphi_2^\bullet \Leftrightarrow (n \in N'$ because φ^\bullet implies ζ and thus by the inductive hypothesis) $T', n \models \varphi_1$ and $T', n \models \varphi_2 \Leftrightarrow T', n \models \varphi_1 \wedge \varphi_2$.

- $\theta = \langle \downarrow \rangle \varphi$. Then $T, n \models (\langle \downarrow \rangle \varphi)^\bullet \Leftrightarrow T, n \models \zeta$ and $T, n \models \langle \downarrow \rangle \varphi^\bullet \Leftrightarrow$ there exists $n' \in N$ such that nEn' and $T, n' \models \varphi^\bullet \Leftrightarrow$ (by the fact $T, n' \models \zeta$ and, thus, $n' \in N'$ and the inductive hypothesis) there exists $n' \in N'$ such that $nE'n'$ and $T', n' \models \varphi \Leftrightarrow T', n' \models \langle \downarrow \rangle \varphi$, as desired.
- $\theta = \langle \theta^+ \rangle \tau$. Assume $T, n \models (\langle \theta^+ \rangle \tau)^\bullet$. Then, by definition of $(\cdot)^\bullet$, $T, n \models \zeta \wedge \langle \theta^{++} \rangle \tau^\bullet$. That means there exists n' in T with nE^+n' such that $T, n' \models \tau^\bullet$ and for all n'' with $nE^+n''E^+n'$ it holds $T, n'' \models \theta^\bullet$. By definition of $(\cdot)^\bullet$, the nodes n, n' and all the nodes between n and n' satisfy ζ and thus belong to T' . By inductive hypothesis, $T', n' \models \tau$ and $T', n'' \models \theta$ for all $nE^+n''E^+n'$, which means $T', n \models \langle \theta^+ \rangle \tau$ holds. Conversely, assume $T', n \models \langle \theta^+ \rangle \tau$. By definition of T' , we have that $T, n \models \zeta$, which is the first conjunct of $(\langle \theta^+ \rangle \tau)^\bullet$. The second conjunct follows by the inductive hypothesis.

As T is a counterexample for $\varphi^\bullet \subseteq \psi^\circ$, (4.4) implies that T' is a counterexample of $\varphi \subseteq \psi$, as desired.

The same argument applies for multi-labeled trees. The only change is to remove the last disjunction (4.3) from AX . \square

4.2. Lower bounds. In this section we show that the containment for TP^\neg is PSPACE hard. This lower bound will carry over to the containment problem for CTP.

Theorem 4.5.

- (i) *The containment problem for CTP is PSPACE-hard.*
- (ii) *The containment problem for TP^\neg is PSPACE-hard,*
- (iii) *Both results also hold for multi-labeled trees.*

Proof. (i) follows from (ii) by Proposition 4.4. (iii) follows from (i) and (ii) by Proposition 4.3. For proving (ii), we reduce the corridor tiling problem, which is known to be hard for PSPACE [7, 18], to the containment problem for TP^\neg . We use the construction from the PSPACE-hardness proof for the containment problem of TP with disjunction over a finite alphabet in [16].

The corridor tiling problem is formalized as follows. Let $\text{Til} = (D, H, V, \bar{b}, \bar{t}, n)$ be a tiling system, where $D = \{d_1, \dots, d_m\}$ is a finite set of tiles, $H, V \subseteq D^2$ are horizontal and vertical constraints, n is a natural number given in unary notation, \bar{b} and \bar{t} are tuples over D of length n . Intuitively, given a corridor of width n , the goal is to construct a tiling of the corridor using the tiles from D so that the horizontal and vertical constraints are satisfied and the bottom and top rows are tiled by \bar{b} and \bar{t} , respectively. We say that a tiling satisfies the horizontal constraints H (respectively the vertical constraints V) if for every $1 \leq i \leq n, j \in \mathbb{N}$ it holds that if d_1 and d_2 are the tiles on the positions (i, j) and $(i + 1, j)$ of the corridor (respectively (i, j) and $(i, j + 1)$), then $(d_1, d_2) \in H$ ($(d_1, d_2) \in V$).

Now we construct in PTIME in the length of Til , two TP^\neg expressions φ and ψ such that the following holds:

$$\varphi \not\subseteq \psi \text{ iff there exists a tiling for } \text{Til}. \quad (4.5)$$

By Proposition 4.2, we may without loss of generality construct ψ as a disjunction of TP^\neg expressions. To this purpose, we use the string representation of a tiling. Each row of the considered tiling is represented by the tiles it consists of. If the tiling of a corridor of width n has k rows, it is represented by its rows separated by the special symbol \sharp . The top row is followed by the symbol $\$$. Thus, a tiling is a word of the form $u_1 \sharp u_2 \sharp \dots \sharp u_k \$$, where

each u_i is the word of length n corresponding to the i -th row in the tiling. In particular $u_1 = \bar{b}$ and $u_k = \bar{t}$. For the sake of readability, for expression r , r^i denotes the path formula $?r; \downarrow; ?r; \dots; \downarrow; ?r$ with i occurrences of r .

Let φ be

$$\langle ?b_1; \downarrow; ?b_2; \dots; \downarrow; ?b_n; \downarrow; ?\sharp; \downarrow^+ ; ?t_1, \downarrow; \dots \downarrow; ?t_n; \downarrow \rangle \$.$$

Intuitively, this expression enforces a tiling to start with a path starting with \bar{b} and finishing with \bar{t} and the final symbol $\$$. Now the formula ψ defines all incorrect tilings and additional constraints. It is the disjunction of the following TP^\top formulas.

- (0) $\langle \downarrow^+ ; ?(-d_1 \wedge \dots \wedge -d_m \wedge -\sharp); \downarrow^+ \rangle \$$. There is a position that is labeled neither by a tile nor a delimiter.
- (1) Incorrect length of a row.
- (1a) $\bigvee_{i=0}^{n-1} \langle \downarrow^+ ; ?\sharp; \top^i; \downarrow \rangle \sharp$, a row is too short;
- (1b) $\langle \downarrow^+ ; (-\sharp)^{n+1} \rangle \top$, a row is too long;
- (2) $\bigvee_{(d_1, d_2) \notin H} \langle \downarrow^+ ; ?d_1; \downarrow; ?d_2 \rangle \top$, some horizontal constraints are violated;
- (3) $\bigvee_{(d_1, d_2) \notin V} \langle \downarrow^+ ; ?d_1; \downarrow; \top^n; \downarrow; ?d_2 \rangle \top$, some vertical constraints are violated.

Note that negated labels are used in (0) and (1b). Also note that the size of φ and ψ is bounded by a polynomial in the size of Til .

We now show (4.5).

(\Leftarrow). Assume that there exists a tiling of the corridor. Let s be the string representation of it. Then, $s = u_1 \sharp u_2 \sharp \dots \sharp u_k \$$, where $|u_i| = n$, $u_i \in D^n$, $u_1 = \bar{b}$ and $u_k = \bar{t}$. Moreover, on the one hand if $x \cdot y$, is an infix of some u_i , then $(x, y) \in H$, and on the other hand for every infix $x \cdot u' \cdot y$ of length $n + 1$ of $u_i \sharp \cdot u_{i+1}$, it holds that $(x, y) \in V$. Let T_s be the corresponding tree, i.e. a single path of $|s|$ nodes $\{v_1, \dots, v_{|s|}\}$ where the labeling is set in accordance with s , i.e. $l(v_i) = s_i$. Clearly, T_s is a model of φ and not of ψ .

(\Rightarrow). Let T be a tree such that $T, r \models \varphi$ and $T, r \not\models \psi$. Since $T, r \models \varphi$, there must exist a path $r = v_1, \dots, v_m$ in T which starts with \bar{b} and finishes with \bar{t} . Moreover, either \sharp or a symbol from D is the label of every node $v_i, 1 \leq i < m$, according to (0).

We define a tiling function $g : \{0, \dots, n-1\} \times \mathbb{N} \rightarrow D$ assigning a tile to every position in the corridor as follows: $g(i, j) = l(v_{(n+1) \times j + i + 1})$, $1 \leq i \leq n$, where l is the labeling function of T . Indeed, this function is well defined, as (1) ensures the correct counting. By formulas (2) and (3) the tiling defined by g satisfies the horizontal and vertical constraints. □

The difference between CTP patterns and TP patterns is that CTP descendent edges can be labeled by patterns. One might ask whether a bound on the degree of such a labeling nesting can lead to a lower complexity of the containment problem. Define the until nesting depth $un : \text{CTP} \rightarrow \mathbb{N}$ as follows.

- $un(p) = un(\top) = 0$,
- $un(\langle \downarrow \rangle \varphi) = un(\langle \downarrow^+ \rangle \varphi) = 1$,
- $un(\varphi_1 \wedge \varphi_2) = \max\{un(\varphi_1), un(\varphi_2)\}$,
- $un(\langle \varphi^+ \rangle \psi) = un(\varphi) + 1$.

Unfortunately, a close examination of the encoding in the lower bound proof and the encodings in Propositions 4.2 and 4.4 gives a negative answer to the question. Thus we obtain

Theorem 4.6. *The CTP containment problem for formulas of until nesting depth one is PSPACE-hard.*

4.3. Upper bounds. In this section, we show a matching PSPACE upper bound for CTP containment.

The complexity of $\text{CTP}^{\neg, \vee}$ containment follows from a translation into existential CTL ($\exists\text{CTL}$), whose containment problem is known to be PSPACE-complete [11]. The only small technical issue is that $\exists\text{CTL}$ formulas are interpreted over *infinite* finitely branching trees. We solve that by relativizing formulas with a new propositional variable s , whose interpretation will provide the desired finite tree.

Theorem 4.7. *The containment problem for $\text{CTP}^{\neg, \vee}$ is decidable in PSPACE. This also holds for multi-labeled trees.*

Proof. We prove the upper bound for the case of multi-labeled trees. The upper bound for single-labeled trees then follows by Proposition 4.3.

We recall from [11] the definition of $\exists\text{CTL}$. Let Prop be a set of propositional variables. $\exists\text{CTL}$ node formulas φ and path formulas α are defined by:

$$\begin{aligned} \varphi &::= \top \mid \perp \mid p \mid \neg p \mid \varphi \wedge \varphi \mid \varphi \vee \varphi \mid \exists \alpha \\ \alpha &::= \varphi \mid \alpha \wedge \alpha \mid \alpha \vee \alpha \mid X\alpha \mid (\alpha U \alpha) \mid (\alpha \hat{U} \alpha) \end{aligned}$$

where $p \in \text{Prop}$, and asking that in node formulas, the operators X (“next”), U (“until”) and \hat{U} (“dual of until”) are always immediately preceded by \exists .

The semantics for $\exists\text{CTL}$ is given by *infinite finitely branching trees*. $\exists\text{CTL}$ path formulas are interpreted at a (possibly infinite) paths of the tree and $\exists\text{CTL}$ node formulas are interpreted at nodes of the tree. Our reduction will translate $\text{CTP}^{\neg, \vee}$ formulas to $\exists\text{CTL}$ node formulas, with occurrences of formulas of the form $\exists X\varphi$ and $\exists\psi U\varphi$. Semantics of such formulas is defined as follows, given an infinite tree $T = (N, E, r, \rho)$ and a node $n \in N$.

$$\begin{aligned} T, n \models \exists X\varphi &\text{ iff there exists a node } m \in N \text{ such that } nEm \text{ and } T, m \models \varphi, \\ T, n \models \exists\psi U\varphi &\text{ iff there exists a node } m \in N \text{ such that } nE^*m \text{ and } T, m \models \varphi, \\ &\text{and for all } n' \in N \text{ with } nE^*n'E^+m \text{ it holds } T, n' \models \psi. \end{aligned}$$

The CTP node formula $\langle \varphi^+ \rangle \psi$ corresponds to the strict until operator, which is expressible in $\exists\text{CTL}$ as $\exists X\exists\psi U\varphi$.

We denote by \subseteq_∞ the containment relation for $\exists\text{CTL}$ node formulas over finitely branching trees where each branch is infinite.

Let φ, ψ be in $\text{CTP}^{\neg, \vee}$ and let s be a new propositional variable not in φ, ψ . The translation $(\cdot)^s$ from $\text{CTP}^{\neg, \vee}$ to $\exists\text{CTL}$ node formulas relativizes every subformula with s and adjusts the syntax.

$$\begin{aligned} (\top)^s &= s \wedge \top \\ (p)^s &= s \wedge p \\ (\neg p)^s &= s \wedge \neg p \\ (\varphi_1 \wedge \varphi_2)^s &= \varphi_1^s \wedge \varphi_2^s \\ (\varphi_1 \vee \varphi_2)^s &= \varphi_1^s \vee \varphi_2^s \\ (\langle \downarrow \rangle \varphi)^s &= s \wedge \exists X\varphi^s \\ (\langle \psi^+ \rangle \varphi)^s &= s \wedge \exists X\exists\psi^s U\varphi^s \end{aligned}$$

We claim that $\varphi \subseteq \psi$ iff $\varphi^s \subseteq_{\infty} \psi^s$.

The proof is by contraposition. First assume that for some finite tree $T = (N, E, r, \rho)$ and node $n \in N$ it holds $T, n \models \varphi$ and $T, n \not\models \psi$. We then construct an infinite tree which is a counterexample. Let s be a propositional variable not occurring in φ or ψ . Let T_{∞}^s be the infinite tree obtained from T by adding to each leaf in T an infinite path. The labeling is changed as follows: all new nodes have the empty label set and we add the label s to the label set of all old nodes. Formally, $T_{\infty}^s = (N_{\infty}, E_{\infty}, r, \rho_{\infty})$, where $N_{\infty} = N \cup \{n_1^m, n_2^m, \dots \mid m \text{ is a leaf in } T\}$ is the set of nodes, $E_{\infty} = E \cup \{(m, n_1^m) \mid n_1^m \in N_{\infty}\} \cup \{(n_i^m, n_{i+1}^m) \mid n_i^m, n_{i+1}^m \in N_{\infty}, i \geq 1\}$ is the set of edges, r is the root and ρ_{∞} is the labeling function defined as follows.

$$\rho_{\infty}(n) = \begin{cases} \rho(n) \cup \{s\} & \text{if } n \in N, \\ \emptyset & \text{if } n \in N_{\infty} \setminus N. \end{cases}$$

We show by induction on the structure of the formula that for every node $n \in N$ and every CTP ^{\neg, \vee} formula θ , $T, n \models \theta$ iff $T_{\infty}^s, n \models \theta^s$.

To show this we will use the following two facts which are easy consequences of the definitions of T_{∞}^s and the translation $(\cdot)^s$:

$$T_{\infty}^s, n \models s \Leftrightarrow n \in N, \quad (*)$$

and

$$T_{\infty}^s, n \models \varphi^s \Rightarrow T_{\infty}^s, n \models s. \quad (**)$$

- $\theta = \top$. We have $T, n \models \top$ iff $n \in N$ iff $n \in N_{\infty}$ and $s \in \rho_{\infty}(n)$ iff $T_{\infty}^s, n \models s \wedge \top$ iff $T_{\infty}^s, n \models (\top)^s$.
- $\theta = p, p \in \Sigma$. We have $T, n \models p$ iff $p \in \rho(n)$ iff (since $n \in N$) $\{p, s\} \subseteq \rho_{\infty}(n)$ iff $T_{\infty}^s, n \models s \wedge p$ iff $T_{\infty}^s, n \models (p)^s$.
- $\theta = \neg p, p \in \Sigma$. We have $T, n \models \neg p$ iff $p \notin \rho(n)$ iff (since $n \in N$ and, thus, $s \in \rho_{\infty}(n)$) $p \notin \rho_{\infty}(n)$ and $s \in \rho_{\infty}(n)$ iff $T_{\infty}^s, n \models s \wedge \neg p$ iff $T_{\infty}^s, n \models (\neg p)^s$.
- $\theta = \varphi_1 \wedge \varphi_2$. We have $T, n \models \varphi_1 \wedge \varphi_2$ iff $T, n \models \varphi_1$ and $T, n \models \varphi_2$ iff (by the induction hypothesis, $(*)$ and $(**)$) $T_{\infty}^s, n \models \varphi_1^s$ and $T_{\infty}^s, n \models \varphi_2^s$ iff $T_{\infty}^s, n \models \varphi_1^s \wedge \varphi_2^s$ iff $T_{\infty}^s, n \models \varphi^s$.
- $\theta = \varphi_1 \vee \varphi_2$. By the same argument as for $\varphi_1 \wedge \varphi_2$.
- $\theta = \langle \downarrow \rangle \varphi$. We have $T, n \models \langle \downarrow \rangle \varphi$ iff there exists $m \in N$ such that nEm and $T, m \models \varphi$ iff (by the induction hypothesis, $(*)$, $(**)$ and the fact that $n \in N$) $T_{\infty}^s, n \models s$ and there exists $m \in N_{\infty}$ such that $nE_{\infty}m$ and $T_{\infty}^s, m \models \varphi^s$ iff $T_{\infty}^s, n \models s \wedge \exists X \varphi^s$ iff $T_{\infty}^s \models (\langle \downarrow \rangle \varphi)^s$.
- $\theta = \langle \psi^+ \rangle \varphi$. We have $T, n \models \langle \psi^+ \rangle \varphi$ iff there exists $m \in N$ such that $T, m \models \varphi$ and for all $n' \in N$ with $nE^+n'E^+m$ it holds $T, n' \models \psi$. Then by the induction hypothesis, $(*)$ and $(**)$, the latter is equivalent to existence of $m \in N_{\infty}$ such that $T_{\infty}^s, m \models \varphi^s$ and for all $n' \in N_{\infty}$ with $nE_{\infty}^+n'E_{\infty}^+m$ it holds $T_{\infty}^s \models \psi^s$. The latter implication holds because for every node n' with $nE_{\infty}^+n'E_{\infty}^+m$ it holds $n' \in N$ and $nE^+n'E^+m$. Equivalently, there exists $n_1 \in N_{\infty}$ and $m \in N_{\infty}$ such that $nE_{\infty}n_1E_{\infty}^*m$, $T_{\infty}^s, m \models \varphi^s$ and for all n' with $n_1E_{\infty}^*n'E_{\infty}^+m$ it holds $T_{\infty}^s, n' \models \psi^s$. This is equivalent to $T_{\infty}^s, n \models s \wedge \exists X \exists \psi^s U \varphi^s$ as desired.

Thus, we obtain both $T_{\infty}^s, n \models \varphi^s$ and $T_{\infty}^s, n \not\models \psi^s$.

For the other direction, let $T_{\infty} = (N_{\infty}, E_{\infty}, r, \rho_{\infty})$ be some infinite tree over Σ such that $T_{\infty}, r \models \varphi^s$ and $T_{\infty}, r \not\models \psi^s$. We then remove from the model all nodes without s and their descendants. Formally, $T = (N, E, r, \rho)$, where $N = N_{\infty} \setminus \{n \mid \exists n' \in N_{\infty}. s \notin \rho_{\infty}(n')\}$

$\rho_\infty(n') \wedge n'E_\infty^*n$ is the set of nodes, $E = E_\infty|_{N \times N}$ is the set of edges, r is the root, and $\rho(n) = \rho_\infty(n)$ is the labeling function.

By induction we can show that for every $\text{CTP}^{\neg, \vee}$ formula θ and node $n \in N$, it holds $T_\infty, n \models \theta^s$ if and only if $T, n \models \theta^s$. The two non-trivial cases are the following.

- $\theta = \langle \downarrow \rangle \varphi$. We have $T_\infty, n \models s \wedge \exists X \varphi^s$ iff $T_\infty, n \models s$ and there exists $m \in N_\infty$ such that $nE_\infty m$ and $T_\infty, m \models \varphi^s$ iff $T, n \models s$ and there exists $m \in N$ such that nEm and $T, m \models \varphi^s$. The direction from right to left is obvious, while the direction from left to right because of the following. Since $T_\infty, m \models \varphi^s$, it follows that $T_\infty, m \models s$, i.e. $s \in \rho_\infty(m)$. Then since $n \in N$, we have that $m \in N$ too by the definition. The latter also implies that nEm holds in T and thus that $T, n \models s \wedge \exists X \varphi^s$.
- $\theta = \langle \psi^+ \rangle \varphi$. We have $T_\infty, n \models s \wedge \exists X \exists \psi^s U \varphi^s$ iff $T_\infty, n \models s$ and there exists $m \in N_\infty$ such that $T_\infty, m \models \varphi^s$, $nE_\infty^+ m$ and for all $n' \in N_\infty$ with $nE_\infty^+ n'E_\infty^+ m$ it holds $T_\infty, n' \models \psi^s$. This is equivalent to $T, n \models s$ and there exists $m \in N$ such that $T, m \models \varphi^s$ and for all n' with $nE^+ n'E^+ m$ it holds $T, m \models \psi^s$, which means $T, n \models s \wedge \exists X \exists \psi^s U \varphi^s$ as desired. The direction from left to write is trivial since T is a substructure of T_∞ . The direction from right to left follows from the fact that $m \in N$ and all the nodes n' such that $nE_\infty^+ n'E_\infty^+ m$ belong to N as well and, moreover, it holds $nE^+ n'E^+ m$.

Thus, we obtain that $T, r \models \varphi^s$ and $T, r \not\models \psi^s$. Because all nodes of T now make s true, we can discard the relativization with s in the formulas. Thus obtaining $T, r \models \varphi$ and $T, r \not\models \psi$. The only problem is that T is infinite. But using the simulations from Theorem 3.3 it is easy to see that we can turn T into a finite counterexample.

Thus, we have shown that $\varphi \subseteq \psi$ if and only if $\varphi^s \subseteq_\infty \psi^s$. The latter containment problem in $\exists\text{CTL}$ is known to be decidable in PSPACE. This concludes the proof. \square

Corollary 4.8. *The containment problem for TP^\neg is PSPACE complete.*

However, restricting negation to a safe negation $p \wedge \neg q_1 \wedge \dots \wedge \neg q_n$ keeps the complexity of the containment problem for tree patterns in CONP . Note that $p \wedge \neg q_1 \wedge \dots \wedge \neg q_n$ only adds expressive power over multi-labeled models, because on single labeled models $p \equiv p \wedge \neg q_1 \wedge \dots \wedge \neg q_n$. The proof of the following result can be found in [14]. Here TP^{\vee, \neg^s} is tree patterns expanded with disjunction and the safe negation construct.

Theorem 4.9. *The containment problem for TP^{\vee, \neg^s} over multi-labeled trees is in CONP .*

5. CONCLUSION

We have shown that adding conditions on the edges of tree patterns gives a boost in expressive power which comes with the price of a higher – PSPACE – complexity for the containment problem than for tree patterns. We located the source of the extra complexity in the fact that unrestricted negations of labels can be coded in Conditional Tree Patterns. Adding negations of labels to tree patterns causes an increase in complexity of the containment problem from CONP to PSPACE.

This paper is a first step in exploring “Regular Tree Patterns” and fragments of it. We mention some directions for future work.

Miklau and Suciu in [15] mention that existence of a homomorphism between tree patterns is a necessary but not sufficient condition for containment in TP . Can we extend the simulations between conditional tree patterns and trees to simulations between queries, partly capturing containment as for tree patterns?

What is the complexity of containment of regular tree patterns, i.e. the positive fragment of Regular XPath without disjunction and union? As the satisfiability (and thus also the containment) problem for Regular XPath is known to be EXPTIME-complete, it must lie in between PSPACE and EXPTIME.

There are also interesting characterization questions: what is the exact **FO** fragment corresponding to CTP or CTP with disjunction and union? Ten Cate [17] showed that full Regular XPath (with all four axis relations) expanded with path equalities is equally expressive as binary **FO*** (First-order logic extended with a transitive closure operator that can be applied to formulas with exactly two free variables). Is every positive forward binary **FO*** formula expressible as a Regular Tree Pattern? Is CTP with disjunction and union equally expressive as **FO** intersected with positive **FO***? Are unions of CTP equivalent to unions of the first order fragment of conjunctive regular path queries [6]?

6. ACKNOWLEDGEMENT

We would like to thank the anonymous referees for their helpful comments.

REFERENCES

- [1] S. Amer-Yahia, S. Cho, L. Lakshmanan, and D. Srivastava. Tree pattern query minimization. *The VLDB Journal*, 11:315–331, 2002.
- [2] V. Bárány, B. ten Cate, and L. Segoufin. Guarded negation. In L. Aceto, M. Henzinger, and J. Sgall, editors, *ICALP (2)*, volume 6756 of *Lecture Notes in Computer Science*, pages 356–367. Springer, 2011.
- [3] M. Benedikt, W. Fan, and G. M. Kuper. Structural properties of XPath fragments. *Theor. Comput. Sci.*, 336(1):3–31, 2005.
- [4] H. Björklund, W. Martens, and T. Schwentick. Conjunctive query containment over trees. *J. Comput. Syst. Sci.*, 77(3):450–472, 2011.
- [5] P. Blackburn, M. de Rijke, and Y. Venema. *Modal Logic*. Cambridge University Press, 2001.
- [6] D. Calvanese, G. D. Giacomo, M. Lenzerini, and M. Y. Vardi. Containment of conjunctive regular path queries with inverse. In *KR*, pages 176–185, 2000.
- [7] B. S. Chlebus. Domino-tiling games. *J. Comput. Syst. Sci.*, 32(3):374–392, 1986.
- [8] E. M. Clarke, E. A. Emerson, and A. P. Sistla. Automatic verification of finite-state concurrent systems using temporal logic specifications. *ACM Trans. Prog. Lang. Syst.*, 8:244–263, 1986.
- [9] G. Gottlob, C. Koch, and R. Pichler. Efficient algorithms for processing XPath queries. *ACM Trans. Database Syst.*, 30(2):444–491, 2005.
- [10] B. Kimelfeld and Y. Sagiv. Revisiting redundancy and minimization in an XPath fragment. In *EDBT’08*, pages 61–72, 2008.
- [11] O. Kupferman and M. Y. Vardi. An automata-theoretic approach to modular model checking. *ACM Trans. Prog. Lang. Syst.*, 22(1):87–128, 2000.
- [12] L. Libkin and C. Sirangelo. Reasoning about XML with temporal logics and automata. *Journal of Applied Logic*, 8(2):210 – 232, 2010. Selected papers from the Logic in Databases Workshop 2008.
- [13] M. Marx. Conditional XPath. *ACM Trans. Database Syst.*, 30(4):929–959, 2005.
- [14] M. Marx and E. Sherkhonov. Containment for queries over trees with attribute value comparisons. *Submitted*, 2015.
- [15] G. Miklau and D. Suciu. Containment and equivalence for a fragment of XPath. *J. ACM*, 51(1):2–45, 2004.
- [16] F. Neven and T. Schwentick. On the complexity of XPath containment in the presence of disjunction, DTDs, and variables. *Logical Methods in Computer Science*, 2(3), 2006.
- [17] B. ten Cate. The expressivity of XPath with transitive closure. In S. Vansummeren, editor, *PODS*, pages 328–337. ACM, 2006.

- [18] P. van Emde Boas. The convenience of tilings. In A. Sorbi, editor, *Complexity, Logic and Recursion Theory*, volume 187 of *Lecture notes in pure and applied mathematics*, pages 331–363. Marcel Dekker Inc., 1997.

APPENDIX A. TRANSLATIONS BETWEEN CTP AND *ctp*

The following constructions are similar to the tree and graph representations for tree patterns and conjunctive queries over trees [15, 4]

We define the translation function $c(\cdot)$ which assigns an equivalent conditional tree pattern with one output node to every CTP formula. The output node of $c(\varphi)$, where φ is a node formula, equals the root.

Let α and φ be path and node CTP formulas. We define $c(\alpha)$ and $c(\varphi)$ by mutual induction on the complexity of the path and node formulas. We take $c(\alpha)$ or $c(\varphi)$ to be the tree $(N, E, r, o, \rho_N, \rho_E)$, where the components are defined according to the cases.

- $\alpha = \downarrow$. Then N consists of two nodes r and n . The edge relation E is defined as $\{\langle r, n \rangle\}$. Moreover, $o := n$, $\rho_N(v) = \emptyset$ for $v \in \{r, n\}$ and $\rho_E(\langle r, n \rangle) = \downarrow$,
- $\alpha = ?\varphi$. Then $c(\alpha) := c(\varphi)$,
- $\alpha = \alpha_1; \alpha_2$. Let $c(\alpha_1)$ and $c(\alpha_2)$ be the conditional tree patterns for α_1 and α_2 . Then $c(\alpha)$ is the tree obtained as follows. We fuse the root of $c(\alpha_2)$ with the output node of $c(\alpha_1)$ and declare the output node of $c(\alpha_2)$ as the output node of $c(\alpha)$. The label of the fusion node is the union of the labels of the output node of $c(\alpha_1)$ and the root of $c(\alpha_2)$. Labels of other nodes and the edges remain the same as in $c(\alpha_1)$ and $c(\alpha_2)$.
- $\alpha = (\downarrow; ?\varphi)^*; \downarrow$. Then N consist of two nodes r and n . The edge relation E is defined as $\{\langle r, n \rangle\}$. Moreover, $o := n$, $\rho_N(v) = \emptyset$ for $v \in \{r, n\}$ and $\rho_E(\langle r, n \rangle) = c(\varphi)$.

For node formulas, the output node of the translation result is always defined as the root r .

- $\varphi = p$. Then N consists of a single node r , E is empty, the labeling $\rho_N(r) = \{p\}$.
- $\varphi = \top$. Similar to the previous case, with the exception that $\rho_N(r) = \emptyset$
- $\varphi = \varphi_1 \wedge \varphi_2$. Then $c(\varphi) := c(\varphi_1) \oplus c(\varphi_2)$, i.e. the fusion of the conditional tree patterns $c(\varphi_1)$ and $c(\varphi_2)$.
- $\varphi = \langle \alpha \rangle \varphi_1$. Let $c(\alpha)$ and $c(\varphi_1)$ be the corresponding conditional tree patterns. Then $c(\varphi)$ is obtained by fusing the output node of $c(\alpha)$ with the root of $c(\varphi_1)$. The labeling of the fusion node is defined as the union of the labels of the root of $c(\varphi_1)$ and the output node of $c(\alpha)$. Labels of the other nodes and edges remain the same as in $c(\alpha)$ and $c(\varphi_1)$.

Translation $f(\cdot)$ works the other way around. Let $t = (N, E, r, o, \rho_N, \rho_E)$ be a conditional tree pattern with one output node. We define $f(t)$ by induction on the nesting depth and the depth of the tree. We first define a mapping φ from nodes $v \in N$ to CTP node formulas. The mapping is defined inductively starting from the leaves as follows. Here, for a finite set $S = \{p_1, \dots, p_n\}$, by $\wedge S$ we denote the finite conjunction $p_1 \wedge \dots \wedge p_n$. We take $\wedge \emptyset$ to be \top . For a conditional tree pattern t , by r_t we denote the root of t . For v a node in pattern t :

$$\begin{aligned} \varphi(v) = & \wedge \rho_N(v) \wedge \\ & \bigwedge_{\langle v, v' \rangle \in E, \rho_E(v, v') = t'} \langle ?\varphi(r_{t'}) \rangle \varphi(v') \wedge \\ & \bigwedge_{\langle v, v' \rangle \in E, \rho_E(v, v') = \downarrow} \langle \downarrow \rangle \varphi(v') \end{aligned}$$

Now let $r = v_1, \dots, v_n = o$ be the path from the root r to the output node o in t . Let the expression $d_i, 1 \leq i \leq n - 1$ be defined by: $d_i = \downarrow$ if $\rho_E(v_i, v_{i+1}) = \downarrow$ and $d_i = (\downarrow; ?\varphi(r_{t'}))^*; \downarrow$ if $\rho_E(v_i, v_{i+1}) = t'$. Then the result of the translation $f(t)$ of the conditional tree pattern t is the CTP path formula:

$$?\varphi(v_1); d_1; ?\varphi(v_2); d_2; \dots; ?\varphi(v_{n-1}); d_{n-1}; ?\varphi(v_n).$$

For the next proposition we need the definition of equivalence between conditional tree patterns. Let t be a conditional tree pattern with output nodes $\bar{o}, |\bar{o}|=k$, and T a tree. Then the *answer set* of t over T is the set $Out(t, T) = \{\langle g(o_1), \dots, g(o_k) \rangle \mid g \text{ is a simulation of } t \text{ in } T\}$. We say that two conditional tree patterns t_1 and t_2 are *equivalent*, denoted as $t_1 \simeq t_2$, if $Out(t_1, T) = Out(t_2, T)$ for every tree T .

Proposition A.1. *Let φ and α be CTP node and path formulas, t a conditional tree pattern. Then it holds that*

- (i) $f(\mathbf{c}(\varphi)) \equiv \varphi$ and $f(\mathbf{c}(\alpha)) \equiv \alpha$,
- (ii) $\mathbf{c}(f(t)) \simeq t$.

The proof of (i) is by mutual induction on α and φ , and the proof of (ii) is by induction on nesting depth and depth of t .