

Tree Pattern Rewriting Systems

B. Genest³, A. Muscholl¹, O. Serre², and M. Zeitoun¹

¹LaBRI, Bordeaux & CNRS; ²LIAFA, Paris 7 & CNRS; ³IRISA, Rennes 1 & CNRS

Abstract. Classical verification often uses abstraction when dealing with data. On the other hand, dynamic XML-based applications have become pervasive, for instance with the ever growing importance of web services. We define here *Tree Pattern Rewriting Systems* (TPRS) as an abstract model of dynamic XML-based documents. TPRS systems generate infinite transition systems, where states are unranked and unordered trees (hence possibly modeling XML documents). Their guarded transition rules are described by means of tree patterns. Our main result is that given a TPRS system (T, \mathcal{R}) , a tree pattern P and some integer k such that any reachable document from T has depth at most k , it is *decidable* (albeit of non elementary complexity) whether some tree matching P is reachable from T .

1 Introduction

Classical verification techniques often use abstraction when dealing with data. On the other hand, dynamic *data-intensive* applications have become pervasive, for instance with the ever growing importance of web services. The format of the exchanged data is based on XML, which is nowadays the standard for semistructured data. XML documents can be seen as unranked trees, *i.e.*, trees in which every node can have an arbitrary (but finite) number of children, not depending on its labels. Very often, the order of siblings in the document is of no importance. In this case, trees are in addition unordered. There is a rich body of results concerning the analysis of fixed XML documents (with or without data), see e.g [11,9] for surveys on this topic.

The analysis of the *dynamics* of XML documents accessed and updated in a multi-peer environment has been considered only very recently [2,3]. Dynamically evolving XML documents are of course crucial, for instance when doing static analysis of XML-based web services. A general framework, Active XML (*AXML* for short), has been defined in [2] for the purpose of unifying data (XML) and control (services), by allowing data to be given implicitly in form of service calls.

In this paper we propose an abstract model for dynamically evolving documents, based on guarded rewriting rules on unranked, unordered tree. We show that basic properties, such as reachability of tree patterns and termination, are decidable for a natural subclass of our rewriting systems.

A standard technique to analyze unranked trees is to encode them as binary trees [11]. However, this encoding does not preserve the depth of the tree, neither

locality, nor path properties. For these reasons, we define our TPRS directly on unranked trees. The rewriting rules are based on tree patterns, occurring in two distinct contexts. First, tree patterns are used for describing how the structure of the tree changes through the rules (*e.g.*, some subtrees might be deleted, new nodes can be added). Second, rules are guarded, and the guard condition is tested via a *Tree Pattern Query* (TPQ). The role of the TPQ is actually twofold: it is used in the pre-condition of the rule, and the results can enhance the information of the new tree. We call such a system *Tree Pattern Rewriting System*, *TPRS* for short. For an easier comparison with other works, we include an example of a Mail-Order System in our presentation, close to the one used in [3].

The main tool we use to show decidability of various properties of TPRS are well-structured systems [1,8]. Such systems cover several interesting classes of infinite-state systems, such as Petri nets and lossy channel systems. Our TPRS are of course not well-structured, in general. We impose two restrictions in order to obtain well-structure. First, guards must be used positively: equivalently, a rule cannot be disabled because of the existence of some tree pattern. Second, we need a uniform bound on the depth of the trees obtained by rewriting. Indeed, we show that if the depth is not uniformly bounded, then TPRS can encode Turing machines. We show that TPRS that satisfy both conditions yield well-structured transition systems, and we show how to apply forward and backward analysis of well-structured systems for obtaining the decidability of pattern reachability as well as of the termination. On the negative side, we show that exact reachability and the finite state property are undecidable for such TPRS. In the decidable cases, we also show that the complexity is at least non elementary.

Related work. We review here other approaches where it is possible to decide interesting properties of transition systems based on unranked trees.

The systems considered in positive AXML [2] are *monotone*: a document is modified by adding subtrees at nodes labeled by service calls. In particular, positive AXML documents can only grow. In the mail order example, it would mean that with positive AXML, one cannot delete a product from the cart. Moreover, there is no deterministic description of the semantics of a service: a service call can create any tree granted that it satisfies some DTD. Such a system is always confluent, and one can decide whether, after some finite number of steps, the system will stabilize.

Guard AXML [3] is very similar to our model, service calls being based on tree pattern queries. The focus of [3] is to be able to cope with potentially unbounded data and to analyze the evolution of sets of documents satisfying a given schema. On an abstract level, this means that trees are labeled by symbols from an infinite alphabet and that tree patterns use data constraints. Moreover, [3] allows negated patterns in guards, which we have to disallow in order to obtain well-structured systems. However, for decidability, [3] needs to impose a severe restriction on the way services are called, which is non-recursiveness. This restriction would mean in our setting that there is a fixed bound on the length of the rewriting chains. So the decidable fragment of Guard AXML is incomparable with our framework.

At last, ground tree rewriting rules on unranked, ordered trees were considered in [10]. In this setting a rule can only be applied at the deepest levels of the tree, and the rewriting is done either as subtree rewriting or as prefix rewriting on a sequence of siblings. For such systems it is shown that reachability is decidable. The rewriting used in [10] is much more restrictive than ours. In the mail order example, one could add and delete products inside an order, but not directly refer to an order as it could contain an unbounded number of products (hence deleting the whole order in one step is not possible).

2 Tree Pattern Rewriting

The tree rewriting model presented in this section is inspired by the Active XML (AXML) system developed at INRIA [4]. Active XML extends the framework of XML to describe semi-structured data by a dynamic component, describing data implicitly via service (function) calls. The evaluation of such calls is based on queries, resulting in extra data that can be added to the document tree. The abstract model is that of XML, i.e., unranked, unordered, labeled trees, together with a specification of the semantics for each service.

Trees considered in this paper are labeled by tags from a finite set \mathcal{T} . We will distinguish a subset $\mathcal{T}_{var} \subseteq \mathcal{T}$ of so-called *tag variables*. In addition, we use the special symbol $\$$ to mark nodes where service calls insert new data. Trees are in the following unranked and unordered, with nodes labeled by $\mathcal{T} \cup \mathcal{T}_{\$}$, where $\mathcal{T}_{\$} = \mathcal{T} \times \{\$\}$. We will not distinguish function/service nodes, since we consider here an abstract model for AXML documents, that is based on tree rewriting. We also do not consider multiple peers actually: their joint behavior can be described as the evolution of a unique document tree.

A *tree* $(V, \text{parent}, \text{root}, \lambda)$ consists of a set of nodes V with a distinguished node called *root*, a mapping $\text{parent} : V \setminus \{\text{root}\} \rightarrow V$ associating a node with its parent, and a mapping $\lambda : V \rightarrow \mathcal{T} \cup \mathcal{T}_{\$}$ labeling each node by a tag. Moreover, for each node $v \in V$, there is some $k \geq 0$ such that $\text{parent}^k(v) = \text{root}$. Such a tree is called a *document* if its labeling satisfies $\lambda(V) \subseteq \mathcal{T} \setminus \mathcal{T}_{var}$, that is, if neither

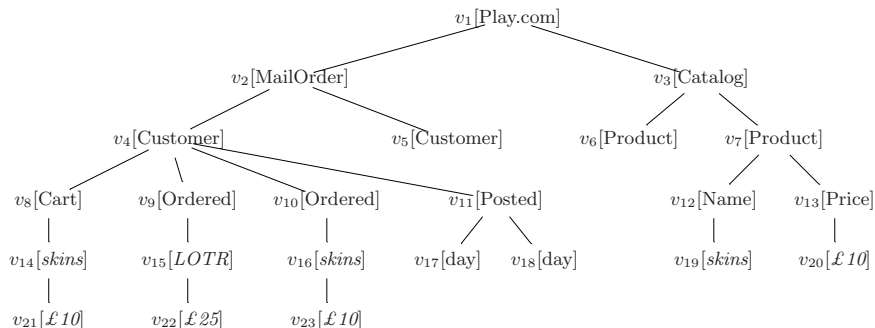


Fig. 1. Tree document representing a catalog of products and customers history.

tag variables nor the \$ sign are used as node labels. A *forest* is a finite multiset of trees.

Consider as an example the tree in Fig. 1. Informally, this tree represents a simplified version of the `Play.com` database, containing several products and information about customers. The subtrees of nodes v_5 and v_6 are not represented in the figure. The document shows two customers, one of which is currently shopping on the website with one product in her cart. This customer has 3 outstanding orders, one of which was posted 2 days ago (the counter *days* is encoded in unary in the tree).

One represents a tree in a term-like way. For example, to denote the empty catalog with no customer, we write $v_1[\text{Play.com}](v_2[\text{MailOrder}], v_3[\text{Catalog}])$, or if node names are irrelevant, $[\text{Play.com}]([\text{MailOrder}], [\text{Catalog}])$. Since trees are unordered, a tree can have several such representations.

The atomic operations in our model are from a set \mathcal{R} of guarded tree rewriting rules, as described below. On an abstract level we view a service s as described for instance by a regular expression $R(s)$ over the set of rewriting rules \mathcal{R} . For example, the following expression describes an order service on `play.com`:

$$(\text{add-product} + \text{delete-product})^* \text{checkout}$$

The tree resulting in the invocation of service s corresponds to the application of some sequence of rewriting rules in $R(s)$. The atomic rewriting rules will use queries based on tree patterns, as described next.

Definition 1 (Tree-Pattern). A tree pattern (TP for short) is a tuple $P = (V, \text{parent}, \text{ancestor}, \text{root}, \lambda)$, where $(V, \text{parent} \uplus \text{ancestor}, \text{root}, \lambda)$ is a tree.

A tree pattern represents a set of trees that have a similar shape. As for trees, a TP can be described in a term-like way, *ancestor*-edges being represented by the symbol $-$ (such edges are represented by a double line in the figures). For instance, the tree pattern `LQBill` shown in Fig. 2 can be written as $w_1(-w_2(w_3(w_4)))$, with $\lambda(w_1) = \text{Play.com}$, $\lambda(w_2) = \text{Ordered}$, $\lambda(w_3) = X$, $\lambda(w_4) = Y$ (here X and Y are tag variables: $X, Y \in \mathcal{T}_{\text{var}}$). This pattern represents trees with root “Play.com”, having a node “Ordered”, having itself a grandchild.

Definition 2 (Matching). A tree $T = (V, \text{parent}, \lambda, \text{root})$ matches a TP $P = (V', \text{parent}', \text{ancestor}', \lambda', \text{root}')$ if there exist two mappings $f : V' \rightarrow V$ and $t : \mathcal{T}_{\text{var}} \rightarrow \mathcal{T} \setminus \mathcal{T}_{\text{var}}$ such that:

- $f(\text{root}') = \text{root}$,
- For all $v \in V'$, $\lambda(f(v)) = \lambda'(v)$ if $\lambda'(v) \notin \mathcal{T}_{\text{var}}$, and else $\lambda(f(v)) = t(\lambda'(v))$,
- For all $v \in V'$ the following holds
 - If $\text{parent}'(v)$ is defined, then $f(\text{parent}'(v)) = \text{parent}(f(v))$;
 - If $\text{ancestor}'(v)$ is defined, then $f(\text{ancestor}'(v))$ is an ancestor of $f(v)$ in T .

Mappings (f, t) as above are called a matching between T and P . Furthermore, if f is injective, then (f, t) is called an injective matching.

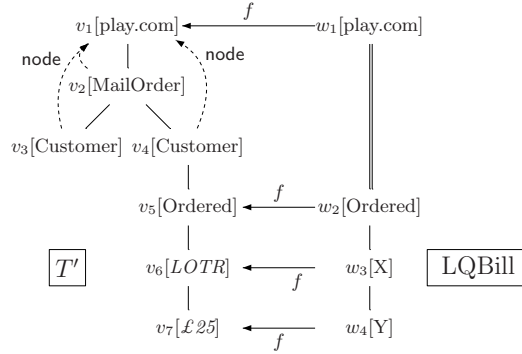


Fig. 2. A tree T' matching the TP LQBill.

Fig. 2 shows an example of an injective matching between a tree T' and the TP LQBill. The only possible matching is $f(w_1) = v_1, f(w_2) = v_5, f(w_3) = v_6, f(w_4) = v_7$ and $t(X) = LOTR, t(Y) = £25$. With a matching $f : V' \rightarrow V$ we associate the mapping $\text{node} : V \setminus f(V') \rightarrow f(V')$ with $\text{node}(v)$ being the lowest ancestor of v in $f(V')$. For instance, for the mapping f matching T to P , $\text{node}(v_2) = \text{node}(v_3) = \text{node}(v_4) = v_1$.

Similarly to [2] we use in our model tree-pattern queries (TPQ for short, also called *positive queries* in [2]). Such queries have the form $Q \rightsquigarrow P$, with Q a TP and P a tree, and the variables used in Q are also used in P . The TP Q selects tags in the tree. The result of a *query* $= Q \rightsquigarrow P$ on T is the forest $\text{query}(T)$ of all instantiations of P by matchings between T and Q . That is, for each matching (f, t) between T and Q we obtain an instance of P in which each \mathcal{T}_{var} -label X has been replaced by the tag $t(X)$. For instance, let RQBill be the tree $\text{Product}(\text{Name}(X), \text{Price}(Y))$. Then the result of the TPQ $(\text{LQBill} \rightsquigarrow \text{RQBill})$ on the tree in Fig. 1 is the forest depicted in Fig. 3.




Fig. 3. Forest produced as the result of the TPQ $\text{LQBill} \rightsquigarrow \text{RQBill}$ on document T .

We now define a generic kind of (guarded) rewriting rules, as a model for active documents. Our rules are based on tree patterns, that occur in two distinct contexts. First, tree patterns are used for describing how the structure of the document tree changes through the rule - some subtrees might be deleted, new nodes can be added. Second, rules are guarded, and the guard condition is tested

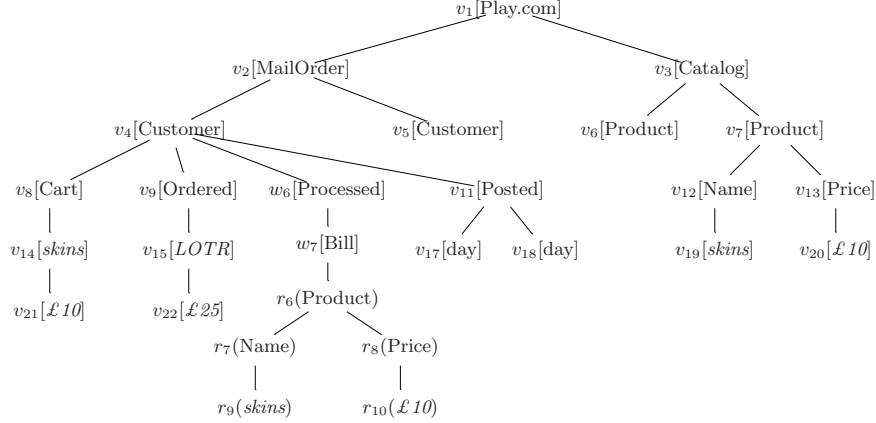


Fig. 5. The tree document T' resulting of the application of the rule Bill.

For instance, in the TP rule Bill, the TP LQBill uses names w_1, w_2 common to left, hence the results of the TPQ (LQBill \rightsquigarrow RQBill) are restricted to the particular order chosen by the matching f between T and left. The result is thus the subtree rooted at node r_6 in Fig. 3, but not the subtree rooted at r_1 , since it would require $f(w_2) = v_9$, while $f(w_2) = v_{10}$. This restriction is desirable, since we want to process a bill only for the products of this particular order. Here, w_7 is $\$$ -marked, and the result forest is defined by nodes r_6, \dots, r_{10} .

More formally, let $\text{query} = Q \rightsquigarrow P$ be a TPQ and let (f, t) be an injective matching between T and left. Moreover, let V_l be the nodes of left and V_Q those of Q . Let S_1, \dots, S_k be the trees composing the resulting forest $\text{query}(T)$, and let g_1, \dots, g_k be the respective associated matchings (that is, $g_i : V_Q \rightarrow V$ and S_i is the instantiation of P by g_i). Then we define $\text{query}_f(T)$ as the forest S_{i_1}, \dots, S_{i_l} of those trees S_j such that g_j agrees with f over $V_l \cap V_Q$. That is, $\text{query}_f(T)$ is the subset of $\text{query}(T)$ that is consistent with the matching f . We now turn to the formal semantics of rules.

Definition 4 (Semantics of rules). Let $T = (V, \text{parent}, \lambda, \text{root})$ be a tree and $R = (\text{left}, \text{query} = Q \rightsquigarrow P, \text{guard}, \text{right})$ be a rule, where $\text{left} = (V_l, \text{parent}_l, \text{ancestor}_l, \lambda_l, \text{root}_l)$ and $\text{right} = (V_r, \text{parent}_r, \text{ancestor}_r, \lambda_r, \text{root}_r)$.

We say that R is enabled if there exists an injective matching (f, t) from left into T , such that $\text{query}_f(T) \in \text{guard}$. The result of the application of R via (f, t) is the tree $T' = (V', \text{parent}', \lambda', \text{root}')$ defined as follows:

- $V' = V_1 \uplus V_2 \uplus V_3 \uplus V_4$ with
 1. $V_1 = f(V_l \cap V_r)$, % in the example on Fig. 5, $V_1 = \{v_1, v_4\}$
 2. $V_2 = \text{node}^{-1}(V_1)$, % in the example on Fig. 5, $V_2 = \{v_i \mid i \notin \{1, 4\}\}$
 3. $V_3 = f(V_r \setminus V_l)$, % in the example on Fig. 5, $V_3 = \{w_6, w_7\}$
 4. V_4 consists of distinct copies of the nodes of the forest $\text{query}_f(T)$, one for each node marked by $\$$ in right.
% in the example on Fig. 5, $V_4 = \{r_i \mid i \in \{6, \dots, 10\}\}$
- Setting $f(u) = u$ for all $u \in V_3$, we extend $f : V_r \cup V_l \rightarrow V_1 \uplus V_3$.

- $root' = f(root_r)$.
- Let $u \in V_1$ and let $\bar{u} = f^{-1}(u)$ be the associated node in $V_l \cap V_r$. If $\lambda_r(\bar{u}) \notin \mathcal{T}_{var}$ then $\lambda'(u) = \lambda_r(\bar{u})$, else $\lambda'(u) = t(\lambda_r(\bar{u}))$. For all $u \neq root'$, if $parent_r(\bar{u})$ is defined, then $parent'(u) = f(parent_r(\bar{u}))$ else $parent'(u) = parent(u)$.
- For all $u \in V_2$, $parent'(u) = parent(u)$ and $\lambda'(u) = \lambda(u)$.
- For all $u \in V_3 \setminus \{root_r\}$, $parent'(u) = f(parent_r(u))$ and $\lambda'(u) = \lambda_r(u)$.
- To each node $u \in V'$ marked by \$, we add a copy of the forest $query_f(T)$ as children of u , and we unmark the node.

Note that if $x \in V_2$, then its parent is in $V_1 \cup V_2$. The same stands if $ancestor_r(x)$ is defined. Note also that we indeed obtain a tree: for instance, if $u \in V_1$ and $parent_r(\bar{u})$ is not defined, then $parent(u)$ is defined. This is because $v = ancestor_r(\bar{u})$ is then defined, so by Def. 3, $v = ancestor_l(\bar{u})$ in left, so that $u = f(\bar{u})$ cannot be the root of T .

We write $T \xrightarrow{R} T'$ if T' can be obtained from T by applying the rule R . More generally, given a set of rules \mathcal{R} we write $T \rightarrow T'$ if there is some rule $R \in \mathcal{R}$ with $T \xrightarrow{R} T'$, and $T \xrightarrow{*} T'$ for the reflexive-transitive closure of the previous relation. Notice that the tree T' matches right, through the matching $f' : V_r \rightarrow V'$ defined by $f'(v) = f(v)$ if $v \in V_r \cap V_l$, and $f'(v) = v$ if $v \in V_r \setminus V_l$.

Remark 1. Positive AXML systems [2] are covered by our model except for one detail: in [2], positive queries can address not only the document tree, but also specific subtrees, such as the subtree of a service call, or the subtree of its parent. However, such queries on subtrees can be expressed easily in our formalism. An AXML service call corresponds to a rule (left, query, guard, right) with left = right = $r(-(v(w)))$, where w is labeled by a function tag. In addition, node v in right is marked by \$ (since it collects the result of query as new subtrees). A positive AXML TP P that is matched against the subtree of w can be rewritten into a TP $root(-P')$, where P' is obtained from P by renaming the root node into w .

Example (Play.com rules). To show how easily rules can be defined, we describe now the other rules of the Play.com system. When the rule does not use a query or a guard, we only describe the left and right components.

- The rule New-Customer adds a new customer and its cart.
 - left = $w_1[Play.com](w_2[MailOrder])$
 - right = $w_1[Play.com](w_2[MailOrder](w_3[Customer](w_4[Cart])))$.
- The rule Delete-Product deletes a product from a cart.
 - left = $w_1[Play.com](-w_7[Cart](w_8[X](w_9[Y])))$
 - right = $w_1[Play.com](-w_7[Cart])$.
- The rule Add-Product adds a new product to the cart of a customer.
 - left = $w_1[Play.com](-w_2[Product](w_3[Name](w_4[X]), w_5[Price](w_6[Y])), -w_7[Cart])$
 - right = $w_1[Play.com](-w_2[Product](w_3[Name](w_4[X]), w_5[Price](w_6[Y])), -w_7[Cart](w_8[X](w_9[Y])))$.

- The rule Checkout can turn the cart into an order.
 - left = $w_1[\text{Play.com}](-w_2[\text{Cart}])$
 - right = $w_1[\text{Play.com}](-w_2[\text{Ordered}])$
- As long as the bill is not processed, the customer can cancel her order with the rule Cancel.
 - left = $w_1[\text{Play.com}](-w_2[\text{Ordered}])$
 - right = $w_1[\text{Play.com}]$
- A processed order can be posted, and we count in unary the number of days since it was sent using *day*.
 - left = $w_1[\text{Play.com}](-w_2[\text{Processed}])$
 - right = $w_1[\text{Play.com}](-w_2[\text{Posted}](w_3[\text{day}]))$
- A posted parcel can be received.
 - left = $w_1[\text{Play.com}](-w_2[\text{Posted}])$
 - right = $w_1[\text{Play.com}](-w_2[\text{Received}])$
- Every new day, if a posted parcel has not yet been received yet, then the *day* counter is incremented.
 - left = $w_1[\text{Play.com}](-w_2[\text{Posted}])$
 - right = $w_1[\text{Play.com}](-w_2[\text{Posted}](w_3[\text{day}]))$
- If after 21 days a posted parcel is still not received, the customer can require a payback. We use the guard to ensure this time limit. Notice that the query is $Q \rightsquigarrow P$, where Q uses the same w_2 as in left, that is the number of days will be counted only for this particular parcel.
 - left = $w_1[\text{Play.com}](-w_2[\text{Posted}])$
 - $Q = w_1[\text{Play.com}](-w_2[\text{Posted}](w_3[\text{day}]))$
 - $P = x_4[\text{day}]$
 - guard: a forest containing at least 21 trees (and possibly more nodes) whose root is labeled *day*.
 - right = $w_1[\text{Play.com}]$.

Notice that we could model a stock by using tokens for each product in stock below its name in the catalog, decrementing the token when a product is added to cart and incrementing the token when it is deleted from cart. Also, a query could cancel a particular provisioning from a manufacturer from some product when the stock is greater than some threshold.

3 Static Analysis of TPRS

We assume from now on that an active document is given by a *tree pattern rewriting systems (TPRS)* (T, \mathcal{R}) , consisting of a set \mathcal{R} of TP rules and a \mathcal{T} -labeled tree T . That is, we assume that each service corresponds to a rule. Our results are easily seen to hold in the more general setting where services are regular expressions over \mathcal{R} .

A tree T with node set V is *subsumed* by a tree T' with node set V' , noted $T \preceq T'$, if there is an injective mapping from V to V' that preserves the labeling, the root, and the parent relation. A forest F is *subsumed* by a forest F' , written $F \preceq F'$, if F is mapped injectively into F' such that each tree in F is subsumed

by its image in F' . Similarly, a TP P with node set V is *subsumed* by a TP P' with node set V' , if there is an injective mapping from V to V' that preserves the labeling, the root, the **parent** and the **ancestor** relations.

With a TPRS (T, \mathcal{R}) we can associate the (infinite-state) transition system $\langle S(T, \mathcal{R}), \rightarrow \rangle$ with $S(T, \mathcal{R}) = \{T' \mid T \xrightarrow{*} T'\}$. We are interested in checking the following properties:

- Termination: Are all derivation chains $T \rightarrow T_1 \rightarrow T_2 \rightarrow \dots$ of (T, \mathcal{R}) finite?
- Finite-state property: Is the set $S(T, \mathcal{R})$ of reachable trees finite?
- Reachability: Given (T, \mathcal{R}) and a tree T' , is T' reachable in (T, \mathcal{R}) ?
- Confluence: For any pair of trees $T_1, T_2 \in S(T, \mathcal{R})$, does there exist some T' such that $T_1 \xrightarrow{*} T'$ and $T_2 \xrightarrow{*} T'$?
- Pattern reachability (coverability): Given (T, \mathcal{R}) and a tree pattern P , does $T \xrightarrow{*} T'$ hold for some T' matching P ?
- Weak confluence: From any pair of trees $T_1, T_2 \in S(T, \mathcal{R})$, does there exist $T'_1 \preceq T'_2$ such that $T_1 \xrightarrow{*} T'_1$ and $T_2 \xrightarrow{*} T'_2$?

For instance, pattern reachability is a key property when we want to talk about the reachability of some pattern. For example, we might ask whether an already cancelled order could still be delivered, which would mean a problem in the system. For this, it suffices to tag cancelled orders with a special symbol $\#$, and check for the pattern $w_1[\text{Play.com}](-w_2[\text{delivered}](w_3[\#]))$. This is the same kind of properties which are checked in [3]. As expected, any of the non-trivial questions above is undecidable in the general case, see Theorem 1 below.

We are thus looking for relevant restrictions yielding decidability of at least some of these problems.

In the next section we consider a subclass of TPRS, which is a special instance of the so-called *well-structured* systems. We say that (T, \mathcal{R}) is *positive* if all guards occurring in the rules from \mathcal{R} , are upward-closed. This means that for every guard G , and all forests F, F' with $F \preceq F'$, $F \in G$ implies $F' \in G$, too. In particular, if a rule R in a positive system is enabled for a tree T , then R is enabled for any tree T' that subsumes T . The reason is that for any TPQ query, we have that for every tree T'_1 in $\text{query}(T')$, there is some tree T_1 in $\text{query}(T)$ such that T_1 is subsumed by T'_1 . Notice that positive TPRS allow deletion of nodes, so they are more powerful than the positive AXML systems considered in [2].

The next theorem shows that upward-closed guards alone do not suffice for obtaining decidability of termination:

Theorem 1. *Any two-counter machine M can be simulated by a positive TPRS (T, \mathcal{R}) in such a way that M terminates iff (T, \mathcal{R}) terminates.*

Proof. Let M be a two-counter machine with control state set Q and counters c and d . In the sequel, a configuration of M , i.e. a tuple $(q, x, y) \in Q \times \mathbb{N} \times \mathbb{N}$ where x and y respectively denote the value of counter c and d , will be represented by a tree whose root is labeled by q , which has two children, respectively labeled

by c and d . The tree rooted at the node labeled c (respectively d) is formed of a single path containing x (respectively y) intermediate nodes labeled by i (as internal) and ending by a leaf labeled \perp . We denote the previous tree by $T_{(q,x,y)}$.

Without loss of generality we may assume that any transition of M requires the value of each counter to be either zero or non-zero (this can be enforced by inserting a zero test before each transition).

Each possible transition of M is modeled by a rewriting rule. For instance a transition enabled in state q with c being non-zero and d being zero that goes to state q' , decrementing the counter c and incrementing the counter d , is simulated by the rule with $\text{left} = \text{root}(u(w(t)), v(t'))$ and $\text{right} = \text{root}'(u(t), v(z(t')))$, where root is labeled by q , root' is labeled by q' , u is labeled by c , v is labeled by d , w and z are labeled by i , t is labeled by $*$ and t' is labeled by \perp . One defines similar rules for other possible transitions in M . Note here that our rules have no query/guard part, hence the system is positive.

It follows directly from the definition of the rules that a configuration (q', x', y') is reachable in M from some configuration (q, x, y) iff $T_{(q,x,y)} \xrightarrow{*} T_{(q',x',y')}$. \square

Remark 2. The TPRS defined in the previous proof uses rules that can erase nodes (and their subtrees). Technically we could avoid deletions by adding a special garbage node to which we attach all nodes of left that do not occur in right.

Theorem 1 shows that any non trivial property is undecidable for positive TPRS without further restrictions. However, notice that the above proof needs trees of unbounded depth. A realistic restriction in the XML setting is to consider only trees of bounded depth: XML documents are usually large, but shallow. A TPRS (T, \mathcal{R}) is called *depth-bounded*, if there exists some fixed integer K such that every tree T' with $T \xrightarrow{*} T'$ has depth at most K . Of course, Theorem 1 implies that it is undecidable to know whether a TPRS is depth-bounded. However, in many real-life examples this property is easily seen to hold (see *e.g.* the Play.com example, which has depth at most 8).

4 Decidability for positive and depth-bounded TPRS

For positive and depth-bounded TPRS, we can apply well-known techniques from the verification of infinite-state systems that are *well-structured*. Well-structured transition systems were considered independently by [1,8] and they cover many interesting models, such as Petri nets or lossy channel systems. We recall first some basics of well-structured systems.

Definition 5. A *well-quasi-order (wqo)* on a set X is a preorder \preceq such that in every infinite sequence $(x_n)_{n \geq 0} \subseteq X$, there exist some indices $i < j$ with $x_i \preceq x_j$.

In general, the “subsumed” relation \preceq on the set X of \mathcal{T} -labeled trees is not a wqo.¹ However, using Higman’s lemma (see, *e.g.*, [6, Chap. 12]), one can show that \preceq is a wqo on the set of trees of depth at most K (for any fixed K):

Proposition 1. *Fix $K \in \mathbb{N}$, and let X_K denote the set of unordered \mathcal{T} -labeled trees of depth at most K . The “subsumed” relation $\preceq \subseteq X_K \times X_K$ is a wqo.*

Proof. It is easy to see that the relation \preceq is a preorder. We show the second property by induction on $K \geq 0$. Let $K > 0$ and consider an infinite sequence of trees T_1, T_2, \dots of depth at most K . We find a subsequence T_{i_1}, T_{i_2}, \dots of trees with the same root label a . Each such tree T_{i_j} consists of the root a , plus a multiset S_{i_j} of trees of depth at most $K - 1$. By induction, (X_{K-1}, \preceq) is a wqo. Higman’s lemma states that whenever (X, \preceq) is a wqo, we have that (X^*, \preceq) is a wqo, too. In particular, the set of multi-sets over X together with \preceq , is a wqo. Thus, we can find some $j < k$ such that $S_{i_j} \preceq S_{i_k}$, thus $T_{i_j} \preceq T_{i_k}$. \square

By the previous statement, a positive and depth-bounded TPRS (T, \mathcal{R}) yields a well-structured transition system $\langle S(T, \mathcal{R}), \rightarrow \rangle$ as defined in [8] (see also² [1]). This follows from the transition relation \rightarrow being upward compatible: whenever $T \xrightarrow{R} T'$ and $T \preceq T_1$, there exists T'_1 with $T_1 \xrightarrow{R} T'_1$ and $T' \preceq T'_1$.

For the next theorem we need first some notation. Given a set X and a preorder \preceq , we denote by $\uparrow X$ the upward closure $\{T' \mid T \preceq T' \text{ for some } T \in X\}$ of X . By $\min(X)$ we denote the set of minimal elements³ of X . Finally, by $\text{Pred}(X)$ we denote the set of immediate predecessors of elements of X . Note that whenever the transition relation \rightarrow is upward compatible and X upward-closed, the set $\text{Pred}(X)$ is upward-closed, too.

Since the subsumed relation \preceq is a wqo, the \preceq relation on forests is a wqo as well. Thus, each guard G in a positive, depth-bounded TPRS (T, \mathcal{R}) can be described by the (finite) set of forests $\min(G)$. The size $|G|$ of G is the maximal size of a forest in $\min(G)$.

Theorem 2. *Termination and pattern reachability are both decidable for positive and depth-bounded TPRS.*

Proof. First, termination is decidable for well-structured systems such that 1) \preceq is decidable, 2) \rightarrow is computable and 3) upward compatible, see [8, Thm. 4.6].

For pattern reachability, it is easy to see that the set of trees of depth bounded by some K which matches a TP P is upward-closed, and that the set of its minimal elements is effectively computable. We can thus use [1], which shows decidability of the reachability of $\uparrow T$ under the assumption that the set $\min(\text{Pred}(\uparrow T))$

¹ Indeed consider the sequence of trees $(T_n)_{n \geq 0}$ where for each $n \geq 0$, T_n is the tree formed by a single branch of length $n + 1$ whose internal nodes are labeled by a and the unique leaf is labeled by b .

² As shown in the proof of Thm. 2, $\langle S(T, \mathcal{R}), \rightarrow \rangle$ is also well-structured as defined in [1], which requires in addition that the set of predecessors of upward-closed sets is effectively computable.

³ For a wqo (X, \preceq) and $Y \subseteq X$, the set $\min(Y)/\sim$ is finite, where $\sim = \preceq \cap \preceq^{-1}$. For the subsumed relation \preceq , note that \sim is the identity.

is computable. This allows to use the obvious backward exploration algorithm. From now on, we fix the tree T and the bound K of the system (T_0, \mathcal{R}) . Let us also fix a rule $R = (\text{left}, \text{query}, \text{guard}, \text{right})$.

It suffices to consider the finite set $\mathcal{S}_R(T)$ of all trees T' of size at most $|T| + |\text{left}| + K|\text{query}||\text{guard}|$ with $T' \xrightarrow{R} T$. Then, defining M as the set of minimal elements of $\bigcup_{R \in \mathcal{R}} \mathcal{S}_R(T)$, we get by definition $M \subseteq \min(\text{Pred}(\uparrow T))$.

Let us now show that $M = \min(\text{Pred}(\uparrow T))$. Let $T_1 \in \min(\text{Pred}(\uparrow T))$. Thus, there exist some rule R and some injective matching (f, t) with $T_1 \xrightarrow{R} T$ via (f, t) . Let also $F \in \min(\text{guard})$ be a forest with $F \preceq F'$, where F' is the result of query on T_1 (compatible with the matching f).

The nodes of T_1 can be partitioned into 4 sets V_1, V_2, V'_3, V'_2 (below V_l are the nodes of left and V_r those of right):

1. $V_1 = f(V_l \cap V_r)$,
2. $V_2 = \text{node}^{-1}(V_1)$,
3. $V'_3 = f(V_l \setminus V_r)$,
4. $V'_2 = \text{node}^{-1}(V'_3)$.

Notice that V_1 and V_2 are common with T (see Def. 4), hence $|T_1| \leq |T| + |V'_3| + |V'_2|$. Now, $|V'_3| \leq |\text{left}|$. We now explain that V'_2 has at most $|\text{query}||\text{guard}|$ leaves, hence $|V'_2| \leq K|\text{query}||\text{guard}|$ which shows that $T_1 \in \mathcal{S}_R(T)$. Otherwise one can delete a leaf from V'_2 and get a tree $T'_1 \preceq T_1$ with $T'_1 \xrightarrow{R} T$ (via (f, t)), and still $F \preceq F''$, where F'' denotes the result of query on T'_1 . This contradicts the minimality of T_1 . \square

On the negative side, depth-bounded well-structured systems can simulate reset Petri nets (i.e., nets with an additional transition that empties a place), hence we can deduce the following from known results:

Theorem 3. *Exact reachability, confluence, weak confluence and the finite-state property are undecidable for positive and depth-bounded TPRS.*

Proof. We encode a reset Petri net with n places $(p_i)_{i \leq n}$ by a well-structured system with $K = 2$. Namely, we consider trees with root labeled r and with n children labeled p_i . Each node p_i has as many children as p_i has tokens and the leaves are all labeled by t .

Each transition is modeled by a rule. For instance, a reset of place p_i is modeled by the rule with $\text{left} = \text{root}(v)$, $\text{right} = \text{root}(w)$, where root is labeled r and both v and w labeled by p_i (the rule has an empty query/guard part). Applying such a rule preserves the root, but the subtree whose root is labeled by p_i is destroyed, and a new child of root is created with the label p_i , and this node has no child. A rule which takes 2 tokens from place p_i and one from place p_j and creates a token in place p_k is modeled as $\text{left} = \text{root}(v(w, x), y(z), s)$, $\text{right} = \text{root}(v, y, s(u))$, with root, v, y, s labeled respectively by r, p_i, p_j, p_k , and w, x, z, u are labeled by t .

It directly follows from the construction that the previous TPRS system is depth-bounded and positive. As both exact reachability and the finite-state

property are undecidable for reset Petri nets [7], it follows that this holds for positive and depth-bounded TPRS as well.

We now show the undecidability of (weak) confluence for a reset Petri net, which hence implies the undecidability for positive and depth-bounded TPRS. For this, we reduce reachability to the confluence of a reset Petri net.

The question is whether a marking M_f can be reached in a reset Petri net P from the initial marking. We create a new reset Petri net P' with the same places as P , plus two new places q, r . The initial marking of P' is the same as for P , and there is one token in place q and zero in place r .

For two markings M, M' , we denote by (M, M') the transition with $M(a)$ tokens taken from place a , and $M'(a)$ tokens added to place a , for all a . For instance, the transition $(\{p, 2 \cdot q\}, \{p, s, t\})$ checks that there are at least one token in place p and two in q , it deletes two tokens from place q , and adds one to both s and t .

For each transition (M, M') of P , there is a transition $(M + \{q\}, M' + \{q\})$ in P' . We fix a place p of P . We add the three transitions $(M_f + \{q\}, \{r\})$, $(\{q\}, \{r, p\})$ and $(\{r, 2 \cdot p\}, \{r, p\})$, plus a transition $(\{r, l\}, \{r, p\})$ for all places $l \notin \{p, q, r\}$.

First, note that in each reachable marking of P' , there is either a token in q or in r (but not both), and no marking with a token in q is reachable from a marking with a token in r . Second, every marking except $\{r\}$ can reach $\{r, p\}$. Indeed, if the marking has a token in q , then it can first execute the transition $(\{q\}, \{r, p\})$. With a token in r , each place other than r, p can be emptied with $(\{r, l\}, \{r, p\})$ (and p is sure to be non-empty). Then all tokens in p but one can be deleted with $(\{r, 2 \cdot p\}, \{r, p\})$. Also, notice that $\{r\}$ is reachable in P' iff M_f is reachable in P . The reason is that from any marking M with $M(r) = 1$ and $M \neq \{r\}$ we cannot reach $\{r\}$.

Therefore P' is confluent iff $\{r\}$ is not reachable in P' iff M_f is not reachable in P . As P' is weakly confluent iff it is confluent, this concludes the proof. \square

On the positive side, we can show that the finite-state property is decidable for positive, depth-bounded TPRS, that are *strict*, i.e., such that for any rule (left, query, guard, right), we require $V_{\text{left}} \subseteq V_{\text{right}}$. One cannot encode reset Petri nets with such systems because deletion is no longer possible (actually one can only relabel an existing node and create new nodes). Strict systems enjoy the additional property that whenever $T \xrightarrow{R} T'$ and $T \prec T_1$, there exists T'_1 with $T_1 \xrightarrow{R} T'_1$ and $T' \prec T'_1$ (notice that for non strict systems, we can only guarantee that $T' \preceq T'_1$). The results from [8] yield the following theorem.

Theorem 4. *The finite-state property and reachability are decidable for TPRS that are positive, depth-bounded, and strict.*

Note that the finite-state property is not very interesting in itself, but if it holds, then the other problems become decidable as we are dealing with a finite-state system. In particular, in order to test for confluence, it suffices to test that $(S(T, \mathcal{R}), \rightarrow)$ has a unique maximal strongly connected component.

Note that reachability is decidable for positive, depth-bounded and strict TPRS simply because $T \rightarrow T'$ implies that $|T| \leq |T'|$, and then it suffices to look for reachability of a tree T_1 in the finite state system $\{T' \mid |T'| \leq |T_1|\}$.

The undecidability of confluence and weak confluence for positive, depth-bounded and strict TPRS follows from Theorem 3 together with a simple technical modification: instead of deleting nodes (and their subtrees) we attach them to a special garbage node $\#$. Additional rewriting rules ensure that at any moment, arbitrary descendants of the garbage node can be created. These rules ensure confluence, since for any subtrees T_1, T_2 of $\#$ (of bounded depth) we can rewrite both T_1, T_2 into some (larger) tree T_3 .

Below is a table that sums up the results we obtained so far. It presents (un)decidability results concerning the various classes of positive TPRS we considered (depth-bounded and strict). The negative results about strict TPRS come from Theorem 1 (see subsequent remark). Term., FS, Reach., P-reach, Confl. and W-confl. stand respectively for termination, finite state property, reachability, pattern reachability, confluence and weak confluence.

Model	Term.	FS	Reach.	P-reach.	Confl.	W-confl.
Strict	U	U	U	U	U	U
Depth-Bounded	D	U	U	D	U	U
Depth-Bounded & Strict	D	D	D	D	U	U

Table 1. Decidability results for positive TPRS.

5 Lower bounds and extensions

Decidability results are obtained with non-constructive proofs coming from Higman's Lemma. This ensures termination of the algorithms, but without yielding complexity bounds. It is thus relevant to obtain lower bounds for these results.

Theorem 5. *The following problems have at least non-elementary complexity:*

- *Input: A TP P , a TPRS system (S, \mathcal{R}) and an integer k such that the depth of (S, \mathcal{R}) is bounded by k .*
- *Problem1: Is the pattern P reachable in (S, \mathcal{R}) ?*
- *Problem2: Does (S, \mathcal{R}) terminate, equiv., does it have an infinite path?*

Proof. Let $\text{tower}(0, n) = n$ and $\text{tower}(k + 1, n) = 2^{\text{tower}(k, n)}$. Let M be an $n \mapsto \text{tower}(k, n)$ -space bounded deterministic Turing machine and x be an input of M . Denote by $\log^* n$ the smallest integer m such that $n \leq \text{tower}(m, 2)$ and let $K = k + \log^* |x|$, so that the computation of M on x uses at most $\text{tower}(k, |x|) \leq \text{tower}(K, 2)$ tape cells. We build a $(K + 1)$ -depth bounded TPRS of size $O(|M| + |x|)$ simulating M on x .

Informally, we encode each configuration of M by a tree. Each cell is encoded by a subtree of the root, labeled at its own root by the cell content, with the forest below it encoding the position of the cell. Since such a position is smaller than $\text{tower}(K, 2)$, it can itself be encoded recursively by a forest of depth at most K (such a recursive encoding of large integers, by words, has already been used in [13]). For instance, one can encode integers from 0 to $15 = \text{tower}(2, 2) - 1$ at depth 2 . The forest of Fig. 6 encodes 13 (**1101** in binary). To recover its position, each bit of the base 2 representation has under itself a forest of depth 1 encoding its position (recursively with the same encoding scheme). For instance, the leftmost **1** is at position 00, which is encoded by the forest $\{[0]([0]), [0]([1])\}$.




Fig. 6. A level 2 counter encoding 13.

Formally, a *level 0 counter* consists of a single node labeled by 0 or 1. We define inductively a *level ℓ counter* as a forest of $\text{tower}(\ell - 1, 2)$ trees, each of depth ℓ , to encode any integer between 0 and $\text{tower}(\ell, 2) - 1$. The level ℓ counter encoding $i \leq \text{tower}(\ell, 2) - 1$ will be noted F_i^ℓ (for $\ell = 0$: $F_0^0 = \{0\}$, $F_1^0 = \{1\}$). Assume inductively that $F_i^{\ell-1}$ has been defined for $0 \leq i \leq \text{tower}(\ell - 1, 2) - 1$. For $0 \leq j \leq \text{tower}(\ell, 2) - 1$, let $a_0 \cdots a_p$ be the binary representation of j , with $p = \text{tower}(\ell - 1, 2) - 1$. The level ℓ counter F_j^ℓ is the forest $\{[a_i](F_i^{\ell-1}) \mid 0 \leq i \leq p\}$.

Let $N = \text{tower}(K, 2) - 1$. We encode the configuration C of M with tape content $a_0 \cdots a_N$, current state q and scanned position m , by the forest $F_C = \bar{a}_0(F_0^K) \cdots \bar{a}_N(F_N^K)$ of depth $K + 1$, with $\bar{a}_m = [a_m, q]$ and $\bar{a}_i = [a_i]$ for $i \neq m$. The head position is thus doubly tagged: by the letter, and by the state. Such a node with a double tag $[\alpha, \beta]$ is said *marked by β* , or a *β -node*.

In order to navigate through the cells, we use for each level $\ell \leq K$ an additional placeholder node, child of the root, named c_ℓ for holding a level ℓ counter below it. The idea is that the counter attached below c_K will be able to count up to N , and hence can pinpoint a tape position. The other counters are needed in the inductive process. During the computation, additional markers will be used either as pebbles, or to guide the control. Figure 7 shows a typical tree reached during the computation. The rules of the TPRS are set up so that it performs successively the following actions:

1. It creates the forest F_{C_0} , and attaches it under the root, leaving c_K labeled by $[c_K, \text{run}]$ and for $\ell < K$, c_ℓ labeled by $[c_\ell, \text{end}]$.
2. It simulates repeatedly transitions of M , stopping if the final state is reached.

We only show how to encode transitions. The generation of the initial configuration, starting from the tree $[M]([c_0, \text{end}], \dots, [c_{K-1}, \text{end}], [c_K, \text{create-init-config}])$,




Fig. 7. The tree coding the tape $bbqba$ of the Turing machine M .

is done using similar routines. We use a finite set of rules without query/guard part. Although the TPRS will be nondeterministic, appropriate tags shall ensure that rules applicable at some step have all the same *left* member. When the TPRS discovers that a nondeterministic guess was wrong, it blocks. Therefore, if M halts on x , then the TPRS always terminate. If M does not halt on x , then the corresponding run of the TPRS where all guesses are correct does not either. This ensures termination iff M halts on x .

To simulate a transition, the TPRS first performs the changes in the configuration, nondeterministically guessing the new head position. To check whether the head has been properly placed, it $*$ -marks the original head position. The node c_K marked by $\{\text{run, check-suc1, check-suc2} \dots, \text{check-pred1, check-pred2} \dots\}$ (for instance, check-suc needs several steps) encodes the current stage of the simulation. For instance, to simulate a transition $p \xrightarrow{a/b/\rightarrow} q$, we use the rule:

- left = $r[M](x[a, p], y[X], z[c_K, \text{run}])$,
- right = $r[M](x[b, *], y[X, q], z[c_K, \text{check-suc1}])$,

slightly abusing notation: we use a double tag involving a variable, to abbreviate a finite set of rules (with the obvious interpretation).

To complete the simulation of the transition, the TPRS checks whether the position written below the node pinpointed by q is a successor of that below the node pinpointed by $*$. If yes, it deletes the mark $*$, and labels c_K back to $[c_K, \text{run}]$. If not, the head position was incorrectly guessed and the system blocks.

The steps of check-suc, which checks that the nodes marked by $*$ and q occur successively, are first a copy under c_K of the level K counter located below the $*$ -node, then an increment of that copy, and a comparison of the result to the counter below the q -node. We use auxiliary markers $*_\ell, *'_\ell$ for each level ℓ , attached to nodes below an ℓ counter: $*_\ell$ in the part of the tree representing the configuration, and $*'_\ell$ under some $c_i, i > \ell$. We define inductively rules to achieve the following tasks for each level $\ell \leq K$:

- copy(ℓ) copies below the $*'_\ell$ -marked node the level ℓ counters found below c_ℓ .
- increment(ℓ) increments the level ℓ counter below c_ℓ .
- compare(ℓ) compares level ℓ counters below c_ℓ and below the $*_\ell$ -marked node.
- test-max(ℓ) tests if the level ℓ counter below c_ℓ has its maximal value.
- zero(ℓ) generates under c_ℓ the level ℓ counter F_0^ℓ .

Each task of level ℓ will be implemented using a sequence of tasks of level $(\ell - 1)$, plus some fresh tags to correctly organize the order of these level $(\ell - 1)$ tasks.

Let us explain how to handle the control. First, each function of level ℓ only performs local tasks, or calls to functions of level $\ell - 1$. Therefore, the underlying call stack size is K -bounded. The first tasks to be called are those of level K (recall that we aimed at checking that the head was correctly placed, thus having to handle level K counters). The call stack is simulated by markers put at the c_i nodes. When the activated function is at level ℓ , nodes $c_0, \dots, c_{\ell-1}$ are all marked `ready`, `true` or `false`, while nodes c_ℓ, \dots, c_K have other marks indicating an “active” state. For instance at the beginning of the simulation, c_K is marked `run` while for $\ell < K$, c_ℓ is marked `ready`. This way, a left side of the form $r[M](u[c_{\ell-1}, \text{ready}], v[c_\ell, x])$ for some tag $x \neq \text{ready}$ selects the counter where the current action is to be performed.

As an example, `compare`(ℓ) works by running a level $\ell - 1$ counter from 0 to its maximal value, using `zero`($\ell - 1$), `increment`($\ell - 1$) and `test-max`($\ell - 1$). For each value k of this counter, it compares the k -th bits of the level ℓ counters under comparison, by nondeterministically marking one bit of each counter, verifying with `compare`($\ell - 1$) and `compare'`($\ell - 1$) that they are at position k , and if so, comparing them. The call `compare`(ℓ) is activated when c_ℓ is tagged `compare`. We use as tags either names of functions to be called (`compare`, `zero`, etc.), “return values” (`true`, `false`, `ready`), or control states (`cmp-guess-1st-bit`, ...). Notice that we use another operator `compare'`(ℓ) to compare bits marked $*'_\ell$ and the counter under c_ℓ . Formally:

- Initialize level $\ell - 1$ counter to 0, by marking $c_{\ell-1}$ by `zero`, to “call” `zero`($\ell - 1$):
 - `left` = $r[M](x[c_{\ell-1}, \text{ready}], y[c_\ell, \text{compare}])$
 - `right` = $r[M](x[c_{\ell-1}, \text{zero}], y[c_\ell, \text{cmp-guess-1st-bit}])$.
- Guess and mark by $*_{\ell-1}$ a bit below $*_\ell$. Check the guess with `compare`($\ell - 1$):
 - `left` = $r[M](x[c_{\ell-1}, \text{ready}], y[c_\ell, \text{cmp-guess-1st-bit}], -t[T, *_\ell](u[U]))$,
 - `right` = $r[M](x[c_{\ell-1}, \text{compare}], y[c_\ell, \text{cmp-guess-2nd-bit}], -t[T, *_\ell](u[U, *_{\ell-1}]))$.
- Guess and mark by $*'_{\ell-1}$ a bit below c_ℓ . Check the guess with `compare'`($\ell - 1$):
 - `left` = $r[M](x[c_{\ell-1}, \text{true}], y[c_\ell, \text{cmp-guess-2nd-bit}](z[Z]))$,
 - `right` = $r[M](x[c_{\ell-1}, \text{compare}'], y[c_\ell, \text{cmp-compare-bits}](z[Z, *'_{\ell-1}]))$.
- If the bits agree: clear the $*_\ell$ and $*'_{\ell-1}$ marks and test if their position was maximal. Notice that bit equality is tested by using the same variable twice:
 - `left` = $r[M](x[c_{\ell-1}, \text{true}], y[c_\ell, \text{cmp-compare-bits}](z[X, *'_{\ell-1}]), -u[X, *_{\ell-1}])$,
 - `right` = $r[M](x[c_{\ell-1}, \text{test-max}], y[c_\ell, \text{cmp-next-bits}](z[X]), -u[X])$.
- Return `true` if done:
 - `left` = $r[M](x[c_{\ell-1}, \text{true}], y[c_\ell, \text{cmp-next-bits}])$,
 - `right` = $r[M](x[c_{\ell-1}, \text{ready}], y[c_\ell, \text{true}])$.
- Proceed to the loop if not, incrementing the value of the level $\ell - 1$ counter:
 - `left` = $r[M](x[c_{\ell-1}, \text{false}], y[c_\ell, \text{cmp-next-bits}])$,
 - `right` = $r[M](x[c_{\ell-1}, \text{increment}], y[c_\ell, \text{cmp-guess-1st-bit}])$.

The rules for other tasks follow the same ideas and are not described here. \square

The bounded depth restriction needed for our decidability results can be relaxed if we forbid the use of the direct parent-child edges in tree patterns.

This leads to the following preorder on unranked, unordered \mathcal{T} -labeled trees, which is a well quasi-order by Kruskal's theorem (see [6, Chap. 12]). For two trees T, T' with sets of nodes V, V' respectively, we write $T \prec T'$, if there is an injective mapping from V to V' that preserves the labeling, the root, and the ancestor relation. So compared to the relation \preceq used previously, we do not require that the parent relation is preserved.

Clearly, we need to restrict the TPRS rules in order to obtain well-structured systems. Namely we require that all TP occurring in `query` and `left` use only `ancestor` edges (`right` can still use `parent` edges, but the `parent` relation cannot be tested for). We call such TPRS *undirected*. Using similar proofs as in Sect. 4, we get the same results as in Table 1. For the lower bound we obtain a stronger result, by encoding reachability for lossy channel systems (LCS). These are finite-state machines communicating over FIFO channels that can loose arbitrary many messages. Reachability for LCS has non primitive recursive complexity [12], already for LCSs made up of two finite-state machines and two channels [5].

Theorem 6. *Termination and pattern reachability have at least non primitive recursive complexity for undirected TPRS.*

Proof. We reduce the reachability of some global control state of an LCS to the pattern reachability of an undirected TPRS. Fix an LCS made up of two finite-state machines A_1, A_2 communicating through two lossy FIFO channels (C_1 from A_1 to A_2 and C_2 in the opposite direction).

A configuration of the LCS is given by (q_1, q_2, w_1, w_2) , where q_i is the current state of A_i and $w_i \in \Sigma^*$ is the content of channel C_i . We encode this configuration by the following tree, where for instance $w_1 = a \cdots c$. Channel tails, where writes occur, are marked by `[end]`.




Fig. 8. The tree encoding the configuration $(q_1, q_2, a \cdots c, b \cdots a)$ of an LCS.

The rules associated with such a system are given below, assuming disjoint state sets for A_1 and A_2 , that do not contain symbols $C1, C2$. Notice that we do not use any query/guard part. One can assume that message losses happen just before read actions, ahead of the message actually read in that channel.

- A message loss possibly occurs at the head of channel C_1 , a message a is then read in C_1 by A_2 , which switches its state from q_2 to q'_2 :
 - left = $w_1[LCS](-w_2[q_2], -w_3[C1](-w_4[a]))$,
 - right = $w_1[LCS](-w_2[q'_2], w_4[C1])$.
- A_1 performs a transition from state q_1 to q'_1 , sending message a into C_1 :
 - left = $w_1[LCS](-w_2[q_1], -w_3[C1](-w_4[end]))$,
 - right = $w_1[LCS](-w_2[q'_1], -w_3[C1](-w_4[a](w_5[end])))$.

Dual rules are defined for A_2 . It should be clear that there is a bisimulation between the undirected TPRS system and the LCS, which yields the result. \square

References

1. P. A. Abdulla, K. Cerans, B. Jonsson, and Y.-K. Tsay. General decidability theorems for infinite-state systems. In *LICS'96*, pp. 313–321. IEEE Comp. Soc., 1996.
2. S. Abiteboul, O. Benjelloun, and T. Milo. Positive Active XML. In *PODS'04*, pp. 35–45. ACM, 2004.
3. S. Abiteboul, L. Segoufin, and V. Vianu. Static Analysis of Active XML Services. In *PODS'08*. ACM, 2008. To appear.
4. Active XML. <http://www.activexml.net/>.
5. P. Chambart and Ph. Schnoebelen. The Ordinal Recursive Complexity of Lossy Channel Systems. In *LICS'08*. IEEE Comp. Soc., 2008. To appear.
6. R. Diestel. *Graph theory*. 2005. <http://www.math.uni-hamburg.de/home/diestel>.
7. C. Dufourd, A. Finkel and Ph. Schnoebelen. Reset Nets between Decidability and Undecidability. In *ICALP'98*, LNCS 1443, pp. 103–115. Springer, 1998.
8. A. Finkel and Ph. Schnoebelen. Well-structured transition systems everywhere! *Theor. Comput. Sci.*, 256(1-2):63–92, 2001.
9. L. Libkin. Logics for unranked trees: An overview. *Log. Meth. Comput. Sci.*, 2(3), 2006. Available at [http://dx.doi.org/10.2168/LMCS-2\(3:2\)2006](http://dx.doi.org/10.2168/LMCS-2(3:2)2006).
10. Ch. Löding and A. Spelten. Transition Graphs of Rewriting Systems over Unranked Trees. In *MFCS'07*, LNCS 4708, pp. 67–77. Springer, 2007.
11. F. Neven. Automata, Logic, and XML. In *CSL'02*, LNCS 2471, pp. 2–26. Springer, 2002.
12. Ph. Schnoebelen. Verifying Lossy Channel Systems has Nonprimitive Recursive Complexity. *Inf. Process. Lett.* 83(5):251–261, 2002.
13. I. Walukiewicz. Difficult Configurations—on the Complexity of LTrL. *Form. Methods Syst. Des.*, 26(1): 27–43. Kluwer, 2005. Short version in *ICALP'98*, LNCS 1443.