

Distributed XML Design

S. Abiteboul^a, G. Gottlob^b, M. Manna^c

^aINRIA Saclay – Île-de-France & University Paris Sud, FR

^bOxford University Computing Laboratory &

Oxford-Man Institute of Quantitative Finance, University of Oxford, UK

^cDepartment of Mathematics, University of Calabria, IT

Abstract

A *distributed XML document* is an XML document that spans several machines. We assume that a distribution design of the document tree is given, consisting of an *XML kernel-document* $T_{[f_1, \dots, f_n]}$ where some leaves are “docking points” for external resources providing XML subtrees (f_1, \dots, f_n , standing, e.g., for Web services or peers at remote locations). The top-down design problem consists in, given a *type* (a schema document that may vary from a DTD to a tree automaton) for the distributed document, “propagating” locally this type into a collection of types, that we call *typing*, while preserving desirable properties. We also consider the bottom-up design which consists in, given a type for each external resource, exhibiting a global type that is enforced by the local types, again with natural desirable properties. In the article, we lay out the fundamentals of a theory of distributed XML design, analyze problems concerning typing issues in this setting, and study their complexity.

Keywords: Semistructured Data, XML Schemas, Distributed Data, Database Design, Distributed XML

1. Introduction

Context and Motivation. With the Web, information tends to be more and more distributed. In particular, the distribution of XML data is essential in many areas such as e-commerce (shared product catalog), collaborating editing (e.g., based on WebDAV [21]), or network directories [24]. (See also the W3C XML Fragment Interchange Working Group [19].) It becomes often cumbersome to verify the validity, e.g., the type, of such a hierarchical structure spanning several machines. In this paper, we consider typing issues raised by the distribution of XML documents. We introduce “nice” properties that the distribution should obey to facilitate type verification based on locality conditions. We propose an automata-based study of the problem. Our theoretical investigation provides a

Email addresses: serge.abiteboul at inria.fr (S. Abiteboul),
georg.gottlob@comlab.ox.ac.uk (G. Gottlob), manna@mat.unical.it (M. Manna)

starting point for the distributed validation of tree documents (verification) and for selecting a distribution for Web data (design). In general, it provides new insights in the typing of XML documents.

A *distributed XML document* $T_{[t_1..t_n]}$ is given by an *XML kernel-document* $T_{[f_1, \dots, f_n]}$, that is stored locally at some site, some of which leaves (the *docking points*) refer to external resources, here denoted by f_1, \dots, f_n , that provide the additional XML data $t_1..t_n$ to be attached, respectively, to T . For simplicity, each node playing the role of docking point is called a *function-node* and it is labeled with the resource that it refers.

The *extension* $\text{ext}_T(t_1..t_n)$ of T is the whole XML document obtainable from the distributed document $T_{[t_1..t_n]}$ by replacing the node referring resource f_i with the forest of XML trees (in left-to-right order) directly connected to the root of t_i , for each i in $[1..n]$.

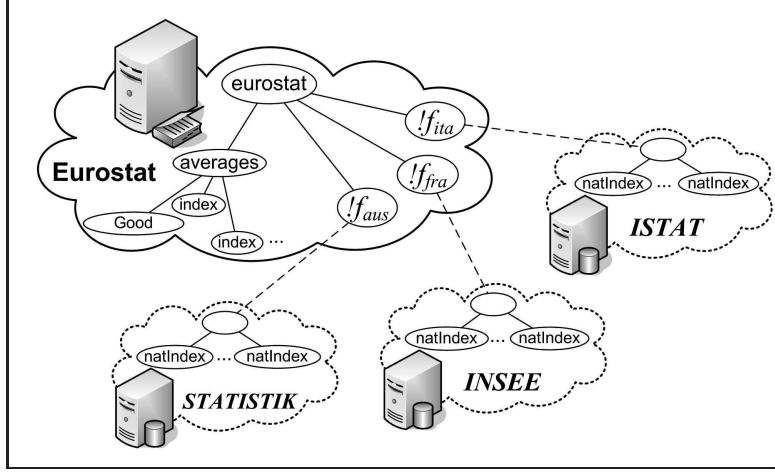


Figure 1: A distributed XML document for the *National Consumer Price Index* where both its kernel and some of its remote XML (sub)documents have been hightailed.

Figure 1 shows a (drastically simplified) possible distributed XML document for the *National Consumer Price Index (NCPI)*¹ maintained by the *Eurostat*². This example is detailed further in this section.

Typically, a global designer first chooses a specific language for constraining the documents of interest. The focus in this paper is on “structural constraints”. Clearly, one could also consider other constraints such as key and referential constraints. So, say the designer has to specify documents using DTDs. Then he specifies a kernel document $T_{[f_1, \dots, f_n]}$ together with either:

- bottom-up design*: types τ_i for each f_i ;
- top-down design*: a global type τ .

¹See <http://epp.eurostat.ec.europa.eu>

²See <http://ec.europa.eu/eurostat>

In the bottom-up case, we are interested by the global type that results from each local source enforcing its local type. *Can such typing be described by specific type languages?*

In the top-down case, we would like the extension of T to satisfy τ . The issue is “*Is it possible to enforce it using only local control?*” In particular, we would like to break down τ into local types τ_i that could be enforced locally. More precisely, we would like to provide each \mathbf{f}_i with a typing τ_i guaranteeing that (i) if each \mathbf{f}_i verifies its type, then the global type is verified (soundness), and (ii) the typing $\tau_1.. \tau_n$ is not more restrictive than the global type (completeness). We call such a typing local typing. We both study (maximal) local typings and an even more desirable notion, namely “perfect typings” (to be defined).

To conclude this introduction, we next detail the Eurostat example. We then present a formal overview of the paper (which may be skipped in a first reading.) Finally, we survey related works.

Working Example. Before mentioning some related works and concluding this section, we further illustrate these concepts by detailing our *Eurostat* example.

The NCPI is a document containing consumer price data for each EC country. We assume that the national data are maintained in local XML repositories by each country’s national statistics bureau (INSEE for France, Statistik for Austria, Istat for Italy, UK Statistics Authority, and so on). Each national data set is under the strict control of its respective statistics bureau. The kernel document T_0 is maintained by Eurostat in Luxembourg and has a docking point for each resource \mathbf{f}_i located in a particular country. In addition, T_0 contains average data for the entire EU zone. Figure 2 shows a possible extension of T_0 , where the actual data values are omitted.

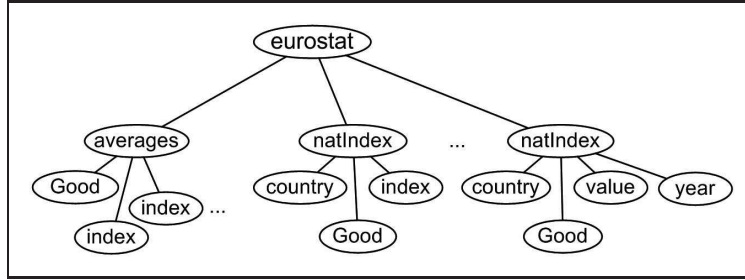


Figure 2: The extension of a possible distributed document having kernel T_0 , and complying with the whole structure showed in Figure 1.

We first assume that Eurostat specifies the global type τ for the distributed NCPI document, where τ is given by the DTD document shown in Figure 3. (In the following, we adopt a more succinct notation for types where the content model of an element name is either left undefined if it is solely “#PCDATA”, or defined by a rule of the form “ $\text{index} \rightarrow \text{value}, \text{year}$ ”, otherwise.) Briefly, DTD τ requires that each possible extension $\text{ext}_{T_0}(t_1..t_n)$ consists of a subtree containing average data for *Goods* (such as food, energy, education, and so on).

Each *Good* item is evaluated in different years by means of an *index*. Moreover, $\text{ext}_{T_0}(t_1..t_n)$ may contain a forest of *nationalIndex*, namely indexes associated to goods in precise countries.

```

<!ELEMENT eurostat (averages, nationalIndex*)>
  <!ELEMENT averages (Good, index+)+>
    <!ELEMENT nationalIndex (country, Good, (index | value, year))>
      <!ELEMENT index (value, year)>
    <!ELEMENT country (#PCDATA)>
    <!ELEMENT Good (#PCDATA)>
    <!ELEMENT value (#PCDATA)>
    <!ELEMENT year (#PCDATA)>

```

Figure 3: W3C DTD τ

To comply with different national databases, two different formats are allowed: $(\text{country}, \text{Good}, \text{index})$ or $(\text{country}, \text{Good}, \text{value}, \text{year})$. It is easy to see that the pair $\langle \tau, T_0 \rangle$ allows a local typing (see Figure 4) that is even perfect (so, can be obtained by the algorithm shown in Section 6), as we will clarify in the next section.

```

rooti → nationalIndex*
nationalIndex → country, Good, (index | value, year)
index → value, year

```

Figure 4: Type τ_i ($1 \leq i \leq n$) in the perfect typing for the top-down design $\langle \tau, T_0 \rangle$

Suppose now that a designer defined instead the DTD τ' shown in Figure 5 as global type. The pair $\langle \tau', T_0 \rangle$ would be a bad design since τ' imposes to all countries to adopt the same format for their indexes (*natIndA* or *natIndB*). But this represents a constraint that cannot be controlled locally. Indeed, this new design does not admit any local typing. The nice locality properties of designs are obvious in such simplistic examples. However, when dealing with a large number of peers with very different desires and complex documents, the problem rapidly starts defeating human expertise. Consider, for instance, the type τ''

```

eurostat → averages, (natIndA* | natIndB*)
averages → (Good, index+)+
natIndA → country, Good, index
natIndB → country, Good, value, year
index → value, year

```

Figure 5: Type τ'

defined in Figure 6 and the kernel $T_1 = \text{eurostat}(\mathbf{f}_1, \text{nationalIndex}(\mathbf{f}_2), \mathbf{f}_3)$ containing only three function calls. Even if this design is as small as $\langle \tau, T_0 \rangle$, it already starts to become hard to manage with no automatic technique. Here, *natIndA* and *natIndB* are different specializations of *nationalIndex* elements

(note that, as detailed in Section 2.2, this feature requires schema languages more expressive than DTDs), while all other elements have no specialization.

```

eurostat → averages, (natIndA, natIndB)+
averages → (Good, index+)+
natIndA  → country, Good, index
natIndB  → country, Good, value, year
index    → value, year

```

Figure 6: Type τ''

In this case, it is not as easy as before to state that the new design has no perfect typing and exactly the two maximal local typings shown below (only the content models of the roots are specified). This is mainly because the functions in T_1 have different depth, but also due to specializations.

$$\begin{aligned}
\tau''_{1,1}: \text{root}_1 &\rightarrow \text{averages}, (\text{natIndA}, \text{natIndB})^* \\
\tau''_{2,1}: \text{root}_2 &\rightarrow \text{country}, \text{Good}, \text{index} \\
\tau''_{3,1}: \text{root}_3 &\rightarrow \text{country}, \text{Good}, \text{value}, \text{year}, (\text{natIndA}, \text{natIndB})^* \\
\\
\tau''_{1,2}: \text{root}_1 &\rightarrow \text{averages}, (\text{natIndA}, \text{natIndB})^*, \text{natIndA} \\
\tau''_{2,2}: \text{root}_2 &\rightarrow \text{country}, \text{Good}, \text{value}, \text{year} \\
\tau''_{3,2}: \text{root}_3 &\rightarrow (\text{natIndA}, \text{natIndB})^*
\end{aligned}$$

The techniques developed in this paper are meant to support experts in designing such distributed document schemas.

Overview of Results. We next precise the formal setting of the paper and its results. From a formal viewpoint, we use Active XML terminology and notation for describing distributed documents [1].

Not surprisingly, our results depend heavily of the nature of the typing that is considered. For types, we consider abstract versions of the conventional typing languages [6, 29, 33, 35], namely \mathcal{R} -DTDs (for *W3C DTDs*), \mathcal{R} -SDTDs (for *W3C XSD*), and \mathcal{R} -EDTDs (for regular tree grammars such as *Relax-NG*) where \mathcal{R} (varying among **nFAs**, **dFAs**, **nREs**, and **dREs**, namely automata and regular expressions both nondeterministic and deterministic) denotes the formalism for specifying content models.

As a main contribution, we initiate a theory of local typing. We introduce and study three main notions of locality: local typing, maximal local typing, and perfect typing. For a given XML schema language \mathcal{S} , we study the following verification problems:

- Given an \mathcal{S} -typing for a top-down \mathcal{S} -design, determine whether the former is local, maximal local, or perfect. We call these problems $\text{LOC}_{[\mathcal{S}]}$, $\text{ML}_{[\mathcal{S}]}$ and $\text{PERF}_{[\mathcal{S}]}$, respectively;
- Given a top-down \mathcal{S} -design, establish whether a local, maximal local, or perfect \mathcal{S} -typing does exist (and, of course, find them). We call these problems $\exists\text{-LOC}_{[\mathcal{S}]}$, $\exists\text{-ML}_{[\mathcal{S}]}$, and $\exists\text{-PERF}_{[\mathcal{S}]}$, respectively;

- Given a bottom-up \mathcal{S} -design, establish whether it defines an \mathcal{S} -type. The problem is called $\text{CONS}_{[\mathcal{S}]}$.

The analysis carried out in this paper provides tight complexity bounds for some of these problems. In particular, for bottom-up designs, we prove that $\text{CONS}_{[\mathcal{S}]}$ is:

- decidable in constant time for \mathcal{R} -EDTDs, for each \mathcal{R} ;
- **PSPACE**-complete both for \mathcal{R} -DTDs and \mathcal{R} -SDTDs, in general;
- **PSPACE**-hard with an **EXPTIME** upper bound for dRE-DTDs and dRE-SDTDs.

For top-down designs, after showing that the problems for trees can be reduced to problems on words, we specialize the analysis to the case of $\mathcal{R} = \text{nFA}$. In particular:

- $\text{LOC}_{[\mathcal{S}]}$, $\text{ML}_{[\mathcal{S}]}$, $\text{PERF}_{[\mathcal{S}]}$, and $\exists\text{-PERF}_{[\mathcal{S}]}$ are **PSPACE**-complete when \mathcal{S} stands for **nFA**-DTD or **nFA**-SDTD, and $\text{LOC}_{[\mathcal{S}]}$ is **EXPTIME**-complete for **nFA**-EDTDs;
- $\exists\text{-LOC}_{[\mathcal{S}]}$ and $\exists\text{-ML}_{[\mathcal{S}]}$ are **PSPACE**-hard with an **EXPSpace** upper bound when \mathcal{S} stands for **nFA**-DTD or **nFA**-SDTD;
- the remaining problems are **EXPTIME**-hard with either **coNEXPTIME** or **2-EXPSpace** upper bounds.

Related Work. Distributed data design has been studied quite in depth, in particular for relational databases [13, 34]. Some previous works have considered the design of Web applications [12]. They lead to the design of Web sites. The design there is guided by an underlying process. It leads to a more dynamic notion of typing, where part of the content evolves in time, e.g., creating a cart for a customer. For obvious reasons, distributed XML has raised a lot of attention recently. Most works focused on query optimization, e.g., [4]. The few that consider design typically assume no ordering or only limited one [9]. This last work would usefully complement the techniques presented here. Also, works on relational database and LDAP³ design focus on unordered collections. Even the W3C goes in this direction with a working group on *XML Fragment Interchange* [19]. The goal is to be able to process (e.g., edit) document fragments independently. Quoting the W3C Candidate Recommendation: “It may be desirable to view or edit one or more [fragments] while having no interest, need, or ability to view or edit the entire document.” This is clearly related to the problem we study here. Finally, the concept of distributed documents, as defined in this paper, is already implemented in Active XML, a declarative framework that harnesses web services for data integration, and is put to work in a peer-to-peer architecture [1, 2]. Moreover, XML documents, XML schemas, and formal languages have been extensively studied and, although all the problems treated in this paper are essentially novel,⁴ the theoretical analysis has got

³Lightweight Directory Access Protocol (LDAP) is a set of open protocols used to access centrally stored information over a network.

⁴Consider that, as highlighted in [31], an interesting problem in Formal Language Theory open for more than ten years, named *Language Primality*, is essentially a special case of our problem $\exists\text{-LOC}_{[\text{dFA}]}$. The complexity of Primality has been also settled in [31].

benefit from a number of existing works. Classical results about formal (string and regular) languages come from [20, 23, 25, 26, 32, 36, 37, 42] and in particular, those about state complexity of these languages can be found in [22, 43], those about one-unambiguous regular languages in [7, 11], and those about alternating finite state machines in [17, 42]. Finally, regarding XML documents and schemas, our abstract presentation builds on ten years of research in this field. In particular, it has been strongly influenced by document typings studied in [6, 8, 10, 14, 15, 27–30, 33, 33, 35, 38] and results on them presented there.

Structure of the Paper. This concludes the introduction. The remaining of the paper is organized as follows. Section 2 fixes some preliminary notation, formally introduces our notions of type, distributed XML document, and defines the decision problems studied. It also provides an overview of the results. Section 3 considers the bottom-up design. Section 4 presents basic results regarding the top-down design. Sections 5 and 6 present the main results for the word case. Section 7 completes the complexity analysis. Section 8 concludes and mentions possible areas for further research.

2. General setting

In this paper, we use a widespread abstraction of XML documents and XML Schemas focusing on document structure [6, 29, 33, 35], and Active XML terminology and notation for describing distributed documents [1, 2]. In particular, for XML Schemas we will consider families of *tree grammars* (called \mathcal{R} -DTDs, \mathcal{R} -SDTDs, and \mathcal{R} -EDTDs) each of which allows different formalisms for specifying content models (\mathcal{R} may vary among **nFAs**, **dFAs**, **nREs**, and **dREs**, respectively, nondeterministic automata, deterministic automata, regular expressions, and deterministic regular expressions). This, could be surprising at first sight because the W3C standards impose stricter limitations. However, as we will informally motivate later, (and has been formally proved in [31]), some of the problems we define and analyze here, have the same complexity independently of whether we use deterministic or nondeterministic string-automata, or even deterministic regular expressions. Informally, it can be observed that the document distribution often erases the benefits of determinism. For this reason, and because this paper intends to be a first fundamental study of XML distribution, we include in our analysis different possibilities for schema languages, even if for some problem we only analyze the most general case (\mathcal{R} is set to **nFAs**) in order to delimit its complexity. Moreover, the typing problems we study hint at the possibility that there could be interesting real world applications (all distributed applications that involve the management of distributed data, such as data integration from databases and other data resources exported as Web services, or managing active views on top of data sources) where W3C recommendations are too strict and thus unsuitable in the context of distributed XML documents.

2.1. Preliminaries

In this paper, we use also the following notation. We always denote, by Σ , a (finite) *alphabet*; by ε , the empty string; by \emptyset , the empty language; by \cdot , the binary relation of concatenation on Σ^* and by \circ , its extension on 2^{Σ^*} ; by \mathcal{A} , an *automaton* for defining a string-language or tree-language over Σ ; by r a *regular expression* over Σ ; by \mathcal{R} , a formalism for defining string languages; by \mathcal{S} , a formalism for defining tree languages; by τ , an \mathcal{R} -*type* or an \mathcal{S} -*type* (a concrete formal structure defining, respectively, a string languages or a tree languages, such as a regular expression or an XML schema document) over Σ ; by $[\tau]$, the language defined by τ .

2.1.1. XML Documents

An *XML document* can be viewed, from a structural point of view, as a *finite ordered, unranked tree* (hereafter just a *tree*) t with nodes labeled over a given alphabet Σ . The *topmost node* in t is denoted by $\text{root}(t)$, while for any node x of t , we denote by

- $\text{parent}_t(x)$ the (unique) parent node of x (if node x is not the root);
- $\text{children}_t(x)$ is the sequence of children (possibly empty) of x in left-to-right order;
- $\text{tree}_t(x)$ the subtree of t rooted at x ;
- $\text{lab}_t(x) \in \Sigma$ the label of x ;
- $\text{anc-str}_t(x) \in \Sigma^+$ is the sequence of labels of the path from the root of t to x ;
- $\text{child-str}_t(x) \in \Sigma^*$ the labels of the children of x in left-to-right order.

In particular, if $\text{child-str}_t(x) = \varepsilon$, then x is called a *leaf node*. The size of t , denoted by $\|t\|$, is the number of its nodes. Also in these predicates we may omit the subscript t when it is clear from the context.

2.1.2. Regular String Languages

A *nondeterministic finite state machine* (**nFA**) over Σ is a quintuple $\mathcal{A} = \langle K, \Sigma, \Delta, q_s, F \rangle$ where K are the *states*, $q_s \in K$ is the *initial state*, $F \subseteq K$ are the *final states*, and $\Delta \subseteq K \times (\Sigma \cup \{\varepsilon\}) \times K$ is the *transition relation*. Each triple $(q, \alpha, q') \in \Delta$ is called a *transition* of \mathcal{A} . Sometimes the notation $q' \in \Delta(q, \alpha)$, where Δ is seen as a function from $K \times (\Sigma \cup \{\varepsilon\})$ to 2^K , is more convenient. By $\Delta^* \subseteq K \times \Sigma^* \times K$ we denote the *extended transition relation* defined as the reflexive-transitive closure of Δ , in such a way that $(q, w, q') \in \Delta^*$ iff there is a sequence of transitions from q to q' recognizing string w . The set of strings $[\mathcal{A}] = \{w \in \Sigma^* : \Delta^*(q_s, w) \in F\}$ is the *language* defined by \mathcal{A} . Such machines can be combined in various ways (see [22] for a comprehensive analysis). In particular, $\bar{\mathcal{A}}$ denotes the complement of \mathcal{A} , and defines the language $\Sigma^* - [\mathcal{A}]$. Given two **nFAs** \mathcal{A}_1 and \mathcal{A}_2 , we denote by $\mathcal{A}_1 \cdot \mathcal{A}_2$, $\mathcal{A}_1 \cup \mathcal{A}_2$, $\mathcal{A}_1 \cap \mathcal{A}_2$, and $\mathcal{A}_1 - \mathcal{A}_2$ the **nFA** defining $[\mathcal{A}_1] \circ [\mathcal{A}_2]$, $[\mathcal{A}_1] \cup [\mathcal{A}_2]$, $[\mathcal{A}_1] \cap [\mathcal{A}_2]$, and $[\mathcal{A}_1] - [\mathcal{A}_2]$, respectively (operators \cdot and \circ are often omitted). Also, for a set $\mathbb{A} = \{\mathcal{A}_1, \dots, \mathcal{A}_m\}$ of **nFAs**, we often write $\cap \mathbb{A}$ (or $\cup \mathbb{A}$) instead of $\mathcal{A}_1 \cap \dots \cap \mathcal{A}_m$ (or $\mathcal{A}_1 \cup \dots \cup \mathcal{A}_m$).

A *deterministic finite automaton* (dFA) over Σ is an nFA where Δ is a function from $K \times \Sigma$ to K .

A (possibly nondeterministic) *regular expression* (nRE or also *regex*, for short) r over Σ is generated by the following abstract syntax:

$$r ::= \varepsilon \mid \emptyset \mid a \mid (r \cdot r) \mid (r + r) \mid r? \mid r^+ \mid r^*$$

where a stands generically for the elements of Σ . When it is clear from the context, we avoid unnecessary brackets or the use of \cdot for concatenation. The language $[r]$ is defined as usual.

A *deterministic regular expression* (dRE) r is an nRE with the following restriction. Let us consider the regex \tilde{r} built from r by replacing each symbol $a \in \Sigma$ with \tilde{a}^i where i is the position from left-to-right of a in r . By definition, r is a dRE if there are no strings $w\tilde{a}^i u$ and $w\tilde{a}^j v$ in $[\tilde{r}]$ such that $i \neq j$. The language $[r]$ of a dRE r is called *one-unambiguous* [11].

A cartesian product of n finite sets is called a “box” [41]. More precisely, fix a positive number n . Let Σ be an alphabet. A *box* B over Σ is any language of the form $\Sigma_1 \dots \Sigma_n$ where n is its *width*, and $\Sigma_i \subseteq \Sigma$ for each i in $[1..n]$. Clearly, each box is a regular language as it is a finite one.

2.1.3. Regular Tree Languages

A *nondeterministic Unranked Tree Automaton* (nUTA) is a quadruple $\mathcal{A} = \langle K, \Sigma, \Delta, F \rangle$ where Σ is the alphabet, K is a finite set of *states*; $F \subseteq K$ is the set of *final states*; Δ is a function mapping pairs from $(K \times \Sigma)$ to nFAs over K . A tree t belongs to $[\mathcal{A}]$ if and only if there is a mapping μ from the nodes of t to K such that (i) $\mu(\text{root}(t)) \in F$, and (ii) for each node x of t , either ε or $\mu(\text{children}(x))$ belongs to $[\Delta(\mu(x), \text{lab}(x))]$ according to whether x is a leaf-node or not, respectively.

A *bottom-up-deterministic Unranked Tree Automaton* (dUTA) over Σ is an nUTA where Δ is a function from $(K \times \Sigma)$ to dFAs over K in such a way that $[\Delta(q, a)] \cap [\Delta(q', a)] = \emptyset$ for each $q \neq q'$.

2.1.4. Known decision problems

In this section we recall some well known decision problems.

Definition 1. $\text{EQUIV}_{[\mathcal{S}]}$ is the following decision problem. Given two \mathcal{S} -types, do they define the same language? \square

In particular, whenever we consider two \mathcal{R} -types instead of \mathcal{S} -types, we still denote by $\text{EQUIV}_{[\mathcal{R}]}$ the equivalence problem defined exactly as above.

Definition 2. $\text{ONE-UNAMB}_{[\mathcal{R}]}$ is the following decision problems. Given a regular language L specified by an \mathcal{R} -type, is L one-unambiguous? \square

2.2. Types

As already mentioned, we consider abstractions of the most common XML Schemas by allowing regular languages, specified by possibly different formalisms for defining content models. More formally, let \mathcal{R} be a mechanism for describing regular languages (**nFAs**, **dFAs**, **nREs**, **dREs**, or even others). We want to define and computationally characterize the problems regarding *Distributed XML design* in a comparative analysis among the three main actual formalisms for specifying XML schema documents: *W3C DTDs*, *W3C XSD* and *Regular Tree Grammars* (like *Relax-NG*). For each of these schema languages, we adopt a class of abstractions that we call \mathcal{R} -DTDs, \mathcal{R} -SDTDs, and \mathcal{R} -EDTDs, respectively, where \mathcal{R} is the particular mechanism for defining content models. We show that a number of properties do not depend on the choice of \mathcal{R} (or even of \mathcal{S}) and for some complexity results we focus our analysis to the case of **nFAs**. Before that, we summarize in Table 1 the relevance of the different tree grammars.

Table 1: Comparison between our abstractions of XML Schemas and existing formalisms.

Schema language	Previously introduced formalism	Our abstraction
<i>W3C DTDs</i>	<i>DTDs and ltds</i>	dRE-DTDs
<i>W3C XSD</i>	<i>Single-Type Tree Grammars</i> and <i>single-type EDTDs</i>	dRE-SDTDs
<i>Relax NG</i>	<i>unranked regular tree languages</i> (<i>specialized ltds</i> and <i>EDTDs</i>)	nRE-EDTDs

2.2.1. \mathcal{R} -DTD types

The following definition generalizes definitions considered in the literature such as *ltds* [6, 35] or *DTDs* [29, 30], and defined for analyzing the properties of *W3C Document Type Definitions*. As we marry these views, we define the following class of abstractions capturing all of them.

Definition 3. An \mathcal{R} -DTD is formalized as a triple $\tau = \langle \Sigma, \pi, s \rangle$ where

- Σ is an alphabet (the **element names**);
- π is a function mapping the symbols of Σ to \mathcal{R} -types still over Σ ;
- $s \in \Sigma$ is the **start symbol**.

A tree t , having labels over Σ belongs to $[\tau]$ if and only if: $\text{lab}(\text{root}(t)) = s$ and $\text{child-str}(x) \in [\pi(\text{lab}(x))]$, for each node x of t . For a given element name a , the regular language $[\pi(a)]$, associated to a , is usually called the **content model** of a . \square

Notice that, due to the above definition, \mathcal{R} -DTDs with useless element names, or even defining the empty language, do exist. This is because the above definition allows to specify \mathcal{R} -DTDs that are, in a sense, “not reduced” (think about finite automata with unreachable states). Since it is much more convenient to deal with types that are not effected by these drawbacks, after giving some more definition, we formalize the notion of **reduced** types.

We introduce the **dFA** $dual(\tau)$. It is the language consisting of the set of paths from the root to a leaf in trees in $[\tau]$ and it is in some sense the *vertical language* of τ .

Definition 4. Let $\tau = \langle \Sigma, \pi, s \rangle$ be an \mathcal{R} -DTD. We build from τ the **dual dFA** $dual(\tau) = \langle K, \Sigma, \delta, q_0, F \rangle$ as follows:

- $K = \{q_0\} \cup \{q_a : a \in \Sigma\}$;
- $\delta(q_0, s) = q_s$;
- for each $a, b \in \Sigma$, $\delta(q_a, b) = q_b$ iff b appears in the alphabet of $\pi(a)$;
- $q_a \in F$ iff $\varepsilon \in [\pi(a)]$. □

Before defining a set of conditions ensuring that all the content models of a given \mathcal{R} -DTD τ are well defined and have no redundancy w.r.t. the language $[\tau]$, we mark the states of $dual(\tau)$ (in a bottom-up style) as follows:

1. Mark each final state of $dual(\tau)$ as *bound*;
2. For each non-*bound* state q_b , consider the set $\Sigma_b \subseteq \Sigma$ where $\delta(q_b, a)$ is *bound* iff $a \in \Sigma_b$. If $[\pi(b)] \cap \Sigma_b^+ \neq \emptyset$, then mark also q_b as *bound*;
3. Repeat step 2 until no more states can be marked.

Definition 5. Let τ be an \mathcal{R} -DTD. We say that τ is **reduced** iff

- Each state of $dual(\tau)$ is in at least a path from q_0 to a final state in F ;
- Each state of $dual(\tau)$ is *bound*;
- $[dual(\tau)]$ is nonempty. □

We consider only *reduced* \mathcal{R} -DTDs where, by the previous definition, it is clear that $[\tau] \neq \emptyset$. Note that for a given \mathcal{R} -DTD τ , it is very easy to build $dual(\tau)$ and for each “unprofitable” state q_a

- remove the element name a from Σ ;
- remove the rule $\pi(a)$ from π ;
- modify the rules containing a in their content models (using standard regular language manipulation) to produce only words not containing a (see [28], for more details.)

Finally, we notice that only the last step of the reducing algorithm may depend on the choice of \mathcal{R} . Clearly, an \mathcal{R} -DTD and its reduced version describe the same language.

From a theoretical point of view, \mathcal{R} -DTDs do not express more than the *local tree languages* [33]. In particular, **nFA**-DTDs, **dFA**-DTDs and **nRE**-DTDs exactly

capture this class of languages while **dRE-DTDs** are less expressive [33, 35]. Nevertheless, the last class of types (using deterministic regular expressions [11] and that does not capture all the local tree languages) is, from a structural point of view, the closest to *W3C DTDs*.

In this paper, for a given \mathcal{R} -DTD where \mathcal{R} stands for **dFAs** or **nFAs** (for shortness, w.l.o.g., and only in examples) we often specify π as a function that maps Σ -symbols to Σ -**nREs** (recall that any regular expression of size n can be transformed into an equivalent ε -free **nFA** with $\mathcal{O}(n \log^2 n)$ transitions in time $\mathcal{O}(n \log^2 n)$ [20, 23].)

Finally, an example of **dRE-DTD** is $\tau_1 = \langle \{s_1, c\}, \pi_1, s_1 \rangle$ with $\pi_1(s_1) = c^*$ and $\pi_1(c) = \varepsilon$. In the rest of the paper, we often omit to specify rules such as $\pi_1(c) = \varepsilon$; i.e., if no rule is given for a label, nodes with this label are assumed to be (solely) leaves.

2.2.2. \mathcal{R} -SDTD types

The following definition generalizes definitions considered in the literature such as *Single-Type Tree Grammars* [33] or *single-type EDTDs* [29], and defined for analyzing the properties of *W3C XML Schema Definitions*. Also here, we define a class of abstractions capturing all of them.

Definition 6. An \mathcal{R} -SDTD (standing for single-type extended \mathcal{R} -DTD) is a quintuple $\tau = \langle \Sigma, \tilde{\Sigma}, \pi, \tilde{s}, \mu \rangle$ where

- $\tilde{\Sigma}$ are the **specialized element names**;
- $\langle \tilde{\Sigma}, \pi, \tilde{s} \rangle$ is an \mathcal{R} -DTD on $\tilde{\Sigma}$ and denoted by $\text{dtd}(\tau)$;
- $\mu : \tilde{\Sigma} \rightarrow \Sigma$ is a mapping from all the specialized element names *onto* the set of element names. For each $a \in \Sigma$, we denote by $\tilde{a}^1, \dots, \tilde{a}^n$ the distinct elements in $\tilde{\Sigma}$ that are mapped to a . This set is denoted $\tilde{\Sigma}(a)$;
- Let $\text{dual}(\text{dtd}(\tau))$ be $\langle K, \tilde{\Sigma}, \tilde{\delta}, q_0, F \rangle$. Build from this **dFA** the possibly **nFA** $\text{dual}(\tau) = \langle K, \Sigma, \delta, q_0, F \rangle$ where for each $q, q' \in K$ and $a \in \Sigma$, $\delta(q, a) = q'$ iff there is an element $\tilde{a} \in \tilde{\Sigma}$ such that $\tilde{\delta}(q, \tilde{a}) = q'$. We require that $\text{dual}(\tau)$ is a **dFA** (this captures the single-type requirement). Also in this case, $\text{dual}(\tau)$ defines the vertical language of τ .

A tree t , labeled over Σ , is in $[\tau]$ if and only if there exists a tree $t' \in [\text{dtd}(\tau)]$ such that $t = \mu(t')$ (where μ is extended to trees). Informally, we call t' a witness for t . Finally, an \mathcal{R} -SDTD τ is **reduced** if and only if $\text{dtd}(\tau)$ is. \square

As for \mathcal{R} -DTDs, we consider only *reduced* \mathcal{R} -SDTDs.

From a theoretical point of view, \mathcal{R} -SDTDs are more expressive than \mathcal{R} -DTDs but do not capture the *unranked regular tree languages* yet.

2.2.3. \mathcal{R} -EDTDs types

The following definition generalizes definitions considered in the literature such as *specialized ltds* [6, 35] or *EDTDs* [29]. Such formalisms (like *Relax-NG*), from a structural perspective, express exactly the *homogeneous unranked regular tree languages* and are as expressive as *unranked tree automata* or *Regular Tree Grammars* [10].

Definition 7. An \mathcal{R} -EDTD (extended \mathcal{R} -DTD) τ is an \mathcal{R} -SDTD without the single-type requirement. More formally, the automaton $dual(\tau)$, built as for \mathcal{R} -SDTD, may be here an **nFA**. The language $[\tau]$ is defined as for \mathcal{R} -SDTDs. \square

2.3. Distributed Documents

In the context of distributed architectures (e.g., P2P architectures), distributed documents (or distributed trees), such as AXML documents, are XML documents that may contain embedded *function calls*. In particular, a *distributed XML document* $T_{[t_1..t_n]}$ can be viewed as a collection of (classical) XML documents $t_1..t_n$ brought together by a unique (special) XML document $T_{[\mathbf{f}_1, \dots, \mathbf{f}_n]}$, the *kernel*, some of whose leaf-nodes, called function-nodes, play the role of “docking points” for the external resources $\mathbf{f}_1, \dots, \mathbf{f}_n$. The “activation” of a node of T having a function as label, say \mathbf{f}_i , consists in a call to resource (or function) \mathbf{f}_i the result of which is still an XML document, say t_i . When \mathbf{f}_i is invoked, its result is used to extend the kernel $T_{[\mathbf{f}_1, \dots, \mathbf{f}_n]}$. Thus, each docking point connects the peer that holds the kernel and invokes the resource \mathbf{f}_i , and the peer that provides the corresponding XML document t_i . For simplicity of notation, for labeling a function-node we use exactly the name of the resource it refers. For instance, the tree $T_0 = s(a \ \mathbf{f}_1 \ b(\mathbf{f}_2))$ is a kernel having s as root, and containing two function-nodes referring the external resource \mathbf{f}_1 and \mathbf{f}_2 .

The *extension* $\text{ext}_T(t_1..t_n)$ of T is the whole XML document (without any function at all) obtained from the distributed document $T_{[t_1..t_n]}$ by replacing each node referring resource \mathbf{f}_i with the forest of XML trees (in left-to-right order) directly connected to the root of t_i . This process is called *materialization*. For instance, the extension of kernel T_0 would be $s(a \ c(dd) \ b(d(ef)))$ in case of resources \mathbf{f}_1 and \mathbf{f}_2 provided trees $s_1(c(dd))$ and $s_2(d(ef))$, respectively.

An interesting task is to associate a type τ_i (e.g., a *W3C XSD* document) to each resource \mathbf{f}_i in such a way that the XML document t_i returned as answer is *valid* w.r.t. this type and any materialization process always produces a document $\text{ext}_T(t_1..t_n)$ valid w.r.t. a given global type τ (still specified by the *W3C XSD* syntax). A global type and a kernel document represent the (top-down) *design* of a given distributed architecture. A collection of types associated to the function calls in such a design is called a *typing*. Given a distributed design, we would like to know whether either a precise typing has some properties or a typing with some properties does exist. But also, we could directly start from a kernel T and a typing (bottom-up design) and analyze the properties of the tree language consisting in each possible extension $\text{ext}_T(t_1..t_n)$.

More formally, let Σ and $\Sigma^{\mathbf{f}}$ be two alphabets, respectively, of *element names* (such as s, a, b, c , etc.) and *function symbols* (such as \mathbf{f}, \mathbf{g} , etc.). A *kernel document* or *kernel tree* $T_{[\mathbf{f}_1, \dots, \mathbf{f}_n]}$ (or also $T(\mathbf{f}_n)$, with (\mathbf{f}_n) denoting a sequence⁵ of length n) is a tree over $(\Sigma \cup \Sigma^{\mathbf{f}})$ where:

- (i) the root is an element node (say s_0);

⁵We denote a finite sequence of objects (x_1, \dots, x_n) over an index set $I = \{1, \dots, n\}$ by (x_n) and we often omit the specification of the index set I .

- (ii) the function nodes $\mathbf{f}_1, \dots, \mathbf{f}_n$ are leaf nodes;
- (iii) no function symbol occurs more than once.

In particular, for each non-leaf node of T , say x , the *kernel string* $\text{child-str}(x)$, with $k \geq 0$ functions, is of the form $w_h \mathbf{f}_{h+1} w_{h+1} \dots \mathbf{f}_{h+k} w_{h+k}$ (for some h in $[1..n]$) where $w_i \in \Sigma^*$ for each $i \in \{h, \dots, h+k\}$, $\mathbf{f}_i \in \Sigma^{\mathbf{f}}$ for each $i \in \{h+1, \dots, h+k\}$, and $\mathbf{f}_i \neq \mathbf{f}_j$ for each $i \neq j$.

We next consider its semantics. It is defined by providing a tree for each function-node. In particular, an *extension* ext maps each i in $[1..n]$ to a tree $t_i = \text{ext}(\mathbf{f}_i)$. The *extension* $\text{ext}_T(t_1..t_n)$ of a kernel $T_{[\mathbf{f}_1, \dots, \mathbf{f}_n]}$ is obtained by replacing each \mathbf{f}_i with the forest of trees (in left-to-right order) directly connected to the root of t_i .

A *type* τ for a kernel tree T is one of an \mathcal{R} -DTD, \mathcal{R} -SDTD, or \mathcal{R} -EDTD. Given an extension $t_1..t_n$, we say that tree $T_{[t_1..t_n]}$ satisfies type τ if and only if $\text{ext}_T(t_1..t_n)$ does. This motivates requirement (iii) to avoid irregularities: For instance, in the kernel $T_1 = s(\mathbf{f} \mathbf{f})$ the children of s in any extension of T_1 are of the form ww for some word w . But since this is not a regular language, the type of T_1 cannot be defined by none of the three adopted formalisms. Although we disallow the same function to appear twice, several functions may share the same type. Also, even if for labeling a function-node we use exactly the name of the resource it refers (for simplicity of notation), this does not prohibit a resource to provide two XML subtrees to be attached to the kernel. In fact, different names (function symbols) can be associated to the same resource still preserving extensions from irregularities.

We introduce *typings* to constrain the types of the function calls of a kernel document. A *typing* for a kernel tree $T(\mathbf{f}_n)$ is a positional mapping from the functions in (\mathbf{f}_n) to a sequence (τ_n) of types (schema documents). Now, as we replace each \mathbf{f}_i (in the extensions of T) with a forest of XML documents then, for each type τ_i associated to \mathbf{f}_i , we actually use a schema document containing an “extra” element name, say s_i , being only the label of the root in all the trees in $[\tau_i]$.

Definition 8. We denote by $\text{ext}_T(\tau_n)$ the tree language consisting of all possible extensions $\text{ext}_T(t_1..t_n)$ where $t_i \models \tau_i$ (t_i is valid w.r.t τ_i) for each i . \square

Definition 9. We denote by $T(\tau_n)$ the nFA-EDTD (or nRE-EDTD) constructed from T and (τ_n) in the obvious way such that $[T(\tau_n)] = \text{ext}_T(\tau_n)$. \square

In Section 3.1 we will show precisely how to build $T(\tau_n)$ in polynomial time, prove that the construction is semantically correct, and establish that the size of $T(\tau_n)$ is purely linear in the size of T and (τ_n) . Let us illustrate for now the issues with an example. Observe, for instance, that for the tree $T = s_0(a(b)\mathbf{f}_1a(c))$, no matter which type τ_1 is, there is no \mathcal{R} -DTD-typing expressing the language $\text{ext}_T(\tau_1)$. Indeed, this is even the case for $T = s_0(a(b)a(c))$ with no function at all. If we consider the tree $T = s_0(a(\mathbf{f}_1)a(\mathbf{f}_2))$, then the typing $[\tau_1] = \{s_1(b)\}$, $[\tau_2] = \{s_2(c)\}$ prohibits that $\text{ext}_T(\tau_1, \tau_2)$ is expressible by an \mathcal{R} -DTD-type because $[T(\tau_n)] = \{s_0(a(b)a(c))\}$ entailing that the content model of b is non-regular;

while the typing $[\tau_1] = \{s_1(b)\}$, $[\tau_2] = \{s_2(b)\}$ allows that, because $[T(\tau_n)] = \{s_0(a(b)a(b))\}$ entailing that all the content models of s_0 , a and b are regular languages, $\{aa\}$, $\{b\}$ and \emptyset , respectively. Such situations motivated Definition 9.

Before concluding this section, we adapt the previous definitions to strings in the straightforward way. (We will often use reductions to strings problems in the paper.) Let $w(\mathbf{f}_n) = w_0\mathbf{f}_1w_1 \dots \mathbf{f}_nw_n$ be a *kernel string*. For typing strings, we use \mathcal{R} -types where $\mathcal{R} \in \{\mathbf{nFA}, \mathbf{dFA}, \mathbf{nRE}, \mathbf{dRE}\}$. A *typing* for $w(\mathbf{f}_n)$ is still a positional mapping from the functions in (\mathbf{f}_n) to a sequence (τ_n) of \mathcal{R} -types. By $\text{ext}_w(\tau_n)$ we still denote the string language consisting of all possible extensions of w , and by $w(\tau_n)$ the \mathbf{nFA} (or \mathbf{nRE}) constructed from w and (τ_n) is such a way that $[w(\tau_n)] = \text{ext}_w(\tau_n)$.

We will use in our proofs a generalization to “Boxes”. A *kernel box* $B(\mathbf{f}_n) = B_0\mathbf{f}_1B_1 \dots \mathbf{f}_nB_n$ is, here, a finite regular language over $(\Sigma \cup \Sigma^f)$ where $\mathbf{f}_1, \dots, \mathbf{f}_n$ are as above, and each B_i is a box (of a fixed width) over Σ . With $B(\tau_n)$ we denote the \mathbf{nFA} (or \mathbf{nRE}) constructed from B and (τ_n) is such a way that $[B(\tau_n)] = \text{ext}_B(\tau_n)$.

2.4. The Typing Problems

In this section, we introduce the notion of distributed XML design, define the design problems that are central to the present work, and give the overview of the complexity results. We consider two different approaches, *bottom-up* and *top-down*, according to whether the distributed design, other than a *kernel tree*, consists of a *typing* or a *target type*, respectively.

Definition 10. Let \mathcal{S} be a schema language, and $T_{[\mathbf{f}_1, \dots, \mathbf{f}_n]}$ be a kernel document. We call \mathcal{S} -**design** (or just design) one of the following:

- $D = \langle (\tau_n), T_{[\mathbf{f}_1, \dots, \mathbf{f}_n]} \rangle$ where (τ_n) is an \mathcal{S} -typing. This is **bottom-up design**.
- $D = \langle \tau, T_{[\mathbf{f}_1, \dots, \mathbf{f}_n]} \rangle$ where τ is a (target) \mathcal{S} -type. This is **top-down design**. \square

Intuitively, given a bottom-up design, one would like to find a global type that captures the typing of the global document. On the other hand, given a top-down design, one would like to find types for the local documents that will guarantee the global type.

With the following definition, we start the bottom-up analysis. Notice that the concepts used for bottom-up design will be also useful when we consider top-down design.

Definition 11. Given an \mathcal{S} -design $D = \langle (\tau_n), T \rangle$, the \mathcal{S} -typing (τ_n) is \mathcal{S} -**consistent** with T (simply consistent when \mathcal{S} is understood) if there exists an \mathcal{S} -type τ such that $[\tau] = \text{ext}_T(\tau_n)$, in other words, if $\text{ext}_T(\tau_n)$ is definable by some \mathcal{S} -type. This problem (deciding whether an \mathcal{S} -typing is \mathcal{S} -consistent with a kernel tree) is called $\text{CONS}_{[\mathcal{S}]}$. \square

We will denote by $\text{type}_{T,\mathcal{S}}(\tau_n)$, or $\text{type}_T(\tau_n)$ when \mathcal{S} is understood, the \mathcal{S} -type when it exists such that $[\text{type}_T(\tau_n)] = \text{ext}_T(\tau_n)$. Notice that if both \mathcal{S} and T are fixed, then $\text{type}_T(\tau_n)$ plays the role of a function from the set of all possible \mathcal{S} -typings of length n to a set of certain \mathcal{S} -types. According to every possible decision-answer of $\text{CONS}_{[\mathcal{S}]}$ (where T is now fixed), such a function might be always definable, never, or only for some \mathcal{S} -typing. Finally, the complexity of deciding $\text{CONS}_{[\mathcal{S}]}$ or computing $\text{type}_T(\tau_n)$ (with an estimation, w.r.t. $T(\tau_n)$, of its possible size), may vary considerably due to \mathcal{S} .

Table 2 summarizes the complexity results of $\text{CONS}_{[\mathcal{S}]}$. We vary \mathcal{S} among \mathcal{R} -DTDs, \mathcal{R} -SDTDs and \mathcal{R} -EDTDs, for various kinds of \mathcal{R} . In all cases but **dRE**, we get tight bounds. For DTDs and SDTDs with **dRE**, we provide nonmatching lower and upper bounds. The table also shows the size that $\text{type}_T(\tau_n)$ may have in the worst case. Again this is given precisely for all cases but **dRE**. For DTDs and SDTDs with **dRE**, we provide nonmatching bounds.

In the next sections, we systematically analyze the complexity of this problem by varying \mathcal{S} among \mathcal{R} -DTDs, \mathcal{R} -SDTDs and \mathcal{R} -EDTDs, and we will consider $\text{type}_T(\tau_n)$ for each of these schema languages. We next give an example to illustrate some of the main concepts introduced.

Table 2: Complexity results of $\text{CONS}_{[\mathcal{S}]}$ compared with the worst-case-optimal size of $\text{type}_T(\tau_n)$ with respect to $m = \|T(\tau_n)\|$.

	-DTDs	-SDTDs	-EDTDs
nFA	PSPACE -complete $\Theta(m)$	PSPACE -complete $\Theta(m)$	DTIME ($\mathcal{O}(1)$) $\Theta(m)$
nRE	PSPACE -complete $\Theta(m)$	PSPACE -complete $\Theta(m)$	DTIME ($\mathcal{O}(1)$) $\Theta(m)$
dFA	PSPACE -complete $\Theta(2^m)$	PSPACE -complete $\Theta(2^m)$	DTIME ($\mathcal{O}(1)$) $\Theta(m^2)$
dRE	PSPACE -hard \rightsquigarrow EXPTIME $\Omega(2^m) \rightsquigarrow \mathcal{O}(2^{2^m})$	PSPACE -hard \rightsquigarrow EXPTIME $\Omega(2^m) \rightsquigarrow \mathcal{O}(2^{2^m})$	DTIME ($\mathcal{O}(1)$) $\Theta(m)$

Example 1. Consider the kernel $T = s_0(a \mathbf{f}_1 c \mathbf{f}_2)$ and the pair $\tau_1 = \langle \{s_1, b\}, \pi_1, s_1 \rangle$ and $\tau_2 = \langle \{s_2, d\}, \pi_2, s_2 \rangle$ of **dRE**-DTD-types, with $\pi_1(s_1) = b^*$ and $\pi_2(s_2) = d^*$. The activation of both \mathbf{f}_1 and \mathbf{f}_2 may return trees $s_1(bb)$ and $s_2(d)$, respectively. These trees can be plugged into T producing the extension $s_0(abbcd)$. The tree language obtained by considering each possible extension of T is $\text{ext}_T(\tau_1, \tau_2) = \{s_0(ab^n cd^m) : n, m \geq 0\}$. Now, we have:

$$\text{type}_T(\tau_1, \tau_2) = \langle \{s_0, a, b, c, d\}, \pi, s_0 \rangle$$

where $\pi(s_0) = a b^* c d^*$ and all the other element names other than s_0 are leaves. Finally, (τ_1, τ_2) is **dRE-DTD-consistent** with T . \square

We now define the top-down design problems. But before, we introduce some straightforward notation. Let τ and τ' be two types. We say that:

- $\tau \equiv \tau'$ (equivalent) iff $[\tau] = [\tau']$
- $\tau \leq \tau'$ (smaller or equivalent) iff $[\tau] \subseteq [\tau']$
- $\tau < \tau'$ (smaller) iff $[\tau] \subset [\tau']$

and also that, given two typings (τ_n) and (τ'_n) :

- $(\tau_n) \equiv (\tau'_n)$ iff $\tau_i \equiv \tau'_i$ for each i
- $(\tau_n) \leq (\tau'_n)$ iff $\tau_i \leq \tau'_i$ for each i
- $(\tau_n) < (\tau'_n)$ iff $(\tau_n) \leq (\tau'_n)$ and $\tau_i < \tau'_i$ for some i

Definition 12. Given an \mathcal{S} -design $D = \langle \tau, T \rangle$, we say that a typing (τ_n) is:

- **sound** if $\text{ext}_T(\tau_n) \subseteq [\tau]$;
- **maximal** if it is sound, and there is no other sound typing (τ'_n) s.t. $(\tau_n) < (\tau'_n)$;
- **complete** if $\text{ext}_T(\tau_n) \supseteq [\tau]$;
- **local** if $\text{ext}_T(\tau_n) = [\tau]$, namely if it is both sound and complete;
- **perfect** if it is local, and $(\tau'_n) \leq (\tau_n)$ for each other sound typing (τ'_n) ;
- **D -consistent** if it is an \mathcal{S} -typing which is \mathcal{S} -consistent as well. \square

Remark 1. It should be clear that for a given \mathcal{S} -design $D = \langle \tau, T \rangle$ we could have sound typings that are not D -consistent. But, note that, it is even possible to have a sound typing where $T(\tau_n)$ does not define a regular tree language. Consider the design D where $T = s_0(\mathbf{f}_1)$ and $\tau = s_0(a^+b^+)$. Clearly, the typing $[\tau_1] = \{s_1(a^n b^n) : n > 0\}$ is sound but $[T(\tau_1)]$ is not regular. Anyway, we prove in Section 6.1 (for strings, but the results generalize to trees due to our reductions) that if an \mathcal{S} -design admits a sound typing (τ_n) , then it also admits a sound **nFA-EDTD**-typing (τ'_n) such that $(\tau_n) \leq (\tau'_n)$.

Also, by definition of maximality, note that for instance, for a given **dRE-DTD**-design D , a **dRE-DTD**-typing (τ_n) is not maximal even if there is a sound **nFA-DTD**-typing (τ'_n) for D such that $(\tau_n) < (\tau'_n)$. One could have some objection to such a definition. Anyway, Martens et al. [31] proved that whenever the illustrated situation happens, then there is also a **dRE-DTD**-typing (τ''_n) such that $(\tau_n) < (\tau''_n)$. \square

Clearly, local typings present the advantage of allowing a local verification of document consistency (soundness and completeness by definition). Also, no consistent document is ruled out (completeness). Maximal locality guarantees that in some sense, no unnecessary constraints are imposed to the participants. Finally, perfect typings are somehow the ultimate one can expect in terms of not imposing constraints to the participants. Many designs will not accept a perfect typing. However, there are maximal sound typings which are not local.

This is not surprising as there are designs that have at least a sound typing but do not allow any local at all, and clearly, if there is a sound typing, then there must also exist a maximal sound one. We will see examples that separate these different classes further. But before, we make an observation on D -consistency and formally state the problems studied in the paper.

Let \mathcal{S} be any schema language among \mathcal{R} -DTDs, \mathcal{R} -SDTDs, and \mathcal{R} -EDTDs, where $\mathcal{R} \in \{\mathbf{nFA}, \mathbf{dFA}, \mathbf{nRE}, \mathbf{dRE}\}$. Whenever we consider a top-down \mathcal{S} -design $D = \langle \tau, T \rangle$, we require that a typing (τ_n) for D has to be D -consistent, namely both $T(\tau_n)$ is \mathcal{S} -consistent (it has an equivalent \mathcal{S} -type) and each τ_i is an \mathcal{S} -type. In order to verify such a condition, we can exploit the techniques that we have developed for bottom-up design. In particular, it is not hard to see that if (τ_n) is not \mathcal{S} -consistent, then it can not be local. Thus, our approach aims at isolating problems concerning locality from those concerning consistency.

Definition 13. $\text{LOC}_{[\mathcal{S}]}$, $\text{ML}_{[\mathcal{S}]}$, $\text{PERF}_{[\mathcal{S}]}$ are the following decision problems. Given an \mathcal{S} -design $D = \langle \tau, T \rangle$ and a D -consistent typing (τ_n) , is (τ_n) a local, or maximal local, or perfect typing for D , respectively? \square

Definition 14. $\exists\text{-LOC}_{[\mathcal{S}]}$, $\exists\text{-ML}_{[\mathcal{S}]}$, $\exists\text{-PERF}_{[\mathcal{S}]}$ are the following decision problems. Given an \mathcal{S} -design $\langle \tau, T \rangle$, does there exist a local, or maximal local, or perfect D -consistent typing for this design, respectively? \square

We similarly define the corresponding *word* problems (\mathcal{S} is simply \mathcal{R}). We have $\text{LOC}_{[\mathcal{R}]}$, $\text{ML}_{[\mathcal{R}]}$, $\text{PERF}_{[\mathcal{R}]}$, $\exists\text{-LOC}_{[\mathcal{R}]}$, $\exists\text{-ML}_{[\mathcal{R}]}$ and $\exists\text{-PERF}_{[\mathcal{R}]}$. Finally, we will use in proofs box versions of the problems, $\exists\text{-LOC}_{[\mathcal{R}]}^B$, $\exists\text{-ML}_{[\mathcal{R}]}^B$ and $\exists\text{-PERF}_{[\mathcal{R}]}^B$.

Remark 2. In this paper, although we analyze all the three defined schema languages (\mathcal{R} -DTDs, \mathcal{R} -SDTDs, and \mathcal{R} -EDTDs) for top-down designs, after providing reductions from trees to strings, we specialize the analysis to the case of $\mathcal{R} = \mathbf{nFA}$. More tractable problems may be obtained by considering deterministic content models or restricted classes of regular expressions [18, 27] as made by Martens et al. [31]. Also, notice that we pay more attention to maximum locality rather than to maximality proper. In fact, for the latter notion, the existence problem is trivial. Moreover, the complexity of the verification problem essentially coincides for both notions. Nevertheless, one could be interested in a maximal sound typing when, for some reason, the design can not be improved and does not admit any local typing. There could be even cases where a local typing does not exist but, there is a unique maximal sound typing comprising any other possible sound typing, a sort of quasi-perfect typing. For instance, the design $T = s(a \mathbf{f}_1)$ and $\tau = s(ab^* + d)$ has such a property. Our techniques can be easily adapted to these cases, too. \square

Table 3 gives an overview of the complexity results for the typings problems previously defined. We will see in Section 4 that, for \mathcal{R} -DTDs and \mathcal{R} -SDTDs, each problem on trees is logspace-reducible to a set of problems on strings (thus, it suffices to prove the results in Table 3 for words) and that, for \mathcal{R} -EDTDs, the problems on trees depend on the problems on boxes in a more complex manner.

In particular, row D includes two problems that are actually the same (they only differ if $\mathcal{R} = \mathbf{dREs}$, as shown in Martens et al. [31]). Each number in brackets refers either to the corresponding statement/proof in the paper (if rounded) or the paper where the particular result has already been proved (if squared).

Table 3: Complexity results in case of top-down design

		[1] nFAs / nFA-DTDs / nFA-SDTDs	[2] nFA-EDTDs
[A]	LOC	PSPACE -complete (5.3)	EXPTIME -complete (4.19)
[B]	ML	in PSPACE [31] \vee (7.1) PSPACE -hard (5.2)	EXPTIME -hard (4.18) in 2- EXSPACE (7.10)
[C]	PERF	PSPACE -complete (6.7)	EXPTIME -hard (4.18) in coNEXPTIME (7.9)
[D]	$\exists\text{-LOC}/\exists\text{-ML}$	PSPACE -hard (5.4) in EXSPACE [31] \vee (6.11)	EXPTIME -hard (4.9) in 2- EXSPACE (7.4)
[E]	$\exists\text{-PERF}$	PSPACE -complete (6.8)	EXPTIME -hard (4.9) in coNEXPTIME (7.8)

We now present examples that separate the different design properties of typings.

Example 2. Let $\tau = \langle \{s, a, b, c\}, \pi, s \rangle$ be an **nRE**-DTD where $\pi(s) = a^*bc^*$, and $T = s(\mathbf{f}_1\mathbf{f}_2)$ be a kernel tree. It is easy to see that both $s_1(a^*bc^*)$, $s_2(c^*)$ and $s_1(a^*)$, $s_2(a^*bc^*)$ are *local typings* as $a^*bc^*c^* \equiv a^*a^*bc^* \equiv a^*bc^*$. In fact, they are also *maximal local typings*, and so there is no perfect typing for this design. Observe that, for instance, $s_1(a?)$, $s_2(a^*bc^*)$ is still a *local typing* that, however, is not maximal because it imposes unnecessary constraints to the local sites. If desired, one could leave them more freedom, e.g., type the first function with a^* . \square

Example 3. Let $\tau = s(a^*bc^*)$ be a type and $T = s(\mathbf{f}_1b\mathbf{f}_2)$ be a kernel tree. The typing $s_1(a^*)$, $s_2(c^*)$ is *perfect*. This has to be an excellent typing since there is no alternative maximal local typing. \square

Example 4. Let $\tau = (ab)^*$ be a type and $T = s(\mathbf{f}_1\mathbf{f}_2)$ be a kernel tree. The typing $s_1((ab)^*)$, $s_2((ab)^*)$ is a unique maximal local but it is not perfect. Consider, in fact, typing $s_1(a)$, $s_2(b)$. It is sound but $(a, b) \leq ((ab)^*, (ab)^*)$ does not hold. Clearly, a perfect typing cannot exist. \square

Example 5. Let $\tau = (ab)^+$ be a type and $T = s(\mathbf{f}_1\mathbf{f}_2)$ be a kernel tree. There are three maximal local typings:

$$s_1((ab)^*), s_2((ab)^+) \quad s_1((ab)^*a), s_2(b(ab)^*) \quad s_1((ab)^+), s_2((ab)^*)$$

according to whether either s_1 , $s_1(a)$, or none of them may belong to each possible $\text{ext}(\mathbf{f}_1)$, respectively. \square

The following theorem completes the comparison of the properties of typing we study. (Note that its converse is not true by Example 4.)

Theorem 2.1. *Every perfect typing is unique maximal local.*

Proof. Consider a perfect typing (τ_n) for $T(\mathbf{f}_n)$ and τ . We observe that (τ_n) is local, by definition. Moreover, by definition, for each other sound (so also local) typing (τ'_n) , we have $(\tau'_n) \leq (\tau_n)$. The typing, (τ_n) is maximal because there exists no other sound typing (τ''_n) such that $(\tau_n) < (\tau''_n)$, and it is unique because there is no another local typing (τ''_n) such that for some index i and some string w , then $w \in [\tau''_i]$ but $w \notin [\tau_i]$. \square

3. Bottom-up design

In this section, we consider bottom-up design.

3.1. \mathcal{R} -EDTDs typing

Let $T(\mathbf{f}_n)$ be a kernel and (τ_n) be an \mathcal{R} -EDTD-typing where each $\tau_i = \langle \Sigma_i, \tilde{\Sigma}_i, \pi_i, \tilde{s}_i, \mu_i \rangle$. We next present the construction of $T(\tau_n)$, that (to be as general as possible) is an **nFA**-EDTD. We use the following notations:

1. Σ_0 contains the element names in T (the labels but not the functions);
2. $\tilde{\Sigma}_0$ contains a specialized element name \tilde{a}_0^x , for each $a \in \Sigma_0$ and each node x of T with label a .
3. s_0 is the root of T ;
4. s_i is the root of trees in $[\tau_i]$ for each i .

We also make without loss of generality the following assumptions:

1. $\tilde{\Sigma}_i \cap \tilde{\Sigma}_j = \emptyset$, for each $i, j, i \neq j$. (Note that $\Sigma_i \cap \Sigma_{j \neq i}$ may be nonempty.)

Consider the **nFA**-EDTD $T(\tau_n) = \langle \Sigma, \tilde{\Sigma}, \pi, \tilde{s}_0, \mu \rangle$ defined as follows:

1. $\Sigma = \Sigma_0 \cup (\Sigma_1 - \{s_1\}) \cup \dots \cup (\Sigma_n - \{s_n\})$;
2. $\tilde{\Sigma} = \tilde{\Sigma}_0 \cup (\tilde{\Sigma}_1 - \{\tilde{s}_1\}) \cup \dots \cup (\tilde{\Sigma}_n - \{\tilde{s}_n\})$;
3. $\tilde{s}_0 = \tilde{s}_0^x$, where x is the root of T ;
4. $\mu(\tilde{a}) = a$ for each $\tilde{a} \in \tilde{\Sigma}$;
5. $\pi(\tilde{a}_0^x) = \mathbf{nFA}(\{\varepsilon\})$ for each leaf-node x of T with label $a \in \Sigma_0$;
6. $\pi(\tilde{a}_i) = \mathbf{nFA}([\pi_i(\tilde{a}_i)])$ for each $\tilde{a}_i \in \tilde{\Sigma}$ with i in $[1..n]$
7. for each node x of T with label a and children $y_1 \dots y_p$, we define $\pi(\tilde{a}_0^x) = \mathbf{nFA}(L_1 \dots L_p)$ where each language L_k is
 - $\{\tilde{b}_0^{y_k}\}$ if y_k has label $b \in \Sigma$;
 - $[\pi_i(\tilde{s}_i)]$ if y_k is labeled by f_i .

The previous algorithm clearly runs in polynomial time by scanning the tree T and performing some easy regular language manipulation. Also, the size of $T(\tau_n)$ is linear in the size of the input pair T and (τ_n) . This is clearly true for $\mathcal{R} \in \{\mathbf{nFA}, \mathbf{dFA}\}$ where only a linear number of ε -transitions is required. If $\mathcal{R} \in \{\mathbf{nRE}, \mathbf{dRE}\}$, it is also true because the translation from regular expressions to \mathbf{nFA} s produce at most an $n \log^2 n$ blow up but because in these cases we might define $T(\tau_n)$ directly as an \mathbf{nRE} -EDTD-type of actual linear size. These considerations immediately yield the following proposition:

Proposition 3.1. *Given $T(\mathbf{f}_n)$ and (τ_n) , the \mathbf{nFA} -EDTD-type $T(\tau_n)$ can be constructed in polynomial time, and its size is linear in the input pair.*

Now we prove that our construction preserves the semantics of $\text{ext}_T(\tau_n)$.

Theorem 3.2. *Given a kernel $T(\mathbf{f}_n)$ and an \mathcal{R} -EDTD-typing (τ_n) , $[T(\tau_n)] = \text{ext}_T(\tau_n)$ holds for each possible \mathcal{R} .*

Proof. By construction of $T(\tau_n)$, we assume the specialized element names in each type τ_i of (τ_n) to be different (in fact, they could always be renamed appropriately before building $T(\tau_n)$). Also, the specialized element names added for giving witnesses to the nodes of T labeled with an element name belong to a fresh set (it is $\tilde{\Sigma}_0$). This means that there is no “competition” among all of these witnesses. So we just create new content models that exactly allow all and only the trees being valid for each τ_i and the non-function nodes that are already in T . But this is exactly the semantic definition of $\text{ext}_T(\tau_n)$. \square

Corollary 3.3. *All the problems $\text{CONS}_{[\mathbf{nFA}\text{-EDTD}]}$, $\text{CONS}_{[\mathbf{dFA}\text{-EDTD}]}$, $\text{CONS}_{[\mathbf{nRE}\text{-EDTD}]}$, and $\text{CONS}_{[\mathbf{dRE}\text{-EDTD}]}$ always have a yes answer. Thus, they are decidable in constant time.*

Proof. For $\text{CONS}_{[\mathbf{nFA}\text{-EDTD}]}$, $\text{CONS}_{[\mathbf{dFA}\text{-EDTD}]}$, and $\text{CONS}_{[\mathbf{nRE}\text{-EDTD}]}$ the decision-answer is always “**yes**” because each content model in $T(\tau_n)$ is, respectively, already an \mathbf{nFA} , expressible by a \mathbf{dFA} , and expressible by an \mathbf{nRE} .

For $\text{CONS}_{[\mathbf{dRE}\text{-EDTD}]}$ the decision-answer is always “**yes**” as well, but the reason is less obvious. In general, there are regular languages not expressible by \mathbf{dRE} s. Anyway, in our case, by considering how π is built in $T(\tau_n)$, we are sure that each content model has an equivalent \mathbf{dRE} . In fact, $\pi(\tilde{a}_0^x) = \varepsilon$ (step 4) is already a \mathbf{dRE} ; $\pi(\tilde{a}_i) = \mathbf{nFA}([\pi_i(\tilde{a}_i)])$ (step 5) has an equivalent \mathbf{dRE} because $\pi_i(\tilde{a}_i)$ is already a \mathbf{dRE} by definition; $\pi(\tilde{a}_0^x) = \mathbf{nFA}(L_1 \dots L_p)$ (step 6) is expressible by a \mathbf{dRE} because each L_k originates itself from a \mathbf{dRE} and does not share any symbol with any $L_{j \neq k}$. \square

By Corollary 3.3, we now give a safe and easy construction of $\text{type}_T(\tau_n)$ from $T(\tau_n)$ according to the schema language \mathcal{S} used for (τ_n) .

- For \mathbf{nFA} -EDTDs, we choose $\text{type}_T(\tau_n) = T(\tau_n)$;

- For **dFA-EDTDs**, we modify $T(\tau_n)$ by computing the ε -closure for each content model. Notice that this can be done in polynomial time and the size of $\text{type}_T(\tau_n)$ is at most quadratic (and there are cases where this could really happen) in the size of $T(\tau_n)$ because each content model originates from **dFAs** that do not share any symbol.
- For **nRE-EDTDs** or **dRE-EDTDs**, we modify the content models of $T(\tau_n)$ as follows: each $\pi(\tilde{a}_i) = \pi_i(\tilde{a}_i)$ and each $\pi(\tilde{a}_0^x) = R_1 \dots R_p$ where the generic R_k is either $\tilde{b}_0^{y_k}$ or $\pi_i(\tilde{s}_i)$ (compare with the $T(\tau_n)$ definition). Also here the size of $\text{type}_T(\tau_n)$ is linear in the size of $T(\tau_n)$ due to the previous corollary.

3.2. \mathcal{R} -SDTDs typing

For \mathcal{R} -SDTDs we also use $T(\tau_n)$ as defined for \mathcal{R} -EDTDs because any \mathcal{R} -SDTD can be seen as a special \mathcal{R} -EDTD and the algorithm for building $T(\tau_n)$ still works with no problem. At this point, it should be clear that $T(\tau_n)$ can easily not be an \mathcal{R} -SDTD because of our assumptions ($\tilde{\Sigma}_i \cap \tilde{\Sigma}_{j \neq i} = \emptyset$). But, in this case it is also possible that $T(\tau_n)$ does not have an equivalent \mathcal{R} -SDTD. Indeed, $T(\mathbf{f}_n)$ may contain some pattern that already prohibits obtaining an \mathcal{R} -SDTD for any possible typing (τ_n) , or it may contain a function-layout that prohibits obtaining an \mathcal{R} -SDTD for some (τ_n) . So, we have to discriminate when this is possible or not. Such a problem (deciding whether an \mathcal{R} -EDTD has an equivalent \mathcal{R} -SDTD) is in general (when \mathcal{R} stands for **nREs** or **nFAs**) an **EXPTIME**-complete problem [29]. Nevertheless, we will show that, in our case, it is almost always “easier” and in particular that, in general, it depends on the complexity of equivalence between tree-languages specified by **nFA-SDTDs**, which, in turn depends on the complexity of equivalence between string-languages specified by **nFAs**. Before giving proofs of that, we illustrate the definition of distributed document using **nRE-SDTD**-types.

Example 6. Let $T = s_0(\mathbf{f}_1 \ a(b \ \mathbf{f}_2) \ c)$ be a kernel tree and τ_1, τ_2 be two **nRE-SDTD**-types describing respectively $b \cdot d^+ \cdot a(b^+)^*$ and b^* . In the **nRE-SDTDs** syntax, $\tau_1 = \langle \{s_1, a, b, d\}, \{\tilde{s}_1, \tilde{a}_1, \tilde{b}_1, \tilde{d}_1\}, \pi_1, \tilde{s}_1, \mu_1 \rangle$ and $\tau_2 = \langle \{s_2, b\}, \{\tilde{s}_2, \tilde{b}_2\}, \pi_2, \tilde{s}_2, \mu_2 \rangle$ two types where:

- $\Sigma_0 = \{s_0, a, b, c\}$; $\tilde{\Sigma}_0 = \{\tilde{s}_0^1, \tilde{a}_0^3, \tilde{b}_0^4, \tilde{c}_0^6\}$, where $\{1, 3, 4, 6\}$ are the nodes of T with label in Σ_0 (based on a preorder traversal of T);
- $\pi_1(\tilde{s}_1) = \tilde{b}_1 \cdot \tilde{d}_1^+ \cdot \tilde{a}_1^*$; $\pi_1(\tilde{a}_1) = \tilde{b}_1^+$; $\pi_2(\tilde{s}_2) = \tilde{b}_2^*$; $\pi_1(\tilde{b}_1) = \pi_1(\tilde{d}_1) = \pi_2(\tilde{b}_2) = \varepsilon$;
- μ_1 and μ_2 are clear;

For instance, the activation of both \mathbf{f}_1 and \mathbf{f}_2 may return trees $s_1(bda(bbb))$ and $s_2(bb)$, respectively. In general, the resulting type is $s_0(b \cdot d^+ \cdot a(b^+)^+ \cdot c)$. It can be described by an **nRE-SDTDs**. Thus, (τ_1, τ_2) is an **nRE-SDTDs**-typing consistent with T . \square

Now we need to introduce some definitions and mention previous results.

Lemma 3.4. *Let $\tau = \langle \Sigma, \tilde{\Sigma}, \pi, \tilde{s}, \mu \rangle$ be an \mathcal{R} -SDTD. For each $\tilde{a} \in \tilde{\Sigma}$, also $\tau(\tilde{a}) = \langle \Sigma, \tilde{\Sigma}, \pi, \tilde{a}, \mu \rangle$ is.*

Proof. By definition of \mathcal{R} -SDTD (in the worst case, if τ is reduced, then $\tau(\tilde{a})$ may be not.) \square

Definition 15 ([29]). A tree language L is **closed under ancestor-guarded subtree exchange** if the following holds. For each $t_1, t_2 \in L$, and for each x_1, x_2 in t_1, t_2 , respectively, with $\text{anc-str}_{t_1}(x_1) = \text{anc-str}_{t_2}(x_2)$, the trees obtained by exchanging $\text{tree}_{t_1}(x_1)$ and $\text{tree}_{t_2}(x_2)$ are still in L . \square

Lemma 3.5 ([29]). *A tree language is definable by a \mathcal{R} -SDTD iff it is “closed under ancestor-guarded subtree exchange” and each content model is defined by an \mathcal{R} -type.*

Remark 3. Intuitively, this means that the witness associated by an \mathcal{R} -SDTD-type τ to a node x of a tree $t \in [\tau]$ only depends on the string $\text{anc-str}_t(x)$. This is consistent with the definition of $\text{dual}(\tau)$ as a **dFA**. In fact, the (unique) sequence of states that $\text{dual}(\tau)$ scans for recognizing $\text{anc-str}_t(x)$ (except the initial one) exactly gives the unique witness to each node of t in the path from the root to x . \square

Proposition 3.6. [11]

1. There is an equivalent **dRE** for each one-unambiguous regular language;
2. Let \mathcal{A} be a minimum **dFA**. There is an algorithm, that runs in time quadratic in the size of \mathcal{A} , deciding whether $[\mathcal{A}]$ is one-unambiguous;
3. There are one-unambiguous regular languages where the smallest equivalent **dRE** is exponential in the size of the minimum equivalent **dFA**. (This is worst-case optimal;)
4. There are one-unambiguous regular languages where some **nRE** is exponentially more succinct than the smallest equivalent **dRE**. In particular, the language $\{(a+b)^m b(a+b)^n : m \leq n, n > 0\}$ has such a property;
5. The set of all one-unambiguous regular languages is not closed under concatenation.

Corollary 3.7.

- (1) Problem $\text{ONE-UNAMB}_{[\text{nRE}]}$ is in **EXPTIME**.
- (2) For each **nRE** defining a one-unambiguous grammar, there exists an equivalent **dRE** which is, at most, doubly exponential in size. (An exact bound is still open.)
- (3) There are pairs of **dREs** the concatenation of which, by a string separator, defines a one-unambiguous language such that the smallest equivalent **dRE** has an exponential size.

Proof. (1): Let r be an **nRE**. Build, in polynomial time from r , an equivalent **nFA** \mathcal{A} . Run the quadratic-time algorithm described in [11] on the minimum **dFA** (at most exponentially larger) equivalent to \mathcal{A} .

(2): By Proposition 3.6, the **dRE** r' that we construct from the **dFA** \mathcal{A} , introduced in (1), has at most size exponential in the size of \mathcal{A} . Thus, the size of r' is at most doubly exponential in the size of r .

(3): Let $r_1 = (a+b)^m$ and $r_2 = (a+b)^n$ be to **nREs**, with $m \leq n$. By definition, it is clear that they are also both **dREs** linear in n . Consider the new **nRE** $r = r_1 b r_2$. By the previous proposition, r defines a one-unambiguous language but its smallest equivalent **dRE** is exponentially larger. \square

Lemma 3.8 ([7]). *Problem $\text{ONE-UNAMB}_{[\text{nRE}]}$ is PSPACE-hard .*

Definition 16. $\text{CONCAT-UNIV}_{[\mathcal{R}]}$ is the following decision problem. Let Σ be an alphabet. Given two \mathcal{R} -types τ_1 and τ_2 over Σ , is $[\tau_1] \circ [\tau_2] = \Sigma^*$. \square

Lemma 3.9 ([25, 31, 32]). *$\text{CONCAT-UNIV}_{[\mathcal{R}]}$ is PSPACE-complete for each $\mathcal{R} \in \{\text{nFA}, \text{nRE}, \text{dFA}, \text{dRE}\}$.*

After introducing some necessary definitions and results, we are ready for proving the following theorem. It is fundamental for pinpointing the complexity of $\text{CONS}_{[\mathcal{R}\text{-SDTDs}]}$, for giving size-bounds about $\text{type}_T(\tau_n)$ and the guidelines for constructing it.

Theorem 3.10. *Let $T(\mathbf{f}_n)$ be a kernel and (τ_n) be an \mathcal{R} -SDTD-typing.*

1. *If $\mathcal{R} \in \{\text{nFA}, \text{nRE}\}$ (nondeterministic and closed under concatenation), then $\text{CONS}_{[\mathcal{R}\text{-SDTD}]}$ is polynomial-time Turing reducible to $\text{EQUIV}_{[\mathcal{R}\text{-SDTD}]}$ and $\text{type}_T(\tau_n)$ is not larger than $T(\tau_n)$;*
2. *If $\mathcal{R} = \text{dFA}$ (deterministic and closed under concatenation), then problem $\text{CONS}_{[\mathcal{R}\text{-SDTD}]}$ is polynomial-time Turing reducible to $\text{EQUIV}_{[\text{nFA}\text{-SDTD}]}$ and $\text{type}_T(\tau_n)$ has unavoidably a single-exponential blow up w.r.t. $T(\tau_n)$ in the worst case;*
3. *If $\mathcal{R} = \text{dRE}$ (not closed under concatenation), then $\text{CONS}_{[\mathcal{R}\text{-SDTD}]}$ is polynomial-space Turing reducible to $\text{ONE-UNAMB}_{[\text{nRE}]}$. There are cases where the size of $\text{type}_T(\tau_n)$ is, at least, exponential in the size of $T(\tau_n)$. A doubly exponential size is sufficient in the worst case. (The exact bound is still open.)*

Proof. First of all we observe that, by construction, the only content models of $T(\tau_n)$ that might not satisfy the single-type requirement are those related to the witnesses of the non-leaf nodes of T . More formally, let x be any non-leaf node of T the label of which is denoted by a , the content model $\pi(\tilde{a}_0^x)$ of its (unique) witness \tilde{a}_0^x is the only one(s) that may contain some conflict. All other content models either refer to leaves (ε is single-type) or come from some τ_i (that is already single-type.)

Case 1. **PROOF IDEA:** Consider (τ_n) being simply an \mathcal{R} -EDTD. Build $T(\tau_n)$ and (from it) $\text{type}_T(\tau_n) = \langle \Sigma, \tilde{\Sigma}, \pi, \tilde{s}_0, \mu \rangle$ (both in polynomial time as described in Section 3.1) and try to “simplify” the latter (in a bottom-up way starting from the nodes of T having only leaves as children and going on to the root) for satisfying the single-type requirement. If the algorithm does not fail during its run ($\text{CONS}_{[\mathcal{R}\text{-SDTD}]}$ admits a “yes” answer), then the resulting $\text{type}_T(\tau_n)$ is now an \mathcal{R} -SDTD. During the proof we only make use of the “ancestor-guarded subtree

exchange” property, and so, by Lemma 3.5, we can conclude that if we cannot simplify $\text{type}_T(\tau_n)$, then it does not have an equivalent \mathcal{R} -SDTD. Moreover, due to the simplification process that does not change the structure of $T(\tau_n)$ but only *merges* some specialized element names, then the resulting type is at most as large as the original one. Finally, to check the subtree exchange property we only use equivalence between \mathcal{R} -SDTDs and the number of performed steps is clearly polynomial in the size of $T(\tau_n)$ witch, by Proposition 3.1, is polynomial in T and (τ_n) .

More formally, for each node x of T having only leaves as children and of course an element name as label, say a , and for each pair of children $y \neq z$ of x , do:

1. If both y and z are not function nodes and have the same label, say b . As $\pi(\tilde{b}_0^y) = \pi(\tilde{b}_0^z) = \varepsilon$ (by definition) we can consider hereafter, by Lemma 3.5, \tilde{b}_0^y and \tilde{b}_0^z the same element.
2. If only one of the two, say y , has an element name as label, say b , while z has a function as label, say \mathbf{f}_i , and $\pi_i(\tilde{s}_i)$ contains in its specification an element \tilde{b}_i (at most one, as τ_i is already an \mathcal{R} -SDTD), by Lemma 3.5, if $[\text{type}_T(\tau_n, \tilde{b}_i)] = \{b()\}$, then we can consider hereafter, \tilde{b}_i and \tilde{b}_0^y the same element; otherwise we can conclude that $\text{type}_T(\tau_n)$ does not have an equivalent \mathcal{R} -SDTD.
3. Finally, if both y and z are function nodes having label \mathbf{f}_i and \mathbf{f}_j , respectively, for each element name in Σ , say b , if both $\pi_i(\tilde{s}_i)$ and $\pi_j(\tilde{s}_j)$ contain in their specifications the elements \tilde{b}_i and \tilde{b}_j (at most one for each of them, as τ_i and τ_j are already \mathcal{R} -SDTD), by Lemma 3.5, if $[\text{type}_T(\tau_n, \tilde{b}_i)] = [\text{type}_T(\tau_n, \tilde{b}_j)]$ (by construction, this can be done by deciding whether the two \mathcal{R} -SDTDs $\tau_i(\tilde{b}_i)$ and $\tau_j(\tilde{b}_j)$ are equivalent), then we can consider hereafter, \tilde{b}_i and \tilde{b}_j the same element; otherwise if for some \tilde{b}_i and \tilde{b}_j this is not true, we can conclude that $\text{type}_T(\tau_n)$ does not have an equivalent \mathcal{R} -SDTD;

If the corresponding condition is satisfied for each y and z , then we can conclude that $\pi(\tilde{a}_0^x)$ complies with the single-type requirement, that $\text{type}_T(\tau_n, \tilde{a}_0^x)$ has an equivalent \mathcal{R} -SDTD (obtained by applying the previous steps), and that it can be used for checking equivalences when we consider the parent of x , its children and (some modifications of) the three previous steps (see further.)

If $\text{type}_T(\tau_n, \tilde{a}_0^x)$ has an equivalent \mathcal{R} -SDTD for each considered node x , then the next iteration considers each node x' of T having only leaves as children or a node already analyzed. We perform Step 3 exactly as above, while Step 1 or Step 2 with the following trivial changes. Let y be, now, a non-leaf node (instead of a leaf one):

- 1'. If both y and z are not function nodes and have the same label, say b . By Lemma 3.5, if $[\text{type}_T(\tau_n, \tilde{b}_0^y)] = \{b()\}$ we can consider hereafter, \tilde{b}_0^y and \tilde{b}_0^z the same element; otherwise we can conclude that $\text{type}_T(\tau_n)$ does not have an equivalent \mathcal{R} -SDTD;
- 2'. If only one of the two, say y , has an element name as label, say b , while z has a function as label, say \mathbf{f}_i , and $\pi_i(\tilde{s}_i)$ contains in its specification an

element \tilde{b}_i (at most one, as τ_i is already an \mathcal{R} -SDTD), by Lemma 3.5, if $[\text{type}_T(\tau_n, \tilde{b}_i)] = [\text{type}_T(\tau_n, \tilde{b}_0^y)]$, then we can consider hereafter, \tilde{b}_i and \tilde{b}_0^y the same element; otherwise we can conclude that $\text{type}_T(\tau_n)$ does not have an equivalent \mathcal{R} -SDTD.

Finally, if we reach the root of T and after checking equivalences on its children we can conclude that $\pi(\tilde{s}_0)$ complies with the single-type requirement, then $\text{type}_T(\tau_n, \tilde{s}_0) = \text{type}_T(\tau_n)$ is now (after merging the prescribed specialized element names) an \mathcal{R} -SDTD.

Case 2. If $\mathcal{R} = \text{dFA}$ then, when we merge some specialized element names in the same content model, we can obtain an **nFA**. So, we can still invoke the $\text{EQUIV}_{[\text{nFA-SDTD}]}$ problem but the size of $\text{type}_T(\tau_n)$ may be exponential as we want it to be a **dFA-SDTD**, and there are cases for which this may happen already by concatenating two **dFAs** [43]. Given that, the blowup cannot be larger than single-exponential, this bound is optimal.

Case 3. If $\mathcal{R} = \text{dRE}$ then, when we merge some specialized element names in the same content model, we can obtain (due to the concatenation and by Proposition 3.6) an **nRE** that may not be expressible by a **dRE**. We can still invoke the $\text{EQUIV}_{[\text{nRE-SDTD}]}$ problem (as necessary condition) but we also have to invoke the $\text{ONE-UNAMB}_{[\text{nRE}]}$ problem (at least as hard as the first one). Notice that this new check does not compromise the soundness of the algorithm. In fact, for each possible **dRE-SDTD** (if any) equivalent to $\text{type}_T(\tau_n)$ the unique witness that can be assigned to x , due to Lemma 3.5, must define the same language as $\pi(\tilde{a}_0^x)$ by applying μ to them. Finally, if both the two decision problems answer yes, then we can consider a new iteration of the previous algorithm. In case that each content model has an equivalent **dRE** specification and we reach the root of T , we can conclude that $\text{type}_T(\tau_n)$ is now a **dRE-SDTD**. By Proposition 3.6, there are cases where $\text{type}_T(\tau_n)$ may require, at least, single-exponential size. By Corollary 3.7, a doubly exponential size is sufficient in the worst case. \square

We now have the following result:

Corollary 3.11.

- (1) Problems $\text{CONS}_{[\text{nRE-SDTD}]}$ and $\text{CONS}_{[\text{nFA-SDTD}]}$ are **PSPACE-complete**;
- (2) Problem $\text{CONS}_{[\text{dFA-SDTD}]}$ is **PSPACE-complete**;
- (3) Problem $\text{CONS}_{[\text{dRE-SDTD}]}$ is both **PSPACE-hard** and in **EXPTIME**;

Proof. Membership. For (1) and (2) consider that both $\text{EQUIV}_{[\text{nRE-SDTD}]}$ and $\text{EQUIV}_{[\text{nFA-SDTD}]}$ are feasible in **PSPACE** [29]. While for (3) we also consider that $\text{ONE-UNAMB}_{[\text{nRE}]}$ is doable in **EXPTIME**, by Corollary 3.7.

Hardness. For (1) we know that both $\text{EQUIV}_{[\text{nRE-SDTD}]}$ and $\text{EQUIV}_{[\text{nFA-SDTD}]}$ are also **PSPACE-hard** [29].

For (2) and (3) we directly consider a reduction from $\text{CONCAT-UNIV}_{[\mathcal{R}]}$ (**PSPACE-hard**, by Lemma 3.9) to problem $\text{CONS}_{[\mathcal{R-SDTD}]}$ ($\mathcal{R} \in \{\text{dFA}, \text{dRE}\}$). In particular, let A_1, A_2 be two \mathcal{R} -types, we consider the consistency problem for the kernel tree

$T = s(a(\mathbf{f}_1 \mathbf{f}_2) a(\mathbf{f}_3))$ and the \mathcal{R} -SDTDs typing (τ_1, τ_2, τ_3) where the trees in τ_1, τ_2 have only one level other than the root, $\pi_1(\tilde{s}_1) = A_1$, $\pi_2(\tilde{s}_2) = A_2$, and $[\pi_3(\tilde{s}_3)] = \Sigma^*$. It is easy to see that (τ_1, τ_2, τ_3) is consistent with T if and only if $[A_1] \circ [A_2] = \Sigma^*$. \square

We conclude this section with a remark.

Remark 4. The exponential blow-up affecting $\text{type}_T(\tau_n)$ may suggest that there are cases for which it may be better to store an XML document in a distributed manner keeping each part valid w.r.t. its local (and small) type τ_i rather than validate the whole document w.r.t. a very large type. \square

3.3. \mathcal{R} -DTDs typing

Even for \mathcal{R} -DTDs we use $T(\tau_n)$ as defined for \mathcal{R} -EDTDs. But here the algorithm we introduced for \mathcal{R} -EDTDs does not work any more because an \mathcal{R} -DTD-typing is structurally different from an \mathcal{R} -SDTD or an \mathcal{R} -EDTD.

Let T be a kernel and (τ_n) be an \mathcal{R} -DTD-typing. Before building $T(\tau_n)$ we construct, from (τ_n) , an equivalent \mathcal{R} -SDTD-typing (τ'_n) as follows. Let $\tau_i = \langle \Sigma_i, \pi_i, s_i \rangle$ be the i^{th} type in (τ_n) . Consider the \mathcal{R} -SDTD-type $\tau_i = \langle \Sigma_i, \tilde{\Sigma}_i, \pi'_i, \tilde{s}_i, \mu_i \rangle$ defined as follows:

- $\tilde{a} \in \tilde{\Sigma}_i$ iff $a \in \Sigma_i$;
- μ_i is a bijection between $\tilde{\Sigma}_i$ and Σ_i ;
- $\pi'_i(\tilde{a}) = \mu_i^{-1}(\pi_i(a))$.

The two types are trivially equivalent. So, we can build the new **nFA**-EDTD-type (or **nRE**-EDTD-type), representing $\text{ext}_T(\tau_n)$, by using (τ'_n) . But since the overhead of constructing (τ'_n) is completely negligible, we still denote it by $T(\tau_n)$ instead of $T(\tau'_n)$.

Also in this case we would like to decide whether $T(\tau_n)$ has an equivalent \mathcal{R} -DTD-type or not, and even the general problem (when \mathcal{R} stands for **nREs** or **nFAs**) of deciding whether an \mathcal{R} -EDTD has an equivalent \mathcal{R} -DTD is **EXPTIME**-complete [29]. As for \mathcal{R} -SDTD, we will show that in our settings we can do better.

Definition 17 ([35]). A tree language L is **closed under subtree substitution** if the following holds. Whenever for two trees $t_1, t_2 \in L$ with nodes x_1 and x_2 , respectively, $\text{lab}_{t_1}(x_1) = \text{lab}_{t_2}(x_2)$, then the trees obtained, from t_1 and t_2 , by exchanging $\text{tree}_{t_1}(x_1)$ and $\text{tree}_{t_2}(x_2)$ are still in L . \square

Lemma 3.12 ([35]). A tree language is definable by an \mathcal{R} -DTD iff it is “closed under subtree substitution” and each content model is defined by an \mathcal{R} -type.

The following theorem (with the related corollary) concludes the set of results for the bottom-up design problem, and gives the last guidelines for constructing $\text{type}_T(\tau_n)$ or evaluating its size.

Theorem 3.13. Let $T(\mathbf{f}_n)$ be a kernel and (τ_n) be an \mathcal{R} -DTD-typing.

1. If $\mathcal{R} \in \{\mathbf{nFA}, \mathbf{nRE}\}$, then $\text{CONS}_{[\mathcal{R}\text{-DTD}]}$ is polynomial-time Turing reducible to $\text{EQUIV}_{[\mathcal{R}\text{-SDTD}]}$ and $\text{type}_T(\tau_n)$ is linear in $T(\tau_n)$;
2. If $\mathcal{R} = \mathbf{dFA}$, then $\text{CONS}_{[\mathcal{R}\text{-DTD}]}$ is polynomial-time Turing reducible to $\text{EQUIV}_{[\mathbf{nFA}\text{-SDTD}]}$ and $\text{type}_T(\tau_n)$ has unavoidably a single-exponential blow up w.r.t. $T(\tau_n)$ in the worst case;
3. If $\mathcal{R} = \mathbf{dRE}$, then $\text{CONS}_{[\mathcal{R}\text{-DTD}]}$ is polynomial-space Turing reducible to $\text{ONE-UNAMB}_{[\mathbf{nRE}]}$ and there are case where $\text{type}_T(\tau_n)$ is, at least, exponentially larger than $T(\tau_n)$. A doubly exponential size is sufficient in the worst case. (The exact bound is still open.)

Proof. Build the \mathcal{R} -SDTD-typing (τ'_n) from (τ_n) as said before. Perform, from T and (τ'_n) , the decision-algorithm defined in the proof of Theorem 3.10 by enforcing, due to Lemma 3.12, the additional constraint at the end of each macro-step when we assert that $\text{type}_T(\tau_n, \tilde{a}_0^x)$ has an equivalent \mathcal{R} -SDTD:

- $[\mu(\pi(\tilde{a}'))] = [\mu(\pi(\tilde{a}''))]$ for each \tilde{a}', \tilde{a}'' already considered;

Finally, notice that the (polynomial number) additional steps are special cases of calls to $\text{EQUIV}_{[\mathcal{R}\text{-SDTD}]}$, and that the same observations made for \mathcal{R} -SDTDs hold for $\text{type}_T(\tau_n)$ as well. \square

Corollary 3.14. *We have the following results:*

- (1) Problems $\text{CONS}_{[\mathbf{nRE}\text{-DTD}]}$ and $\text{CONS}_{[\mathbf{nFA}\text{-DTD}]}$ are **PSPACE**-complete;
- (2) Problem $\text{CONS}_{[\mathbf{dFA}\text{-DTD}]}$ is **PSPACE**-complete;
- (3) Problem $\text{CONS}_{[\mathbf{dRE}\text{-DTD}]}$ is both **PSPACE**-hard and in **EXPTIME**;

Proof. Membership. As for \mathcal{R} -SDTDs (see Corollary 3.11).

Hardness. For (1) we know that both $\text{EQUIV}_{[\mathbf{nRE}\text{-DTD}]}$ and $\text{EQUIV}_{[\mathbf{nFA}\text{-DTD}]}$ are also **PSPACE**-hard [29, 32].

For (2) and (3) we use the same reduction (from $\text{CONCAT-UNIV}_{[\mathcal{R}]}$) that we have used in Corollary 3.11, where the problem $\text{CONS}_{[\mathcal{R}\text{-SDTD}]}$ is replaced now by $\text{CONS}_{[\mathcal{R}\text{-DTD}]}$. In particular, we just notice that also in this case (τ_1, τ_2, τ_3) is consistent with $T = s(a(\mathbf{f}_1 \mathbf{f}_2) a(\mathbf{f}_3))$ if and only if $[\pi_1(\tilde{s}_1)] \circ [\pi_2(\tilde{s}_2)] = [\pi_3(\tilde{s}_3)] = \Sigma^*$. \square

4. Top-down design

In this section, we consider design problems where we start from a kernel and a given global type, and we show how to reduce each of these problems on trees to a set of typing problems on strings. In the next section, we will show how to solve the problems for strings.

4.1. \mathcal{R} -DTDs

We briefly present some obvious results on equivalence of \mathcal{R} -DTDs. The proof of the next result is obvious and thus omitted.

Proposition 4.1. *Two reduced \mathcal{R} -DTDs τ_1 and τ_2 are equivalent if and only if the following are true:*

1. *They have the same root;*
2. *They use the same element names;*
3. *For each element name a , the content models of a in both are equivalent.*

Theorem 4.2. *Let $D = \langle \tau, T(\mathbf{f}_n) \rangle$ be a distributed design where $\tau = \langle \Sigma, \pi, s \rangle$ is an \mathcal{R} -DTD. The following are equivalent:*

- (1) *D admits a local \mathcal{R} -DTD-typing;*
- (2) *The \mathcal{R} -design $D^x = \langle \pi(\mathbf{lab}(x)), \mathbf{child-str}(x) \rangle$ admits a local \mathcal{R} -typing for each node x in T where $\mathbf{lab}(x) \in \Sigma$.*

Proof. (1) \Rightarrow (2): Let (τ_n) be a local typing for D , then $\mathbf{type}_T(\tau_n) \equiv \tau$ holds. This means (by Proposition 4.1) that for each node x in T such that $\mathbf{lab}(x) \in \Sigma$, the content model $\pi(\mathbf{lab}(x))$ of x has an equivalent specification in $\mathbf{type}_T(\tau_n)$. But this means that the subset of types in (τ_n) in bijection with the functions of $\mathbf{child-str}(x)$ represents a local typing for D^x as well.

(2) \Rightarrow (1): Also in this case, by Proposition 4.1, since each node x such that $\mathbf{lab}(x) \in \Sigma$ has a local typing, this means that such a typing allows describing exactly the content model $\pi(\mathbf{lab}(x))$. Thus, by combining all the local typings of the various string-designs with the content models of τ we obtain a D -consistent typing also local for D . To be more precise, we now show how to exploit the local string-typings for building a local typing for D . First of all we observe that, for each i in $[1..n]$, there exists only one node x of T such that \mathbf{f}_i is in the kernel string $\mathbf{child-str}(x)$ of D^x . Since each D^x admits a local typing, then there is a sequence, say $(\tau_1^{str}, \dots, \tau_n^{str})$, of string-types (one for each function) allowing that. In particular, if for some x , $\mathbf{child-str}(x)$ has no function, then this necessarily means that D^x admits a trivial local typing, namely $[\pi(\mathbf{lab}(x))] = \{\mathbf{child-str}(x)\}$ must hold. Let i be an index in $[1..n]$, and x be the parent of \mathbf{f}_i . The new type (not necessarily reduced) $\tau_i = \langle \Sigma_i, \pi_i, s_i \rangle$ is defined as follows:

- $\Sigma_i = \Sigma \cup \{s_i\}$;
- π_i contains all the rules of π and the extra rule $\pi_i(s_i) = \tau_i^{str}$.

Finally, it is very easy to see that, $T(\tau_n)$ is structurally equivalent to τ . \square

Corollary 4.3. *The problems $\text{LOC}_{[\mathcal{R}\text{-DTD}]}$, $\text{ML}_{[\mathcal{R}\text{-DTD}]}$, $\text{PERF}_{[\mathcal{R}\text{-DTD}]}$, $\exists\text{-LOC}_{[\mathcal{R}\text{-DTD}]}$, $\exists\text{-ML}_{[\mathcal{R}\text{-DTD}]}$ and $\exists\text{-PERF}_{[\mathcal{R}\text{-DTD}]}$ are logspace Turing reducible to $\text{LOC}_{[\mathcal{R}]}$, $\text{ML}_{[\mathcal{R}]}$, $\text{PERF}_{[\mathcal{R}]}$, $\exists\text{-LOC}_{[\mathcal{R}]}$, $\exists\text{-ML}_{[\mathcal{R}]}$ and $\exists\text{-PERF}_{[\mathcal{R}]}$, respectively.*

Proof. Let $D = \langle \langle \Sigma, \pi, s \rangle, T(\mathbf{f}_n) \rangle$ be a top-down \mathcal{R} -DTD-design.

Consider, firstly, the $\exists\text{-LOC}_{[\mathcal{R}\text{-DTD}]}$ problem. Scan T in document order, which is well known to be feasible in logarithmic space [16]. For each node x in T such that $\text{lab}(x) \in \Sigma$, solve the problem $\exists\text{-LOC}_{[\mathcal{R}]}$ for the design D^x .

If we consider, instead, the problem $\text{LOC}_{[\mathcal{R}\text{-DTD}]}$, as (τ_n) is D -consistent, then $\text{type}_T(\tau_n)$ exists and there are no different content models for the same element name. So it is also enough to scan T in document order, and for each node x in T such that $\text{lab}(x) \in \Sigma$, solve the problem $\text{LOC}_{[\mathcal{R}]}$ for the design D^x and the subset of types from (τ_n) in bijection with the functions in D^x .

For the maximal and perfect requirements, as they are specializations of the local requirement, it is enough to observe that, by Theorem 4.2, they only depend on the structure of the various D^x . \square

4.2. \mathcal{R} -SDTDs

Before proving that a similar reduction still holds for \mathcal{R} -SDTDs, we need a proposition and a new definition.

Proposition 4.4. *Let $\tau_1 = \langle \Sigma_1, \tilde{\Sigma}_1, \pi_1, \tilde{s}_1, \mu_1 \rangle$ and $\tau_2 = \langle \Sigma_2, \tilde{\Sigma}_2, \pi_2, \tilde{s}_2, \mu_2 \rangle$ be two reduced \mathcal{R} -SDTDs. If they are equivalent, then for each i, j in $[1..2]$, and each $\tilde{a}_i \in \tilde{\Sigma}_i$ there is $\tilde{a}_j \in \tilde{\Sigma}_j$ such that $\mu_i(\pi_i(\tilde{a}_i)) = \mu_j(\pi_j(\tilde{a}_j))$.*

Proof. As τ_1 is reduced, for each specialized element name $\tilde{a}_1 \in \tilde{\Sigma}_1$ there is a tree $t \in [\tau_1]$ such that its unique witness t' contains at least one node having \tilde{a}_1 as label. So, let us fix \tilde{a}_1 , t , and a node x of t such that $\text{anc-str}(x)$ ends with the element name a . As τ_1 and τ_2 are equivalent, there must exist also a unique witness t'' for t produced by τ_2 . Let us denote by \tilde{a}_2 the specialized element name associated to x in t'' . As $[\tau_1] = [\tau_2]$, if $\mu_1(\pi_1(\tilde{a}_1)) \neq \mu_2(\pi_2(\tilde{a}_2))$, then we would violate the ancestor-guarded subtree exchange property. \square

Definition 18. Let $D = \langle \tau, T(\mathbf{f}_n) \rangle$ be a distributed design where the type $\tau = \langle \Sigma, \tilde{\Sigma}, \pi, \tilde{s}, \mu \rangle$ is an \mathcal{R} -SDTD. For each node x in T such that $\text{lab}(x) \in \Sigma$ we denote by $D^x = \langle \pi(\tilde{a}), w_x \rangle$ the unique string-design induced by D , where \tilde{a} is the (unique) witness assigned by τ to x . Moreover, $w_x = \varepsilon$ if x is a leaf, and it is the string obtained from $\text{children}(x)$ by changing each non-function node with the corresponding (unique) witness assigned by τ , otherwise. \square

Theorem 4.5. *Let $D = \langle \tau, T(\mathbf{f}_n) \rangle$ be a distributed design where the type $\tau = \langle \Sigma, \tilde{\Sigma}, \pi, \tilde{s}, \mu \rangle$ is an \mathcal{R} -SDTD. The following are equivalent:*

- (1) D admits a local \mathcal{R} -SDTD-typing;
- (2) Each \mathcal{R} -design induced by D admits a local \mathcal{R} -typing.

Proof. (1) \Rightarrow (2): Since D admits a local \mathcal{R} -SDTD-typing, say (τ_n) , then (τ_n) is \mathcal{R} -SDTD-consistent with T , and $\text{type}_T(\tau_n) \equiv \tau$ holds. For each node x of T having an element name as label, say a , consider the unique witness associated by τ and $\text{type}_T(\tau_n)$ to x , say \tilde{a}_τ and \tilde{a}_T , respectively. By hypothesis, both $\pi_\tau(\tilde{a}_\tau)$ and $\pi_T(\tilde{a}_T)$ satisfy the single-type requirement, and by Proposition 4.4,

$\mu_\tau(\pi_\tau(\tilde{a}_\tau)) \equiv \mu_T(\pi_T(\tilde{a}_T))$. Thus, as $\pi_\tau(\tilde{a}_\tau)$ and the children of x have a local decomposition induced by (τ_n) , then also $\pi_T(\tilde{a}_T)$ and the children of x (namely D^x) have.

(2) \Rightarrow (1): We show, if the premise is true, how to build (τ_n) in such a way that $\text{type}_T(\tau_n) \equiv \tau$ holds and in particular that, from a structural point of view, $\text{type}_T(\tau_n)$ is equivalent to τ . We denote by $(\tau_1^{str}, \dots, \tau_n^{str})$ one possible sequence of string-types satisfying contemporarily all the string-designs. In particular, if for some node x of T the kernel string of D^x has no function call this necessarily means that D^x admits a trivial local typing, namely $\mu(\pi(\tilde{a})) = \text{child-str}(x)$, where \tilde{a} is the unique witness assigned by τ to x . In particular, for each function \mathbf{f}_i consider its parent node in T , say x . The new type $\tau_i = \langle \Sigma_i, \tilde{\Sigma}_i, \pi_i, \tilde{s}_i, \mu_i \rangle$ is defined as follows:

- $\Sigma_i = \Sigma \cup \{s_i\}$;
- $\tilde{\Sigma}_i = \tilde{\Sigma} \cup \{\tilde{s}_i\}$;
- π_i contains all the rule of π and the extra rule $\pi_i(\tilde{s}_i) = \tau_i^{str}$;
- \tilde{s}_i is the usual extra witness for the root of any tree in $[\tau_i]$;
- μ_i is defined as μ and also $\mu_i(\tilde{s}_i) = s_i$.

It is very easy to see that, if we build $T(\tau_n)$ without renaming the specialized element names we obtain exactly τ , and so $\text{type}_T(\tau_n) \equiv \tau$. In particular, when we assign the witnesses to the non-function nodes of T we choose exactly those assigned by τ . The only difference may be in the specification of the content models because the recomposition after a decomposition may produce a different structure (for instance a different **nFA**) being, anyway, equivalent to the original one. Notice that, the “ancestor-guarded subtree exchange” property is guarantied because we also require that all the designs without any function call admit local typings. For instance, if $D^x = \langle \pi(\tilde{a}), \varepsilon \rangle$ admits a local typing, where x is a leaf node of T and \tilde{a} is its witness assigned by τ , this necessarily means that $\pi(\tilde{a}) = \{\varepsilon\}$, and we automatically take it into account when we build each τ_i . \square

Corollary 4.6. *The problems $\text{LOC}_{[\mathcal{R}\text{-SDTD}]}$, $\text{ML}_{[\mathcal{R}\text{-SDTD}]}$, $\text{PERF}_{[\mathcal{R}\text{-SDTD}]}$, $\exists\text{-LOC}_{[\mathcal{R}\text{-SDTD}]}$, $\exists\text{-ML}_{[\mathcal{R}\text{-SDTD}]}$ and $\exists\text{-PERF}_{[\mathcal{R}\text{-SDTD}]}$ are logspace Turing reducible to $\text{LOC}_{[\mathcal{R}]}$, $\text{ML}_{[\mathcal{R}]}$, $\text{PERF}_{[\mathcal{R}]}$, $\exists\text{-LOC}_{[\mathcal{R}]}$, $\exists\text{-ML}_{[\mathcal{R}]}$ and $\exists\text{-PERF}_{[\mathcal{R}]}$, respectively.*

Proof. Exactly the same as for \mathcal{R} -DTDs. \square

4.3. \mathcal{R} -EDTDs

Although \mathcal{R} -EDTDs have nice properties simplifying the $\text{CONS}_{[\mathcal{R}\text{-EDTD}]}$ problem and the construction of $\text{type}_T(\tau_n)$ (when we start from a kernel and an \mathcal{R} -EDTD-typing), things dramatically change when we consider the problems concerning locality. The freedom of using, in the same content model, various specialized element names for the same element name has a price. Consider the following example.

Example 7. Let $D = \langle \tau, T \rangle$ be a dRE-EDTD-design where $T = s_0(\mathbf{f}_1 \mathbf{f}_2)$ and $\tau = \langle \Sigma, \tilde{\Sigma}, \pi, \tilde{s}_0, \mu \rangle$. In particular, $\pi(\tilde{s}_0) = \tilde{a}^1(\tilde{b}^1)^* + \tilde{a}^2(\tilde{b}^2)^*$; $\pi(\tilde{a}^1) = \tilde{c}^1$; $\pi(\tilde{a}^2) = \tilde{d}^1$; $\pi(\tilde{b}^1) = \tilde{e}^1 + \tilde{g}^1$; $\pi(\tilde{b}^2) = \tilde{g}^1 + \tilde{h}^1$. It is not hard to see that the string-design $\langle \pi(\tilde{s}_0), \mathbf{f}_1 \mathbf{f}_2 \rangle$ admits only two maximal local typings:

$$(\varepsilon, \tilde{a}^1(\tilde{b}^1)^* + \tilde{a}^2(\tilde{b}^2)^*) \quad (\tilde{a}^1(\tilde{b}^1)^* + \tilde{a}^2(\tilde{b}^2)^*, \varepsilon)$$

But, only the first one is also maximal for D , while the actual second one is $(\tilde{a}^1(\tilde{b}^1)^* + \tilde{a}^2(\tilde{b}^2)^*, (\tilde{b}^3)^*)$ where $[\tau_2(\tilde{b}^3)] = b(g)$. \square

The problem highlighted by the previous example originates from the fact that \tilde{b}^1 and \tilde{b}^2 can not be considered completely distinct as \tilde{a}^1 and \tilde{a}^2 (notice that $[\tau(\tilde{a}^1)] \cap [\tau(\tilde{a}^2)] = \emptyset$), and as we naturally do for two different symbols of an alphabet in string languages, yet they are witnesses for two sets of trees with a nonempty intersection. In fact, $[\tau(\tilde{b}^1)] \cap [\tau(\tilde{b}^2)] = b(g)$ can be part of τ_2 in the second maximal local typing for D .

From this, it is unclear whether, by only analyzing content models (such as $\pi(\tilde{s}_0)$, in the previous example), we can decide whether a given design admits at least a local typing. Clearly, if we apply μ to both $(\tilde{a}^1(\tilde{b}^1)^* + \tilde{a}^2(\tilde{b}^2)^*) \cdot (\tilde{b}^3)^*$ and to $\pi(\tilde{s}_0)$ we obtain the same string-language, namely ab^* , but unfortunately, this is only a necessary condition and even if (ab^*, b^*) is a maximal local typing for ab^* , it is not clear how to assign the witnesses for obtaining $(\tilde{a}^1(\tilde{b}^1)^* + \tilde{a}^2(\tilde{b}^2)^*, (\tilde{b}^3)^*)$.

The following theorems, give a further idea of the higher complexity of locality when we consider \mathcal{R} -EDTD-designs.

Theorem 4.7 ([36, 37]). *Problems $\text{EQUIV}_{[\text{nFA-EDTD}]}$ and $\text{EQUIV}_{[\text{nRE-EDTD}]}$ are EXPTIME-complete.*

Theorem 4.8. *Problems $\exists\text{-LOC}_{[\mathcal{R}\text{-EDTD}]}$, $\exists\text{-ML}_{[\mathcal{R}\text{-EDTD}]}$, and $\exists\text{-PERF}_{[\mathcal{R}\text{-EDTD}]}$ are at least as hard as $\text{EQUIV}_{[\mathcal{R}\text{-EDTD}]}$.*

Proof. We define a logspace transformation φ from $\text{EQUIV}_{[\mathcal{R}\text{-EDTD}]}$ to $\exists\text{-LOC}_{[\mathcal{R}\text{-EDTD}]}$. Afterwards, we show that the statement also holds for the other two problems by using exactly the same reduction. Let τ', τ'' be two arbitrary \mathcal{R} -EDTDs. The application of φ to this pair produces the design $D = \langle \tau, T \rangle$, where

- $T = s_0(\mathbf{f}_1 \ c \ \mathbf{f}_2)$
- $\pi(\tilde{s}_0) = \mathcal{R}(\tilde{a}_1 \tilde{c}_1 \tilde{d}_1 + \tilde{b}_1 \tilde{c}_1 \tilde{d}_2)$
- $\pi(\tilde{d}_1) = \mathcal{R}(\tilde{s}'_0)$, where \tilde{s}'_0 is the root of the trees in $[\tau']$
- $\pi(\tilde{d}_2) = \mathcal{R}(\tilde{s}''_0)$, where \tilde{s}''_0 is the root of the trees in $[\tau'']$
- $\pi(\tilde{a}_1) = \pi(\tilde{b}_1) = \pi(\tilde{c}_1) = \mathcal{R}(\varepsilon)$
- \tilde{c}_1 does not appear in any other content model of τ and c appears exactly once in any tree in $[\tau]$

Informally, $[\tau] = s_0(acd([\tau']) + bcd([\tau'']))$. First of all, we observe that all the new content models (other than those being already in τ' and τ'') can be represented by \mathcal{R} -types, even dREs. Now, it is easy to see that D admits a local typing iff $[\tau(\tilde{d}_1)] = [\tau(\tilde{d}_2)]$ iff $\tau' \equiv \tau''$. It is $[\tau_1] = s_1(a + b)$ and $[\tau_2] = s_2(d([\tau']))$. Finally, we just notice that if $\tau' \equiv \tau''$ holds, then (τ_1, τ_2) is the unique maximal local typing for D which is even perfect. \square

Corollary 4.9. *Problems $\exists\text{-LOC}_{[\mathcal{R}\text{-EDTD}]}$, $\exists\text{-ML}_{[\mathcal{R}\text{-EDTD}]}$, and $\exists\text{-PERF}_{[\mathcal{R}\text{-EDTD}]}$ are EXPTIME-hard if $\mathcal{R} \in \{\mathbf{nFA}, \mathbf{nRE}\}$.*

The equivalence between $[\tau(\tilde{d}_1)] = [\tau(\tilde{d}_2)]$, in the previous reduction, is necessary because we do not know, a priori, whether \mathbf{f}_1 is imposing a constraint on \mathbf{f}_2 or not. In particular, this is an extreme case of the fact that $[\tau(\tilde{d}_1)] \cap [\tau(\tilde{d}_2)] \neq \emptyset$.

What we really need is to be able to consider *completely distinct*, in the same content model, each pair of different specialized element names of the form \tilde{a} and \tilde{a}' , namely $[\tau(\tilde{a})] \cap [\tau(\tilde{a}')] = \emptyset$. To do that, given an $\mathcal{R}\text{-EDTD}$, we construct an equivalent \mathbf{nUTA} [30], we transform it into an \mathbf{dUTA} [15], and finally we try to derive a new $\mathcal{R}\text{-EDTD}$ satisfying our requirement. If $\mathcal{R} = \mathbf{dRE}$ the last step could not be always possible.

Given an $\mathcal{R}\text{-EDTD}$ $\tau = \langle \Sigma, \tilde{\Sigma}, \pi, \tilde{s}, \mu \rangle$, an equivalent \mathbf{nUTA} $\mathcal{A} = \langle K, \Sigma, \Delta, F \rangle$ can be constructed as follows: $K = \tilde{\Sigma}$; $\Delta(\tilde{a}, a) = \mathbf{nFA}(\pi(\tilde{a}))$, for each $\tilde{a} \in \tilde{\Sigma}$; $F = \{\tilde{s}_0\}$. Now we want to transform \mathcal{A} into an equivalent \mathbf{dUTA} \mathcal{A}^d (that may be exponential in size). Notice that \mathcal{A}^d will have only one final state as well. Finally, we convert again \mathcal{A}^d (whenever it is possible) into an $\mathcal{R}\text{-EDTD}$ τ^d as follows: $\tilde{\Sigma} = K$; $\pi(\tilde{a}) = \mathcal{R}(\Delta(\tilde{a}, a))$, for each $\tilde{a} \in K$.

Lemma 4.10. *Let τ^d be an $\mathcal{R}\text{-EDTD}$ built as above. For each element name, say a , and each pair \tilde{a}, \tilde{a}' of different specialized element names in $\tilde{\Sigma}^d(a)$, then $[\tau^d(\tilde{a})] \cap [\tau^d(\tilde{a}')] = \emptyset$.*

Proof. It is easy to see that by a (bottom-up) run of \mathcal{A}^d over each tree $t \in [\tau^d(\tilde{a})] \cup [\tau^d(\tilde{a}')]$, there is only one possible state (between \tilde{a} and \tilde{a}') that can be associated to the root of t , and the states of \mathcal{A}^d coincide with the specialized element names of τ^d . \square

Now we are ready for handling $\mathcal{R}\text{-EDTDs}$ (that we call **normalized**) satisfying the above property. But before we introduce a general property of $\mathcal{R}\text{-EDTDs}$.

Proposition 4.11. [29] *Let τ be an $\mathcal{R}\text{-EDTD}$. Whenever for two trees $t_1, t_2 \in [\tau]$ with nodes x_1 and x_2 , respectively, there are witnesses \tilde{t}_1 and \tilde{t}_2 assigning the same specialized element name to both x_1 and x_2 , then the trees obtained, from t_1 and t_2 , by exchanging $\text{tree}_{t_1}(x_1)$ and $\text{tree}_{t_2}(x_2)$ are still in $[\tau]$.*

The following lemma holds for general $\mathcal{R}\text{-EDTDs}$ but is it also useful for normalized $\mathcal{R}\text{-EDTDs}$. Consider the design $D = \langle \tau^d, T \rangle$ where $T = s_0(a(\mathbf{f}_1) \mathbf{f}_2)$ and τ^d is a normalized $\mathbf{nRE}\text{-EDTD}$ having $\pi(\tilde{s}_0) = (\tilde{a}^1 + \tilde{a}^2)^+$ (we ignore the other content models). As $[\tau^d(\tilde{a}^1)] \cap [\tau^d(\tilde{a}^2)] = \emptyset$, it is clear that the unique maximal local typing (τ_1, τ_2) for D has $\pi_1(\tilde{s}_1) = \pi(\tilde{a}^1) + \pi(\tilde{a}^2)$ and $\pi_2(\tilde{s}_2) = (\tilde{a}^1 + \tilde{a}^2)^*$. Thus the node under the root labeled by a may have either \tilde{a}^1 or \tilde{a}^2 as witness depending on the tree replacing \mathbf{f}_1 .

Lemma 4.12. *Let $D = \langle \tau, T \rangle$ be an $\mathcal{R}\text{-EDTD}\text{-design}$ and (τ_n) be a local typing for D . For each node x of T having an element name as label, say a , there is a set of specialized element names $\tilde{\Sigma}^x \subseteq \tilde{\Sigma}(a)$ such that $\bigcup_{\tilde{a} \in \tilde{\Sigma}^x} [\tau(\tilde{a})] = [T(\tau_n, \tilde{a}_0^x)]$.*

Proof. By hypothesis, $[\tau] = [T(\tau_n)]$. We recall that this equivalence is obtained when we consider all the possible extensions of T . Let x be a node of T having an element name as label, say a , and k be the cardinality of $\tilde{\Sigma}(a)$. First of all, we observe that if $\tilde{\Sigma}^x = \tilde{\Sigma}(a)$, then

$$[T(\tau_n, \tilde{a}_0^x)] \subseteq [\tau(\tilde{a}^1)] \cup \dots \cup [\tau(\tilde{a}^k)]$$

is trivially true. In fact, in such a case, the first member has to be necessarily a subset of the second one because, otherwise, in some extension of T there would be a subtree rooted in x which τ cannot produce any witness for.

Starting from $\tilde{\Sigma}^x = \tilde{\Sigma}(a)$ we **claim** that each \tilde{a}^j , with $1 \leq j \leq k$, is either a “friend” to keep in $\tilde{\Sigma}^x$ or an “intruder” to remove from $\tilde{\Sigma}^x$. Finally, the resulting $\tilde{\Sigma}^x$ will prove the statement. Consider now each \tilde{a}^j . We distinguish two cases:

1. We say that \tilde{a}^j is a *friend* if and only if $[\tau(\tilde{a}^j)] \subseteq [T(\tau_n, \tilde{a}_0^x)]$ because it then clearly contributes to prove the statement. We leave it in $\tilde{\Sigma}^x$.
2. We say that \tilde{a}^j is an *intruder* (and we remove it from $\tilde{\Sigma}^x$) if and only if one of the following is true:
 - $[\tau(\tilde{a}^j)] \cap [T(\tau_n, \tilde{a}_0^x)] = \emptyset$ because even if we remove it from $\tilde{\Sigma}^x$, then $[T(\tau_n, \tilde{a}_0^x)] \subseteq \bigcup_{\tilde{a} \in \tilde{\Sigma}^x} [\tau(\tilde{a})]$ is still true.
 - $[\tau(\tilde{a}^j)] \cap [T(\tau_n, \tilde{a}_0^x)] \neq \emptyset$, $[\tau(\tilde{a}^j)] \not\subseteq [T(\tau_n, \tilde{a}_0^x)]$, and for each (sub)tree st_1 in the intersection there is not a tree $t_1 \in [T(\tau_n)]$ having st_1 as subtree rooted in x , such that at least a witness \tilde{t}_1 from τ associates \tilde{a}^j to x . It is still an intruder because for each (sub)tree st_1 in the intersection and for each possible trees $t_1 \in [T(\tau_n)]$ having st_1 as subtree rooted in x , since tree t_1 must necessarily have a witness \tilde{t}_1 from τ , then \tilde{t}_1 associates $\tilde{a}^{i \neq j}$ to x entailing that each st_1 is contained also in $[\tau(\tilde{a}^i)]$. Finally, even if we remove \tilde{a}^j from $\tilde{\Sigma}^x$, then $[T(\tau_n, \tilde{a}_0^x)] \subseteq \bigcup_{\tilde{a} \in \tilde{\Sigma}^x} [\tau(\tilde{a})]$ is still true.
 - $[\tau(\tilde{a}^j)] \cap [T(\tau_n, \tilde{a}_0^x)] \neq \emptyset$, $[\tau(\tilde{a}^j)] \not\subseteq [T(\tau_n, \tilde{a}_0^x)]$, and for some (sub)tree st_1 in the intersection there is a tree $t_1 \in [T(\tau_n)]$ having st_1 as subtree rooted in x and there is a witness \tilde{t}_1 from τ associating \tilde{a}^j to x . Anyway, this last case is not possible because it would contradict the hypothesis that (τ_n) is local. In fact, consider a tree $t_2 \in \tau$ where some node y has also witness \tilde{a}^j and the subtree rooted in y is $st_2 \in [\tau(\tilde{a}^j)] - [T(\tau_n, \tilde{a}_0^x)]$. For each node z in t_1 such that $\text{anc-str}_{t_1}(z) = \text{anc-str}_{t_1}(x)$ and the witness for z in \tilde{t}_1 is \tilde{a}^j , by Proposition 4.11, we may replace (in t_1) each subtree rooted in z with st_2 . The new tree is still in τ but cannot be obtained by any possible extension of T as $[T(\tau_n, \tilde{a}_0^x)]$ contains all the possible trees obtainable by all the extensions of the functions under x and st_2 is not among them.

Thus, if we consider all the feasible cases, then the claim is true as well as the theorem: $\tilde{\Sigma}^x$ is exactly the set of all friends (or a possible subset if some $[\tau(\tilde{a}^j)]$ can be obtained by the union of some $[\tau(\tilde{a}^i)]$ with \tilde{a}^i still in $\tilde{\Sigma}^x$). \square

Definition 19. Let $D = \langle \tau, T(\mathbf{f}_n) \rangle$ be an \mathcal{R} -EDTD-design where the type $\tau = \langle \Sigma, \tilde{\Sigma}, \pi, s, \mu \rangle$ is normalized. We denote by

- κ any function associating to each node x of T either a set $\tilde{\Sigma}^x \subseteq \tilde{\Sigma}(a)$ if a is the label of x , or the set $\{\mathbf{f}\}$ if \mathbf{f} is the label of x .
- $D_\kappa^x = \langle \pi(\kappa(x)), B^x \rangle$, for each node x in T with $\text{lab}(x) \in \Sigma$, the box-design induced by D and κ where either $B^x = \{\varepsilon\}$ if x is a leaf node, or $B^x = \kappa(y_1) \dots \kappa(y_k)$ if $\text{children}(x) = y_1 \dots y_k$.

Given a sound typing (τ_n) for D , we say that

- κ is induced by the pair (τ_n) and T if, for each non-function node x of T , $\kappa(x)$ contains exactly all the specialized element names associated to x by validating each possible tree in $\text{ext}_T(\tau_n)$.
- $\kappa' \leq \kappa$ iff $\kappa'(x) \subseteq \kappa(x)$, for each x . \square

The intention is to relate locality properties about D with locality properties about each D_κ^x similarly as we made for \mathcal{R} -SDTDs, with the difference that here D_κ^x depends on the choice of κ . Unfortunately, although τ is normalized, if D admits local typings, then κ may not be unique. Consider the following example.

Example 8. Let $D = \langle \tau, T \rangle$ be a normalized dRE-EDTD-design where $T = s_0(\mathbf{f}_1 a(\mathbf{f}_2) \mathbf{f}_3)$, $\tau = \langle \Sigma, \tilde{\Sigma}, \pi, \tilde{s}_0, \mu \rangle$, $\pi(\tilde{s}_0) = (\tilde{a}^1 \tilde{a}^2)^+$, $\pi(\tilde{a}^1) = \tilde{b}^1$, and $\pi(\tilde{a}^2) = \tilde{c}^1$. We have two successful mappings κ^1, κ^2 such that

- $\kappa^1(x_1) = \tilde{s}_0$, $\kappa^1(x_3) = \tilde{a}^1$, $D_{\kappa^1}^{x_1} = \langle (\tilde{a}^1 \tilde{a}^2)^+, \mathbf{f}_1 \tilde{a}^1 \mathbf{f}_3 \rangle$, and $D_{\kappa^1}^{x_3} = \langle \tilde{b}^1, \mathbf{f}_2 \rangle$
- $\kappa^2(x_1) = \tilde{s}_0$, $\kappa^2(x_3) = \tilde{a}^2$, $D_{\kappa^2}^{x_1} = \langle (\tilde{a}^1 \tilde{a}^2)^+, \mathbf{f}_1 \tilde{a}^2 \mathbf{f}_3 \rangle$, and $D_{\kappa^2}^{x_3} = \langle \tilde{c}^1, \mathbf{f}_2 \rangle$

From them we have two different maximal local typings for D :

$$((\tilde{a}^1 \tilde{a}^2)^*, \tilde{b}^1, \tilde{a}^2(\tilde{a}^1 \tilde{a}^2)^*) \quad ((\tilde{a}^1 \tilde{a}^2)^* \tilde{a}^1, \tilde{c}^1, (\tilde{a}^1 \tilde{a}^2)^*)$$

Notice that they are substantially different and also that from the other possible mapping κ^3 , where $\kappa^3(x_3) = \{\tilde{a}^1, \tilde{a}^2\}$, we cannot derive any local typing because if \mathbf{f}_2 is replaced by b , then \mathbf{f}_3 must start with $a(c)$, and if \mathbf{f}_2 is replaced by c , then \mathbf{f}_1 must start with $a(b)$. But $((\tilde{a}^1 \tilde{a}^2)^* \tilde{a}^1, \tilde{b}^1 + \tilde{c}^1, \tilde{a}^2(\tilde{a}^1 \tilde{a}^2)^*)$ is neither local (even) nor sound. \square

Now we prove the main results of this section.

Theorem 4.13. *Let $D = \langle \tau, T(\mathbf{f}_n) \rangle$ be a distributed design where τ is a **normalized** \mathcal{R} -EDTD. The following are equivalent:*

- (1) D admits a local typing;
- (2) There is a function κ , as defined above, such that each box-design D_κ^x admits a local typing.

Proof. (1) \Rightarrow (2): Let (τ_n) be a local typing for D , then $T(\tau_n) \equiv \tau$ holds. Consider the function κ induced by (τ_n) and T (the choice is consistent with Lemma 4.12). As τ is normalized, there is only one possibility for validating (in a bottom-up way) each tree in $\text{ext}_T(\tau_n)$. If for some node x of T the box-design D_κ^x did not admit any local typing, then there would be no possibility of generating all the strings in $\pi(\kappa(x))$. Contradiction.

(2) \Rightarrow (1): If for some κ each box-design D_κ^x admits a local typing, then we can construct each type τ_i as made for \mathcal{R} -SDTDs, in such a way that $T(\tau_n)$ is structurally equivalent to τ . \square

Corollary 4.14. *Problem $\exists\text{-LOC}_{[\mathcal{R}\text{-EDTD}]}$ (or $\exists\text{-ML}_{[\mathcal{R}\text{-EDTD}]}$ but $\mathcal{R} \neq \text{dRE}$) for **normalized** $\mathcal{R}\text{-EDTDs}$ is decidable by an oracle machine in $\text{NP}^{\mathcal{C}}$ where \mathcal{C} is the complexity class of solving $\exists\text{-LOC}_{[\mathcal{R}]}^B$ (or $\exists\text{-ML}_{[\mathcal{R}]}^B$).*

Proof. Let $\tau = \langle \Sigma, \pi, s \rangle$ be a type and $T(!f_n)$ be a kernel. Consider the $\exists\text{-LOC}_{[\mathcal{R}\text{-EDTD}]}$ problem and the following algorithm:

1. **Guess:** the function κ ;
2. **Check:** call $\exists\text{-LOC}_{[\mathcal{R}]}^B$ over D_κ^x for each node x of T with $\text{lab}(x) \in \Sigma$.

For $\exists\text{-ML}_{[\mathcal{R}\text{-EDTD}]}$ we use the same algorithm since, in general ($\mathcal{R} \neq \text{dRE}$), a maximal local typing always exists if there is a local one. \square

Problem $\exists\text{-ML}_{[\text{dRE-EDTD}]}$ will be discussed in Section 7.

Theorem 4.15. *Let $D = \langle \tau, T(\mathbf{f}_n) \rangle$ be a distributed design where τ is a **normalized** $\mathcal{R}\text{-EDTD}$. The following are equivalent:*

- (1) D admits a perfect typing;
- (2) There is a function κ such that each D_κ^x admits a perfect typing, and for each sound typing (τ'_n) for D , $\kappa' \leq \kappa$ where κ' is induced by (τ'_n) .

Proof. (1) \Rightarrow (2): Let (τ_n) be the perfect typing of D and κ be the function induced by (τ_n) . By Theorem 4.14, each box-design D_κ^x admits a local typing, and clearly it is perfect as (τ_n) is. Finally, we observe that since $(\tau'_n) \leq (\tau_n)$, then (τ'_n) can not induce in κ' more elements than (τ_n) .

(2) \Rightarrow (1): As we made for $\mathcal{R}\text{-SDTDs}$, the typing (τ_n) that we can construct by the local typings of the various D_κ^x (without renaming the specialized element names) together with the needful content models already in τ produces a type $T(\tau_n)$ structurally equivalent to τ . \square

Corollary 4.16. *Problem $\exists\text{-PERF}_{[\mathcal{R}\text{-EDTD}]}$ for **normalized** $\mathcal{R}\text{-EDTDs}$ is polynomial time reducible to $\exists\text{-PERF}_{[\mathcal{R}]}^B$.*

Proof. Let $\tau = \langle \Sigma, \pi, s \rangle$ be a type and $T(!f_n)$ be a kernel.

PROOF IDEA: Build κ in polynomial time and in a top-down style (this is the technical core of the proof) by assuming that a perfect typing exists. Thus, call $\exists\text{-PERF}_{[\mathcal{R}]}^B$ over D_κ^x for each node x of T with $\text{lab}(x) \in \Sigma$.

More formally, let x be the root of T and m be the number of its children. Consider the following steps:

1. Build from $\text{child-str}(x)$ a **dRE**, that we call $r(x)$, as follows. For each j in $[1..m]$,
 - if $\text{child-str}(x)[j]$ is an element name, say a , then replace it with the set $\tilde{\Sigma}_j(a)$ (where the subscript means that all the specialized element names are renamed with j as subscript);
 - else, if $\text{child-str}(x)[j]$ is a function, then replace it with $\tilde{\Sigma}_j^*$ (j has the same meaning as above);

2. Build from $\pi(\tilde{s}_0)$ an \mathcal{R} -type $\tau(x)$ by replacing, in the alphabet of $\pi(\tilde{s}_0)$, each symbol of the form \tilde{a}^ℓ with $\tilde{a}_1^\ell, \dots, \tilde{a}_m^\ell$.
3. Perform the intersection $L = [r(x)] \cap [\tau(x)]$.
4. For each child y of x having a as label and position j , then $\kappa(y)$ contains all the elements of the form \tilde{a} such that \tilde{a}_j is in the alphabet of L .

As we know $\kappa(y)$ for each child y of x , then we repeat the previous steps for the children of y by considering $\pi(\kappa(y))$ instead of $\pi(\tilde{s}_0)$. We will stop when we reach the leaves of T .

The algorithm is correct because if we have a look at the alphabet of L , we see that it contains, for each j in $[1..m]$, exactly the specialized element names that we need to associate to the j^{th} child of x because are induced by all possible local typings. Intuitively, if the alphabet of L contains, for instance, \tilde{a}_1^ℓ this means that there is a sound typing for D that induces \tilde{a}^ℓ for the first child of x , and if the alphabet of L does not contain, for instance, \tilde{b}_3^ℓ there is no sound typing for D inducing \tilde{b}^ℓ for the third child of x . \square

Now we consider the remaining complexity result that does not require any reduction to strings.

Theorem 4.17. *Problems $\text{LOC}_{[\mathcal{R}\text{-EDTD}]}$, $\text{ML}_{[\mathcal{R}\text{-EDTD}]}$, and $\text{PERF}_{[\mathcal{R}\text{-EDTD}]}$ are at least as hard as $\text{EQUIV}_{[\mathcal{R}\text{-EDTD}]}$.*

Proof. We define a logspace transformation φ from $\text{EQUIV}_{[\mathcal{R}\text{-EDTD}]}$ to $\text{LOC}_{[\mathcal{R}\text{-EDTD}]}$. Afterwards, we show that the statement also holds for the other two problems. Let τ', τ'' be two arbitrary \mathcal{R} -EDTDs. The application of φ to the this pair produces the design $D = \langle \tau, T \rangle$ and the typing τ_1 , where $[\tau] = s_0([\tau'])$, $T = s_0(\mathbf{f}_1)$, and $[\tau_1] = s_1([\tau''])$. Since $T(\tau_1)$ is exactly $s_0([\tau''])$, it is clear that $\tau \equiv T(\tau_1)$ if and only if $\tau' \equiv \tau''$. Finally, we just notice that $\tau' \equiv \tau''$ iff τ_1 is both perfect and maximal local as T consists of just a function node other than the root. \square

Corollary 4.18. *Problems $\text{LOC}_{[\text{nFA-EDTD}]}$, $\text{ML}_{[\text{nFA-EDTD}]}$, and $\text{PERF}_{[\text{nFA-EDTD}]}$ are **EXPTIME-hard**.*

Theorem 4.19. *Problem $\text{LOC}_{[\text{nFA-EDTD}]}$ is **EXPTIME-complete**.*

Proof. (Membership) Let $D = \langle \tau, T \rangle$ be an **nFA-EDTD**-design and (τ_n) be a D -consistent typing. Build $T(\tau_n)$ in polynomial time (by Proposition 3.1) and check in exponential time if $T(\tau_n) \equiv \tau$ (by Theorem 4.7).

(Hardness) By Corollary 4.18. \square

5. The typing problems for words

We study in this section the typing problems for words. (Recall that most of our problems for trees has been reduced to problems for words.) We present a number of complexity results. We leave for the next section, two issues, namely $\text{PERF}_{[\text{nFA}]}$ and $\exists\text{-PERF}_{[\text{nFA}]}$, for which we will need a rather complicated automata construction. We start by recalling a definition and a result that we will use further.

Theorem 5.1 ([32]). $\text{EQUIV}_{[\text{nFA}]}$ is *PSPACE-complete*.

The hardness of the $\text{EQUIV}_{[\text{nFA}]}$ problem is used to show some hardness results of our problems.

Theorem 5.2. Problems $\text{LOC}_{[\text{nFA}]}$, $\text{ML}_{[\text{nFA}]}$, $\text{PERF}_{[\text{nFA}]}$ are *PSPACE-hard*.

Proof. We define a logspace transformation φ , in such a way that

$$\text{EQUIV}_{[\text{nFA}]} \leq_m^{\text{L}} \text{LOC}_{[\text{nFA}]}$$

Afterwards, we show that the statement also holds for the other two problems. Let $\mathcal{A}, \mathcal{A}_1$ be two arbitrary **nFAs**. The application of φ to the pair $\mathcal{A}, \mathcal{A}_1$ produces the design $\langle \tau, w \rangle$ and the typing τ_1 , where $\tau = \mathcal{A}$, $w = \mathbf{f}_1$ and $\tau_1 = \mathcal{A}_1$. Since $w(\tau_1) = \mathcal{A}_1$, it is clear that $\tau \equiv w(\tau_1)$ if and only if $\mathcal{A} \equiv \mathcal{A}_1$. Finally, we just notice that $\mathcal{A} \equiv \mathcal{A}_1$ if and only if τ_1 is both perfect and maximal local as w consists of just a function. \square

We now consider upper bounds. Section 6 will show that $\text{PERF}_{[\text{nFA}]}$ is in **PSPACE**. We next show that $\text{LOC}_{[\text{nFA}]}$ is.

Theorem 5.3. $\text{LOC}_{[\text{nFA}]}$ is in *PSPACE* (so it is *PSPACE-complete*).

Proof. Let $w(\mathbf{f}_n)$ be a kernel string, τ be an **nFA**, and (τ_n) be a typing. Since the new automaton $w(\tau_n)$ has size $\mathcal{O}(\|w\| + |(\tau_n)|)$, we can check in polynomial space if $w(\tau_n) \equiv \tau$. \square

The proof that also $\text{ML}_{[\text{nFA}]}$ is in **PSPACE** requires more technical insights and it is deferred to Section 7.

Let us turn to the hardness of the \exists -versions of the problems.

Theorem 5.4. $\exists\text{-LOC}_{[\text{nFA}]}$, $\exists\text{-ML}_{[\text{nFA}]}$, and $\exists\text{-PERF}_{[\text{nFA}]}$ are *PSPACE-hard*.

Proof. We define a logspace transformation φ , in such a way that the following relations hold:

- (1) $\text{EQUIV}_{[\text{nFA}]} \leq_m^{\text{L}} \exists\text{-LOC}_{[\text{nFA}]}$;
- (2) $\text{EQUIV}_{[\text{nFA}]} \leq_m^{\text{L}} \exists\text{-ML}_{[\text{nFA}]}$;
- (3) $\text{EQUIV}_{[\text{nFA}]} \leq_m^{\text{L}} \exists\text{-PERF}_{[\text{nFA}]}$.

Let $\mathcal{A}_1 = \langle K_1, \Sigma_1, \Delta_1, s_1, F_1 \rangle$, $\mathcal{A}_2 = \langle K_2, \Sigma_2, \Delta_2, s_2, F_2 \rangle$ be two **nFAs**. The application of φ to the pair $(\mathcal{A}_1, \mathcal{A}_2)$ produces the design $D = \langle \mathcal{A}, w \rangle$ where

- $w = \mathbf{f}_1 \ c \ \mathbf{f}_2$, with c being a fresh terminal symbol which does not belong to $(\Sigma_1 \cup \Sigma_2)$;
- while automaton $\mathcal{A} = \langle K, \Sigma, \Delta, s, F \rangle$ is defined as follows: (i) $K = K_1 \cup K_2 \cup \{s, p_c, q_c\}$; (ii) $\Sigma = \Sigma_1 \cup \Sigma_2 \cup \{a, b, c\}$; (iii) $\Delta = \Delta_1 \cup \Delta_2 \cup \{(s, a, p_c), (s, b, p_c), (p_c, c, q_c), (q_c, \varepsilon, s_1), (q_c, \varepsilon, s_2)\}$; (iv) $F = F_1 \cup F_2$.

Intuitively, if we consider \mathcal{A}_1 and \mathcal{A}_2 as **nREs**, then \mathcal{A} is $(ac\mathcal{A}_1 + bc\mathcal{A}_2)$.

We claim that there is a local typing (similarly, maximal local, or perfect) for D if and only if $\mathcal{A}_1 \equiv \mathcal{A}_2$. First of all, we observe that transformation φ is extremely simple and it is clearly in logspace. In fact, string w is a constant, while the choice of a terminal symbol which does not appear in \mathcal{A}_1 nor in \mathcal{A}_2 can be done in logspace, and also \mathcal{A} can be obtained by merging \mathcal{A}_1 and \mathcal{A}_2 with a constant number of transitions. We prove the statement for (1) and we just notice that whenever there is a local typing for D , then the typing $((a+b), \mathcal{A}_1)$ is perfect (thus, also maximal).

(\Rightarrow) *If there is a local typing for D then $\mathcal{A}_1 \equiv \mathcal{A}_2$.* Since $\mathcal{A} = (ac\mathcal{A}_1 + bc\mathcal{A}_2)$, then $[ac\mathcal{A}_1]$ and $[bc\mathcal{A}_2]$ form a partition of $[\mathcal{A}]$. In this case, any local typing must have the following form $((aX_1 + bX_2), Y)$ where X_1, X_2, Y are **nFAs**. Clearly, all the strings accepted by w are obtained by aX_1cY and bX_2cY . Then $c\mathcal{A}_1 \equiv X_1cY$ and $c\mathcal{A}_2 \equiv X_2cY$ must hold. But since any string in $[\mathcal{A}_1]$ or $[\mathcal{A}_2]$ does not start with c , then necessarily $[X_1] = [X_2] = \varepsilon$. This way, $\mathcal{A}_1 \equiv Y$ and $\mathcal{A}_2 \equiv Y$ and then $\mathcal{A}_1 \equiv \mathcal{A}_2$.

(\Leftarrow) *If $\mathcal{A}_1 \equiv \mathcal{A}_2$, there is a local typing for D .* This part of the proof is trivial because $((a+b), \mathcal{A}_1)$ always represents a local typing for D . \square

We now have lower bounds for all these problems and some upper bounds. We will derive missing upper bounds using the construction of automata that we call “perfect” for given design problems.

6. Perfect automaton for words

We next present the construction of the *perfect automaton* for a design word problem. The perfect automaton has the property that if a perfect typing exists for this problem, it is “highlighted” by the automaton. This will provide a **PSPACE** procedure for finding this perfect typing if it exists.

Let $\mathcal{A} = \langle K, \Sigma, \Delta, s, F \rangle$ be an **nFA**. We can assume w.l.o.g. that it has no ε -transition. Given two states q_i, q_f in K , a string w in Σ^* is said to be *delimited* in \mathcal{A} by q_i and q_f if $(q_i, w, q_f) \in \Delta^*$. By exploiting this notion, the sets of all the states delimiting w in \mathcal{A} are defined as follows:

$$Ini(\mathcal{A}, w) = \{q_i \in K : \exists q_f \in K \text{ s.t. } (q_i, w, q_f) \in \Delta^*\}$$

$$Fin(\mathcal{A}, w) = \{q_f \in K : \exists q_i \in K \text{ s.t. } (q_i, w, q_f) \in \Delta^*\}$$

In particular, if $w = \varepsilon$, these two sets are $Ini(\mathcal{A}, \varepsilon) = Fin(\mathcal{A}, \varepsilon) = K$. $Ini(\mathcal{A}, w)$ is called the set of *initial states* while $Fin(\mathcal{A}, w)$ is the set of *final states* for the word w . Given two states q_i, q_f in K , the *local automaton* $\mathcal{A}(q_i, q_f) = \langle K' \subseteq K, \Sigma, \Delta', q_i, \{q_f\} \rangle$ induced from \mathcal{A} by q_i, q_f is a portion of \mathcal{A} containing all those transitions of \mathcal{A} leading from q_i to q_f . More precisely, for each pair of states q, q' in K and for each symbol a in Σ , $(q, a, q') \in \Delta'$ if and only if there are two strings u, v in Σ^* such that: $(q_i, u, q) \in \Delta^*$, $(q, a, q') \in \Delta$, and $(q', v, q_f) \in \Delta^*$. Finally, given two strings w_1, w_2 in Σ^+ , then $\mathcal{A}(w_1, w_2)$ is

the set of all local automata induced by w_1 and w_2 . It is formally defined as $\mathcal{A}(w_1, w_2) = \{\mathcal{A}(q_i, q_f) : q_i \in \text{Fin}(\mathcal{A}, w_1), q_f \in \text{Ini}(\mathcal{A}, w_2)\}$. In particular, if $w_i = \varepsilon$ for some i in $[1..n]$, the kernel string contains consecutive functions. In particular for the previous definitions we have:

$$\mathcal{A}(w_1, \varepsilon) = \{\mathcal{A}(q_i, q_f) : q_i \in \text{Fin}(\mathcal{A}, w_1) \text{ and } q_f \in K\}$$

$$\mathcal{A}(\varepsilon, w_2) = \{\mathcal{A}(q_i, q_f) : q_i \in K \text{ and } q_f \in \text{Ini}(\mathcal{A}, w_2)\}$$

$$\mathcal{A}(\varepsilon, \varepsilon) = \{\mathcal{A}(q_i, q_f) : q_i, q_f \in K\}.$$

Similarly, given a string w in Σ^* , $\mathcal{A}(w)$ is the set of all local automata induced by w . It is defined as $\mathcal{A}(w) = \{\mathcal{A}(q_i, q_f) : (q_i, w, q_f) \in \Delta^*\}$ and in particular $\mathcal{A}(\varepsilon) = \{\mathcal{A}(q, q) : q \in K\}$ is a set of $|K|$ automata, one for each state in K .

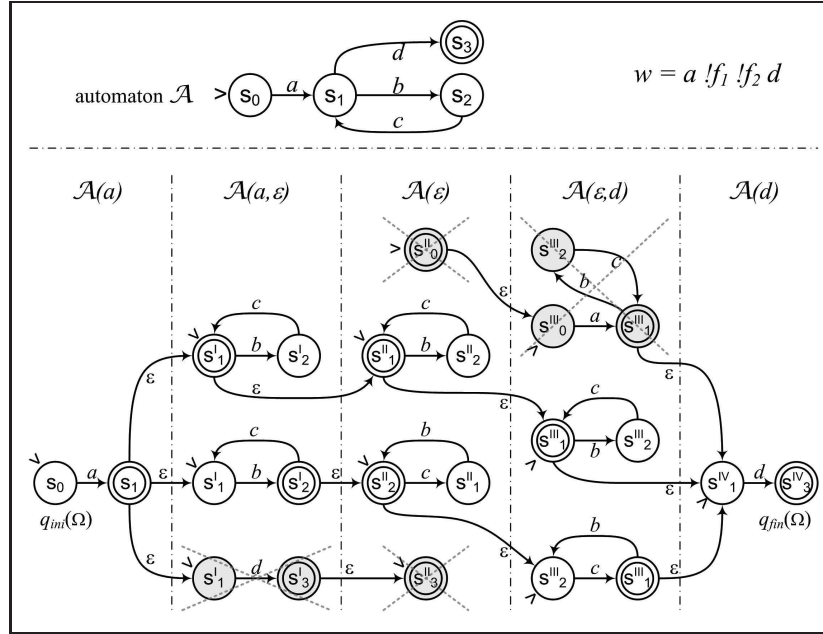


Figure 7: A perfect automaton (construction)

Let $w(\mathbf{f}_n)$ be a kernel string and \mathcal{A} be an nFA. The *perfect automaton* w.r.t. \mathcal{A} and w consists of several local automata suitably joined together by ε -transitions. It is denoted by $\Omega(\mathcal{A}, w)$ (or Ω when it is clear from the context who are \mathcal{A} and w). Algorithm 1 describes how to build the perfect automaton (assume that any pair of local automata have disjoint sets of states labeled as in \mathcal{A}), while Figure 7 shows the perfect automaton obtained by a given finite state machine and a kernel string. We say that \mathcal{A} is *compatible* with w if the set of all (legal) local automata in Ω is not empty after correction steps, or equivalently, if there exists at least a sound typing. Moreover,

Algorithm 1 PERFECTAUTOMATON(w, \mathcal{A})

```

1. Input:  $w(\mathbf{f}_n) = w_0 \mathbf{f}_1 w_1 \dots \mathbf{f}_n w_n$ ,  $\mathcal{A} = \langle K, \Sigma, \Delta, s, F \rangle$ 
2. Output:  $\Omega(\mathcal{A}, w) := \emptyset$ 
3. for each automaton  $W \in \mathcal{A}(w_0)$  do
   $\triangleright$  add  $W$  to  $\Omega$ 
4. for each  $i$  in  $[1..n]$  do
   $\triangleright$  for each automaton  $X \in \mathcal{A}(w_{i-1}, w_i)$  do
    a. add  $X$  to  $\Omega$ 
    b. for each automaton  $W \in \mathcal{A}(w_{i-1})$  do
      - if  $\text{label}(q_{fin}(W)) = \text{label}(q_{ini}(X))$ 
         $\cdot$  add the transition  $(q_{fin}(W), \varepsilon, q_{ini}(X))$  to  $\Omega$ 
    c. for each automaton  $W \in \mathcal{A}(w_i)$  do
      - add  $W$  to  $\Omega$ 
      - if  $\text{label}(q_{fin}(X)) = \text{label}(q_{ini}(W))$ 
         $\cdot$  add the transition  $(q_{fin}(X), \varepsilon, q_{ini}(W))$  to  $\Omega$ 
  //Correction steps:
5. for each automaton  $W \in \mathcal{A}(w_0)$  do
  - if  $\text{label}(q_{ini}(W)) \neq s$  //if  $w_0 = \varepsilon$ 
     $\cdot$  remove  $W$  from  $\Omega$  //it is illegal
6. merge all automata in  $\Omega$  being in  $\mathcal{A}(w_0)$  according to their
   labels and use the (unique) initial state as initial state for  $\Omega$ 
7. for each automaton  $W \in \mathcal{A}(w_n)$  do
  - if  $\text{label}(q_{fin}(W)) \in F$ 
     $\cdot F(\Omega) = F(\Omega) \cup \{q_{fin}(W)\}$ 
  else //if  $w_n = \varepsilon$ 
     $\cdot$  remove  $W$  from  $\Omega$  //it is illegal
8. for each automaton  $A \in \Omega$  do
  - if (there is no path from  $q_{ini}(\Omega)$  to  $A$  or
    there is no path from  $A$  to any final state of  $\Omega$ )
     $\cdot$  remove  $A$  from  $\Omega$  //it is illegal

```

- $Seq(\Omega)$ denotes the set of all the sequences $W_0, X_1, W_1, \dots, X_n, W_n$ of connected automata in Ω such that: W_0 is an automaton in $\mathcal{A}(w_0)$, while W_i and X_i are, respectively, in $\mathcal{A}(w_i)$ and $\mathcal{A}(w_{i-1}, w_i)$ for any i in $[1..n]$;
- $Typ(\Omega) = \{(X_i) : W_0, X_1, W_1, \dots, X_n, W_n \in Seq(\Omega)\}$ is the set containing all different typings (X_1, \dots, X_n) from any sequence in $Seq(\Omega)$;
- $Aut(\Omega_i) = \{X_i : (X_1, \dots, X_n) \in Typ(\Omega)\}$ is the set of all legal automata in $\mathcal{A}(w_{i-1}, w_i)$;
- $\Omega_i = \cup Aut(\Omega_i)$ is the type obtained by the union of all automata $Aut(\Omega_i)$;
- (Ω_n) is the typing for w and \mathcal{A} obtained from Ω .

Let (\mathcal{A}_n) be a sequence of automata. We define the *direct extension* of (\mathcal{A}_n) as the set of string defined as $[(\mathcal{A}_n)] = \{u_1 \dots u_n \mid \text{for each } i \ u_i \in [\mathcal{A}_i]\}$.

Lemma 6.1. *For any nFA \mathcal{A} , then $\Omega \leq \mathcal{A}$ holds. On the other hand, $\mathcal{A} \leq \Omega$ does not hold in general.*

Proof. Given a string u in $[\Omega]$, then there exists a sequence (τ_{2n+1}) of automata in $Seq(\Omega)$ accepting u and expressible as $\mathcal{A}(s, q_0)$, $\mathcal{A}(q_0, s_1)$, $\mathcal{A}(s_1, q_1)$, \dots , $\mathcal{A}(q_{n-1}, s_n)$, $\mathcal{A}(s_n, q_n)$ for some states $q_0, s_1, q_1 \dots, s_n, q_n$. Moreover, by definition of direct extension, for each string $u_0 \sigma_1 u_1 \dots \sigma_n u_n$ in $[(\tau_{2n+1})]$ we have that $u_0 \in [\mathcal{A}(s, q_0)]$, $\sigma_i \in [\mathcal{A}(q_{i-1}, s_i)]$ and $u_i \in [\mathcal{A}(s_i, q_i)]$, for each i in $[1..n]$. But, by definition of local automata, the following sequence of transitions (each of which belongs to Δ^*): (s, w_0, q_0) , (q_0, σ_1, s_1) , (s_1, w_1, q_1) , \dots , (q_{n-1}, σ_n, s_n) , (s_n, w_n, q_n) , where $q_n \in F$, is also derivable by \mathcal{A} .

For the second part of the proof consider the string $w = a f c$ and the dRE $abc + d$. \square

Lemma 6.2. *Let $w(\mathbf{f}_n)$ be a string compatible with an nFA \mathcal{A} . Any typing in $Typ(\Omega)$ is sound for w and \mathcal{A} .*

Proof. Given any typing (X_n) in $Typ(\Omega)$, by definition, there is a sequence (τ_{2n+1}) of automata such that $X_i = \tau_{2i}$ for each i in $[1..n]$. By Lemma 6.1 $(\tau_{2n+1}) \leq \mathcal{A}$ holds. Moreover as, by definition, the extension of $w(X_n)$ is $[w(X_n)] = \{w_0 \sigma_1 w_1 \dots \sigma_n w_n : \sigma_i \in [X_i], 1 \leq i \leq n\}$. Then $w(X_n) \leq (\tau_{2n+1})$ as well since all strings w_0, \dots, w_n are accepted by $\tau_1, \tau_3 \dots, \tau_{2n+1}$, respectively, by definition of local automata induced by a single string. Therefore, $w(X_n) \leq \mathcal{A}$. \square

Theorem 6.3. *Let $w(\mathbf{f}_n)$ be a kernel string compatible with a given nFA \mathcal{A} , and (τ_n) be a sound typing for them. Then, both $w(\tau_n) \leq \Omega$ and $(\tau_n) \leq (\Omega_n)$ hold.*

Proof. Since (τ_n) is sound for w and \mathcal{A} , then $w(\tau_n) \leq \mathcal{A}$ holds. In particular, for each string $\chi = w_0 \sigma_1 w_1 \dots \sigma_n w_n$ in $[w(\tau_n)]$, where each $\sigma_i \in [\tau_i]$, there is a sequence of states $q_0, s_1, q_1 \dots, s_n, q_n$ proving the membership of χ in $[\mathcal{A}]$ by the following sequence of transitions $(s, w_0, q_0) \in \Delta^*$, $(q_0, \sigma_1, s_1) \in \Delta^*$, $(s_1, w_1, q_1) \in \Delta^*$, \dots , $(q_{n-1}, \sigma_n, s_n) \in \Delta^*$, $(s_n, w_n, q_n) \in \Delta^*$ where $q_n \in F$ holds as well. But, this means that the sequence $\mathcal{A}(s, q_0)$, $\mathcal{A}(q_0, s_1)$, $\mathcal{A}(s_1, q_1)$, \dots , $\mathcal{A}(q_{n-1}, s_n)$, $\mathcal{A}(s_n, q_n)$ of automata belongs to $Seq(\Omega)$, so $w(\tau_n) \leq \Omega$ holds. Moreover, since each $\mathcal{A}(q_{i-1}, s_i) \in Aut(\Omega_i)$, it follows that $\tau_i \leq \Omega_i$ for each i , that is $(\tau_n) \leq (\Omega_n)$. \square

Corollary 6.4. *Let $w(\mathbf{f}_n)$ be a kernel string compatible with a given nFA \mathcal{A} , and (τ_n) be a local typing for them. Then, $w(\tau_n) \equiv \Omega \equiv \mathcal{A}$ holds.*

Proof. By Lemma 6.1 and Theorem 6.3. \square

Theorem 6.5. *Let $w(\mathbf{f}_n)$ be a kernel string and \mathcal{A} be an nFA compatible with w . There is a perfect typing for w and \mathcal{A} if and only if $w(\Omega_n) \equiv \mathcal{A}$. If so, the perfect typing is exactly (Ω_n) .*

Proof. (\Rightarrow) if there is a perfect typing for w and \mathcal{A} then $w(\Omega_n) \equiv \mathcal{A}$. If w and \mathcal{A} admit a perfect typing, say (τ_n) , then (as it is also sound), by Theorem 6.3, $(\tau_n) \leq (\Omega_n)$. Suppose that $(\tau_n) < (\Omega_n)$ held. There would be (at least) an i in $[1..n]$ such that $\tau_i < \Omega_i$. In other words, there would be an automaton $\tau'_i \in \text{Aut}(\Omega_i)$ accepting some strings rejected by τ_i . Consider the typing $(\tau'_n) \in \text{Typ}(\Omega)$ containing τ'_i in position i . By Lemma 6.2, (τ'_n) is sound and then $\tau'_i \leq \tau_i$, by definition. But this is a contradiction. Therefore $(\tau_n) \equiv (\Omega_n)$ and then $w(\Omega_n) \equiv \mathcal{A}$, as (τ_n) is also local.

(\Leftarrow) if $w(\Omega_n) \equiv \mathcal{A}$ then there is a perfect typing for w and \mathcal{A} . This is true since (Ω_n) is local and because, by Theorem 6.3, $(\tau_n) \leq (\Omega_n)$ for any sound typing (τ_n) . \square

The following two examples show that if there exists a local typing (τ_n) for w and \mathcal{A} , then $(\tau_n) < (\Omega_n)$ might hold. This can happen even if (τ_n) is a unique maximal local.

Example 9. Consider the string $w = a \mathbf{f}_1 c \mathbf{f}_2 e$, and the regular expression $\tau = abccde$ compatible with w . Clearly, the typing (b, cd) is local (sound and complete) for w and τ because $w(b, cd) \equiv \tau$. Nevertheless, $(\Omega_2) = (bc?, c?d)$ is (strictly) greater than (b, cd) since $[bc?] = \{b, bc\} \supset \{b\}$ and $[c?d] = \{d, cd\} \supset \{cd\}$. \square

Example 10. Let $w = a \mathbf{f}_1 \mathbf{f}_2 d$ be a kernel string and τ be the regular expression $a(bc)^*d$. Clearly, the typing $((bc)^*, (bc)^*)$ is local (also unique maximal local but not perfect). But, as consequence of the construction of perfect automaton, we have: $\text{Aut}(\Omega_1) = \{(bc)^*, (bc)^*b\}$ and $\text{Aut}(\Omega_2) = \{(bc)^*, c(bc)^*\}$. Consequently, $\Omega_1 \equiv ((bc)^*b?)$ and $\Omega_2 \equiv (c?(bc)^*)$ do not represent a sound (and hence local) typing since they allow strings such as $abccbcd$ or $abcbcd$ that are not accepted by τ . \square

The following example shows that even if there is no local typing for w and τ , then $\Omega \equiv \tau$ may hold.

Example 11. Let a τ be the regular expression $ab + ba$ and $w = \mathbf{f}_1 \mathbf{f}_2$. There are two sound typings: (a, b) and (b, a) , but there is no local typing. However, $\Omega \equiv \tau$. \square

We can now use the perfect automata construction to characterize the complexity of PERF_{nFA} . We use the next lemma:

Lemma 6.6. *Let $w(\mathbf{f}_n)$ be a kernel string and \mathcal{A} be a k -state nFA. The algorithm for building the perfect automaton $\Omega(\mathcal{A}, w)$ works in polynomial time.*

Proof. Any set $\mathcal{A}(w_i)$ or $\mathcal{A}(w_{i-1}, w_i)$ contains at most k^2 automata each of which having size $\mathcal{O}(k)$. Therefore, the number of macro-iterations of the algorithm are $\mathcal{O}(nk^2)$, while the size of Ω is $\mathcal{O}(nk^3)$. For each w_i , the sets $\text{Ini}(\mathcal{A}, w_i)$ and $\text{Fin}(\mathcal{A}, w_i)$ can be obtained in nondeterministic logarithmic space (thus in polynomial time) because for any pair of states q_1, q_2 in \mathcal{A} , we check if the

string w_i is in the language $[\mathcal{A}(q_1, q_2)]$. Finally, all the automata in $\mathcal{A}(w_i)$ and $\mathcal{A}(w_{i-1}, w_i)$ are nothing else but different copies of \mathcal{A} having different initial and final states. \square

Now, we have:

Theorem 6.7. $\text{PERF}_{[\text{nFA}]}$ is in PSPACE . So it is also PSPACE -complete by Theorem 5.2.

Proof. Let $w(\mathbf{f}_n)$ be a kernel string, τ be an nFA , and (τ_n) be a typing. Construct the perfect automaton $\Omega(\tau, w)$. By Lemma 6.6, Ω can be built in polynomial time w.r.t. $|\tau| + \|w\|$. Then, check in polynomial space if $w(\Omega_n) \equiv \tau \equiv w(\tau_n)$. \square

And w.r.t. finding a perfect typing (if it exists), we have:

Theorem 6.8. $\exists\text{-PERF}_{[\text{nFA}]}$ is in PSPACE . So it is also PSPACE -complete by Theorem 5.4.

Proof. Let $\langle \tau, w(\mathbf{f}_n) \rangle$ be a (string) design. Construct the perfect automaton $\Omega(\tau, w)$. By Lemma 6.6, Ω can be built in polynomial time w.r.t. $|\tau| + \|w\|$. Then, check if $w(\Omega_n) \equiv \tau$, which is feasible in polynomial space. \square

6.1. Additional properties

We now show how to exploit perfect automaton properties to find (maximal) sound typings when a design does not allow any perfect. Clearly, this technique can be used for seeking (maximal) local typings as well. Let $w(\mathbf{f}_n)$ be a kernel string and \mathcal{A} be an nFA -type compatible with w . All the automata belonging to $\text{Aut}(\Omega_i)$ can be decomposed in at most $2^{|\text{Aut}(\Omega_i)|} - 1$ different automata such that there are no two of them accepting the same string. In particular, this new set is denoted by $\text{Dec}(\Omega_i)$ and defined as follows:

$$\text{Dec}(\Omega_i) = \{\cap \mathbb{A}_1 - \cup \mathbb{A}_2 : \emptyset \neq \mathbb{A}_1 \subseteq \text{Aut}(\Omega_i), \mathbb{A}_2 = \text{Aut}(\Omega_i) - \mathbb{A}_1\}$$

An example for three automata is given in Figure 8. Finally, $\text{Dec}(\Omega) = \{(D_1, \dots, D_n) : D_i \in \text{Dec}(\Omega_i)\}$ is the set of all different typings from $\text{Dec}(\Omega_1) \times \dots \times \text{Dec}(\Omega_n)$. Given a typing (τ_n) , we say that $(\tau_n) \in \text{Dec}(\Omega)$ if there exists a sequence $(D_n) \in \text{Dec}(\Omega)$ such that $\tau_i \equiv D_i$, for each i .

Given a type $\tau \leq \Omega_i$ for some i in $[1..n]$, $\text{Dec}(\tau, i) = \{\tau \cap \tau' : \tau' \in \text{Dec}(\Omega_i)\}$ denotes the partition of τ , namely $\cup \text{Dec}(\tau, i) \equiv \tau$, obtained by its projection on $\text{Dec}(\Omega_i)$. Let (τ_n) be any typing for a kernel string $w(\mathbf{f}_n)$. Given a string $u \in \Sigma^*$ and an i in $[1..n]$, then $(\tau_n)_{[\tau_i|u]}$ denotes the new typing obtained from (τ_n) by replacing τ_i with the minimal dFA accepting only the string u . In particular $[w(\tau_n)_{[\tau_i|u]}]$ is defined as $\{w_0\sigma_1w_1\dots\sigma_nw_n : \sigma_i = u, \sigma_j \in [\tau_j] \forall j \neq i\}$ and clearly,

$$w(\tau_n) \equiv \bigcup_{u \in [\tau_i]} w(\tau_n)_{[\tau_i|u]}$$




Figure 8: Partitioning of (three) sets and enumeration of the parts

We now define an extension of (τ_n) as the new typing obtained from (τ_n) by replacing τ_i with the new type $(\tau_i \cup \tau)$, and denoted by $(\tau_n)_{[\tau_i \cup \tau]}$. In particular,

$$w(\tau_n)_{[\tau_i \cup \tau]} \equiv \bigcup_{u \in [\tau_i \cup \tau]} w(\tau_n)_{[\tau_i | u]}$$

Clearly, if $\tau \leq \tau_i$, then $(\tau_n) \equiv (\tau_n)_{[\tau_i \cup \tau]}$. Otherwise $(\tau_n) < (\tau_n)_{[\tau_i \cup \tau]}$.

Definition 20. A type τ **extends** another type τ' if $[\tau] - [\tau'] \neq \emptyset$ holds. Moreover, the extension is called **partial** or **total** depending on whether $[\tau] \cap [\tau'] \neq \emptyset$ or not, respectively. \square

Lemma 6.9. Let $D = \langle \mathcal{A}, w \rangle$ be an **nFA**-design, (τ_n) be a consistent sound typing for D , and $\tau \in \text{Dec}(\Omega_i)$ be an **nFA** belonging to the decomposition of Ω_i , for some i in $[1..n]$. If τ partially extends τ_i , then the extension $(\tau_n)_{[\tau_i \cup \tau]}$ of (τ_n) is still sound.

Proof. By definition 20, $[\tau]$ contains at least a string that does not belong to $[\tau_i]$ but also a string, say u' , accepted by both τ_i and τ . In order to prove the statement, we show that $w(\tau_n)_{[\tau_i | u]} \leq \Omega$ holds for each $u \in [\tau] - [\tau_i]$ (recall that, by Lemma 6.1, $\Omega \leq \mathcal{A}$).

Since, by Theorem 6.3, $\tau_i \leq \Omega_i$, then there is a nonempty set $\mathbb{A} \subseteq \text{Aut}(\Omega_i)$ containing all-and-only the automata accepting u' . Clearly, since $\tau \in \text{Dec}(\Omega_i)$ and $u' \in [\tau]$, then τ is also in $\text{Dec}(\tau', i)$ for each $\tau' \in \mathbb{A}$. This means that each string $u \in [\tau] - [\tau_i]$ is accepted by all-and-only the automata in \mathbb{A} as well. By Theorem 6.3, $w(\tau_n) \leq \Omega$, and in particular $w(\tau_n)_{[\tau_i | u']} \leq \Omega$, as $u' \in [\tau_i \cap \tau]$. In other words, any string in $[w(\tau_n)_{[\tau_i | u]}]$ is accepted by (at least) a sequence of automata in $\text{Seq}(\Omega)$. Finally, as both u and u' are recognized by all-and-only the automata in \mathbb{A} , then each string $w_0 \sigma_1 w_1 \dots \sigma_n w_n$ in $[w(\tau_n)_{[\tau_i | u]}]$ (with $\sigma_i = u$) has a twin in $[w(\tau_n)_{[\tau_i | u']}]$ (with $\sigma_i = u'$) and both of them are accepted by exactly the same sequences in $\text{Seq}(\Omega)$. \square

Theorem 6.10. Let (τ_n) be a maximal typing for a kernel string $w(\mathbf{f}_n)$ and an **nFA** \mathcal{A} compatible with w . Then for each i , $\text{Dec}(\tau_i, i) \subseteq \text{Dec}(\Omega_i)$.

Proof. Let i be an index arbitrarily fixed in $[1..n]$. As (τ_n) is maximal then, by definition, it is sound and, by Theorem 6.3, $\tau_i \leq \Omega_i$. Let D_i be a copy of $\text{Dec}(\Omega_i)$. Then $\tau_i \leq \cup D_i$. Remove now, from D_i , each automata τ_{D_i} (if any) such that $[\tau_{D_i}] \cap [\tau_i] = \emptyset$. Still, $\tau_i \leq \cup D_i$ holds. Hence, consider the two possible (and alternative) cases: (1) $\tau_i \equiv \cup D_i$, or (2) $\tau_i < \cup D_i$. In the first case the theorem is already proved. While, in the latter case, there is (at least) an automaton $\tau \in D_i$ that partially extends τ_i entailing relation $(\tau_n) < (\tau_n)_{[\tau_i \cup \tau]}$. But since $(\tau_n)_{[\tau_i \cup \tau]}$ is still sound (see Lemma 6.9), then there is a contradiction because (τ_n) is assumed to be maximal. \square

We are now ready to prove a main results of the section. In our original paper, we showed a **2-EXPSpace** upper bound for $\exists\text{-LOC}_{[\text{nFA}]}$ and $\exists\text{-ML}_{[\text{nFA}]}$. This was improved to **EXPSpace** in [31]. We present here an alternative proof of that results using the previous decomposition.

Theorem 6.11. *Problems $\exists\text{-LOC}_{[\text{nFA}]}$ and $\exists\text{-ML}_{[\text{nFA}]}$ are in **EXPSpace**.*

Proof. By Lemma 6.9 and Theorem 6.10, if an **nFA**-design $D = \langle \tau, w \rangle$ admits a (maximal) local typing, say (τ_n) , then for each τ_i there exists a subset of $\text{Dec}(\Omega_i)$, say D_i , such that $\cup D_i \equiv \tau_i$.

Let m be the number of states of τ , and $\nu + n$ be the length of w where n is clearly the number of functions and ν is the length of the non-function symbols in w . By definition, for each i in $[1..n]$, each automaton in $\text{Aut}(\Omega_i)$ has size at most m and the cardinality of $\text{Aut}(\Omega_i)$ is at most m^2 . Thus, the cardinality of $\text{Dec}(\Omega_i)$ is no more than 2^{m^2} , as well as the cardinality of D_i . In the worst case, an automaton in $\text{Dec}(\Omega_i)$ is obtained as $\cap \mathbb{A}_1 - \cup \mathbb{A}_2$ where both $|\mathbb{A}_1| = |\mathbb{A}_2| = \mathcal{O}(m^2)$. So, the size of $\cap \mathbb{A}_1$ is no more than $(m^2)^{m^2}$ [22], that is clearly lower than 2^{m^3} . The size of $\cup \mathbb{A}_2$ is at most m^3 [22]. Now, for computing $\cap \mathbb{A}_1 - \cup \mathbb{A}_2$ we perform the following intersection $(\cap \mathbb{A}_1) \cap (\overline{\cup \mathbb{A}_2})$. The complement of $\cup \mathbb{A}_2$ may have 2^{m^3} states [22]. Finally $\cap \mathbb{A}_1 - \cup \mathbb{A}_2$ require no more than 2^{2m^3} states, and the size of $\cup D_i \equiv \tau_i$ is at most $2^{2m^3} * 2^{m^2}$ being clearly $2^{\mathcal{O}(m^3)}$.

Now, we are ready for computing the size of the **nFA** $w(\tau_n)$. It is exactly $\nu + n * 2^{\mathcal{O}(m^3)}$. So, for deciding whether $w(\tau_n) \equiv \tau$ we need no more than exponential space w.r.t. the input size $\nu + n + m$. The only problem we still have is that we do not know a priori how to choose D_i . There are $2^{2^{m^2}}$ possible subsets. But as **NEXPSpace** = **EXPSpace** (by Savitch's theorem), then we can simply guess each D_i .

About $\exists\text{-ML}_{[\text{nFA}]}$, we must find a maximal D_i . But in **EXPSpace** we can still guess the sequence D_1, \dots, D_n and prove (by Theorem 6.10), for each D_i , that none of the automata in $\text{Dec}(\Omega_i) - D_i$ can be added to D_i because the resulting typing would loose its soundness. After the guess, the number of checks (each of which may require exponential space) is at most $n * 2^{m^2}$. \square

7. Complexity for trees

Based on Theorem 6.11, we now obtain complexity bounds for the tree problems. This completes results obtained in [3, 31] on this topic. The next result first appeared in [3]. However, the sketch of proof given there was not correct. A proof was then presented in [31]. We next present a new proof based on perfect automata.

Theorem 7.1. $ML_{[nFA]}$ is in **PSPACE** (so the problem is **PSPACE**-complete).

Proof. Let $D = \langle \tau, w \rangle$ be an **nFA**-design, and (τ_n) be a D -consistent typing. First of all, we check if (τ_n) is local (and we have already proved that $LOC_{[nFA]}$ is doable in **PSPACE**). If so, then $\bar{\tau} \cap w(\tau_n) \equiv \emptyset$ (where $\bar{\tau}$ is the **nFA** of possibly exponential size accepting the complement of language $[\tau]$). Subsequently, we check if (τ_n) is *not* maximal. In particular, by Lemma 6.9 and Theorem 6.10, (τ_n) is not maximal if there is an **nFA** $\mathcal{A} \in Dec(\Omega_i)$ for some i in $[1..n]$ such that at least one of the following is true:

- \mathcal{A} totally extends τ_i and $w(\tau_n)_{[\tau_i \cup \mathcal{A}]}$ is still sound, namely $[\mathcal{A}] \cap [\tau_i] = \emptyset$ and $\bar{\tau} \cap w(\tau_n)_{[\tau_i \cup \mathcal{A}]} \equiv \emptyset$.
- \mathcal{A} partially extends τ_i , namely $[\mathcal{A}] - [\tau_i] \neq \emptyset$ and $[\mathcal{A}] \cap [\tau_i] \neq \emptyset$;

So we proceed as follows:

1. Guess an index i , and a nonempty set of automata $\mathbb{A}_1 \subseteq Aut(\Omega_i)$
2. Compute $\mathbb{A}_2 = Aut(\Omega_i) - \mathbb{A}_1$
3. Let \mathcal{A} denote the automaton $\cap \mathbb{A}_1 - \cup \mathbb{A}_2$ (we do not really build it);
4. If $[\mathcal{A}] \cap [\tau_i] = \emptyset$ then,
 - if $\bar{\tau} \cap w(\tau_n)_{[\tau_i \cup \mathcal{A}]} \equiv \emptyset$, then (τ_n) is not maximal
 - else if $[\mathcal{A}] - [\tau_i] = [\mathcal{A}] \cap [\bar{\tau}_i] \neq \emptyset$, then (τ_n) is not maximal

Observe that even if \mathcal{A} , $\bar{\tau}$, or $\bar{\tau}_i$ may be exponential in size, we only use them for *intersection nonemptiness* or *intersection emptiness* problems that are both **NL**-complete problems [26]. Intuitively, we could avoid the materialization of such automata with “on-the-fly” constructions. Hence, an **NL** algorithm on a non-materialized (single) exponential automaton leads to **PSPACE**. More formally, we consider alternating finite state machines **aFAs** (for more details see [17, 42]). We do not completely define them but we just recall what we need:

- given an **aFA** \mathcal{A} , deciding whether $[\mathcal{A}] = \emptyset$ is **PSPACE**-complete;
- Any **nFA** is trivially a special kind of **aFA**;
- Given two **aFAs** \mathcal{A} and \mathcal{A}' , a new **aFA** for $\bar{\mathcal{A}}$, $\mathcal{A} \cup \mathcal{A}'$, and $\mathcal{A} \cap \mathcal{A}'$, can be constructed in polynomial time and its size is linear.

Finally, we observe that all the above emptiness decisions deal with **nFAs** of polynomial size and can be checked in **PSPACE** as well as the nonemptiness decisions as **PSPACE** is closed under complement. \square

Now, we show how to reduce locality problems on boxes to locality problems on strings.

Definition 21. Let $D = \langle \tau, B \rangle$ be an \mathcal{R} -design where $B = B_0 \mathbf{f}_1 B_1 \dots \mathbf{f}_n B_n$ is a kernel box. Consider the k^{th} sequence of strings w_0, \dots, w_n (with $1 \leq k \leq |B_0| * \dots * |B_n|$) built from B_0, \dots, B_n by varying w_i among the strings in $[B_i]$ (in some fixed order) for each $i \in [0..n]$. We denote by $D^k = \langle \tau, w^k \rangle$ the k^{th} \mathcal{R} -design built from D where $w^k(\mathbf{f}_n)$ is the kernel string $w_0 \mathbf{f}_1 w_1 \dots \mathbf{f}_n w_n$. \square

Lemma 7.2. Let $D = \langle \tau, B \rangle$ be an \mathcal{R} -design and (τ_n) be a D -consistent typing. We have that:

- (1) If (τ_n) is local for D , then it is sound for each D^k ;
- (2) If (τ_n) is sound for each D^k , then it is sound for D as well.

Proof. (1): If (τ_n) is local for D , then $B_0 \tau_1 B_1 \dots \tau_n B_n \equiv \tau$. This means that $w_0 \tau_1 w_1 \dots \tau_n w_n \leq \tau$ for each $w_i \in [B_i]$. Thus, (τ_n) is sound for each D^k .

(2): If for each design D^k we have that $w_0 \tau_1 w_1 \dots \tau_n w_n \leq \tau$ holds, then $B_0 \tau_1 B_1 \dots \tau_n B_n \leq \tau$ as well. \square

A direct consequence of the above theorem is that if a typing is not sound for some D^k , then it can not be local for D . So a local typing candidate for D is a typing being sound for each D^k . Now suppose that (τ_n) is a maximal sound typing for D^{k_1} but it is not sound for D^{k_2} . This means that at least one $[\tau_i]$ contains some extra string such that $[w^{k_2}(\tau_n)]$ is not fully contained in $[\tau]$. So we could remove such strings to obtain a typing sound for both D^{k_1} and D^{k_2} but not maximal for D^{k_1} any more. So we can guess a maximal sound typing for each D^k and then, remove the exceeding strings. This is equivalent to keeping the componentwise intersection of these maximal typings. Let $\beta = |B_0| * \dots * |B_n|$, we should build β (it is an exponential number) perfect automata. For each i in $[1..n]$ we should consider the sets of automata $Aut(\Omega_i^1), \dots, Aut(\Omega_i^\beta)$ and from these the respective decompositions $Dec(\Omega_i^1), \dots, Dec(\Omega_i^\beta)$. Now we can guess β subsets D_i^1, \dots, D_i^β and finally compute τ_i as $(\cup D_i^1) \cap \dots \cap (\cup D_i^\beta)$. But this is equivalent to consider directly $Aut(\Omega_i) = Aut(\Omega_i^1) \cup \dots \cup Aut(\Omega_i^\beta)$, compute the decomposition $Dec(\Omega_i)$ and guess a subset D_i from $Dec(\Omega_i)$. This is much more convenient because $Aut(\Omega_i)$ contains at most a quadratic number of automata w.r.t. the states of τ . Now, we show how to extend the construction of Ω to a box-design for obtaining this new $Aut(\Omega_i)$. Let \mathcal{A} be an nFA and $B(\mathbf{f}_n)$ a kernel box, we have that:

- $Ini(\mathcal{A}, B_i) = \{q_i \in K : \exists q_f \in K \text{ s.t. } (q_i, w, q_f) \in \Delta^*, w \in [B_i]\}$
- $Fin(\mathcal{A}, B_i) = \{q_f \in K : \exists q_i \in K \text{ s.t. } (q_i, w, q_f) \in \Delta^*, w \in [B_i]\}$
- $\mathcal{A}(B_{i-1}, B_i) = \{\mathcal{A}(q_i, q_f) : q_i \in Fin(\mathcal{A}, B_{i-1}), q_f \in Ini(\mathcal{A}, B_i)\}$

$Aut(\Omega_i)$ is the set of all legal automata in $\mathcal{A}(B_{i-1}, B_i)$ as for string. Note that, due to the structure of each B_i , it is very easy to build $Ini(\mathcal{A}, B_i)$ and $Fin(\mathcal{A}, B_i)$ without enumerating all the strings in $[B_i]$.

Theorem 7.3. Problems $\exists\text{-LOC}_{[\text{nFA}]}^B$ and $\exists\text{-ML}_{[\text{nFA}]}^B$ are in EXSPACE .

Proof. Let $D = \langle \tau, B \rangle$ be an **nFA**-design where B is a kernel box. We guess, for each i , a subset of automata in $Dec(\Omega_i)$, the decomposition of the new set $Aut(\Omega_i)$ built as above. Thus, we check if it is a (maximal) local typing for D as made in the proof of Theorem 6.11. \square

Corollary 7.4. $\exists\text{-LOC}_{[\text{nFA-EDTD}]}$ and $\exists\text{-ML}_{[\text{nFA-EDTD}]}$ are in 2-EXPSPACE .

Proof. Let $D = \langle \tau, T \rangle$ be an **nFA**-EDTD-design. We build from τ its equivalent normalized version τ^d that, after all, is a **dFA**-EDTD of exponential size. So the oracle machine discussed in Corollary 4.14 actually works in NEXPTIME^C where C is the complexity class of solving $\exists\text{-LOC}_{[\text{dFA}]}^B$ (or $\exists\text{-ML}_{[\text{dFA}]}^B$). By Theorem 7.3, both of these problems are in **EXPSPACE**. Thus, the whole algorithm works in 2-EXPSPACE . (Note that, **EXPSPACE** is the best known upper bound even for $\exists\text{-LOC}_{[\text{dFA}]}$ [31].) \square

The following analysis makes use of a technique introduced in [31] for building the perfect automaton for **dFA**-designs.

Definition 22. Let $D = \langle \tau, B \rangle$ be an \mathcal{R} -design where $B = B_0 \mathbf{f}_1 B_1 \dots \mathbf{f}_n B_n$ is a kernel box. Together with D^k we consider the string design \hat{D}^k defined as follows. Let $\hat{\Sigma} = \Sigma \uplus \{\sigma_0, \dots, \sigma_n\}$ be an extension of Σ and $\sigma(\mathbf{f}_n) = \sigma_0 \mathbf{f}_1 \sigma_1 \dots \mathbf{f}_n \sigma_n$ be the kernel string built by combining the new symbols with the functions of B . We denote by $\hat{D}^k = \langle \hat{\Omega}^k, \sigma \rangle$ the k^{th} **dFA**-design built from D where $\hat{\Omega}^k = \hat{\Omega}^k(\tau, w^k)$ is the perfect automaton built as described in [31]. \square

The following lemma is a direct consequence of the definition of $\hat{\Omega}$ in [31].

Lemma 7.5. A typing (τ_n) is sound for D^k iff it is sound for \hat{D}^k .

Theorem 7.6. Let $D = \langle \tau, B \rangle$ be a **dFA**-design and (τ_n) be a D -consistent typing. The following are equivalent:

- (1) (τ_n) is perfect for D ;
- (2) (τ_n) is both local for D and perfect for each \hat{D}^k .

Proof. (1) \Rightarrow (2): If (τ_n) is perfect for D , then it is sound for each D^k , and by Lemma 7.5, sound for \hat{D}^k as well. Suppose that (τ_n) is not local for some \hat{D}^k , there is a string $\sigma_0 u_1 \sigma_1 \dots u_n \sigma_n \in [\hat{\Omega}^k]$ (all the strings have this form by definition) not captured by $\sigma(\tau_n)$. By Lemma 7.5, the string $w_0 u_1 w_1 \dots u_n w_n$ belongs to $[\tau]$ and as (τ_n) is perfect, then each $u_i \in [\tau_i]$: contradiction. Suppose that (τ_n) is not perfect for some \hat{D}^k . There is a sound typing (τ'_n) for \hat{D}^k not contained in (τ_n) , but by Lemma 7.5, (τ'_n) is also sound for D^k , so $w(\tau'_n) \leq \tau$: again a contradiction because (τ_n) is perfect.

(2) \Rightarrow (1): If (τ_n) is perfect for each \hat{D}^k then, by Lemma 7.5, it is sound for each D^k . Suppose that it is not perfect for D , then there is a sound typing (τ'_n) not contained in (τ_n) such that, for some k , $w(\tau'_n) \leq [\tau]$ for the k^{th} string w_0, \dots, w_n . So (τ'_n) is sound for D^k and also for \hat{D}^k . Contradiction. \square

Lemma 7.7. $\exists\text{-PERF}_{[\text{dFA}]}^B$ is in *coNP*.

Proof. Let $D = \langle \tau, B \rangle$ be a **dFA**-design where $B = B_0 \mathbf{f}_1 B_1 \dots \mathbf{f}_n B_n$ is a kernel box. We can decide in **NP** whether D does not admit any perfect typing by performing the following steps:

1. **Guess:** four string-designs D^{k_1} , D^{k_2} , D^{k_3} , and D^{k_4} .
2. **Check:** answer “yes” (D does not admit any perfect typing) if at least one of the following holds
 - (a) D^{k_1} does not admit any perfect typing;
 - (b) D^{k_2}, D^{k_3} have different perfect typings;
 - (c) D^{k_4} admits a perfect typing, say (τ_n) , but it is not local for D .

Each of check (a), (b), and the first part of (c) require polynomial time [31]. For the second part of check (c) we build $B(\tau_n)$ and prove that $B(\tau_n) < \tau$, namely $B(\tau_n) \cap \bar{\tau} = \emptyset$. Notice that if the yes answer only depends on step (c) this means that (τ_n) is sound for each D^k , and so it is not possible that $B(\tau_n) > \tau$. Thus, as τ is a **dFA**, its complement has the same size and the intersection emptiness can be done in polynomial time as well. \square

Corollary 7.8. $\exists\text{-PERF}_{[\text{nFA-EDTD}]}$ is in **coNEXPTIME**.

Proof. Let $D = \langle \tau, T \rangle$ be an **nFA-EDTD**-design. We build from τ its equivalent normalized version τ^d that, after all, is a **dFA-EDTD** of exponential size. By Corollary 4.16 we polynomial-time reduce $\exists\text{-PERF}_{[\mathcal{R}\text{-EDTD}]}$ (for **normalized** \mathcal{R} -EDTDs) to $\exists\text{-PERF}_{[\mathcal{R}]}^B$. So in our case we call $\exists\text{-PERF}_{[\text{dFA}]}^B$. But, as τ^d may be exponentially larger, then, by adapting the upper bound of Lemma 7.7, the whole algorithm works in **coNEXPTIME**. \square

Theorem 7.9. $\text{PERF}_{[\text{nFA-EDTD}]}$ is in **coNEXPTIME**.

Proof. Let $D = \langle \tau, T \rangle$ be an **nFA-EDTD**-design, and (τ_n) be a D -consistent typing. Compute in **coNEXPTIME** a perfect typing (τ'_n) if there is one. Transform (τ_n) into a **dFA-EDTD**-typing of exponential size. As $\text{EQUIV}_{[\text{AUTAs}]}$ is in **P**TIME then we can decide in **EXPTIME** whether (τ_n) and (τ'_n) are equivalent. \square

Unfortunately, for $\text{ML}_{[\text{nFA-EDTD}]}$ we do not have any good algorithm. Let $D = \langle \tau, T \rangle$ be an **nFA-EDTD**-design, (τ_n) be a maximal local typing for D , and κ be the function induced by (τ_n) and T . At the moment, we do not even know whether there could be a (non-maximal) local typing $(\tau'_n) < (\tau_n)$ such that $\kappa' < \kappa$. If there is none, given a local typing (τ_n) and its induced function κ , then each maximal local typing that extends (τ_n) has to induce the same κ as well. So we could compare the various D_κ^x with (τ_n) . But, the only known upper bound is given by the following theorem.

Theorem 7.10. $\text{ML}_{[\text{nFA-EDTD}]}$ is in **2-EXPSpace**.

Proof. Let $D = \langle \tau, T \rangle$ be an **nFA-EDTD**-design and (τ_n) be a D -consistent typing. We can check whether it is not maximal. Check in **EXPTIME** whether it is local or not. So, build the normalized type τ^d from τ . Guess a function κ and check whether each D_κ^x admits a local typing. This is in **2-EXPSpace** by Corollary 7.4.

So, build the typing (τ'_n) induced by the box-designs. It may be an **nFA-EDTD** typing exponentially larger. Check whether $(\tau_n) < (\tau'_n)$. This can be done in **2-EXPTIME**. So the algorithm works in **2-EXPSpace** and as this class is closed under complementation we also can decide $\text{ML}_{[\text{nFA-EDTD}]}$ in it. \square

Finally, we consider the reduction from trees to boxes for $\exists\text{-ML}_{[\text{dRE-EDTD}]}$. The difficulties affecting $\text{ML}_{[\text{nFA-EDTD}]}$ (as we do not know whether there could be a local typing $(\tau'_n) < (\tau_n)$ such that $\kappa' < \kappa$) concern also the existential problem in case of **dREs**.

Theorem 7.11. $\exists\text{-ML}_{[\text{dRE-EDTD}]}$ for *normalized \mathcal{R} -EDTDs* is decidable by an oracle machine in $\text{PSPACE}^{\mathcal{C}}$ where \mathcal{C} is the complexity class of solving the most difficult problem between $\exists\text{-ML}_{[\text{dRE}]}^B$ $\exists\text{-LOC}_{[\text{dFA}]}^B$.

Proof. In this case we have to check two sources of maximality depending on the choice of κ and on the related box-designs. To do that, we guess a function κ (the candidate for a maximal local typing) and we check whether each induced box-design (i) admits a local typing, (ii) is maximal and (iii) is **dRE**-definable. So, we have to prove that each $\kappa' > \kappa$ does not lead to any local typing. In particular:

1. Guess a functions κ ;
2. Prove that, for each node x of T with $\text{lab}(x) \in \Sigma$, the answer of $\exists\text{-ML}_{[\text{dRE}]}^B$ over D_κ^x is “yes”;
3. Prove that, for each $\kappa' > \kappa$, there is at least a node x of T with $\text{lab}(x) \in \Sigma$ such that the answer of $\exists\text{-LOC}_{[\text{dFA}]}^B$ over $D_{\kappa'}^x$ is “no”.

We just notice that there could be an exponential number of κ' to be enumerated and checked, as well as the number of calls to $\exists\text{-LOC}_{[\text{dFA}]}^B$. \square

8. Conclusion

As explained in the introduction, this work can serve as a basis for designing the distribution of a document. It would be interesting to extend our definitions and methods to richer types of web data. First, this would involve graph data and not just tree data. Then one should consider unordered collections and functional dependencies as in the relational model [5, 39]. Other dependencies and in particular inclusion dependencies would also clearly make sense in this setting [40]. More specific design methodology would also extend the techniques presented in this paper by considering concrete network configurations; this is left for future research.

Database design has a long history, see most database text book. Distributed database design has also been studied since the early days of databases, but much less, because distributed data management was limited by the difficulty to deploy distributed databases. The techniques that were developed, e.g., vertical and horizontal partitioning, are very different from the ones presented here because we focus on ordered trees and collections are not ordered in relational

databases. We believe that traditional database studies even on mainly theoretical topics such as normal forms are also relevant in a Web setting. An interesting direction of research is to introduce some of these techniques in our setting.

In the paper, the focus was on local typing that forces verification to be purely local. More generally, it would be interesting to consider typings of the resources that would minimize the communications needed for type checking (and not completely avoid them). Moreover, it would be interesting to analyze cases where a kernel document may change from time to time by adhering to some global type which uses function symbols in the specification itself. We are investigating this direction. Let us give a short example exhibiting some of the new difficulties that would arise in case kernel document changes were taken into account. Consider the kernel string $w = af$ and the type $\tau = af?ba^+$. By directly applying the techniques proposed in this paper, it seems clear that $f?ba^+$ would be the perfect typing for this design. So, one extension of w may be $afba$ (by attaching the tree fba complying with the perfect typing) which, in turn, represents a new kernel. But, this extension might still be extended, by attaching again tree fba , to form $afbaba$, since the first extension still contains a function call and the perfect typing defined for the remote resource should not vary. This last step could be performed several time. The language obtained by all possible extensions is defined by the type $af?(ba^+)^+$, being clearly different from τ . The problem here is that τ does not express directly a set of trees without taking into account a specific typing. New interesting questions might be: *How to look for typings that are, in a sense, fixpoints w.r.t. the original type with functions?* or *How to avoid irregularities?* or even *Is the perfect typing still unique?* Finally, interesting issues may also come from studying the impact of distributed typing (as studied here) on query optimization.

9. Acknowledgments

This work is a co-operation in the context of the ICT-2007.8.0 FET Open project Nr. 233599: Foundations of XML - safe processing of dynamic data over the Internet (FOX). Serge Abiteboul's work was supported by the European Research Council Advanced Grant Webdam and by the French ANR Grant Docflow. Georg Gottlob's work was supported by EPSRC grant EP/E010865/1 "Schema Mappings and Automated Services for Data Integration and Exchange". Gottlob also gratefully acknowledges a Royal Society Wolfson Research Merit Award. Marco Manna acknowledges the support and hospitality of the Oxford-Man Institute of Quantitative Finance, where he worked as an Academic Visitor. Finally, the authors also want to thank Thomas Schwentick, Wim Martens, and Matthias Niewerth for useful discussions on the problem.

References

- [1] Abiteboul, S., Benjelloun, O., Milo, T., 2008. The Active XML project: an overview. The VLDB Journal 17 (5), 1019–1040.

- [2] Abiteboul, S., Bonifati, A., Cobéna, G., Manolescu, I., Milo, T., 2003. Dynamic XML documents with distribution and replication. In: Proceedings of SIGMOD '03, San Diego, California. ACM, New York, NY, USA, pp. 527–538.
- [3] Abiteboul, S., Gottlob, G., Manna, M., 2009. Distributed XML Design. In: Proceedings of PODS '09, Providence, Rhode Island, USA. ACM, New York, NY, USA, pp. 247–258.
- [4] Abiteboul, S., Manolescu, I., Taropa, E., 2006. A Framework for Distributed XML Data Management. In: Advances in Database Technology - EDBT 2006. Vol. 3896 of LNCS. Springer Berlin / Heidelberg, pp. 1049–1058.
- [5] Arenas, M., Libkin, L., 2004. A normal form for XML documents. ACM Trans. Database Syst. 29 (1), 195–232.
- [6] Balmin, A., Papakonstantinou, Y., Vianu, V., 2004. Incremental validation of XML documents. ACM Trans. Database Syst. 29 (4), 710–751.
- [7] Bex, G. J., Gelade, W., Martens, W., Neven, F., 2009. Simplifying XML schema: effortless handling of nondeterministic regular expressions. In: Proceedings of SIGMOD '09, Providence, Rhode Island, USA. ACM, New York, NY, USA, pp. 731–744.
- [8] Bray, T., Paoli, J., Sperberg-McQueen, C. M., Maler, E., Yergeau, F., Cowan, J., August 2006. Extensible Markup Language (XML) 1.1 (Second Edition). W3C Recommendation 16 August 2006, edited in place 29 September 2006.
URL <http://www.w3.org/TR/2006/REC-xml11-20060816/>
- [9] Bremer, J.-M., Gertz, M., 2003. On Distributing XML Repositories. In: Proceedings of WebDB '03, San Diego, California. pp. 73–78.
- [10] Brüggemann-Klein, A., Murata, M., Wood, D., 2001. Regular Tree and Regular Hedge Languages over Unranked Alphabets: Version 1. Tech. rep., HKUST Institutional Repository (Hong Kong).
URL <http://repository.ust.hk/dspace/handle/1783.1/738>
- [11] Brüggemann-Klein, A., Wood, D., 1998. One-Unambiguous Regular Languages. Inf. Comput. 140 (2), 229–253.
- [12] Ceri, S., Fraternali, P., Bongio, A., 2000. Web Modeling Language (WebML): a modeling language for designing Web sites. Computer Networks 33 (1-6), 137–157.
- [13] Ceri, S., Pernici, B., Wiederhold, G., 1984. An Overview of Research in the Design of Distributed Databases. IEEE Database Eng. Bull. 7 (4), 46–51.

- [14] Clark, J., Murata, M., December 2001. RELAX NG Specification. Committee Specification.
URL <http://relaxng.org/spec-20011203.html>
- [15] Comon, H., Dauchet, M., Gilleron, R., Löding, C., Jacquemard, F., Lugiez, D., Tison, S., M., T., October 2007. Tree Automata Techniques and Applications. Release October, 12th 2007.
URL <http://www.grappa.univ-lille3.fr/tata>
- [16] Cook, S. A., McKenzie, P., 1987. Problems complete for deterministic logarithmic space. *Journal of Algorithms* 8 (3), 385–394.
- [17] Fellah, A., Jürgensen, H., Yu, S., 1990. Constructions for alternating finite automata. *Int. J. Comp. Math* 35 (1), 117–132.
- [18] Ghelli, G., Colazzo, D., Sartiani, C., 2007. Efficient Inclusion for a Class of XML Types with Interleaving and Counting. In: *Database Programming Languages*. Vol. 4797 of LNCS. Springer Berlin / Heidelberg, pp. 231–245.
- [19] Grosso, P., Veillard, D., February 2001. XML Fragment Interchange. W3C Candidate Recommendation.
URL <http://www.w3.org/TR/xml-fragment>
- [20] Hagenah, C., Muscholl, A., 1998. Computing ε -Free NFA from Regular Expressions in $O(n \log^2(n))$ Time. In: *Mathematical Foundations of Computer Science 1998*. Vol. 1450 of LNCS. Springer Berlin / Heidelberg, pp. 277–285.
- [21] Hernández, L. O., Pegah, M., 2003. WebDAV: what it is, what it does, why you need it. In: *Proceedings of SIGUCCS '03*, San Antonio, TX, USA. ACM, New York, NY, USA, pp. 249–254.
- [22] Holzer, M., Kutrib, M., 2003. State Complexity of Basic Operations on Nondeterministic Finite Automata. In: *Implementation and Application of Automata*. Vol. 2608 of LNCS. Springer Berlin / Heidelberg, pp. 61–79.
- [23] Hromkovič, J., Seibert, S., Wilke, T., 1997. Translating regular expressions into small ε -free nondeterministic finite automata. In: *Proceedings of STACS '97*, Lbeck, Germany. Vol. 1200 of LNCS. Springer Berlin / Heidelberg, pp. 55–66.
- [24] Jagadish, H. V., Lakshmanan, L. V. S., Milo, T., Srivastava, D., Vista, D., 1999. Querying network directories. *SIGMOD Record* 28 (2), 133–144.
- [25] Jiang, T., Ravikumar, B., 1993. Minimal NFA problems are hard. *SIAM J. Comput.* 22 (6), 1117–1141.
- [26] Jones, N. D., 1975. Space-bounded reducibility among combinatorial problems. *J. Comput. Syst. Sci.* 11 (1), 68–85.

- [27] Martens, W., Neven, F., Schwentick, T., 2004. Complexity of Decision Problems for Simple Regular Expressions. In: Proceedings of MFCS '04, Prague, Czech Republic. Vol. 3153 of LNCS. Springer Berlin / Heidelberg, pp. 889–900.
- [28] Martens, W., Neven, F., Schwentick, T., 2009. Complexity of Decision Problems for XML Schemas and Chain Regular Expressions. *SIAM J. Comput.* 39 (4), 1486–1530.
- [29] Martens, W., Neven, F., Schwentick, T., Bex, G. J., 2006. Expressiveness and complexity of XML Schema. *ACM Trans. Database Syst.* 31 (3), 770–813.
- [30] Martens, W., Niehren, J., 2007. On the minimization of XML Schemas and tree automata for unranked trees. *J. Comput. Syst. Sci.* 73 (4), 550–583.
- [31] Martens, W., Niewerth, M., Schwentick, T., 2010. Schema design for XML repositories: complexity and tractability. In: Proceedings of PODS '10, Indianapolis, Indiana, USA. ACM, New York, NY, USA, pp. 239–250.
- [32] Meyer, A. R., Stockmeyer, L. J., 1973. Word problems requiring exponential time(Preliminary Report). In: Proceedings of STOC '73, Austin, Texas, USA. ACM, New York, NY, USA, pp. 1–9.
- [33] Murata, M., Lee, D., Mani, M., Kawaguchi, K., 2005. Taxonomy of XML schema languages using formal language theory. *ACM Trans. Internet Technol.* 5 (4), 2005.
- [34] Özsu, M. T., Valduriez, P., 1991. Distributed Database Systems: Where Are We Now? *Computer* 24 (8), 68–78.
- [35] Papakonstantinou, Y., Vianu, V., 2000. DTD inference for views of XML data. In: Proceedings of PODS '00, Dallas, Texas, US. ACM, New York, NY, USA, pp. 35–46.
- [36] Seidl, H., 1990. Deciding Equivalence of Finite Tree Automata. *SIAM J. Comput.* 19 (3), 424–437.
- [37] Suciu, D., 2002. Typechecking for Semistructured Data. In: Database Programming Languages. LNCS. Springer Berlin / Heidelberg, pp. 1–20.
- [38] Thompson, H. S., Beech, D., Maloney, M., Mendelsohn, N., October 2004. XML Schema Part 1: Structures Second Edition. W3C Recommendation. URL <http://www.w3.org/TR/xmlschema-1/>
- [39] Vincent, M., Liu, J., 2003. Functional Dependencies for XML. In: Web Technologies and Applications. Vol. 2642 of LNCS. Springer Berlin / Heidelberg, pp. 600–600.

- [40] Vincent, M. W., Schrefl, M., Liu, J., Liu, C., Dogen, S., 2004. Generalized Inclusion Dependencies in XML. In: Advanced Web Technologies and Applications. Vol. 3007 of LNCS. Springer Berlin / Heidelberg, pp. 224–233.
- [41] Winkler, P., 2004. Mathematical Puzzles: A Connoisseur’s Collection. AK Peters.
- [42] Yu, S., 1997. Handbook of Formal Languages: word, language, grammar. Vol. 1. Springer-Verlag, New York, NY, USA, Ch. Regular Languages, pp. 41–110.
- [43] Yu, S., 2001. State complexity of regular languages. J. Autom. Lang. Comb. 6 (2), 221–234.