

Attribute Grammars for Scalable Query Processing on XML Streams

Christoph Koch*

Database Group
LFCS, University of Edinburgh
Edinburgh EH9 3JZ, UK
koch@dbai.tuwien.ac.at

Stefanie Scherzinger

Lehrstuhl für Dialogorientierte Systeme
FMI, Universität Passau
D-94030 Passau, Germany
scherzin@fmi.uni-passau.de

Abstract

We introduce the new notion of XML Stream Attribute Grammars (XSAGs). XSAGs are the first scalable query language for XML streams (running strictly in linear time with bounded memory consumption independent of the size of the stream) that allows for actual data transformations rather than just document filtering. XSAGs are also relatively easy to use for humans. Moreover, the XSAG formalism provides a strong intuition for which queries can or cannot be processed scalably on streams. We introduce XSAGs together with the necessary language-theoretic machinery, study their theoretical properties such as their expressiveness and complexity, and discuss their implementation.

1 Introduction

In recent years, XML has become a standard format for document exchange and now seems to develop into a preeminent representation language for streaming data as well. This development calls for flexible query languages for processing streams which support data transformations.

In [11, 14, 6], fragments of the standard XML Query language [15] are evaluated on XML streams. These fragments tend to support powerful data transformations, with the consequence that query processing neither scales in terms of runtime nor memory consumption. Indeed, in these works, memory buffers are required that can grow arbitrarily large, depending on the amount of data communicated via the stream.

This problem is due to the nature of XML Query, which renders it ill-suited for stream processing: Fea-

tures such as nested for-loops with transitive paths (e.g., using the descendant axis), which may lead to a nonlinearly-sized output, and nonlocal computations such as joins and the reordering and sorting of data cannot be handled scalably on streams. In addition, the syntax of XML Query makes it difficult to tell for a user whether a query can – at least in principle – be evaluated scalably, in linear time using little memory.

Query languages that require unbounded memory buffers constitute a scalability issue on streams and are not in the spirit of the database community’s quest for tailored formalisms that provide the appropriate tradeoffs between expressiveness and complexity for the data management challenge at hand.

XML streams by definition may be *very* long, or should even be assumed to be infinite. For query processing to be feasible on streams, there is a need for special-purpose query languages and evaluation algorithms which *scale to streams*, i.e.,

- (a) which can be evaluated strictly in linear time in the size of the input,
- (b) which work in the streaming fashion, by one linear forward scan of the data, and
- (c) for which, at any time during query evaluation, memory consumption is bounded¹.

Among the models of computation that allow for better control of complexity than languages such as XML Query, there are various forms of automata/transducers and certain attribute grammars. The former are, however, unsuitable as query languages used by humans because their specifications

¹Note that a stack of memory proportional to the maximum depth of the XML tree is necessary for even the most basic sequential navigation and parsing tasks (see e.g. [7, 9]). To be precise, in (c) we thus call for memory consumption that is bounded w.r.t. the length of the stream but not the depth of the XML tree (an indication of its structural complexity). XML trees tend to be very shallow but wide, therefore such a stack is not considered a bottleneck to scalability.

* Work sponsored by Erwin Schrödinger Grant J2169 of the Austrian Research Fund (FWF).

tend to be large, technical, and hard to read. The latter approach is developed in the present paper.

Attribute Grammars for Stream Processing

In this work, we develop and investigate a formalism for processing XML streams called *XML Stream Attribute Grammars* (XSAGs), a new class of attribute grammars specifically designed for scalable XML stream processing. XSAGs can be evaluated strictly in linear time in streaming fashion, consuming only a stack of memory bounded by the depth of the XML tree being streamed. Thus, XSAGs satisfy our desiderata (a) through (c).

XSAGs are based on *extended regular tree grammars*, i.e. regular tree grammars in which the right-hand sides of productions may contain regular expressions, allowing to specify nodes in the parse tree that have an unbounded number of children. Extended regular tree grammars are thus well-suited for specifying classes of unranked trees denoting XML documents. We assume that extended regular tree grammars are often available for XML streams in the dialect of *Document Type Definitions* (DTDs). This adds to the relevance of the present formalism.

An XSAG is obtained by annotating a given extended regular tree grammar with attribution functions that describe the output to be produced from the input stream. In the tradition of L-attributed grammars [1], XSAGs are assumed to perform a single scan of the XML stream, which amounts to effecting a single depth-first left-to-right traversal of the document tree. Also in the tradition of L-attributed grammars, right-hand sides of productions can be annotated with *two* attribution functions, one – placed at the beginning of the right-hand side – that is executed when reaching the opening tag of a node in the XML document (or equivalently, when descending into a subtree), and the second – placed at the end of the right-hand side – which is executed when reaching the corresponding closing tag (or, equivalently, when returning from the depth-first-traversal of the subtree).

Example 1.1 Consider the extended regular tree grammar $G = (Nt, T, P, bib)$ with nonterminals

$$Nt = \{bib, book, article, title, author\},$$

start nonterminal bib , terminals

$$T = \{bib, book, article, title, author, PCDATA\},$$

and the productions P

$$\begin{aligned} bib & ::= bib((book \cup article)^*) \\ book & ::= book(title.author.author^*) \\ article & ::= article(title.author.author^*) \\ title & ::= title(PCDATA) \\ author & ::= author(PCDATA) \end{aligned}$$

which defines an XML bibliography database.

By changing the first production to

$$bib ::= \{ECHO\} bib((book \cup article)^*)$$

we obtain an XSAG that simply echoes the input stream. Indeed, the start production matches the root node of the document, and *ECHO* writes the entire subtree of the current node to the output as XML.

If we are instead only interested in books arriving on the stream, we can use the XSAG obtained by changing the bib and $book$ productions to

$$\begin{aligned} bib & ::= \{\text{print } \langle \text{books} \rangle\} \\ & \quad bib((book \cup article)^*) \\ & \quad \{\text{print } \langle / \text{books} \rangle\} \\ book & ::= \{ECHO\} book(title.author.author^*) \end{aligned}$$

Here we apply *ECHO* to the “book” subtrees, but not to articles. We output the opening and closing tags of the root node explicitly, and label the root node of the output produced by this XSAG “books”, rather than “bib”. \square

Based on the basic notion of XSAGs (bXSAGs) exemplified so far, we introduce the *easy XSAGs* (yXSAGs). These allow to annotate the regular expressions inside productions with attribution functions as well, which adds to the flexibility of the formalism.

Example 1.2 The yXSAG with production

$$\begin{aligned} article & ::= article(\{\text{print } \langle \text{article} \rangle\} \\ & \quad ((\{ECHO\}title). \\ & \quad (\{\text{print } \langle \text{authors} \rangle; ECHO\} \\ & \quad \quad (author.author^*) \\ & \quad \quad \{\text{print } \langle / \text{authors} \rangle\})) \\ & \quad \{\text{print } \langle / \text{article} \rangle\}) \end{aligned}$$

outputs articles basically as they arrive on the stream, but groups the authors of each article under a common *authors* node. Here, the second appearance of *ECHO* in the production applies to the tree region matched by the regular expression $author.author^*$, i.e., to the subtrees below *article* nodes that are rooted by *author* nodes. \square

Being attribute grammars, bXSAGs and yXSAGs of course support attributes. In order to assure scalability in the strictest sense, we require that attributes range over a finite domain fixed with the XSAG.

Example 1.3 Assume that our grammar assures that books in addition have a *year* of publication as a first child:

$$\begin{aligned} book & ::= book(year.title.author.author^*) \\ year & ::= year(PCDATA) \end{aligned}$$

Then, for instance, the yXSAG

```

bib ::= {print <books>}
      bib((book ∪ article)*
          {print </books>})
book ::= book(( {MATCH("2003", c)} year).
              ( {if c = true then
                 begin
                   print <book>; ECHO
                 end}
                (title.author.author*)
                {if c = true then print </book>})))

```

outputs books whose year of publication is 2003 with their title and author children, but without the years.

In this example yXSAG, we use a boolean-valued condition attribute² c whose value is set to *true* by *MATCH* if the string value of the *year* child of a book node matches “2003”, otherwise to *false*.³ This condition attribute is passed on through the document tree during its traversal. Just before we first visit the part of the tree below v that matches the regular expression *title.author.author** (i.e., the list of subtrees rooted by children of v that match the regular expression), if the value of attribute c is “true”, we output the opening tag $\langle \text{book} \rangle$ and echo the tree region matching that regular expression. On leaving the tree region, if c is true, we output the closing tag $\langle / \text{book} \rangle$.

Note that this yXSAG is equivalent to XML Query

```

<books>
{ for $x in //book where $x/year = 2003
  return </book> { $x/title } { $x/author } </book> }
</books>

```

on documents conforming to our grammar. □

Attribute grammars are well known in the field of compilers and have recently been revisited in the context of XML, for instance for grammar-directed XML publishing [3, 2]. Some of their theory relevant in the context of structured documents has been studied in [13, 12].

Our emphasis is on designing a *practical* formalism for query processing that is *relatively easy to use*. Attribute grammars are widely agreed to carry a strong intuition for specifying syntax-directed translations. In our setting, they provide a metaphor for strictly linear-time one-pass XML transformations that can be grasped very intuitively. This renders it relatively easy for a user to recognize or design queries which can be executed (scalably) on a stream, even if this intuition is paid for by our formalism being more operational than languages such as XML Query. While ease of use

²Note that in the technical sections of this paper, we will use a somewhat more explicit syntax when employing attributes (see e.g. Example 3.6).

³A *MATCH*(ρ, c) command applied to a tree region matches a regular expression ρ against the character text encountered while traversing the tree region. That is, when *MATCH* is used on entering a certain tree region – here, that of the *year* node – the value of attribute c is assigned upon leaving the tree region.

cannot be conclusively asserted based only on our own observations and the examples we provide, alternative formalisms such as deterministic pushdown transducers (DPDTs) are unsuitable as query languages to be used by humans; Query processors for languages such as XML Query, on the other hand, do not scale to streams. We can therefore argue that XSAGs achieve our goal of relative ease of use. Already bXSAGs are much more practical than DPDTs. yXSAGs permit very convenient and elegant nested attributions, which, as can be seen in Example 1.2 and others throughout the paper, allow to specify many interesting data transformations conveniently.

Contributions

The technical contributions of this paper are as follows.

- We examine the framework of extended regular tree grammars and craft grammar classes appropriate for attribution and stream processing.
- In order to be able to characterize yXSAGs properly, we develop the new notion of *strongly one-unambiguous regular expressions*, as well as some of its theory. These expressions are precisely those for which the *parse tree* of a word (analogously to the derivation tree of a grammar) can be unambiguously constructed *online*, with just a one-symbol lookahead, while processing the stream.

yXSAGs allow for attributions to be nested inside regular expressions by only permitting strongly one-unambiguous regular expressions in the right-hand sides of productions.

- We introduce and formally define our two notions of attribute grammars, bXSAGs and yXSAGs, and compare them with respect to usability.
- We introduce *XML-DPDTs*, deterministic pushdown transducers with a natural stack discipline that assures that the size of the stack remains strictly proportional to the depth of the XML tree and which can only accept well-formed XML documents. *XML-DPDTs* in a sense capture the intuition of scalable XML stream processing and serve as an expressiveness yardstick for XSAGs.
- We show that both bXSAGs and yXSAGs are precisely as expressive as *XML-DPDTs*. XSAGs provide the same quasi-optimal trade-off between expressiveness and evaluation cost as do *XML-DPDTs*.
- Finally, we study the complexity of XSAG query evaluation and their implementation.

The *structure* of this paper basically follows the order of contributions described above.

2 Regular Tree Grammars

Let Tag be a set of node labels (“tags”) and let $Char$ be a set of characters. An *extended regular tree grammar* is a grammar $G = (Nt, T, P, s)$ where

1. Nt is a set of nonterminals,
2. $T = Tag \cup Char$ is a set of terminals,
3. P is a set of productions $p ::= t(\rho)$ where $p \in Nt$, $t \in T$, and ρ is a regular expression over alphabet Nt such that if $t \in Char$ then $\rho = \epsilon$, and
4. $s \in Nt$ is the start production.

We assume the standard meaning of grammars and their derivations for which we refer to [8] for basic and to [10] for extended grammars.

Extended regular tree grammars (and DTDs, which are a dialect of extended regular tree grammars) are a convenient way to specify a class of unranked labeled trees and thus XML documents.

Let $Char = \{c_1, \dots, c_n\}$. As a shortcut, we define the regular expression macro

$$PCDATA ::= (c'_1 \cup \dots \cup c'_n)^*$$

which, using new nonterminals c'_1, \dots, c'_n and productions $c'_i ::= c_i(\epsilon)$ for each $1 \leq i \leq n$, can be used just like a terminal in right-hand sides of grammar productions. PCDATA accepts all character strings. As a further notational convenience, we will allow ourselves to be somewhat imprecise below whenever we only use PCDATA, but no characters, in our grammar definitions. Then we will list “PCDATA” among the terminals and will keep the nonterminals c'_1, \dots, c'_n unmentioned.

Given a production $n ::= t(\rho)$ of an extended regular tree grammar, let $\tau(\rho)$ denote the regular expression in which each nonterminal n' with production $n' ::= t'(\rho')$ occurring in ρ is replaced by the terminal (tag or character) t' .

Each extended regular tree grammar can be alternatively considered as an extended context-free word grammar (CFG), which is obtained by simply rewriting each right-hand side $tag(\rho)$ into $\langle tag \rangle \rho \langle /tag \rangle$. In such a CFG, the enclosing start- and end-tags act as determining *guards*, rendering the CFG deterministic (at least if $\tau(\rho)$ is unambiguous). The relevance of this stems from the fact that the deterministic context-free languages are precisely those recognizable by the deterministic pushdown automata (DPDA, see e.g. [8]). These DPDAs run comfortably on streams requiring only a stack of memory bounded by the depth of the input tree. It follows that we can validate all extended regular tree grammars (and thus DTDs, which are special regular tree grammars) on streams using only the inexpensive DPDA model.

This degree of determinism is sufficient for tree recognition on streams but not so for the evaluation of

attribute grammars. There, we need to be able to unambiguously refer to the atomic *symbols* (from $N \cup T$) in the regular expressions ρ to be able to access or assign attributes. In extended attribute grammars, a straightforward solution [12] is to require regular expressions to be *unambiguous*.

Example 2.1 Consider the grammar

$$\begin{aligned} bib & ::= \text{bib}((book_1 \cup book_2)^*) \\ book_1 & ::= \text{book}(\rho) \\ book_2 & ::= \text{book}(\rho) \end{aligned}$$

where ρ is some regular expression. The regular expression $(book_1 \cup book_2)^*$ is unambiguous, but

$$\tau(book_1 \cup book_2)^* = (\text{book} \cup \text{book})^*$$

is not. Therefore, when considering the tags of the children of the root “bib” node, we cannot determine where to apply which of the two “book” productions with their possibly different attributions. \square

We cannot look ahead in the stream beyond a non-terminal (which may stand for a large subtree that we do not want to buffer) when parsing the input. Thus, we will require the stronger notion of *one-unambiguity* for regular expressions $\tau(\rho)$. That is, we will require that $\tau(\rho)$ can be unambiguously parsed with just one symbol of lookahead.

2.1 One-unambiguity and TDLL(1)

By a *marking* of a regular expression ρ over alphabet Σ , we denote a regular expression ρ' such that each occurrence of an atomic symbol in ρ is replaced by the symbol with its position among the atomic symbols of ρ added as subscript. That is, the i -th occurrence of a symbol $a \in \Sigma$ in ρ is replaced by a_i . For instance, the marking of $(a \cup b)^*.a.a^*$ is $(a_1 \cup b_2)^*.a_3.a_4^*$. The reverse of a marking (indicated by $\#$) is obtained by dropping the subscripts.

Let ρ be a regular expression, ρ' its marking, and Σ' the marked alphabet used by ρ' . Then, ρ is called *one-ambiguous* iff there are words u, v, w over Σ' and symbols $x, y \in \Sigma'$ such that $uxv, uyw \in L(\rho')$, $x \neq y$, and $x^\# = y^\#$. A regular expression is called *one-unambiguous* if it is not one-ambiguous.

Example 2.2 Consider the regular expression $\rho = a^*.a$ and its marking $\rho' = a_1^*.a_2$. Let $u = a_1$, $x = a_2$, $y = a_1$, $v = \epsilon$, and $w = a_2$. Clearly, $uxv = a_1.a_2$ and $uyw = a_1.a_1.a_2$ are both words of $L(\rho')$, thus ρ is one-ambiguous. On the other hand, the equivalent regular expression $a.a^*$ is one-unambiguous. \square

Definition 2.3 ([10]) A *TDLL(1)-Grammar* is an extended regular tree grammar in which for each regular expression ρ in the right-hand side of a production, $\tau(\rho)$ is one-unambiguous. \square




Figure 1: Parse trees of regular expressions $a^* \cup b^*$ (a) and $a.a^*$ (b) with nodes annotated with markings.

Example 2.4 The grammar of Example 1.1 is TDLL(1). On the other hand, since the grammar of Example 2.1 contains a regular expression ρ such that $\tau(\rho)$ is not even unambiguous, that grammar is not TDLL(1). \square

Remark 2.5 (One-unambiguity in DTDs) For XML elements that exclusively have elements as children (but no character data), the W3C recommendation [4]⁴ explicitly requires a one-unambiguous *content model* (that is, right-hand side regular expression). As mentioned in Appendix E of [4], this is to assure compatibility with SGML.

Productions defining elements with mixed content (also containing character data) must be constructed according to either the pattern

$$n_0 ::= (\text{PCDATA} \cup n_1 \cup \dots \cup n_m)^*$$

or $n_0 ::= \text{PCDATA}$ (where n_0, \dots, n_m are DTD element names, i.e., nonterminals). Clearly, all regular expressions such constructed are also one-unambiguous.

Thus, since DTDs also contain at most one production $n ::= t(\rho)$ for each “element” t (and thus if ρ is one-unambiguous, $\tau(\rho)$ is as well), DTDs are TDLL(1) grammars (see also [10]). \square

2.2 Strong One-unambiguity and STDLL(1)

One-unambiguity and TDLL(1) grammars allow us to use attributed regular tree grammars on XML streams. However, as we will show below, the ability to add attribution functions into the regular expressions at the right-hand sides of productions will allow us to write many practical queries in a much more user-friendly fashion. Our machinery for achieving this is the notion of *strongly one-unambiguous* regular expressions.

By a *bracketing* of a regular expression ρ , we intuitively refer to a *marking of the nodes in the parse tree* of ρ using distinct indexes. Here, we assume that the indexes are assigned as by a depth-first left-to-right traversal of the parse tree, that is, in document order (see Figure 1 for two examples). The bracketing ρ^\square is obtained by inductively mapping each subexpression π of ρ with index i to $[i.\pi.]_i$. Thus, a bracketing is a regular expression

over the alphabet $\Sigma \cup \Gamma$, where $\Gamma = \{[i.]_i \mid i \in \{1, 2, 3, \dots\}\}$. (We assume Σ and Γ disjoint.) For example, $[1.((([2.([3.a.]_3)^*.]_2) \cup ([4.([5.b.]_5)^*.]_4)).]_1]$ is the bracketing of $a^* \cup b^*$ and $[1.((([2.a.]_2).([3.([4.a.]_4)^*.]_3)).]_1]$ is the bracketing of $a.a^*$.

Definition 2.6 Let ρ be a regular expression and ρ^\square be its bracketing. A regular expression ρ is called *strongly one-unambiguous* iff there do not exist words u, v, w over $\Sigma \cup \Gamma$, words $\alpha \neq \beta$ over Γ , and a symbol $x \in \Sigma$ such that $u\alpha xv, u\beta xw \in L(\rho^\square)$ or $u\alpha, u\beta \in L(\rho^\square)$. \square

Example 2.7 The regular expression $a^* \cup b^*$ (see Figure 1 (a)) is one-unambiguous but not strongly one-unambiguous. Consider the bracketing $[1.((([2.([3.a.]_3)^*.]_2) \cup ([4.([5.b.]_5)^*.]_4)).]_1]$. The empty word can be matched in two possible ways, namely as $u.\alpha = [1.[2.]_2.]_1]$ and as $u.\beta = [1.[4.]_4.]_1]$. The equivalent regular expression $a.a^*$ is strongly one-unambiguous. \square

Example 2.8 The regular expression $(a^*)^*$ with the bracketing $[1.([2.([3.a.]_3)^*.]_2)^*.]_1]$ is not strongly one-unambiguous. For instance, for the word $a.a$, there are the bracketings $[1.[2.[3.a.]_3.[3.a.]_3.]_2.]_1]$ and $[1.[2.[3.a.]_3.]_2.[2.[3.a.]_3.]_2.]_1]$ ($u = [1.[2.[3.a.]_3.]_2.]_1]$, $\beta = [3.]_3.]_2.]_2.$, $x = a$, $v = w = [3.]_3.]_2.]_1]$). \square

Definition 2.9 An *STDLL(1) Grammar* is an extended regular tree grammar in which for each regular expression ρ in the right-hand side of a production, $\tau(\rho)$ is strongly one-unambiguous. \square

Obviously, all STDLL(1) grammars are also TDLL(1) grammars.

Remark 2.10 We suspect that most practical DTDs use only strongly one-unambiguous regular expressions in productions and are thus STDLL(1) grammars. Strong one-unambiguity is only a short way from one-unambiguity, and many of the most widely used forms of regular expressions are actually strongly one-unambiguous (e.g., regular expressions of the form $(e_1 \cup \dots \cup e_m)^*$, where e_1, \dots, e_m are element names). In particular, the syntactic restriction on mixed-content models mentioned in Remark 2.5 ensures that such regular expressions are guaranteed to be strongly one-unambiguous. \square

Throughout the paper, we will assume a generalized notion of parse trees for STDLL(1) grammars, which reflects the structure of the regular expressions occurring in the grammars. For TDLL(1) grammars, the parse trees are simply the usual document trees associated to XML documents. However, for STDLL(1) grammars, the parse trees incorporate the derivation trees of the regular expressions occurring in productions in a natural way, which we illustrate in the next

⁴Sections 3.2.1, 3.2.2, and Appendix E.

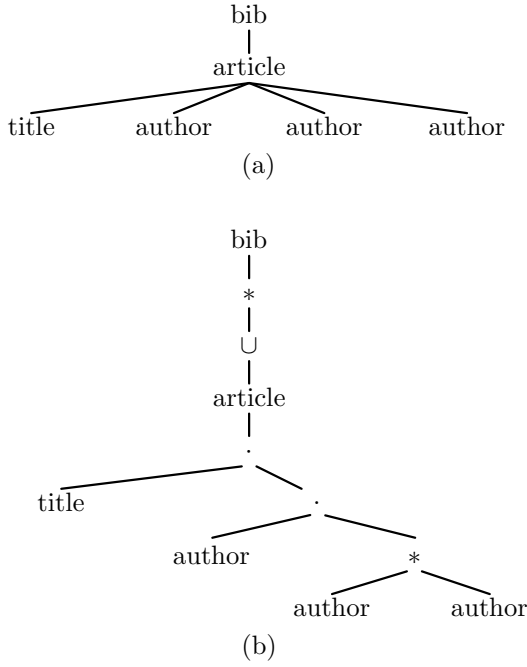


Figure 2: Parse trees of Example 2.11.

example. Notably, in STDLL(1) parse trees, nodes labeled “∪” have precisely one, nodes labeled “.” have precisely two, and nodes labeled “*” may have an arbitrary number of children (possibly none).

Example 2.11 The extended regular tree grammar G of Example 1.1 is an STDLL(1) grammar. Consider the XML document

```

<bib> <article>
  <title/> <author/> <author/> <author/>
</article> </bib>

```

Using G as a TDLL(1) grammar, the document parses into the tree depicted in Figure 2 (a). To be precise, we assume here that the operation “.” associates to the right and that the production

$$article ::= article(title.author.author^*)$$

of G is thus equivalent to

$$article ::= article(title.(author.author^*)).$$

The parse tree for G viewed as an STDLL(1) grammar is shown in Figure 2 (b). \square

3 XML Stream Attribute Grammars

We are now in the position to define our main attribute grammar formalism, XML Stream Attribute Grammars (XSAGs).

3.1 XSAGs in the Abstract

Definition 3.1 (Syntax) Let $Att = \{a_1, \dots, a_k\}$ be a set of attributes and Dom be a finite set of *domain values* (to invoke an alternative intuition, of *states*).

Let F_I denote the class of partial functions

$$f_I : Dom^k \rightarrow Dom^k \times string,$$

and let F_{II} denote the partial functions

$$f_{II} : Dom^{2k} \rightarrow Dom^k \times string.$$

We will refer to the functions of F_I and F_{II} as *attribution functions*. (We will introduce a language for implementing these partial functions in Section 3.2).

A basic XSAG (bXSAG) is an attributed extended regular tree grammar $G = (Nt, T, P, s)$ with nonterminals Nt , terminals $T = Tag \cup Char$, and productions in P that each are of one of the four forms

$$n ::= t(\rho) \quad n ::= \{f_I\} t(\rho)$$

$$n ::= t(\rho) \{f_{II}\} \quad n ::= \{f_I\} t(\rho) \{f_{II}\}$$

where $n \in Nt$, $t \in T$, $f_I \in F_I$, $f_{II} \in F_{II}$, and ρ is either ϵ or a regular expression over Nt such that $\tau(\rho)$ is one-unambiguous.

The abstract syntax of an *attributed regular expression* over symbols Σ can be specified by the EBNF

$$\begin{aligned}
aregex & ::= (“\{” F_I “\}”)? aregex_0 (“\{” F_{II} “\}”)?; \\
aregex_0 & ::= \Sigma \mid aregex “.” aregex \mid \\
& \quad aregex “\cup” aregex \mid aregex “^*”.
\end{aligned}$$

An easy XSAG (yXSAG) is an attributed extended regular tree grammar $G = (Nt, T, P, s)$ where each production in P is of one of the four forms

$$n ::= t(\alpha) \quad n ::= \{f_I\} t(\alpha)$$

$$n ::= t(\alpha) \{f_{II}\} \quad n ::= \{f_I\} t(\alpha) \{f_{II}\}$$

where α is either ϵ or an attributed regular expression over symbols Nt such that for the regular expression ρ obtained from α by removing the attributions (enclosed in curly braces), $\tau(\rho)$ is *strongly* one-unambiguous. \square

The main purpose of the grammar component of an XSAG is to unambiguously map XML documents to parse trees.⁵ The only differences between bXSAGs and yXSAGs are that the former use TDLL(1) grammars while the latter use STDLL(1) grammars, and that in yXSAGs, right-hand side regular expressions may be attributed⁶

It remains to specify how our attribute grammars are evaluated on such parse trees. Obviously, there is a natural method of assigning the attribution functions from F_I and F_{II} to nodes of the parse tree. For the sake of simplicity, we assume that each node v of the parse tree is assigned two functions $f_I^v \in F_I$ and $f_{II}^v \in F_{II}$

⁵However, for evaluating XSAGs it will not at any time be necessary to maintain entire parse trees in memory.

⁶And indeed, precisely the restriction to STDLL(1) grammars makes it safe to attribute regular expressions in XSAGs.

through the attribute grammar definition. Where this has not been the case (because for a nonterminal of an bXSAG resp. a right-hand side subexpression of a yXSAG either the left or right attribution function were not provided), the defaults are

$$f_I^v(a_1, \dots, a_k) := \langle a_1, \dots, a_k, \epsilon \rangle$$

and

$$f_{II}^v(a_1, \dots, a_k, a_{k+1}, \dots, a_{2k}) := \langle a_{k+1}, \dots, a_{2k}, \epsilon \rangle.$$

bXSAGs and yXSAGs are (attributed) *extended* regular tree grammars. For such grammars, nodes of the parse tree may have an arbitrary number of children. When dealing with streams, we generally cannot store the attribute values of all these children in memory. We thus have to introduce special restrictions to be able to deal with streams on one hand and at the same time assure ease of use and expressiveness to cover practical queries on the other.

We define XSAGs as L-attributed grammars, i.e., attribute grammars whose attributes are evaluated by a single depth-first left-to-right traversal of the document tree. Each node v of the parse tree is visited twice (the visits are referred to as I and II), (I) from the previous sibling or the parent of v (if v has no previous sibling) and (II) on returning from the rightmost child of v .

Definition 3.2 (Semantics) Let $q_{\perp} \in Dom$ be a special “uninitialized” value. We evaluate an XSAG on a parse tree T by a depth-first traversal of T in which we compute each attribute $a_i \in Att$ for each node v as follows.

$$v.in[1].a_i := \begin{cases} q_{\perp} & \dots v \text{ is the root node} \\ v_0.out[1].a_i & \dots v \text{ is the first child} \\ & \text{of } v_0 \\ v_0.out[2].a_i & \dots v \text{ is the right} \\ & \text{sibling of } v_0 \end{cases}$$

$$v.in[2].a_i := \begin{cases} v.out[1].a_i & \dots v \text{ has no children} \\ w.out[2].a_i & \dots w \text{ is the rightmost} \\ & \text{child of } v \end{cases}$$

In the first visit of node v , we compute

$$\langle v.out[1].a_1, \dots, v.out[1].a_k, \sigma \rangle := f_I^v(v.in[1].a_1, \dots, v.in[1].a_k)$$

and write σ to the stream. In the second, we compute

$$\langle v.out[2].a_1, \dots, v.out[2].a_k, \sigma \rangle := f_{II}^v(v.in[1].a_1, \dots, v.in[1].a_k, v.in[2].a_1, \dots, v.in[2].a_k)$$

and write σ to the stream. In case f_I^v or f_{II}^v is undefined on its input, the evaluation terminates and the input is rejected. \square

Even though this semantics definition may seem involved, we believe that its application is natural; examples will be provided later on in this section.




Figure 3: bXSAG parse tree and traversal of Example 3.3.

3.2 Concrete XSAGs

We use a simple programming language for defining attribution functions. This language is basically a fragment of Pascal, comprising the following constructs: (1) if-then-else statements, (2) blocks of multiple commands starting with the keyword “begin” and ending with “end”, (3) boolean formulas – using “and”, “or”, and “not” – over equality conditions comparing two r-values⁷ (used in if-statements), (4) assignments of r-values to l-values, (5) the keyword “reject” for terminating the computation and rejecting the input, and (6) “print” statements taking a constant string as argument.

We assume that all l- and r-values range over the same finite domain Dom . For (partial) functions in F_I ,

$$lval_I := \{out[1].a \mid a \in Att\}$$

$$rval_I := \{in[1].a \mid a \in Att\} \cup Dom$$

and for (partial) functions in F_{II} ,

$$lval_{II} := \{out[2].a \mid a \in Att\}$$

$$rval_{II} := \{in[1].a, in[2].a \mid a \in Att\} \cup Dom.$$

As a default, attributes $out[1].a$ (resp., $out[2].a$) that do not appear on the left-hand side of an assignment obtain the value $in[1].a$ (resp., $in[2].a$).

Such a program defines functions in F_I resp. F_{II} in the obvious way, with the notable fact that the functions are assumed undefined for inputs for which the “reject” statement is called.

Example 3.3 Consider the bXSAG G with productions

$$bib ::= \{ \text{print } \langle bib \rangle \} bib((book \cup article)^*)$$

$$\{ \text{print } \langle /bib \rangle \}$$

$$book ::= \{ \text{if } in[1].prev = q_{article} \text{ then} \\ \text{print } \langle book/ \rangle ; \\ out[1].prev := q_{book} \} book(\epsilon)$$

$$article ::= \{ \text{print } \langle article/ \rangle ; out[1].prev := q_{article} \} \\ article(\epsilon)$$

⁷We follow the widely used convention of calling constructs of our language that may appear on the right side of an assignment *r-values* and those that may appear on the left side of an assignment *l-values*.

F	Attribute value		Output
	before	after	
$f_I^{v_1}$	q_{\perp}	q_{\perp}	$\langle \text{bib} \rangle$
$f_I^{v_2}$	q_{\perp}	q_{article}	$\langle \text{article} / \rangle$
$f_{II}^{v_2}$	q_{article}	q_{article}	ϵ
$f_I^{v_3}$	q_{article}	q_{book}	$\langle \text{book} / \rangle$
$f_{II}^{v_3}$	q_{book}	q_{book}	ϵ
$f_I^{v_4}$	q_{book}	q_{book}	ϵ (!!)
$f_{II}^{v_4}$	q_{book}	q_{book}	ϵ
$f_{II}^{v_1}$	q_{book}	q_{book}	$\langle / \text{bib} \rangle$

Table 1: Run of bXSAG G of Example 3.3.

G uses a single attribute $prev$ that ranges over $Dom = \{q_{\perp}, q_{\text{book}}, q_{\text{article}}\}$ and which is initialized with q_{\perp} . The grammar requires the input to consist of a dummy bibliography database containing book and article nodes without children. As output, the XSAG writes a root node labeled “bib”, to which it assigns bachelor nodes labeled “book” and “article” as children, filtering out books which are not right neighbors of articles⁸.

The parse tree of XML document

$\langle \text{bib} \rangle \langle \text{article} / \rangle \langle \text{book} / \rangle \langle \text{book} / \rangle \langle / \text{bib} \rangle$

is shown in Figure 3. Of course,

$$\begin{aligned}
f_I^{v_1} &= \{\text{print } \langle \text{bib} \rangle\} \\
f_{II}^{v_1} &= \{\text{print } \langle / \text{bib} \rangle\} \\
f_I^{v_2} &= \{\text{print } \langle \text{article} / \rangle; \text{out}[1].prev := q_{\text{article}}\} \\
f_I^{v_3} &= \{\text{if } in[1].prev = q_{\text{article}} \text{ then print } \langle \text{book} / \rangle; \\
&\quad \text{out}[1].prev := q_{\text{book}}\}
\end{aligned}$$

Stated differently,

$$\begin{aligned}
f_I^{v_1} &: x \mapsto (x, \langle \text{bib} \rangle) \\
f_{II}^{v_1} &: x \mapsto (x, \langle / \text{bib} \rangle) \\
f_I^{v_2} &: x \mapsto (q_{\text{article}}, \langle \text{article} / \rangle) \\
f_I^{v_3} &: x \mapsto \begin{cases} (q_{\text{book}}, \langle \text{book} / \rangle) & \dots \quad x = q_{\text{article}} \\ (q_{\text{book}}, \epsilon) & \dots \quad \text{otherwise} \end{cases}
\end{aligned}$$

Moreover, $f_I^{v_4} = f_I^{v_3}$ and

$$f_{II}^{v_2}, f_{II}^{v_3}, f_{II}^{v_4} : (x_1, x_2) \mapsto (x_2, \epsilon),$$

i.e. is the default.

G is evaluated on the parse tree as shown in Table 1. The four columns have the following meaning. The first column shows which attribution function is applied in the respective step. The second and third columns show the value of our attribute $prev$ before and after the application of the attribution function, respectively, and the rightmost column shows which

⁸This is a somewhat contrived example but illustrates a number of important points related to the evaluation of XSAGs.

output is produced and written to the output stream. Clearly, G outputs

$\langle \text{bib} \rangle \langle \text{article} / \rangle \langle \text{book} / \rangle \langle / \text{bib} \rangle$

and accepts its input. \square

Example 3.4 We continue the previous example. Alternatively, if we want to reject the stream⁹ if two books arrive in sequence, we define f_I^{book} as

{if $in[1].prev = q_{\text{article}}$ then
begin print $\langle \text{book} / \rangle$; $out[1].prev := q_{\text{book}}$ end
else reject}

The meaning of this attribution function is

$$f_I^{\text{book}} : x \mapsto \begin{cases} (q_{\text{book}}, \langle \text{book} / \rangle) & \dots \quad x = q_{\text{article}} \\ \text{undefined} & \dots \quad \text{otherwise.} \end{cases}$$

This new XSAG rejects the input of Example 3.3. \square

3.3 Built-in Macros

For the convenient definition of queries using XSAGs, we introduce three standard built-in macros, *ECHO*, *ECHO OFF*, and *MATCH*. These are redundant with the formalism presented so far, but allow to define queries in a more concise way.

Echo and Echo Off

If macro *ECHO* is used in an attribution function (which must be of type F_I), the subtree of the parse tree to which the attribution function applies is copied to the output. Correspondingly, macro *ECHO OFF* can be used to override *ECHO* and suppress the output of certain XML subtrees.

To realize *ECHO*, we define a boolean attribute $echo \in Att$, initialized with *false*. For every production $\{f_I\}t(\rho)\{f_{II}\}$,

- if $t \in Tag$, we append

if $in[1].echo = true$ then print $\langle t \rangle$

to f_I and

if $in[1].echo = true$ then print $\langle /t \rangle$

to f_{II} .

- If $t \in Char$ we append

if $in[1].echo = true$ then print t

to f_I .

⁹This could be alternatively achieved by modifying the grammar rather than the attributions as done in this example, but the goal here is to illustrate the use of the “reject” statement, resp. partially undefined attribution functions.

- In both cases, we then append

$$out[2].echo := in[1].echo$$

to f_{II} .¹⁰

- Moreover, occurrences of *ECHO* are replaced by $out[1].echo := true$ and occurrences of *ECHO OFF* are replaced by $out[1].echo := false$.

Example 1.1 illustrates the use of the macro *ECHO*. Below, we show an example that combines *ECHO* and *ECHO OFF*.

Example 3.5 The bXSAG

$$\begin{aligned} bib &::= \{ECHO\} bib(book^*) \\ book &::= book(title.author. \\ &\quad (\{ECHO\ OFF\} (author^*))) \\ title &::= title(PCDATA) \\ author &::= author(PCDATA) \end{aligned}$$

outputs books with their titles and first authors (dropping all further authors). \square

String Matching

Macro $MATCH(\rho, c)$ matches the string value of an XML node v (i.e., the string obtained by concatenating all the character data encountered in a document-order traversal¹¹ of the subtree of v) against a regular expression ρ , yielding *true* or *false* as a value for the user-defined condition attribute c . $MATCH$ may only be used in attribution functions of type F_I . The evaluation result becomes available in the corresponding attribution function of type F_{II} .

Example 3.6 We slightly extend Example 1.3. The yXSAG production

$$\begin{aligned} book &::= book(\{\{MATCH(2003, c)\}year\}. \\ &\quad (\{\text{if } in[1].c = true \text{ then} \\ &\quad \quad \text{begin} \\ &\quad \quad \quad \text{print } \langle book \rangle; ECHO \\ &\quad \quad \quad \text{end}\} \\ &\quad (title.author.author^*) \\ &\quad \{\text{if } in[2].c = true \text{ then} \\ &\quad \quad \text{print } \langle year \rangle 2003 \langle /year \rangle \langle /book \rangle\})) \end{aligned}$$

selects those books for which child “year” has string value “2003”; moreover, the year is output as the rightmost child of book, rather than as the leftmost as required for the input. (We omit productions defining bib, year, title, and author, which are as in Example 3.7.) \square

MATCH can be easily implemented by compiling ρ into a DFA and running it on the characters visited while traversing a subtree.

¹⁰Thus, even though *ECHO* may be overshadowed by *ECHO OFF* within a subtree, we may not accidentally create malformed documents, as on leaving a node on which *ECHO* or *ECHO OFF* is used, attribute *echo* is reset to its former value.

¹¹That is, a depth-first left-to-right traversal of the subtree.

3.4 bXSAGs vs. yXSAGs

As we show in the next section, bXSAGs and yXSAGs have the same expressive power. However, yXSAGs are more convenient to use. In particular, it is often necessary to introduce more attributes and more complicated attribution functions to encode a given query as a bXSAG than to encode it as a yXSAG.

Example 3.7 Consider the yXSAG with productions

$$\begin{aligned} bib &::= bib(\{\text{print } \langle bib \rangle\} article^* \{\text{print } \langle /bib \rangle\}) \\ article &::= article(\{\{\text{print } \langle article_short \rangle; ECHO\} \\ &\quad (title.author.author^*) \\ &\quad \{\text{print } \langle /article_short \rangle\}\} \\ &\quad \cup \\ &\quad (\{\text{print } \langle article_long \rangle; ECHO\} \\ &\quad (year.title.author.author^*.pub) \\ &\quad \{\text{print } \langle /article_long \rangle\})) \\ title &::= title(PCDATA) \\ author &::= author(PCDATA) \\ year &::= year(PCDATA) \\ pub &::= pub(PCDATA) \end{aligned}$$

Article entries can appear either in a short version with title and authors only, or in a long version which also contains a year and a publisher. In the former case, article nodes are relabeled as *article_short*, and in the latter case, article nodes are relabeled *article_long*.

The following bXSAG is equivalent to the above yXSAG:

$$\begin{aligned} bib &::= bib(\{\text{print } \langle bib \rangle\} article^* \{\text{print } \langle /bib \rangle\}) \\ article &::= \{\text{out}[1].state := } q_{unknown}\} \\ &\quad article((title.author.author^*) \cup \\ &\quad (year.title.author.author^*.pub)) \\ &\quad \{\text{if } in[2].state = } q_{short}\} \text{ then} \\ &\quad \quad \text{print } \langle /article_short \rangle \\ &\quad \text{else if } in[2].state = } q_{long}\} \text{ then} \\ &\quad \quad \text{print } \langle /article_long \rangle\} \\ title &::= \{\text{if } in[1].state = } q_{unknown}\} \text{ then} \\ &\quad \text{begin } out[1].state := } q_{short}; \\ &\quad \quad \text{print } \langle article_short \rangle \text{ end; } ECHO\} \\ &\quad title(PCDATA) \\ author &::= \{ECHO\} author(PCDATA) \\ year &::= \{\text{if } in[1].state = } q_{unknown}\} \text{ then} \\ &\quad \text{begin } out[1].state := } q_{long}; \\ &\quad \quad \text{print } \langle article_long \rangle \text{ end; } ECHO\} \\ &\quad year(PCDATA) \\ pub &::= \{ECHO\} pub(PCDATA) \end{aligned}$$

While there are other ways of encoding our query as a bXSAG, it does not seem to be possible to represent the query as a bXSAG without using attributes. \square

It is easy to verify that bXSAGs equivalent to the yXSAGs of Examples 1.2 and 3.6 are also much more complicated.

4 Expressive Power of XSAGs

We introduce deterministic pushdown transducers (DPDTs) as deterministic pushdown automata with output which accept by empty stack. As with pushdown automata [8], the DPDTs accepting by empty stack are equivalent to the DPDTs accepting by final state.

Definition 4.1 (DPDT) A *deterministic pushdown transducer* is a tuple

$$\mathcal{T} = (Q, \Sigma, \Gamma, \Omega, \delta, q_0, Z_0)$$

where Q is a finite set of states, Σ , Γ , and Ω are the alphabets for input tape, stack, and output tape respectively, δ is the transition function

$$\delta : Q \times (\Sigma \cup \{\epsilon\}) \times \Gamma \rightarrow Q \times \Omega^* \times \Gamma^*,$$

$q_0 \in Q$ denotes the initial state, and Z_0 the initial stack symbol.

We define a run of \mathcal{T} by means of instantaneous descriptions (IDs): An ID is a quadruple

$$(q, w, \alpha, o) \in Q \times \Sigma^* \times \Gamma^* \times \Omega^*,$$

where q is a state, w a string of input symbols, α a string of stack symbols, and o a string of output symbols. We make a transition

$$(q, aw, X\alpha, o) \vdash (q', w, \gamma\alpha, o\sigma)$$

if $\delta(q, a, X) = (q', \sigma, \gamma)$ where $a \in \Sigma \cup \{\epsilon\}$, $X \in \Gamma$, $\sigma \in \Omega^*$, and $q' \in Q$.

Here, $\gamma \in \Gamma^*$ is the string of stack symbols which replace X on top of the stack. For $\gamma = \epsilon$, the stack is popped, whereas for $\gamma = X$, the stack remains unchanged. If $\gamma = YX$, then stack symbol Y is pushed on top of X .

Let \vdash^* be the reflexive and transitive closure of \vdash . \mathcal{T} accepts by empty stack if

$$(q_0, w, Z_0, \epsilon) \vdash^* (q, \epsilon, \epsilon, o)$$

for $w \in \Sigma^*$, $q \in Q$, and $o \in \Omega^*$.

Two conditions must hold for \mathcal{T} to be deterministic:

1. For each $q \in Q$ and $X \in \Gamma$, if $\delta(q, \epsilon, X)$ is nonempty, then $\delta(q, a, X)$ is empty for all $a \in \Sigma$;
2. For no $q \in Q$, $X \in \Gamma$, and $a \in \Sigma \cup \{\epsilon\}$ does $\delta(q, a, X)$ contain more than one element. \square

Definition 4.2 (Well-formedness) An XML document is called *well-formed* iff it conforms to an extended regular tree grammar $G = (Nt, T, P, s)$ where the production for the start symbol is $s ::= t(\rho)$, $t \in \text{Tag}$. \square

A well-formed document contains at least one element (i.e. the root element) and has element start-tags and end-tags properly nested within each other. Furthermore, the first symbol in the document must be the opening tag for the root node. An XML document is *malformed* if it is not well-formed.

Definition 4.3 (XML-DPDT) Let input alphabet $\Sigma = \{\langle t \rangle, \langle /t \rangle \mid t \in \text{Tag}\} \cup \text{Char}$ consist of matching opening and closing tags and characters. An XML-DPDT is a DPDT $\mathcal{T} = (Q, \Sigma, \Gamma, \Omega, \delta, q_0, Z_0)$ which rejects malformed XML documents and for which the transition function δ is restricted as follows:

$$\delta : Q \times (\Sigma \cup \{\epsilon\}) \times \Gamma \rightarrow Q \times \Omega^* \times \Gamma^d, d \leq 2$$

- In the very first transition the initial stack symbol Z_0 is replaced: $\delta(q_0, \langle t \rangle, Z_0) = (p, \sigma, Y)$ for $\langle t \rangle \in \Sigma$, $p \in Q$, $\sigma \in \Omega^*$, and $Y \in \Gamma$.
- For all other configurations of $q \in Q$ and $X \in \Gamma$, a symbol is only pushed on the stack when an opening tag is read from the input stream: $\delta(q, \langle t \rangle, X) = (p, \sigma, YX)$ for $\langle t \rangle \in \Sigma$, $p \in Q$, $\sigma \in \Omega^*$, and $Y \in \Gamma$.
- A symbol is only popped from the stack when a closing tag is encountered in the input stream, i.e. $\delta(q, \langle /t \rangle, X) = (p, \sigma, \epsilon)$ for $q, p \in Q$, $\langle /t \rangle \in \Sigma$, $X \in \Gamma$, and $\sigma \in \Omega^*$. \square

The conditions required in the definition of XML-DPDT are only natural in the context of XML stream processing: The size of the stack is bounded by the maximum depth of the incoming document tree. Moreover, the input has to start with the root element of the XML document being read. Due to acceptance by empty stack, only well-formed XML-documents are accepted.

Lemma 4.4 For each XML-DPDT with ϵ -transitions there exists an equivalent ϵ -free XML-DPDT.

We call XSAG G and XML-DPDT \mathcal{T} output-equivalent if for any given input w , the outputs of G and \mathcal{T} are equal (denoted $G(w) = \mathcal{T}(w)$). XSAG G and XML-DPDT \mathcal{T} are called *equivalent* if they are output-equivalent and accept the same inputs, i.e. $L(G) = L(\mathcal{T})$.

Both bXSAGs and yXSAGs possess exactly the same expressiveness as XML-DPDTs¹².

Theorem 4.5 For each bXSAG, there is an equivalent XML-DPDT.

¹²Note, however, that there are bXSAGs that are no yXSAGs and vice versa, so Theorems 4.5 and 4.7 are not redundant.

Theorem 4.6 *For each XML-DPDT, there is an equivalent XSAG which is both a bXSAG and a yXSAG.*

Theorem 4.7 *For each yXSAG, there is an equivalent XML-DPDT.*

The lengthy proofs for these three theorems will be provided in the long version of this paper.

5 Efficient Evaluation of XSAGs

The proofs of Theorems 4.5 and 4.7 are based on a translation to DPDTs which, as a method for evaluating XSAGs, has the strong point that once the DPDT has been created, the query evaluation time is in principle independent of the size of the XSAG/DPDT and only depends on the input data.

Corollary 5.1 *An XSAG G can be evaluated on a tree T in time $O(f(|G|)+|T|)$ using only a stack of memory of size $O(\text{depth}(T))$.*

However, the DPDTs of the construction of the proof of Theorem 4.5 are of size *exponential* in the number k of attributes in the XSAG, i.e., f is $O(2^k)$.

The exponential-time compilation phase can be avoided by using a simple hybrid evaluation method in which the grammars (and in particular the regular expressions appearing in the grammar productions) are compiled into transducers which however *interpret* the attribution functions (rather than materializing the *graphs* of the attribution functions as is done in our proofs). Thus one obtains an XSAG evaluation method which runs scalably on streams and which is strictly polynomial in the size of the XSAG.

Theorem 5.2 *A bXSAG G can be evaluated on a tree T in time $O(|G|^2 + |T| \cdot |G|)$ using a stack of size $O(\text{depth}(T))$.*

Theorem 5.2 also makes use of the fact that DFAs for one-unambiguous regular expressions can be computed in polynomial (actually, quadratic) time¹³ [5].

For yXSAGs, the construction of the proof of Theorem 4.7 is in addition exponential in the maximum depth of the parse trees of the regular expressions used (which only depend on the XSAG). This can be resolved by pushing attributes onto the stack at yXSAG regular expression nodes as well. The stack consumption of course remains proportional to the depth of the input tree.

The main technical challenge we have to deal with when evaluating yXSAGs is the matching of attributed regular expressions on the stream and the invocation of attribution functions at the right time.

¹³This construction – that of the Glushkov automaton of a regular expression – also provides a procedure for deciding one-ambiguity with the same complexity. A regular expression is known to be one-unambiguous precisely if its Glushkov automaton is deterministic.

A finite-state transducer (FST) is an NFA with output, which in each transition $q \xrightarrow{a/w} q'$ from state q to q' on input symbol a outputs a fixed word w . A deterministic finite-state transducer (DFT) is an FST that is deterministic, i.e., which is a DFA if the output is ignored and for which no two transitions $q \xrightarrow{a/v} q'$ and $q \xrightarrow{a/w} q'$ exist such that $v \neq w$.

Below, in regular expressions of the form $\rho \odot$, let \odot be a new end-marker symbol that does not occur in ρ .

Theorem 5.3 *Let ρ be a regular expression. Then, there is an FST $\mathcal{A}^{\square}(\rho)$ which*

1. recognizes $L(\rho \odot)$,
2. is deterministic iff ρ is strongly one-unambiguous,
3. if ρ is strongly one-unambiguous, outputs the bracketing of word w for each $w \odot \in L(\rho \odot)$, and
4. can be computed in time $O(|\rho|^3)$.

Using the DFT construction of Theorem 5.3, the preprocessing phase for yXSAGs takes time cubic in the size of each of the productions.

Theorem 5.4 *A yXSAG G can be evaluated on a tree T in time $O(|G|^3 + |T| \cdot |G|)$ using a stack of size $O(\text{depth}(T) * |G|)$.*

6 Discussion and Conclusions

The goal of this paper was to develop a framework for query formulation which

1. satisfies our criteria for scalable query processing on streams,
2. has a good and well-justified foundation, and
3. is user-friendly, i.e. allows to state many common queries quickly and easily.

We can argue that XSAGs satisfy these three desiderata.

(1) Each XSAG can be translated into a DPDT with a stack discipline that assures that the size of the stack remains proportional to the depth of the XML tree. This is known to be the minimum amount of memory required to do any meaningful (sequential) processing of XML data [9]. Of course queries are evaluated strictly in linear time.

(2) Throughout the paper, we have explained and justified our design choices. Regular tree grammars are a commonly accepted grammar formalism for XML (as are DTDs, which are restricted regular tree grammars). In the right-hand sides of the productions of such grammars, we use regular expressions to be able to parse nodes with an unbounded number of children. Our restriction of these regular expressions to strongly one-unambiguous ones in the case of yXSAGs allows for precisely those expressions for which the parse trees of words can be unambiguously generated

using a lookahead of only one symbol (a necessity in stream processing). Having the regular expressions inside grammar productions available for attribution allows to conveniently define attribute grammars for unranked trees, and to approach the usability of XML Query with a formalism that allows for much better control of complexity.

We have precisely characterized the expressive power of XSAGs relative to deterministic pushdown transducers.

Note that our formalism fully fits into the classical framework of attribute grammars (and more precisely, L-attributed grammars), even if we did not introduce, say, the distinction between synthesized and inherited attributes.

(3) A number of examples in this paper and our experiences with many more demonstrate that XSAGs are of practical value, and that they fill an important void in the design space of tailored query languages.

Earlier in this paper, we defined XSAGs with attributes ranging exclusively over a finite domain to be able to assure scalability and memory bounds in the strongest sense. However, it is desirable and often justified to generalize this framework to make certain uniformity assumptions and to allow for values from an infinite domain. In the future, we plan to carry out a more detailed study of conservative extensions of our formalism with small buffers (using uniformity assumptions for numbers, small strings, and small subtrees).

Acknowledgments

We thank Peter Buneman, Bernhard Stegmaier, Val Tannen, and the anonymous reviewers of DBPL 2003 for their suggestions, which helped to substantially improve the presentation of the paper.

References

- [1] A. V. Aho, R. Sethi, and J. D. Ullman. *Compilers – Principles, Techniques, and Tools*. Addison-Wesley, 1986.
- [2] M. Benedikt, C. Y. Chan, W. Fan, J. Freire, and R. Rastogi. “Capturing both Types and Constraints in Data Integration”. In *Proceedings of the 2003 ACM SIGMOD International Conference on Management of Data (SIGMOD’03)*, pages 277–288, 2003.
- [3] M. Benedikt, C. Y. Chan, W. Fan, R. Rastogi, S. Zheng, and A. Zhou. “DTD-Directed Publishing with Attribute Translation Grammars”. In *Proceedings of the 28th International Conference on Very Large Data Bases (VLDB’02)*, pages 838–849, 2002.
- [4] T. Bray, J. Paoli, and C. Sperberg-McQueen. “Extensible Markup Language (XML) 1.0”. Technical report, W3C, Feb. 1998.
- [5] A. Brüggemann-Klein and D. Wood. “One-Unambiguous Regular Languages”. *Information and Computation*, **142**(2):182–206, 1998.
- [6] L. Fegaras, D. Levine, S. Bose, and V. Chaluvadi. “Query Processing of Streamed XML Data”. In *Proc. 11th ACM International Conference on Information and Knowledge Management (CIKM)*, pages 126–133, 2002.
- [7] T. J. Green, G. Miklau, M. Onizuka, and D. Suciu. “Processing XML Streams with Deterministic Automata”. In *Proc. of the 9th International Conference on Database Theory (ICDT’03)*, 2003.
- [8] J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory, Languages, and Computation*. Addison-Wesley Publishing Company, Reading, MA USA, 1979.
- [9] C. Koch. “Efficient Processing of Expressive Node-Selecting Queries on XML Data in Secondary Storage: A Tree Automata-based Approach”. In *Proceedings of the 29th International Conference on Very Large Data Bases (VLDB’03)*, Berlin, Germany, 2003.
- [10] D. Lee, M. Mani, and M. Murata. “Reasoning about XML Schema Languages using Formal Language Theory”. Technical Report RJ 10197 Log 95071, IBM Research, Nov. 2000.
- [11] B. Ludäscher, P. Mukhopadhyay, and Y. Papakonstantinou. “A Transducer-Based XML Query Processor”. In *Proceedings of the 28th International Conference on Very Large Data Bases (VLDB’02)*, pages 227–238, 2002.
- [12] F. Neven. “Extensions of Attribute Grammars for Structured Document Queries”. In *Proc. 8th International Workshop on Database Programming Languages (DBPL)*, pages 99–116, 1999.
- [13] F. Neven and J. van den Bussche. “Expressiveness of Structured Document Query Languages Based on Attribute Grammars”. *Journal of the ACM*, **49**(1):56–100, Jan. 2002.
- [14] D. Olteanu, T. Kiesling, and F. Bry. “An Evaluation of Regular Path Expressions with Qualifiers against XML Streams”. In *Proceedings of the 19th IEEE International Conference on Data Engineering (ICDE)*, Bangalore, Mar. 2003. Poster Session.
- [15] World Wide Web Consortium. “XML Query”. <http://www.w3c.org/XML/query/>.