XPath: Looking Forward

Dan Olteanu[1], Holger Meuss[2], Tim Furche[1], and François Bry[1]

[1] Insitute for Computer Science, University of Munich, Germany
{olteanu, timfu, bry}@informatik.uni-muenchen.de
[2] Center for Information and Language Processing, University of Munich, Germany
meuss@cis.uni-muenchen.de

Abstract. The location path language XPath is of particular importance for XML applications since it is a core component of many XML processing standards such as XSLT or XQuery. In this paper, based on axis symmetry of XPath, equivalences of XPath 1.0 location paths involving reverse axes, such as `anc` and `prec`, are established. These equivalences are used as rewriting rules in an algorithm for transforming location paths with reverse axes into equivalent reverse-axis-free ones. Location paths without reverse axes, as generated by the presented rewriting algorithm, enable efficient SAX-like streamed data processing of XPath.

1 Introduction

Query languages for XML and semistructured data rely on location paths for selecting nodes in data items. In particular, XQuery [1] and XSLT [2] are based on XPath [3]. XPath takes a navigational approach for specifying the nodes to be selected, described by a large number of navigational axes, e.g. `child`, `descendant`, `preceding`. The number as well as the relevance of these navigational axes for querying XML has been challenged in [4,1,5].

The random access to XML data that is enabled by the various navigational axes of XPath has proven particularly difficult for an efficient stream-based processing of XPath queries. Processing of XML has seen the widespread use of the W3C document object model (DOM) [6], where an in-memory representation of the entire XML data is used. As DOM has been developed with focus on document processing in user agents (e.g. browsers), this approach has several shortcomings for other application areas:

First, a considerable amount of XML applications, in particular data-centric applications, handle documents too large to be processed in memory. Such documents are often encountered in natural language processing [7], in biological [8] and astronomical [9] projects.

Second, the need for progressive processing (also referred to as sequential processing) of XML has emerged: Stream-based processing generating partial results as soon as they are available gives rise to a more efficient evaluation in certain contexts, e.g.:

- For selective dissemination of information (SDI), documents have to be filtered according to complex requirements specified as XPath queries before

A.B. Chaudhri et al. (Eds.): EDBT 2002 Workshops, LNCS 2490, pp. 109–127, 2002.

being distributed to the subscribers [10,11]. The routing of data to selected receivers is also becoming increasingly important in the context of web service architectures.
- To integrate data over the Internet, in particular from slow sources, it is desirable to progressively process the input before the full data is retrieved [12,13].
- As a general processing scheme for XML, several solutions for pipelined processing have been suggested, where the input is sent through a chain of processors each of which taking the output of the preceding processor as input, e.g. Apache Cocoon [14].
- For progressive rendering of large documents, e.g. by means of XSL(T), cf. Requirement 19 of [5]. There have been attempts to solve this problem [15].

There is a great interest in the identification of a subset of XPath that allows efficient streamed and progressive processing, cf. [4] and Requirement 19 of [5].

For stream-based processing of XML data, the Simple API for XML (SAX) [16] has been specified. Of particular concern for progressive SAX-like processing are the reverse axes of XPath, i.e. those navigational axes (e.g. `par`, `prec`) that select nodes occuring before the context node in document order. A restriction to forward axes (i.e. axes selecting only nodes after the context node) in location paths is a straightforward consideration for an efficient stream-based evaluation of XPath queries [4].

There are three principal options how to evaluate reverse axes in a stream-based context:

- Storing in memory sufficient information that allows to access past events when evaluating a reverse axis. This amounts to keeping in memory a (possibly pruned) DOM representation of the data [15].
- Evaluating an XPath expression in more than one run. With this approach, it is also necessary to store additional information to be used in successive runs. This information can be considerably smaller than what is needed in the first approach.
- Replacing XPath expressions by equivalent ones without reverse axes.

In this paper it is shown that the third approach is possible. It is less time consuming than the second approach and does not require the in-memory storage of fragments of the input as the first approach does. Hence, XPath can be evaluated without restriction on the use of reverse axes.

Section 2 specifies the location path language considered in the rest of the paper. Then, the notion of equivalence between location paths is defined in Section 3 using a formal model and a denotational semantics for XPath based on [17,18]. Furthermore, two sets of equivalences with rather different properties are established. These equivalences are used as rewriting rules in an algorithm, called "rare", for transforming absolute XPath location paths with reverse axes into equivalent reverse-axis-free ones, as presented in Section 4. Two rewritings, based on the two rule sets, are considered. In Section 5, related work is discussed and Section 6 concludes the paper.

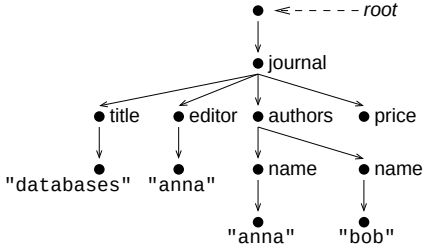


```
<journal>
  <title>databases</title>
  <editor>anna</editor>
  <authors>
    <name>anna</name>
    <name>bob</name>
  </authors>
  <price />
</journal>
```

Fig. 1. Tree and XML data it represents

Due to space limitations, parts of this work have been omitted, most notably the equivalence proofs. They can be found in the full version [19] of this paper.

2 Preliminaries

In this paper, specificities of XML that are irrelevant to the issue of concern are left out. Thus, namespaces, comments, processing instructions, attributes, schema types, and references are not considered. The results given in this paper extend straightforwardly to unrestricted XML documents.

The root node of a document corresponds to the document node of DOM and of the XQuery 1.0 and XPath 2.0 Data Model [20] – i.e. it is none of the document elements. A leaf is an empty element or a text node – cf. Figure 1.

The mathematical model used in this paper is adapted from [17,18]. The full formal model as well as the denotational semantics can be found in the full version [19] of this paper. It consists of mathematical functions that can be seen as (formal specifications of) elementary procedures.

2.1 Location Path Language

The location path language considered in the following is XPath without constructs, such as attribute selection and functions, irrelevant to the issue of concern. Note that while functions are not considered in the following sections, the results almost immediately apply to XPath 1.0 [3] location paths with functions. The only class of functions that needs special treatment are functions for accessing the context position or size of a node.

For convenience, this language will be referred to as xPath, as defined in Figure 2. Note that the axes specified in xPath are shorthands of XPath axes.

$\perp$ is used as a canonical equivalent path to the xPath expressions that select no nodes whatever the context node and document are, e.g. `/par::*`.

p_1 `==` p_2 expresses node equality based on identity. Thus, if p_1 and p_2 are two paths, then p_1 `==` p_2 holds if there is a node selected by p_1 which is *identical* to a node selected by p_2. `==` corresponds to built-in node equality operator (`==`) in XPath 2.0 and XQuery 1.0, but it can also be used for comparing

$$
\begin{aligned}
\mathit{path} &::= \mathit{path}\ \texttt{|}\ \mathit{path} \mid \texttt{/}\ \mathit{path} \mid \mathit{path}\ \texttt{/}\ \mathit{path} \mid \mathit{path}\ \texttt{[}\ \mathit{qualif}\ \texttt{]} \mid \mathit{axis}\ \texttt{::}\ \mathit{nodetest} \mid \bot\ .\\
\mathit{qualif} &::= \mathit{qualif}\ \texttt{and}\ \mathit{qualif} \mid \mathit{qualif}\ \texttt{or}\ \mathit{qualif} \mid \texttt{(}\ \mathit{qualif}\ \texttt{)} \mid\\
&\quad\ \mathit{path}\ \texttt{=}\ \mathit{path} \mid \mathit{path}\ \texttt{==}\ \mathit{path} \mid \mathit{path}\ .\\
\mathit{axis} &::= \mathit{reverse\_axis} \mid \mathit{forward\_axis}\ .\\
\mathit{reverse\_axis} &::= \texttt{par} \mid \texttt{anc} \mid \texttt{anc-or-self} \mid \texttt{prec} \mid \texttt{prec-sibl}\ .\\
\mathit{forward\_axis} &::= \texttt{self} \mid \texttt{child} \mid \texttt{desc} \mid \texttt{desc-or-self} \mid \texttt{foll} \mid \texttt{foll-sibl}\ .\\
\mathit{nodetest} &::= \mathit{tagname} \mid \texttt{*} \mid \texttt{text()} \mid \texttt{node()}\ .
\end{aligned}
$$

Fig. 2. Grammar for xPath

node sets similar to general comparisons in XPath 2.0. As XPath 1.0 has built-in support only for equality based on node values, the XPath 1.0 expression $\texttt{count(}p_1\ \texttt{|}\ p_2\texttt{)}\ \texttt{<}\ \texttt{count(}p_1\texttt{)}\ \texttt{+}\ \texttt{count(}p_2\texttt{)}$ can be used for expressing `==`.

A *path* expression will be called a "location path", or "path" for short. A *qualif* expression is a "qualifier" (or condition). Expressions $\mathit{axis}\texttt{::}\mathit{nodetest}$ and $\mathit{axis}\texttt{::}\mathit{nodetest}\texttt{[}\mathit{qualif}\texttt{]}$ are "location steps", or "steps". The length of a location path is the number of location steps it contains outside and inside qualifiers. Note that every location path is a qualifier, but the converse is false.

Absolute location paths are recursively defined as follows: A disjunctive path, i.e. a path of the form $p_1\ \texttt{|}\ \ldots\ \texttt{|}\ p_i\ \texttt{|}\ \ldots\ \texttt{|}\ p_k$, is an absolute path if for all $i = 1, \ldots, k$, p_i is an absolute path. A non-disjunctive path is an absolute path if it is of the form $\texttt{/}p$, where p is a path. A location path, which is not an absolute path, is a "relative path". A step is a "forward step", if its axis is a forward axis, or a "reverse step", if its axis is a reverse axis.

The axes of the following pairs are "symmetrical" of each other: `par` – `child`, `anc` – `desc`, `desc-or-self` – `anc-or-self`, `prec` – `foll`, `prec-sibl` – `foll-sibl`, and (useful in proofs) `self` – `self`.

[17] and [18] give a denotational semantics for XPath, which is slightly modified for our purpose in [19]. The semantics defines a function $\mathcal{S}$ that assigns a set of nodes to a location path and a context node: $\mathcal{S}[\![p]\!]x$ is the set of nodes selected by p from node x.

3 Location Path Equivalences

A set of simple equivalences is first established. These are then used to prove equivalences of paths with reverse axes. We distinguish between general equivalences that can be applied to remove any reverse axis, and specific equivalences, each of them being applicable to a certain case. Making use of the semantics of xPath given in the full version [19] of this paper, the equivalence of location paths can be formally defined as follows.

Definition 1 (Path equivalence). *Two location paths p_1 and p_2 are* equivalent, *noted $p_1 \equiv p_2$, if $\mathcal{S}[\![p_1]\!] = \mathcal{S}[\![p_2]\!]$, i.e. if $\mathcal{S}[\![p_1]\!]x = \mathcal{S}[\![p_2]\!]x$ for all nodes x (from any document).*

Intuitively, two location paths are equivalent if they select the same set of nodes for every document and every context node in this document.

Lemma 1. *Let p, p_1, and p_2 be location paths, q, q_1, and q_2 qualifiers, n a node test, and $\theta \in \{\texttt{==}, \texttt{=}\}$.*

1. Right step adjunction: *If $p_1 \equiv p_2$ and p relative, then $p_1/p \equiv p_2/p$.*
2. Left step adjunction: *If $p_1 \equiv p_2$ and p_1, p_2 relative, then $p/p_1 \equiv p/p_2$.*
3. Qualifier adjunction: *If $p_1 \equiv p_2$, then $p_1\texttt{[}q\texttt{]} \equiv p_2\texttt{[}q\texttt{]}$ and $p\texttt{[}p_1\texttt{]} \equiv p\texttt{[}p_2\texttt{]}$.*
4. Relative/absolute path conversion: *If $p_1 \equiv p_2$, then $/p_1 \equiv /p_2$.*
5. Qualifier flattening: $p\texttt{[}p_1/p_2\texttt{]} \equiv p\texttt{[}p_1\texttt{[}p_2\texttt{]]}$.
6. `anc-or-self` decomposition: $\texttt{anc-or-self::}n \equiv \texttt{anc::}n \mid \texttt{self::}n$.
7. `desc-or-self` decomposition: $\texttt{desc-or-self::}n \equiv \texttt{desc::}n \mid \texttt{self::}n$.
8. Qualifiers with joins: $p\texttt{[}p_1\ \theta\ /p_2\texttt{]} \equiv p\texttt{[}p_1\texttt{[self::node()}\ \theta\ /p_2\texttt{]]}$.

Recall that $\bot$ is a location path never selecting any node whatever the context node and document are. Since the root node has no parents and no siblings, the following holds:

Lemma 2. *Let m and n be node tests, i.e. m and n are tag names or one of the xPath constructs* `*`, `node()`, *or* `text()`.

– *Let a be one of the axes* `par`, `anc`, `prec`, `prec-sibl`, `self`, `foll`, *or* `foll-sibl`. *The following holds:*

$$/a\texttt{::}n \equiv \begin{cases} / & \text{if } a = \texttt{self},\ n = \texttt{node()} \\ \bot & \text{otherwise} \end{cases} \tag{1}$$

– *Let a be the* `prec` *or* `anc` *axis. The following holds:*

$$\texttt{/child::}m/a\texttt{::}n \equiv \begin{cases} \texttt{/self::node()[child::}m\texttt{]} & \text{if } a = \texttt{anc},\ n = \texttt{node()} \\ \bot & \text{otherwise} \end{cases} \tag{2}$$

$$\texttt{/child::}m\texttt{[}a\texttt{::}n\texttt{]} \equiv \begin{cases} \texttt{/child::}m & \text{if } a = \texttt{anc},\ n = \texttt{node()} \\ \bot & \text{otherwise} \end{cases} \tag{3}$$

3.1 General Equivalences

The nodes selected by a reverse step within a location path are necessarily descendants of the document root. The following equivalences show how for any reverse axis only those descendants of the root can be selected that are also matched by the original reverse step.

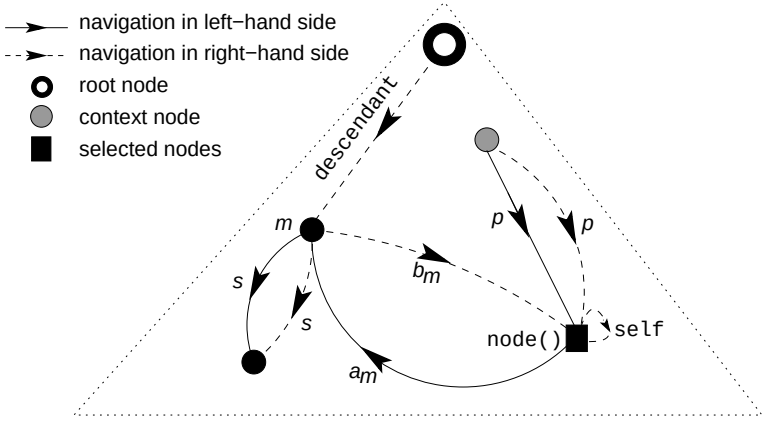

Fig. 3. Tree navigation used in Equivalence (4)

Proposition 1. *Let p and s be relative location paths, n and m node tests, a_m a reverse axis, a_n a forward axis, and b_m the symmetrical axis of a_m. Then the following holds*

$$p[a_m::m/s] \equiv p[\texttt{/desc::}m[s]/b_m\texttt{::node() == self::node()}] \tag{4}$$

$$\texttt{/}p/a_n\texttt{::}n/a_m\texttt{::}m \equiv \texttt{/desc::}m[b_m\texttt{::}n\texttt{ == /}p/a_n\texttt{::}n] \tag{5}$$

$$\texttt{/}a_n\texttt{::}n/a_m\texttt{::}m \equiv \texttt{/desc::}m[b_m\texttt{::}n\texttt{ == /}a_n\texttt{::}n] \tag{5a}$$

Equivalence (4) shows that it is possible to remove the first step in a location path within a qualifier. With help of Lemma 1.5 this result is generalized to reverse steps having an arbitrary position within a qualifier.

Figure 3 illustrates the key idea of Equivalence (4), where solid arrows are used for the navigation in the left-hand side of the equivalence, and dashed for the right-hand side: Instead of looking back from the context node specified by path p for matching a certain node ($a_m\texttt{::}m$), one can look forward from the beginning of the document for matching the node (`/desc::`m) and then, still forward, for reaching the initial context node ($b_m\texttt{::node()}$). Hence, e.g. instead of checking whether the context node specified by path p has a preceding m (p`[prec::`m`]`), one rather looks for an m node and then for a following node that is identical to the context node:

$$p\texttt{[/desc::}m\texttt{/foll::node() == self::node()]}\ .$$

Equivalence (5) removes the first reverse step from an absolute location path using the same underlying idea.

Note that the equality occurring in these equivalences is based on node identity. The equivalent paths might remain expensive to evaluate, but no evaluation of the $a_m\texttt{::}m$ reverse step is needed anymore.

Example 1. Consider the example of Figure 1 and a query asking for all `name`s that appear before a `price`. A way to select these nodes is using the following location path:

```
/desc::price/prec::name .
```

By Equivalence (5a), the previous path can be translated into the following equivalent location path:

```
/desc::name[foll::price == /desc::price] .
```

While the initial location path selects all `name` nodes preceding a `price` node, the equivalent location path selects all `name` nodes, that have a following `price` node, if that node is also a descendant of the root. It is obvious, that there is a considerably simpler equivalent location path (dropping the join), `/desc::name[foll::price]`. The need for the join arises, as the location path selecting the context nodes, relative to which the reverse step is evaluated (in this case the `price` nodes), can be arbitrarily complex:

Consider a slightly modified case of the previous one, where only `price`s that are inside a `journal` with a `title` should be considered. A possible location path for this query with reverse axis is:

```
/desc::journal[child::title]/desc::price/prec::name .
```

Again, by Equivalence (5) this is equivalent to

```
/desc::name[foll::price == /desc::journal[child::title]/desc::price] .
```

In this case it is impossible to remove the introduced join. Note that the join in the first example can be removed by additional equivalence rules for simplifying location paths, rules that are outside the scope of this paper.

Using the equivalences above, it is possible to replace reverse steps in xPath expressions. Nonetheless, in the following section specific equivalences for reverse axes are given, that yield to location paths without joins.

3.2 Specific Equivalences

In this section the interaction of the reverse axes (`anc`, `anc-or-self`, `par`, `prec`, and `prec-sibl`) with forward axes is treated, i.e. equivalences are given, that (if read as rewriting rules from left to right), depending on the location step L_f before a reverse location step L_r, either replace the reverse location step L_r or rewrite the location path into one, where the reverse step L_r is "pushed leftwise". For every reverse step the interaction with every forward step is shown.

In general, the equivalences have the following structure

$$p/L_f/L_r \equiv p' \quad \text{or} \quad p/L_f[L_r] \equiv p',$$

where p is an absolute path, L_f a forward location step, L_r a reverse location step, and p' the equivalent location path. Sometimes the equivalences can be formulated without the leading path p.

Note that interaction with reverse axes, e.g. interaction of `par` with `prec-sibl`, is not necessary to investigate in these equivalences due to the way our algorithm works (removing reverse steps from left to right of the location path in question). Also, equivalences involving `anc-or-self` and `desc-or-self` are not necessary since these location steps can be replaced using Equivalences (1.6) and (1.7).

Some of the following equivalences do still contain reverse steps on the right-hand side, but these reverse steps are either more on the left of the location path, or the right-hand side is of a form, where other equivalences can be applied to fully remove the reverse location steps as elaborated in Section 4.

Parent. The equivalences in the following proposition are divided in two sets. The first set (Equivalences (6) to (10)) covers the case of `par` location steps outside, the second inside a qualifier. Note that there is a strong structural similarity between the equivalences of the two sets.

Proposition 2 (`par` axis). *Let m and n be node tests and p a location path. Then the following holds:*

$$\texttt{desc::}n\texttt{/par::}m \equiv \texttt{desc-or-self::}m\texttt{[child::}n\texttt{]} \tag{6}$$

$$\texttt{child::}n\texttt{/par::}m \equiv \texttt{self::}m\texttt{[child::}n\texttt{]} \tag{7}$$

$$p\texttt{/self::}n\texttt{/par::}m \equiv p\texttt{[self::}n\texttt{]/par::}m \tag{8}$$

$$p\texttt{/foll-sibl::}n\texttt{/par::}m \equiv p\texttt{[foll-sibl::}n\texttt{]/par::}m \tag{9}$$

$$p\texttt{/foll::}n\texttt{/par::}m \equiv p\texttt{/foll::}m\texttt{[child::}n\texttt{]} \mid p\texttt{/anc-or-self::*[foll-sibl::}n\texttt{]/par::}m \tag{10}$$

$$\texttt{desc::}n\texttt{ [par::}m\texttt{]} \equiv \texttt{desc-or-self::}m\texttt{/child::}n \tag{11}$$

$$\texttt{child::}n\texttt{[par::}m\texttt{]} \equiv \texttt{self::}m\texttt{/child::}n \tag{12}$$

$$p\texttt{/self::}n\texttt{[par::}m\texttt{]} \equiv p\texttt{[par::}m\texttt{]/self::}n \tag{13}$$

$$p\texttt{/foll-sibl::}n\texttt{[par::}m\texttt{]} \equiv p\texttt{[par::}m\texttt{]/foll-sibl::}n \tag{14}$$

$$p\texttt{/foll::}n\texttt{[par::}m\texttt{]} \equiv p\texttt{/foll::}m\texttt{/child::}n \mid p\texttt{/anc-or-self::*[par::}m\texttt{]/foll-sibl::}n \tag{15}$$

Example 2. Consider the data of Figure 1. The following location path selects all `editor`s of `journal`s:

```
/desc::editor[par::journal].
```

According to Equivalence (11), this path is equivalent to:

```
/desc-or-self::journal/child::editor.
```

Ancestor. The following proposition gives equivalences that either move an `anc` step to the left of a path or remove it completely. Equivalences (16a) and (21a) are special cases of Equivalences (16) and (21), respectively.

Proposition 3 (`anc` axis). *Let m and n be node tests and p a location path. Then the following holds:*

$$p/\texttt{desc::}n/\texttt{anc::}m \equiv p[\texttt{desc::}n]/\texttt{anc::}m \mid p/\texttt{desc-or-self::}m[\texttt{desc::}n] \tag{16}$$

$$/\texttt{desc::}n/\texttt{anc::}m \equiv /\texttt{desc-or-self::}m[\texttt{desc::}n] \tag{16a}$$

$$p/\texttt{child::}n/\texttt{anc::}m \equiv p[\texttt{child::}n]/\texttt{anc-or-self::}m \tag{17}$$

$$p/\texttt{self::}n/\texttt{anc::}m \equiv p[\texttt{self::}n]/\texttt{anc::}m \tag{18}$$

$$p/\texttt{foll-sibl::}n/\texttt{anc::}m \equiv p[\texttt{foll-sibl::}n]/\texttt{anc::}m \tag{19}$$

$$p/\texttt{foll::}n/\texttt{anc::}m \equiv p/\texttt{foll::}m[\texttt{desc::}n] \mid p/\texttt{anc-or-self::*}[\texttt{foll-sibl::*/desc-or-self::}n]/\texttt{anc::}m \tag{20}$$

$$p/\texttt{desc::}n[\texttt{anc::}m] \equiv p[\texttt{anc::}m]/\texttt{desc::}n \mid p/\texttt{desc-or-self::}m/\texttt{desc::}n \tag{21}$$

$$/\texttt{desc::}n[\texttt{anc::}m] \equiv /\texttt{desc-or-self::}m/\texttt{desc::}n \tag{21a}$$

$$p/\texttt{child::}n[\texttt{anc::}m] \equiv p[\texttt{anc-or-self::}m]/\texttt{child::}n \tag{22}$$

$$p/\texttt{self::}n[\texttt{anc::}m] \equiv p[\texttt{anc::}m]/\texttt{self::}n \tag{23}$$

$$p/\texttt{foll-sibl::}n[\texttt{anc::}m] \equiv p[\texttt{anc::}m]/\texttt{foll-sibl::}n \tag{24}$$

$$p/\texttt{foll::}n[\texttt{anc::}m] \equiv p/\texttt{foll::}m/\texttt{desc::}n \mid p/\texttt{anc-or-self::*}[\texttt{anc::}m]/\texttt{foll-sibl::*/desc-or-self::}n \tag{25}$$

Preceding-sibling. In the following proposition the `prec-sibl` axis is treated. Note that the right-hand side of equivalences for `prec-sibl` (and `prec`) contains more union terms than the other equivalences, since there is no `-or-self` variant of these axes.

Proposition 4 (`prec-sibl` axis). *Let m and n be node tests and p a location path. Then the following holds:*

$$\texttt{desc::}n\texttt{/prec-sibl::}m \equiv \texttt{desc::}m\texttt{[foll-sibl::}n\texttt{]} \tag{26}$$

$$\texttt{child::}n\texttt{/prec-sibl::}m \equiv \texttt{child::}m\texttt{[foll-sibl::}n\texttt{]} \tag{27}$$

$$p\texttt{/self::}n\texttt{/prec-sibl::}m \equiv p\texttt{[self::}n\texttt{]/prec-sibl::}m \tag{28}$$

$$p\texttt{/foll-sibl::}n\texttt{/prec-sibl::}m \equiv p\texttt{[self::}m\texttt{/foll-sibl::}n\texttt{]} \tag{29}$$
$$\mid p\texttt{[foll-sibl::}n\texttt{]/prec-sibl::}m$$
$$\mid p\texttt{/foll-sibl::}m\texttt{[foll-sibl::}n\texttt{]}$$

$$p\texttt{/foll::}n\texttt{/prec-sibl::}m \equiv p\texttt{/foll::}m\texttt{[foll-sibl::}n\texttt{]} \tag{30}$$
$$\mid p\texttt{/anc-or-self::*[foll-sibl::}n\texttt{]}$$
$$\texttt{/prec-sibl::}m$$
$$\mid p\texttt{/anc-or-self::}m\texttt{[foll-sibl::}n\texttt{]}$$

$$\texttt{desc::}n\texttt{[prec-sibl::}m\texttt{]} \equiv \texttt{desc::}m\texttt{/foll-sibl::}n \tag{31}$$

$$\texttt{child::}n\texttt{[prec-sibl::}m\texttt{]} \equiv \texttt{child::}m\texttt{/foll-sibl::}n \tag{32}$$

$$p\texttt{/self::}n\texttt{[prec-sibl::}m\texttt{]} \equiv p\texttt{[self::}n\texttt{]/foll-sibl::}m \tag{33}$$

$$p\texttt{/foll-sibl::}n\texttt{[prec-sibl::}m\texttt{]} \equiv p\texttt{[self::}m\texttt{]/foll-sibl::}n \tag{34}$$
$$\mid p\texttt{/foll-sibl::}m\texttt{/foll-sibl::}n$$
$$\mid p\texttt{[prec-sibl::}m\texttt{]/foll-sibl::}n$$

$$p\texttt{/foll::}n\texttt{[prec-sibl::}m\texttt{]} \equiv p\texttt{/foll::}m\texttt{/foll-sibl::}n \tag{35}$$
$$\mid p\texttt{/anc-or-self::*[prec-sibl::}m\texttt{]}$$
$$\texttt{/foll-sibl::}n$$
$$\mid p\texttt{/anc-or-self::/foll-sibl::}n$$

Preceding. The following proposition describes the interaction of `prec` with other axes.

Proposition 5 (`prec` axis). *Let m and n be node tests and p a location path. Then the following holds:*

$$p\texttt{/desc::}n\texttt{/prec::}m \equiv p\texttt{[desc::}n\texttt{]/prec::}m \tag{36}$$
$$\mid p\texttt{/child::*[foll-sibl::*/desc-or-self::}n\texttt{]}$$
$$\texttt{/desc-or-self::}m$$

$$\texttt{/desc::}n\texttt{/prec::}m \equiv \texttt{/desc::}m\texttt{[foll::}n\texttt{]} \tag{36a}$$

$$p\texttt{/child::}n\texttt{/prec::}m \equiv p\texttt{[child::}n\texttt{]/prec::}m \tag{37}$$
$$\mid p\texttt{/child::*[foll-sibl::}n\texttt{]/desc-or-self::}m$$

$$p\texttt{/self::}n\texttt{/prec::}m \equiv p\texttt{[self::}n\texttt{]/prec::}m \tag{38}$$

$$p\texttt{/foll-sibl::}n\texttt{/prec::}m \equiv p\texttt{[foll-sibl::}n\texttt{]/prec::}m \tag{39}$$
$$\mid p\texttt{/foll-sibl::*[foll-sibl::}n\texttt{]/desc-or-self::}m$$
$$\mid p\texttt{[foll-sibl::}n\texttt{]/desc-or-self::}m$$

$$p\texttt{/foll::}n\texttt{/prec::}m \equiv p\texttt{[foll::}n\texttt{]/prec::}m \tag{40}$$
$$\mid p\texttt{/foll::}m\texttt{[foll::}n\texttt{]}$$
$$\mid p\texttt{[foll::}n\texttt{]/desc-or-self::}m$$

$$p/\texttt{desc::}n\texttt{[prec::}m\texttt{]} \equiv p\texttt{[prec::}m\texttt{]/desc::}n \quad (41)$$
$$\mid p\texttt{/child::*[desc-or-self::}m\texttt{]}$$
$$\texttt{/foll-sibl::*/desc-or-self::}n$$
$$\texttt{/desc::}n\texttt{[prec::}m\texttt{]} \equiv \texttt{/desc::}m\texttt{/foll::}n \quad (41a)$$
$$p\texttt{/child::}n\texttt{[prec::}m\texttt{]} \equiv p\texttt{[prec::}m\texttt{]/child::}n \quad (42)$$
$$\mid p\texttt{/child::*[desc-or-self::}m\texttt{]/foll-sibl::}n$$
$$p\texttt{/self::}n\texttt{[prec::}m\texttt{]} \equiv p\texttt{[prec::}m\texttt{]/self::}n \quad (43)$$
$$p\texttt{/foll-sibl::}n\texttt{[prec::}m\texttt{]} \equiv p\texttt{[prec::}m\texttt{]/foll-sibl::}n \quad (44)$$
$$\mid p\texttt{/foll-sibl::*[desc-or-self::}m\texttt{]/foll-sibl::}n$$
$$\mid p\texttt{[desc-or-self::}m\texttt{]/foll-sibl::}n$$
$$p\texttt{/foll::}n\texttt{[prec::}m\texttt{]} \equiv p\texttt{[prec::}m\texttt{]/foll::}n \quad (45)$$
$$\mid p\texttt{/foll::}m\texttt{/foll::}n$$
$$\mid p\texttt{[desc-or-self::}m\texttt{]/foll::}n$$

Example 3. Consider the location path

```
/desc::price/prec::name
```

of Example (1). With Rule 36a it can be rewritten to

```
/desc::name[foll::price].
```

This result is more compact and closer to the original than the result of Example (1) using Equivalence (5a).

4 Location Path Rewriting

Each Equivalence (i) $p_1 \equiv p_2$ of Section 3 gives rise to a rewriting rule: A path matching with the left-hand side p_1 can be rewritten into a path corresponding to the right-hand side p_2. In the following, Rule (i) denotes the rewriting rule $p_1 \to p_2$ induced by Equivalence (i) $p_1 \equiv p_2$. The equivalences of Lemma 1 induce rewriting rules, denoted Rules (1.1) to (1.8).

The equivalences of Section 3 are split in two sets of rules for use in a rewriting algorithm:

1. RuleSet_1, containing Rules (1) to (3) and the general Rules (4) to (5a).
2. RuleSet_2, containing Rules (1) to (3) and the specific Rules (3) to (42).

A rule can be applied to a location path in the following manner:

Definition 2 (Rule application). *Let p be a non-disjunctive location path, and let $p_l \to p_r$ be a rule either from RuleSet_1 or RuleSet_2. If p is of the form p_l/p', then let q denote the path p_r/p'. If p_l is a relative path and if p is of the form $p_1/p_l/p_2$, then let q denote the path $p_1/p_r/p_2$. In both cases q is called the* result *of the application of rule $p_l \to p_r$ to p.*

An algorithm, called "rare" (sketched in Figure 4) for computing a reverse-axis-free path equivalent to an absolute path is considered below. The input for the algorithm is restricted to paths without qualifiers containing "RR joins":

Definition 3 (RR join). *An RR join is an XPath expression of the form* $p_1\ \theta\ p_2$ *where* $\theta \in \{==, =\}$*, and both* p_1 *and* p_2 *are* ***R****elative paths such that at least one of them contains a* ***R****everse step.*

For the consideration of termination and correctness of the algorithm, some important properties of the application of the rewriting rules to a location path are required:

Lemma 3 (Properties of rule application). *Let* p *be an absolute location path with no qualifier containing RR joins.*

1. *If* p *contains a reverse step, then a rule from* $RuleSet_1$ *and a rule from* $RuleSet_2$ *is applicable to* p*. Possibly, Rules (1.1) to (1.8) have to be applied first.*
2. *The result of a rule application to the first reverse step in* p *is an absolute path with no qualifiers containing RR joins.*
3. *If* q *is the result of a rule application to* p*, then* $p \equiv q$*.*

Proof. **(1):** Let L be the first reverse location step.

First consider RuleSet_1: If L occurs outside a qualifier, Rules (5), (5a) or (1) to (3) can be applied, since p is an absolute location path. If L occurs as the first location step inside a qualifier, Rule (4) can be applied. If L appears at any other position inside a qualifier, Rule (1.5) can be applied in order to construct a qualifier with L as first location step. Rule (4) can be applied now.

RuleSet_2 provides rules for interaction between each reverse step and an arbitrary forward step s, there is always a rule, that can be applied to the first reverse step in p.

(2) Only Rules (4), (5), and (5a) introduce a binary relation (namely ==), if they are applied to a location path. But always one of the operands is an absolute path. Hence, in any case the result of the rule application contains no RR join. Furthermore, since p is an absolute path, the result of applying a rule to p is also an absolute path.

(3) This holds due to Lemma 1.1−4.

"rare" Algorithm. The "rare" Algorithm, outlined in Figure 4, can be used for RuleSet_1 as well as for RuleSet_2. The algorithm takes as input a location path which is absolute, since some rules from RuleSet_1 and RuleSet_2 are only applicable to absolute location paths.

Theorem 1 (Removal of reverse location steps using RuleSet_1). *Let* p *be an absolute path without qualifiers in which RR joins occur. There exists an absolute path* p' *with no reverse steps such that* $p \equiv p'$*. Using "rare" and* $RuleSet_1$*, this path* p' *has a length and can be computed in a time linear in the length of* p*.*

Let ξ = RuleSet$_1$ or RuleSet$_2$.

Auxiliary functions:

match(p): returns the result of a rule application from ξ to the first reverse location step in p.
apply-lemmas(p): returns p if Rules (1.1-8) are not applicable to p. Otherwise, it returns the result of the repeated application of Rules (1.1-8) to p.
union-flattening(p): returns a path equivalent to p with unions at top level only.

***rare*(p)**
Input: p {absolute location path without qualifiers containing RR joins}.

$p \leftarrow$ *apply-lemmas*(p).
$p \leftarrow$ *union-flattening*(p) $= U_1 \mid \ldots \mid U_n$ ($n \geq 1$).
$S \leftarrow$ empty stack.
for $i \leftarrow 1$ to n **do**
  push(U_i, S).
end for
$p' \leftarrow \perp$. {initialization}
while *not*(*empty*(S)) **do**
  $U \leftarrow$ *pop*(S).
  while U contains a reverse step **do**
    $U \leftarrow$ *match*(U).
    $U \leftarrow$ *apply-lemmas*(U).
    $U \leftarrow$ *union-flattening*(U) $= V_1 \mid \ldots \mid V_n$ ($n \geq 1$).
    for $i \leftarrow 2$ to n **do**
      push(V_i, S).
    end for
    $U \leftarrow V_1$.
  end while
  $p' \leftarrow p' \mid U$.
end while

Output: p' {location path without reverse axes equivalent to p}.

Fig. 4. Algorithm ***rare*** (**r**everse **a**xis **re**moval)

Proof. A path equivalent to p is constructed as sketched in Figure 4. All reverse location steps are rewritten, one after another. Lemma 3 guarantees that a rule of RuleSet$_1$ can be applied to any path containing a reverse location step. The resulting path p' contains no reverse location steps and is equivalent to p.

The location path p' is of linear size and constructed in linear time, since each rule application removes one reverse step, adds at most two forward location steps and no reverse ones.

Theorem 2 (Removal of reverse location steps using RuleSet$_2$). *Let p be an absolute path with no qualifiers in which RR joins occur. There exists an absolute path p' with no reverse steps such that $p \equiv p'$. Using "rare" and*

RuleSet$_2$, this path p' has a length and can be computed in a time exponential in the length of p.

Proof. An application of a rule from RuleSet$_2$ can have three different result types:

1. removes completely a reverse step, e.g. Rules (1) to (3)) or (6);
2. pushes the reverse step from right to left in the path, e.g. Rule (8);
3. for the interaction between a `foll` and a reverse step L_r, i.e. `foll::`n/L_r or `foll::`n`[`L_r`]`, a union of several other paths is obtained, e.g. Rule (10); the resulting union terms have reverse steps at positions less than or equal in the original path and they do not contain anymore the interaction between the initial `foll` step and a reverse step.

Since the path has a finite length, the procedure of pushing reverse steps leftwise terminates. Also, the number of interactions between `foll` and reverse steps is finite. Hence, the algorithm terminates.

Each rule application having the first result type removes a reverse step and does not change the number of union terms. Hence, in the best case, i.e. using only rule applications with the first result type, the algorithm has a linear time complexity in the length of the input path p.

The last two result types are significant for the worst-case complexity of the algorithm, since each rule application can produce intermediate rewritten paths with more than one union term (up to three union terms). Hence, each rule application can increase the order of the input for the next rule application, yielding an exponential time complexity in the length of the input path p.

Example runs of the algorithm for both set of rules are presented in Figure 5.

Comparison. Both RuleSet$_1$ and RuleSet$_2$ have advantages and it is still an open issue which one is preferable. The path rewriting using RuleSet$_2$ has in the worst case an exponential time complexity and output size in the length of the input location path. As location paths are in practice small (less than ten steps), the exponential worst-case complexity of RuleSet$_2$ does not necessarily generate longer paths than RuleSet$_1$. In addition, since they do not contain joins, the location paths generated using RuleSet$_2$ are simpler (as can be seen in the examples), hence more convenient to evaluate, than those generated using RuleSet$_1$, which contain the same number of joins as there are reverse steps in the input location path.

For further comparison, practical tests of the rewritten location paths, as generated by the two rule sets, have been performed using a Java prototype of a streamed XPath processor [21] against the MONDIAL geographical database [22], a highly structured XML document of comparatively small size.

Three types of location paths with reverse steps have been considered for comparing the efficiency of the rule sets:

Q1: simple location paths without qualifiers, e.g.

```
/desc::name/anc::country .
```

Consider the example of Figure 1 and a query asking for all `titles` that appear before a `name` and are inside `journals`. This query can be expressed as the following location path:

```
/desc::name/prec::title[anc::journal]
```

Note that p is an absolute path without qualifiers containing RR-joins.

RuleSet$_1$

Step 1: $p \leftarrow$ *apply-lemmas*$(p) = p$.
Step 2: $U_1 \leftarrow$ `/desc::name/prec::title[anc::journal]`.
Step 3: *push*(U_1, S).
Step 4: $p' \leftarrow \perp$, $U \leftarrow$ *pop*(S).
Step 5: U contains a reverse step (`prec::title`).
Step 6: $U \leftarrow$ *match*$(U) =$
`/desc::title[anc::journal]`
`[foll::name == /desc::name]` {Rule (5)}
Step 7: $U \leftarrow$ *apply-lemmas*$(U) = U$.
Step 8: U contains a reverse step (`anc::journal`).
Step 9: $U \leftarrow$ *match*$(U) =$
`/desc::title[/desc::journal/desc::node() == self::node()]`
`[foll::name == /desc::name]`. {Rule (4)}
Step 10: $U \leftarrow$ *apply-lemmas*$(U) = U$.
Step 11: U does not contain reverse steps.
Step 12: $p' \leftarrow U$, S is empty.
Output: $p =$ `/desc::title[/desc::journal/desc::node() == self::node()]`
`[foll::name == /desc::name]`.

RuleSet$_2$

Step 1: $p \leftarrow$ *apply-lemmas*$(p) = p$.
Step 2: $U_1 \leftarrow$ `/desc::name/prec::title[anc::journal]`.
Step 3: *push*(U_1, S).
Step 4: $p' \leftarrow \perp$, $U \leftarrow$ *pop*(S).
Step 5: U contains a reverse step (`prec::title`).
Step 6: $U \leftarrow$ *match*$(U) =$
`/desc-or-self::title[foll::name][anc::journal]` {Rule (36a)}
Step 7: $U \leftarrow$ *apply-lemmas*$(U) =$ `/desc::title[foll::name][anc::journal]`
`| /self::title[foll::name][anc::journal]`.
Step 8: U contains a reverse step (`anc::journal`).
Step 9: $U \leftarrow$ *match*$(U) =$
`/desc::title[foll::name][anc::journal]` $\mid \perp$. {Rules (1) to (3)}
Step 10: U contains a reverse step (`anc::journal`).
Step 11: $U \leftarrow$ *match*$(U) =$
`/desc::journal/desc::title[foll::name]`. {Rule (21a)}
Step 12: $U \leftarrow$ *apply-lemmas*$(U) = U$.
Step 13: U does not contain reverse steps.
Step 14: $p' \leftarrow U$, S is empty.
Output: $p' =$ `/desc::journal/desc::title[foll::name]`.

Fig. 5. Example runs of ***rare*** algorithm

Q2: location paths with a single qualifier, e.g.

```
/desc::name/par::city[anc::country] .
```

Q3: complex location paths with several qualifiers, potentially rewritten by the second RuleSet$_2$ to location paths with exponential size, e.g.

```
/desc::province[par::country]/child::city[anc::mondial] .
```

The queries are rewritten using RuleSet_1 and RuleSet_2 respectively. Figure 6 shows the results of processing the given queries against the MONDIAL database on a Pentium III 1 GHz, 512 MB system running SuSE Linux 7.3.

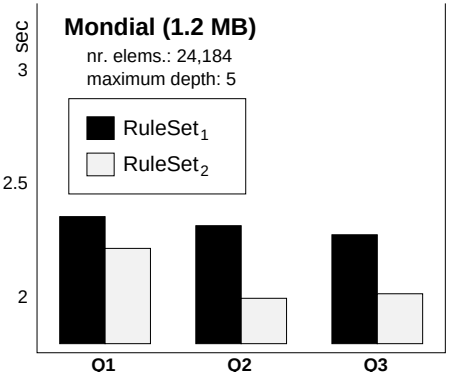

Fig. 6. Experimental evaluation of the rule sets

The results show that the two rule sets are very competitive for all three types of queries. As expected, RuleSet_2 generating simpler rewritten queries performs slightly better than RuleSet_1.

Rewriting location paths using variables. There are two classes of location paths not covered by the rules given so far: relative location paths and location paths with RR joins (cf. Definition 3), e.g. p`[self::*` = `prec::*]`. Any attempt to remove the reverse location steps in these cases results in losing the context given by p.

In the full version [19] of this paper an approach is proposed to solve this problem by remembering the context in a variable. It is based on a `for-return` construct for variable binding, as provided by XPath 2.0, XQuery, and XSLT. Using this approach *every* location path can be rewritten to an equivalent reverse-axis-free one.

5 Related Work

Several methods have been proposed for rewriting XPath expressions taking integrity constraints or schemas into account [23,24], and the equivalence and containment problems for XPath expressions have been investigated [25,26]. Furthermore, a growing interest in query optimization for XML databases, including optimization of XPath expressions, recently emerged. To the best of our knowledge, however, no other approach has been proposed for removing reverse steps from XPath expressions relying upon XPath symmetry. Note that using

equivalence preserving rewriting rules for removing reverse steps from XPath expressions, as it is proposed in the present paper, is not closely related to the general equivalence problem for XPath expressions.

In [27] redundancies in XPath expressions based on a "model-oriented" approach are investigated. Such an approach relies on an abstract model of XPath that views XPath expressions as tree patterns. [27] shows that redundant branches of a tree pattern can be eliminated in polynomial time. Tree patterns are more abstract than XPath expressions in a way which is relevant to the work described in the present paper: A same tree pattern represents multiple equivalent XPath expressions. In particular, the symmetries in XPath exploited in the present paper are absent from tree patterns. Tree patterns do not consider the document order and therefore the concept of forward and reverse steps. In some sense the present work shows in which cases this simplified view upon an XPath expression can be justified.

Stream-based query processing has gained considerable interest in the past few years, e.g. due to its application in data integration [13,12] and in publish-subscribe architectures [11,10]. They all consider a navigational approach (XML-QL or XPath) consisting of a restricted subset of forward axes from XPath. This fact contrasts with the present work, which enables the use of the unrestricted set of XPath axes in a stream-based context.

6 Conclusion

The main result of this paper consists in two rule sets, RuleSet_1 and RuleSet_2, used in an algorithm for transforming XPath 1.0 expressions containing reverse axes into reverse-axis-free equivalents. Both RuleSet_1 and RuleSet_2 have advantages and it is still an open issue which one is preferable. The location paths generated using RuleSet_2 do not contain joins, in contrast with those generated using RuleSet_1, which contain the same number of joins as there are reverse steps in the input location path. However, the path rewriting using RuleSet_2 has an exponential complexity in the length of the input location path, in contrast with rewriting using RuleSet_1 which has only a linear complexity. Preliminary experimental results show that the rewritten queries generated by RuleSet_2 are evaluated slightly faster than those generated by RuleSet_1.

Closely related to the comparison of RuleSet_1 and RuleSet_2 we plan to investigate the notion of "minimality" or "simplicity" of XPath expressions. We are focusing on defining a notion of a minimal XPath expression that can be evaluated more efficiently in a stream-based context than its equivalents. A notion of minimality will allow for well-founded optimization techniques for XPath expressions.

The equivalences proposed in this paper are drawn from and represent prerequisites for an efficient streamed evaluation of unrestricted XPath, as considered in [21].

References

[1] W3C, "XQuery 1.0: An XML query language," W3C Working Draft, 2002.

[2] W3C, "XSL Transformations (XSLT) Version 1.0," W3C Recommendation, 1999.

[3] W3C, "XML Path Language (XPath) Version 1.0," W3C Recommendation, 1999.

[4] A. Desai, "Introduction to Sequential XPath," in *Proc. of IDEAlliance XML Conference*, 2001.

[5] W3C, "XSL Transformations (XSLT) Version 2.0," W3C Working Draft, 2002.

[6] W3C, "Document Object Model (DOM) Level 2 Core Specification," W3C Recommendation, 2000.

[7] N. Ide, P. Bonhomme, and L. Romary, "XCES: An XML-based standard for linguistic corpora," in *Proc. of the Second Annual Conference on Language Resources and Evaluation*, 2000.

[8] F. Bry and P. Kr ger, "A Computational Biology Database Digest: Data, Data Analysis, and Data Management," Tech. Rep. PMS-FB-2002-8, University of Munich, 2002.

[9] "Astronomical Data Center," homepage `http://adc.gsfc.nasa.gov`.

[10] C. Chan, P. Felber, M. Garofalakis, and R. Rastogi, "Efficient Filtering of XML Documents with XPath Expressions," in *Proc. of International Conference on Data Engineering (ICDE)*, 2002.

[11] M. Altinel and M. Franklin, "Efficient Filtering of XML Documents for Selective Dissemination of Information," in *Proc. of 26th Conference on Very Large Databases (VLDB)*, 2000.

[12] A. Levy, Z. Ives, and D. Weld, "Efficient Evaluation of Regular Path Expressions on Streaming XML Data," Tech. Rep., University of Washington, 2000.

[13] T. J. Green, M. Onizuka, and D. Suciu, "Processing XML Streams with Deterministic Automata and Stream Indexes," Tech. Rep., University of Washington, 2001.

[14] Apache Project, "Cocoon 2.0: XML publishing framework," available at `http://xml.apache.org/cocoon/index.html`.

[15] Apache Project, "Xalan-Java Version 2.2," available at `http://xml.apache.org/xalan-j/index.html`.

[16] D. Megginson, "SAX: The Simple API for XML," 1998.

[17] P. Wadler, "A formal semantics of patterns in XSLT," in *Proc. of Conference on Markup Technologies*, 1999.

[18] P. Wadler, "Two semantics of XPath," Tech. Rep., 2000.

[19] D. Olteanu, H. Meuss, T. Furche, and F. Bry, "XPath: Looking Forward," Tech. Rep. PMS-FB-2001-17, University of Munich, 2001.

[20] W3C, "XQuery 1.0 and XPath 2.0 data model," W3C Working Draft, 2001.

[21] "XPath Evaluation Project," University of Munich, homepage `http://www.pms.informatik.uni-muenchen.de/forschung/xpath-eval.html`.

[22] W. May, "Information Extraction and Integration with FLORID: The MONDIAL Case Study," Tech. Rep. 131, University of Freiburg, Institut for Computer Science, 1999, available at `www.informatik.uni-freiburg.de/~may/Mondial/`.

[23] K. Boehm, K. Gayer, T. Oezsu, and K. Aberer, "Query Optimization for Structured Documents Based on Knowledge on the Document Type Definition," in *Proc. of the Advances in Digital Libraries Conference*, 1998.

[24] P. T. Wood, "Optimising Web Queries Using Document Type Definitions," in *2nd ACM Workshop on Web Information and Data Management (WIDM'99)*, 1999.

[25] A. Deutsch and V. Tannen, "Containment for Classes of XPath Expressions Under Integrity Constraints," in *Knowledge Representation meets Databases (KRDB)*, 2001.

[26] P. T. Wood, "On the Equivalence of XML Patterns," in *Proc. 6th Int. Conf. on Rules and Objects in Databases (DOOD)*, 2000.

[27] Sihem Amer-Yahia, SungRan Cho, Laks V. S. Lakshmanan, and Divesh Srivastava, "Minimization of Tree Pattern Queries," in *SIGMOD*, 2001.