HAL
open science

Logic-Based XPath Optimization

Pierre Genevès, Jean-Yves Vion-Dury

To cite this version:

Pierre Genevès, Jean-Yves Vion-Dury. Logic-Based XPath Optimization. First International Workshop on High Performance XML Processing, May 2004, New York, United States. inria-00423381

HAL Id: inria-00423381

https://inria.hal.science/inria-00423381v1

Submitted on 9 Oct 2009

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Logic-Based XPath Optimization

Pierre Genevès[1] and Jean-Yves Vion-Dury[1,2]

Abstract

XPath [16] was introduced by the W3C as a standard language for specifying node selection, matching conditions, and for computing values from an XML document. XPath is now used in many XML standards such as XSLT [15] or the forthcoming XQuery [14] database access language. Since efficient XML content querying is crucial for the performance of almost all XML processing architectures, a growing need for studying high performance XPath-based querying has emerged. Our approach aims at optimizing XPath performance through static analysis and syntactic transformation of XPath expressions.

1 Introduction

In this paper, we first describe a formal architecture for static analysis that could be implemented independently from any particular XPath engine. Second, as an application, we show how the containment relation over XPath expressions can be used for general optimization purposes. The important advantage of our approach is that logic-based optimizations can be applied at syntactic level and thus remain compatible with any XPath engine. Hence, one does not have to modify an XPath engine to optimize XPath queries.

Related Work Abundant literature on query rewriting for optimization can be found in database systems. As XML is becoming the de facto standard for representing structured content, the mapping between heritage of database theory and structured documents querying is being worked out. Optimization results on tree patterns [2, 6] or regular paths [1] often rely on a notion of query *equivalence* or query *containment* with respect to the considered model. Our approach also relies on a similar relation but aims at studying and taking XPath semantics peculiarities into account, in order to be extensible to a large XPath fragment. The containment over XPath expressions is defined using XPath semantics, usually described by a formal semantics function $\mathcal{S}$ (as found in [11]). Thus, the containment relation between two XPath expressions p_1 and p_2 holds when, for any XML tree t and any context node x of t, the set-theoretic inclusion relation holds between the sets of nodes respectively returned by the evaluation of p_1 and p_2:

$$p_1 \leq p_2 \quad \text{iff} \quad \forall t, \forall x \in t, \mathcal{S}[\![p_1]\!]_x \subseteq \mathcal{S}[\![p_2]\!]_x$$

Containment for XPath expressions is being actively studied [7, 4, 8, 12, 10], but none of these approaches explains how the containment can be practically used for XPath optimization. A rewriting approach [9] proposed to rewrite XPath backward axes into forward ones, in order to optimize XPath queries for stream-based processing. However, this rewriting technique mostly relies on XPath symmetry and does not involve containment-based optimization. We believe the present work is a first step toward the use of the containment for XPath optimization.

Outline In section 2 we present the abstract syntax and semantics of XPath expressions we consider. The next sections detail our syntactic transformations: our global architecture is described in section 3 and we explain how the containment can be used for XPath optimization in section 4. We then describe and analyze preliminary experiments in section 5 before we conclude in section 6.

2 XPath Syntax and Semantics

In this paper, we consider a subset of the XPath specification [16], mainly composed of forward axes and qualifiers. Our intent is to extend this subset to more features, such as “count()” and “position()” in qualifiers. Our XPath

[1]INRIA Rhône-Alpes, Zirst, 655 avenue de l’Europe, Montbonnot, 38334 Saint Ismier Cedex - France
[2]Xerox Research Centre Europe, 6 chemin de Maupertuis, 38240 Meylan - France

fragment includes several variants: the void path $\perp$ and the explicit root node $\wedge$ (respectively proposed and defined in [9] and [10]) to ease formal analysis and to make the XPath syntax fully compositional. In addition, we chose to include two XPath 2.0 [13] extensions: qualified paths (e.g. $(p)[q]$) instead of qualified steps (e.g. $a{::}N[q]$) and path intersection ($p_1 \cap p_2$). Our fragment also includes an important extension with respect to qualifiers: the node set inclusion constraint $p_1 \sqsubseteq p_2$, defined in [10], which brings extra - yet tractable - expressive power. Note that the usual form $p_1[p_2]$ is a syntactic sugar for $p_1[\textbf{not}\ \ (p_2 \sqsubseteq \perp)]$. In the XPath abstract syntax shown below, a node test n denotes any element name:

Path	p	::=	$\wedge \mid \perp \mid p_1 \mid p_2 \mid p_1 \cap p_2 \mid p_1/p_2 \mid (p) \mid p[q] \mid a{::}N$
Qualifier	q	::=	**true** | **false** | (q) | **not** q | q_1 **or** q_2 | q_1 **and** q_2 | $p_1 \sqsubseteq p_2$
Axis	a	::=	child | descendant | self | descendant-or-self | attribute | namespace
NodeTest	N	::=	$n \mid *$ | **node**() | **text**() | **element**() | **processing-instruction**() | **comment**()

The denotational semantics of a path p provided a context node x, noted $\mathcal{S}[\![p]\!]_x$, is given in appendix A. An originality comes from the addition of the inclusion constraint between two paths inside qualifiers: $\mathcal{Q}[\![p_1 \sqsubseteq p_2]\!]_x = \mathcal{S}[\![p_1]\!]_x \subseteq \mathcal{S}[\![p_2]\!]_x$. Note that this definition allows us to define terms such as $p[p_1 \sqsubseteq p_2$ **and** $p_2 \sqsubseteq p_1]$, which can be used to give a clear semantics to constructs of the form $p[p_1 == p_2]$ usually found in the literature.

3 Formal Architecture for Static Analysis

Formal methods for static analysis of XPath expressions have to face the combinatorial complexity of any significant fragment of XPath syntax. A way to ease analysis is to simplify the form of XPath expressions. To this end, our work currently focuses on rewriting an XPath expression into an equivalent but simplified form, using semantic-preserving rules. This normalization aims at easing whatever subsequent analysis operations. As a particular application, we show in section 4 how to use the containment relation to perform generic optimizations. Figure 2 gives an applicative view of our architecture.

Normalization Phase More formally, a path p is rewritten into its normal form $\overline{p}$ using two sets of rules $\mathcal{N}$ and $\mathcal{C}$. This transformation is noted $p \xrightarrow{\mathcal{N}^*}\xrightarrow{\mathcal{C}^*} \overline{p}$ and first uses the set of rules $\mathcal{N}$ (see [10]) mainly composed of distributivity rules such as $(p_1 \mid p_2)/p \rightarrow p_1/p \mid p_2/p$ and other structural rules needed to reach the disjunctive normal form structure. The second step of the normalization phase involves $\mathcal{C}$ (shown on figure 1) that performs void path elimination: rule ra24 detects a possible contradiction; void paths are propagated to the top level of the path structure (using r4f, r4e, r4a, and r4b) and then eliminated using r4c and r4cp. At the end of the normalization

$$
\begin{array}{lcll}
p[\textbf{false}] & \rightarrow & \perp & \text{(r4f)} \\
\perp[q] & \rightarrow & \perp & \text{(r4e)} \\
\perp/p & \rightarrow & \perp & \text{(r4a)} \\
p/\perp & \rightarrow & \perp & \text{(r4b)}
\end{array}
\qquad
\begin{array}{lcll}
 & \perp \mid p & \rightarrow\ p & \text{(r4c)} \\
 & p \mid \perp & \rightarrow\ p & \text{(r4cp)} \\
n_1 \neq n_2, & a{::}n_1[q_1]/\text{self}{::}n_2[q_2] & \rightarrow\ \perp & \text{(ra24)}
\end{array}
$$

Figure 1: Detection, Propagation, and Elimination of Void Paths ($\mathcal{C}$).

phase, a normal path $\overline{p}$ is either the void path $\perp$ or a very constrained disjunction, as captured by the right graph on figure 3. Intuitively, the set of normal paths is isomorphic to the transitive closure of the graph. A normal path also verifies additionnal properties not captured by structural constraints. For example, union subterms are syntactically different, and the root node $\wedge$ only occurs either in the first or in the second step of the expression. To prove the existence of such a normal form, we adopted a logical approach which model the rewriting system in the calculus of inductive constructions, beyond the scope of this paper. We are currently working on the normalization issues (in particular the termination and confluence proofs; see [3] for general issues) with the Coq proof assistant [17].

Optimization Phase The normal path $\overline{p}$ is transformed into a path p' using a set of optimization rules $\mathcal{O}$. This transformation is noted $\overline{p} \xrightarrow{\mathcal{O}^*}\xrightarrow{\mathcal{C}^*} p'$. Note that the set of rules $\mathcal{C}$ is applied again in order to eliminate void paths potentially introduced by $\mathcal{O}$. The following section details how a path can be optimized using the containment relation over XPath expressions.

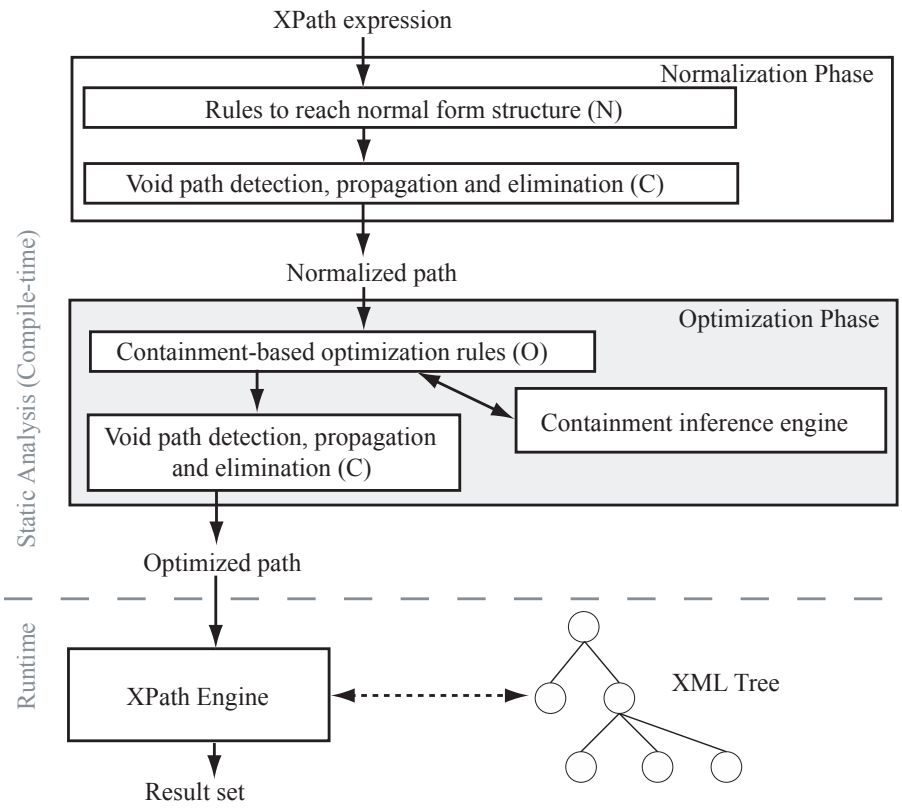

Figure 2: Applicative View of the Transformation Architecture.

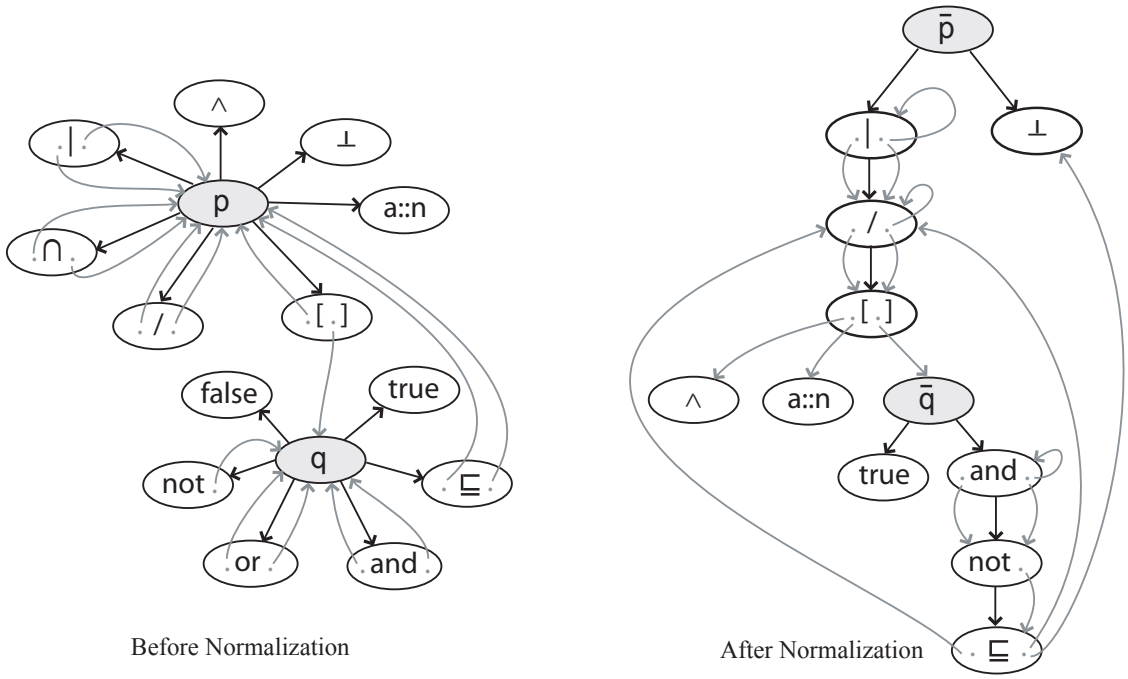

Figure 3: Stratification and Associativity of Operators.

4 Using Containment for XPath Optimization

Deciding Containment The XPath fragment we consider in this paper is quite close to the one found in [8], for which the containment has been shown to be decidable. In order to assert containment facts, we rely on an inference and rewriting based approach described in [10]. The latter uses logical rules à la Hilbert (see [5] for a comprehensive introduction to logical formalisms):

$$\frac{A_1 ... A_n}{B}\,\mathrm{r}$$

where A_i and B are judgments over paths and qualifiers, and r is the rule name. Such a logical rule means that if all judgments A_i are true, then B is true. For example, the containment rule c1 states that the void path is contained in any other path, c2 constructs the reflexivity directly, d1 addresses the comparison of steps and d2 captures the general behavior of containment relation w.r.t. the slash operator:

$$\frac{}{\bot \leq p}\,\mathrm{c1} \qquad \frac{}{p \leq p}\,\mathrm{c2} \qquad \frac{a_1 \leq a_2 \quad N_1 \leq N_2}{a_1{::}N_1 \leq a_2{::}N_2}\,\mathrm{d1} \qquad \frac{p_1[p_2] \leq p_3[p_4] \quad p_2 \leq p_4}{p_1/p_2 \leq p_3/p_4}\,\mathrm{d2} \qquad \frac{p_1 \leq p_2 \quad q_1 \Rightarrow q_2}{p_1[q_1] \leq p_1[q_1]}\,\mathrm{d3}$$

The comparison of steps (rule d1) relies on a partial ordering of XPath axes and nodetests:

$$\text{self} \leq \text{descendant-or-self} \qquad \text{child} \leq \text{descendant} \leq \text{descendant-or-self} \qquad n \leq * \leq \textbf{node}()$$

The containment relation $\leq$ is also defined using a dual relation $\Rightarrow$, the logical implication between qualifiers. Hence, the comparison of qualified paths through rule d3 involves the qualifier implication, and rules e_i handles connectives:

$$\frac{p_3 \leq p_1 \quad p_2 \leq p_4}{p_1 \sqsubseteq p_2 \Rightarrow p_3 \sqsubseteq p_4}\,\mathrm{e1} \qquad \frac{q_2 \Rightarrow q_1}{\textbf{not}\ q_1 \Rightarrow \textbf{not}\ q_2}\,\mathrm{e2} \qquad \frac{q_1 \Rightarrow q}{q_1\ \textbf{and}\ q_2 \Rightarrow q}\,\mathrm{e4} \qquad \frac{q \Rightarrow q_1 \quad q \Rightarrow q_2}{q \Rightarrow q_1\ \textbf{and}\ q_2}\,\mathrm{e5}$$

Two techniques allow us to reduce the size of the axiomatic system: the first one is an equivalence relation $\equiv$ that typically captures commutativity and associativity, and which is fully (left and right) congruent w. r. t. the containment and implication[1] , and the fundamental rule h which involves the normalization process $\mathcal{N}$ in order to transform the operands:

$$\frac{p_1 \xrightarrow{\mathcal{N}^*} p_1' \quad p_2 \xrightarrow{\mathcal{N}^*} p_2' \quad p_1' \leq p_2'}{p_1 \leq p_2}\,\mathrm{h}$$

For more details on the approach, with a larger XPath fragment, the reader can refer to [10]. We are currently working on a full characterization of this axiomatic system[2], using the Coq proof assistant [17] and semi-automated strategies in order to tackle the combinatorial complexity. An algorithm for deciding the containment is considered as a proof tree computation.

Containment-Based Optimization Rules Our system is basically the set $\mathcal{O}$ of conditional term rewriting rules. Numerators are logical conditions that must be satisfied by terms or subterms in order to apply the rules. Most of the rules make use of the containment relation, and thus, each reduction step must be justified by using the inference system for containment. We distinguishes two sets of rules:

1. *redundancy elimination.* Rules ru1, ru2 (and their symmetric variants not detailed in figure 4), ru3, ru4 and ru5 handle union operator; rules rs1 and rs2 (figure 5) eliminate qualifier conditions induced by path composition (through the / operator), such as $a[*]/b$. This latter case is formally optimized through application of rs1 as follow

$$\cfrac{\cfrac{\cfrac{\cfrac{b \leq * \quad \bot \leq \bot}{* \sqsubseteq \bot \Rightarrow b \sqsubseteq \bot}\,\mathrm{e1}}{\textbf{not}\ b \sqsubseteq \bot \Rightarrow \textbf{not}\ * \sqsubseteq \bot}\,\mathrm{e2}}{\textbf{not}\ b \sqsubseteq \bot \Rightarrow *}\,\equiv}{a[*]/b \rightarrow a/b}\,\mathrm{rs1}$$

   The rules rs3, rs4 manage cases where the qualifier is naturally induced by the path, as in $a[\text{self}{::}*]$ and the last subset of figure 7 is dedicated to qualifier simplification.

[1] left congruence: for all p_1, p_2, p if $p_1 \leq p$ and $p_1 \equiv p_2$, then $p_2 \leq p$

[2] We are currently investigating properties such as soundness and completeness.

$$\frac{p_1 \leq p_2}{p_1|p_2 \rightarrow p_2}\ \text{ru1} \qquad \frac{p_1 \leq p_2}{p_1/p_3|p_2/p_4 \rightarrow p_2/(p_3|p_4)}\ \text{ru2} \qquad \frac{p_3 \leq p_4 \quad p_4 \leq p_3}{p_1/p_3|p_2/p_4 \rightarrow (p_1|p_2)/p_4}\ \text{ru3}$$

$$\frac{q_1 \Rightarrow q_2 \quad q_2 \Rightarrow q_1}{p_1[q_1] \ \mid\ p_2[q_2] \rightarrow (p_1|p_2)[q_1]}\ \text{ru4} \qquad \frac{a \in \{\text{desc}, \text{child}\}}{p_1/\text{desc}::*/\text{a}::N \ \mid\ p_2/\text{a}::N \rightarrow (p_1 \mid p_2)/\text{desc}::N}\ \text{ru5}$$

Figure 4: Containment-Based Optimization Rules ($\mathcal{O}$): union operator .

2. *void paths elimination.* Some implicit contradictions (that result in void paths) are not captured in $\mathcal{N}$ and $\mathcal{C}$ rewriting systems, because it requires some significant inference power. Rules of figure 6 aims to detect these cases using the path containment or qualifier implication proof system.

$$\frac{\textbf{not}\ (p_2 \sqsubseteq \bot) \Rightarrow q}{p_1[q]/p_2 \rightarrow p_1/p_2}\ \text{rs1} \qquad \frac{\textbf{not}\ (p_2 \sqsubseteq \bot) \Rightarrow q_1}{p_1[q_1\ \textbf{and}\ q_2]/p_2 \rightarrow p_1[q_2]/p_2}\ \text{rs2} \qquad \frac{\textbf{not}\ (p \sqsubseteq \bot) \Rightarrow q}{p[q] \rightarrow p}\ \text{rs3} \qquad \frac{\textbf{not}\ (p \sqsubseteq \bot) \Rightarrow q_1}{p[q_1\ \textbf{and}\ q_2] \rightarrow p[q_2]}\ \text{rs4}$$

Figure 5: Containment-Based Optimization Rules ($\mathcal{O}$): paths vs qualifiers.

$$\frac{q \Rightarrow (p_2 \sqsubseteq \bot)}{p_1[\textbf{not}\ q]/p_2 \rightarrow \bot}\ \text{rv1} \qquad \frac{q \Rightarrow (p \sqsubseteq \bot)}{p[\textbf{not}\ q] \rightarrow \bot}\ \text{rv2} \qquad \frac{q_1 \Rightarrow (p_2 \sqsubseteq \bot)}{p_1[(\textbf{not}\ q_1)\ \textbf{and}\ q_2]/p_2 \rightarrow \bot}\ \text{rv1p} \qquad \frac{q_1 \Rightarrow (p \sqsubseteq \bot)}{p[(\textbf{not}\ q_1)\ \textbf{and}\ q_2] \rightarrow \bot}\ \text{rv2p}$$

Figure 6: Containment-Based Optimization Rules ($\mathcal{O}$): contradictions.

In any axiomatic system construction, the selection of axioms is based on "intuition" and therefore conveys some arbitrary decisions; then, the effort to reach consistency and completeness brings light on the underlying choices behind the axioms, and often provide rationale in order to reorganize, simplify or extend the theory. After developing a short illustration, we will propose some hints in this direction. Let us consider the following example (the "descendant" axis is abbreviated into "desc")

Example 1 $\quad A \ = \ \wedge//a[*/b/c\ \textbf{and}\ \textit{desc::}b]$

The first normalization step expands the syntactic sugar (here $\wedge//a$). Then the desc-or-self axis is transformed into a disjunctive term and finally, qualifiers are introduced at each path step. In the notation below, $\xrightarrow{\mathcal{N}}$ denotes a one-step derivation from the rewriting system $\mathcal{N}$ and $\xrightarrow{\mathcal{N}^+}$ denotes the transitive closure of this relation.

$$\begin{array}{lll} A & \xrightarrow{\mathcal{N}} & \wedge/\text{desc-or-self}::*/a[*/b/c\ \textbf{and}\ \text{desc}::b] \\ & \xrightarrow{\mathcal{N}^+} & \wedge/\text{desc}::*/a[*/b/c\ \textbf{and}\ \text{desc}::b] \ \mid\ \wedge/\text{self}::*/a[*/b/c\ \textbf{and}\ \text{desc}::b] \\ & \xrightarrow{\mathcal{N}^+} & \wedge[\textbf{true}]/\text{desc}::*[\textbf{true}]/\text{child}::a[Q] \ \mid\ \wedge[\textbf{true}]/\text{child}::a[Q] \quad = \overline{A} \end{array}$$

$$\text{with} \quad Q = \text{child}::*[\textbf{true}]/\text{child}::b[\textbf{true}]/\text{child}::c[\textbf{true}]\ \textbf{and}\ \text{desc}::b[\textbf{true}]$$

At this stage, several optimizing rewriting steps can be applied using rules from $\mathcal{O}$, described on figures 4,5,7:

$$\begin{array}{lll} \overline{A} & \xrightarrow{\text{rs3,rs3,rs3}} & \wedge/\text{desc}::*/\text{child}::a[Q] \ \mid\ \wedge/\text{child}::a)[Q] \\ & \xrightarrow{\text{ru4}} & (\wedge/\text{desc}::*/\text{child}::a \ \mid\ \wedge/\text{child}::a)[Q] \\ & \xrightarrow{\text{ru5}} & (\wedge \ \mid\ \wedge)/\text{desc}::a[Q] \\ & \xrightarrow{\text{ru1}} & \wedge/\text{desc}::a[Q] \\ & \xrightarrow{\text{rq1}} & \wedge/\text{desc}::a[\text{child}::*[\textbf{true}]/\text{child}::b[\textbf{true}]/\text{child}::c[\textbf{true}]] \\ & \xrightarrow{\text{rs3}^*} & \wedge/\text{desc}::a[\text{child}::*/\text{child}::b/\text{child}::c] = \wedge/\text{desc}::a[*/b/c] \end{array}$$

So what do we gain? First the descendant-or-self axis is gone, replaced by a descendant axis; second, the desc::b qualifier is gone, and this is certainly the most interesting -and difficult- point, we have to detail here after. To

$$\frac{q_1 \Rightarrow q_2}{q_1 \textbf{ and } q_2 \rightarrow q_1}\,\text{rq1} \qquad \frac{q_1 \Rightarrow q_2}{q_1 \textbf{ and } (\textbf{not } q_2) \rightarrow \textbf{false}}\,\text{rq2} \qquad \frac{p_1 \leq p_2}{p_1 \sqsubseteq p_2 \rightarrow \textbf{true}}\,\text{rq3} \qquad \frac{p_1 \cap p_2 \leq \bot}{p_1 \sqsubseteq p_2 \rightarrow \textbf{false}}\,\text{rq4}$$

Figure 7: Containment-Based Optimization Rules ($\mathcal{O}$): qualifier simplification.

achieve this step, we must use the inference system, as the rule $rq1$ requires to prove that if $*/b/c$ is not empty, then so is desc::b :

$$\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\dfrac{\cdot \leq \cdot \quad \dfrac{\text{child} \leq \text{desc} \quad * \leq *}{\text{child::}* \leq \text{desc::}*}\,\text{d1}}{./\text{child::}* \leq ./\text{desc::}*}\,\text{d2} \quad \dfrac{\text{child} \leq \text{desc} \quad b \leq b}{\text{child::}b \leq \text{desc::}b}\,\text{d1}}{*/b \leq \text{desc::}b}\,\text{g1} \quad \dfrac{}{c \Rightarrow \textbf{true}}\,\text{e3a}}{*/b[c] \leq \text{desc::}b}\,\text{d3} \quad \dfrac{}{\bot \leq \bot}\,\text{c1}}{\text{desc::}b \sqsubseteq \bot \;\Rightarrow\; */b[c] \sqsubseteq \bot}\,\text{e1}}{\text{desc::}b \sqsubseteq \bot \;\Rightarrow\; */b/c \sqsubseteq \bot}\,\text{f2}}{\textbf{not } */b/c \sqsubseteq \bot \;\Rightarrow\; \textbf{not } \text{desc::}b \sqsubseteq \bot}\,\text{e2}}{*/b/c \Rightarrow \text{desc::}b}\,\equiv}{p[*/b/c \textbf{ and } \text{desc::}b] \rightarrow p[*/b/c]}\,\text{rq1}$$

Discussion Through example 1, we outlined the interest of inferring containment facts. A first requirement would be to mathematically characterize the approach:

1. *completeness*. Do we have enough rules in $\mathcal{O}$? Up to now, we only followed a mathematical intuition. In order to state any completeness property, we need an objective complexity measure. A kind of relation, e.g. $p_2 \succ p_1$ saying that p_1 is more optimized than p_2. Then, we would be able to reason on the existence of an optimal form for any path p, and also that $\mathcal{O}$ is complete.

$$\forall x, p_1, p_2 \qquad \mathcal{S}[\![p_1]\!]_x = \mathcal{S}[\![p_2]\!]_x \;\wedge\; p_1 \succ p_2 \quad \Rightarrow \quad p_1 \xrightarrow{\mathcal{O}^+} p_2$$

2. *soundness*. Each rewriting rules in $\mathcal{O}$ must preserve the semantics of paths. This can be most probably proved through a case-by-case analysis of each rule.

$$\forall x, p_1, p_2 \qquad p_1 \xrightarrow{\mathcal{O}} p_2 \quad \Rightarrow \quad \mathcal{S}[\![p_1]\!]_x = \mathcal{S}[\![p_2]\!]_x \;\wedge\; p_1 \succ p_2$$

3. *termination*. As a rewriting system, termination is an essential property. However, the current system seems easy to handle, for instance through a trivial syntactic complexity measure, shown to be monotonic decreasing (see [3]). It could become more difficult, depending on forthcoming rules we could add into the current system. Note that this issue is tightly related to the first point, since the existence of a optimal form requires the finiteness of the computation.

$$\forall p_1, p_3 \;\; \exists p_2 \qquad p_1 \xrightarrow{\mathcal{O}^*} p_2 \Rightarrow \quad \neg(p_2 \xrightarrow{\mathcal{O}} p_3)$$

4. *confluence*. Such a joinability property is quite useful in order to reason on uniqueness of optimal forms. At least, uniqueness up to an equivalence relation $\approx$ is required, otherwise our system would be fragile, as the result would arbitrarily depend on the choice and the application order of rules.

$$\forall p_1, p_2, p_3 \qquad p_1 \xrightarrow{\mathcal{O}^*} p_2 \;\wedge\; p_1 \xrightarrow{\mathcal{O}^*} p_3 \quad \Rightarrow \quad p_2 \approx p_3$$

Beyond characterization, the reader may object that any realistic optimization is tightly bound to evaluation algorithms and data structures, and thus can not be stated on the basis of an abstract description. Our goal is to focus on *general* optimization rules for *any* implementation. However, this notion is controversial, as the boundaries might be hard to shape. Let us consider the following optimization step

$$\text{child::}a \;\mid\; \text{child::}*/\text{desc::}a \rightarrow \text{descendant::}a$$

Do we actually always perform an optimization? On one hand, we probably gain the evaluation cost related to the union operator, and also the memory cost of storing intermediate results; On the other hand, the operator might be not directly implemented (e.g. could be a pipe between two concurrent processes), and the memory gain could be even much fuzzy, as we cannot state any algorithmic hypothesis about storage. More interestingly, what is undoubtedly gained in such a rewriting step is redundancy: child::a and child::$*$ both require scanning all children of the context node, and this has an execution and/or a memory cost, whatever used algorithm and architecture. This notion is roughly captured in [6] through the notion of minimal pattern. The authors define pattern minimality as the non-existence of recovering sub-patterns. We propose to adapt this notion to our approach through the inductive definition of an embedding relation:

Definition 1 *Embedding relation* $\preceq$ *for all paths* p, p_1, p_2 *and qualifiers* q, q_1, q_2

$$\begin{array}{lllllll} p & \preceq & p_1 \mid p_2 & if & p \preceq p_1 & \vee & p \preceq p_2 \\ p & \preceq & p_1/p_2 & if & p \preceq p_1 & \vee & p \preceq p_2 \\ p[q] & \preceq & p_1[q_1] & if & p \preceq p_1 & \wedge & q \preceq q_1 \\ p & \preceq & p_1 & if & p \leq p_1 & & \end{array} \qquad \begin{array}{lllllll} q & \preceq & q_1 \textbf{ and } q_2 & if & q \preceq q_1 & \vee & q \preceq q_2 \\ q & \preceq & q_1 \textbf{ or } q_2 & if & q \preceq q_1 & \vee & q \preceq q_2 \\ \textbf{not } q & \preceq & \textbf{not } q_1 & if & q \preceq q_1 & & \\ q & \preceq & q_1 & if & q \Rightarrow q_1 & & \end{array}$$

Now, this abstract notion of embedding allows us to quantify the internal redundancy of a path p as the number of overlapping embedded pairs:

Definition 2 *Redundancy Measure.*

$$|p| = \textbf{card}\,\{\langle p_1, p_2\rangle \mid p_1 \neq p_2 \ \wedge\ p_1 \preceq p \ \wedge\ p_2 \preceq p \ \wedge\ p_1 \leq p_2\}$$

Now we are able to define the ordering relation required for the characterization of our rewriting system $\mathcal{O}$

Definition 3 *Minimality Order.* $\forall p_1, p_2 \qquad p_1 \succ p_2 \quad$ *iff* $\quad |p_1| > |p_2|$

5 Preliminary Empirical Analysis

In this section, we describe a preliminary set of experiments. Our goal is to check if a particular containment-based optimization remains useful when evaluated by a real world XPath engine, that probably applies internally various optimization techniques, e.g. caching or case based simplification. The experiments compute the XPath expression of example 1:

$$_\wedge//a[*/b/c \text{ and descendant::}b]$$

and its optimized variant on the same set of documents. This one is randomly generated (but complying to precise construction rules) in order to cover a significant variety of trees. The reader may find further details on our experimental protocol in the appendix B.

We defined three different factors: maximum branching width, maximal depth of the tree, and maximal cardinal of the node label set. Each of the three experiments measures the time response depending on the variation of one particular factor. We ran both queries five times on each document sample and the results were averaged (standard deviation was observed as a criterion of the experiment's apparatus quality). The final result is built from the average of five runs on five random trees. Figures 8 and 9 summarize the test results for two different XPath engines: "Xalan C++" [18] and "Libxml" [19]. In the first experiment, we varied the depth factor and fixed the maximum possible width factor to be 11[2] and the number of element names to be 20[2]. The results are captured on the leftmost graph of each figure. The middle graphs capture the results of the second experiment, when we varied the maximum width factor and fixed the depth factor to be 12 and the number of element names to be 20. Finally, the rightmost graphs show results of the third experiment, when we varied the number of element names and fixed the maximum depth and width factor to be 7 and 6, respectively.

The results show that the optimization is working best for deeper and wider trees. In the first and second experiments, although query performance degrades exponentially (as the size of the tree increases exponentially), the non-optimized query time increases much faster than the optimized query time. For very deep trees, when the depth reaches 8, the optimization provides a three-fold increase in query performance using Xalan C++, as well as for very wide trees, when the width reaches 12. The third experiment confirms that the optimization improves performance when the query retrieves either small or large result sets.

[2] See Appendix B for the definitions of factor units.

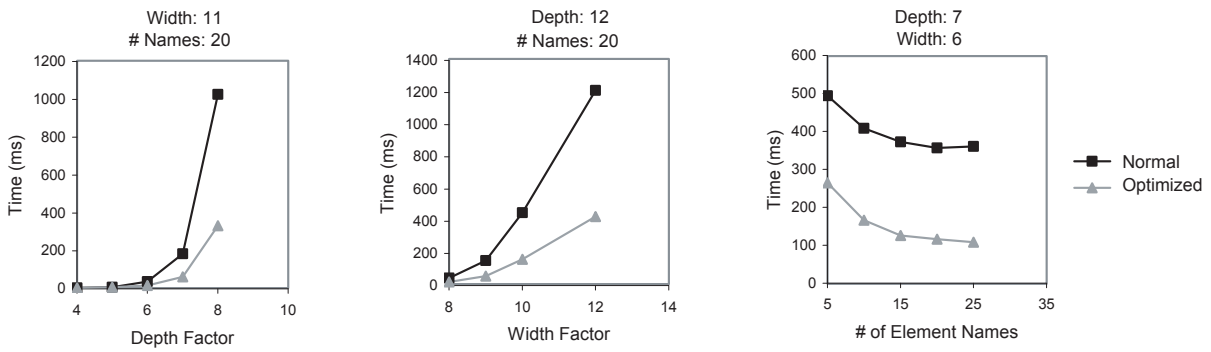

Figure 8: Varying Factors Independently (Xalan C++)

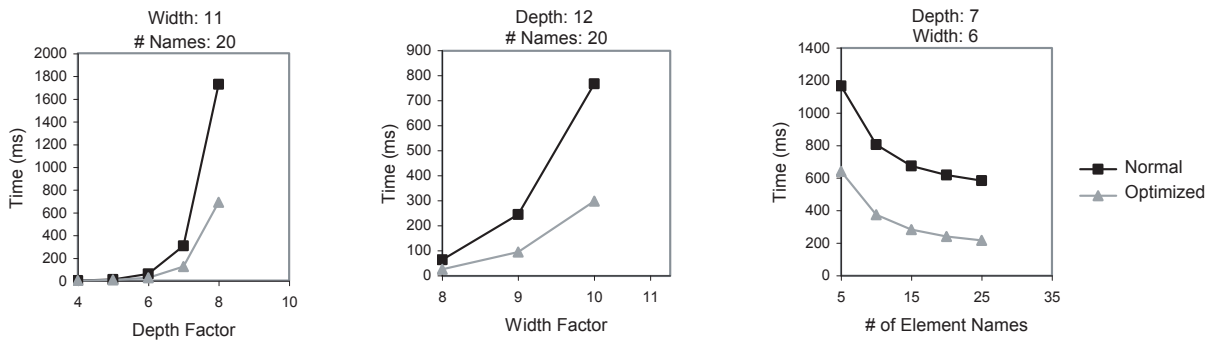

Figure 9: Varying Factors Independently (Libxml)

Experiments show that containment-based optimization can improve query performance for large, deep, and bushy trees, wether the query has a high match probability or not. The experiments confirm the interest of the static optimization we propose - at least for the particular case we considered - in the sense that XPath engines do not achieve at runtime any equivalent optimization.

6 Conclusion and Perspectives

In this paper, we have described a static analysis framework for XPath optimization, in which we have proposed to use conditional rewriting rules based on the containment relation for XPath expressions. The optimization aims at eliminating implicit redundancies and contradictions found in user-defined or generated XPath queries. Preliminary experiments show that this framework could provide an interesting basis for a static optimization layer, that could be either implemented on top of an XPath engine, or built inside a compiler.

Further on-going work on the issues explored in this paper is worth mentioning. First, we are working on the completeness issue for the containment inference system. Proving this important property would allow us to state that we can determine whether the containment relation holds or not between any two XPath expressions. Note that even an incomplete containment could be of practical interest as well, especially if the causes are well understood. Second, our optimization technique deserves further investigation for characterizing an optimality notion for XPath expressions (as briefly mentioned in section 4), and ensuring that all kinds of redundancies and contradictions can be detected. In addition, the continuity of this work would require looking at decision issues and algorithmic complexity.

Finally, the real world application of this work relies on our ability to extend the XPath fragment we consider. The next step - our underlying goal - is to address potential scalability issues. The introduction of the **count**(p), **position**() and **last**() primitives in qualifiers are of very first importance, both for the expressive power and for showing the potential of the logic-based framework we are working on. Handling quantitative assertions seems

quite tractable using a Presburger-like arithmetic (known to be decidable) together with rules such as:

$$\frac{p_1 \leq p_2}{\textbf{count}\ (p_1) \leq \textbf{count}\ (p_2)} \qquad \frac{}{0 \leq \textbf{count}\ (p)}$$

Another promising direction would be to consider structural constraints on documents (so-called *schemas*) in order to handle more precise containment assertions in the same conceptual framework.

References

[1] S. Abiteboul and V. Vianu, *Regular path queries with constraints*, In Proc. 16th PODS, pages 122–133, Tucson (Arizona, USA), May 1997. ACM Press.

[2] S. Amer-Yahia, S. Cho, L. V. S. Lakshmanan, and D. Srivastava, *Minimization of Tree Patterns Queries*, In Proc. ACM SIGMOD Conf., pages 497-508, 2001.

[3] N. Dershowitz and J-P. Jouannaud, *Rewrite Systems*, Handbook of Theoretical Computer Science, Volume B: Formal Models and Semantics (B) 1990: 243-320.

[4] A. Deutsch and V. Tannen, *Containment and Integrity Constraints for XPath Fragments*, KRDB 2001.

[5] G. Dowek, *Introduction to Proof Theory*, http://www.lix.polytechnique.fr/~dowek/

[6] S. Flesca, F. Furfaro; E. Masciari *On the minimization of XPath Queries*, Proceedings of the 29th VLDB Conf., Berlin, January 2000, http://wwwdb.informatik.uni-rostock.de/vldb2003/papers/S06P02.pdf.

[7] G. Miklau and D. Suciu, *Containment and equivalence for an XPath fragment*, In 21st ACM Symposium on Principles of Databases Systems (PODS), Madison, Wisconsin, pages 65-76, 2002.

[8] F. Neven and T. Schwentick,*XPath Containment in the Presence of Disjunction, DTDs, and Variables*, International Conference on Database Theory, 2003.

[9] D. Olteanu, H. Meuss, T. Furche, F. Bry, *XPath: Looking Forward*, In Proc. of the EDBT Workshop on XML Data Management (XMLDM), 2002.

[10] J-Y. Vion-Dury, N. Layaïda, *Containment of XPath expressions: an Inference and Rewriting based approach*, Extreme Markup Languages, August 4-8, 2003.

[11] P. Wadler, *Two semantics for XPath*, January 2000, http://www.research.avayalabs.com/user/wadler/papers/xpath-semantics/xpath-semantics.pdf.

[12] P. Wood, *Containment for XPath Fragments under DTD Constraints*, In Proc. of the 9th International Conference on Database Theory, Pages: 300 - 314, 2003.

Recommendations & Working Drafts

[13] A. Berglund, S. Boag, D. Chamberlin, M. Fernández, M. Kay, J. Robie, and J. Siméon, *XML Path Language (XPath) 2.0*, W3C Working Draft, August, 2003, http://www.w3.org/TR/2003/WD-xpath20-20030822.

[14] D. Draper, P. Fankhauser, M. Fernández, A. Malhotra, K. Rose, M. Rys, J. Siméon, P. Wadler, *XQuery 1.0 and XPath 2.0 Formal Semantics*, W3C Working Draft, August, 2003, http://www.w3.org/TR/2003/WD-xquery-semantics-20030822/.

[15] J. Clark, *XSL Transformations (XSLT) Version 1.0*, W3C Recommendation, November 1999, http://www.w3.org/TR/1999/REC-xslt-19991116.

[16] J. Clark, S. DeRose, *XML Path Language (XPath) Version 1.0*, W3C Recommendation, November 1999, http://www.w3.org/TR/1999/REC-xpath-19991116.

Software

[17] The Coq proof assistant, http://coq.inria.fr

[18] Xalan-C++ version 1.7, http://xml.apache.org/xalan-c/

[19] Libxml2, http://www.xmlsoft.org/

A XPath Semantics

Two functions $\mathcal{S}$ and $\mathcal{Q}$ respectively define the semantics of paths and qualifiers, provided a context node x in the tree:

$$
\begin{array}{lcl}
\multicolumn{3}{l}{\mathcal{S} : \textit{Path} \longrightarrow \textit{Node} \longrightarrow \text{Set}(\textit{Node})} \\
\mathcal{S}[\![\wedge]\!]_x & = & \{x_1 \mid x_1 \rightarrowtriangle^+ x \wedge \text{root}(x_1)\} \\
\mathcal{S}[\![\bot]\!]_x & = & \emptyset \\
\mathcal{S}[\![p_1 \mid p_2]\!]_x & = & \mathcal{S}[\![p_1]\!]_x \cup \mathcal{S}[\![p_2]\!]_x \\
\mathcal{S}[\![p_1 \cap p_2]\!]_x & = & \{x_1 \mid x_1 \in \mathcal{S}[\![p_1]\!]_x \wedge x_1 \in \mathcal{S}[\![p_2]\!]_x\} \\
\mathcal{S}[\![p_1/p_2]\!]_x & = & \{x_2 \mid x_1 \in \mathcal{S}[\![p_1]\!]_x \wedge x_2 \in \mathcal{S}[\![p_2]\!]_{x_1}\} \\
\mathcal{S}[\![(p)]\!]_x & = & \mathcal{S}[\![p]\!]_x \\
\mathcal{S}[\![p[q]]\!]_x & = & \{x_1 \mid x_1 \in \mathcal{S}[\![p]\!]_x \wedge \mathcal{Q}[\![q]\!]_{x_1}\} \\
\mathcal{S}[\![\text{a::}N]\!]_x & = & \{x_1 \mid x_1 \in f_a(x) \wedge \mathcal{T}_a(x_1, N)\}
\end{array}
$$

$$
\begin{array}{lcl}
\multicolumn{3}{l}{\mathcal{Q} : \textit{Qualifier} \longrightarrow \textit{Node} \longrightarrow \textit{Boolean}} \\
\mathcal{Q}[\![\textbf{true}]\!]_x & = & \text{true} \\
\mathcal{Q}[\![\textbf{false}]\!]_x & = & \text{false} \\
\mathcal{Q}[\![q_1 \textbf{ and } q_2]\!]_x & = & \mathcal{Q}[\![q_1]\!]_x \wedge \mathcal{Q}[\![q_2]\!]_x \\
\mathcal{Q}[\![q_1 \textbf{ or } q_2]\!]_x & = & \mathcal{Q}[\![q_1]\!]_x \vee \mathcal{Q}[\![q_2]\!]_x \\
\mathcal{Q}[\![(q)]\!]_x & = & \mathcal{Q}[\![q]\!]_x \\
\mathcal{Q}[\![\textbf{not } q]\!]_x & = & \neg\mathcal{Q}[\![q]\!]_x \\
\mathcal{Q}[\![p_1 \sqsubseteq p_2]\!]_x & = & \mathcal{S}[\![p_1]\!]_x \subseteq \mathcal{S}[\![p_2]\!]_x
\end{array}
$$

The navigational semantics of axes relies on the relation $\rightarrowtriangle$ that maps a node to its children in the XML tree, and on the transitive closure $\rightarrowtriangle^+$ of this relation; the function $\mathcal{T}$ performs a node test (see the tables below)

a	$f_a(x)$
self	$\{x\}$
child	$\{x_1 \mid x \rightarrowtriangle x_1\}$
descendant	$\{x_1 \mid x \rightarrowtriangle^+ x_1\}$
descendant-or-self	$\{x\} \cup \{x_1 \mid x \rightarrowtriangle^+ x_1\}$
attribute	$\{x_1 \mid x \rightarrowtriangle x_1 \wedge \text{attribute}(x_1)\}$
namespace	$\{x_1 \mid x \rightarrowtriangle x_1 \wedge \text{namespace}(x_1)\}$

N	a	$\mathcal{T}_a(N, x)$
n		name(x)=n
$*$	attribute	attribute(x)
$*$	namespace	namespace(x)
$*$	other	element(x)
text()		text(x)
comment()		comment(x)
processing-instruction()		pi(x)
element()		element(x)
node()		true

B Experimental Environment

We conducted experiments on a laptop PC (IBM Thinkpad T22). It has an Intel(R) Pentium(R) 3 CPU 900Mhz, 256MB RAM and a 56GB (7200 RPM) hard drive. The PC runs Windows XP Professional version 2002 SP1. We installed the binaries versions of Xalan C++ 1.7.0 and Libxml 2.6.4 for win32.

Test Programs Our test program uses static linking to call the original Libxml2 and Xalan DLLs, found in the distributions. The test programs basically initialize the XPath engine, and then determines the time spent to evaluate an XPath expression by recording the system time before and after a call:

- to the method "xmlXPathEvalExpression()", in charge of the evaluation when using Libxml;
- to the method "evaluate ()" of the class "XPathEvaluator" when using Xalan.

The test programs were compiled with Microsoft Visual C++ 6.0 in release mode (debug mode switched off). We isolated the machine for testing. Only the test program and normal operating system background processes are running during the testing period.

Random Experiment We generated random XML documents for testing with the following configurable parameters:

- The *depth factor* represents the level of nesting of elements in the XML document, in order to model the recursive structure of a document model. The depth factor controls the depth of subtrees and can be fixed or randomly chosen from a range.
- The *width factor* describes the number of children of a non-leaf node in the tree, in order to model the bushiness of an XML instance. The width factor can be fixed or randomly chosen from a range.

- The *number of element names* describes the number of possible element names, in order to model the cardinality of a namespace allowed by a document model. This factor is fixed.

The tree is made random in two ways:

1. In order to capture a broad range of different topologies, the depth and width of any subtree is made random to test with short, bushy trees, or deep, skinny trees or some combination thereof. Because of limited memory, we considered depth and width factors up to 12. The sizes of the smallest and largest XML documents used for testing were respectively 1KB and 15MB.

2. Each node in the tree can become any element from the allowed range with the same probability, e.g. if the number of element names is 26, a node in the tree will either be named “a” or “b”... or “z” in an equiprobable way. This allows us to vary the query match probability from low to high, in order to test when XPath expressions have a higher chance to match and retrieve large result sets, or when queries have a lower chance to match and retrieve small result sets, or some combination thereof.