

# Type-Based XML Projection

Véronique Benzaken<sup>1</sup>

Giuseppe Castagna<sup>2</sup>

Dario Colazzo<sup>1</sup>

Kim Nguyễn<sup>1</sup>

<sup>1</sup>LRI, Université Paris-Sud 11, Orsay - France

<sup>2</sup>École Normale Supérieure de Paris - France

## ABSTRACT

XML data projection (or pruning) is one of the main optimization techniques recently adopted in the context of main-memory XML query-engines. The underlying idea is quite simple: given a query  $Q$  over a document  $D$ , the subtrees of  $D$  not necessary to evaluate  $Q$  are pruned, thus obtaining a smaller document  $D'$ . Then  $Q$  is executed over  $D'$ , hence avoiding to allocate and process nodes that will never be reached by navigational specifications in  $Q$ .

In this article, we propose a new approach, based on types, that greatly improves current solutions. Besides providing comparable or greater precision and far lesser pruning overhead our solution, unlike current approaches, takes into account backward axes, predicates, and can be applied to multiple queries rather than just to single ones. A side contribution is a new type system for XPath able to handle backward axes, which we devise in order to apply our solution.

The soundness of our approach is formally proved. Furthermore, we prove that the approach is also complete (i.e., yields the best possible type-driven pruning) for a relevant class of queries and DTDs, which include nearly all the queries used in the XMark and XPathMark benchmarks. These benchmarks are also used to test our implementation and show and gauge the practical benefits of our solution.

## 1. MOTIVATIONS AND CONTRIBUTION

As explained by Marian and Siméon [14], main-memory XML query engines are often the primary choice for applications that do not wish or cannot afford to build secondary storage indexes or load a database before query processing. One of the main optimisation techniques recently adopted in this context is XML data projection (or pruning) [14, 9].

The basic idea behind document projection is very simple and powerful at the same time. Given a query  $Q$  over a document  $D$ , sub-trees of  $D$  that are not necessary to evaluate  $Q$  are pruned, thus yielding a smaller document  $D'$ . Then  $Q$  is executed over  $D'$ , hence avoiding to allocate and process nodes that will never be reached by navigational specifications in  $Q$ . This ensures that evaluation over  $D'$  is equivalent to and more efficient than the evaluation over  $D$ .

Permission to copy without fee all or part of this material is granted provided that the copies are not made or distributed for direct commercial advantage, the VLDB copyright notice and the title of the publication and its date appear, and notice is given that copying is by permission of the Very Large Data Base Endowment. To copy otherwise, or to republish, to post on servers or to redistribute to lists, requires a fee and/or special permission from the publisher, ACM.

VLDB '06, September 12-15, 2006, Seoul, Korea.

Copyright 2006 VLDB Endowment, ACM 1-59593-385-9/06/09.

As shown in [14, 9], XML navigation specifications expressed in queries tend to be very selective, especially in terms of document structure. Therefore, pruning may yield significant improvements both in terms of execution time and in terms of memory usage (for main-memory XML query engines, very large documents can not be queried without pruning).

### 1.1 State of the art

Marian and Siméon[14] propose that the actual data-needs of a query  $Q$  (that is, the part of data that is necessary to the execution of the query) is determined by statically extracting all paths in  $Q$ . These paths are then applied to  $D$  at load time, in a SAX-event based fashion, in order to prune unneeded parts of data. The technique is powerful since: (i) it applies to most of XQuery core, (ii) it can be applied to a set of queries over the same document, and (iii) it does not require any *a priori* knowledge of the structure of  $D$ . However, this technique suffers some limitations. First, the document loader-pruner is not able to manage *backward axes* nor path expressions with predicates (sometimes called “qualifiers”) which, especially the latter, can contain precious information to optimise pruning. Also, as a consequence of (iii), the technique does not behave efficiently in terms of loading time and pruning precision (hence, memory allocation) when  $//$  occurs in paths. Indeed, when  $//$  is present in a projection path, the pruning process requires to visit all descendants of a node in order to decide whether the node contains a useful descendant. What is worst is that pruning time tends to be quite high and it drastically increases (together with memory consumption) when the number of  $//$  augments in the pruning path-set. As a matter of facts, in this technique pruning corresponds to computing a further query, whose time and memory occupation may be comparable to those required to compute the original query. In particular, in this technique every occurrence of  $//$  may yield a full exploration of the tree (e.g. see in [14] the test for the XMark [17] query Q7 which only contains three  $//$  steps and for which just computing the pruning takes longer than executing the query on the original document). Therefore, pruning execution overhead and its high memory footprint may jeopardise the gains obtained by using the pruned document. Finally, as we explain in Section 5, the precision of pruning drastically degrades (even nullified) for queries containing the XPath expressions  $descendant::node[cond]$ , which are very useful and used in practice.

Bressan *et al.* [9] introduce a different and quite precise XML pruning technique for a subset of XQuery FLWR expressions. The technique is based on the *a priori* knowledge of a data-guide for  $D$ . The document  $D$  is first matched against an abstract representation of  $Q$ . Pruning is then performed at run time, it is very precise, and, thanks to the use of some indexes over the data-guide, it ensures good improvements in terms of query execution time. However,

the technique is one-query oriented, in the sense that it cannot be applied to multiple queries, it does not handle XPath predicates, and cannot handle backward axes (recall that the encodings of [15] are defined for XPath, and no extension to XQuery-like languages is known). Also, the approach requires the construction and management of the data-guide and of adequate indexes.

## 1.2 Our contribution

In this article, we present a new pruning approach which is applicable in the presence of typed XML data. This is often the case, as most applications require that data are valid with respect to some external schema (e.g. DTD or XML Schema).

Our technique combines the advantages of the previously mentioned works while relaxing their limitations. Unlike [14, 9], our approach accounts for backward axes, performs a fine-grained analysis of predicates, allows (unlike [9]) for dealing with bunches of queries, and (unlike [14]) cannot be jeopardised by pruning overhead. Our solution provides comparable or greater precision than the other approaches, while it requires always negligible or no pruning overhead. Moreover, contrary to [14, 9], our approach is formally proved to be *sound* (pruning does not alter the result of queries) and, furthermore, we can also prove it to be *complete* (it produces the best possible type-driven pruning) for a substantial class of queries and DTDs.

For the sake of presentation we introduce our framework in three steps. In the first step, we consider a simplified version of XPath, we dub  $\text{XPath}^\ell$ , which includes only upward/downward axes and unnested disjunctive predicates. We define for  $\text{XPath}^\ell$  a static analysis that determines a set of type names, a *type projector*, that is then used to prune the document(s). One of the particular features of this approach is that our pruning algorithm is characterised by a constant (and low) memory consumption and by an execution time linear in the size of the document to prune. More precisely, a pruning based on type projectors is equivalent to a single bufferless one-pass traversal of the parsed document (it simply discards elements not generated by any of the names in the projector). So if embedded in query processors, pruning can be executed during parsing and/or validation and brings no overhead, while if used as an external tool it requires a time always smaller than or equal to the time used to parse the queried document. Soundness and (partial) completeness results for the static analysis are stated.

The second step consists of extending the analysis to the whole XPath (more precisely, to XPath 1.0), that is, we need to show how to deal with missing axes and with general predicates as defined in the XPath specification. This is done by associating to each XPath query  $Q$  a  $\text{XPath}^\ell$  query  $P$  which soundly approximates  $Q$ , in the sense that the projector inferred for  $P$  is also a sound projector for  $Q$ .

The final step is to extend the approach to XQuery (hence, to XPath 2.0). This is obtained by defining a path extraction algorithm as done in [14]. Our path extraction algorithm improves in several aspects (in particular, in terms of extracted paths' selectivity) the one of [14]. It also computes the  $\text{XPath}^\ell$  approximation of the extracted paths so that the static analysis of the first step can be directly applied to them.

We gauged and validated our approach by testing it both on the XPathMark [12] and on the XMark [17] benchmarks. This validation confirmed expected results: thanks to the handling of backward axes and of predicates the precision of our pruning is in general noticeably higher than for current approaches; the pruning time is linear in the size of the queried document and has a very low memory footprint; the time of the static analysis is always negligible (lower than half a second) even for complex queries and DTDs. But bench-

marks also brought unexpected (and pleasant) results. In particular, they showed that type-based pruning brings benefits that go beyond those of the reduced size of the pruned document: by excluding a whole set of data structures (those whose type names are not included in the type projector), the pruning may drastically reduce the resources that must be allocated at run-time by the query processor. For instance, our benchmarks show that for several XMark and XPathMark queries our pruning yields a document whose size is two thirds of the size of the original document, but the query can then be processed using three times less memory than when processed on the original document. This is a very important gain, especially for DOM-based processors, or memory sensitive processors as Galax [1]. As an aside we want to stress that our technique relies on the definition of a new type system for XPath able to handle backward axes, which constitutes a contribution on its own.

The article is organised as follows. Section 2 introduces basic definitions and notations: data model, DTD, validation, projection, type projector. In Section 3 we define  $\text{XPath}^\ell$  and its semantics, and formally describe how general XPath predicates can be soundly approximated in it. In Section 4 we present our type projectors inference algorithm for  $\text{XPath}^\ell$ , state its formal properties, and deal with the missing XPath axes. In Section 5 we extend our approach to XQuery. Section 6 discusses our implementation and reports the results of our benchmarks. We finally conclude in Section 7 by presenting the perspectives of this work.

For space reasons all proofs of properties are omitted from this presentation. They can be found in the extended version of this work.

## 2. NOTATIONS

### 2.1 Data Model

For the sake of concision we present our solution for a simplified version of the XQuery data model where we do not consider node attributes. The extension of our approach to attributes is straightforward (and included in our implementation, see Section 6). An instance of the XQuery data model can then be generated by the following grammar:

$$\begin{array}{ll} \mathbf{Trees} & t ::= s_{\mathbf{i}} \mid l_{\mathbf{i}}[f] \\ \mathbf{Forest} & f ::= () \mid f, f \mid t \end{array}$$

Essentially, it is an ordered sequence of labelled ordered *trees* (ranged over by  $t$ ), that is an ordered *forest* (ranged over by  $f$ ), where each node has a unique *identifier* (ranged over by  $\mathbf{i}$ ) and where  $()$  denotes the empty forest. Tree nodes are labelled by *element tags* (ranged over by  $l$ ) while, without loss of generality, we consider only leaves that are text nodes (that is, strings, ranged over by  $s$ ) or empty trees (that is, elements that label the empty forest).

We define a complete partial order  $\preceq$  on forests (and thus on trees) by relating a forest with the forests obtained either by adding or by deleting subforests:

**DEFINITION 2.1 (PROJECTION ( $\preceq$ )).** *Given two forests  $f$  and  $f'$  we say that  $f'$  is a projection of  $f$ , noted as  $f' \preceq f$ , if  $f'$  is obtained by replacing some subforests of  $f$  by the empty forest.*

**DEFINITION 2.2 (GOOD FORMATION).** *A forest is well formed if every identifier  $\mathbf{i}$  occurs in it at most once. Given a well-formed forest  $f$  and an identifier  $\mathbf{i}$  occurring in it, we denote by  $f@_{\mathbf{i}}$  the unique subtree  $t$  of  $f$  such that  $t = s_{\mathbf{i}}$  or  $t = l_{\mathbf{i}}[f']$ . The set of identifiers of a forest  $f$  is then defined as  $\text{Ids}(f) = \{\mathbf{i} \mid \exists t. f@_{\mathbf{i}} = t\}$*

Henceforth we will consider only well-formed forests and confound the notions of a node with that of the identifier of the node.

DEFINITION 2.3 (ROOT ID). *Given a tree  $t$ , if  $t = s_1$  or  $t = l_1[f]$  then we define  $\text{RootId}(t) = \mathbf{i}$ .*

## 2.2 DTDs and validation

In this work we present the approach for DTDs, but the treatment for XML Schema is similar.<sup>1</sup> Following [13] we define a DTD as a *local tree grammar*, namely a pair  $(X, E)$  where  $X$  is a distinguished *name* (actually, a non-terminal meta-variable) and  $E$  is a set of productions (or *edges*) of the form  $\{X_1 \rightarrow R_1, \dots, X_n \rightarrow R_n\}$  such that

1. the  $X_i$ 's are pairwise distinct;
2. each  $R_i$  is of the form  $a_i[r_i]$  or *String*, where  $a_i$  is an element tag, and each  $r_i$  is a regular expression over *names*  $\{X_1, \dots, X_n\}$ ;
3. for each pair  $X_i \rightarrow a_i[r_i]$  and  $X_j \rightarrow a_j[r_j]$ ,  $i = j$  if and only if  $a_i = a_j$ ;
4.  $X$  is in  $\{X_1, \dots, X_n\}$  (it denotes the root element type).

In the following we write  $\text{Names}(r)$  for the set of all names used in  $r$  and  $DN(E)$  for the set of names defined in  $E$  (that is,  $\{X_1 \dots X_n\}$ ). We also say that  $r$  is a regular expression over  $(X, E)$ , if  $r$  is a regular expression over names in  $DN(E)$ . We will use  $W, X, Y, Z$  to range over *names*. We use Greek letters to range over sets of names (in particular we use  $\pi$  to stress that the set of names is a *type projector* [cf. Def 2.6] and  $\kappa$  and  $\tau$  to stress that the set is used as a context or as a type, respectively [cf. Section 4.1]) and  $S$  to range over sets of (node) identifiers. When speaking of DTDs we will often identify them with their set of edges  $E$ , leaving the root  $X$  as implicit.

DEFINITION 2.4 (VALID TREES). *A tree  $t$  is valid with respect to a DTD  $(X, E)$ , if there exists a mapping (interpretation)  $\mathfrak{S}$  from  $\text{Ids}(t)$  to  $DN(E)$  such that:*

1.  $\mathfrak{S}(\text{RootId}(t)) = X$
2. for each  $\mathbf{i}$  in  $\text{Ids}(t)$ , if  $t@i = s_1$  then  $\mathfrak{S}(\mathbf{i}) = Y$  and  $(Y \rightarrow \text{String}) \in E$
3. for each  $\mathbf{i}$  in  $\text{Ids}(t)$ , if  $t@i = l_1[t_1, \dots, t_n]$ , then  $\mathfrak{S}(\mathbf{i}) \rightarrow l[r] \in E$  and  $\mathfrak{S}(\text{RootId}(t_1)), \dots, \mathfrak{S}(\text{RootId}(t_n))$  is generated by  $r$ .

*In this case we say that  $t$  is  $\mathfrak{S}$ -valid with respect to  $(X, E)$  and write  $t \in_{\mathfrak{S}}(X, E)$  to indicate it.*

Algorithms to validate XML trees are well known (see [13]). Every validation algorithm produces, as a side effect, an interpretation for the validated tree. Note that if  $t$  is valid with respect to a DTD, then there is a unique interpretation  $\mathfrak{S}$  from  $t$  to the DTD. This is a direct consequence of the fact that, in DTDs, element tags determine their content (as stated by the third condition on local tree grammars).

## 2.3 Type projectors

Given a tree  $t$  valid with respect to a DTD  $(X, E)$ , we can use subsets of  $DN(E)$  to project that tree. Essentially, only nodes that are associated with names in the projecting subset of  $DN(E)$  are kept in the projection. Of course not every subset of names can be used to project a tree, since we want to delete whole subtrees (not nodes in the middle of a tree), thus if we discard some name, we must also discard all the names it generates. In order to define formally this notion we need to define the reachability relation  $\Rightarrow_E$ , that we introduce below together with several other definitions that we use later in the paper.

<sup>1</sup>The extension of our approach to XML Schema simply needs some special treatment of local elements. More difficult instead is to modify it so as to obtain efficient pruning also for the new XPath 2.0 tests that check the schema of nodes. See the discussion in our conclusion.

DEFINITION 2.5 (FORWARD REACHABILITY). *Given a DTD  $(X, E)$  and  $Z \in DN(E)$ , we write  $Z \Rightarrow_E Y$  if and only if  $Z \rightarrow a[r] \in E$  and  $Y \in \text{Names}(r)$ . We use  $\Rightarrow_E^+$  and  $\Rightarrow_E^*$  to denote respectively the transitive closure and the transitive and reflexive closure of  $\Rightarrow_E$ .*

Strings of names are called *chains* and ranged over by  $c, c_i, c', \dots$ . In particular we use  $\text{Chains}_{(X, E)}(Y)$  to denote the set of all chains rooted at  $Y$ , defined as  $\{Y X_1 \dots X_n \mid Y \Rightarrow_E X_1 \Rightarrow_E \dots \Rightarrow_E X_n, n \geq 0\}$ . We use  $\text{Names}(c)$  to denote the set of all names occurring in a chain  $c$ .

DEFINITION 2.6 (TYPE-PROJECTORS). *Given a DTD  $(X, E)$ , a (possibly empty) set of names  $\pi \subseteq DN(E)$  is a type projector for  $(X, E)$  if and only if there exists  $C \subseteq \text{Chains}_{(X, E)}(X)$  such that*

$$\pi = \bigcup_{c \in C} \text{Names}(c)$$

A type projector is thus a set of names generated (i.e. reached) by a suite of productions starting from the root of the DTD. A type projector can be used to prune a valid tree as follows:

DEFINITION 2.7 (TYPE DRIVEN PROJECTIONS). *Let  $\pi$  be a type projector for  $(X, E)$  and  $t$  a forest or tree such that  $t \in_{\mathfrak{S}}(X, E)$ . The  $\pi$ -projection of  $t$ , noted as  $t \setminus_{\mathfrak{S}} \pi$ ,*

*is defined as follows:*

$$\begin{aligned} l_1[f] \setminus_{\mathfrak{S}} \pi &= l_1[f \setminus_{\mathfrak{S}} \pi] & \mathfrak{S}(\mathbf{i}) \in \pi \\ l_1[f] \setminus_{\mathfrak{S}} \pi &= () & \mathfrak{S}(\mathbf{i}) \notin \pi \\ s_i \setminus_{\mathfrak{S}} \pi &= s_i & \mathfrak{S}(\mathbf{i}) \in \pi \\ s_i \setminus_{\mathfrak{S}} \pi &= () & \mathfrak{S}(\mathbf{i}) \notin \pi \\ (f, f') \setminus_{\mathfrak{S}} \pi &= (f \setminus_{\mathfrak{S}} \pi, (f' \setminus_{\mathfrak{S}} \pi)) \end{aligned}$$

In words, pruning erases (by replacing it by an empty forest) every node that corresponds to a name not in  $\pi$ .

LEMMA 2.8. *Let  $\pi$  be a type projector for  $(X, E)$ . Then for every tree  $t \in_{\mathfrak{S}}(X, E)$  it holds  $(t \setminus_{\mathfrak{S}} \pi) \preceq t$ .*

## 3. XPATH AND XPATH<sup>ℓ</sup>

In XPath, queries are expressed by defining a path of steps separated by  $/$ . For instance,

$$\begin{aligned} Q &= /descendant::author \\ &\quad /child::text[self::node="Dante"] \\ &\quad /parent::book/child::title \end{aligned}$$

is the query that returns all titles of books whose author is “Dante”. First, the navigational part instructs to descend to all text nodes whose parent is an author ( $/descendant::author/child::text$ ), then the predicate selects those nodes that are the string “Dante” ( $[self::node="Dante"]$ ), and finally the navigation ascends to the book element and descends to the title.

The inference rules we define in Section 4 do not work directly on queries such as  $Q$ . The rules are defined for XPath<sup>ℓ</sup> a subset of XPath that we introduce in this section. XPath<sup>ℓ</sup> includes downward and upward axes and a special kind of predicates. In order to statically analyse  $Q$  (or any other XPath query that is not in XPath<sup>ℓ</sup>), we will find a XPath<sup>ℓ</sup> query that approximates  $Q$  soundly with respect to the pruning inferred by the rules (Section 3.3), and use it to deduce the pruning for  $Q$ .<sup>2</sup> Of course, these approximations, as well as those we introduce later on, will only be used to determine the pruning: the pruned document will be queried by the original query.

For the sake of presentation, we first deal with “simple paths”, that is, path expressions with upward and downward axes in which

<sup>2</sup>For instance, the approximation of our sample query  $Q$  is obtained by replacing in  $Q$  the predicate  $[self::node="Dante"]$  for the current one.

no predicate occurs. Then, in Section 3.2 we add XPath<sup>ℓ</sup> predicates, i.e. disjunctions of simple predicates, and finally in Section 3.3 we show how to approximate generic XPath conditions into XPath<sup>ℓ</sup>. The missing axes are dealt with in Section 4.3.

### 3.1 Simple paths

Simple paths are defined by the following grammar:

$$\begin{aligned} SPath & ::= Step \mid SPath/SPath \mid /SPath \\ Step & ::= Axis :: Test \\ Axis & ::= self \mid child \mid descendant \\ & \quad \mid parent \mid ancestor \mid ancestor-or-self \\ & \quad \mid descendant-or-self \\ Test & ::= tag \mid node \mid text \end{aligned}$$

where *tag* is a meta-variable ranging over element tags. Henceforward, we omit the treatment of leading / (i.e., absolute paths) and of **descendant-or-self** and **ancestor-or-self** axes: their handling would blur definitions and can be easily deduced from the rest.

The formal semantics of paths is given in three definitions. First, we formalise *Test* filtering, then *Axis* selections, and finally we combine the two notions to define the semantics of a single step *Axis :: Test*. The definitions comply with the W3C XPath semantics [2].

**DEFINITION 3.1 (FILTERING).** *Given a tree  $t$  and a set of nodes  $S \subseteq Ids(t)$  we define*

$$\begin{aligned} S ::_t l & = \{ \mathbf{i} \in S \mid t @ \mathbf{i} = l_i[f] \} \\ S ::_t \text{node} & = S \\ S ::_t \text{text} & = \{ \mathbf{i} \in S \mid \exists s . t @ \mathbf{i} = s_1 \} \end{aligned}$$

**DEFINITION 3.2 (AXES SELECTION).** *Given a tree  $t$  and a set of nodes  $S \subseteq Ids(t)$  (called context nodes), we define  $\llbracket Step \rrbracket_t(S)$  as the set of nodes resulting by applying *Step* to each node in  $S$*

$$\begin{aligned} \llbracket self \rrbracket_t(S) & = S \\ \llbracket child \rrbracket_t(S) & = \bigcup_{\mathbf{i} \in S} \{ \mathbf{i}' \mid (\mathbf{i}, \mathbf{i}') \in \mathcal{E}(t) \} \\ \llbracket parent \rrbracket_t(S) & = \bigcup_{\mathbf{i} \in S} \{ \mathbf{i}' \mid (\mathbf{i}', \mathbf{i}) \in \mathcal{E}(t) \} \\ \llbracket descendant \rrbracket_t(S) & = \bigcup_{\mathbf{i} \in S} \{ \mathbf{i}' \mid (\mathbf{i}, \mathbf{i}') \in \mathcal{E}(t)^+ \} \\ \llbracket ancestor \rrbracket_t(S) & = \bigcup_{\mathbf{i} \in S} \{ \mathbf{i}' \mid (\mathbf{i}', \mathbf{i}) \in \mathcal{E}(t)^+ \} \end{aligned}$$

where  $\mathcal{E}(t)$  is the edge relation of  $t$ , that is,  $\mathcal{E}(t) = \{ (\mathbf{i}, \mathbf{i}') \mid t @ \mathbf{i} = l_i[f, t', f'] \wedge RootId(t') = \mathbf{i}' \}$ , and  $\mathcal{E}(t)^+$  is its transitive closure.

**DEFINITION 3.3 (SIMPLE PATH SEMANTICS).** *Given  $t$ , a set  $S \subseteq Ids(t)$  and a path *SPath*, we define the evaluation of path *SPath* over  $S$  nodes as follows:*

$$\begin{aligned} \llbracket Axis :: Test \rrbracket_t(S) & = (\llbracket Axis \rrbracket_t(S)) ::_t Test \\ \llbracket SPath_1 / SPath_2 \rrbracket_t(S) & = \llbracket SPath_2 \rrbracket_t(\llbracket SPath_1 \rrbracket_t(S)) \end{aligned}$$

### 3.2 Predicates

XPath queries use predicates to express some filtering conditions that cannot be expressed by simple paths. Predicates mix *structural conditions* (directly expressed by means of paths) with *non-structural conditions* (expressed by functions, operators, values, etc.).

We have seen an example of a non-structural condition in the query *Q* extracting all book titles of books written by Dante, defined at the beginning of the section. The best pruning for the *Q* query is the one that deletes all books whose authors do not include Dante. To implement such a pruning, one should extract from the query value-based conditions (e.g. being equal to “Dante”). This would drastically complicate the treatment without bringing a significant gain: previous experiments have shown that navigational

specifications are already sufficient to obtain important improvements in memory reduction and query execution time [14]. Hence we’d rather abstract out non-structural conditions and only retain structural ones. More precisely, our analysis will have to work only on conditions defined as follows:

$$Cond ::= SPath \mid Cond \text{ or } Cond$$

XPath<sup>ℓ</sup> is then defined by the following grammar:

$$Path ::= Step \mid Step[Cond] \mid Path/Path$$

We will use meta-variables *Path* and *P* to range over these paths, and reserve *SPath* for simple paths and *Q* for general XPath queries. Note that the definition of *Cond* uses simple paths, therefore in XPath<sup>ℓ</sup> conditions are not nested.

Semantics of XPath<sup>ℓ</sup>’s paths is defined by substituting in Definition 3.3 *Path* for *SPath* and by adding the following cases

$$\begin{aligned} \llbracket self :: node[C] \rrbracket_t(S) & = \{ \mathbf{i} \in S \mid Check_t[C](\mathbf{i}) \} \\ \llbracket Axis :: Test[C] \rrbracket_t(S) & = \llbracket Axis :: Test/self :: node[C] \rrbracket_t(S) \end{aligned}$$

where  $Check_t[Cond](\mathbf{i})$  is the following boolean function:

$$\begin{aligned} Check_t[Path](\mathbf{i}) & = \llbracket Path \rrbracket_t(\{ \mathbf{i} \}) \neq \emptyset \\ Check_t[C_1 \text{ or } C_2](\mathbf{i}) & = Check_t[C_1](\mathbf{i}) \vee Check_t[C_2](\mathbf{i}) \end{aligned}$$

### 3.3 Handling XPath predicates

The predicates of the previous section cover only a small part of XPath. If we want to apply our analysis to XPath and XQuery we must be able to deal with the more general expressions used in conditions.

In this section we show how to rewrite every predicate *Exp* expressible in XPath to a simple condition *Cond* such that *Cond* is a sound approximation of *Exp* with respect to data needs: the pruning determined for *Cond* preserves the semantics for *Exp*. In other words, if we take a generic XPath query *Q* and approximate all its predicates to infer a projector  $\pi$ , then the execution of (the original) *Q* on a given document or on the document pruned by  $\pi$  yield the same result. This rewriting, together with the treatment of missing axes of Section 4.3, allows us to deal with a large subset of XQuery and XPath queries, covering those in XPathMark [12] and XMark [17] benchmarks.

More formally, we show how to rewrite an expression *Exp* into a condition *Cond*, where *Exp* is defined as

$$Exp ::= Q \mid Exp \text{ op } Exp \mid f(Exp_1, \dots, Exp_n) \mid AExp$$

where *op* ∈ {eq, ne, lt, le, gt, ge, =, !=, <, <=, >, >=, is, <<, >>, or, and} is an operator, *AExp* ranges over arithmetic expressions (see [2]) and base values (PCDATA), *f* ranges over XPath and XQuery functions and operators [5] such as **count**, **contains**, **is-zero**, **not**, **empty**, etc., and *Q* is a generic XPath query, that is:

$$Q ::= Step \mid Step[Exp] \mid Step/Q \mid Step[Exp]/Q$$

The rewriting is obtained by a path-extracting function *P* that applied to an expression *Exp* returns a set of *simple* paths whose “or” constitutes the approximation of *Exp*.<sup>3</sup>

Let us outline the rewriting by an example. Consider the predicate  $\llbracket position() > 1 \text{ and } parent :: node/book/author = \text{“Dante”} \rrbracket$

<sup>3</sup>For lack of space we cannot present the full treatment of predicates that we have implemented in our prototype. In particular, we do not consider absolute paths (although they need special treatment they do not introduce any significant problem) nor we formally define the approximation for each XPath and XQuery function.

and `year>1313`]. In our system this predicate is approximated by `[ self::node or parent::node/book/author or year ]`. Essentially, given a predicate `Exp` we obtain a condition `Cond` that soundly approximates it by retaining the disjunction of all structural conditions (like `parent::node/book/author` and `year` in the previous example), plus either `descendant-or-self::node` or `self::node` if some non-structural condition is present (for instance, `position(>1)`). The choice between `self::node` and `descendant-or-self::node` depends on the functions and operators used in the condition: for instance functions like `position` or `count` require `self::node` since their execution requires only the root nodes; instead a function such as `string` needs the whole tree. Therefore we suppose to have a predefined function `F` that for each `f` returns either `descendant-or-self::node` or `self::node`. For the sake of generality we suppose that this function depends on the position of the argument in  $n$ -ary function. Thus, for, say, `count(SPath)` and `string(SPath)`, we have  $P(\text{count}(SPath)) = SPath/F(\text{count}, 1) = SPath/\text{self}::\text{node}$ , and  $P(\text{string}(SPath)) = SPath/F(\text{string}, 1) = SPath/\text{descendant-or-self}::\text{node}$ . Formally, we have:

$$\begin{aligned}
P(\text{Step}) &= \{\text{Step}\} \\
P(\text{Step}[Exp]) &= \text{Step}/P(Exp) \\
P(\text{Step}/Q) &= \text{Step}/P(Q) \\
P(\text{Step}[Exp]/Q) &= \text{Step}/(P(Q) \cup P(Exp)) \\
P(Exp \text{ op } Exp') &= P(Exp) \cup P(Exp') \\
P(f(Exp_1, \dots, Exp_n)) &= \bigcup_{i=1..n} (P(Exp_i)/F(f, i)) \cup \{\text{self}::\text{node}\}
\end{aligned}$$

where we used the notation  $Step/A$  as a shorthand to denote the set  $\{Step/SPath \mid SPath \in A\}$  when  $A$  is a set of simple paths (similarly for  $A/Step$ ).

The presence of  $\{\text{self}::\text{node}\}$  in the last line is motivated by the fact that when we have a non structural condition, paths must not be used to restrict the inferred projectors, since this would not yield a sound approximation. More precisely, when `Exp` is purely structural, that is it only involves paths in (possibly nested) conditions, then these paths are extracted to refine the projection. For instance, in `descendant::node[child::a]` we can use the condition `[child::a]` to refine projection inference: we select only element types having an `a` child. On the other hand, when `Exp` is not purely structural, as in `descendant::node[not(child::a)]` or `descendant::node[count(child::a)<5]`, we can not use the same projector as for `descendant::node[child::a]`: if we use `[child::a]` to restrict the projection, we would alter the result of the last two queries, so the projector would be unsound. To guarantee soundness, we extract paths from the arguments `not` and `count` and add the condition  $\{\text{self}::\text{node}\}$  to ensure that we do not prune nodes necessary to the evaluation of the functions. So, for the two queries, after condition rewriting, we have the approximating query `descendant::node[child::a or self::node]`, yielding a sound projector.

To resume, to indicate the fact that, in the presence of not purely structural conditions, paths must not be used to restrict inferred projectors, we add the always true condition  $\{\text{self}::\text{node}\}$ . Of course, we could have adopted more precise (and complex) techniques, but we preferred this solution as we consider it a good compromise between precision and simplicity.

We want also to stress that here we reach the limits of XQuery and XPath type systems. If we had worked on more advanced XML languages such as CDuce [6] or CQL [7] their richer type system (it includes union, intersection, negation, and singleton types) would allow us to precisely capture more predicates and use them for a much finer pruning (as it is done in CQL query optimisation).

## 4. STATIC ANALYSIS

In this section we define deduction rules to statically infer from a XPath<sup>ℓ</sup> path  $P$  and a DTD  $E$  a type-projector for an input document validating  $E$ . We show that the analysis is sound, and that it enjoys completeness for a large class of queries when  $E$  is a \*-guarded and non-recursive DTD (see Definition 4.3 below). Soundness means that executing the query on the original document and on the document pruned by the inferred projector yields the same result. Completeness means that if we take a type projector smaller (i.e., more selective) than the inferred one, then there exists a document validating  $E$  for which the result of the two executions is not the same. When the conditions on DTDs or on queries are relaxed the analysis is still sound but it may be not complete. Nevertheless, as we will illustrate, it still is very precise.

In order to define our static type inference we proceed in two steps.

1. Given a path  $P$  and a DTD  $E$  we type  $P$  by the set of all elements that may appear in the result of applying  $P$  to a document validating  $E$ . This is done in Section 4.1 (actually, we will be more precise and type  $P$  by the set of all names of  $E$  that generate the elements in the result).
2. We use the type inference at the previous point to define the inference of type projectors. In particular we will use the cases in which the previous type inference returns the empty set to determine the points in which pruning must be performed. This is done in Section 4.2.

### 4.1 Type inference

Given a path  $Path$  and a DTD  $E$  we want to find a set of names of  $E$  that generates elements that can be found in the result of  $P$ . Formally, we want to infer a set  $\tau \subseteq DN(E)$  such that

$$\forall t \in \mathfrak{S} E. \mathfrak{S}(\llbracket Path \rrbracket_t(\text{RootId}(t))) \subseteq \tau \quad (1)$$

which states the soundness of the analysis.

Moreover, we aim at an analysis which is precise enough to guarantee, on a large class of types and for a large class of queries, that whenever the path semantics is empty over all possible instances of the input DTD, then the inferred type  $\tau$  is empty, as well:

$$\forall t \in \mathfrak{S} E. \mathfrak{S}(\llbracket Path \rrbracket_t(\text{RootId}(t))) = \emptyset \Rightarrow \tau = \emptyset \quad (2)$$

(the converse is a consequence of (1)). The precision described by (2) will then be used during the inference of type-projectors to discard elements that are useless in the evaluation of  $Path$ .

We start by inferring types for single-step paths.

DEFINITION 4.1 (SINGLE STEP TYPING). *Let  $E$  be a DTD and  $\tau \subseteq DN(E)$ , then:*

$$\begin{aligned}
A_E(\tau, \text{ancestor}) &= \bigcup_{Y \in \tau} \{Z \mid Z \Rightarrow_E^+ Y\} \\
A_E(\tau, \text{child}) &= \bigcup_{Y \in \tau} \{Z \mid Y \Rightarrow_E Z\} \\
A_E(\tau, \text{parent}) &= \bigcup_{Y \in \tau} \{Z \mid Z \Rightarrow_E Y\} \\
A_E(\tau, \text{descendant}) &= \bigcup_{Y \in \tau} \{Z \mid Y \Rightarrow_E^+ Z\} \\
A_E(\tau, \text{self}) &= \tau \\
T_E(\tau, a) &= \{Y \mid Y \in \tau, E(Y) = a[r]\} \\
T_E(\tau, \text{node}) &= \tau \\
T_E(\tau, \text{text}) &= \{Y \mid Y \in \tau, E(Y) = \text{String}\}
\end{aligned}$$

The type of a single step query  $Axis::Test$  for the DTD  $(X, E)$  is then given by  $T_E(A_E(\{X\}, Axis), Test)$ . Soundness of this definition, i.e. property (1), is given by the following lemma.

LEMMA 4.2. *Let  $t$  be a tree  $\mathfrak{S}$ -valid with respect to the DTD  $E$ . For every  $S \subseteq Ids(t)$  and type  $\tau$ , if  $\mathfrak{S}(S) \subseteq \tau$  then*

1.  $\mathfrak{S}(\llbracket Axis \rrbracket_t(S)) \subseteq A_E(\tau, Axis)$
2.  $\mathfrak{S}(S::t \text{ Test}) \subseteq T_E(\tau, Test)$

<b>Primitive Single Step</b>	$\frac{}{\Sigma \vdash_E \text{Axis} :: \text{node} : (\mathbf{A}_E(\Sigma_\tau, \text{Axis}), \Sigma_\kappa \cup \mathbf{A}_E(\Sigma_\tau, \text{Axis}))} \text{Axis} \in \{\text{self}, \text{child}, \text{descendant}\}$ $\frac{}{\Sigma \vdash_E \text{Axis} :: \text{node} : (\mathbf{A}_E(\Sigma_\tau, \text{Axis})) \cap \Sigma_\kappa, \mathbf{A}_E(\Sigma_\kappa, \text{Axis}) \cap \Sigma_\kappa} \text{Axis} \in \{\text{parent}, \text{ancestor}\}$ $\frac{}{\Sigma \vdash_E \text{self} :: \text{Test} : (\mathbf{T}_E(\Sigma_\tau, \text{Test}), (\Sigma_\kappa \cap \mathbf{A}_E(\mathbf{T}_E(\Sigma_\tau, \text{Test}), \text{ancestor})) \cup \mathbf{T}_E(\Sigma_\tau, \text{Test}))} \text{Test} \neq \text{node}$ $\frac{\forall X_i \in \Sigma_\tau, P_j \in \text{Cond}, (\{X_i\}, \Sigma_\kappa) \vdash_E P_j : \Sigma^{ij}}{\Sigma \vdash_E \text{self} :: \text{node}[\text{Cond}] : (\tau, (\Sigma_\kappa \cap \mathbf{A}_E(\tau, \text{ancestor})) \cup \tau)} \tau = \{X_i \mid \exists j. \Sigma_\tau^{ij} \neq \emptyset\}$
<b>Encoded Single Step</b>	$\frac{\Sigma \vdash_E \text{Axis} :: \text{node}/\text{self} :: \text{Test} : \Sigma'}{\Sigma \vdash_E \text{Axis} :: \text{Test} : \Sigma'} \begin{array}{l} \text{Test} \neq \text{node} \\ \wedge \\ \text{Axis} \neq \text{self} \end{array}$ $\frac{\Sigma \vdash_E \text{Axis} :: \text{Test}/\text{self} :: \text{node}[\text{Cond}] : \Sigma'}{\Sigma \vdash_E \text{Axis} :: \text{Test}[\text{Cond}] : \Sigma'} \begin{array}{l} \text{Test} \neq \text{node} \\ \wedge \\ \text{Axis} \neq \text{self} \end{array}$
<b>Composed paths</b>	$\frac{\Sigma \vdash_E \text{Step} : \Sigma'' \quad \Sigma'' \vdash_E \text{Path} : \Sigma'}{\Sigma \vdash_E \text{Step}/\text{Path} : \Sigma'}$

**Figure 1: Inference rules for single step queries**

The presence of upward axes makes the typing of composed paths much more difficult. To ensure precision, i.e. property (2), we have to be careful in dealing with DTDs in which an element may occur in the content of different elements. The naive solution consisting of inferring a type for composed paths by composing the functions we just defined for single steps, works only in the absence of upward axes. This can be illustrated by an example. Consider the following DTD rooted at  $X$ :

$$\{X \rightarrow c[Y, Z], Y \rightarrow a[W, \text{String}], Z \rightarrow b[\text{String}], W \rightarrow d[Y?]\}$$

and observe that  $Y$  occurs in two different element content definitions. If we consider the path  $\text{self} :: c/\text{child} :: a/\text{parent} :: \text{node}$  over documents of the above DTD, then the precise type that this path should have is  $\{X\}$ . However, by using Definition 4.1 we end up with  $\{X, W\}$ . This is because the first step selects  $\{Y\}$  and then, according to Definition 4.1, the second step selects  $\{X, W\}$ , as  $Y$  is in the content definition of these two names.

To solve this problem we introduce particular types, called *contexts*, to be updated at each step and containing names already encountered in previous steps. We then use them to refine type inference for upward axes. In the previous example, when typing the first step we build a *context*  $\{X, Y\}$  indicating that for the moment the two names are the only ones visited by the traversal. Then, we use Definition 4.1 to type  $\text{parent}$  thus obtaining  $\{X, W\}$ , as before, but this time we intersect it with the context thus obtaining the precise answer  $\{X\}$ .

This idea is formalised by the (deterministic) type system of Figure 1. We use the meta-variables  $\tau$  to range over types and  $\kappa$  over contexts, both denoting sets of names defined by the input DTD  $E$ . An environment, ranged over by  $\Sigma$ , is a pair  $(\tau, \kappa)$ ; we use  $\Sigma_\tau$  and  $\Sigma_\kappa$  to denote the first and second projection of  $\Sigma$ , respectively.

**Environments**  $\Sigma ::= (\tau, \kappa)$

**Judgements**  $J ::= \Sigma \vdash_E \text{Path} : \Sigma'$

The judgement  $(\tau_c, \kappa_c) \vdash_E \text{Path} : (\tau_r, \kappa_r)$  means that given a DTD  $E$ , starting from the names in  $\tau_c$  and the current context  $\kappa_c$ , the path  $\text{Path}$  generates the names  $\tau_r$  in an updated context  $\kappa_r$ .

An environment  $(\tau, \kappa)$  is well-formed with respect to  $E$ , if  $\tau \subseteq DN(E)$ , and  $\kappa \subseteq \tau \cup \mathbf{A}_E(\tau, \text{ancestor})$ , that is, if the context contains only names that occur in chains ending with names in  $\tau$ . A judgement  $\Sigma \vdash_E \text{Path} : \Sigma'$  is well formed if both  $\Sigma$  and  $\Sigma'$  are well formed with respect to  $E$ . It is easy to see that the type inference rules of Figure 1 preserve well-formedness.

The rules are relatively simple to understand. The first two rules implement our main idea: when we follow an axis  $\text{Axis}$ , we compute the type by  $\mathbf{A}_E(\Sigma_\tau, \text{Axis})$ ; if the axis is a downward one, then we add this type to the current context, otherwise if the axis is an upward one, then we intersect it with the current context (both for the type part and for the context part). The rule for  $\text{self} :: \text{Test}$  is slightly more difficult since it discards from the current set of nodes those that do not satisfy the test: the type is computed by  $\mathbf{T}_E(\Sigma_\tau, \text{Test})$ , while the context is obtained by erasing all the names that were in there just because they generated one of the discarded nodes; to do it it generates (the type of) all ancestors of the nodes satisfying the test, and intersects them with the current context. These first three rules are enough to type all the paths of the form  $\text{Axis} :: \text{Test}$  since, as stated by the fifth typing rule, all remaining cases are encoded as  $\text{Axis} :: \text{node}/\text{self} :: \text{Test}$ . The fourth rule is the most difficult one: recall that  $\text{Cond}$  is a disjunction of *simple* paths; the type  $\tau$  is obtained by discarding from  $\Sigma_\tau$  all (names of) nodes for which  $\text{Cond}$  never holds; thus for each  $X_i$  in  $\Sigma_\tau$  we compute the type of all the paths in  $\text{Cond}$ , and keep in  $\tau$  only names for which at least one path may yield a non-empty result; the context then is computed as in the third rule, by discarding from the context all names that generated only names discarded from  $\Sigma_\tau$ . Once more, all the remaining cases of conditional steps are encoded by this one, as stated by the sixth rule. Finally, step composition is dealt as a logical cut.

The type system is sound. It is also complete for DTDs that are \*-guarded, non-recursive, and parent-unambiguous. Intuitively, a DTD is \*-guarded when every union occurring in its productions is guarded by \* (or by +), it is non recursive if the depth of all documents validating it is bound, while it is parent-unambiguous if no name types both the parent and a strict ancestor of the parent of another name. Formally, we have the following definition

DEFINITION 4.3. *Let  $(X, E)$  be a DTD.*

1.  *$E$  is \*-guarded if for each  $Y \rightarrow l[r]$  in  $E$ , the regular expression is a product  $r = r_1, \dots, r_n$  and whenever  $r_i$  contains a union, then  $r_i = (r')^*$ ;*
2.  *$E$  is non-recursive if it is never the case that  $Y \Rightarrow_E^+ Y$ , for any name  $Y \in DN(E)$ ;*
3.  *$E$  is parent-unambiguous if for all chains  $c$  and names  $Y, Z$  such that  $cYZ \in \text{Chains}_{(X, E)}(X)$  the following implication  $cYc'Z \in \text{Chains}_{(X, E)}(X) \implies c' = \varepsilon$  holds ( $\varepsilon$  denotes the empty chain).*

Non-recursive and \*-guardedness are properties enjoyed by a large number of commonly used DTDs. As an example, the reader can consider the DTDs of the XML Query Use Cases [3]: among the ten DTDs defined in the Use Cases, seven are both non-recursive and \*-guarded, one is only \*-guarded, one is only non-recursive, and just one does not satisfy either property. Furthermore our personal experience is that most of the DTDs available on the web are \*-guarded. Concerning the parent-unambiguous property, although DTDs satisfying this property are less frequent (five on the ten DTDs in [3]), its absence is in practice not very problematic since, as we will see, only the presence of the `parent` axis may hinder completeness.

THEOREM 4.4 (SOUNDNESS AND COMPLETENESS). *Let  $(X, E)$  be a DTD and  $P$  a path. If  $(\{X\}, \{X\}) \vdash_E P : (\tau, \kappa)$  then (soundness):*

$$\tau \supseteq \bigcup_{t \in \mathfrak{S}E} \mathfrak{S}(\llbracket P \rrbracket_t(\text{RootId}(t)))$$

*Furthermore, if  $(X, E)$  is \*-guarded and non-recursive, and parent-unambiguous, then we also have (completeness):*

$$\tau \subseteq \bigcup_{t \in \mathfrak{S}E} \mathfrak{S}(\llbracket P \rrbracket_t(\text{RootId}(t)))$$

To see why completeness does not hold in general consider the following DTD rooted at  $X$  and which is recursive and not \*-guarded

$$\{X \rightarrow c[Y \mid Z], Y \rightarrow a[Y^*, \text{String}], Z \rightarrow b[\text{String}]\}$$

and the following two queries `self :: c[child :: a]/child :: b` and `self :: c/child :: a/parent :: node`. The type inferred for the first query contains both  $Y$  and  $Z$ . These are useless since the query is always empty. This is due to the non \*-guarded union  $Y \mid Z$ : if we had  $(Y \mid Z)^*$  instead, then the query might yield a non-empty result, therefore  $Y$  and  $Z$  must correctly (and completely) be in the query type. The second query shows the reason why completeness does not hold in presence of recursion and backward axes (recursion with only forward axes does not pose any problem for completeness). The type of the second query should be  $\{X\}$ , but instead the type  $\{X, Y\}$  is inferred. This is due to the recursion  $Y \rightarrow a[Y^*, \dots]$ : since  $Y \Rightarrow_E Y$ , once  $Y$  is reached it is kept in the inferred type for every backward step.<sup>4</sup>

For queries over parent-ambiguous DTDs, completeness does not hold because the fourth rule in Figure 1—the one defined for `self` ::

<sup>4</sup>The techniques developed in [11, 10] can be adapted to recover completeness for cases like the first query, while a more sophisticated type analysis could solve the problem with the second. In view of the precision of the current approach this is not a priority and we leave this investigation as future work.

`node[Cond]`—is not precise for the parent axis. For instance, consider the following DTD rooted at  $X$

$$\{X \rightarrow a[Y, Z], Y \rightarrow b[Z], Z \rightarrow c[\ ]\}$$

and the query `self :: a/child :: b/child :: c/parent :: node`. The precise type of this query should be  $\{Y\}$ . However, the inferred type is  $\{X, Y\}$ . This is because the last step `parent :: node` is typed with the context  $\{X, Y, Z\}$  and this contains  $\mathbf{A}_E(\{Z\}, \text{parent}) = \{X, Y\}$ . Here  $Z$  is the type for the  $c$  node selected by `child :: c` and the  $\mathbf{A}_E(\cdot)$  operator assigns it  $\{X, Y\}$  as parent type, even if the real parent type for  $Z$  in this case should be  $\{Y\}$ . Hence, the intersections operated by the type rule for `parent` are not powerful enough to guarantee precision for cases like this one. In a nutshell, this happens because in the presence of parent-ambiguous DTDs the type analysis may produce contexts containing false parent types (with respect the current type  $\tau$ ). This suggests that to be extremely precise, instead of sets of names, contexts should rather be sets of chains of names, computed and opportunely managed by the type analysis. However (i) managing sets of chains instead of simple sets of names dramatically complicates the treatment, due to recursive axes like `descendant`, (ii) the problem may arise only for queries that use parent axis and the concomitance of parent-ambiguity make the event rare in practice, and (iii) the loss of precision looks in most cases negligible. Therefore we considered that such a small gain (remember that completeness is just some icing on the cake since while it helps to gauge the precision of the approach its absence does not hinder its application) did not justify the dramatic increase in complexity needed to handle this case.

Note also that the type system, hence the completeness result, is stated for predicates of the form described in Section 3.2, therefore it does not account for the approximations introduced in Section 3.3. However very few non-structural conditions can be expressed at the level of types, so the impact of these approximations on completeness is very light.

## 4.2 Type-Projection inference

In this section we use the type inference of the previous section to infer type-projectors. Once more naive solutions do not work. For instance, for simple paths  $\text{Step}_1 / \dots / \text{Step}_n$ , we may consider as type projector with respect to  $(X, E)$  the set  $\bigcup_{i=1 \dots n} \tau_i \cup \{X\}$ , where for  $i = 1 \dots n$ :

$$(\{X\}, \{X\}) \vdash_E \text{Step}_1 / \dots / \text{Step}_i : (\tau_i, -)$$

(we use “ $-$ ” as a placeholder for uninteresting parameters). This definition is sound but not precise at all, as can be seen by considering `descendant :: node/Path`: the use of the above union yields a set containing  $\tau_1$  defined as

$$(\{X\}, \{X\}) \vdash_E \text{descendant :: node} : (\tau_1, -)$$

that is, all descendants of the root  $X$  (no pruning is performed). Instead, we would like to discard, at least, all names that are descendants of  $X$  but that are not ancestors of a node matching `Path`. These are the names  $Y \in \mathbf{T}_E(\mathbf{A}_E(\{X\}, \text{descendant}), \text{node})$  such that

$$(\{Y\}, \kappa) \vdash_E \text{descendant :: node/Path} : (\emptyset, -)$$

for some appropriate context  $\kappa$ . A similar reasoning applies to `ancestor`.

Such a selection is performed by the inference rules of Figure 2. For paths formed by a single step, if the step has no condition (first rule), then the type inference of the previous section is enough; otherwise (second rule) the step is transformed into a complex path (a simple trick to avoid the definition of several rules). Thanks to the third rule the type inference can work on just one node at a time, and thanks to the fourth and fifth rule, it just analyses paths whose

<b>Base and induction</b>
$\frac{\Sigma \Vdash_E \text{Step} : (\tau, \kappa)}{\Sigma \Vdash_E \text{Step} : \tau \cup \kappa} \quad \frac{\Sigma \Vdash_E \text{Step}[\text{Cond}] / \text{self} :: \text{node} : \tau}{\Sigma \Vdash_E \text{Step}[\text{Cond}] : \tau} \quad \frac{(\{X_1\}, \kappa) \Vdash_E P : \tau_1 \quad \dots \quad (\{X_n\}, \kappa) \Vdash_E P : \tau_n}{(\{X_1, \dots, X_n\}, \kappa) \Vdash_E P : \bigcup_{i=1..n} \tau_i} \text{ if no other rule applies}$
<b>Encoded Rules</b>
$\frac{\Sigma \Vdash_E \text{Axis} :: \text{node} / \text{self} :: \text{Test} / P : \tau}{\Sigma \Vdash_E \text{Axis} :: \text{Test} / P : \tau} \quad \text{Test} \neq \text{node} \quad \text{Axis} \neq \text{self} \quad \frac{\Sigma \Vdash_E \text{Axis} :: \text{Test} / \text{self} :: \text{node}[\text{Cond}] / P : \tau}{\Sigma \Vdash_E \text{Axis} :: \text{Test}[\text{Cond}] / P : \tau} \quad \text{Test} \neq \text{node} \quad \text{Axis} \neq \text{self}$
<b>Primitive Rules</b>
$\frac{(\{Y\}, \kappa) \Vdash_E \text{self} :: \text{Test} : \Sigma \quad \Sigma \Vdash_E P : \tau}{(\{Y\}, \kappa) \Vdash_E \text{self} :: \text{Test} / P : \{Y\} \cup \tau} \quad \frac{(\{Y\}, \kappa) \Vdash_E \text{self} :: \text{node}[P_1 \text{ or } \dots \text{ or } P_n] : \Sigma \quad \Sigma \Vdash_E P : \tau \quad \Sigma \Vdash_E P_i : \tau_i}{(\{Y\}, \kappa) \Vdash_E \text{self} :: \text{node}[P_1 \text{ or } \dots \text{ or } P_n] / P : \{Y\} \cup \tau \cup \tau_1 \cup \dots \cup \tau_n} \quad n \geq 1$
$\frac{(\{Y\}, \kappa) \Vdash_E \text{Axis} :: \text{node} : (\{X_1, \dots, X_n\}, \kappa') \quad (\{X_i\}, \kappa') \Vdash_E P : \Sigma^i \quad (\tau, \kappa') \Vdash_E P : \tau'}{(\{Y\}, \kappa) \Vdash_E \text{Axis} :: \text{node} / P : \{Y\} \cup \tau \cup \tau'} \quad \text{Axis} \in \{\text{parent}, \text{child}\} \quad \tau = \{X_i \mid \Sigma_i^i \neq \emptyset\}$
$\frac{(\{Y\}, \kappa) \Vdash_E \text{desc} :: \text{node} : (\{X_1, \dots, X_n\}, \kappa') \quad (\{X_i\}, \kappa') \Vdash_E \text{desc} :: \text{node} / P : \Sigma^i \quad (\tau, \kappa') \Vdash_E \text{child} :: \text{node} / P : \tau'}{(\{Y\}, \kappa) \Vdash_E \text{desc} :: \text{node} / P : \tau \cup \tau'} \quad \tau = \{X_i \mid \Sigma_i^i \neq \emptyset\} \cup \{Y\}$
$\frac{(\{Y\}, \kappa) \Vdash_E \text{ancs} :: \text{node} : (\{X_1, \dots, X_n\}, \kappa') \quad (\{X_i\}, \kappa') \Vdash_E \text{ancs} :: \text{node} / P : \Sigma^i \quad (\tau, \kappa') \Vdash_E \text{parent} :: \text{node} / P : \tau'}{(\{Y\}, \kappa) \Vdash_E \text{ancs} :: \text{node} / P : \tau \cup \tau'} \quad \tau = \{X_i \mid \Sigma_i^i \neq \emptyset\} \cup \{Y\}$

**Figure 2: Projectors inference rules (where ancs and desc are shorthands for ancestor and descendant)**

components have one of the following three forms: (i)  $\text{self}::\text{Test}$ , (ii)  $\text{self}::\text{node}[\text{Cond}]$ , or (iii)  $\text{Axis}::\text{node}$ . These three cases are handled by the ‘‘Primitive Rules’’ of Figure 2: The first rule handles the case (i) simply by collecting the current context. The second rule handles the case (ii), by collecting besides the context also all the parts that are necessary to compute the condition (which in the rule is expanded in its more general form); the case (iii) is handled by the last three rules which are nothing but slight variations of the same rule according to the particular axis taken into account: each rule infers the type  $\tau$  obtained by discarding from the type  $\{X_1, \dots, X_n\}$  of the step, all names that are useless for the rest of the path, and then uses this  $\tau$  to continue the inference of the projector.

**THEOREM 4.5 (SOUNDNESS OF PROJECTOR INFERENCE).**

Let  $(X, E)$  be a DTD and  $P$  a path. If  $(\{X\}, \{X\}) \Vdash_E P : \tau$ , then  $\tau$  is a type projector for  $(X, E)$  and for every  $t \in \mathfrak{S} E$

$$\llbracket P \rrbracket_{t \setminus \mathfrak{S} \tau}(\text{RootId}(t)) = \llbracket P \rrbracket_t(\text{RootId}(t))$$

The above theorem states that executing the query  $P$  on a tree  $t$  returns the same set of nodes as executing it on  $t \setminus \mathfrak{S} \tau$  the tree  $t$  pruned by the inferred projector. From a practical perspective it is important to notice that according to standard XPath semantics, the semantics of a query contains *only* the nodes of the result of the query not their sub-trees. The latter may thus be pruned by the inferred projector. Therefore, if we want to *materialise* the result of a query we must not cut these nodes, and rather use the projection  $\tau = \tau' \cup \mathbf{A}_E(\tau', \text{descendant})$  where  $(\{X\}, \{X\}) \Vdash_E P : \tau'$  and  $(\{X\}, \{X\}) \Vdash_E P : (\tau', -)$ .

Completeness requires not only completeness of the type system (thus, \*-guarded, non-recursive, and parent-unambiguous DTDs), but also the following condition on queries:

**DEFINITION 4.6.** An XPath query  $Q$  is strongly-specified if (i) its predicates do not use backward axes, (ii) along  $Q$  and along each path in the predicates of  $Q$  there are no two consecutive (possibly conditional) steps whose Test part is node, and (iii) each predicate in  $Q$  contains at most one path and this does not terminate by a step whose Test is node.

For instance, among the following queries, only the first two are strongly-specified.

1.  $\text{descendant}::\text{node}/\text{self}::\text{a}/\text{ancestor}::\text{node}$
2.  $\text{descendant}::\text{node}[\text{child}::\text{b}]/\text{self}::\text{a}/\text{parent}::\text{node}$
3.  $\text{descendant}::\text{node}/\text{ancestor}::\text{node}/\text{self}::\text{a}$
4.  $\text{descendant}::\text{node}[\text{child}::\text{b}/\text{child}::\text{node}]/\text{self}::\text{a}$
4.  $\text{child}::\text{a}[\text{descendant}::\text{node}/\text{parent}::\text{b}]/\text{child}::\text{c}$

Once more, we are in presence of a very common class of queries: for instance, almost all paths in the XMark and XPathMark benchmarks are strongly specified.

**THEOREM 4.7 (COMPLETENESS OF PROJECTOR INFERENCE).**

Let  $(X, E)$  be a \*-guarded, non-recursive, and parent-unambiguous DTD, and  $P$  a strongly-specified path. If  $(\{X\}, \{X\}) \Vdash_E P : \tau$ , then there exists  $t \in \mathfrak{S} E$  such that for each  $Y \in \tau$ , if  $\pi = \tau \setminus (\{Y\} \cup \mathbf{A}_E(\{Y\}, \text{descendant}))$ , then

$$\llbracket P \rrbracket_{t \setminus \mathfrak{S} \pi}(\text{RootId}(t)) \neq \llbracket P \rrbracket_t(\text{RootId}(t))$$

The fact that completeness may not hold for not \*-guarded, non-recursive, or parent-ambiguous DTDs, is a consequence of the analogous property of the type system. To see that also strong-specification is a necessary condition consider documents valid with respect to the following DTD rooted at  $X$ :

$$\{X \rightarrow a[Y, W], W \rightarrow c[], Y \rightarrow b[Z], Z \rightarrow d[]\}.$$



Query them by the following query which not strongly-specified since it does not satisfy condition (ii) of Definition 4.6

`self :: a[child :: node].`

$\{X, Y\}$  is an optimal projector for this query, but the presence of the condition `self :: node` makes the system to include also  $W$  in the inferred projector, thus breaking completeness. Concerning the presence of backward axes in predicates, consider the query `self :: a[descendant :: node/ancestor :: a]` which does not satisfy condition (i). An optimal projector for this query on the same DTD is  $\{X, Y\}$ . However, since the `ancestor` condition is true for all descendants of  $a$  nodes,  $\{W, Z\}$  is included in the projector as well. Finally, it is straightforward to check that the query `self :: a[child :: b or child :: c]`, which does not satisfy condition (iii), is not complete for the same DTD.

Of course, it is possible to state completeness for other classes of queries but, once more, this seems an excellent compromise between simplicity and generality.

**THEOREM 4.8 (DECIDABILITY).** *Given a path  $P$ , a DTD  $E$ , and an environment  $\Sigma$  well-formed with respect to  $E$ , the inference of a context  $\Sigma'$  and a type  $\tau$  such that  $\Sigma \vdash_E P : \Sigma'$  and  $\Sigma \Vdash_E P : \tau$  is decidable.*

### 4.3 Adding sibling, preceding and following axes.

We could deal with the missing XPath axes by adding specific inference rules. Instead we opt to use an approximation of these axes in term of the previous ones, since it appears as the best compromise between simplicity and efficiency.

The approximation is performed by two logical rewriting passes. In the first pass we rewrite preceding and following axes as specified in the W3C specifications [4]. Namely, we substitute each step `Axis :: Test` with `Axis ∈ {preceding, following}` by the following equivalent path `ancestor-or-self :: node/(Axis-sibling) :: node/descendant-or-self :: Test`

The second pass is the one which introduces the approximation since it replaces all steps of the form `Axis::Test` with `Axis ∈ {preceding-sibling, following-sibling}` by the path `parent::node/child::Test`.

Clearly, the static analysis of the approximation yields a less precise projection than the one we could obtain by working directly on the original query. However, we still achieve good precision of pruning in practice as we will show in Section 6. For instance, by applying the above rewriting to XPathMark queries Q9 and Q11, we were able to prune a document down to 7.5% of its original size.

## 5. EXTENSION TO XQUERY

In this section we extend the technique to XQuery. More precisely to the FLWR core of XQuery described by the following grammar:

$$\begin{aligned} q ::= & () \mid q, q \mid \langle \text{tag} \rangle q \langle /\text{tag} \rangle \mid \text{Exp} \\ & \mid \text{for } x \text{ in } q \text{ return } q \mid \text{let } x = q \text{ return } q \\ & \mid \text{if } q \text{ then } q \text{ else } q \end{aligned}$$

where the definition of *Exp* (given in Section 3.3) is extended with variables, and with generic XPath expressions  $Q$  of Section 3.3 that can be rooted at a variable or at  $:$

$$\text{Exp} ::= x \mid Q \mid x/Q \mid /Q \mid \text{Exp op Exp} \mid f(\text{Exp}, \dots, \text{Exp}) \mid A\text{Exp}$$

Without loss of generality, we assume that FLWR expressions do not occur in `if`-conditions nor in predicates (every query can be put

into this form by adding appropriate `let`-expressions). Also, we do not consider either queries which first construct new elements and then navigate on them (these are rarely used in practice), nor those containing XQuery clauses like `order_by`, `switch_case`, etc.: our approach can be easily extended to both cases.

In order to apply the previous analysis to infer a projector for  $q$ , we first extract a set of XPath<sup>ℓ</sup> expressions from  $q$ , denoting the data needs for  $q$ . This set of paths is extracted from the query by the extraction function  $E$ , whose definition is given in Figure 3. The extraction function has the form  $E(q, \Gamma, m)$ . The first parameter is the query at issue. The second parameter  $\Gamma$  is an environment that keeps track of bindings of the form  $(x; \text{for } P)$  or  $(x; \text{let } P)$ , whose scope  $q$  is in (see the definition of  $\Gamma'$  in the last two lines of Figure 3, and observe, by a simple induction reasoning, that environments contain paths already in XPath<sup>ℓ</sup>). Finally,  $m$  is a flag indicating whether  $q$  is a query that serves to materialise a partial or final result ( $m = 1$ ), or that just selects a set of nodes whose descendants are not needed ( $m = 0$ ). Thus, the set of path expressions (possibly containing qualifiers) extracted from a top-level query  $q$  is  $E(q, \emptyset, 1)$ .

Once the set of paths are extracted from a query  $q$ , we use it to infer a projector for  $q$  according to rules in Section 4.2. Formally, for each  $P_i$  extracted from  $q$  we deduce a projector  $\pi_i$ , and use for the whole  $q$  the union of these projectors (projectors are closed by union). Also, note that the extracted path of a closed query will not contain free variables since possible free variables are persistent roots that must be solved before the analysis.

Most of the rules in Figure 3 are not difficult to understand, therefore only few of them deserve further commentary. The flag is needed since each path determining the result ( $m = 1$ ) must be extended with `descendant-or-self`, in order to project on all nodes needed in the query result. This is done by the lines 6, 8, and 10 of the definition. Expressions are dealt in a way similar to the path extractor  $P$  of Section 3.3; the extractor  $P$  itself is used in line 12 to produce simple paths (where we used the notation  $\text{or}(\{P_1, \dots, P_n\})$  for  $P_1 \text{ or } \dots \text{ or } P_n$ , and omitted the—straightforward—rules for single step paths). Also note that when a result is computed (lines 2 and 5) paths in “for”-environments are added (“let” are added only if their binding variable is used).

These rules subsume and enhance the whole Marian and Siméon’s technique [14]. In particular, (i) the technique we use to exclude useless intermediate paths is simpler and more compact, (ii) we do not need to distinguish between two kinds of extracted paths but, more simply, we always manage a unique set of path expressions, and (iii) last but not least, our path extractor can be used even if the user cannot access an XQuery to XQuery-Core compiler, which is necessary for [14].

Before applying the extraction function  $E$  to a query  $q$  we apply some heuristics that rewrite  $q$  so to improve the pruning capability of the inferred paths. Among these heuristics the most important is the one that rewrites

```
for y in Q/descendant-or-self::node
  return if C(y) then q else ()
```

into

```
for y in
  Q/descendant-or-self::node[C(self::node)]
  return q
```

whenever  $C(y)$  is a condition referring only to  $y$  and does not use external functions ( $C(\text{self}::\text{node})$  is obtained by replacing `self :: node` for all occurrences of  $y$  free in  $C$ ). If we apply  $E$  to the first query, then a path ending by `descendant-or-self::node` is extracted thus annulling further pruning: the entire forest selected

1.	$E((), \Gamma, m)$	$= \emptyset$
2.	$E(AExp, \Gamma, 1)$	$= \{P \mid (x; \text{for } P) \in \Gamma\}$
3.	$E(AExp, \Gamma, 0)$	$= \emptyset$
4.	$E((q_1, q_2), \Gamma, m)$	$= E(q_1, \Gamma, m) \cup E(q_2, \Gamma, m)$
5.	$E(\langle \text{tag} \rangle q \langle / \text{tag} \rangle, \Gamma, m)$	$= \{P \mid (x; \text{for } P) \in \Gamma\} \cup E(q, \Gamma, 1)$
6.	$E(x, \Gamma, 1)$	$= \{P/\text{descendant-or-self} :: \text{node} \mid (x; -P) \in \Gamma\}$
7.	$E(x, \Gamma, 0)$	$= \{P \mid (x; -P) \in \Gamma\}$
8.	$E(/P, \Gamma, 1)$	$= \{/P/\text{descendant-or-self} :: \text{node}\}$
9.	$E(/P, \Gamma, 0)$	$= \{/P\}$
10.	$E(x/P, \Gamma, 1)$	$= \{P'/P/\text{descendant-or-self} :: \text{node} \mid (x; -P') \in \Gamma\}$
11.	$E(\text{Step}/q, \Gamma, m)$	$= \text{Step}/E(q, \Gamma, m)$
12.	$E(\text{Step}[Exp]/q, \Gamma, m)$	$= \text{Step}[\text{or}(P(Exp))]/E(q, \Gamma, m)$
13.	$E(Exp_1 \text{ op } Exp_2, \Gamma, m)$	$= E(Exp_1, \Gamma, m) \cup E(Exp_2, \Gamma, m)$
14.	$E(f(Exp_1, \dots, Exp_n), \Gamma, m)$	$= \bigcup_{i=1..n} (E(Exp_i, \Gamma, 0)/F(f, i)) \cup \{\text{self} :: \text{node}\}$
15.	$E(\text{if } q \text{ then } q_1 \text{ else } q_2, \Gamma, m)$	$= E(q, \Gamma, 0) \cup E(q_1, \Gamma, 1) \cup E(q_2, \Gamma, 1) \cup \{P \mid (x; -P) \in \Gamma\}$
16.	$E(\text{for } x \text{ in } q_1 \text{ return } q_2, \Gamma, m)$	$= E(q_1, \Gamma, 0) \cup E(q_2, \Gamma \cup \Gamma', m)$ (where $\Gamma' = \{(x; \text{for } P) \mid P \in E(q_1, \Gamma, 0)\}$ )
17.	$E(\text{let } x = q_1 \text{ return } q_2, \Gamma, m)$	$= E(q_1, \Gamma, 0) \cup E(q_2, \Gamma \cup \Gamma', m)$ (where $\Gamma' = \{(x; \text{let } P) \mid P \in E(q_1, \Gamma, 0)\}$ )

Figure 3: XQuery path extraction

by  $Q$  is loaded in main memory. This also happens with the approaches of Bressan *et al.* [9] and of Marian and Siméon [14]. In our and Marian and Siméon's approach the query can be rewritten as above (this is not possible in [9] since their subset of XQuery does not include predicates). However, Marian and Siméon's path based pruning degenerates (no further pruning is performed) also for the second query, since the `descendant-or-self::node` ends up in the set of pruner paths, thus selecting all nodes. This is because their approach cannot manage predicates. In our approach instead predicates are taken into account and therefore only nodes satisfying  $C(y)$  are kept by the projector, thus yielding a very precise pruning.

It is important to stress that despite their specific form the first kind of queries is very common in practice since they are generated from  $XQuery \rightarrow XQuery\text{-Core}$  compilation of a non negligible class of queries (for instance Q13 of the XPathMark) or when rewriting upward axes into downward ones. This latter observation shows that the application of rewriting rules of [15] to extend Marian and Siméon's approach to upward axes is not feasible since the rewriting may completely compromise pruning.

## 6. EXPERIMENTS

We have implemented a complete version of the algorithm defined for full XPath. The code (available at <http://www.lri.fr/~kn>) is written in OCaml, uses the PXP library for parsing XML documents, and its correctness was verified for all tests. After the path extraction of Section 5, it performs the rewriting presented in Sections 3.3 and 4.3, and the static analysis defined in Section 4. The latter is extended to deal with attributes, with the wildcard test `element()`, with `{descendant, ancestor}-or-self` and `{preceding, following}-siblings` axes, and with absolute paths. It also uses a couple of heuristics. One heuristic rewrites the DTD  $E$  so that every name  $Y$  defined as  $Y \rightarrow \text{String}$  occurs exactly once in the right hand side of an edge of  $E$ ; this enhances the precision of pruning by reducing the number of conflicts on the leaves of the tree. The other heuristic keeps track of the depth of elements in the paths in order to improve pruning, especially in presence of recursive DTDs (this latter heuristics could be embedded in the formal treatment, but we preferred to keep it simpler). Pruning is then performed *in streaming* and merely consists of a

one-pass traversal of the document. We also added an optional validation option, that makes it possible to prune the document while validating it. Programs that use an external validator can therefore prune their document without any overhead.

We performed our tests on a GNU/Linux desktop, with 3GHz processor, 512 MB of RAM and a single S-ATA hard-drive, using DTDs, document generator, and queries of XMark and XPathMark (the latter is interesting because its queries use all the available axes). Queries were processed by the latest version of Galax (that is, the 0.5.0). Swap was disabled to test memory limits.

For what concerns the overhead of the optimisation, tests confirmed that it is always negligible, both in memory and time consumption: the only noticeable overhead is pruning time, which is linear in the size of the pruned document, but can be embedded in document parsing and/or validation (e.g., for 60MB documents computing the projector took around 0.5s while pruning and saving the pruned document to disk was always below 10s). These results were confirmed by further experiments on large DTDs (e.g. XHTML) and long XPath expressions (twenty steps or so).

In Table 1 we report part of the results of our tests. For space reasons just a selection of XMark (QM) and XPathMark (QP) queries are presented.


**Projector efficiency.** The fourth line of Table 1 reports the effect of inferred projectors and it is an indicator of the selectivity of the query. For several XMark queries the size of the pruned document is around 70-80% of the size of the original document. This is due to the fact that XMark documents contain mixed-content `<description>` elements which account for about 70% of the total size. Thus, queries whose execution requires the whole content of `<description>` elements, preserve a large part of the file. On the contrary, for very selective queries like QM06, 99.7% of the document is discarded. Finally, for queries that are very little selective, like QP13, the whole document has to be kept. It should be noted in Table 1, fourth line, that for all XMark queries but QM14 we could prune more than 95% of the original document.

**Execution time and memory occupation.** The comparison of performances of the Galax query engine on an original document and its pruned version is given in Figures 4 and 5, which respectively report the processing times and main memory occupation for documents of 56MB. They show that time and memory gains are

	QM03	QM06	QM07	QM14	QM15	QM19	QP01	QP02	QP03	QP04	QP05	QP06	QP07	QP08	QP09	QP10	QP11	QP12	QP13	QP21	QP23
Original Document Size (MB)	930	2048*	1100	202	2048*	964	112	313	258	291	123	190	168	123	459	123	369	134	79	224	403
Pruned Document Size(MB)	25	5.3	42	139	24	24	89	50	46	50	98	133	123	99	35	98	28	107	78	152	42
Main Memory Usage (MB)	374	90	380	512	245	512	391	399	433	434	418	485	467	466	466	483	456	460	504	459	465
Gain in Size (% of original)	2.5	0.3	3.4	69.6	1.15	2.5	80.4	15.7	17.5	16.8	80.4	69.6	73.2	80.4	7.5	80.4	7.5	80.4	98.2	67.9	10.4
Gain in Speed ( $\times$ faster)	17.8	110.1	28.2	3.9	62.6	7.5	1.5	3.6	3.7	4.3	1.5	2.9	2.6	1.1	4.9	1.6	4.2	1.6	1.0	3.6	3.6

\*: biggest file the XMark generator was able to produce.

**Table 1: Sizes (in MBytes) of the biggest document processed thanks to pruning, size of its pruned version, and memory used to process the latter. Percent of the pruned document and speedup of the execution time for a 56MB document.**




**Figure 4: Processing time of a query on original (56MB) and pruned documents**

similar.

These gains translate in practice into much faster executions and the possibility to process much larger documents. The improvement can be measured by looking at the first and last lines of Table 1. The first line reports the size of the largest document it was possible to process thanks to pruning. This must be compared with the fact that, for all queries, the largest document that can be processed without pruning is 68MBytes large. The last line reports how many times the execution on a pruned document is faster than the execution on the original document. It is important to note that, depending on the nature of the query, the gain can be much higher than the proportion given by the percent of the size of the pruning. For instance, for queries such as QM14, QP6, and QP21 the size of the pruned document is two-thirds of the size of the original document, but they can then be processed from three to four times faster and, as Figure 5 shows, using three times less memory than when processed on the original. The latter is a huge gain when one knows that memory usage is one of the main bottlenecks for real life query processing (e.g., in DOM-based implementations of XPath or XSLT processors).

Quite informative, as well, is the data in the second line of Table 1 which reports, for each query, the size in MB of the maximum pruned document. It is interesting to see that, while the maximum size for an unpruned document is 68MB, we can process documents for which the projection has a size of 152MB (on disk). This is due to the fact that projecting a document not only reduces its size but also its *complexity* by reducing the number of types of nodes. This simplification of the document reduces the amount of extra-information the query engine has to keep for each node and, consequently, its memory usage. More precisely, the benefit of pruning



**Figure 5: Memory used to process a query on original (56MB) and pruned documents**

out some (types of) nodes is twofold: first, the fan out of the document is reduced and this may impact memory usage for engines that chase sibling pointers and, second, the number of element names is reduced, which may reduce memory occupation when shredding.

These results are a clear-cut improvement over current technology. While we cannot directly compare processing performances since no implementation of the other pruning approaches is publicly available, we want to stress two points: (i) with one exception (QM14) the amount of pruning on common experiments is always equal or better with our approach than the others and (ii) performing pruning never is a bottleneck in our case thanks to fact that our solution consists of a single bufferless one pass traversal of the input document (on our 512MB machine we were able to efficiently prune arbitrary large documents, while in case of [14] pruning can end up using as much memory as the execution of the query).

## 7. CONCLUSION AND FUTURE WORK

The benchmarks show the clear advantages of applying our optimisation technique to query XML documents, and the characteristics of our solution make it profitable in all application scenarios. We discussed several aspects for which our approach improves the state of the art: for performances (better pruning, more speedup, less memory consumption), for the analysis techniques (linear pruning time, negligible memory and time consumption), for its generality (handling of all axes and of predicates), and, last but not least, for the formal foundation it provides (correctness formally proved, limits of the approach formally stated).

Future work will be pursued in three distinct areas: formal developments, database integration, and implementation issues.

For what concerns the formal treatment, we have to integrate in it the heuristics used in the implementation of the static analysis and to formally state the soundness and completeness of some approximations presented in the work. Also, it should be easy to adapt the approach to work in the absence of DTDs, by using data-guides/path-summaries instead. We intend also to adapt our technique to optimise queries written in CQL [7] the query language of CDuce [6]: as we said at the end of Section 3, their rich type system will allow us to assign more precise types to queries (for instance, it will be possible to capture by types many XPath predicates, since disjunction, conjunctions and negations can be handled by the corresponding type operators and the value of attributes and element contents can be expressed by singleton types) and thus to perform more selective pruning. Finally, we want to modify our approach so that it can yield efficient pruning also in the presence of XPath 2.0 predicates that test the XML Schema of nodes. Note indeed that such predicates are blockers for pruning: we have to leave the entire subtree intact so that the engine can verify that it has the specified schema. But since the projector inference algorithm already statically checks this property, the idea is to make the inference algorithm also rewrite predicates so as to push the schema tests down where they are strictly necessary, thus making further pruning possible.

From a database perspective we want to study the integration of our optimisation technique with classical database ones. Our technique must be viewed as a preliminary step that can be further combined with more traditional database optimisations. More precisely, as our technique is able to take into account the workload, in the line of [8], it could help the database administrator to deduce relevant clustering strategies of XML data on disk and to define well-adapted indexes and/or materialised views. Second, our pruning technique can also be used for pruning indexes. For example, if indexes over element tags are present before query processing (like in the TIMBER system), the index can be pruned as well. In TIMBER, for a 472 MB document, such an index can reach a 241MB size [16], thus it is worth being pruned, in order to improve buffer management and concurrent query evaluations.

Finally, implementation-wise, the natural extension of our work is to interface our pruning system with a query processing engine. This would bring several advantages: (i) the pruning overhead would be diluted in the parsing/validation phase and (ii) an interaction between the query engine and the loading module would provide a way not only to prune the document but to start answering the query in streaming, when possible.

**Acknowledgements.** We would like to thank Haiming Chen for pointing us an error in the two typing systems of a preliminary version of this work. This work benefitted from several discussions with and suggestions from Ioana Manolescu and Carlo Sartiani. Two of the three VLDB anonymous referees provided very useful feedback. This work was partially funded by the French ACI project “Transformation Languages for XML: Logics and Applications” (TraLaLA) and the French ACI young researcher project “WebStand”.

## 8. REFERENCES

- [1] Galax. <http://www.galaxquery.org>.
- [2] XML Path Language (XPath) 2.0. <http://www.w3.org/TR/xpath20>.
- [3] XML Query Use Cases. <http://www.w3.org/TR/xquery-use-cases/>.
- [4] XQuery 1.0 and XPath 2.0 Formal Semantics. <http://www.w3.org/TR/xquery-antics>.
- [5] XQuery 1.0 and XPath 2.0 Functions and Operators. <http://www.w3.org/xquery-operators>.
- [6] V. Benzaken, G. Castagna, and A. Frisch. CDuce: an XML-centric general-purpose language. In *ICFP '03, 8th ACM Int. Conf. on Functional Programming*, pages 51–63, 2003.
- [7] V. Benzaken, G. Castagna, and C. Miachon. A full pattern-based paradigm for XML query processing. In *PADL '05, the 7th Int. Symp. on Practical Aspects of Declarative Languages*, number 3350 in LNCS. Springer, 2005.
- [8] V. Benzaken, C. Delobel, and G. Harrus. Clustering strategies in O<sub>2</sub>: an overview. In *Building an Object-Oriented Database System: the Story of O<sub>2</sub>*. Morgan Kaufman, 1992.
- [9] S. Bressan, B. Catania, Z. Lacroix, Y-G Li, and A. Maddalena. Accelerating queries by pruning XML documents. *Data Knowl. Eng.*, 54(2):211–240, 2005.
- [10] D. Colazzo. *Path Correctness for XML Queries: Characterization and Static Type Checking*. PhD thesis, Dip. di Informatica, Università di Pisa, 2004.
- [11] D. Colazzo, G. Ghelli, P. Manghi, and C. Sartiani. Types for Path Correctness for XML Queries. In *ICFP '04, 9th ACM Int. Conf. on Functional Programming*, 2004.
- [12] M. Franceschet. XPathMark - An XPath benchmark for XMark generated data. In *XSym 2005, 3rd Int. XML Database Symposium*, LNCS n. 3671, 2005.
- [13] D. Lee, M. Mani, and M. Murata. Reasoning about XML Schema Languages using Formal Language Theory. Technical report, IBM Almaden Research, 2000.
- [14] A. Marian and J. Siméon. Projecting XML documents. In *VLDB '03*, pages 213–224, 2003.
- [15] D. Olteanu, H. Meuss, T. Furché, and F. Bry. XPath: Looking forward. In *Proc. EDBT Workshop (XMLDM)*, volume 2490 of LNCS, pages 109–127. Springer, 2002.
- [16] S. Pappas and H.V. Jagadish. Pattern tree algebras: Sets or sequences? In *VLDB*, 2005.
- [17] A. Schmidt, F. Waas, M. L. Kersten, M. J. Carey, I. Manolescu, and R. Busse. XMark: A benchmark for XML data management. In *VLDB '02*, pages 974–985, 2002.