An Algorithm for Streaming XPath Processing with Forward and Backward Axes

Charles Barton, Philippe Charles
Deepak Goyal, Mukund Raghavachari
IBM T.J. Watson Research Center

Marcus Fontoura, Vanja Josifovski
IBM Almaden Research Center

Abstract

We present a novel streaming algorithm for evaluating XPath expressions that use backward axes (*parent* and *ancestor*) and forward axes in a single document-order traversal of an XML document. Other streaming XPath processors, such as YFilter, XTrie, and TurboXPath handle only forward axes. We show through experiments that our algorithm significantly outperforms (by more than a factor of two) a traditional non-streaming XPath engine. Furthermore, since our algorithm only retains relevant portions of the input document in memory, it scales better than traditional XPath engines. It can process large documents; we have successfully tested documents over 1GB in size. On the other hand, the traditional XPath engine degrades considerably in performance for documents over 100 MB in size and fails to complete for documents of size over 200 MB.

1 Introduction

XPath 1.0 [8], a language for addressing parts of XML [4] documents, is an integral component of languages for XML processing such as XSLT [7] and XQuery [10]. The performance of implementations of these languages depends on the efficiency of the underlying XPath engine. XPath expressions have also been used as a general-purpose mechanism for accessing portions from XML documents, for example, an XPath-based API is provided in DOM 3 [15] for traversing DOM [11] trees. XPath expressions have found use in publish-subscribe systems as a mechanism for specifying content-based subscriptions [5]. Given the central role that XPath plays in the XML stack, algorithms for improving the performance of evaluating common XPath expressions are essential.

In many environments, a natural way of processing documents is to *stream* over them, *i.e.*, evaluate the query on the input document as it is parsed, storing only portions of the document relevant to the result of the query [1, 9, 12]. A streaming XPath engine is structured as shown in Figure 1. An XPath expression is analyzed

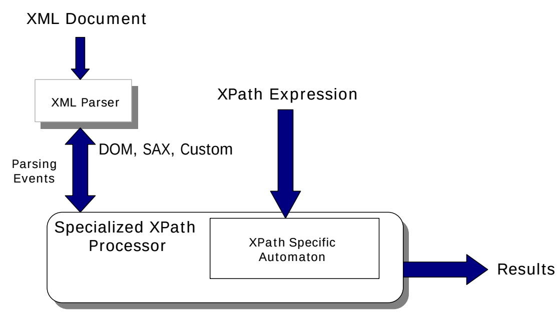

Figure 1: Structure of a streaming XPath processor.

and represented as an automaton. The XPath engine consumes events (for example, SAX events) produced by a parser, and for each event, the automaton may make state transitions, and if necessary, store the element. At the end of processing the document, the XPath engine returns the list of elements that are the result of the evaluation of the XPath expression.

Most current XPath engines, for example, the one provided with Xalan [2], require that the entire document be in memory before evaluating an XPath expression. For large documents, this approach may result in unacceptable overhead. Furthermore, the XPath engine in Xalan evaluates XPath expressions in a naive manner, and may perform unnecessary traversals of the input document. For example, consider an expression such as /descendant::x/ancestor::y, which selects all y ancestors of x elements in the document. The Xalan XPath engine evaluates this expression by using one traversal over the entire document to find all the x elements, and for each x element, a visit to each of its ancestors to find appropriate y elements. As a result, some elements in the document are visited more than once.

The premise of streaming XPath is that, in many instances, XPath expressions can be evaluated in one depth-first, document-order traversal of an XML document. The benefits of streaming XPath are twofold. First, rather than storing the entire document in memory, only the portion of the document relevant to the evaluation of the XPath is stored. Second, the algorithm visits each node in the document exactly once, avoiding unnecessary traversals.

In this paper, we present the $\chi\alpha o\varsigma$[1] algorithm, which can evaluate XPath expressions containing both backward (such as parent and ancestor) and forward axes in a streaming fashion. Other streaming XPath processors, such as YFilter [9], XTrie [5], and TurboXPath [12] handle only forward axes. This paper makes the following contributions:

1. A novel streaming algorithm for handling both backward and forward axes.

2. A concise representation of an XPath expression, called X-dag (Section 3.2), where all backward constraints are converted into forward constraints, a key step in making streaming XPath processing in the presence of backward axes possible.

3. A data structure called the matching structure (Section 3.4) that represents compactly all matchings (Section 3.3) of an XPath expression in a document.

1.1 Related Work

Our work is most closely related to the *XFilter* [1], *YFilter* [9], *XTrie* [5], and *TurboXPath* [12] systems, all of which involve evaluation of XPath/XQuery-based queries on streaming XML documents. XFilter, YFilter, and XTrie are XML filtering systems where documents are routed and filtered based on subscriptions that are expressed as queries. The TurboXPath system has been used for XML-enabled data integration where user queries can operate over a mixture of locally stored data in a relational database and data streamed from external sources. The XFilter system handles simple XPath location path expressions (straight-line path expressions without any branching and predicates) by transforming them into a Deterministic Finite Automaton. The YFilter system is an extension of XFilter in which a group of simple XPath location path expressions are combined into a single Nondeterministic Finite Automaton (NFA), which corresponds to the union of these path expressions. Both XTrie and TurboXPath can handle tree-shaped path expressions involving predicates (which are internally represented as trees called the XTrie and ParseTree respectively). In addition, TurboXPath can also handle multiple output nodes. However, all of these systems are limited to handling location path expressions that only contain forward axes (e.g. child, descendant, forward-sibling). The $\chi\alpha o\varsigma$ system improves upon these systems by adding the ability to handle both backward (e.g. parent, ancestor) and forward axes in the context of streaming XML.

[1] $\chi\alpha o\varsigma$ (Xaos, pronounced Chaos) is an acronym for XML Analysis, Optimization, and Stuff

Tozawa and Murata [14] describe a method for converting an XPath expression into modal logic formulas with past modalities. They present an algorithm for converting such formulas into tree automata, which can be used to evaluate XPath expressions on an input document. Their paper describes a theoretical approach that can handle all XPath axes. The current status of the implementation of their algorithm is unclear. It would be interesting to compare the performance of their implementation with that of $\chi\alpha o\varsigma$.

The *NiagaraCQ* [6] system is a continuous query system that supports querying of distributed XML datasets using an XML query language. Continuous queries allow users to receive new results as they become available. The focus of the NiagaraCQ project is on exploiting similarities in structure of queries to share computation across groups of queries, and use of incremental group optimization and incremental evaluation techniques. However, the queries that they focus on involve simple structural pattern matching rather than XPath/XQuery-based queries that we deal with in this paper.

2 Background

We describe the tree model of XML documents that is the basis of the definition of XPath. We then describe the event stream that drives the $\chi\alpha o\varsigma$ algorithm. Finally, we present the subset of XPath that we focus on in this paper.

2.1 Tree Model for XML Documents

An XML document can be represented as a tree, whose nodes represent the structural components of the document — elements, text, attributes, comments, and processing instructions. Parent-child edges in the tree represent the inclusion of the child component in its parent element, where the scope of an element is bounded by its start and end tags. The tree corresponding to an XML document is rooted at a virtual element, Root, which contains the document element. We will, henceforth, discuss XML documents in terms of their tree representation; $\mathcal{D}$ represents an XML document, and $V_{\mathcal{D}}$ and $E_{\mathcal{D}}$ denote its nodes and edges respectively. Figure 2 illustrates the tree representation of an XML document.

For simplicity of exposition, we focus on elements in this paper, and ignore attributes, text nodes, etc. The tree, therefore, consists of the virtual root and the elements of the document. To avoid confusion between the XML document tree and the tree representation of the XPath (described later), we use *elements* to refer to the nodes of the XML tree. We assume that the following functions are defined on the elements of an XML document:

- $id : V_{\mathcal{D}} \rightarrow Integer$: Returns a unique identifier for each element in a document.
- $tag : V_{\mathcal{D}} \rightarrow String$: Returns the tag name of the element.
- $level : V_{\mathcal{D}} \rightarrow Integer$: Returns the distance of the element from the root, where $level(\mathsf{Root}) = 0$.

We use $\mathsf{x}_{i,l}$ to denote an element with $tag = \mathsf{x}, id = i, level = l$. For example, the element U in Figure 2(b) is denoted by $\mathsf{U}_{8,3}$.

2.2 Event-Based Parsing

An event-based parser, for example, a SAX parser, scans an XML document, producing events as it recognizes element tags and other components of the document. We register functions that are invoked by the parser on start and end element events. Each event conveys the name and level of the corresponding element. The production of events is equivalent to that of a depth-first, pre-order traversal of the document tree, where for each element, a start element event is generated, then its subtree is processed in depth-first order, and finally, an end element event is generated.

2.3 XPath

The XPath language defines expressions for addressing parts of an XML document. We focus on *location path* expressions which evaluate to a set of elements in the document. A location path is a structural pattern composed of sub-expressions called *Step*, joined by the '/' character. Each step consists of an *axis specifier*, a *nodetest*, and zero or more predicates. Location paths are *absolute* if they begin with a '/'; otherwise they are *relative*. Table 1 provides the BNF for the XPath subset that we shall use in this paper (we refer to expressions satisfying this grammar as Restricted XPaths – $\mathcal{R}$*xp*).[2]

XPath expressions are evaluated relative to a context node in the document tree. The context node for an absolute location path is always the root element. To evaluate a relative location path, *Step / RelLocPath*, with respect to a context node, c, one first computes *Step* relative to c, yielding a set of elements, $\mathcal{N}$. The meaning of *Step / RelLocPath* is the union of the sets of elements obtained by evaluating *RelLocPath* in context d, where d ranges over $\mathcal{N}$.

The set of elements searched in the evaluation of a *Step* at a context node, c, depends on its axis specifier. For example, the result of evaluating descendant::section is the subset of the proper descendants of the context node that match section. While the $\chi\alpha o\varsigma$ algorithm is extensible to handle all thirteen axis specifiers in XPath 1.0, we focus on four: child, descendant, parent, and ancestor.

[2] We will not use abbreviated XPath expressions in this paper.

Table 1: XPath subset addressed in paper.

$AbsLocPath$	:=	$'/'\ RelLocPath$
$RelLocPath$	:=	$Step\ '/'\ RelLocPath \mid Step$
$Step$	:=	$Axis :: NodeTest \mid$
		$Step\ '['\ PredExpr\ ']'$
$PredExpr$	:=	$RelLocPath\ and\ PredExpr \mid$
		$AbsLocPath\ and\ PredExpr \mid$
		$RelLocPath \mid AbsLocPath$
$Axis$	:=	$ancestor \mid parent \mid child \mid$
		$descendant$
$NodeTest$	:=	$String$

Steps may contain predicates, which restrict the set of elements selected. For example, descendant::chapter[ancestor::book and child::table] selects all chapter descendants of the context node that have a book element as an ancestor and a table element as a child. Note that each chapter element is used as a context node in evaluating the subexpressions, ancestor::book and child::table.

3 X-tree, X-dag, Matchings, and Matching Structures

The $\chi\alpha o\varsigma$ algorithm operates on two representations of an XPath expression called X-tree and X-dag. The x-dag is a key construct in our algorithm since it converts backward constraints, such as parent, into forward constraints, thus making streaming processing possible. We use an alternate semantics of XPath expressions defined on x-trees based on the notion of *matchings*. It can be shown that our semantics is equivalent to the semantics provided in the XPath 1.0 specification. The construction of the *matching-structure*, a compact representation of matchings, is the main goal of the algorithm. We describe these concepts in this section.

3.1 X-tree

We represent an $\mathcal{R}$*xp* expression as a rooted tree, called X-tree, with labeled vertices and edges, $\mathcal{T} = (V_{\mathcal{T}}, E_{\mathcal{T}})$, where the root is labeled Root. We use the term x-node to refer to the vertices of an x-tree. For each *NodeTest* in the expression, there is an x-node in the x-tree labeled with the nodetest. Each x-node (except Root) has a unique incoming edge, which is labeled with the *Axis* specified before the *NodeTest*. One of the x-nodes is designated to be the output x-node. There are functions, $label : V_{\mathcal{T}} \rightarrow String$, and $axis : E_{\mathcal{T}} \rightarrow \{ancestor, parent, child, descendant\}$ that return the

```
<X>
  <Y>
    <Z>
      <V/>
      <V/>
      <W>
        <W/>
      </W>
    </Z>
    <U/>
  </Y>
  <Y>
    <Z>
      <W/>
    </Z>
  </Y>
</X>
```

(a)

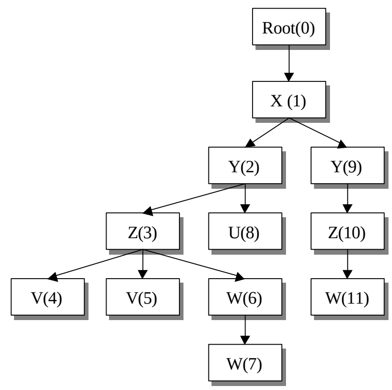

(b)

Figure 2: (a) An XML Document (b) Tree representation of the same document. The number in parentheses next to the tag of each element is the id of the element.

labels associated with the x-nodes and edges respectively. Rules for building an x-tree from an $\mathcal{R}xp$ are provided in Appendix A for the interested reader. Figure 3a provides an example of an x-tree.

3.2 X-dag

We also use a directed, acyclic graph representation of an $\mathcal{R}xp$ called an X-dag. The x-dag is obtained from the x-tree by reformulating the ancestor and parent constraints in the tree as descendant and child constraints. More precisely, it is a directed, labeled graph, $\mathcal{G} = (V_\mathcal{G}, E_\mathcal{G})$, with the same set of vertices as $\mathcal{T}$, and edges defined as follows:

1. Edges in $\mathcal{T}$ labeled child or descendant are also edges of $\mathcal{G}$.

2. For each edge in $\mathcal{T}$ labeled parent, there is an edge joining the same nodes but with direction reversed and label changed to child. Similarly, ancestor edges are reversed and relabeled as descendant edges.

3. For any non-root x-node $v \in \mathcal{G}$ having no incoming edges, a descendant edge is added from Root to v.

Figure 3b gives the x-dag corresponding to the x-tree in Figure 3a.

3.3 Matchings

Let v_1 and v_2 be two x-nodes in an x-tree $\mathcal{T}$ connected by an edge e, and let d_1 and d_2 be two elements in a document $\mathcal{D}$. We say that the pair (v_1, d_1) is *consistent* with (v_2, d_2) (relative to x-tree $\mathcal{T}$ and document $\mathcal{D}$) if d_1 and d_2 satisfy the relation $axis(e)$. For example, if v_1 and v_2 are connected by an edge labeled ancestor, then d_2 must be an ancestor of d_1 in $\mathcal{D}$. A matching, $m : V_\mathcal{T} \rightarrow V_\mathcal{D}$, is a partial mapping from x-nodes of x-tree $\mathcal{T}$ to elements of document $\mathcal{D}$ such that the following conditions hold.

1. For all x-nodes $v \in domain(m)$, $label(v) = tag(m(v))$, *i.e.* all mapped vertices satisfy the node-test.

2. For all x-nodes v_1 and v_2 connected by an edge in $\mathcal{T}$ such that $v_1, v_2 \in domain(m)$, $(v_1, m(v_1))$ is consistent with $(v_2, m(v_2))$.

A matching is *at an x-node* v if and only if its domain is contained in the sub-tree rooted at v. A matching at v is total if its domain contains all the vertices of the subtree rooted at v. It is easy to show that a document element n is in the result of an $\mathcal{R}xp$ r, if and only if, there exists a total matching at Root for the x-tree $\mathcal{T}$ of r, where the output x-node of $\mathcal{T}$ is mapped to n. $\chi\alpha o\varsigma$ computes the result defined by an $\mathcal{R}xp$ precisely in this manner. It finds all matchings from $V_\mathcal{T}$ to $V_\mathcal{D}$, and emits the document elements that correspond to the output x-node.

3.4 Matching-Structure

The algorithm constructs a data structure called a *matching-structure* which is a compact representation of all total matchings at Root of the $\mathcal{R}xp$ relative to the input document. A matching-structure, $\mathcal{M}_{v,e}$, is associated with x-node v, and represents a set of matchings at v in which v is mapped to the document element e. The matching-structure $\mathcal{M}_{v,e}$ additionally contains a submatching for every child of v in the x-tree. A submatching at child w of v is a (possibly empty) set of matching-structures at w. For any matching-structure $\mathcal{M}_{w,e'}$ in the submatching of $\mathcal{M}_{v,e}$ at w, we require that (v, e) be consistent with (w, e'). A matching-structure $\mathcal{M}_{v,e}$ is said to

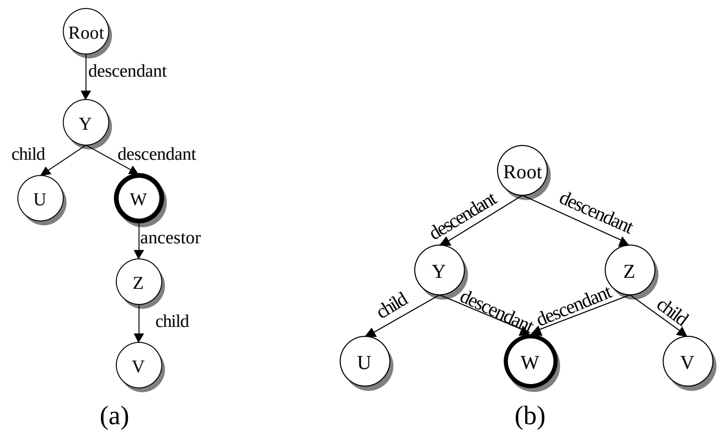

Figure 3: (a) X-tree representation of /descendant::y[child::u]/descendant::w[ancestor::z/child::v] (b) X-dag representation of the same XPath expression. The circle corresponding to W has a thick edge to represent the fact that it is the output node.

be a *parent-matching* of a matching-structure $\mathcal{M}_{w,e'}$ if v is a parent of w in x-tree $\mathcal{T}$ and (v, e) is consistent with (w, e'). If $\mathcal{M}_{v,e}$ is a parent-matching of $\mathcal{M}_{w,e'}$, then we say also that $\mathcal{M}_{w,e'}$ is a *child-matching* of $\mathcal{M}_{v,e}$.

Figure 4 shows the matching structure at the end of processing the XPath of Figure 3 on the document in Figure 2, and the four total matchings at Root. The result is obtained by taking the W projection, that is $\{W_{6,4}, W_{7,5}\}$.

4 The $\chi\alpha o\varsigma$ Algorithm

$\chi\alpha o\varsigma$ processes events as they are generated by an event-based parser by operating over both the x-dag and x-tree views of the input $\mathcal{R}$*xp*. At the end of processing the document, the result of the $\mathcal{R}$*xp* expression is encoded in $\mathcal{M}_{\mathsf{Root},\mathsf{Root}}$. For efficiency, $\chi\alpha o\varsigma$ filters out events that do not contribute to any matchings. Relevant events are processed to build matching-structures. Finally, $\mathcal{M}_{\mathsf{Root},\mathsf{Root}}$ is used to emit the appropriate output. We shall describe these three stages in this section. A walk through of the execution of the algorithm on the $\mathcal{R}$*xp* of Figure 3 and the document of Figure 2 is provided in Table 2.

4.1 Filtering Events

At any point during execution, $\chi\alpha o\varsigma$ has processed a prefix of the input document. An infinite number of XML documents share the same prefix, and $\chi\alpha o\varsigma$ cannot predict the future sequence of events that will be generated by the parser. An element, e, is *relevant* if there exists some document completion where e participates in a matching. All relevant elements must be processed. As events are processed, new relevant elements may be seen, or elements that were earlier deemed relevant may no longer be relevant. The x-dag representation of the $\mathcal{R}$*xp* is used to determine if an element is relevant.

An element that does not match any x-node is not relevant since it cannot participate in any matching. Let *open* elements be those elements for which we have seen a start element event, but not an end element event. By virtue of the depth-first manner in which events are generated, at a start element event for element e that matches a x-node v, the open elements are ancestors of e in the document. If for some parent of v' of v in x-dag, there is no open, relevant element, e', such that (v, e) and (v', e') are consistent, then e cannot be relevant. There is no sequence of events that a parser can generate for which that constraint will become true, and therefore, e cannot contribute to any matching. For the x-dag in Figure 3b, no element that matches W will be relevant unless there are open relevant elements that match Y and Z.

At every step, we maintain a *looking-for* set, $\mathcal{L}$, which allows us to evaluate whether the element associated with the next start element event is relevant. The members of $\mathcal{L}$ are $(v \in V_{\mathcal{P}}, level)$ pairs, where level may be an integer or $*$. The element e associated with a start element event is relevant if and only if there exists $(v, level) \in \mathcal{L}, label(v) = tag(e),\ and\ (level = level(e)\ or\ level = *)$. Integer levels are used to enforce the constraint that if (v, e) and (v', e') are consistent and if $axis(v, v')$ = child, then $level(v') = 1 + level(v)$.

4.2 Building Matching Structures

We assume from now on that all events corresponding to elements that are not relevant have been discarded. When $\chi\alpha o\varsigma$ processes a start element event for an element e

Table 2: Walk through of evaluation of XPath of Figure 3 on document of Figure 2. **Start (End)**: $A_{x,y}$ denotes the start (end) element event for an element, $A_{x,y}$. The Looking-for set column shows $\mathcal{L}$ at the end of processing the event.

Step	Event	Matches	Comments	Looking-for Set
1	**Start**: $\mathsf{Root}_{0,0}$	$(\mathsf{Root}, 0)$	Add $(Y, *)$ and $(Z, *)$ to $\mathcal{L}$, since Root is an open, relevant element matching their ancestors in the x-dag.	$\{(Y, *), (Z, *)\}$
2	**Start**: $\mathsf{X}_{1,1}$		Discarded.	$\{(Y, *), (Z, *)\}$
3	**Start**: $\mathsf{Y}_{2,2}$	$(Y, *)$	Start looking for U at level 3 since U is connected to Y by a child edge in the x-dag, and Y is matched at level 2. Do not add W to $\mathcal{L}$ because there is no open element that matches its Z parent in the x-dag. Continue looking for $(Y, *)$ because any element with tag Y in the subtree of this element will also be a candidate for matching Y.	$\{(Y, *), (Z, *), (U, 3)\}$
4	**Start**: $\mathsf{Z}_{3,3}$	$(Z, *)$	Start looking for $(V, 4)$ since we have open relevant elements matching Z and Root in the x-dag. We look for it at level 4 because the (Z, V) edge is labeled child.	$\{(Y, *), (Z, *), (W, *), (V, 4)\}$
5	**Start**: $\mathsf{V}_{4,4}$	$(V, 4)$	We stop looking for $(V, 4)$ because until we see the end of this element, $level > 4$.	$\{(Y, *), (Z, *), (W, *)\}$.
6	**End**: $\mathsf{V}_{4,4}$	$(V, 4)$	There is a total matching at V, represented as $\mathcal{M}_{V,4}$. This matching-structure is propagated to the appropriate submatching of $\mathcal{M}_{Z,3}$, the only parent-matching of $\mathcal{M}_{V,4}$.	$\{(Y, *), (Z, *), (W, *), (V, 4)\}$.
7	**Start**: $\mathsf{V}_{5,4}$	$(V, 4)$		$\{(Y, *), (Z, *), (W, *)\}$
8	**End**: $\mathsf{V}_{5,4}$	$(V, 4)$	As before, $\mathcal{M}_{V,5}$ is added to the appropriate submatching of $\mathcal{M}_{Z,3}$.	$\{(Y, *), (Z, *), (W, *), (V, 4)\}$
9	**Start**: $\mathsf{W}_{6,4}$	$(W, *)$		$\{(Y, *), (Z, *), (W, *)\}$
10	**Start**: $\mathsf{W}_{7,5}$	$(W, *)$		$\{(Y, *), (Z, *), (W, *)\}$
11	**End**: $\mathsf{W}_{7,5}$	$(W, *)$	W in the x-dag has an outgoing ancestor edge. All child-matchings of $\mathcal{M}_{W,7}$, in this case, $\mathcal{M}_{Z,3}$, are propagated into the appropriate submatching of $\mathcal{M}_{W,7}$. All submatchings of $\mathcal{M}_{W,7}$ are now non-empty. $\mathcal{M}_{W,7}$ is propagated to $\mathcal{M}_{Y,2}$	$\{(Y, *), (Z, *), (W, *)\}$
12	**End**: $\mathsf{W}_{6,4}$	$(W, *)$	As above, $\mathcal{M}_{W,6}$ is propagated to $\mathcal{M}_{Y,2}$.	$\{(Y, *), (Z, *), (W, *)(V, 4)\}$
13	**End**: $\mathsf{Z}_{3,3}$	$(Z, *)$	Z has an incoming edge labeled ancestor. Since $\mathcal{M}_{Z,3}$ is satisfied, no clean up is necessary.	$\{(Y, *), (Z, *)(U, 3)\}$
14	**Start**: $\mathsf{U}_{8,3}$	$(U, 3)$		$\{(Y, *), (Z, *)\}$
15	**End**: $\mathsf{U}_{8,3}$	$(U, 3)$	The total matching at U, $\mathcal{M}_{U,8}$ is propagated to $\mathcal{M}_{Y,2}$.	$\{(Y, *), (Z, *), (U, 3)\}$
16	**End**: $\mathsf{Y}_{2,2}$	$(Y, *)$	$\mathcal{M}_{Y,2}$ is satisfied since both submatchings, corresponding to U and W are non-empty. We propagate $\mathcal{M}_{Y,2}$, and we have a total matching at Root.	$\{(Y, *), (Z, *)\}$
17	**Start**: $\mathsf{Y}_{9,2}$	$(Z, *)$		$\{(Y, *), (Z, *), (U, 3)\}$
18	**Start**: $\mathsf{Z}_{10,3}$	$(Z, *)$		$\{(Y, *), (Z, *), (V, 4), (W, *)\}$
19	**Start**: $\mathsf{W}_{11,4}$	$(W, *)$		$\{(Y, *), (Z, *), (W, *)\}$
20	**End**: $\mathsf{W}_{11,4}$	$(W, *)$	Again, since W has an outgoing edge labeled ancestor, we add $\mathcal{M}_{Z,10}$ optimistically to the appropriate submatching of $\mathcal{M}_{W,11}$. Since this matching is now satisifed, it is propagated to $\mathcal{M}_{Y,9}$.	$\{(Y, *), (Z, *), (W, *), (V, 4)\}$
21	**End**: $\mathsf{Z}_{10,3}$	$(Z, *)$	$\mathcal{M}_{Z,10}$ is not satisfied — the submatching for V is empty. We undo the propagation of $\mathcal{M}_{Z,10}$ to $\mathcal{M}_{W,11}$. Since $\mathcal{M}_{W,11}$ now is no longer satisfied, we undo the propagation from $\mathcal{M}_{W,11}$ to $\mathcal{M}_{Y,9}$.	$\{(Y, *), (Z, *), (U, 3)\}$
22	**End**: $\mathsf{Y}_{9,2}$	$(Y, *)$	$\mathcal{M}_{Y,9}$ is not satisfied. Nothing is propagated.	$\{(Y, *), (Z, *)\}$
23	**End**: $\mathsf{X}_{1,1}$		Discarded.	$\{(Y, *), (Z, *)\}$
24	**End**: $\mathsf{Root}_{0,0}$	$(\mathsf{Root}, 0)$	There is one entry in the submatching corresponding to Y, $\mathcal{M}_{Y,2}$. $\mathcal{M}_{\mathsf{Root},0}$ is satisfied.	$\{(\mathsf{Root}, 0)\}$

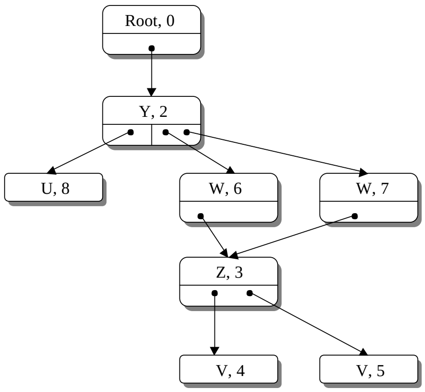

Total Matchings at Root

$[\mathsf{Root} \mapsto 0, Z \mapsto 3, Y \mapsto 2, U \mapsto 8, V \mapsto 4, W \mapsto 6]$
$[\mathsf{Root} \mapsto 0, Z \mapsto 3, Y \mapsto 2, U \mapsto 8, V \mapsto 4, W \mapsto 7]$
$[\mathsf{Root} \mapsto 0, Z \mapsto 3, Y \mapsto 2, U \mapsto 8, V \mapsto 5, W \mapsto 6]$
$[\mathsf{Root} \mapsto 0, Z \mapsto 3, Y \mapsto 2, U \mapsto 8, V \mapsto 5, W \mapsto 7]$

Solution: $\{W_{6,4}, W_{7,5}\}$

Figure 4: Matching Structure at the end of processing the XPath of Figure 3. The boxes represent matching-structures. For a matching-structure, $\mathcal{M}_{v,e}$, the top half of the box shows $(v, id(e))$. Each slot in the bottom half of the box corresponds to a submatching, which is represented as a list of pointers to the child matchings.

that matches a x-node, v, it creates a matching-structure, $\mathcal{M}_{v,e}$, to represent the match. Note that e may match more than one x-node in the x-tree; a matching-structure is created for each such match. The submatchings for these matching-structures are initially empty. As $\chi\alpha o\varsigma$ processes events, it stitches together these matching-structures, so that when the end of the document is seen, $\mathcal{M}_{\mathsf{Root},\mathsf{Root}}$ encodes all total matchings at Root in the document.

The key step in this process is *propagation*. At an end element event for an element e that matches x-node v, we attempt to determine if $\mathcal{M}_{v,e}$ represents a total matching at v. If there is a total matching, we insert $\mathcal{M}_{v,e}$ into the appropriate submatching of its parent-matchings. This propagation may be optimistic in that one may have to undo the propagation as more events are processed. Let us first, however, consider the simpler situation where no cleanup of propagation is necessary, when the x-tree does not contain any edges labeled ancestor or parent. This corresponds to $\mathcal{R}xp$'s that use only the child and descendant axes.

When the x-tree contains only child and descendant constraints, any total matching m at v, where $m(v) = e$ maps all x-nodes in the subtree of v to elements that lie in the document subtree of e. Since the total matching is contained within the subtree of e, by the time the end element event for e is seen, we can determine conclusively if $\mathcal{M}_{v,e}$ represents a total matching at v. This leads naturally to an inductive approach to building matchings. For an end element event e, where $\mathcal{M}_{v,e}$ is a matching-structure:

1. If v is a leaf in the x-tree, $\mathcal{M}_{v,e}$ represents a total matching at v by definition (v has no subtrees). We propagate $\mathcal{M}_{v,e}$ to the appropriate parent-matchings.

2. If v is not a leaf, $\mathcal{M}_{v,e}$ represents a total matching at v, if and only if, all submatchings are non-empty. Otherwise, no total matching exists. If we had found appropriate total matchings for each of the children of v in the x-tree, they would have been propagated to $\mathcal{M}_{v,e}$ by the time the end element event for e is processed. As above, if $\mathcal{M}_{v,e}$ represents a total matching, we propagate it to all appropriate parent-matchings.

If at the end of processing the document (when we receive the end element event for Root), $\chi\alpha o\varsigma$ finds that all the submatchings of $\mathcal{M}_{\mathsf{Root},\mathsf{Root}}$ are non-empty, we have a total matching at Root.

The presence of ancestor and parent edges in the x-tree complicates this process because one may not be able to determine conclusively whether a total matching exists for a $\mathcal{M}_{v,e}$ by the end of element e. For example, in Figure 3a, one might not find a total matching for the subtree rooted at Z, until after one sees the end of an element matching W. The propagation process remains the same, except for a x-node that has an incoming or an outgoing edge labeled ancestor or parent. For a $\mathcal{M}_{v,e}$, the modified steps are as follows:

- If there is an outgoing edge (v, v') labeled ancestor or parent, and the submatching for v' is empty, we cannot assert that there exists no total matching at v. We, optimistically, propagate each child-matching, $\mathcal{M}_{v',e'}$, into the appropriate submatching of $\mathcal{M}_{v,e}$. We then proceed as before. If all submatchings are satisfied, $\mathcal{M}_{v,e}$ is propagated to its parent-matchings.

- If there is an incoming edge (v', v) labeled ancestor or parent, then $\mathcal{M}_{v,e}$ may have been propagated optimistically to its parent-matchings. If we can determine conclusively that $\mathcal{M}_{v,e}$ cannot represent a total matching at v, we undo the propagation of $\mathcal{M}_{v,e}$. The removal of $\mathcal{M}_{v,e}$ from a submatching of a parent-matching $\mathcal{M}_{v',e'}$ may result in that submatching becoming empty — $\mathcal{M}_{v',e'}$ is no longer a total matching at v'. We then recursively undo the propagation of $\mathcal{M}_{v',e'}$ from its parent-matchings.

4.3 Emitting Output

At the end of processing the document, if the submatchings of $\mathcal{M}_{\mathsf{Root},\mathsf{Root}}$ are all non-empty, we have at least one total matching at Root. The output is emitted by traversing the matching structure, and emitting an element e when we visit $\mathcal{M}_{v,e}$, where v is the output x-node of the $\mathcal{R}xp$. For example, in Figure 4, we output $W_{6,4}$ when we first visit $\mathcal{M}_{W,6}$ and $W_{7,5}$, when we first visit $\mathcal{M}_{W,7}$.

5 Experimental Results

The $\chi\alpha o\varsigma$ algorithm examines each element event exactly once and the processing of an event involves only constant-time operations. The execution time of the $\chi\alpha o\varsigma$ algorithm is, therefore, linear in terms of the input document size. Furthermore, $\chi\alpha o\varsigma$ stores only those elements relevant to the calculation of the final solution. We would, therefore, expect the $\chi\alpha o\varsigma$ algorithm to show better memory utilization than Xalan [2], which stores the whole document in memory. In this section, we provide experimental results that validate these claims. We, first, provide results using documents generated by XMark [13]. To gain further insight into the relative performance of $\chi\alpha o\varsigma$ and Xalan, we also run experiments using a custom XPath and XML document generator.

All experiments were run on a 550 MhZ, 256 MB, Pentium III box, running Linux 2.4. $\chi\alpha o\varsigma$ was written in C++, and we use Xalan-C++ 1.3.1. Both $\chi\alpha o\varsigma$ and Xalan were compiled using gcc -O (version 2.92).

5.1 Experiments using XMark

Using XMark, we generated documents with scale factors .03125, .0625, .125, .25, .5, 1, 2, and 4, respectively. These correspond to documents ranging in size from 3.5 MB to 446 MB. We then evaluate the XPath expression, //listitem/ancestor::category//name on these documents, using both $\chi\alpha o\varsigma$ and Xalan. Figure 5 reports the results of these experiments.

Note that Xalan fails to complete on the two largest documents (approx. 222 MB and 446 MB), and furthermore, that there is a sharp spike in going from 55 MB to 111 MB. These effects can be attributed to the memory requirements of Xalan (the spike is the region where Xalan exhibits thrashing behavior in memory). On the other hand, $\chi\alpha o\varsigma$ scales linearly, as is expected. Table 3 reports the number of elements discarded by the algorithm

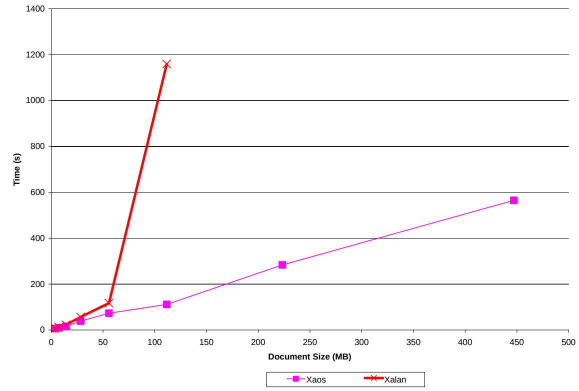

Figure 5: Time in seconds on XMark-generated documents: $\chi\alpha o\varsigma$ versus Xalan. The XPath expression executed is //listitem/ancestor::category//name

Table 3: Number of elements discarded by $\chi\alpha o\varsigma$ in processing of XMark-generated documents

Scale	Doc. Size	Elements	% Discarded
.03125	3.49 MB	52069	99.8 %
.625	6.88 MB	103999	99.8 %
.125	13.86 MB	210538	99.8 %
.25	27.87 MB	417160	99.8 %
.5	55.32 MB	832911	99.8 %
1	111.12 MB	166311	99.8 %
2	222.90 MB	3337649	99.8 %
4	446.71 MB	6688651	99.8 %

as not being relevant. As can be seen from the results, a very small percentage of elements in a document (less than .2 %) is stored and processed, resulting in a signficant reduction in storage requirements.

5.2 Custom XPath generator

We use a custom XPath generator to generate a set of random XPath expressions (of size 6 – six node tests in the expression), and for each XPath expression, we generate a random XML document based on the XPath expression. The generated XML document has the characteristic that, for large document sizes, the XPath expression will have many matches (and near matches) in the document.

We use two versions of $\chi\alpha o\varsigma$ in our comparison. The first, $\chi\alpha o\varsigma$(SAX), uses the Xerces SAX parser [3], which is also used by Xalan. To factor out the costs of parsing and building a tree from the time to evaluate an expression, we also implemented a version of $\chi\alpha o\varsigma$ on top of Xalan. $\chi\alpha o\varsigma$(DOM) builds an internal version of the input document in the same way that Xalan does. We then traverse this tree in a depth-first fashion and gener-

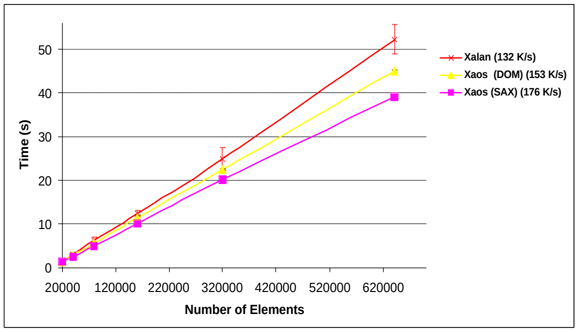

Figure 6: Overall Time in seconds: $\chi\alpha o\varsigma$ versus Xalan

ate events that a SAX parser would. By subtracting the parsing and tree-building time from the overall time, we get an accurate measure of the time spent in evaluating the expression.

We vary the XML document size from 20,000 elements to 640,000 elements (200K - 6.7 MB). At each document size, we execute 10 runs of the following: 1) generate an XPath expression, 2) generate an XML document from the XPath expression, and 3) evaluate the XPath expression using $\chi\alpha o\varsigma$ and Xalan. We report the average execution time and the standard deviation of the 10 runs at each XML document size.

5.2.1 Overall Execution Time

We first compare the performance of $\chi\alpha o\varsigma$ to that of using the Xalan XPath engine (SimpleXPathAPI). Figure 6 plots the average execution time (average over the 10 runs at each document size) versus document size (in number of elements). The error bars represent the standard deviation from the mean. All times include the cost of parsing.

As can be seen from the graph, $\chi\alpha o\varsigma$(SAX) is roughly 25% faster than the Xalan XPath engine. With documents of size 640,000 elements (6.7 MB) the average times are $\chi\alpha o\varsigma$: 39.0 seconds, Xalan XPath: 52.28 seconds. Note the difference in the standard deviations between the two lines (the error bars in the plot). Whereas the standard deviation for $\chi\alpha o\varsigma$ is relatively constant, that of Xalan XPath is fairly high. We shall discuss the cause of this behavior in the next section.

5.2.2 Comparison Excluding Parsing Times

Excluding parsing costs, the performance of our XPath engine is more than twice that of the Xalan engine (Figure 7). This is mainly due to avoiding unnecessary traversals of the tree. Note that the difference in standard deviation is much more apparent in this graph. The cause of this high variance is the bimodal behavior of the Xalan XPath engine. On "good" XPath expressions, where it

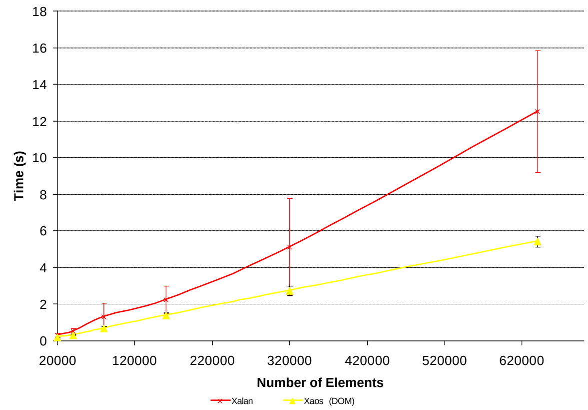

Figure 7: Searching Time in seconds: $\chi\alpha o\varsigma$ versus Xalan

does not perform many unnecessary traversals, the performance of the Xalan XPath engine is similar to that of ours. On "bad" XPath expressions, such as those involving the use of the descendant axes, its performance can be four times worse. Our XPath engine's performance, however, is linear in the size of the XML document and shows little variance.

6 Summary

We have presented a novel algorithm for handling backward and forward XPath axes in a streaming fashion. Our experiments reveal that significant performance benefits can be obtained by using the $\chi\alpha o\varsigma$ algorithm for evaluating XPath expressions on XML documents in a streaming fashion. We are working on extending the $\chi\alpha o\varsigma$ engine to handle more of XPath, building on the framework we have described in this paper.

References

[1] M. Altinel and M. Franklin. Efficient filtering of XML documents for selective dissemination of information. In *Proceedings of VLDB*, Cairo, Egypt, September 2000.

[2] Apache.org. Xalan XSLT stylesheet processor. http://xml.apache.org.

[3] Apache.org. Xerces XML parser. http://xml.apache.org.

[4] T. Bray, J. Paoli, C. Sperberg-McQueen, and E. Maler. Extensible Markup Language (XML) 1.0 (second edition). Technical report, W3C Recommendation, October 6 2000. http://www.w3.org/TR/REC-xml.

[5] C.-Y. Chan, P. Felber, M. Garofalakis, and R. Rastogi. Efficent filtering of XML documents with XPath expressions. In *18th International Conference on Data Engineering*. IEEE, February 2002.

[6] J. Chen, D. DeWitt, F. Tian, and Y. Wang. NiagaraCQ: A scalable continuous query system for internet databases. In *Proceedings of SIGMOD 2000*, pages 379–390, 2000.

[7] J. Clark. XSL transformations (XSLT) version 1.0. Technical report, W3C Recommendation, November 16 1999. http:/www.w3.org/TR/xslt.

[8] J. Clark and S. DeRose. XML path language (XPath) version 1.0. Technical report, W3C Recommendation, November 16 1999. http://www.w3.org/TR/xpath.

[9] Y. Diao, P. Fischer, M. Franklin, and R. To. Yfilter: Efficient and scalable filtering of XML documents. Demo at ICDE, February 2002.

[10] M. Fernandez, J. Marsh, and M. Nagy. XQuery 1.0 and XPath 2.0 data model. Technical report, W3C Working Draft, April 30 2002. http://www.w3.org/TR/xquery-datamodel.

[11] A. L. Hors et al. Document object model (DOM) level 3 core specification. Technical report, W3C Working Draft, April 9 2002. http://www.w3.org/TR/DOM-Level-3-Core.

[12] V. Josifovski, M. Fontoura, and A. Barta. Enabling relational engines to query XML streams. IBM Internal publication.

[13] A. Schmidt et al. The XML benchmark project. Technical Report INS-R0103, CWI, Amsterdam, Netherlands, April 2001.

[14] A. Tozawa and M. Murata. Tableau construction of tree automata from queries on structured documents. IBM Internal publication.

[15] R. Whitmer. Document Object Model (DOM) level 3 XPath specification version 1.0. Technical report, W3C Working Draft, March 28 2002. http://www.w3.org/TR/DOM-Level-3-XPath.

A Rules for Building an X-tree

We represent an $\mathcal{R}xp$ expression as a rooted tree $\mathcal{T} = (V_\mathcal{T}, E_\mathcal{T})$ called X-tree, with labeled vertices and edges. The root is labeled Root. For each *NodeTest* in the expression, there is an x-node in the x-tree labeled with the nodetest. Each x-node (except Root) has a unique incoming edge labeled with the *Axis* specified before the *NodeTest*. One of the x-nodes is designated to be the output x-node. There are functions, $label : V_\mathcal{T} \rightarrow String$, and $axis : E_\mathcal{T} \rightarrow \{ancestor, parent, child, descendant\}$ that return the labels associated with the x-nodes and edges respectively. An x-tree-like structure is also defined for a *RelLocPath*, which is called an x-forest. It consists of two rooted trees, one rooted at Root, and the other rooted at a special x-node labeled context, which, like Root, has no incoming edges. The structure corresponding to a *PredExpr* may either be an x-tree or an x-forest, but none of the x-nodes is designated as an output x-node. The following rules can be used inductively (based on the structure of the $\mathcal{R}xp$) to build a x-tree from an $\mathcal{R}xp$.

$Step ::= Axis :: NodeTest$ The x-forest for $Step$ contains three x-nodes labeled Root, context, and $NodeTest$ (designated as the output node), and an edge from context to $NodeTest$ labeled $Axis$.

$Step ::= Step_1 \,'['\, PredExpr \,']'$ Let T_1 refer to the x-forest resulting from $Step_1$, and T_2 refer to the x-forest or x-tree resulting from $PredExpr$. The x-forest for $Step$ is obtained by merging the output x-node of T_1 with the context x-node of T_2 (if any), and merging the root x-nodes of T_1 and T_2. The output x-node of T_1 is designated as the new output x-node.

$RelLocPath ::= Step \,'/'\, RelLocPath_1$ Let T_1 and T_2 refer to the x-forests obtained from $Step$ and $RelLocPath_1$ respectively. The x-forest for $RelLocPath$ is obtained by merging the output x-node of T_1 with the context x-node of T_2, merging the root x-nodes of T_1 and T_2, and designating the output x-node of T_2 as the new output x-node.

$PredExpr ::= RelLocPath \; and \; PredExpr_1$ Let T_1 and T_2 refer to the structures obtained from *RelLocPath* and $PredExpr_1$ respectively. The x-forest for $PredExpr$ is obtained by merging the context of T_1 with the context of T_2 (if any), and merging the root x-nodes of T_1 and T_2. There is no output vertex.

$PredExpr ::= AbsLocPath \; and \; PredExpr_1$ similar to the previous case.

$AbsLocPath ::= '/'\, RelLocPath$ The x-tree is obtained by merging Root and context x-nodes of the x-forest obtained from $RelLocPath$.