

The Michigan Benchmark: Towards XML Query Performance Diagnostics

Kanda Runapongsa Jignesh M. Patel H. V. Jagadish Yun Chen Shurug Al-Khalifa

University of Michigan
1301 Beal Avenue; Ann Arbor, MI 48109-2122; USA
{krunapon, jignesh, jag, yunc, shurug}@eecs.umich.edu

Abstract

We propose a *micro-benchmark* for XML data management to aid engineers in designing improved XML processing engines. This benchmark is inherently different from application-level benchmarks, which are designed to help users choose between alternative products. We primarily attempt to capture the rich variety of data structures and distributions possible in XML, and to isolate their effects, without imitating any particular application. The benchmark specifies a single data set against which carefully specified queries can be used to evaluate system performance for XML data with various characteristics.

We have used the benchmark to analyze the performance of three database systems: two native XML DBMS, and a commercial ORDBMS. The benchmark reveals key strengths and weaknesses of these systems. We find that commercial relational techniques are effective for XML query processing in many cases, but are sensitive to query rewriting, and require better support for efficiently determining indirect structural containment.

1 Introduction

XML query processing has taken on considerable importance recently, and several XML databases have been constructed on a variety of platforms. There has naturally been an interest in benchmarking the performance of these systems, and a number of benchmarks have been proposed [4, 5, 16]. The focus of currently proposed benchmarks is to assess the performance of a given XML database in performing a variety of representative tasks. Such benchmarks are valuable to potential users of a database system in providing an indication of the performance that the user can

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, requires a fee and/or special permission from the Endowment.

**Proceedings of the 29th VLDB Conference,
Berlin, Germany, 2003**

expect on their specific application. The challenge is to devise benchmarks that are sufficiently representative of the requirements of “most” users. The TPC series of benchmarks accomplished this, with reasonable success, for relational database systems. However, no benchmark has been successful in the realm of ORDBMS and OODBMS which have extensibility and user defined functions that lead to great heterogeneity in the nature of their use. It is too soon to say whether any of the current XML benchmarks will be successful in this respect - we certainly hope that they will.

One aspect that current XML benchmarks do not focus on is the performance of the basic query evaluation operations such as selections, joins, and aggregations. A “micro-benchmark” that highlights the performance of these basic operations can be very helpful to a database developer in understanding and evaluating alternatives for implementing these basic operations. A number of questions related to performance may need to be answered: What are the strengths and weaknesses of specific access methods? Which areas should the developer focus attention on? What is the basis to choose between two alternative implementations? Questions of this nature are central to well-engineered systems. Application-level benchmarks, by their nature, are unable to deal with these important issues in detail. For relational systems, the Wisconsin benchmark [8] provided the database community with an invaluable engineering tool to assess the performance of individual operators and access methods. The work presented in this paper is inspired by the simplicity and the effectiveness of the Wisconsin benchmark for measuring and understanding the performance of relational DBMSs. The goal of this paper is to develop a comparable benchmark for XML DBMSs. The benchmark that we propose to achieve this goal is called the Michigan benchmark.

A challenging issue in designing any benchmark is the choice of the benchmark’s data set. If the data is specified to represent a particular “real application”, it is likely to be quite uncharacteristic for other applications with different data characteristics. Thus, holistic benchmarks can succeed only if they are able to find a real application with data characteristics that are reasonably representative for a large class of different applications.

For a micro-benchmark, the challenges are different. The benchmark data set must be *complex* enough to incorporate data characteristics that are likely to have an impact

on the performance of query operations. However, at the same time, the benchmark data set must be *simple* so that it is not only easy to pose and understand queries against the data set, but also easy to pinpoint the component of the system that is performing poorly. We attempt to achieve this balance by using a data set that has a simple schema but carefully orchestrated structure.

The main contributions of this paper are:

- The identification of XML data characteristics that may impact the performance of XML query processing engines.
- A single heterogeneous data set against which carefully specified queries can be used to evaluate system performance for XML data with various characteristics.
- Insights from running this benchmark on three database systems: a commercial native XML database system, a native XML database system that we have been developing at the University of Michigan, and a commercial object-relational DBMS.

The remainder of this paper is organized as follows: In Section 2, we discuss related work. In Section 3, we present the rationale for the benchmark data set design. In Section 4, we describe the benchmark queries. In Section 5, we present results from using this benchmark on three systems. We conclude with some final remarks in Section 6.

2 Related Work

Several proposals for generating synthetic XML data have been proposed [1, 3]. Abounaga et al. [1] proposed a data generator that accepts as many as 20 parameters to allow a user to control the properties of the generated data. Such a large number of parameters adds a level of complexity that may interfere with the ease of use of a data generator. Furthermore, this data generator does not make available the schema of the data which some systems could exploit. Most recently, Barbosa et al. [3] proposed a template-based data generator for XML, ToXgene, which can generate multiple tunable data sets. The ToXgene user can specify the distribution of different element values in these data sets. In contrast to these previous data generators, the data generator in this proposed benchmark produces an XML data set designed to test different XML data characteristics that may affect the performance of XML engines. In addition, the data generator requires only a few parameters to vary the scalability of the data set. The schema of the data set is also available to exploit.

Four benchmarks [4, 5, 16, 21] have been proposed for evaluating the performance of XML data management systems. XMach-1 [4] and XMark [16] generate XML data that models data from particular Internet applications. In XMach-1 [4], the data is based on a web application that consists of text documents, schema-less data, and structured data. In XMark [16], the data is based on an Internet auction application that consists of relatively structured and data-oriented parts. XOO7 [5] is an XML ver-

sion of the OO7 Benchmark [7], which is a benchmark for OODBMSs. The OO7 schema and instances are mapped into a Document Type Definition (DTD), and the eight OO7 queries are translated into three respective languages for query processing engines: Lore [10, 13], Kweelt [14], and an ORDBMS. Recognizing that different applications requires different benchmarks, Yao et al. [21] have recently proposed, Xbench, which is a family of a number of different application benchmarks.

While each of these benchmarks provides an excellent measure of how a test system would perform against data and queries in their targeted XML application, it is difficult to extrapolate the results to data sets and queries that are different from ones in the targeted domain. Although the queries in these benchmarks are designed to test different performance aspects of XML engines, they cannot be used to perceive the system performance change as the XML data characteristics change. On the other hand, we have different queries to analyze the system performance with respect to different XML data characteristics, such as tree fanout and tree depth; and different query characteristics, such as predicate selectivity.

Finally, we note that [15] presents desiderata for an XML database benchmark, identifies key components and operations, and enumerates ten challenges that the XML benchmark should address. The central focus of this work is application-level benchmarks, rather than micro-benchmarks of the sort we propose.

3 Benchmark Data Set

In this section, we first discuss the characteristics of XML data sets that can have a significant impact on the performance of query operations. Then, we present the schema and the generation algorithm for the benchmark data.

3.1 A Discussion of the Data Characteristics

In a relational paradigm, the primary data characteristics are the selectivity of attributes (important for simple selection operations) and the join selectivity (important for join operations). In an XML paradigm, there are several complicating characteristics to consider, as discussed in Section 3.1.1 and Section 3.1.2.

3.1.1 Depth and Fanout

Depth and fanout are two structural parameters important to tree-structured data. The depth of the data tree can have a significant performance impact, for instance, when computing indirect containment relationships between ancestor and descendant nodes in the tree. Similarly, the fanout of nodes can affect the way in which the DBMS stores the data, and answers queries that are based on selecting children in a specific order (for example, selecting the last child of a node).

One potential way of evaluating the impact of fanout and depth is to generate a number of distinct data sets with different values for each of these parameters and then

Level	Fanout	Nodes	% of Nodes
1	2	1	0.0
2	2	2	0.0
3	2	4	0.0
4	2	8	0.0
5	13	16	0.0
6	13	208	0.0
7	13	2,704	0.4
8	1/13	35,152	4.8
9	2	2,704	0.4
10	2	5,408	0.7
11	2	10,816	1.5
12	2	21,632	3.0
13	2	43,264	6.0
14	2	86,528	11.9
15	2	173,056	23.8
16	-	346,112	47.6

Figure 1: Distribution of the Nodes in the Base Data Set

run queries against each data set. The drawback of this approach is that the large number of data sets makes the benchmark harder to run and understand. Instead, our approach is to fold these into a single data set.

We create a base benchmark data set of a depth of 16. Then, using a “level” attribute, we can restrict the scope of the query to data sets of certain depth, thereby, quantifying the impact of the depth of the data tree. Similarly, we specify high (13) and low (2) fanouts at different levels of the tree as shown in Figure 1. The fanout of 1/13 at level 8 means that every thirteenth node at this level has a single child, and all other nodes are childless leaves. This variation in fanout is designed to permit queries that focus isolating the fanout factor. For instance, the number of nodes is the same (2,704) at levels 7 and 9. Nodes at level 7 have a fanout of 13, whereas nodes at level 9 have a fanout of 2. A pair of queries, one against each of these two levels, can be used to isolate the impact of fanout. In the rightmost column of Figure 1, “% of Nodes” is the percentage of the number of nodes at each level to the number of total nodes in a document.

3.1.2 Data Set Granularity

To keep the benchmark simple, we choose a single large document tree as the default data set. If it is important to understand the effect of document granularity, one can modify the benchmark data set to treat each node at a given level as the root of a distinct document. One can compare the performance of queries on this modified data set against queries on the original data set.

A good benchmark needs to be able to scale in order to measure the performance of databases on a variety of platforms. In the relational model, scaling a benchmark data set is easy - we simply increase the number of tuples. However, with XML, there are many scaling options, such as increasing number of nodes, depth, or fanout. We would like to isolate the effect of the number of nodes from the effects due to other structural changes, such as depth and fanout. We achieve this by keeping the tree depth constant

for all scaled versions of the data set and changing the number of fanouts of nodes at only a few levels, namely levels 5–8. In the design of the benchmark data sets, we deliberately keep the fanout of the bottom few levels of the tree constant. This design implies that the percentage of nodes in the lower levels of the tree (levels 9–16) is nearly constant across all the scaled data sets. This allows us to easily express queries that focus on a specified percentage of the total number of nodes in the database. For example, to select approximately 1/16 of all the nodes, irrespective of the scale factor, we use the predicate `aLevel = 13`.

Due to space limitation, more details regarding the scaling of the benchmark data set are suppressed here. An interested reader can find more details at [20].

3.2 Schema of Benchmark Data

The construction of the benchmark data is centered around the element type `BaseType`. Each `BaseType` element has the following attributes:

1. `aUnique1`: A unique integer generated by traversing the entire data tree in a breadth-first manner. This attribute also serves as the element identifier.
2. `aUnique2`: A unique integer generated randomly.
3. `aLevel`: An integer set to store the level of the node.
4. `aFour`: An integer set to `aUnique2 mod 4`.
5. `aSixteen`: An integer set to `aUnique1 + aUnique2 mod 16`. This attribute is generated using *both* the unique attributes to avoid a correlation between the value of this attribute and other *derived* attributes.
6. `aSixtyFour`: An integer set to `aUnique2 mod 64`.
7. `aString`: A string approximately 32 bytes in length.

The content of each `BaseType` element is a long string that is approximately 512 bytes in length. The generation of the element content and the string attribute `aString` is described in Section 3.3.

In addition to the attributes listed above, each `BaseType` element has two sets of subelements. The first is of type `BaseType`. The number of repetitions of this subelement is determined by the fanout of the parent element, as described in Figure 1. The second subelement is an `OccasionalType`, and can occur either 0 or 1 time. The presence of the `OccasionalType` element is determined by the value of the attribute `aSixtyFour` of the parent element. A `BaseType` element has a nested (leaf) element of type `OccasionalType` if the `aSixtyFour` attribute has the value 0. An `OccasionalType` element has content that is identical to the content of the parent but has only one attribute, `aRef`. The `OccasionalType` element refers to the `BaseType` node with `aUnique1` value equal to the parent’s `aUnique1–11` (the reference is achieved by assigning this value to `aRef` attribute.) In the case where there is no `BaseType` element has the parent’s `aUnique1–11` value (e.g., top few nodes in the tree), the `OccasionalType` element refers to the root node of the tree.

The XML Schema specification of the benchmark data set is available at [20].

3.3 String Attributes and Element Content

The element content of each `BaseType` element is a long string. Since this string is meant to simulate a piece of text in a natural language, it is not appropriate to generate this string from a uniform distribution. Selecting pieces of text from real sources, however, involves many difficulties, such as how to maintain roughly constant size for each string, how to avoid idiosyncrasies associated with the specific source, and how to generate more strings as required for a scaled benchmark. Moreover, we would like to have benchmark results applicable to a wide variety of languages and domain vocabularies.

To obtain string values that have a distribution similar to the distribution of a natural language text, we generate these long strings synthetically, in a carefully stylized manner. We begin by creating a pool of $2^{16} - 1$ (over sixty thousands)¹ synthetic words. The words are divided into 16 buckets, with exponentially growing bucket occupancy. Bucket i has 2^{i-1} words. For example, the first bucket has only one word, the second has two words, the third has four words, and so on. Each made-up word contains information about the bucket from which it is drawn and the word number in the bucket. For example, “15twenty9B14” indicates that this is the 1,529th word from the fourteenth bucket. To keep the size of the vocabulary in the last bucket at roughly 30,000 words, words in the last bucket are derived from words in the other buckets by adding the suffix “ing” (to get exactly 2^{15} words in the sixteenth bucket, we add the dummy word “oneB0ing”).

The value of the long string is generated from the template shown in Figure 2, where “PickWord” is actually a placeholder for a word picked from the word pool described above. To pick a word for “PickWord”, a bucket is chosen, with each bucket equally likely, and then a word is picked from the chosen bucket, with each word equally likely. Thus, we obtain a discrete Zipf distribution of parameter roughly 1. We use the Zipf distribution since it seems to reflect word occurrence probabilities accurately in a wide variety of situations. The value of `aString` attribute is simply the first line of the long string that is stored as the element content.

4 Benchmark Queries

In creating the data set above, we make it possible to tease apart data with different characteristics, and to issue queries with well-controlled yet vastly differing data access patterns. We are more interested in evaluating the cost of individual pieces of core query functionality than in evaluating the composite performance of queries that are of application-level. Knowing the costs of individual basic

¹Roughly twice the number of entries in the second edition of the Oxford English Dictionary. However, half the words that are used in the benchmark are “derived” words, produced by appending “ing” to the end of a word.

```
Sing a song of PickWord,  
A pocket full of PickWord  
Four and twenty PickWord  
All baked in a PickWord.  
  
When the PickWord was opened,  
The PickWord began to sing;  
Wasn't that a dainty PickWord  
To set before the PickWord?  
  
The King was in his PickWord,  
Counting out his PickWord;  
The Queen was in the PickWord  
Eating bread and PickWord.  
  
The maid was in the PickWord  
Hanging out the PickWord;  
When down came a PickWord,  
And snipped off her PickWord!
```

Figure 2: Generation of the String Element Content

operations, we can estimate the cost of any complex query by just adding up relevant piecewise costs (keeping in mind the pipelined nature of evaluation, and the changes in sizes of intermediate results when operators are pipelined).

We find it useful to refer to simple queries as “selection queries”, “join queries” and the like, to clearly indicate the functionality of each query. A complex query that involves many of these simple operations can take time that varies monotonically with the time required for these simple components.

In the following subsections, we describe the benchmark queries in detail. In these query descriptions, the types of the nodes are assumed to be `BaseType` unless specified otherwise.

4.1 Selection

Relational selection identifies the tuples that satisfy a given predicate over its attributes. XML selection is both more complex and more important because of the tree structure. Consider a query, against a bibliographic database, that seeks books, published in the year 2002, by an author with name including the string “Blake”. This apparently straightforward selection query involves matches in the database to a 4-node “query pattern”, with predicates associated with each of these four (namely `book`, `year`, `author`, and `name`). Once a match has been found for this pattern, we may be interested in returning only the `book` element, all the nodes that participated in the match, or various other possibilities. We attempt to organize the various sources of complexity in the following.

4.1.1 Returned Structure

In a relation, once a tuple is selected, the tuple is returned. In XML, as we saw in the example above, once an element

is selected, one may return the element, as well as some structure related to the element, such as the sub-tree rooted at the element. Query performance can be significantly affected by how the data is stored and when the returned result is materialized.

To understand the role of returned structure in query performance, we use the query, “Select all elements with aSixtyFour = 2.” The selectivity of this query is $1/64$ (1.6%)².

This query is run in the following cases:

- **QR1.** Return only the elements in question, not including any subelements.
- **QR2.** Return the elements and all their immediate children.
- **QR3.** Return the entire sub-tree rooted at the elements.
- **QR4.** Return the elements and their selected descendants with aFour = 1.

The remaining queries in the benchmark simply return the unique identifier attributes of the selected nodes (aUnique1 for BaseType and aRef for OccasionalType), except when explicitly specified otherwise. This design choice ensures that the cost of producing the final result is a small portion of the query execution cost.

4.1.2 Simple Selection

Even XML queries involving only one element and few predicates can show considerable diversity. We examine the effect of this simple selection predicate in this set of queries.

- **Exact Match Attribute Value Selection**

Value-based selection on a string attribute.

QS1. Low selectivity. Select nodes with aString = “Sing a song of oneB4”. Selectivity is 0.8%.

QS2. High selectivity. Select nodes with aString = “Sing a song of oneB1”. Selectivity is 6.3%.

Value-based selection on an integer attribute.

These following queries have almost the same selectivities as the above string attribute queries.

QS3. Low selectivity. Select nodes with aLevel = 10. Selectivity is 0.7%.

QS4. High selectivity. Select nodes with aLevel = 13. Selectivity is 6.0%.

Selection on range values.

QS5. Select nodes with aSixtyFour between 5 and 8. Selectivity is 6.3%.

Selection with sorting.

QS6. Select nodes with aLevel = 13 and have the returned nodes sorted by aSixtyFour attribute. Selectivity is 6.0%.

Multiple-attribute selection.

QS7. Select nodes with attributes aSixteen = 1 and aFour = 1. Selectivity is 1.6%.

- **Element Name Selection**

QS8. Select nodes with the element name eOccasional. Selectivity is 1.6%.

- **Order-based Selection**

QS9. High fanout. Select the second child of every node with aLevel = 7. Selectivity is 0.4%.

QS10. Low fanout. Select the second child of every node with aLevel = 9. Selectivity is 0.4%.

Since the fraction of nodes in these two queries are the same, the performance difference between them is likely to be on account of fanout.

- **Element Content Selection**

QS11. Low selectivity. Select OccasionalType nodes that have “oneB4” in the element content. Selectivity is 0.2%.

QS12. High selectivity. Select nodes that have “oneB4” as a substring of element content. Selectivity is 12.5%.

- **String Distance Selection**

QS13. Low selectivity. Select all nodes with element content that the distance between keyword “oneB5” and keyword “twenty” is not more than four. Selectivity is 0.8%.

QS14. High selectivity. select all nodes with element content that the distance between keyword “oneB2” and keyword “twenty” is not more than four. Selectivity is 6.3%.

4.1.3 Structural Selection

Selection in XML is often based on patterns. Queries should be constructed to consider multi-node patterns of various sorts and selectivities. All queries listed in this section return only the root of the selection pattern, unless specified otherwise. In these queries, the selectivity of a predicate is noted following the predicate.

- **Order-Sensitive Parent-Child Selection**

QS15. Local ordering. Select the second element below *each* element with aFour = 1 (sel=1/4) if that second element also has aFour = 1 (sel=1/4). Selectivity is 3.1%.

QS16. Global ordering. Select the second element with aFour = 1 (sel=1/4) below *any* element with aSixtyFour = 1 (sel=1/64). This query returns at most one element, whereas the previous query returns one for each parent.

QS17. Reverse ordering. Among the children with aSixteen = 1 (sel=1/16) of the parent element with aLevel = 13 (sel=6.0%), select the last child. Selectivity is 0.7%.

²Detailed computation of the query selectivities can be found in [20].

- **Parent-Child Selection**

QS18. Medium selectivity of both parent and child. Select nodes with `aLevel = 13` (`sel=6.0%`, approx. 1/16) that have a child with `aSixteen = 3` (`sel=1/16`). Selectivity is approximately 0.7%.

QS19. High selectivity of parent and low selectivity of child. Select nodes with `aLevel = 15` (`sel=23.8%`, approx. 1/4) that have a child with `aSixtyFour = 3` (`sel=1/64`). Selectivity is approximately 0.7%.

QS20. Low selectivity of parent and high selectivity of child. Select nodes with `aLevel = 11` (`sel=1.5%`, approx. 1/64) that have a child with `aFour = 3` (`sel=1/4`). Selectivity is approximately 0.7%.

- **Ancestor-Descendant Selection**

QS21. Medium selectivity of both ancestor and descendant. Select nodes with `aLevel = 13` (`sel=6.0%`, approx. 1/16) that have a descendant with `aSixteen = 3` (`sel=1/16`). Selectivity is 3.5%.

QS22. High selectivity of ancestor and low selectivity of descendant. Select nodes with `aLevel = 15` (`sel=23.8%`, approx. 1/4) that have a descendant with `aSixtyFour = 3` (`sel=1/64`). Selectivity is 0.7%.

QS23. Low selectivity of ancestor and high selectivity of descendant. Select nodes with `aLevel = 11` (`sel=1.5%`, approx. 1/64) that have a descendant with `aFour = 3` (`sel=1/4`). Selectivity is 1.5%.

- **Ancestor Nesting in Ancestor-Descendant Selection**

In the ancestor-descendant queries above (QS21-QS23), ancestors are never nested below other ancestors. To test the performance of queries when ancestors are recursively nested below other ancestors, we have three other ancestor-descendant queries. These queries are variants of QS21-QS23.

QS24. Medium selectivity of both ancestor and descendant. Select nodes with `aSixteen = 3` (`sel=1/16`) that have a descendant with `aSixteen = 5` (`sel=1/16`).

QS25. High selectivity of ancestor and low selectivity of descendant. Select nodes with `aFour = 3` (`sel=1/4`) that have a descendant with `aSixtyFour = 3` (`sel=1/64`).

QS26. Low selectivity of ancestor and high selectivity of descendant. Select nodes with `aSixtyFour = 9` (`sel=1/64`) that have a descendant with `aFour = 3` (`sel=1/4`).

QS27. Similar to query QS26, but return both the root node and the descendant node of the selection pattern. Thus, the returned structure is a pair of nodes with an inclusion relationship between them.

The overall selectivities of these queries (QS24-QS26) cannot be the same as that of the “equivalent” unnested

queries (QS21-QS23) for two situations – first, the same descendants can now have multiple ancestors they match, and second, the number of candidate descendants is different (fewer) since the ancestor predicate can be satisfied by nodes at any level (and will predominantly be satisfied by nodes at levels 15 and 16, due to their large numbers). These two effects may not necessarily cancel each other out. We focus on the local predicate selectivities and keep these the same for all of these queries (as well as for the parent-child queries considered before).

- **Complex Pattern Selection**

Complex pattern matches are common in XML databases, and in this section, we introduce a number of *chain* and *twig* queries that we use in this benchmark. Figure 3 shows an example for these query types. In the figure, each node represents a predicate such as an element tag name predicate, or an attribute value predicate, or an element content match predicate. A structural parent-child relationship in the query is shown by a single line, and an ancestor-descendant relationship is represented by a double-edged line. The chain query shown in the Figure 4(i) finds all nodes matching condition A, such that there is a child matching condition B, such that there is a child matching condition C, such that there is a child matching condition D. The twig query shown in the Figure 4(ii) matches all nodes that satisfy condition A, and have a child node that satisfies condition B, and also have a descendant node that satisfies condition C.




Figure 3: Samples of Chain and Twig Queries

The benchmark uses the following complex queries:

- **Parent-Child Complex Pattern Selection**

QS28. One chain query with three parent-child joins with the selectivity pattern: high-low-low-high. The query is to test the choice of join order in evaluating a complex query. To achieve the desired selectivities, we use the following predicates: `aFour=3` (`sel=1/4`), `aSixteen=3` (`sel=1/16`), `aSixteen=5` (`sel=1/16`) and `aLevel=16` (`sel=47.6%`).

QS29. One twig query with two parent-child joins with the selectivity pattern: low-high, low-low. Select parent nodes with `aLevel = 11` (`sel=1.5%`) that

have a child with `aFour = 3` ($\text{sel}=1/4$), and another child with `aSixtyFour = 3` ($\text{sel}=1/64$).

QS30. One twig query with two parent-child joins with the selectivity pattern: high-low, high-low. Select parent nodes with `aFour = 1` ($\text{sel}=1/4$) that have a child with `aLevel = 11` ($\text{sel}=1.5\%$) and another child with `aSixtyFour = 3` ($\text{sel}=1/64$).

- **Ancestor-Descendant Complex Pattern Selection**

QS31-QS33. Repeat queries QS28-QS30, but using ancestor-descendant in place of parent-child.

QS34. One twig query with one parent-child join and one ancestor-descendant join. Select nodes with `aFour = 1` ($\text{sel}=1/4$) that have a child of nodes with `aLevel = 11` ($\text{sel}=1.5\%$) and a descendant with `aSixtyFour = 3` ($\text{sel}=1/64$).

- **Negated Selection**

In XML, some elements are optional and some queries test the existence of these optional elements. Negated selection query selects elements which does not contain a descendant that is an optional element.

QS35. Find all `BaseType` elements below where there is no `OccasionalType` element.

4.2 Value-Based Join

A value-based join involves comparing values at two different nodes that need not be related structurally. In computing the value-based joins, one would naturally expect *both* nodes participating in the join to be returned. As such, the return structure is the pair of the `aUnique1` attributes of nodes joined.

QJ1. Low selectivity. Select nodes with `aSixtyFour = 2` ($\text{sel}=1/64$) and join with themselves based on the equality of `aUnique1` attribute. The selectivity of this query is approximately 1.6%.

QJ2. High selectivity. Select nodes based on `aSixteen = 2` ($\text{sel}=1/16$) and join with themselves based on the equality of `aUnique1` attribute. The selectivity of this query is approximately 6.3%.

4.3 Pointer-Based Join

These queries specify joins using references that are specified in the DTD or XML Schema, and the implementation of references may be optimized with logical OIDs in some XML databases.

QJ3. Low selectivity. Select all `OccasionalType` nodes that point to a node with `aSixtyFour = 3` ($\text{sel}=1/64$). Selectivity is 0.02%.

QJ4. High selectivity. Select all `OccasionalType` nodes that point to a node with `aFour = 3` ($\text{sel}=1/4$). Selectivity is 0.4%.

Both of these pointer-based joins are semi-join queries. The returned elements are only the `eOccasional` nodes, not the nodes pointed to.

4.4 Aggregation and Update Queries

The benchmark also contains three value-based aggregate queries, four structure-based aggregate queries, and seven update queries, which include point update and delete, bulk insert and delete, bulk load, bulk reconstruction, and bulk restructuring. In the interest of space, we omit these queries in this paper and refer an interested reader to [20].

5 The Benchmark in Action

In this section, we present and analyze the performance of different databases using the Michigan benchmark. We conducted experiments using one native commercial XML DBMS, a university native XML DBMS, and one leading commercial ORDBMS. Due to the nature of the licensing agreement for the commercial systems, we can not disclose the actual names of the system, and will refer to the commercial native system as **CNX**, and the commercial ORDBMS as **COR**.

The native XML database is Timber [12], a university native XML DBMS system that we are developing at the University of Michigan. Timber uses the Shore storage manager [6], and implements various join algorithms, query size estimation, and query optimization techniques that have been developed for XML DBMSs.

COR is provided by a leading database vendor, and we used the Hybrid inlining algorithm to map the data into a relational schema [19]. To generate good SQL queries, we adopted the algorithm presented in [9]. The queries in the benchmark were converted into SQL queries (sanitized to remove any system-specific keywords in the query language) which can be found at [20].

The commercial native XML system provides an XPath interface and a recently released XQuery interface. We started by writing the benchmark queries in XQuery. However, we found that the XQuery interface was unstable and in most cases would hang up either the server or the Java client, or run out of memory on the machine. Unfortunately, no other interface is available for posing XQueries to this commercial system. Consequently, we reverted to writing the queries using XPath expression, which implies that join queries are written as nested XPath expressions, that get evaluated using a nested-loops paradigm. In all the cases that we could run queries using the XQuery interface, the XPath approach was faster. Consequently, all the numbers reported here are written as XPath queries. The actual queries for all these systems (sanitized to remove any system-specific keywords in the query language) are available at the website for this benchmark [20].

5.1 Experimental Platform and Methodology

All experiments were run on a single-processor 550 MHz Pentium III machine with 256 MB of main memory. The benchmark machine was running the Windows2000 and

was configured with a 20 GB IDE disk. All three systems were configured to use a 64 MB buffer pool size.

5.1.1 System Setup

For both commercial systems, we used default settings that the system chooses during the software installation. The only setting that we changed for COR was to enable the use of hash joins, as we found that query response times generally improved with this option turned on. For COR, after loading the data we update all statistics to provide the optimizer with the most current statistical information.

For CNX, we tried running the queries with and without indices. CNX permits building both structure (i.e. path indices), and value indices. Surprisingly, we found that indexing did not help the performance of the benchmark queries, and in fact in most cases actually reduced the performance of the queries. In very few cases, the performance improved but by less than 20% over the non-indexed case. The reason for the ineffectiveness of the index is that CNX indexing can not handle the “//” operator, which is used often in the benchmark. Furthermore, the index is not effective on `BaseType` element as it is recursively nested below other `BaseType` elements. In the interest of space, we only present the queries using the non-indexed version, and refer the interested reader to [20] for a more detailed analysis.

5.1.2 Data Sets

For this experiment we loaded the base data set (739K nodes and 500MB of raw data). Although we wanted to load larger scaled up data sets, we found that in many cases the parsers are fragile and break down with large documents. Consequently, for this study, we decided to load another *scaled down* version of the data. The scaled down data set, which we call `ds0.1x`, is produced by changing the fanouts of the nodes at levels 5, 6, 7 and 8 to 4, 4, 4 and 1/4 respectively. This scaled down data set is approximately 1/10th the size of the `ds1x` data set. Note that because of the nature of the document tree, the percentage of nodes at the levels close to the leaves remains the same, hence the query selectivities stay roughly constant even in this scaled down data set.

5.1.3 Measurements

In our experiments, each query was executed five times, and the execution times reported in this section is an average of the middle three runs. Queries were always run in “cold” mode, so the query execution times do not include side-effects of cached buffer pages from previous runs.

5.1.4 Benchmark Results

In our own use of the benchmark, we have found it useful to produce two kinds of tables: a *summary* table which

presents a single number for a group of related queries, and a *detail* table that shows the query execution time for each individual query. The summary table presents a high-level view of the performance of the benchmark queries. It contains one entry for a *group* of related queries, and shows the geometric mean of the response times of the queries in that group. Figure 4 shows the summary table for the systems we benchmarked and also indicates the sections in which the detailed numbers are presented and analyzed. In the figures *N/A* indicates that queries could not be run with the given configuration and system software.

From Figure 4, we observe that Timber is very efficient at processing XML structural queries. The implementation of “traditional” relational-style queries such as value-based joins is not highly tuned in Timber. This is primarily because Timber is a research prototype and most of the development attention has been paid to the XML query processing issues that are not covered by traditional relational techniques.

COR can execute almost all queries well, except for the ancestor-descendant relationship queries. COR performs very well on the parent-child queries, which are evaluated using foreign key joins.

From Figure 4, we observe that CNX is usually slower than the other two systems. A large query overhead, of about 2 secs, is incurred by CNX, even for very small queries. CNX is considerably slower than Timber, and faster than COR only on the ancestor-descendant queries.

5.2 Detailed Performance Analysis

In this section, we analyze the impact of various factors on the performance that was observed. The details of the performance of the queries on each of the systems is shown in Figure 5. In the interest of space, we only present the detailed numbers for a subset of the queries in the benchmark. A full analysis of these systems using the entire benchmark can be found at [20].

5.2.1 Returned Structure (QR1-QR4)

Examining the performance of the returned structure queries, QR1-QR4 in Figure 5, we observe that the returned structure has an impact on all systems. Timber performs the worst when the whole subtree is returned (QR3). This is surprising since Timber stores elements in depth-first order, so that retrieving a sub-tree should be a fast sequential scan. It turns out that Timber uses SHORE [6] for low level storage and memory management, and the initial implementation of the Timber data manager makes one SHORE call per element retrieved. The poor performance of QR3 helped Timber designers identify this implementation weakness, and to begin taking steps to address it.

COR takes more time in selecting and returning descendant nodes (QR3 and QR4), than returning children nodes

Discussed Section	Query Group (Queries in Group)	Geometric Mean Response Times (seconds)					
		ds0.1x			ds1x		
		CNX	Timber	COR	CNX	Timber	COR
5.2.1	Returned structure (QR1-QR4)	2.53	0.06	0.27	9.40	0.37	3.87
5.2.2	Exact match attribute value Selection (QS1-QS7)	2.25	0.04	0.04	6.73	0.48	0.28
See [20]	Element name selection (QS8)	2.12	0.01	0.02	5.98	0.08	0.15
See [20]	Order-based selection (QS9-QS10)	2.18	0.00	0.06	6.25	0.05	0.61
See [20]	Element content selection (QS11-QS12)	2.47	0.07	0.12	7.01	0.69	2.75
See [20]	String distance selection (QS13-QS14)	N/A	3.00	1.12	N/A	32.92	39.52
See [20]	Order-sensitive selection (QS15-QS17)	2.28	0.02	0.08	6.75	0.54	0.26
5.2.3	Parent-child (P-C) selection (QS18-QS20)	2.36	0.13	0.04	6.89	1.38	0.34
5.2.3	Ancestor-descendant (A-D) selection (QS21-QS23)	2.68	0.14	2.14	7.74	1.41	17.92
5.2.3	Ancestor nesting in A-D selection (QS24-QS26)	2.76	0.13	1.16	8.02	1.39	12.65
See [20]	P-C complex pattern selection (QS27-QS30)	3.26	0.25	0.03	7.31	2.56	0.50
See [20]	A-D complex pattern selection (QS31-QS34)	3.64	0.27	2.49	10.63	2.75	24.74
5.2.4	Negated selection (QS35)	2.84	1.29	2.10	66.15	12.58	23.38
5.2.5	Value-based join (QJ1-QJ2)	359.17	1.72	0.05	2537.92	18.82	0.42
5.2.5	Pointer-based join (QJ3-QJ4)	161.80	3.15	0.02	1339.54	19.73	0.14

Figure 4: Benchmark Numbers for Three DBMSs. CNX - a commercial native XML DBMS, Timber - a university native XML DBMS, and COR - a commercial ORDBMS. N/A indicates that the queries could not be run on that system.

(QR2). This is because COR exploits the indices on the primary keys and the foreign keys when it retrieves the children nodes. On the other hand, COR needs to call recursive SQL statements in retrieving the descendant nodes.

CNX produces results in reasonably short times. Surprisingly, returning the result element itself (QR1) takes a little more time than returning the element and its immediate children (QR2). We suspect that CNX returns the entire subtree by default, and then will carry out post-processing selection to return the element itself, which takes more time. This also explains why CNX incurs more query processing time than the other DBMSs in most queries.

5.2.2 Exact Attribute Value Selection (QS1-QS4)

Selectivity has an impact on both Timber and COR. The response times of the high selectivity queries (QS2 and QS4) are more than those of the low selectivity queries (QS1 and QS3), with the response times growing linearly with the increasing selectivity for both systems. In both systems, selection on short strings is as efficient as selection on integers.

Overall, CNX does not perform as well as the other two DBMSs. It is interesting to notice that although selectivity does have some impact on CNX, it's not as strong as on the other two DBMSs (this is true even with indexing in CNX [20]). Although the response times of the high selectivity queries (QS2, QS4) are higher than the response times of the low selectivity queries (QS1, QS3), the difference does not reflect a linear growth. Selection on short strings takes a little more time than that on integers, although the difference is negligible.

Discussion of the queries QS5-QS7 is omitted here, and can be found at [20].

Structural Selection (QS18-QS35)

Figure 5 shows the performance of selected structural selection queries, QS18-QS26, and QS35. In this figure, “P-C:low-high” refers to the join between a parent with low selectivity and a child with high selectivity, whereas, “A-D:high-low” refers to the join between an ancestor with high selectivity and a descendant with low selectivity.

5.2.3 Simple Containment Selection (QS18-QS26)

As seen from the results for the direct containment queries (QS18-QS20) in Figure 5, the COR processes direct containment queries better than Timber, but Timber handles indirect containment queries (QS21-QS26) better.

CNX underperforms as compared to the other systems on direct containment queries (QS18-QS20). However, on indirect containment queries (QS21-QS26), it often performs better than COR. CNX only has slightly performance on direct containment queries (QS18-QS20) than on indirect containment queries (QS21-QS27). Examining the effect of query selectivities on CNX query execution (see QS21-QS23 as an example), we notice that the execution times are relatively immune to the query selectivities, implying that the system does not effectively exploit the differences in query selectivities in picking query plans.

In Timber, structural joins [2] are used to evaluate both types of containment queries. Each structural join reads both inputs (ancestor/parent and descendant/child) once from indices. It keeps potential ancestors in a stack and joins them with the descendants as the descendants arrive. Therefore, the cost of the ancestor-descendant queries is not necessarily higher than the parent-child queries. From the performance of these queries, we can deduce that the higher selectivity of ancestors, the greater the delay in the

Query	Query Description	Sel.(%)	Response Times (seconds)					
			ds0.1x			ds1x		
			CNX	Timber	COR	CNX	Timber	COR
QR1	Return result element	1.6	2.18	0.01	0.02	6.19	0.08	0.16
QR2	Return element and immediate children	1.6	1.68	0.02	0.31	4.83	0.27	2.59
QR3	Return entire sub-tree	1.6	3.63	0.26	1.09	10.17	> 1 hr	26.09
QR4	Return element and selected descendants	1.6	2.37	0.19	0.97	7.14	2.44	20.57
QS1	Selection on string attribute value (low sel.)	0.8	1.99	0.003	0.02	6.06	0.05	0.08
QS2	Selection on string attribute value (high sel.)	6.3	2.05	0.03	0.06	6.21	0.34	0.63
QS3	Selection on integer attribute value (low sel.)	0.7	2.25	0.01	0.02	6.76	0.04	0.08
QS4	Selection on integer attribute value (high sel.)	6.0	2.28	0.03	0.06	6.82	0.30	0.57
QS18	P-C: medium-medium	0.7	2.30	0.08	0.02	6.73	0.85	0.25
QS19	P-C: high-low	0.7	2.55	0.17	0.05	7.40	1.79	0.44
QS20	P-C: low-high	0.7	2.23	0.17	0.05	6.58	1.73	0.34
QS21	A-D:medium-medium	3.5	2.73	0.09	2.22	7.77	0.93	20.15
QS22	A-D:high-low	0.7	2.57	0.18	0.94	7.40	1.72	5.64
QS23	A-D:low-high	1.5	2.73	0.16	4.69	8.06	1.74	50.65
QS24	Ancestor nesting in A-D:medium-medium	1.0	2.56	0.08	2.16	7.32	0.87	20.03
QS25	Ancestor nesting in A-D:high-low	1.7	3.68	0.19	0.95	10.44	1.94	8.83
QS26	Ancestor nesting in A-D:low-high	0.5	2.23	0.15	0.92	6.74	1.61	11.73
QS35	Negated selection	93.2	2.84	1.29	2.10	66.15	12.58	23.38
QJ1	Value-based join (low sel.)	1.6	187.5	0.69	0.03	1268.4	18.82	0.21
QJ2	Value-based join (high sel.)	6.3	687.9	4.29	0.08	> 1 hr	> 1 hr	0.83
QJ3	Pointer-based join (low sel.)	0.02	160.1	0.73	0.01	1307.6	20.08	0.05
QJ4	Pointer-based join (high sel.)	0.4	163.5	13.63	0.05	1354.2	19.38	0.42

Figure 5: Detailed Benchmark Numbers for Selected Queries. (See [20] for a full list)

query performance (QS19 and QS22). Although the selectivities of ancestors in QS20 and QS23 are lower than those of QS18 and QS21, QS20 and QS23 perform worse because of the high selectivities of descendants.

COR is very efficient for processing parent-child queries (QS18-QS20), since these translate into foreign key joins, which the system can evaluate very efficiently using indices. On the other hand, the COR has much longer response times for the ancestor-descendant queries (QS21-QS23). The only way to answer these queries is by using recursive SQL statements, which are expensive to evaluate.

We also found that the performance of the ancestor-descendant queries was very sensitive to the SQL query that we wrote. To answer an ancestor-descendant query with predicates on both the ancestor and descendant nodes, the SQL query must perform the following three steps: 1) Start by selecting the ancestor nodes, which could be performed by using an index to quickly evaluate the ancestor predicate, 2) Use a recursive SQL to find all the descendants of the selected ancestor nodes, and 3) Finally, check if the selected descendants match the descendant predicate specified in the query. Note that one cannot perform step 3 before step 2 since it is possible that a descendant in the result may be below another descendant that does not match the predicate on the descendant node in the query. Another alternative is to start step 1 by selecting the descendant nodes and following these steps to find the matching ancestors. In

general, it is more effective to start from the descendants if the descendant predicate has a lower selectivity than the ancestor predicate. However, if the ancestor predicate has a lower selectivity, then one needs to pick the strategy that traverses fewer number of nodes. Traversing from the descendants, the number of visited nodes grows proportional to the distance between the descendants and the ancestors. However, traversing from the ancestors, the number of visited nodes can grow exponentially at the rate of the fanout of the ancestors.

The effect of the number of visited nodes in COR is clearly seen by comparing queries QS23 and QS26. Both queries have similar selectivities on both ancestors and descendants, and are evaluated by starting from the ancestors. However, QS23 has a much higher response time. In QS23, starting from the ancestors, the number of visited nodes grow significantly since the ancestors are the nodes at level 11 – each node at this level, and its expanded non-leaf descendant nodes, has a fanout of 2. In contrast, in QS26, the ancestor set is the set of nodes that satisfy the predicate $aSixtyFour = 9$. Since half of these ancestors are at the leaf level, when finding the descendants, the number of visited nodes does not grow as quickly as it did in the case of query QS23.

One may wonder whether QS23 would perform better if the query was coded to start from the descendants. We found that for the ds1x data set, using this option nearly

doubles the response time to 126.01 seconds. Starting from the ancestors results in better performance since the selectivity of the descendants ($sel=1/4$) is higher than the selectivity of the ancestors ($sel=1/64$), which implies that the descendent candidate list is much larger than the ancestor candidate list. Consequently, starting from the descendants results in visiting more number of nodes.

All systems are immune to the recursive nesting of ancestor nodes below other ancestor nodes; the queries on recursively nested ancestor nodes (QS24-QS26) have the same response times as their non-recursive counterparts (QS21-QS23), except QS26 and QS23 that have different response times in COR.

5.2.4 Irregular Structure (QS35)

Since some parts of an XML document may have irregular data structure, such as missing elements, queries such as QS35 are useful when looking for such irregularities. Query QS35 looks for all `BaseType` elements below which there is no `OccasionalType` element.

While looking for irregular data structures, CNX performs reasonably well on the small scale database, but as one might notice, it does not scale very well like with other queries. The selectivity of this query is fairly high (93.2%), and as the database size increases, the return result grows dramatically. CNX seems to spend a large part of its execution time in processing the results at the client, and this part does not seem to scale very well.

In Timber, this operation is very fast because it uses a variation of the structural joins used in evaluating containment queries. This join outputs ancestors that *do not* have a matching descendant.

In COR, there are two ways to implement this query. A naive way is to use a set difference operation which results in a very long response time (1517.4 seconds for the `ds1x` data set). This long response time is because the COR first needs to find a set of elements that contain the missing elements (using a recursive SQL query), and then find elements that are not in that set. The second alternative of implementing this query is to use a left outer join. That is first create a view that selects all `BaseType` elements have some `OccasionalType` descendants (this requires a recursive SQL statement). Then compute a left-outer join between the view and the relation that holds all `BaseType` elements, selecting only those `BaseType` elements that are not present in the view (this can be accomplished by checking for a null value). Compared to the response time of the first implementation (1517.4 seconds), this rewriting query results in much less response time (23.38 seconds) as reported in Figure 5.

5.2.5 Value-Based and Pointer-Based Joins (QJ1-QJ4)

The performance of the value-based join queries, QJ1-QJ4, is shown in Figure 5. Both CNX and Timber show poor

performance on these “traditional” join queries. In Timber, a simple, unoptimized nested loop join algorithm is used to evaluate value-based joins. Both QJ1 and QJ2 perform poorly because of the high overhead in retrieving attribute values through random accesses.

COR performs well on this class of queries, which are evaluated using foreign-key joins which are very efficiently implemented in traditional commercial database systems.

5.3 Performance Analysis on Scaling Databases

In this Section, we discuss the performance of the three systems as the data set is scaled from `ds0.1x` to `ds1x`. Please refer to Figure 4 for the performance comparisons between these two data sets.

5.3.1 Scaling Performance on CNX

In almost all of the queries, the ratios of the response times when using `ds0.1x` over `ds1x` are less than or around 10, except for QS35, which consists of nested aggregate `count()` function.

5.3.2 Scaling Performance on Timber

Timber scales linearly for all queries, with a response time ratio of approximately 10, with two exceptions. Where large return result structures have to be constructed, Timber is inefficient, and scales poorly, as discussed above in Sec. 5.2.1. Also, the value-based join implementation is naive, and scales poorly.

5.3.3 Scaling Performance on COR

Once more, with two exceptions, the ratios of the response times when using `ds0.1x` over `ds1x` are approximately 10, showing linear scale-up.

QR3 and QR4 require result XML reconstruction with descendant access, and have response times grow about 20 times as data size increases about 10 times. Recently, Shanmugasundaram et al. [17, 18] have addressed this problem as they proposed techniques for efficiently publishing and querying XML view of relational data. However, these techniques were not implemented in COR.

6 Conclusions

We proposed a benchmark that can be used to identify individual data characteristics and operations that may affect the performance of XML query processing engines. With careful analysis of the benchmark queries, engineers can diagnose the strengths and weaknesses of their XML databases. In addition, engineers can try different query processing implementations and evaluate these alternatives with the benchmark. Thus, this benchmark is a simple and effective tool to help engineers improve system performance.

We have used the benchmark to evaluate three XML systems: a commercial XML system, Timber, and a commercial Object-Relational DBMS. The results show that the commercial native XML system has substantial room for performance improvement on most of the queries. The benchmark has already become an invaluable tool in the development of the Timber native XML database, helping us identify portions of the system that need performance tuning. Consequently, on most benchmark queries Timber outperforms the other systems. A notable exception to this behavior is the poor performance of Timber on traditional value-based join queries.

This benchmarking effort also shows that the ORDBMS is sensitive to the method used to translate an XML query to SQL. While this has been shown to be true for some XML queries in the past [9, 18], we show that this is also true for simple indirect containment queries, and queries that search for irregular structures. We also demonstrate that using recursive SQL one can evaluate any structural query in the benchmark, however, this is much more expensive in the ORDBMS than the implementations in Timber, which use efficient XML structural join algorithms.

Finally, we note that the proposed benchmark meets the key criteria for a successful domain-specific benchmark that have been proposed in [11]. These key criteria are: relevant, portable, scalable, and simple. The proposed Michigan benchmark is *relevant* to testing the performance of XML engines because proposed queries are the core basic components of typical application-level operations of XML application. Michigan benchmark is *portable* because it is easy to implement the benchmark on many different systems. In fact, the data generator for this benchmark data set is freely available for download from the Michigan benchmark's web site [20]. It is *scalable* through the use of a scaling parameter. It is *simple* since it comprises only one data set and a set of simple queries, each designed to test a distinct functionality.

References

- [1] A. Aboulmaga and J. Naughton and C. Zhang. Generating Synthetic Complex-structured XML Data. In *WebDB*, 2001.
- [2] S. Al-Khalifa, H. V. Jagadish, N. Koudas, J. Patel, D. Srivastava, and Y. Wu. Structural Joins: A Primitive for Efficient XML Query Processing Pattern Matching. In *ICDE*, 2002.
- [3] D. Barbosa, A. Mendelzon, J. Keenleyside, and K. Lyons. ToXgene: An Extensible Templated-based Data Generator for XML. In *WebDB*, 2002.
- [4] T. Böhme and E. Rahm. XMach-1: A Benchmark for XML Data Management. In *BTW*, 2001.
- [5] S. Bressan and G. Dobbie and Z. Lacroix and M. L. Lee and Y. G. Li and U. Nambiar and B. Wadhwa. XOO7: Applying OO7 Benchmark to XML Query Processing Tools. In *CIKM*, 2001.
- [6] M. Carey, D. J. DeWitt, M. J. Franklin, N. E. Hall, M. McAuliffe, J. F. Naughton, D. T. Schuh, and M. H. Solomon. Shoring up Persistent Applications. In *SIGMOD*, 1994.
- [7] M. J. Carey and D. J. DeWitt and J. F. Naughton. The OO7 Benchmark. *SIGMOD Record*, 22(2):12–21, 1993.
- [8] D. J. DeWitt. The Wisconsin Benchmark: Past, Present, and Future. In *The Benchmark Handbook for Database and Transaction Systems*, editor J. Gray. Morgan Kaufmann, 2nd edition, 1993.
- [9] M. F. Fernandez, A. Morishima, and D. Suciu. Efficient Evaluation of XML Middle-ware Queries. In *SIGMOD*, 2001.
- [10] R. Goldman, J. McHugh, and J. Widom. From Semistructured Data to XML: Migrating to the Lore Data Model and Query Language. In *International Workshop on the Web and Databases*, 1999.
- [11] J. Gray. Introduction. In J. Gray, editor, *The Benchmark Handbook for Database and Transaction Systems*. Morgan Kaufmann, second edition, 1993.
- [12] H. V. Jagadish, S. Al-Khalifa, A. Chapman, L. V. S. Lakshmanan, A. Nierman, S. Pappas, J. M. Patel, D. Srivastava, Y. Wu, and C. Yu. TIMBER: A Native XML Database, 2003. To appear in *VLDB Journal*.
- [13] J. McHugh, S. Abiteboul, R. Goldman, D. Quass, and J. Widom. Lore: A Database Management System for Semistructured Data. *SIGMOD Record*, 26(3):54–66, September 1997.
- [14] A. Sahuguet, L. Dupont, and T. L. Nguyen. Querying XML in the New Millennium. <http://db.cis.upenn.edu/KWEELT/>.
- [15] A. Schmidt and F. Wass and M. Kersten and D. Florescu and M. J. Carey and I. Manolescu and R. Busse. Why And How To Benchmark XML Databases. *SIGMOD Record*, 30(3), September 2001.
- [16] A.R. Schmidt, F. Wass, M.L. Kersten, D. Florescu, I. Manolescu, M.J. Carey, and R. Busse. The XML Benchmark Project. Technical report, CWI, Amsterdam, The Netherlands, April 2001.
- [17] J. Shanmugasundaram, J. Keirnan, E. J. Shekita, C. Fan, and J. Funderburk. Querying XML Views of Relational Data. In *VLDB*, 2001.
- [18] J. Shanmugasundaram and E. J. Shekita and R. Barr and M.J. Carey and B.G. Lindsay and H. Pirahesh and B. Reinwald. Efficiently Publishing Relational Data as XML Documents. *The VLDB Journal*, 10(2-3):133–154, 2001.
- [19] J. Shanmugasundaram, K. Tufte, G. He, C. Zhang, D. DeWitt, and J. Naughton. Relational Databases for Querying XML Documents: Limitations and Opportunities. In *VLDB*, 1999.
- [20] The Michigan Benchmark Team. The Michigan Benchmark: Towards XML Query Performance Diagnostics, Feb 2003. <http://www.eecs.umich.edu/db/mbench>.
- [21] B. B. Yao and M. Tamer Özsu and J. Keenleyside. XBenCh – A Family of Benchmarks for XML DBMSs. In *VLDB EEXTT Workshop*, 2002.