



Centrum voor Wiskunde en Informatica

REPORT*RAPPORT*

INS

Information Systems



Information Systems

An empirical evaluation of XQuery processors

S. Manegold

REPORT INS-E0704 MARCH 2007

Centrum voor Wiskunde en Informatica (CWI) is the national research institute for Mathematics and Computer Science. It is sponsored by the Netherlands Organisation for Scientific Research (NWO). CWI is a founding member of ERCIM, the European Research Consortium for Informatics and Mathematics.

CWI's research has a theme-oriented structure and is grouped into four clusters. Listed below are the names of the clusters and in parentheses their acronyms.

Probability, Networks and Algorithms (PNA)

Software Engineering (SEN)

Modelling, Analysis and Simulation (MAS)

Information Systems (INS)

Copyright © 2007, Stichting Centrum voor Wiskunde en Informatica
P.O. Box 94079, 1090 GB Amsterdam (NL)
Kruislaan 413, 1098 SJ Amsterdam (NL)
Telephone +31 20 592 9333
Telefax +31 20 592 4199

ISSN 1386-3681

An empirical evaluation of XQuery processors

ABSTRACT

This paper presents an extensive and detailed experimental evaluation of XQuery processors. The study consists of running five publicly available XQuery benchmarks --- the Michigan benchmark (MBench), XBench, XMach-1, XMark and X007 --- on six XQuery processors, three stand-alone (file-based) XQuery processors (Galax, Qizx/Open, Saxon-B) and three XML/XQuery database systems (BerkeleyDB/XML, MonetDB/XQuery, X-Hive/DB). Next to assessing and comparing the functionality, performance and scalability for the various systems, the major focus of this work is to report in detail about the experiences made while performing such an exhaustive study, to discuss all the problems that we encountered and how we solved them, and hence to hopefully provide some guidelines (or even a recipe) for performing reproducible large-scale experimental research and system evaluation.

1998 ACM Computing Classification System: H.3.4: Performance evaluation; D.2.8: Performance measures; C.4: Measurement techniques

Keywords and Phrases: XML; XQuery; XQuery processors; Database Systems; Benchmarks; Performance Evaluation

An Empirical Evaluation of XQuery Processors

Stefan Manegold
CWI
Kruislaan 413
1098 SJ Amsterdam
The Netherlands
Stefan.Manegold@cwi.nl

ABSTRACT

This paper presents an extensive and detailed experimental evaluation of XQuery processors. The study consists of running five publicly available XQuery benchmarks — the Michigan benchmark (MBench), X Bench, XMach-1, XMark and X007 — on six XQuery processors, three stand-alone (file-based) XQuery processors (Galax, Qizx/Open, Saxon-B) and three XML/XQuery database systems (BerkeleyDB/XML, MonetDB/XQuery, X-Hive/DB). Next to assessing and comparing the functionality, performance and scalability for the various systems, the major focus of this work is to report in detail about the experiences made while performing such an exhaustive study, to discuss all the problems that we encountered and how we solved them, and hence to hopefully provide some guidelines (or even a recipe) for performing reproducible large-scale experimental research and system evaluation.

1. INTRODUCTION

Experimental evaluation and comparison of (new) techniques, algorithms and/or complete systems is a vital means to assess the practical impact and benefit of research results, especially in applied domains such as data management systems. While many publications present experimental results, the extent of the presentation — or even the experiments themselves — are often very limited due to space, time and/or other resource constraints. The major focus of most research publications is on the (so-called) *scientific contributions*.

In this study, we shift the focus. Performing the actual experimental evaluation becomes the primary subject. Presenting the experimental setup in detail, we do not hesitate to reveal all the “nasty details” and “minor problems” that give us headaches and cause sleep-less nights. Although choosing XML data management using XQuery as the sample scenario for this study, we believe that the proposed techniques can easily be adapted to other data management scenarios. The major contribution is hence a detailed cook-

book about how to conduct an experimental comparison and assessment of data management systems. However, though focusing on the actual experimentation techniques, we also present the “hard facts”: the detailed performance results we gathered.

The remainder of this paper is organized as follows. We start with presenting our experimentation environment in Section 2. Section 3 lists the benchmarks that we are using and discusses adaptations of the benchmarks, their queries and documents that are necessary to perform our experiments. In Section 4, we introduce the six XQuery systems we use, and explain in detail how we compile, install, configure and use them. Section 5 reveals how we run our experiments, validate results and timings, and collect the performance data that is presented in detail in Section 6. We draw final conclusions in Section 7.

2. THE SCENARIO

Given the popularity of XML in the data management world — both in research and in commercial environments — we chose XML data management using XQuery as the sample scenario for our experimental study. In particular, the idea is to run the five most popular publicly available XQuery benchmarks — the Michigan benchmark (MBench), X Bench, XMach-1, XMark and X007 — on the most popular publicly available XQuery processors. We picked three stand-alone (file-based) XQuery processors (Galax, Qizx/Open, Saxon-B) and three XML/XQuery database systems (BerkeleyDB/XML, MonetDB/XQuery, X-Hive/DB). For simplicity, we will use *engine* as a unified term to refer to both stand-alone (file-based) XQuery processors and XML/XQuery database systems in the remainder of this text.

Despite using XML/XQuery as sample scenario, none of the general approaches and techniques presented and used in this work are limited to XML and/or XQuery. With obvious and straightforward adaptations, they can likewise be applied to other data management scenarios.

The chosen setup is very complex. First of all, there are 5×6 combinations of benchmarks with engines (actually 8×6 , as X Bench comes in four flavors, TC/SD, DC/SD, TC/MD, DC/MD). Secondly, each benchmark consists of 8 up to 46 individual queries (in total 163 queries for all 5 benchmarks). Thirdly, each benchmark provides (at least) 3 to 6 different document sizes. Hence, we face two major challenges: (1) Running a huge set of experiments, and (2)

collecting, analyzing, and presenting a huge amount of experimental results. To tackle these challenges, we decided to use the *XCheck benchmark platform* [2].

2.1 Benchmark tool: XCheck

XCheck [2] provides a convenient integrated platform to run various benchmarks (or experiments in general), each consisting of several individual queries and varying input documents, using different engines. Per benchmark, all experiments are performed by a single invocation of XCheck.

For each benchmark, XCheck iterates over all engines, and for each engine over all document sizes, executing each query of the respective benchmark on the given document size using the given engine. Each individual experiment is repeated $n+1$ times. The execution times of the first run are neglected (“*warm-up*”). XCheck then calculates the average and standard deviation of the execution times for the remaining n runs¹. Next to the overall execution time, XCheck allows to collect the breakdown times for document processing, query translation, query execution, and result serialization — as far as provided by the various engines (see Section 4.3.3 for details). Additionally, XCheck collects the sizes of the produced results and (pre-defined) error messages.

2.2 Customizing XCheck

XCheck can be extended to use new benchmarks and/or new engines by providing the respective information, e.g., the benchmark queries and documents (or the respective generators), in XML configuration files.

While being very convenient and working reliably, XCheck (Version 0.1.3) comes with two limitations that are quite relevant for our experiments.

The first limitation is that XCheck can only handle a single document per query. XCheck replaces the URI in the `fn:doc()` call in the queries to use the requested document (size), and only one URI is replaced. Hence, multi-document experiments/benchmarks cannot easily be run with XCheck. To circumvent this limitation while retaining the multi-document characteristics, we gather the URIs of all documents of each multi-document benchmark (XMach-1, XBench TC/MD & DC/MD) in a single XML file

```
<uri_collection>
  <uri><!-- URI of document 1 --></uri>
  <!-- ... -->
  <uri><!-- URI of document n --></uri>
</uri_collection>
```

and obtain the sequence of documents on-the-fly via the following preamble to each query:

```
let $doc :=
  for $d in doc("uri_collection.xml")//uri/text()
  return doc($d)
```

¹In our experiments, we use the default: $n = 3$.

The second limitation is the way how XCheck calls the engines to run a query. This is done via a single command line call. While this is sufficient for stand-alone (file-based) XQuery processors, it does not allow for a “natural” usage of XML/XQuery database systems. Firstly, most (XML/XQuery) database systems have a client-server architecture, where the server is running in the background, and each query is executed by calling of the client program that then connects to the server. Secondly, XML/XQuery database systems allow that the files that contain the XML documents need to be read only once, loading the documents into the database. All queries then only need to access the documents as stored in the database, not requiring the parsing of the original document with each individual query.

While having access to the source code of XCheck (it is basically a collection of perl scripts), we decided not to change the code, but rather exploit (mis-use?) some features of the *adapters* that specify for each individual engine how XCheck should execute a query. Next to the actual call to execute the query, XCheck allows for both a pre-processing- and a post-processing-call. All three interfaces are basically simple command line calls of arbitrary executables, parametrized with (the name/location of) the query-file and the document-file. Hence, a simple straightforward solution would be to start the database server in the pre-call, run the query via the client, and finally stop the server with the post-call. However, we think it is not “natural” to start and stop the database server for each individual query execution — let alone the extra overhead/delay that extends the overall execution time of the whole benchmark considerably.

Alternatively, we could start *all* required engines and load *all* used documents prior to starting the XCheck runs. However, this would mean that we (1) require enough disk space to store all (possibly huge) documents of each benchmark in all databases concurrently, and (2) all database servers are running concurrently — though only one is active at a time, the “idle” ones consume vital memory, which might compromise the behavior of the active one.

Instead, we want to start each database server individually only once per benchmark, engine and document, and leave it running while all queries of the given benchmark are executed on the given document with the given engine. To implement this, we added two empty queries to each benchmark, containing only the comments “(:StartServer:)” and “(:StopServer:)”. Used as first and last query for each benchmark, these queries trigger the pre- and post-calls to start and stop the respective database server. For all other queries, the pre- and post-calls do nothing.

Likewise, we added a query Q0 to all benchmarks, that loads the respective document into the database using the respective database’s document loading functionality. Thus, XCheck automatically uses these queries to measure and collect the document loading times. All other benchmark queries then access only the pre-loaded document, just like a database scenario is supposed to work. To save disk space, our post-calls remove the document from the database once they see the respective “(:StopServer:)”-query.

X007 [8] (cfg)		XMark [15] (sf)		MBench [13] (sf)		XMach-1 [6] (#docs)		XBench [17]						
								(size)	TC/SD	TC/MD (#docs)		DC/SD	DC/MD (#docs)	
small3	2x 4.5 MB	0.001	110 KB			100	2.3 MB	small	11 MB	26	9.1 MB	11 MB	2597	9.9 MB
small6	2x 8.7 MB	0.01	1.1 MB			1000	18 MB	normal	104 MB	266	97 MB	104 MB	25925	100 MB
small9	2x 13 MB	0.1	11 MB	0.1	46 MB	10000	174 MB	large	1.1 GB	2666	1.1 GB	1.1 GB	259205(?)	(?)
med3	2x 44 MB	1.0	110 MB	1.0	496 MB			huge	11 GB	26666	16 GB	(?)	2592005(?)	(?)
med6	2x 86 MB	10.0	1.1 GB	10.0	4.8 GB							[document generator failed/crashed]		
med9	2x 129 MB	100.0	11 GB											
22		20		46		8		#queries	17	19		16	15	

Table 1: Benchmarks, their document sizes and number of queries

3. THE BENCHMARKS

In this work, we consider the five most popular publicly available XML/ XQuery related benchmarks: MBench [13], X007 [8], XBench [17], XMach-1 [6], XMark [15], as listed in Table 1. All benchmarks consist of a set of queries and provide document generators that allow to generate documents of various sizes.

3.1 Documents

Compiling, installing and running the document generators was no problem with most of the benchmarks. However, some generators required small fixes. For X007 and MBench, the generators required minor (quite obvious) changes to get the source code compiled with gcc/g++ 4.0.2 on Fedora Core 4 (details omitted here). The XBench document generator worked fine for the text-centric documents (TC/SD, TC/MD), but kept crashing with some “obscure” Java exception when trying to generate the document-centric documents (DC/SD, DC/MD) on our Fedora Core 4 (64-bit) systems using Java 1.5.0. Luckily, Loredana Afanasiev could provide us with the generated documents — at least the smaller sizes.

3.2 Queries

Except from XMark, not all queries of the benchmarks were publicly available in a form that complies with the latest XQuery syntax requirements. In fact, 62 out of the total 163 benchmark queries were not available as valid XQuery queries. Our thanks go to Loredana Afanasiev [4] for making these queries compliant with the latest XQuery requirements², so that they can be processed by most of the considered XQuery processors. Some queries still give syntactical or runtime errors with some of the engines, mainly due to limitations of the respective engines (see Section 6.1 for details).

4. THE SYSTEMS

To be included in our evaluation, XQuery processors need to fulfill the following criteria:

- free public availability, either in open source, or at least as a binary (evaluation) version;
- running under Linux on an x86_64 (AMD_64) or x86 (Intel i686) platform;
- supporting (a reasonable subset of) XQuery.

²cf., <http://staff.science.uva.nl/~lafanasi/xcheck/queries.html>

While there might be more systems fulfilling these criteria, we limit our evaluation to the following ones:

Engine	Version	Code	Bits
M MonetDB/XQuery [7]	0.14.0	C	64
X X-Hive/DB [16]	7.2.2	Java	64
B BerkeleyDB/XML [5]	2.2.13	C/C++	64
S Saxon-B [14]	8.7.1	Java	64
G Galax [9]	0.6.10	OCaml	64
Q Qizx/Open [12]	1.0	Java	64

The first three are XML/XQuery database systems, the last three are stand-alone (file-based) XQuery processors.

4.1 Compilation

While all chosen XQuery engines are publicly available, there is no unified way to install all of them. Neither are all engines available as ready-to-run packages suitable for our experimentation platform (cf., Section 6.2), nor are all of them available in open source. In the following, we briefly describe how we installed each engine.

X-Hive/DB, Saxon-B, Qizx/Open: For the 3 Java-based engines, we use the pre-compiled .jar packages, even if the source code is available as well.

The remaining three engines — MonetDB/XQuery, BerkeleyDB/XML, Galax — are all available in open source. Hence, we compiled them optimized for our experimentation platform.

Galax: As recommended on the Galax web-site, we used the GODI installation and configuration tool to automatically compile and install a 64-bit version of Galax from the sources. We used the default optimization switches for the OCaml compiler, as GODI does not allow to change them (easily).

MonetDB/XQuery: We use the 64-bit, 32-bit OIDs binary RPMs as available from MonetDB’s SourceForge site. These packages have been compiled with full optimization (‘configure --enable-optimize’³).

BerkeleyDB/XML: We compiled a 64-bit version of the sources using gcc/g++ 4.0.2 and the same optimization switches³ as with MonetDB/XQuery.

³gcc -O6 -fexpensive-optimizations -falign-loops=4 -frerun-loop-opt -funroll-loops -falign-jumps=4 -finline-functions -frerun-cse-after-loop -falign-functions=4 -fomit-frame-pointer

(In fact, we planned to analyze the impact of compiler optimization switches — mainly the difference between the default `-g -O2` and the excessive list used here — as well as the impact of using different compilers, e.g., the Intel compiler, on the various open source systems. However, time and resource limitations forced us to postpone this analysis. We plan to perform it and report the results as soon as time and resources allow it.)

4.2 Configuration

The basic idea is to use the default “*out-of-the-box*” configuration of the systems. We only applied some minor configurations related to memory. Given that we are using a 64-bit machine with 8 GB of main memory (see Section 6.2 for details), we allow the Java virtual machine for all Java-based systems to allocate up to 6 GB of main memory (`'java -mx6144m'`). Using the complete 8 GB is not possible, as some system processes use some memory (probably less than 2 GB, but we did not spend time on finding the maximum we could safely assign to Java; 6 GB seemed to work and be sufficient for moderate document sizes), and our machine is configured to not allow memory allocations that exceed the physical available memory. The latter is an attempt to avoid instable behavior of our machine under excessive virtual memory allocations.

In particular, we did *not* explicitly create any indices with the database systems, though especially X-Hive and BerkeleyDB/XML might benefit significantly from creating the right indices for the right queries (cf., [7]). The main reason for this decision was that we did not have enough in-depth knowledge of all systems to tune all of them equally well, and thus ensure a *fair* comparison. Hence, treating all systems equally with “no tuning at all” seems the fairest approach for now. In fact, investigating the impact/benefit of using indices and further tuning would provide enough material for a separate study of this kind.

4.3 Adapters

XCheck uses simple XML configuration files (“*adapters*”) to specify the details how to call each individual engine and how to collect the detailed timing information (if available) from their output.

Saxon-B, Galax, Qizx/Open: For the stand-alone file-based processors, we use the default adapters that come with XCheck 0.1.3. In the following, we describe our new adapters for the three database engines used here.

4.3.1 Starting Servers and Loading Documents

To model a “realistic” database usage scenario, we decided to start the database server and pre-load the XML documents from the benchmarks’ document XML files into the database only once per benchmark and document, not repetitively for each query. To implement this, we exploit the concept of pre- and post-calls in the XCheck adapters. The principle mechanism is sketched in Section 2.2. We now discuss the details for each database engine.

MonetDB/XQuery: Our pre-call script starts the server in daemon mode via

```
'Mserver --set monet_daemon=yes
--dbinit='module(pathfinder)';'
```

Loading the document simply requires execution of query Q0: `count(doc("<doc.uri>")[*])`, exploiting the document caching feature of MonetDB/XQuery: `fn:doc()` reads the document from the given URI and stores it in the database.⁴ Subsequent `fn:doc()` calls with the same URI avoid reloading the document, unless its timestamp has changed. The respective post-call script stops the server by killing it.

X-Hive/DB: We start and stop the server via

```
'XHStartServer' and 'XHStopServer',
```

respectively. To load the documents, the execution script extracts the document URI from Q0 and calls a small Java program (an adaptation of the sample `StoreDocuments.java` that comes with X-Hive/DB) that stores the document in the database and assigns it the filename from the URI for later reference.

BerkeleyDB/XML: Though also a database system, BerkeleyDB/XML does not have a client/server architecture, but simply works as a stand-alone program that uses a persistent database storage. Hence, our pre-call script only needs to create a database, using the `createContainer` command of the `dbxml` console application. To load documents, we use the `putDocument` command of the `dbxml` console application. Like with X-Hive/DB, we extract the URI from Q0 and assign the filename from the URI for later reference.

4.3.2 Running Queries

We use the following commands and respective commandline options to execute XQuery queries (given in a file) with the various engines.

Saxon-B:

```
'java -cp saxon8.jar net.sf.saxon.Query'
```

Galax:

```
'galax-run -output-xml'
```

Qizx/Open:

```
'qizxopen.batch -serial'
```

MonetDB/XQuery:

```
'MapiClient -lxquery -oxml'
```

X-Hive/DB:

We adapted the sample `XQuery.java` that comes with the distribution according to our needs.

BerkeleyDB/XML:

We use the `query` and `print` commands of the `dbxml` console application to execute queries.

⁴By default, when loaded implicitly via `doc("<doc.uri>")`, MonetDB/XQuery keeps only documents up to 100 MB persistent in the database; for our experiments, we raise the limit to 20 GB (`--set xquery_cacheMB=20000`).

	M	X	B	S	G	Q
query translation	+	*	+	+	+	
query execution	+	*	+	+	+	+
result serialization	+	*	+		+	+
document processing	+	*	+	+	+	
communication	$T - (tran + exec + seri + docu)$					
Total time	+	*	XCheck			

+: default

*: self-made

Table 2: Which systems provides which timing

4.3.3 Measuring Time

While XCheck measures the overall evaluation time for each query execution itself, it can also collect more detailed timing information from the engines. The various engines provide different means to get detailed timing information.

Saxon-B: The default adapter calls `net.sf.saxon.Query` from `saxon8.jar` with the `-t` switch to get information about tree build time (“Tree built in”), “Compilation time”, and “Execution time”. The former two are taken as document processing time and query translation time, respectively. Query execution time is calculated as difference of Execution time and Compilation time. Result serialization time is not available for Saxon-B.

Galax: The default adapter calls `galax-run` with the `-monitor-time` option to get the detailed timing for document processing, query translation, query execution and result serialization. (See also Section 5.2.)

Qizx/Open: The default adapter calls `qizxopen_batch` with the `-tex` switch to get information about “evaluation time” and “display time” which are used as query execution time and result serialization time, respectively. Document processing time and query translation time are not available for Qizx/Open.

MonetDB/XQuery: We call `MapiClient` with `-T` to get the detailed timings for document processing (“Shred”), query translation (“Trans”), query execution (“Query”), and results serialization (“Print”). We also get the total execution time (“Timer”), overruling XCheck’s measurements.

X-Hive/DB: We add timing statements to the modified sample applications (`StoreDocuments.java` & `XQuery.java`) that measure the times taken by document loading, query translation (`rootLibrary.executeXQuery(theQuery);`), query execution (`result.next();`) and result serialization (`System.out.println(value.toString());`). Like with MonetDB/XQuery, we also measure the overall evaluation time, overruling the measurement done by XCheck.

BerkeleyDB/XML: We call `dbxml` with `-vv` to get the query translation time (“Optimizer - Finished parse, time taken”) and query execution time (“Query - Finished query execution, time taken”). Additionally, we use the `time` prefix with the `putDocument` and `print` commands to get the document processing time and result serialization time, respectively. Alternatively, we also could use the `time` prefix with the `query` command to get the sum of query

translation time and query execution time. However, in many cases, it turned out that none of these timings can be trusted (see Section 5.2).

Table 2 summarizes, which engine provides which timing information. For the database engines (MonetDB/XQuery, X-Hive/DB, BerkeleyDB/XML), the document processing time represents only the time required to load, parse and process the original XML document in order to store it in the database. Any costs for accessing the document once it is stored in the database are included in the query processing times and/or communication times (cf., Section 5.2). Please note that the classification of detailed timings given here is based on the semantics we could easily derive from the respective systems’ documentation (as far as available). There is no guarantee that all systems use the same semantics and/or definition for these detailed timings.

5. THE EXPERIMENTS

Finally, we are ready to run our experiments. To ensure that the measured performance results do indeed make sense, we need to make sure that all engines work properly, with respect to both the actual XQuery processing and the collection of detailed timing information.

5.1 Checking/Validating Results

Given our extended (purely practical) experience in software testing and validation⁵, our first concerns are not the actual performance results, but rather the question whether the various engines indeed produce the correct results. Unfortunately, the benchmarks do not come with correct results. In fact, this is hardly feasible, given the randomness built into most document generators for good reasons. Moreover, even with correct results provided, validating the actual results requires more than a simple `diff`, as both XQuery semantics and XML specifications allow for some variation that cannot easily be recognized as equivalent. For the smallest document size of each benchmark, we did “by hand” verify “consistency” among the engines, i.e., all engines yield the same result for each query. For the remaining document sizes, we are left with what XCheck offers us. First, XCheck detects engine-specific error messages, using regular expressions given in the adapters. Second, XCheck collects the sizes of the produced query results. Though we do not have the resources to analyze this in detail, a quick comparison reveals that the all engines produce results “of similar size” for all document sizes and each query processed without errors.

5.2 Checking/Validating Timings

Already shortly after starting the experiments, we noticed that there were some inconsistencies with the detailed timings of some of the systems. Basically, the breakdown timings (document processing, query translation, query execution, result serialization) did not always add up to the total evaluation time. In most cases, the sum was less than the total; we assume that this is due to start-up and communication costs that are not included in the detailed timings. Once the execution time (per query) exceeds one minute (60 seconds), the detailed timings reported by Galax are so small, that their sum is only a minor fraction of the total

⁵cf., <http://monetdb.cwi.nl/TestWeb/>

execution time. We guess that the respective code is not correct, and hence, consider these timings as unreliable. In our graph in the next section, we hence depict these “missing times” as **communication** times.

More severe are the cases where the sum of the detailed timings exceed the total times (often significantly). This is mainly the case with BerkeleyDB/XML, regardless which of the alternatives we use to get the query execution times (cf., Section 4.3.3). Apparently, there is some bug in the code that measures the respective times in BerkeleyDB/XML. In the breakdown graphs in the next section, we mark these “excess times” as **(void)**. In the scalability graphs, we use the total times as measured by XCheck.

6. EXPERIMENTAL RESULTS

Figures 1 through 18 present a subset of the performance results that we collected from our exhaustive experiments, running all benchmarks with various document sizes, as depicted in Table 1. To avoid “endless” runs, we limit the execution time for each single query to at most *one hour* (3600 seconds). Only for loading documents (query Q0), we allow up to *one day* (24 hours).

The goal of this work is not to crown a single best system nor to advise users which system to use for their purpose. Rather, we want to provide detailed information and insights such that in particular developers can draw their own conclusions as to whether, where and how to improve their systems. Of course, also users are welcome to draw their own conclusions from the detailed results that we provide.

On the y-axis of all graphs, we list all queries of the respective benchmark, identified by their number on the right-hand side of each graph. “load doc.” identifies the document loading query Q0 as introduced in Sections 2.2 & 4.3. For each query, we list all six engines, identified by their first letter on the left-hand side of each graph. For each benchmark, we show two plots.

Execution time breakdown. For one sample document size per benchmark, the graphs on the left-hand side (“odd” Figures 1, 3, ..., 17⁶) depict the relative contribution of the detailed timings to the total evaluation time per engine and query. The different sections of the horizontal bars represent the various detailed timings, provided they are reported by the respective engine (cf., Section 4.3.3):

tran: query translation
 exec: query execution
 seri: result serialization
 comm: communication (cf., Section 5.2)
 docu: document processing
 (void): wrongly reported times that exceed the actual total times (cf., Section 5.2).

As mentioned in Section 4.3.3, the detailed definition and semantics of these breakdown times are not standardized and can vary between the systems.

⁶A full-color version of this paper is available on-line at <http://www.cwi.nl/htbin/ins1/publications?request=abstract&key=Ma:TR-CWI:07>

E01: preceding axis not supported	(see Section 6.1)
E02: parsing / (static) typing	
- <i>Qizx/Open</i> : MBench QA2; X007 Q23; XBench-DC/MD Q4,17; XBench-DC/SD Q17,20; XMach-1 Q3; XMark Q3,11,12,18.	
E03: invalid variable reference	(see Section 6.1)
E04: fatal XQuery compiler error	(see Section 6.1)
E05: materialization out of bounds	(see Section 6.1)
E06: out of memory	(see Section 6.1)
E07: out of Java heap space	
- <i>Saxon-B</i> : MBench QS6,J1-4 (496 MB), Q0-A6 (4.8 GB); XBench-TC/MD Q0-19 (≥ 2666 docs / 1.1 GB); XBench-TC/SD Q0-19 (≥ 1.1 GB); XMark Q1,4,6,7,18,19 (1.1 GB), Q0-20 (11 GB).	
- <i>Qizx/Open</i> : MBench QA4 (496 MB), Q0-J2 (4.8 GB); XBench-TC/MD Q0-19 (26666 docs / 16 GB); XBench-TC/SD Q0-19 (11 GB); XMark Q0-20 (11 GB).	
E08: segmentation fault	
- <i>BerkeleyDB/XML</i> : MBench QS12 (≥ 496 MB).	
- <i>Galax</i> : X007 Q19 ($\geq 2 \times 44$ MB); MBench QS35 (≥ 496 MB); XBench-DC/SD Q0-20 (1.1 GB); XMark Q0-20 (≥ 1.1 GB).	
E09: abort	
- <i>BerkeleyDB/XML</i> : XBench-TC/SD Q2,17,18 (≥ 1.1 GB); XMark Q7 (≥ 1.1 GB), Q20 (11 GB).	
E10: unknown error/crash (cf., Sec. 6.1)	
- <i>X-Hive/DB</i> : MBench QS12 (≥ 496 MB); X007 Q19 (2x 129 MB).	
- <i>BerkeleyDB/XML</i> : MBench QS1,2,A2,6,QR1-A6 (4.8 GB); X007 Q14,15, Q19 (2x 129 MB); XBench-TC/MD Q19; XBench-TC/SD Q3,6,9,10 (≥ 1.1 GB); XMach-1 Q3,7; XMark Q6,14 (≥ 1.1 GB), Q2-4,13,15-19 (11 GB).	
E11: cast not applied to a single atomic value	
- <i>Galax</i> : XMark Q3 (≥ 1.1 MB). (new since [11])	
doc: document loading failed	
- <i>MonetDB/XQuery</i> : XBench-TC/MD Q1-19 (26666 docs / 16 GB); XBench-TC/SD Q1-19 (11 GB).	
DNF: timeout (>1 h)	
- <i>MonetDB/XQuery</i> : MBench QS12,A4,J1-4 (≥ 496 MB), QS33 (4.8 GB); XBench-DC/MD Q19 (25925 docs / 100 MB); XBench-TC/MD Q0 (26666 docs / 16 GB; >24 h); XBench-TC/SD Q0 (11 GB; >24 h); XMark Q10,11,12,19,20 (11 GB).	
- <i>X-Hive/DB</i> : MBench QA4 (≥ 496 MB); XBench-TC/SD Q3 (≥ 1.1 GB); XMach-1 Q7 (10000 docs / 174 MB); XMark Q8,9,11,12 (≥ 110 MB), Q10,19 (11 GB).	
- <i>BerkeleyDB/XML</i> : MBench QJ2, QS35,A4,J1,3,4 (≥ 496 MB); XBench-TC/MD Q18 (26666 docs / 16 GB); XBench-TC/SD Q2-19 (11 GB); XMark Q8,9,11,12 (≥ 110 MB), Q10 (≥ 1.1 GB).	
- <i>Saxon-B</i> : MBench QA2,4 (≥ 496 MB); XMark Q8-12 (1.1 GB).	
- <i>Galax</i> : MBench QA4,J1-4, QS12,25,33 (≥ 496 MB), Q0-A6 (4.8 GB); X007 Q0,6 (2x 129 MB); XBench-TC/MD Q0,2-19 (≥ 2666 docs / 1.1 GB), Q1 (26666 docs / 16 GB); XBench-TC/SD Q0-19 (≥ 1.1 GB); XMach-1 Q2,7 (10000 docs / 174 MB).	
- <i>Qizx/Open</i> : MBench QJ3,4 (≥ 496 MB); XBench-TC/SD Q3 (≥ 1.1 GB); XMark Q10 (1.1 GB).	
lic: 30-days evaluation license expired	
- <i>X-Hive/DB</i> : MBench Q0-A6 (11 GB), QJ1-4; XBench-TC/MD Q0-19 (26666 docs / 16 GB); XBench-TC/SD Q0-19 (11 GB).	

Table 3: Errors and Error Codes used in Fig. 1–18

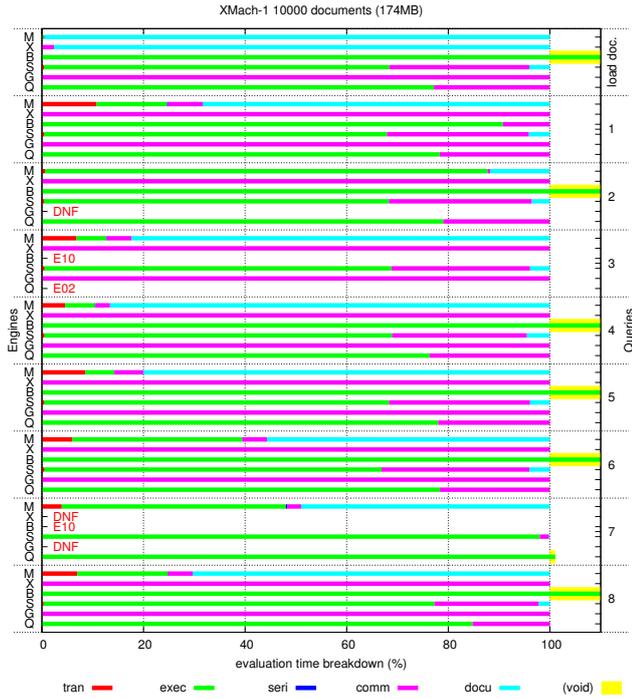


Figure 1: XMach-1: Execution time breakdown

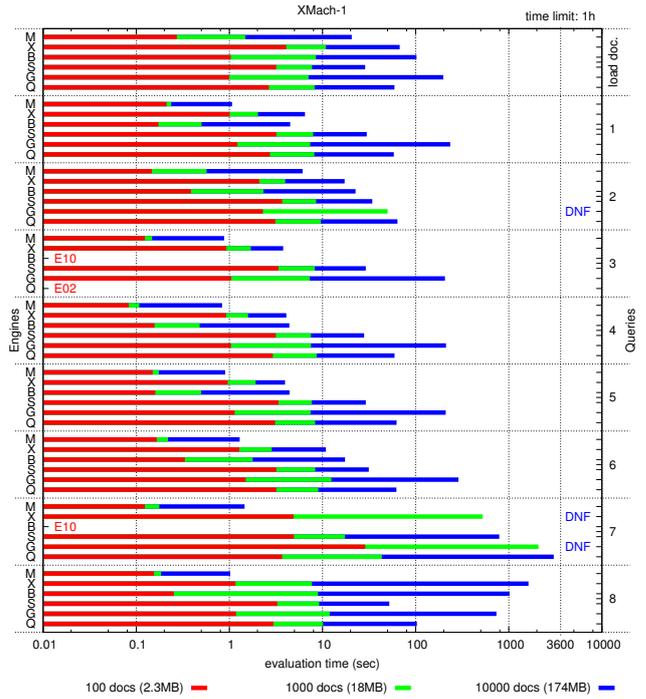


Figure 2: XMach-1: Scalability

Scalability. The graphs on the right-hand side (“even” Figures 2, 4, ..., 18⁶) depict the total times (wall-clock) for all document sizes of each benchmark. The execution time per engine, query and document is the length of the respective horizontal bar taken from the left margin of the graph, i.e., bars for smaller documents cover the left part of the bars of larger documents. To accommodate the results for all document sizes of one benchmark in one graph, we use a logarithmic scale (decimal base) for the x-axis of the scalability graphs.

All reported results are collected using XCheck as described in the previous sections. In particular, the measured times represent the average of the last three of four consecutive runs, i.e., “hot” results, neglecting the first “warm-up” run (cf., Section 2.1).

6.1 Errors

Some queries fail to execute successfully. Table 3 lists all errors that occur with all our experiments. We use the error codes (E01–E11, doc, DNF) from Table 3 to indicate the errors in Figures 1–18. In the “Scalability” figures, the error codes are depicted in the color of the smallest document size that the respective error occurs with. In various cases, the “unknown error/crash” (E10) could actually be caused by the fact that we kill the respective engine (or client) due to a timeout, in which case they should rather read DNF. However, we did not check this by hand in all cases.

For ease of comparison, we use the same error codes as in [11], although errors E01, E03–E06 do not occur any more. The respective problems with MonetDB/XQuery 0.10.2 and Galax 0.5.0 (used in [11]) have been fixed in the newer version used here (MonetDB/XQuery 0.14.0 and Galax 0.6.10).

6.2 Hardware & Operating System

Our experimentation platform is a dual 1.6 GHz AMD Opteron 242 (1 MB L2 cache) processor with 8 GB RAM and a RAID-5 disk subsystem (3ware 7810, configured with eight 250 GB IDE disks of 7200 RPM). The operating system is Fedora Core 4 (Linux 2.6.14 kernel), using a 64-bit address space. We use gcc/g++ 4.0.2 and Java 1.5.0 (64-bit).

7. CONCLUSIONS

First of all, we can conclude that our exercise demonstrates the feasibility of such an extensive and detailed experiment — though it requires quite some work, time, and resources. We describe our experimental setup in detail and explain, how we tackle various problems to reach our ambitious goal. We hope, we provide all information that is necessary and sufficient to reproduce our results.

As expected, the actual performance results do not crown a single winner. However, some general trends can be observed. In a realistic database scenario, i.e., with the documents pre-loaded in the database, the database engines perform considerably better (up to two orders of magnitude) than the file-based stand-alone systems. Even if we add the initial document loading times, they are often still faster, but hardly ever slower. In pure document loading performance (query Q0), MonetDB/XQuery, Saxon-B and Qizx/Open lead the race neck-to-neck; X-Hive/DB and BerkeleyDB/XML follow within a factor 2–3; Galax runs about factor 10 behind the leaders. While all systems have their strengths and weaknesses, MonetDB/XQuery seems to be ahead of the pack in most cases, usually followed (in that order) by BerkeleyDB/XML, X-Hive/DB, Qizx/Open, Saxon-B, and finally Galax. Join recognition and processing (still) seems to be the biggest challenge to be solved.

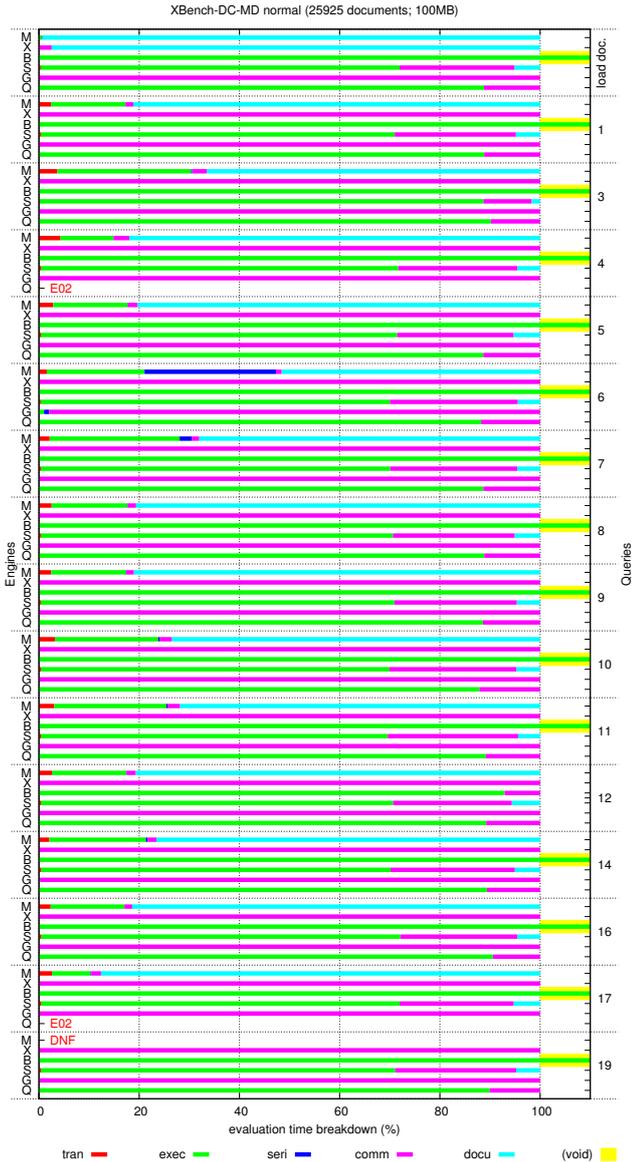


Figure 3: XBench-DC/MD: Execution time breakdown

Compared to [11], we upgraded two systems to newer versions. For MonetDB/XQuery, version 0.10.2 has been replaced by version 0.14.0, and for Galax, version 0.5.0 (pre-compiled 32-bit binary) has been replaced by version 0.6.10, which we compiled from the source into a 64-bit binary (cf., Section 4.1). For single-document benchmarks, the performance of MonetDB/XQuery 0.14.0 is within 20% of the performance of MonetDB/XQuery 0.10.2. With multi-document benchmarks, however, MonetDB/XQuery 0.14.0 shows an improvement of up to three orders of magnitude. The new (64-bit) version of Galax runs much more stable and with much less memory problems than the old (32-bit) one. In particular, it is now able to successfully run the XMark join queries (Q8,9,11,12) on documents of up to 110 MB ($sf=1$). However, overall scalability has hardly improved, and overall performance has dropped by a factor 2–5.

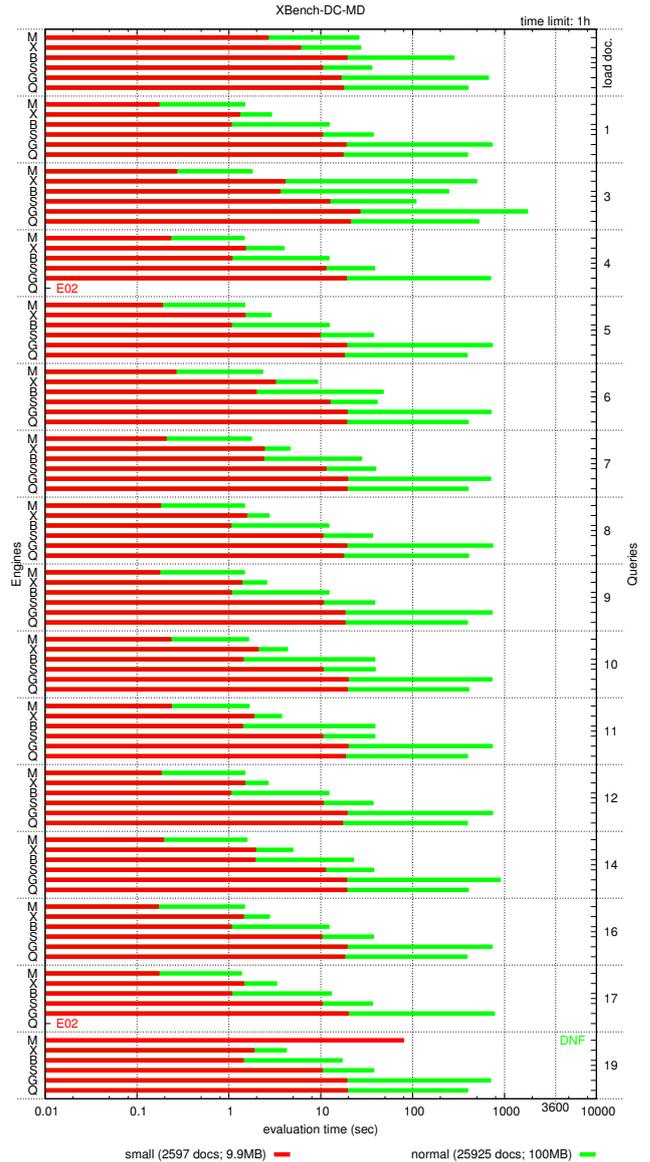


Figure 4: XBench-DC/MD: Scalability

For the future, we plan to extend the scenario in various dimensions, e.g., including more systems (e.g., eXist [1]) and benchmarks (e.g., XPathMark [10], MemBeR [3]), considering other compilers and optimization flags, using different hardware and operating systems, etc.. The goal is not to find the single best setup, but rather to show that all these often neglected factors can influence experimental results considerably. Hence, all reports of experimental results should reveal all these information explicitly in order to (1) provide all information to make the results reproducible, and (2) to put them in the right perspective. Finally, we hope that our experiences with XCheck will help to improve and extend this convenient experimentation tool.

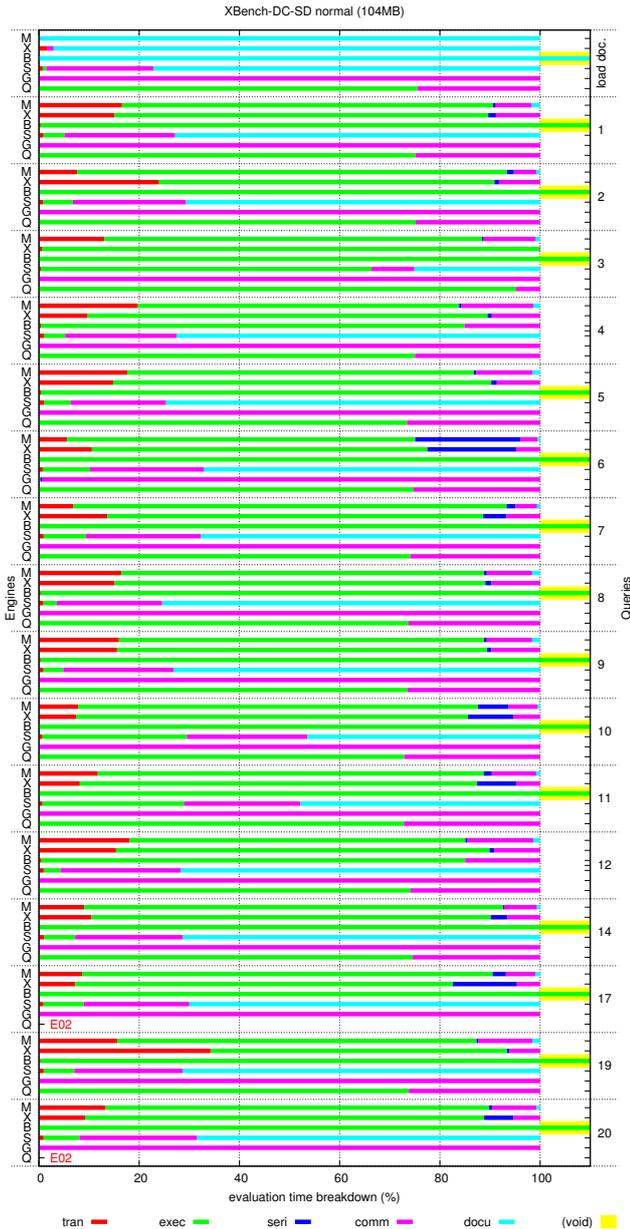


Figure 5: Xbench-DC/SD: Execution time breakdown

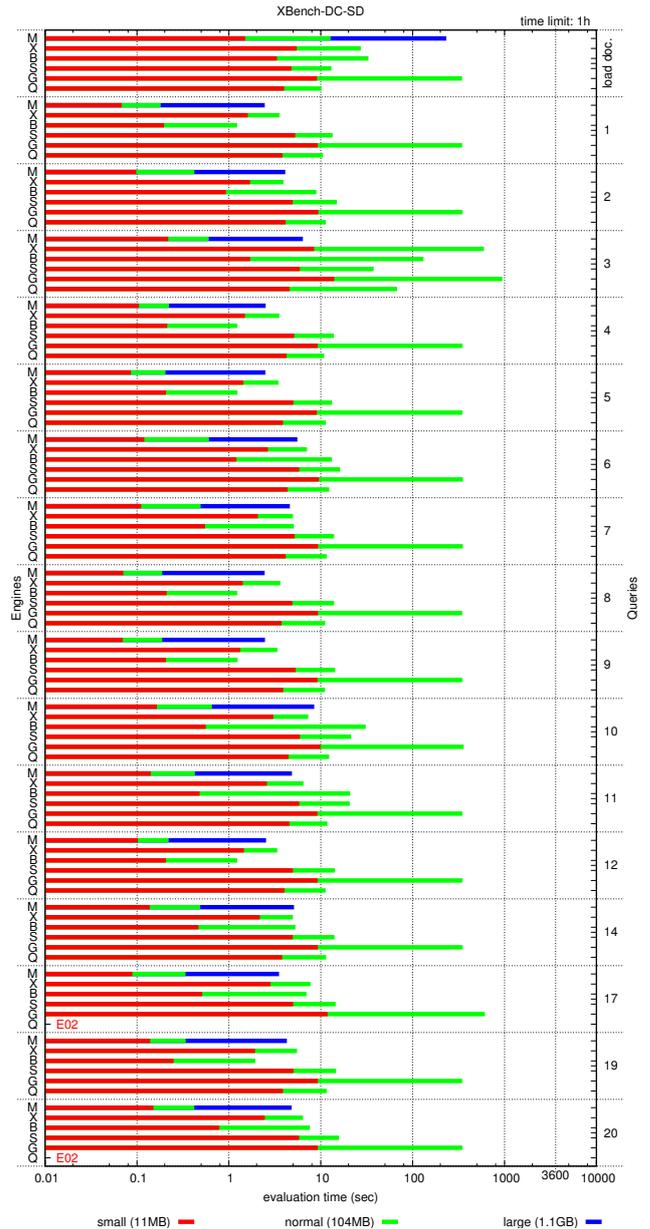


Figure 6: Xbench-DC/SD: Scalability

8. REFERENCES

- [1] eXist. <http://exist.sourceforge.net/>.
- [2] L. Afanasiev, M. Franceschet, M. Marx, and E. Zimuel. XCheck: a Platform for Benchmarking XQuery Engines. In *VLDB*, 2006. Demo. <http://ilps.science.uva.nl/Resources/XCheck/>.
- [3] L. Afanasiev, I. Manolescu, and P. Michiels. MemBeR: A Micro-benchmark Repository for XQuery. In *XSym*, 2005.
- [4] L. Afanasiev and M. Marx. An Analysis of the Current XQuery Benchmarks. In *ExpDB*, 2006.
- [5] Berkeley DB XML. <http://www.sleepycat.com/products/bdbxml.html>.
- [6] T. Böhme and E. Rahm. XMach-1: A benchmark for XML data management. In *BTW*, 2001. <http://dbs.uni-leipzig.de/de/projekte/XML/XmlBenchmarking.html>.
- [7] P. A. Boncz, T. Grust, M. van Keulen, S. Manegold, J. Rittinger, and J. Teubner. MonetDB/XQuery: A Fast XQuery Processor Powered by a Relational Engine. In *SIGMOD*, 2006. <http://monetdb-xquery.org/>.
- [8] S. Bressan, G. Dobbie, Z. Lacroix, M. Lee, Y. Li, U. Nambiar, and B. Wadhwa. X007: Applying 007 benchmark to XML query processing tool. In *CIKM*, 2001. <http://www.comp.nus.edu.sg/~ebh/X007.html>.

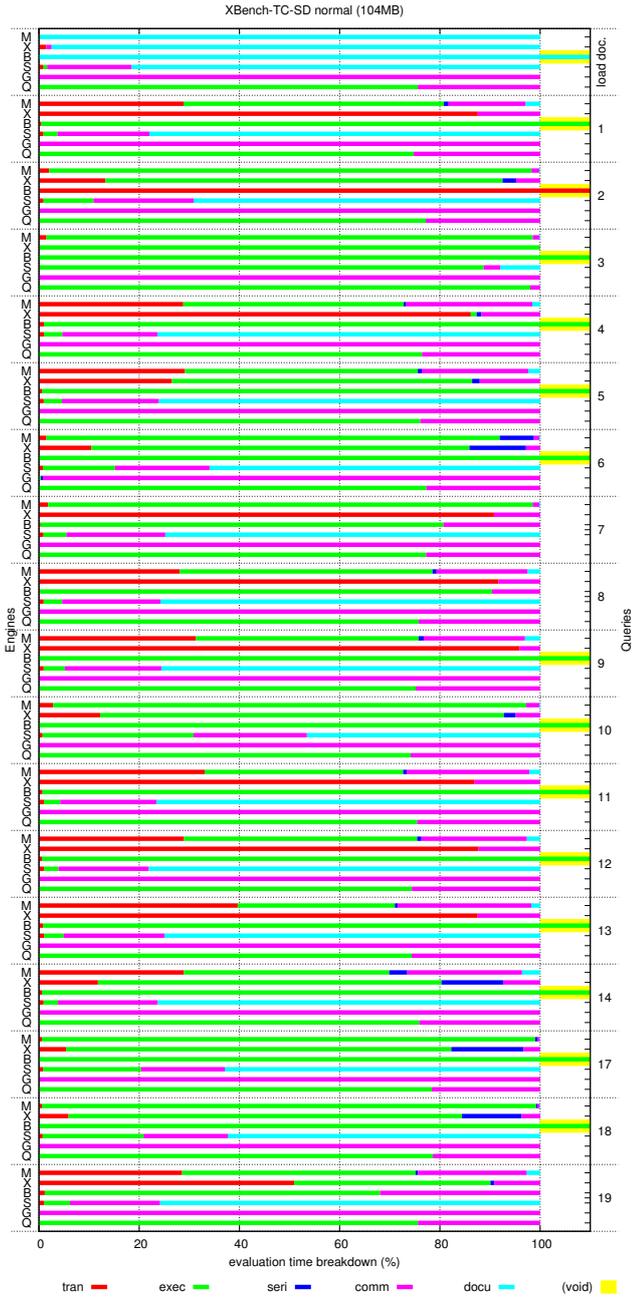


Figure 7: Xbench-TC/SD: Execution time breakdown

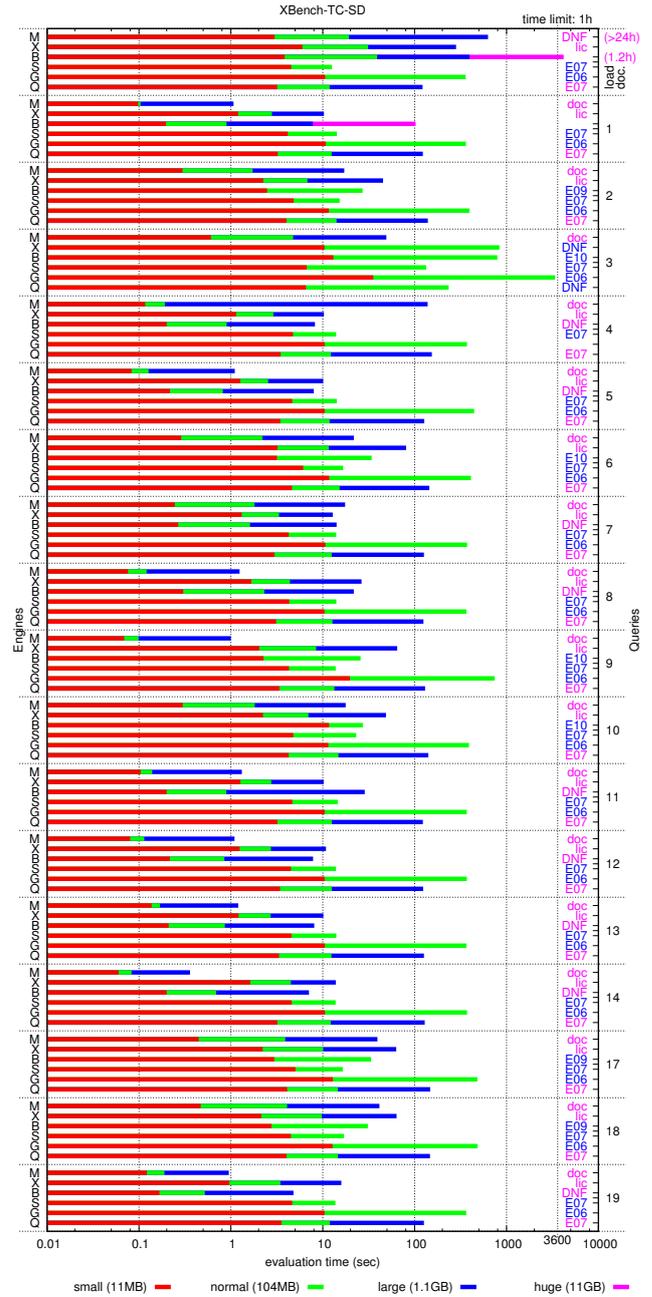


Figure 8: Xbench-TC/SD: Scalability

- [9] M.F. Fernández, J. Siméon, B. Choi, A. Marian, and G. Sur. Implementing XQuery 1.0: The Galax Experience. In *VLDB*, 2003. <http://www.galaxquery.org/>.
- [10] M. Franceschet. XPathMark: An XPath Benchmark for the XMark Generated Data. In *XSym*, 2005.
- [11] S. Manegold. An Empirical Evaluation of XQuery Processors. In *ExpDB*, 2006. <http://www.cwi.nl/htbin/ins1/publications?request=abstract&key=Ma:EXPDB:06>.
- [12] Qizx/Open. <http://www.axyana.com/qizxopen/>.
- [13] K. Runapongsa, J. Patel, H. Jagadish, Y. Chen, and S. Al-Khalifa. The Michigan Benchmark: A Microbenchmark for XML Query Processing Systems. In *EEXTT*, 2002. <http://www.eecs.umich.edu/db/mbench/>.
- [14] Saxon-B. <http://saxon.sourceforge.net/>.
- [15] A. Schmidt, F. Waas, M.L. Kersten, M.J. Carey, I. Manolescu, and R. Busse. XMark: A Benchmark for XML Data Management. In *VLDB*, 2002. <http://xml-benchmark.org/>.
- [16] X-Hive/DB. <http://www.x-hive.com/products/db/>.

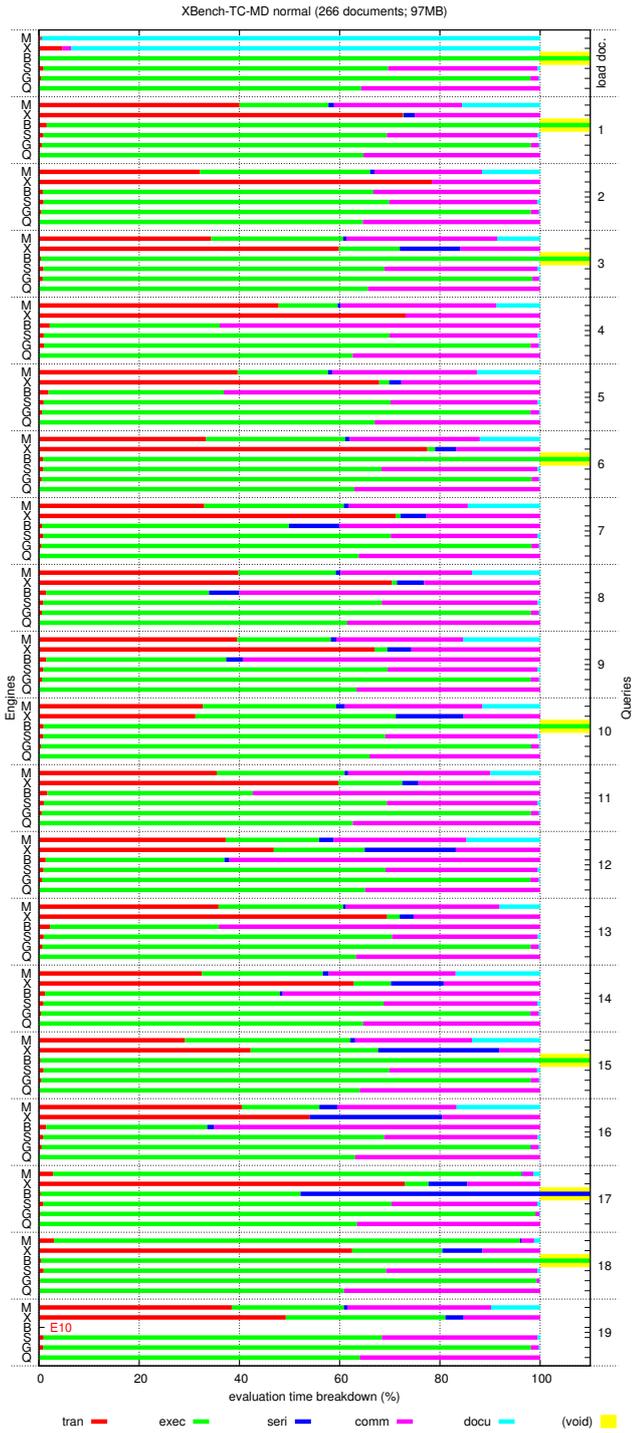


Figure 9: XBench-TC/MD: Execution time breakdown

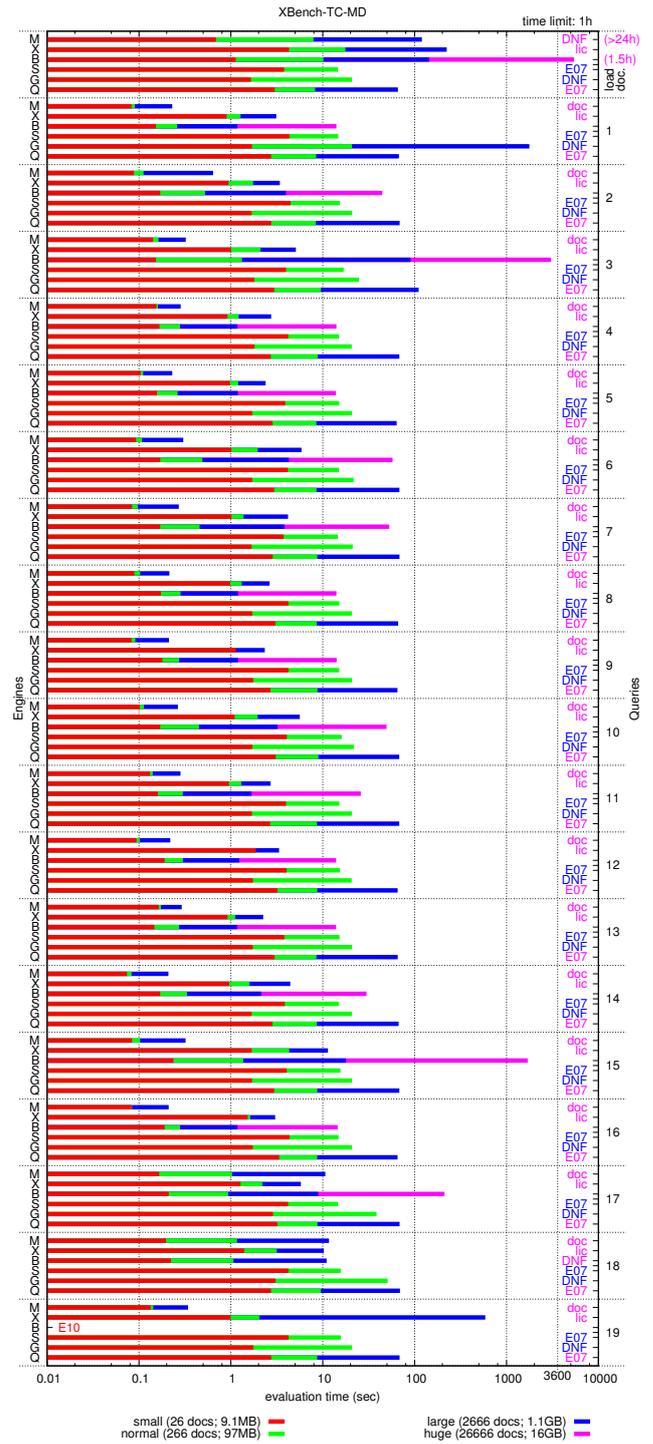


Figure 10: XBench-TC/MD: Scalability

- [17] B. Yao, T. Özsu, and N. Khandelwal. Xbench benchmark and performance testing of XML DBMSs. In *ICDE*, 2004.
<http://se.uwaterloo.ca/~ddbms/projects/xbench/>.

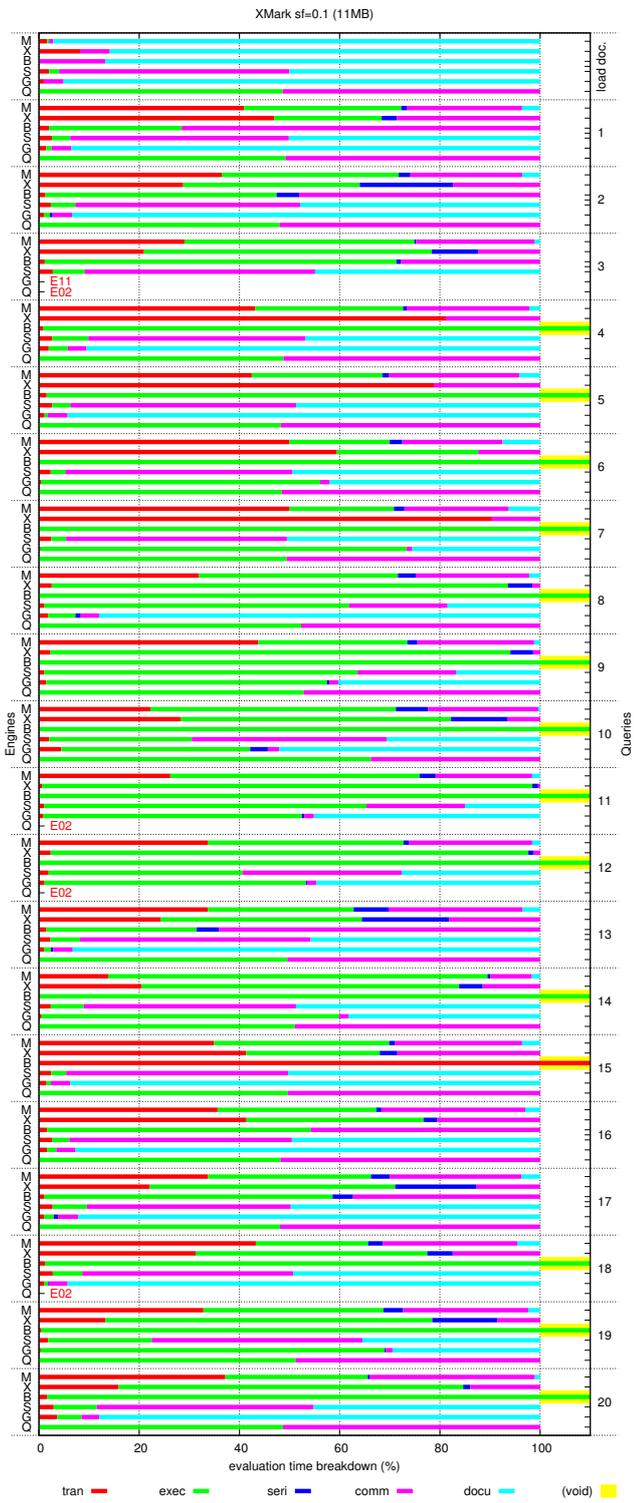


Figure 11: XMark: Execution time breakdown

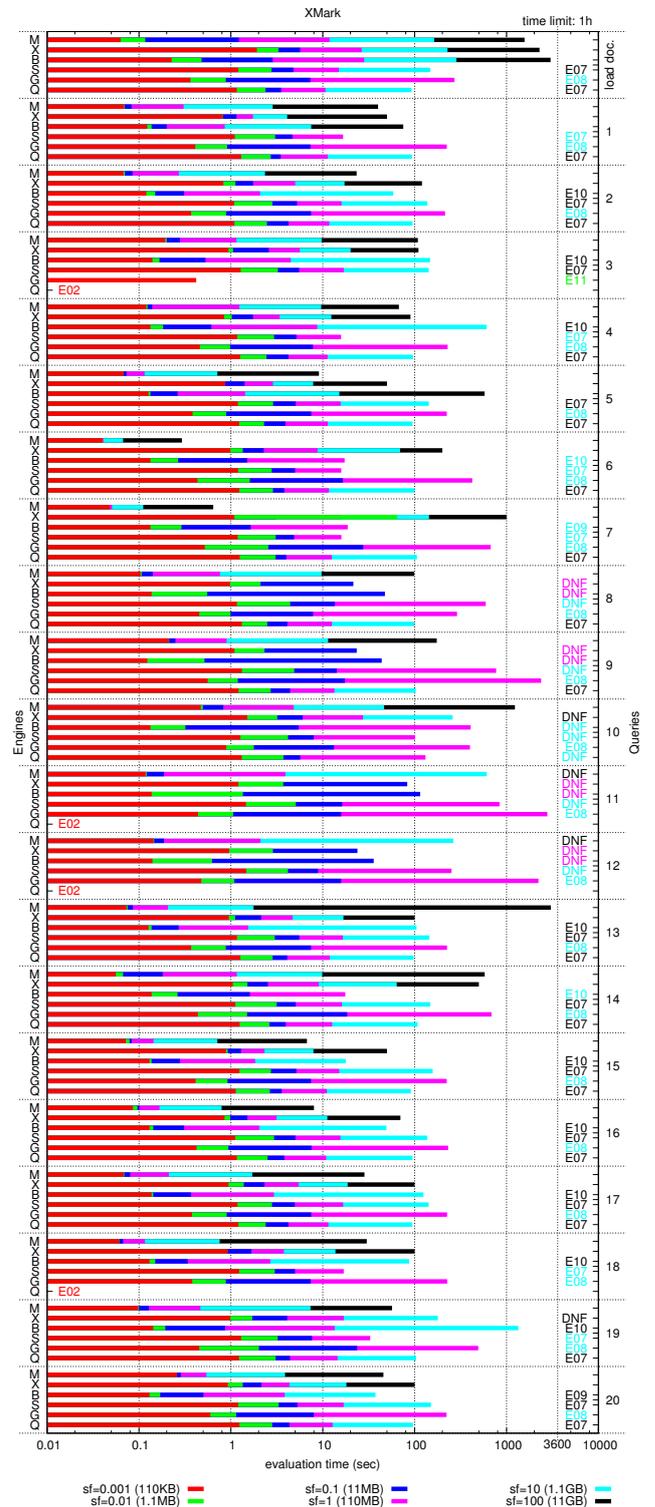


Figure 12: XMark: Scalability

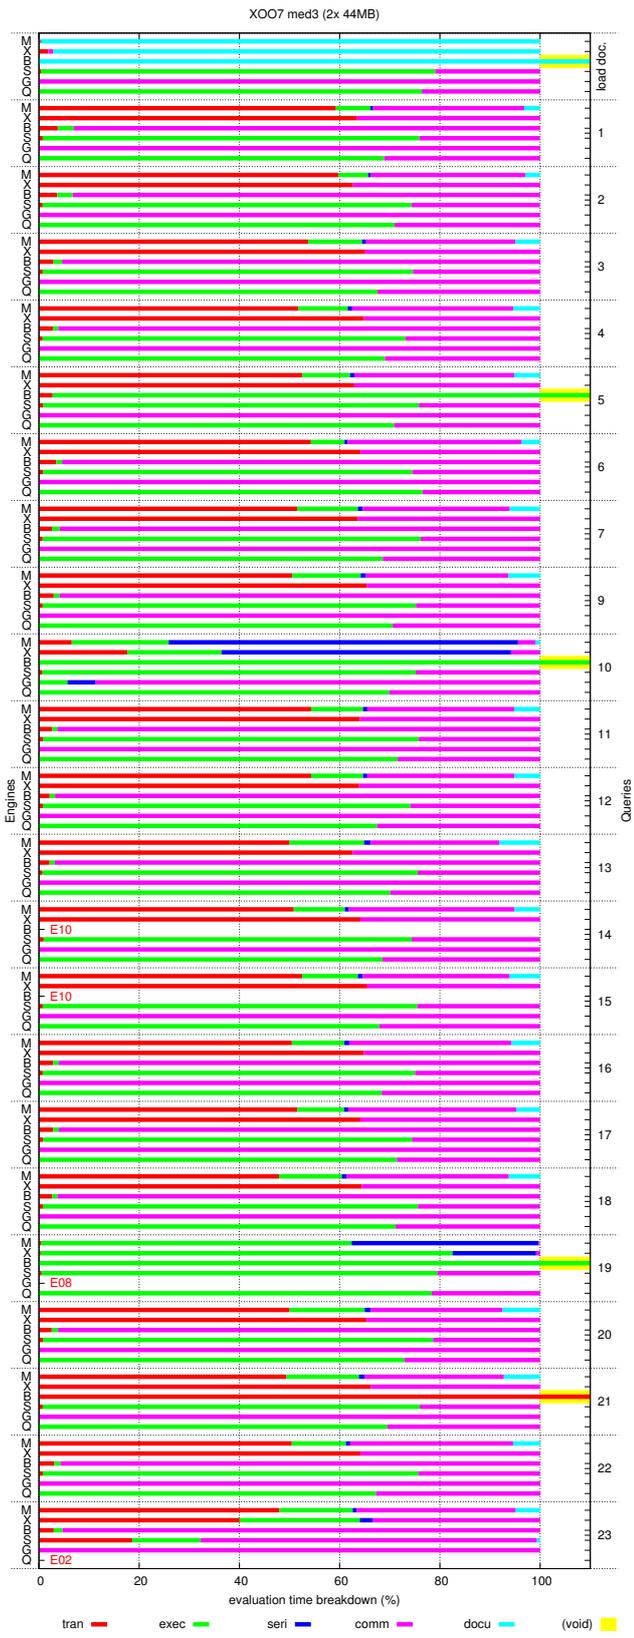


Figure 13: XOO7: Execution time breakdown

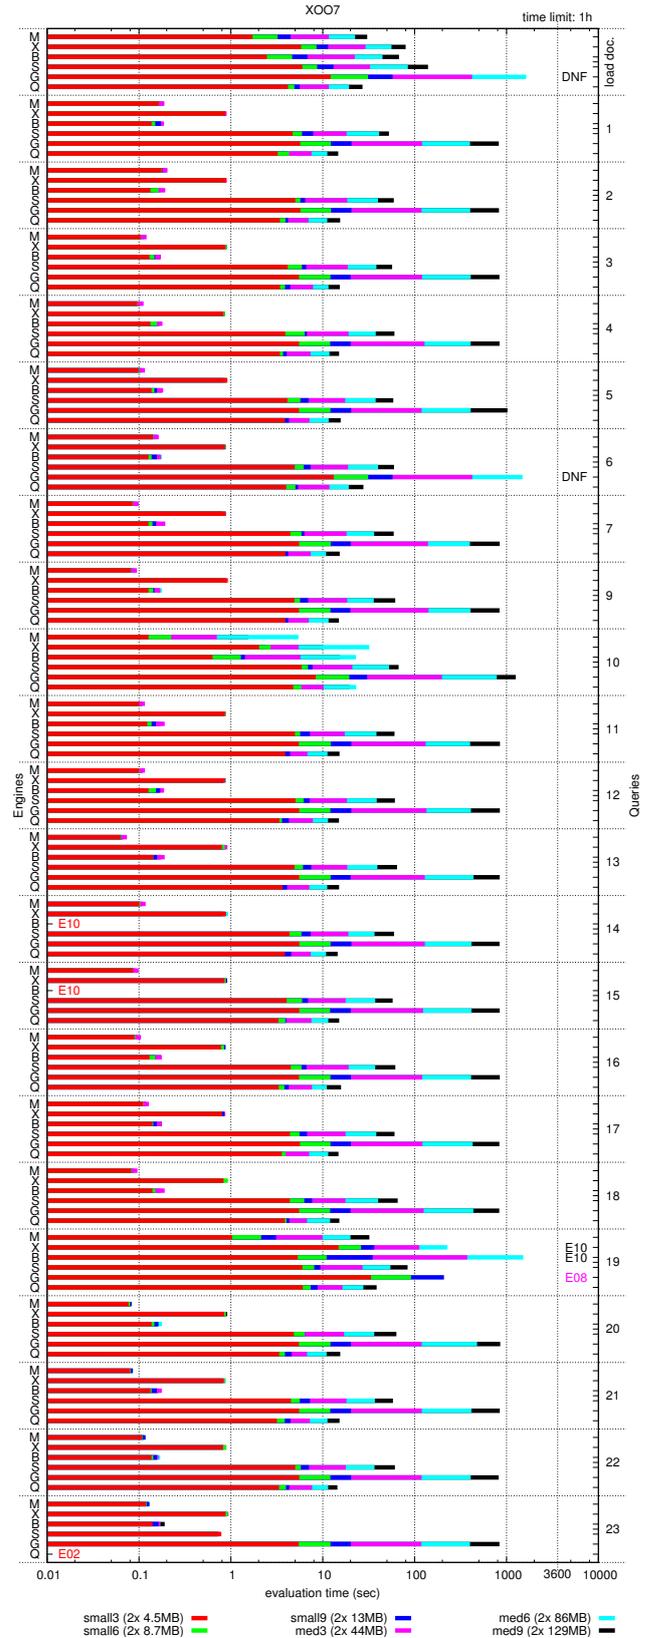


Figure 14: XOO7: Scalability

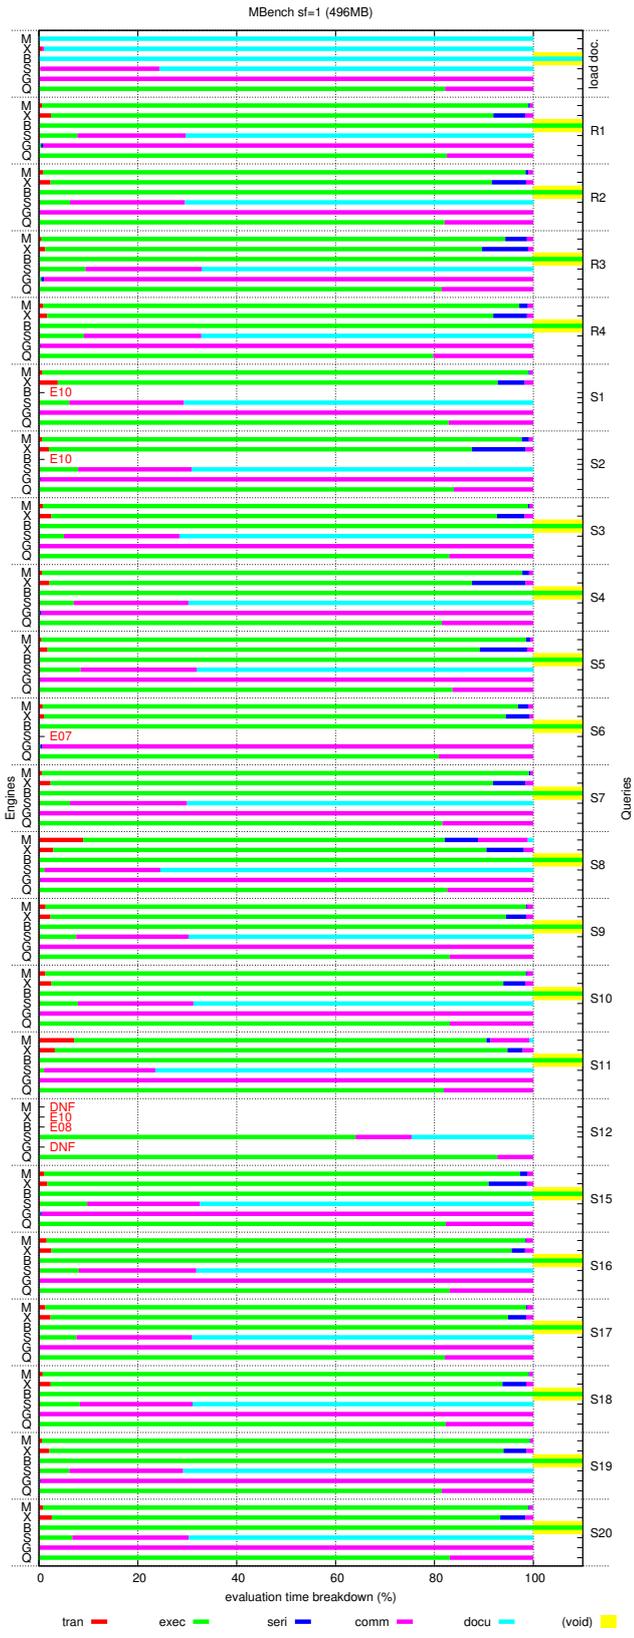


Figure 15: MBench: Execution time breakdown (1/2)

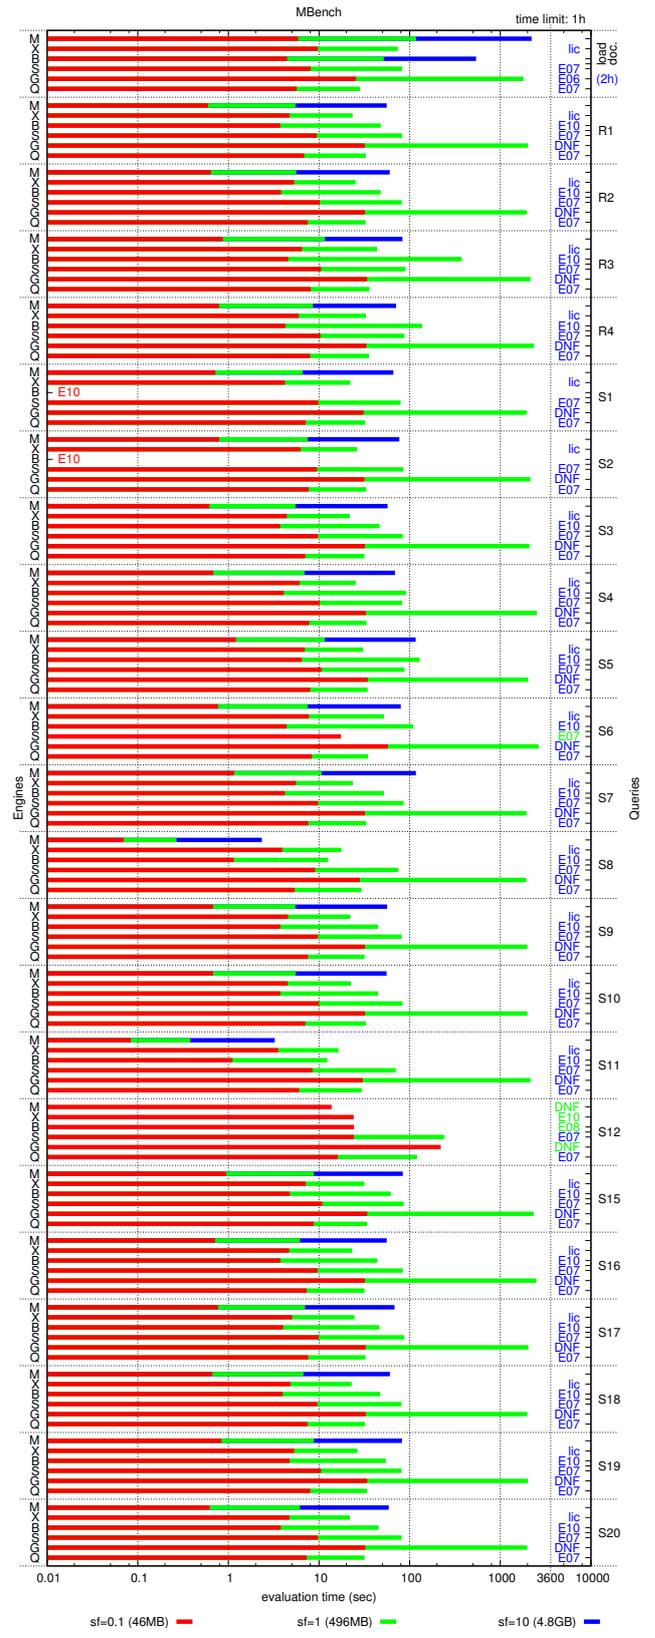


Figure 16: MBench: Scalability (1/2)

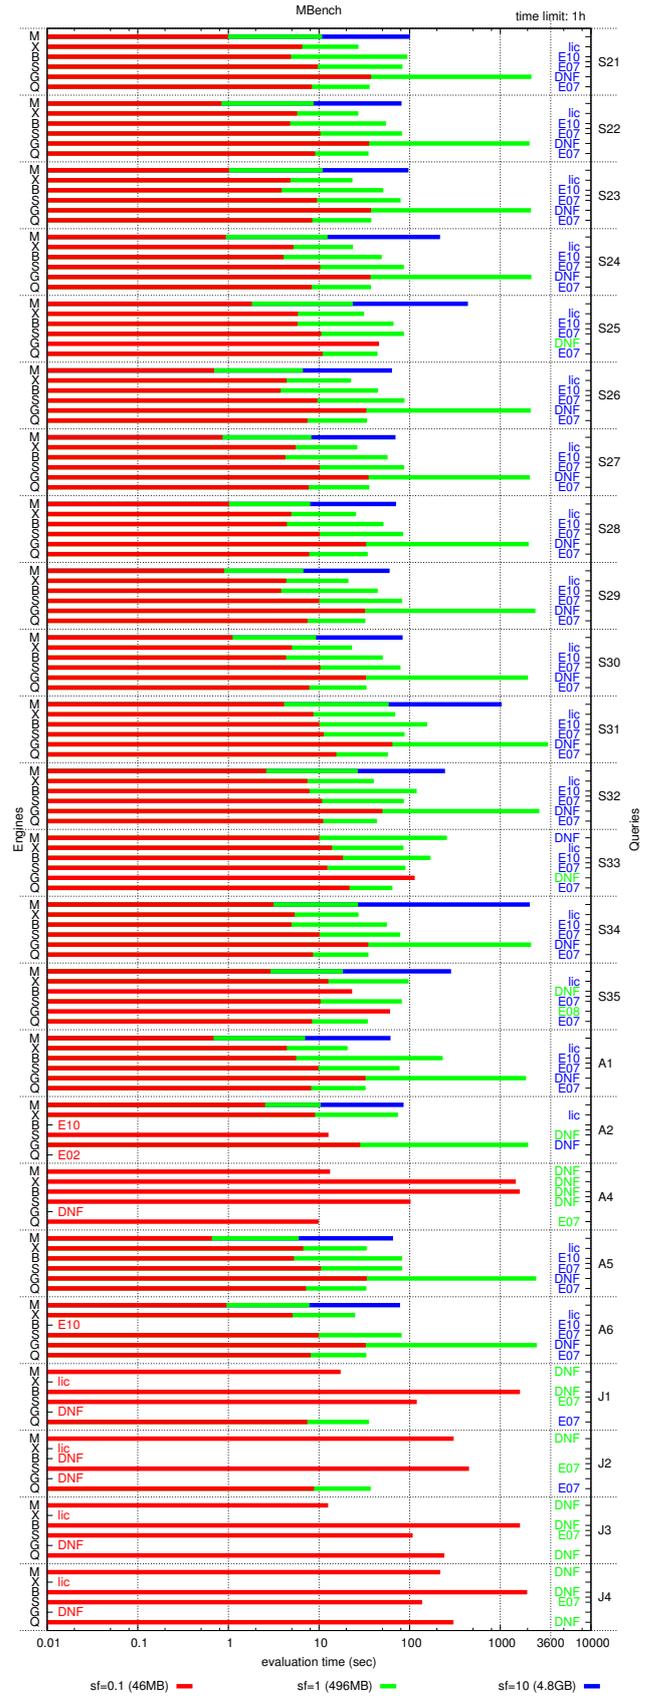
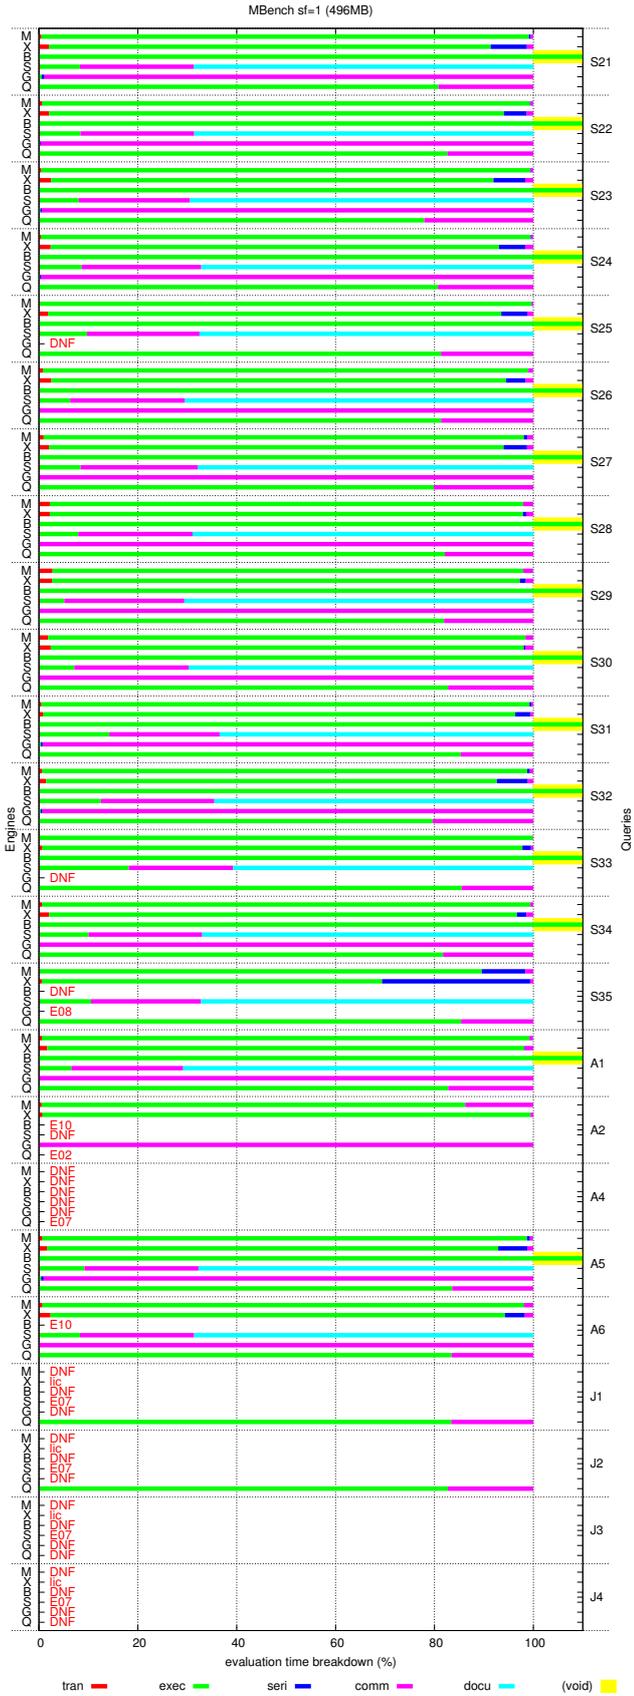


Figure 17: MBench: Execution time breakdown (2/2)

Figure 18: MBench: Scalability (2/2)