

ISSN: 2326-4411 (Print)
ISSN: 2326-442X (Online)

International Journal of Big Data

**Big
DATA**



A Technical Publication of the Services Society

2015

VOL. NO.
0201

COMPARING NoSQL DATABASES WITH A RELATIONAL DATABASE: PERFORMANCE AND SPACE

João Ricardo Lourenço, Bruno Cabral, Jorge Bernardino, Marco Vieira
CISUC, ISEC

joaoml@student.dei.uc.pt, bcabral@dei.uc.pt, jorge@isec.pt, mvieira@dei.uc.pt

Abstract

The continuous information growth in current organizations has created a need for adaptation and innovation in the field of data storage. Alternative technologies such as NoSQL have been heralded as the solution to the ever-growing data requirements of the corporate world, but these claims have not been backed by many real world studies. Current benchmarks evaluate database performance by executing specific queries over mostly synthetic data. These artificial scenarios, then, prevent us from easily drawing conclusions for the real world and appropriately characterize the performance of databases in a real system. To counter this, we used a real world enterprise system with real corporate data to evaluate the performance and space characteristics of popular NoSQL databases and compare them to SQL counterparts. We present one of the first write-heavy evaluations using enterprise software and big data. We tested Cassandra, MongoDB, Couchbase Server and MS SQL Server, comparing their performance and total used space while handling demanding and large write requests from a real company with an electrical measurement enterprise system.

Keywords: NoSQL, Big Data, Enterprise, Write-Heavy, MongoDB, Couchbase, Cassandra, SQL Server

1. INTRODUCTION

Big Data topic has recently taken more prominence in the industry, with very high amounts of data in need of constant, quick, and always-available processing (Chen et al., 2014). Traditional relational systems are known for their ACID (Atomicity, Consistency, Isolation, Durability) properties and consistency guarantees, but this design choice may often limit their availability and scalability (Leavitt, 2014). To counter this scenario, NoSQL systems have been developed. They sacrifice some of the ACID properties, namely consistency, favoring availability (Moniruzzaman et al., 2013), resulting in BASE (Basically Available Soft-state services with Eventual-consistency) systems. Indeed, these systems are Basically Available, have a Soft state during which consistency is not yet assured, but then are Eventually consistent (Brewer, 2000). Trading off consistency for availability has resulted in many different NoSQL systems which have different architectures and different use-case

scenarios (Hecht et al., 2011). They have been evaluated multiple times in recent years, but there seems to be a lack of real-world studies with enterprise data or systems (Lourenço et al., 2015, August). Rather, the NoSQL literature is mostly based on artificial benchmarks such as the Yahoo! Cloud Serving Benchmark (YCSB) (Cooper et al., 2010).

We were given a real enterprise system with a write-heavy workload and big datasets, and asked to evaluate the feasibility of replacing part of its storage backend with a NoSQL based one. This allowed us to test the performance of NoSQL systems in the real world, as opposed to the aforementioned artificial benchmarks, with respect to throughput and storage requirements. To further our analysis, we also tested space usage for each of these databases, comparing them and drawing conclusions regarding the more space-effective solution. As we will show in the rest of this paper, although the workloads of the enterprise system are not representative of big data itself, the limits of the system were being reached, with bottlenecks in the database operation becoming

clear. To solve this problem, techniques for handling big data were experimented, in particular NoSQL technology. It is the data model associated with the system that justifies our experiments and research into big data technologies.

In this paper we present the problem, our findings and the conclusions we have drawn from our experiments. Several popular NoSQL systems were tested (MongoDB, Cassandra, Couchbase Server) and compared with the relational model (MS SQL Server) to assert if they presented any performance gains in a real enterprise scenario. MS SQL Server is representative of the relational world, and our results intend to be broader, comparing NoSQL database performance and space usage with relational equivalents. To the best of our knowledge, our work is one the first to shed light on a real world write-heavy enterprise scenario, not relying on an artificial system and data. Our results may serve as a bridge to fill in the literature gap between theory and practice, artificial and real, by depicting the impact of NoSQL on a real world enterprise application that handles large amounts of data. Henceforth, our main contribution consists of showing how NoSQL systems – a typical big data solution – behave in a write-heavy enterprise scenario, which made use of relational technology, with regards to performance and space. We once again note that this research was aimed at asserting how big data related technologies behaved in a real-world scenario, and to which degree they could help overcome bottlenecks in relational technology in such scenarios.

The remaining of this paper is structured as follows. Section 2 presents a state of the art on NoSQL systems which supports our choice of tested databases, together with related work. Section 3 presents Cassandra, MongoDB and Couchbase Server, highlighting their features and characteristics. Section 4 introduces the enterprise system we worked with, the problems at hand, the current architecture and our proposed architecture. Section 5 explains the experiments we ran and their setup. Section 6 presents our results and their discussion. Finally, in Section 7 we draw the conclusions of our work and future work opportunities, also focusing on the lessons learned from our experiments and suggesting a set of defining characteristics that NoSQL systems need to better support our enterprise scenario.

2. STATE OF THE ART

With a large amount of available NoSQL solutions, it is important to be able to distinguish the fittest database solution for a particular system with respect to its distinctive characteristics. All the continuous development and evolution of non-relational technology, over the past years, has contributed to the constant interest in evaluating NoSQL databases. Moreover, until now, all the available studies were highly focused on the performance testing using standard benchmarks (Zhong et al., 2014), (Lourenço et al., 2015a). Although those evaluations provide a basic knowledge of the database behavior, there is no performance guarantee while working in a real enterprise environment, where data and interaction are much more random, unpredictable and hard to model (Zhong et al., 2014), (Abramova et al., 2013).

In recent NoSQL evaluations, the authors focus on different possibilities of adapting NoSQL solutions and using those databases along with other domain-specific technologies, such as clinical decision support systems (Chen et al., 2015), (Mazurek, 2014). In these studies, the authors evaluated different possibilities of integrating NoSQL databases in existing systems where flexibility or big data handling capabilities were needed. They concluded that, according to each system characteristic, NoSQL, in fact, could be a good possibility for data management. In particular, its flexibility and scalability were seen as fitting for the needs of each of these works. One of the drawbacks of NoSQL databases is the learning curve (Lourenço et al., 2015, February). Another recent trend has been the development of different approaches in terms of data queries. While most developers and DBAs are comfortable and accustomed with SQL, the querying and management of non-relational databases requires more time to learn (Stonebraker, 2010). More than that, NoSQL technology is known for the non-existence of a querying standard, with different databases having different querying languages (Stonebraker, 2010), (Stonebraker, 2011). Therefore Bach et al., 2014) and (Sellami et al., 2014) describe querying in NoSQL databases and a possibility of a modeling system that is capable of executing queries regardless of the database. Finally, other authors propose efficient data mining and Big Data processing (Lomotey et al., 2014). Their frameworks are capable of providing better data querying and analysis. In this paper we evaluate the writing performance of NoSQL databases in a real working system. This evaluation will generate new insights on how NoSQL databases perform on systems with such a specific set of requirements.

Further reviewing the literature, one can gather that although there have been many evaluations, with a specific focus on synthetic data, there are few evaluations focused on write-heavy datasets (we note, however, that there are some read-heavy evaluations, such as (Dory et al., 2011) and (Silva et al., 2014)). Regarding write-heavy workloads, the authors of (Datastax, online) focus on many scenarios, one of them being “mostly writes”, but use the artificial YCSB benchmark. In much the same way, in (Tudorica et al., 2011), a write intensive analysis is performed for various NoSQL Databases, but, again, the YCSB framework is used. In (Floratos et al., 2012), an in-depth analysis of various factors with varying datasets, one of them write-heavy, is made. However, the authors also use artificial frameworks for their tests. When YCSB was first introduced, the authors published results with the benchmark (Cooper et al., 2010), analysing read, write and other characteristics. The results among these papers, for the same databases, tend to vary – a typical example of how NoSQL is constantly evolving. Thus, while there is indeed data on the write-heavy performance of NoSQL, it is mostly artificial, disregarding the enterprise applicability and its context. In that sense, similarly to our work, Zhong et al. (Zhong et al., 2014) present three enterprise scenarios where NoSQL is used. However, unlike our work, their scenarios do not represent mature enterprise systems, and are very much far apart from a write-heavy scenario. Their work, and future developments, might provide valuable contributions to the enterprise evaluation of NoSQL, but it neglects write-heavy scenarios.

3. OVERVIEW OF EVALUATED DATABASES

After evaluating several databases, a final group of three different databases was selected for testing. These were the only that met all our criteria, running in the Windows operating system, allowing for transparent sharding and having .NET bindings and drivers. In this section, we present an overview of Cassandra, MongoDB and Couchbase Server, highlighting their features and characteristics.

3.1 CASSANDRA

Cassandra is an open-source shared-nothing NoSQL column-store database developed and used in Facebook (Gudivada et al., 2014), (Haughian, 2014), (Kuznetsov et al., 2014). It is based on the ideas behind Google BigTable (Chang et al., 2008) and Amazon Dynamo (Decandia et al., 2007).

Cassandra is similar to BigTable in what concerns the data model. The minimal unit of storage is a column, with rows consisting of columns or super columns (nested columns). Columns themselves consist of the name, value and timestamp, all of which are provided by the client. Since it is column-based, rows need not have the same number of columns (Kuznetsov et al., 2014).

Cassandra supports a SQL-like language called CQL, together with other protocols (Kuznetsov et al., 2014). Indexes and secondary indexes are supported, and atomicity is guaranteed at the level of one table row. Persistence is ensured by logging. Consistency is highly tunable according to the desired operation – the application developer can specify the desired level of consistency, trading off latency and consistency. Conflicts are resolved based on timestamps (the newest record is kept). The database operates in master-master mode (Gudivada et al., 2014), where no node is different from another, and combines disk-persistence with in-memory caching of results, resulting in high write throughput operations (Haughian, 2014), (Gudivada et al., 2014). The master-master architecture makes it easy for horizontal scalability to happen (Haughian, 2014). There are several different partitioning techniques and replication can be automatically managed by the database (Haughian, 2014).

3.2 MONGODB

MongoDB is an open-source document-oriented database written in C++ and developed by the 10gen company. It uses JSON (data is stored and transferred in a binary, more compact form named BSON), allowing for a schemaless data model where the only requirement is that an id is always present (Kuznetsov et al., 2014), (Haughian, 2014).

MongoDB's horizontal scalability is mainly provided through the use of automatic sharding (Haughian, 2014). Replication is also supported using locks and the asynchronous master-slave model, meaning that writes are only processed by the master node and reads can be made from both the master

node and from one of the slave nodes. Writes are propagated to the slave nodes by reading from the master's *oplog* (operation log) (Haughian, 2014). Database clients can choose the kind of consistency models they wish, by defining whether reads from secondary nodes are allowed and from how many nodes the confirmation must be obtained.

Document manipulation is a strong focus of MongoDB, as the database provides different frameworks (e.g. MapReduce and Aggregation Framework) and ways of interacting with documents (Kuznetsov et al., 2014). These can be queried, sorted, projected, iterated with cursors, aggregated, among other operations. The changes to a document are guaranteed to be atomic. Indexing can be used on one or several fields (implemented using B-trees), with the possibility of using two-dimensional spatial indexes for geometry-based data (Kuznetsov et al., 2014). There are many different programming interfaces supported by MongoDB, with most popular programming languages having native bindings. A REST interface is also supported (Kuznetsov et al., 2014).

3.3 COUCHBASE

Couchbase is a combination of Membase (a key-value system with memcached compatibility) and CouchDB. It can be used in key-value fashion, but is considered a document store working with JSON documents (similarly to CouchDB) (Kuznetsov et al., 2014).

Documents, in Couchbase, have an intrinsic unique id and are stored in what are called data buckets. Like CouchDB, queries are built using MapReduce views in Javascript. The optimistic locking associated with an append-only B-tree is also implemented like in CouchDB. The default consistency level is eventual consistency (due to MapReduce views being constructed asynchronously). There is also the option of specifying that data should be indexed immediately (Kuznetsov et al., 2014).

A major difference when comparing Couchbase with CouchDB regards sharding (Kuznetsov et al., 2014). Whereas CouchDB does not natively support sharding (there are projects, such as CouchDB Lounge (Kuznetsov et al., 2014) which enable this), Couchbase comes with transparent sharding off-the-shelf, with application transparency. Replication is also a major point of difference between the two databases, as couchbase supports intercluster and

intracluster replication. The latter is performed within a cluster, guaranteeing immediate consistency. The former kind of replication ensures eventual consistency and is made asynchronously between geographically distributed clusters (conflict resolution is performed in the same way CouchDB does it). This database is mostly intended to run in-memory, so as to hold the whole dataset in RAM (Cudré-Mauroux et al., 2013), (Kuznetsov et al., 2014).

4. A REAL ENTERPRISE DISTRIBUTED SYSTEM

The enterprise system we worked with deals with storing several measurements coming from electrical components, obtained at high sample rates. Data is made up of tuples containing mostly simple floating point values, a timestamp and some additional metadata. This system, as it is currently deployed, is hindered by a performance bottleneck during large batches of write-only operations, which are issued to a central database.

By itself, the system uses a distributed architecture. The data are generated by specific equipment and stored in an accessible way – be it XML or a private database. Another set of independent processes eventually gather the data, do some very simple processing and then write said data to the central server. Since there are so many records being generated per minute in some of these systems (as many as 3 million values in every sampling timeslice), very high loads of data can be gathered at the same time and quickly sent to be written to the SQL server, resulting in large batches of write operations. These are the operations which cause the bottleneck and performance hit in the system.

The experiments consisted, then, of modifying the database backend of the system to reduce the bottleneck, as well as to reduce disk space usage, making use of NoSQL systems. A more detailed view of the currently implemented solution and our proposed changes is presented in the next sections.

4.1 ORIGINAL SYSTEM – RELATIONAL

The original system uses a centralized MS SQL Server database where all the gathered data is stored.

Furthermore, additional data and metadata needed for other application functionality is stored in other tables in the same database. Data is gathered from all the necessary sources and is written, in parallel, using batch operations, resulting in hundreds of thousand write operations per second. The data itself is very lightly preprocessed from the original source and, indeed, is merely dumped into a table in this database. While considering a relational system for a large number of inserts it is important to disable indexing and some integrity constraints, something which we took into account when we were performing the experiment.

This approach favours the use of well established relational technology with all ACID properties. However, as we have shown, it is prone to bottlenecks and performance problems, suggesting that an alternative might provide an appropriate trade-off between ACID compliance and performance. In this sense, we proposed a system using NoSQL that is now detailed.

4.2 PROPOSED SYSTEM – NOSQL

We sought to minimize the changes to the current system, and only touch areas that really had to be changed. Furthermore, since the system's bottleneck was in a specific table with many write operations, and not in the rest of the system itself, we decided that the original SQL database should be kept in order to store metadata and other application control data. However, we would take the tables where most of the data was being written to and place them in a NoSQL system. Thus, the data causing the bottlenecks, and the code that handled it, had to be moved and adapted to NoSQL. Given that we were looking for ideal picks for a write-heavy scenario, and that NoSQL's strength lies in large part in its high horizontal scalability (Pokorny, 2009), it was also decided that we would run a cluster with sharded/partitioned NoSQL databases in order to maximize write throughput. One of the basic NoSQL considerations is the possibility to achieve both performance and throughput by parallelizing the execution of requests. Therefore, we exploited this in the setup and, ensuring the network could handle all the data, and that there were no limitations from the source, the system should be able to handle the demand and overall performance would increase.

The original system was already finished and in deployment, which meant that our database choice

was conditioned by the system architecture. In this sense, we were limited to databases which would run in a Windows environment and that had .NET bindings. We would also not consider any non-qualified and non-robust code since it could compromise the entire evaluation. Therefore, we chose some of the most popular databases with high support and documentation. To this end, we tested Cassandra, MongoDB and Couchbase Server. One of the drawbacks in satisfying this database criterion was that we were not able to choose any Key-Value Store database – the final chosen solutions were two document stores (MongoDB, Couchbase) and one wide column store (Cassandra). These three choices are among the most popular NoSQL solutions (Datastax, online), (Strauch et al., 2011) and make up for what we believe is a reasonable sample of the best NoSQL technology.

5. EXPERIMENTAL PROCEDURE

The goal of the experiment was that of measuring the performance of the databases under a real write-heavy scenario. This was the most important quality attribute of the system, so throughput was measured to assess it.

5.1 EXPERIMENTAL SETUP

A homogeneous cluster using four machines with similar hardware was chosen to host the databases. This choice was cost-effective, adequate to the budget that we were allowed, and also mimicked the real-world scenarios that the system is used in. Two additional machines were used to perform the write-heavy queries, simulating real data sources. Figure 1 depicts the experimental setup, and Table 1 summarizes the machine characteristics.

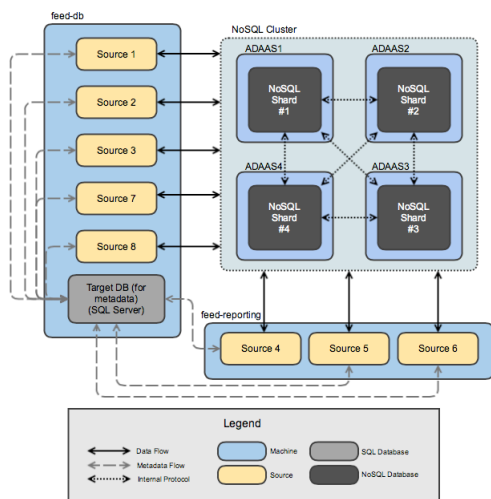


Figure 1. Experimental Setup

CPU + OS	RAM	HDD	# Nodes
Intel Core i3 3.10GHz Windows Server 2012 64 bit	4GB	WD5000AAKX 7200 RPM 16MB	4
Intel Core i7 2.9GHz Windows Server 2008 64	8GB	N/A	1
Intel core 2Duo 2.53GHz Windows Server 2008 32 bit	4GB	ST9320325AS 5400 RPM 8MB	1

Table 1. Summary of the Machine Specifications

In order to ensure that there was no bottleneck in the experimental setup, several tests were made. In particular, we tested network throughput and stability by utilizing all the available bandwidth with traffic generated by an application between the machines, analyzing the throughput and ensuring it did not have spikes. The data sources were also tested for potential

bottlenecks by running read-intensive workloads in the source machines and comparing the read throughput to the write throughput of the cluster. Upon comparison, read throughput was higher than write throughput, eliminating the possibility of a bottleneck. When measuring the write throughput with local batch inserts on the clusters and comparing it to network bandwidth, we confirmed that the latter was higher than the former, eliminating possible network bottlenecks. Finally, to make sure that any bottleneck would only be in the database itself, we ran increasingly write-intensive workloads on the database until performance peaked, showing that an increase in workload had a corresponding increase in throughput, without exceeding network bandwidth or disk write throughput.

Since the last two machines (responsible for providing the load to the first four machines) had different characteristics, the number of load generating processes they could run in parallel was also different. We ran initial tests in “ad-hoc” fashion to find the best number for each of these machines, culminating with the total of 8 sources, 5 in one and 3 in another one. These tests involved varying the number of data sources in each machine and evaluating throughput performance. The total of 8 sources, divided in groups of 5 and 3, made the system usable and didn't hinder its performance (using lower or higher values than eight would lead to system performance decrease for our initial tests). Lastly, it is of relevance noting that the machines in the cluster did not have a Solid State Drive (SSD) storage medium.

5.2 EXPERIMENTAL METHODOLOGY

The experiments were performed by running the enterprise system in controlled fashion. This enterprise system, described in the previous section was setup to mimic real enterprise distributed systems. In each run of the experiment, the source machines had a number of source processes executing write operations on the cluster, as in the real-world scenario (in essence, this simulated the large batches of data being generated on the fly by electrical equipment). All of these sources inserted the same number of records concurrently, putting stress on the databases. Henceforth, the total amount of records being inserted per workload was the sum of all records inserted by these sources. When the insertions were finished, the total insertion time was then measured so that throughput could be calculated. In addition, the space used by the database was

measured, taking into account data and metadata. A record structure is presented in Figure 2.

Field	1	2	3	4	5	6	7	8	9	Total
Type	String	int	int	int	Date	String	int	double	Date	
Size (bytes)	36	4	4	4	8	50	4	8	8	126

Figure 2. Record structure of the enterprise system we tested

Upon measuring the total insertion time and space-related data, the system was stopped and the databases reset to their initial state, allowing for repetitions of the test or for another test with a different workload to be performed.

During the adaptation of the code for the several NoSQL solutions, we tried to use database best practices, for instance using batch inserts wherever possible. However, we noticed that current NoSQL databases don't seem to have a wide support for large batch insert operations, something which needs to be improved if they are to ever be used with this kind of dataset. While the operations themselves are supported, performance gains are lower than expected, documentation was scarce and code maturity seemed low overall. Nevertheless, they did provide speedup to the operations.

5.3 WORKLOADS AND DATABASE CONFIGURATIONS

The testing procedure described in the previous section was used for a single run. In the experiment, multiple runs were executed to attain statistical significance. In particular, 10 tests were used and the final result is the average. This value was chosen as the best trade-off between the time to run all tests and the significance of the results – with 10 tests, the confidence interval for the throughput had an amplitude of roughly 5% of the total throughput, with some exceptions for higher throughput values, where the confidence interval amplitude peaked at approximately 17% (see Figure 3 where these results can be seen in the form of error bars). If a higher number of tests had been used, there would be little to gain in statistical significance, as the confidence interval amplitude was already quite low. In contrast, a lower number of tests would increase the amplitude and hinder the test significance.

In spite of using 10 tests for averaging the final result, the effective number of tests run for each experiment was 13. The first three tests were discarded to simulate a real world environment and, thus, remove cache side-effects (resulting from a cold-boot), meaning that the warm-up time was given by the first three executions.

Regarding the workloads, each of the thirteen tests was executed for 10.000 records per source (i.e., in total, 80.000 records). After that, this procedure was repeated for 20.000 records, 30.000, etc, until 330.000 per source. This way the impact of varying database load on throughput was measured, as intended. We chose 10.000 records based on real-world data and scenarios, as recommended by the system developers. The increment was also 10.000 because it was concluded, from initial testing, that a lower value did not produce significant throughput change and was, thus, not of much importance to analyse. Furthermore, throughout our tests there seemed to be no need to adjust this uniform increment or focus on any particular domain of records per source, as the throughput showed no irregular behavior and, instead, seemed to be modeled by a linear relationship with the number of records. The 330.000 limit was a consequence of our setup limitations – it was the highest load we could cope with for most tested databases.

NoSQL databases are known for their configurability and tuning (Gajendran, 2012). To maximize their performance and adapt them to this scenario, we adjusted their configuration settings.

We used Cassandra version 2.0.9, adjusting all timeouts to 90 seconds, with the exception of the range timeout, whose value was unchanged. This value allowed all tests to finish successfully. Row caching was disabled to maximize write throughput. The concurrent reads and writes settings were adjusted as suggested by the documentation, using the $8 \times \text{number_of_cores}$ rule of thumb. The memtable flush writer threads setting was set to 1, given that we had no SSDs and only had one data directory. The remaining settings were left at their default values for the results presented in this paper. However, additional tests were made by varying other parameters such as the RPC server type and other memtable settings. Since these changes produced no measurable difference in performance, their default values were used.

MongoDB version 2.6 Standard was tested with the default options and the highest write concern settings. A setup with two config servers and one

query router was chosen because it provided the minimum amount of MongoDB instances that had to be executed, allowing focus to shift to the sharding itself. For sharding to take place, the hashing of the id was used, as suggested by the MongoDB guidelines.

Couchbase Server was tested with version 3.0.1-1444 and configured to minimize the use of compaction, in order to delay performance degradation when performing IO. The number of writer threads was also increased in comparison to the number of reader threads. However, these configuration settings had little effect in performance. Instead, to achieve a 3-fold performance increase, when compared to our initial testing, the batch insertion method was adapted by using multi-threaded REST requests. The number of threads and records per thread was chosen to maximize performance, after testing several values to see which had the best performance. The resulting settings were 25 threads of 1000 records each (per source).

Regarding the relational engine, SQL Server 2014 Standard edition was used. The server was used with full ACID guarantees, as is the case in real deployments, but with all indexing facilities turned off to maximize write throughput. This server was only executed on one node, as we shall discuss in the next section (Section 6).

All databases were assigned 2GB of the available memory. This was the most memory that could be assigned without damaging normal system operation. Furthermore, since write-performance was the focus of the experiment, all indexing capabilities were disabled wherever possible. Finally, SQL Server's in-memory functionality was not used, as it was not a particularly useful, since both full ACID-compliance and high write performance were goals of the experiment. On the other hand, NoSQL systems which have such in-memory functionality as the basis of their operation (e.g. Cassandra) could not have this feature turned off (or turning it off would deeply impact performance).

6. EXPERIMENTAL RESULTS

We now present the results of the experiment, as well as our analysis. The limitations of the experiment are also presented and discussed.

6.1 PERFORMANCE ANALYSIS

The results (see Figures 3 and 4) show that the relational/original solution provided a much higher throughput than the NoSQL solution with sharding. The throughput of MS SQL Server, whilst running on a single node, is up to five times that of the best NoSQL solution. Indeed, Cassandra performs the best out of the three tested NoSQL systems by a fairly large margin (see Figure 4). These results were predictable, since the other two tested databases are document stores, fine-tuned for read operations (Gajendran, 2012). Cassandra, on the other hand, is write-optimized (Cooper et al., 2010), (Gajendran, 2012). Nevertheless, in spite of outperforming its NoSQL counterparts by a factor 4, it is itself outperformed by the original system running MS SQL server on a single node. Although one might be tempted to think that the tested workloads were not adequate for the NoSQL systems being tested, we note that MongoDB, for instance, was unable to handle the larger workloads, indicating that the choice of workload, for our experimental setup, was indeed adequate.

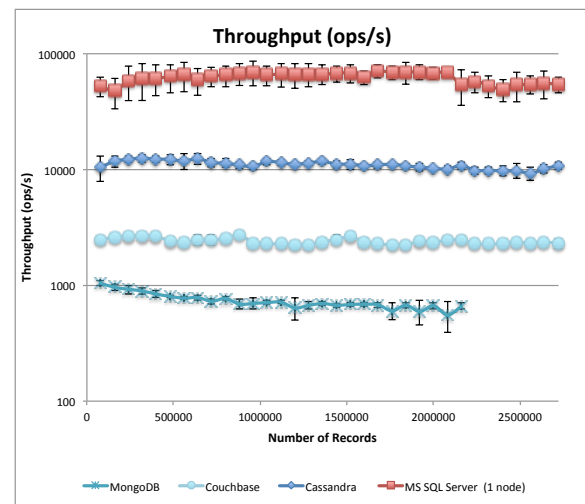


Figure 3. Throughput results for all tested databases. Error bars show standard deviation. There are missing values for MongoDB because it could not finish the test and handle higher workloads

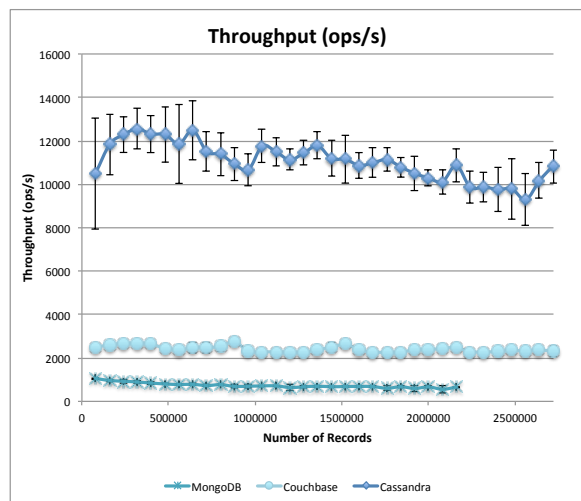


Figure 4. Throughput only for NoSQL databases. Error bars show standard deviation. There are missing values for MongoDB because it could not finish the test and handle higher workloads.

Cassandra's throughput is more irregular than that of the other NoSQL solutions, showing a decreasing trend with increasing workloads. While the other NoSQL databases also show this trend (although to a lesser extent than Cassandra), MS SQL Server does not seem to suffer as much from performance degradation with increasing workloads, although such degradation can still be seen.

MongoDB, as mentioned, was unable to cope with the highest workloads, which can be seen in Figures 3 and 4. When a certain high workload was reached, it began displaying errors due to internal problems because of the large overload of data. This result, although somewhat surprising (in the sense that a database crash is not expected behavior), can be attributed to the exhaustion of all machine resources – MongoDB ran out of resources, whereas other databases did not. Of the three NoSQL databases, it showed the lowest performance, with a very low throughput when compared to SQL Server. Even taking into account that this database is document-based and, thus, read-oriented, this result was somewhat surprising.

Couchbase has the most stable throughput among all the databases tested, although it is much lower than Cassandra and SQL Servers'. As we have already seen, this is expected due to the fact that it is a document-based database. However, contrary to MongoDB, it can cope with the high workloads.

The “flat-line” results shown in Figures 3 and 4 could be interpreted as some bottleneck in the experiment other than a database bottleneck. We dismiss this idea for two reasons: Firstly, because, as explained in Section 5, extensive tests were executed to assure there was no other possible bottleneck; secondly, because the results make sense if we take into account that our initial workload is already stressing the system – this was intended, because it reflects the real-world enterprise system, where lower workloads do not exist and would thus be meaningless to test.

We conclude that a cluster with sharding enabled, made up of a small number of machines (in our case, 4) running any of these popular NoSQL solutions cannot outperform the existing SQL Server based solution with ACID guarantees. It might be the case that there is a high overhead due to internal protocol communications and that disabling sharding would actually improve performance, instead of doing the reverse, but this is matter for future work.

6.2 SPACE USED ANALYSIS

Regarding space usage (see Figures 5 and 6), the relational solution, again, proved better. Throughout our workloads, SQL Server required less space for the same amount of data. Nevertheless, Cassandra's results were not too different from SQL Server's. This NoSQL database occupies 1.5 times the space of SQL Server for storing the data.

The document-based solutions have far higher disk requirements, as much as five times more than the relational solution. This can be explained by the document structure, where attribute names are repeated for each existent value, as well as by the internal architecture of the database which is not geared towards saving space. The sudden increases in MongoDB's disk usage can be explained by its internal chunk-based architecture that pre-allocates space in large chunks. The space necessary per record was estimated.




Figure 5. Database used space per workload for all databases. There are missing values for MongoDB because it could not finish the test and handle higher workloads.

Figure 6 shows the space needed per record. Those values represent the amount of disk (in KB) that each record requires (including the attribute values and all indexes and metadata). Therefore, the presented values contain metadata, primary (key) indexes that are generated automatically, etc. The initial variation in this value reflects both the fact that some data structures are created only once per each database and, secondly, that we may have included some additional information in our measured data. If the latter is the reason for this variation, then the fact that for increasing workloads there is a convergence to a fixed value shows that additional information remained constant (or sufficiently small when compared to the increase in data coming from the different workloads).

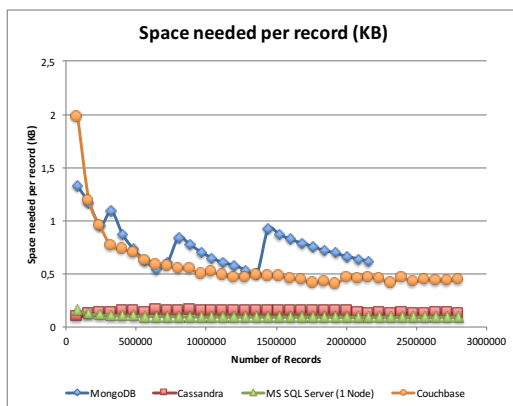


Figure 6. Space per record for each workload for all databases. There are missing values for MongoDB because it could not finish the test and handle higher workloads.

6.3 RESULT LIMITATIONS

While our results intend to forward the study of NoSQL in the real-world, they are not without their limitations. Most NoSQL databases are optimized for Solid State Drives (SSDs) and might perform poorly on systems which do not have them (Menon et al., 2014), (Baron et al., 2013). In this particular scenario, our cluster did not have SSDs, meaning it wasn't one for which these NoSQL systems were fine-tuned. The real-world application we are studying is not deployed over SSDs, so to keep our experiments as close to reality as possible, we decided to apply the same restriction on our experimental setup.

Similarly, it is a well-known fact that NoSQL systems tend to be more optimized for UNIX-like operating systems, in particular Linux (Menon et al., 2014), (Baron et al., 2013). Due to our application requirements, we had to use a Windows system. We believe that this might impact the generality of our results but provides important information regarding performance of NoSQL systems on this platform. It would no doubt be interesting to rerun our experiment on other operating systems and perform a comparative performance analysis.

Another key factor in this study regards the relatively small size of our cluster. As we have discussed previously, for such a small number of nodes, the protocol overhead might have hindered the performance of the NoSQL databases. Further tests which we do not present in this paper hinted that running these NoSQL systems on a single-node system increased their throughput, strengthening this idea that there was a protocol overhead, or some other kind of overhead related to sharding. There seems to be a trade-off inherent to the architecture of the NoSQL database, in the nature of its internal protocol and storage, between the performance gain in adding nodes to a cluster, and the protocol overhead resulting from that. These claims are in line with the findings and reasoning of (Kaviani et al., 2013) and (Martin et al., 1997), but need further research. A larger cluster size could have rendered better results if the performance gains of more nodes outweighed the protocol overhead, but this experiment alone is not sufficient to study this, leaving room for future work.

Still on the topic of the cluster size, one should also question what the outcome of running MS SQL Server in a clustered environment would be, and how our experiment does not inspect this scenario. It may be the case that the whole reason that NoSQL showed

poor results when compared with relational technologies is simply due to the aforementioned protocol overhead. By testing SQL Server in cluster mode, it would at least be able to achieve a “fair ground comparison”, although we highlight that we did not test this due to the nature of the problem we had at hands – the experiment only compared these NoSQL databases with MS SQL Server to assert the performance benefits of using NoSQL solutions in the enterprise system. In that sense, then, the single-node SQL Server was the baseline test, which would ideally improve by using NoSQL technology.

In spite of the testing setup being based in a real world scenario, contrasting with artificial benchmarks such as YCSB, the specific nature of our records might have an influence in the tests. The records at hand were very simple with no particular schema associated other than a set of integers, doubles and dates. In other scenarios, NoSQL might provide more performance, for instance by exploiting the nature of web-based data (Hecht et al., 2011), which is the basis for MongoDB and Couchbase's document oriented model (Hecht et al., 2011). The lack of a fixed schema, common to most NoSQL databases, could give performance benefits, and this characteristic is not taken into account into this work. Additionally, due to the limitations in the resources we had for our experiments, the number of concurrent applications sending records was fixed at 8. Lastly, these results apply only to the tested NoSQL database versions, which are no longer the most up-to-date ones. In particular, MongoDB 3.0 has introduced a new compression mechanism which might influence our results. We leave this issue for future work.

7. CONCLUSIONS AND FUTURE WORK

NoSQL is constantly evolving, rendering past evaluations obsolete in short time spans. This work not only fills in a gap in what concerns the link between theory and practice, but it also serves as a new up-to-date evaluation of NoSQL systems which, due to their constant evolution, quickly render past evaluations obsolete. In particular, it is, to our knowledge, the first real world test of a write-heavy enterprise application, especially one involving big data.

Our results show that a relational database engine, running on a single node, can outperform popular NoSQL solutions running in a cluster of 4 nodes with sharding by a factor of five when dealing with a

write-heavy scenario. Whether this result stems from protocol overhead in the cluster or due to the more developed nature and age-tested technology of relational databases is a question for future work. If, indeed, protocol overhead is at the root of these results, then increasing the cluster size might also increase performance without adding so much overhead, whereas reducing the cluster size might reduce protocol overhead and also provide better results. When faced with a scenario such as the one in our enterprise application, however, these results would recommend to keep the SQL Server and not opt by a NoSQL cluster.

The second part of our experiments allows us to draw conclusions regarding space usage in the relational and non-relational world. It becomes clear that relational solutions are still the best option in what concerns space. This difference might not be significant when comparing SQL Server and a column-based database such as Cassandra, but it is noticeable when we compare the relational database with document-based database. The latter can use up to five times more space. These results are not surprising due to the different nature and design goals of NoSQL databases: they are meant for a world where data replication is used to increase performance and availability, hence, also inherently increasing the space required for storing data. Nevertheless, these databases are continuously improving to lower their space requirements.

From our work, we might extract some valuable lessons regarding NoSQL and write-heavy datasets. Indeed, if one is looking for a good NoSQL solution for their write-heavy dataset, it becomes clear that a column-family database such as Cassandra will provide the best results, in spite of sacrificing some ACID properties. The protocol overhead involved in the communication inside the cluster, which we believe is one of defining factors of our results, should also be minimized and fine-tuned for large batches of write requests, but we note that this must be the subject of future work. Furthermore, NoSQL databases should have a wider support for large batch insert operations if they are to ever be used with this kind of dataset.

As future work, we propose to complement the original experiment by using Solid State Drives and machines with more resources, as well test SQL Server itself in a clustered environment. Further testing regarding cluster size might also advance the state of the art on NoSQL's practical applications. The impact of the record size is a factor not yet thoroughly reviewed in literature, and has high

practical applications, thus being a good candidate for future work. The number of sources is an interesting parameter that should be studied with other experimental setups. It is our belief that increasing the number of concurrent applications will only decrease the throughput by saturating the network connection and the databases' request handling mechanisms. In much the same way, decreasing the number of concurrent applications should increase the throughput, a direct consequence of decreasing the number of concurrent records inserted per second. We leave this experiment for future work. The issue of protocol overhead has been mentioned before and should be subject of future work – is there some cluster setup where protocol overhead outweighs the performance speedup? New versions of the tested databases are already available and it we would like to rerun our tests with them. Lastly, it would be interesting to compare the results of our real-world benchmark with those of YCSB under the same configuration.

8. ACKNOWLEDGMENTS

This research would not have been made possible without support and funding of the FEED – Free Energy Data and iCIS – Intelligent Computing in the Internet Services (CENTRO-07 - ST24 – FEDER – 002003) projects, to which we are extremely grateful.

9. REFERENCES

- Abramova, V., & Bernardino, J. (2013, July). NoSQL databases: MongoDB vs cassandra. In Proceedings of the International C* Conference on Computer Science and Engineering (pp. 14-22). ACM.
- Bach, M., & Werner, A. (2014). Standardization of NoSQL database languages. In *Beyond databases, architectures, and structures* (pp. 50-60). Springer International
- Baron, J., & Kotecha, S. (2013). Storage options in the aws cloud. Amazon Web Services, Washington DC, Tech. Rep.
- Bastiao Silva, L., Beroud, L., Costa, C., & Oliveira, J. L. (2014, June). Medical imaging archiving: A comparison between several NoSQL solutions. In IEEE-EMBS International Conference on and Health Informatics (BHI) (pp. 65-68). IEEE.
- Brewer, Eric A. "Towards robust distributed systems." PODC. Vol. 7. 2000.
- Chang, F., Dean, J., Ghemawat, S., Hsieh, W. C., Wallach, D. A., Burrows, M., ... & Gruber, R. E. (2008). Bigtable: A distributed storage system for data. *ACM Transactions on Computer Systems (TOCS)*, 26 (2), 4.
- Chen, C. P., & Zhang, C. Y. (2014). Data-intensive applications, challenges, techniques and technologies: A survey on Big Data. *Information Sciences*, 275, 314-347.
- Chen, Z., Yang, S., Tan, S., He, L., Yin, H., & Zhang, G. (2015). A new fragment re-allocation strategy for NoSQL database systems. *Frontiers of Computer Science*, 9 (1), 127.
- Cooper, B. F., Silberstein, A., Tam, E., Ramakrishnan, R., & Sears, R. (2010, June). Benchmarking cloud serving systems with YCSB. In Proceedings of the ACM symposium on Cloud computing (pp. 143-154). ACM.
- Cudré-Mauroux, P., Enchev, I., Fundatureanu, S., Groth, P., Haque, A., Harth, A., ... & Wylot, M. (2013). NoSQL databases for RDF: an empirical evaluation. In *Semantic Web—ISWC 2013* (pp. 310-325). Springer Berlin Heidelberg.
- Datastax. Benchmarking Top NoSQL Databases: A Performance Comparison for Architects and IT Managers. White paper.
- DeCandia, G., Hastorun, D., Jampani, M., Kakulapati, G., Lakshman, A., Pilchin, A., ... & Vogels, W. (2007, October). Dynamo: amazon's highly available key-store. In *ACM SIGOPS Operating Systems Review* (Vol. 41, No. 6, pp. 205-220). ACM.
- Dory, T., Mejías, B., Van Roy, P., & Tran, N. L. (2011). Comparative elasticity and scalability measurements of cloud databases. In *Proc of the 2nd ACM on cloud computing (SoCC)* (Vol. 11).
- Featherston, D. (2010, August). Cassandra: Principles and application. In *International Conference on Computing, Engineering and Information, of Illinois at Urbana-Champaign*.
- Floratou, A., Teletia, N., DeWitt, D. J., Patel, J. M., & Zhang, D. (2012). Can the elephants handle

the NoSQL onslaught?. Proceedings of the VLDB Endowment, 12), 1712-1723.

Gajendran, S. K. (2012). A survey on nosql databases. Technical report.

Gudivada, V. N., Rao, D., & Raghavan, V. V. (2014, June). Nosql systems for big data management. In IEEE World Congress on Services (SERVICES) (pp. 197). IEEE.

Hammes, D., Medero, H., & Mitchell, H. (2014). Comparison of NoSQL and SQL Databases in the Cloud. Southern Association for Information Systems (SAIS) Paper, 12.

Haughian, G. (2014). Benchmarking Replication in NoSQL Data Stores.

Hecht, R., & Jablonski, S. (2011). NoSQL evaluation: A use case oriented survey.

Kaviani, N., Wohlstadter, E., & Lea, R. (2013). Cross-tier application and data partitioning of web applications for hybrid cloud deployment. In *Middleware* (pp. 226-246). Springer Berlin Heidelberg.

Kuznetsov, S. D., & Poskonin, A. V. (2014). NoSQL data management systems. *Programming and Computer Software*, 40 (6), 323-332.

Leavitt, N. (2010). Will NoSQL databases live up to their promise?. *Computer*, 43 (2), 12-14.

Lomotey, R. K., & Deters, R. (2014, June). Terms Mining in Document-Based NoSQL: Response to Unstructured Data. In IEEE BigData Congress (pp. 661-668). IEEE.

Lourenço, J. R., Abramova, V., Vieira, M., Cabral, B., & Bernardino, J. (2015a). NoSQL Databases: A Software Engineering Perspective. In *New in Information Systems and Technologies* (pp. 741-750). Springer International Publishing.

Lourenço, J. R., Abramova, V., Cabral, B., Bernardino, J., Carreiro, P., & Vieira, M. (2015b). No SQL in Practice: A Write-Heavy Enterprise Application. In IEEE International Congress on Big Data (BigData Congress) (pp. 584-591). IEEE.

Lourenço, J. R., Cabral, B., Carreiro, P., Vieira, M., & Bernardino, J. (2015c). Choosing the right NoSQL database for the job: a quality attribute *Journal of Big Data*, 2 (1), 1-26.

Martin, R. P., Vahdat, A. M., Culler, D. E., & Anderson, T. E. (1997). Effects of communication latency, overhead, and bandwidth in a cluster architecture (25, No. 2, pp. 85-97). ACM.

Mazurek, M. (2014). Applying NoSQL Databases for Operationalizing Clinical Data Mining Models. In *Beyond Databases, Architectures, and Structures* (pp. 527-536). Springer Publishing.

Menon, P., Rabl, T., Sadoghi, M., & Jacobsen, H. A. (2014, March). CaSSanDra: An SSD boosted key-value store. In IEEE 30th Conference on Data Engineering (ICDE) (pp. 1162-1167). IEEE.

Moniruzzaman, A. B. M., & Hossain, S. A. (2013). Nosql database: New era of databases for big data analytics-classification, characteristics and arXiv preprint arXiv:1307.0191.

Ordonez, C., & Cereghini, P. (2000, May). SQLEM: Fast clustering in SQL using the EM algorithm. In *ACM SIGMOD Record* (Vol. 29, No. 2, pp. 559-570). ACM.

Pokorny, J. (2013). NoSQL databases: a step to database scalability in web environment. *International Journal of Web Information Systems*, 9 (1), 69-82.

Sellami, R., Bhiri, S., & Defude, B. (2014, June). ODBAPI: a unified REST API for relational and NoSQL data stores. In IEEE International Congress on Big Data (BigData Congress) (pp. 653-660). IEEE.

Stonebraker, M. (2010). SQL databases v. NoSQL databases. *Communications of the ACM*, 53 (4), 10-11.

Stonebraker, M. (2011). Stonebraker on NoSQL and enterprises. *Commun. ACM*, 54 (8), 10-11.

Strauch, C., Sites, U. L. S., & Kriha, W. (2011). NoSQL databases. *Lecture Notes*, Stuttgart Media University.

Tudorica, B. G., & Bucur, C. (2011, June). A comparison between several NoSQL databases with comments and notes. In 10th Roedunet International Conference (pp. 1-5). IEEE.

Van der Veen, J. S., Van der Waaij, B., & Meijer, R. J. (2012, June). Sensor data storage performance: Sql or nosql, physical or virtual. In IEEE 5th International Conference on Cloud Computing (pp. 431-438). IEEE.

Zhong, T., Doshi, K., Tang, X., Lou, T., Lu, Z., & Li, H. (2014). Big Data Workloads Drawn from Real-Time Analytics Scenarios Across Three Deployed Solutions. In *Advancing Big Data Benchmarks* (pp. 97-104). Springer International.

Authors



João Ricardo Lourenço is currently a researcher at CISUC – Centre for Informatics and Systems of the University of Coimbra. He received his bachelor degree in Informatics Engineering from the University of Coimbra in

2014. He is finishing his Master's in the same university. He has published works related to big data and NoSQL, and is currently working on exception handling models. His main research interests are distributed systems, operating systems, big data, NoSQL, and programming languages.



Bruno Cabral is an Assistant Professor at the University of Coimbra (UC), Portugal. He has more than 15 years of experience in teaching of Software Engineering, Operating Systems, and Databases, among other subjects. Bruno is currently

the coordinator of the joint Masters Program in Software Engineering, offered by the UC in partnership with Carnegie Mellon University (CMU), USA. Bruno was Adjunct Associate Teaching Professor at CMU in 2011/2012. His main research interests are Concurrent Programming and Programming Languages, Exception Handling Models and Code instrumentation. He is the author of several publications in top-tier peer reviewed conferences and journals in topics related with concurrency and programming languages, such as language support for software transactional memory, automatic exception handling models, and run-time instrumentation and optimization of code. He has participated and acted as PI in many research and software projects in cooperation with institutions such as the European Space Agency, Carnegie Mellon University, and the Portuguese Government. Bruno is also the founder and CTO of Sentilant, a spinoff company of the UC.



Jorge Bernardino received the degree in computer engineering in 1987, the masters degree in systems and information technologies in 1994, and the PhD degree in computer science from the University of Coimbra in 2002. He is a Coordinator Professor

at ISEC (Instituto Superior de Engenharia de Coimbra) of the Polytechnic Institute of Coimbra, Portugal. His main research fields are big data, data warehousing, business intelligence, and open source tools, subjects in which he has authored or co-authored dozens of papers in refereed conferences and journals. He has served on program committees of many conferences and acted as a referee for many international conferences and journals in data warehousing and databases. He was President of ISEC from 2005 to 2010. During 2014 he was Visiting Professor at CMU – Carnegie Mellon University.



Marco Vieira is an Associate Professor at the University of Coimbra, Portugal. He is an expert on experimental dependability and security assessment and benchmarking and his research interests also include fault injection, intrusion detection, software

development processes, and software quality assurance, subjects in which he has authored or co-authored more than 150 papers in refereed conferences and journals. He has participated in many research projects, both at national and European level. Marco Vieira has served on program committees of the major conferences of the dependability area and acted as referee for many international conferences and journals in the dependability and databases areas.



hipore.com

DATA

A Technical Publication of the Services Society



ISSN: 2326-4411 (Print)
ISSN: 2326-442X (Online)