

Towards an Extensible Middleware for Database Benchmarking

David Bermbach¹, Jörn Kuhlenskamp², Akon Dey³, Sherif Sakr⁴, and Raghunath Nambiar⁵

¹ TU Berlin, Information Systems Engineering Group
david.bermbach@tu-berlin.de,

² TU Berlin, Information Systems Engineering Group
j.kuhlenskamp@tu-berlin.de

³ University of Sydney, School of Information Technologies
akon.dey@sydney.edu.au

⁴ NICTA and UNSW
ssakr@cse.unsw.edu.au

⁵ Cisco Systems, Inc.
rnambiar@cisco.com

Abstract. Today’s database benchmarks are designed to evaluate a particular type of database. Furthermore, popular benchmarks, like those from TPC, come without a ready-to-use implementation requiring database benchmark users to implement the benchmarking tool from scratch. The result of this is that there is no single framework that can be used to compare arbitrary database systems. The primary reason for this, among others, being the complexity of designing and implementing distributed benchmarking tools.

In this paper, we describe our vision of a middleware for database benchmarking which eliminates the complexity and difficulty of designing and running arbitrary benchmarks: workload specification and interface mappers for the system under test should be nothing but configuration properties of the middleware. We also sketch out an architecture for this benchmarking middleware and describe the main components and their requirements.

1 Introduction

Relational database management systems (RDBMS) have been around since the 1960s and have long been considered to be a one-size-fits-all solution to data persistence. However, over the last few years, a plethora of new data storage solutions – typically referred to as NoSQL (Not Only SQL) systems – have been developed to step in where RDBMS have previously been unable to fulfill certain complex application requirements, e.g., elastic scalability. Today’s data storage systems are primarily categorized by their supported data model and their application data access interface into *column stores* (e.g., Bigtable [14])

or Cassandra [31]), *key-value stores* (e.g., Dynamo [19] or Voldemort¹), *document stores* (e.g., MongoDB² or CouchDB³), and *RDBMS* (e.g., MySQL⁴ or PostgreSQL⁵) [13, 44]. In addition to these, there are other database systems targeting special use cases, e.g., caching storage for objects (e.g., Memcached⁶ or Redis⁷) and graph-oriented data (e.g., Neo4j⁸), and the so-called NewSQL⁹ databases (e.g., VoltDB¹⁰ or NuoDB¹¹). In essence, there are hundreds of different database systems available today and the number is increasing everyday.

Choosing a single database system from this large set of available database systems for a concrete use case is a non-trivial task [41]. From an application developer’s perspective, there are certain functional requirements for a database system based on the application’s needs. From the subset of database systems fulfilling the demanded functional requirements, an application developer typically wants to select the “best” database system based on non-functional quality attributes like performance, availability, data consistency, security, cost, etc. This shows the clear need for an ability to compare different database systems in terms of their non-functional quality characteristics which is usually addressed by benchmarking.

For such a benchmark to measure these quality attributes in a meaningful way, it requires running a workload that is comparable to the workload which will eventually be generated by the application. Therefore, measurement results obtained while running different workloads have little meaning. Most benchmarks available today can either be used with a subset of database systems (e.g., TPC¹² benchmarks for RDBMS) or do not use realistic, application-driven workloads (e.g., YCSB [15] or YCSB++ [36]). This prevents a fair comparison of database systems from different categories, e.g., a column store and an RDBMS.

Furthermore, such database benchmarks should also be easy to use, i.e., it should consist of both a measurement method and a ready-to-use toolkit. Again, existing benchmarks fall short by either not including a toolkit (e.g., see the TPC Express initiative [35]) or implementing the toolkit based on design decisions which limit the toolkit’s applicability to only a subset of existing databases (e.g., YCSB [15]).

¹ project-voldemort.com

² mongodb.org

³ couchdb.apache.org

⁴ mysql.com

⁵ postgresql.org

⁶ memcached.org

⁷ redis.io

⁸ neo4j.org

⁹ NewSQL is a term used to refer to a new generation of RDBMS that attempt to provide the same scalable performance of NoSQL systems for OLTP applications while maintaining the full ACID guarantees provided by traditional RDBMS.

¹⁰ voltdb.com

¹¹ nuodb.com

¹² tpc.org

We argue that a middleware for the execution of arbitrary database benchmarks is missing. When designing a benchmark, the benchmark designer should concentrate on his core competences – namely, specifying a realistic, application-driven workload profile and means to analyze obtained measurements. Instead of having the middleware take care of the hassle of distributed benchmarking and managing the measurement infrastructure, today, a benchmark designer has to implement this from scratch for every single benchmark. The result of this is obvious; we end up with either application-driven benchmarks without a toolkit or benchmarking toolkits with flawed implementations and limited features. We are of the opinion that workload specifications and mappers should be treated as configuration parameters only.

In this vision paper, we present the first steps towards a middleware for the execution of arbitrary database benchmarks which:

- should offer an execution environment for diverse workloads,
- should not make any assumptions regarding the underlying architecture and implementation of the database under test, and
- should incorporate the measurement of performance as well as of more complex quality of service (QoS) levels.

For this purpose, we first identify the general requirements for benchmarks in section 2 and describe why these requirements are difficult to maintain without a benchmarking middleware. Based on this, we discuss the requirements such a middleware should, hence, fulfill in Section 3. Finally, we sketch out the high-level architecture we envision for this middleware (Section 4) and discuss related work (Section 5) before coming to a conclusion in section 6.

Please, note that some of the requirements from sections 2 and 3 may overlap as requirements related to the execution of benchmarks will inevitably also be requirements for a middleware solution targeting the execution of benchmarks. Still, the focus might differ depending on the section.

2 Requirements for Database Benchmarks

The process of benchmarking database systems is typically repetitive, time-consuming and tiresome [42]. To exacerbate the situation, without a good benchmark, the results of the benchmarking process can be confusing and misleading and may result in making wrong design or database system choices. The task of performing the benchmark in a new application domain or using newly developed database systems, like NoSQL systems, can be even more drawn-out due to the lack of a good benchmark. In order, to be suitable for testing the performance and usability of a wide range of database systems and to be useful in simulating a wide variety of application use cases, we define the following desirable characteristics of the benchmark.

Easy to use: In order to be suitable for a wide variety of application and cater to different types of database users, the benchmark should be easy to configure,

run, use, and extend. The results of the benchmarking process must be easy to understand so that they are suitable for making objective decisions.

Distributed Application Aware: Increasingly, typical applications are deployed as a set of distributed database clients, spread across an often large geographical area that may span continents. To successfully emulate such application scenarios, the benchmark itself must be distributed in nature to simulate these geographically distributed database clients. The framework must make it easy to define a workload simulating a distributed application; distribution and coordination of the workload should be handled efficiently and correctly. The results of the execution of the benchmark should be gathered across all the benchmark workload instances in a correct and efficient manner as though they were running on a single machine.

Negligible Impact on Results: The benchmark itself and the infrastructure needed to perform the distribution must be sufficiently light-weight so that it neither becomes a performance bottleneck nor adversely skews the recorded measurements. This means that the benchmark itself should be able to scale and that changing the implementation of the benchmark should not affect the measurement results.

Fine-Grained Measurements: Measurements obtained during the benchmark should be collected and stored in a suitably fine-grained manner so that they can be easily sliced and diced for further analysis. This does not preclude collection of aggregate metrics as long as raw, unprocessed data is captured as well. Therefore, the size of collected measurement data can become large, e.g., for high request numbers per second or long running benchmarks. Thus, in combination with the requirement for negligible impact on results, efficient data structures are needed to persist measurements.

Repeatability: The workload and operations performed during the execution of the benchmark should be repeatable. This will enable an identical workload to be run against different database systems in various configurations in order to perform comparative analysis and A-B testing. For instance, it may be needed to perform an analysis based on whether encryption is enabled or not. This requires being able to replay exactly the same sequence of operation of another benchmarking run.

Wide Applicability to Application Domains: The benchmark should closely emulate a wide variety of application use cases in the form of predefined workloads so that it can be used by application designers and architects to objectively pick a suitable database system that meets their needs. Extending the predefined workloads to incorporate application-specific operations or scenarios in order to enable an apples-to-apples comparison of the database systems under test should also be easy. There should be an ability to define workloads in the form of a mix of a variety of operations to be executed in different execution patterns to simulate a wide variety of application scenarios. This should include the ability to simulate cases of increasing and decreasing load and differing periodicities of workload intensity such as sudden spikes and troughs. Workloads should not be limited to either OLTP or OLAP applications.

Provide Suitable Abstractions: The benchmark should not make any assumptions about the specific capabilities of the database under test and be unaware of the specific database implementation. For instance, it must not know whether the database supports transactions with ACID guarantees. The interface to the database should be abstracted to enable this behavior. However, a too high abstraction level can result in workloads that only use limited software features provided by a database system under test. Thus, the implementation of specific workload scenarios becomes more difficult. Similarly, there should not be any assumptions about other QoS guarantees provided by the underlying database. Modern database systems make explicit but frequently also implicit QoS trade-offs decisions in their implementation. The benchmark should track both sides of the trade-off, i.e., different QoS dimensions. For instance, instead of making assumptions about transactional capabilities as prescribed by ACID guarantees it should track transaction isolation violations during the execution of the workload.

Support Micro-Analysis: Many evaluations of database systems use micro-benchmarks to target specific database system features, e.g., index structures, using a subset of an application scenario. To be useful in such situations, the benchmark must allow the user to define targeted patterns of operations in the benchmark that can be used to define micro-benchmarks which analyze specific database features in isolation. Based on the more high-level application workload, micro-benchmarks should be a direct result of a drill-down in the workload stack, i.e., the benchmark should allow to enable and disable all operations of the high-level, application-driven workload individually.

Support Different Deployment Topologies: The benchmark should be able to simulate scenarios in which database applications as well as the database itself are deployed in various deployment topologies including geo-distributed deployments. Increasingly, applications also use one or more database technologies simultaneously working in concert. Benchmarking these combined setups in parallel should be supported in order to simulate such distributed application scenarios.

In our opinion, a benchmarking middleware is a suitable architecture for building a benchmarking framework that fulfills the requirements above. Due to today's speed of innovation in database systems, such a middleware may even have become a necessity. This is largely because of the following reasons:

- A middleware-based architecture will enable reuse which mitigates the risk of having to rebuild the infrastructure needed for new benchmarking applications. This in turn reduces the risk of mistakes in the benchmarking application leading to distorted measurement results.
- The application developer or database administrator can focus on the actual application-specific workload instead of worrying about the infrastructure needed to run a benchmark (e.g., distributed execution and coordination or the mapping between operations and database interface).
- The middleware abstracts the measurement and collection of a wide variety of metrics and performance characteristics across the various QoS dimen-

sions. For instance, it is a non-trivial task to either detect transaction isolation anomalies as a result of transactional ACID property violations or to measure degrees of database (in-)consistency.

3 Requirements for a Benchmarking Middleware

Building on the requirements for distributed database benchmarks which we identified in section 2, we will now analyze middleware features necessary for creating middleware-supported benchmarks according to the mentioned benchmarking requirements. We will use this to identify requirements that the middleware for database benchmarking should fulfill.

In order to increase ease of use and allow for wide applicability to application domains, a ready-to-use toolkit which does not make any assumptions about the application domain or application scenario is required. This is closely linked to providing suitable abstractions as well as supporting distributed workloads and deployments, i.e., the benchmarking middleware itself should be completely unaware of anything happening above or below the middleware layer; it should, hence, be entirely *database- and application-agnostic* and be able to handle *distribution and coordination*. Ease of use also calls for the middleware tool to come bundled with *built-in basic workloads* that are ready to use.

Fine-grained measurements will create large amounts of data so that a benchmarking middleware must be able to handle such a data stream in terms of persistence, i.e., *storage of fine-grained results* is required. Due to the complexity of implementing such measurements, the measurement process itself should also be handled by the middleware – not only for well-explored QoS dimensions like latency or throughput but also for more complex dimensions like consistency or transaction isolation. Therefore, a benchmarking middleware should come bundled with *support for advanced QoS dimensions*. Finally, a benchmarking implementation with little or no impact on repeatable measurement results obviously requires the same of an underlying middleware solution, that is, an *efficient implementation* offering a *trace-based execution* of experiments.

Building on this connection between requirements for benchmarks and the corresponding requirements for a benchmarking middleware, we will now discuss each of the middleware requirements in more detail:

Database- and Application-Agnostic: The middleware should provide a one-size-fits-all framework that supports all kinds of workloads with a variety of database access patterns (e.g., changing workload intensity, OLTP and OLAP, transactional and non-transactional, complex queries and key-based access, etc.) without any assumptions on the application scenario or the application workload. In addition to this, the middleware needs to be independent of the underlying database and should be equipped with a flexible set of mappers that map and facilitate the execution of its workloads on the different data models supported by different database systems (e.g., relational, column-oriented, document-oriented, and key-value) with the flexibility to support the definition of new data models with associated mappers.

Support for Advanced QoS Dimensions: In addition to measuring the standard evaluation metrics for database systems like response time and throughput, the middleware should facilitate the measurement of advanced QoS metrics such as consistency [4, 9–12, 26, 48, 50], availability and elasticity [43], where applicable. Ideally, the measurement component of the middleware should be extensible so that it provides the ability to define new metrics and to plug-in new measurement modules for these newly defined metrics. Furthermore, for cloud-based deployments, the middleware should be able to track concrete usage of cloud resources (e.g., storage, CPU, network, etc.) in order to facilitate monetary cost benchmarking which is an important aspect of cloud environments [23].

Storage of Fine-Grained Results: The middleware needs to track and store all the raw measurement data which introduces a certain degree of complexity as raw measurement results can potentially become very large. To ease further analysis and to also offer convenient access to aggregated results, the middleware should also come bundled with a data analysis module that provides the ability to gain insights from the collected large amount of measurement data by suitably aggregating, transforming, correlating, and analyzing the raw measurement data.

Efficient Implementation: The middleware itself should not have a significant negative impact on the performance of the benchmark and the measurements of the evaluation metrics, i.e., choosing a different middleware implementation should not notably affect the measurement results.

Distribution and Coordination: The middleware should manage distribution and coordination of (potentially geographically) distributed measurement clients and do so in a way transparent to the benchmark developer.

Trace-Based Execution: A repeatable benchmark execution, e.g., for A-B testing, requires a trace-driven implementation so that a priori instead of ad hoc workload generation is the logical choice. Obviously, such an operation trace should also contain information on the measurement client issuing the operation according to the desired level of distribution. This will assert that repeated executions will issue the same operations after the same test duration from the same geo-location. Depending on whether the load balancer is considered to be part of the application or the database system, this may even require – in the case of replication – to issue these same requests also to the same replicas. This is especially important when measuring consistency behavior [8].

Built-in basic workloads: The middleware should come bundled with a basic set of built-in basic workloads which are ready to use. These can then also be used as building blocks for advanced workloads more closely resembling the actual workload of a concrete application use case. This will also serve as a way to perform an apples-to-apples comparison between competing database systems for general purpose use without a concrete application in mind. The TPC suite of benchmarks and the default workloads provided by YCSB [15] have played a similar role in the past.

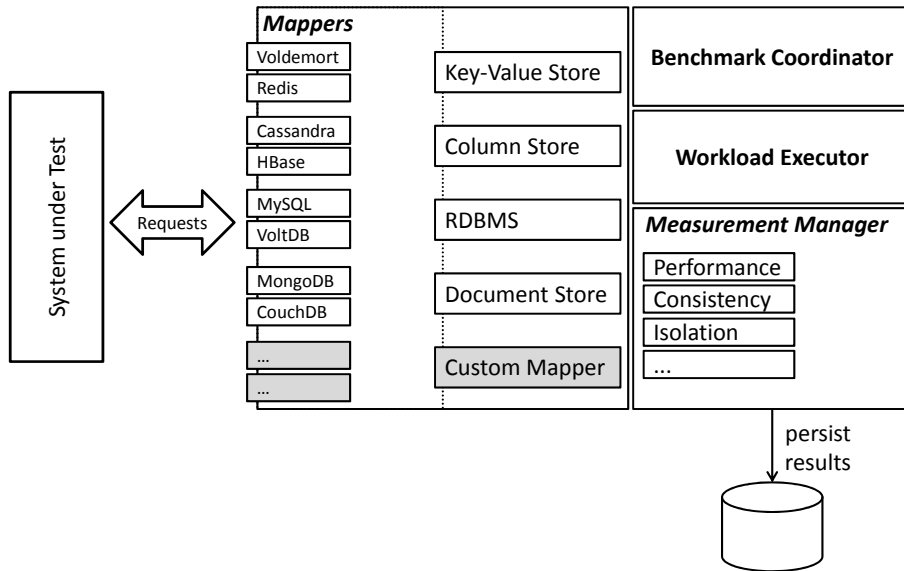


Fig. 1. High Level Architecture of the Proposed Benchmarking Middleware

4 An Architecture for a Benchmarking Middleware

To fulfill the requirements identified in the last section, a benchmarking middleware needs to be extensible in several dimensions. We aim to address these with the following high-level architecture (see also figure 1 which gives an overview):

Mappers are responsible for mapping queries to a particular database instance. This is implemented in two stages. In the first stage, queries are mapped to a particular type of database, e.g., a column store, but are not yet bound to a specific database system, e.g., Cassandra [31]. In the second stage, the abstract interface is mapped to a concrete database system. This way, standard mappers for different types of databases can be supported along with custom mappers. Furthermore, this also enables the ability to support additional database systems as well as entirely new types of database systems in the future.

For example, the first stage may choose to store an object as well as all objects referenced by this object through one-to-many relationships under the same key in a key-value store. It may also choose to do so in the JSON format. The second stage, in contrast, will then map the abstract CRUD interface of the key-value store to, for instance, Voldemort. This staged mapper approach keeps the middleware entirely agnostic of both the application workload and the underlying database system under test.

As another example, a query retrieving a data item based on three filter criteria might be mapped to an RDBMS with a query like `SELECT * FROM table WHERE a AND b AND c`. For a column store, in contrast a standard mapping

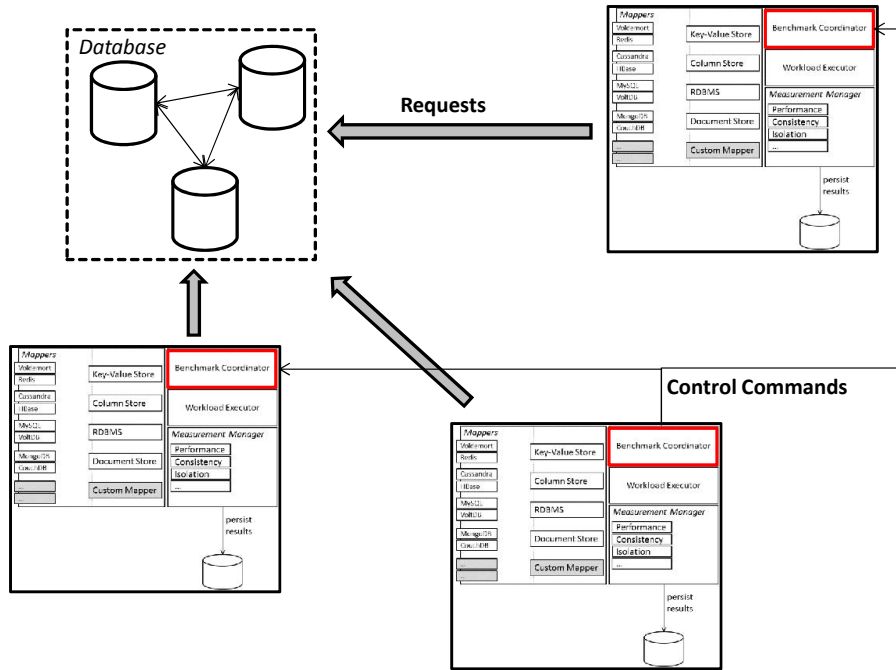


Fig. 2. Coordination Between Multiple Benchmarking Instances

might be to use the concatenated values of the fields referenced by a , b and c as row key and to, therefore, issue a get query for this row key.

The *Workload Executor* is responsible for executing all requests of the workload according to the workload specification. For this purpose, we propose to use a precomputed operation trace, i.e., a priori workload generation, to remove much of the necessary coordination effort. We imagine that a secondary tool will be used to create operation traces – either based on real-world traces or the corresponding workload description – in a standardized format, e.g., as tuples containing a relative timestamp, a SQL query and an origin location. Through this a priori workload generation, the Workload Executor will maintain the repeatability requirement. The requirement of having an efficient implementation is largely independent of the proposed architecture and mainly depends on the concrete implementation but an a priori workload generation through an external tool will certainly help the middleware layer to remain relatively thin by moving most coordination overhead to a time before the actual benchmark execution.

The *Measurement Manager* keeps track of all the individual measurement modules (e.g., performance or consistency behavior) and their individual measurement outputs. For this purpose, the manager provides detailed information to the modules (e.g., start and end time of requests or business transactions, operation results, etc.) and persists the metric output in a local database for

post-processing. This component will include predefined measurement modules to determine latency, throughput, availability, elasticity, scalability, and consistency behavior. Therefore, the requirements of being able to store fine-grained results as well as support for advanced QoS metrics will not be violated.

As a typical benchmark run will be distributed, there will be different benchmarking instances which need to be coordinated. The *Benchmark Coordinator* is, therefore, responsible for communication across instances to assert, for instance, that all benchmarking instances have the correct information available on operation traces, start times, etc. We propose a master-based approach as a single point of failure is, in this particular case, actually desirable, since failures will render benchmarking results unusable and fault-tolerance mechanisms would hide this from the user. The coordination of benchmarking instances is illustrated in Figure 2. This approach requires sufficiently precise clock synchronization as provided by NTP¹³ as well as the use of reliable messaging protocols for communication. Implementing the Benchmark Coordinator component will fulfill the requirement of handling coordination and distribution within the middleware layer.

5 Related Work

Related work on benchmarking middleware for distributed databases is scarce. Difallah et al. [21] propose an extensible testbed for the execution of benchmarks against relational databases. In particular, they argue for encapsulation of recurring functionality within a universal benchmarking infrastructure. Therefore, the work is closely related to our proposed benchmarking middleware. To our knowledge, this is the sole publication addressing not only a subset of our use case.

In the following, we discuss related work in three groups: (i) identified *benchmarking requirements* for database systems in different contexts, e.g., cloud environments, (ii) existing *benchmarking approaches* that address subsets of requirements outlined in sections 2 and 3 and (iii) specialized *foundational works* that address single functionalities of a benchmarking middleware, e.g., distributed generation of synthetic data sets.

Benchmarking Requirements: A good benchmarking middleware must identify requirements for building meaningful and useful benchmarks. Huppler [29] discusses five general characteristics of a good benchmark, i.e., *relevant*, *repeatable*, *fair*, *verifiable* and *economical*. Nambiar and Poess argue that database technologies are changing at such a rapid pace that deployment of benchmarks have become increasingly complicated. Therefore, easier means to develop and execute benchmarks are required. Smith [46] and Folkerts et al. [24] discuss requirements for benchmarks in cloud environments. Furthermore, Poess [37] and Baru et al. [6] discuss requirements for Big Data environments, respectively. Bermbach [8] discusses requirements for (consistency) metrics. Specifically, these

¹³ ntp.org

have to be *meaningful, fine-grained*, have a *high resolution* and allow *reproducible* measurement results.

Benchmarking Approaches: Benchmarking approaches are *application-driven* or *system-driven*. Application-driven approaches, i.e., end-to-end benchmarks, focus on providing realistic and meaningful workloads for an application domain. System-driven approaches are often micro-benchmarks and typically build on synthetic workload generation; they focus on a broad coverage of workloads to measure isolated database features.

Examples of application-driven approaches grouped by application domains are: OLTP [17,18], decision support [16,25], social media [5], CEP [32], ETL [49]. Examples of system-driven approaches are the Yahoo! Cloud Serving Benchmark (YCSB) [15] and Rain [7]. YCSB provides an extensible benchmarking tool with adapters for a number of distributed database systems, e.g., HBase, Cassandra and MongoDB. Two extensions to YCSB are YCSB++ [36] and YCSB+T [20]. YCSB++ adds new features to YCSB: bulk loading for databases based on B-trees, Zookeeper-based [28] start of distributed YCSB clients and distributed monitoring based on Ganglia [33]. YCSB+T extends the original workload by providing the ability to define multi-item transactions and a data validation and anomaly detection phase that can be used to classify and quantify database anomalies introduced by the workload. Rain [7] is a workload generation toolkit that provides a load scheduling mapper that can be extended with application-specific workload generators.

Different TPC benchmarks do not provide an implementation and workload models provide limited flexibility, thus, setup and customization requires additional effort. YCSB allows the generation of synthetic workloads based on a stochastic workload model. For a large number of application workloads, an accurate emulation based on the workload model is impossible or difficult to instantiate. Furthermore, support for important features, e.g., benchmark distribution and customized collection of metrics, are limited by existing benchmarks and frameworks. Our proposed benchmarking middleware closes the gap between application-driven and system-driven workload models and provides frequently used benchmarking functionals in a transparent way.

Foundational Work: To realize the different components of our envisioned benchmarking middleware, related work from different areas must be taken into account. To address this, we discuss related work with selected examples, which include, distributed data set generation for the Benchmark Coordinator, workload scheduling for the Workload Executor and complex QoS-metrics for the Measurement Manager. Gray et al. [27] discuss the generation of synthetic data sets. This work has been extended with a focus on distributed generation of data sets by Alexandrov et al. [1–3] and Rabl et al. [38,40]. In essence, both these approaches aim to provide a speed-up through parallelization with regard to the preparation time of database benchmarks. Schroeder et al. [45] propose additional parameters to be explicitly considered during workload generation,

namely a scheduling model for requests. The model differentiates between workload generators that do or do not schedule new requests independent of received responses to the preceding request. Since, trade-offs exist between QoS metrics, it is not sufficient to characterize database systems based on a single QoS dimension, e.g., performance. Among others, the following examples of benchmarking approaches address specific system qualities: (i) *consistency* [4, 9–12, 26, 48, 50], (ii) *dependability* [47], (iii) *scalability* [15, 39], (iv) *elasticity* [22, 30], and (v) *security* [34].

6 Conclusion

Historically, industry standard database benchmarks have enabled healthy competition among rival vendors that resulted in improved product offerings and was also a significant contributing factor towards the evolution of database systems themselves. While these benchmarks continue to serve both the industry and research community well, they lack some of the flexibility and extensibility required by modern cloud-based application systems in which different types of data management systems often coexist and complement each other. Further, these applications often make QoS and performance trade-offs based on a much wider set of requirements and criteria – yet, in current database benchmarks the ability to study these trade-offs is missing. In addition to this, today’s database benchmarks either do not come with a read-to-use toolkit or are limited to certain kinds of database systems and based on synthetic workloads.

In this paper, we describe our vision and architecture of a middleware for benchmarking different databases and workloads. The authors plan to further extend as well as actually implement this framework to provide the necessary infrastructure for benchmarking database systems with regards to arbitrary QoS dimensions and trade-offs, to help in determining price-performance trade-offs, and to enable modern benchmarks for studying QoS behavior of multi-tenant, frequently cloud-based, federated multi-database environments.

References

1. A. Alexandrov, C. Brücke, and V. Markl. Issues in Big Data Testing and Benchmarking. In *Proceedings of the Sixth International Workshop on Testing Database Systems*, DBTest ’13, pages 1:1–1:5, New York, NY, USA, 2013. ACM.
2. A. Alexandrov, B. Schiefer, J. Poelman, S. Ewen, T. O. Bodner, and V. Markl. Myriad: Parallel Data Generation on Shared-nothing Architectures. In *Proceedings of the 1st Workshop on Architectures and Systems for Big Data*, ASBD ’11, pages 30–33, New York, NY, USA, 2011. ACM.
3. A. Alexandrov, K. Tzoumas, and V. Markl. Myriad: Scalable and Expressive Data Generation. *Proceedings of the VLDB Endowment*, 5(12):1890–1893, Aug. 2012.
4. E. Anderson, X. Li, M. A. Shah, J. Tucek, and J. J. Wylie. What Consistency Does Your Key-value Store Actually Provide? In *Proceedings of the 6th Workshop on Hot Topics in System Dependability (HOTDEP)*, HotDep’10, pages 1–16, Berkeley, CA, USA, 2010. USENIX Association.

5. T. G. Armstrong, V. Ponnkanti, D. Borthakur, and M. Callaghan. LinkBench: A Database Benchmark Based on the Facebook Social Graph. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*, SIGMOD '13, pages 1185–1196, New York, NY, USA, 2013. ACM.
6. C. Baru, M. Bhandarkar, R. Nambiar, M. Poess, and T. Rabl. Setting the direction for big data benchmark standards. In R. Nambiar and M. Poess, editors, *Selected Topics in Performance Evaluation and Benchmarking*, volume 7755 of *Lecture Notes in Computer Science*, pages 197–208. Springer Berlin Heidelberg, 2013.
7. A. Beitch, B. Liu, T. Yung, R. Griffith, A. Fox, and D. A. Patterson. Rain: A workload generation toolkit for cloud computing applications. Technical report, University of California at Berkeley, 2010.
8. D. Bermbach. *Benchmarking Eventually Consistent Distributed Storage Systems*. PhD thesis, Karlsruhe Institute of Technology, Germany, February 2014. to be published.
9. D. Bermbach and J. Kuhlenkamp. Consistency in distributed storage systems: An overview of models, metrics and measurement approaches. In V. Gramoli and R. Guerraoui, editors, *Networked Systems*, volume 7853 of *Lecture Notes in Computer Science*, pages 175–189. Springer Berlin Heidelberg, 2013.
10. D. Bermbach and S. Tai. Eventual consistency: How soon is eventual? an evaluation of amazon s3's consistency behavior. In *Proceedings of the 6th Workshop on Middleware for Service Oriented Computing (MW4SOC)*, MW4SOC '11, pages 1:1–1:6, New York, NY, USA, 2011. ACM.
11. D. Bermbach and S. Tai. Benchmarking eventual consistency: Lessons learned from long-term experimental studies. In *Proceedings of the 2nd International Conference on Cloud Engineering (IC2E)*. IEEE, 2014.
12. D. Bermbach, L. Zhao, and S. Sakr. Towards comprehensive measurement of consistency guarantees for cloud-hosted data storage services. In R. Nambiar and M. Poess, editors, *Performance Characterization and Benchmarking*, volume 8391 of *Lecture Notes in Computer Science*, pages 32–47. Springer International Publishing, 2013.
13. R. Cattell. Scalable SQL and NoSQL data stores. *SIGMOD Record*, 39(4):12–27, 2010.
14. F. Chang, J. Dean, S. Ghemawat, W. C. Hsieh, D. A. Wallach, M. Burrows, T. Chandra, A. Fikes, and R. E. Gruber. Bigtable: A distributed storage system for structured data. In *Proceedings of the 7th Symposium on Operating Systems Design and Implementation (OSDI)*, OSDI '06, pages 205–218, Berkeley, CA, USA, 2006. USENIX Association.
15. B. F. Cooper, A. Silberstein, E. Tam, R. Ramakrishnan, and R. Sears. Benchmarking Cloud Serving Systems with YCSB. In *Proceedings of the 1st Symposium on Cloud Computing (SOCC)*, SOCC '10, pages 143–154, New York, NY, USA, 2010. ACM.
16. T. P. P. Council. Tpc benchmark ds: Standard specification version 1.1.0. Technical report, Transaction Processing Performance Council, 2012.
17. T. P. P. Council. Tpc benchmark e: Standard specification version 1.13.0. Technical report, Transaction Processing Performance Council, 2014.
18. T. P. P. Council. Tpc benchmark c: Standard specification revision 5.11. Technical report, The Transaction Processing Council, 2010.
19. G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Voshall, and W. Vogels. Dynamo: Amazon's highly avail-

- able key-value store. In *Proceedings of 21st Symposium on Operating Systems Principles (SOSP)*, SOSP '07, pages 205–220, New York, NY, USA, 2007. ACM.
20. A. Dey, A. Fekete, R. Nambiar, and U. Röhm. YCSB+T: Benchmarking web-scale transactional databases. In *Data Engineering Workshops (ICDEW), 2014 IEEE 30th International Conference on*, pages 223–230, March 2014.
 21. D. Difallah and A. Pavlo. OLTP-Bench: An Extensible Testbed for Benchmarking Relational Databases. *Proceedings of the VLDB Endowment*, 7(4):277–288, 2013.
 22. T. Dory, B. Mej, and P. V. Roy. Measuring Elasticity for Cloud Databases. In *Proceedings of the Second International Conference on Cloud Computing, GRIDs, and Virtualization (CLOUD COMPUTING'2011)*, pages 154–160, 2011.
 23. D. Florescu and D. Kossmann. Rethinking cost and performance of database systems. *SIGMOD Record*, 38(1):43–48, 2009.
 24. E. Folkerts, A. Alexandrov, K. Sachs, A. Iosup, V. Markl, and C. Tosun. Benchmarking in the Cloud: What It Should, Can, and Cannot Be. In R. Nambiar and M. Poess, editors, *Selected Topics in Performance Evaluation and Benchmarking*, volume 7755 of *Lecture Notes in Computer Science*, pages 173–188. Springer Berlin Heidelberg, 2013.
 25. A. Ghazal, T. Rabl, M. Hu, F. Raab, M. Poess, A. Crolotte, and H.-A. Jacobsen. BigBench: Towards an Industry Standard Benchmark for Big Data Analytics. In *Proceedings of the 2013 ACM SIGMOD International Conference on Management of Data*, SIGMOD '13, pages 1197–1208, New York, NY, USA, 2013. ACM.
 26. W. Golab, X. Li, and M. A. Shah. Analyzing Consistency Properties for Fun and Profit. In *Proceedings of the 30th Symposium on Principles of Distributed Computing (PODC)*, PODC '11, pages 197–206, New York, NY, USA, 2011. ACM.
 27. J. Gray, P. Sundaresan, S. Englert, K. Baclawski, and P. J. Weinberger. Quickly generating billion-record synthetic databases. In *ACM SIGMOD Record*, volume 23, pages 243–252. ACM, 1994.
 28. P. Hunt, M. Konar, F. Junqueira, and B. Reed. ZooKeeper: Wait-free coordination for Internet-scale systems. *USENIX ATC*, 2010.
 29. K. Huppler. The Art of Building a Good Benchmark. In R. Nambiar and M. Poess, editors, *Performance Evaluation and Benchmarking*, volume 5895 of *Lecture Notes in Computer Science*, pages 18–30. Springer Berlin Heidelberg, 2009.
 30. J. Kuhlenkamp, M. Klems, and O. Röss. Benchmarking Scalability and Elasticity of Distributed Database Systems. *Proceedings of the VLDB Endowment*, 7(11):to appear, 2014.
 31. A. Lakshman and P. Malik. Cassandra: A Decentralized Structured Storage System. *SIGOPS Operating Systems Review*, 44(2):35–40, Apr. 2010.
 32. C. Li and R. Berry. CEPBen: A Benchmark for Complex Event Processing Systems. In R. Nambiar and M. Poess, editors, *Performance Characterization and Benchmarking*, volume 8391 of *Lecture Notes in Computer Science*, pages 125–142. Springer International Publishing, 2014.
 33. M. L. Massie, B. N. Chun, and D. E. Culler. The ganglia distributed monitoring system: design, implementation, and experience. *Parallel Computing*, 30(7):817–840, July 2004.
 34. S. Müller, D. Bermbach, S. Tai, and F. Pallas. Benchmarking the Performance Impact of Transport Layer Security in Cloud Database Systems. In *Proceedings of the 2nd International Conference on Cloud Engineering (IC2E)*. IEEE, 2014.
 35. R. Nambiar and M. Poess. Keeping the TPC relevant! *Proceedings of the VLDB Endowment*, 6(11):1186–1187, 2013.

36. S. Patil, M. Polte, K. Ren, W. Tantisiroj, L. Xiao, J. López, G. Gibson, A. Fuchs, and B. Rinaldi. YCSB++: Benchmarking and Performance Debugging Advanced Features in Scalable Table Stores. In *Proceedings of the 2nd Symposium on Cloud Computing (SOCC)*, SOCC '11, pages 9:1–9:14, New York, NY, USA, 2011. ACM.
37. M. Poess. TPCs Benchmark Development Model: Making the First Industry Standard Benchmark on Big Data a Success. In T. Rabl, M. Poess, C. Baru, and H.-A. Jacobsen, editors, *Specifying Big Data Benchmarks*, volume 8163 of *Lecture Notes in Computer Science*, pages 1–10. Springer Berlin Heidelberg, 2014.
38. T. Rabl, M. Frank, H. Sergieh, and H. Kosch. A Data Generator for Cloud-Scale Benchmarking. In R. Nambiar and M. Poess, editors, *Performance Evaluation, Measurement and Characterization of Complex Systems*, volume 6417 of *Lecture Notes in Computer Science*, pages 41–56. Springer Berlin Heidelberg, 2011.
39. T. Rabl, S. Gómez-Villamor, M. Sadoghi, V. Muntés-Mulero, H.-A. Jacobsen, and S. Mankovskii. Solving Big Data Challenges for Enterprise Application Performance Management. *Proceedings of the VLDB Endowment*, 5(12):1724–1735, Aug. 2012.
40. T. Rabl and H.-A. Jacobsen. Big Data Generation. In T. Rabl, M. Poess, C. Baru, and H.-A. Jacobsen, editors, *Specifying Big Data Benchmarks*, volume 8163 of *Lecture Notes in Computer Science*, pages 20–27. Springer Berlin Heidelberg, 2014.
41. S. Sakr. Cloud-hosted databases: technologies, challenges and opportunities. *Cluster Computing*, 17(2):487–502, 2014.
42. S. Sakr and F. Casati. Liquid Benchmarks: Towards an Online Platform for Collaborative Assessment of Computer Science Research Results. In *TPCTC*, pages 10–24, 2010.
43. S. Sakr and A. Liu. Is Your Cloud-Hosted Database Truly Elastic? In *SERVICES*, pages 444–447, 2013.
44. S. Sakr, A. Liu, D. M. Batista, and M. Alomari. A Survey of Large Scale Data Management Approaches in Cloud Environments. *IEEE Communications Surveys and Tutorials*, 13(3):311–336, 2011.
45. B. Schroeder, A. Wierman, and M. Harchol-Balter. Open Versus Closed: A Cautionary Tale. In *Proceedings of the 3rd Conference on Networked Systems Design & Implementation - Volume 3*, NSDI'06, pages 18–18, Berkeley, CA, USA, 2006. USENIX Association.
46. W. Smith. Characterizing Cloud Performance with TPC Benchmarks. In R. Nambiar and M. Poess, editors, *Selected Topics in Performance Evaluation and Benchmarking*, volume 7755 of *Lecture Notes in Computer Science*, pages 189–196. Springer Berlin Heidelberg, 2013.
47. M. Vieira and H. Madeira. A Dependability Benchmark for OLTP Application Environments. In *Proceedings of the 29th International Conference on Very Large Data Bases - Volume 29*, VLDB '03, pages 742–753. VLDB Endowment, 2003.
48. H. Wada, A. Fekete, L. Zhao, K. Lee, and A. Liu. Data Consistency Properties and the Trade-offs in Commercial Cloud Storages: the Consumers' Perspective. In *Proceedings of the 5th Conference on Innovative Data Systems Research (CIDR)*, pages 134–143, January 2011.
49. L. Wyatt, B. Caufield, and D. Pol. Principles for an ETL Benchmark. In R. Nambiar and M. Poess, editors, *Performance Evaluation and Benchmarking*, volume 5895 of *Lecture Notes in Computer Science*, pages 183–198. Springer Berlin Heidelberg, 2009.
50. K. Zellag and B. Kemme. How Consistent is Your Cloud Application? In *Proceedings of the 3rd Symposium on Cloud Computing (SOCC)*, SOCC '12, pages 6:1–6:14, New York, NY, USA, 2012. ACM.