

BigDataBench-MT: A Benchmark Tool for Generating Realistic Mixed Data Center Workloads

Rui Han¹(✉), Shulin Zhan³, Chenrong Shao⁴, Junwei Wang⁵,
Lizy K. John⁶, Jiangtao Xu¹, Gang Lu^{1,2}, and Lei Wang¹

¹ Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China
`{hanrui,xujiangtao,lugang,wl}@ict.ac.cn`

² University of Chinese Academy of Sciences, Beijing, China

³ ICarsclub, Beijing, China
`zhanshulin@ppzuche.com`

⁴ Xi'an Jiaotong University, Xi'an, China
`scr1994@stu.xjtu.edu.cn`

⁵ Kingsoft Cloud, Beijing, China
`wangjunwei@kingsoft.com`

⁶ Department of Electrical and Computer Engineering,
The University of Texas, Austin, TX, USA
`ljohn@ece.utexas.edu`

Abstract. Long-running service workloads (e.g. web search engine) and short-term data analysis workloads (e.g. Hadoop MapReduce jobs) co-locate in today's data centers. Developing realistic benchmarks to reflect such practical scenario of mixed workload is a key problem to produce trustworthy results when evaluating and comparing data center systems. This requires using actual workloads as well as guaranteeing their submissions to follow patterns hidden in real-world traces. However, existing benchmarks either generate actual workloads based on probability models, or replay real-world workload traces using basic I/O operations. To fill this gap, we propose a benchmark tool that is a first step towards generating a mix of actual service and data analysis workloads on the basis of real workload traces. Our tool includes a combiner that enables the replaying of actual workloads according to the workload traces, and a multi-tenant generator that flexibly scales the workloads up and down according to users' requirements. Based on this, our demo illustrates the workload customization and generation process using a visual interface. The proposed tool, called BigDataBench-MT, is a multi-tenant version of our comprehensive benchmark suite BigDataBench and it is publicly available from <http://prof.ict.ac.cn/BigDataBench/multi-tenancyversion/>.

Keywords: Data center · Benchmark · Workload trace · Mixed workloads

1 Introduction

In modern cloud data centers, a large number of tenants are consolidated to share a common computing infrastructure and execute a diverse mix of workloads. Benchmarking and understanding these workloads is a key problem for system designers, programmers and researchers to optimize the performance and energy efficiency of data center systems and to promote the development of data center technology. This work focuses on two classes of popular data center workloads [46]:

- *Long-running services.* These workloads offer online services such as web search engines and e-commerce sites to end users and the services usually keep running for months and years. The *tenants* of such workloads are *service end users*.
- *Short-term data analysis jobs.* These workloads process input data of many scales using relatively short periods (e.g. in Google and Facebook data centers, a majority (over 90 %) of analytic jobs complete within a few minutes [30, 46]). The *tenants* of such workloads are *job submitters*.

As data analysis systems such as Hadoop and Spark mature, both types of workloads widely co-locate in today’s data centers, hence the pressure to benchmark and understand these mixed workloads rises. Within this context, we believe that it will be of interest to the data management community and a large user base to generate realistic workloads such that trustworthy benchmarking reflecting the practical data center scenarios can be conducted. Considering the heterogeneity and dynamic nature of data center workloads and their aggregated resource demands and arrival patterns, this requires overcoming two major challenges.

Benchmarking Using Actual Workloads Based on Real-World Workload Races. Data analysis jobs usually have various computation semantics (i.e. implementation logics or source codes) and input data sizes (e.g. ranging from KB to PB), and their behaviors also heavily rely on the underlying software stacks (such as Hadoop or MPI). Hence it is difficult to emulate the behaviors of such highly diverse workloads just using synthetic workloads such as I/O operations. On the other hand, generating workloads whose arrival patterns follow real-world traces is an equally important aspect of realistic workloads. This is because these traces are the most realistic data sources including both explicit and implicit arrival patterns (e.g. sequences of time stamped requests or jobs).

Benchmarking Using Scalable Workloads with Realistic Mixes. A good benchmark needs to flexibly adjust the scale of workloads to meet the requirements of different benchmarking scenarios. Based on our experience, we noticed that in many cases, obtaining real workload traces is difficult due to confidential issues. The limited trace data also restrict the scalability of benchmark. It is therefore challenging to produce workloads at different scales while still guaranteeing their realistic mix corresponding to real-world scenarios.

In this paper, we propose a benchmark tool that is a first step towards generating realistic mixed data center workloads. This tool, called BigDataBench-MT,

is a multi-tenancy version of our open-source project BigDataBench, which is a comprehensive benchmark suite including 14 real-world data sets and 33 actual workloads covering five application domains [8, 51]. The goal of BigDataBench-MT is not only supporting the generation of service and data analysis workloads based on real workload traces, but also providing a multi-tenant framework to enable the scaling up and down of such workloads with guarantee of their realistic mixes. Considering our community may feel interest in using these workloads to evaluate new system designs and implementations, our tool and the corresponding workload traces are publicly available from <http://prof.ict.ac.cn/BigDataBench/multi-tenancyversion/>.

2 Related Work

We now review existing data center benchmarks from three perspectives, as shown in Table 1.

Evaluated Platform. First of all, we classify data center benchmarks according to their targeted systems. We consider three popular camps of systems in today’s data centers: (1) *Hadoop-related systems*: the great prosperity of the Hadoop-centric systems in industry brings a wide diversity of systems (e.g. Spark [6], HBase [1], Hive [2] and Impala [3]) on top of Hadoop MapReduce and HDFS [11] as well as a wide range of benchmarks specifically designed for these systems. (2) *Data stores*: parallel DBMSs (e.g. MySQL [12] and Oracle [15]) and NoSQL data stores (e.g. Amazon Dynamo [33], Cassandra [41] and LinkedIn Voldemort [50]) also widely exist in data centers. (3) *Web services*: long-running web services such as Nutch search engine [5] and multi-tier cloud applications [36] are another important type of data center applications. These services usually have stringent response time requirement [37] and their request processing is distributed into a large number of service components for parallel processing, thus the service latency is determined by the tail latency of these components [32, 38].

Workload Implementation Logic. Consider the complexity and diversity of workload behaviors in current data center systems, the implementation logic of existing data center benchmarks can be classified into three categories. The first category of benchmarks implement their workloads with algorithms. For example, HiBench [39] include workloads implemented with machine learning algorithms in Mahout [4]. The second category of benchmarks implement workloads using database operations such as reading, loading, joining, grouping, unifying, ordering, aggregating and splitting data. The third category of benchmarks implement workloads as I/O operations. For example, NNBench [13] and TestDFSIO [23] emulate I/O operations on Hadoop HDFS; GridMix [10] provides two workloads: LoadJob that performs I/O operations and SleepJob that sleeps the jobs; and SWIM [30] provides four workloads that stimulate the operations of Hadoop jobs to read, write, shuffle and sort data. We view the first two categories of workloads as *actual workloads*, because these workloads have semantics and they

consume resources of processors, memories, caches and I/O bandwidths in execution. By contrast, workloads belonging to third category only consume I/O resources.

Workload Mix. Finally, we classify data center benchmarks into three categories from the perspective of workload mix. The first type of data center benchmarks either generate single workloads (e.g. WordCount [26], Grep [9] and Sort [19]) or generate multiple workloads individually (e.g. CALDA [44], AMPLab benchmark [7] and CloudSuite [34]). That is, these benchmarks donot consider workload mix. The second category of benchmarks generate synthetic mixes of workloads. Many benchmarks (e.g. PigMix [17], HcBench [47] and BigBench [35]) generate mixes of workloads by manually determining their proportions. Similarly, TPC benchmarks [24] design a query set as a synthetic mix of queries with different proportions. YCSB [31] uses a package to include a set of related workloads. MRBS decides the frequencies of different workloads using probability distributions such as a random distribution. Finally, third category of benchmarks generate a realistic mix of synthetic workloads whose arrival patterns faithfully follow real-world traces. For example, GridMix [10] and SWIM [21, 30] first build a job trace to describe the realistic job mix by mining production loads, and then run synthetic I/O operations according to the trace. However, how to generate actual workloads on the basis of real workload traces is still an open question.

Table 1. Overview of data center benchmarks

Workload mix	Workload implementation logic		
	Algorithms	Database operations	I/O operations
No mix	WordCount ¹ [26], Grep ¹ [9], Sort ¹ [19], Terasort ¹ [22], HiBench ¹ [39], TPCx-HS ¹ [43], Graphalytics ¹ [29], CloudSuite ⁴ [34]	MRBench ¹ [40], CALDA ² [44], AMPLab benchmark ² [7], YCSB ² [31], BG benchmark ² [28], CloudSuite ⁴ [34]	NNBench ¹ [13], TestDFSIO ¹ [23], HiBD ¹ [49]
Synthetic mix	HcBench ¹ [47], MRBS ¹ [48]	PigMix ¹ [17], HcBench ¹ [47], MRBS ¹ [48], BigBench ² [35], LinkBench ² [27], TPC benchmarks ² [24], TPC-W ³ [25], BigDataBench ⁴ [51]	HiBench ¹ [39], SPECWeb99 ³ [20]
Realistic mix			Gridmix ¹ [10], SWIM ¹ [21, 30]

¹ Hadoop-related systems

² Data stores

³ Web services

⁴ All three types of systems

3 System Overview

Figure 1 shows the framework of our benchmark tool. It consists of three main modules. In the *Benchmark User Portal*, users can first specify their benchmarking requirements, including the machine type and number to be tested, and the types of workload to use. A set of workload traces following these requirements are then selected. The next step of *Combiner of Workloads and Traces* is to match the real workload and the selected workload traces, and outputs workload replaying scripts to guide the workload generation. Finally, the *Multi-tenant Workload Generator* extracts the tenant information from the scripts and constructs a multi-tenant framework to generate a mix of service and data analysis workloads.

In BigDataBench-MT, we employ the Sogou user query logs [18] as the basis to generate the service workload (i.e. the Nutch search engine [5]) and the Google cluster workload trace as the basis to generate data analysis workloads (i.e. Hadoop and Shark workloads). The Sogou trace records logs from 50 days and it includes over 9 million users and 43 million queries. The Google trace records logs from 29 days and 12,492 machines and it includes over 5 K users, 40 K workload types, 1000 K jobs and 144 million tasks. As a preprocessing step, we converted both traces into Impala databases (full version) and MySQL database (24-hour version) to facilitate the customization of benchmarking scenarios. In the following subsections, we describe the last two modules of our tool.

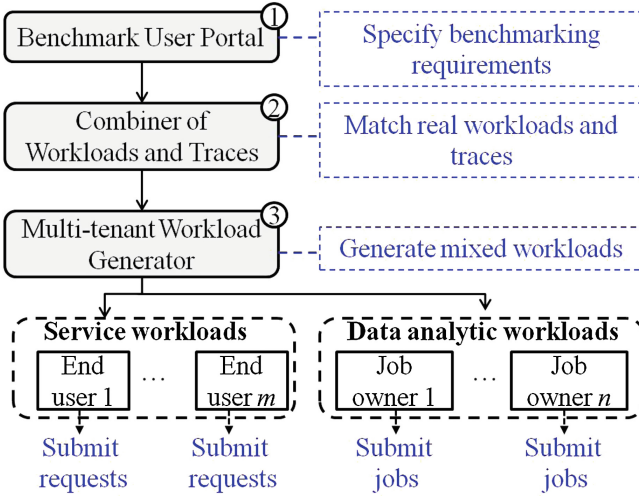


Fig. 1. The BigDataBench-MT framework

3.1 Combiner of Workload and Traces

The goal of the combiner is to extract the request/job arrival patterns from real-world traces and combine them with actual workloads. The combiner applies differentiated combination techniques to the service and data analysis workloads because their workload generations have different features.

Service Workloads. The generation of a service workload is determined by three factors: the request submitting time, the sequence of requests and the content of each request query. Take the web search engine for example, the combiner implements a request submitting service that automatically derives these factors from the Sogou trace and uses them to determine the request submission process.

Data Analysis Workloads. The generation of a data analysis workload is determined by four factors: the job submitting time, the workload type (i.e. the computation semantics and the software stack) and the input data (i.e. the data source and size). The current workload traces usually show the information of job submitting time but only provide *anonymous jobs* whose workload types and/or input data are unknown. Hence the basic idea of the combiner is to derive the workload characteristics of both actual and anonymous jobs and then match jobs whose workload characteristics are sufficiently similar. Table 2 lists the metrics used to represent workload characteristics, which reflect both jobs' performance (execution time and resource usage) and micro-architectural behaviors (CPI and MAI).

Table 2. Metrics to represent workload characteristics of data analysis jobs

Metric	Description
Execution time	Measured in seconds
CPU usage	Total CPU time per second
Total memory size	Measured in GB
CPI	Cycles per instruction
MAI	The number of memory accesses per instruction

Figure 2 shows the process of matching actual data analysis jobs and traces' anonymous jobs and it consists of two parallel sub-processes. First, the actual jobs with different input data sizes are tested and their metrics of workload characteristics are collected. In BigDataBench-MT, we provide auto-running scripts to collect performance metrics and hardware performance counters (Perf [16] and Oprofile [14] for Linux 2.6+ based systems) to obtain micro-architectural metrics. Using the testing results as samples, the combiner trains the multivariate regression model to describe the relationship between an actual job (including both its workload type and input size as the independent variables) and its workload characteristic metrics (one metric is a dependent variable). Second, the

combiner views each anonymous job as an entity and the five workload characteristic metrics as its attributes, and employs the Bayesian Information Criterion (BIC)-based k-means clustering algorithm [45] to group anonymous jobs in the trace into different clusters.

Based on the constructed regression models and clusters, the combiner further matches each cluster to one actual job with a specific input data. In the matching, the coefficient of variation (CV) measure, defined as the ratio of the standard deviation σ to the mean μ , is used to describe the dispersion of jobs in the same cluster. The combiner iteratively tests actual jobs of different workload types and input sizes, and matches an actual job with a cluster under two conditions: (i) the CV of the cluster is smaller than a specified threshold (e.g. 0.5), which indicates the anonymous jobs in this cluster are closely similar to each other; (ii) the change in this CV is smaller than a threshold (e.g. 0.1) after the actual job is added to the cluster. This means the workload characteristics of the added job are sufficiently similar to those of the anonymous jobs in the cluster. If multiple matched actual jobs are found for one cluster, the combiner selects the job resulting the smallest CV change. Finally, the combiner produces workload replaying scripts as the output.

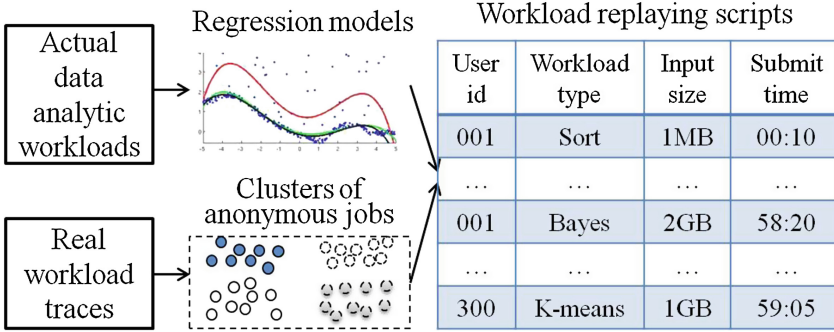


Fig. 2. The matching process of real and synthetic data analysis jobs

Note that BigDataBench-MT provides two ways of using the above combiner. First, it directly provides some workload replaying scripts, which are the combination results of representative actual workloads (e.g. Hadoop Sort and WordCount) and the Google workload trace. Second, it also supports benchmark users to directly use the above combination technique to match their own data analysis jobs with Google anonymous jobs.

3.2 Multi-tenant Workload Generator

Based on workload replaying scripts, the workload generator applies a multi-tenant mechanism to generate a mix of workloads using two steps. First, the generator extracts the tenant information from the scripts. For the service and

data analysis workloads, this tenant information represents the number of concurrent end users and submitters of analytic jobs, respectively. Second, the generator creates a client for each tenant and emulates the scenarios that a number of end users/job submitters concurrently submit requests/jobs to the system. This multi-tenant framework allows the flexible adjustment of workload scales with guarantee of their realistic mixes. For example, benchmark users can double or halve the size of concurrent tenants, after which the distributions of requests/jobs submitted by these tenants still correspond to those in real workload traces.

4 Demonstration Description

4.1 Chosen Workloads and Workload Traces

In our demonstration, benchmark users want to evaluate their data center systems using a mix of service and data analysis workloads. The Nutch web search engine [5] is used as the example service workload and four Hadoop workloads are used as the example data analysis workloads. The chosen Hadoop workloads have a variety of workload characteristics: WordCount and Naïve Bayes classification are typical CPU-intensive workloads with integer and float point calculations; Sort is the typical I/O-intensive workload and PageIndex is the workload having similar demands for CPU and I/O resources. Both the data generators [42] and workloads in the demo can be obtained from BigDataBench [8].

Our demo uses a 24-h user query logs from Sogou, which include 1,724,264 queries from 519,876 end users, as the basis to generate realistic search engine service; and uses a 24-h cluster workload trace from Google, which includes 37,842 anonymous jobs from 2,261 job submitters, as the basis to generate realistic Hadoop jobs.

4.2 System Demonstration

BigDataBench-MT provides a visual interface in the *Benchmark User Portal* to help benchmark users make appropriate benchmarking decisions. This portal provides users necessary information, allows them input benchmarking requirements and executes system evaluations on their behalf. The whole process consists of three steps, as shown in Figs. 3, 4 and 5, respectively.

Step 1. Specification of tested machines and workloads. The first step of the demo presents an overview of workload traces (i.e. Sogou and Google traces) and the data center status, including the six types of machines, their machine number and configurations, and the user, job and task statistics in these machines. This information assists benchmark users to select the type and number of machines to be evaluated, and the workloads they want to use. Suppose users select *Type Four* of the machines with 2 process cores and 4 GB memory and 100 machines to be tested, the workload traces belonging to these machines are extracted and forwarded to the next step.



Fig. 3. System demonstration screenshots: step 1

Step 2. Selection of benchmarking period and scale. At this step, users have the option to select the period and scale for their benchmarking scenarios. To facilitate this selection, BigDataBench-MT shows the statistic information of both the service workload (including its number of requests and end users per second) and the data analysis workloads (including their number of jobs and average CPU, memory resource usages) at each of the 24 hs. Suppose users select the benchmarking period of 12:00 to 13:00 and the scale factor is 1 (that is, no scaling is needed). The workload traces belonging to this period are selected for step 3.



Fig. 4. System demonstration screenshots: Step 2

Step 3. Generation of mixed workloads. After both workloads and traces have been selected, the final step employs the combiner described in Sect. 3.1 to generate workload replaying scripts for both the service and data analysis workloads, and sends these scripts as feedback to users. In the matching of actual Hadoop jobs with anonymous ones, we tested each Hadoop workload type using 20 different input sizes to build the regression models. Based on the replaying scripts,

Time	Workload ID	Description
12:00:00	SW00003	合成
12:00:00	SW00004	哈尔滨大学
12:00:00	SW00005	如何安装qq摄像头
12:00:00	SW00006	音乐播放软件
12:00:00	SW00007	倪萍的几次婚姻
12:00:00	SW00008	安全义务下的侵权责任
12:00:00	SW00009	本田飞度最新价格
12:00:00	SW00010	长春房地产报

Submit time	User id	Workload type	Input Size(MB)
12:00:49	DAW07	Bayes	70
12:01:41	DAW04	Wordcount	69
12:02:02	DAW06	Bayes	25
12:03:17	DAW01	Bayes	22
12:03:28	DAW08	Sort	342
12:05:49	DAW07	Index	22
12:06:13	DAW00	Sort	236

Fig. 5. System demonstration screenshots: step 3

benchmark users can press the “Generate mixed workload” button to trigger the multi-tenant workload generator, in which each tenant is an independent workload generator and multiple tenants generate a mix of realistic workloads.

5 Future Work

There are multiple avenues for extending the functionality of our benchmark tool. A first step will be to support more actual workloads. Given that there are 33 actual workloads in the BigDataBench and many workloads (e.g. Bayes classification and WordCount) have three versions of implementations (Hadoop, Spark and MPI), adding more workloads to BigDataBench-MT will be helpful to support wider benchmarking scenarios. We also plan to extend our multi-tenant workload generator to support different classes of tenants and allow users to apply different priority disciplines in workload generation.

Acknowledgements. This work is supported by the National High Technology Research and Development Program of China (Grant No. 2015AA015308), the National Natural Science Foundation of China (Grant No. 61502451), and the Key Technology Research and Development Programs of Guangdong Province, China (Grant No. 2015B010108006).

References

1. Apache hbase. <http://hbase.apache.org/>
2. Apache hive. <https://cwiki.apache.org/confluence/display/Hive/Home>
3. Apache impala. <http://impala.io/>
4. Apache Mahout. <http://mahout.apache.org/>
5. Apache Nutch. <http://nutch.apache.org/>
6. Apache spark. <https://spark.apache.org/>
7. Big data benchmark by amplab of UC berkeley. <https://amplab.cs.berkeley.edu/benchmark/>
8. BigDataBench. <http://prof.ict.ac.cn/BigDataBench/>
9. Grep. <http://wiki.apache.org/hadoop/Grep>
10. Gridmix. <https://hadoop.apache.org/docs/r1.2.1/gridmix.html>

11. Hadoop ecosystems. <https://hadoopecosystemtable.github.io/>
12. MySQL database. <https://www.mysql.com/>
13. Nnbench. <http://grepcode.com/file/repo1.maven.org/maven2/org.apache.hadoop/hadoop-mapred-test/0.22.0/org/apache/hadoop/hdfs/NNBench.java/>
14. Oprofile. <http://oprofile.sourceforge.net/>
15. Oracle database. <http://www.oracle.com/>
16. Perf. <https://perf.wiki.kernel.org/>
17. PigMix. <https://cwiki.apache.org/confluence/display/PIG/PigMix>
18. Sogou user query logs. <http://www.sogou.com/labs/dl/q-e.html>
19. Sort. <http://wiki.apache.org/hadoop/Sort>
20. Specweb99 benchmark. <https://www.spec.org/web2009/>
21. Swim. <https://github.com/SWIMProjectUCB/SWIM/wiki>
22. Terasort. <https://hadoop.apache.org/docs/current/api/org/apache/hadoop/examples/terasort/package-summary.html>
23. Testdfsio. <https://support.pivotal.io/hc/en-us/articles/200864057-Running-DFSIO-mapreduce-benchmark-test/>
24. TPC benchmarks. <http://www.tpc.org/>
25. TPC-W benchmark. <http://www.tpc.org/tpcw/>
26. WordCount. <http://wiki.apache.org/hadoop/WordCount>
27. Armstrong, T.G., Ponnemanti, V., Borthakur, D., Callaghan, M., Linkbench: a database benchmark based on the facebook social graph. In: SIGMOD 2013, pp. 1185–1196. ACM (2013)
28. Barahmand, S., Ghandeharizadeh, S., BG: A benchmark to evaluate interactive social networking actions. In: CIDR 2013. Citeseer (2013)
29. Capotă, M., Hegeman, T., Iosup, A., Prat-Pérez, A., Erling, O., Boncz, P., Graphalytics: A big data benchmark for graph-processing platforms. In: Proceedings of the GRADES 2015, pp. 7. ACM (2015)
30. Chen, Y., Alspaugh, S., Katz, R.: Interactive analytical processing in big data systems: A cross-industry study of mapreduce workloads. VLDB **5**(12), 1802–1813 (2012)
31. Cooper, B.F., Silberstein, A., Tam, E., Ramakrishnan, R., Sears, R.: Benchmarking cloud serving systems with YCSB. In: SoCC 2010, pp. 143–154. ACM (2010)
32. Dean, J., Barroso, L.A.: The tail at scale. Commun. ACM **56**(2), 74–80 (2013)
33. DeCandia, G., Hastorun, D., Jampani, M., Kakulapati, G., Lakshman, A., Pilchin, A., Sivasubramanian, S., Voshall, P., Vogels, W., Dynamo: amazon’s highly available key-value store. In: ACM SIGOPS Operating Systems Review, vol. 41, pp. 205–220. ACM (2007)
34. Ferdman, M., Adileh, A., Kocberber, O., Volos, S., Alisafae, M., Jevdjic, D., Kaynak, C., Popescu, A.D., Ailamaki, A., Falsafi, B.: Clearing the clouds: A study of emerging workloads on modern hardware. Technical report (2011)
35. Ghazal, A., Rabl, T., Hu, M., Raab, F., Poess, M., Crotte, A., Jacobsen, H.-A.: Bigbench: Towards an industry standard benchmark for big data analytics. In: SIGMOD 2013, pp. 1197–1208. ACM (2013)
36. Han, R., Ghanem, M.M., Guo, L., Guo, Y., Osmond, M.: Enabling cost-aware and adaptive elasticity of multi-tier cloud applications. Future Gener. Comput. Syst. **32**, 82–98 (2014)
37. Han, R., Wang, J., Ge, F., Vazquez-Poletti, J.L., Zhan, J.: SARP: producing approximate results with small correctness losses for cloud interactive services. In: CF 2015, pp. 22. ACM (2015)

38. Han, R., Wang, J., Huang, S., Shao, C., Zhan, S., Zhan, J., Vazquez-Poletti, J.L., SARP: producing approximate results with small correctness losses for cloud interactive services. In: ICPP 2015, pp. 490–499. IEEE (2015)
39. Huang, S., Huang, J., Dai, J., Xie, T., Huang, B., The hibenx benchmark suite: Characterization of the mapreduce-based data analysis. In: IEEE 26th International Conference on Data Engineering Workshops (ICDEW 2013), pp. 41–51. IEEE (2010)
40. Kim, K., Jeon, K., Han, H., Kim, S.-G., Jung, H., Yeom, H.Y.: Mrbench: A benchmark for mapreduce framework. In: ICPADS 2008, pp. 11–18. IEEE (2008)
41. Lakshman, A., Malik, P.: Cassandra: a decentralized structured storage system. *ACM SIGOPS Oper. Syst. Rev.* **44**(2), 35–40 (2010)
42. Ming, Z., Luo, C., Gao, W., Han, R., Yang, Q., Wang, L., Zhan, J.: BDGS: A scalable big data generator suite in big data benchmarking. In: Rabl, T., Raghunath, N., Poess, M., Bhandarkar, M., Jacobsen, H.A., Baru, C. (eds.) *Advancing Big Data Benchmarks*. LNCS, vol. 8585, pp. 138–154. Springer, Heidelberg (2014)
43. Nambiar, R.: A standard for benchmarking big data systems. In: BigData 2014, pp. 18–20. IEEE (2014)
44. Pavlo, A., Paulson, E., Rasin, A., Abadi, D.J., DeWitt, D.J., Madden, S., Stonebraker, M.: A comparison of approaches to large-scale data analysis. In: SIGMOD 2009, pp. 165–178. ACM (2009)
45. Pelleg, D., Moore, A.W., X-means, et al.: Extending k-means with efficient estimation of the number of clusters. In: ICML 2000, pp. 727–734 (2000)
46. Reiss, C., Tumanov, A., Ganger, G.R., Katz, R.H., Kozuch, M.A.: Heterogeneity, dynamics of clouds at scale: Google trace analysis. In: *Proceedings of the Third ACM Symposium on Cloud Computing*, pp. 7. ACM (2012)
47. Saletore, V., Krishnan, K., Viswanathan, V., Tolentino, M.E., Hcbench, et al.: Methodology, development, and characterization of a customer usage representative big data/hadoop benchmark. In: IISWC 2013, pp. 77–86. IEEE (2013)
48. Sangroya, A., Serrano, D., Bouchenak, S.: MRBS: towards dependability benchmarking for hadoop mapreduce. In: Caragiannis, I., Alexander, M., Badia, R.M., Cannataro, M., Costan, A., Danelutto, M., Desprez, F., Krammer, B., Sahuquillo, J., Scott, S.L., Weidendorfer, J. (eds.) *Euro-Par Workshops 2012*. LNCS, vol. 7640, pp. 3–12. Springer, Heidelberg (2013)
49. Shankar, D., Lu, X., Wasi-ur-Rahman, M., Islam, N., Panda, D.K.D.K.: A Micro-benchmark suite for evaluating hadoop mapreduce on high-performance networks. In: Zhan, J., Rui, H., Weng, C. (eds.) *BPOE 2014*. LNCS, vol. 8807, pp. 19–33. Springer, Heidelberg (2014)
50. Sumbaly, R., Kreps, J., Gao, L., Feinberg, A., Soman, C., Shah, S.: Serving large-scale batch computed data with project voldemort. In: FAST 2012, pp. 18–18. USENIX Association (2012)
51. Wang, L., Zhan, J., Luo, C., Zhu, Y., Yang, Q., He, Y., Gao, W., Jia, Z., Shi, Y., Zhang, S., et al.: Bigdatabench: a big data benchmark suite from internet services. In: HPCA 2014, pp. 488–499. IEEE (2014)

Big Data Benchmarks, Performance Optimization, and
Emerging Hardware

6th Workshop, BPOE 2015, Kohala, HI, USA, August 31 -
September 4, 2015. Revised Selected Papers

Zhan, J.; Han, R.; Zicari, R.V. (Eds.)

2016, X, 147 p. 53 illus. in color., Softcover

ISBN: 978-3-319-29005-8