

Does Big Data Require Big Systems? A Study of Complex Versus Lean Distributed Data Processing Systems

Duarte M. Nascimento

INESC-ID, Instituto Superior Técnico, Universidade de Lisboa
Lisbon, Portugal
duarte.miguel@tecnico.ulisboa.pt

ABSTRACT

Spark is a widely used big data processing system that can be integrated with Hadoop to provide powerful abstractions to developers, such as distributed storage through HDFS and resource management through YARN. When all the required configurations are made, Spark can also provide quality attributes, such as scalability, fault tolerance, and security. However, all of these benefits come at the cost of complexity, with high memory requirements, and additional latency in processing. An alternative approach is to use a lean software stack, like Unicage, that delegates most control back to the developer.

In this work we benchmarked the big data processing performance of Spark versus Unicage, in a cluster environment hosted in the IBM Cloud. Two sets of experiments were performed: batch processing of unstructured data sets, and query processing of structured data sets. The input data sets were of significant size, ranging from 64 GB to 8192 GB in volume. The results show that the performance of Unicage shell scripts is superior to Spark for search workloads like *grep* and *select*, but that the abstractions of distributed storage and resource management from the Hadoop stack enable Spark to execute workloads with inter-record dependencies, such as *sort* and *join*, with correct outputs.

1 INTRODUCTION

The paradigm of big data has been growing with the rise of the Internet of Things (IoT). Sensors and actuators now join the already established sources of data, to produce even larger quantities, at faster speeds. Juniper Research estimates an increase of 137% in the number of IoT connections from 2020 to 2024 [22]. Statista estimates an increase of 129% in the amount of generated data from 2020 to 2024 [14].

Big data is characterized by a set of *V properties*: *volume*, *velocity* and *variety* of sources and formats characterize the dimensions of data, while *veracity* and *value* characterize its quality. Common data becomes big data once handling these attributes transcends the capabilities of a single machine or of a small cluster [4]. The cloud is a resourceful environment for the deployment of big data systems and the processing of workloads in production. Big data systems seek to handle the dimensions of data sets while preserving their quality. *Hadoop MapReduce* was the popular big data system of yore, but as the paradigm evolved, so did the necessity for more capable systems, like *Spark*¹. Most recent big data systems offer a plethora of abstractions to allow the developer to focus strictly on the data and how to deal with it. These abstractions include distributed storage (e.g. HDFS in the Hadoop software stack) and distributed

resource management (e.g. YARN in the Hadoop software stack). Additionally, the systems seek to optimize a variety of quality attributes including performance, scalability, fault tolerance, security, usability and manageability. As a result, the systems harboring these responsibilities have inevitably complex software stacks and are hereafter addressed as *complex big data systems*. The complexity of these software stacks affects the performance of the big data processing workloads, with additional latency. Furthermore, this complexity hinders the deployment in low-powered devices with low hardware specifications, such as IoT devices.

As an alternative to the shortcomings of complex big data systems, *lean big data systems* are characterized by an absence of many abstraction layers, and put the responsibility to implement the needed mechanisms mostly on the developer. This way, these solutions can be tailored to the needs of the use case and optimized only for the required set of quality attributes. Unicage² enables the implementation of lean big data systems through shell scripting.

Figure 1 illustrates the software stack difference between complex and lean big data systems.

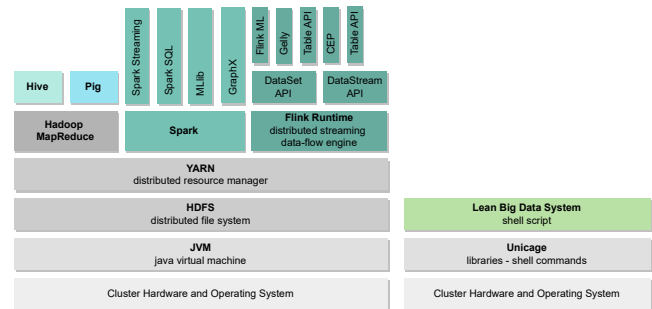


Figure 1: Software stacks of complex big data systems provided by the Apache Software Foundation (left); and of a lean big data system implemented with Unicage (right).

This study presents the first benchmark that compares complex and lean big data systems directly in a cloud environment, and is focused strictly on the performance evaluation of data loading, batch processing and query processing workloads. Spark was benchmarked as the main representative of complex big data systems, as it is widely used and has been, in recent years, one of the most popular projects hosted by the Apache Software Foundation [23, 24]. Hadoop MapReduce and Hive (over MapReduce) are not as widely used nowadays, but were also benchmarked for historical reference, and as tiebreakers for the verification of results. Unicage was benchmarked as the representative of lean big data systems.

¹<https://spark.apache.org/>

²<https://unicage.eu/>

2 BACKGROUND

This section provides a brief description of the big data systems benchmarked in this study.

2.1 Complex Big Data Systems

The Hadoop stack runs on top of the Java Virtual Machine (JVM), featuring HDFS [25] and YARN [29] as the underlying distributed file system and distributed resource manager, respectively. Hadoop MapReduce [5, 6] is a big data system for batch processing. Hive [26, 27] originally enabled query processing on top of MapReduce, but has been extended to use Spark as the underlying processing engine instead. Spark [32] and Flink [3] are multipurpose big data systems that can use the distributed storage and resource management capabilities of the Hadoop stack.

2.1.1 MapReduce. MapReduce is a programming model designed for batch processing of large volumes of data, represented as key/value pairs. The typical MapReduce application is comprised of two user-defined operations: the *Mapper* operation that processes the input key/value pairs and produces intermediate key/value pairs; and the *Reducer* operation that processes the intermediate key/value pairs output by the Mapper and produces output key/value pairs. A *Split* operation often precedes the Map operation to split the input data set in records of a user-defined size, to be processed by individual Mappers. A *Shuffle* operation often succeeds the Map operation to group the intermediate key/value pairs output by the Mappers in logic key/value pairs, sorted by key, to be processed by individual Reducers. The Split and Shuffle operations allow Mappers and Reducers to be parallelized.

2.1.2 HDFS. The Hadoop Distributed File System enables distributed storage across a cluster with multiple nodes, presenting it as a single global namespace to both the user and applications. HDFS provides other quality attributes to the Hadoop software stack, such as fault tolerance through self-managed file block replication, and scalability through a leader/follower architecture. The architecture of HDFS provides these capabilities through one or more leader nodes called *NameNodes*, where file metadata is persisted; and follower nodes called *DataNodes*, where the raw data of files is persisted and replicated.

One of the shortcomings of HDFS is the latency associated with file block management and lookup. To minimize this, file block size is set to 128 MB, and the maximum number of blocks per file is set to 10000, both by default. When there is a large number of files whose size is significantly smaller than the set file block size, causing a large number of very small file blocks to exist within the file system, the performance of HDFS is known to deteriorate [7].

2.1.3 YARN. “Yet Another Resource Negotiator” is the distributed resource manager of the Hadoop software stack, providing load balancing and cluster scalability through the dynamic negotiation of resources to each Hadoop application running in the hosting cluster. YARN enables *multi-tenancy*, allowing multiple logically separated Hadoop applications to share physical resources in a shared environment, each translating to an isolated YARN application. Managed resources include allocated memory, CPU, and disk space. Their negotiation is handled by: a *ResourceManager* service that runs in the NameNode and is responsible for the resource allocation in the

form of bundled application containers; one *NodeManager* service running in each DataNode, responsible for updating the Resource-Manager with the status of the corresponding DataNode; and one *ApplicationMaster* service running per-application, responsible for the negotiation of resources with the ResourceManager and the execution, management and monitoring of the YARN application containers.

2.1.4 Hive. Hive is a data warehousing solution that enables query processing of structured data in Hadoop. Programming in MapReduce is notoriously difficult, so Hive offers an SQL-based querying language called *HiveQL*, whose statements are converted to MapReduce jobs. A structured data set to be processed with HiveQL queries needs to be stored in HDFS, and imported to the Hive database in tabular format, along with its schema. This does not imply copying the data set over to the database, as only the schema and information about the tables is persisted in a metadata repository called *Metastore*. Modern implementations of Hive use Spark as the underlying processing engine, instead of MapReduce.

2.1.5 Spark. Oftentimes, Hadoop applications perform iterative algorithms that require MapReduce jobs to load data from disk repeatedly, which is an inefficient use of resources. Spark abandons the MapReduce paradigm in favor of partitioning data sets into specialized immutable data structures called *resilient distributed data sets* (RDDs) [31]. Jobs are compiled into direct acyclic graphs of operations to be performed on RDDs. RDDs are partitioned across the nodes of a Spark cluster and their data lineage information is used to rebuild them if nodes fail or partitions are lost. RDD partitions are stored in memory and loaded on demand, providing a significant performance increase when compared with the MapReduce approach. Caching may result in memory becoming full, to which Spark reacts by ceasing to cache RDD partitions and instead recomputing them on demand, allowing the job to still proceed but with reduced performance. Alternatively, the caching level can also be configured so RDD partitions that do not fit in memory are stored and read from disk instead.

Spark has a leader/follower architecture that allows its integration with HDFS and YARN. Every Spark application runs in the context of a *SparkContext*. A typical Spark cluster features two types of processes: one *Driver* process, responsible for the creation, execution and management of the SparkContext, the scheduling of the direct acyclic graphs and planning of tasks, and the communication with YARN for resource negotiation; and one or more *Executor* processes, responsible for the execution of the tasks planned by the Driver.

Cluster mode enables the Driver process to execute in one of the worker nodes. If Spark is deployed with YARN, the NameNode runs the ResourceManager service, the Driver process runs in one of the DataNodes along with the ApplicationMaster service, and the Executors also run in the DataNodes as YARN application containers.

The architecture of Spark, combined with the immutable nature of RDDs allowed it to expand beyond its original purpose of batch processing to feature various domain-specific libraries, including *Spark SQL* [2] for query processing; *Spark Streaming* [33] and *Spark Structured Streaming* [1] for stream processing; *GraphX* [12] for graph processing; and *MLlib* [15] for machine learning.

2.2 Lean Big Data Systems

Lean big data systems delegate most responsibilities to the developer. Unicage embraces the UNIX fundamentals and offers a set of proprietary commands written in the C language, that can be used alongside typical shell commands to implement data processing pipelines within shell scripts. The C language uses less memory and consumes less energy than other languages [19], and therefore saves energy in the servers and is more suitable for deployment in low-powered devices with low hardware specifications. Unicage features a base library for single-machine data processing called *Tukubai*, and an extension of this library with additional commands for parallel and distributed data processing called *BOA*.

2.2.1 Unicage Tukubai. Tukubai features 200 commands designed for single-machine data processing. Table 1 describes some of the most commonly used commands in the library.

Command	Description
<code>msort</code>	Performs in-memory sort of records using the merge-sort algorithm.
<code>lcnt</code>	Counts the number of records in a file.
<code>dmerge</code>	Merges two sorted files based on specified columns.
<code>self</code>	Selects and reorders the columns in a file.
<code>sm2</code>	Aggregates records grouped by specified columns.

Table 1: Examples of Unicage Tukubai commands.

Ferreira et al. [8] demonstrated a use case with Tukubai, applied to processing data generated from a smart grid.

2.2.2 Unicage BOA. “BigData Oriented Architecture” features parallel and distributed counterparts for the commonly used commands in the Tukubai library. The typical BOA cluster has a leader/follower architecture, and the followers that partake in the computation are specified per-command. BOA features three types of commands: the `para-*` commands, aimed at parallel data processing in a single machine; the `clust-*` commands, aimed at parallel and distributed data processing in a cluster of machines, and are typically issued from the leader node; and the `distr-*` commands, similar to the `clust-*` commands, but the leader node does not take part in the computation. Table 2 lists some commonly used commands in the BOA library.

Tukubai base command	BOA commands		
	<code>para-*</code>	<code>clust-*</code>	<code>distr-*</code>
<code>msort</code>	n/a	<code>clust-msort</code>	<code>distr-msort</code>
<code>dmerge</code>	n/a	n/a	<code>distr-dmerge</code>
<code>self</code>	<code>para-self</code>	<code>clust-self</code>	<code>distr-self</code>
<code>sm2</code>	<code>para-sm2</code>	<code>clust-sm2</code>	<code>distr-sm2</code>

Table 2: Examples of Unicage BOA commands.

3 RELATED WORK

The study discussed in this paper borrowed from previous studies in the field of big data benchmarking, all of which are strictly focused on benchmarking complex big data systems and do not consider lean big data systems, with the exception of *LeanBench*. This section highlights two of the more well-known big data benchmark suites, as well as *LeanBench*, the only benchmark which has used Unicage.

Big Data Benchmark Suites follow and expand on the standards defined by the Transaction Processing Performance Council (TPC) and the Standard Performance Evaluation Corporation (SPEC) benchmarking consortia. The scope of big data benchmarks typically fits one of three categories [13]: *micro-benchmarks* are designed to evaluate individual system components through small operations; *end-to-end benchmarks* are designed to evaluate whole systems through typical application scenarios; *benchmark suites* are combinations of micro and end-to-end benchmarks targeted at various big data systems.

3.1 BigDataBench

BigDataBench [9, 30] is a big data and artificial intelligence benchmark suite. It features micro-benchmarks for workloads comprised of single small units of computation, component benchmarks for workloads comprised of sets of small units of computation, and end-to-end benchmarks that mimic realistic application scenarios with combinations of component benchmarks. BigDataBench provides a suite of implementations of batch, query and stream processing workloads, among others, for the state-of-the-art big data systems including Spark. BDGS [16], which stands for Big Data Generation Suite, is the data generator of BigDataBench.

3.2 BigBench

BigBench [11] is an end-to-end benchmark targetted at data warehousing and query processing frameworks and mimics the use-case of a product retailer. It was first introduced in 2013 but has since been adapted into an industry standard by TPC as the TPCx-BB benchmark. PDGF [20], which stands for Parallel Data Generation Framework, is suitable for the generation of structured data. BigBench extends PDGF to generate semi-structured and unstructured data. In 2017, BigBench V2 [10] was proposed as a rework of the original data model and data generator to better reflect real-life scenarios.

3.3 LeanBench

Moreira et al. [18] benchmarked the performance of the *grep*, *sort*, *wordcount*, *select*, *join* and *aggregation* workloads implemented in Unicage Tukubai, Hadoop and Hive, deployed in a single machine. Figure 2 highlights the conclusions of LeanBench: in a single-machine environment, Unicage performs better than Hadoop if the task does not involve sorting large data sets (over 400 MB), and performs better than Hive for data sets of all volumes, with the latter being the preferred system only if using a querying language is a requirement.

The main limitation of LeanBench was the deployment in a single machine which prevented tests with larger data sets.

4 EXPERIMENTAL SET-UP

This section provides a detailed view of the experimental set-up, followed by an overview of the processes involved.

4.1 Environment

The experimental environment was divided in three sub-clusters: the *Hadoop cluster* for data processing with complex big data systems; the *Unicage cluster* for data processing with lean big data

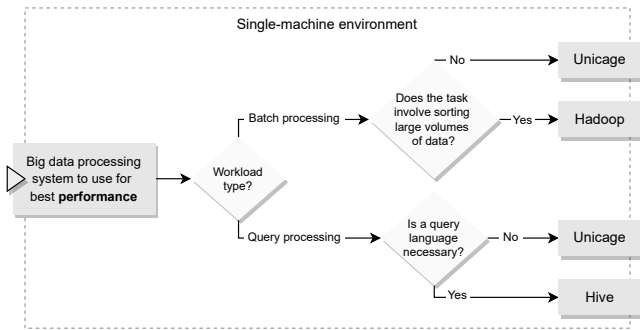


Figure 2: Summary of conclusions from the LeanBench study.

systems; and the *Producer cluster* for data generation and loading into the Hadoop and Unicage clusters. The Hadoop and Unicage clusters had the same specifications, as shown in Table 3, which lists the specifications of every node, hosted as a VSI (Virtual Server Instance) in the IBM Cloud VPC infrastructure.

Hostname	Software			Hardware				
	Hadoop	Hive	Spark	Unicage	CPU	RAM	Network bandwidth	Disk
Hadoop cluster								
<i>namenode</i>	3.3.1	3.1.2	3.2.1	n/a	8 vCPU, 2.0 GHz	32 GB	12 Gbps	1024 GB, 412 Mbps
<i>datanode1...5</i>	3.3.1	n/a	n/a	n/a	4 vCPU, 2.0 GHz	16 GB	6 Gbps	4096 GB, 1607 Mbps
Unicage cluster								
<i>unicageleader</i>	n/a	n/a	n/a	TKB, BOA	8 vCPU, 2.0 GHz	32 GB	12 Gbps	1024 GB, 412 Mbps
<i>unicageworker1...5</i>	n/a	n/a	n/a	TKB, BOA	4 vCPU, 2.0 GHz	16 GB	6 Gbps	4096 GB, 1607 Mbps
Producer cluster								
<i>producer1...4</i>	3.3.1	n/a	n/a	TKB, BOA	8 vCPU, 2.0 GHz	32 GB	12 Gbps	4096 GB, 1607 Mbps

Table 3: Specifications of the experimental cluster environment.

The nodes in the Producer cluster were configured as HDFS clients to allow data loading into the Hadoop cluster; and the relevant Unicage commands were also installed to allow data loading into the Unicage cluster.

The benchmarks did not contemplate fault tolerance, so replication was turned off in the Hadoop cluster, by setting the `dfs.replication` property to 0 in the configuration file of HDFS. Some of the largest input data set volumes used in the benchmarks exceeded the maximum number of file blocks per file allowed in HDFS, of 10000, by a small margin, so to load these files into the Hadoop cluster, this limit was lifted by setting the `dfs.namenode.fs-limits.max-blocks-per-file` property to a sufficiently large number in the configuration file of HDFS. Spark was configured to take advantage of the four CPU cores of the five DataNodes in the Hadoop cluster, by setting the `spark.executor.instances` and `spark.executor.cores` configuration properties to 5 and 4, respectively. The caching level of Spark was left with the default configuration: when it runs out of memory, it uses the lineage information of RDDs to recompute them on demand, allowing the workloads to proceed with reduced performance. Spark was also configured to use the Metastore of Hive for access to the tables and schema of the structured data sets.

4.2 Experimental Overview

Two sets of benchmarks were performed: a set of batch processing workloads over unstructured data sets, and a set of query processing

workloads over structured data sets. Each benchmark had a similar set of steps, as represented by the data-flow diagram in Figure 3.




Figure 3: Data-flow diagram of a big data benchmarking experiment.

4.2.1 Data Generation. Data was generated in the Producer cluster. The unstructured data sets simulating *Wikipedia text entries* were generated with BDGS and were used for the batch processing benchmarks. The structured data sets simulating *e-commerce item/order tables* were generated with the demo version of PDGF³ and were used for the query processing benchmarks. The planned volumes of the input data sets ranged from 64 GB to 8192 GB, but the generated volumes had some variability in size.

The structured e-commerce item/order tables data sets featured two tables, whose sizes are listed in Table 4.

	Real volumes (GB)							
Data set	68.16	138.84	280.79	564.68	1144.61	2320.33	4671.93	9416.87
<i>Item table</i>	42.50	85.84	172.71	346.45	698.00	1406.42	2823.43	5671.80
<i>Order table</i>	25.66	53.00	108.08	218.23	446.61	913.91	1848.50	3745.07

Table 4: Data generation: the real generated table volumes of the e-commerce tables structured data sets.

The two data sets also had distinct layouts in the file system. The *Wikipedia text entries* data sets were divided in files with a fixed average volume of 500 MB each. A sufficient increase in the volume of these data sets resulted in an increased number of files. The *e-commerce item/order tables* data sets were divided in a fixed number of files. Any increase in the volume of these data sets resulted in increased volumes for each file. Table 5 highlights this dichotomy for the largest tested volumes.

Data set	Real volume	Number of files	Volume per file
Wikipedia text entries	8019 GB	16380 files	0.490 GB
E-commerce item tables	5672 GB	4 files	1418 GB
E-commerce order tables	3745 GB	4 files	936 GB

Table 5: Number of files and estimated volume per file of the largest tested input data sets.

³The demo version of PDGF allows parallel generation of structured data in a single-machine, while the full version allows distributed data generation in a cluster. Unfortunately, the full version was not available, so we purposefully extended the demo version to allow distributed data generation.

4.2.2 Data Loading. The input data sets were loaded from the Producer cluster into the the Hadoop cluster with the put operation of HDFS, and loaded into the Unicage cluster with the `distr-distr`⁴ BOA command.

After the data loading process, both the Hadoop and Unicage clusters had the same data sets loaded and ready for the data processing benchmarks.

4.2.3 Data Processing. *Batch processing* workloads handle unstructured data sets. *Query processing* workloads handle structured data sets through the use of a querying language. Both are performed on fully available bounded data sets. Table 6 lists the benchmarked combinations of data processing workload and big data system.

Workload	Big data systems				Description
	Hadoop	Hive	Spark	Unicage	
Batch processing - Wikipedia text entries data set					
<i>grep</i>	Yes	No	Yes	Yes	Count the number of occurrences of a substring in the text data set.
<i>sort</i>	Yes	No	Yes	Yes	Sort a single-column rendition of the text data set, with one word per row.
<i>wordcount</i>	Yes	No	Yes	Yes	Count the number of occurrences of each unique word in the text data set.
Query processing - e-commerce item/order tables data set					
<i>select</i>	No	Yes	Yes	Yes	Select rows of the <i>item</i> table whose target column exceeds a given value.
<i>join</i>	No	Yes	Yes	Yes	Combine rows of the <i>item</i> and <i>order</i> tables based on a related column.
<i>aggregation</i>	No	Yes	Yes	Yes	Sums the values of column A of the <i>item</i> table, grouped by column B.

Table 6: Benchmarked workloads and big data systems.

4.2.4 Monitoring. The data loading and data processing workloads were monitored through the collection of a variety of metrics. *Execution time* was the main performance measure, and *loading/processing rate* was calculated to capture performance fluctuations between benchmarks with different input data set volumes. Additionally, Netdata⁵ was used to collect the system resource usage metrics that allowed a better understanding of the resource usage patterns and aided in the configuration and optimization of the big data systems and workloads.

4.2.5 Verification. The output of each combination of data processing workload and big data system was verified alongside the outputs of peer combinations. For any given workload, to efficiently assess the agreement between the tested big data systems, the output of each big data system was converted to its MD5 hash and the resulting hashes were compared. If the hash of the output of Unicage was the same as the hash of the output of Hadoop and Spark, then there was *agreement*, and the output was assumed to be correct. Singular disagreements were deemed incorrect outputs.

4.2.6 Validation. Every benchmark was repeated thrice, and both the sampled execution times and loading/processing rates were statistically validated through the estimation of the population mean and confidence intervals, as described by Montgomery [17].

⁴`distr-distr` splits a file in a number of chunks equal to the number of Unicage worker nodes, and places one chunk in each worker, in a specified directory.

⁵Netdata is a tool for cluster observability. <https://www.netdata.cloud/>

5 EXPERIMENTAL RESULTS

This section lists and discusses the results produced by the benchmarks.

5.1 Batch Processing Benchmarks

The Wikipedia text entries were used as input data sets for the batch processing workloads in this benchmark set.

5.1.1 Data Loading - Unstructured Data Set. Figures 4 and 5 show the execution times and loading rates, respectively, of loading the Wikipedia text entries into both the Hadoop and Unicage clusters, for the tested input data set volumes.




Figure 4: Data loading execution times for the Wikipedia text entries unstructured data set.

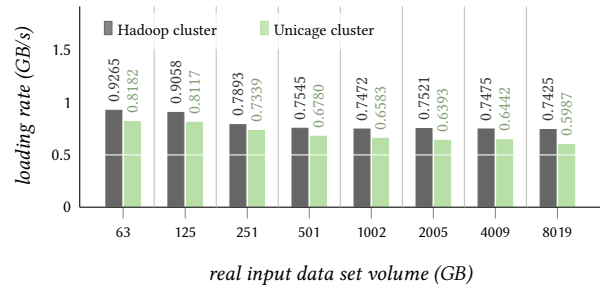


Figure 5: Data loading rates for the Wikipedia text entries unstructured data set.

These two types of bar charts are recurrent for the upcoming results. Bar charts pertaining to the *execution times* highlight the performance differences between the big data systems, for the set of tested input data set volumes. In these charts: lower time bars mean better performance. The vertical axis shows the execution time, in hours, and the horizontal axis shows the set of tested input data set volumes, in gigabytes.

Bar charts pertaining to the *loading* or *processing rates* provide an understanding of the performance fluctuations of each big data system, across the set of tested input data set volumes. In these charts: higher rate bars mean better performance. The vertical axis shows the loading rate (for data loading workloads) or processing rate (for data processing workloads) in gigabytes per second, and the horizontal axis shows the set of tested input data set volumes, in gigabytes.

In both types of chart, for easy reading, grey vertical lines mark the divisions between data set volumes.

Loading the unstructured Wikipedia text entries data sets into the Unicage cluster was slightly slower than loading the same data sets into the Hadoop cluster, across all tested input data set volumes.

5.1.2 *Grep*. Figures 6 and 7 show the execution times and processing rates, respectively, of each big data system when performing the *grep* workload for the tested input data set volumes.

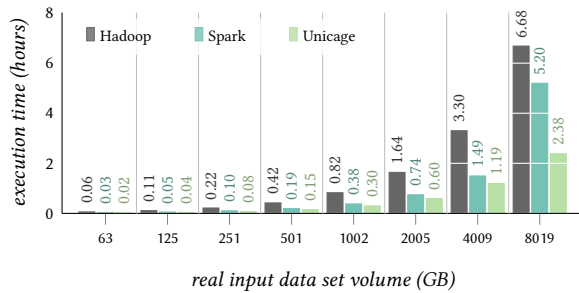


Figure 6: Data processing execution times for *grep*.




Figure 7: Data processing rates for *grep*.

The performance of the *grep* implementation with Unicage was superior to those of Hadoop and Spark, across all tested input data set volumes.

A recurrent pattern in the processing rate charts is the reduced processing rate of Spark for smaller data sets. This is caused by the warm-up period of Spark, when the Driver process is started, along with the SparkContext and when the resource negotiation with YARN takes place. This is more perceivable in workloads that take very little time to complete, such is the case of the *grep* workload for the 63 GB and 125 GB input data sets.

The processing rate of Spark when performing the *grep* workload on the 8019 GB input data set is visibly reduced. This was caused by the performance degradation of Spark when it starts recomputing RDD partitions on demand. A similar performance degradation was observed in the *select*, *join* and *aggregation* workloads.

5.1.3 *Sort*. Figures 8 and 9 show the execution times and processing rates, respectively, of each big data system under test, when performing the *sort* workload for the tested input data set volumes.




Figure 8: Data processing execution times for *sort*.




Figure 9: Data processing rates times for *sort*.

The performance of the *sort* implementation with Unicage was superior to those of Hadoop and Spark for the 63 GB, 125 GB and 251 GB input data sets, but produced arbitrarily incorrect outputs for the 501 GB and larger input data sets.

The output of the *sort* implementation with Spark for the 501 GB input data set was verified through the comparison with the output of a single run of the *sort* implementation with Hadoop, which took approximately 38 hours to complete, and thus was not validated statistically. The *sort* workload was not attempted for larger data sets as the experimental setup did not allow the efficient verification of the outputs of Spark.

5.1.4 *Wordcount*. Figures 10 and 11 show the execution times and processing rates, respectively, of each system performing the *wordcount* workload for the tested input data set volumes.




Figure 10: Data processing execution times for *wordcount*.

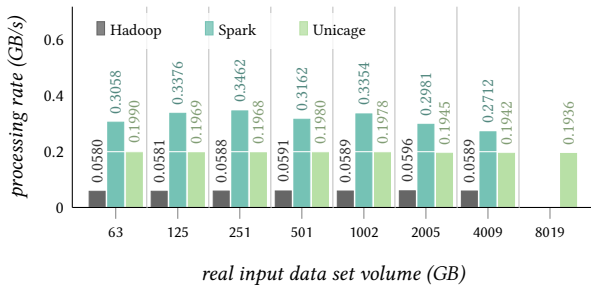


Figure 11: Data processing rates for wordcount.

The performance of the wordcount implementation with Unicage was superior to that of Hadoop and inferior to that of Spark for the data sets with volumes ranging from 63 GB to 4009 GB. The wordcount implementation with Spark failed to produce correct outputs for the 8019 GB data set, as the JVM ran out of space in the Java heap during the execution of the workload, as indicated by the throwing of a java.lang.OutOfMemoryError exception.

The output of the wordcount implementation with Unicage for the 8019 GB input data set was verified through the comparison with the output of a single run of the wordcount implementation with Hadoop, which took approximately 38 hours to complete, and thus was not validated statistically.

5.2 Query Processing Benchmarks

The structured e-commerce item/order tables data sets were used as input data sets for the query processing workloads in this benchmark set.

5.2.1 Data Loading - Structured Data Set. Figures 12 and 13 show the execution times and loading rates, respectively, of loading the structured e-commerce item/order tables data sets into both the Hadoop and Unicage clusters, for the tested input data set volumes.




Figure 12: Data loading execution times for the e-commerce tables structured data set.

The stacked bar in Figure 12 indicates loading a structured data set into the Hadoop cluster required the step of loading it into HDFS, and the additional step of creating the Hive tables to support its schema. Loading the same data set into the Unicage cluster was significantly faster, as Unicage handles structured and unstructured data equally, as long as the data sets can be represented as text files

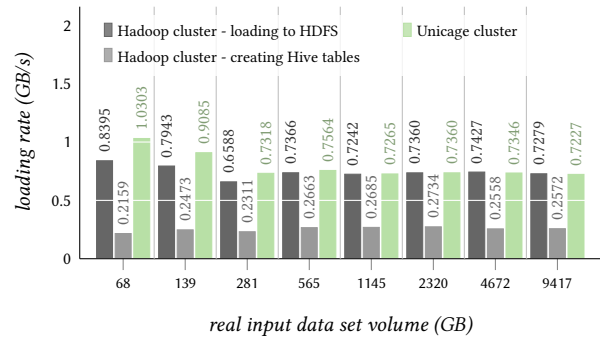


Figure 13: Data loading rates for the e-commerce tables structured data set.

in the file system. The trade-off is that in the absence of formal schema, it is of the responsibility of the developer to preserve the structure and veracity of the data set.

5.2.2 Select. Figures 14 and 15 show the execution times and processing rates, respectively, of each big data system under test, when performing the select workload for the tested input data set volumes.




Figure 14: Data processing execution times for select.




Figure 15: Data processing rates for select.

The select workload used the structured e-commerce tables, but only performed the search on the item table. Therefore, the horizontal axis of the execution time and processing rate charts reflect the real volumes of the item table only, as shown previously in Table 4.

The performance of the *select* implementation with Unicage was superior to those of Hive and Spark, across all tested input data set volumes.

5.2.3 *Join*. Figures 16 and 17 show the execution times and processing rates, respectively, of each system under test, when performing the *join* workload for the input data set volumes.




Figure 16: Data processing execution times for *join*.




Figure 17: Data processing rates for *join*.

The performance of the *join* implementation with Unicage was vastly inferior to those of Hive and Spark for the 68 GB and 139 GB input data sets, and produced arbitrarily incorrect outputs for the 281 GB and larger data sets.

5.2.4 *Aggregation*. Figures 18 and 19 show the execution times and processing rates for the systems performing the *aggregation* workload.

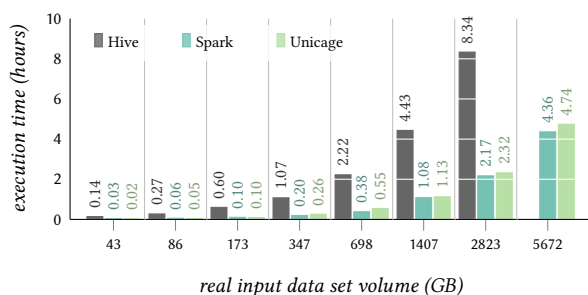


Figure 18: Data processing execution times for *aggregation*.




Figure 19: Data processing rates for *aggregation*.

Similarly to the *select* workload, the *aggregation* workload only performed the grouping on the *item* table. The horizontal axis of the execution time and processing rate charts reflect the real volumes of this table. The performance of the *aggregation* implementation with Unicage was superior to those of Hive and Spark for the 63 GB, 86 GB and 173 GB data sets, but inferior to that of Spark for the 347 GB and larger data sets.

Many Unicage commands, including the `sm2` grouping operation used in the *aggregation* implementation, require the input files to be sorted. Sorting with Unicage is most efficient with the `msort` command, which performs in-memory sorting. Each file of the structured e-commerce item/order tables data set can grow quite large, as shown previously in Table 5. This caused the premature failure of the *aggregation* implementation with Unicage, as `msort` ran out of memory for the 347 GB data set. To overcome this, the implementation was refactored to split each input file in chunks small enough to be sorted by `msort`. This allowed the benchmarks to proceed, but caused the performance degradation of the *aggregation* implementation with Unicage, observed for the 347 GB and larger data sets.

The *aggregation* implementation with Hive failed to produce correct outputs for the 5672 GB data set, with a failing MapReduce.

5.3 Discussion

Complex big data systems offer a plethora of abstractions to ease the implementation and execution of big data processing workloads. In our experimental setup, HDFS provided distributed storage and YARN provided distributed resource management to the Hadoop cluster.

5.3.1 *Distributed File System*. A sorted data set has inter-record dependencies, meaning each record has a placement dependency on other records. Unicage requires output data sets with this property to be stored in a single machine. For this reason, even though it failed prematurely, the *sort* implementation with Unicage was predicted to fail for data sets larger than 1024 GB, since that was the maximum storage capacity of the *unicageleader* node (see Table 3), where the outputs of the workloads were consolidated. In this aspect, Unicage can scale horizontally to improve the performance of the workloads, but is hindered by the capacity of singular machines, being highly dependent on vertical scaling to address larger volumes. HDFS offers the abstraction of a global namespace not only to the user but also to the applications, allowing large inter-record dependent outputs to be stored across multiple nodes. This

is not immediately possible in lean big data systems without the integration of additional systems to support a similar abstraction to that of a distributed file system.

A similar comment can be made regarding the *join* workload, which explores the inter-record dependencies of the input data set. To join two records from different tables, Unicage required these to be stored in the same node. As such, the *join* implementation with Unicage required intermediate files to be moved around the cluster, severely crippling the performance. Since HDFS allows applications to access the distributed file system as a single global namespace, having joinable records in different machines was not a challenge for the tested complex big data systems.

5.3.2 Distributed Resource Manager. Another challenge of lean big data systems was the optimization of the resource usage, which was a manual endeavor that was only possible with the assistance of observability tools. This process was also error prone, as the resource usage patterns of shell script implementations often became difficult to predict with the increase in volume of the input data sets. Optimizing resource usage was not a challenge with the tested complex big data systems, as that is the main function of YARN.

5.3.3 Data Loading Performance - the Impact of Data Sets with Different File System Layouts. The Wikipedia text entries data set was comprised of a growing number of 513 MB files, while the e-commerce item/order tables data set was comprised of a fixed number of 8 files with growing volumes. This was previously highlighted in Table 5. This dichotomy translated into a small performance impact for both data loading into the Hadoop cluster with HDFS and into the Unicage cluster.

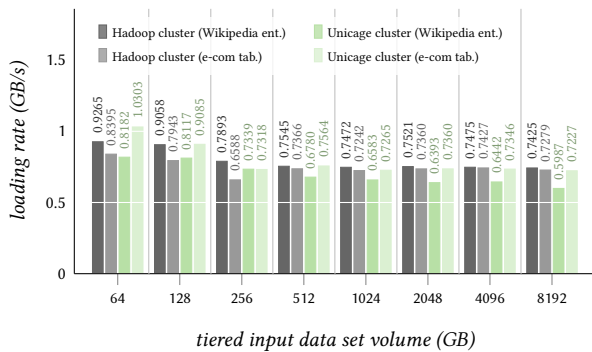


Figure 20: Composite chart of data loading rates.

Figure 20 shows a composite chart of the loading rates of all benchmarked data loading workloads. The vertical axis shows the loading rate, in gigabytes per second, and the horizontal axis shows the tiered input data set volumes, in gigabytes. For the structured e-commerce item/order tables data set, the process of creating the Hive tables in the Hadoop cluster was excluded from the chart because that is a consequence of the structural properties of the data set, and not of its layout in the file system and as far as HDFS is aware, both data sets are just files.

The performance of loading data into the Hadoop cluster with the put operation of HDFS is slightly increased for data sets with a large number of small files, such as the Wikipedia text entries data

sets, providing each file is large enough to fill file blocks in the file system (128 MB each, by default). The performance of loading data into the Unicage cluster with the `distr-distr` Unicage command is slightly increased for data sets with a small number of large files, such as the e-commerce item/order tables data sets.

6 CONCLUSION

This paper presented a performance benchmark of two systems with two distinct approaches to big data processing in cluster environments: Spark using a complex software stack, and Unicage shell scripts using a lean software stack. We benchmarked the performance of batch processing workloads over unstructured data sets, and query processing workloads over structured data sets. The tiered volumes of the data sets ranged from 64 GB to 8192 GB. We also benchmarked the data loading performance.

We have shown that there are big data workloads that do not require the use of complex big data systems: for search workloads (*grep* and *select*), Unicage shell scripts were preferred to Spark, across all tested input data set volumes; for grouping workloads (*wordcount* and *aggregation*), Unicage was only preferred to Spark for large data sets (8192 GB); for the *sort* workload, Unicage was only preferred to Spark for smaller data sets (from 64 GB to 256 GB); for the *join* workload, Spark was preferred to Unicage, across all tested input data set volumes. Figure 21 summarizes these findings, as they apply in the tested experimental environment.

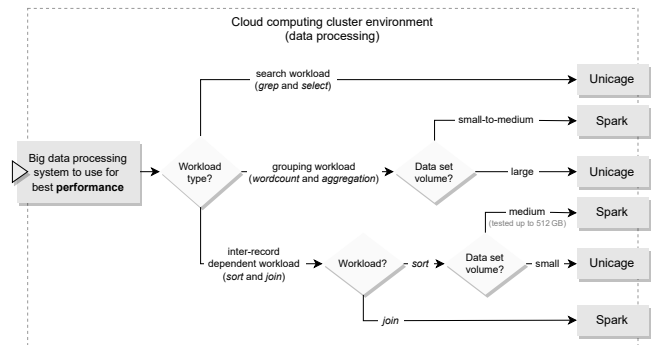


Figure 21: Summary of conclusions from the data processing benchmarks.

Regarding the data loading, we have shown that the performance of loading unstructured data into a Unicage cluster was slightly inferior to that of loading the same data into a Hadoop cluster with HDFS, and the performance of loading structured data into a Unicage cluster was superior to that of loading the same data into a Hadoop cluster with HDFS and Hive, across all tested input data set volumes.

Our findings favor further work on complex and lean big data systems. Future studies should address more complex sets of workloads in application-specific scenarios. Other types of data processing can also be benchmarked, including stream processing and machine learning. The obstacles found in the *sort* and *join* workloads motivate the future integration of distributed storage and resource usage abstractions with lean big data systems. The integration of recent technologies such as PASH [28] and PosH [21] with lean big

data systems could also provide further performance optimizations to the shell script-based approach.

REFERENCES

- [1] Michael Armbrust, Tathagata Das, Joseph Torres, Burak Yavuz, Shixiong Zhu, Reynold Xin, Ali Ghodsi, Ion Stoica, and Matei Zaharia. 2018. Structured streaming: A declarative API for real-time applications in Apache Spark. In *Proceedings of the International Conference on Management of Data*. 601–613.
- [2] Michael Armbrust, Reynold S Xin, Cheng Lian, Yin Huai, Davies Liu, Joseph K Bradley, Xiangrui Meng, Tomer Kaftan, Michael J Franklin, Ali Ghodsi, et al. 2015. Spark SQL: Relational data processing in Spark. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*. 1383–1394.
- [3] Paris Carbone, Asterios Katsifodimos, Stephan Ewen, Volker Markl, Seif Haridi, and Kostas Tzoumas. 2015. Apache Flink: Stream and batch processing in a single engine. *Bulletin of the IEEE Computer Society Technical Committee on Data Engineering* 36, 4 (2015).
- [4] Andrea De Mauro, Marco Greco, and Michele Grimaldi. 2016. A formal definition of Big Data based on its essential features. *Library Review* (2016).
- [5] Jeffrey Dean and Sanjay Ghemawat. 2004. MapReduce: Simplified data processing on large clusters. (2004).
- [6] Jeffrey Dean and Sanjay Ghemawat. 2008. MapReduce: Simplified data processing on large clusters. *Commun. ACM* 51, 1 (2008), 107–113.
- [7] Bo Dong, Qinghua Zheng, Feng Tian, Kuo-Ming Chao, Rui Ma, and Rachid Anane. 2012. An optimized approach for storing and accessing small files on cloud storage. *Journal of Network and Computer Applications* 35, 6 (2012), 1847–1862.
- [8] Miguel Ferreira, André Neves, Rodrigo Gorjao, Carlos Cruz, and Miguel L Pardal. 2021. Smart meter data processing: A showcase for simple and efficient textual processing. [*Private communication*] (2021).
- [9] Wanling Gao, Jianfeng Zhan, Lei Wang, Chunjie Luo, Daoyi Zheng, Xu Wen, Rui Ren, Chen Zheng, Xiwen He, Hainan Ye, et al. 2018. BigBataBench: A scalable and unified big data and AI benchmark suite. *arXiv preprint arXiv:1802.08254* (2018).
- [10] Ahmad Ghazal, Todor Ivanov, Pekka Kostamaa, Alain Crolotte, Ryan Voong, Mohammed Al-Kateb, Waleed Ghazal, and Roberto V Zicari. 2017. BigBench v2: The new and improved BigBench. In *33rd IEEE International Conference on Data Engineering*. IEEE, 1225–1236.
- [11] Ahmad Ghazal, Tilmann Rabl, Mingqing Hu, Francois Raab, Meikel Poess, Alain Crolotte, and Hans-Arno Jacobsen. 2013. BigBench: Towards an industry standard benchmark for big data analytics. In *Proceedings of the ACM SIGMOD International Conference on Management of Data*. 1197–1208.
- [12] Joseph E Gonzalez, Reynold S Xin, Ankur Dave, Daniel Crankshaw, Michael J Franklin, and Ion Stoica. 2014. GraphX: Graph processing in a distributed dataflow framework. In *11th USENIX Symposium on Operating Systems Design and Implementation*. 599–613.
- [13] Rui Han, Lizy Kurian John, and Jianfeng Zhan. 2017. Benchmarking big data systems: A review. *IEEE Transactions on Services Computing* 11, 3 (2017), 580–597.
- [14] Arne Holst. 2021. *Volume of data/information created, captured, copied, and consumed worldwide from 2010 to 2025*. Technical Report. Statista.
- [15] Xiangrui Meng, Joseph Bradley, Burak Yavuz, Evan Sparks, Shivaram Venkataraman, Davies Liu, Jeremy Freeman, DB Tsai, Manish Amde, Sean Owen, et al. 2016. MLlib: Machine learning in Apache Spark. *The Journal of Machine Learning Research* 17, 1 (2016), 1235–1241.
- [16] Zijian Ming, Chunjie Luo, Wanling Gao, Rui Han, Qiang Yang, Lei Wang, and Jianfeng Zhan. 2013. BDGS: A scalable big data generator suite in big data benchmarking. In *Advancing Big Data Benchmarks*. Springer, 138–154.
- [17] Douglas C Montgomery and George C Runger. 2010. *Applied statistics and probability for engineers*. John Wiley & Sons.
- [18] Joao MP Moreira, Helena Galhardas, and Miguel L Pardal. 2018. LeanBench: Comparing software stacks for batch and query processing of IoT data. *Procedia Computer Science* 130 (2018), 448–455.
- [19] Rui Pereira, Marco Couto, Francisco Ribeiro, Rui Rua, Jácome Cunha, João Paulo Fernandes, and João Saraiva. 2017. Energy efficiency across programming languages: How do energy, time, and memory relate?. In *Proceedings of the 10th ACM SIGPLAN International Conference on Software Language Engineering*. 256–267.
- [20] Tilmann Rabl, Michael Frank, Hatem Mousselly Sergieh, and Harald Kosch. 2010. A data generator for cloud-scale benchmarking. In *Technology Conference on Performance Evaluation and Benchmarking*. Springer, 41–56.
- [21] Deepthi Raghavan, Sadjad Fouladi, Philip Levis, and Matei Zaharia. 2020. POSH: A data-aware shell. In *USENIX Annual Technical Conference*. 617–631.
- [22] Markus Rothmuller and Sam Barker. 2020. *IoT - The Internet of Transformation 2020*. Technical Report. Juniper Research.
- [23] Sally. 2021. *Apache in 2020 - By the digits*. <https://news.apache.org/foundation/entry/apache-in-2020-by-the> (Visited on 09/27/2022).
- [24] Sally. 2022. *Apache in 2021 - By the digits*. <https://news.apache.org/foundation/entry/apache-in-2021-by-the> (Visited on 09/27/2022).
- [25] Konstantin V. Shvachko, Hairong Kuang, Sanjay R. Radia, and Robert J. Chansler. 2010. The Hadoop Distributed File System. *26th IEEE Symposium on Mass Storage Systems and Technologies* (2010), 1–10.
- [26] Ashish Thusoo, Joydeep Sen Sarma, Namit Jain, Zheng Shao, Prasad Chakka, Suresh Anthony, Hao Liu, Pete Wyckoff, and Raghotham Murthy. 2009. Hive: A warehousing solution over a Map-Reduce framework. *Proceedings of the VLDB Endowment* 2, 2 (2009), 1626–1629.
- [27] Ashish Thusoo, Joydeep Sen Sarma, Namit Jain, Zheng Shao, Prasad Chakka, Ning Zhang, Suresh Antony, Hao Liu, and Raghotham Murthy. 2010. Hive - A petabyte scale data warehouse using Hadoop. In *26th IEEE International Conference on Data Engineering*. IEEE, 996–1005.
- [28] Nikos Vasilakis, Konstantinos Kallas, Konstantinos Mamouras, Achilles Benetopoulos, and Lazar Cvetković. 2021. PASH: Light-touch data-parallel shell processing. In *Proceedings of the Sixteenth European Conference on Computer Systems*. 49–66.
- [29] Vinod Kumar Vavilapalli, Arun C Murthy, Chris Douglas, Sharad Agarwal, Mahadev Konar, Robert Evans, Thomas Graves, Jason Lowe, Hitesh Shah, Siddharth Seth, et al. 2013. Apache Hadoop YARN: Yet another resource negotiator. In *Proceedings of the 4th annual Symposium on Cloud Computing*. 1–16.
- [30] Lei Wang, Jianfeng Zhan, Chunjie Luo, Yuqing Zhu, Qiang Yang, Yongqiang He, Wanling Gao, Zhen Jia, Yingjie Shi, Shujie Zhang, et al. 2014. BigDataBench: A big data benchmark suite from internet services. In *20th IEEE International Symposium on High Performance Computer Architecture*. IEEE, 488–499.
- [31] Matei Zaharia, Mosharaf Chowdhury, Tathagata Das, Ankur Dave, Justin Ma, Murphy McCauly, Michael J Franklin, Scott Shenker, and Ion Stoica. 2012. Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing. In *9th USENIX Symposium on Networked Systems Design and Implementation*. 15–28.
- [32] Matei Zaharia, Mosharaf Chowdhury, Michael J Franklin, Scott Shenker, Ion Stoica, et al. 2010. Spark: Cluster computing with working sets. *HotCloud* 10, 10–10 (2010), 95.
- [33] Matei Zaharia, Tathagata Das, Haoyuan Li, Timothy Hunter, Scott Shenker, and Ion Stoica. 2013. Discretized streams: Fault-tolerant streaming computation at scale. In *Proceedings of the 24th ACM Symposium on Operating Systems Principles*. 423–438.