

A Big Data analyzer for large trace logs

Alkida Balliu¹ · Dennis Olivetti¹ ·
Ozalp Babaoglu² · Moreno Marzolla² ·
Alina Sirbu²

Received: 4 August 2015 / Accepted: 10 December 2015 / Published online: 28 December 2015
© Springer-Verlag Wien 2015

Abstract Current generation of Internet-based services are typically hosted on large data centers that take the form of warehouse-size structures housing tens of thousands of servers. Continued availability of a modern data center is the result of a complex orchestration among many internal and external actors including computing hardware, multiple layers of intricate software, networking and storage devices, electrical power and cooling plants. During the course of their operation, many of these components produce large amounts of data in the form of event and error logs that are essential not only for identifying and resolving problems but also for improving data center efficiency and management. Most of these activities would benefit significantly from data analytics techniques to exploit hidden statistical patterns and correlations that may be present in the data. The sheer volume of data to be analyzed makes uncovering these correlations and patterns a challenging task. This paper presents Big Data analyzer (*BiDAI*), a prototype Java tool for log-data analysis that incorporates several Big Data technologies in order to simplify the task of extracting information from data traces

✉ Alina Sirbu
alina.sirbu@unibo.it

Alkida Balliu
alkida.balliu@gssi.infn.it

Dennis Olivetti
dennis.olivetti@gssi.infn.it

Ozalp Babaoglu
ozalp.babaoglu@unibo.it

Moreno Marzolla
moreno.marzolla@unibo.it

¹ Gran Sasso Science Institute (GSSI), L'Aquila, Italy

² Department of Computer Science and Engineering, University of Bologna, Bologna, Italy

produced by large clusters and server farms. *BiDAI* provides the user with several analysis languages (SQL, R and Hadoop MapReduce) and storage backends (HDFS and SQLite) that can be freely mixed and matched so that a custom tool for a specific task can be easily constructed. *BiDAI* has a modular architecture so that it can be extended with other backends and analysis languages in the future. In this paper we present the design of *BiDAI* and describe our experience using it to analyze publicly-available traces from Google data clusters, with the goal of building a realistic model of a complex data center.

Keywords Big Data · Log analysis · Workload characterization · Google cluster trace · Model · Simulation

Mathematics Subject Classification 68N01 · 68P20 · 68U20

1 Introduction

Large data centers are the engines of the Internet that run a vast majority of modern Internet-based services such as cloud computing, social networking, online storage and media sharing. A modern data center contains tens of thousands of servers and other components (e.g., networking equipment, power distribution, air conditioning) that may interact in subtle and unintended ways, making management of the global infrastructure a nontrivial task. Failures are extremely costly both for data center operators and their customers, since the services provided by these huge infrastructures have become vital to society in general. In this light, monitoring and managing large data centers to keep them running correctly and continuously become critical tasks.

The amount of log data produced by modern data centers is growing steadily, making log management itself technically challenging. For instance, a 2010 Facebook study reports 60 TB of log data being produced by its data centers each day [34]. For live monitoring of its systems and analyzing their log data, Facebook has developed a dedicated software tool called Scuba [2] that uses a large in-memory database running on hundreds of servers with 144 GB of RAM each. This infrastructure needs to be upgraded every few weeks to keep up with the increasing computational power and storage requirements that Scuba generates.

Making sense of these huge data streams is a task that continues to rely heavily on human judgement, and is therefore error-prone, time-consuming and potentially inefficient. Log analysis falls within the class of Big Data applications: the data sets are so large that conventional storage and analysis techniques are not appropriate to process them. There is a real need to develop novel tools and techniques for analyzing logs, possibly incorporating data analytics to uncover hidden patterns and correlations that can help system administrators avoid critical states, or to identify the root cause of failures or performance problems. The “holy grail” of system management is to render data centers fully autonomic; ideally, the system should be capable of analyzing its state and use this information to identify performance or reliability problems and correct them or alert system managers directing them to the root causes of the problem. Even better, the system should be capable of anticipating situations that may lead to

performance problems or failures, allowing for proactive countermeasures to be put in place in order to steer the system away from undesirable states towards desired operational states. These challenging goals are still far from being realized [28].

Numerous studies have analyzed trace data from a variety of sources for different purposes (see the related work in Sect. 4), but typically without relying on an integrated software framework developed specifically for log analysis [7, 22, 26]. Reasons for this are several fold: first, the amount, content and structure of logs are often system- and application-specific, requiring ad-hoc solutions that are difficult to port to other contexts. Furthermore, log trace data originating from commercial services are highly sensitive and need to be kept strictly confidential. All these facts lead to fragmentation of analysis frameworks and difficulty in porting them to traces from other sources. One isolated example of analysis framework is the Failure Trace Archive Toolkit [19], limited however to failure traces. Lack of a more general framework for log data analysis results in time being wasted “reinventing the wheel”—developing software for parsing, interpreting and analyzing the data, repeatedly for each new trace [19].

As a first step towards realizing the above goals, we present Big Data analyzer (*BiDAL*), a prototype software tool implementing a general framework for statistical analysis of very large trace data sets. *BiDAL* is built around two main components: a storage backend and an analysis framework for data processing and reduction. Currently, *BiDAL* supports HDFS and SQLite as storage backends, and SQL, R, and Hadoop MapReduce as analysis frameworks. However, *BiDAL* is extensible so that additional backends and analysis frameworks can be easily added, and multiple types can coexist and be used at the same time.

After describing the architecture of *BiDAL*, we illustrate how it has been used to analyze publicly-available Google cluster trace data [39]. These are *real trace data* describing the workload and machine status for a Google cluster consisting of 12,453 nodes, monitored over a 29-day period starting from 19:00 EDT on May 1st 2011. Over 1.3 billion records are present, which include job, task and machine events, as well as resource usages totaling over 40GB of compressed data (just under 200 GB raw). We have analyzed these data in order to extract parameters of a cluster model which we have implemented. The source code for both the *BiDAL* prototype and the model are freely available (see Sect. 5).

The contributions of this work are several fold. First, we present *BiDAL* and describe its architecture incorporating several Big Data technologies that facilitate efficient processing of large datasets for data analytics. Then, we describe an application scenario where we use *BiDAL* to extract workload parameters from Google cluster traces. We introduce a model of the Google cluster which allows for simulation of the Google system. Depending on the input to the model, several types of simulations can be performed. Using the exact workload from the Google trace as input, our model is able to faithfully reproduce many of the behaviors that are observed in the traces. By providing the model with distributions of the various parameters that are obtained using *BiDAL* more general workloads can also be simulated; in this scenario, simulation results show that our model is able to approximate average behavior, although variability is lower than in the real counterpart.

This paper is organized as follows. In Sect. 2 we provide a high level overview of the framework followed by a detailed description of its components. In Sect. 3

we apply the framework to characterize the workload from a public Google cluster trace, and use this information to build a model of the Google cluster and perform simulations. In Sect. 4 we discuss related work, and conclude with new directions for future research in Sect. 5.

2 The Big Data analyzer (*BiDAL*) prototype

2.1 General overview

The typical *BiDAL* workflow consists of three steps: instantiation of a storage backend (or opening an existing one), data selection and aggregation, and data analysis; Fig. 1 shows the overall data flow within *BiDAL*.

For storage creation, *BiDAL* is designed to import CSV files (comma separated values, the typical format for trace data) into an SQLite database or to a Hadoop file system (HDFS) storage, depending on the user's preference; HDFS is the preferred choice for handling large amounts of data using the Hadoop framework. Except for the CSV format, no other restrictions on the data type exist, so the platform can be easily used for data from various sources, as long as they can be viewed as CSV tables. Even though the storages currently implemented are based on the the concept of tables (stored in a relational database by SQLite and CSV files by Hadoop), other storage types can be supported by *BiDAL*. Indeed, Hadoop supports HBase, a non-relational database that works with $\langle \text{key}, \text{value} \rangle$ pairs. Since Hadoop is already supported by *BiDAL*, a new storage that works on this type of non-relational databases can be easily added.

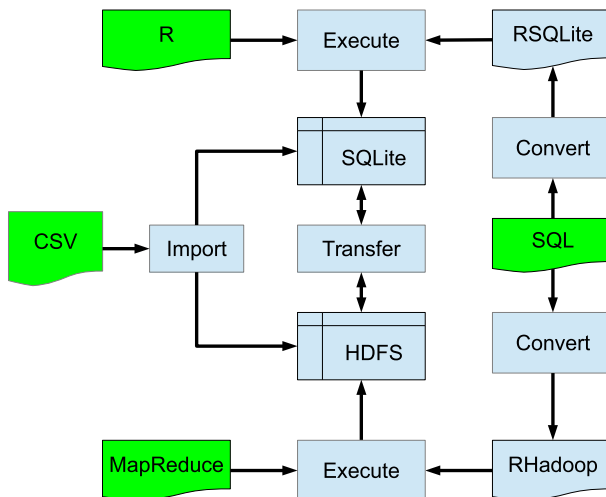


Fig. 1 Data flow in *BiDAL*. Raw data in CSV format is imported into the selected storage backend, and can be selected or aggregated into new tables using SQL queries. Commands using R or MapReduce can then be applied both to the original imported data and to the derived tables. Data can be automatically and transparently moved between the storage backends

Selections and aggregations can be performed through queries expressed using a subset of SQL, for example to create new tables or to filter existing data. SQL queries are automatically translated into the query language supported by the underlying storage system (SQLite or Hadoop). At the moment, the supported statements in the SQL subset are SELECT, FROM, WHERE and GROUP BY. Queries executed on the SQL storage do not require any processing, since the backend (SQLite) already supports a larger subset of SQL. For the Hadoop backend, GROUP BY queries are mapped to MapReduce operations. The WHERE and SELECT clauses are implemented in the Map function, which generates keys corresponding to the attributes in the GROUP BY clause. Then Reduce applies the aggregate function to the result. At the moment, only one column can be selected with the SELECT clause.

BiDAL can perform statistical data analysis using both R [25] and Hadoop MapReduce [9, 31] by offering a set of predefined commands. Commands implemented in R are typically applied to the SQLite storage, while those in MapReduce to the Hadoop storage. However, the system allows mixed execution of both types of commands regardless of the storage used, being able to switch between backends (by exporting data) transparent to the user. For instance, after a MapReduce command, it is possible to analyze the outcome using commands implemented in R; in this case, the software automatically exports the result obtained from the MapReduce step, and imports it to the SQLite storage where the analysis can continue using commands implemented in R. This is particularly useful for handling large datasets, since the volume of data can be reduced by applying a first processing step with Hadoop/MapReduce, and then using R to complete the analysis on the resulting (smaller) dataset. The drawback is that the same data may end up being duplicated into different storage types so, depending on the size of the dataset, additional storage space will be consumed. However, this does not generate consistency issues, since log data does not change once it is recorded.

2.2 Design

BiDAL is a modular application designed for extensibility and ease of use. It is written in Java, to facilitate portability across different operating systems, and uses a graphical user interface (GUI) based on the standard Model–View–Controller (MVC) architectural pattern [13]. The View provides a Swing GUI, the model manages different types of storage backends, and the Controller handles the interaction between the two. Fig. 2 outlines the architecture using the UML class diagram.

The Controller class connects the GUI with the other components of the software. The Controller implements the Singleton pattern, with the one instance accessible from any part of the code. The interface to the different storage backends is given by the GenericStorage class, that has to be further specialized by any concrete backend developed. In our case, the two existing concrete storage backends are represented by the SqliteStorage class to support SQLite, and the HadoopStorage class, to support HDFS. Neither the Controller nor the GUI elements communicate directly with the concrete storage backends, but only with the abstract class GenericStorage. This sim-




Fig. 2 UML diagram of *BiDAI* classes. This shows the modular structure where the storage is separated from the user interface, facilitating addition of new types of storage backends

plifies the implementation of new backends without the need to change the Controller or GUI implementations.

The user can inspect and modify the data storage using a subset of SQL; the *SqliteStorage* and *HadoopStorage* classes use the open source SQL parser Akiban to convert the queries inserted by users into SQL trees that are further mapped to the native language (RSQLite or RHadoop) using the Visitor pattern. The *HadoopStorage* uses also a *Bashexecuter* that allows to load files on the HDFS using bash shell commands. A new storage class can be implemented by providing a suitable specialization of the *GenericStorage* class, including the mapping of the SQL tree to specific commands understood by the backend. In particular, the simple *Sqlite* backend can be replaced by a more scalable storage service such as *Hive* [33] or *Impala* [21] (both based on Hadoop), and *Spark SQL* [3] (based on Apache Spark). These systems all support the SQL language to some extent. *Pig Latin* [24] (still based on Hadoop) uses a dialect of SQL and would require a simple translator to be used within our tool. *Hive*, *Impala* and *Pig Latin* are of particular interest since they are all based on Hadoop; therefore, they allow users to execute SQL (or SQL-like) queries directly on the *BiDAI* Hadoop data store, without the need to transfer them to a separate SQL database.

Although the SQL parser supports the full SQL language, the developer must define a mapping of the SQL tree into the language supported by the underlying storage; this often limits the number of SQL statements that can be supported due to the difficulty of realizing such a mapping.

2.3 Using R with *BiDAI*

BiDAI provides a list of predefined commands, implemented in R, that can be selected by the user from a graphical interface (see Fig. 3 for a screenshot and Table 1 for a

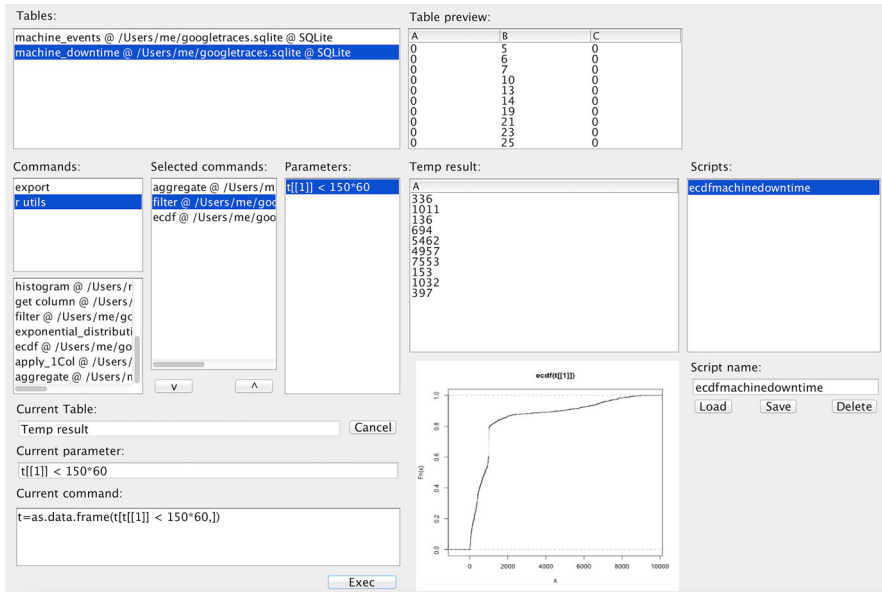


Fig. 3 Screenshot of the BiDAL analysis console. In the *upper-left corner*, we see the list of available tables in the current storage backend (SQLite in this case), and the list of available commands (implemented in R). The results of running the selected command (ecdf) on the selected table (Machine_Downtime) are shown in the *plot at the bottom*. The command implementation can be edited in the *lower-left panel*. New commands can be saved, with a list of existing custom commands displayed in the “Scripts” *panel to the right*

partial list of the available commands). When a command is selected, an input box appears asking the user to provide the parameters needed by that specific command. Additionally, a text box (bottom-left corner of Fig. 3) allows the user to modify on the fly the R code to be executed.

All commands are defined in an external text file. New operations can therefore be added quite easily by simply including them in the file.

2.4 Using Hadoop/MapReduce with BiDAL

BiDAL allows computations to be distributed across many machines through the Hadoop/MapReduce abstractions. The user can access any of the builtin commands implemented in RHadoop, or create new ones. Usually, the Mapper and Reducer are implemented in Java, generating files that need to be compiled and then executed. However, BiDAL abstracts from this approach by using the RHadoop library which handles MapReduce job submission and permits to interact with Hadoop’s file system HDFS using R. This allows for reuse of the BiDAL R engine for the Hadoop backend. Once the dataset of interest has been chosen, the user can execute the Map and Reduce commands implemented in RHadoop or create new ones. Again, the commands and corresponding RHadoop code are saved in an external text file, using the same format

Table 1 A partial list of *BiDAL* commands implemented in R

<i>BiDAL</i> command	Description
<code>get_column</code>	Selects a column
<code>apply_1Col</code>	Applies the desired R function to each element of a column
<code>aggregate</code>	Takes as input a column to group by; among all rows selects the ones that satisfies the specified condition; the result obtained is specified from the R function given to the third parameter
<code>difference_between_rows</code>	Calculates the differences between consecutive rows
<code>filter</code>	Filters the data after the specified condition
<code>exponential_distribution</code>	Plots the fit of the exponential distribution to the data
<code>lognormal_distribution</code>	Plots the fit of the lognormal distribution to the data
<code>polynomial_regression</code>	Plots the fit of the n-grade polynomial regression to the data in the specified column
<code>ecdf</code>	Plots the cumulative distribution function of the data in the specified column
<code>spline</code>	Divides the data in the specified column in n intervals and for each range plots spline functions. Also allows to show a part of the plot or all of it
<code>log_histogram</code>	Plots the histogram of the data in the specified column, using a logarithmic y-axis

described above, so the creation of new commands does not require any modification to *BiDAL* itself. At the moment, one Map command is implemented in *BiDAL*, which groups the data by the values of a column. A Reduce command is also available, which counts the elements of each group. Other commands can be added by the user, similar to those implemented in R.

3 Case study

The development of *BiDAL* was motivated by the need to process large data from cluster traces, such as those publicly released by Google [39]. Our goal was to extract workload parameters from the traces in order to instantiate a model of the compute cluster capable of reproducing the most important features observed in the real data. The model, then, could be used to perform “what-if analyses” by simulating different scenarios where the workload parameters are different, or several types of faults are injected into the system.

In this section we first present the structure of the model, then describe the use of *BiDAL* for analyzing the Google traces and extracting parameters for the model.

3.1 Modeling the Google compute cluster

We built a model of the Google compute cluster corresponding to that from which the traces were obtained. According to available information, the Google cluster is basically a large system where computational tasks of different types are submitted and




Fig. 4 Simple model of a Google compute cluster. This includes active entities that exchange messages among each other (Machine Arrival, Job Arrival, Scheduler, Network and Machine), and passive entities that are silent (Jobs and Tasks). The *arrows* show the flow of information

executed on a large server pool. Each job may describe constraints for its execution (e.g., a minimum amount of available RAM on the execution host); a scheduler is responsible for extracting jobs from the waiting queue, and dispatching them to a suitable execution host. As can be expected on a large infrastructure, jobs may fail and can be resubmitted; moreover, execution hosts may fail as well and be temporarily removed from the pool, or new hosts can be added. The Google trace contains a list of timestamped events such as job arrival, job completion, activation of a new host and so on; additional (anonymized) information on job requirements is also provided.

The model, shown in Fig. 4, consists of several active and passive interacting entities. The passive entities (i.e., those that do not exchange any message with other entities) are Jobs and Tasks. The active entities are those that send and receive messages: Machine, Machine Arrival, Job Arrival, Scheduler and Network. The model was implemented using C++ and Omnet++ [35], a discrete-event simulation tool.

A Task represents a process in execution, or ready to be executed. Each task is characterized by its Id, the information regarding the *requested* and *used* resources (CPU and RAM), its priority, duration, termination cause and other information regarding the execution constraints. Note that the total duration of a task, the termination cause and the effective use of resources are *not* used to take decisions (for example on which machine to execute a task). This choice is necessary in order to simulate a real scenario, where one does not know in advance the length, exit code and resource usage of a task.

A Job is identified by a unique ID, and can terminate either because all of its tasks complete execution, or because it is aborted. In the real traces, this latter outcome

occurs through either a KILL or FAIL event. KILL events are triggered by the user while FAIL events are generated by the scheduler (when job tasks fail multiple times). Due to the fact that FAIL job events were few compared to the other event types, we considered KILL and FAIL as a single event category in the model. Note that in the Google cluster, tasks from the same job do not necessarily have to be executed at the same time. The Job Arrival entity generates events that signal new jobs being submitted. At each event, a Job entity is created and sent to the scheduler.

The Machine entity represents an execution node in the compute cluster. Each machine is characterized by an Id and its maximum amount of free resources. Machine Arrival is the entity in charge of managing all machines, and generates events related to addition or removal of machines to the cluster, as well as update events (when the maximal resources of the machine are changed). In all of these cases, the Machine entity is notified of these changes. At regular intervals, each Machine will notify the Scheduler about the free resources owned. Free resources are computed as the difference between the total resources and those used by all tasks running on that machine. The resources used by each task are considered to be equal to the *requested* amount for the first 5 min, then equal to the *average used* amount extracted from the traces. This strategy was adopted after careful analysis of the traces, without knowing any details about the system producing the traces (Borg). Recent publication of Borg details [36] confirms that our scheduling strategy is very similar to the real system. In the Google cluster, the requested resources are initially reserved (just like in our case), and at 5 min intervals the reservation is adjusted based on the real usage and a safety margin through a so-called *resource reclamation* mechanism.

The Scheduler implements a simple job scheduling mechanism. Each time a job is created by the Job Arrival entity, the scheduler inserts its tasks in the ready queue. For each task, the scheduler examines which execution nodes (if any) match the task constraints; the task is eventually sent to a suitable execution node. Due to the fact that the Scheduler does not know in real time the exact amount of free resources for all machines, it may happen that it sends a task to a machine that can not host it. In this case, the machine selects a task to interrupt (evict) and sends it back to the scheduler. Similar to the scheduling policies implemented by the Google cluster, we allow a task with higher priority to evict a running task with lower priority. The evicted task will be added back to the ready queue.

When the machine starts the execution of a task, it generates a future event: the termination event, based on the duration generated/read from the input data. The system does not differentiate between tasks that terminate normally or because they are killed or they fail; the only distinction is for evicted tasks, as explained previously. When the termination event will be handled, the scheduler will be notified by a message. Note that the duration is used only to generate the event, it is not used to make decisions. This is necessary in order to simulate a scenario in which the execution time of a task is not known a priori. In case the task is evicted, this event is deleted and will be recreated when the task will be restarted.

Finally, the Network entity is responsible for exchanging messages between the other active entities. In this way, it is possible to use a single gate to communicate with every other entity. Messages include notifications of new jobs arriving, tasks being submitted to a machine, machines reporting their status to the scheduler, etc.

Each message holds two different IDs: the sender and the receiver, and the network will be responsible to correctly route the messages by interfacing with the Omnet framework. This scenario reflects the real configuration of Google datacenter where there is a common shared network and the storage area is uniformly accessible from each machine. It was not possible to give a limit to the bandwidth while the latency of the channels is considered to be null. This does not affect the simulation since in Google clusters, the internal network does not seem to be a bottleneck. However it is possible to extend the Network entity in order to implement a latency and a maximal bandwidth between channels.

The Google traces contain information about both *exogenous* and *endogenous* events. Exogenous events are those originating outside the system, such as jobs and machines arrivals or job characteristics; endogenous events are those originating inside the system, such as jobs starting/finishing execution, failure events and similar.

In order to instantiate the model and perform simulations, several parameters concerning endogenous and exogenous events have to be provided as input. The implementation provides two input options:

- *Synthetic-trace-driven simulation* In this mode, the simulator is provided with distributions of the various job characteristics and event probabilities and inter-arrival times. For instance, one can specify distributions for number of tasks per job, required and used resources, priorities, and others. During simulation, these distributions are used by the Job Arrival and Machine Arrival entities to generate new Job entities and machine events, obtaining in this way a synthetic trace. Distributions can be specified in two ways. One is by providing CDFs extracted from real traces. We will demonstrate this case in Sect. 3.2, when we will extract distributions from the Google trace using *BiDAL* and we will perform simulations. The second option is to specify in a configuration file known distributions for the parameters. For instance, one can use a Gaussian distribution for resource utilization. Synthetic-trace-driven simulation is useful for exploring the behavior of the Google cluster under arbitrary conditions, e.g., under heavy load or massive failures, that may not occur in the traces; this is an example of “what-if analysis”.
- *Real-trace-driven simulation* In this mode, information regarding jobs and machine arrivals is contained in a file that is provided at the beginning of the simulation. This includes all properties of each incoming job as described by a trace, and the exact times when machines are added, removed or updated. The data is used by the Job Arrival and Machine Arrival entities to reproduce exactly the same workload during simulation. Trace-driven simulation is used to validate the model, since we expect the output of the simulation runs to match the Google cluster behavior observed in the traces. In Sect. 3.3 we show results from simulation using the Google traces.

3.2 Synthetic-trace-driven simulation

To generate realistic synthetic traces, we used *BiDAL* to extract distributions from the Google data to characterize the workload of the cluster and other relevant endogenous events. It is worth observing that the traces consist of over 2000 large CSV files

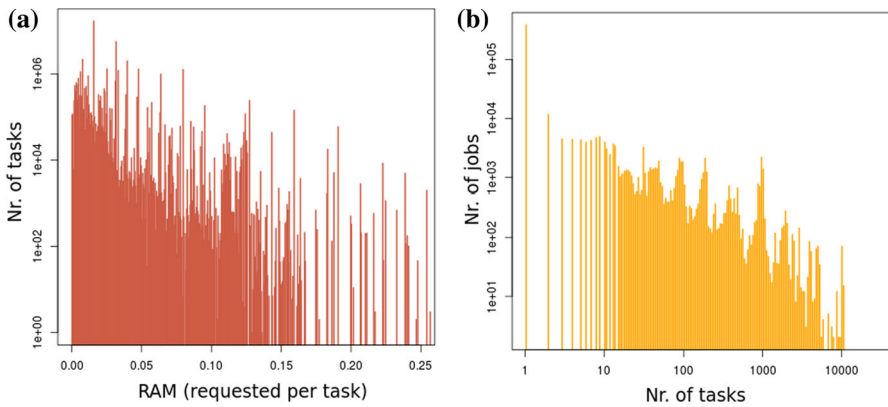


Fig. 5 Examples of distributions obtained with BiDAL. Panels: **a** RAM requested by tasks. Values are normalized by the maximum RAM available on a single node in the Google cluster; **b** number of tasks per job. These do not appear to follow any known distribution

containing records about job and task events, resources used by tasks, task constraints, and so on. In the following we first describe the distribution obtained, then we show simulation results.

Workload characterization of the Google cluster We extracted the arrival time distribution of each job, the distribution of the number of tasks per job, and the distributions of execution times of different types of tasks (e.g., jobs that successfully completed execution, jobs that are killed by the users, and so on). These distributions are used by the model to generate jobs into the system. Additionally, we analyzed the distribution of machines downtime and of the time instants when servers are added/removed from the pool.

Some of the results obtained with *BiDAL* are shown in the following figures (these are the actual plots that were produced by *BiDAL*). Figure 5a shows the the amount of RAM requested by tasks, while Fig. 5b shows the distribution of number of tasks per job.

To generate the graph in Fig. 5b, we first extracted the relevant information from the trace files. Job and task IDs were required, therefore we generated a new table, called *job_task_id*, from the *task_events.csv* files released by Google [39]. The query generation is automated by *BiDAL* which allows for simple selection of columns using the GUI. Since the *DISTINCT* clause is not yet implemented in *BiDAL*, we added it manually in the generated query. The final query used was:

```
SELECT DISTINCT V3 AS V1, V4 AS V2 FROM task_events
```

Here *V3* is the *job_id* column while *V4* represents the *task_id*. On the resulting *job_task_id* table, we execute another query to estimate how many tasks each job has, generating a new table called *tasks_per_job*:

```
SELECT V1 AS V1, COUNT(V2) AS V2 FROM job_task_id GROUP BY V1
```

Three *BiDAL* commands were used on the *tasks_per_job* table to generate the graph. The first extracts the second column (job id), the second filters out some uninteresting data and the third plots the result. The *BiDAL* commands used are shown in Table 2.

Table 2 Commands used to generate Fig. 5b

Command	Parameter type	Parameter value
get_column	Column number	2
filter	Condition	$t[[1]] < 11,000$
log_histogram	Column number, log step, log axis	1, 0.06, xy

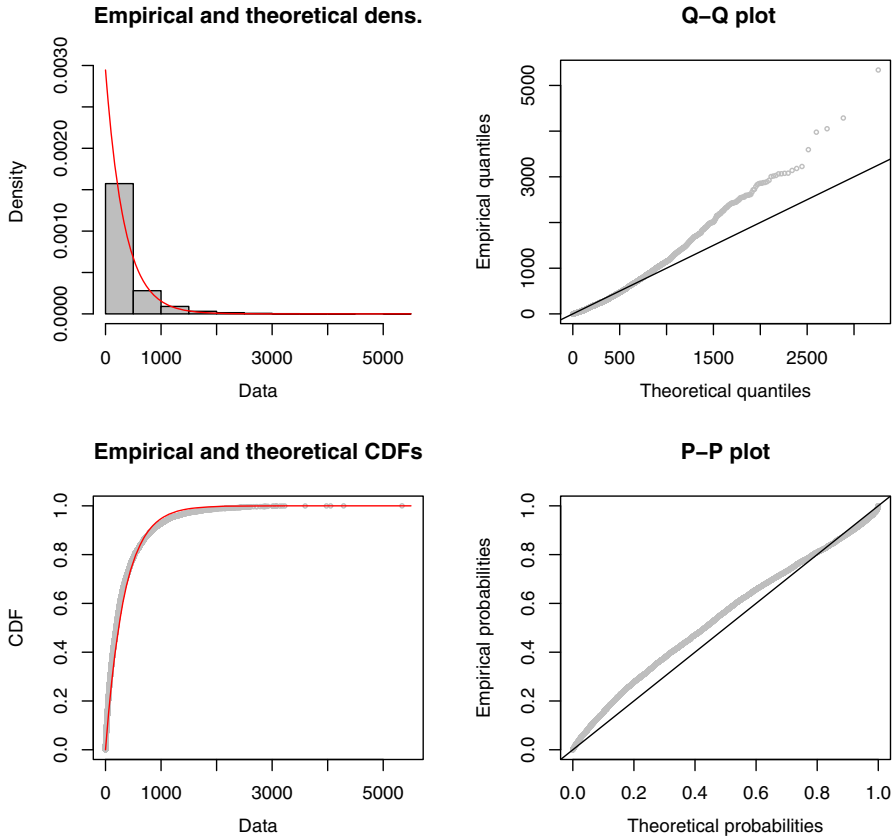


Fig. 6 Machine update inter-event times, fitted with an exponential distribution. The *left panels* show the density and cumulative distribution functions, with the *lines* representing exponential fitting and the *bars/circles* showing real data. The *right panels* show goodness of fit in *Q-Q* and *P-P* plots (*straight lines* show perfect fit)

The analysis was performed on a 2.7 GHz i7 quad core processor with 16 GB of RAM and a hard drive with simultaneous read/write speed of 60 MB/s. For the example above, importing the data was the most time consuming step, requiring 11 min to load 17 GB of data into the SQLite storage (the load time is determined for the most part by the disk speed). However, this step is required only once. The first SQL query took about 4 min to complete, while the second query and the *BiDAI* commands were almost instantaneous.

In Fig. 6 we fit the time between consecutive machine update events (i.e., events that indicate that a machine has changed its list of resources) with an exponential

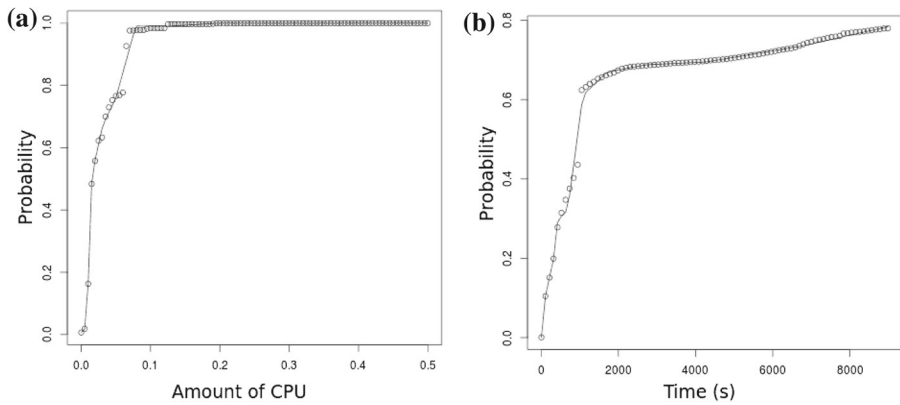


Fig. 7 Examples of CDFs fitted by sequences of splines, obtained with BiDAL. Panels: **a** CPU tasks requirements; **b** machine downtime. The *circles* represent the data, while the *lines* show the fitted splines. The CDFs are employed to produce synthetic traces to be used as input to our model

distribution. We use four standard plots for the goodness of fit: the probability density distribution, the cumulative distribution, the quantile–quantile (Q–Q) plot, and probability–probability (P–P) plot [14]. The P–P plot displays the values of the cumulative distribution function (CDF) of the data (empirical probabilities) versus the CDF of the fitted exponential distribution (theoretical probabilities). Specifically, for each value i of the inter-event time, the x -axis shows the percentage of values in the *theoretical exponential distribution* that fall below i while the y -axis shows the percentage of points in the *data* that fall below i . If the two values are equal, i.e. the entire plot follows the diagonal, then the fit between the data and theoretical distributions is good. The Q–Q plot, on the other hand, displays the quantiles of the data (empirical quantiles) versus those in the fitted exponential distribution (theoretical quantiles). Again, perfect fit means the Q–Q plot follows the diagonal, i.e. quantiles coincide. All plots show that the observed data is in good agreement with the fitted distribution.

Cumulative distribution functions have also been computed from the data and fitted with sequences of splines, in those cases where the density functions were too noisy to be fitted with a known distribution. For instance, Fig. 7a shows the distribution of CPU required by tasks while Fig. 7b shows machine downtime, both generated with BiDAL. Several other distributions that are listed in Table 3 were generated in a similar way to enable simulation of the Google cluster. Once distributions were generated, integration in the model was straightforward since BiDAL is able to generate C code related to the different distributions found. In our study, the distributions, hence the C code related to them, represent empirical CDFs. We extracted several other numerical parameters with BiDAL to be used by the model, as described in Table 4.

Job constraints were simplified in the synthetic traces compared to real data. For this purpose, we analyzed the traces and studied the influence of the constraints. We calculated the percentage of tasks with constraints and the mean satisfiability s_{c_i} of each constraint c_i as the average fraction of machines that satisfy c_i . To simulate the constraint system and assign the same mean satisfiability to each constraint, each machine is associated a numerical value x in an interval $I = [a, b]$. Each constraint

Table 3 Parameters extracted from the data and used as input for the model in the form of CDFs

Parameters extracted as CDFs
CPU required by tasks
Machine downtime
RAM required by tasks
Task priority
Duration of tasks that end normally
Duration of killed tasks
Tasks per job
Job inter-arrival time
Machine failure inter-arrival time
Machine CPU
Machine RAM

Table 4 Additional numerical parameters extracted from the data and used as input for the model

Numerical parameters
Probability of submitting tasks with different constraints
Probability that a machine satisfies a constraint
Amount of initial tasks running
Probability of submitting long running tasks (executing from the beginning until the end of the simulation)
Amount of RAM available on the machines
Probability that a task terminates normally or is killed

c_i is assigned a subinterval $I_{c_i} = [c, d] \subseteq I$ so that $\frac{d-c}{b-a} = s_{c_i}$. A machine satisfies a constraint c_i if $x \in I_{c_i}$. In this way, each constraint is satisfied with the same probability detected from the traces.

Simulation results using synthetic workload The parameters and distributions described above have been used to instantiate the model. We performed ten simulation runs and the results were analyzed in terms of number of running and completed tasks, the length of the ready queue and the number of evicted processes. It is important to note that all evaluation criteria used depend on the model logic, while input parameters are external to the model. The distribution of these values, compared to the original data, are shown in Fig. 8; Table 5 reports means and standard deviations along with the difference between the means of the distributions (real vs. simulated).

The number of tasks in execution, the length of the ready queue and finished tasks are on average similar to the real traces, indicating that the parameters used are fairly good approximations of the real data. However the distributions of these quantities produced by the simulator have a lower variability than those observed in the real data. This is probably due to the fact that resource usage for tasks is averaged over the entire length of the task, rather than being variable in time, as in the real system.




Fig. 8 Distribution of number of tasks in different categories for each 450 s time window, for the *synthetic-trace-driven simulation* compared to the *real data*. The y-axis shows the fraction of time windows with the corresponding number of tasks running, waiting, finished or evicted. Average behavior is similar between simulation and data for all categories except tasks evicted. However, variability is larger in the data for all cases

In terms of the number of evicted tasks, differences among average behaviors are much larger. The model tends to evict twice as many tasks as the real system. The mean of the simulation output still falls within a standard deviation from the mean of the real data; however, the simulation never generates low numbers of evicted jobs as are observed in the traces. This can be due, again, to the fact that the simulator is fed with average values for the resource usage and other parameters. Indeed, the same problem is observed, to a smaller extent, also in the real-trace-driven simulation described in the next section. Indeed, resource usage is averaged in real-trace-driven simulation as well.


The number of submitted tasks needs a separate discussion. This metric is different from the other ones because the number of submitted tasks is derived directly from the input distribution, and therefore does not depend on the model; in other words, this is derived from an *input* parameter, rather than the simulation output, so it shows how well *BiDAL* is capable of producing *accurate synthetic traces*. The number of

Table 5 Differences between means of the distributions from Fig. 8

	Running tasks	Ready tasks	Finished tasks	Evicted tasks	Submitted tasks
Mean value obtained in the simulation	136,037	5726	2317	2165	4323
Mean value in the data	126,375	5991	3271	1053	4540
Standard deviation obtained in the simulation	3116	1978	756	3482	8344
Standard deviation in the data	11,620	6274	2716	2376	4535
Mean difference	7 %	4 %	29 %	105 %	4 %

For most measures, averages are very similar. Larger differences are observed for finished and evicted tasks, with our system evicting more and finishing less jobs in each time window, compared to the real system

Fig. 9 Distribution of number of submitted tasks for the synthetic workload (*simulation*), compared to the real workload (*data*). The synthetic workload shows less variation than the real workload



submitted tasks depends on the distributions of the job inter-arrival time and of the number of tasks per job.

Figure 9 compares the distribution of the number of submitted tasks as seen during simulation and in the real data. The two distributions are very similar; the synthetic trace appears slightly more narrow than the real data, which partly explains why the simulation output has lower variability as well (see Fig. 8). Similar lower variability for synthetic data has also been reported by other authors [40]. Figure 9 also displays a few time intervals where a large number of tasks are submitted. This explains the larger standard deviation for synthetic data reported in Table 5. The mean error of 4 % is comparable to existing methods [41] that report differences from 2 to 36 % between real and synthetic logs, depending on the cluster, application structure and criterion used.

The results indicate that some fine tuning of the model is necessary to get more accurate results. First, the input distributions should better reflect the real data, especially for the arrival rate of tasks. To obtain wider distributions of the number of tasks in the different states, resource usage should be allowed to change over time (as happens in the real data). Furthermore, other system parameters, such as the resource usage limit, should be studied in more detail to get better fits.

3.3 Real-trace-driven simulation

In the real-trace-driven simulation we provide the simulation model with the real workload extracted from the traces. The purpose is to validate the model by comparing the simulation results with the real system behavior inferred from the traces.

The Google trace have been collected on a running system; therefore, some jobs were already in the queue, and others were being executed at the beginning of the trace. To properly account for these jobs, we bootstrap the simulation by inserting all the tasks already in execution at the beginning of the trace into the ready queue. These jobs are processed by the scheduler and assigned to machines. At the same time, new Job and Machine events are generated, according to the trace itself. It takes several minutes of wallclock time for the simulation to stabilize and reach a configuration similar to the system state at the beginning of the trace. This phase represents the initial transient and has been removed from the results. The model takes as input the events of the first 40 h of the original traces, with the first 5 h considered as part of the initial transient phase.

Running our simulation, we observed that all jobs were scheduled very quickly, with no evicted tasks. However, the Google trace contains many task evictions. The description of the Google data indicates that some machine resources are reserved by the scheduler for itself and for the operating system, so not all resources are available to tasks [27]. This reserved amount is however not specified. We can account for the unknown reserved resources by decreasing the amount of resources available to tasks within the model. We decided to decrease the amount of available memory to a fraction f_m of the total. After several simulations for fine tuning, the value $f_m = 0.489$ produced the best fit to the data. The accuracy of our simulation is highly sensitive to this parameter and small variations result in large differences. For instance, for values slightly different from 0.489, the number of jobs evicted during simulation is very different from the real traces. The value obtained for f_m may seem rather large, since it is unlikely that the scheduler reserves half the memory for itself. However, this accounts also for the slight difference in allocating resources in our model compared to the real system. In our case, we reserve exactly the used resources, while the Google cluster, within its resource reclamation mechanism described in Sect. 3.1, uses a safety margin which is not specified. Our chosen value $f_m = 0.489$ includes both the unknown resource reclamation margins and operating system reservation.

To assess the accuracy of the simulation results we perform a transient analysis, comparing the output of the simulator with the real data from the traces. Specifically, four metrics were considered: number of running tasks (Fig. 10a), number of completed tasks (Fig. 10b), number of waiting tasks (ready queue size, Fig. 10c) and number of evicted tasks (Fig. 10d). Comparison of the real and simulated values can bring important evidence whether the model is able to reproduce the behavior of the real Google cluster.

All plots show the time series extracted from the trace data (green lines) and those produced by our model (red lines), with the additional application of exponential smoothing (to both) to reduce fluctuations. The figures show a very good agreement between the simulation results and the actual data from the traces. This means that the model provides a good approximation of the Google cluster.

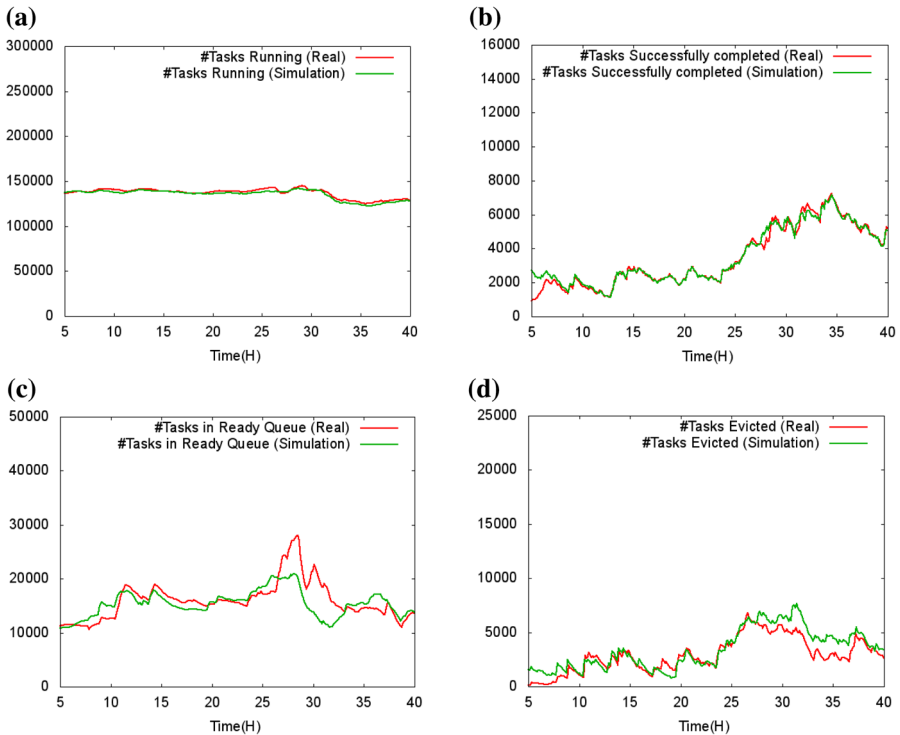


Fig. 10 Simulation and real data for four different metrics: **a** number of running tasks; **b** number of tasks completed; **c** number of tasks waiting; **d** number of tasks evicted. All show good agreement between the behavior of our model and that of the real system

We executed ten simulation runs; due to the fact that the model is deterministic (the only variation is in the choice of the machine where to execute a certain process), there are small differences across the runs. We report in Table 6 several statistics regarding the running, completed, waiting and evicted tasks. These results are collected at intervals of 450 s. It is clear that all measures are very close between real and simulated data. Standard deviations are in general slightly smaller for the simulation, showing that the model tends to flatten the data, as also seen previously for synthetic data [40].

4 Related work

With the public availability of the two cluster traces [39] generated by the Borg system at Google [36], numerous analyses of different aspects of the data have been reported. These provide general statistics about the workload and node state for such clusters [22, 26] and identify high levels of heterogeneity and dynamicity of the system, especially in comparison to grid workloads [10]. Heterogeneity at user level—large variations between workload submitted by the different users—is also observed [1]. Prediction is

Table 6 Statistics of four evaluation criteria at intervals of 450 s

Evaluation criterion	Running tasks	Completed tasks	Waiting tasks	Evicted tasks
Mean value obtained from the simulation	134,476	3671.3	15,400.6	3671.32
Mean value shown in the real traces	136,152	3654.6	15,893.9	2895.76
Standard deviation obtained from the simulation	6644.6	2375	3645.8	3336.9
Standard deviation shown in the real traces	6913.4	2452.3	5490.3	3111.6
Maximum error (absolute value)	4622	1974	9318	2639
Maximum error (in percentage w.r.t. the mean value)	3.40 %	56.00 %	59.00 %	92 %
Mean error (absolute value)	1858	246	1944	755
Mean error (in percentage w.r.t. the mean value)	0.01 %	7.00 %	12.20 %	26 %

attempted for job [15] and machine [32] failures and also for host load [11]. However, no unified tool for studying the different traces were introduced. *BiDAL* is one of the first such tools facilitating Big Data analysis of trace data, which underlines similar properties of the public Google traces as the previous studies. Other traces have been analyzed in the past [7, 8, 20], but again without a general-purpose tool available for further study.

BiDAL can be very useful in generating synthetic trace data. In general synthesizing traces involves two phases: characterizing the process by analyzing historical data and generation of new data. The aforementioned Google traces and log data from other sources have been successfully used for workload characterization. In terms of resource usage, classes of jobs and their prevalence can be used to characterize workloads and generate new ones [8, 23, 37], or real usage patterns can be replaced by the average utilization [41]. Placement constraints have also been synthesized using clustering for characterization [30]. Other traces have also been analyzed for the purpose of building synthetic workload data using hyperexponential distributions [40]. Our tool enables workload and cloud structure characterization through fitting of distributions that can be further used for trace synthesis. The analysis is not restricted to one particular aspect, but the flexibility of our tool allows the the user to decide what phenomenon to characterize and then simulate. Furthermore, by using empirical CDFs, our tool is not restrictive in terms of the data distributions.

Traces (either synthetic or the exact events) can be used for validation of various workload management algorithms. The Google trace has been used recently in [17] to evaluate consolidation strategies, in [4, 5] to validate over-committing (overbooking), in [42] to perform provisioning for heterogeneous systems and in [12] to investigate checkpointing algorithms. Again, data analysis is performed individually by the research groups and no specific tool was published. *BiDAL* is very suitable for extend-

ing these analyses to synthetic traces, to evaluate algorithms beyond the exact timeline of the Google dataset.

Recently, the Failure Trace Archive (FTA) has published a toolkit for analysis of failure trace data [19]. This toolkit is implemented in Matlab and enables analysis of traces from the FTA repository, which consists of about 20 public traces. It is, to our knowledge, the only other tool for large scale trace data analysis. However, the analysis is only possible if traces are stored in the FTA format in a relational database, and is only available for traces containing failure information. *BiDAL* on the other hand provides two different storage options, including HDFS, with transfer among them transparent to the user, and is available for any trace data, regardless of what process it describes. Additionally, usage of FTA on new data requires publication of the data in their repository, while *BiDAL* can be used also for sensitive data that cannot be made public.

Although public tools for analysis of general trace data are scarce, several large corporations reported to have built in-house custom applications for analysis of logs. These are, in general, used for live monitoring of the system, and analyze in real time large amounts of data to provide visualization that help operators make administrative decisions. While Facebook use Scuba [2], mentioned before, Microsoft have developed the Autopilot system [18], which helps with the administration of their clusters. Autopilot has a component (Cockpit) that analyzes logs and provides real time statistics to operators. An example from Google is CPI2 [43] which monitors cycles per instruction (CPI) for running tasks to determine job performance interference; this helps in deciding task migration or throttling to maintain high performance of production jobs. All these tools are, however, not open, apply only to data of the corresponding company and sometimes require very large computational resources (e.g., Scuba). Our aim in this paper is to provide an open research tool that can be used also by smaller research groups that have more limited resources.

In terms of simulation, numerous modeling tools for computer systems have been introduced, ranging from queuing models to agent-based and other statistical models. The systems modeled range from clusters to grids, and more recently, to clouds and data centers [44]. CloudSim is a recent discrete event simulator that allows simulation of virtualized environments [6]. Google has recently released a lightweight simulator of Omega [29] that allows comparison of various scheduler architectures. More specialized simulators such as MRPerf have been designed for MapReduce environments [38]. In general, these simulators are used to analyze the behavior of different workload processing algorithms (e.g., schedulers) and different networking infrastructures. A comprehensive model is Green Data Centre Simulator (GDCCSim), a very detailed simulator that takes into account computing equipment and its layout, data center physical structure (such as raised floors), resource management and cooling strategies [16]. However the level of detail limits scalability of the system. Our simulator is more similar to the former examples and allows for large scale simulations of workload management (experiments with 12 k nodes).

5 Conclusions

In this paper we presented *BiDAL*, a framework that facilitates use of Big Data tools and techniques for analyzing large cluster traces. We discussed a case study where we successfully applied *BiDAL* to analyze Google trace data in order to derive workload parameters required by an event-based model of the cluster. Based on a modular architecture, *BiDAL* currently supports two storage backends based on SQLite and Hadoop, while other backends can be easily added. It uses a subset of SQL as a common query language that is automatically translated to the appropriate commands supported by each backend. Additionally, data analysis using R and Hadoop MapReduce is possible.

Analysis of the Google trace data consisted of extracting distributions of several relevant quantities, such as number of tasks per job, resource consumption by tasks, etc. These parameters were easily computed using our tool, showing how this facilitates Big Data analysis even to users less familiar with R or Hadoop.

The model was analyzed under two scenarios. In the first scenario we performed a real-trace-driven simulation, where the input data were taken directly from the real traces. The results produced by the simulation in this scenario are in good agreement with the real data. The fidelity was obtained by fine tuning the model in terms of available resources, which accounts for unknown policies in the real cluster. Our analysis showed that reducing available memory to 48.9 % produces a good estimate of the actual data. In the second scenario we used *BiDAL* to produce synthetic inputs by fitting the real data to derive their distribution. In this scenario the average values of the output parameters are in good agreement with the average values observed in the traces; however, the general shape of the output distributions are quite different. These differences could be due to over-simplifications of the model, such as the fact that only average values for resource consumption are used, or that the task arrival process is not modeled accurately. Improvements of the accuracy of the model will be the subject of future work as well as implementation and comparison with existing models and workload synthesizers.

At the moment, *BiDAL* can be used for pre-processing and initial data exploration; however, in the future we plan to add new commands to support machine learning tools for predicting abnormal behavior from log data. This could provide new steps towards achieving self-* properties for large scale computing infrastructures in the spirit of autonomic computing.

In its current implementation, *BiDAL* is useful for batch analysis of historical log data, which is important for modeling and initial training of machine learning algorithms. However, live log data analysis is also of interest, so we are investigating the addition of an interface to streaming data sources to our platform. Future work also includes implementation of other storage systems, especially to include non-relational models. Improvement of the GUI and general user experience will also be pursued.

Source code availability The source code for *BiDAL* and the Google cluster simulator is available under the terms of the GNU General Public License (GPL) on GitHub at <https://github.com/alkida/bidal> and <https://github.com/alkida/clustersimulator>, respectively.

References

1. Abdul-Rahman OA, Aida K (2014) Towards understanding the usage behavior of Google cloud users: the mice and elephants phenomenon. In: 2014 IEEE 6th international conference on cloud computing technology and science (CloudCom). IEEE, pp 272–277. doi:[10.1109/CloudCom.2014.75](https://doi.org/10.1109/CloudCom.2014.75)
2. Abraham L, Allen J, Barykin O, Borkar V, Chopra B, Gerec C, Merl D, Metzler J, Reiss D, Subramanian S, Wiener JL, Zed O (2013) Scuba: diving into data at facebook. Proc VLDB Endow 6(11):1057–1067. doi:[10.14778/2536222.2536231](https://doi.org/10.14778/2536222.2536231)
3. Armbrust M, Xin RS, Lian C, Huai Y, Liu D, Bradley JK, Meng X, Kaftan T, Franklin MJ, Ghodsi A, Zaharia M (2015) Spark SQL: relational data processing in spark. In: Proceedings of the 2015 ACM SIGMOD international conference on management of data (SIGMOD'15). ACM, New York, pp 1383–1394. doi:[10.1145/2723372.2742797](https://doi.org/10.1145/2723372.2742797)
4. Breitgand D, Dubitzky Z, Epstein A, Feder O, Glikson A, Shapira I, Toffetti G (2014) An adaptive utilization accelerator for virtualized environments. In: 2014 IEEE international conference on cloud engineering (IC2E). IEEE, pp 165–174. doi:[10.1109/IC2E.2014.63](https://doi.org/10.1109/IC2E.2014.63)
5. Caglar F, Gokhale A (2014) iOverbook: intelligent resource-overbooking to support soft real-time applications in the cloud. In: Proceedings of the 2014 IEEE international conference on cloud computing (CLOUD'14). IEEE Computer Society, Washington, DC, pp 538–545. doi:[10.1109/CLOUD.2014.78](https://doi.org/10.1109/CLOUD.2014.78)
6. Calheiros RN, Ranjan R, Beloglazov A, De Rose CAF, Buyya R (2011) Cloudsim: a toolkit for modeling and simulation of cloud computing environments and evaluation of resource provisioning algorithms. Softw Pract Exp 41(1):23–50. doi:[10.1002/spe.995](https://doi.org/10.1002/spe.995)
7. Chen Y, Alspaugh S, Katz RH (2012) Design insights for MapReduce from diverse production workloads. Tech. Rep. UCB/EECS-2012-17, EECS Department, University of California, Berkeley. <http://www.eecs.berkeley.edu/Pubs/TechRpts/2012/EECS-2012-17.html>. Accessed Dec 2015
8. Chen Y, Ganapathi A, Griffith R, Katz RH (2011) The case for evaluating MapReduce performance using workload suites. In: 2011 IEEE 19th annual international symposium on modelling, analysis, and simulation of computer and telecommunication systems, pp 390–399. doi:[10.1109/MASCOTS.2011.12](https://doi.org/10.1109/MASCOTS.2011.12)
9. Dean J, Ghemawat S (2010) Mapreduce: a flexible data processing tool. Commun ACM 53(1):72–77. doi:[10.1145/1629175.1629198](https://doi.org/10.1145/1629175.1629198)
10. Di S, Kondo D, Cirne W (2012) Characterization and comparison of Google cloud load versus grids. In: 2012 IEEE international conference on cluster computing (CLUSTER), Beijing, pp 230–238. doi:[10.1109/CLUSTER.2012.35](https://doi.org/10.1109/CLUSTER.2012.35)
11. Di S, Kondo D, Cirne W (2012) Host load prediction in a Google compute cloud with a bayesian model. In: Proceedings of the international conference on high performance computing, networking, storage and analysis. IEEE Computer Society Press, USA, pp 1–11. doi:[10.1109/SC.2012.68](https://doi.org/10.1109/SC.2012.68)
12. Di S, Robert Y, Vivien F, Kondo D, Wang CL, Cappello F (2013) Optimization of cloud task processing with checkpoint-restart mechanism. In: 2013 international conference for high performance computing, networking, storage and analysis (SC). IEEE, pp 1–12
13. Gamma E, Helm R, Johnson R, Vlissides J (1994) Design patterns: elements of reusable object-oriented software. Addison-Wesley Professional, Boston
14. Gibbons JD, Chakraborti S (2010) Nonparametric statistical inference. Chapman and Hall/CRC, London
15. Guan Q, Fu S (2013) Adaptive anomaly identification by exploring metric subspace in cloud computing infrastructures. In: Proceedings of the 2013 IEEE 32nd international symposium on reliable distributed systems (SRDS'13). IEEE Computer Society, Washington, DC, pp 205–214. doi:[10.1109/SRDS.2013.29](https://doi.org/10.1109/SRDS.2013.29)
16. Gupta SKS, Banerjee A, Abbasi Z, Varsamopoulos G, Jonas M, Ferguson J, Gilbert RR, Mukherjee T (2014) Gdcsim: a simulator for green data center design and analysis. ACM Trans Model Comput Simul 24(1):3:1–3:27. doi:[10.1145/2553083](https://doi.org/10.1145/2553083)
17. Iglesias JO, Murphy L, De Cauwer M, Mehta D, O'Sullivan B (2014) A methodology for online consolidation of tasks through more accurate resource estimations. In: Proceedings of the 2014 IEEE/ACM 7th international conference on utility and cloud computing (UCC'14). IEEE Computer Society, Washington, DC, pp 89–98. doi:[10.1109/UCC.2014.17](https://doi.org/10.1109/UCC.2014.17)
18. Isard M (2007) Autopilot: automatic data center management. SIGOPS Oper Syst Rev 41(2):60–67. doi:[10.1145/1243418.1243426](https://doi.org/10.1145/1243418.1243426)

19. Javadi B, Kondo D, Iosup A, Epema D (2013) The failure trace archive: enabling the comparison of failure measurements and models of distributed systems. *J Parallel Distrib Comput* 73(8):1208–1223. doi:[10.1016/j.jpdc.2013.04.002](https://doi.org/10.1016/j.jpdc.2013.04.002)
20. Kavulya S, Tan J, Gandhi R, Narasimhan P (2010) An analysis of traces from a production MapReduce cluster. In: *Proceedings of the 2010 10th IEEE/ACM international conference on cluster, cloud and grid computing (CCGRID'10)*. IEEE Computer Society, Washington, DC, pp 94–103. doi:[10.1109/CCGRID.2010.112](https://doi.org/10.1109/CCGRID.2010.112)
21. Kornacker M, Behm A, Bittorf V, Bobrovitsky T, Ching C, Choi A, Erickson J, Grund M, Hecht D, Jacobs M, Joshi I, Kuff L, Kumar D, Leblang A, Li N, Pandis I, Robinson H, Rorke D, Rus S, Russell J, Tsirogiannis D, Wanderman-Milne S, Yoder M (2015) Impala: a modern, open-source SQL engine for Hadoop. In: *CIDR 2015, seventh biennial conference on innovative data systems research, Asilomar*
22. Liu Z, Cho S (2012) Characterizing machines and workloads on a Google cluster. In: *2012 41st international conference on parallel processing workshops (ICPPW)*, pp 397–403. doi:[10.1109/ICPPW.2012.57](https://doi.org/10.1109/ICPPW.2012.57)
23. Mishra AK, Hellerstein JL, Cirne W, Das CR (2010) Towards characterizing cloud backend workloads: insights from Google compute clusters. *SIGMETRICS Perform Eval Rev* 37(4):34–41. doi:[10.1145/1773394.1773400](https://doi.org/10.1145/1773394.1773400)
24. Olston C, Reed B, Srivastava U, Kumar R, Tomkins A (2008) Pig latin: a not-so-foreign language for data processing. In: *Proceedings of the 2008 ACM SIGMOD international conference on management of data (SIGMOD'08)*. ACM, New York, pp 1099–1110. doi:[10.1145/1376616.1376726](https://doi.org/10.1145/1376616.1376726)
25. R Development Core Team (2008) R: a language and environment for statistical computing. R Foundation for Statistical Computing, Vienna. <http://www.R-project.org>. Accessed Dec 2015
26. Reiss C, Tumanov A, Ganger GR, Katz RH, Kozuch MA (2012) Heterogeneity and dynamicity of clouds at scale: Google trace analysis. In: *Proceedings of the third ACM symposium on cloud computing (SoCC'12)*. ACM, New York, pp 7:1–7:13. doi:[10.1145/2391229.2391236](https://doi.org/10.1145/2391229.2391236)
27. Reiss C, Wilkes J, Hellerstein JL (2011) Google cluster-usage traces: format + schema. Technical report, Google Inc., Mountain View. <http://code.google.com/p/googleclusterdata/wiki/TraceVersion2>. Accessed 20 March 2012
28. Salfner F, Lenk M, Malek M (2010) A survey of online failure prediction methods. *ACM Comput Surv* 42(3):10:1–10:42. doi:[10.1145/1670679.1670680](https://doi.org/10.1145/1670679.1670680)
29. Schwarzkopf M, Konwinski A, Abd-El-Malek M, Wilkes J (2013) Omega: flexible, scalable schedulers for large compute clusters. In: *Proceedings of the 8th ACM European conference on computer systems (EuroSys'13)*. ACM, New York, pp 351–364. doi:[10.1145/2465351.2465386](https://doi.org/10.1145/2465351.2465386)
30. Sharma B, Chudnovsky V, Hellerstein JL, Rifaat R, Das CR (2011) Modeling and synthesizing task placement constraints in Google compute clusters. In: *Proceedings of the 2nd ACM symposium on cloud computing (SOCC'11)*. ACM, New York, pp 3:1–3:14. doi:[10.1145/2038916.2038919](https://doi.org/10.1145/2038916.2038919)
31. Shvachko K, Kuang H, Radia S, Chansler R (2010) The Hadoop distributed file system. In: *Proceedings of the 2010 IEEE 26th symposium on mass storage systems and technologies (MSST'10)*. IEEE Computer Society, USA, pp 1–10. doi:[10.1109/MSST.2010.5496972](https://doi.org/10.1109/MSST.2010.5496972)
32. Sirbu A, Babaoğlu O (2015) Towards data-driven autonomies in data centers. In: *IEEE international conference on cloud and autonomic computing (ICCAC)*. IEEE
33. Thusoo A, Sarma JS, Jain N, Shao Z, Chakka P, Zhang N, Anthony S, Liu H, Murthy R (2010) Hive—a petabyte scale data warehouse using Hadoop. In: *Proceedings of the 26th international conference on data engineering (ICDE)*, Long Beach, pp 996–1005. doi:[10.1109/ICDE.2010.5447738](https://doi.org/10.1109/ICDE.2010.5447738)
34. Thusoo A, Shao Z, Anthony S, Borthakur D, Jain N, Sen Sarma J, Murthy R, Liu H (2010) Data warehousing and analytics infrastructure at facebook. In: *Proceedings of the 2010 ACM SIGMOD international conference on management of data (SIGMOD'10)*. ACM, New York, pp 1013–1020. doi:[10.1145/1807167.1807278](https://doi.org/10.1145/1807167.1807278)
35. Varga A et al (2001) The OMNeT++ discrete event simulation system. In: *Proceedings of the European simulation multiconference (ESM'01)*, Prague
36. Verma A, Pedrosa L, Korupolu M, Oppenheimer D, Tune E, Wilkes J (2015) Large-scale cluster management at Google with borg. In: *Proceedings of the tenth European conference on computer systems (EuroSys'15)*. ACM, New York, pp 18:1–18:17. doi:[10.1145/2741948.2741964](https://doi.org/10.1145/2741948.2741964)
37. Wang G, Butt AR, Monti H, Gupta K (2011) Towards synthesizing realistic workload traces for studying the Hadoop ecosystem. In: *Proceedings of the 2011 IEEE 19th annual international symposium on modelling, analysis, and simulation of computer and telecommunication systems (MASCOTS'11)*. IEEE Computer Society, Washington, DC, pp 400–408. doi:[10.1109/MASCOTS.2011.59](https://doi.org/10.1109/MASCOTS.2011.59)

38. Wang G, Butt AR, Pandey P, Gupta K (2009) A simulation approach to evaluating design decisions in MapReduce setups. In: IEEE international symposium on modeling, analysis simulation of computer and telecommunication systems (MASCOTS'09), pp 1–11 (2009). doi:[10.1109/MASCOT.2009.5366973](https://doi.org/10.1109/MASCOT.2009.5366973)
39. Wilkes J (2011) More Google cluster data. Google research blog. <http://googleresearch.blogspot.com/2011/11/more-google-cluster-data.html>. Accessed Dec 2015
40. Wolski R, Brevik J (2014) Using parametric models to represent private cloud workloads. *IEEE Trans Serv Comput* 7(4):714–725. doi:[10.1109/TSC.2013.48](https://doi.org/10.1109/TSC.2013.48)
41. Zhang Q, Hellerstein JL, Boutaba R (2011) Characterizing task usage shapes in Google's compute clusters. In: Proceedings of the 5th international workshop on large scale distributed systems and middleware
42. Zhang Q, Zhani MF, Boutaba R, Hellerstein JL (2014) Dynamic heterogeneity-aware resource provisioning in the cloud. *IEEE Trans Cloud Comput* 2(1):14–28. doi:[10.1109/TCC.2014.2306427](https://doi.org/10.1109/TCC.2014.2306427)
43. Zhang X, Tune E, Hagmann R, Jnagal R, Gokhale V, Wilkes J (2013) CPI2: CPU performance isolation for shared compute clusters. In: Proceedings of the 8th ACM European conference on computer systems (EuroSys'13). ACM, New York, pp 379–391. doi:[10.1145/2465351.2465388](https://doi.org/10.1145/2465351.2465388)
44. Zhao W, Peng Y, Xie F, Dai Z (2012) Modeling and simulation of cloud computing: a review. In: 2012 IEEE Asia Pacific cloud computing congress (APCloudCC), pp 20–24. doi:[10.1109/APCloudCC.2012.6486505](https://doi.org/10.1109/APCloudCC.2012.6486505)