

# Medusa: An Efficient Cloud Fault-Tolerant MapReduce

Pedro A. R. S. Costa<sup>\*</sup>, Xiao Bai<sup>†</sup>, Fernando M. V. Ramos<sup>\*</sup>, Miguel Correia<sup>‡</sup>

LaSIGE, Faculdade de Ciências, Universidade de Lisboa, Portugal<sup>\*</sup>, Yahoo Labs, Sunnyvale, USA<sup>†</sup>,

INESC-ID, Instituto Superior Técnico, Universidade de Lisboa, Portugal<sup>‡</sup>

palcosta@fc.ul.pt, xbai@yahoo-inc.com, fvramos@ciencias.ulisboa.pt, miguel.p.correia@tecnico.ulisboa.pt

**Abstract**—Applications such as web search and social networking have been moving from centralized to decentralized cloud architectures to improve their scalability. MapReduce, a programming framework for processing large amounts of data using thousands of machines in a single cloud, also needs to be scaled out to multiple clouds to adapt to this evolution. The challenge of building a multi-cloud distributed architecture is substantial. Notwithstanding, the ability to deal with the new types of faults introduced by such setting, such as the outage of a whole datacenter or an arbitrary fault caused by a malicious cloud insider, increases the endeavor considerably.

In this paper we propose Medusa, a platform that allows MapReduce computations to scale out to multiple clouds and tolerate several types of faults. Our solution fulfills four objectives. First, it is transparent to the user, who writes her typical MapReduce application without modification. Second, it does not require any modification to the widely used Hadoop framework. Third, the proposed system goes well beyond the fault-tolerance offered by MapReduce to tolerate arbitrary faults, cloud outages, and even malicious faults caused by corrupt cloud insiders. Fourth, it achieves this increased level of fault tolerance at reasonable cost. We performed an extensive experimental evaluation in the ExoGENI testbed, demonstrating that our solution significantly reduces execution time when compared to traditional methods that achieve the same level of resilience.

## I. INTRODUCTION

MapReduce [1] has been attracting a lot of interest as a convenient tool for processing massive datasets using computer clusters. MapReduce was developed and is widely used by Google, whereas one of its open source implementations, Apache Hadoop [2], is currently used by several other companies, such as Yahoo and Facebook.

Originally this framework targeted a single datacenter, so currently used implementations are not designed to work across multiple datacenters<sup>1</sup>. However, the last few years have witnessed the requirements for data- and compute-intensive analysis to grow significantly, increasing the need to scale-out computation across clouds. As an indication of that direction, Baeza-Yates et al. [3] have assessed the feasibility of building distributed Web search engines comprising geographically dispersed sites. The increasingly larger data sets used in bioinformatics have also led to the execution of multi-cloud computations for this type of applications [4]. In addition, for a growing number of applications the data is gathered and stored in different datacenters – i.e., it is inherently geo-distributed – but the analyses required are global. A well-known example includes the analysis of scientific data such as those originated from the Large Hadron Collider (LHC). The tens of petabytes of

data produced every year by this particle accelerator are stored in more than 140 datacenters distributed across 34 countries [5].

Acknowledging this trend, the research community has recently proposed MapReduce-based platforms that scale out to multiple clouds. Wang et al. proposed G-Hadoop [6], a new framework that enables large-scale distributed MapReduce computations distributed across multiple clusters. The core of this system is the replacement of Hadoop’s native file system, HDFS, with Gfarm, a distributed file system that scales to multiple clusters. Another instance of this trend is G-MR [7], another Hadoop-based framework that processes geo-distributed data across multiple datacenters. These examples are demonstrative of the challenge to build multi-cloud distributed architectures. Notwithstanding, the ability to deal with – and to tolerate – the new types of faults that are introduced by this setting makes building such system significantly harder.

Indeed, at scales of thousands of computers, switches, routers, power units and other components, failures are frequent. Therefore, both the original MapReduce and Hadoop use two mechanisms to tolerate such faults: they monitor the execution of tasks and reinitialize them in case of failure (thus tolerating crash faults); and they add checksums to files that contain data to detect file corruptions [2], [8]. Despite the use of such fault tolerance mechanisms, these MapReduce implementations do not consider three harder to tolerate failure modes. First, some hardware faults may lead to the *corruption of the processing*, leading to wrong outputs. For instance, data may be corrupted while stored in DRAM or due to core chipset errors, two issues recently found to be reasonably frequent [9], [10]. This problem is, naturally, amplified when computations are scaled out to multiple datacenters. Second, *malicious attacks* perpetrated by corrupt cloud insiders (or by external hackers that attack a specific cloud) can also cause corruption of the processing. The original MapReduce fault tolerance mechanisms cannot deal with such arbitrary [11] or malicious faults [12]. Third, cloud outages may lead to the *unavailability of MapReduce instances and their data*. Experience shows that these events are also frequent, with cases of unavailability of hours to days in services like Windows Azure, Google Drive or Amazon EC2, to name just a few. Again, current MapReduce systems cannot deal with cloud outage as they are restricted to work in a single datacenter.

To deal with such faults, one needs to add redundancy to the computation. As both malicious faults and cloud outage can impair a complete cloud, handling them involves resorting to more than one cloud. The terms *cloud federation* [13] and *cloud-of-clouds* [14] have been used recently to denote such

<sup>1</sup>We use the terms *cloud*, *cluster* and *datacenter* interchangeably.

virtual environments composed of multiple clouds. We explore this idea to replicate MapReduce jobs in different clouds to avoid their incorrectness or unavailability due to the three kinds of faults explained before.

Tolerating cloud faults through replication is commonly considered expensive as one expects these faults to be rare. The reality contradicts this expectation: cloud outages are becoming increasingly common [15], [16]. More importantly, for very critical applications such as bioinformatics or finance, any kind of errors and unavailability are unacceptable. For example, a malicious insider in a cloud that hosts an epidemiological surveillance system that tampers with the diagnosis of patients may lead to disastrous consequences. This particular problem is, today, a significant concern. A recent report from the Cloud Security Alliance states malicious insiders as one of the top threats in cloud computing [17], and alarming instances of this problem have recently occurred in companies such as Google [18]. Moreover, temporary unavailability of the financial system in one cloud (due to cloud outage) may dramatically influence the overall investment decision and cause huge financial loss. Motivated by these scenarios, we believe the replication cost to be acceptable for such critical applications, in order to guarantee that rare faults with devastating consequences do not occur. Interestingly, cloud providers seem to share this concern: Amazon S3 recently launched Cross-Region Replication to automatically replicate data across regions [19].

#### A. Our proposal

We propose a novel approach that allows MapReduce to scale out to multiple clouds to tolerate arbitrary and malicious faults, as well as cloud outages, for critical applications. As per above, the use of multiple clouds for MapReduce is not in itself new. The novelty of this work arises from the use of a multi-cloud environment to not only parallelize computation, but also to transparently tolerate different types of faults at the minimum cost. Our solution addresses several non-trivial challenges for this purpose. First, it aims to be a transparent solution for the user. The MapReduce API is not changed, and the user simply writes her typical MapReduce application without modification. Second, it does not require any modification to the Hadoop framework. Third, it tolerates not only crash faults, as the original MapReduce, but also arbitrary faults, cloud outages, and malicious faults caused by corrupt cloud insiders. Fourth, it achieves this level of fault tolerance at the minimum replication cost and guaranteeing acceptable performance.

Our approach relies on a *proxy*, Medusa, that runs in the client and that interacts with (unchanged) MapReduce runtimes in different clouds to tolerate the three kinds of faults above. The basic idea is to replicate each MapReduce job in more than one cloud and to compare the outputs of the replicated jobs to tolerate faults. The challenge is to perform this efficiently, and doing so without changing the framework and keeping the whole process transparent to the user. Achieving efficiency requires (i) replicating each job the minimum number of times; and (ii) assigning each replicated job to the cloud that ensures the best performance. To this end, instead of replicating each job at least  $2f + 1$  times to ensure a majority of correct results and tolerate  $f$  faults, as is common [20], [21], our approach is

crafted to run only  $f + 1$  replicas for each job when there is no fault, and  $2f + 1$  replicas when there are  $f$  faults.

Importantly, our approach can tolerate not only  $f$  but any number of faulty replicas or clouds as long as no more than  $f$  faulty replicas return the same wrong output. This includes the possibility that up to  $f$  clouds maliciously collude and the system remains able to reach a correct output. Moreover, we introduce a novel scheduling algorithm that takes into account the heterogeneity of the individual clouds to schedule the replicated jobs among them in order to reduce data communication and job completion time.

We performed an extensive experimental evaluation of our approach in a real testbed (ExoGENI). The results demonstrate that our solution significantly reduces the execution time when compared to traditional methods that achieve the same level of resilience. As an example, in certain scenarios we achieve a gain of up to 3 in efficiency when compared with a conventional round-robin approach that tolerates cloud faults.

In summary, our proposal is practical and transparent to the user by only involving a new software module running in the client (Medusa), not requiring any modification to the MapReduce framework or to user applications. Simultaneously, it tolerates arbitrary and malicious faults, and cloud outages, efficiently.

#### B. Outline

The remainder of the paper is organized as follows. In the next section we describe the system model and define the problem. Then, in Section III we present the detailed design of Medusa, the cloud fault-tolerant MapReduce system we propose. Section IV reports on the experimental evaluation. In Section V we discuss related work and, finally, we conclude this paper in Section VI.

## II. PROBLEM STATEMENT

To cope with the unprecedented data growth, applications such as web search, social networking and bioinformatic applications can be distributed in geographically-distant datacenters to leverage data locality and improve data processing efficiency [3], [4]. In these cases, multiple MapReduce clusters are set up in distributed datacenters. Each cluster collects and processes the data that is close to its datacenter. The final data processing results are obtained by aggregating their respective outputs in an aggregation MapReduce job. In this work, we aim at providing cloud fault tolerance while minimizing the overall execution time for the MapReduce computations running in such multi-cloud systems.

#### A. Hadoop MRv2

MapReduce is a programming model for processing large-scale datasets with parallel and distributed algorithms in computer clusters. We focus in this work on Hadoop, which is an open source implementation of the original MapReduce. The Apache Hadoop platform includes the Hadoop kernel, Hadoop MapReduce, and Hadoop Distributed File System (HDFS).

Fig. 1 depicts the simplified execution process of MapReduce 2.0<sup>2</sup>. Specifically, a client requests the execution of a job from

<sup>2</sup>Hadoop MapReduce has undergone a complete overhaul in Hadoop 0.23 and is now designated MRv2, or YARN.




Fig. 1: Simplified MapReduce execution process.

the *resource manager* (previously called job tracker). The role of the resource manager is to define which (map and reduce) tasks will be executed in which server. The execution of a (map or reduce) task in a server is managed by the *node manager* (previously called task tracker).

The Hadoop Distributed File System (HDFS) is the main distributed storage used by MapReduce applications. The input files for a MapReduce job typically reside in HDFS and are divided into logical blocks called splits. Each split is executed by a map task that produces an intermediate output in the form of key-value pairs. The content of the intermediate outputs is sorted by key and then consumed by the reduce tasks. The reduce tasks produce the final output.

Hadoop tolerates faults by (i) monitoring and restarting tasks when servers, node managers or the tasks themselves crash through the resource manager; and (ii) adding checksums to the files in HDFS to detect data corruption in disks. However, these mechanisms only work in a single cloud. Additional techniques are needed to tolerate faults for MapReduce jobs running in multi-cloud systems.

### B. System model

In a multi-cloud system, the MapReduce job runs in a federation of *clouds*. Each cloud has an HDFS instance to store the initial inputs and final outputs of the jobs running in that cloud. The entire data to be processed by the job is distributed across the clouds in the system, *i.e.*, each cloud has a subset of the data stored in its HDFS. The data can be either collected by the cloud itself or assigned by some external processes, but we ignore this detail as it is application dependent and is orthogonal to the MapReduce execution.

The MapReduce job is composed of a set of distributed *processes* (Fig. 2). The *client* submits the job through the process *proxy*, Medusa, that controls the execution of the job in all the clouds. Each cloud first runs its own MapReduce instance to process the data it has: the resource manager controls the execution of the part of job assigned to that cloud; and the node manager in each server runs the map and reduce tasks assigned to that server. Finally, the proxy is in charge of assigning the outputs from these clouds to an aggregation MapReduce job in one of the clouds to obtain the final output for this job.

The messages between the proxy and the clouds are mediated by a message queuing service (MQ), which uses reliable channels so that no messages are lost, duplicated or corrupted. In practice, this is provided by establishing TCP/IP connections. A message is only lost if the cloud is unreachable.




Fig. 2: MapReduce in a multi-cloud system.

The system is asynchronous, *i.e.*, we make no assumptions about bounds on processing and communication delays in each stage of job execution (Section III-A). We assume the use of authenticated reliable channels for communication.

We assume the existence of one correct dataset in each cloud at the beginning of job execution. We further assume the existence of a collision-resistant hash function, *i.e.*, a hash function for which it is impossible to let two different inputs produce the same output (*e.g.*, SHA-256). We use the digests produced by this function to verify the integrity of the data replicated between clouds, and to validate the correctness of the job outputs.

The proxy is the key component in the multi-cloud system for tolerating cloud faults. We focus, in this work, on the design of the proxy as well as on a scheduling mechanism that guarantees good system performance.

### C. Fault model

We say that a process is *correct* if it follows the algorithm, otherwise we say that it is *faulty*. We assume that the client is correct, because it is the party interested in getting a correct job output. We also assume that the proxy is correct, similarly to the Hadoop assumption that the resource manager does not fail. The proxy is responsible for scheduling jobs among clouds to tolerate cloud faults and cloud outages.

Resource managers and node managers can fail arbitrarily: they can return wrong results (*e.g.*, processing corruption) or even stop (*e.g.*, cloud outage). Malicious processes can also produce wrong results. There is no limit on the number of faulty components of these kinds. However, we assume that no more than  $f$  resource managers – *i.e.*,  $f$  clouds – return to the proxy identical wrong outputs of the execution of a job. This assumption allows the proxy to know that a result is correct by getting  $f + 1$  identical results. The minimum number of clouds necessary to ensure termination is  $2f + 1$ , as there may be  $f$  faulty clouds.

### D. Problem formulation

We aim at tolerating (i) arbitrary and malicious faults, and (ii) cloud outages, when running MapReduce jobs in multi-cloud systems. To tolerate  $f$  faults, a basic approach is to create  $2f + 1$

replicas of each job (*i.e.*, the job running in each cloud and the aggregation job), spread them in  $2f + 1$  different clouds, and compare the  $2f + 1$  outputs of each job. If at least  $f + 1$  outputs are identical, their corresponding MapReduce jobs are correct and the identical output is the correct output of this job.

This basic approach has two major problems. First, it is expensive in terms of computation, communication, and storage. Even if there is no fault, each job is executed  $2f + 1$  times. This requires replicating the data initially hosted by each cloud to  $f + 1$  other clouds, which can be expensive in geographically distributed clouds. The same data also have to be stored in the HDFS of  $2f + 1$  clouds. In fact, if there is no fault, executing each job  $f + 1$  times is enough (there is no need for additional computations). Second, the basic approach does not take into account the difference among clouds for data replication and data processing. Intuitively, the data initially hosted by a cloud should be replicated to other clouds with which the original cloud has a high pair-wise bandwidth and, simultaneously, has high computational power. This would ensure the efficiency of the entire job execution in the multi-cloud system.

Therefore, our objective in this work is to design a MapReduce proxy that ensures the MapReduce job running in multiple distributed clouds to tolerate cloud faults while (*i*) minimizing the amount of data replication and processing; and (*ii*) ensuring efficient completion of the entire MapReduce job.

### III. MEDUSA: A CLOUD FAULT-TOLERANT MR

As mentioned in Section II-B, a full job execution in a multi-cloud system is comprised of two phases. The first phase runs a *vanilla MapReduce job* in each cloud that holds a subset of data initially owned by that cloud. The second phase runs a *global MapReduce job* that aggregates the outputs from all clouds to generate the final results. To tolerate arbitrary, malicious faults, and cloud outage, the MapReduce jobs in each phase need to be replicated to other clouds for ensuring the existence of  $f + 1$  identical outputs and thus the correctness of the results.

We propose in this paper a MapReduce proxy that works as a middleware in a multi-cloud system (*i.e.*, a federation of clouds). We refer to this proxy as *Medusa*<sup>3</sup>.

As explained, we assume that no more than  $f$  replicas of a job return to the proxy identical wrong outputs from their executions. This assumption allows the proxy to know that a result is correct by getting  $f + 1$  identical results from different clouds / resource managers. Given the expected low probability of arbitrary and malicious faults, it is too expensive to always execute  $2f + 1$  replicas of a job as is done in typical approaches. Therefore, instead of replicating each job  $2f + 1$  times, Medusa first replicates each job  $f + 1$  times in  $f + 1$  different clouds (*i.e.*, once in the cloud where the data initially are and  $f$  times in other clouds that do not have the corresponding data initially). If the executions of these replicas do not produce identical outputs, one more replicated job is launched in a different cloud until the proxy gets  $f + 1$  identical outputs. This *deferred job execution* avoids the redundant data transmission, storage and processing when no fault happens. As per the point above, each

<sup>3</sup>Medusa is the mythological figure that has living snakes in place of hair. Metaphorically, the connections of our proxy to the clouds are the snakes, and these follow orders given by Medusa's brain (our proxy).

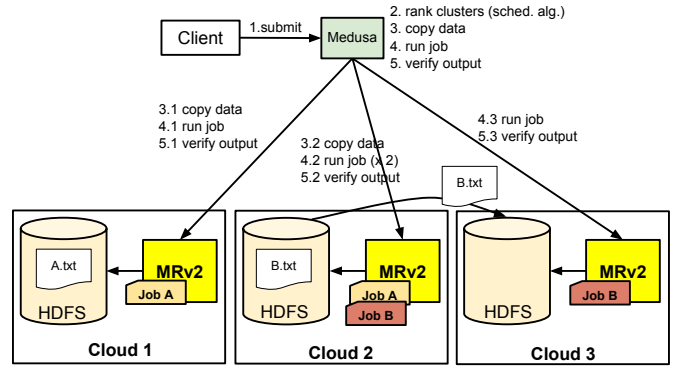


Fig. 3: Phase 1: Vanilla MapReduce execution.

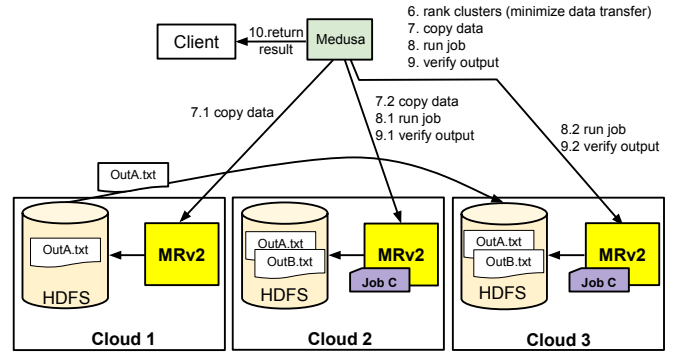


Fig. 4: Phase 2: Global MapReduce execution.

subset of data initially available in one cloud is replicated (at least)  $f + 1$  times, instead of  $2f + 1$  times.

We detail in Section III-A how Medusa works and explain in Section III-B how the replicated jobs are scheduled among the clouds in the system to ensure efficient completion of the entire job.

#### A. Medusa proxy in a nutshell

Each of the two phases of a full job execution is composed of three stages: replicating the data from the local cloud that initially holds them to other clouds; running the replicated jobs in all the clouds having the same data; and agreeing on the outputs of the replicated jobs. In each stage, all processes wait until every other process finishes, to move to the next stage. Fig. 3 and Fig. 4 depicts the two-phase execution by means of an example, where the full job is running in 3 clouds and is set to tolerate 1 fault ( $f = 1$ ).

The proxy interacts with the clouds with the help of a message queuing service (MQ). In Fig. 3 and Fig. 4, MQ is tightly coupled with the proxy, but in practice it can be running in a different host.

In the first phase (Fig. 3), the client submits a job (Step 1) with input data *A.txt* and *B.txt* (representing the two subsets of the entire data to be processed). As mentioned earlier, we assume that the cloud collected or generated the data itself as we anticipate this to be the most common scenario. Hence, the input data is already stored in *cloud 1* and *cloud 2*, respectively. After receiving the request from the client, Medusa runs a scheduling algorithm (Section III-B) to select the best cloud(s)

to run the replicated jobs (Step 2). Medusa selects the clouds that offer the best performance in terms of data transmission and data processing to ensure low job makespan (*i.e.*, the time it takes for the whole job to finish). In this example, Medusa chooses to copy *A.txt* to *cloud 2* and *B.txt* to *cloud 3*. Note that to guarantee data integrity, a digest of each data is computed using a collision-resistant hash function (*e.g.*, SHA-256) and is sent with the original data for validation at reception.

If the data transmission is successful and the communication is not tampered, the execution of the job will start (Step 4). After the executions of the  $f + 1$  job replicas finish, digests of their outputs are computed using the same collision-resistant hash function, and the correctness of these outputs are verified (Step 5). If all the digests are identical, the vanilla MapReduce job finishes successfully. Otherwise, the data is copied to a different cloud, selected by Medusa’s scheduling algorithm, and the job is executed again to obtain  $f + 1$  identical outputs. Once a majority of correct MapReduce jobs have finished, the second phase can start.

In the second phase (Fig. 4), Medusa chooses the  $f + 1$  clouds that minimize the data transmission time to gather the outputs of the first phase and to run the global MapReduce jobs to produce the final output. In this example, only *OutA.txt* is replicated to *cloud 3* (Step 7). Then, after verifying the integrity of the data, the  $f + 1$  replicated jobs are launched (Step 8). Once the  $f + 1$  replicated jobs finish, the digests of their outputs are computed and compared (Step 9). As in the first phase, if the digests are identical, the job finishes successfully, and the correct result is returned to the client. Otherwise, another replica of the global MapReduce job is launched until  $f + 1$  equal results from different clouds are obtained.

### B. The Medusa scheduler

When a client submits a job, Medusa needs to instantiate  $f + 1$  replicas of the job. These replicas will be launched in different clouds to tolerate cloud faults as explained in the previous section. Whenever there is a disagreement on the output of a job, Medusa needs to launch an extra execution of that job.

Deciding in which cloud a replicated job should be executed is crucial for the performance of the system. Specifically, the client is interested in minimizing the time it takes for the submitted job to complete: the *makespan* of the job. Intuitively, if a job, as well as the data it needs to process, is replicated to a particular cloud with high computational power and to which it is connected by high-bandwidth links, it should take relatively shorter time for the job to complete, when compared to the available alternatives. In contrary, if the data has to be copied to a cloud using low bandwidth links, or to a cloud that is already overloaded, it may take a very long time for the job to complete. Therefore, to ensure a small makespan of the entire job, it is important for Medusa to choose the appropriate clouds to replicate each job.

To this end, we propose a scheduler to distributes the replicated jobs across different clouds based on the predicted data transmission time and data processing time in each cloud. The prediction takes into account the historical performance as well as the current status of each cloud, allowing us to

incorporate the *heterogeneity* of the clouds into the scheduling decision.

The overall makespan is determined by the longest time to complete the vanilla MapReduce jobs running in all the clouds and the time to complete the global MapReduce job. Each job is replicated and executed between  $f + 1$  times and  $2f + 1$  times, depending on the number of faults  $f$ . The Medusa scheduler follows a greedy approach and selects, for each job, the clouds that minimize the time to complete the executions of all its replicas. Specifically, the Medusa scheduler estimates, for each job, the time to replicate and process the corresponding data in each cloud that does not initially holds the data, and selects the  $f + 1$  clouds in increasing order of the estimated time. In case of a fault, the Medusa scheduler makes another estimation and chooses the cloud with the shortest estimated time to run the extra replica. Note that the scheduler will deal in accordance with the type of fault. When the system is set (by the client) to assume any fault can be malicious, it will not consider the cloud where the fault occurred as an option to run the extra replica. That cloud cannot be trusted again since the malicious insider may corrupt the results once more. When the system is set to only tolerate arbitrary faults and cloud outage, any alive cloud can be an option to launch the extra replica.

Formally, the estimated time  $t_1(i)$  for completing a (replicated) vanilla MapReduce job (*i.e.*, phase 1) in cloud  $i$  can be written as

$$t_1(i) = t_{trans}(j, i) + t_{1,proc}(i), \quad (1)$$

where  $t_{trans}(j, i)$  is the estimated time to transfer the data to be processed by the job from cloud  $j$  to cloud  $i$ , and  $t_{1,proc}(i)$  is the estimated time to execute the vanilla MapReduce job in cloud  $i$ .

The estimated time  $t_2(i)$  for completing a (replicated) global MapReduce job (*i.e.*, phase 2) in cloud  $i$  can be written as

$$t_2(i) = \max_{j \in C \wedge j \neq i} \{t_{trans}(j, i)\} + t_{2,proc}(i). \quad (2)$$

This equation requires further explanation. For the global MapReduce job, all the outputs produced by the vanilla MapReduce jobs need to be copied to cloud  $i$  except the ones that already exist in cloud  $i$ . As the data transmission can take place in parallel, the time for transmitting the outputs of the vanilla MapReduce jobs is bounded by the maximum time for transmitting the outputs from any cloud  $j$  to cloud  $i$  ( $t_{trans}(j, i)$ ).  $t_{2,proc}(i)$  is the estimated time to execute the global MapReduce job in cloud  $i$ .

Therefore, for each job the Medusa scheduler selects the cloud  $i$  that has the shortest estimated time  $t_1(i)$  or  $t_2(i)$  to run a replica among the clouds that have not been selected until the correct output is obtained. We explain in the following how the data transmission time  $t_{trans}(j, i)$  and the data processing time  $t_{1,proc}(i)$  and  $t_{2,proc}(i)$  are estimated, respectively.

a) *Estimating data transmission time:* The data transmission time between two clouds depends on (i) the network distance between them, (ii) the network throughput between them, and (iii) the size of the data to transfer. Hence, we estimate the time to transfer data of size  $S$  between cloud  $j$  and cloud  $i$  as

$$t_{trans}(j, i) = l(j, i)/2 + S/T(j, i), \quad (3)$$

where  $l(j, i)$  is the round-trip time between cloud  $j$  and  $i$ , and  $T(j, i)$  is the estimated throughput between them.

The value of  $l(j, i)$  depends on the geographical distance between clouds  $j$  and  $i$ , the speed of transmission in the different propagation media connecting them, among other variables, and is usually stable. Since this value is small (typically in the milliseconds range), accurate estimation of the data transmission time depends largely on the estimation of the network throughput  $T(j, i)$  between cloud  $j$  and cloud  $i$ . Nevertheless, we decide to include both variables in Equation 3 for the sake of model completeness.

Considering that the throughput varies depending on the traffic load of other connections, Medusa keeps track of the throughput for each pair of clouds in the system. The measures are taken sequentially and periodically by running the network tool *Iperf*, a tool that measures maximum TCP bandwidth. Specifically, a script is running in each cloud to measure its throughput to other clouds. Before scheduling a job, Medusa will access this information to estimate the throughput from one cloud to the others. The throughput between two clouds is estimated as the average throughput between them over a window of size  $k$ , where  $k$  is the number of the most recent throughput measurements.

*b) Estimating data processing time:* The time for completing a given MapReduce job mainly depends on (i) the capacity of the cloud running this job and (ii) the configuration of the job. Obviously, if a cloud has high computational power and large amounts of free resources, the job running on it will finish within a short time. In addition, a high level of parallelization (i.e., large number of map and reduce tasks) for the same job in the same cloud implies shorter data processing times. Considering this, the Medusa scheduler relies on a linear regression model to predict the data processing time for a MapReduce job to complete in a cloud.

Specifically, the Medusa scheduler trains one linear regression model for each cloud in the system to predict its time to complete a MapReduce job, in the form

$$\hat{y} = \beta_1 x_1 + \dots + \beta_n x_n + \beta_0, \quad (4)$$

where  $\hat{y}$  is the data processing time to predict (i.e.,  $t_{1,proc}(i)$  or  $t_{2,proc}(i)$  in cloud  $i$ ) and  $x_1, \dots, x_n$  are the  $n$  features we use to make the prediction. We estimate the parameters  $\beta_0, \dots, \beta_n$  using the least squares approach [22].

The Medusa scheduler relies on three types of features to make the prediction: (i) job configuration; (ii) cloud capacity; and (iii) cloud overhead. We describe the representative features for each type in the following.

- *Job configuration features.* We consider the size of the input data, the number of map tasks and the number of reduce tasks as features in this type. Clearly, large input data and a small number of map and reduce tasks imply long job completion times. The values of these features are always known to the Medusa scheduler.
- *Cloud capacity features.* We consider the clock speed (MHz) and the number of cores of the CPU and the total memory capacity (MBs) as features in this type. These are variables that define the capacity of cloud, but they do not tell the load of the cloud in a specific time.

- *Cloud overhead features.* In addition to the computational capacity, the overhead in the cloud also has an impact on the completion time of the job to schedule. For instance, if a cloud is overloaded and there are already a number of jobs queued to be launched, the scheduled job will not finish in a short time even if very small. In contrast, if a cloud has more free resources the scheduled job can finish early even if its capacity is relatively low. We use the number of MapReduce jobs that are currently running in the cloud, the percentage of completion of the running MapReduce jobs, the number of MapReduce jobs that are queued to run, and the size of the input data of the running jobs as features in this type. These are part of the filesystem information that MRv2 can provide.

For each vanilla MapReduce job to schedule, the Medusa scheduler estimates the data processing time  $t_{1,proc}(i)$  in cloud  $i$  using this linear regression model. For the global MapReduce job, since the data to be processed is typically orders of magnitude smaller than the original data (thanks to the vanilla MapReduce jobs), it can be efficiently executed in any cloud. We thus ignore the time for running the global MapReduce job (i.e.,  $t_{2,proc}(i)$ ) when making the estimation of job completion time (i.e.,  $t_2(i)$ ) for cloud  $i$ . This helps to improve the efficiency of the Medusa scheduler.

## IV. EVALUATION

This section evaluates the performance of our system. Section IV-A describes the experimental setup as well as the implementation and configuration of Medusa. Section IV-B reports on its performance, considering both the presence and absence of faults during job execution.

### A. Experimental setup

*a) Applications:* We evaluate the performance of Medusa using the *WordCount*, *WebdataScan*, and *Monsterquery* benchmarks from Hadoop’s Gridmix benchmark [23], as examples of applications commonly used in real-world scenarios.

Running *WordCount* in a multi-cloud system can be considered as building the inverted indexes of a multi-site web search engine for each search site (i.e., cloud). Specifically, in a multi-cloud web search engine each site is in charge of collecting and indexing a subset of the entire document collection. To build the search index (that supports the term frequency-inverse document frequency style ranking functions, TF-IDF), each search site runs a local MapReduce job to parse the documents it has and to count the occurrences of each term in a document (i.e., TF) and the number of documents in that search site containing each term (i.e., partial IDF). This can be achieved by running *WordCount* as a vanilla MapReduce job in each cloud. To ensure the same search results can be retrieved as in a single-cloud search engine, the local outputs from previous executions need to be aggregated to obtain the number of documents in the search engine that contain each term (i.e., global IDF). To this end, we implemented a *WordCountAggregator*. This corresponds to the global MapReduce job described in Section III. *WebdataScan* is a benchmarking application that extracts samples from a large data set, which is a common form of processing in many systems. *Monsterquery*

is another benchmarking application that queries part of the data from a large data set. The MapReduce framework divides a query into steps and the dataset into chunks, and then runs those step/chunk pairs in separate physical hosts. The mappers perform the data collection phase and the reducers take care of data processing.

For evaluating these applications we used up to 6GB of data generated by Gridmix, equally partitioned and stored in all clouds. For a small subset of the experiments we tested larger files. The results obtained using these larger datasets confirm the general trend we report next.

*b) Implementation and configuration:* We evaluated the system in the ExoGENI testbed [24], a distributed networked infrastructure-as-a-service spread across the USA that allows setting up virtual topologies across sites and servers in each site.

We set up four clouds located in different sites for each experiment. For the WordCount benchmark, we used clouds located in the East and West coasts of the USA: California, Chicago, and West-Virginia. In the WebdataScan and Monsterquery benchmarks the clouds were geographically closer: Pittsburgh, Massachusetts, and Texas. For the sake of heterogeneity we set for each experiment two clouds in the same state, with the other two in different states (e.g., in the WordCount experiment we set two clouds in Chicago). The hardware used in all applications was diversified in terms of CPU characteristics and RAM size. We have set one specific cloud in each experiment with better resources than the remaining ones in order to maximize heterogeneity and to demonstrate the benefits of the Medusa scheduler. In the WordCount experiments the best cloud was in Chicago, whereas in the WebdataScan case it was in Pittsburgh.

Each cloud is composed of 4 hosts with a MRv2 runtime: one *resource manager* (master) and 3 *node managers* (slaves). The MRv2 framework is not modified, which leads to each Resource Manager being a single point of failure *in its cloud*. However, Medusa has the capability to detect if the MRv2 in a cloud is running. If not, that cloud is considered faulty (a type of fault our system tolerates). Medusa is installed in the client machine because we assume that the client is always correct. If a client was faulty, the job output and the proxy results would be compromised.

Medusa is implemented in Python 2.7. Medusa is the key component to schedule the jobs and tolerate cloud faults. It submits and coordinates the execution of jobs, makes scheduling decisions, and verifies the integrity of the replicated data and job outputs (checking if  $f + 1$  replicas are identical and launching new replicas accordingly). These operations require the proxy to communicate with the Resource Manager in each cloud. The proxy is logically located outside the system (in the sense that the clouds are oblivious to it). In our experiments, it runs in the same machine as the client, an Amazon AWS host located in Oregon. The message queuing service used is RabbitMQ [25]. The RabbitMQ server runs in the same machine as the proxy.

Each experiment is configured to first execute the WordCount, WebdataScan, or Monsterquery job and then aggregate the results to obtain the final output. Each experiment is repeated 40 times. The Medusa scheduler relies on a history of 30 executions to train the linear regression model and estimate

the throughput between clouds. We use Hadoop’s DistCp3 tool<sup>4</sup> to copy data between the HDFS of two clouds. We inject random workloads by running random numbers of extra jobs in each cloud, to simulate background overhead added by other users.

*c) Baseline:* We compare the performance of the Medusa scheduler in terms of job makespan – the time it takes to complete the entire job – and system workload against a baseline that also tolerates cloud faults but uses a simple strategy to schedule the replicated jobs – a *Round-robin* scheduler. This scheduler selects the clouds to run the replicas of a job in a circular order, assuming the clouds are numbered sequentially.

## B. Experimental performance

*a) Performance without faults:* We first evaluate the performance of Medusa when no fault occurs during the entire job execution. We consider  $f = 1$  in our experiments, so each job will be executed twice, *i.e.*,  $f + 1$  times. This choice of  $f = 1$  is based on the observation that the faults we consider in this paper, despite potentially having devastating consequences, are assumed not to be frequent.


*Job efficiency.* Fig. 5 compares the job makespan of the Medusa scheduler against the baseline with different job sizes. We observe that Medusa outperforms Round-robin in all cases. In particular, notice that for the WordCount experiment (Fig. 5a) *the Medusa scheduler is up to 3 times faster to complete the job when compared with Round-robin for the larger input sizes*. Moreover, from the bars depicting the first and third quartiles of the makespan in both figures, we conclude that *the Medusa scheduler offers more stable performance* as the variance of the Round-robin scheduler is much more evident (in fact, the variance of the others is so low that the bars are barely perceptible).

The performance gain of Medusa in the other two experiments exists but is not so pronounced because the throughput between clouds is similar. However, in terms of stability Medusa still offers an advantage: the variation in performance of the Round-robin is perceptible in all plots, contrary to our scheme.

*Discussion.* We now focus on trying to understand the results a bit further. For this purpose we measured the load distribution among clouds with the WordCount application to understand the improved makespan of the Medusa scheduler in more detail. Fig. 6 illustrates the usage of each cloud for executing an entire job when different schedulers are used. The Round-robin scheduler distributes the workload evenly across the clouds, ignoring cloud and network performance. In contrast, with the Medusa scheduler the jobs are launched more often in the Chicago clouds, followed by West-Virginia, and only rarely in California. The reason is twofold. First, as can be seen in Fig. 7, which shows the throughput measured between each pair of clouds in the system, the Chicago clouds have better network connections between them when compared to the other clouds. Thus, the transfer time to these clouds is, on average, smaller, and this advantage results in the Chicago clouds being

<sup>4</sup>Distributed Copy, or DistCp, is a tool that uses MapReduce for large inter/intra-cluster data copying.





**Fig. 5:** Job makespan of WordCount, WebdataScan, and MonsterQuery executions (no faults).

chosen more frequently to run replicated jobs. We observe, however – and that is the second reason –, that although West-Virginia has slightly lower network throughput when compared to California, more jobs are still scheduled to that cloud. This is because the Medusa scheduler considers *both* network transfer time and available computational capacity of the cloud when making its decision. So, in this case, as the relative difference of network throughput is not very high, the cloud with higher computational power (West-Virginia in this case) is capable of offering a lower makespan of the entire job.

Why does Medusa outperform Round-robin so considerably in the WordCount experiment, in particular, when compared with the other two? Recall from Section IV-A that we ran the two sets of experiments in different sets of clouds. Contrary to Fig. 7, the throughput measured between each pair of clouds in the WebdataScan and Monsterquery (Figure 8) experiments was very stable. The average throughput was around 1 Gbps.

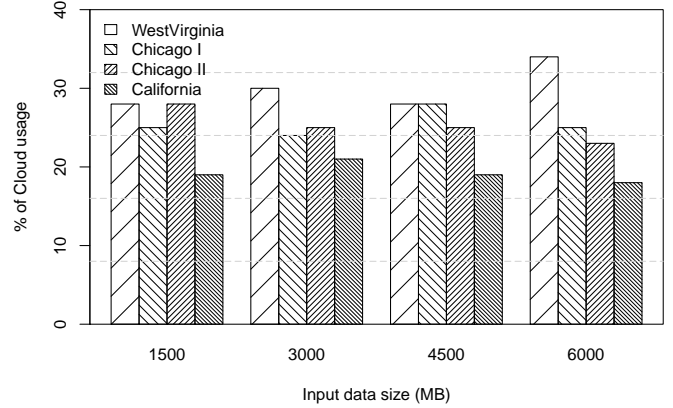
As is clear, this scenario is more homogeneous in terms of network throughput between clouds, so a Round-robin scheduler performs reasonably well in this setting. The reason why Medusa shows a very significant improvement for the WordCount experiment and less so for the others is therefore related to the heterogeneity of the scenario.

In summary, as we have shown, *the combination of the cloud characteristics and network throughput both influence the scheduling decision explaining the best performance of Medusa over a traditional scheduler.*


*b) Performance with faults:* In this section, we evaluate the performance of the Medusa scheduler when there are faults. We consider  $f = 1$  in this experiment, as before.

We first evaluate the performance of the Medusa scheduler when a malicious fault occurs. To this end, we inject a fault that corrupts the digest of a job output at the end of the vanilla WordCount execution, forcing the scheduler to launch an extra replica of this job in a different cloud. Fig. 9 shows the performance of the Medusa scheduler when one such fault occurs (Medusa w/ malicious faults). In this case, the job makespan doubles when compared to the case with no faults (Medusa w/o fault) as one extra job has to be executed in a different cloud. This is due to the fact that the scheduler can not consider as options the cloud where the job has just run. As explained in Section III-B, that cloud cannot be trusted again.

We also evaluate the performance of the Medusa scheduler by assuming that the fault is arbitrary but not malicious (Medusa w/ arbitrary faults). In this scenario, we also inject a fault that



**(a) Round-robin**



**(b) Medusa**

**Fig. 6:** Percentage of cloud usage.

tampers the digest of one job output at the end of the vanilla WordCount execution. The difference, in this case, is that the Medusa scheduler has the possibility to launch the extra replica of this job in the same cloud where the fault occurred (as we assume it is not malicious). As we observe from Fig.9, only tolerating arbitrary faults allows reducing the job makespan. This is mainly because the extra replica of the faulty job will with high probability be scheduled to the same cloud, so there is no need to copy the input of this job to another cloud.

Finally, we evaluate the performance of the Medusa scheduler when one cloud outage happens (Medusa w/ cloud outage). To this end, we simulate a cloud outage by crashing the resource manager of a random cloud. This forces the scheduler to launch






Fig. 7: Throughput between each pair of clouds (WordCount).

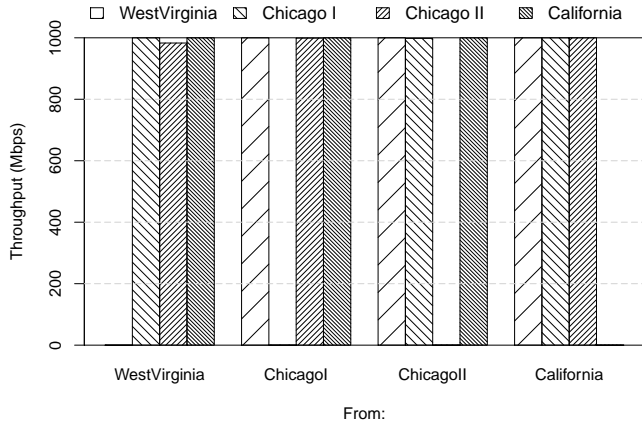


Fig. 8: Throughput between each pair of clouds (WebdataScan and Monsterquery).

an extra replica of this job to a different cloud. As shown in Fig. 9, tolerating a cloud outage requires higher job makespan than tolerating arbitrary or malicious faults. This is due to the scheduler taking some time to detect that the cloud has failed, looking for a copy of the data in another cloud, making the necessary data transfers and running the job.

Interestingly, when the input data is large, *the Medusa scheduler in the presence of faults performs better than Round-robin with no faults occurring*. We find it therefore unnecessary to report the performance of the Round-robin scheduler when a fault occurs.

To summarize, our experimental evaluation shows that the Medusa scheduler is more efficient than the conventional round-

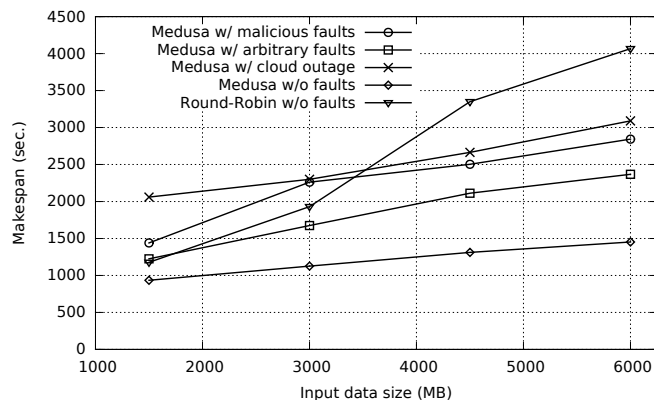


Fig. 9: Job makespan with one fault injected (WordCount).

robin alternative, and achieves a significant gain in realistic scenarios of high cloud heterogeneity.

## V. RELATED WORK

There has been much research on MapReduce since this framework was originally proposed in 2004 [26]. Just to give a few examples, there has been significant work on running MapReduce efficiently in different environments: multi-core/multiprocessor systems [27], heterogeneous environments [28], mixed CPU-GPU environments [29], high-latency eventually-consistent environments [30], among others. As the related work on this subject is vast, we invite the reader for the survey by Lee et al. [31].

With the requirements for compute-intensive analysis growing significantly, the need to scale-out MapReduce computation across clouds has lately given rise to interesting research on the subject. G-Hadoop [6] and G-MR [7] are two frameworks recently proposed to enable large-scale MapReduce computations to be distributed across clusters. Our work shares the goal of scaling out to multiple clouds but has several differentiating points. First, and most important, neither G-Hadoop nor G-MR deal with the types of faults our system tolerates: arbitrary and malicious faults, and cloud outages. They deal solely with crash faults, similar to the original MapReduce. In addition, they either require significant changes to Hadoop (G-Hadoop) or are a completely new Hadoop-based platform (G-MR). In contrast, our proposal uses Hadoop and does not require any changes to this widely used platform.

These works show the importance of MapReduce and Hadoop. Yet, neither advances the original platform from the fault tolerance point of view, as we propose in this paper. We leverage the vast literature on dependability in distributed systems that appeared since the algorithms to tolerate Byzantine or arbitrary faults were introduced, 30 years ago [32]. State machine replication, for instance, is a generic solution to build systems that are crash or Byzantine fault-tolerant [20]. The practicality of implementing efficient Byzantine fault-tolerant replication was demonstrated in [33], and from then onwards other efficient algorithms have been proposed [21], [34], [35].

We have proposed in previous work a Byzantine fault-tolerant MapReduce [36]. Contrary to the system we propose in the current paper, the target of [36] was a single datacenter. Hence, the system did not tolerate cloud outages or malicious faults caused by malicious insiders. Notably, our previous solution required changes to Hadoop, contrary to the proposal we make here. ClusterBFT is also a system for Byzantine fault-tolerant data-flow processing in clouds [37]. Contrary to our work, it was not designed for MapReduce.

## VI. CONCLUDING REMARKS

In this paper we proposed a platform, Medusa, for scaling out MapReduce to multiple clouds and, simultaneously, tolerating several kind of faults introduced by such multi-cloud environment.

Our proposal fulfilled the objectives we set forth. First, Medusa scales MapReduce computations to multiple clouds. Second, it extends the fault-tolerance offered by MapReduce to tolerate arbitrary and malicious faults, as well as cloud outages.

Third, it does so transparently to the user. The Hadoop API is not touched and the existing Hadoop MapReduce programs run without modification. Forth, Medusa is a proxy in the client, and thus the system does not require any modification to the Hadoop framework. The clouds just need to run “vanilla” Hadoop (e.g., Amazon Elastic MapReduce [38]). Finally, as demonstrated by our extensive experimental evaluation in the ExoGENI testbed, it achieves this increased level of fault tolerance at a reasonable cost when compared with common alternatives. It does so by minimizing replication and by running a novel scheduling algorithm that judiciously chooses the best clouds to perform the necessary replicated jobs.

We made our code available as open source<sup>5</sup>. We invite the community to use it and improve upon our code. As future work, we plan to investigate techniques to improve the performance of our system by replicating at a more fine-grained level than at a job level.

#### ACKNOWLEDGMENT

This work was partially supported by the EU FP7 project LEADS (ICT-318809), by the EC through project H2020-643964 (SUPERCLOUD), and by national funds through Fundação para a Ciência e a Tecnologia (FCT) with reference UID/CEC/00408/2013 (LaSIGE) and UID/CEC/50021/2013 (INESC-ID). This work was partially conducted during the first author’s internship at Yahoo Labs Barcelona.

#### REFERENCES

[1] J. Dean and S. Ghemawat, “MapReduce: a flexible data processing tool,” *Communications of the ACM*, vol. 53, pp. 72–77, Jan. 2010.

[2] T. White, *Hadoop: The Definitive Guide*. O’Reilly, 1st ed., June 2009.

[3] R. Baeza-Yates, A. Gionis, F. Junqueira, et al., “On the feasibility of multi-site web search engines,” in *Proceedings of the 18th ACM Conference on Information and Knowledge Management*, pp. 425–434, 2009.

[4] A. Matsunaga, M. Tsugawa, and J. Fortes, “Cloudblast: Combining mapreduce and virtualization on distributed resources for bioinformatics applications,” in *Proceedings of the 4th IEEE International Conference on eScience*, pp. 222–229, 2008.

[5] G. Brumfiel, “High-energy physics: Down the petabyte highway,” *Nature*, vol. 469, pp. 282–283, Jan. 2011.

[6] L. Wang, J. Tao, R. Ranjan, et al., “G-Hadoop: MapReduce across distributed data centers for data-intensive computing,” *Future Generation Computer Systems*, vol. 29, no. 3, pp. 739 – 750, 2013.

[7] C. Jayalath, J. J. Stephen, and P. Eugster, “From the cloud to the atmosphere: Running MapReduce across data centers,” *IEEE Transactions on Computers*, vol. 63, no. 1, pp. 74–87, 2014.

[8] S. Ghemawat, H. Gobioff, and S.-T. Leung, “The Google file system,” in *Proceedings of the 19th ACM Symposium on Operating Systems Principles*, pp. 29–43, 2003.

[9] B. Schroeder, E. Pinheiro, and W.-D. Weber, “DRAM errors in the wild: A large-scale field study,” in *Proceedings of the 11th International Joint Conference on Measurement and Modeling of Computer Systems*, pp. 193–204, 2009.

[10] E. B. Nightingale, J. R. Douceur, and V. Orgovan, “Cycles, cells and platters: an empirical analysis of hardware failures on a million consumer PCs,” in *Proceedings of the ACM SIGOPS/EuroSys European Conference on Computer Systems*, pp. 343–356, 2011.

[11] A. Avizienis, J.-C. Laprie, B. Randell, and C. Landwehr, “Basic concepts and taxonomy of dependable and secure computing,” *IEEE Transactions on Dependable and Secure Computing*, vol. 1, pp. 11–33, Mar. 2004.

[12] P. Verissimo, N. F. Neves, and M. P. Correia, “Intrusion-tolerant architectures: Concepts and design,” in *Architecting Dependable Systems*, vol. 2677 of *LNCS*, Springer-Verlag, 2003.

[13] T. Kurze, M. Klems, D. Bermbach, et al., “Cloud federation,” in *Proceedings of the 2nd International Conference on Cloud Computing, Grids, and Virtualization*, Sept. 2011.

[14] A. Bessani, M. Correia, B. Quaresma, F. André, and P. Sousa, “DepSky: Dependable and secure storage in a cloud-of-clouds,” in *Proceedings of the 6th ACM SIGOPS/EuroSys European Conference on Computer Systems*, pp. 31–46, 2011.

[15] C. Cerin et al., “Downtime statistics of current cloud solutions.” The International Working Group on Cloud Computing Resiliency, June 2013.

[16] G. Clarke, “Microsoft azure was most fail-filled cloud of 2014.”

[17] Cloud Security Alliance, “The notorious nine cloud computing top threats in 2013.” <https://cloudsecurityalliance.org/download/the-notorious-nine-cloud-computing-top-threats-in-2013/>.

[18] M. Kandias, N. Virvilis, and D. Gritzalis, “The insider threat in cloud computing,” in *Critical Information Infrastructure Security*, vol. 6983 of *LNCS*, pp. 93–103, Springer Berlin Heidelberg, 2013.

[19] Amazon Web Services Inc, “Amazon S3 introduces cross-region replication,” Mar 2015. <http://aws.amazon.com/pt/about-aws/whats-new/2015/03/amazon-s3-introduces-cross-region-replication/>.

[20] F. B. Schneider, “Implementing fault-tolerant service using the state machine approach: A tutorial,” *ACM Computing Surveys*, vol. 22, pp. 299–319, Dec. 1990.

[21] G. S. Veronese, M. Correia, A. N. Bessani, L. C. Lung, and P. Verissimo, “Efficient Byzantine fault tolerance,” *IEEE Transactions on Computers*, vol. 62, no. 1, pp. 16–30, 2013.

[22] A. Björck, *Numerical Methods for Least Squares Problems*. Siam Philadelphia, 1996.

[23] “MapReduce 0.22 Documentation – GridMix.” <http://hadoop.apache.org/mapreduce/docs/current/gridmix.html>.

[24] I. Baldine, Y. Xin, A. Mandal, P. Ruth, C. Heerman, and J. Chase, “Exogeni: A multi-domain infrastructure-as-a-service testbed,” in *Testbeds and Research Infrastructure. Development of Networks and Communities* (T. Korakis, M. Zink, and M. Ott, eds.), vol. 44 of *Lecture Notes of the Institute for Computer Sciences, Social Informatics and Telecommunications Engineering*, pp. 97–113, Springer Berlin Heidelberg, 2012.

[25] A. Videla and J. Williams, *RabbitMQ in Action: Distributed Messaging for Everyone*. Manning Publications Company, 2012.

[26] J. Dean and S. Ghemawat, “MapReduce: simplified data processing on large clusters,” in *Proceedings of the 6th Symposium on Operating Systems Design & Implementation*, Dec. 2004.

[27] C. Ranger, R. Raghuraman, A. Penmetsa, G. Bradski, and C. Kozyrakis, “Evaluating MapReduce for multi-core and multiprocessor systems,” in *Proceedings of the 13th IEEE International Symposium on High Performance Computer Architecture*, pp. 13–24, 2007.

[28] M. Zaharia, A. Konwinski, A. D. Joseph, R. Katz, and I. Stoica, “Improving MapReduce performance in heterogeneous environments,” in *Proceedings of 8th USENIX Symposium on Operating Systems Design and Implementation*, pp. 29–42, 2008.

[29] W. Jiang and G. Agrawal, “MATE-CG: A Map Reduce-like framework for accelerating data-intensive computations on heterogeneous clusters,” in *Proceedings of the 2012 IEEE International Parallel & Distributed Processing Symposium*, pp. 644–655, 2012.

[30] T. Gunarathne, T.-L. Wu, J. Qiu, and G. Fox, “MapReduce in the clouds for science,” in *Proceedings of the 2nd IEEE International Conference on Cloud Computing Technology and Science*, pp. 565–572, 2010.

[31] K.-H. Lee, Y.-J. Lee, H. Choi, Y. D. Chung, and B. Moon, “Parallel data processing with MapReduce: A survey,” *SIGMOD Rec.*, vol. 40, pp. 11–20, Jan 2012.

[32] L. Lamport, R. Shostak, and M. Pease, “The Byzantine generals problem,” *ACM Transactions on Programming Languages and Systems*, vol. 4, pp. 382–401, July 1982.

[33] M. Castro and B. Liskov, “Practical Byzantine fault-tolerance and proactive recovery,” *ACM Transactions Computer Systems*, vol. 20, pp. 398–461, Nov. 2002.

[34] A. Clement, M. Kapritsos, S. Lee, Y. Wang, L. Alvisi, M. Dahlin, and T. Rich, “UpRight cluster services,” in *Proceedings of the 22nd ACM Symposium on Operating Systems Principles*, Oct. 2009.

[35] Y. Amir et al., “Scaling Byzantine fault-tolerant replication to wide area networks,” in *Proceedings of the IEEE/IFIP 36th International Conference on Dependable Systems and Networks*, pp. 105–114, June 2006.

[36] P. Costa, M. Pasin, A. N. Bessani, and M. P. Correia, “On the performance of Byzantine fault-tolerant MapReduce,” *IEEE Transactions on Dependable and Secure Computing*, vol. 10, no. 5, pp. 301–313, 2013.

[37] J. J. Stephen and P. Eugster, “Assured cloud-based data analysis with ClusterBFT,” in *Proceedings of the 14th ACM/IFIP/USENIX International Conference on Middleware*, vol. 8275 of *LNCS*, pp. 82–102, Springer, 2013.

[38] Amazon Web Services Inc, “Amazon EMR.” <http://aws.amazon.com/elasticmapreduce/>.

<sup>5</sup>[https://bitbucket.org/pcosta\\_pt/medusa\\_hadoop](https://bitbucket.org/pcosta_pt/medusa_hadoop)