

Resource Provisioning Framework for MapReduce Jobs with Performance Goals ^{*}

Abhishek Verma¹, Ludmila Cherkasova², and Roy H. Campbell¹

¹ University of Illinois at Urbana-Champaign
{verma7, rhc}@illinois.edu
² HP Labs, Palo Alto
{lucy.cherkasova}@hp.com

Abstract. Many companies are increasingly using MapReduce for efficient large scale data processing such as personalized advertising, spam detection, and different data mining tasks. Cloud computing offers an attractive option for businesses to rent a suitable size Hadoop cluster, consume resources as a service, and pay only for resources that were utilized. One of the open questions in such environments is the amount of resources that a user should lease from the service provider. Often, a user targets specific performance goals and the application needs to complete data processing by a certain time deadline. However, currently, the task of estimating required resources to meet application performance goals is solely the users' responsibility. In this work, we introduce a novel framework and technique to address this problem and to offer a new resource sizing and provisioning service in MapReduce environments. For a MapReduce job that needs to be completed within a certain time, the job profile is built from the job past executions or by executing the application on a smaller data set using an automated profiling tool. Then, by applying scaling rules combined with a fast and efficient capacity planning model, we generate a set of resource provisioning options. Moreover, we design a model for estimating the impact of node failures on a job completion time to evaluate worst case scenarios. We validate the accuracy of our models using a set of realistic applications. The predicted completion times of generated resource provisioning options are within 10% of the measured times in our 66-node Hadoop cluster.

1 Introduction

Private and public clouds offer a new delivery model with virtually unlimited computing and storage resources. Many companies are following the new trend of using MapReduce [1] and its open-source implementation Hadoop for large-scale, data intensive processing and for mining petabytes of unstructured information. However, setting up a dedicated Hadoop cluster requires a significant capital expenditure that can be difficult to justify. Cloud computing offers a compelling alternative and allows users to rent resources in a “pay-as-you-go” fashion. A list of supported services by Amazon (Amazon Web Services) includes MapReduce environments for rent. It is an attractive and cost-efficient option for many users because acquiring and maintaining complex, large-scale infrastructures is a difficult and expensive decision. Hence, a typical practice among

^{*} This work was largely completed during A. Verma's internship at HP Labs. R. Campbell and A. Verma are supported in part by NSF CCF grants #0964471, IIS #0841765 and Air Force Research grant FA8750-11-2-0084.

MapReduce users is to develop their applications in-house using a small development testbed and test it over a small input dataset. They can lease a MapReduce cluster from the service provider and subsequently execute their MapReduce applications on large input datasets of interest. Often, the application is a part of a more elaborate business pipeline, and the MapReduce job has to produce results by a certain time deadline, i.e., it has to achieve certain performance goals and service level objectives (SLOs). Thus, a typical performance question in MapReduce environments is “*how to estimate the required resources (number of map and reduce slots) for a job so that it achieves certain performance goals and completes data processing by a given time?*” Currently, there is no available methodology to easily answer this question, and businesses are left on their own to struggle with the resource sizing problem: they need to perform adequate application testing, performance evaluation, capacity planning estimation, and then request appropriate amount of resources from the service provider.

In this work, we propose a novel framework to solve this problem and offer a new resource sizing and provisioning service in MapReduce environments. First, we introduce an automated profiling tool that extracts a compact job profile from the past application execution(s) in the production Hadoop cluster. Alternatively, profiling can be done by executing a given application with a smaller input dataset than the original one. The power of the designed technique is that it offers a compact job profile that is comprised of performance *invariants* which are independent of the amount of resources assigned to the job (i.e., the size of the Hadoop cluster) and the size of the input dataset. The job profile accurately reflects the application performance characteristics during all phases of a given job: map, shuffle/sort, and reduce phases. Our automated profiling technique does not require any modifications or instrumentation of neither the application nor of the underlying Hadoop execution engine. All this information can be obtained from the counters at the job master during the job’s execution or alternatively parsed from the job execution logs written at the job tracker.

For many applications, increasing input dataset while keeping the same number of reduce tasks leads to an increased amount of data shuffled and processed per reduce task. Using linear regression, we derive *scaling factors* for shuffle and reduce phases to estimate their service times as a function of input data.

We design a *MapReduce performance model* that predicts the job completion time based on the job profile, input dataset size, and allocated resources. We enhance the designed model to evaluate the performance impact of failures on job completion time. This model helps in evaluating worst case scenarios and deciding on the necessity of additional resources or program changes as a means of coping with potential failure scenarios. Finally, we propose a fast and efficient, fully automated capacity planning procedure for estimating the required resources to meet a given application SLO. The output of the model is a set of plausible solutions (if such solutions exist for a given SLO) with a choice of different numbers of map and reduce slots that need to be allocated for achieving performance goals of this application.

We validate the accuracy of our approach and performance models using a set of realistic applications. First, we build the application profiles and derive scaling factors using small input datasets for processing. Then we perform capacity planning and generate plausible resource provisioning options for achieving a given application SLO for processing a given (large) dataset. The predicted completion times of these generated options are within 10% of the measured times in the 66-node Hadoop cluster.

This paper is organized as follows. Section 2 provides a background on MapReduce. Section 3 introduces our approach towards profiling MapReduce jobs. Section 4 presents a variety of MapReduce performance models and the SLO-based resource provisioning. The efficiency of our approach and the accuracy of designed models is evaluated in Section 5. Section 6 describes the related work. Section 7 summarizes the paper and outlines future directions.

2 MapReduce Background

This section provides an overview of the MapReduce [1] abstraction, execution, scheduling, and failure modes. In the MapReduce model, computation is expressed as two functions: map and reduce. The map function takes an input pair and produces a list of intermediate key/value pairs. The intermediate values associated with the same key k_2 are grouped together and then passed to the reduce function. The reduce function takes intermediate key k_2 with a list of values and processes them to form a new list of values.

$$\begin{aligned} \text{map}(k_1, v_1) &\rightarrow \text{list}(k_2, v_2) \\ \text{reduce}(k_2, \text{list}(v_2)) &\rightarrow \text{list}(v_3) \end{aligned}$$

MapReduce jobs are distributed and executed across multiple machines: the map stage is partitioned into *map tasks* and the reduce stage is partitioned into *reduce tasks*.

Each map task processes a logical split of input data that generally resides on a distributed file system. Files are typically divided into uniform sized blocks (default size is 64 MB, which is a file system parameter) and distributed across the cluster nodes. The map task reads the data, applies the user-defined map function on each record, and buffers the resulting output. This data is sorted and partitioned for different reduce tasks, and written to the local disk of the machine executing the map task. The reduce stage consists of two phases: *shuffle* and *reduce* phase. In the *shuffle phase*, the reduce tasks fetch the intermediate data files from the already completed map tasks, thus following the “pull” model. The intermediate files from all the map tasks are sorted. An external merge sort is used in case the intermediate data does not fit in memory as follows: the intermediate data is shuffled, merged in memory, and written to disk. After all the intermediate data is shuffled, a final pass is made to merge all these sorted files. Thus, the shuffling and sorting of intermediate is interleaved: we denote this by *shuffle/sort* or simply *shuffle* phase. Finally, in the *reduce phase*, the sorted intermediate data is passed to the user-defined reduce function. The output from the reduce function is generally written back to the distributed file system.

Job scheduling in Hadoop is performed by a master node, which manages a number of worker nodes in the cluster. Each worker has a fixed number of *map slots* and *reduce slots*, which can run tasks. The number of map and reduce slots is statically configured. The slaves periodically send heartbeats to the master to report the number of free slots and the progress of tasks that they are currently running. Based on the availability of free slots and the scheduling policy, the master assigns map and reduce tasks to slots in the cluster.

In the real world, user code is buggy, processes crash, and machines fail. MapReduce is designed to scale to a large number of machines and to yield a graceful performance degradation in case of failures. There are three types of failures that can occur.

First, a map or reduce task can fail because of buggy code or runtime exceptions. The worker node running the failed task detects task failures and notifies the master. The master reschedules the execution of the failed task, preferably on a different machine. Secondly, a worker can fail, e.g., because of OS crash, faulty hard disk, or network interface failure. The master notices a worker that has not sent any heartbeats for a specified time interval and removes it from its worker pool for scheduling new tasks. Any tasks in progress on the failed worker are rescheduled for execution. The master also reschedules all the completed map tasks on the failed worker that belong to running jobs, since the intermediate data of these maps may not be accessible to reduce tasks of these jobs. Finally, the failure of the master is the most serious failure mode. Currently, Hadoop has no mechanism for dealing with the failure of the job master. This failure is rare and can be avoided by running multiple masters and using a Paxos consensus protocol to decide the primary master.

3 Profiling MapReduce Jobs

In this section, we discuss different executions of the same MapReduce job in the Hadoop cluster as a function of the job's map and reduce tasks and the allocated map and reduce slots for executing it. Our goal is to extract a single *job profile* that uniquely captures critical performance characteristics of the job execution in different stages. We introduce a fully automated profiling technique that extracts a compact job profile from the past application execution(s) in the production Hadoop cluster.

3.1 Job Execution as a Function of Allocated Resources

Let us consider two popular MapReduce applications (described below) and demonstrate the differences between their job executions and job completion times as a function of the amount of resources allocated to these jobs.

- The first application is the *Sort benchmark* [2], which involves the use of identity map/reduce function: The output of the map and reduce task is the same as its input. Thus, the entire input of map tasks is shuffled to reduce tasks and then written as output.
- The second application is *WikiTrends* application that processes Wikipedia article traffic logs that were collected (and compressed) every hour. *WikiTrends* counts the number of times each article has been visited in the given input dataset, i.e. access frequency or popularity count of each Wikipedia article over time.

First, we run the Sort benchmark with 8GB input on 64 machines each configured with a single map and a single reduce slot, i.e., with 64 map and 64 reduce slots overall. Figure 1 shows the progress of the map and reduce tasks over time (on the x-axis) vs the 64 map slots and 64 reduce slots (on the y-axis). Since we use blocksize of 128MB, we have $8\text{GB}/128\text{MB} = 64$ input splits. As each split is processed by a different map task, the job consists of 64 map tasks. This job execution results in a single map and reduce *wave*. We split each reduce task into its constituent shuffle/sort and reduce phases. As seen in the figure, since the shuffle phase starts immediately after the first map task is completed, the shuffle phase overlaps with the map stage.

Next, we run the Sort benchmark with the same 8GB input dataset on the same testbed, except this time, we provide it with fewer resources: 16 map slots and 22 reduce




Fig. 1. Sorting with 64 map and 64 reduce slots




Fig. 2. WikiTrends with 64 map and 64 reduce slots




Fig. 3. Sorting with 16 map and 22 reduce slots




Fig. 4. WikiTrends with 16 map and 16 reduce slots

slots. As shown in Figure 3, since the number of map tasks is greater than the number of provided map slots, the map stage proceeds in multiple rounds of slot assignment, viz. 4 waves ($\lceil 64/16 \rceil$)¹. These waves are not synchronized with each other, and the scheduler assigns the task to the slot with the earliest finishing time. Similarly, the reduce stage proceeds in 3 waves ($\lceil 64/22 \rceil$).

While the executions of four map waves resemble each other, note the difference between the first reduce wave and the following two reduce waves in this job execution. As we mentioned earlier, the shuffle phase of the first reduce wave starts immediately after the first map task completes. Moreover, this first shuffle phase continues until all the map tasks are complete, and their intermediate data is copied to the active reduce tasks. Thus the first shuffle phase overlaps with the entire map stage. After the shuffle/sort phase is done, the reduce computation can be performed and the generated data

¹ Note, that for multiple map and reduce waves, there is an extra overhead for starting map and reduce tasks. In our ongoing work, we design a set of micro-benchmarks to automatically assess these overheads in different MapReduce environments for incorporating them in the performance model.

written back to HDFS. After that, the released reduce slots become available to the next reduce tasks. As shown in Figure 3, there is a drastic difference between the execution of the first reduce wave, but the executions of the remaining reduce waves bear a strong resemblance to each other.

Figures 2 and 4 present our second example with WikiTrends application. In this example, we process a subset of logs collected during a few days in April, 2011. There are 71 files in the set that correspond to 71 map tasks and 64 reduce tasks in this application. First, we execute WikiTrends with 64 map and 64 reduce slots. The job execution consists of two map waves ($\lceil 71/64 \rceil$) and a single reduce wave as shown in Figure 2. The second map wave processes only 7 map tasks. However, the shuffle phase of the reduce stage can be completed only when all the map tasks are done, and overlaps with both preceding map waves. Figure 4 shows the WikiTrends execution with 16 map and 16 reduce slots. The job execution has 5 map and 4 reduce waves. Again, we can see a striking difference between the first reduce wave and the remaining 3 reduce waves (which resemble each other).

As observed from Figures 1- 4, it is difficult to predict the completion time of the same job when different amount of resources are given to the job. Traditionally, a simple rule of thumb states [3], that if T is a completion time of a MapReduce job with X map and Y reduce slots then by using a smaller Hadoop cluster with $X/2$ map and $Y/2$ reduce slots the same job will be processed twice as slow, i.e., in $2 \cdot T$. While it is clear, that the job execution time is a function of allocated resources, the scaling rules are more complex, and the simple example with WikiTrends shows this. The completion time of WikiTrends in 64x64 configuration is approx. 200 sec. However, the completion time of WikiTrends in 16x16 configuration (4 times smaller cluster) is approx. 570 sec, which is far less than 4 times (naively expected) completion time increase. Evidently, more elaborate modeling and job profiling techniques are needed to capture the unique characteristics of MapReduce applications and to predict their completion time.

3.2 Job Performance Invariants as a Job Profile

Our goal is to create a compact job profile comprising of performance *invariants* that are independent of the amount of resources assigned to the job over time and that reflects all phases of a given job: map, shuffle/sort, and reduce phases. Metrics and timing of different phases, that we use below, can be obtained from the counters at the job master during the job's execution or parsed from the logs.

Map Stage: The map stage consists of a number of map tasks. To compactly characterize the distribution of the map task durations and other invariant properties, we extract the following metrics:

$$(M_{min}, M_{avg}, M_{max}, AvgSize_M^{input}, Selectivity_M), \text{ where}$$

- M_{min} – the minimum map task duration. Since the shuffle phase starts when the first map task completes, we use M_{min} as an estimate for the beginning of the shuffle phase.
- M_{avg} – the average duration of map tasks to summarize the duration of a map wave.

- M_{max} – the maximum duration of the map tasks². Since the shuffle phase can finish only when the entire map stage completes, i.e. all the map tasks complete, M_{max} is an estimate for a worst map wave completion time.
- $AvgSize_M^{input}$ – the average amount of input data for the map task. We use it to estimate the number of map tasks to be spawned for processing a new dataset.
- $Selectivity_M$ – the ratio of the map output size to the map input size. It is used to estimate the amount of intermediate data produced by the map stage as input to the reduce stage.

Reduce Stage: As described earlier, the reduce stage consists of the shuffle/sort and reduce phases. The shuffle phase begins only after the first map task has completed. The shuffle phase (of any reduce wave) completes when the entire map stage is complete and all the intermediate data generated by the map tasks has been shuffled to the reduce tasks and has been sorted. After shuffle/sort completes, the reduce phase is performed. Thus the profiles of shuffle and reduce phases are represented by the *average* and *maximum* of their tasks durations. In addition, for the reduce phase, we compute the *reduce selectivity*, denoted as $Selectivity_R$, which is defined as the ratio of the reduce output size to its input.

The shuffle phase of the first reduce wave may be significantly different from the shuffle phase that belongs to the next reduce waves (illustrated in Figure 3, 4). This happens because the shuffle phase of the first reduce wave overlaps with the entire map stage and depends on the number of map waves and their durations. Therefore, we collect two sets of measurements: (Sh_{avg}^1, Sh_{max}^1) for shuffle phase of the first reduce wave (called, *first shuffle*) and $(Sh_{avg}^{typ}, Sh_{max}^{typ})$ for shuffle phase of the other waves (called, *typical shuffle*). Since we are looking for performance invariants that are independent of the amount of allocated resources to the job, we characterize shuffle phase of the first reduce wave in a special way and include only the non-overlapping portions of the first shuffle in (Sh_{avg}^1) and (Sh_{max}^1) . Thus the job profile in the shuffle phase is characterized by two pairs of measurements: $(Sh_{avg}^1, Sh_{max}^1, Sh_{avg}^{typ}, Sh_{max}^{typ})$.

The reduce phase begins only after the shuffle phase is complete. The profile of the reduce phase is represented by the *average* and *maximum* of the reduce tasks durations and the *reduce selectivity*, denoted as $Selectivity_R$, which is defined as the ratio of the reduce output size to its input: $(R_{avg}, R_{max}, Selectivity_R)$.

The extracted job profile is independent of the scheduler that is used for executing the jobs: the job profile metrics represent the task durations and the amount of processed data. The collected job profile is “scheduler-agnostic”.

4 MapReduce Performance Model

In this section, we design a MapReduce performance model that is based on *i*) the job profile and *ii*) the performance bounds of completion time of different job phases. This model can be used for predicting the job completion time as a function of the input dataset size and allocated resources.

² To avoid outliers and improve the robustness of the measured maximum durations, one can use the mean of a few top values instead of the maximum alone.

4.1 General Theoretical Bounds

First, we establish the performance bounds for a makespan (completion time) of a given set of n tasks that is processed by k servers (or by k slots in MapReduce environments).

Let T_1, T_2, \dots, T_n be the duration of n tasks of a given job. Let k be the number of slots that can each execute one task at a time. The assignment of tasks to slots is done using a simple, online, *greedy* algorithm, i.e., assign each task to the slot with the earliest finishing time.

Let $avg = (\sum_{i=1}^n T_i)/n$ and $max = \max_i \{T_i\}$ be the *average* and *maximum* durations of the n tasks respectively.

Makespan Theorem: The makespan of the greedy task assignment is at least $n \cdot avg / k$ and at most $(n - 1) \cdot avg/k + max$.

The lower bound is trivial, as the best case is when all n tasks are equally distributed among the k slots (or the overall amount of work $n \cdot avg$ is processed as fast as possible by k slots). Thus, the overall makespan is at least $n \cdot avg/k$.

For the upper bound, let us consider the worst case scenario, i.e., the longest task $\hat{T} \in \{T_1, T_2, \dots, T_n\}$ with duration max is the last processed task. In this case, the time elapsed before the final task \hat{T} is scheduled is at most the following: $(\sum_{i=1}^{n-1} T_i)/k \leq (n - 1) \cdot avg/k$. Thus, the makespan of the overall assignment is at most $(n - 1) \cdot avg/k + max$.³ ■

These bounds are particularly useful when $max \ll n \cdot avg/k$, i.e., when the duration of the longest task is small as compared to the total makespan. The difference between lower and upper bounds represents the range of possible job completion times due to non-determinism and scheduling.

4.2 Bounds-Based Completion Time Estimates of a MapReduce Job

Let us consider job J with a given profile extracted from the past job executions. Let J be executed with a new dataset that is partitioned into N_M^J map tasks and N_R^J reduce tasks. Let S_M^J and S_R^J be the number of map and reduce slots allocated to job J respectively.

Let M_{avg} and M_{max} be the average and maximum durations of map tasks (defined by the job J profile). Then, by Makespan Theorem, the lower and upper bounds on the duration of the entire map stage (denoted as T_M^{low} and T_M^{up} respectively) are estimated as follows:

$$T_M^{low} = N_M^J \cdot M_{avg}/S_M^J \quad (1)$$

$$T_M^{up} = (N_M^J - 1) \cdot M_{avg}/S_M^J + M_{max} \quad (2)$$

The reduce stage consists of shuffle (which includes the interleaved sort phase) and reduce phases. Similarly, Makespan Theorem can be directly applied to compute the lower and upper bounds of completion times for reduce phase (T_R^{low} , T_R^{up}) since we have measurements for average and maximum task durations in the reduce phase, the numbers of reduce tasks N_R^J and allocated reduce slots S_R^J .⁴

³ Similar ideas were explored in the classic papers on scheduling, e.g., to characterize makespan bounds in [4].

⁴ For simplicity of explanation, we omit the normalization step of measured durations in job profile with respect to $AvgSize_M^{input}$ and $Selectivity_M$. We will discuss it next in Section 4.3.

The subtlety lies in estimating the duration of the shuffle phase. We distinguish the non-overlapping portion of the *first shuffle* and the task durations in the *typical shuffle* (see Section 3 for definitions). The portion of the typical shuffle phase in the remaining reduce waves is computed as follows:

$$T_{Sh}^{low} = \left(\frac{N_R^J}{S_R^J} - 1 \right) \cdot Sh_{avg}^{typ} \quad (3)$$

$$T_{Sh}^{up} = \left(\frac{N_R^J - 1}{S_R^J} - 1 \right) \cdot Sh_{avg}^{typ} + Sh_{max}^{typ} \quad (4)$$

Finally, we can put together the formulae for the lower and upper bounds of the overall completion time of job J :

$$T_J^{low} = T_M^{low} + Sh_{avg}^1 + T_{Sh}^{low} + T_R^{low} \quad (5)$$

$$T_J^{up} = T_M^{up} + Sh_{max}^1 + T_{Sh}^{up} + T_R^{up} \quad (6)$$

Note that we can re-write Eq. 5 for T_J^{low} by replacing its parts with more detailed Eq. 1 and Eq. 3 and similar equations for sort and reduce phases as it is shown below:

$$T_J^{low} = \frac{N_M^J \cdot M_{avg}}{S_M^J} + \frac{N_R^J \cdot (Sh_{avg}^{typ} + R_{avg})}{S_R^J} + Sh_{avg}^1 - Sh_{avg}^{typ} \quad (7)$$

This presentation allows us to express the estimates for completion time in a simplified form shown below:

$$T_J^{low} = A_J^{low} \cdot \frac{N_M^J}{S_M^J} + B_J^{low} \cdot \frac{N_R^J}{S_R^J} + C_J^{low}, \quad (8)$$

where $A_J^{low} = M_{avg}$, $B_J^{low} = (Sh_{avg}^{typ} + R_{avg})$, and $C_J^{low} = Sh_{avg}^1 - Sh_{avg}^{typ}$. Eq. 8 provides an explicit expression of a job completion time as a function of map and reduce slots allocated to job J for processing its map and reduce tasks, i.e., as a function of (N_M^J, N_R^J) and (S_M^J, S_R^J) . The equation for T_J^{up} can be written similarly.

4.3 Scaling Factors

In the previous section, we showed how to extract the job profile and use it for predicting job completion time when different amounts of resources are used. When the job is executed on a larger dataset the number of map tasks and reduce tasks may be scaled proportionally if the application structure allows it. In some cases, the number of reduce tasks is statically defined, e.g., 24 hours a day, or the number of categories (topics) in Wikipedia, etc. When the job is executed on a larger dataset while the number of reduce tasks is kept constant, the durations of the reduce tasks naturally increase as the size of the intermediate data processed by each reduce task increases. The duration of the map tasks is not impacted because this larger dataset is split into a larger number of map tasks but each map task processes a similar portion of data. The natural attempt might be to derive a single scaling factor for reduce task duration as a function of the amount of processed data, and then use it for the shuffle and reduce phase duration scaling as well. However, this might lead to inaccurate results. The reason is that the shuffle phase

performs data transfer and its duration is mainly defined by the network performance. The reduce phase duration is defined by the application specific computation of the user supplied reduce function and significantly depends on the disk write performance. Thus, the duration scaling in these phases might be different. Consequently, we derive two scaling factors for shuffle and reduce phases separately, each one as a function of the processed dataset size.

Therefore in the staging environment, we perform a set of k experiments ($i = 1, 2, \dots, k$) with a given MapReduce job for processing different size input datasets (while keeping the number of reduce tasks constant), and collect the job profile measurements. We derive scaling factors with linear regression in the following way. Let D_i be the amount of intermediate data for processing per reduce task. Note that we can always estimate this amount of intermediate data from the average amount of input data per map task, the number of job map and reduce tasks, and the selectivity metric in the job profile. Let $Sh_{i,avg}^{typ}$ and $R_{i,avg}$ be the job profile measurements for shuffle and reduce phases respectively. Then, using linear regression, we solve the following sets of equations:

$$C_0^{Sh} + C_1^{Sh} \cdot D_i = Sh_{i,avg}^{typ}, (i = 1, 2, \dots, k) \quad (9)$$

$$C_0^R + C_1^R \cdot D_i = R_{i,avg}, (i = 1, 2, \dots, k) \quad (10)$$

Derived scaling factors (C_0^{Sh}, C_1^{Sh}) for shuffle phase and (C_0^R, C_1^R) for reduce phase are incorporated in the job profile. When job J processes an input dataset that leads to a different amount of intermediate data D_{new} per reduce task, its profile is updated as $Sh_{avg}^{typ} = C_0^{Sh} + C_1^{Sh} \cdot D_{new}$ and $R_{avg} = C_0^R + C_1^R \cdot D_{new}$. Similar scaling factors can be derived for maximum durations Sh_{max}^{typ} and R_{max} as well as for the first shuffle phase measurements. The proposed scaling factors support a general MapReduce performance model where the ratio of map to reduce tasks in the job can vary (or be changed) for different job executions.

For jobs that are routinely executed in the production cluster on new datasets, scaling factors may be derived in an *on-line* fashion by using the same method described above.

4.4 Impact of Failures on the Completion Time Bounds

The performance implications of failures depend on the type of failures (discussed in Section 2). For example, *disk failures* are typical, but their performance implications depend on the amount of data that needs to be reconstructed: for each data block with a number of copies less than the default replication level, Hadoop will reconstruct the additional copies. In this work, we consider a *worker failure* – another typical type of failure which has direct performance implications for a MapReduce job. If the failure happens while the job was running, the failed worker might have completed or in-progress map tasks. All of these map tasks need to be recomputed, since the intermediate data generated by these tasks might be unavailable to current or future reduce tasks. The same applies to the reduce tasks which were in progress on the failed worker: they need to be restarted on a different node. There are reported cases of up to 50% job completion time increase when a single worker failure happens during the reduce stage [5]. The application designer might be interested in evaluating worst case scenarios and deciding on the necessity of additional resources or program changes (e.g., accepting the partial results) as a means of coping with potential failure scenarios.

In order to understand the performance impact of a worker failure on job completion time, we need to consider not only *when* the failure happened, but also whether additional resources in the system can be allocated to the job to compensate for the failed worker. For example, if a worker failure happens in the very beginning of the map stage and the resources of the failed worker are immediately replenished with additional ones, then the lower and upper bounds of job completion time remain practically the same. However, if the failed worker resources are not replenished then the performance bounds are higher. On the other hand, if a worker failure happens during the job's last wave of reduce tasks then all the completed map tasks that reside on the failed node as well as the reduce tasks that were in-progress on this node have to be re-executed, and even if the resources of the failed node are immediately replenished there are serious performance implications of this failure so late during the job execution. The latency for recomputing the map and reduce tasks of the failed node can not be hidden: this computation time is explicitly on the critical path of the job execution and is equivalent of adding entire map and reduce stage latency: $M_{max} + Sh_{max}^{typ} + R_{max}$.

Given the time of failure t_f , we need to quantify the failure impact on the job completion time bounds. Let us consider job J with a given profile, which is partitioned into N_M^J map tasks and N_R^J reduce tasks. Let the worker failure happen at some point of time t_f . There are two possibilities for the job J execution status at the time of failure, it is either in the map or the reduce stage. Using the job profile metrics, we can analyze whether the failure happened during the map or reduce stage. We can do it by estimating the completion time of the map stage (using low or upper bounds of a completion time, or its average). For example, if $t_f \leq T_M^{low}$ then the failure happened during the map stage, otherwise the map stage has been completed and the failure occurred during the job reduce stage. For both scenarios, we need to approximate the number of map and reduce tasks yet to be completed.

- *Case (1)*: Let us assume that the job execution is in the map stage at time t_f (i.e., $t_f \leq T_M^{low}$). In order to determine the number of map tasks yet to be processed, we approximate the number of completed ($N_{M_{done}}^J$) and failed ($N_{M_{fail}}^J$) tasks as follows:

$$N_{M_{done}}^J \cdot M_{avg} / S_M^J = t_f \implies N_{M_{done}}^J = \lfloor t_f \cdot S_M^J / M_{avg} \rfloor$$

If there are W worker nodes in the Hadoop cluster for job J processing and one of them fails, then the number of failed map tasks is:

$$N_{M_{fail}}^J = \lfloor N_{M_{done}}^J / W \rfloor$$

Thus, the number of map and reduce tasks yet to be processed at time t_f (denoted as N_{M,t_f}^J and N_{R,t_f}^J) are determined as follows:

$$N_{M,t_f}^J = N_M^J - N_{M_{done}}^J + N_{M_{fail}}^J \quad \text{and} \quad N_{R,t_f}^J = N_R^J$$

- *Case (2)*: Let us now assume that the map stage is complete, and the job execution is in the reduce stage at time t_f , $t_f \geq T_M^{low}$ and all the map tasks N_M^J are completed. The number of completed reduce tasks $N_{R_{done}}^J$ at time t_f can be evaluated using Eq. 8:

$$B_J^{low} \cdot \frac{N_{R_{done}}^J}{S_R^J} = t_f - C_J^{low} - A_J^{low} \cdot \frac{N_M^J}{S_M^J}$$

Then the number of failed map and reduce tasks can be approximated as:

$$N_{M_{fail}}^J = \lfloor N_M^J / W \rfloor \quad \text{and} \quad N_{R_{fail}}^J = \lfloor N_{R_{done}}^J / W \rfloor$$

The remaining map and reduce tasks of job J yet to be processed at time t_f are determined as follows:

$$N_{M,t_f}^J = N_{M_{fail}}^J \quad \text{and} \quad N_{R,t_f}^J = N_R^J - N_{R_{done}}^J + N_{R_{fail}}^J$$

Let S_{M,t_f}^J and S_{R,t_f}^J be the number of map and reduce slots allocated to job J after the node failure. If the failed resources are not replenished, then the number of map and reduce slots is correspondingly decreased. The number of map and reduce tasks yet to be processed are N_{M,t_f}^J and N_{R,t_f}^J as shown above. Then the performance bounds on the processing time of these tasks can be computed using Eq. 5 and Eq. 6 introduced in Section 4.2. The worker failure is detected only after time δ depending on the value of the *heart beat interval*. Hence, the time bounds are also increased by δ .

This model is easily extensible to estimate the impact of multiple failures.

4.5 SLO-Based Resource Provisioning

When users plan the execution of their MapReduce applications, they often have some *service level objectives* (SLOs) that the job should complete within time T . In order to support the job SLOs, we need to be able to answer a complementary performance question: given a MapReduce job J with input dataset D , how many map and reduce slots need to be allocated to this job so that it finishes within T ?

We observe a *monotonicity property* for MapReduce environments. Let job J complete within time T when we allocate some number of map and reduce slots respectively. Clearly, by allocating a higher number of map and reduce slots to a job (i.e., \hat{M}), one can only decrease the job completion time. In the light of this monotonicity property, we reformulate the problem as follows. Given a MapReduce job J with input dataset D identify *minimal combinations* (S_M^J, S_R^J) of map and reduce slots that can be allocated to job J so that it finishes within time T ? We consider three design choices for answering this question:

1) T is targeted as a *lower bound* of the job completion time. Typically, this leads to the least amount of resources allocated to the job for finishing within deadline T . The lower bound corresponds to an ideal computation under allocated resources and is rarely achievable in real environments.

2) T is targeted as an *upper bound* of the job completion time. Typically, this leads to a more aggressive resource allocations and might lead to a job completion time that is much smaller than T because worst case scenarios are also rare in production settings.

3) Given time T is targeted as the *average* between lower and upper bounds on job completion time. This more balanced resource allocation might provide a solution that enables the job to complete within time T .

Algorithm 1 finds the minimal combinations of map/reduce slots (S_M^J, S_R^J) for one of design choices above, e.g., when T is targeted as a *lower bound* of the job completion time. The algorithm sweeps through the entire range of map slot allocations and finds the corresponding values of reduce slots that are needed to complete the job within time T using a variation of Eq. 8 introduced in Section 4.2. The other cases when T is targeted as the upper bound and the average bound are handled similarly.

Algorithm 1. Resource Allocation Algorithm

Input:Job profile of J $(N_M^J, N_R^J) \leftarrow$ Number of map and reduce tasks of J $(S_M, S_R) \leftarrow$ Total number of map and reduce slots in the cluster $T \leftarrow$ Deadline by which job must be completed**Output:** $P \leftarrow$ Set of plausible resource allocations (S_M^J, S_R^J)

for $S_M^J \leftarrow \text{MIN}(N_M^J, S_M)$ **to 1 do**

 Solve the equation $\frac{A_J^{low} \cdot N_M^J}{S_M^J} + \frac{B_J^{low} \cdot N_R^J}{S_R^J} = T - C_J^{low}$ for S_R^J
if $0 < S_R^J \leq S_R$ **then** $P \leftarrow P \cup (S_M^J, S_R^J)$ **else***// Job cannot be completed within deadline T* *// with the allocated map slots***Break** out of the loop**end if****end for**

Note, that the complexity of the proposed Algorithm 1 is $O(\min(N_M^J, S_m))$ and thus linear in the number of map slots.

5 Evaluation

We perform our experiments on 66 HP DL145 GL3 machines. Each machine has four AMD 2.39MHz cores, 8 GB RAM and two 160GB hard disks (7200rpm SATA). The machines are set up in two racks. The 1Gb network interfaces of the machines in the same rack are connected to a Gigabit Procurve 2650 switch. The racks are interconnected using a ProCurve 2900 switch. We use Hadoop 0.20.2 with two machines as job master and the DFS master. The remaining 64 machines are used as worker nodes, each configured with a single map and reduce slot. The blocksize of the file system is set to 64MB and the replication level is set to 3. We disabled speculation in all our experiments as it did not lead to any significant improvements.

In order to validate our model, we use four representative MapReduce applications:

1. **Twitter:** This application uses the 25GB twitter dataset created by Kwak et. al. [6] containing an edge-list of twitter users. Each edge (i, j) means that user i follows user j . The Twitter application counts the number of asymmetric links in the dataset, that is, $(i, j) \in E$, but $(j, i) \notin E$.

2. **Sort:** The Sort application sorts 64GB of random data generated using random text writer in GridMix2⁵. It uses identity map and reduce tasks, since the framework performs the actual sorting.
3. **WikiTrends:** We use the data from Trending Topics (TT)⁶: Wikipedia article traffic logs that were collected (and compressed) every hour in the months of April to August 2010. Our MapReduce application counts the number of times each article has been visited according to the given input dataset, which is very similar to the job that is run periodically by TT.
4. **WordCount:** It counts the word frequencies in 27 GB of Wikipedia article text corpus. The map task tokenizes each line into words, while the reduce task counts the occurrence of each word.

5.1 Performance Invariants

In our first set of experiments, we aim to validate whether the metrics, that we chose for the inclusion in the job profile, indeed represent performance invariants across different executions of the job on the same input dataset. To this end, we execute our MapReduce jobs on the same datasets and the same Hadoop cluster but with a variable number of map and reduce slots: *i*) 64 map and 32 reduce slots, *ii*) 16 map and 16 reduce slots. The collected job profile metrics are summarized in Table 1. The average input size for WikiTrends is 59.15MB and 64MB for the other applications. We observe that the average duration metrics are within 10% of each other. The maximum durations show slightly higher variance. Each experiment is performed 10 times, and again, collected metrics exhibit less than 10% variation. From these measurements, we conclude that job profile indeed accurately captures application behavior characteristics and reflect the job performance invariants.

Table 1. Job profiles of the four MapReduce applications

Job	Map slots	Reduce slots	Map Task duration (s)			Map Selectivity	1st Shuffle (s)		Typ. Shuffle (s)		Reduce (s)		Reduce Selectivity
			Min	Avg	Max		Avg	Max	Avg	Max	Avg	Max	
Twitter	64	32	26	30	42	3.24	8	11	37	40	22	44	3.2×10^{-8}
	16	16	26	29	43	3.24	7	10	37	41	21	41	3.2×10^{-8}
Sort	64	32	2	5	16	1.00	7	13	30	50	53	75	1.00
	16	16	2	4	14	1.00	8	11	30	51	54	73	1.00
WordCount	64	32	5	34	40	1.31	8	11	24	30	11	14	0.46
	16	16	5	34	41	1.31	7	10	23	28	10	14	0.46
WikiTrends	64	32	66	99	120	9.98	13	27	115	142	26	34	0.37
	16	16	65	98	121	9.98	14	27	113	144	26	32	0.37

5.2 Scaling Factors

We execute WikiTrends and WordCount applications on gradually increasing datasets with a fixed number of reduce tasks for each application. Our intent is to measure the trend of the shuffle and reduce phase durations (average and maximum) and validate the linear regression approach proposed in Section 4.3. The following table below gives the details of the experiments and the resulting co-efficients of linear regression, i.e., scaling factors of shuffle and reduce phase durations derived for these applications.

⁵ <http://hadoop.apache.org/mapreduce/docs/current/gridmix.html>

⁶ <http://trendingtopics.org>

Parameters	WikiTrends	WordCount
Size of input dataset	4.3GB to 70GB	4.3GB to 43GB
Number of map tasks	70 to 1120	70 to 700
Number of reduce tasks	64	64
Number of map, reduce slots	64, 32	64, 32
$C_{0,avg}^{Sh}, C_{1,avg}^{Sh}$	16.08, 2.44	6.92, 0.66
$C_{0,max}^{Sh}, C_{1,max}^{Sh}$	10.75, 2.29	11.28, 0.71
$C_{0,avg}^R, C_{1,avg}^R$	11.45, 0.56	4.09, 0.22
$C_{0,max}^R, C_{1,max}^R$	7.96, 0.43	7.26, 0.24

Figure 5 shows that the trends are indeed linear for WikiTrends and WordCount. While we show multiple points in Fig. 5, typically, at least two points are needed for deriving the scaling factors. The accuracy will improve with accumulated job execution points. Note that the lines do not pass through the origin and hence the durations are not directly proportional to the dataset size. We observe similar results for Twitter and Sort applications but do not include them in the paper due to lack of space.

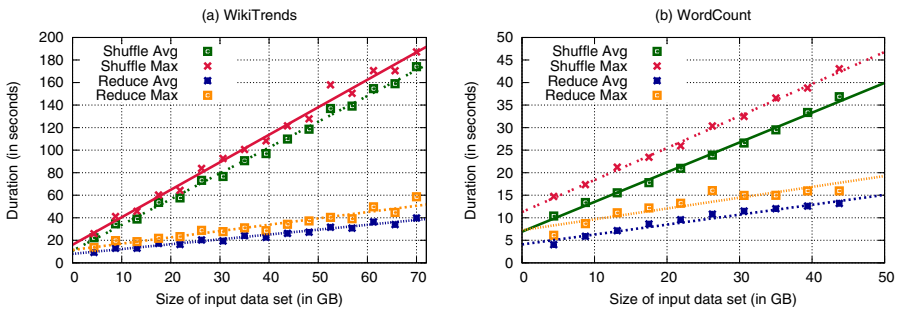


Fig. 5. Linear scaling of shuffle and reduce durations for WikiTrends and WordCount

5.3 Performance Bounds of Job Completion Times

In section 4, we designed performance bounds that can be used for estimating the completion time of MapReduce application with a given job profile. The expectations are that the job profile can be built using a set of job executions for processing small size input datasets, and then this job profile can be used for predicting the completion time of the same application processing a larger input dataset. Therefore, in these experiments, first, the job profiles are built using the three trials on small datasets (e.g., 4.3, 8.7 and 13.1 GB for WordCount) with different numbers of map and reduce slots. After that, by applying linear regression to the extracted job profiles from these runs, we determine the scaling factors for shuffle and reduce phases of our MapReduce jobs. The derived scaling factors are used to represent the job performance characteristics and to extrapolate the duration of the shuffle and reduce phases when the same applications are used for processing larger input datasets with parameters shown in the following table:

Parameters	Twitter	Sort	WikiTrends	WordCount
# of map tasks	370	1024	168	425
# of reduce tasks	64	64	64	64
# of map slots	64	64	64	64
# of reduce slots	16	32	8	8

Finally, by using the updated job profiles and applying the formulae described in Section 4, we predict the job completion times.

The results of these experiments are shown in Figure 6. We observe that our model accurately bounds the measured job completion time, and if we use for prediction the average of lower and upper bounds (denoted T_J^{avg}) then the relative error between T_J^{avg} and the measured job completion time is less than **10%** in all cases. The predicted upper bound on the job completion time T_J^{up} can be used for ensuring SLOs.




Fig. 6. Comparison of predicted and measured job completion times

The solid fill color within the bars in Figure 6 represent the reduce stage duration, while the pattern portion reflects the duration of the map stage. For Sort and WordCount, bounds derived from the profile provide a good estimate for map and reduce stage durations. For WikiTrends, we observe a higher error in the estimation of the durations, mostly, due to the difference in processing of the unequal compressed files as inputs. For Twitter, we observe that even though the predicted minimum and maximum are farther apart because of the difference in average and maximum durations of it phases, the average of the two bounds would provide a good estimate of the measured duration.

The power of the proposed approach is that it offers a compact job profile that can be derived from past executions (while processing smaller dataset) and then used for completion time prediction of the job on a large input dataset while also using different amount of resources assigned to the job.

5.4 SLO-Based Resource Provisioning

In this section, we perform experiments to validate the accuracy of the SLO-based resource provisioning model introduced in Section 4.5. It operates over the following inputs *i*) a job profile built in the staging environment using smaller datasets, *ii*) the targeted amount of input data for processing, *iii*) the required job completion time. We aim to evaluate the accuracy of resource allocations recommended by the model for completing the job within a given deadline.

Figure 7 shows a variety of plausible solutions (the outcome of the SLO-based model) for WordCount, WikiTrends and Twitter with a given deadline $D= 8, 9,$ and

8 minutes respectively. The X and Y axes of the graph show the number of map and reduce slots respectively that need to be allocated in order to meet the job’s deadline. Figure 7 presents three curves that correspond to three possible design choices for computing the required map/reduce slots as discussed in Section 4.5: when the given time T is targeted as the lower bound, upper bound, or the average of the lower and upper bounds. As expected, the recommendation based on the upper bound (worst case scenario) suggests more aggressive resource allocations with a higher number of map and reduce slots as compared to the resource allocation based on the lower bound. The difference in resource allocation is influenced by the difference between the lower and upper bounds. For example, WordCount has very tight bounds which lead to more similar resource allocations based on them. For Twitter, the difference between the lower and upper bounds of completion time estimates is wider, which leads to a larger difference in the resource allocation options.




Fig. 7. Different allocation curves based on bounds for different deadlines




Fig. 8. Do we meet deadlines using the bounds?

Next, we perform a set of experiments with the applications on our 66-node Hadoop cluster. We sample each curve in Figure 7, and execute the applications with recommended allocations of map and reduce slots in our Hadoop cluster to measure the actual job completion times. Figure 8 summarizes the results of these experiments. If we base our resource computation on the lower bound of completion time, it corresponds to the “ideal” scenario. The model based on lower bounds suggests insufficient resource allocations: almost all the job executions with these allocations have missed their deadline. The closest results are obtained if we use the model that is based on the average of lower and upper bounds of completion time. However, in many cases, the measured completion time can exceed a given deadline (by 2-7%). If we base our computation on the upper bounds of completion time, the model over provisions resources. While all the job executions meet their deadline, the measured job completion times are lower

than the target SLO, sometimes by as much as 40%. The resource allocation choice will depend on the user goals and his requirements on how close to a given SLO the job completion time should be. The user considerations might also take into account the service provider charging schema to evaluate the resource allocation alternatives on the curves shown in Figure 7.

5.5 Prediction of Job Completion Time with Failures

In this section, we validate the model for predicting the job completion time with failures introduced in Section 4.4. For this experiment, we set the heartbeat interval to 3s. If a heartbeat is not received in the last 20s, the worker node is assumed to have failed. We use the WikiTrends application which consists of 720 map and 120 reduce tasks. The application is allocated 60 map and 60 reduce slots. The WikiTrends execution with given resources takes $t = 1405s$ to complete under normal circumstances. Figure 9 shows a set of two horizontal lines that correspond to lower and upper bounds of the job completion time under normal case.

Then, using the model with failures introduced in Section 4.4, we compute the lower and upper bounds for job completion time when a failure happens at time t_f (time is represented by X-axes). The model considers two different scenarios: when resources of the failed node are 1) replenished and 2) not replenished. These scenarios are reflected in Figure 9 (a) and (b) respectively. Figure 9 shows the predicted lower and upper bounds (using the model with failures) along with the measured job completion time when the worker process is killed at different points in the course of the job execution.

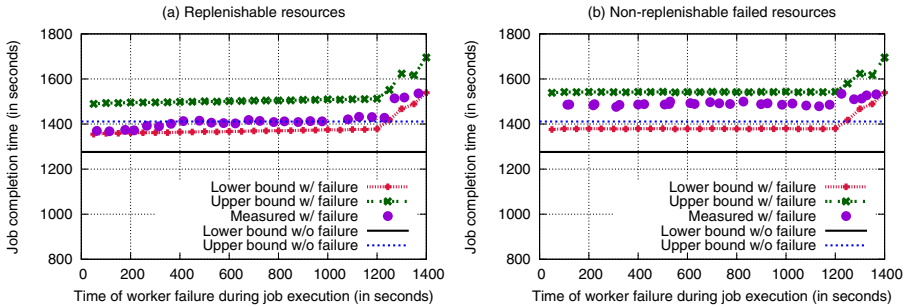


Fig. 9. Model with failures: two cases with replenishable resources and non-replenishable failed resources

The shape of the lines (lower and upper bounds) for job completion time with failures is quite interesting. While the completion time with failures increases compared to the regular case, but this increase is practically constant until approximately $t = 1200s$. The map stage completes at $t = 1220(\pm 10)s$. So, the node failure during the map stage has a relatively mild impact on the overall completion time, especially when the failed resources are replenished. However, if the failure happens in the reduce stage (especially towards the end of the job processing) then it has a more significant impact on the

job completion time even if the failed node resources are replenished. Note that the measured job completion time with failures stays within predicted bounds, and hence the designed model can help the user to estimate the worst case scenario.

6 Related Work

Originally, MapReduce (and its open source implementation Hadoop) was designed for periodically running large batch workloads. With a primary goal of minimizing the job makespan the simple *FIFO* scheduler was very efficient and there was no need for special resource provisioning since the entire cluster resources could be used by the submitted job. As the number of users sharing the same MapReduce cluster increased, a new *Capacity* scheduler [7] was introduced to support more efficient cluster sharing. Capacity scheduler partitions the resources into pools and provides separate queues and priorities for each pool. By enabling partitioning of the cluster resources, the users and system administrators do need to answer an additional question: how much resources do the time-sensitive jobs require and how to translate these requirements in the capacity scheduler settings? This question is still open: there are many research efforts discussed below that aim to design a MapReduce performance model for resource provisioning and predicting the job completion time.

In order to maintain fairness between different users, the *Hadoop Fair Scheduler* (HFS) [8] allocates equal shares to each of the users running the MapReduce jobs. It also tries to maximize data locality by delaying the scheduling of the task, if no local data is available. Similar fairness and data locality goals are pursued in *Quincy* scheduler [9] proposed for the Dryad environment [10]. However, both HFS and Quincy do not provide any special support for achieving the application performance goals and the service level objectives (SLOs). Dynamic proportional share scheduling [11] allows users to bid for map and reduce slots by adjusting their spending over time. While this approach allows dynamically controlled resource allocation, it is driven by economic mechanisms rather than a performance model and/or application profiling.

FLEX [12] extends HFS by proposing a special slot allocation schema that aims to optimize explicitly some given scheduling metric. FLEX relies on the speedup function of the job (for map and reduce stages) that produces the job execution time as a function of the allocated slots. This function aims to represent the application model, but it is not clear how to derive this function for different applications and for different sizes of input datasets. FLEX does not provide a technique for job profiling and detailed MapReduce performance model, but instead uses a set of simplifying assumptions about the job execution, tasks durations and job progress over time.

Polo et al. [13] introduce an online job completion time estimator which can be used for adjusting the resource allocations of different jobs. However, their estimator tracks the progress of the map stage alone and has no information or control over the reduce stage. Ganapathi et al. [14] use Kernel Canonical Correlation Analysis to predict the performance of MapReduce workloads. However, they concentrate on Hive queries and do not attempt to model the actual execution of the MapReduce job. The authors discover the feature vectors through statistical correlation.

Morton et al. [15] propose *ParaTimer* for estimating the progress of parallel queries expressed as Pig scripts [16] that can translate into directed acyclic graphs (DAGs) of MapReduce jobs. In their earlier work [17], they designed *Parallax* – a progress

estimator that aims to predict the completion time of a limited class of Pig queries that translate into a sequence of MapReduce jobs. In both papers, instead of a detailed profiling technique that is designed in our work, the authors rely on earlier debug runs of the same query for estimating throughput of map and reduce stages on the input data samples provided by the user. The approach is based on precomputing the expected schedule of all the tasks, and therefore identifying all the pipelines (sequences of MapReduce jobs) in the query. The approach relies on a simplistic assumption that map (reduce) tasks of the same job have the same duration.

Phan et al. [18] aim to build an optimal schedule for a set of MapReduce jobs with given deadlines. The authors investigate different factors that impact job performance and its completion time such as ratio of slots to core, the number of concurrent jobs, data placement, etc. MapReduce jobs with a single map and reduce waves are considered, and the scheduling problem is formulated as a constraint satisfaction problem (CSP). There are some other simplifications in MapReduce job processing where the data transfer (shuffle and sort) is considered as a separate (intermediate) phase between map and reduce tasks while in reality the shuffle phase overlaps significantly with map stage. All these assumptions and the CSP complexity issues, make it difficult to extend the proposed approach for a general case.

Originally, Hadoop was designed for homogeneous environments. There has been recent interest [19] in heterogeneous MapReduce environments. Our approach and the proposed scaling technique will efficiently work in heterogeneous MapReduce environments. In a heterogeneous cluster, the slower nodes would be reflected in the longer tasks durations, and they all would contribute to the average and maximum task durations in the job profile. While we do not explicitly consider different types of nodes, their performance is reflected in the job profile and used in the future prediction.

In our earlier work, we proposed a framework, called ARIA [20], for a Hadoop deadline-based scheduler which extracts and utilizes the job profiles from the past executions. The shortcoming of the earlier work is that it does not have scaling factors to adjust the extracted profile and lacks the ability for job profiling on smaller datasets. Our earlier workshop paper [21] proposed the idea of scaling factors. This work builds on it providing a detailed evaluation and also models the impact of failures.

Much of the recent work also focuses on anomaly detection, stragglers and outliers control in MapReduce environments [19, 22–24] as well as on optimization and tuning cluster parameters and testbed configuration [25, 26]. While this work is orthogonal to our research, the results are important for performance modeling in MapReduce environments. Providing more reliable, well performing, balanced environment enables reproducible results, consistent job executions and supports more accurate performance modeling and predictions.

7 Conclusion

We have designed a novel framework that aims to enrich private and public clouds offering with an automated SLO-driven resource sizing and provisioning service in MapReduce environments. While there are several companies that offer Hadoop clusters for rent, they do not provide additional performance services to answer a set of typical questions: How much resources the user application needs in order to achieve certain performance goals and complete data processing by a certain deadline. What is the

impact of failures on job completion time? To answer these questions, we introduced a novel automated profiling technique for MapReduce applications by building a compact but representative job profile in a staging environment. The approach allows executing a given application on the set of small input datasets. Then by applying a special scaling technique and designed performance models, one can estimate the resources required for processing a targeted large dataset while meeting given SLOs. We also designed a performance model for estimating the impact of failures on MapReduce applications.

We validated the accuracy of our approach and designed performance models using a set of realistic applications in the 66-node Hadoop cluster. The accuracy of the results depends on resource contention, especially, the network contention in the production Hadoop cluster. In our testbed, the network was not a bottleneck, and it led to the accurate prediction for job completion time. Typically, service providers tend to over provision network resources to avoid undesirable side effects of network contention. At the same time, it is an interesting modeling question whether such a network contention factor can be introduced, measured, and incorporated in the proposed performance models. Another interesting future work is the resource provisioning of more complex applications (e.g., Pig queries) that are defined as a composition of MapReduce jobs and meeting SLO requirements for a given set of MapReduce jobs.

References

1. Dean, J., Ghemawat, S.: MapReduce: Simplified data processing on large clusters. *Communications of the ACM* 51(1), 107–113 (2008)
2. O'Malley, O., Murthy, A.: *Winning a 60 second dash with a yellow elephant* (2009)
3. White, T.: *Hadoop: The Definitive Guide*, Page 6. Yahoo Press
4. Graham, R.: Bounds for certain multiprocessing anomalies. *Bell System Tech. Journal* 45, 1563–1581 (1966)
5. Ko, S., Hoque, I., Cho, B., Indranil Gupta, I.: On Availability of Intermediate Data in Cloud Computations. In: *Proceedings of HotOS 2009* (2009)
6. Kwak, H., Lee, C., Park, H., Moon, S.: What is Twitter, a social network or a news media?. In: *Proc. of Intl. Conference on World Wide Web*, pp. 591–600. ACM (2010)
7. Apache, *Capacity Scheduler Guide* (2010), http://hadoop.apache.org/common/docs/r0.20.1/capacity_scheduler.html
8. Zaharia, M., Borthakur, D., Sen Sarma, J., Elmeleegy, K., Shenker, S., Stoica, I.: Delay scheduling: A simple technique for achieving locality and fairness in cluster scheduling. In: *Proc. of EuroSys*, pp. 265–278. ACM (2010)
9. Isard, M., Prabhakaran, V., Currey, J., Wieder, U., Talwar, K., Goldberg, A.: Quincy: fair scheduling for distributed computing clusters. In: *SOSP. ACM* (2009)
10. Isard, M., Budiu, M., Yu, Y., Birrell, A., Fetterly, D.: Dryad: distributed data-parallel programs from sequential building blocks. *ACM SIGOPS OS Review* 41(3) (2007)
11. Sandholm, T., Lai, K.: Dynamic Proportional Share Scheduling in Hadoop. In: Frachtenberg, E., Schwiegelshohn, U. (eds.) *JSSPP 2010. LNCS*, vol. 6253, pp. 110–131. Springer, Heidelberg (2010)
12. Wolf, J., Rajan, D., Hildrum, K., Khandekar, R., Kumar, V., Parekh, S., Wu, K.-L., Balmin, A.: FLEX: A Slot Allocation Scheduling Optimizer for MapReduce Workloads. In: Gupta, I., Mascolo, C. (eds.) *Middleware 2010. LNCS*, vol. 6452, pp. 1–20. Springer, Heidelberg (2010)
13. Polo, J., Carrera, D., Becerra, Y., Torres, J., Ayguadé, E., Steinder, M., Whalley, I.: Performance-driven task co-scheduling for mapreduce environments. In: *12th IEEE/IFIP Network Operations and Management Symposium* (2010)

14. Ganapathi, A., Chen, Y., Fox, A., Katz, R., Patterson, D.: Statistics-driven workload modeling for the cloud. In: Proceedings of SMDDB (2010)
15. Morton, K., Balazinska, M., Grossman, D.: ParaTimer: a progress indicator for MapReduce DAGs. In: Proceedings of SIGMOD, pp. 507–518. ACM (2010)
16. Olston, C., Reed, B., Srivastava, U., Kumar, R., Tomkins, A.: Pig latin: a not-so-foreign language for data processing. In: Proceedings of SIGMOD, pp. 1099–1110. ACM (2008)
17. Morton, K., Friesen, A., Balazinska, M., Grossman, D.: Estimating the progress of MapReduce pipelines. In: Proceedings of ICDE, pp. 681–684. IEEE (2010)
18. Phan, L., Zhang, Z., Loo, B., Lee, I.: Real-time MapReduce Scheduling. In: Technical Report No. MS-CIS-10-32, University of Pennsylvania (2010)
19. Zaharia, M., Konwinski, A., Joseph, A., Katz, R., Stoica, I.: Improving mapreduce performance in heterogeneous environments. In: OSDI (2008)
20. Verma, A., Cherkasova, L., Campbell, R.: ARIA: Automatic Resource Inference and Allocation for MapReduce Environments. In: Proc. of ICAC (2011)
21. Verma, A., Cherkasova, L., Campbell, R.: SLO-Driven Right-Sizing and Resource Provisioning of MapReduce Jobs. In: Proc. of Workshop on Large Scale Distributed Systems and Middleware, LADIS (2011)
22. Tan, J., Pan, X., Kavulya, S., Marinelli, E., Gandhi, R., Narasimhan, P.: Kahuna: Problem Diagnosis for MapReduce-based Cloud Computing Environments. In: NOMS (2010)
23. Tan, J., Kavulya, S., Gandhi, R., Narasimhan, P.: Visual, Log-Based Causal Tracing for Performance Debugging of MapReduce Systems. In: ICDCS, pp. 795–806. IEEE (2010)
24. Ananthanarayanan, G., Kandula, S., Greenberg, A., Stoica, I., Lu, Y., Saha, B., Harris, E.: Reining in the Outliers in Map-Reduce Clusters using Mantri. In: OSDI (2010)
25. Intel, Optimizing Hadoop* Deployments, (2010), <http://communities.intel.com/docs/DOC-4218>
26. Kambatla, K., Pathak, A., Pucha, H.: Towards optimizing hadoop provisioning in the cloud. In: Proc. of the First Workshop on Hot Topics in Cloud Computing (2009)