

Deadline-Aware Cost Optimization for Spark

Subhajit Sidhanta, Wojciech Golab, and Supratik Mukhopadhyay

Abstract—We present OptEx, a closed-form model of job execution on Apache Spark, a popular parallel processing engine. To the best of our knowledge, OptEx is the first work that analytically models job completion time on Spark. The model can be used to estimate the completion time of a given Spark job on a cloud, with respect to the size of the input dataset, the number of iterations, and the number of nodes comprising the underlying cluster. Experimental results demonstrate that OptEx yields a mean relative error of 6% in estimating the job completion time. Furthermore, the model can be applied for estimating the cost-optimal cluster composition for running a given Spark job on a cloud under a completion deadline specified in the *SLO* (i.e., Service Level Objective). We show experimentally that OptEx is able to correctly estimate the required cluster composition for running a given Spark job under a given SLO deadline with an accuracy of 98%. We also provide a tool which can classify Spark jobs into job categories based on bisimilarity analysis on lineage graphs collected from the given jobs.

Index Terms—distributed systems, parallel processing, distributed file systems, middleware, performance evaluation, reliability, availability, and serviceability

1 INTRODUCTION

OPTIMIZING the cost of usage of cloud resources for running data-intensive jobs on large-scale parallel processing engines is an important, yet relatively less explored problem. Cloud service providers, like Amazon, Rackspace, Microsoft, etc., allow users to outsource the hosting of jobs and services to a cloud using clusters of *virtual machine instances*. The cloud service providers charge a *service usage cost* to the user on the basis of the hourly usage [1] of the virtual machine instances. The cloud service providers present the users with a variety of virtual machine instance types to choose from, such as micro, small, large, etc., for Amazon ec2 [1]. Each virtual machine instance type has a different specification, in terms of CPU, I/O, etc., and different hourly usage cost. The *cost-optimal cluster composition* specifies a number of virtual machine instances (of different virtual machine instance types), that enable execution of the given job under the SLA (i.e., Service Level Agreement) deadline, while minimizing the service usage cost. However most of the current state-of-the-art [2, 3, 4, 5] resource provisioning solutions are designed specifically for Hadoop mapreduce jobs. Other researchers [6, 7] have proposed resource allocation strategies that do not deal with the dual objective of meeting a given SLA deadline for job completion and ensuring that the service usage cost is minimized.

We present OptEx [8], a closed-form job execution model for Apache Spark [9], the most popular parallel processing engine [10, 11]. OptEx can be used to determine a cost-optimal cluster composition, comprising virtual machine instances provided by cloud service providers for executing a given Spark job under an SLA deadline. As far as we know, OptEx is the first work that analytically models job execution on Spark. We consider time-sensitive Spark applications that are developed for popular real world industry use cases such as abandonment of virtual shopping carts [12], search retargeting on web search platforms,

and near real-time content personalization in web content hosting services like Netflix [13]. Such applications often work under strict deadlines such as real-time decision making applications that need to execute near real-time ETL queries [14]. OptEx decomposes the execution of a target Spark job into smaller phases, and models the completion time of each phase in terms of: 1) the cluster size, the number of iterations, the input dataset size, and 2) certain model parameters estimated using *job profiles*. OptEx categorizes Spark applications into job categories, and generates separate job profiles for each job category by executing specific *representative jobs*. The model parameters for the target job are estimated from the components of the job profile corresponding to the job category of the target job. Experimental results demonstrate that OptEx yields a mean relative error of 6% in estimating the job completion time. Using the model of job completion time (OptEx), we derive an objective function for minimizing the service usage cost for running a given Spark job under an SLA deadline. The cost-optimal cluster composition for running the target Spark job under the SLA deadline is obtained using constrained optimization on the above objective function. Experimental results demonstrate that OptEx is able to correctly estimate the cost-optimal cluster composition for running a given Spark job under an SLA deadline with an accuracy of 98%. We also show experimentally that OptEx can be used to provision a cluster to finish a given Spark job within an SLA deadline under the constraints of a given budget. We also provide a tool that can compare a pair of Spark jobs based on bisimilarity analysis of lineage graphs [9]. For lack of space, we have included an extended version of the paper in github [15].

2 SPARK JOB EXECUTION PHASES




Fig. 1: Phases in a Spark Job Execution Flow

- S. Sidhanta is currently with Indian Institute of Technology Jodhpur. E-mail: subhajit@iitj.ac.in
- S. Mukhopadhyay is with Louisiana State University, USA. Email: supratik@csc.lsu.edu
- W. Golab is with University of Waterloo, Canada. Email: wgo-lab@uwaterloo.ca

Manuscript received March 5, 2018; revised May 5, 2018.

TABLE 1: An Example Job Profile: Profile for Data Mining jobs, like movie recommendation, on m4.xlarge instances

App	$T_{mit}(\text{sec})$	$T_{prep}(\text{sec})$	$T_{us}^{baseline}(\text{sec})$	$coeff$	$T_{comm}^{baseline}(\text{sec})$	cf_{comm}	T_{exec}	
							RDD task	$M_a^k(\text{ms})$
ALS	20	13	15	0.004	11	0.070	mean	100
							map	98
							flatmap	72
							first	5
							count	124
							distinct	300

We decompose a typical Spark job execution flow into logically distinct phases illustrated in Figure 1. The first phase in a Spark job is the *initialization* phase, which performs activities like class loading, symbol table creation, object initialization, function loading, and logger initialization. The second phase is the *preparation phase*, which is responsible for job scheduling, resource allocation, and context creation. The initialization and preparation phases are relatively invariant to changes in input variables [9]. The next phase is the *variable sharing phase* that deals with broadcasting or accumulating blocks of data from the Spark master to the workers.

Internally, Spark processes data using a novel in-memory data structure called the *RDD* (i.e., resilient distributed dataset) for fast and fault tolerant computation [9]. Spark provides a wide range of built-in unit RDD operations, packaged within several library modules [16], like MLlib, Spark SQL modules, etc. During the last phase, i.e., the *computation phase* (Figure 1), the given job makes calls to methods from the above library modules, which in turn triggers the respective unit RDD operations on the workers. The computation phase comprises: 1) the *communication phase* that communicates the intermediate variables among the workers, and 2) the *execution phase* that involves the actual execution of the unit RDD operations on the workers. The loops in the rectangles corresponding to the broadcast phase and execution phase represent iterations in iterative jobs. The lengths of the variable sharing phase and the computation phase monotonically increase with the input variables, i.e., the number of iterations, and the dataset size [9]. In particular, the variable sharing phase and the computation phase are repeated under iterations, and the lengths of the above phases increase with respect to number of iterations.

3 JOB CATEGORIZATION AND JOB PROFILES

OptEx categorizes Spark jobs, and following in the footsteps of [3] applies profiling to generate separate job profiles for each job category with representative jobs for each category. As discussed in the extended version, one of the most distinguishing feature of Spark is that it stores the lineage information for the RDDs generated and processed during the execution of a job. We leverage this lineage information to categorize Spark jobs into job categories. Internally, the lineage information for a job is processed in the form of a directed acyclic graph (DAG). Each RDD action invoked from the source code of a given Spark job accepts one or more RDDs, and transforms it into another RDD. Thus, each RDD action corresponds to two nodes in the DAG, a parent node and a child node, connected by a labelled directed edge representing the dependency between the two RDDs. The parent node represents the RDD on which the RDD action is performed. The child node represents the RDD which is created as a result of the action. The parent node, the child node, and the edge between them, are labelled by the name of the parent RDD, name of the child RDD, and the name of the RDD action, respectively. The lineage graphs

are of finite size, bounded by the number of RDD actions in the job, which is typically a finite number. The comment about finite size is applicable even in the case of iterative jobs.

In our case the lineage graph is a particularly interesting artefact as it captures the information about how the data is manipulated in different stages of job execution. We can apply well known techniques from graph theory to analyse the structural equivalence of the lineage graphs of different Spark jobs. Particularly, we can check if the lineage graphs of given jobs are bisimilar, i.e., there exist bisimulations¹ among the graphs, to cluster them into different categories. Checking for bisimulations have been used by researchers [17, 18, 19] as an effective technique to determine equivalence between behavioural characteristics of two programs. Bisimulations have been successfully applied on graphical representations of concurrent program behaviours in the past [20, 21, 22]. Checking for strong and weak bisimulations are known to be *co-NP-hard* and Π_2^P -*hard* problems, respectively [22, 23]. However, Spark DAGs are of finite size, and in most cases the size is not very large. Thus even strong bisimulations can be performed on these DAGs with a considerably low overhead.

We provide a categorization tool (described in Section 7.1) that computes bisimulations on a sample of spark DAGs collected from execution logs of example Spark jobs. This gives rise to a set of bisimilar equivalence classes of spark DAGs. These equivalence classes form the job categories into which we segregate any given Spark job. Each job category is associated with a group of jobs characterised by respective DAGs that are bisimilar to one another. For a given job category, we choose one job from the group of sample jobs associated with it, and label it as the representative job for that category. We can choose any job from the group of sample jobs associated with a given job category as the representative job for that job category. Alterations in this choice would not have caused any changes in our results, since all jobs belonging a given category exhibit the same program behaviour characterised by their corresponding DAG. Then, we can configure the chosen representative job with a standard profiling tool to generate the job profile for the respective job category. We can configure the profiler to execute as a background process on the JVM while running the representative Spark job for a given job category in the foreground (refer to Section 7.3). In the next section, we discuss how we use these job profiles to estimate the OptEx model parameters and coefficients. Before running the bisimulation algorithm, the categorization tool pre-processes the spark DAGs using various techniques which we elaborate in Section 7.1.

1. A bisimulation B is a binary relation $B \subseteq S \times S$ over a labelled transition system $\langle S, \Delta, \rightarrow \rangle$ such that both B and B^{-1} satisfy the following conditions. For each pair of elements $(p, q) \in B, \forall \alpha \in \Delta, \text{ and } \forall p' \in S, p \xrightarrow{\alpha} p'$ implies $\exists q' \in S$ such that $q \xrightarrow{\alpha} q'$ and $(p', q') \in B$. The same condition must be satisfied if we replace the occurrence of B with B^{-1} in the above condition.

TABLE 2: Glossary of symbols and terms

T_{vs}	Estimated completion time for the variable sharing phase	T_{Est}	Estimated job completion time
n	The cluster size	T_{init}	Estimated completion time for the initialization phase
M_a^k	Execution time of the k^{th} RDD operation for job A	T_{prep}	Estimated completion time for the preparation phase
T_{Rec}	Recorded execution time	\mathcal{I}	Number of iterations
$m T_{comp}$	Estimated completion time for the computation phase	T_{commn}	Estimated completion time for the communication phase in T_{comp}
$T_{vs}^{baseline}$	Baseline value of T_{vs}	$coeff$	Coefficient of T_{vs} in T_{Est}
$T_{commn}^{baseline}$	Baseline value of T_{commn}	cf_{commn}	Coefficient of T_{commn} in T_{Est}

4 OPTEx MODEL BASED ON JOB PROFILES

A Job profile consists of values of different parameters and coefficients collected during execution of a given job. Components of the job profile are used as estimates for the model parameters of the target job. The length of the initialization phase (refer to Figure 1)) T_{init} and the length of the preparation phase T_{prep} remain constant to variations in the input variables [9]. The length of the execution phase (T_{exec}) and the length of the variable sharing phase (T_{vs}) increase monotonically with respect to the input variables [9]. Thus, the length of these phases in the execution of a representative job (contained in the job profile) can act as the point of reference, i.e., *baseline*, for measuring the length of the corresponding phases in the target job [3]. In this section, we elaborate how these baseline values in the job profile can be used for estimating the parameters of the model for each job phase.

During profiling, the representative job a is run on a single node, and the length of the initialization phase, the preparation phase, the variable sharing phase, and the communication phase (Figure 1) is recorded in the job profile. The lengths of the above phases in the job profile act as baseline values for estimating the lengths of the corresponding phases in a given target job. The length of the initialization phase T_{init} (Table 1) and the length of the preparation phase T_{prep} for a given job are directly estimated from the lengths of the corresponding phases in the job profile (since, as discussed in Section 2, these phases remain constant with respect to the variations in the input variables). As elaborated in Section 2, the length of the variable sharing phase T_{vs} increases monotonically [9] with respect to the cluster size and the number of iterations. Hence, T_{vs} is expressed as a function of:

- The input variable n represents the number of nodes.
- The input variable \mathcal{I} represents the number of iterations.
- The baseline value $T_{vs}^{baseline}$, contained in the job profile, representing the length of the variable sharing phase of the representative job. It is the baseline for estimating the length of the variable sharing phase T_{vs} of a given target job. ble

OptEx computes the total duration of the variable sharing phase across all iterations. The length of the variable sharing phase T_{vs} is a function of the baseline $T_{vs}^{baseline}$, the input variables n , and number of iterations \mathcal{I} . Thus T_{vs} is expressed as:

$$T_{vs} = coeff \times \mathcal{I} \times n \times T_{vs}^{baseline}, \quad (1)$$

where \mathcal{I} is the number of iterations, n is the number of nodes, $T_{vs}^{baseline}$ is the baseline value, and $coeff$ is a coefficient term. The

coefficient term $coeff$ is empirically estimated during job profiling using curve fitting on the results of repetitive experiments with the representative job. The length of the computation phase T_{comp} (Table 2) is made up of two logical components: the length of the communication phase T_{commn} , and the length of the execution phase T_{exec} (Figure 1).

The communication phase is responsible for fetching the values of the intermediate variables computed by operations in the earlier stages of the given job. While profiling, the length of the communication phase of a representative job a on a single node is recorded in the job profile as $T_{commn}^{baseline}$. It serves as the baseline measure against which the length of the communication phase T_{commn} is estimated. The size s of the input dataset is given in bytes (for example, the size of the input for the wordcount job is given as the combined size of the input files). Since the length of the communication phase T_{commn} increases with respect to the input variable s [9], T_{commn} is expressed as a product of the input dataset size s , a coefficient cf_{commn} , and the baseline value $T_{commn}^{baseline}$, where cf_{commn} and $T_{commn}^{baseline}$ are obtained from the job profile. Again, the coefficient cf_{commn} is empirically estimated in the profiling stage by applying curve fitting on the outputs of experiments with the representative jobs. Thus,

$$T_{commn} = cf_{commn} \times T_{commn}^{baseline} \times s \quad (2)$$

As discussed in Section 3, the execution phase of a Spark job comprises a permutation of unit RDD operations. The OptEx job profile records the average running time M_a^k (Table 2) of each unit RDD operation component k comprising the representative Spark job a . If there are multiple occurrences of an RDD operation i in a , we consider the average running time for all occurrences of the operation i . By design, the representative job for a job category contains all the unit RDD operations comprising any given job in that category. Hence, the length of the execution phase of the given Spark job is estimated as a function of the average running time M_a^k of each unit RDD operation k in the job profile.

5 DERIVATION OF THE OPTEx MODEL

Input Variables: OptEx accepts the following input variables: the size s of the input dataset in bytes, the number of nodes n in the cluster, and the number of iterations \mathcal{I} in the given job [9]. OptEx only considers those iterative jobs for which the number of iterations is either passed as a runtime argument by the developer or can be determined from the source code [9]. Moreover, Spark jobs typically have only few lines of code. Hence if we need to determine $iter$ from the code, we do not require sophisticated techniques involving static analysis [24]. The other input variables, i.e., number of nodes n and input dataset size s , are also directly passed to the model as runtime arguments.

While modelling the estimated total completion time for the target job, the user provides an estimated upper bound for the number of iterations \mathcal{I} for the target job, as an input to the model. During the actual running time of the target job, the user provides the number of iterations \mathcal{I}_{exec} as a runtime argument to the job [9]. The number of iterations \mathcal{I}_{exec} provided in the running time may differ from the number of iterations \mathcal{I} provided in the modelling phase. The difference between \mathcal{I}_{exec} and \mathcal{I} may cause: 1) unpredicted wastage of cluster resources, and 2) the failure to satisfy the SLO. In that case, the estimations need to be redone, with a new input value for the number of iterations. For multiple runs of the target job with different values of the runtime argument

\mathcal{I}_{exec} supplied by the user in each run, the maximum of the \mathcal{I}_{exec} values, i.e., \mathcal{I}_{exec}^{max} , is supplied as the new input for the estimation. The estimation using the new value \mathcal{I}_{exec}^{max} amounts to computing the value of T_{Est} from the Equation 8 with a time complexity of $\Theta(1)$ (since the degree of T_{Est} is 1 [25]), thus incurring negligible overhead.

Formulation of the Model: OptEx decomposes the job completion time into four phases (Figure 1), and models the total job completion time T_{Est} as the sum of the lengths of the component phases. Thus,

$$T_{Est} = T_{Init} + T_{prep} + T_{vs} + T_{comp}, \quad (3)$$

where T_{Init} , T_{prep} , T_{vs} , and T_{comp} are the lengths of the initialization phase, preparation phase, variable sharing phase, and computation phase, respectively. As discussed in Section 7.3, the execution phase of a given Spark job comprises a permutation of low-level unit RDD operations. The number of unit RDD operations n_{unit} increases monotonically with increasing input dataset size s and number of iterations \mathcal{I} [9]. Hence, the number of unit RDD operations n_{unit} can be expressed as a function of the size of the input dataset denoted by s , the number of iterations in the job denoted by \mathcal{I} , and the baseline term for the number of unit RDD operations denoted by $n_{unit}^{baseline}$. The value of the term $n_{unit}^{baseline}$ is obtained from the job profile (Section 4) as follows. Spark enables parallel execution by dividing the input dataset into partitions, and distributing the partitions/slices among the worker nodes [9]. $n_{unit}^{baseline}$ directly corresponds to the number of partitions that the input dataset is comprised of. The number of partitions can be: 1) computed from the size s of the input dataset and the number of iterations \mathcal{I} [9], or 2) programmatically provided as a parameter to the built-in transformation method used to create the RDDs from the input dataset [9].

For example, the Spark ChiSqTest program, working on input files from a HDFS backend, divides the input dataset into as many partitions as the number of HDFS blocks comprising the input files. Consider a portion of the Popular Kids dataset in HDFS [26] consisting of 164 files, where the size of each file is less than the HDFS block size. Hence the number of partitions, and in turn the number of unit RDD operations is 164. Thus, the baseline $n_{unit}^{baseline}$ is 164. Thus, the increase of n_{unit} , with respect to the above baseline $n_{unit}^{baseline}$, in terms of the parameters s and \mathcal{I} , is expressed as

$$n_{unit} = n_{unit}^{baseline} \times s \times \mathcal{I}. \quad (4)$$

As discussed already in Section 4, the length of the initialization phase T_{Init} and the length of the preparation phase T_{prep} are directly estimated from the corresponding components in the job profile (Section 4). As discussed in Section 2, the length of the variable sharing phase T_{vs} and the length of the computation phase T_{comp} vary with the input variables. Hence the length of the variable sharing phase T_{vs} (Equation 1) and the length of the computation phase T_{comp} are estimated as functions of the job profile components, and the input variables (Section 4). The expression for the length of the variable sharing phase T_{vs} , comprising the baseline value $T_{vs}^{baseline}$ and coefficient $coeff$ obtained from the job profile, is given by Equation 1.

The length of the computation phase T_{comp} in Equation 3 can be further decomposed into the following two logical components: **A)** T_{commn} : the length of the communication phase T_{commn} is obtained from the Equation 2, and **B)** T_{exec} : the length of the execution phase T_{exec} in Equation 3 comprises the actual

execution of k RDD operations comprising the job on the worker nodes (Section 4). T_{exec} depends on various factors [9]: 1) the running times of the unit RDD operations comprising the given job, 2) the number of iterations \mathcal{I} , 3) the number of stages in the job, 4) parallelization of the job across the worker nodes, and 5) sharing of the RDD variables across the cluster. Hence, execution phase length T_{exec} is expressed as the sum over the estimated computation times of all unit RDD operations comprising j , along with coefficients accounting for the above factors. Thus, the length of the execution phase T_{exec} of job a , without taking into account the parallelization factor n , is given as:

$$T_{exec} = \mathcal{I} \times \sum_{k=1}^{n_{unit}} M_a^k, \quad (5)$$

where n_{unit} is the number of unit RDD operations given in Equation 4, M_a^k is the average job execution time of a unit RDD operation k comprising the job a , and \mathcal{I} is the number of iterations in the job.

Following prior work on modelling execution of parallel operations [3], the overall length of the computation phase T_{comp} is divided by the factor n , taking into account the parallelization of the given job across the n worker nodes. Thus, the computation phase is rewritten as the sum of its two components, divided by n :

$$T_{comp} = \frac{T_{commn} + T_{exec}}{n}. \quad (6)$$

Combining the Equations 2, 5, and 6, we get

$$T_{comp} = \mathcal{I} \times \sum_{k=1}^{n_{unit}} M_a^k / n + \frac{A \times s}{n}, \quad (7)$$

where n_{unit} is the number of unit RDD operations given in Equation 4, and $A = \frac{cf_{commn} \times T_{commn}^{baseline}}{S_{baseline}}$. Finally, combining the Equations 1 and 7 in Equation 3, the estimated total completion time for the target job is given as

$$T_{Est} = T_{Init} + T_{prep} + n \times \mathcal{I} \times C + \mathcal{I} \times B / n + \frac{A \times s}{n}, \quad (8)$$

where n_{unit} is the number of unit RDD operations given in Equation 4, $A = \frac{cf_{commn} \times T_{commn}^{baseline}}{S_{baseline}}$, $B = \sum_{k=1}^{n_{unit}} M_a^k$, and $C = coeff \times T_{vs}^{baseline}$.

6 COST OPTIMAL CLUSTER COMPOSITION

OptEx models the completion time T_{Est} (Equation 8) of a Spark job on a cluster comprising virtual machine instances provisioned from a cloud service provider, like Amazon (EC2), RackSpace, Microsoft, etc. The OptEx model is further used to estimate the cost optimal cluster composition for running a given job on virtual machine instances provided by any cloud provider, under the job completion deadline specified in the SLO, while minimizing the service usage cost. Let the optimal cluster size be given as $n = \sum_{t=1}^m n_t$, where n_t is the number of virtual machine instances of type t , and m is the total number of possible machine instance types (m depends on the instance offerings of the chosen cloud provider). Let total service usage cost of running the given job on the cloud be denoted by \mathcal{C} . Let c_t be the hourly cost of each machine instance of type t (c_t depends on the current rates charged by the chosen cloud provider), and T_{Est} be the estimated completion time of the given job (Equation 8). Our objective is to determine the cost optimal cluster composition for finishing the

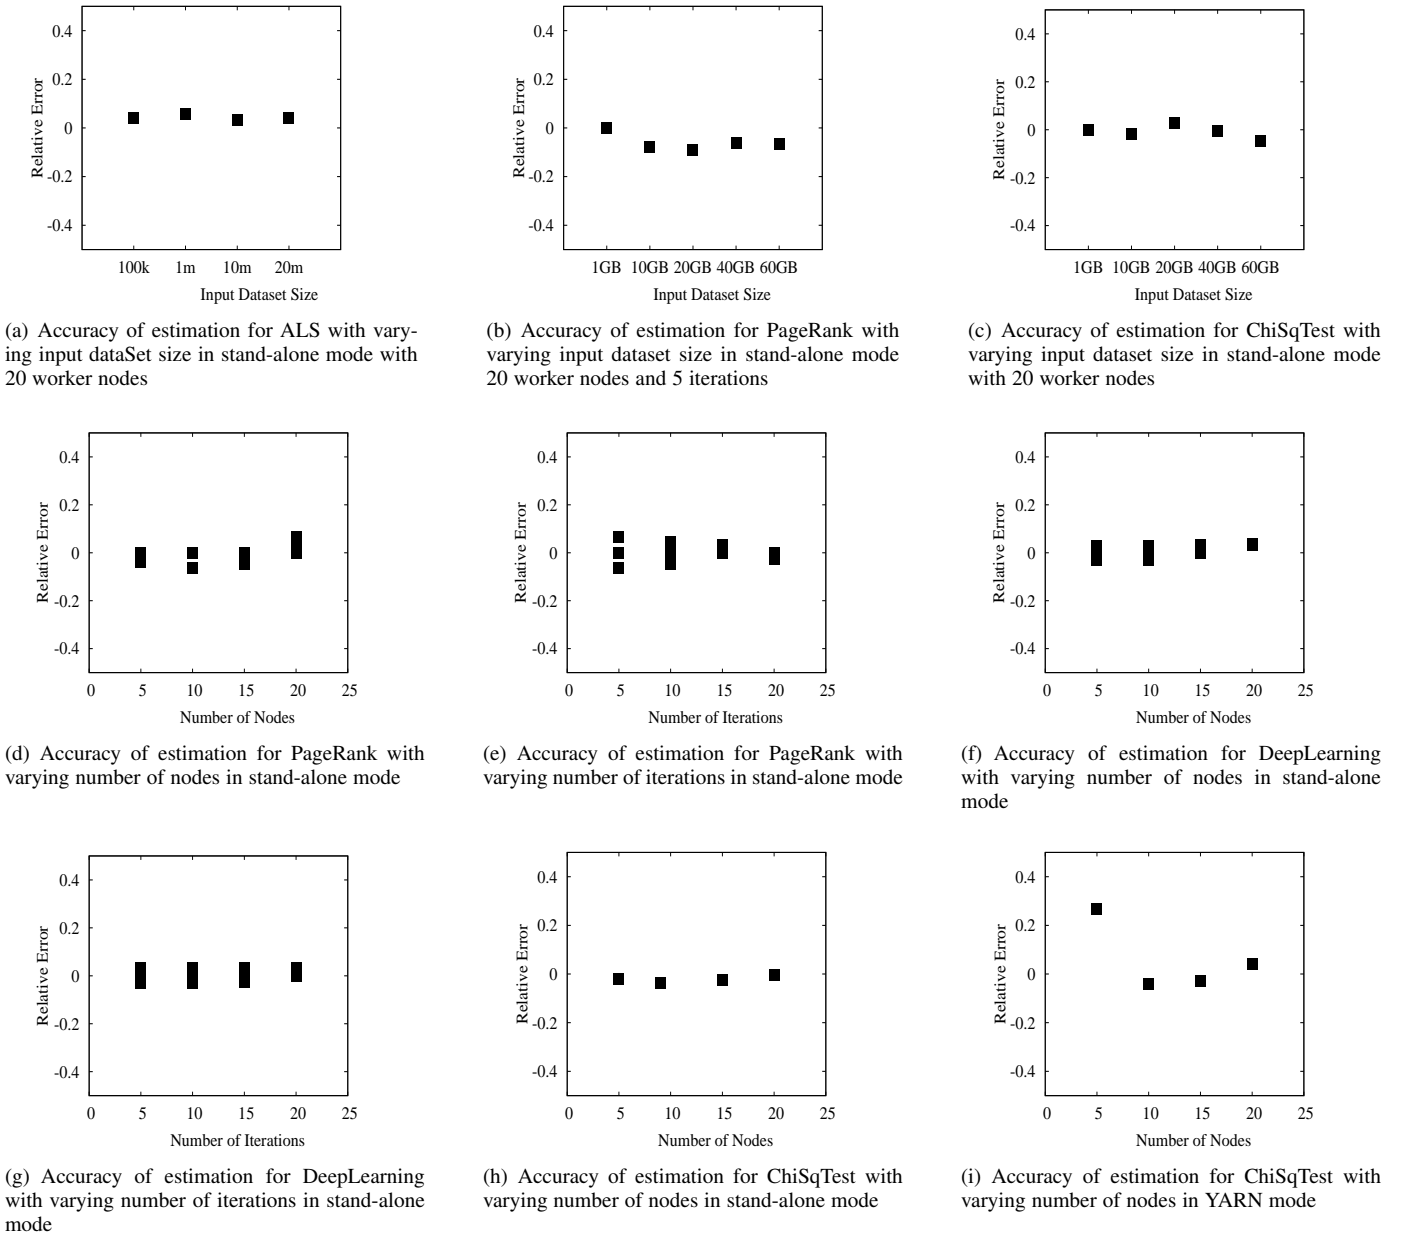


Fig. 2: Accuracy Of estimations against varying input dataset size, number of nodes and iterations

TABLE 3: Relative Error of Estimation of Job Execution Times With the Version of the SparkRDDAnalyze Tool

SparkRDDAnalyze Version	Standalone				YARN			
	Data Mining	Statistical Evaluation	Quantitative Analysis	Miscellaneous	Data Mining	Statistical Evaluation	Quantitative Analysis	Miscellaneous
V1	0.7	0.5	0.5	0.65	0.6	0.55	0.45	0.65
V2	0.3	0.2	0.1	0.25	0.2	0.15	0.2	0.3

given Spark job within an SLO deadline with minimum service usage cost. This goal can be mathematically stated as: optimize the objective function

$$C = \sum_{t=1}^m c_t \times n_t \times T_{Est}, \quad (9)$$

and obtain the cost optimal cluster composition, given as $N_t =$

$$\{n_t \mid 1 \leq t \leq m\},$$

under the constraint $T_{Est} < SLO$, where SLO is the given deadline, and T_{Est} is estimated using Equation 8.

We optimize the above objective function (Equation 9) and determine an optimal cluster configuration given by N_t , under the constraint $T_{Est} < SLO$. The above constraint involving T_{Est} is a convex nonlinear function [15] over n and \mathcal{I} (Equation 8). The above optimization problem of minimizing the cost C under the nonlinear constraint $T_{Est} < SLO$ is solved using the Interior Point algorithm [25] which is one of the most efficient linear programming solvers till date. The solution to the above optimization problem enables: 1) estimating whether a given job will finish under the deadline SLO , 2) optimal job scheduling

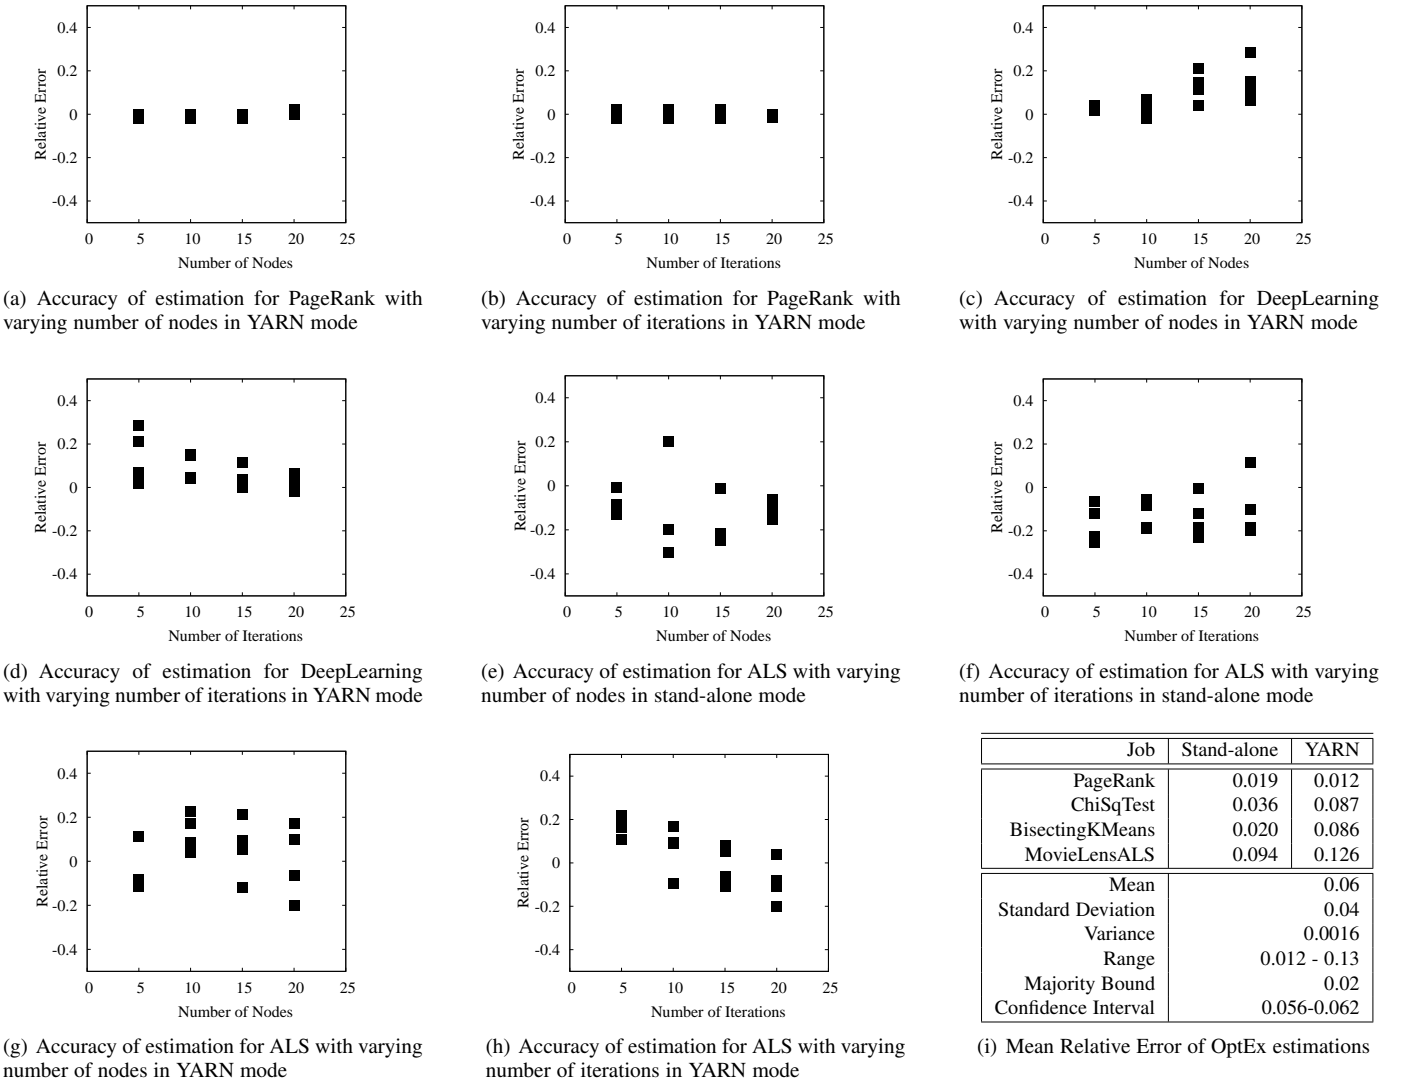


Fig. 3: More accuracy results and the observed Mean Relative Error statistics

under the given deadline SLO , while minimizing cost C , and 3) estimating optimal cluster composition, given a cost budget \mathcal{C} and an SLO .

7 IMPLEMENTATION AND EVALUATION

7.1 Implementation: Job Categorization Tool

We provide SparkRDDAnalyze, an open-source tool for determining equivalence relationship between a pair of given Spark applications based on their respective DAGs [27]. This tool is used to categorize Spark applications into job categories, and select representative jobs for each category. The tool uses the Construction and Analysis of Distributed Processes (CADP) toolbox [28]. The BISIMULATOR tool [29] comprised in CADP is used to check equivalence of pairs of DAGs. The tool is widely used by researchers in the programming languages and verification community [30]. We encode the DAGs in Binary Coded Graphs (BCG) file format, which is a popular format for representing Labelled State Transition (LTS) systems such as the Spark DAGs. Then we use the bcg_cmp method of BISIMULATOR to compare the two BCG files and check for bisimilarity between the two

respective DAGs. Depending on the boolean output of bcg_cmp , we segregate given sample of Spark jobs into equivalence classes, which we refer to as job categories in this paper.

We have implemented two versions of the SparkRDDAnalyze tool, namely SparkRDDAnalyze-unlabelled and SparkRDDAnalyze-labelled. In SparkRDDAnalyze-unlabelled, we omit the edge labels of the DAGs in the pre-processing step before running the bcg_cmp method from the CADP tool. In SparkRDDAnalyze-labelled, we relabel the DAG edges in the pre-processing step according to the following rule. The Spark programming guide lists all available RDD transformations available in a current release of Spark, and includes a brief description of each transformation [31]. If a pair of RDD transformations corresponding to a given pair of DAG edge labels perform an identical task on a given input RDD, and produces as output a particular RDD dataset, according to the programming guide, these RDD transformations are considered equivalent, and thus the two edge labels can be placed interchangeably in the two DAGs. For example, the transformations map , $flatMap$, $mapPartitions$, and $mapPartitionsWithIndex$ are equivalent because inherently they perform the same core task, i.e., they process each

element in a given dataset using a function $func$, and returns a new distributed dataset.

7.2 Experimental Setup and Procedure

The experimental setup consists of Apache Spark version 2.3.1, built-in within the Cloudera CDH 6.0.1 package, on a cluster of m4.xlarge Amazon EC2 machine instances, each comprising 4 cores, 16 GB of RAM, 10 GB EBS with a dedicated EBS bandwidth of 750 Mbps. Each VM runs RedHat Enterprise Linux version 7. Since Spark performs in-memory data processing and Spark applications typically are compute-intensive, we choose VM instances that provide a healthy balance of compute, memory and network resources. However since minimizing the hourly usage in the public cloud is the primary objective of this paper, we choose m4.xlarge instances which is the cheapest VM type that satisfies the above criteria [1]. We configure the instances with general purpose SSD volumes which are able to provide burst throughput upto 3000 IOPS. Instead of EBS storage, we use the local HDFS storage as the backend for storing and processing the input dataset in the local persistent storage of the EC2 instances. Each Spark worker is configured with an executor memory of 2 GB. The above configurations result in minimizing the degree of network I/O and throughput interference.

We use the Interior point algorithm [25] from the Optimization toolbox of the Matlab version 2013b for solving the given non-linear convex optimization problem (Section 6), and estimating the cost-optimal cluster composition. We use the 10-M MovieLens dataset obtained from grouplens.org [32] as the input workload for the MovieLensALS job. PageRank is evaluated with the social network dataset collected from LiveJournal [33], an online community comprising roughly 10 million members. The LiveJournal dataset has 48,47571 nodes and 6,8,993773 edges. The input workload for the ChiSqTest job is the dataset "Popular Kids" available in the Statlib Data and Story Library (DASL) [26]. The input workload for the Deep Learning job is the popular MNIST dataset for hand written digits available at yann.lecun.com comprising 60,000 training samples and 10,000 testing samples.

7.3 Evaluation of the Job Categorization Scheme

We run the SparkRDDAnalyze tool over a sample of 30 Spark jobs packaged with the Spark v2.2.1 source code to generate different equivalent classes, which correspond to different job categories. Following [3], we obtain job profile for each category by running a representative job for each category. We estimate the components of the job profile from the snapshots of the execution flow of the representative job obtained using YourKit Java Profiler [34]. Table 3 illustrates the accuracy of estimation of job execution times using the two versions of the tool SparkRDDAnalyze, namely SparkRDDAnalyze-unlabelled (referred to as V1) and SparkRDDAnalyze-labelled (referred to as V2), respectively. From the estimated completion time T_{Est} and the recorded (i.e., observed) completion time T_{Rec} , we compute the relative error of estimation $RE = (T_{Est} - T_{Rec})/T_{Rec}$ with a particular example jobs chosen as the representative job for that job category. We estimate job completion times of 10 target jobs for each job category with representative jobs chosen according to a scheme discussed later in this section. We record the relative error RE for each target job, and finally compute the mean relative error for all target jobs in a particular job category. Subsequently, we express the accuracy in terms of mean relative error for each respective

job category in Table 3. Mean relative error measures overall how close the job completion time estimated with the tool is to the observed job completion time.

Due to the absence of a non-trivial labelling scheme in V1 (i.e., SparkRDDAnalyze-unlabelled), the tool considers edges in a pair of DAGs as dissimilar if and only if their labels are distinct from a lexicographical sense, although essentially those edges might perform an identical set of tasks. According to this scheme, each job in the sample has a unique RDD DAG, and there are as many job categories as there are jobs in the sample. Thus V1 groups Spark jobs into 30 job categories since we have considered a sample size of 30 for the purpose of our experiments. Hence a given Spark job is rarely similar to a given representative job in any job category with SparkRDDAnalyze-unlabelled, which results in considerably high relative error in estimating the execution time of a given job using the job profile of a chosen representative job. On the other hand, SparkRDDAnalyze-labelled produces much more accurate categorization on the same group of DAGs by relabelling the edge labels appropriately as discussed in Section 7.1. V2 (i.e., SparkRDDAnalyze-labelled) produces 4 job categories by relabelling the edge labels in the RDD DAGs according to the scheme described in Section 7.1. Since the number of equivalent classes is smaller with V2, the chances are higher for a target job being correctly estimated with job profile of a representative job chosen from a particular job category. Due to the above reason, in Table 3 the relative errors with SparkRDDAnalyze-unlabelled are significantly higher than relative errors with SparkRDDAnalyze-labelled.

Specifically, V2 of the tool generates the following equivalence classes: 1) Data Mining, comprising jobs that perform typical data mining tasks like recommendation engines or jobs that run mining queries on logs, 2) Statistical Evaluation, comprising jobs that perform typical statistical hypothesis tests on given input datasets based on different statistical evaluation techniques and standard test metrics, such as Chi Square test, etc., 3) Quantitative Analysis, comprising jobs like pagerank and logistic regression, which perform typical quantitative analysis tasks on a given dataset, and 4) Miscellaneous, comprising any other jobs that can not be grouped into any of the above categories like word count. For the job category Statistical Evaluation we use the job ChiSqTest from the group of example applications packaged with the spark source code. The representative job for the Data Mining group of jobs is the movie rating job MovieLensALS [16]. The Deep Learning Pipelining framework [35] provides a natural construct for applying distributed learning on a large-scale dataset. As a representative of heavily compute intensive MLlib applications, we ran a handwritten digit recognition application based on the Deep Learning Pipelining on the MNIST dataset from yann.lecun.com. The application trains an InceptionV3 convolutional neural network, and applies the DeepImageFeaturizer algorithm to filter off the last layer of the InceptionV3 network. The output from the remaining layers of the network acts as features for training a logistic regression model with elastic net regularization. PageRank is the representative job for the Quantitative Analysis group of jobs [16]. We could have chosen any other job from the group of sample jobs associated with the corresponding equivalent class as the representative job for the respective job category without any effect on the mean relative error. In Table 3 the relative error with SparkRDDAnalyze-labelled varies from 0.1 to 0.3, with a mean relative error of 0.21 and a standard deviation of 0.18. The difference in relative error with the two versions V1 and V2 ranges

from 0.25 with Quantitative Analysis jobs to the extent of 0.4 with Data Mining and Statistical Evaluation jobs. In the case of V2, mean relative error of 0.21 demonstrates that the estimation with V2 is fairly accurate. The fact that the relative error is smaller with version V2 of the tool than with version V1 signifies that the profile parameters estimated using the representative job chosen with V2 produces a more accurate estimate of the job completion time than the estimates obtained with V1.

7.4 Accuracy of the Estimations Using OptEx

Being the first work in modelling Spark jobs, OptEx has no prior baseline to compare with. However, we demonstrate (see Figures 2 and 3) that OptEx provides accurate (i.e., average relative error 0.06) estimations of the job completion time against variations in all the input parameters of the model (i.e, against increasing size of dataset, number of nodes, and the number of iterations), and on jobs of different categories. The Figures 2(a), 2(b), and 2(c) illustrate the variations in the relative error RE , for the MovieLensALS, PageRank, and ChiSqTest jobs, with increasing size s of the input dataset. The figures 2(d), 2(f), 2(h), and 3(e) illustrate the variations in the relative error RE , for PageRank, DeepLearning, ChiSqTest, and MovieLensALS, against varying cluster size n , in the stand-alone mode. The figures 2(e), 2(g), and 3(f) illustrate the variations in the relative error RE for the same jobs, i.e., PageRank, DeepLearning, and MovieLensALS, respectively, against varying number of iterations \mathcal{I} , in the stand-alone mode. Figures 2(i), 3(a), 3(c), and 3(g) illustrate the variations in the relative error RE for ChiSqTest, PageRank, DeepLearning, and MovieLensALS with varying n in the YARN mode. Figures 3(b), 3(d), and 3(h) illustrate the variations in the relative error RE for PageRank, DeepLearning, and MovieLensALS against varying \mathcal{I} , in YARN mode. We evaluate the OptEx model with the mean relative error metric [25] $\delta = \frac{\sum_{j=1}^k |T_{Est} - T_{Rec}| / T_{Rec}}{k}$, where k is the total number of jobs submitted. The absolute differences between T_{Est} and T_{Rec} eliminate the signs in the error, and gives the magnitudes of the errors. The error values δ are given in the Table 3(i). The average δ score for all the cases is 0.06, i.e., 6%, which is comparable to prior work on performance modelling of state-of-the-art data analytics engines [3].

7.5 Analysis of the Results

The magnitude of the relative error RE for the experiments with ChiSqTest, PageRank, and DeepLearning jobs (Figure 2 is strictly within 0-0.06, bounded by a 95% confidence interval of 0.056-0.062 (refer to Table 3(i)) for the Statistical Evaluation and Data Mining job categories (refer to Section 3), in stand-alone mode. The only outlier is the ChiSqTest run in YARN mode in Figure 2(i) (reasons given later in this section). The experiments with increasing dataset size yield relative error of magnitude between 0.007 to 0.05 (Figures 2(a), 2(b), 2(c)).

The estimated execution time T_{Est} of a Spark job comprises two components X_1 and X_2 , where $X_1 = T_{Init} + T_{prep}$ and $X_2 = n \times \mathcal{I} \times C + \mathcal{I} \times B/n + \frac{A \times s}{n}$ (refer to Equation 8). The first component X_1 is independent of variations in the values of the input variables. The second component X_2 comprises the last three phases of the Spark job execution (refer to Section 5), each phase varying differently with respect to the input variables n , \mathcal{I} , and s (refer to Section 7.7). Hence, X_2 accounts for the observed random variations in the relative error, with respect to variations in the input variables (Figures 2 and 3). The execution

phase in X_2 encompass the execution of the job stages, comprising unit RDD operations, on the worker nodes (Section 5). The execution of the job stages on the workers is inherently non-deterministic (unpredictable) in nature, due to the dependency on various components of the Spark cluster, like the driver, the cluster manager, the workers, etc., and their configuration [9]. The job stages may get unpredictably delayed, i.e., stages can fail and be retried by the master repeatedly, due to various factors like momentary unavailability of required resources, delays in allocation of resources by the master, communication delays among the workers, etc., [9]. The above unpredictable delays in the job stages, however small, can cause the observed values of X_2 to deviate randomly from the estimated values of X_2 , while X_1 stays constant (Section 5). This, in turn, causes the overall observed completion time T_{Rec} of the job, to vary unpredictably with respect to the estimated job completion time T_{Est} , estimated from X_1 and X_2 (Section 5). This causes the observed random variations in the values of the relative error (i.e., $RE = T_{Est} - T_{Rec}$), though still bounded by the confidence interval of 0.056-0.062 (Figures 2 and 3). The relative error increases slightly with increasing number of nodes (Figures 2(a), 2(d), 2(f), 2(h), 2(i), 3(a), 3(c), 3(e), and 3(g)). Increase in number of worker nodes augments the chances of unpredictable failures of the job stages due to dependency on communication between a larger number of nodes, causing unpredictable variations in the component X_2 of the overall job completion time T_{Est} . This, in turn, causes the observed job completion time T_{Rec} to deviate more unpredictably from the estimated completion time T_{Est} , estimated from X_1 and X_2 . The result is greater variation in the relative error (i.e., $RE = T_{Est} - T_{Rec}$) with increasing number of nodes. Our goal is to provide correct estimations for SLA-driven user-facing jobs. Few user-facing jobs, that work under an SLA deadline, will require more than 50 nodes [36, 37]. OptEx can provide estimations for such typical SLA-driven user-facing jobs with a relative error close to 0 (Figures 2 and 3). Jobs that do not meet this criteria are batch processing jobs, like bioinformatics, genomics, data analytics jobs, etc., which typically do not work under a deadline [38].

For experiments run in YARN mode (Figures 3(a), 3(b), 3(c), 3(d), 3(g), and 3(h)), the variations in the observed relative error, with respect to the variations in the input variables, are noticeably larger than the experiments run in stand-alone mode. In YARN mode, the submitted jobs are additionally dependent on the YARN resource manager to allocate resources, and to execute the jobs on the workers [39]. Hence, the chances of unpredictable delays in the intermediate stages of a job are greater in YARN mode due to unreliability of the communication channel between the YARN resource manager and the Spark master [39]. Thus, the chances of observing randomness in the relative error is greater for jobs run in the YARN mode, though the magnitude of the average relative error does not exceed 0.04. Further, the relative error, for YARN mode, is even closer to 0 for jobs with number of iterations larger than 10 (Figures 3(a), 3(c), and 3(g)), representing real world use cases reported by various enterprises and implemented in several widely used web services [36, 37]. OptEx cannot account for the non-deterministic delays in communicating the intermediate RDD objects among the worker nodes during the execution of an iterative Spark job on the workers [9]. The above delays result in deviations in the observed length of the later phases of job execution, comprising the component X_2 , from the estimated completion time [9]. For experiments with large number of iterations, the job stages in the initial iterations cache the

TABLE 4: Optimal Scheduling With Estimated Optimal Cluster Size Under Varying SLA

SLA(sec)	Mode	App	Iterations											
			5			10			15			20		
			n	$T_{Est}(sec)$	$T_{Rec}(sec)$	n	$T_{Est}(sec)$	$T_{Rec}(sec)$	n	$T_{Est}(sec)$	$T_{Rec}(sec)$	n	$T_{Est}(sec)$	$T_{Rec}(sec)$
300	Standalone	ALS	3	203.2	207.0	5	188.0	193.0	6	189.7	92.0	7	188.8	195.0
340	YARN	ALS	2	240.0	430.0	3	328.4	335.0	4	311.5	315.0	4	331.0	337.0
380	Standalone	ChiSqTest	1	332.8	345.0									
800	YARN	ChiSqTest	1	794.8	821.0									
350	Standalone	ALS	3	200.0	297.0	4	299.0	297.0	4	340.5	343.0	6	303.2	113.0
400	YARN	ALS	3	174.2	315.0	4	381.0	385.0	5	378.6	386.0	5	398.0	395.0
430	Standalone	ChiSqTest	2	421.7	425.0									
790	YARN	ChiSqTest	2	783.7	781.0									
200	Standalone	ALS	4	79.5	275.0	6	176.1	178.0	6	189.7	191.0	7	188.8	194.0
260	YARN	ALS	4	250.5	257.0	5	259.0	258.0	6	260.0	255.0	7	2159.8	258.0

intermediate RDD objects locally in the worker nodes, resulting in a decrease in the time spent in communicating the RDD objects among the workers during the later iterations [9]. This results in a decrease in the deviation in the observed length of the job phases comprising X_2 from their estimated lengths. This, in turn, reduces the deviations in the overall observed completion time T_{Rec} with respect to the estimated overall completion time T_{Est} , estimated from X_1 and X_2 . Indeed, with increasing number of iterations, a decreasing trend is observed in the relative error (Figures 2 and 3). Hence we believe that OptEx can provide more accurate estimations with typical production level use cases, which typically involve number of iterations larger than 10 [36, 37].

Bottleneck: Prior to us, Ousterhout et al. [40] presented a technique to understand performance bottlenecks in Spark job execution, and how execution time is affected by network, disk I/O, stragglers, etc., which are difficult to control, especially in a public cloud. To capture the impact of bottlenecks in our model, we would have to control the above non-deterministic factors, namely network and disk I/O, while running our profiler. In that case, the model would be dependent on a comprehensive profiling technique where all possible network and I/O conditions are simultaneously modulated, which, in turn, complicates the proposed model. Further, since Ousterhout et al. rely on fine-grained instrumentation which is specifically based on a small sample of Spark jobs, their approach can not be readily applied to a wider group of Spark jobs.

Interference: The degree of interference due to contention in shared network I/O can be controlled in an in-house private datacenter. However, such solutions incur considerable capital expenditure as well as operational expenditure at the cost of less interference. Most organizations trade off this reduction in interference in favor of a pay-as-you-go cost model in an uncontrolled virtualized environment, namely a public cloud. Predicting the performance degradation due to interference in such an environment would require extensive experimentation with varying degrees of interference which is reliant on an accurate instrumentation of the target application. Contrastingly, OptEx uses a simple analytical model to estimate the job completion times of applications hosted in a public cloud. Despite not considering interference explicitly in the model, OptEx achieves a low relative error of 6% on an average which is comparable to claims by state-of-the-art performance prediction tools [3]. The variation of the estimation error using the OptEx model plotted in Figures 2 and 3 quantify the non deterministic deviations of the recorded job completion times from the estimated job completion times using our model. The above non determinism in the variation of the errors in estimations occur due to many reasons including interference in a public cloud. Thus our model indirectly captures the effect of interference to certain extent. Due to lack of space and for the sake of keeping

the model simple, evaluation of the model against varying degrees of interference is not covered this paper. Still the Figures show that our model suffers relatively unbiased variation in relative error despite the presence of uncontrolled network I/O in a public cloud. Spark applications typically are compute-intensive by nature, and require larger VMs which experience lower degree of throughput interference due to lower chances of sharing the same physical resources with other co-located VMs. Direct comparison among network I/O interferences with different VM sizes is difficult because it is also dependent on other factors like distribution of the VMs across different server racks and across datacenters. Hence even with experiments performed with varying VM size, drawing conclusions regarding performance degradation resulting from interference would be difficult. So we did not perform experiments with varying VM size in this paper. The obvious approach to eliminate interference would be to use the largest VM instance available but that would use up more cloud resources than what is actually required. Also, it would unnecessarily increase the cost of cloud usage which goes against our original goal of provisioning a cost optimal cluster. Instead, we choose VMs belonging to an instance type which provides sufficiently large amount of compute, memory, and network resources for running a given Spark application, while incurring the least hourly cost. With large VMs the major cause of interference is network I/O between workers located in different datacenters [41]. Since most analytics applications are hosted in a cluster comprising VMs that are typically located in the same datacenter, network I/O interference is negligibly less. Typically, the VMs are placed in the same rack, placed in the same subnet, or within a common VPC to minimize the interference due to inter datacenter network I/O. Hence in reality, the error component in our estimations that accounts for the effect of interference comprises only a very small fraction of the overall relative error. In support of our decision of not considering interference, we further argue that there is documented proof (refer to Figure 1 in [42]) of non deterministic fluctuations in the performance degradation due to interference even in a tightly controlled cloud environment. To accurately estimate such fluctuations in interference, researchers [42] resort to black box approaches where they have to perform intense experiments with actual target applications to collect extensive performance statistics. The above approach incurs large amount of resources and depends too much on an accurate instrumentation of the target applications without adding much value to the model.

Heterogeneity: Existing cloud schedulers, like the internal Spark stand-alone scheduler or external resource managers such as Mesos or Yarn, identify a virtual machine instance by the number of CPU cores and the size of its memory. But these schedulers do not consider other factors like network and I/O although these vary widely in different cluster setups. They schedule a task to run on a

node in a cluster solely based on data locality, required number of cores and memory size. Hence our model is sufficient in estimating job completion time on a public cloud comprising heterogeneous resources with above schedulers. Further, our model is intended to serve the needs of not just the large enterprises, like Google, Alibaba, etc., which typically require clusters in the scale of hundreds of machines [43], but also for small scale enterprises who dominate the modern market for cloud computing. In most of the day-to-day use cases of Spark in such enterprises, heterogeneity of underlying resources is not much of a concern. A popular example can be a video surveillance application, where users typically need to perform object detection on a large image dataset under stringent deadlines with the goal of sending real-time alerts. For the above task, users need to quickly provision a Spark cluster on a public cloud while incurring a minimum resource usage cost (since operational expenditure is a big concern for small enterprises), often for performing just that one job at a time. The cluster will be relatively small and homogeneous since it is serving one user and one job compared to the large clusters operated internally by high-profile private sector entities like Alibaba. Therefore the clusters that our model would normally provision typically execute a single job, and would experience interference predominantly with respect to other tenants in a virtualized public cloud environment.

7.6 Optimal Resource Provisioning

The work closest to OptEx is Elastisizer [2], which predicts optimal cluster composition for Hadoop, but does not address Spark. Moreover, Elastisizer over predicts, on an average by 20% and in the worst case by 59% [2]. Since OptEx uses a closed-form to estimate the completion time, it does not suffer from over-prediction. Table 4 demonstrates the effectiveness of the constrained optimization technique of OptEx (Section 6) in designing optimal scheduling strategies. For each job (refer to the column labelled App in Table 4) running in Standalone or YARN mode (refer to the column labelled Mode), Table 4 gives the cost-optimal cluster composition for executing a given job under given SLA deadlines (refer to the SLA column), while minimizing the cost of usage of the virtual machine instances. The optimal cluster size n for 5, 10, 15, and 20 iterations are given in separate group of columns each labelled n . The group of columns labelled T_{Est} in Table 4 gives the completion times T_{Est} estimated using OptEx, for 5, 10, 15, and 20 iterations, respectively. The other group of columns labelled T_{Rec} gives the recorded completion times T_{Rec} with the estimated cluster composition. Following [3], we propose a statistic \mathcal{S} to measure

TABLE 5: Resource Provisioning Under Cost Budget

Budget(\$)	Mode	App	n	T_{Est} (sec)	T_{Rec} (sec)
0.30	Standalone	ALS	53	49.2	48.0
0.80	YARN	ALS	58	120.2	115.0
1.00	Standalone	ChiSqTest	27	318.0	321.0
1.50	YARN	ChiSqTest	16	780.1	775.0
0.20	Standalone	ALS	35	49.4	50.0
0.50	YARN	ALS	36	120.0	119.0
0.80	Standalone	ChiSqTest	22	318.0	321.0
1.20	YARN	ChiSqTest	13	780.1	780.0
0.20	Standalone	ALS	26	49.7	50.0
0.40	YARN	ALS	29	120.6	117.0
0.60	Standalone	ChiSqTest	16	318.1	311.0
1.00	YARN	ChiSqTest	11	780.1	757.0
0.10	Standalone	ALS	17	50.7	52.0
0.30	YARN	ALS	21	121.1	125.0
0.40	Standalone	ChiSqTest	11	318.1	315.0

the effectiveness of OptEx in estimating whether a given job will satisfy the SLA deadline, while minimizing the cost. \mathcal{S} gives the percentage of cases which did not violate the SLA deadline, in the experiments recorded in the Table 4. \mathcal{S} evaluates approximately to 98% on cluster compositions comprising representative virtual machine instances, using job profiles generated by running example Spark applications on benchmark datasets (Section 7). Our evaluation follows closely with benchmark results reported by prior researchers [44]. Table 5 demonstrates that OptEx can be used effectively in project planning, i.e., for optimal cluster provisioning under given budget, while minimizing job execution times. Table 5 records the optimal cluster size (refer to the 4th column) required to run a given job (refer to the 3rd column) in Standalone or YARN mode (the 2nd column) estimated using Equation 8 under different values of the cost budget (refer to the 1st column), while optimizing the completion times. The 5th column of Table 5 gives the completion times T_{Est} estimated using OptEx, and the last column gives the recorded completion time T_{Rec} with the estimated cluster composition.

7.7 Error in Estimation of Execution Times

To ensure that the model is effective in estimating execution time of a given job, the model must accurately reflect the variation of job completion time with the input parameters. Here we demonstrate that the OptEx model exhibits a high degree of accuracy in the estimation of completion times of a given job which processes a given input dataset on a given cluster composition. First, we measure the model's error (which indirectly reflects models's accuracy) in terms of the difference between the standard deviations of the estimated and recorded (observed) job execution times, respectively. The error is computed as follows. $error = |\sigma_{T_{Est}} - \sigma_{T_{Rec}}|$. According to Equation 3 the total execution time is estimated, i.e., T_{Est} is computed, using the equation $T_{Est} = T_{Init} + T_{prep} + T_{vs} + T_{comp}$. We compute the standard deviation of T_{Est} as follows.

$$\begin{aligned} \sigma_{T_{Est}}^2 &= \sigma_{T_{Init}}^2 + \sigma_{T_{prep}}^2 + \sigma_{T_{vs}}^2 + \\ &\sigma_r^2 T_{comp} + 2 \times \text{Cov}(T_{Init}, T_{prep}) + 2 \times \text{Cov}(T_{prep}, T_{vs}) + \\ &2 \times \text{Cov}(T_{vs}, T_{Init}) + 2 \times \text{Cov}(T_{comp}, T_{Init}) + \\ &2 \times \text{Cov}(T_{comp}, T_{prep}) + 2 \times \text{Cov}(T_{comp}, T_{vs}). \end{aligned} \quad (10)$$

The parameters T_{Init} and T_{prep} are independent of the input variables n , s , and \mathcal{I} . Hence the covariances evaluate to 0. Using Equations 11 and 8 in the above error formula we get the following equation.

$$\begin{aligned} error &= \left| \sqrt{\sigma_{T_{Est}}^2} - \sqrt{\sigma_{T_{Rec}}^2} \right| \\ &= \left| \sqrt{\sigma_{T_{Init} + T_{prep} + n \times \mathcal{I} \times C + \mathcal{I} \times B/n + \frac{A \times s}{n}}^2} - \sigma_{T_{Rec}} \right| \\ &= |C^2 * \{\sigma_{\mathcal{I}}^2 * \sigma_n^2 + \sigma_{\mathcal{I}} * E^2(n) + E^2(\mathcal{I}) * \sigma_n^2\} + \\ &B^2 * \{\sigma_{\mathcal{I}}^2 * \sigma_{1/n}^2 + \sigma_{\mathcal{I}} * E^2(1/n) + E^2(s) * \sigma_{1/n}^2\} + \\ &A * \{\sigma_s^2 * \sigma_{1/n}^2 + \sigma_s * E^2(1/n) + E^2(s) * \sigma_{1/n}^2\} |^{1/2} - \sigma_{T_{Rec}}. \end{aligned} \quad (11)$$

Figures 4(a), 4(b), and 4(c) illustrate standard deviation of estimated job completion times and recorded execution times, for PageRank, ChiSqTest, and MovieLensALS applications, respectively, computed as $\sigma_{T_{Est}}$ and $\sigma_{T_{Rec}}$, respectively. These figures demonstrate that the distribution of estimated job completion time




Fig. 4: Standard Deviation of Estimations with OptEx

computed from the OptEx model closely echoes the distribution of recorded job completion time, i.e., the OptEx estimation follows the variation pattern of the recorded job completion times. The *error* between estimated standard deviation for recorded and estimated job completion times falls within 0.5, i.e., $error \leq \epsilon$ where the error threshold $\epsilon = 0.5$ seconds. This demonstrates the effectiveness of the OptEx model under the error bound of 0.5 seconds.

8 RELATED WORK

Verma et al. [3] present the design of ARIA, a framework for optimal resource allocation for Hadoop MapReduce. Herodotou et al. [2] present Elastisizer, a system that predicts the optimum cluster configuration for running a given Hadoop MapReduce job, using a model built by exploring a search space consisting of job profiles. Though Spark [9] is fast surpassing Hadoop in popularity and usage, there has not been much work in modelling Spark jobs yet [45, 4, 5]. [46] presents an approach to determine equivalence of Spark programs by converting the user defined functions comprised in the programs to equivalent first order formulae. But unlike OptEx, the approach described in this paper is restricted to programs that use only map, filter, cartesian product, fold, and foldByKey methods from the Apache Spark library. Huang et al. [44] presents DtCraft, a high performance system for compute intensive applications using parallel programming model. On the other hand, we are interested in data-intensive computing which typically deal with I/O bound data-centric workloads. DtCraft uses abstractions based on StreamGraph which is similar to the lineage graph in Spark. Hence theoretically OptEx can be extended to model DtCraft applications as well but this is out of the scope of this paper. Nguyen et al. [47] presents Galois, a system that provides radical performance improvement while running graph analytics tasks using Domain Specific Languages (DSLs). However, OptEx is designed for cost optimal cluster provisioning for a wider range of tasks, not just graph based DSLs. Ehsan et al. [48] presents AffordHadoop, a system built to process data-intensive computations which is designed specifically for scheduling map-reduce jobs. It deals with optimal resource scheduling for map-reduce jobs already running on a cloud whereas OptEx deals with cloud provisioning, i.e., estimating the composition of a cluster to run a given Spark job. Chen et al. [49] presents a performance model of internet services which is based on application level parameters like login rate and number of active connections. They provide power and CPU utilization model for traditional

web-based applications and database servers based on login rate and number of active connections. They donot consider large-scale bigdata applications based on modern distributed computing paradigms. Kansal et al. [50] presents a model for power metering and provisioning of virtual machines that can assist in management of power and reduce wastage of power capacity in the virtual machines comprised in the datacenters. Chaisiri et al. [51] provides a technique to reduce custer resource cost by leveraging reservation plans instead of on-demand provisioning plans. Here, the execution time on each virtual machine class is known beforehand, whereas with OptEx the execution time of the target job is unknown. Uргаonkar et al. [52] provides a dynamic resource provisioning model based on the request-response model which is developed from queuing theory. But this model is dependent on the parameters of the queuing model such as arrival rate and request rate which vary with time, hence it is not applicabe in our case. Islam et al. [53] assumes a black box approach that leverages machine learning techniques whereas OptEx provides an empirical model of job completion times.

9 CONCLUSIONS

OptEx models Spark job execution using analytical techniques. OptEx provides a mean relative error of 6% in estimating job completion time. OptEx yields a success rate of 98% in completing Spark jobs under a given SLA deadline with cost-optimal cluster compositions. OptEx can be used to estimate whether a given job will finish under a given deadline with the given resources on the cloud. It can be used to devise optimal scheduling strategy for Spark. The project is partially funded by NSERC, Canada and Army Research Office (ARO), USA.

REFERENCES

- [1] A. Inc, *Amazon Elastic Compute Cloud (Amazon EC2)*. <http://aws.amazon.com/ec2/#pricing>: Amazon Inc.
- [2] H. Herodotou, F. Dong, and S. Babu, "No one (cluster) size fits all: Automatic cluster sizing for data-intensive analytics," in *SOCC '11*.
- [3] A. Verma, L. Cherkasova, and R. H. Campbell, "Aria: Automatic resource inference and allocation for mapreduce environments," in *ICAC '11*.
- [4] W. Zhang, S. Rajasekaran, T. Wood, and M. Zhu, "Mimp: Deadline and interference aware scheduling of hadoop virtual machines," *CCGrid '14*.
- [5] K. Kc and K. Anyanwu, "Scheduling hadoop jobs to meet deadlines," ser. CLOUDCOM '10.
- [6] S. Imai, T. Chestna, and C. A. Varela, "Accurate resource prediction for hybrid IAAS clouds using workload-tailored elastic compute units," ser. UCC '13.

- [7] G. Singer, I. Livenson, M. Dumas, S. N. Srirama, and U. Norbirsath, "Towards a model for cloud computing cost estimation with reserved instances," *CloudComp 2010*.
- [8] S. Sidhanta, W. M. Golab, and S. Mukhopadhyay, "Optex: A deadline-aware cost optimization model for spark," in *CCGrid'16*.
- [9] M. Zaharia, M. Chowdhury, T. Das, A. Dave, J. Ma, M. McCauley, M. J. Franklin, S. Shenker, and I. Stoica, "Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing," in *NSDI'12*.
- [10] J. G. Shanahan and L. Dai, "Large scale distributed data science using apache spark," in *KDD '15*.
- [11] S. Salloum, R. Dautov, X. Chen, P. X. Peng, and J. Z. Huang, "Big data analytics on apache spark," *International Journal of Data Science and Analytics*, 2016.
- [12] L. Centenaro, "How to track shopping cart abandonment in google analytics," 2017.
- [13] F. Siddiqi, "Netflix at spark+ai summit 2018," 2018.
- [14] dezyre.com, "Tutorial: Introduction to apache spark," 2018.
- [15] S. Sidhanta, S. Mukhopadhyay, and W. Golab, "OptEx: Deadline-Aware Cost Optimization for Spark," Available at https://github.com/ssidhanta/OptEx/blob/master/optex_technical.pdf, Tech. Rep., 01 2018.
- [16] A. S. Foundation, "Apache spark libraries," 2015.
- [17] R. Milner, "An algebraic definition of simulation between programs," in *IJCAI'71*.
- [18] D. Park, "Concurrency and automata on infinite sequences," in *5th GI-Conference*.
- [19] R. Milner, *Communicating and Mobile Systems: The π -calculus*. New York, NY, USA: Cambridge University Press, 1999.
- [20] D. Latella, M. Massink, and E. P. de Vink, "Bisimulation of labelled state-to-function transition systems coalgebraically," *Logical Methods in Computer Science*, vol. 11, no. 4, 2015.
- [21] I. Castellani, "Bisimulations and abstraction homomorphisms," *J. Comput. Syst. Sci.*, vol. 34, no. 2/3, pp. 210–235, 1987.
- [22] S. Abriola, P. Barceli, D. Figueira, and S. Figueira, "Bisimulations on data graphs," in *KR'16*.
- [23] R. Mayr, "On the complexity of bisimulation problems for basic parallel processes," in *Automata, Languages and Programming*. Springer Berlin Heidelberg, 2000.
- [24] F. Zuleger, S. Gulwani, M. Sinn, and H. Veith, "Bound analysis of imperative programs with the size-change abstraction," in *SAS'11*.
- [25] J. Leader, *Numerical Analysis and Scientific Computation*. Pearson Addison Wesley, 2004.
- [26] M. Chase and G. Dummer, "Statlib data and story library (dasl)," <http://www.stat.yale.edu/Courses/1997-98/101/chisq.htm>, Jun. 1996.
- [27] S. Sidhanta, "Optex job classifier and model generator," 2018. [Online]. Available: <https://github.com/ssidhanta/SparkRDDAnalyze>
- [28] H. Garavel, F. Lang, R. Mateescu, and W. Serwe, "CADP 2011: a toolbox for the construction and analysis of distributed processes," *STTT*, vol. 15, no. 2, pp. 89–107, 2013.
- [29] D. Bergamini, N. Descoubes, C. Joubert, and R. Mateescu, "Bisimulator: A modular tool for on-the-fly equivalence checking," in *TACAS'05*.
- [30] X. Yang, J. Katoen, H. Lin, and H. Wu, "Proving linearizability via branching bisimulation," *CoRR*, vol. abs/1609.07546, 2016.
- [31] A. S. Foundation, "Spark programming guide," 2017.
- [32] *MovieLens dataset*, <http://www.grouplens.org/data/>.
- [33] J. Leskovec and A. Krevl, "SNAP Datasets: Stanford large network dataset collection," <http://snap.stanford.edu/data>, Jun. 2014.
- [34] "The YourKit Java Profiler," <http://www.yourkit.com>, last 2008.
- [35] Databricks, "Deep learning pipelines for apache spark," 2017.
- [36] A. S. Foundation, "Powered by spark," 2015.
- [37] Cloudera, "Category archives: Use case," 2015.
- [38] S. Pappas, "9 super-cool uses for supercomputers," 2015.
- [39] Apache Software Foundation. (2014) Apache hadoop nextgen mapreduce (yarn).
- [40] K. Ousterhout, R. Rasti, S. Ratnasamy, S. Shenker, and B.-G. Chun, "Making sense of performance in data analytics frameworks," in *NSDI'15*.
- [41] X. Pu, L. Liu, Y. Mei, S. Sivathanu, Y. Koh, C. Pu, and Y. Cao, "Who is your neighbor: Net i/o performance interference in virtualized clouds," *IEEE Transactions on Services Computing*, vol. 6, no. 3, pp. 314–329, 2013.
- [42] D. Novaković, N. Vasić, S. Novaković, D. Kostić, and R. Bianchini, "Deepdive: Transparently identifying and managing performance interference in virtualized environments," in *USENIXATC'13*.
- [43] Z. Ren, X. Xu, J. Wan, W. Shi, and M. Zhou, "Workload characterization on a production hadoop cluster: A case study on taobao," in *2012 IEEE International Symposium on Workload Characterization (IISWC)*, Nov 2012, pp. 3–13.
- [44] T. Huang, C. Lin, and M. D. F. Wong, "Dtcraft: A distributed execution engine for compute-intensive applications," in *ICCAD'17*.
- [45] N. B. Rizvandi, J. Taheri, R. Moraveji, and A. Y. Zomaya, "On modelling and prediction of total cpu usage for applications in mapreduce environments," in *ICA3PP'12*.
- [46] S. Grossman, S. Cohen, S. Itzhaky, N. Rinetzky, and M. Sagiv, "Verifying equivalence of spark programs," in *CAV '17*.
- [47] D. Nguyen, A. Lenharth, and K. Pingali, "A lightweight infrastructure for graph analytics," in *SOSP*. ACM, 2013, pp. 456–471.
- [48] M. Ehsan, K. Chandrasekaran, Y. Chen, and R. Sion, "Cost-efficient tasks and data co-scheduling with affordhadoop," *IEEE Transactions on Cloud Computing*, pp. 1–1, 2018.
- [49] G. Chen, W. He, J. Liu, S. Nath, L. Rigas, L. Xiao, and F. Zhao, "Energy-aware server provisioning and load dispatching for connection-intensive internet services," in *NSDI'08*.
- [50] A. Kansal, F. Zhao, J. Liu, N. Kothari, and A. A. Bhattacharya, "Virtual machine power metering and provisioning," in *SoCC '10*.
- [51] S. Chairisri, B.-S. Lee, and D. Niyato, "Optimization of resource provisioning cost in cloud computing," *IEEE Trans. Serv. Comput.*
- [52] B. Urgaonkar, P. Shenoy, A. Chandra, and P. Goyal, "Dynamic provisioning of multi-tier internet applications," in *ICAC'05*.
- [53] S. Islam, J. Keung, K. Lee, and A. Liu, "Empirical prediction models for adaptive resource provisioning in the cloud," *Future Gener. Comput. Syst.*

Subhajit Sidhanta Subhajit Sidhanta received his PhD in Computer Science from Louisiana State University, USA, in 2016. He spent one and a half year as a Postdoctoral Researcher with the Distributed Systems Research Group at INESC-ID research lab, affiliated with Instituto Superior Tecnico at University of Lisbon, Portugal, and then joined the Indian Institute of Technology Jodhpur in 2018 as an Assistant Professor in Computer Science and Engineering. His major research interest encompasses the areas of consistency in distributed systems, performance modelling of distributed storage and parallel computing systems.



Supratik Mukhopadhyay Supratik Mukhopadhyay is a Professor in the Computer Science Department at Louisiana State University, USA. He received his PhD from the Max Planck Institut fuer Informatik and the University of Saarland in 2001, MS from the Indian Statistical Institute, India and his BS from Jadavpur University, India. His areas of specialization include formal verification of software, inference engines, bigdata analytics distributed systems, parallel computing framework and program synthesis.



Wojciech Golab Wojciech Golab received his PhD in Computer Science from the University of Toronto in 2010. After a post-doctoral fellowship at the University of Calgary, he spent two years as a Research Scientist at Hewlett-Packard Labs in Palo Alto. He then joined the University of Waterloo in 2012, where he is presently an Associate Professor in Electrical and Computer Engineering. Dr. Golab's research agenda focuses on algorithmic problems in distributed computing with applications to the design, optimization, and



verification of software systems for data storage and analytics. His research publications have won two best papers awards, and his doctoral work on shared memory algorithms was distinguished by the ACM Computing Reviews as one of 91 notable computing items published in 2012".