

Elastic Executor Scheduling in Data Analytics Systems

Libin Liu¹, Member, IEEE, ACM, and Hong Xu², Senior Member, IEEE, Member, ACM

Abstract—Modern data analytics systems use long-running executors to run an application’s entire DAG. Executors exhibit salient time-varying resource requirements. Yet, existing schedulers simply reserve resources for executors statically, and use the peak resource demand to guide executor placement. This leads to low utilization and poor application performance. We present Elasecutor, a novel executor scheduler for data analytics systems. Elasecutor dynamically allocates and explicitly sizes resources to executors over time according to the predicted time-varying resource demands. Rather than placing executors using their peak demand, Elasecutor strategically assigns them to machines based on a concept called *dominant remaining resource* to minimize resource fragmentation. Elasecutor further adaptively reprovisions resources in order to tolerate inaccurate demand prediction and reschedules tasks to deal with inadequate reprovisioning resources on one machine. Testbed evaluation on a 35-node cluster with our Spark-based prototype implementation shows that Elasecutor reduces makespan by more than 36% on average, and improves cluster utilization by up to 55% compared to existing work.

Index Terms—Data analytics systems, elastic scheduling, executor, distributed systems.

I. INTRODUCTION

DATA analytics systems are widely used to process big data [1]–[3], [5], [25], [36], [39], [41], [53], [66], [78], [80]. The workflow of an analytics application can be expressed as a directed acyclic graph (DAG), which is composed of different stages of processing. Each stage runs a number of tasks on worker machines, and each task performs the same computation on different partitions of data [23], [25], [33], [34], [51], [55]–[57], [80].

Resource scheduling is critical in data analytics systems. Many resource schedulers have been developed for various objectives, such as fairness, cluster utilization, application completion time, etc. [22], [27], [29], [30], [33], [37], [40], [42], [48], [52], [58], [62], [67], [68], [70], [81]. Most are developed for a task-based execution model. They assume

Manuscript received January 25, 2019; revised May 28, 2020; accepted December 10, 2020; approved by IEEE/ACM TRANSACTIONS ON NETWORKING Editor Y. Guan. Date of publication January 22, 2021; date of current version April 16, 2021. This work was supported in part by the Hong Kong Research Grants Council under Award C7036-15G and Award 11216317, and in part by The Chinese University of Hong Kong under Grant 4937007, Grant 4937008, Grant 5501329, and Grant 5501517. This article was presented at ACM SoCC’18. (Corresponding author: Hong Xu.)

Libin Liu is with the Theory Lab, Huawei Hong Kong Research Center, Hong Kong, China (e-mail: liulibinsdu@gmail.com).

Hong Xu is with the Department of Computer Science and Engineering, The Chinese University of Hong Kong, Hong Kong, China (e-mail: hongxu@cuhk.edu.hk).

Digital Object Identifier 10.1109/TNET.2021.3050927




Fig. 1. A toy example to demonstrate the drawbacks of static allocation. The host’s capacity is 8, and applications are submitted at time t_0 and their resource demand time series are shown. In (a), static allocation runs all three applications sequentially, since their peak demands prevent them from running concurrently. In (b), elastic allocation allows them to run as long as their demand time series can be packed together, and reduces the makespan by 50%.

a task, which corresponds to one stage of the application DAG and is an individual process, is the basic execution unit and therefore the basic scheduling unit. This holds in general for batch processing systems like Hadoop [2] and cluster schedulers such as Yarn [68]. In Yarn for instance, tasks run in individual “containers” where Java virtual machines (JVMs) are spawned. A “container” runs one task only and is shut down after the task finishes.

However, in-memory analytics systems such as Spark [3] and Storm [4] rely on a different executor-based model [46], [72]. An executor is a long-running JVM process that executes an application’s entire DAG [7], [80]. Once an executor is launched, the scheduler can dispatch different tasks to it. This enables data reuse across tasks and significantly reduces the overhead of launching tasks which is critical for fast in-memory processing [57], [80]. Executor-based systems usually adopt task-based resource schedulers for simplicity. This, however, leads to various performance and efficiency problems as a result of the unfitting assumptions.

First, since the resource usage of a task is roughly constant, most schedulers use static allocation [22], [24], [32], [33], [37], [52], [58], [59], [62], [68]. However, executors naturally exhibit time-varying resource usages since they run the application’s entire DAG. As we will show in §II, an executor’s peak-to-trough resource usage can be as high as 409x for extended periods of time. Existing schedulers then have to use peak demands to reserve resources [24], [33], [37], [68] which leads to cluster underutilization and degraded makespan.¹ Applications may also have to wait longer for enough resources to become available, resulting in prolonged completion times and poor user experience. Figure 1 illustrates this using a toy example.

Recent work such as Morpheus [67] addresses this by dynamically reserving resources and adjusting the number

¹Makespan measures the total time used to complete all applications.

of executors as the application progresses. Also, some work [32]–[35], [42] tries to increase task parallelism within an executor. That is, we keep the executor resource allocation fixed, and increase or decrease the number of tasks run in parallel within the executors. Both methods are shown to improve resource utilization. However, they do not fundamentally address the problem, because each executor is still allocated a fixed bundle of resources during the entire time, which still incurs underutilization or overallocation of different resources in face of time-varying multi-resource demands across different stages.

Besides resource allocation, a scheduler also needs to carefully assign executors to machines in a cluster. Usually this is done using the executor’s peak demands and machine’s remaining capacity [37], [68], [70]. Such an approach does not precisely capture the time-series utilization of the worker machines and executors and creates fragmentation, when resources are idle but cannot be used to schedule executors that are ready. We give two examples of fragmentation: (1) An application’s peak demand may exceed a machine’s remaining capacity at the moment, and it cannot be placed on this machine. Yet, it is possible that the peak demand only happens in a later time, at which point the machine will have enough capacity to run it (because some applications will have finished by then); (2) An application with a very short period of high peak demand is selected to run first, preventing applications with stable demand from being scheduled on the same machine.

To address these problems, we propose Elasecutor, a novel executor scheduler that exploits time-varying resource demands for resource allocation and executor assignment. Elasecutor considers multiple resources: CPU, memory, network, and disk I/O. It exploits the recurring nature of many applications in production [17], [19], [33], [52], [67], and predicts the resource demands of applications over its lifetime. Elasecutor elastically allocates resources to executors according to the predicted time series of demands in order to reduce underutilization. It then packs executors strategically onto machines to minimize fragmentation among multiple resources and improve makespan.

We make the following contributions in this article.

- We make a case for elastic executor scheduling (§II). We show through measurements with real workloads that Spark executors exhibit significant time-varying resource usage patterns (§II-A). We further experimentally establish the predictability of executor’s demand time series (§II-B).
- We design a new scheduler called Elasecutor (§III), that allocates time-varying resources to executors and assigns them to machines based on the predicted demand time series (§III-B). To do so, Elasecutor relies on a concept called *dominant remaining resource* to search for executors whose demand time series best matches the time series of a machine’s available resources. Elasecutor also reprovisions resource dynamically at runtime to compensate for possible prediction errors (§III-C) and timely reschedules tasks when machine’s resources are unavailable for reprovisioning (§III-D).
- We implement Elasecutor on Spark (§IV) and present a realistic performance evaluation on a 35-machine cluster (§V). Experiments with real workloads show that Elasecutor substantially improves performance. Compared to existing solutions such as Tetris [33], Elasecutor

TABLE I
INPUT DATASET SIZE FOR PROFILING THE FOUR SPARK WORKLOADS FROM THE HIBENCH BIGDATA BENCHMARK SUITE [10]. EACH WORKLOAD RUNS WITH 20 EXECUTORS, EACH USING AT MOST 3 CORES AND 8GB MEMORY

Workloads	Terasort	K-means	LR	PageRank
Dataset Size	32.0GB	37.4GB	37.3GB	2.8GB

reduces makespan by up to 63%, and improves resource utilization by up to 55%.

II. MOTIVATION

We motivate our work by highlighting the limitations of peak based static executor allocation and assignment commonly used in current systems (§II-A). We also show that the executor resource usage can be fairly well predicted when the workloads are recurring (§II-B).

A. Need for an Elastic Scheduler

Executors are the basic scheduling unit in Spark and similar systems. Each application has dedicated executors to run its tasks [7]. Current executor schedulers [3], [24], [37], [68] work in virtually the same way as task schedulers for systems like Hadoop [2]. Users need to specify the resource demands of an executor, so the scheduler can make resource reservations. The resources allocated to an executor are static and released only after the application finishes. Thus peak resource demands have to be used for allocation.

We argue that current schedulers can lead to severe resource underutilization because application resource usage varies greatly in different stages of data processing [33], [56], [67], [75], [76]. To demonstrate this, we profile the resource usage of executor processes using several typical Spark workloads as shown in Table I, including Terasort, K-means, Logistic Regression (LR), and Pagerank. These workloads are commonly used in existing work [10], [35], [38], [56], [57]. We develop a monitoring module on Spark to collect CPU, memory, network I/O, and disk I/O usage of an executor. The measurements are done on our 35-machine testbed described in §V-A.

From Figure 2, we observe that the executor resource usage exhibits significant temporal variations. CPU usage varies all the way from only 4% to 100% for all applications. Similarly, memory usage ranges from 500MB to nearly 6.8GB for K-means. The network I/O varies from almost 0Gbps to ~4Gbps for LR, and disk I/O from almost 0MB/s to ~190MB/s for Terasort, respectively. More than half of the time an executor actually uses very little resources. Table II further shows our detailed analysis for executor’s resource usage. We can see that the peak-to-trough ratio is high, ranging from 3.3 to 409.6. The period of peak resource usage takes up at most 22.4% of total runtime, and more than half of the runtime the resource usage actually falls below 50% of the peak except for memory.

Therefore static allocation using peak demands would clearly cause severe resource wastage and performance issues.

Some recent work [37], [67] has considered dynamically adjusting the number of executors according to the workload in order to improve utilization. Yet each executor still gets a fixed bundle of resources (CPU and memory) during the application’s entire runtime. As observed in Figure 2, the usages

TABLE II

STATISTICAL ANALYSIS OF THE TIME-VARYING RESOURCE USAGE OF EXECUTORS FOR THE FOUR SPARK WORKLOADS (P/A = PEAK USAGE/AVERAGE USAGE, P/T = PEAK USAGE/TROUGH USAGE, DP/T = DURATION OF PEAK USAGE/TOTAL RUNTIME, DHP/T = DURATION FOR EXCEEDING HALF OF PEAK USAGE/TOTAL RUNTIME)

Applications	Terasort				K-means				LR				Pagerank			
	CPU	Memory	Net	Disk	CPU	Memory	Net	Disk	CPU	Memory	Net	Disk	CPU	Memory	Net	Disk
P/A	1.8	1.7	6.2	1.5	1.7	1.2	11.5	5.6	2.1	1.4	5.5	6.1	3.9	1.3	20.2	9.1
P/T	60	3.3	237	6.1	75	6	53	100	50	12.0	409.6	42.5	50	11.5	119	50
DP/T (%)	22.4	1.5	1.5	10.7	15.4	3.9	1.9	1.9	15.0	3.7	0.6	6.8	1.5	21.8	0.2	0.8
DHP/T (%)	36.5	67	10.3	73.2	51.9	90.4	5.8	13.5	40.4	91.0	19.5	11.3	23	84	2.5	7.1

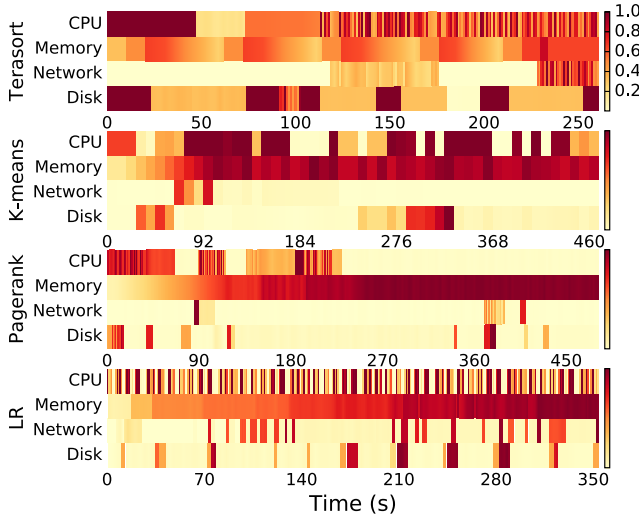


Fig. 2. Heat maps of resource usages for four Spark workloads. The resource usages are normalized to the highest value of the executor. For all applications, the highest CPU usage is 3 cores; Terasort, K-means, Pagerank, and LR have 5.9GB, 6.8GB, 6.1GB, and 6.5GB highest memory usage; 2.4Gbps, 0.53Gbps, 0.62Gbps, and 4Gbps highest network usage; and 181MB/s, 103MB/s, 96MB/s, 170MB/s highest disk I/O usage.

of different resources do not correlate strongly. Thus such an approach does not fundamentally solve the over-allocation issue. Moreover, they only consider CPU and memory, and lack control over shared network and disk I/O resources. This results in possibly severe contention of shared resources which may then lead to underutilization of other resources.

B. Predictability of Resource Usage Time Series

To design an elastic executor scheduler, we need to have prior information about the time series of resource demands for executors. We now show that such information can be fairly easily predicted in practice.

Recent studies report that many production analytics workloads are recurring, such as running the same queries periodically on new data [17], [19], [33], [52], [67]. Further, the running times of the workloads are mostly constant given the same amount of input data and resources [17], [29], [33], [55], [67], [77]. Hence we can predict an executor's future resource requirements based on profiling its previous runs.

To see this, we measure eight workloads each with three input dataset sizes shown in Table III. We vary the number of CPU cores from one to five and memory from 2GB to 10GB correspondingly [11], creating five different resource profiles. Then for each application (a workload with certain dataset size and CPU and memory setting), we run a resource profile five times, each time with different datasets of the same size.




Fig. 3. The CDFs of coefficient of variations (CoVs) for per-stage execution time (SET) and resource peak usage of each executor over five runs for each of the eight workloads with different settings on CPU core and input data size shown in Table III. Each run uses a different input dataset.

Other settings are the same as in §II-A. We collect the completion times of each stage and executor's peak resource usage in each stage, then calculate the coefficient of variation (CoV) over the five runs. Figure 3 shows the cumulative distribution of CoVs for per-stage execution time and per-stage peak resource usage. To make it clearer, we also use Table IV to show the corresponding statistical analysis of CoVs. We can see that CoVs are in general smaller than 14% and each stage's execution time is quite stable (90%ile CoV is less than 5.5%), as a CoV value less than 1 is considered to be low variance [29].

Therefore, for most recurring workloads it is accurate enough to use the profiling results from previous runs with same resource setting to represent the resource demands. For workloads with new settings, we can also build a prediction model [19], [26], [29], [69], [76] to infer their time-series resource usage, which we detail in §IV. Note that for new applications, we need to collect data for several runs before we can predict the resource demand time series. Again given that most workloads are recurring, we believe this does not affect the usability of our system.

III. DESIGN

We now present the design of Elasecutor in this section. We start by presenting the system overview in §III-A. We then explain in detail the elastic executor scheduling algorithm in §III-B, which is the core contribution of the design. Lastly, we discuss in §III-C how Elasecutor uses dynamic re-provisioning at runtime to minimize the impact of prediction errors in inferring the executor's resource demand.

TABLE III
THREE TYPES OF INPUT DATASET SIZES FOR THE EIGHT WORKLOADS OF THE HIBENCH BIGDATA BENCHMARK SUITE [10]

Workloads	Sort	WordCount	Terasort	Bayes	K-means	LR	PageRank	NWeight
Dataset 1	3.1G	3.1G	3.2G	1.8G	3.7G	7.5G	1.7M	37.5M
Dataset 2	30.6G	30.6G	32G	3.5G	18.7G	22.4G	247.9M	294.5M
Dataset 3	286.8G	305.9G	320G	70.1G	37.4G	37.3G	2.8G	2.7G

TABLE IV
STATISTICAL ANALYSIS OF COVS

CoV Statistics (%)	Percentiles			
	10th	50th	90th	99th
SET	0.7	2.6	5.5	9.1
CPU	0	0.3	0.6	0.7
Memory	3.1	5.6	8.6	11.0
Network	2.4	4.2	7.9	13.4
Disk I/O	2.5	2.9	6.8	12.9

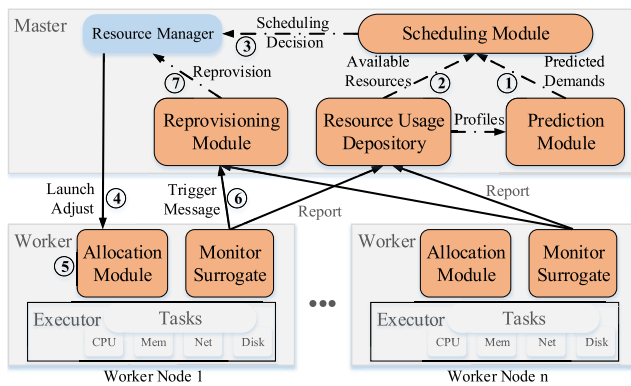


Fig. 4. Elasecutor overview (key components are highlighted).

A. Overview

Elasecutor is an executor resource scheduler for data analytics systems. It predicts executors’ time-varying resource demands (step 1 in Figure 4), collects workers’ available resources (step 2), assigns executors to machines to minimize fragmentation (steps 3 and 4), elastically allocates resources (step 5), and leverages dynamic reprovisioning for better application QoS (steps 6 and 7).

We explain several key components here.

1) *Monitor Surrogate*: Elasecutor employs a *monitor surrogate* at each worker node to continuously monitor the resource usage of executors in real-time. It collects the process-level CPU, memory, network I/O, and disk I/O usage, and reports the time series profiles to the resource usage depository (RUD) at the master node via RPC. The information is then used to build machine learning models to predict executor resource time series. The monitor surrogate also reports the node’s future available resources to the RUD. Moreover, it monitors executor progress to see whether reprovisioning should be triggered due to significant prediction errors.

2) *Resource Usage Depository (RUD)*: The RUD runs as a background process at the master node communicating with monitor surrogates and collecting information at each heartbeat of 3s. For simplicity we use a single master node and one RUD process, which is sufficient in our testbed evaluation.

We can scale the RUD to multiple cores or multiple masters for large-scale deployment following many similar designs in distributed control plane [45], which is beyond the scope of this article.

3) *Scheduling Module*: The scheduling module decides how resources should be allocated to executors and which executors should be assigned to machines. It obtains an application’s demand time series from the prediction module which we will introduce shortly. It then packs executors to machines across multiple resource types, in order to avoid overallocation and minimize fragmentation throughout the executor’s lifetime. For this purpose, we design a scheduling algorithm based on a novel metric called *dominant remaining resource (DRR)* which is detailed in §III-B. Once a scheduling decision is made, the selected worker IDs along with the executor IDs are sent to Spark’s resource manager [7], which instructs the corresponding workers to launch the executors.

4) *Allocation Module*: This module explicitly and dynamically sizes the resource bundles to the executor process according to the resource manager’s instructions. Through this, Elasecutor implements elastic allocation based on time-varying demands, which is illustrated in detail in §IV.

5) *Reprovisioning Module*: Dynamic reprovisioning mainly deals with cases when the executor’s actual resource usage deviates significantly from the predicted time series, which is unavoidable in practice. When an executor’s progress is detected by the monitor surrogate to be slower than expected, the reprovisioning module is activated to calculate extra resources needed to make up for the slowdown. The corresponding algorithm is discussed in §III-C.

6) *Prediction Module*: Finally, the prediction module runs as a background process at the master node. It continuously fetches executor resource profiles from the RUD to train a prediction model for application’s resource demand time series. Many machine learning and time series analysis techniques can be used for this purpose, which is not the focus of this article. §IV provides more information about the prediction algorithm we currently use.

B. Elastic Executor Scheduling

The foremost challenge Elasecutor faces is how to elastically schedule executors with their multi-resource demand time series. We explain our solution to this challenge here.

We focus on recurring applications which are common in production settings [19], [33], [52], [67], [69]. When an application request is submitted, it specifies the number of executors and their configurations. Elasecutor predicts the per-executor resource demand time series based on such information and the past runs of this application. Elasecutor strives to satisfy the application’s request completely, instead of scaling up or down the number of executors based on some fairness criteria. This is because users value performance consistency or predictability much more than fairness in practice [67].

In cases when resources are insufficient applications will simply wait.

1) *An Analytic Model*: We begin with an analytic model to capture the problem.

In our model, we consider four resources: CPU, memory, network I/O, and disk I/O. For each resource r , we denote its capacity on machine j as C_j^r . The per-executor demand of application i on resource r is $D_i^r(t')$ when it is running at t' into its lifetime. Once started, application i runs for T_i time slots, and the number of executors required is N_i . All these are known to the scheduler. Let $x_i(t)$ be the decision variable for scheduling, i.e. $x_i(t) = 1$ if application i is running at time slot t , and let t_i denote the time when application i starts. Let y_{ij} be the decision variable for executor assignment. That is, y_{ij} indicates the number of executors assigned on machine j for application i .

a) *Constraints*: First, we assume that applications cannot be paused or preempted once scheduled, which is consistent with prior work [33], [34], [67]. Thus,

$$x_i(t) = \begin{cases} 1, & t_i \leq t \leq t_i + T_i, \\ 0, & \text{otherwise,} \end{cases} \quad \forall i. \quad (1)$$

Second, the cumulative resource usage on a machine at any given time t cannot exceed its capacity. Each executor's resource allocation is exactly equal to the predicted demand $D_i^r(t - t_i)$ for resource r when it has been running since t_i (without interruption). Thus,

$$\sum_i x_i(t) y_{ij} D_i^r(t - t_i) \leq C_j^r, \quad \forall r, t, j. \quad (2)$$

The scheduler always allocates exactly N_i executors for application i as stated in the beginning of §III-B.

$$\sum_j y_{ij} = N_i, \quad \forall i. \quad (3)$$

b) *Objective function and analysis*: The objective of the scheduling algorithm is to minimize the makespan across all applications. Under a schedule $\{t_i\}$, application i finishes at $t_i + T_i$, and the makespan is $\max_i(t_i + T_i)$. Thus the scheduling problem can be formulated as the following:

$$\begin{aligned} & \min_{t_i} \max_i (t_i + T_i) \\ & \text{s.t. (1), (2), (3).} \end{aligned} \quad (4)$$

It is also possible to use other objective functions, such as application completion time, in our formulation.

Finding an optimal schedule to the above problem (4) is difficult. The objective function and constraints (2) are nonlinear, which makes the problem computationally expensive to solve. In spite of ignoring the objective function and the time-varying nature of resource demand, the problem of packing multi-dimensional balls (executors) to minimum number of bins (machines) is APX-Hard [65], [73]. Moreover, what we have here is an offline setting. The online version where applications arrive dynamically is even more difficult to solve with reasonable competitive ratio. Therefore, most prior work for the packing problems relies on heuristics.

Clearly, makespan would be minimized if all available resources along time could be utilized by applications seamlessly. Naturally, the basis for minimizing makespan is to avoid resource underutilization and minimize machine-level




Fig. 5. An illustration example for DRR. We just use two kinds of resources as an example. Three executors running on machine 1 consume CPU and memory variously over their lifetime. They all start to run at time t_0 and stop at time t_1 , t_2 , and t_3 , respectively, and we assume machine 1's CPU and memory capacity are 8 and 6. Elasecutor calculates DRR at t_1 , which is $1/4$ in this case.

resource fragmentation. Therefore, Elasecutor aims to schedule executors to the best fitted machine in order to minimize multi-resource fragmentation.

We now introduce our heuristic scheduling algorithm.

2) *Packing Executors With MinFrag*: A well-known heuristic to one-dimensional packing problem is Best Fit Decreasing (BFD). BFD proceeds by repeatedly matching the largest ball that can fit in the current bin until no more balls fit, then open a new bin. Intuitively, this approach reduces fragmentation and thus the number of bins used. BFD requires no more than $\frac{11}{9}OPT + 1$ bins, where OPT is the optimal number of bins [31].

Our heuristic *MinFrag* extends BFD, by transforming the multi-dimensional (and time-varying) bin packing problem into the classic one-dimensional problem. Such a transformation relies on a metric called *dominant remaining resource (DRR)*, which we illustrate first.

a) *DRR*: The DRR of a machine is defined similar to dominant resource of a job in [32]. For a given machine, we find the earliest time t at which an executor, among those running on this machine, will finish according to the predicted demand time series. We then compute machine j 's average remaining resource over time for the period up to t , which can be denoted as $R(*, j, t)/C(*, j, t)$ for resource $*$. $R(*, j, t)$ is the integral of the amount of remaining resource along the time dimension up to t , and $C(*, j, t)$ is the integral of total capacity up to t . The machine j 's DRR is then the maximum remaining resource among all types.

Figure 5 shows an example. We select t_1 as the time point to calculate DRR for machine 1. $\frac{R(\text{CPU}, 1, t_1)}{C(\text{CPU}, 1, t_1)} = 1/4$ and $\frac{R(\text{Mem}, 1, t_1)}{C(\text{Mem}, 1, t_1)} = 1/6$, and its DRR is simply $1/4$.

Note we use the maximum remaining resource, not the minimum, because it better reflects machine utilization. If a machine's maximum remaining resource is 10%, then utilization of all resources is at least 90%. However if a machine's minimum remaining resource is 10%, utilization of some resources may still be lower than 90%. Minimum remaining resource reflects a machine's ability or potential to run executors, which is important if our objective is to minimize machine utilization.

b) *MinFrag*: As explained, *MinFrag* is based on BFD. On a high level, *MinFrag* works by iteratively assigning the "largest" executor to a machine that yields the *minimum* DRR in order to maximize utilization and improve makespan. We illustrate it in detail now.

Algorithm 1 shows *MinFrag*, and Table V lists the notations used here. When a heartbeat is received from a machine j , *MinFrag* updates its available resources $AR(j)$ and then repeatedly does the following. It identifies if there is any

Algorithm 1 *MinFrag* Pseudo Code

```

1: when a heartbeat received from machine  $j$ 
2: update  $AR(j)$ 
3: while there are pending executors and  $AR(j) > 0$  do
4:   for each pending executor  $i$  do
5:     if  $RD(i) < AR(j)$  then
6:       compute  $\theta(i, j)$ 
7:       select  $i^* = \arg \min_i \theta(i, j)$ 
8:       launch executor  $i^*$  on  $j$ 
9:       update  $AR(j)$  and  $\theta(j)$ 

```

TABLE V
NOTATIONS IN *MinFrag*

Notation	Explanation
$AR(j)$	Time-series of future available resource at machine j
$RD(i)$	Time-series of resource demands of executor i
$\theta(j)$	DRR of machine j
$\theta(i, j)$	Updated DRR of machine j when executor i is placed on it

executors that can run with enough resources on the machine. If yes, it computes the machine's DRR $\theta(i, j)$ if the executor was placed on it. Then among all eligible executors, *MinFrag* chooses i^* that minimizes $\theta(i, j)$, i.e. the largest executor. It updates the placement result, the machine's DRR $\theta(j)$, and available resource $AR(j)$. It then repeats the process until all executors are scheduled or there are no more available resources on the machine for any pending executor to run.

We use an example in Figure 6 to illustrate how *MinFrag* works. Figure 6a shows the machine's remaining capacity up to time T , when a running executor will complete its execution. There are two executors to schedule, and their demand time series are shown in Figures 6b and 6c, respectively. If executor 1 is assigned to the machine, $t_0 + 0.875T$ is the time point for calculating DRR, and the DRR $\theta(1, j) = \max\{\frac{53}{112}, \frac{165}{448}, \frac{3}{70}, \frac{6}{35}\} = \frac{53}{112}$. If executor 2 is assigned to the machine, $t_0 + T$ is the time point for calculating DRR, and the DRR $\theta(2, j) = \max\{\frac{13}{32}, \frac{43}{128}, \frac{1}{10}, \frac{1}{20}\} = \frac{13}{32}$. *MinFrag* then schedules executor 2 to run which minimizes DRR and thus maximizes utilization in this case. After taking executor 2, this machine does not have any network bandwidth and thus cannot take any more executors at the moment.

Some may argue that we can compute a score for each resource based on its remaining capacity, and convert the vector into a scalar value for comparison across candidate executors (say based on the Euclidean norm). However, the values for different resources have different units which makes such comparison irrelevant. Considering the remaining resource at its face value does not faithfully represent the degree of fragmentation as machines may have different capacities. The ratio between remaining resource and its capacity can represent actual fragmentation. Our evaluation results in §V-F corroborate our argument.

C. Dynamic Reprovisioning

It is difficult to perfectly predict the time-varying resource demands due to many exogenous factors, such as infrastructure issues (e.g., hardware replacements, driver updates, etc.) and application issues (changes in size and skew of input data, changes in code/functionality, etc.). Thus we design

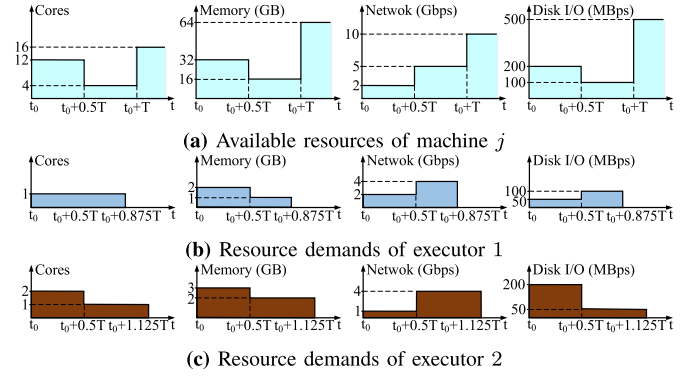


Fig. 6. An illustration example of *MinFrag*. (a) The remaining resources of machine j until time $t_0 + T$, when a running executor will complete its execution. (b) The time-series resource usage of executor 1 which is expected to finish at time $t_0 + 0.875T$. (c) The time-series resource usage of executor 2 which is expected to finish at time $t_0 + 1.125T$. The capacities of machines are: 16 cores for CPU, 64 GB for memory, 10 Gbps for network, and 500MB/s for disk I/O.

a dynamic reprovisioning mechanism that adjusts resource allocation online to tolerate prediction errors in Elasecutor.

Reprovisioning is triggered when an executor's progress is longer than ρ times the expected execution time of a processing stage of the application's DAG. We set ρ to 1.1, which is experimentally determined to balance application performance and resource wastage.

Given reprovisioning is required, we wish to find out the actual resource demand of the application now. Elasecutor does so by temporarily allocating all remaining resource of the machine to this executor for one monitoring period (3s), and observe the executor's resource usage in this period. Suppose the actual usage of the executor i is $u(i, *)$ during this period for resource $*$, and the predicted demand is $p(i, *)$. Elasecutor then scales the allocation proportional to $u(i, *)/p(i, *)$ across all resources for all the remaining processing stages of the executor, and returns the remaining resources to the machine. This heuristic allows Elasecutor to quickly correct resource allocation without causing many missed scheduling opportunities at the machine.

When the machine has little remaining resources, the executor's actual usage observed in the reprovisioning period may be obscured and do not reflect its true demand. This is one limitation of our heuristic. We propose a task rescheduling mechanism to address the issue in the following.

D. Task Rescheduling

When a machine does not have enough resources for reprovisioning of an executor, the corresponding DAG stage cannot be fully accelerated to compensate for the predicted completion time. To deal with this we leverage the speculative execution mechanism built into Spark's task scheduler [15]. Since the tasks of the executor would be delayed (tasks are put in the pending queue or running slowly), Spark's speculative execution would detect their slower progress and launch duplicates of these tasks on other executors. Our rescheduling mechanism picks the executors that have the most available resources for speculative execution. This would then trigger reprovisioning on the chosen executors to obtain more resources to run the newly arrived tasks.

IV. IMPLEMENTATION

We implement a prototype of Elasecutor with $\sim 1K$ LOC in Python, Java, and Scala based on Spark 2.1.0. We open

source our prototype here [8]. The monitor surrogate and allocation module at each worker machine communicate with the master node via RPC. We use lightweight system-level tools such as `psutil` and `jvmtop` in Linux to implement the monitor surrogate. We use `resourceRef`, which is a data structure that includes worker ID, application ID, executor ID, and corresponding resource time series. The scheduling module implements the *MinFrag* algorithm and dispatches the scheduling decisions to Spark’s resource manager, which launches executors on worker machines. We modify the `launchExecutor()` function at the resource manager to send the predicted resource demand time series to the corresponding workers, along with other information. The allocation module then uses the modified `cgroups` and `OpenJDK` [13] to configure resources of the executor process based on the prediction results.

A. Allocation Module

Once a worker node receives the message for launching executors or resizing them from the resource manager, the allocation module adopts subsystems of `cgroups` to configure the time period and the limits of CPU, network, and disk I/O the executor process is entitled to. However, for Java-based systems, the maximum heap size of a JVM stays constant during its lifetime. Dynamically throttling memory outside the JVMs is difficult. Inspired by [71], we adopt a method to enable dynamic memory limits at runtime by modifying `OpenJDK` [13]. As in operating systems, the virtual address space does not have physical memory until it is actually used, and the allocation module leverages this to reserve and commit specified address spaces of a fixed maximum heap size dynamically at runtime. We implement an API `JVMmanage()` inside a `OpenJDK`’s JVM which listens to instructions from the allocation module for such dynamic memory commitment. As a result, each JVM knows the correct maximum memory size it can use at any time.

One may wonder that in case of prediction errors, out of memory (OOM) error [18], [71], [74] may happen due to insufficient memory. In fact, we find that applications in Spark 2.1.0 do not just fail when memory is less than demanded since they can spill data to disk as a remedy. This of course slows down executors and would trigger reprovisioning, which then would fix the problem.

B. Reprovisioning Module

This is implemented as a long-running process at the master node. It continuously collects reports about executor progress via monitor surrogates at each worker, and triggers reprovisioning by invoking the resource manager with the corresponding worker ID, application ID, executor ID, and corresponding resource time series. Subsequently, allocation module is instructed to adjust resource limits at runtime.

C. Prediction Module

Our current implementation simply uses the average resource time series of the latest 3 runs as the prediction result for recurring workloads with the same settings. Our analysis in §II-B demonstrates it is fairly accurate. For applications with new settings Elasecutor has not seen, we rely on a prediction model based on SVR to infer the demand time series. Elasecutor can also leverage other prediction

methods [19], [26], [29], [69], [76] which are beyond the scope of this article. Note that for new applications which are never seen by the system before, we need to collect data for several runs before we can predict the resource demand time series. Until then, the applications are allowed to run with peak demanded resources.

SVR Prediction: We cast our prediction problem as a regression problem. We have $x_i, i = 1, 2, \dots, n$, where x_i is the i -th multidimensional input vector that represents the application type, dataset size, and CPU and memory configurations, and n is the number of training samples. We also have the ground truth y_i which is the actual resource usage time series of the i -th run. As in §II-B, the executor’s time-series resource usage is stable for the same application type and settings. The goal is to learn the relationship between x_i and y_i so that when a application with new settings submitted, we can predict its demand time series based on the model.

To do so, we rely on support vector regression (SVR) [28], [54], [64]. We use the radial basis function (RBF) kernel [44] which we find to have better results than other kernels. We select ϵ -SVR [61] as the optimization model, and find its optimal parameters ϵ based on k -fold crossvalidation and grid search [21] until the prediction reach the accuracy which we show in §V-F.

The prediction model is continuously trained online by successively collecting profiling results from RUD. Once an application with new settings is submitted, the model makes prediction about its resource usage time series and outputs the results for the scheduling module to consume. In our experiments, we found that our prediction process for applications with new settings can be done within 1s.

V. EVALUATION

We now evaluate Elasecutor using testbed experiments. Our evaluation answers the following questions:

- How much overall performance benefit can Elasecutor provide compared to existing solutions? (§V-B)
- How efficiently does Elasecutor utilize resources? (§V-C)
- How well does Elasecutor perform with TPC-H workload and changing mixes of recurring applications? (§V-D, §V-E)
- How well do the prediction, scheduling, reprovisioning, and task rescheduling modules work? (§V-F)
- How much overhead does Elasecutor add? (§V-G)

A. Setup

Our testbed cluster consists of 35 machines connected with a 10 GbE switch. Each machine has two 2.4 GHz Intel Xeon E5-2630 v3 processors, 64 GB DDR4 RAM, a quad-port Intel X710 10 GbE NIC, and two 7200 RPM disks. All machines run Ubuntu 16.04.2 LTS with kernel version 4.4.0, Scala 2.10.4, and HDFS 2.6.0. We deploy our Elasecutor implementation on top of Spark 2.1.0.

1) *Methodology:* To test our prototype, we use eight workloads described in §II-B. The workloads are from the HiBench bigdata benchmarking suite [10], which are commonly used in existing work [35], [38], [56], [57]. For each workload we use different input data sizes listed in Table III and different CPU and memory upper limit configurations, ranging from one to five cores for CPU and 2GB to 10GB for memory as in §II-B. Besides, to make the evaluation more general, in §V-D

and §V-E, we further evaluate Elasecutor’s performance using TPC-H [16] workloads.

We generate 120 recurring applications with different workloads, input data sizes, and resource settings. In each experiment run, the same 120 applications are used. They arrive according to a Poisson process with mean inter-arrival time of 25 seconds for a period of 3200 seconds in each run. This is consistent with existing work [33], [35], [38], [56], [57]. Elasecutor’s prediction module is fed with 3 runs of each application as explained in §IV and is used as the prediction results. Its SVR method has been trained offline with 30 runs of the recurring applications, and the same trained model is used for each experiment run. Besides the recurring applications, we prepare 12 applications with new input data size and new resource settings that has not seen by Elasecutor before. The same 12 applications are used in each run, but they arrive randomly within the period of 3200 seconds (§V-B and §V-C). In §V-E, we change the mix of recurring applications and re-train our prediction model with each unique application mix. Generally we repeat each experiment for five runs unless stated otherwise.

2) *Schemes Compared*: We compare Elasecutor to three existing resource scheduling strategies which are deployed in production systems [33]–[35], [37], [68]. (1) *Static*: This policy statically reserves CPU and memory for each executor according to the peak demand, and launches a fixed number of executors according to user request. (2) *Dynamic*: This uses the built-in dynamic allocation policy in Spark to scale the number of executors dynamically based on the workload. Each executor is allocated a multiple of $\langle 1 \text{ core}, 2\text{GB RAM} \rangle$ [14]. This is similar to Morpheus [67] without preplanned time-varying reservations. We do not compare to Morpheus as it optimizes for SLO and load balance and is very different from Elasecutor. Both *static* and *dynamic* policies are implemented in Spark, and they only consider CPU and memory. (3) *Tetris*: This policy considers network and disk I/O in addition to CPU and memory. Following [33], it greedily chooses an executor that has the highest dot product value between the vectors of machine available resources and executor peak demands, and allocates the peak resource demands to the executor. For all three schemes, we explicitly size related resources using `cgroups`.

3) *Task Scheduling Within an Application*: For a given application, Elasecutor and all other schemes adopt the fair scheduler [9] for task scheduling. Other popular task schedulers [6], [22], [32]–[35] can also be used by Elasecutor.

B. Makespan

We first investigate Elasecutor’s makespan improvement.

We look at the makespan reduction provided by Elasecutor as shown in Figure 7a. Observe that Elasecutor reduces the average and median makespan by over 40% compared to existing schemes (60% over Tetris). Elastic resource allocation and scheduling in Elasecutor ensures that resources are not over-allocated or under-allocated during the executor’s lifetime and machines are better utilized, thus translating to the smallest makespan performance. Also Elasecutor takes into account network I/O and disk I/O, which are not considered by *Static* and *Dynamic*. *Tetris* performs worse than *Static* and *Dynamic* because it has to wait until all four resources are available for the executor’s peak demand, while *Static* and *Dynamic* only wait on CPU and memory. In fact, we observe in

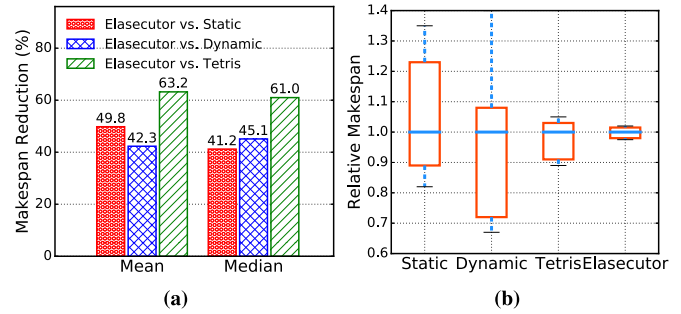


Fig. 7. (a) Makespan reduction of Elasecutor; (b) Boxwhisker plot of makespans which are normalized to the median value of each scheme. Whiskers represent the maximum and minimum values. Each experiment run takes more than 5.9 hours, and we repeat 30 runs here.

our experiments that applications run faster *after they start* in *Tetris* compared to *Static* and *Dynamic*, but they spend longer time waiting on resources.

Tetris’s poor makespan performance is worth more discussion here. Actually its task assignment algorithm optimizes makespan, and is shown to have smaller makespan than DRF [32] and the capacity scheduler [6] in the article [33]. The discrepancy of the results here is caused by two factors. First, *Tetris* along with DRF and capacity scheduler all use peak demand. Second, they are all designed for task-based systems like Hadoop, where tasks are short in duration and using peak demand is fine. When applying *Tetris* to executor-based systems in our systems, as executors last over the application’s lifetime, *Tetris* has to wait until the peak demand can be satisfied for all four resources, and its performance degrades.

Figure 7b further shows the stability of makespan. Here we normalize the makespan to the median values under different policies over 30 runs. Recall that the same applications are submitted in each run with *random* arrival times as explained in §V-A. The result shows that Elasecutor delivers much more stable makespan than other policies with a much smaller box. In other words, Elasecutor is more consistent with respect to the arrival order of the applications. On the other hand, *Static* and *Dynamic* have fluctuating makespan that depends heavily on the arrival order of the applications. This is because their peak-demand based allocation results in fragmentation that very much depends on application’s arrival order, and their random executor assignment adds more inconsistency. *Tetris* has more consistent makespan by using a multi-resource packing heuristic to reduce fragmentation, but it still has similar problems due to the use of peak demand in the heuristic.

C. Resource Utilization

Now we investigate cluster resource utilization with Elasecutor. We show (estimated) resource usage with different policies in Figure 8 for one example run. For resources considered by a scheduling policy (CPU and memory for *Static* and *Dynamic*, all four for *Tetris* and Elasecutor), we measure and plot the actual utilization; for other resources, we derive the utilization based on our predicted demand time series. The same methodology is used in prior work [33], [35]. Observe that as in Figures 8a and 8b, *Static* and *Dynamic* are unable to fully utilize CPU and memory which they consider for executor allocation and placement. They also over-allocate

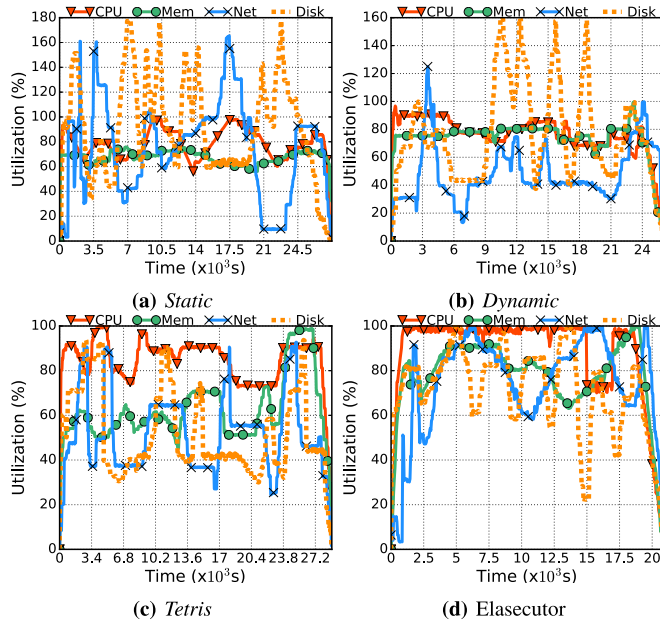


Fig. 8. Resource utilizations under different policies. Note the time unit is 10^3 s.

TABLE VI

FREQUENCIES THAT A MACHINE'S UTILIZATION IN A RESOURCE EXCEEDS A THRESHOLD. WE USE THREE THRESHOLDS 70%, 90%, AND 100% HERE

Policies	CPU (%)	Memory (%)	Network (%)	Disk I/O (%)
<i>Static</i>	61.3, 25.0, -	37.0, -, -	30.7, 25.1, 18.6	49.3, 37.2, 30.6
<i>Dynamic</i>	56.8, 26.6, -	57.7, -, -	25.3, 7.1, 5.0	51.2, 32.4, 26.9
<i>Tetris</i>	89.7, 70.8, -	49.1, 21.1, -	43.4, 13.3, -	37.2, 18.8, -
Elasecutor	97.0, 89.2, -	70.7, 39.4, -	68.7, 41.2, -	59.2, 37.0, -

network and disk resources most of the time, resulting in 180% utilization sometimes. *Tetris* in Figure 8c performs slightly better in that it avoids over-allocation. Elasecutor in Figure 8d improves the utilization of all resources with elastic resource scheduling. The cluster is bottlenecked on different resources at different times.

We also note that although *Dynamic* utilizes CPU and memory more, compared to Elasecutor it still loses 40% makespan performance as observed in §V-B. In addition to the more efficient *MinFrag* heuristic, the reason is that *Dynamic* incurs additional CPU and memory cost as it launches executors over time frequently. This also confirms our argument on the overhead of adjusting executor numbers over time in §I.

Table VI depicts the fraction of time when a machine's utilization exceeds a threshold for each resource. We obtain the usage statistics from all machines with one example run. We can see that Elasecutor utilizes resources much more efficiently: for example 97% of the time CPU utilization is over 70%. On the other hand, *Static* and *Dynamic* sometimes over-allocate network and disk I/O, and waste CPU and memory that they statically reserved for executors.

Lastly, Table VII shows the average utilization improvement with Elasecutor for all four types of resources. Compared to other policies, Elasecutor improves utilization by at least 27.2% for CPU, 22.6% for memory, 33.4% for network, and 25.4% for disk I/O. Clearly, this demonstrates Elasecutor can utilize resources efficiently and saves cost for cluster operators.

TABLE VII
ELASECUTOR'S AVERAGE UTILIZATION IMPROVEMENT OVER OTHER POLICIES

Utilization Improvement (%)	CPU	Memory	Network	Disk I/O
Elasecutor vs. <i>Static</i>	43.4	29.5	40.8	25.4
Elasecutor vs. <i>Dynamic</i>	27.2	22.6	33.4	40.0
Elasecutor vs. <i>Tetris</i>	28.6	25.2	55.6	43.9

TABLE VIII

IMPROVEMENT OF ELASECUTOR OVER *Static*, *Dynamic*, AND *Tetris* WITH HiBENCH AND TPC-H WORKLOADS AND UNDER CHANGING MIXTURES OF RECURRING APPLICATIONS

Improvement (%)	Makespan			
	TPC-H		Mixtures of Applications	
	Average	Median	Average	Median
Elasecutor vs. <i>Static</i>	47.3	44.4	45.2	44.6
Elasecutor vs. <i>Dynamic</i>	41.6	35.9	38.5	31.4
Elasecutor vs. <i>Tetris</i>	54.2	62.7	47.8	55.7

TABLE IX

ELASECUTOR'S AVERAGE RESOURCE UTILIZATION IMPROVEMENTS WITH DIFFERENT MIXES OF RECURRING APPLICATIONS

Utilization Improvement (%)	CPU	Memory	Network	Disk I/O
Elasecutor vs. <i>Static</i>	35.4	34.3	30.2	19.0
Elasecutor vs. <i>Dynamic</i>	18.7	15.9	29.8	30.3
Elasecutor vs. <i>Tetris</i>	27.5	24.5	38.1	33.5

D. Makespan Using HiBench and TPC-H

In addition, we investigate how well Elasecutor performs with the TPC-H benchmark [16] introduced as recurring workloads. We use 22 TPC-H query applications together with the 120 recurring applications from HiBench and train the prediction model to drive the experiments. These applications also arrive according to a Poisson process with mean inter-arrival time of 25 seconds for a period of 3500 seconds in each run. Table VIII shows the makespan improvements. We see Elasecutor improves average and median makespan by 47.3% and 44.4% against *Static*, 41.6% and 35.9% against *Dynamic*, and 54.2% and 62.7% against *Tetris*. This demonstrates that Elasecutor obtains consistent performance improvements in makespan using HiBench and TPC-H workload.

E. Changing Mixtures of Recurring Applications

We now study how Elasecutor performs under changing mixtures of recurring applications. We generate ten different mixes of the 142 recurring applications from HiBench and TPC-H workloads. For each run, we change the arrival orders of recurring applications and randomly increase or decrease their input dataset size between 0 to 1%. Table VIII shows Elasecutor's makespan reduction against other policies. Clearly, Elasecutor reduces average and median makespan over varying application mixtures by at least 38.5% and 31.4%, respectively.

Next, we look at resource utilization with different application mixtures. Table IX summarizes the results. Observe that Elasecutor improves utilization by at least 18.7% for CPU, 15.9% for memory, 29.8% for network, and 19.0% for disk I/O. The results here demonstrate that Elasecutor produces consistent performance improvement with respect to application mixtures.




Fig. 9. (a) CDFs of prediction effectiveness with more runs. (b) CDFs of reductions in application completion time (ACT) and application execution time (AET) by comparing Elasecutor with and without reprovisioning.

TABLE X
IMPROVEMENT OF ELASECUTOR WITH TASK RESCHEDULING (TR)
OVER WITHOUT TASK RESCHEDULING

Statistics (%)	Makespan		ACT		
	Average	Median	50th	90th	99th
With TR vs. Without TR	2.7	2.2	0.6	0.9	4.7

F. Microbenchmarks

We now evaluate individual components of Elasecutor.

1) *Effectiveness of Prediction*: For the prediction module, we define prediction effectiveness as the ratio between the predicted total amount an executor is going to use and the actual total amount used during the execution for each resource. The average effectiveness across all 4 resources of each executor is then one data point. Note here we only concern executors of the 12 new applications with new settings. Figure 9a plots the CDFs of prediction effectiveness of the 15th and 30th run (with continuous training), respectively. We can see that the SVR method in §IV provides fairly accurate estimations: more than 80% of the time prediction is effective with less than 20% difference. Also with more training samples, the prediction improves with less than 15% errors more than 90% of the time (the curve for the 30th run).

2) *Effectiveness of Reprovisioning*: Now we estimate the effectiveness of the reprovisioning module. Figure 9b compares the CDF of the reductions in application completion time (ACT) and application execution time (AET) for Elasecutor with and without reprovisioning. Here AET is the time the application takes to complete after it is scheduled to run. The setup is the same as in §V-A with 132 applications in each run for 5 runs. We can see that by using reprovisioning, Elasecutor’s median ACT and AET decrease by 4.6% and 5.0%, respectively, and the 90%ile ACT and AET decrease by 6.0% and 14.5%, respectively, compared to not using it. The results demonstrate that reprovisioning is important for prediction based resource schedulers to improve application QoS performance.

3) *Effectiveness of Task Rescheduling*: We then investigate the effectiveness of task rescheduling. Table X shows the comparison in makespan and ACT for Elasecutor with and without task rescheduling. The setup is the same as in §V-A. We observe that Elasecutor’s average and median makespan decrease by 2.7% and 2.2%, respectively, and top 90%ile and 99%ile ACT decrease by 0.9% and 4.7%, respectively, with

TABLE XI
IMPROVEMENT OF DRR OVER TRC AS AN ALTERNATIVE
METRIC FOR EXECUTOR PLACEMENT

Design Choices (%)	Makespan			ACT		
	Average	Median	Stdev.	50th	90th	99th
DRR vs. TRC	11.5	13.4	5.7	2.3	6.1	11.0

TABLE XII
RESOURCE CONSUMPTION OF A MONITOR SURROGATE. CPU AND
MEMORY ARE AVERAGED OVER ALL MACHINES. TOTAL EXECUTOR
PROFILE SIZE IS THE SIZE OF ALL EXECUTOR PROFILES
AVERAGED ACROSS ALL MACHINES AT EACH HEARTBEAT

Resource Consumption	CPU	Memory	Total Executor Profile Size
Monitor surrogate	0.3%	0.1%	12.1 KB

task rescheduling. As task rescheduling takes effect only at certain scenarios, the performance benefit is limited.

4) *Effectiveness of DRR*: Our scheduling heuristic *MinFrag* uses DRR. We now investigate its effectiveness by comparing with an alternative metric for multi-resource executor placement. Particularly, we consider to sum up the relative remaining capacity of each resource, i.e. $R(\text{CPU}, m)/C(\text{CPU}, m) + R(\text{Mem}, m)/C(\text{Mem}, m) + R(\text{Net}, m)/C(\text{Net}, m) + R(\text{Disk}, m)/C(\text{Disk}, m)$ for machine m , as the total remaining capacity (TRC). The scheduling module then applies TRC instead of DRR as the metric in the *MinFrag* heuristic in Algorithm 1. Effectively, in each iteration it chooses an executor that minimizes its TRC when placed on the current machine to run.

We compare DRR with TRC in Table XI in terms of makespan and ACT. Observe that DRR reduces the average and median makespan by 11.5% and 13.4%, respectively, and have more stable makespan with 5.7% smaller standard deviation. DRR also improves ACT moderately in median and high percentiles. The results show that DRR works better than other alternative metrics to minimize resource fragmentation.

G. Overhead

We now evaluate the overhead of Elasecutor. We first consider the overhead of monitor surrogate. Table XII shows the average CPU and memory consumption of a monitor surrogate and the average size of executor profiles of all 35 machines of our cluster at each heartbeat. We can see that the overheads are very small, and the additional network overhead of sending profiles is negligible. We also find that the CPU and memory overhead of other modules of Elasecutor at the master node is less than that of the monitor surrogate.

We also evaluate Elasecutor’s latency overhead to the analytics system, particularly the resource manager. The prediction, scheduling, and reprovisioning modules are independent processes running concurrently with the resource manager, and we wish to ensure they do not negatively impact resource manager’s responsiveness. We measure the time taken by the resource manager to process a heartbeat from both a worker node and the application driver in Elasecutor and Spark. The application driver in Spark is responsible for creating context for executing applications, and we do not modify it in Elasecutor. Table XIII shows the results. Elasecutor takes about the same time as Spark to process both types of heartbeat messages.

TABLE XIII

AVERAGE TIME TO PROCESS HEARTBEATS FROM WORKERS AND APPLICATION DRIVER WITH AND WITHOUT ELASECUTOR OVER 100 HEARTBEATS

Time to process (ms)	Unmodified Spark	Elasecutor
Worker heartbeat	~0.031	~0.035
Application driver heartbeat	~0.153	~0.155

VI. DISCUSSIONS

After exhibiting the performance benefit of Elasecutor, we briefly discuss what we have done and the future work on Elasecutor in this section.

A. Model the Resources Jointly

In our current model, we do treat the executors' various resources (CPU, memory, network, and disk I/O) separately. Yet, in practice, they are often correlated. For instance, in Spark, the lack of available memory causes more spills and increases CPU (due to serialization overheads) and network or disk I/O (due to remote/local writes). Even though Elasecutor has dynamic reprovisioning mechanism (§III-C) to compensate for prediction errors caused by not considering such cases, we need to add such correlations between resources into our model to make it work efficiently. Thus one future work for Elasecutor is to revise its resource demand prediction model and take correlations between resources into account.

B. Scalability

Due to the small scale of our testbed, the resource usage of the master node is very small. According to the average file size of an executor's profile on each server, the bandwidth used for transmitting them at each heartbeat is 32.27Kbps per server. Assuming 10Gbps network interfaces, the master node can support ~309K machines. Besides, the CPU and memory costs are fairly low: both less than 1.5% to handle the current testbed scale. Correspondingly, if we fully utilize the CPU of a server, it can support 2.33K machines, which should be able to handle the typical production clusters [60], [63]. The scalability can be further improved in the following ways: (1) In production data centers 40G or 100G NICs are not uncommon [60], [63], which implies the bandwidth overhead of our system is even smaller compared to the 10Gbps links we use in our testbed; (2) We expect some nodes to have GPU, FPGA, or other hardware accelerators [43], [49] that can offload the computation from CPU and support larger clusters; (3) We can reduce the bandwidth requirement of updating executor resource profiles by adopting compression and/or sampling methods.

C. SLOs/SLAs

Elasecutor provides significant application completion time improvements and strives to provide better QoS via reprovisioning, but does not aim at guaranteeing *strict* SLOs/SLAs. To do the latter, Elasecutor would need to (1) have a dedicated component for SLO inference, (2) make resource reservations for executors ahead, and (3) dynamically adjusts resource allocations at runtime so that strict deadlines can be met [29], [67]. On the other hand, Elasecutor may be able to provide *statistical* SLO/SLA guarantees as it inherently predict applications' execution time based on historical data (§IV). We are exploring this as future work.

D. Differences Between DRR and DRF

DRR (dominant remaining resource) used in Elasecutor is inspired by DRF [32] and share some similarities, in particular the fact that they both use "dominant resource" to convert multi-dimensional metrics into scalars. Their differences, on the other hand, are distinct. DRR as defined in §III-B2 concerns the maximum *remaining* resource of a machine *over time* and defines "dominant resource" based on it. DRF represents the maximum *time-invariant* resource *requirement* of a task with respect to the machine's capacity. Further, compared to DRF that aims to achieving fairness between tasks, DRR is used to reduce resource fragmentation and minimize makespan in *MinFrag*.

E. Other Executor-Based Frameworks

Our current Elasecutor implementation is based on Spark. Streaming systems like Storm or Flink also have long running jobs, which potentially can also use Elasecutor. In the next step, we plan to explore other systems Elasecutor can be applied to, and make it an extension in resource managers like Yarn [68], Mesos [37], and Kubernetes [12] to serve all executor-based frameworks.

F. Workload Usecase Analysis

The ideal workloads for Elasecutor to obtain high performance are ones that either have a mix of I/O- and computing-intensive stages like Sort, Terasort, and Wordcount, or exhibit time-varying computation and I/O usage ratios like Pagerank, K-means, and Bayes. The executors used to run them can be packed together by Elasecutor to fully utilize various resources. In addition, deep learning training also often consists of long running tasks. When many training jobs co-exist in the same cluster, they progress in different paces and stress different resources (CPU, GPU, network, etc.), which makes it a good fit for Elasecutor as well. Another future work is then to run experiments for production deep learning training jobs to measure their actual resource usage patterns and investigate the potential benefit of Elasecutor there.

G. Differences Between Tetris, Carbyne, Graphene, and Elasecutor

Tetris [33] is a task scheduler, and Carbyne [34] and Graphene [35] are job schedulers, rather than executor scheduler like Elasecutor. We compare Elasecutor with the variant of Tetris (called *Tetris* in §V), because Tetris is the first cluster scheduler to explicitly consider multi-resource packing with a best fit decreasing (BFD) algorithm [20] like Elasecutor and the BFD packing algorithm in Tetris is natively designed to optimize makespan as Elasecutor. Thus, we modify Tetris to schedule executors with considering peak demands as executors' resource demands and adopting its built-in BFD algorithm. However, Carbyne and Graphene are designed to optimize job completion time, rather than makespan. Thus, comparing them against Elasecutor does not make sense. In addition, Elasecutor may cooperate with Carbyne and Graphene to make adjust executors' time-series resource demands, as executors' time-series demands depend on how DAG jobs and tasks are scheduled.

TABLE XIV
SUMMARY OF PREVIOUS APPROACHES COMPARED TO ELASECUTOR

Existing Work	Granularity	Scheduling			Objective
		Multiple Resources	Elastic Sizing	Machine Assignment	
Jockey [29], Quasar [27]	Task	CPU, memory	✗	✗	SLO
Sparrow [58], Apollo [22]	Task	CPU, memory	✗	✓	ACT and fairness
Tetris [33]	Task	✓	✗	✓	Makespan
Yarn [24], [68], Mesos [37], DRF [32], Omega [62]	Framework / Task	CPU, memory	✗	✗	Fairness
KOALA-F [48]	Framework	Servers	✗	✗	Utilization
Prophet [77]	Executor	Network, disk	✗	✓	Utilization
Morpheus [67]	Executor	CPU, memory	# executors	✓	SLO and load balance
Medea [30]	Executor	✓	✗	✓	Depend on objective functions
Elasecutor	Executor	✓	✓	✓	Makespan

H. Avoiding OOM When Resizing Executors

Theoretically, Elasecutor cannot completely avoid prediction errors. Therefore, Elasecutor needs to deal with cases where predicted memory usages are less than the executors' actual demand. Yet, as we stated in Elasecutor's allocation module implementation part in §IV, applications do not fail when memory is less than demanded since the framework can spill data to disk as a remedy. As a result, this slows down executors and would trigger reprovisioning, which then would fix the problem.

I. Static and Dynamic Executor Scheduler in §V

Static and *Dynamic* randomly pack executors to worker machines whose available resources can meet the executors' resource requirements. *Dynamic* performs worse compared to Elasecutor. The reasons mainly are three-fold: (1) It only considers CPU and Memory; (2) Each executor is allocated fixed resources, which is still inflexible compared to Elasecutor; (3) It lacks an efficient executor placement algorithm.

J. How About Changing the Number of Tasks per Executor Dynamically?

Dynamic scales the number of executors dynamically over time and each executor is allocated a multiple of (1 core, 2GB RAM). Besides, each executor can only run one task and the executor is withdrawn once the task is finished. Therefore, *Dynamic* compared to Elasecutor is similar with dynamically changing the number of tasks per executor across stages of the application. Besides, dynamically changing the number of tasks within one executor can change the executor's time-series demands. Thus Elasecutor may cooperate with dynamically changing number of tasks per executor to improve the flexibility of executor placement and improve the executor packing efficiency, which may be explored as future work.

VII. RELATED WORK

There has been a substantial body of work on scheduling and resource allocation in data analytics systems. We present Elasecutor first in the conference paper [50]. Here we extend the design with task rescheduling, and evaluate it with changing mixtures of recurring applications. We compare Elasecutor to related work along several dimensions as summarized in Table XIV.

A. Granularity

Most of prior work [6], [9], [22], [27], [29], [32]–[35], [40], [58], [79] assumes a task based system such as Hadoop. As discussed in §I and §II, they do not work well for executors which run the application's entire DAG and have time-varying demands. Cluster schedulers such as Yarn [24], [68] and

Mesos [37] manage the resource allocation for co-existing frameworks. They are also commonly used for task scheduling within frameworks in practice. Yet, they do not dig deeply to explore how to schedule the long-running executors within a computing framework and have the same problems as task schedulers when applied to executor scheduling. Prophet [77], Morpheus [67], and Medea [30] consider executor scheduling and are more related to our work.

B. Scheduling Considerations

Existing work can also differ in terms of their scheduling considerations. Most work focuses on CPU and memory only. Tetris [33], like Elasecutor, considers network and disk I/O in addition. Most task and cluster schedulers mentioned above adopt static allocation and do not elastically size the per-task resource allocation. Prophet [77] performs executor scheduling, but it only considers network and disk I/O and cannot support elastic resource sizing. Morpheus [67], another executor scheduler, dynamically adjusts the number of executors to meet application's demand, without changing the per-executor allocation though. This does not cope well with the multi-resource time-varying executor demand as we explained in §I and brings large executor launching cost as we analyze in §V-C. Medea focuses on making good placement decisions but does not consider variations of resource utilization of long-running executors. As far as we know, Elasecutor is the first scheduler that elastically sizes executor resource allocation.

Lastly, for machine assignment, cluster schedulers [24], [32], [37], [62], [68] usually just use random assignment. Sparrow [58] uses randomized sampling to choose machines quickly. Other schedulers typically use packing heuristics but in different ways. Prophet [77] favors machines with the least sum of fragmentation and over-allocation. Morpheus [67] essentially uses the Worst Fit heuristic for SLO, and Tetris [33] uses the BFD heuristic based on the dot product of task's peak demand and machine's available resource capacity. Elasecutor adopts BFD with DRR inspired by DRF [32] and is shown to work more effectively than alternative metrics. Medea applies a mathematical optimization approach that accounts for constraints and global objectives. Clearly, Elasecutor's executor assignment problem can be formulated as Medea's constraints. A technical difficulty is that, as we showed in §III-B1 our objective function and constraints are non-linear, while Medea employs linear programming to formulate its placement constraints.

C. Objective

In this regard, cluster schedulers [24], [32], [37], [62], [68] are the easiest to analyze: they optimize for fairness among

co-existing frameworks. Task and executor schedulers optimize for various objectives: mostly makespan and ACT [58], and SLO or utilization. Elasecutor focuses on makespan and also improves ACT and utilization as experimentally shown in §V. Medea supports various objectives for long-running executors and it provides low placement latency for short-running executors.

VIII. CONCLUSION

We have presented a novel executor scheduler Elasecutor. Elasecutor builds on the following two key ideas: elastically allocating resources to an executor to avoid over-allocation, and placing executors strategically to minimize multi-resource fragmentation. We prototype Elasecutor on Spark and evaluate it on a medium-scale testbed. Compared to existing approaches, Elasecutor reduces makespan by more than 36% on average, while improving cluster resource utilization by up to 55%.

Going further, we are exploring some interesting directions. Placement constraints are common in practice [35], [47] and can be added to Elasecutor’s scheduling heuristic. Elasecutor does not guarantee SLOs in terms of deadlines [29], [67] currently. We seek to better understand the impact of elastic multi-resource scheduling on application performance and SLO in order to support SLO guarantees. Finally, it is possible to develop a general executor scheduler based on Elasecutor so it can be integrated into Yarn and other cluster schedulers and benefit other executor-based systems.


REFERENCES

- [1] *Apache Flink*. Accessed: May 16, 2018. [Online]. Available: <http://flink.apache.org>
- [2] *Apache Hadoop*. Accessed: May 16, 2018. [Online]. Available: <http://hadoop.apache.org>
- [3] *Apache Spark*. Accessed: May 16, 2018. [Online]. Available: <https://spark.apache.org>
- [4] *Apache Storm*. Accessed: May 16, 2018. [Online]. Available: <http://storm.apache.org>
- [5] *Apache Tez*. Accessed: May 16, 2018. [Online]. Available: <http://tez.apache.org>
- [6] *Capacity Scheduler*. Accessed: May 16, 2018. [Online]. Available: <http://bit.ly/1tGpbDN>
- [7] *Cluster Mode Overview*. Accessed: May 16, 2018. [Online]. Available: <https://spark.apache.org/docs/2.1.0/cluster-overview.html>
- [8] *Elasecutor*. Accessed: May 16, 2018. [Online]. Available: <https://github.com/NetX-lab/Elasecutor>
- [9] *Fair Scheduler*. Accessed: May 16, 2018. [Online]. Available: <https://spark.apache.org/docs/2.1.0/job-scheduling.html#fair-scheduler-pools>
- [10] *HiBench*. Accessed: May 16, 2018. [Online]. Available: <https://github.com/intel-hadoop/HiBench>
- [11] *How-to: Tune Your Apache Spark Jobs*. Accessed: May 16, 2018. [Online]. Available: <http://blog.cloudera.com/blog/2015/03/how-to-tune-your-apache-spark-jobs-part-2>
- [12] *Kubernetes*. Accessed: May 16, 2018. [Online]. Available: <https://kubernetes.io/>
- [13] *OpenJDK*. Accessed: May 16, 2018. [Online]. Available: <http://openjdk.java.net>
- [14] *Resource Allocation Policy in Spark 2.1.0*. Accessed: May 16, 2018. [Online]. Available: <https://spark.apache.org/docs/2.1.0/job-scheduling.html#resource-allocation-policy>
- [15] *Spark Configuration*. Accessed: May 16, 2018. [Online]. Available: <https://spark.apache.org/docs/2.1.0/configuration.html>
- [16] *TPC-H Benchmark*. Accessed: May 28, 2020. [Online]. Available: <https://github.com/cartershanklin/hive-testbench/tree/master/sample-queries-tpch>
- [17] S. Agarwal, S. Kandula, N. Bruno, M.-C. Wu, I. Stoica, and J. Zhou, “Re-optimizing data parallel computing,” in *Proc. USENIX Symp. Netw. Syst. Design Implement. (NSDI)*, 2012, pp. 281–294.
- [18] M. K. Aguilera *et al.*, “Remote memory in the age of fast networks,” in *Proc. ACM Symp. Cloud Comput. (SoCC)*, 2017, pp. 121–127.
- [19] O. Alipourfard, H. H. Liu, J. Chen, S. Venkataraman, M. Yu, and M. Zhang, “CherryPick: Adaptively unearthing the best cloud configurations for big data analytics,” in *Proc. USENIX Symp. Netw. Syst. Design Implement. (NSDI)*, 2017, pp. 469–482.
- [20] N. Bansal and A. Khan, “Improved approximation algorithm for two-dimensional bin packing,” in *Proc. 25th Annu. ACM-SIAM Symp. Discrete Algorithms*, Jan. 2014, pp. 13–25.
- [21] J. Bergstra, R. Bardenet, Y. Bengio, and B. Kegl, “Algorithms for hyperparameter optimization,” in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 2011, pp. 2546–2554.
- [22] E. Boutin *et al.*, “Apollo: Scalable and coordinated scheduling for cloud-scale computing,” in *Proc. USENIX Symp. Operating Syst. Design Implement. (OSDI)*, 2014, pp. 285–300.
- [23] M. Chowdhury and I. Stoica, “Efficient coflow scheduling without prior knowledge,” in *Proc. ACM Conf. Special Interest Group Data Commun.*, Aug. 2015, pp. 393–406.
- [24] C. Curino, D. E. Difallah, C. Douglas, S. Krishnan, R. Ramakrishnan, and S. Rao, “Reservation-based Scheduling: If you’re late don’t blame us!” in *Proc. ACM Symp. Cloud Comput. (SoCC)*, 2014, pp. 1–14.
- [25] J. Dean and S. Ghemawat, “MapReduce: Simplified data processing on large clusters,” in *Proc. Symp. Operating Syst. Design Implement. (OSDI)*, 2004, pp. 1–13.
- [26] C. Delimitrou and C. Kozyrakis, “Paragon: QoS-aware scheduling for heterogeneous datacenters,” in *Proc. ACM ASPLOS*, 2013, pp. 1–25.
- [27] C. Delimitrou and C. Kozyrakis, “Quasar: Resource-efficient and QoS-aware cluster management,” in *Proc. ACM ASPLOS*, 2014, pp. 1–17.
- [28] H. Drucker, C. J. C. Burges, L. Kaufman, A. Smola, and V. Vapnik, “Support vector regression machines,” in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 1996, pp. 155–161.
- [29] A. D. Ferguson, P. Bodik, S. Kandula, E. Boutin, and R. Fonseca, “Jockey: Guaranteed job latency in data parallel clusters,” in *Proc. 7th ACM Eur. Conf. Comput. Syst. (EuroSys)*, 2012, pp. 99–112.
- [30] P. Garefalakis, K. Karanasos, P. Pietzuch, A. Suresh, and S. Rao, “MEDEA: Scheduling of long running applications in shared production clusters,” in *Proc. 13th EuroSys Conf.*, Apr. 2018, pp. 1–13.
- [31] M. R. Garey, R. L. Graham, and J. D. Ullman, “Worst-case analysis of memory allocation algorithms,” in *Proc. 4th Annu. ACM Symp. Theory Comput. (STOC)*, 1972, pp. 143–150.
- [32] A. Ghodsi, M. Zaharia, B. Hindman, A. Konwinski, S. Shenker, and I. Stoica, “Dominant resource fairness: Fair allocation of multiple resource types,” in *Proc. USENIX Symp. Netw. Syst. Design Implement. (NSDI)*, 2011, p. 24.
- [33] R. Grandl, G. Ananthanarayanan, S. Kandula, S. Rao, and A. Akella, “Multi-resource packing for cluster schedulers,” in *Proc. ACM Conf. SIGCOMM*, Aug. 2014, pp. 455–466.
- [34] R. Grandl, M. Chowdhury, A. Akella, and G. Ananthanarayanan, “Altruistic scheduling in multi-resource clusters,” in *Proc. USENIX Symp. Operating Syst. Design Implement. (OSDI)*, 2016, pp. 65–80.
- [35] R. Grandl, S. Kandula, S. Rao, A. Akella, and J. Kulkarni, “Graphene: Packing and dependency-aware scheduling for data-parallel clusters,” in *Proc. USENIX Symp. Operating Syst. Design Implement. (OSDI)*, 2016, pp. 81–97.
- [36] M. Grzegorz *et al.*, “Pregel: A system for large-scale graph processing,” in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2010, pp. 135–146.
- [37] B. Hindman *et al.*, “Mesos: A platform for fine-grained resource sharing in the data center,” in *Proc. USENIX Symp. Netw. Syst. Design Implement. (NSDI)*, 2011, p. 22.
- [38] C. Iorgulescu, F. Dinu, A. Raza, W. U. Hassan, and W. Zwaenepoel, “Don’t cry over spilled records: Memory elasticity of data-parallel applications and its application to cluster scheduling,” in *Proc. USENIX Annu. Tech. Conf. (ATC)*, 2017, pp. 97–109.
- [39] M. Isard, M. Budi, Y. Yu, A. Birrell, and D. Fetterly, “Dryad: Distributed data-parallel programs from sequential building blocks,” in *Proc. ACM SIGOPS/EuroSys Eur. Conf. Comput. Syst. (EuroSys)*, 2007, pp. 59–72.
- [40] M. Isard, V. Prabhakaran, J. Currey, U. Wieder, K. Talwar, and A. Goldberg, “Quincy: Fair scheduling for distributed computing clusters,” in *Proc. ACM SIGOPS 22nd Symp. Operating Syst. Princ. (SOSP)*, 2009, pp. 261–276.
- [41] E. G. Joseph, S. X. Reynold, D. Ankur, C. Daniel, J. F. Michael, and S. Ion, “GraphX: Graph processing in a distributed dataflow framework,” in *Proc. USENIX Symp. Operating Syst. Design Implement. (OSDI)*, 2014, pp. 599–613.

- [42] K. Karanasos *et al.*, “Mercury: Hybrid centralized and distributed scheduling in large shared clusters,” in *Proc. USENIX Annu. Tech. Conf. (ATC)*, 2015, pp. 485–497.
- [43] S. Kim *et al.*, “GPUnet: Networking abstractions for GPU programs,” in *Proc. USENIX Symp. Operating Syst. Design Implement. (OSDI)*, 2014, pp. 201–216.
- [44] R. Kondor and T. Jebara, “A kernel between sets of vectors,” in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2003, pp. 361–368.
- [45] T. Koponen *et al.*, “Onix: A distributed control platform for large-scale production networks,” in *Proc. OSDI*, 2010, pp. 1–6.
- [46] M. Kornacker *et al.*, “Impala: A modern, open-source SQL engine for Hadoop,” in *Proc. CIDR*, 2015, p. 9.
- [47] M. Korupolu, A. Singh, and B. Bamba, “Coupled placement in modern data centers,” in *Proc. IEEE Int. Symp. Parallel Distrib. Process. (IPTPS)*, May 2009, pp. 1–12.
- [48] A. Kuzmanovska, R. H. Mak, and D. Epema, “KOALA-F: A resource manager for scheduling frameworks in clusters,” in *Proc. 16th IEEE/ACM Int. Symp. Cluster, Cloud Grid Comput. (CCGrid)*, May 2016, pp. 80–89.
- [49] B. Li *et al.*, “ClickNP: Highly flexible and high performance network processing with reconfigurable hardware,” in *Proc. ACM SIGCOMM Conf.*, Aug. 2016, pp. 1–14.
- [50] L. Liu and H. Xu, “Elasticator: Elastic executor scheduling in data analytics systems,” in *Proc. ACM Symp. Cloud Comput.*, Oct. 2018, pp. 107–120.
- [51] Y. Lu, A. Chowdhery, and S. Kandula, “Optasia: A relational platform for efficient large-scale video analytics,” in *Proc. 7th ACM Symp. Cloud Comput.*, Oct. 2016, pp. 57–70.
- [52] H. Mao, M. Alizadeh, I. Menache, and S. Kandula, “Resource management with deep reinforcement learning,” in *Proc. 15th ACM Workshop Hot Topics Netw.*, Nov. 2016, pp. 50–56.
- [53] M. Zaharia, T. Das, H. Li, T. Hunter, S. Shenker, and I. Stoica, “Discretized streams: Fault-tolerant streaming computation at scale,” in *Proc. 24th ACM Symp. Operating Syst. Princ.*, Nov. 2013, pp. 423–438.
- [54] C. Michele, R. Yan, and L. Zheng, “Adaptive kernel approximation for large-scale non-linear SVM prediction,” in *Proc. Int. Conf. Mach. Learn. (ICML)*, 2011, pp. 409–416.
- [55] K. Morton, M. Balazinska, and D. Grossman, “ParaTimer: A progress indicator for MapReduce DAGs,” in *Proc. Int. Conf. Manage. Data (SIGMOD)*, 2010, pp. 507–518.
- [56] K. Ousterhout, C. Canel, S. Ratnasamy, and S. Shenker, “Monotasks: Architecting for performance clarity in data analytics frameworks,” in *Proc. 26th Symp. Operating Syst. Princ.*, Oct. 2017, pp. 184–200.
- [57] K. Ousterhout, R. Rasti, S. Ratnasamy, S. Shenker, and B.-G. Chun, “Making sense of performance in data analytics frameworks,” in *Proc. USENIX Symp. Netw. Syst. Design Implement. (NSDI)*, 2015, pp. 293–307.
- [58] K. Ousterhout, P. Wendell, M. Zaharia, and I. Stoica, “Sparrow: Distributed, low latency scheduling,” in *Proc. 24th ACM Symp. Operating Syst. Princ.*, Nov. 2013, pp. 69–84.
- [59] J. Rasley, K. Karanasos, S. Kandula, R. Fonseca, M. Vojnovic, and S. Rao, “Efficient queue management for cluster scheduling,” in *Proc. ACM Eur. Conf. Comput. Syst. (EuroSys)*, 2016, pp. 1–15.
- [60] A. Roy, H. Zeng, J. Bagga, G. Porter, and A. C. Snoeren, “Inside the social network’s (datacenter) network,” in *Proc. ACM Conf. Special Interest Group Data Commun.*, Aug. 2015, pp. 123–137.
- [61] B. Schölkopf, P. Bartlett, A. Smola, and R. Williamson, “Shrinking the tube: A new support vector regression algorithm,” in *Proc. Adv. Neural Inf. Process. Syst. (NIPS)*, 1999, pp. 330–336.
- [62] M. Schwarzkopf, A. Konwinski, M. Abd-El-Malek, and J. Wilkes, “Omega: Flexible, scalable schedulers for large compute clusters,” in *Proc. 8th ACM Eur. Conf. Comput. Syst. (EuroSys)*, 2013, pp. 351–364.
- [63] A. Singh *et al.*, “Jupiter rising: A decade of Clos topologies and centralized control in Google’s datacenter network,” in *Proc. ACM SIGCOMM*, 2015, pp. 1–15.
- [64] A. J. Smola and B. Schölkopf, “A tutorial on support vector regression,” *Statist. Comput.*, vol. 14, no. 3, pp. 199–222, 2004.
- [65] J. Son, Y. Xiong, K. Tan, P. Wang, Z. Gan, and S. Moon, “Protego: Cloud-scale multitenant IPsec gateway,” in *Proc. USENIX Annu. Tech. Conf. (ATC)*, 2017, pp. 473–485.
- [66] A. Toshniwal *et al.*, “Storm@Twitter,” in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, 2014, pp. 147–156.
- [67] A. Toshniwal *et al.*, “Morpheus: Towards automated SLOs for enterprise clusters,” in *Proc. USENIX Symp. Operating Syst. Design Implement. (OSDI)*, 2016, pp. 117–134.
- [68] V. K. Vavilapalli *et al.*, “Apache Hadoop YARN: Yet another resource negotiator,” in *Proc. ACM Symp. Cloud Comput. (SoCC)*, 2013, pp. 1–16.
- [69] S. Venkataraman, Z. Yang, M. Franklin, B. Recht, and I. Stoica, “Ernest: Efficient resource prediction for large-scale advanced analytics,” in *Proc. USENIX Symp. Netw. Syst. Design Implement. (NSDI)*, 2016, pp. 363–378.
- [70] A. Verma, L. Pedrosa, M. Korupolu, D. Oppenheimer, E. Tune, and J. Wilkes, “Large-scale cluster Management at Google with Borg,” in *Proc. ACM Eur. Conf. Comput. Syst. (EuroSys)*, 2015, pp. 1–17.
- [71] J. Wang and M. Balazinska, “Elastic memory management for cloud data analytics,” in *Proc. USENIX Annu. Tech. Conf. (ATC)*, 2017, pp. 745–758.
- [72] M. Weimer *et al.*, “REEF: Retainable evaluator execution framework,” in *Proc. ACM SIGMOD Int. Conf. Manage. Data*, May 2015, pp. 1343–1355.
- [73] G. J. Woeginger, “There is no asymptotic PTAS for two-dimensional vector packing,” *Inf. Process. Lett.*, vol. 64, no. 6, pp. 293–297, Dec. 1997.
- [74] W. Xia, H. Jiang, D. Feng, and Y. Hua, “SiLo: A similarity-locality based near-exact deduplication scheme with low RAM overhead and high throughput,” in *Proc. USENIX Annu. Tech. Conf. (ATC)*, 2011, pp. 26–30.
- [75] D. Xie, N. Ding, Y. C. Hu, and R. Kompella, “The only constant is change: Incorporating time-varying network reservations in data centers,” in *Proc. ACM SIGCOMM*, 2012, pp. 199–210.
- [76] G. Xu and C.-Z. Xu, “Prometheus: Online estimation of optimal memory demands for workers in in-memory distributed computation,” in *Proc. ACM Symp. Cloud Comput. (SoCC)*, 2017, p. 655.
- [77] G. Xu, C.-Z. Xu, and S. Jiang, “Prophet: Scheduling executors with time-varying resource demands on data-parallel computation frameworks,” in *Proc. IEEE Int. Conf. Autonomic Comput. (ICAC)*, Jul. 2016, pp. 45–54.
- [78] Y. Yan, Y. Gao, Y. Chen, Z. Guo, B. Chen, and T. Moscibroda, “TR-spark: Transient computing for big data analytics,” in *Proc. 7th ACM Symp. Cloud Comput.*, Oct. 2016, pp. 484–496.
- [79] M. Zaharia, D. Borthakur, J. S. Sarma, K. Elmeleegy, S. Shenker, and I. Stoica, “Delay scheduling: A simple technique for achieving locality and fairness in cluster scheduling,” in *Proc. ACM Eur. Conf. Comput. Syst. (EuroSys)*, 2010, pp. 265–278.
- [80] M. Zaharia *et al.*, “Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing,” in *Proc. USENIX Symp. Netw. Syst. Design Implement. (NSDI)*, 2012, pp. 15–28.
- [81] H. Zhang, L. Stafman, A. Or, and M. J. Freedman, “SLAQ: Quality-driven scheduling for distributed machine learning,” in *Proc. Symp. Cloud Comput.*, Sep. 2017, pp. 390–404.



Libin Liu (Member, IEEE) received the B.E. degree in software engineering from Shandong University and the Ph.D. degree from the Department of Computer Science, City University of Hong Kong. He is currently a Researcher with the Theory Lab, Huawei Hong Kong Research Center, Hong Kong. His current research interests include data analytics systems and machine learning for networking. He is a member of the ACM.



Hong Xu (Senior Member, IEEE) received the B.Eng. degree from The Chinese University of Hong Kong in 2007, and the M.A.Sc. and Ph.D. degrees from the University of Toronto in 2009 and 2013, respectively. From 2013 to 2020, he was with City University of Hong Kong. He is currently an Associate Professor with the Department of Computer Science and Engineering, The Chinese University of Hong Kong. His research interests include computer networking and systems, particularly big data systems and data center networks. He is a member of the ACM. He was a recipient of an Early Career Scheme Grant from the Hong Kong Research Grants Council in 2014. He received three best paper awards, including the IEEE ICNP 2015 Best Paper Award.