



HORIZON 2020 FRAMEWORK PROGRAMME

IOStack

(H2020-644182)

**Software-Defined Storage for Big Data
on top of the OpenStack platform**

D5.2 System Deployment Strategies

Due date of deliverable: 31-12-2016
Actual submission date: 31-12-2016

Start date of project: 01-01-2015

Duration: 36 months

Summary of the document

Document Type	Deliverable
Dissemination level	Public
State	v1.0
Number of pages	56
WP/Task related to this document	WP5 / T5.2
WP/Task responsible	EUR
Leader	Francesco Pace (EUR)
Technical Manager	Pietro Michiardi (EUR)
Quality Manager	Yosef Moatti (IBM)
Author(s)	Francesco Pace (EUR), Daniele Venzano (EUR), Marco Mianesio (EUR), Pietro Michiardi (EUR)
Partner(s) Contributing	URV, GDP, BSC and IBM
Document ID	IOSTACK_D5.2_Public.pdf
Abstract	<p>This work addresses the problem of scheduling user-defined analytic applications, which we define as high-level compositions of frameworks, their components, and the logic necessary to carry out work. The key idea in our application definition, is to distinguish classes of components, including rigid and elastic types: the first being required for an application to make progress, the latter contributing to reduced execution times. We show that the problem of scheduling such applications poses new challenges, which existing approaches address inefficiently. Thus, we present the design and evaluation of a novel, flexible heuristic to schedule analytic applications, that aims at high system responsiveness, by allocating resources efficiently. Our algorithm is evaluated using trace-driven simulations, with large-scale real system traces: our flexible scheduler outperforms a baseline approach across a variety of metrics, including application turnaround times, and resource allocation efficiency.</p> <p>We also present the design and evaluation of a full-fledged system, which we have called <i>Zoe</i>, that incorporates the ideas presented in this work, and report concrete improvements in terms of efficiency and performance, with respect to prior generations of our system.</p>
Keywords	scheduling, core, elastic, analytics applications, distributed

History of changes

Version	Date	Author	Summary of changes
0.1	18/11/2016	F. Pace	Initial text
0.2	21/11/2016	D. Venzano	Review structure, Zoe details
0.9	25/11/2016	F. Pace	Add appendix
1.0	07/12/2016	F. Pace	Moved appendix into appropriate sections and implemented fixes from first review

Table of Contents

I Zoe: Analytics as a Service	2
1 History and motivations	2
2 Zoe applications	2
3 Implementation and architecture	3
3.1 Internal architecture	3
3.1.1 State	4
3.1.2 Scheduling and placement	5
3.2 Back-ends	5
4 User interaction	5
5 Road Map	6
5.1 Research topics	6
5.1.1 Scheduling	6
5.1.2 Dynamic resource allocation	7
5.2 Industrialization	7
5.2.1 Standard distributions	7
5.2.2 User interfaces	7
5.2.3 Advanced back-ends and long running applications	8
5.2.4 Authentication and user management	8
6 IOStack and the Zoe community	8
6.1 Zoe and IOStack	8
6.2 Community and partners	8
6.2.1 Air France	9
6.2.2 KPMG	9
II Analysis of FIFO and size-based scheduling policies	10
7 Introduction	10
8 Definitions and Problem statement	11
8.1 Definitions	11
8.2 Problem Statement	12
9 A Flexible Scheduling Algorithm	14
9.1 Design guidelines	14
9.2 Algorithm Details	15
9.3 Preemptive policies	15
10 Numerical evaluation	15
10.1 Methodology	15
10.2 Comparison with the baseline	18
10.3 Comparison between different definitions of size	29
10.4 Impact of different workloads	45
10.5 Preemption	45

11 Experiments with Zoe	46
12 Related Work	51
III Conclusions	53
13 Conclusions	53

Executive summary

This deliverable presents Zoe, the system built to deploy analytics applications on distributed clusters. Distributed analytics applications are taking more and more space in data centers all over the world. Zoe provides the capability of executing analytics jobs, batch or interactive, in a simple way, optimizing placement and resource usage.

In the first part of the document we describe the software implementation of Zoe, the motivations behind its development and the interest it quickly generated in several well-known companies. In the second part of the deliverable we present a novel approach to scheduling elastic applications and its implementation in Zoe. Optimization objectives, the algorithms and the scheduling policies are explained, together with the evaluation, performed with simulations and a prototype implementation.

Zoe plays a fundamental role in IOStack, by providing the scheduling and placement layer needed to efficiently deploy analytic applications on top of the various storage solutions offered by the SDS controller.

Part I

Zoe: Analytics as a Service

1 History and motivations

Zoe was born in August 2015 as an open source project[1] to satisfy the need of having an easy and integrated way to deploy distributed analytic applications on a cluster of physical or virtual machines. Users can define analytic applications starting from a number of ready-made building blocks, Zoe will schedule and deploy them matching resources requests and availability.

Zoe is developed in Python and is conceived as a thin layer that builds on top of an existing low-level cluster management system, which is used as a back-end to provision resources to applications. Raising the level of abstraction to manipulate analytic applications is beneficial for users and ultimately to the system design itself: application scheduling decisions can be taken with a small amount of state information, and do not happen at the same (extremely fast) pace at which low-level task scheduling does.

As such Zoe is not in competition with cluster management systems like Mesos or Kubernetes, but builds on top of them (see section 3.2), to provide efficient application-level scheduling and a more accessible user experience.

2 Zoe applications

Zoe schedules applications. Each application is made of one or more components, that are run each in its own Linux container. For example, the Spark Notebook application that a user submits to Zoe is made of one Jupyter Notebook[2, 3] component, one Apache Spark[4, 5] master and one or more Apache Spark workers.

In order to produce useful work, in this case for the application to be useful to the user, there is a core set of components that can be identified: the notebook, the master and just one worker. The application can mark additional workers as “optional” (elastic in Zoe’s terminology). Zoe will start them only if there are free and unused resources.

To simplify application descriptions and build a library of building blocks, an intermediate concept of frameworks (groups of components that work together) has been introduced.

Zoe applications are described by users via a simple JSON description that follows a high-level configuration language (CL) to specify applications, frameworks and components with their classes (core or elastic), resource reservations and constraints. The CL is simple and extensible: it aims at conciseness and, with framework templates, can be also used by “casual”, in addition to “power” users. An example of the simplicity and effectiveness of the Zoe CL, building a batch application for the distributed version of Tensorflow[6, 7], required less than 25 lines of CL.

A typical Zoe application description contains a number of application-global metadata items, such as a name, a version (used for input format validation), size/priority¹ and others. Then it contains a list of components, each with its own metadata, that includes:

- image: the Docker image location for the component
- environment: the environment variables that are used to configure the component
- volumes: external storage to mount inside the container
- resources: the resource reservations required to run one instance of this component
- total count: the total number of instances of this component that can be started
- essential count: the minimum number of instances (out of total count) required for the application to produce useful work

¹The size-based scheduler interprets this field as a size hint: the expected duration of the application in minutes. Part of future work is building an automated size estimation module.

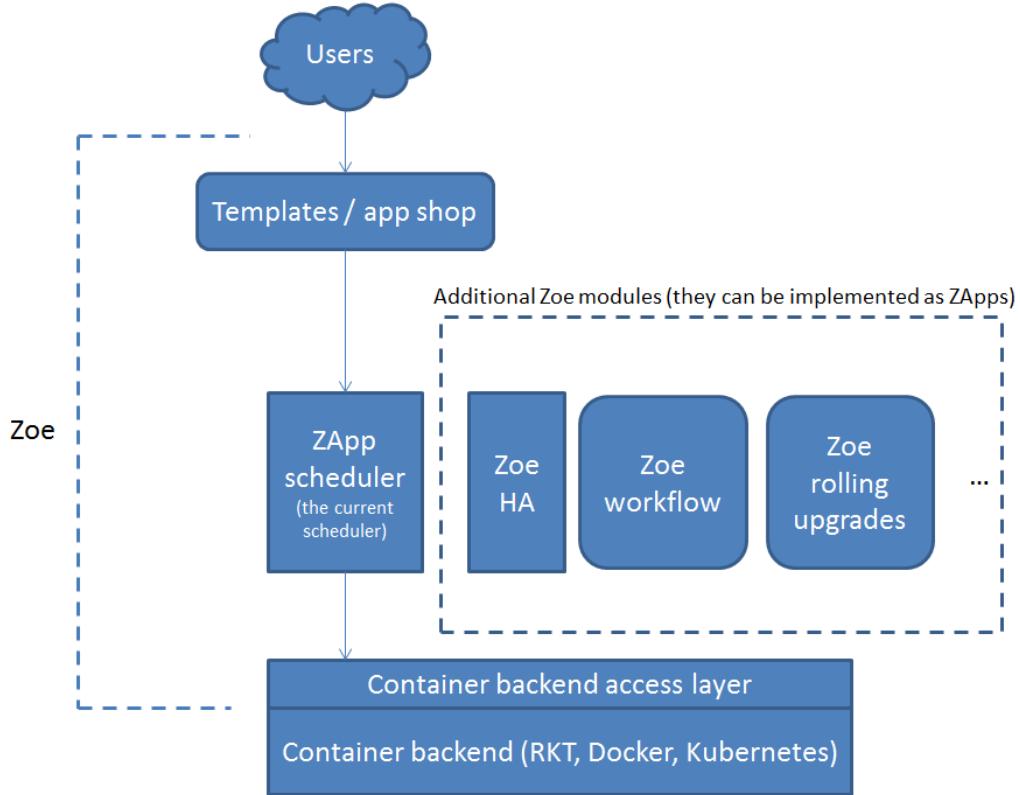


Figure 1: Zoe high-level architecture.

Currently a number of scripts help users create these application descriptions, but we are planning and developing a web-based tool in the style of an app-shop where users are free to compose and configure their own applications in a more intuitive way.

3 Implementation and architecture

Zoe has been implemented in Python, a language that provides a very fast development cycle, with a rich library of ready-made components and an active community. Python has already been used to build big, complex systems used in the industry: a good example being OpenStack.

At high level the architecture of Zoe can be seen in figure Fig. 1. Users interact with the system by submitting application descriptions generated from a set of templates. These descriptions are processed by the main Zoe processes and queued by the scheduler. When enough resources are available, the applications are executed by generating appropriate API commands to the underlying container back-end.

Optional modules can augment the capabilities of the Zoe application description language to:

- Expose functionality provided by advanced back-ends (high availability, fault tolerance, load balancing)
- Build full workflow functionality on top of Zoe applications
- Perform maintenance operations on long-running applications, for example rolling upgrades on a distributed storage cluster

More details on modules and back-end support are in sections 3.2 and 5

3.1 Internal architecture

Zoe is divided into two multi-threaded processes. The Zoe Master and the Zoe API. Both store state information into an external Postgres database, that has well-known reliability and fault tolerance

characteristics.

Users interact with the Zoe API process. It offers the web interface and the REST API, its main task are:

- validate user application descriptions
- provide user interfaces to the state stored in the database
- perform clean-ups of terminated executions
- re-submit executions in case of crashes of the Master process

The Master process contains the scheduler and other threads to process various events. Its tasks are:

- Schedule executions according to resource availability
- Monitor the container back-end for events that happen asynchronously from Zoe, for example components that end (normally or because of an error) by themselves
- Record per-node free/busy resources
- Manage execution termination and save application logs

The two processes use a Zero-MQ based protocol, that has been developed using the Lazy Pirate Pattern [8] to provide client-side communication reliability in the event of crashes or disconnections.

The Master and the API processes do not store any state internally and can be restarted at will in case of upgrades or crashes without any consequence. The API process can be scaled horizontally behind a http-based load balancer.

3.1.1 State

Three main tables are maintained in the state store, the platform table, the executions table and the service table. In this context executions are instances of Zoe applications, while services are instances of components.

The platform table stores information about the number of containers and the amount of free resources for each node available in the cluster. This information is taken as-is from the back-end API and stored into the database for caching reasons. Each entry describes the resources (memory, cores and containers) free and total for one single node. Data is updated whenever the scheduler is triggered.

Each entry in the execution table refers to a single application execution submitted by the user. It records information such as:

- identifier of the user who submitted the application
- timestamps of submission, start and termination events
- current status (submitted, running, error, etc.)
- error message in case the execution failed due to an error

Each entry in the services table refers to a single instance of a component. As was explained in the previous sections, application descriptions contain a list of component with total and essential counts for each of them. These descriptions are exploded into the state table in as many services as is the total number of instances requested. The information recorded for each instance is:

- status: the service status from Zoe point of view
- backend status: the container status from the back-end point of view

- backend identifier: the unique ID the back-end uses to identify this container
- service group: the component name as given into the application description
- error message: filled in when there is an error during the container lifetime (for example image not found or out-of-memory termination)

3.1.2 Scheduling and placement

The scheduler thread is triggered by several events:

- an execution being added to the scheduling queue
- an execution that terminates, either by itself (batch) or by the user
- a timer, to account for resources that are used outside of Zoe's control.

When the scheduler is triggered and selects an application execution to start, according to the configured policy, it performs a placement simulation, trying to find the best fit of core and elastic services in order to maximize the number of running executions. Executions for which all core services cannot be started at the same time are left in the queue, in order to start only executions that can produce useful work.

The queue is ordered according to the policy chosen by the administrator. The stable version of Zoe (0.10.x) implements a FIFO policy with back-filling, where all executions are started in order of arrival and if there are not enough resources for an execution, it is put back at the end of the queue.

We implemented also strict FIFO and size-based policies, that take into account the elasticity of Zoe applications. A detailed explanation, with experimental results, is available in part II of this deliverable.

These policies are currently available in an experimental branch. They will be released in the stable branch in the next few months, after internal interfaces and administration tools are updated to support the new policies.

The placement simulation uses a simple filter-and-sort algorithm. For each container to be placed it filters out nodes that do not match the required resource constraints (amount of memory, but also special hardware needs can be taken into account). The container will be placed in the node with the least amount of running containers and the biggest amount of free memory.

The same filter-and-sort algorithm is used to take placement decisions in most modern cluster manager systems, like OpenStack Nova and Docker Swarm.

3.2 Back-ends

The main design idea of Zoe is to hide the complexities of low-level resource provisioning from application scheduling and use an existing cluster management system, for which many alternatives exist, instead. Historically, Zoe builds on top of Docker Swarm and uses it for orchestration, dependency management, resource isolation, naming and networking.

As part of Zoe's roadmap we have planned to expand the number of supported backends, adding RKT, Docker Engine and Kubernetes. Mesos also is a possibility that we will evaluate during the next year.

As can be seen in figure 1, together with support for new back-ends, a number of additional (and optional) Zoe components will be developed and released. These components will expose to Zoe applications the capabilities offered by more advanced back-ends, such load balancing and fault tolerance.

4 User interaction

Users can interact with Zoe through a web interface or a command-line tool. The web interface provides a high level overview of the application executions for each user and their status. The command-line tool is more advanced and, for example, can be used to script the execution of multiple

batch applications, creating simple automated workflows. Zoe has also a REST client API that is being used by our external partners to develop additional tools and services.

When an application is submitted, via REST, command-line or web interface, Zoe creates an entry in the application state store, and adds it to a pending queue. The scheduler strives to making sure the application selected for execution can make progress as soon as resources are allocated to it: to this end, it relies on the back-end to place all core components according to the simulated plan. Elastic components are scheduled when possible and contribute to decreased application run-time. As a consequence, batch applications (either rigid like Tensorflow or MPI, or elastic such as Spark) can make progress as soon as core components start; similarly, interactive applications – which can be given precedence to reduce queuing times and improve user experience – can also start being used even if not all elastic components (if any) are scheduled. See section 9 for more details.

5 Road Map

The future of Zoe has been planned along two main directions: one is academical research and the other is industrialization, in collaboration with our partners (see section 6.2. Zoe is a strong and innovative project that in just one year of development has seen a lot of positive feedback that helped focus its scope. The roadmap described below is the result of many meetings, in person, on the phone and via the Internet and a substantial work of summarization. IOStack will benefit from all the points described below.

5.1 Research topics

Here we describe general research topics that can be studied with the help of Zoe, where internal components can be easily swapped with experimental ones, in the same way as the elastic scheduler described in part II of this document has been created.

5.1.1 Scheduling

The new scheduler described in part II, with full support for elastic applications, opens up several very interesting research topics, described below.

Scheduling policies Currently Zoe supports only three scheduling policies (six with non-elastic variants):

- Strict FIFO
- FIFO with backfilling
- SRPT (without preemption)

There are many policies, each with multiple variants, described in the literature for traditional schedulers, that can be applied to the concept of elastic applications. One example that we find particularly relevant for analytics applications is deadline scheduling, where the user specifies a time window (or a latency for streaming jobs) in which the job must be completed. If the system cannot satisfy the request it can try preempting other jobs with lower priority to free additional resources.

Size estimation Size estimation plays a fundamental role in all size-based scheduling policies. The current study leaves the task of estimating sizes to an external module, not specified. This research topic will explore the strategies that can be used to estimate job size, off-line and on-line, and the consequences of estimation errors in the quality of the final scheduling. The EUR research team has already important works published in this domain[9, 10, 11].

Preemption All advanced scheduling policies need some form of preemption to dynamically free resources. Here interesting research topics include which processes to preempt and how to preempt them, by killing and potentially losing valuable work already done, or by checkpointing and saving the process state, at the cost of higher latency?

5.1.2 Dynamic resource allocation

In all modern virtualization systems there is the concept of resource reservations: when a user requests a virtual machine or a container, she has to specify the amount of memory, disk and cores to reserve. These reservations are used to take placement decisions in cluster managers, but in reality we observe on the EUR cluster that the effective resource usage is much lower than the reservation, in the general case.

Dynamic resource allocation would use real resource usage to take placement decisions, instead of a user-provided reservation. The biggest challenge is how to react quickly enough when, for example, a process makes a big memory allocation and the node on which it is running is already full.

5.2 Industrialization

The Zoe deployments in Air France and KPMG (see section 6.2) have highlighted a number of features where further development will enhance the usefulness of Zoe in the industrial setting.

5.2.1 Standard distributions

Zoe depends on a number of technologies that can be easily exchanged according to the scale of the deployment, high availability requirements, etc. These are:

- Virtualization back-end
- State store
- Monitoring and logging pipeline
- Shared filesystem

We are planning three main "Zoe Distributions" that will come with full installation instructions:

1. demo: can be installed on a laptop and showcases Zoe and a few very simple analytics applications
2. basic: for small deployments that have no particular needs in terms of data retention or fault tolerance
3. advanced: for enterprise-grade deployments

Each of these distributions will be mostly in the form of documentation and scripts. They will explain and, whenever possible, automate the installation of the set of software dependencies for Zoe more suited for the deployment. For example in the demo distribution a simple non-distributed Docker engine will be used as the virtualization backend, while in the advanced distribution Kubernetes will be proposed.

5.2.2 User interfaces

As in any software, the user interface is the most visible component of Zoe and, as such, it received the biggest number of suggestions for new features and improvements.

Application builder We are planning a Zoe web application builder that will let users manage and build Zoe applications quickly and easily. This builder will have two main areas, that can be enabled or disabled at install time. One area will be an app-store like interface with ready-made applications built by system administrators, where users cannot change any system-related option. The other area will be dedicated to "power users": they will be able to compose frameworks (pre-packaged building blocks) together, set memory limits and, in general, access lower-level options.

Web interface The new Zoe web interface will cover three main user types:

- Administrator: this section of the interface gives access to administrative tools to monitor the activity of the whole Zoe system and the applications running in it. In this area there are also tools for user administration and access control, settings related to the scheduler policies, etc.
- Normal user: a normal user can only start, stop and monitor its own executions
- Advanced analytic user: with this interface we want to offer a very different approach, to target the way data scientists usually work. With this interface it will be possible to create workflows, define pre-processing and post-processing steps and reason in terms of data and operations to perform on them.

5.2.3 Advanced back-ends and long running applications

Currently Zoe supports only Docker Swarm as a container backend. Adding support for more back-ends will let Zoe users define more complex applications (or run Zoe on a single laptop).

We are planning support for:

- Docker Engine: a single PC can be used as the Zoe backend. It comes almost for free since the API is almost the same as the Swarm one. Very useful for demos, development environments and for running small Zoe applications on the go.
- Kubernetes: for large deployments, where long running applications (i.e. a data layer like HDFS) have to be managed with rolling upgrades, load balancing and fault tolerance. Zoe can, through dedicated modules, expose the functionality already available in Kubernetes to high level Zoe applications.

5.2.4 Authentication and user management

Companies run a large number of different authentication and authorization suites. KPMG already built upon the existing LDAP support, contributing Kerberos and OpenAuth modules for Zoe.

We expect to have more schemes supported, together with the necessary administration interfaces to manage users, groups, capabilities and quotas.

6 IOStack and the Zoe community

6.1 Zoe and IOStack

IOStack promotes the separation of compute and storage layers allowing increased flexibility in meeting the variable demand of computation resources, while keeping the data storage system in a stable state. Zoe targets the compute layer of IOStack, by providing a simple way to define and deploy arbitrary analytic applications.

Zoe has been integrated into the IOStack dashboard to give users a single point of access to all IOStack services. Zoe applications can use the storage services offered by the other IOStack components directly (Swift) or through the back-end (volumes). In the last year of the project we envision an even tighter integration with information flowing from the Zoe scheduler to the SDS controller and vice-versa, opening the possibility of automatically taking informed decisions about application scheduling and storage features at run-time.

6.2 Community and partners

Since its beginning in August 2015 Zoe has generated interest in the industry. We have contacts and ongoing discussions with a number of companies and public entities. People feel the need of a simpler interface on top of existing virtualization solutions that lets them concentrate on the problems they face, typically data analysis, and forget all about the lower layers of resource management, networks and virtualization.

Below we go into the details of the two major collaborations we have started thanks to Zoe.

6.2.1 Air France

Air France started experimenting with the earlier prototypes already at the end of 2015, providing bug reports, insights and improvement suggestions. They officially adopted Zoe in September 2016 for running applications for their data analysis teams located in Paris and Amsterdam. During one of the meetings with Air France, they reported that the data scientists involved are enthusiastic of Zoe and the ways they are able to work with it.

Air France is currently keeping a private fork of Zoe and having internal discussions on how to contribute their changes into the public repository. In the meantime they keep providing feedback and suggestions to make Zoe more attractive to production deployments in the industry.

6.2.2 KPMG

After evaluating Zoe, Mesos and Kubernetes, KPMG decided to invest in Zoe to make it the primary software product used to deploy their services for all their customers. In September 2016, KPMG hired two software engineers that work full-time on Zoe, in strict collaboration with Eurecom. We have established a collaboration contract and are in the process of writing a shared project specification that will set priorities and assign tasks to developers.

Part II

Analysis of FIFO and size-based scheduling policies

7 Introduction

The last decade has witnessed the proliferation of numerous distributed frameworks to address a variety of large-scale data analytics and processing projects. First, MapReduce [12] has been introduced to facilitate the processing of bulk data. Subsequently, more flexible tools, such as Dryad [13], Spark [14], Flink [15] and Naiad [16], to name a few, have been conceived to address the limitations and rigidity of the MapReduce programming model. Similarly, specialized libraries [17] and systems like TensorFlow [6] have seen the light to cope with large-scale machine learning problems. In addition to a fast growing ecosystem, individual frameworks are driven by a fast-pace development model, with new releases every few months, introducing substantial performance improvements. Since each framework addresses specific needs, users are left with a wide choice of tools and combination thereof, to address the various stages of their projects.

The context depicted above has driven a lot of research [18, 19, 20, 21, 22, 23, 24, 25, 26, 27, 28, 29, 30, 31] in the area of resource allocation and scheduling, both from academia and the industry. These efforts materialize in cluster management systems that offer simple mechanisms for users to request the deployment of the framework they need. The general underlying idea is that of sharing cluster resources among a heterogeneous set of frameworks, as a response to static partitioning, which has been dismissed for it entails low resource utilization [18, 19, 20]. Existing systems divide the resources at different levels. Some of them, e.g. Mesos and YARN, target low-level orchestration of distributed computing frameworks: to this aim, they require non-trivial modifications of such frameworks to operate correctly. Others, e.g. Kubernetes [32] and Docker Swarm [33], focus on provisioning and deployment of containers, and are thus oblivious to the characteristics of the frameworks running in such containers. To the best of our knowledge, no existing tool currently addresses the problem of scheduling analytic applications as a whole, leveraging the intrinsic properties of the frameworks such applications use, but without requiring substantial modification of such frameworks.

The endeavor of this work is to fill the gap that exists in current approaches, and raise the level of abstraction at which scheduling works. We introduce a general and flexible definition of applications, how they are composed, and how to execute them. For example, a user application addressing the training of a statistical model involves: a user-defined program implementing a learning algorithm, a framework (e.g., Spark) to execute such a program together with information about its resource requirements, the location for input and output data and possibly hyper-parameters exposed as application arguments. Users should be able to express, in a simple way, how such an application must be packaged and executed, submit it, and expect results as soon as possible.

We show that scheduling such applications represents a departure from what has been studied in the scheduling literature, and present the design of a new algorithm to address the problem. A key insight of our approach is to exploit the properties of the frameworks used by an application, and distinguish their components according to classes, core and elastic: the first being required for an application to produce work, the latter contributing to reduced execution times. Our heuristic focuses cluster resources to few applications, and uses the class of application components to pack them efficiently. Our scheduler aims at high cluster utilization and a responsive system. It can easily accommodate a variety of scheduling policies, beyond the traditional “first-come-first-served” or “processor sharing” strategies, that are currently used by most existing approaches. We study the performance of our scheduler using realistic, large-scale workload traces from Google [34, 35], and show it consistently outperforms the existing baseline approach which ignores component classes: application turnaround times are more than halved, and queuing times are drastically reduced. This induces fewer applications waiting to be served, and increases resource allocation up to 20% more

than the baseline.

Finally, we present a full-fledged system, called Zoe, that schedules analytic applications according to our original algorithm and that can use sophisticated policies to determine application priorities. Our system exposes a simple and extensible configuration language that allows application definition. We validate our system with real-life experiments, and report conspicuous improvements when compared to a baseline scheduler, when using a representative workload: median turnaround times are reduced by up to 37% and median resource allocation is 20% higher.

In summary, the contributions of our work are as follows:

- We define, for the first time, a high-level construct to represent analytic applications, focusing on their heterogeneity, and their end-to-end life-cycle;
- We establish a new scheduling problem, and propose a flexible heuristic capable of handling heterogeneous requests, as well as a variety of scheduling policies, with the ultimate objective of improving system responsiveness under heavy loads;
- We evaluate our scheduling policy using realistic, large-scale workload traces and show it consistently outperforms the baseline approach;
- We build a system prototype which materializes the ideas of analytic applications and their scheduling. Our system has been in use for over one year, serving a variety of analytic application workloads. Using our new heuristic, we were able to achieve substantial improvements in terms of system responsiveness and cluster utilization.

We start by clarifying what analytic applications are, give examples and formulate our problem statement in Section 8. We then describe the details of our flexible scheduling heuristic, in Section 9, which we evaluate using simulations in Section 10. The system implementation is described in Part I, and its evaluation is presented in Section 11. Finally, in Section 12 and Section 13 we discuss related work and conclude, hinting at our current research agenda.

8 Definitions and Problem statement

8.1 Definitions

We define a data analytics **framework** as a set of one or more software **components** (executable binaries) to accomplish some data processing tasks. Distributed frameworks are generally composed by a controller, a master and a number of worker components. Examples of distributed frameworks are Apache Spark [5], Google TensorFlow [7] and MPI [36]. Another example of simple data analytics framework we consider is an interactive Notebook [2].

Distributed frameworks require a scheduler to orchestrate their work: they execute jobs, each of which consists of one or more tasks that run in parallel the same program. Such schedulers operate at the task level: they assign tasks to workers, and they are highly specialized to take into account the peculiarities of each framework.

Framework schedulers such as Mesos [18] and Yarn [31] introduce an additional scheduling component to share cluster resources among concurrent frameworks: sharing policies are based on simple variations of Processor Sharing. Similarly, cluster management systems such as Docker Swarm [33] and Kubernetes [32] use a scheduler that assigns resources to generic frameworks. The problem to solve is the efficient allocation of resources by placing framework components and their tasks on cluster machines that satisfy a set of constraints.

We are now ready to define **analytics applications**, which are the elements we schedule in our work. Our main objective is to raise the level of abstraction by manipulating an abstract entity encompassing one or more analytics frameworks, their components and the necessary logic for them to cooperate toward producing useful work by running user-defined jobs. What sets apart our work from the state of the art is that our scheduler takes into account the notion of **component classes**, which allows modeling the specificity of each framework. We have found two distinct component classes to be sufficient to model existing analytic frameworks: thus, framework components either

belong to a **core** or to an **elastic** class. Core components are compulsory for a framework to produce useful work; elastic components, instead, optionally contribute to a job, e.g. by decreasing its runtime². Consider, for example, Spark. To produce work, it needs some core components: a controller (the spark client running the DAG scheduler), a master (in a standalone deployment), and one worker (running executors). We treat additional workers as elastic components. An alternative example is an application using TensorFlow, which only works with core components: one or more parameter servers and a number of workers. These two frameworks have substantially different runtime behavior: Spark is an elastic framework that can dynamically integrate workers to dispatch tasks. TensorFlow is rigid, and uses only core components to make progress.

To summarize, the nature of an application is that of raising the level of abstraction and an application is considered as being a collection of frameworks and their heterogeneous components as a single entity to schedule and allocate in a cluster of computers.

8.2 Problem Statement

We now treat the applications defined above as abstract entities that we call requests: they include one or more components, which belong to a given class, either core or elastic. In the literature, the classical problem of scheduling generic requests to be served by a distributed system has been extensively studied [37, 38, 39]. Requests composed solely by core components are usually referred to as rigid, while requests composed solely by elastic components are referred to as moldable (if the assigned resources are decided when the request is served and they do not change for the whole execution) or malleable (if the resources can vary during the execution³). A key difference with respect to previous work is that we consider heterogeneous requests, composed by both core and elastic components.

For simplicity of exposition, we assume system resources that can be measured in units, and that there are R available units overall to satisfy the requests. Each request i specifies the amount of units for its core and elastic components, labeled C_i and E_i respectively. Ideally, with enough available resources, a request is allocated all of its components: in this case, we define the service (or execution) time as T_i . The amount of work to satisfy a request is the area of the square $W_i = T_i \times (C_i + E_i)$. More generally, a request is allocated at least $C_i + x_i(t)$ resources, where $0 \leq x_i(t) \leq E_i$. Then, the service time is $T'_i = \frac{W_i}{C_i + x_i(t)}$. This simple model allows updating the service time T'_i when a scheduling decision modifies $x_i(t)$, by measuring the amount of work accomplished so far, and by computing the remaining amount of work to be done. While more complex models to describe T'_i can be conceived, for example taking into account the multi-dimensional nature of system resources, our simple approximation doesn't affect the nature of the scheduling problem we are studying.

Essentially, the problem of scheduling the execution of an incoming workload of requests amounts to: *i*) sorting requests to decide in which order to serve them; *ii*) allocating distributed resources to requests selected for service. The sorting phase can be solved using naive approaches, e.g. FIFO ordering, or more sophisticated policies, that use request size information. Even more generally, requests can be placed into "pools" and be assigned priorities, to mimic the hierarchical organization of the users, for example. The allocation phase is more tricky: in the abstract, it is a "packing" problem that has to decide how to shape requests being served. Even assuming service times to be known a-priori (e.g., T_i is given as an input), it is well known that the on-line scheduling problem is NP-hard [37]. Therefore, we need to find a suitable heuristic to approximate a solution to the scheduling optimization problem. In our case, it amounts to minimizing the application turnaround times, which is the interval of time between request i submission and its completion. In the context we consider, optimizing the average turnaround time represents a meaningful performance metric, as it caters system responsiveness. Next, we motivate our problem with a simple illustrative example.

²We currently assume that the contribution of elastic components is linear; this might not always be the case depending on the load of the system and the elasticity defined by the user. For now we trust that the user will correctly build her application without adding elastic components that will not be used. We plan to tackle this problem in our future work.

³An example of malleable framework is Spark[4]. Worker can be added or removed without destroying the application execution.

Illustrative example. We consider a system with 10 available resource units, and four requests waiting to be served, as shown in Figure 2. Each request needs 3 units for the core components, and different units for the elastic components. For each request, $T_i = 10$.

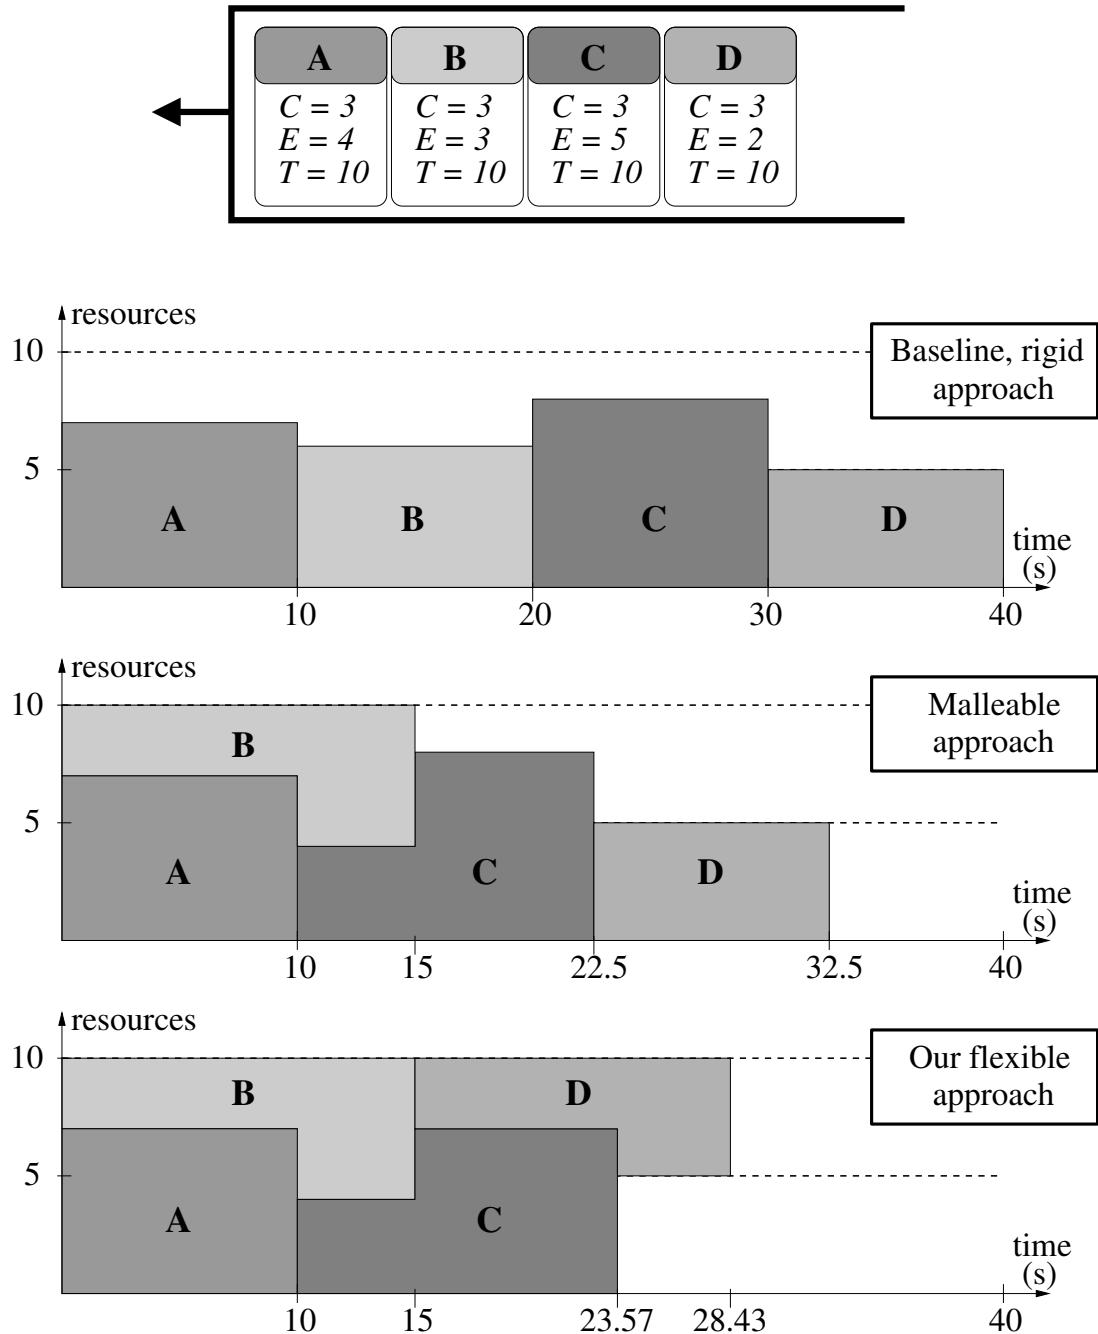


Figure 2: Illustrative examples of request scheduling: (top) rigid, (middle) malleable, (bottom) flexible approaches.

Given these requests, a traditional, rigid approach to scheduling – which does not make the distinction between component classes – assigns all required resources to each request. Since all requests need at least 5 units ($C_i + E_i \geq 5$), and since any pair of requests have an aggregated need that exceeds 10 units, the scheduler serves one request at a time (Figure 2, top): the average turnaround time is 25s. Note that, in this case, backfilling is not possible, i.e., even by changing the order in which requests are served the situation does not change.

Another scheduling approach comes from the literature of malleable job scheduling. The sched-

uler assigns all resources to the first request in the waiting line, then assigns the remaining resources (if any) to the next request, and so on, until no more free resources are available. This heuristic has been shown to be close to optimal [38]. Figure 2, middle, illustrates the idea: request B can be served along with request A. When request A has completed, the scheduler first assigns more resources to request B, and then tries to serve the next request. Similarly, when request B has completed, the scheduler first assigns more resources to request C, then attempts at serving request D. However, since request D needs at least $C_i = 3$ units, the scheduler is blocked (note that request C uses 8 units), so request D needs to wait, and some system resources remain unused. The average turnaround time is 20s.

In this work we advocate the need for a new approach to scheduling, which distinguishes component classes. The idea is to exploit the flexibility of elastic components and use system resources more efficiently. Intuitively, a solution to the problems of existing heuristics is to reclaim some resources assigned to elastic components of a running request and assign them to a pending request. This is shown in the bottom of Figure 2: the scheduler reclaims just one unit from request C so that it can provide 3 units to request D, which are sufficient for starting its core components and produce useful work. With this approach, the average turnaround is 19.25s.

While the above solution seems simple, it poses many challenges: how many units assigned to elastic components can be sacrificed for serving the next request? How many requests should be served concurrently? Should the scheduler focus on core components alone, to make sure many requests are served concurrently? How can scheduling take into account the priorities assigned by the sorting phase?

The last point introduces an additional challenge, related to preemptive scheduling policies. If a high priority request arrives, since it is not possible to interrupt core components – for this would kill the request – how can we select and preempt elastic components to accommodate the new request?

Given heterogeneous, composite requests, which are neither rigid, nor malleable (but both), available scheduling heuristics in the literature fall short in addressing the sorting and allocation problems: a new approach is thus truly desirable.

9 A Flexible Scheduling Algorithm

9.1 Design guidelines

We characterize a request by its arrival time, its priority (to decide the order in which the requests should be served), the resources it asks for (core and elastic) and the execution time (in isolation, *i.e.*, when all required resources are granted to the application). Given an incoming workload, our goal is to optimize the sum of the turnaround times τ_i , that is:

$$\min \sum_i \tau_i \Rightarrow \min \sum_i (\text{queuing}_i + \text{execution}_i)$$

The actual execution time depends on the amount of resources assigned over time to the request. Now, recall that the scheduling problem can be broken into sorting and allocation phases. Sorting determines when a request is served, thus it has an impact on its queuing time. The allocation phase contributes both to queuing and actual execution times. Depending on allocation granularity [19], a request might need to wait for a number of resources to be available before occupying them, thus increasing – albeit indirectly – the queuing time. The execution time is directly related to the allocation algorithm and to the workload characteristics.

In this work we decouple request sorting from allocation:⁴ our scheduler maintains the request ordering, as imposed by an external component, and only focuses on resource allocation. Sorting can be simply based on arrival times (which amounts to implement a FIFO queuing discipline), or can use additional information, such as request size (thus implementing a variety of size-based disciplines).

Overall, we optimize request turnaround times through careful resource allocation, and design an algorithm that strives at allocating all available cluster resources, by serving the least number of

⁴This approach is similar to the one used in the SLURM scheduler [40], where the order of the pending jobs is given by an external, pluggable, component, and the scheduler processes the jobs following that order.

requests at a time. Intuitively, by “focusing” resources to few requests, we expect their execution times to be small. Consequently, queued requests also enjoy smaller wait times, because resources are freed more quickly.

9.2 Algorithm Details

Although we support preemptive scheduling policies, to simplify exposition, we first consider the case with no preemption: resources assigned to a request can only increase, and a new request can be placed, at most, at the head of the waiting line, depending on the sorting component. We stress that the output of our scheduling algorithm is a virtual assignment, i.e., the mechanism to physically allocate resources according to the computed assignment (core and elastic components for running applications) is separate from the scheduling logic, and considered as an implementation detail.

Our resource allocation procedure is called REBALANCE, and it is triggered by two events: request arrivals and departures – see Algorithm 1, ignoring highlighted lines. When a new request arrives (procedure ONREQUESTARRIVAL), the resource assignment is done only if such a request is placed at the head of the waiting line and there are unused resources that are sufficient for running its core components. When a request is completed (procedure ONREQUESTDEPARTURE), the released resources are always reassigned.

The scheduler maintains two ordered sets: the requests waiting to be served (\mathcal{L}), and the requests in service (\mathcal{S}). Each request req needs $req.C$ core components and $req.E$ elastic components; depending on the allocation, request req is granted $0 \leq req.G \leq req.E$ elastic components. The core of the procedure REBALANCE (lines 27-30) operates as follows: each request req in the serving set \mathcal{S} has always **at least** $req.C$ resources assigned. Excess resources are assigned to the requests in \mathcal{S} following the request order. The scheduler assigns as many elastic components as possible to the first request, then to the second, and so on, in cascade.

Following the design guidelines, the set \mathcal{S} should only contain the requests that are strictly necessary to use all the available resources. This is accomplished by the first part of the procedure REBALANCE (lines 17-22): a request is added to \mathcal{S} if the current requests in \mathcal{S} are not able to saturate the total resources (*total*, line 17). Note that we add a request to \mathcal{S} only if there is room to allocate all of its core components.

9.3 Preemptive policies

We now consider preemptive policies: request arrivals can trigger (partial) preemption of running requests, e.g. if new requests have higher priority than that of the last request in service. In this case, the tuple describing a request also stores its priority, $req.P$. It is important to note that, in this work, the preemption mechanism only operates on elastic components of running applications, whereas core components (that are vital for an application) cannot be preempted.

The highlighted lines in Algorithm 1 show the modifications to the procedures ONREQUESTARRIVAL and ONREQUESTDEPARTURE to support preemption. When a new request arrives, if its priority is higher than the requests in service, we check if its core components can be allocated using the resources occupied by the elastic components of currently running requests. If so, we insert the request into the set \mathcal{S} and call REBALANCE. Otherwise, we insert the request into an auxiliary waiting line \mathcal{W} , which is given priority when resources become available. Indeed, procedure ONREQUESTDEPARTURE indicates that we first consider the waiting requests in \mathcal{W} , and we add to the set \mathcal{S} as many of them as possible, considering solely the core components. In other words, requests in \mathcal{W} have higher priority than those in \mathcal{L} . Finally, the call of REBALANCE assigns the remaining resources to the elastic components of high priority requests.

10 Numerical evaluation

10.1 Methodology

We evaluate our algorithm using an event-based, trace-driven discrete simulator developed to study the scheduler Omega [19], which we extended in order to make it work with applications, instead of low-level jobs and to use the concept of component classes. Our scheduler implementation sup-

Algorithm 1: Resource assignment procedures

```

1 procedure ONREQUESTARRIVAL(req)
2   if req.P > S.tail.P then
3     if req.C  $\leq \sum_{j \in \mathcal{S}} \text{req}_j.E$  then
4       INSERT(req, S)
5       REBALANCE()
6     else
7       INSERT(req, W)
8   else
9     INSERT(req, L)
10    if req == L.head and req.C  $\leq \text{avail}$  then
11      REBALANCE()
12
13 procedure ONREQUESTDEPARTURE()
14   while W.head.C +  $\sum_{j \in \mathcal{S}} \text{req}_j.C < \text{total}$  and (W not  $\emptyset$ ) do
15     INSERT(POP(W), S)
16   REBALANCE()
17
18 procedure REBALANCE()
19   while  $\sum_{j \in \mathcal{S}} (\text{req}_j.C + \text{req}_j.E) < \text{total}$  and (L not  $\emptyset$ ) do
20     req  $\leftarrow$  L.head
21     if req.C +  $\sum_{j \in \mathcal{S}} \text{req}_j.C < \text{total}$  then
22       INSERT(POP(L), S)
23     else
24       break
25
26     avail  $\leftarrow$  total -  $\sum_{j \in \mathcal{S}} \text{req}_j.C$ 
27     forall the req  $\in \mathcal{S}$  do
28       req.G  $\leftarrow$  0
29     req  $\leftarrow$  S.head
30     while avail > 0 and (req not NULL) do
31       req.G  $\leftarrow \min(\text{req.E}, \text{avail})$ 
32       avail  $\leftarrow (\text{avail} - \text{req.G})$ 
33       req  $\leftarrow \text{req.next}$ 

```

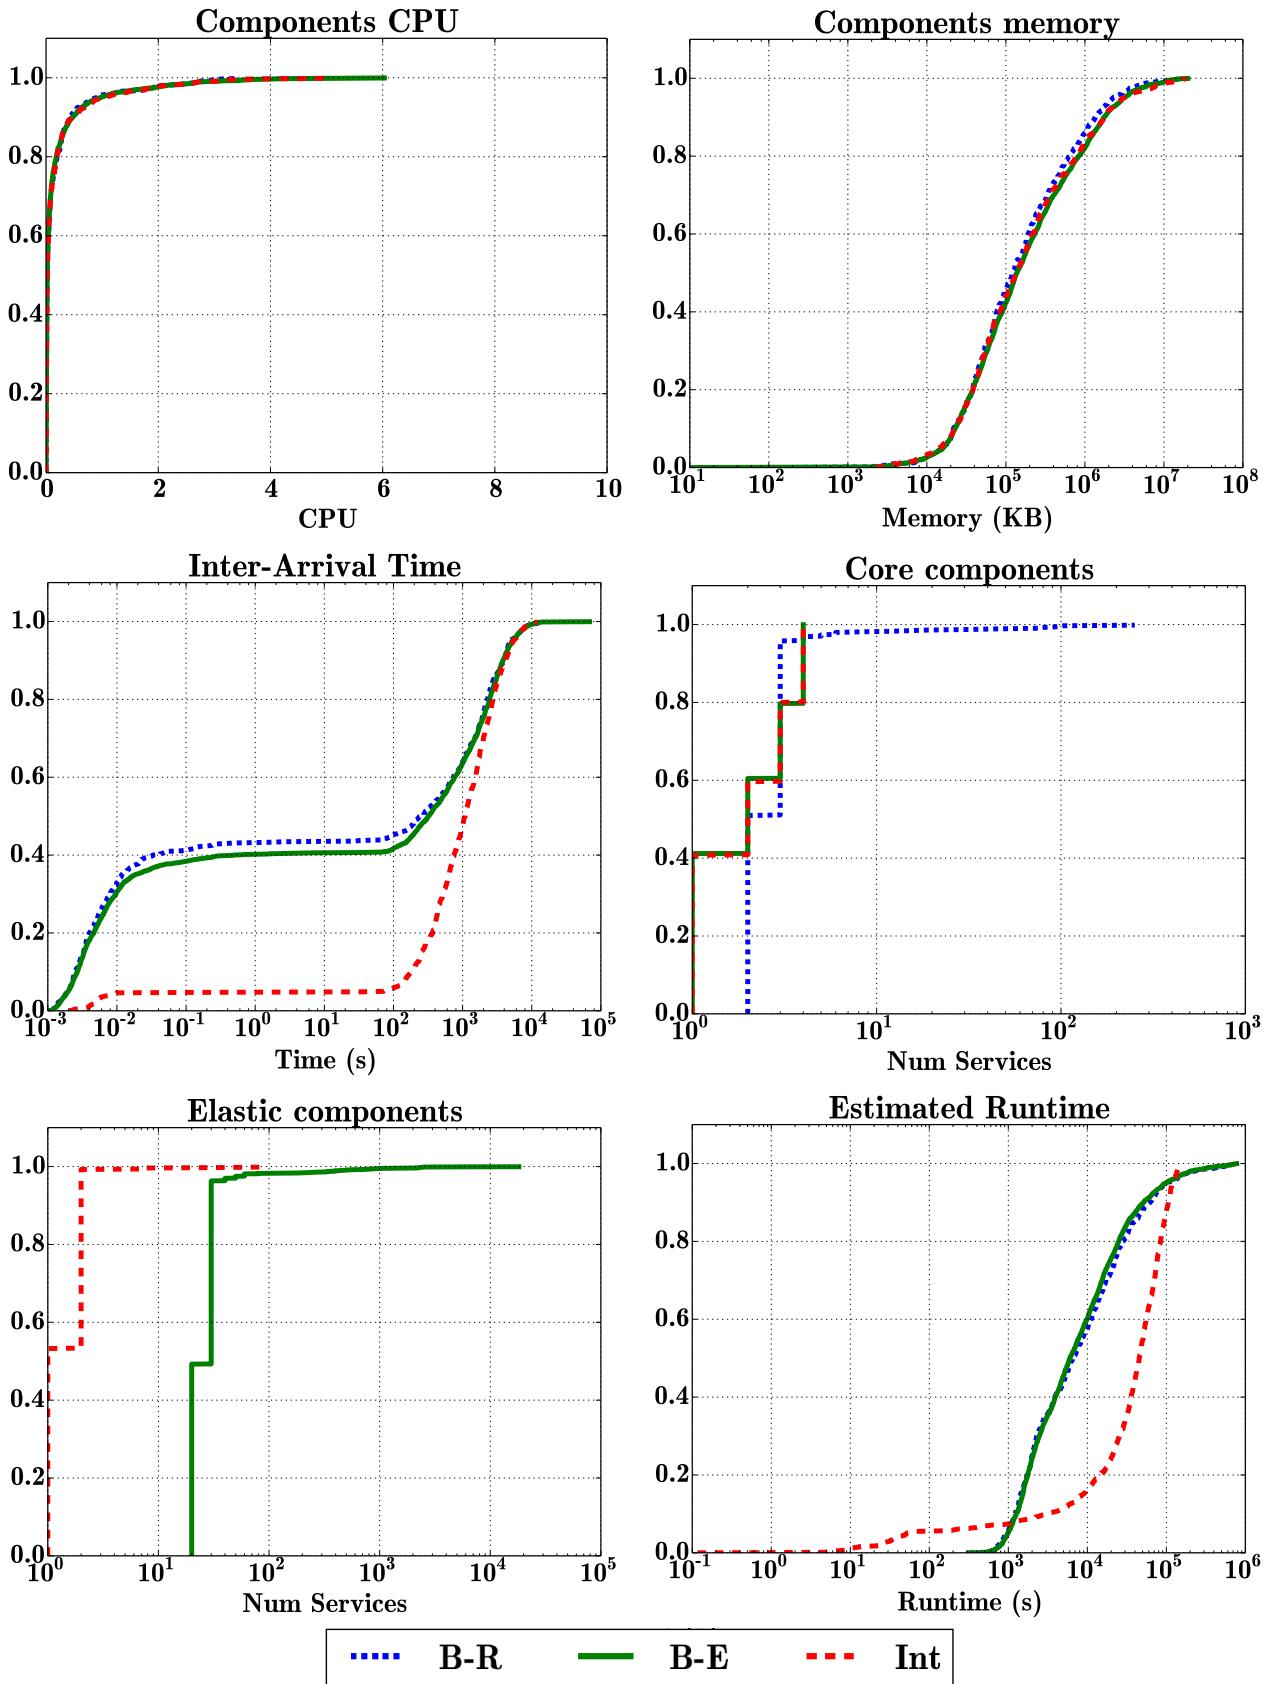


Figure 3: Workload Definition.

ports a variety of policies⁵: we present results for the FIFO and the shortest job first (SJF) policies, which further optimizes system responsiveness. Our implementation first obtains a “virtual assignment” with Algorithm 1, then fulfills it by allocating resources accordingly, which happens instantaneously. Additionally, we have implemented a baseline, consisting of a rigid scheduler that does not distinguish component classes, which is representative of current cluster management systems. In our simulations, we consider two-dimensional resources, including definitions of CPU and RAM requirements. We would like to stress that the “virtual assignment” can take into consideration other constraints as well (e.g., GPU).

Our scheduler currently accepts **application workloads** of two kinds. The first is **batch applications**, that take from a few seconds to a few days to complete: these are delay-tolerant applications, with a very simple life-cycle. Core components must first start to produce useful work, by executing user-defined jobs that are “passed” to the application; eventually, elastic components can contribute to the application progress. Once the user programs are concluded, the application finishes, releasing resources occupied by its frameworks and components. The second is **interactive applications**, which involve a “human in the loop”: these are latency-sensitive applications, with a life-cycle triggered by human activity. In this case, core components must start as soon as possible, to allow user interaction with the application (e.g., a Notebook).

For our performance evaluation, we use publicly available traces [35, 34], and generate a workload by sampling the empirical distributions we compute from such traces. First, we focus on batch applications alone, and simulate both rigid (e.g. TensorFlow) and elastic (e.g. Spark) variants: the label **B-R** represents rigid applications with only core components; the label **B-E** stands for elastic applications, with both core and elastic components. Then, we evaluate the benefit of preemption by complementing the above workload with (simulated) interactive applications.

Figure 3 shows the definition of the workload: in particular we can see the CDFs for requested CPU and memory, then the inter-arrival time and estimated run time and finally, the number of core and elastic components. Batch applications are assigned a number ranging from a few to tens of thousands of components. Instead, interactive applications are smaller, and use up to hundreds of elastic components. The resource requirements of application components follow that of the input traces, ranging from few MB to a few dozens GB of memory, and up to 6 cores. Application runtime is generated according to the input traces, and range from a few dozen seconds to several weeks (of simulated time). Application inter-arrival times are drawn from the empirical distributions of the input traces, and exhibit a bi-modal distribution with fast-paced bursts, as well as longer intervals between application submissions. In summary, our workload consists of 80,000 applications, with 80% batch and 20% interactive applications. Batch applications include 80% elastic and 20% rigid variants.

We simulate a cluster consisting of 100 machines, each with 32 cores and 128GB of memory. All results shown here include 10 simulation runs, for a total of roughly 3 months of simulation time for each run.

Finally, the metrics we use to analyze the results include: application **turnaround**, which allow to reason about the scheduling objective function, and **queuing time**, which is an important factor contributing to the turnaround time. Additionally, we measure the **queue sizes** that hold pending and running applications, and **resource allocation**, measured as the percentage of CPU and memory the scheduler allocates to each application.

10.2 Comparison with the baseline

We now perform a comparative analysis between our flexible scheduler and the baseline, and start by disabling preemption: we omit interactive applications from the workload. Figure 4 (left) illustrates the most important percentiles (in a box-plot) of the distribution of turnaround times, where we compare two policies (namely, FIFO and SJF) used by both the baseline and our scheduler. The benefits of our approach are noticeable, irrespectively of the scheduling discipline: the median turnaround is

⁵We support size-based disciplines in the family of SMART policies [41]. Here we assume application size information to be known a-priori.

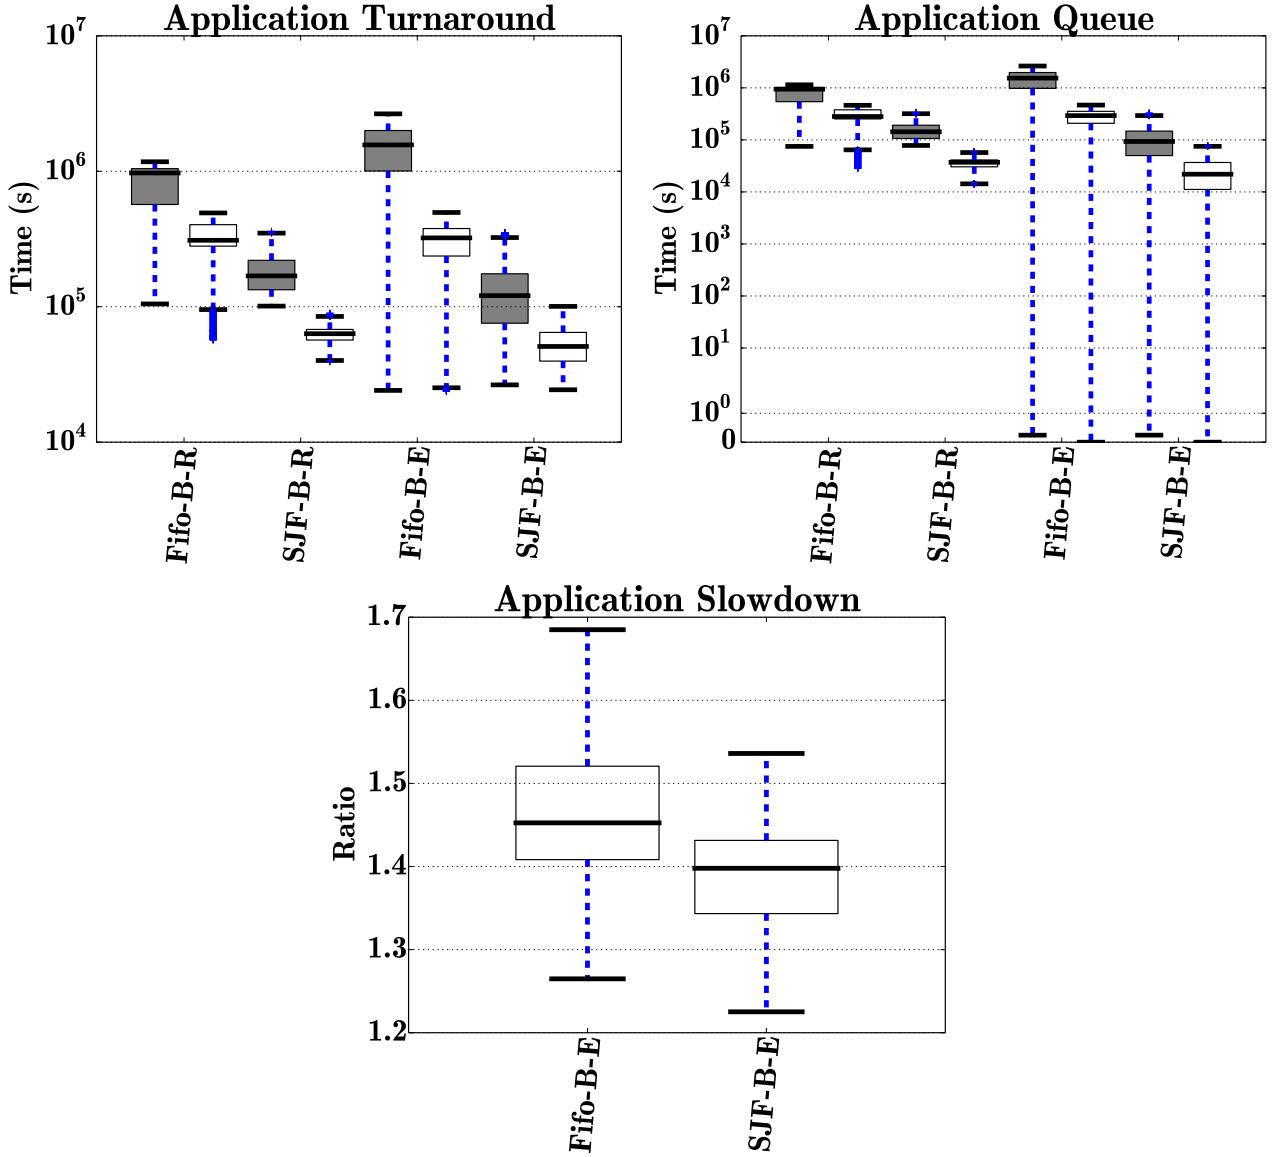


Figure 4: Comparison of turnaround and queue time distributions, and application slowdown distributions for FIFO and SJF policies. White boxes (right box of every pair) corresponds to our flexible scheduler, gray boxes correspond to the baseline. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

halved when compared to the baseline, indicating superior system responsiveness. Additionally, we observe the benefits of a size-based policy in further decreasing turnaround times. We note that our approach is beneficial for both rigid and elastic batch applications: Figure 4 (center) shows a box-plot of application queuing times, which contribute to their turnaround. With our approach, both kinds of applications spend less time waiting in a queue to be served. By differentiating classes of components, applications can execute as soon as enough resources to produce work are available. Finally, Figure 4 (right) focuses on application runtime: we report the **slowdown** computed as the ratio between the nominal application runtime (i.e., the time required for an application to complete in an empty system, with all application components allocated their requested resources) and the effective application runtime obtained with the simulation. Values above one indicate that applications run slower in a system absorbing a given workload when compared to applications running in an empty system. Overall, these results show that our scheduling approach does not impose a high toll on

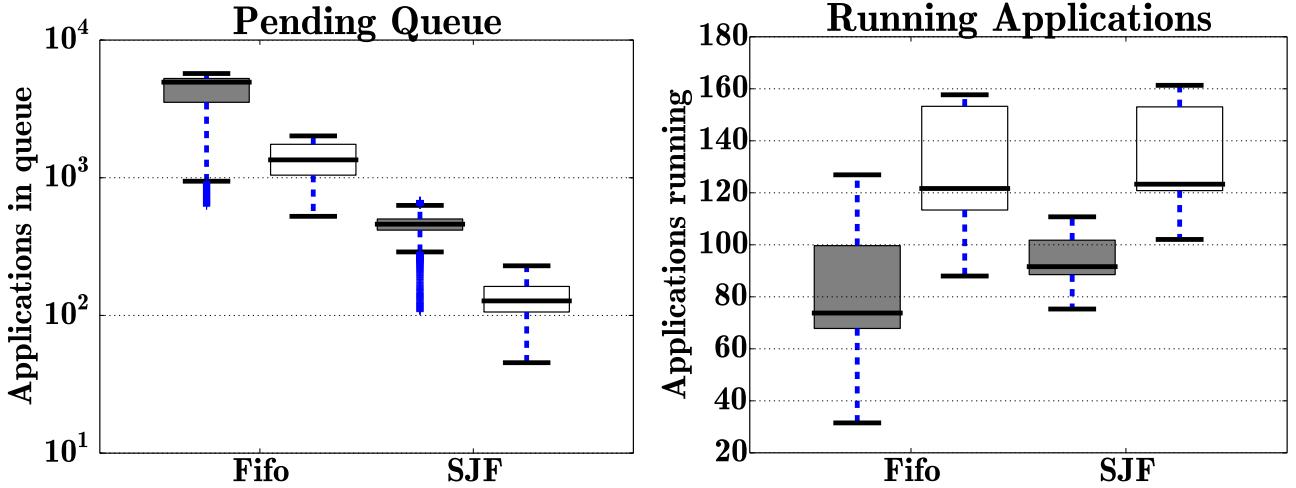


Figure 5: Comparison of queues size for FIFO and SJF between our flexible scheduler and the baseline. The white boxes (right box of every group) correspond to our flexible algorithm, gray boxes to the baseline.

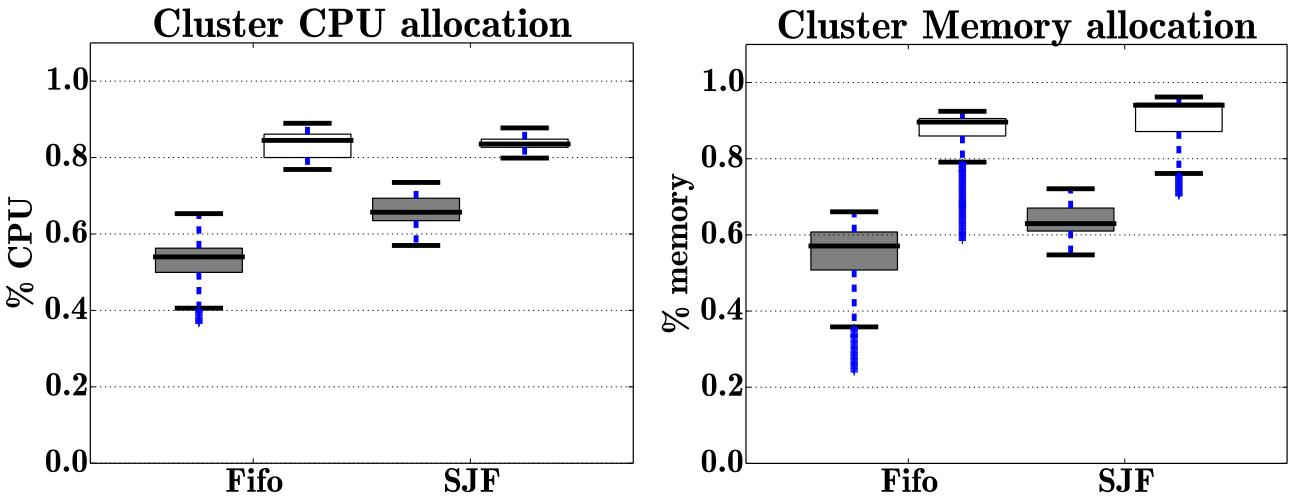


Figure 6: Comparison of resource allocation distributions for FIFO and SJF policies, between our flexible scheduler and the baseline. White boxes (right box of every pair) correspond to our approach, dashed boxes to the baseline.

application runtime, while globally contributing to improved turnaround times.

Next, we support the general results discussed above with additional details. Figure 5 shows the box-plots of the distribution of queue sizes, for both the pending and the running queues. Our approach induces a smaller number of applications waiting to be served, as well as a larger number of applications running in the system, compared to the baseline and across different policies. Indeed, our flexible scheduler achieves a better packing of applications, which means they can start sooner. Additionally, the benefits of a size-based discipline are clear: the number of applications waiting is almost one order of magnitude smaller compared to a FIFO policy, while the number of running applications is similar.

Figure 6 shows metrics from the cluster perspective: our approach (for both disciplines) induces a far better resource allocation compared to the baseline, achieving more than 20% gains in both CPU

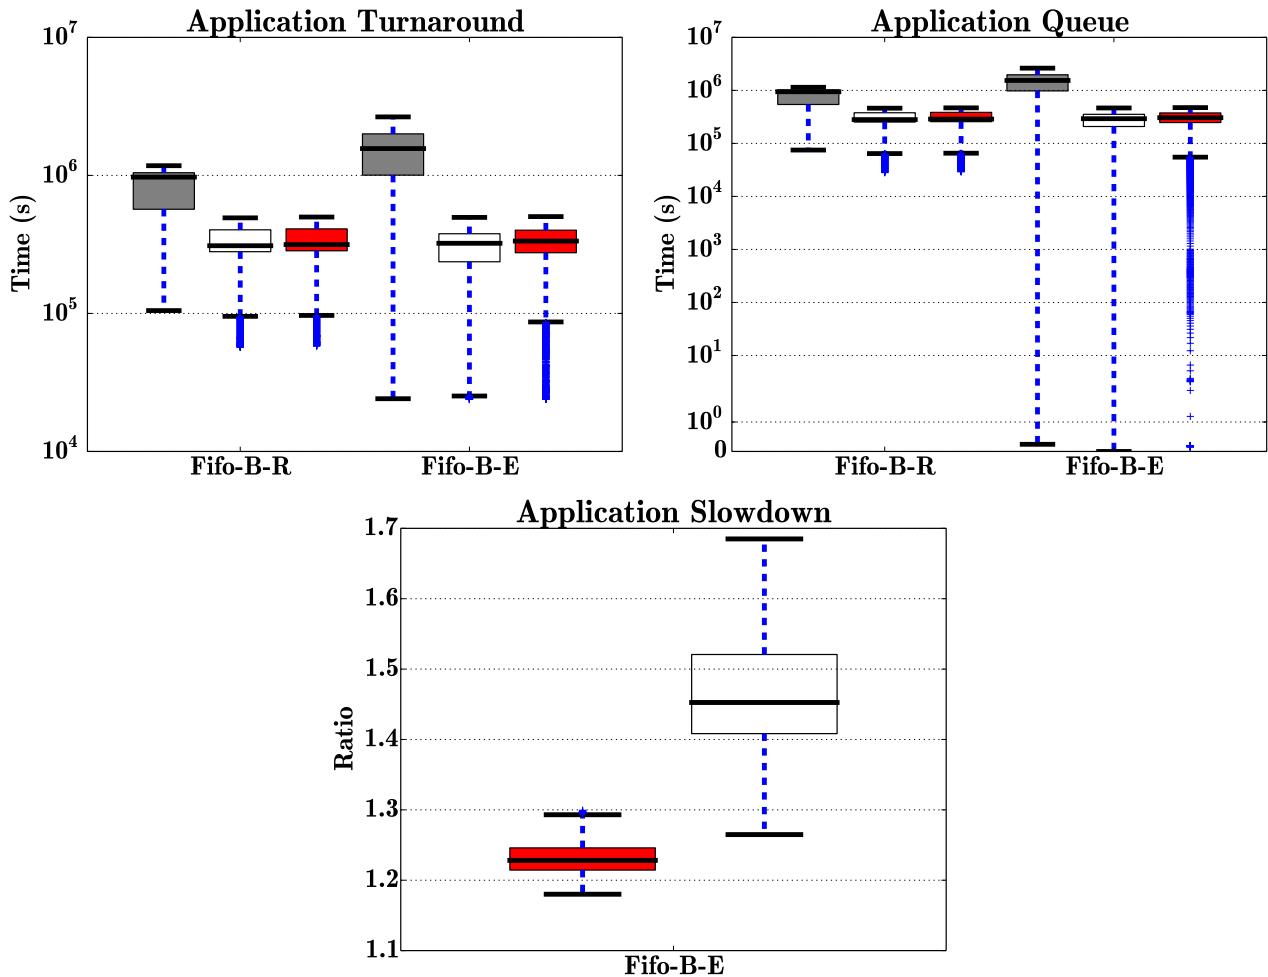


Figure 7: Comparison using FIFO between the rigid, the elastic and our system. For every group, the boxes in the middle correspond to our system, the boxes on the left correspond to the rigid system while the boxes on the right are for the elastic system. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

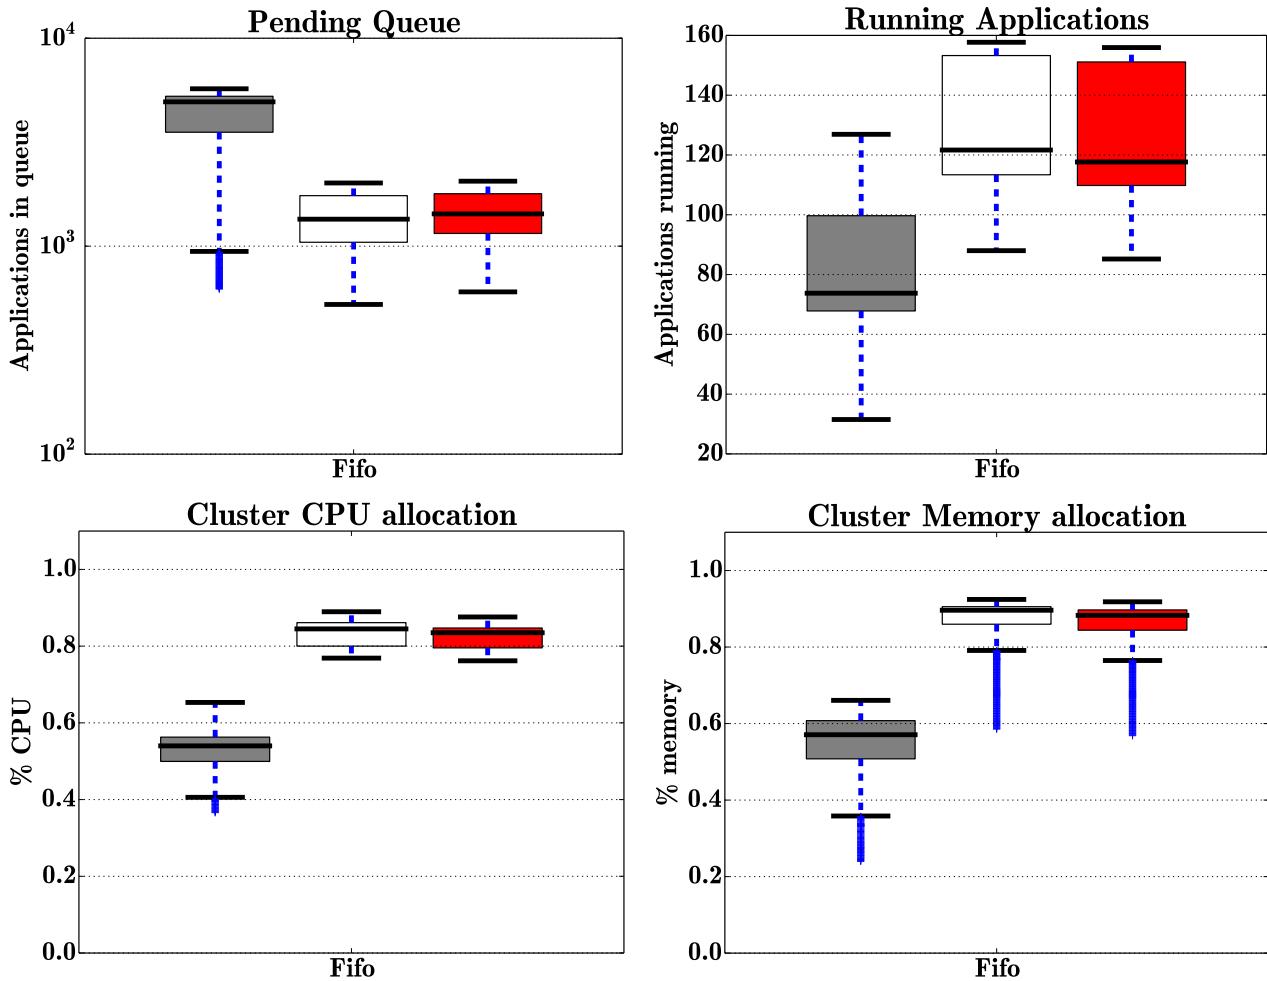


Figure 8: Comparison using FIFO between the rigid, the elastic and our system. For every group, the boxes in the middle correspond to our system, the boxes on the left correspond to the rigid system while the boxes on the right are for the elastic system. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

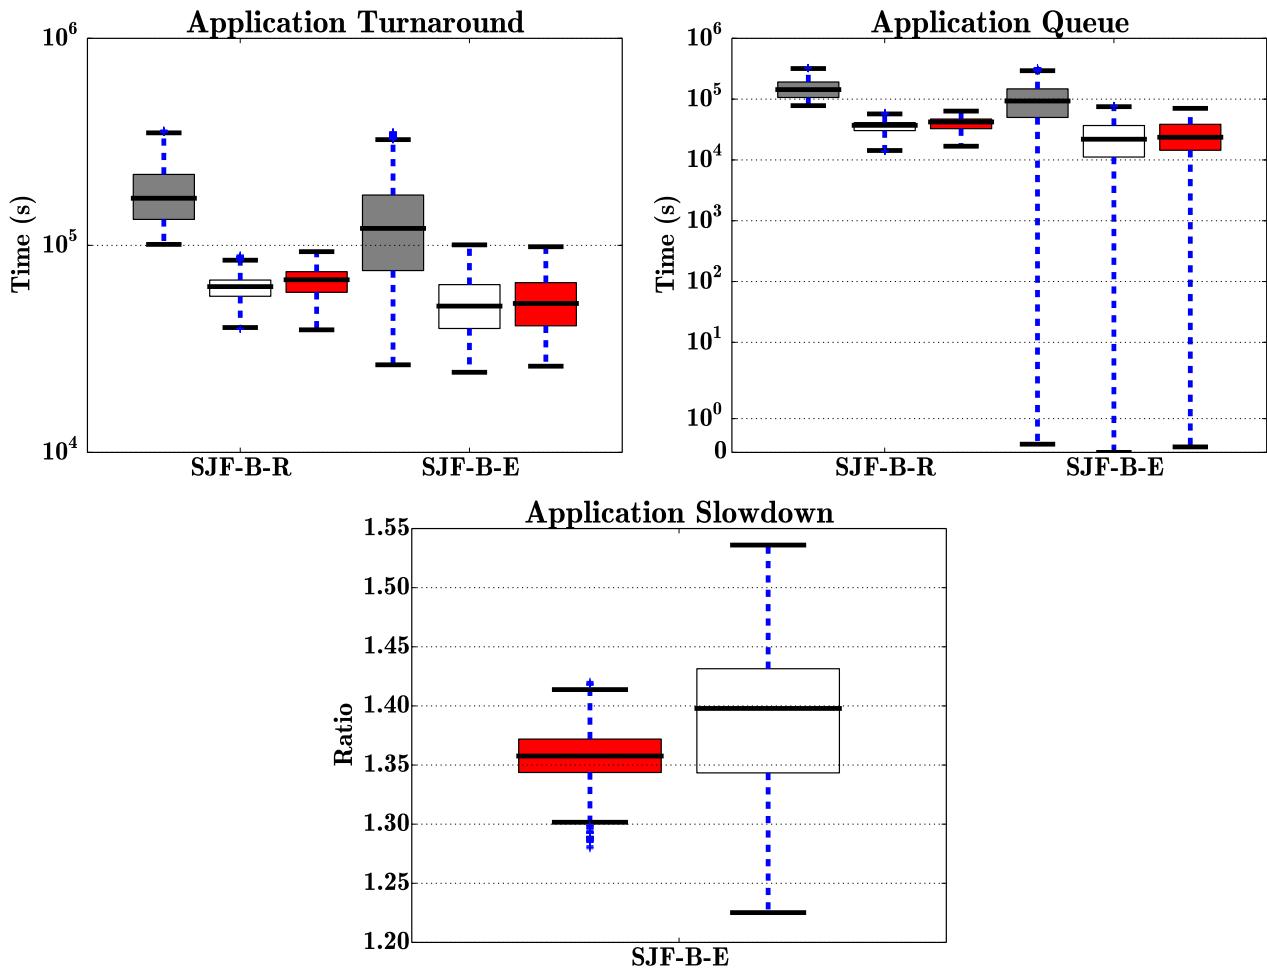


Figure 9: Comparison using SJF between the rigid, the elastic and our system. For every group, the boxes in the middle correspond to our system, the boxes on the left correspond to the rigid system while the boxes on the right are for the elastic system. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

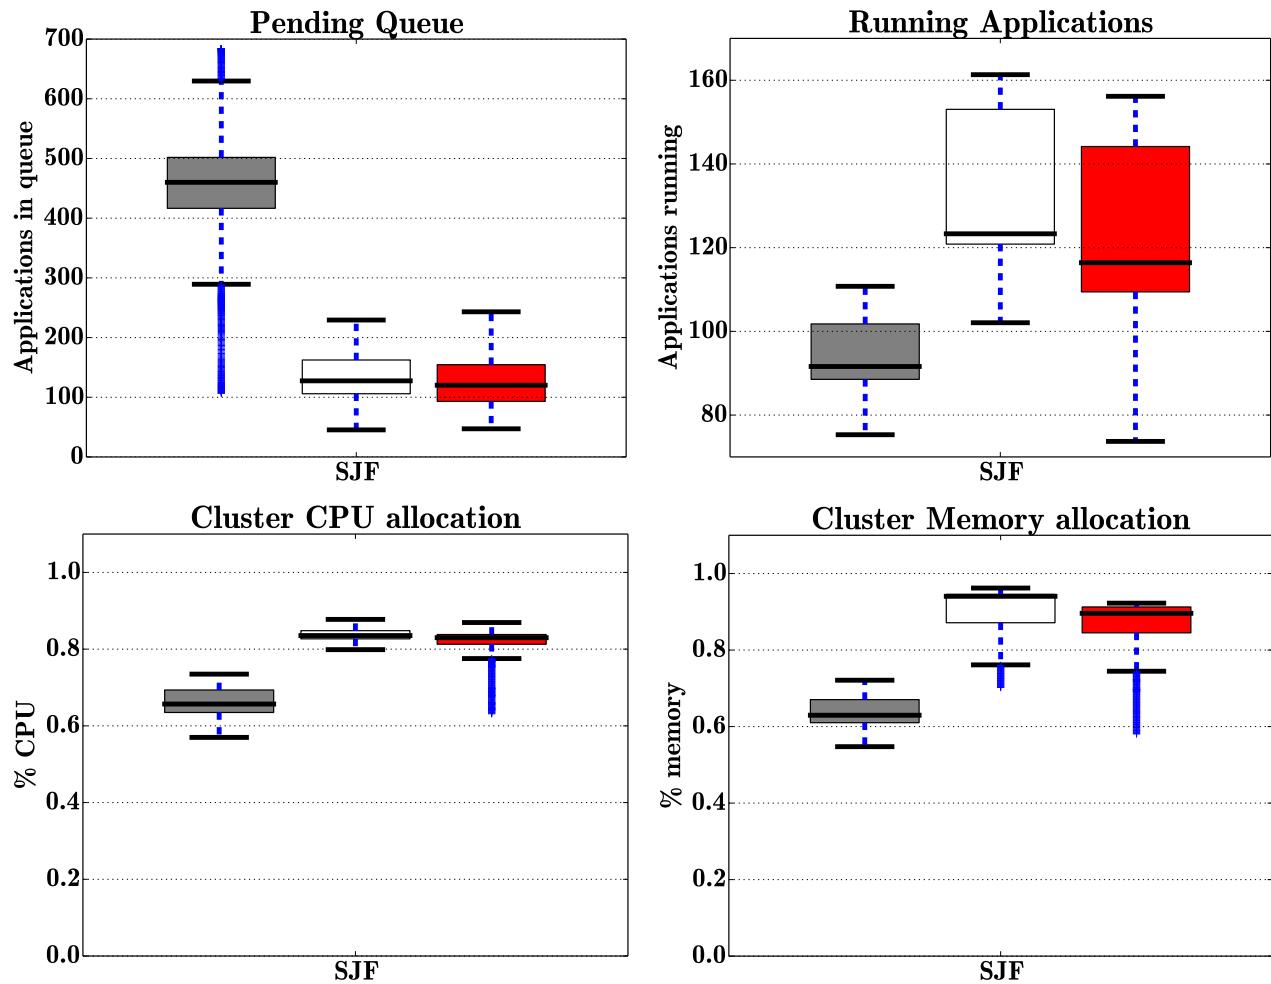


Figure 10: Comparison using SJF between the rigid, the elastic and our system. For every group, the boxes in the middle correspond to our system, the boxes on the left correspond to the rigid system while the boxes on the right are for the elastic system. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

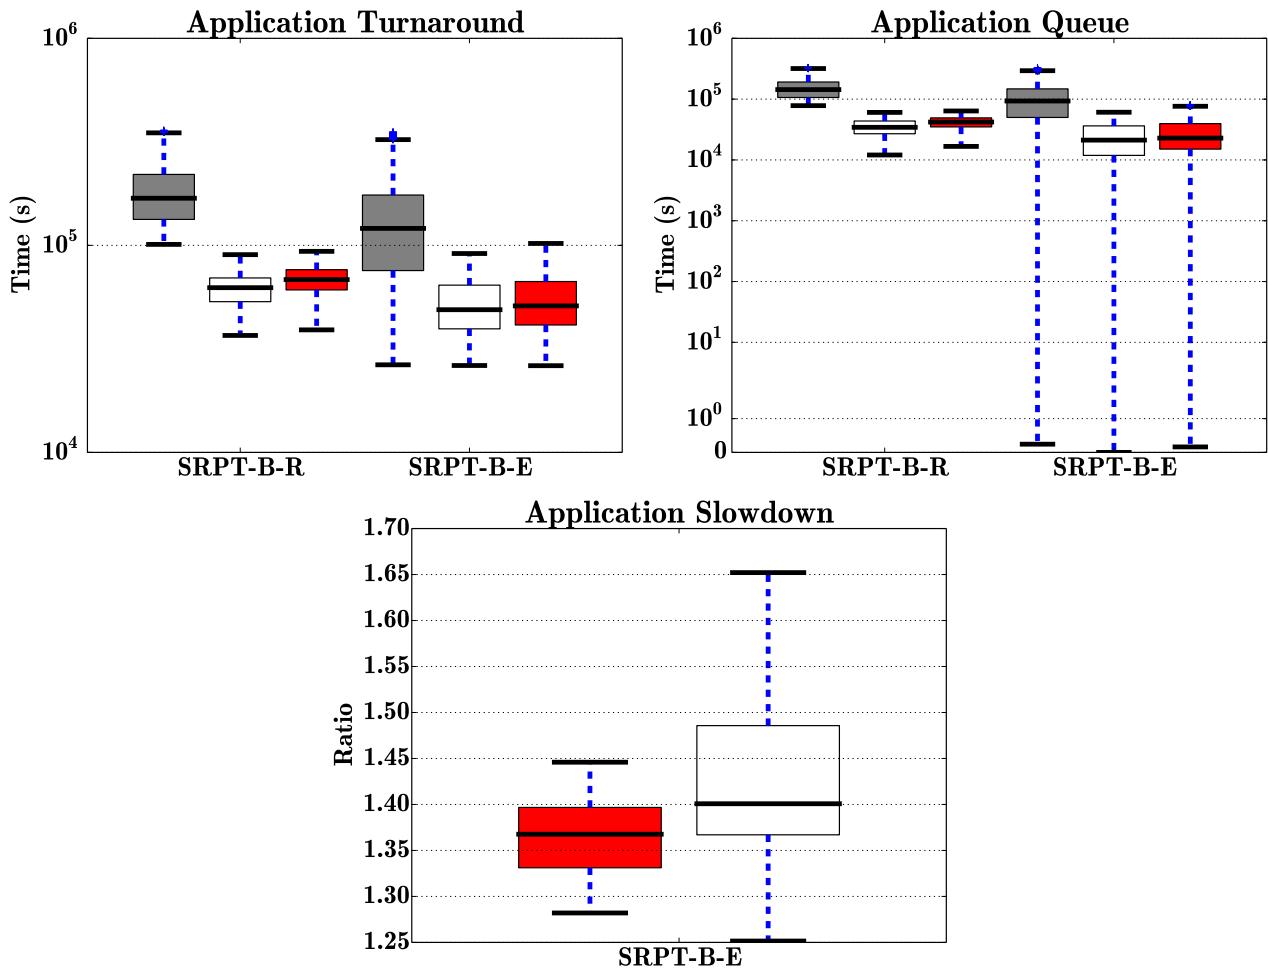


Figure 11: Comparison using SRPT between the rigid, the elastic and our system. For every group, the boxes in the middle correspond to our system, the boxes on the left correspond to the rigid system while the boxes on the right are for the elastic system. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

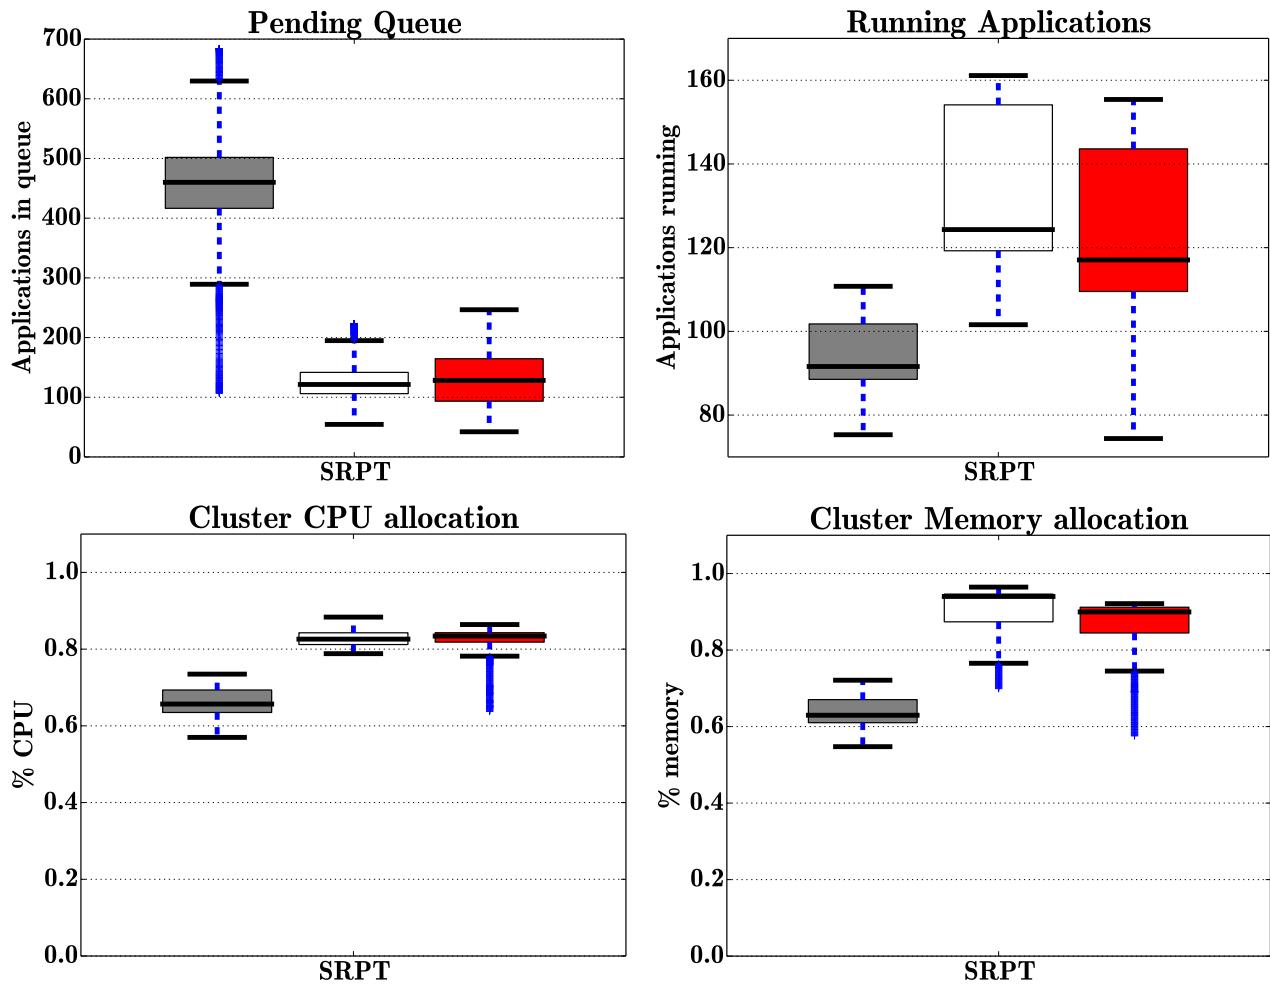


Figure 12: Comparison using SRPT between the rigid, the elastic and our system. For every group, the boxes in the middle correspond to our system, the boxes on the left correspond to the rigid system while the boxes on the right are for the elastic system. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

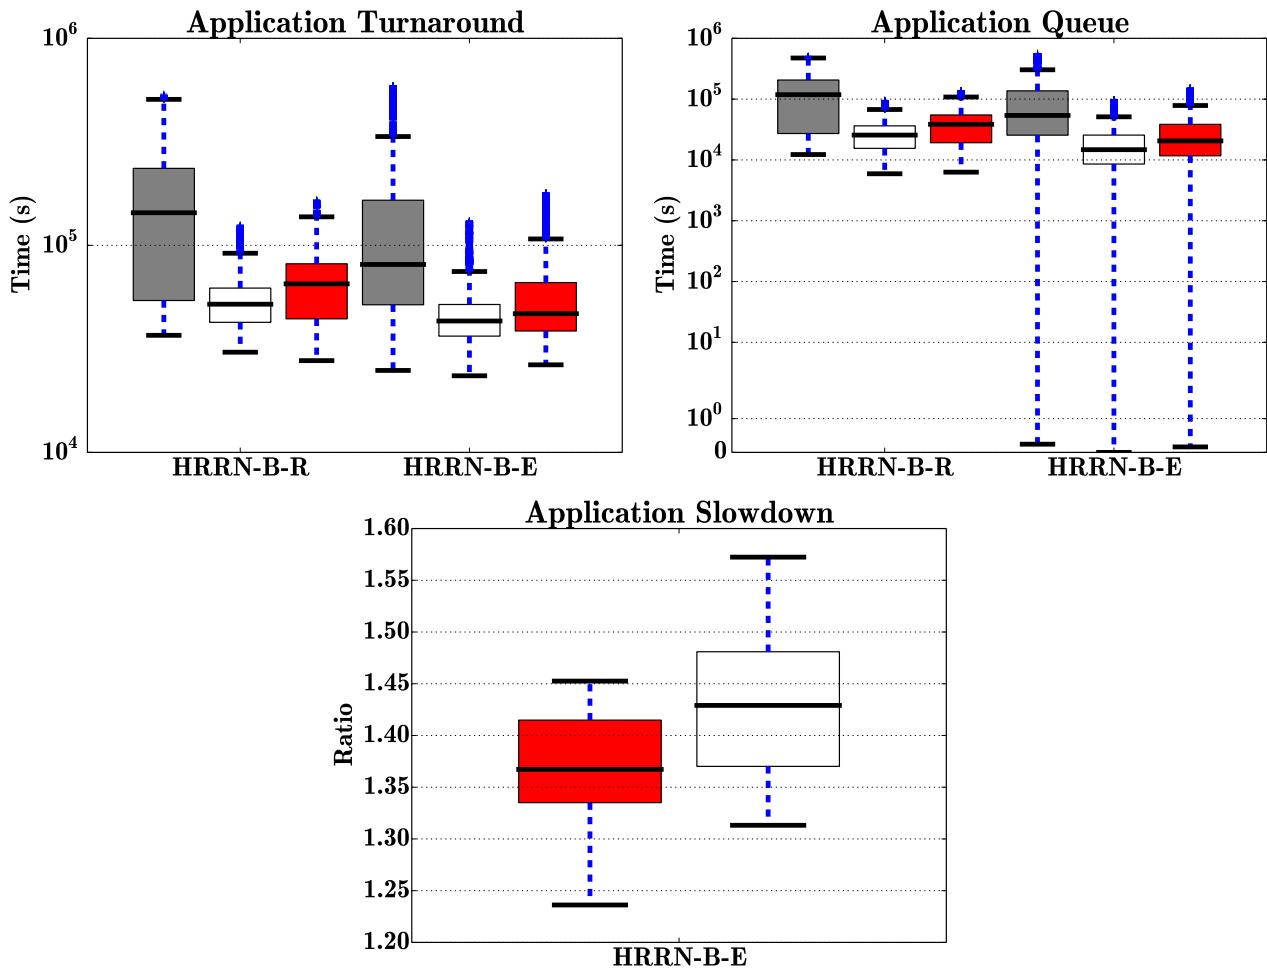


Figure 13: Comparison using HRRN between the rigid, the elastic and our system. For every group, the boxes in the middle correspond to our system, the boxes on the left correspond to the rigid system while the boxes on the right are for the elastic system. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

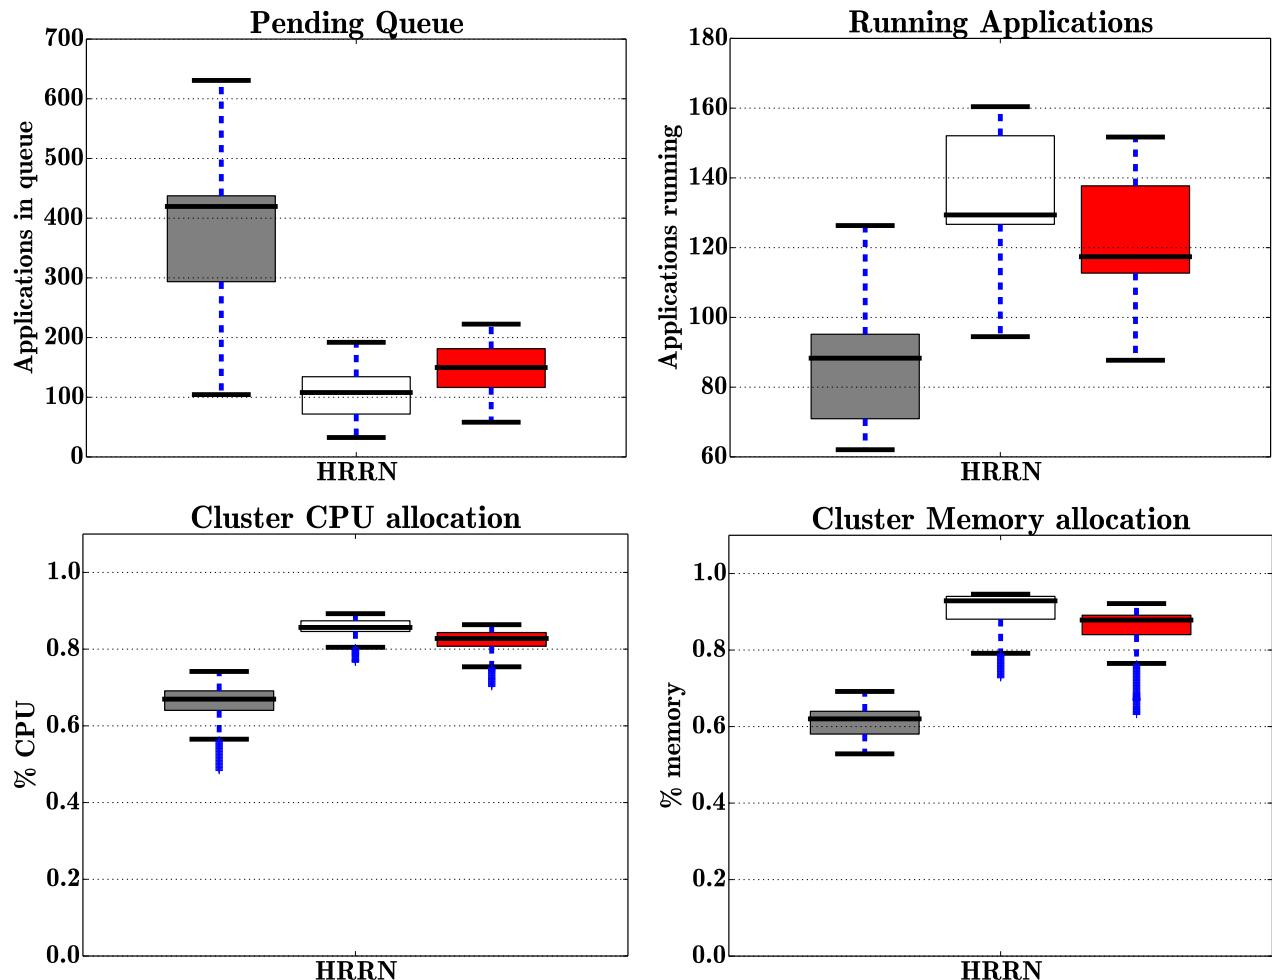


Figure 14: Comparison using HRRN between the rigid, the elastic and our system. For every group, the boxes in the middle correspond to our system, the boxes on the left correspond to the rigid system while the boxes on the right are for the elastic system. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

Name	Definition
SJF-2D	$runTime * \#RequestedServices$
SRPT-2D1	$remainingRunTime * \#RequestedServices$
SRPT-2D2	$remainingRunTime * \#ServicesYetToBeScheduled$
HRRN-2D	$(1 + \frac{waitTime}{runTime}) * \#RequestedServices$
SJF-3D	$runTime * \sum_{i=1}^{services} CPU_i * RAM_i$
SRPT-3D1	$remainingRunTime * \sum_{i=1}^{services} CPU_i * RAM_i$
SRPT-3D2	$remainingRunTime * \sum_{i=1}^{servicesToSchedule} CPU_i * RAM_i$
HRRN-3D	$(1 + \frac{waitTime}{runTime}) * \sum_{i=1}^{services} CPU_i * RAM_i$

Table 10.3a: Definition of size used in the evaluation

and RAM allocation.⁶

Finally, we are going to show more results, that compare a rigid, a malleable and our flexible scheduler with different policies. It is worth mentioning that currently no solution support a malleable scheduler as we presented it in Section 8.2.

In single-server systems, the policies that are optimal for minimizing the average turn around time are called SMART [41]; they prioritize short application over longer one. Two example of SMART policies are SJF (Shorted-Job First) and SRPT (Shortest-Remaining-Processing-Time). However, even if SRPT is considered optimal, it is rarely used because it leads to starvation of long running application. For this reason policies like HRRN (Higher-Response-Ration Next) can be used; they calculate a virtual size that is updated the longer the application reside in queue. In this evaluation we are going to use four different policies: FIFO, SJF, SRPT and HRRN.

Figures 7 to 14 confirm that our solution performs far better than a rigid scheduler and slightly better than a malleable regarding of the policy used to sort the queue.

In conclusion using our flexible scheduler can greatly reduce the turn around time while improving resources allocation regardless of the policy that is used to sort the applications in the pending queue.

10.3 Comparison between different definitions of size

In this section we will focus on size-based policies and in particular on their definition of size; they were meant to be used in a single-server system or an unidimensional world, where there is only one type of resource. In a data center, however, there is a cluster of machines, and each one has at least two type of resources: CPU and Memory. For this reason we compare different definitions of size that will move from an unidimensional to a multidimensional world; similarly to [42], we want to “improve” the definition of size for every single policies that we introduced in the previous section by, progressively, adding more information to them. Table 10.3a shows the definition of size that we use in this section of the evaluation.

SJF-2D	SRPT-2D1	SRPT-2D2	HRRN-2D	SJF-3D	SRPT-3D1	SRPT-3D2	HRRN-3D
38,340.68	38,767.29	39,425.82	89,596.90	33,451.60	33,418.51	33,413.45	229,691.64

Table 10.3b: Turn around time (in second) between different definitions of size in our flexible scheduler.

Table 10.3b shows the results obtained by running the previous workload with these definitions of size. We see that there is an improvement when we add more information to the definition of

⁶Allocation is different from utilization: the simulator does not account for real application execution, so we cannot report utilization figures.

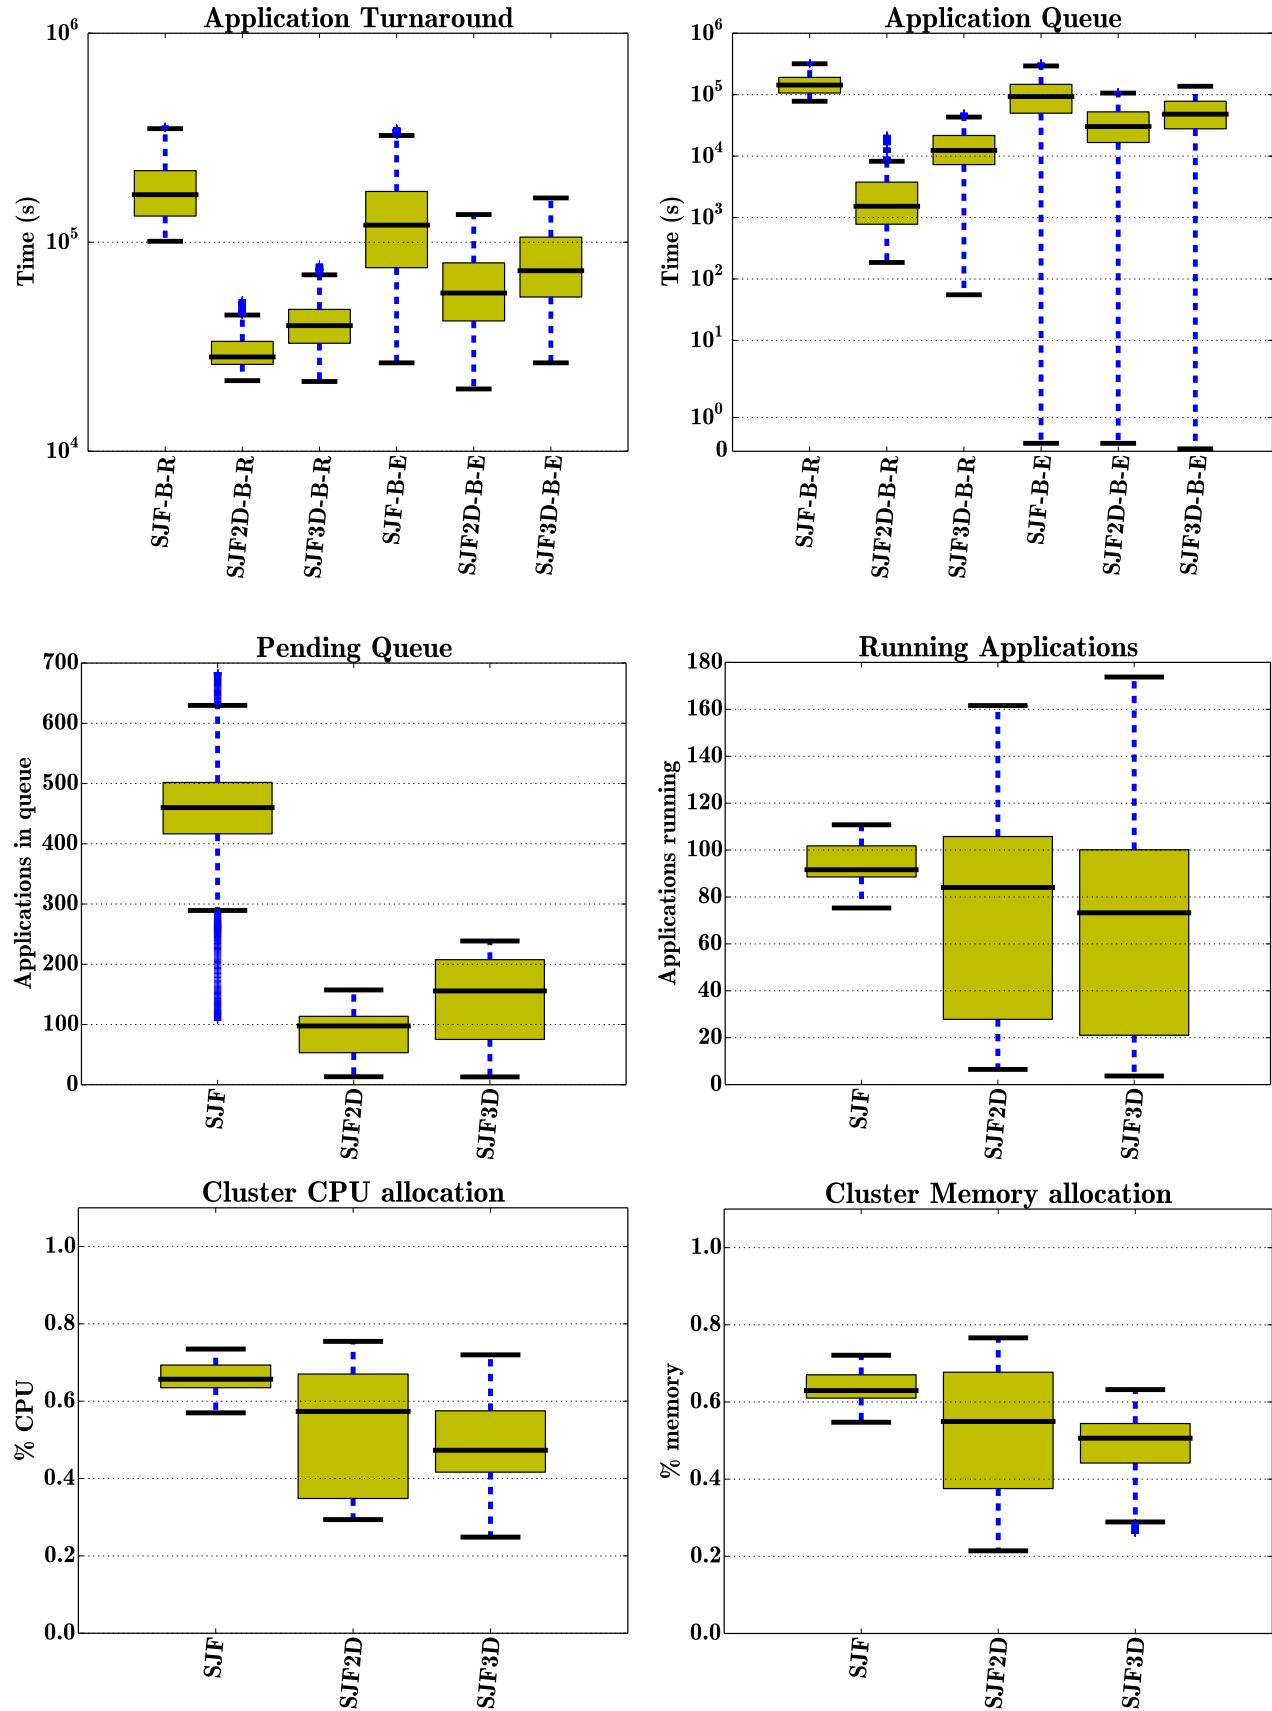


Figure 15: Comparison using different definition of size for SJF with a rigid scheduler. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

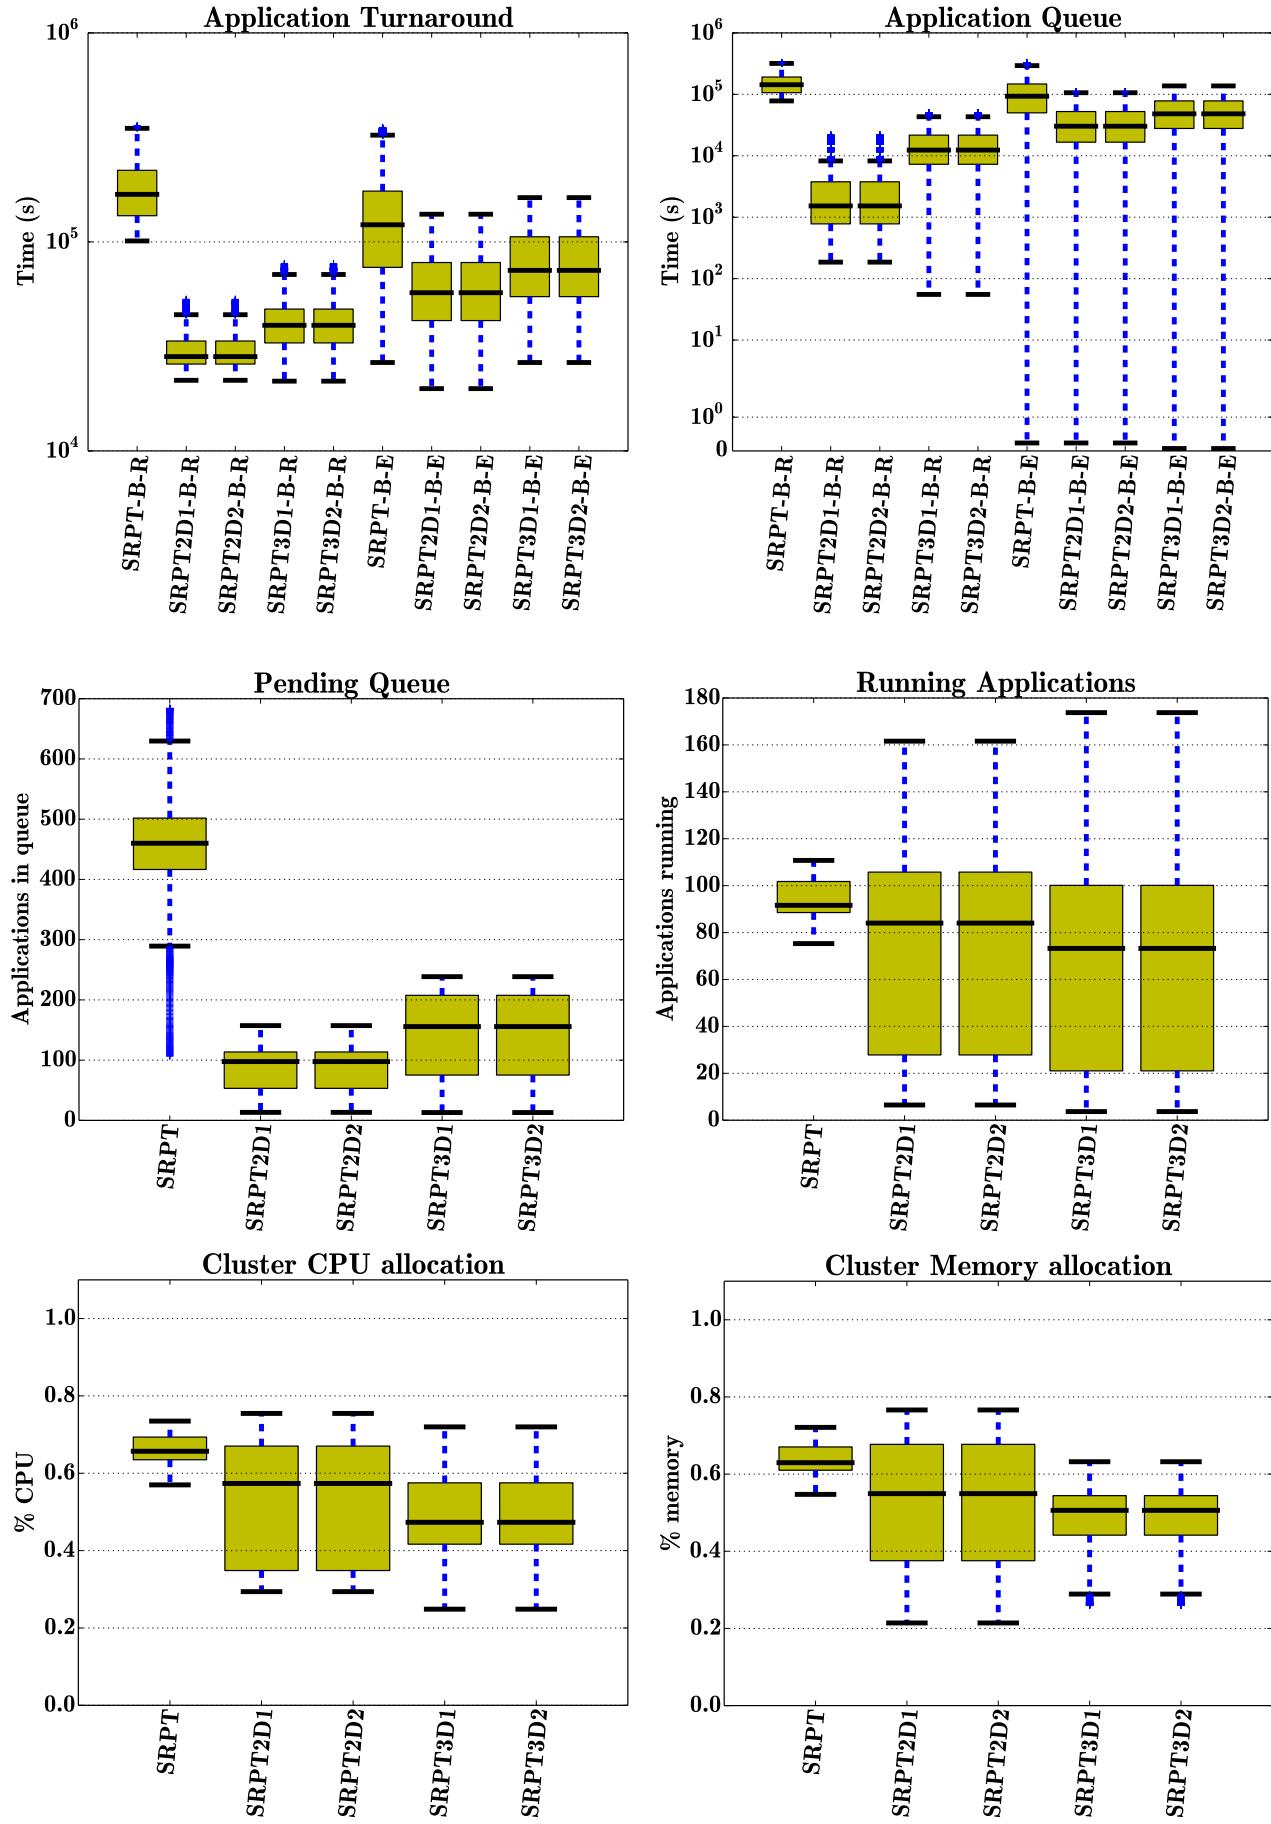


Figure 16: Comparison using different definition of size for SRPT with a rigid scheduler. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

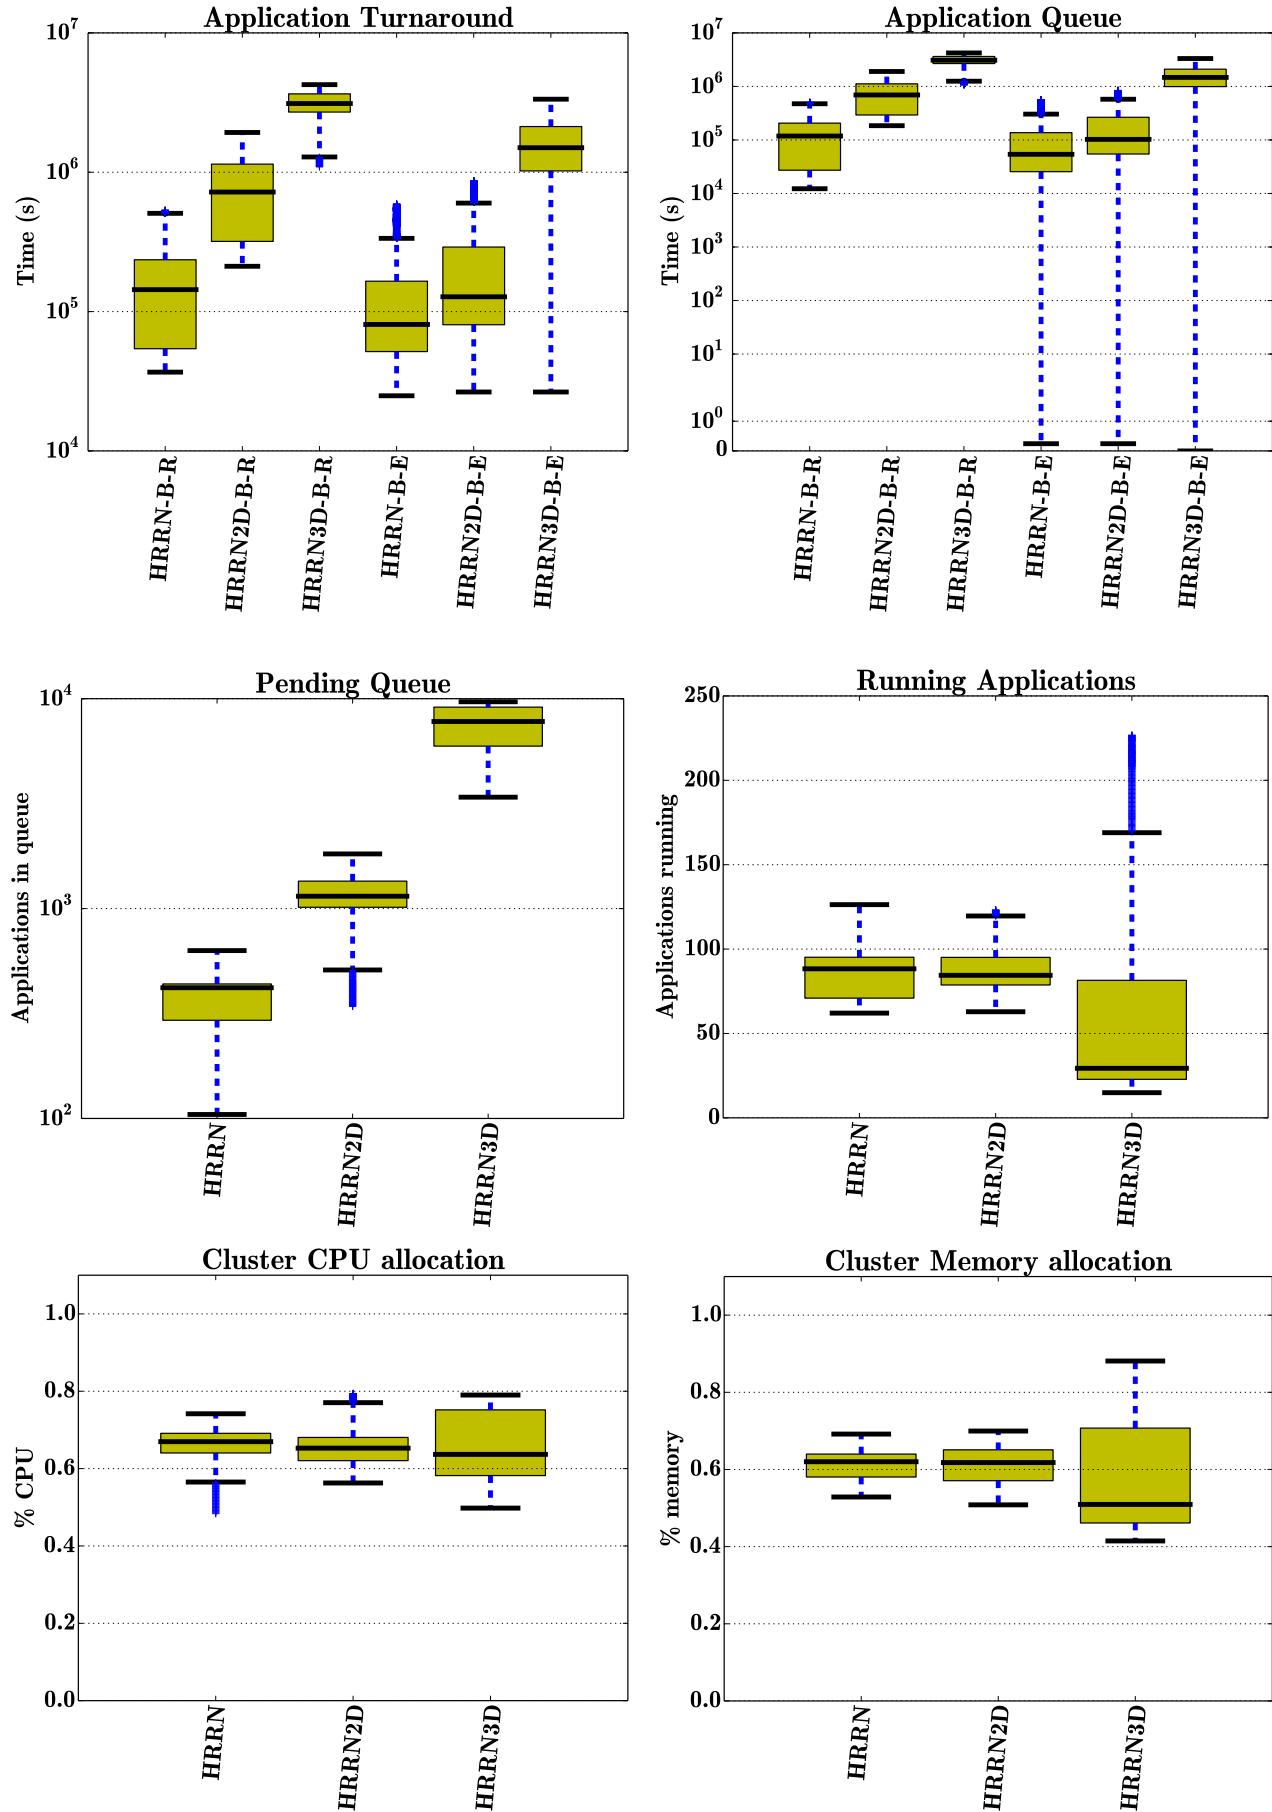


Figure 17: Comparison using different definition of size for HRRN with a rigid scheduler. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

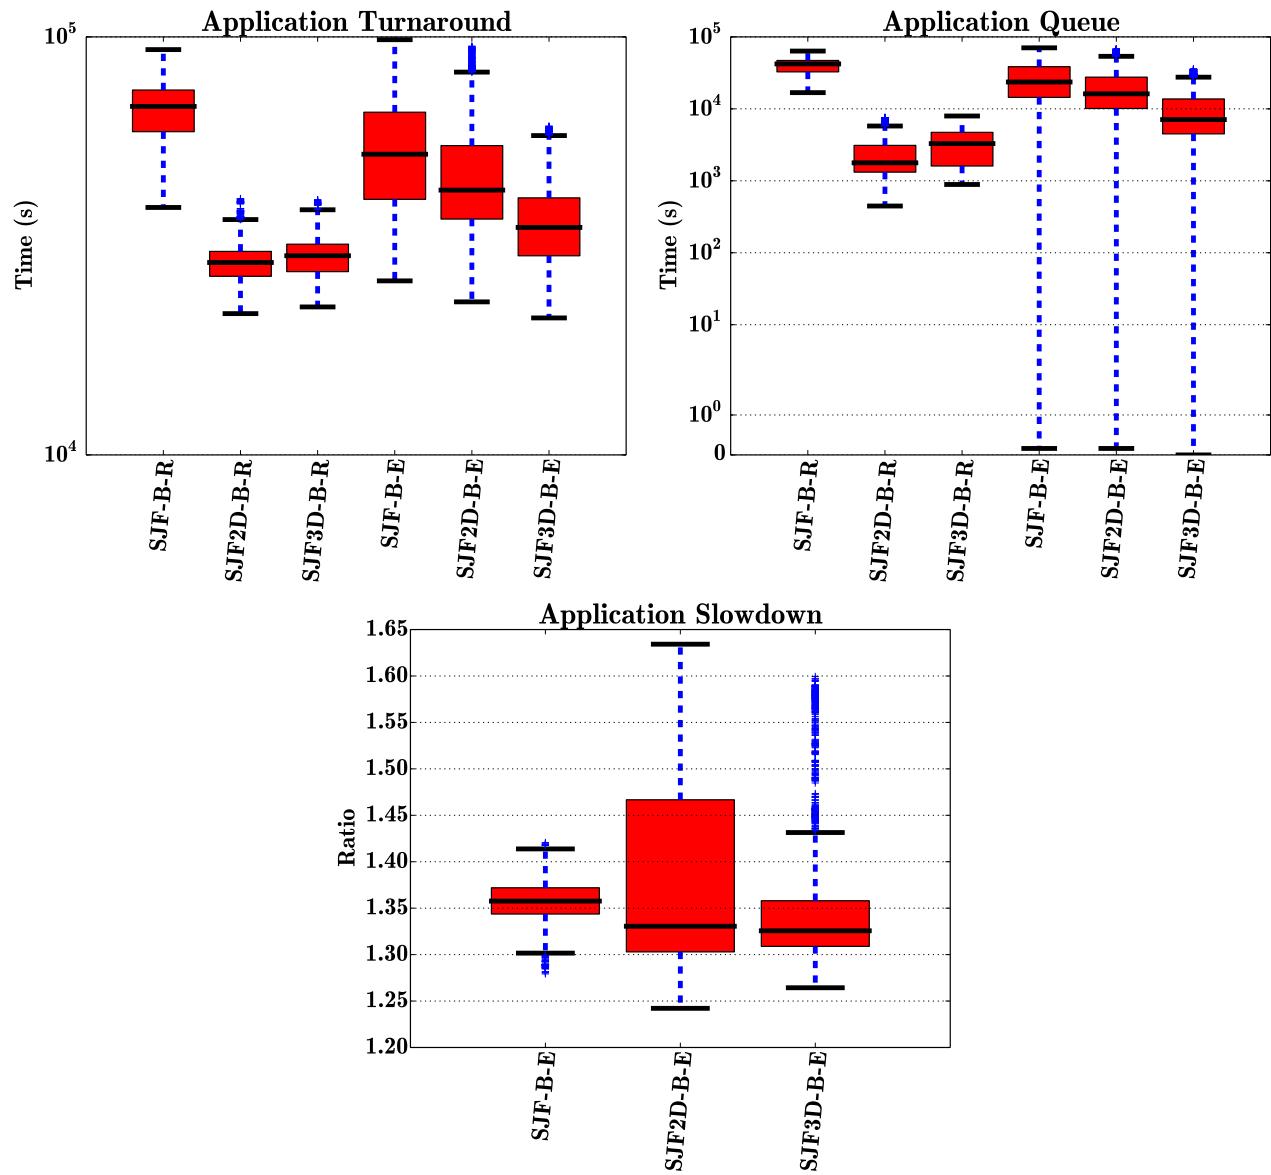


Figure 18: Comparison using different definition of size for SJF with a malleable scheduler. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

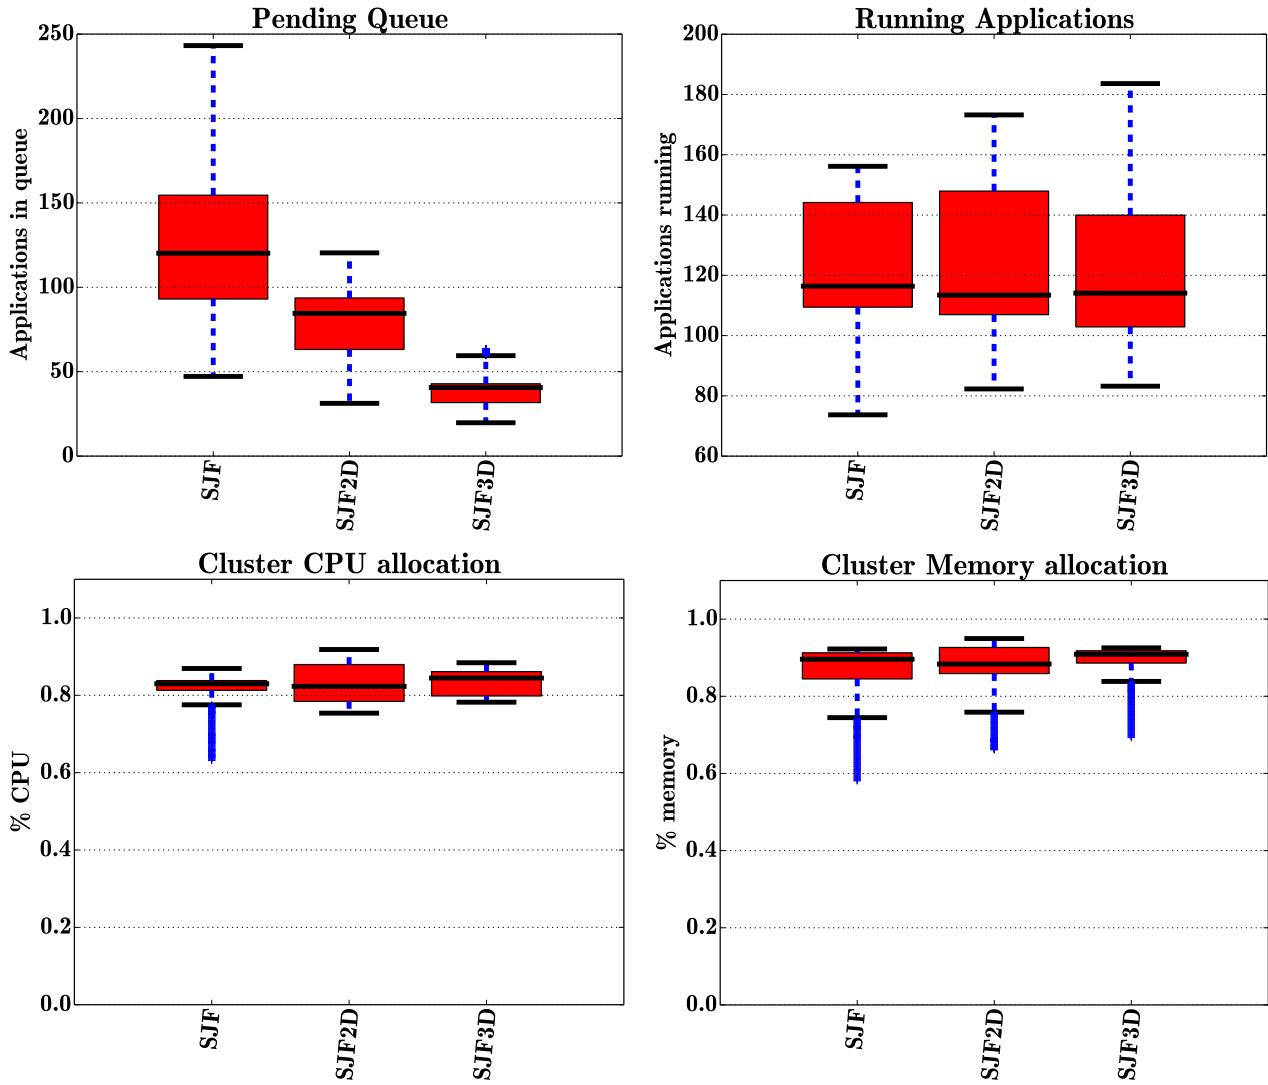


Figure 19: Comparison using different definition of size for SJF with a malleable scheduler. B-E stands for batch elastic and B-R stands for batch rigid applications.

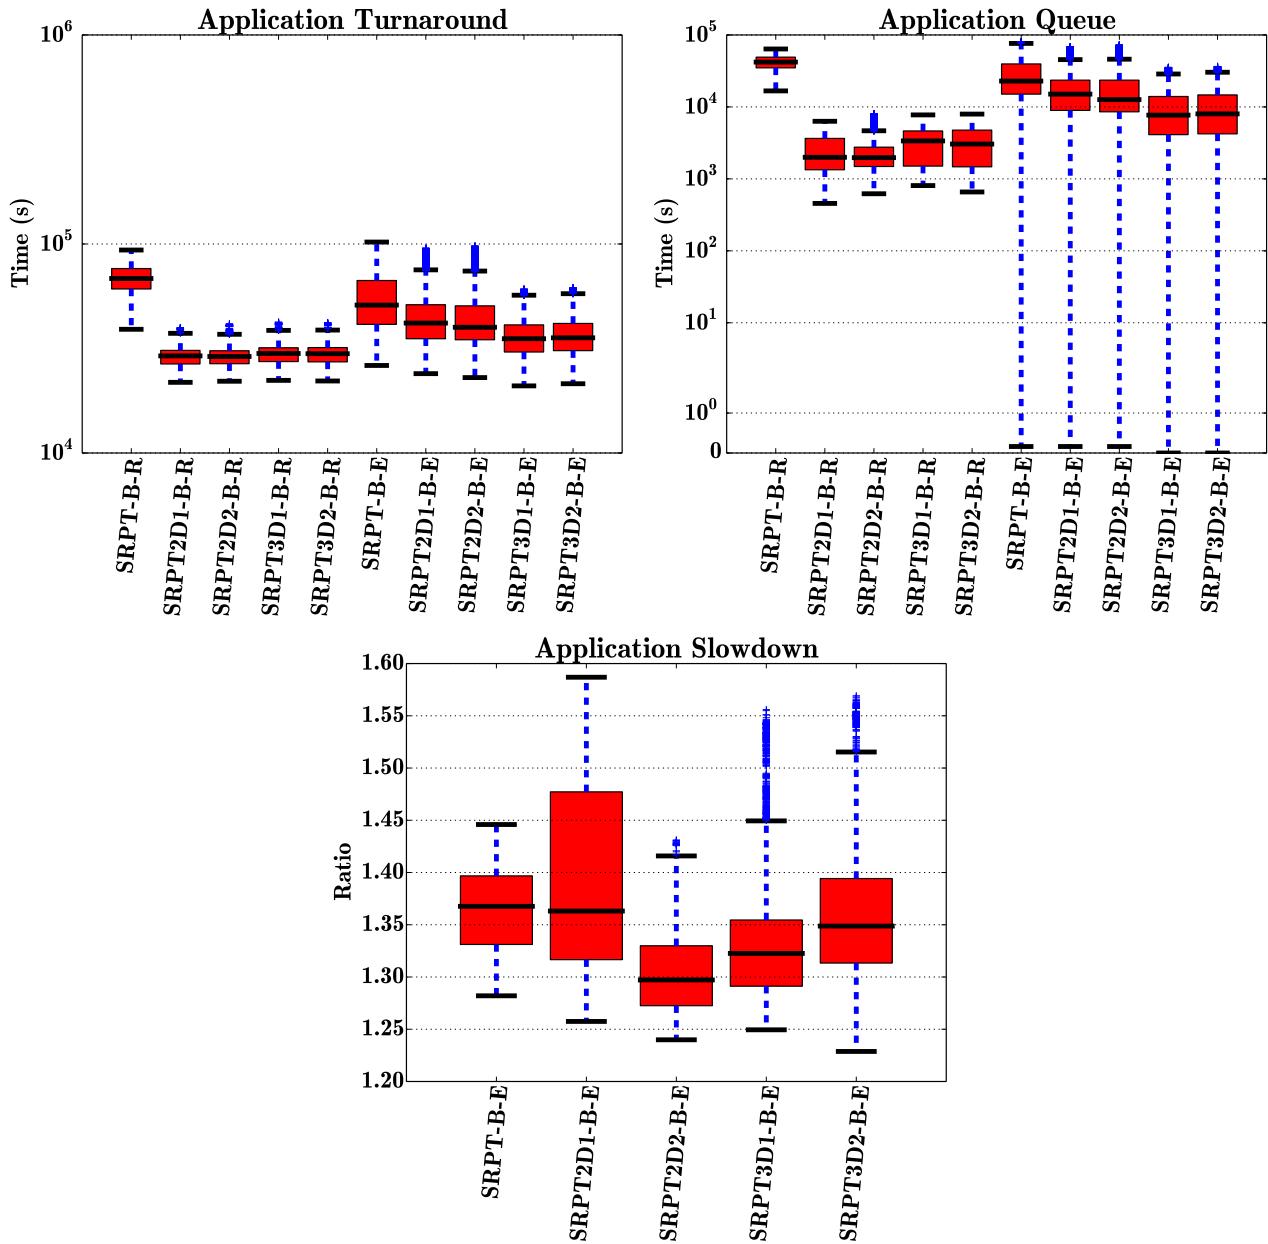


Figure 20: Comparison using different definition of size for SRPT with a malleable scheduler. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

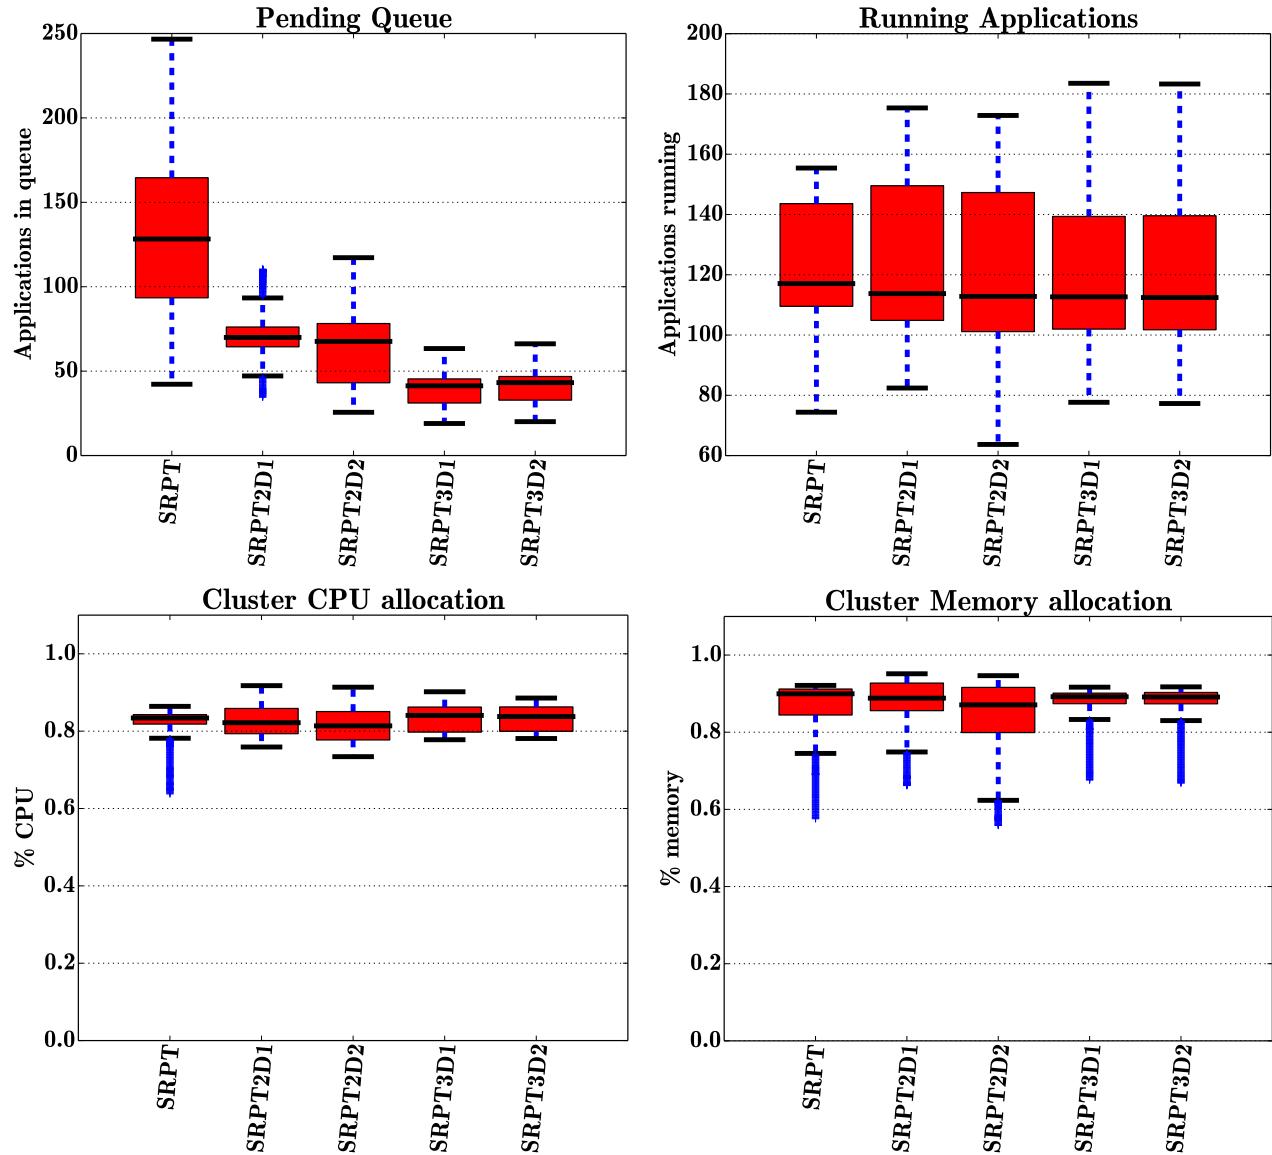


Figure 21: Comparison using different definition of size for SRPT with a malleable scheduler. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

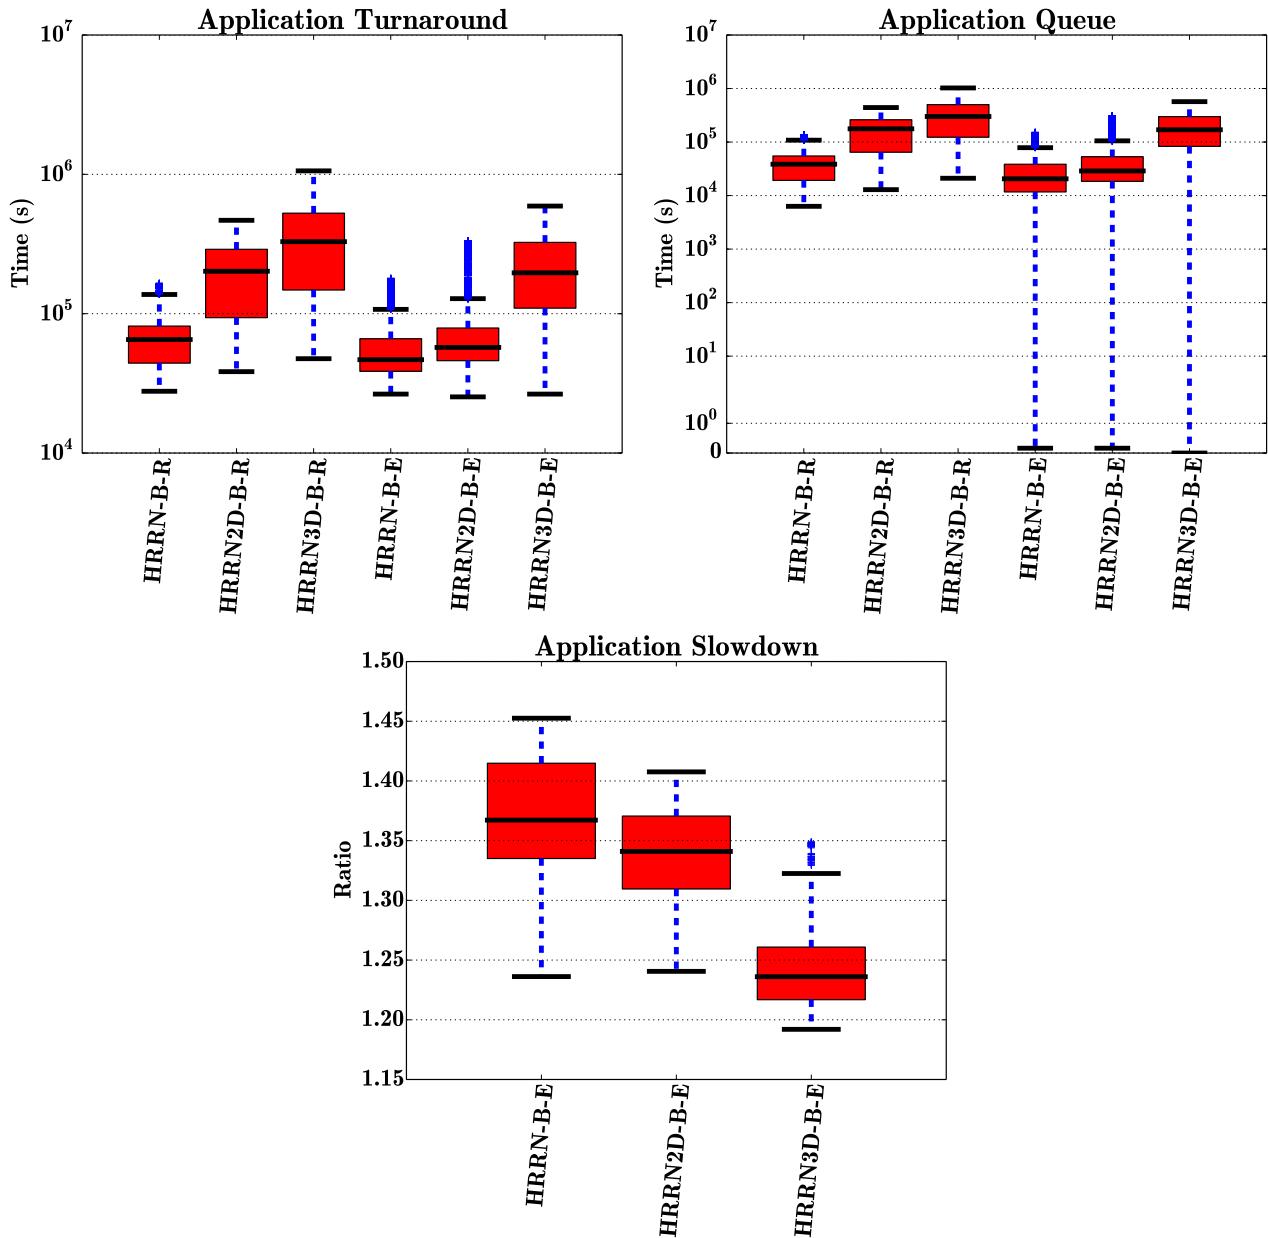


Figure 22: Comparison using different definition of size for HRRN with a malleable scheduler. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

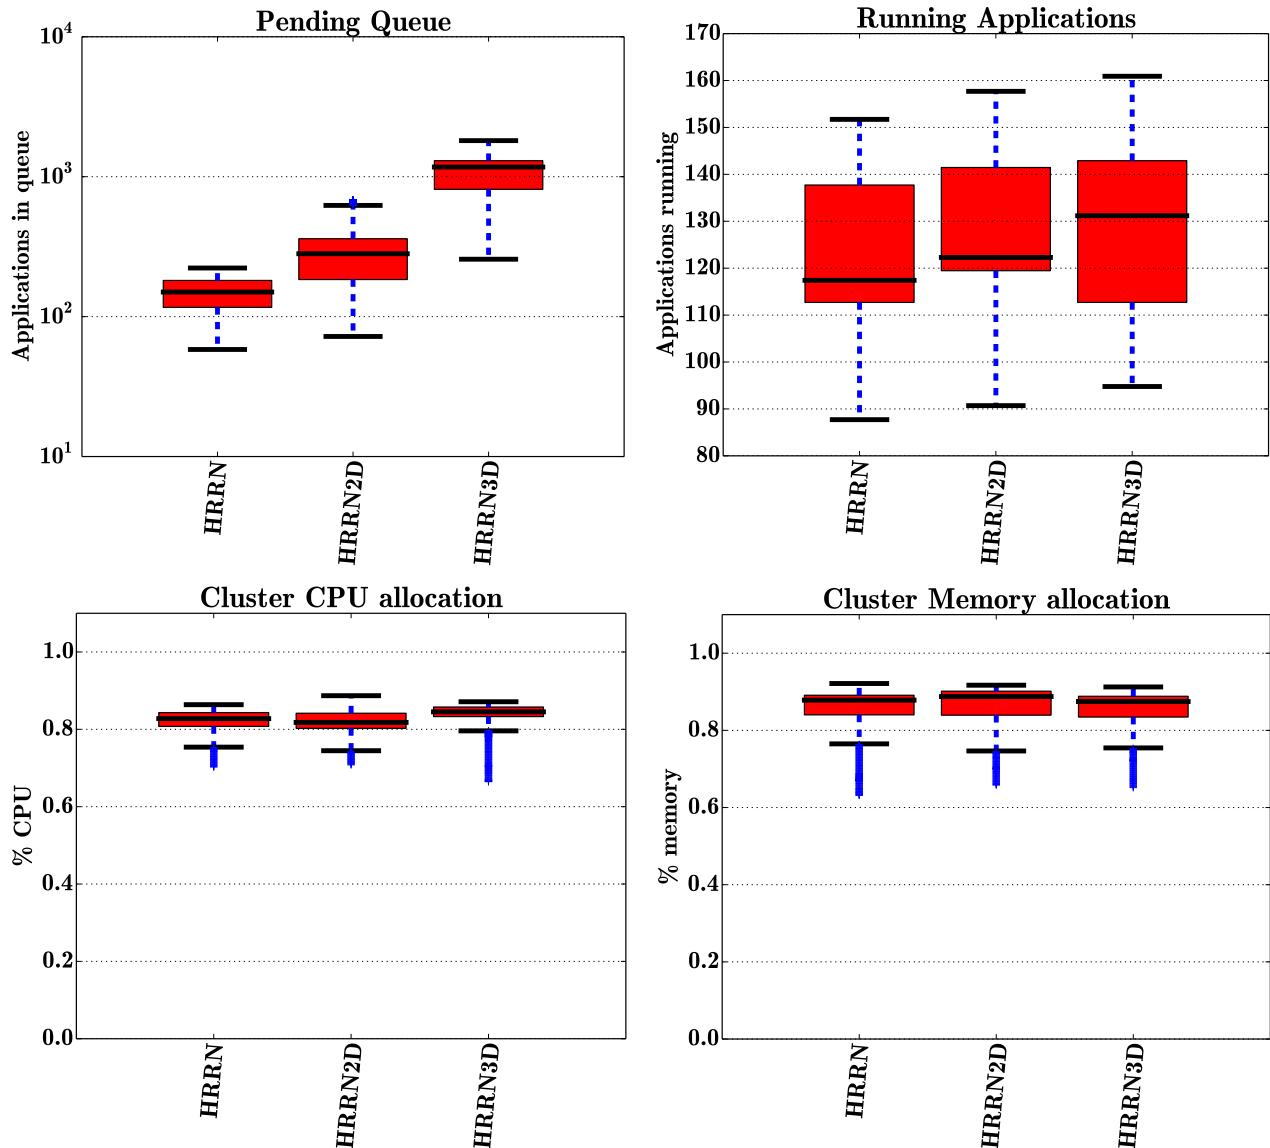


Figure 23: Comparison using different definition of size for HRRN with a malleable scheduler. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

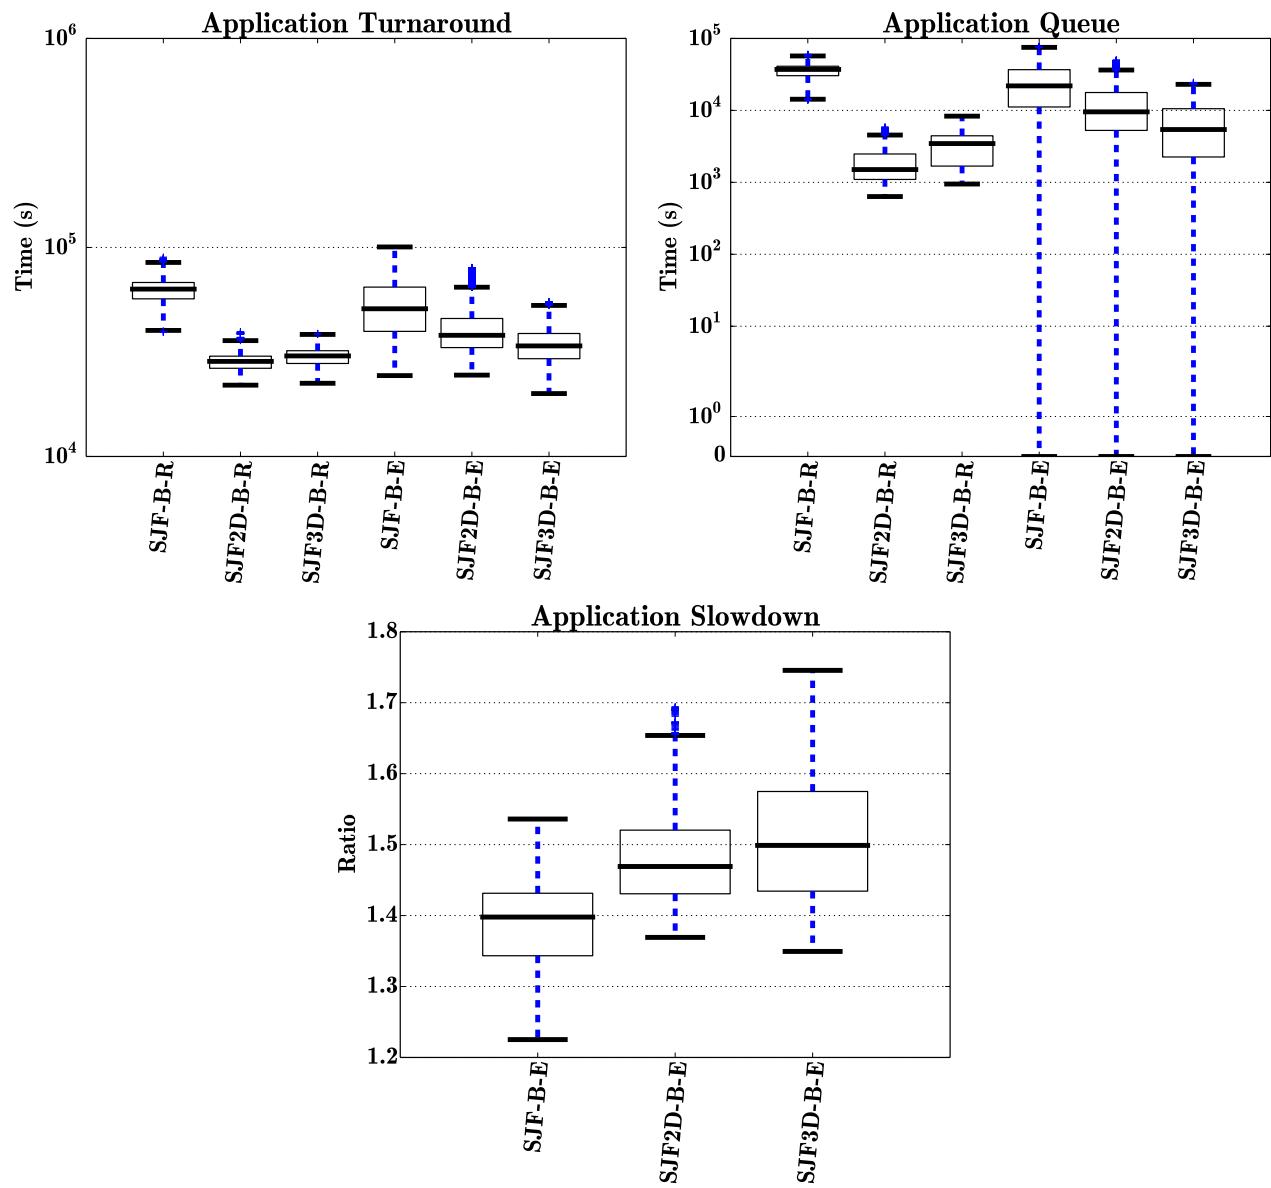


Figure 24: Comparison using different definition of size for SJF with our flexible scheduler. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

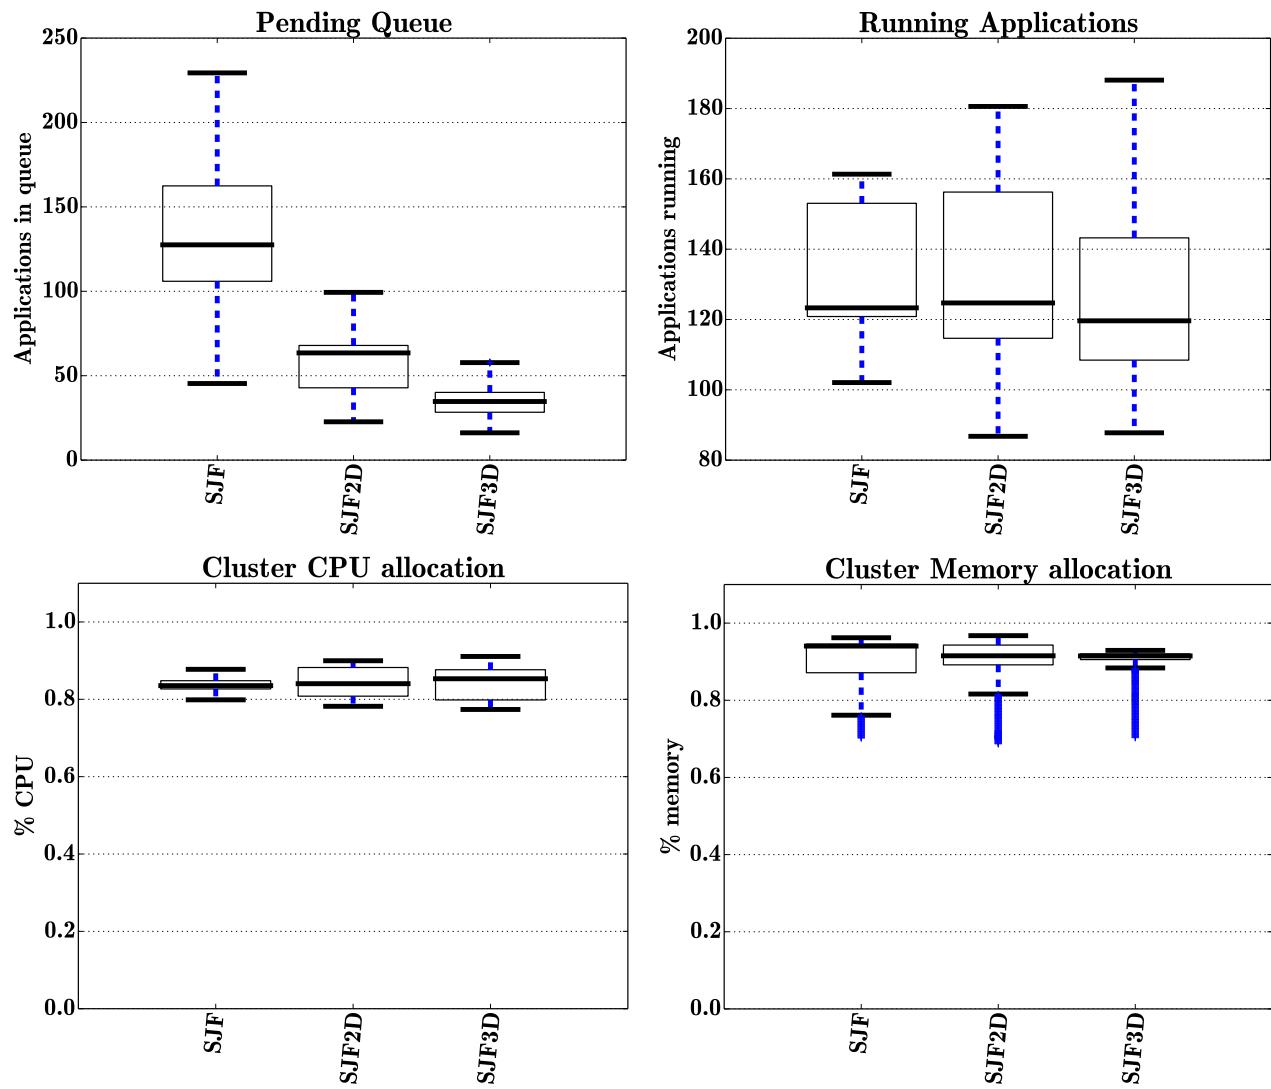


Figure 25: Comparison using different definition of size for SJF with our flexible scheduler. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

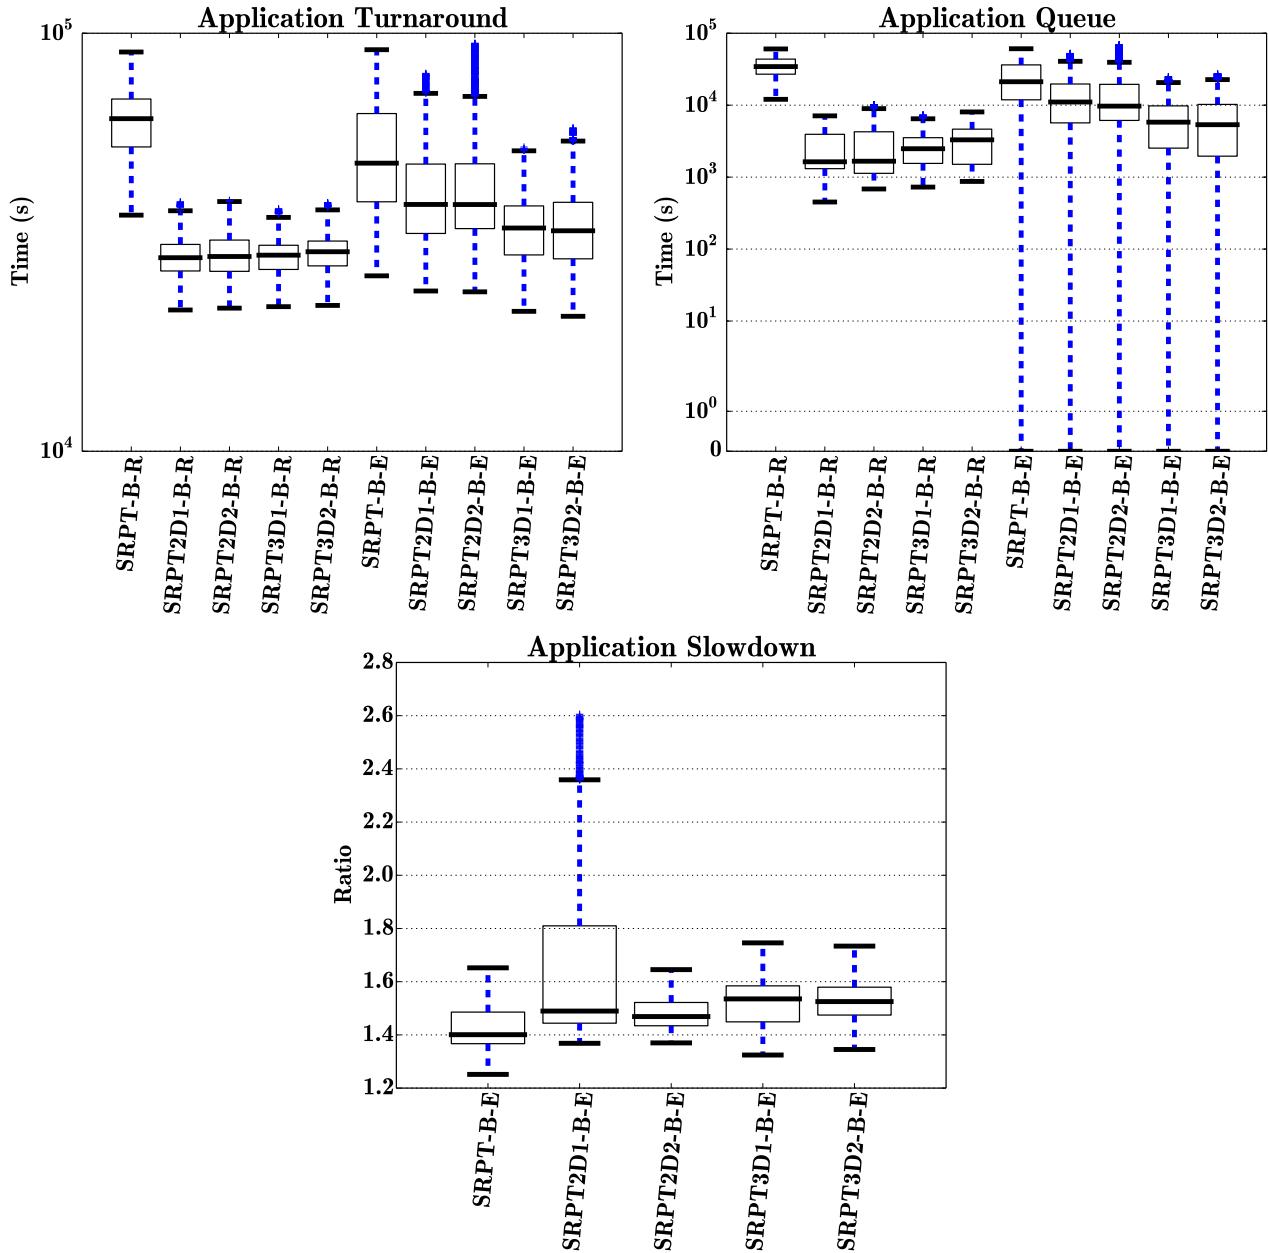


Figure 26: Comparison using different definition of size for SRPT with our flexible scheduler. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

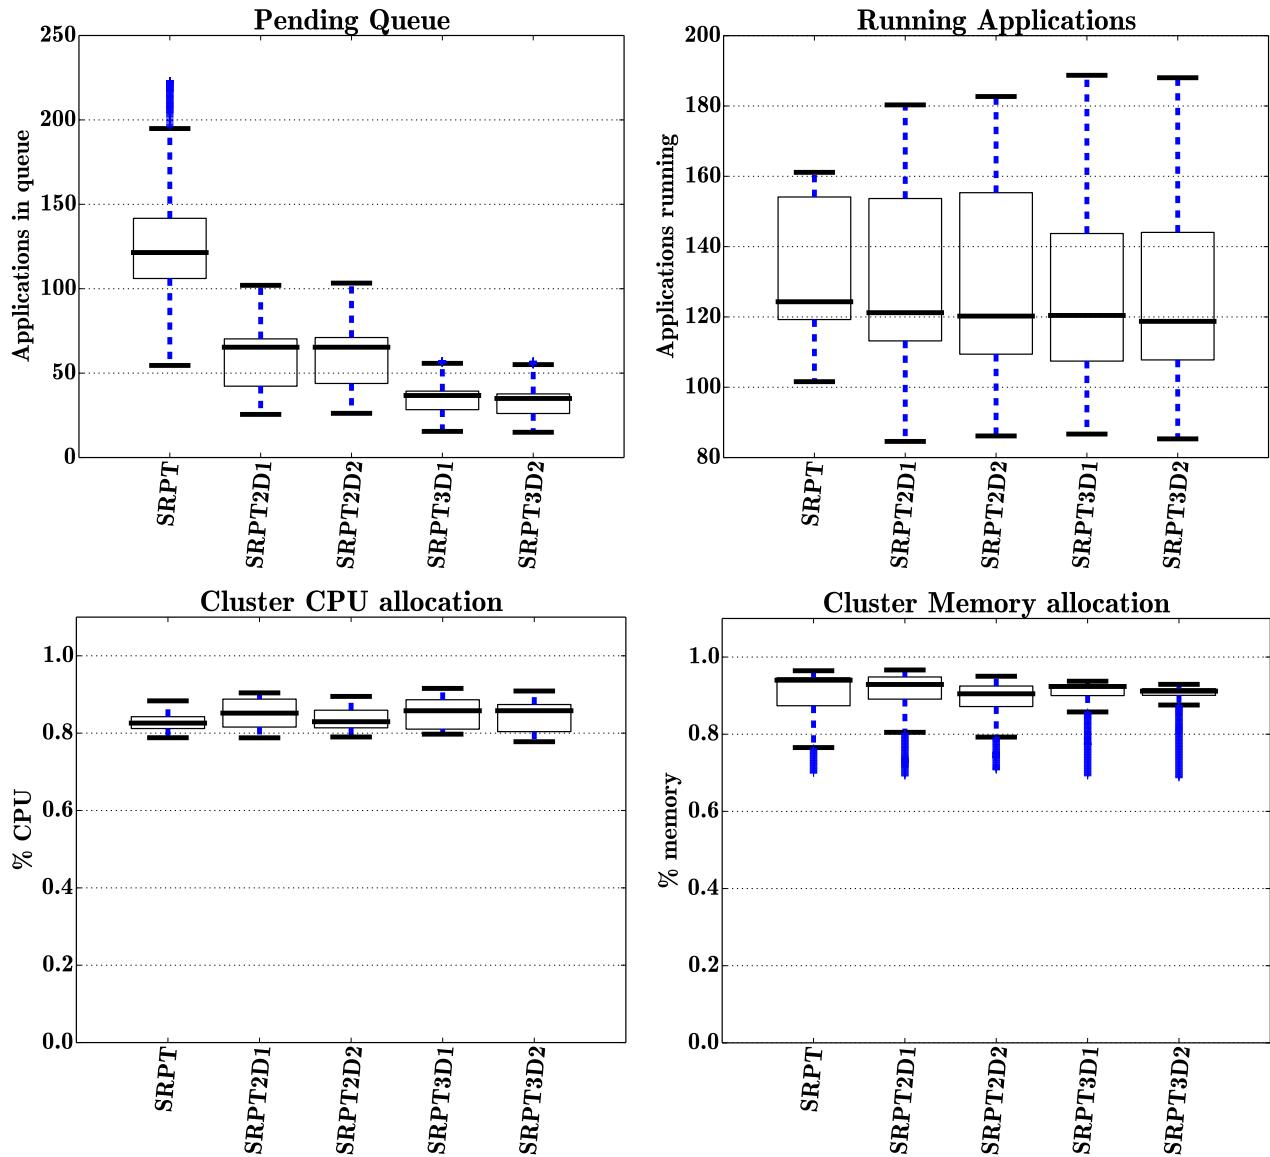


Figure 27: Comparison using different definition of size for SRPT with our flexible scheduler. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

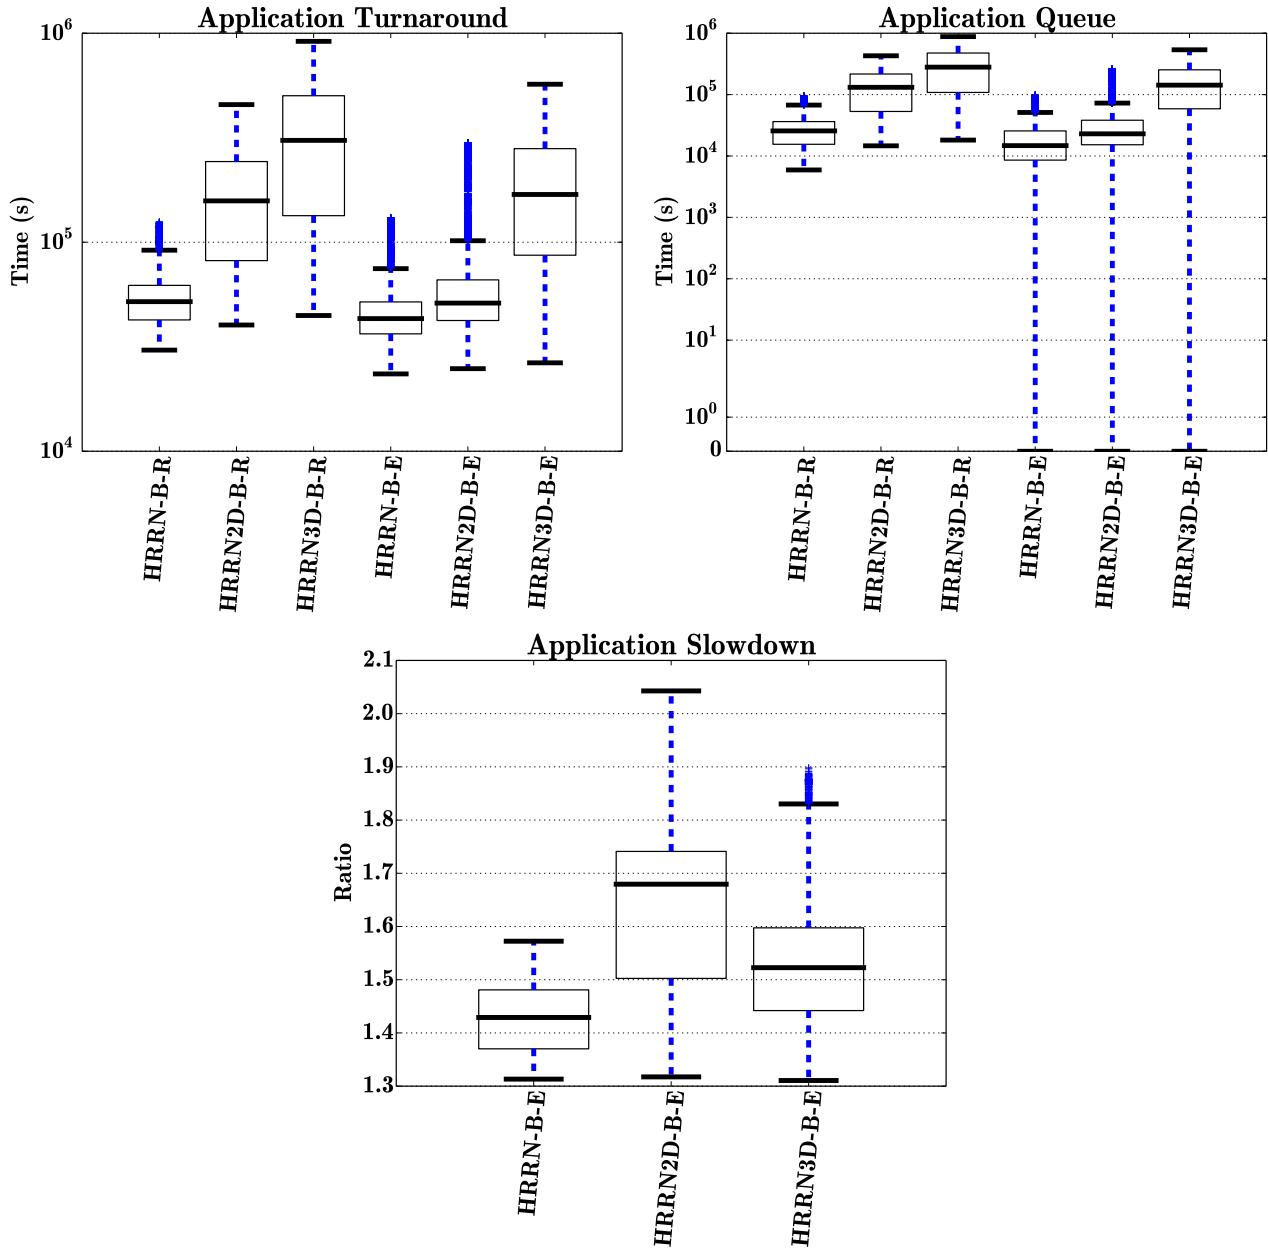


Figure 28: Comparison using different definition of size for HRRN with our flexible scheduler. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

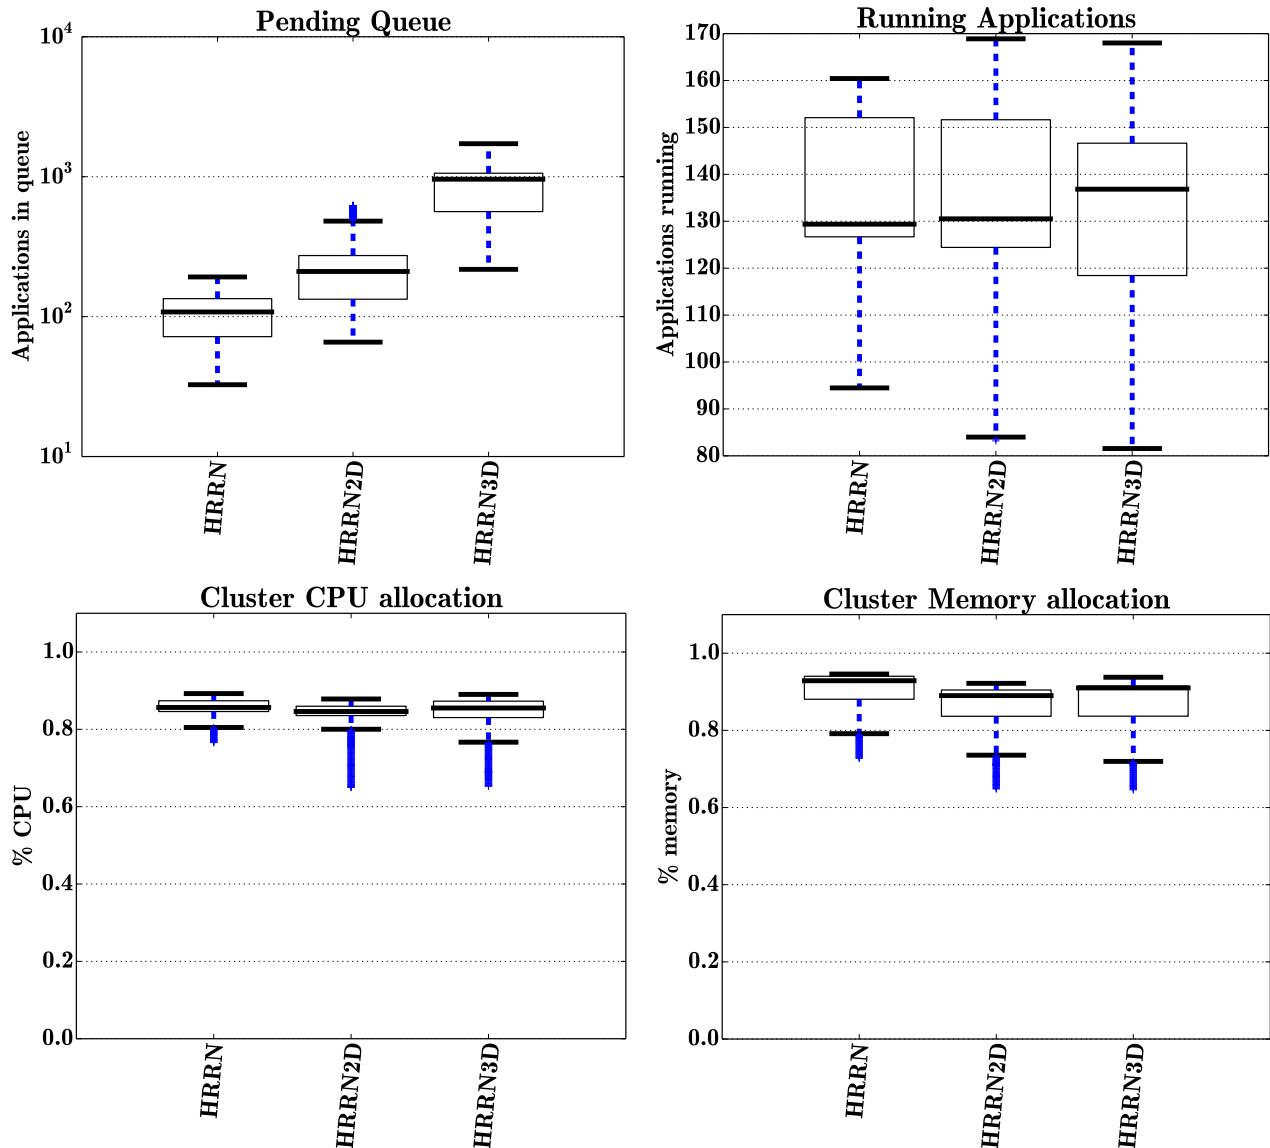


Figure 29: Comparison using different definition of size for HRRN with our flexible scheduler. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

size. The only policy that perform worst in all cases is HRRN; the results are expected since the policy avoid starvation by increasing the priority of applications while they remain in the queue, thus big and long running applications can start before short ones. Figures 15 to 29 show the difference between all the size definitions that we took in exam under a rigid, malleable and our flexible scheduler.

These results are possible because, in a unidimensional world, a short-living application can block the queue due to the high amount of resources requested. To avoid this situation we can add the information of the number of services required to actually allow small applications to be executed before large ones; as we see in Figures 18 to 29, this information improves the turn around time in all policies but HRRN. We would expect that adding more information to the size could further improve the turn around time, instead we see that this is not true for rigid schedulers (Figs. 15 to 17); having a three dimensional size definition may improve the turn around of some policies compared to a unidimensional definition, but in all cases having a two dimensional definition is sufficient to obtain good results. The resource allocation, however, decreases a bit in a rigid system as more information is added; it will allows big-but-short (big resource requirement but short lived) applications to block the queue thus, increase queue times and reduce resources utilization in the cluster. Since our flexible scheduler allows only the indispensable (inelastic) services to be launched, it can overcome the problems that a rigid has with a three dimensional definition and, as a result, we can see that it performs better than a one or two dimensional definitions.

In conclusion, adding more information to the size does not always improve the turn around time in a rigid scheduler because some big-but-short applications can block the queue. However, with a malleable and with our flexible scheduler this can be avoided and we can have lower turn around time and higher cluster resources utilization with a three dimensional definition of size. We would like to stress more the fact that the malleable scheduler is something widely adopted in theory, but not in use real systems.

10.4 Impact of different workloads

In Sections 10.2 and 10.3 we show the benefits on turn around time of a flexible scheduler solutions – that split the services of an application in core and elastic – compared with a *rigid* scheduler that does not. It is intuitive to think that the results highly depend on the workload; on the one hand, more percentage of elastic applications lead to higher benefit for our flexible scheduler; on the other hand, more inelastic (composed by only core components) applications will move the results much closed to the rigid. However, in the worst case our flexible scheduler will perform exactly the same as the rigid; the flexible scheduler will have to allocate the requested resources at one, therefore its behavior is the same as the rigid. Table 10.4a shows the results of a simulation with a completely inelastic workload; the results show that the average turnaround is exactly the same.

In conclusion, our flexible scheduler does not introduce any overhead and, in the worst case, will not perform worst than a rigid.

10.5 Preemption

We turn now our attention to the full workload we defined in Section 10.1, including interactive applications. Preemption is used when a high-priority, interactive application requires resources: this applies both to manually set priorities (e.g., in a FIFO policy) and to size-based priorities. In particular, we report results for the preemptive version of the SRPT policy.

	FIFO	PSJF	SRPT	HRRN
Rigid	1,317,315.92	135,964.50	135,964.50	407,952.05
Hybrid	1,317,315.92	135,964.50	135,964.50	407,952.05

Table 10.4a: Turn around time (in second) between a rigid and our flexible scheduler with a workload composed by only inelastic applications.

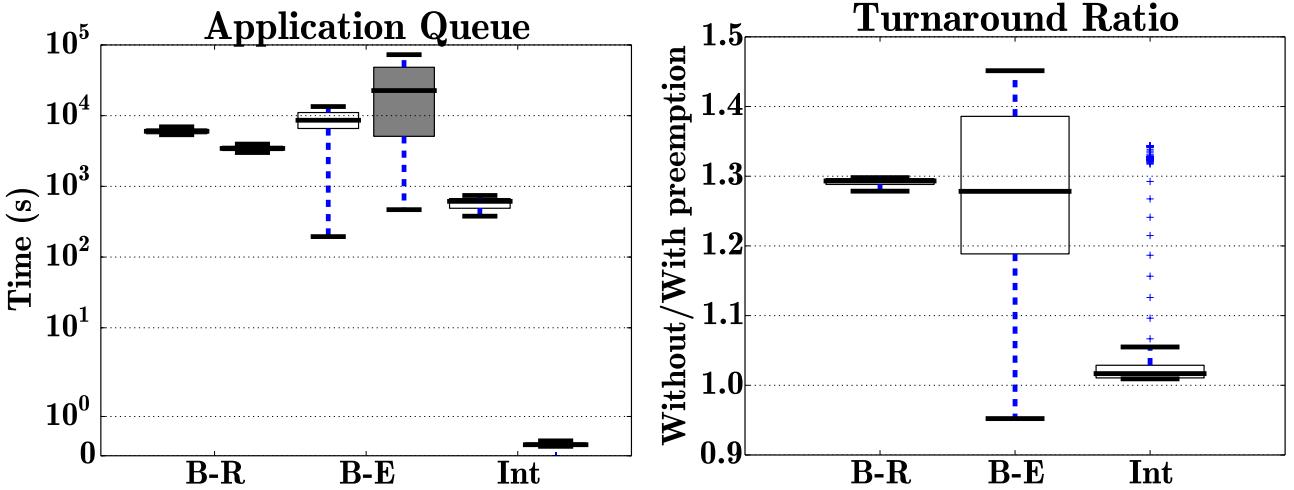


Figure 30: On the left, comparison of queuing time distributions between scheduling with and without preemption. White boxes (left box of every pair) correspond to a non-preemptive system, gray boxes to our preemptive algorithm. On the right, turnaround ratio distributions between scheduling with and without preemption. *B-E* stands for batch elastic applications, *B-R* stands for batch rigid applications and *Int* is for interactive applications.

Figure 30 shows the most relevant percentiles of the distribution of application queuing times, grouped by application type (both cases of batch and interactive applications), with and without our preemption mechanism. Globally, preemption does not subvert the perceived system responsiveness. However, interactive applications under preemptive scheduling enjoy roughly two orders of magnitude less queuing times. Users do not wait for few dozens minutes but only few seconds, for their interactive application to start. As a consequence, elastic batch applications pay with more variability (but stable for the median case) in queuing times.

Since our simulator does not account for real work done by applications, the preemption mechanism does not have any effect on the work that has been done by preempted components. In practice, our current preemption mechanism would instead suppress work done by elastic services, if preempted. Studying new preemption primitives, e.g. by suspending Linux containers, is part of our research agenda: this is the main reason why our current prototype implementation lacks support for preemption.

Finally with Figs. 31 to 33 we show the results for all the policies with their different definition of size that we identified in Section 10.3; globally, the considerations made before for the SRPT policy can be applied for the other policies as well. However, we can see a drop in the resource utilization and an increase in the size of the pending queue when using a preemptive system; this is because we are no longer respecting the order imposed by the policy, but we are prioritizing interactive applications that have a much longer duration compared to the batch, thus a lot of shorter applications are forced to wait.

11 Experiments with Zoe

Our goal now is to perform a comparative analysis of two generations of Zoe: the first, implementing a rigid scheduler, as for the baseline, the second with the flexible scheduler we present here. In our experiments, we replay the exact same workload trace for both generations. Each trace takes about 3 hours from the first submission to the last. During our experiments, no other user was allowed to submit jobs to Zoe.

Workload. We use two representative batch application templates, including: 1) an elastic application using the Spark framework; 2) a rigid application using the TensorFlow framework. Following the statistical distribution of our historical traces, we set our workload to include 80% of elastic and

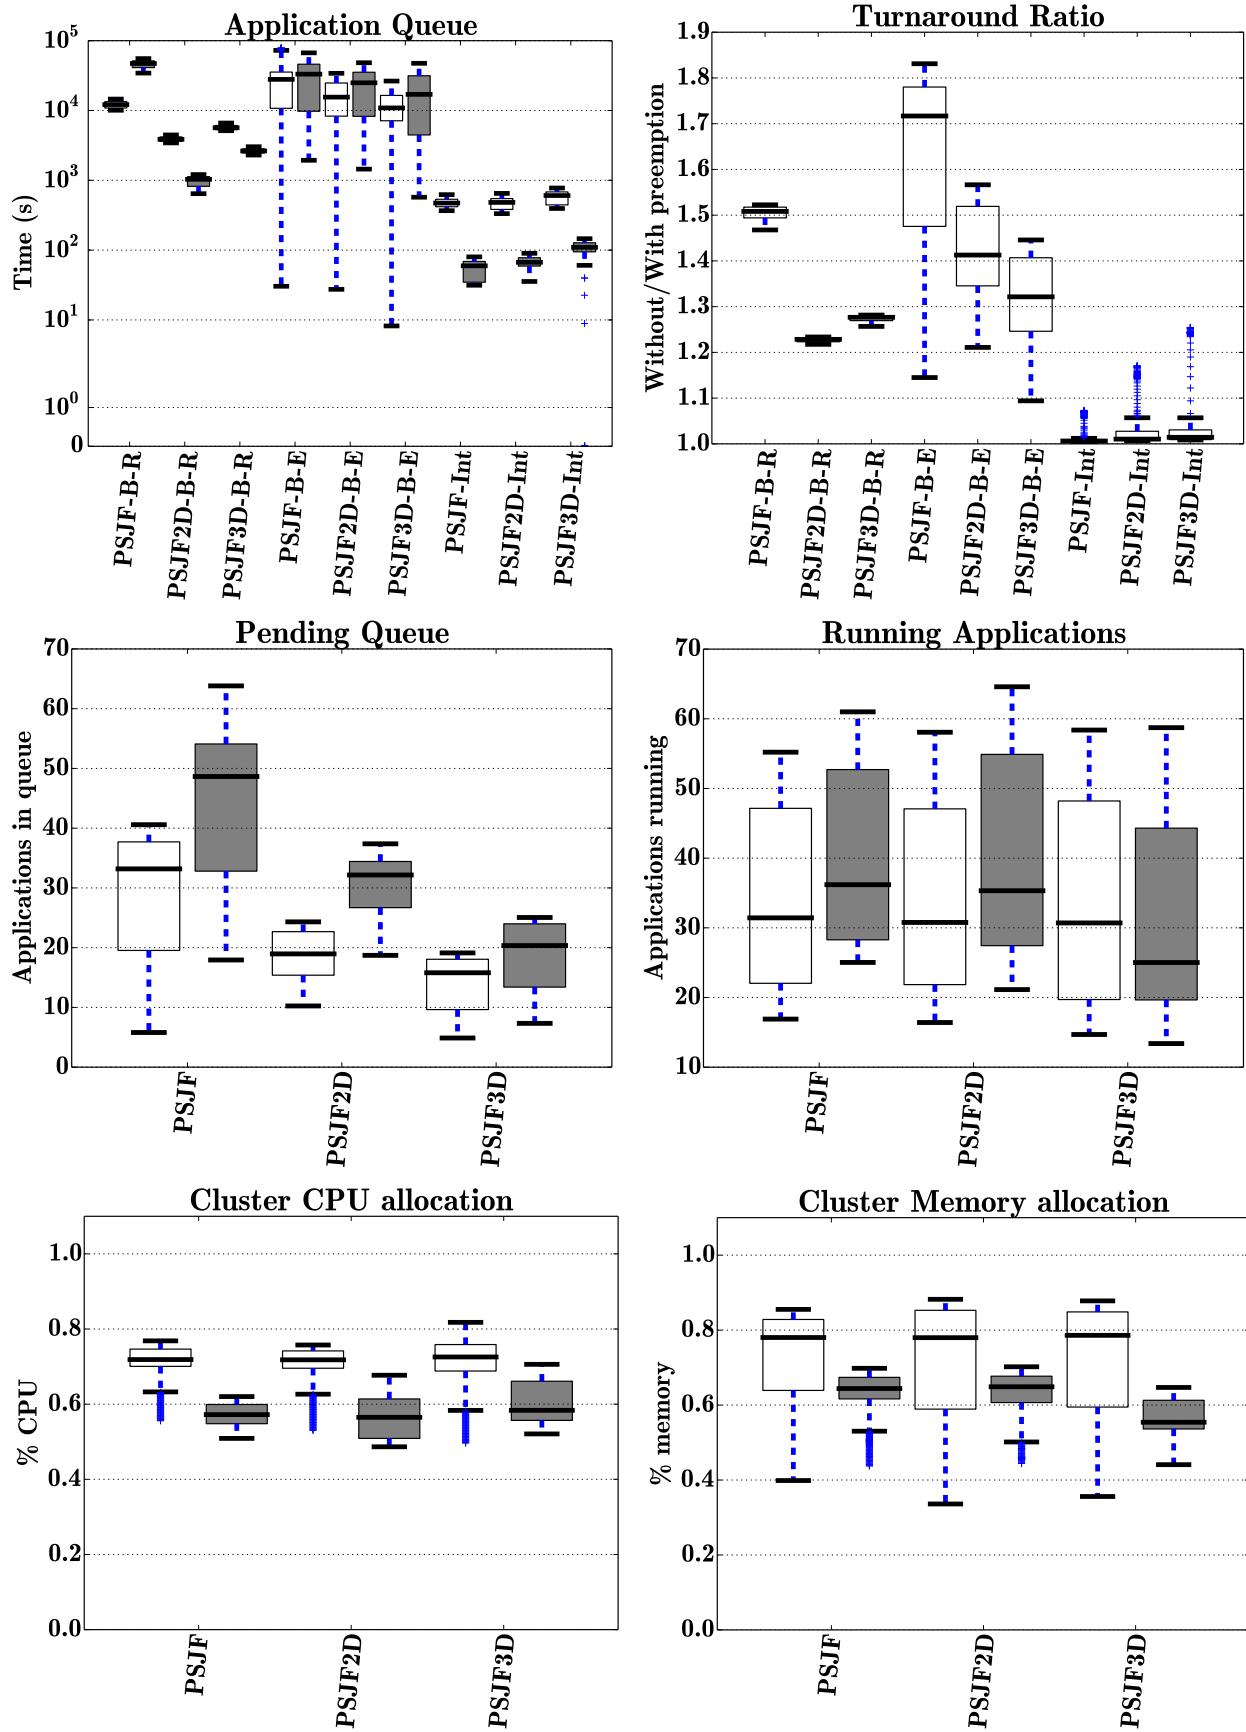


Figure 31: Comparison between a system with and without preemption for PSJF and its different definition of size. White boxes (left box of every pair) correspond to a non-preemptive system, gray boxes to our preemptive algorithm. *B-E* stands for batch elastic applications, *B-R* stands for batch rigid applications and *Int* is for interactive applications.

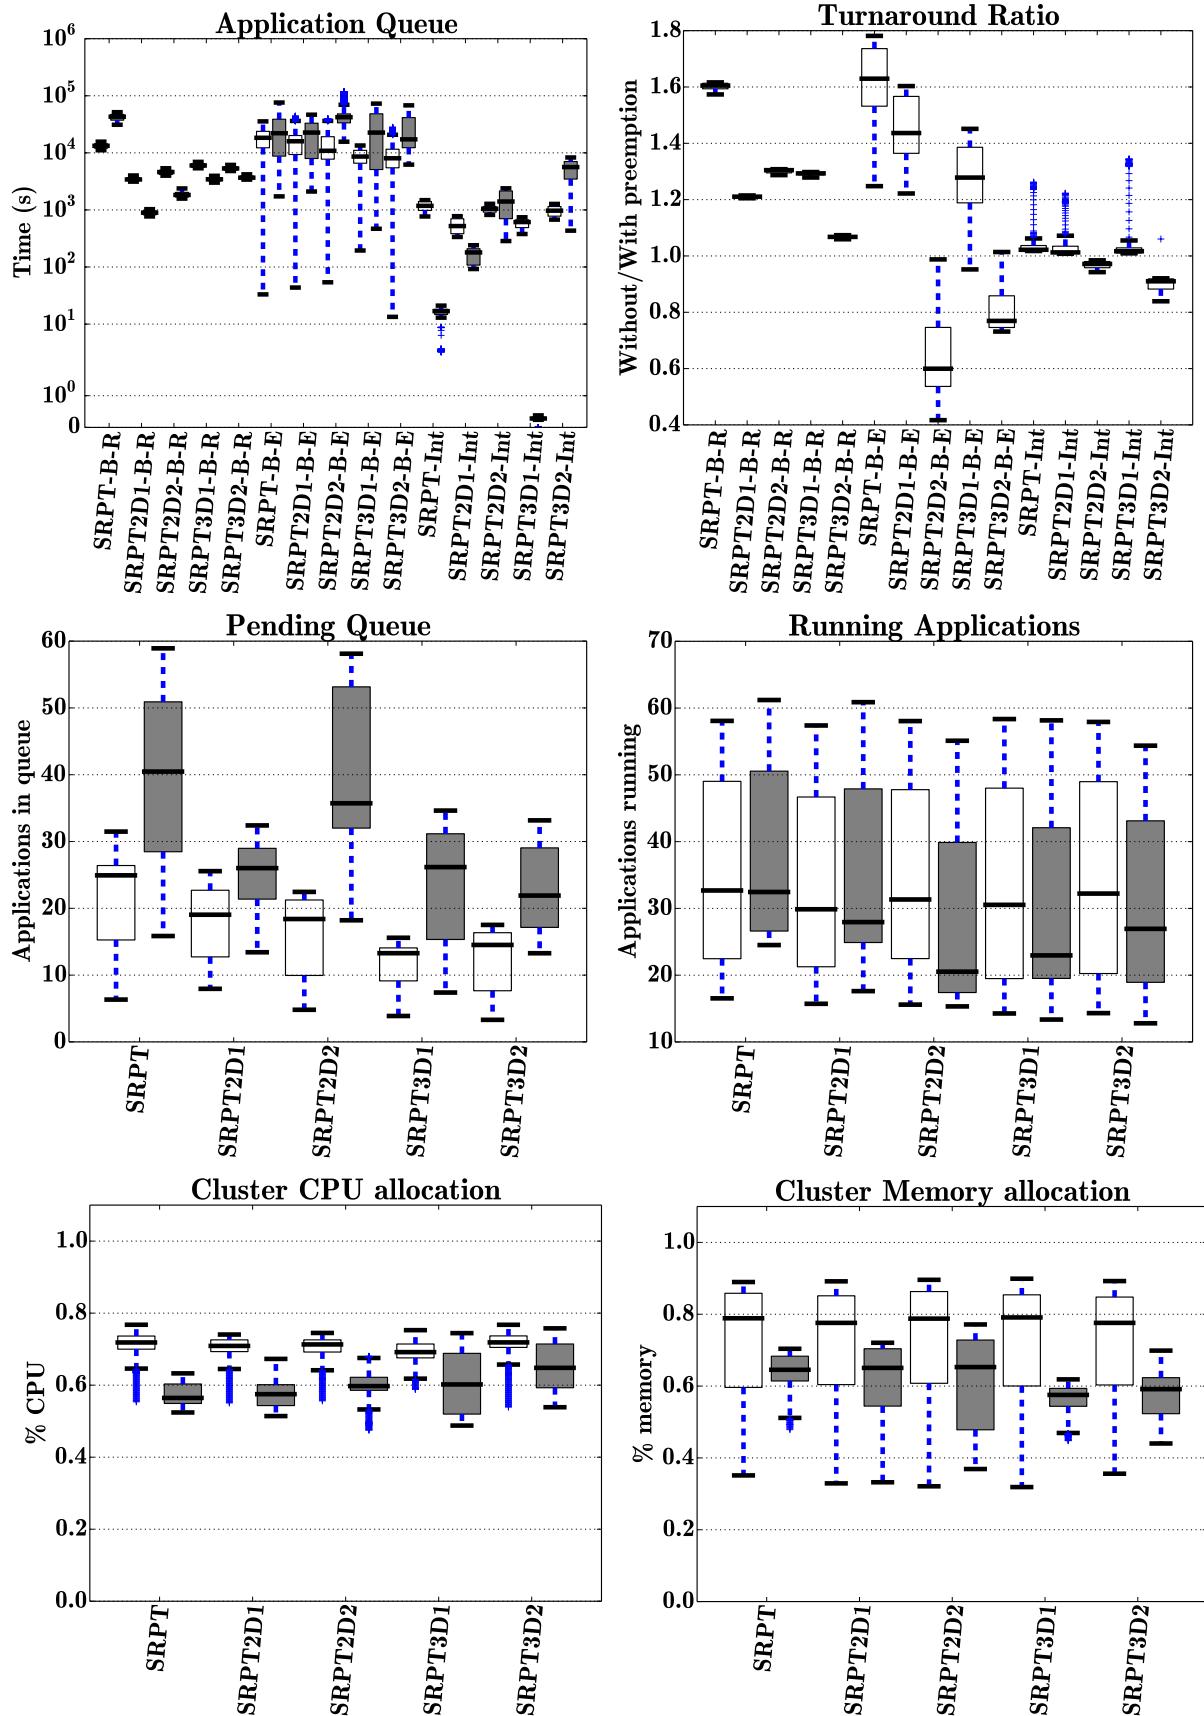


Figure 32: Comparison between a system with and without preemption for SRPT and its different definition of size. White boxes (left box of every pair) correspond to a non-preemptive system, gray boxes to our preemptive algorithm. *B-E* stands for batch elastic applications, *B-R* stands for batch rigid applications and *Int* is for interactive applications.

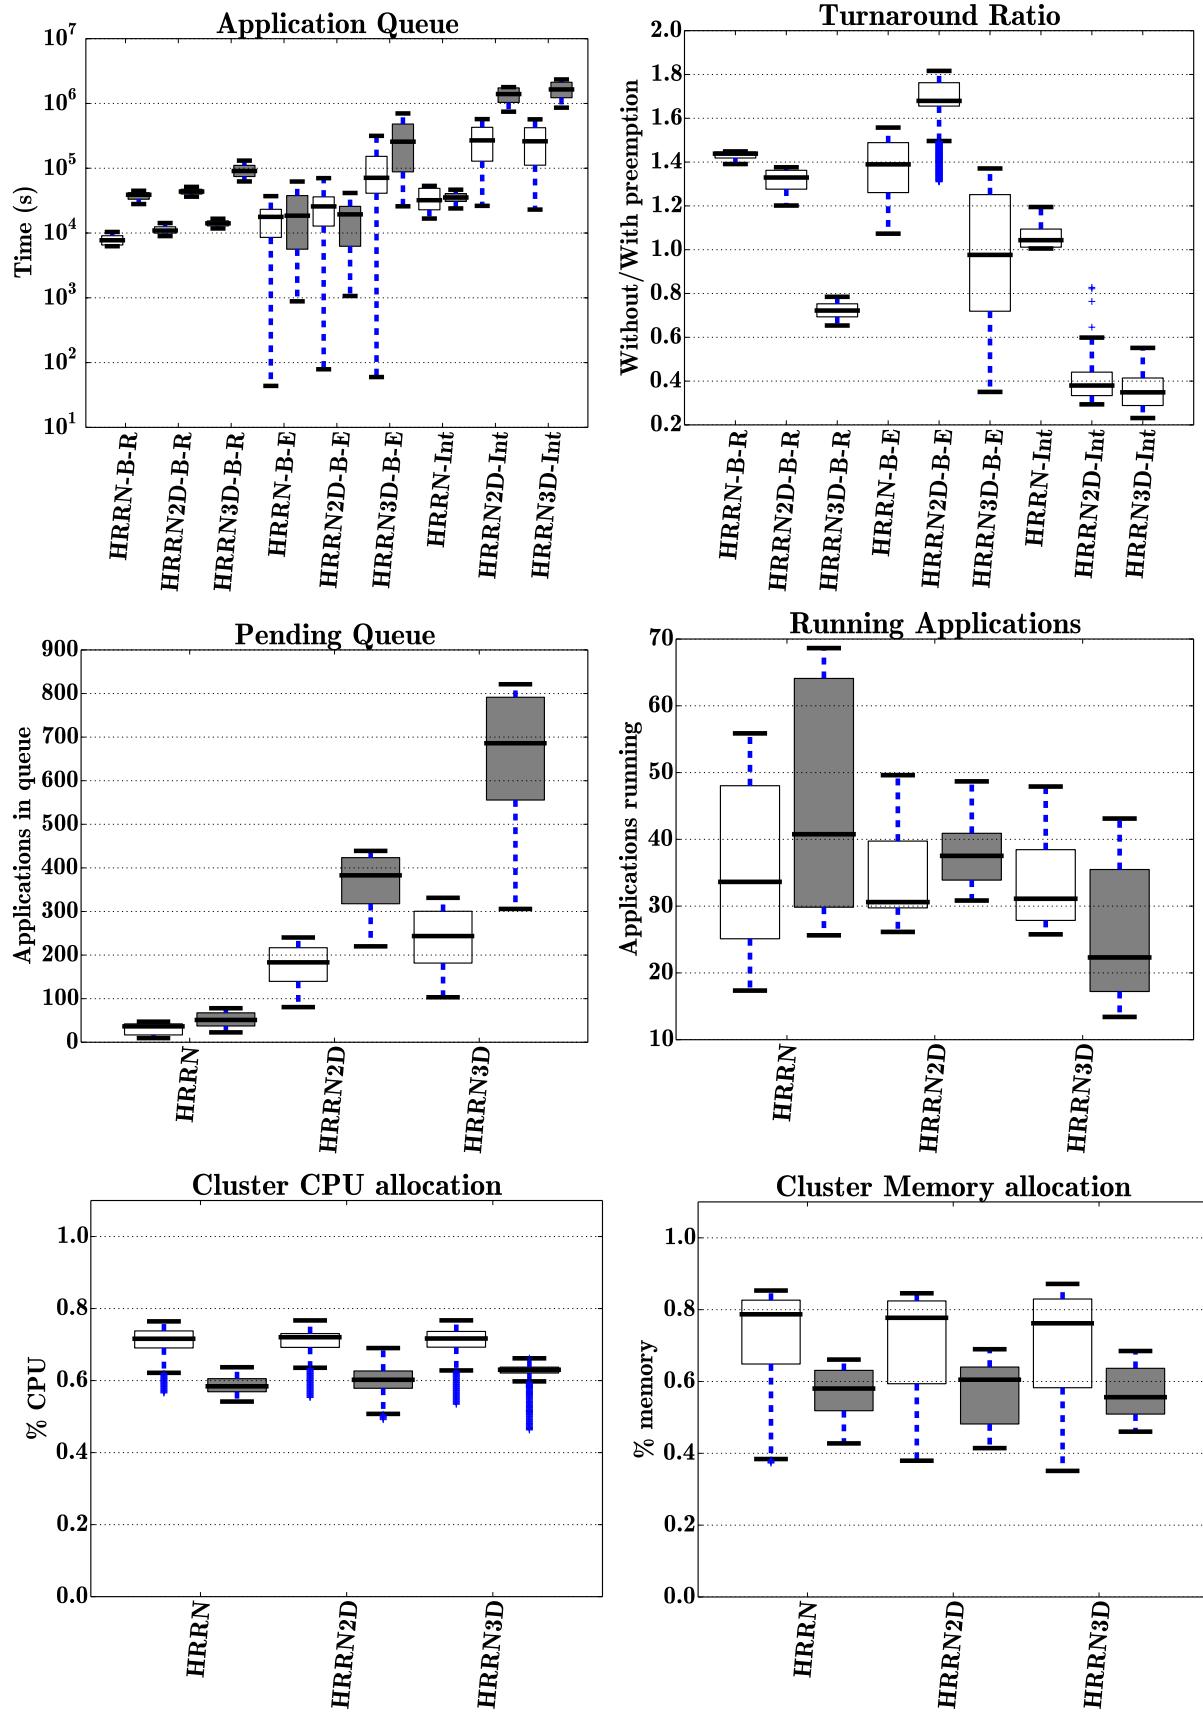


Figure 33: Comparison between a system with and without preemption for HRRN and its different definition of size. White boxes (left box of every pair) correspond to a non-preemptive system, gray boxes to our preemptive algorithm. *B-E* stands for batch elastic applications, *B-R* stands for batch rigid applications and *Int* is for interactive applications.

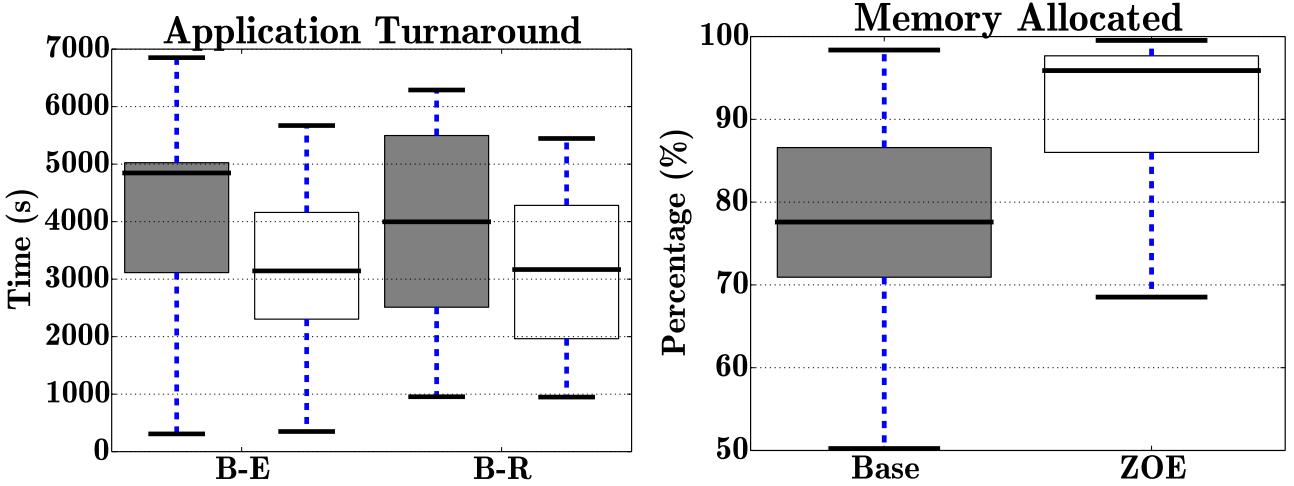


Figure 34: Comparison of turnaround time distributions using the FIFO discipline. White boxes (right box of every pair) correspond to the second generation of Zoe that implements our algorithm. *B-E* stands for batch elastic and *B-R* stands for batch rigid applications.

20% of rigid applications, for a total of 100 applications. Application inter-arrival times follow a Gaussian distribution with parameters $\mu = 60$ sec, and $\sigma = 40$ sec, which is compatible with our historical data. More specifically, using the elastic application templates, we run two use cases. First, an application to induce a random-forest regression model to predict flight delays, using publicly available data from the US DoT.⁷ Second, a music recommender system based on alternating least squares algorithm, using publicly available data from Last.fm⁸. Both applications have two different requirements in term of memory for each elastic component. The random-forest regression model has 3 core components and then 32 elastic components of 16GB or 8GB RAM each; every elastic component uses 1 CPU. The music recommender system has 3 core components and then 24 elastic components of 16GB RAM or 8GB each; every elastic component uses 6 CPU. Instead, using the rigid application template, we train a deep Gaussian Process model [43], and use both a single-node and a distributed TensorFlow program, requiring 1 and 10 workers (and 5 parameter servers) each with 16GB of RAM.

Experimental setup. We run our experiment on a platform with ten servers, each with two 16-core Intel E5-2630 CPU running at 2.40GHz (total of 32 cores with hyper-threading enabled), 128GB of memory, 1Gbps Ethernet network fabric and ten 1TB hard drives. No GPU-enabled machines are available in our platform, at the moment. The servers use Ubuntu 14.04, Docker 1.11 and the standalone Swarm manager. Docker images for the applications are preloaded on each machine to prevent container startup delays and network congestion.

Summary of results. Using the FIFO scheduling policy, we compare the two generations of Zoe according to the distributions of application turnaround times, as shown in Figure 34 (left). The behavior of the two systems indicate a clear advantage for our approach: the median turnaround times are 37% and 22% lower, for elastic and rigid applications respectively. Note also that the tails of the distributions are in favor of our approach.

Overall, the new generation of Zoe that implements the flexible scheduler is more efficient, with a 20% improvement, in allocating and packing applications, as illustrated in Figure 34 (right), where we show the ratio of the distribution of allocated over available resources.

Finally, we present results concerning a low-level metric that measure the application ramp-up time, *i.e.*, the time it takes for applications scheduled for running, to receive their allocations and start producing work. Zoe achieves a container startup time, including placement decisions, of $0.90 \pm$

⁷<http://stat-computing.org/dataexpo/2009/the-data.html>

⁸http://www-etud.iro.umontreal.ca/~bergstrj/audioscrobbler_data.html

0.25ms. Full-fledged applications, made by several containers, only take few seconds to start, which is a compelling property, especially when compared to existing solutions such as Amazon EMR.

12 Related Work

While we cannot do justice to the richness of the scheduling literature, in this section we organize related work in three groups. The “competitors” are existing, mature systems that could be considered sufficient to address our problem statement, at a first sight. The second group includes recent works in the systems research literature, while the third covers works on scheduling at the task level.

Many “competitor” systems have been designed to cope with the problem of sharing cluster resources across a heterogeneous set of applications, some of which can be tweaked to achieve the goals we set in this work. For example, Yarn [31] and Mesos [18] have been among the first to enable multiple frameworks to coexist in the same cluster: usage of these “two-level” schedulers yield a big improvement as compared with monolithic approaches to resource scheduling. Originally designed for analytic frameworks, such systems deal with the scheduling of low-level processing tasks. Recently, more general approaches address the problem of cluster-wide resource management: Omega [19], then Borg [20] (and Kubernetes [32]) reason at the “container” level, and are optimized to achieve efficient placement and utilization of cluster resources, when absorbing a very heterogeneous workload. This latter includes a majority of long-running services, which power Web-scale, latency-sensitive applications. Additionally, container orchestration frameworks, such as Docker Swarm [33], also provide efficient and scalable solutions to the problem of scheduling (that is, placing and provisioning) containers in a cluster. Our work relies on many of the above systems, and can use them as a back-end to support scheduling of high-level applications rather than provisioning low-level containers. Existing auxiliary deployment tools such as Aurora [44] and Docker Compose [45], do not address scheduling problems.

In the systems research literature, we find several inspiring system designs, with ideas that can be “borrowed” to further extend our system prototype and research scope. For example, Koala-f [24] tackles dynamic resource allocation problems, which manifest in our case with “idle” interactive applications. Similarly to HCloud [23], we believe important to break the reservation paradigm, and allow users to express performance requirements rather than machine-level resource counts. This can be easily included in an application description, but requires developing additional components to infer statistical models of application properties based on systems observations, as done for example in Tarcil [29], Paragon [25] and Quasar [26]. Overall, by focusing on a higher-level of abstraction, the focus of our work is to address a rather abstract scheduling problem: our implementation indicates that our ideas work in practice and also bring tangible benefits.

Finally, many works address the problem of low-level task scheduling. Such schedulers are designed to support a specific “data-flow” programming model, but many of their design choices can also be used at a higher level. For example, Tyrex [30] and HFSP [46, 47] are a sample of size-based schedulers, which is a family of policies known to drastically improve turnaround times, as we also have verified with our experiments. Similarly, Quincy [27] and DRF [48] study max-min fair, task-level resource allocation, specifically working on multi-dimensional resources. Although our prototype currently consider a one-dimensional packing problem, due to the characteristics of the back-end we use, which does not yet support CPU-level partitioning, ideas presented in [48] can be extended to our work, considering alternative back-ends supporting multi-resource partitioning. Recently, schedulers supporting complex directed acyclic graphs representing low-level, parallel computations have also appeared: Graphene [49], for example, addresses the problem of complex dependencies and heterogeneous demands among the various stages of the computational graph. The work in [50] indicate substantial improvements in terms of resource utilization (and not only allocation) thanks to worker queues, that independently schedule tasks. Bistro [21] employs a novel hierarchical model of data and computational resources, which enable efficient scheduling of data-parallel workloads. Firmament [51] is a centralized scheduler that has been shown to scale to over ten thousand machines at sub-second task placement latency, using a min-cost max-flow optimization approach. Issues related to scheduling scalability, due to the sheer number of low-level

tasks that are typically required by analytic jobs, have been addressed through a distributed design, such as in Sparrow [28] and in Condor [22]. Although working at the application-level as we do in our work imposes a low toll on the scheduler, distributed designs are interesting also from the failure tolerance point of view, which is why they represent a valid option for our future work.

Part III

Conclusions

13 Conclusions

Efficient resource management of computer clusters has been a long-lasting area of research, with peaks of attention happening in conjunction to improvements in computing machinery, e.g. lately with cloud computing and big data. A new breed of cluster management systems, aiming at becoming “data-center operating systems”, are currently been confronted with problems of efficiency and performance at scale.

Despite recent advances, there exists a gap between the goal of low-level resource management, and that of manipulating high-level, heterogeneous, distributed (analytic) applications running in such cluster environments. In this work we presented a first possible step to fill this gap, in the form of a new application scheduler that interacts with a cluster management back-end, to schedule and allocate resources to applications defined with a simple language and semantics. In addition to careful engineering, required us to design and implement our system we call *Zoe*, our research identified a more fundamental problem, that required to design a novel scheduling heuristic capable of manipulating composite applications, while contributing to system responsiveness.

We validated our algorithm to address our scheduling problem along two lines. We used a numerical approach to simulate large-scale deployments and workloads. We showed our scheduling algorithm to be highly effective in reducing turnaround times, in particular by reducing applications queuing times. Consequently, cluster resources were better allocated. In addition, we reported an overview of the evaluation of *Zoe*, that indicates superior performance and efficiency related to our flexible scheduling heuristic.

Our road-map includes the development of a method to redeem untapped resources from idle but running applications, which calls for a substantial rethinking of the resource reservation paradigm; the design and implementation of application fault tolerance mechanisms; and a long list of pending “tickets” stemming from our open-source *Zoe* project.

References

- [1] "Zoe GitHub." <https://github.com/DistributedSystemsGroup/zoe>.
- [2] M. Ragan-Kelley *et al.*, "The jupyter/ipython architecture: a unified view of computational research, from interactive exploration to communication and publication.," in AGU Fall Meeting Abstracts 2014, vol. 1.
- [3] "Jupyter." <http://jupyter.org/>.
- [4] M. Zaharia *et al.*, "Spark: cluster computing with working sets.," HotCloud, vol. 10, pp. 10–10, 2010.
- [5] "Spark." <http://spark.apache.org/>.
- [6] M. Abadi *et al.*, "Tensorflow: Large-scale machine learning on heterogeneous distributed systems," arXiv preprint arXiv:1603.04467, 2016.
- [7] "TensorFlow." <https://www.tensorflow.org/>.
- [8] "Zero-MQ Guide." <http://zguide.zeromq.org/php:chapter4>.
- [9] M. Dell'Amico, D. Carra, and P. Michiardi, "Psbs: Practical size-based scheduling," IEEE Transactions on Computers, Volume 65, N7, July 2016 Also on ArXiv:1410.6122v2, 2016.
- [10] M. Dell'Amico, D. Carra, M. Pastorelli, and P. Michiardi, "Revisiting size-based scheduling with estimated job sizes," CoRR, vol. abs/1403.5996, 2014.
- [11] M. PASTORELLI, D. Carra, M. Dell'Amico, and P. Michiardi, "Hfsp: Bringing size-based scheduling to hadoop," IEEE Transactions on Cloud Computing, vol. PP, no. 99, pp. 1–1, 2015.
- [12] J. Dean *et al.*, "Mapreduce: simplified data processing on large clusters," Communications of the ACM, vol. 51, no. 1, pp. 107–113, 2008.
- [13] M. Isard *et al.*, "Dryad: distributed data-parallel programs from sequential building blocks," in ACM SIGOPS Operating Systems Review, vol. 41, pp. 59–72, 2007.
- [14] M. Zaharia *et al.*, "Resilient distributed datasets: A fault-tolerant abstraction for in-memory cluster computing," in Proc. of the USENIX NSDI 2012.
- [15] "Flink." <https://flink.apache.org/>.
- [16] D. G. Murray *et al.*, "Naiad: a timely dataflow system," in Proc. of the ACM SOSP 2013.
- [17] X. Meng *et al.*, "Mllib: Machine learning in apache spark," JMLR, vol. 17, no. 34, pp. 1–7, 2016.
- [18] B. Hindman *et al.*, "Mesos: A platform for fine-grained resource sharing in the data center," in Proc. of the USENIX NSDI 2011.
- [19] M. Schwarzkopf *et al.*, "Omega: Flexible, scalable schedulers for large compute clusters," in Proc. of the ACM EuroSys 2013.
- [20] A. Verma *et al.*, "Large-scale cluster management at Google with Borg," in Proc. of the EuroSys 2015.
- [21] A. Goder *et al.*, "Bistro: Scheduling data-parallel jobs against live production systems," in Proc. of the USENIX ATC 2015.
- [22] T. Tannenbaum *et al.*, "Beowulf cluster computing with linux," ch. Condor: A Distributed Job Scheduler, 2002.

- [23] C. Delimitrou et al., "Hcloud: Resource-efficient provisioning in shared cloud systems," in Proc. of the ACM ASPLOS 2016.
- [24] A. Kuzmanovska et al., "Koala-f: A resource manager for scheduling frameworks in clusters," in Proc. of the CCGrid 2016.
- [25] C. Delimitrou et al., "Paragon: Qos-aware scheduling for heterogeneous datacenters," in Proc. of the ACM ASPLOS 2013.
- [26] C. Delimitrou et al., "Quasar: Resource-efficient and qos-aware cluster management," in Proc. of the ACM ASPLOS 2014.
- [27] M. Isard et al., "Quincy: fair scheduling for distributed computing clusters," in Proc. of the ACM SOSP 2009.
- [28] K. Ousterhout et al., "Sparrow: distributed, low latency scheduling," in Proc. of the ACM SOSP 2013.
- [29] C. Delimitrou et al., "Tarcil: Reconciling Scheduling Speed and Quality in Large Shared Clusters," in Proc. of the ACM SOCC 2015.
- [30] B. Ghit et al., "Tyrex: Size-based resource allocation in mapreduce frameworks," in Proc. of the CCGrid 2016.
- [31] V. K. Vavilapalli et al., "Apache hadoop yarn: Yet another resource negotiator," in Proc. of the ACM SoCC 2013.
- [32] "Kubernetes." <http://kubernetes.io/>.
- [33] "Docker Swarm." <https://docs.docker.com/swarm/>.
- [34] C. Reiss et al., "Heterogeneity and dynamicity of clouds at scale: Google trace analysis," in Proc. of the SoCC 2012.
- [35] "Google Public Traces." <https://github.com/google/cluster-data>.
- [36] "Open MPI." <https://www.open-mpi.org/>.
- [37] K. Pruhs et al., "Online scheduling," Handbook of scheduling: algorithms, models, and performance analysis, pp. 15–1, 2004.
- [38] P.-F. Dutot et al., "Scheduling parallel tasks: Approximation algorithms," Handbook of scheduling: Algorithms, models, and performance analysis, pp. 26–1, 2004.
- [39] J. Sgall, "Online preemptive scheduling on parallel machines," 2015.
- [40] "Slurm workload manager." <https://slurm.schedmd.com/>.
- [41] A. Wierman et al., "Nearly insensitive bounds on smart scheduling," in ACM SIGMETRICS Performance Evaluation Review, vol. 33, pp. 205–216, 2005.
- [42] U. Schliegelshohn and R. Yahyapour, "Analysis of first-come-first-serve parallel job scheduling," in SODA, vol. 98, pp. 629–638, Citeseer, 1998.
- [43] K. Cutajar et al., "Practical learning of deep gaussian processes via random fourier features," [arXiv:1610.04386](https://arxiv.org/abs/1610.04386), 2016.
- [44] "Aurora." <http://aurora.apache.org/>.
- [45] "Docker Compose." <https://docs.docker.com/compose/>.

- [46] M. Pastorelli *et al.*, "HFSP: size-based scheduling for hadoop," in Proc. of the IEEE BigData 2013.
- [47] M. Dell'Amico *et al.*, "Revisiting size-based scheduling with estimated job sizes," in Proc. of the IEEE MASCOTS 2014.
- [48] A. Ghodsi *et al.*, "Dominant resource fairness: Fair allocation of multiple resource types," in Proc. of the USENIX NSDI 2011.
- [49] R. Grandl *et al.*, "GRAPHENE: packing and dependency-aware scheduling for data-parallel clusters," in Proc. of the USENIX OSDI 2016.
- [50] J. Rasley *et al.*, "Efficient queue management for cluster scheduling," in Proc. of the ACM EuroSys 2016.
- [51] I. Gog *et al.*, "Firmament: Fast, centralized cluster scheduling at scale," in Proc. of the USENIX OSDI 2016.