

BrownoutCon: A Software System based on Brownout and Containers for Energy-Efficient Cloud Computing

Minxian Xu*, Rajkumar Buyya

Cloud Computing and Distributed Systems (CLOUDS) Laboratory, School of Computing and Information Systems, The University of Melbourne, Australia.

E-mail: minxianx@student.unimelb.edu.au; rbuyya@unimelb.edu.au

Abstract

VM consolidation and Dynamic Voltage Frequency Scaling approaches have been proved to be efficient to reduce energy consumption in cloud data centers. However, the existing approaches cannot function efficiently when the whole data center is overloaded. An approach called brownout has been proposed to solve the limitation, which dynamically deactivates or activates optional microservices or containers. In this paper, we propose a brownout-based software system for container-based clouds to handle overloads and reduce power consumption. We present its design and implementation based on Docker Swarm containers. The proposed system is integrated with existing Docker Swarm without the modification of their configurations. To demonstrate the potential of BrownoutCon software in offering energy-efficient services in brownout situation, we implemented several policies to manage containers and conducted experiments on French Grid'5000 cloud infrastructure. The results show the currently implemented policies in our software system can save about 10% to 40% energy than the existing baselines while ensuring quality of services.

Keywords: Cloud Data Centers, Energy Efficiency, Quality of Service, Containers, Microservices, Brownout

*Corresponding author

1. Introduction

Cloud computing has been regarded as a new paradigm for resource and service provisioning, which has offered vital benefits for IT industry by lowering operational costs and human expenses. However, the huge amount of energy consumption and carbon emissions resulted from cloud data centers have become a significant concern of researchers. Nowadays, data centers contain thousands of servers and their sizes range from 300-4500 square meters, which can consume more than 27,000 kWh energy per day [1]. It is estimated that, in 2010, the energy consumption of data centers consumed 1.1% to 1.5% of total electricity worldwide [1]. Moreover, the excessive usage of brown energy to generate power increases the carbon emission. It is also reported that about 2% carbon emissions of total carbon amount released into the atmosphere worldwide are from data centers [2]. Recently, some dominant service providers have established a community, called Green Grid, to promote energy-efficient techniques to minimize the environmental impact of data centers [3].

Unfortunately, reducing energy consumption is a challenging mission as applications and data are growing complex and consuming more computational resources [4]. The applications and data are generally required to be processed within the required time, and to meet this requirement large and powerful servers are provisioned. To ensure the sustainability of future growth, cloud data centers are required to utilize the resource computing infrastructure efficiently and minimize energy consumption. To address this problem, the concept of green cloud was proposed, which aimed to reduce power consumption, energy cost, carbon emissions and also optimize renewable energy usage [5][6]. Therefore, in addition to resource provisioning and Quality of Service (QoS) assurance, data centers are required to be energy-efficient.

The dominant methods to improve resource utilization and reduce energy consumption are Virtual Machine (VM) consolidation [3] and Dynamic Voltage Frequency Scaling (DVFS) [7]. The VM consolidation method migrates VMs from underutilized hosts to minimize the number of active hosts, and the idle hosts are switched to low-power mode to save energy consumption. The DVFS method reduces energy usage by dynamically scaling voltage frequency. When the host is underutilized, the voltage frequency scales to a lower frequency to reduce power. These approaches have been proved to be efficient to save data center's power consumption, however, when the whole

data center is overloaded, both of them cannot function efficiently. For example, the VMs can not be migrated if all the hosts are overloaded.

In data centers, another reason for high energy consumption lies in that computing resources are inefficiently utilized by applications. Thus, applications are currently built with microservice paradigm in order to utilize resources more efficiently. Microservice is referred as a set of self-contained application components. The components encapsulate their logic and expose their functionality via interfaces to enable flexible deployment and replacement. With microservices or components, developers and users can gain technological heterogeneity, resilience, scalability, ease of deployment, organizational alignment, composability and optimization for replaceability [8]. In addition, microservices also brings the benefits of more fine-grained utilization control over the application resource.

To overcome the limitations of VM consolidation and DVFS, as well as improve the utilization of applications, we take advantage of brownout, a paradigm inspired from voltage shutdown to cope with emergency cases, in which the light bulbs emit fewer lights to save power [9]. Brownout is also applied to cloud scenarios, especially for microservices or application components that are allowed to be shortly deactivated to enhance system robustness. In brownout-compliant microservices, a control knob called dimmer is used to show the probability that whether a microservice should be executed or not [10]. When requests are bursting and the system becomes overloaded, the brownout is triggered to temporally degrade the user experience, so that relieving the overloaded situation as well as saving energy consumption.

Microservices can be featured with brownout characteristic. An example of online shopping system with a recommendation engine is introduced in [10]. The recommendation engine enhances the function of the system and increases profits via recommending products to users. However, because the engine is not the necessary component and it requires more resources in comparison to other components, it is not mandatory to keep running all the time, especially under the overloaded situation when requests have a long delay or even not served. Deactivating the engine enables service providers to serve more requests with essential requirements or QoS constraints. Apart from this example, brownout paradigm can also be applied to other systems that allow application components to be deactivated, especially for the container-based system that applications are built with microservice paradigm. With container technology, the microservices can be functionally isolated, thus the deactivation of some microservices will not influence other microservices. In

addition, as microservices are light-weight, they can be deactivated/activated quickly to support the brownout approach.

In this paper, we propose and develop a software system, called **Brownout-Con**, which is inspired by brownout-based approach to deliver energy-efficient resource scheduling. The implementation of BrownoutCon is based on Docker Swarm [11] that provides the management of container cluster. The software system is designed and implemented as an add-on for Docker Swarm, which has no necessity to modify the configurations of Docker Swarm. The system also applies the public APIs of Grid'5000 [12], which is a real testbed that provides power measurement for hosts. The aims of BrownoutCon are twofold: 1) providing an open-source software system based on brownout and Docker Swarm to manage containers; 2) offering an extensible software system for conducting research on reducing energy consumption and handling overloads in cloud data centers.

The BrownoutCon is designed and implemented by following the brownout enabled system model in our previous works [13][14]. Mandatory containers and optional containers are introduced in the system model, which are identified according to whether the containers can be temporarily deactivated or not. The brownout controller is the key part of the system model to manage brownout, which also provides the scheduling policies for containers. The problem of designing the brownout controller splits into several sub-problems:

1. Predicting the future workloads, so that the system can avoid overloads to foster the system robustness.
2. Determining whether a host is overloaded or not, so that the brownout controller will be triggered to relieve the overloads.
3. Deciding when to disable the containers, so that the system can relieve overloads and reduce energy consumption while ensuring QoS constraints.
4. Selecting the containers to be disabled, so that a better trade-off can be achieved between the reduced energy and QoS constraints.
5. Deciding when to turn the hosts on or into the low-power mode, so that the idle hosts can be switched into low-power mode to save power consumption.

Compared with VM consolidation approaches, the software system based on brownout and containers has two advantages: 1) a container can be stopped or restarted in seconds, while VM migration may take minutes. Thus, scheduling with containers is more light-weight and flexible than VMs. 2) the brownout-based approach provides another optional energy-efficient approach apart from VM consolidation and DVFS, which is also available

to be combined with VM consolidation to achieve better energy efficiency, especially for the situation when the whole data center is overloaded.

To evaluate the proposed system in practice, we conduct our experiments on Grid’5000 [12] real testbed. We also evaluate the performance of proposed system with real traces derived from Wikipedia¹ workloads.

The main **contributions** of our work are as follows:

- Proposed an effective system model that enables brownout approach to manage the containers and resources in a fine-grained manner;
- Designed and developed a software system based on Docker Swarm to provide energy-efficient approaches for cloud data centers;
- Experimental evaluations of our proposed software system on French Grid’5000 infrastructure for service providers to deploy microservices in an energy-efficient manner while ensuring QoS constraints.

The rest of the paper is organized as follows. Section 2 discusses the related work, followed by the system design and implementation in Section 3. Brownout-based policies implemented in BrownoutCon are presented in Section 4. In Section 5, we introduce our experiments setup and evaluate the performance of implemented policies under Grid’5000 testbed. Conclusions along with future work are presented in Section 6.

2. Related Work

It is estimated that U.S. data centers will consume 140 billion kWh of electricity annually by the year 2020, which equals to the annual output of about 50 brown power plants that have high carbon emissions [15][16]. To minimize the operational expenses and impacts on the environment, a variety of state-of-the-art works have been conducted to reduce data center energy consumption.

There is a close relationship between resource utilization and energy consumption, as inefficient utilization of resource contributes to more power consumption [17]. Virtualization is an import technique in Clouds and it can improve resource utilization. Therefore, numerous energy-efficient resource

¹See <http://www.wikibench.eu/wiki/2007-10/> for more details.

scheduling approaches based on VM consolidation have been proposed. Consolidating VMs on fewer physical machines and turning the unused machines into the low-power mode reduce the number of active machines. Beloglazov et al. [3] proposed several VM consolidation algorithms to save data center energy consumption. The VM consolidation process has been modeled as a bin-packing problem, where VMs are regarded as items and hosts are regarded as bins. The objective of these VM consolidation algorithms is mapping the VMs to hosts in an energy-efficient manner. Based on the VM consolidation approaches in this work, other works like [18][19][20] have extended them to improve algorithm performance. Zhang et al. [21] proposed VM allocation algorithm based on evolution algorithm to achieve energy efficiency in cloud data centers for reserved services. Experiments under both simulation and realistic environments showed that the proposed approach can effectively reduce energy consumption for a set of reserved VMs. Li et al. [22] developed a Bayesian network-based estimation model for VM consolidation and took nine data center factors into consideration. The proposed approach can reduce energy consumption while ensuring QoS by avoiding inefficient VM migrations.

Another dominant approach to reduce energy consumption is Dynamic Voltage Frequency Scaling (DVFS). The DVFS approaches achieve energy reduction by adjusting frequencies of processors rather than using less active servers in VM consolidation. The DVFS approach investigates a trade-off between energy consumption and computing performance, where processors lower their frequency when they are lightly loaded and utilize full frequency when loads are heavy.

Kim et al. [7] modeled real-time service as real-time VM requests, and proposed several DVFS algorithms to reduce energy consumption for the DVFS-enabled cluster. Arroba et al. [23] proposed an approach combines DVFS and VM consolidation techniques by considering energy consumption and performance degradation together. Teng et al. [24] presented several heuristic algorithms combining DVFS and VM consolidation together for batch-oriented scenarios. Fan et al. [25] presented an online energy management approach by dynamically configuring voltage frequencies to minimize the power consumption for single processor scheduling while ensuring reliability requirement. All these approaches are focusing on developing algorithms for energy efficiency purposes.

Some research taking both energy consumption and QoS into account have been conducted, which is also the consideration of our proposed soft-

ware prototype system. Dou et al. [26] introduced an energy-aware dynamic VM scheduling approach for QoS enhancement in Clouds for big data, which aimed to benefit users with discount prices and reduce the execution time of tasks. Adhikary et al. [27] developed a QoS-aware and energy-aware cloud resource management approach for multimedia applications, and proposed two distributed and localized resource management algorithms based on energy consumption and resource demands.

VM consolidation and DVFS have been proven to be efficient to reduce energy consumption in both theory and practice, however, both of them cannot function well when the whole data center is overloaded. Thus, brownout is applied to handle data center overloads and reduce energy consumption. Klein et al. [10] applied brownout to design more robust applications under the overloaded or unpredicted situation. In our previous work, brownout was applied to save energy consumption in data centers. In [13], we presented the brownout enabled system model and proposed several heuristic policies to find the microservices or application components that should be deactivated for energy saving purpose. The results showed that a trade-off existed between energy consumption and discount, and in [14], we adopted approximate Markov Decision Process to improve the trade-off.

Compared to the existing energy-efficient approaches based on VMs, our software system is based on containers. Container technology is derived from the Linux LXC techniques [28], which provides mechanism to isolate processes on a shared operating system. Compared with VMs, containerization provides a fine-grained control on microservice resource usage and is more light-weight. Kozhimbayev et al. [29] compared several existing container-based technologies for Clouds and evaluated their strength and weakness. They concluded that containers can give almost the same performance of native systems. As the main reason of the energy consumption issue in clouds is due to the inefficient resource usage, and containers can provide a more fine-grained control on resources compared with VMs, we consider to apply container technology for energy efficiency purposes.

Currently, container technology is mostly focused on the orchestration of construction and deployment for containers [30][31], and it has been applied for various purposes, such as scalability [32], high availability [33], high utilization [34], high throughput [35], and QoS [36]. For example, Liu et al. [37] proposed a flexible container-based computing platform for scientific workflow. Barresi et al. [38] introduced MicroCloud, which is a container-based solution for managing cloud resource efficiently. Santos et al. [39] evalu-

Table 1: Comparison of related work

| Approach | Key Technique | | | Management Unit | | | | Optimization Objective | | Focus | |
|-----------------------------|---------------|------------------|----------|-----------------|------|----|-----------|------------------------|---------|------------------|-----------------|
| | DVFS | VM Consolidation | Brownout | Processor | Host | VM | Container | Energy | QoS/SLA | Algorithm Design | Software System |
| Belogolazov et al. [3] [18] | | ✓ | | | ✓ | ✓ | | ✓ | ✓ | ✓ | |
| Chen et al. [19] | | ✓ | | | ✓ | ✓ | | ✓ | ✓ | ✓ | |
| Han et al. [20] | | ✓ | | | ✓ | ✓ | | ✓ | ✓ | ✓ | |
| Zhang et al. [21] | | ✓ | | | ✓ | ✓ | | ✓ | ✓ | ✓ | |
| Li et al. [22] | | ✓ | | | ✓ | ✓ | | ✓ | ✓ | ✓ | |
| Kim et al. [7] | ✓ | | | ✓ | | | | ✓ | ✓ | ✓ | |
| Arroba et al. [23] | ✓ | ✓ | | ✓ | ✓ | ✓ | | ✓ | ✓ | ✓ | |
| Teng et al. [24] | ✓ | ✓ | | ✓ | ✓ | ✓ | | ✓ | ✓ | ✓ | |
| Fan et al. [25] | ✓ | | | ✓ | | | | ✓ | ✓ | ✓ | |
| Dou et al. [26] | | ✓ | | | ✓ | ✓ | | ✓ | ✓ | ✓ | |
| Adhikary et al. [27] | | ✓ | | | ✓ | ✓ | | ✓ | ✓ | ✓ | |
| Klein et al. [10] | | | ✓ | | ✓ | | | | | ✓ | |
| Xu et al. [13] [14] | | ✓ | ✓ | | ✓ | ✓ | ✓ | ✓ | | ✓ | |
| Liu et al. [37] | | | | | | | ✓ | | ✓ | | ✓ |
| Barresi et al. [38] | | | | | | | ✓ | | ✓ | | ✓ |
| Xu et al. [40] | | | ✓ | | ✓ | | ✓ | ✓ | ✓ | ✓ | |
| Belogolazov et al. [41] | | ✓ | | | ✓ | ✓ | | ✓ | ✓ | | ✓ |
| Goiri et al. [42] | | | | | ✓ | | | ✓ | ✓ | | ✓ |
| BrownoutCon | | | ✓ | | ✓ | | ✓ | ✓ | ✓ | | ✓ |

ated the energy consumption of different applications executed in Docker and bare metal. However, the energy efficient scheduling is not considered in these container-based work.

In our previous work [40], we have proposed an approach for managing energy in container-based clouds while focusing on scheduling algorithms design. Whereas, in this paper, we focus on the design and development of a new software system supporting brownout-based energy-efficient management of clouds. Some of the software systems supported energy-efficient resource management in Clouds, including OpenStack Neat [41], Parasol [42] and vGreen [43]. However, none of them support brownout.

Table 1 shows the comparison of related work. To the best of our knowledge, BrownoutCon is the first software system developed to reduce energy consumption with brownout based on containers, which also considers both energy consumption and QoS.

3. System Architecture, Design and Implementation

The purpose of BrownoutCon is to provide a software system based on brownout and containers for energy-efficient cloud data centers. The system takes advantage of public APIs of Docker Swarm and is evaluated under Grid’5000 testbed. The system is designed to be extensible, which means new components can be added without the necessity to modify the original codes or configurations of Docker Swarm and Grid’5000.

Our software system is deployed on Docker Swarm master and worker nodes. Docker Swarm provides a platform for managing container cluster,

monitoring status of swarm master and worker nodes, deploying containers on nodes, collecting resource usage of containers, controlling the lifecycle of containers, sending messages and commands between the master and worker nodes. Docker Swarm needs to be deployed on physical machines or virtual machines. Therefore, we adopt Grid’5000, a real testbed that provides access to ample resources for Docker Swarm deployment. We also take advantage of the Grid’5000 APIs to collect the energy consumption data of the machines. In the following sections, we discuss the system requirements, assumptions, system design and its implementation.

3.1. Requirements and assumptions

The components of the proposed software prototype system are running on the Docker Swarm master and worker nodes. Our current implementation assumes that a single instance of each key components is invoked on the master node, such as components for controlling brownout, monitoring system status, managing deployment policies and managing models. On each worker node, a single instance of a component that collects node information is running. When new nodes are joining the Docker Swarm as worker nodes, the master node is responsible for deploying containers to the nodes.

BrownoutCon saves the energy consumption and handles overloads via temporarily disabling some containers, therefore, we assume that the services in the target system (e.g. e-commerce system) are implemented with microservice paradigm and some services (e.g. recommendation engine service) are not necessary to keep running all the time.

The main optimization objective of our software system is reducing energy consumption, a precise power probe to collect energy usage is required. Container scheduling policies may use the energy usage data to make decisions on controlling containers.

Another requirement is that a manager is needed to control all the hosts to turn them into low-power mode or active. This manager is used by the brownout controller on master node to connect with other worker nodes via the communication protocol. Grid’5000 has provided the APIs to switch the status of hosts, and more details will be introduced in the following sections.

3.2. BrownoutCon architecture

Figure 1 depicts the architecture of BrownoutCon, and the details of the main components are introduced as below:

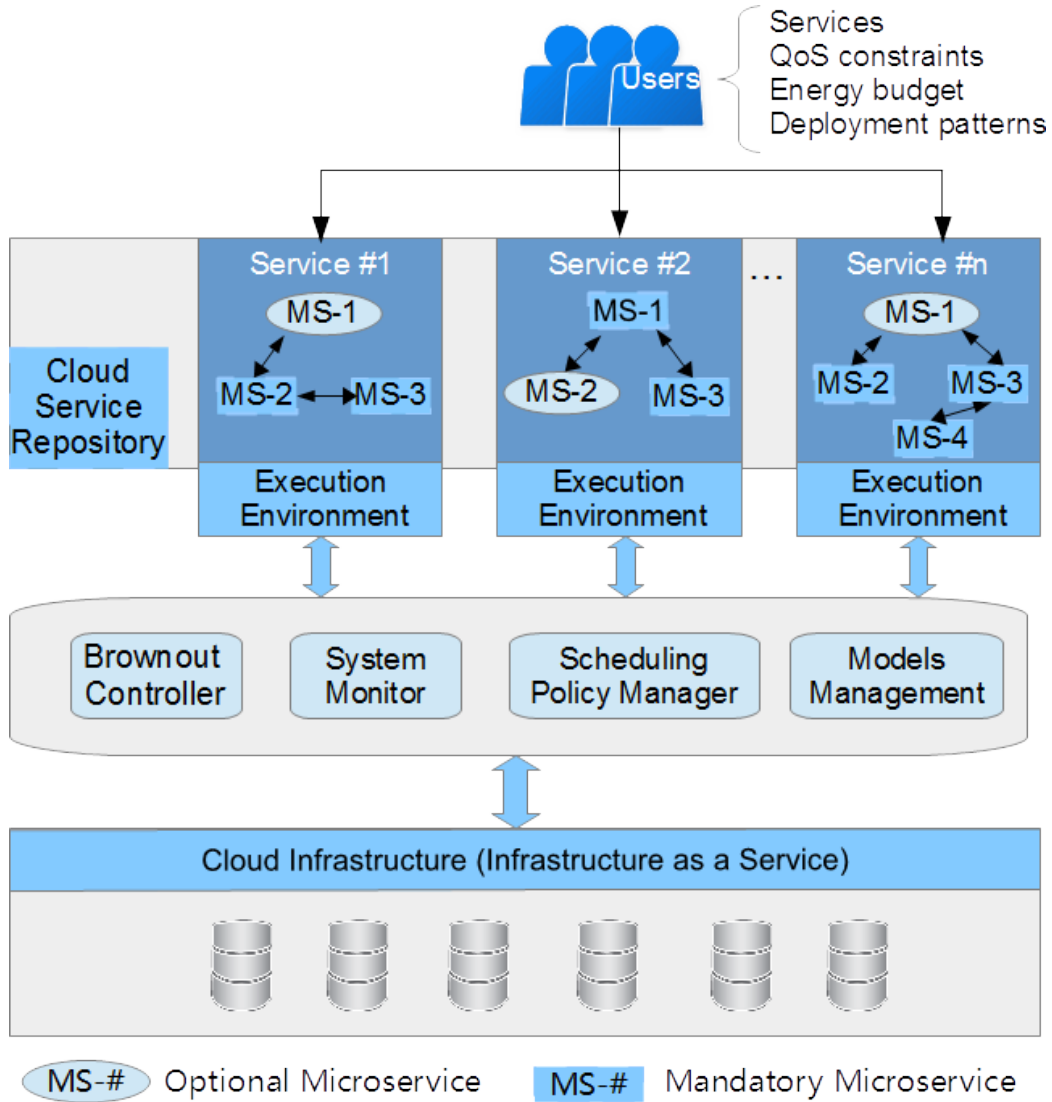


Figure 1: BrownoutCon architecture

- 1) **Users:** This component contains user and requests information. It also captures system configurations such as predefined QoS constraints (e.g. average response time and SLA violations), energy budget and service deployment patterns according to users' demand.
- 2) **Cloud Service Repository:** This component manages the services offered to users, including service information, such as service name and

image. Each service may be constructed via a set of microservices. In order to manage microservices with brownout, the microservices are identified as mandatory or optional.

a. Mandatory Microservices: These microservices keep running all the time when they are started and cannot be temporarily stopped, like database related microservices.

b. Optional Microservices: These microservices can be deactivated temporarily depending on system status. Microservices are connected if there are communications between them. We consider that if one optional microservice is deactivated, then other connected microservices should also be deactivated.

Notes: A microservice can be identified as optional if the service/content it provides is defined as optional by its creators. For instance, the online recommendation engine in the online shopping system and the spell checker in the online editor system can be identified as optional microservices under resource constrained situations.

3) Execution Environment: This component provides the container-based environment for microservices or containers. The dominant execution environments for microservices or containers are Docker, Kubernetes, and Mesos. In BrownoutCon, we use Docker as the execution environment for microservices.

4) Brownout Controller: This component controls optional microservices or containers based on system status. It applies policies introduced in Section 4 to provide an efficient solution for managing brownout and containers. As noted in Section 1, brownout has a control knob called dimmer that represents the probability to execute microservices. We make some adjustments to make the dimmer of brownout to be adapted to this component as well as our architecture. Our dimmer is only applied to the optional microservices and its value is computed according to the severity of system overloads (the number of overloaded hosts in the data center).

5) System Monitor: It is a component that monitors the health of nodes and collects hosts resource consumption status. It uses the third-party toolkit to support its function, such as Grid'5000 public APIs that provide real-time data on infrastructure metrics, including host health, CPU utilization, and power consumption.

6) Scheduling Policy Manager: This component provides and manages the policies for Brownout Controller to schedule microservices/containers. In order to ensure the energy budget and QoS constraints, different

policies designed for different preferences are required. For instance, if the service provider wants to balance the trade-off between energy and QoS, then a policy that considers the trade-off is preferred.

7) Models Management: This component maintains the energy consumption and QoS models in the system. In BrownoutCon, the power consumption model is closely related to the utilization of microservice or container, and the QoS model is applied to define the QoS constraints.

8) Cloud Infrastructure: Under Infrastructure as a Service model, it is a component that offers physical resources to users, where microservices or containers are deployed. In our experiments, we use Grid’5000 as our infrastructure. More details are given in Section 5.

3.3. Energy-efficient scheduling architecture

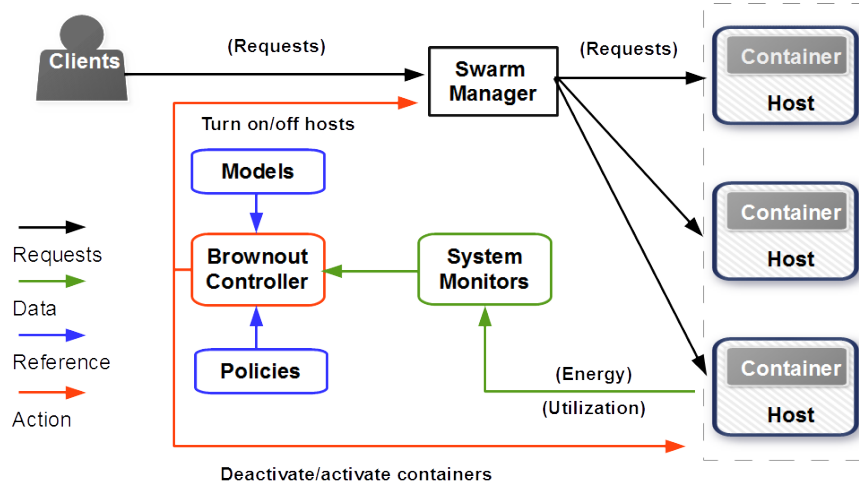


Figure 2: Energy-efficient scheduling architecture

The main purpose of our software system is energy efficiency, and the main approach to achieve this goal is through energy-efficient scheduling policies. Deriving from BrownoutCon architecture, Figure 2 shows the energy-efficient scheduling architecture based on brownout, which depicts the BrownoutCon from the energy-efficient scheduling perspective.

In this scheduling architecture, clients submit their requests to the system, and Docker Swarm Manager dispatches the requests to containers and hosts. The System Monitors collect the energy and utilization information

from hosts, and then send the information to Brownout Controller. With the information from System Monitors, the Brownout Controller refers to the host power consumption or utilization models to compute how much utilization/energy should be reduced. Then the Brownout Controller makes decisions based on scheduling policies to switch the states of hosts and containers, such as turning the hosts into low-power mode or deactivating containers.

3.4. Integration with Docker Swarm

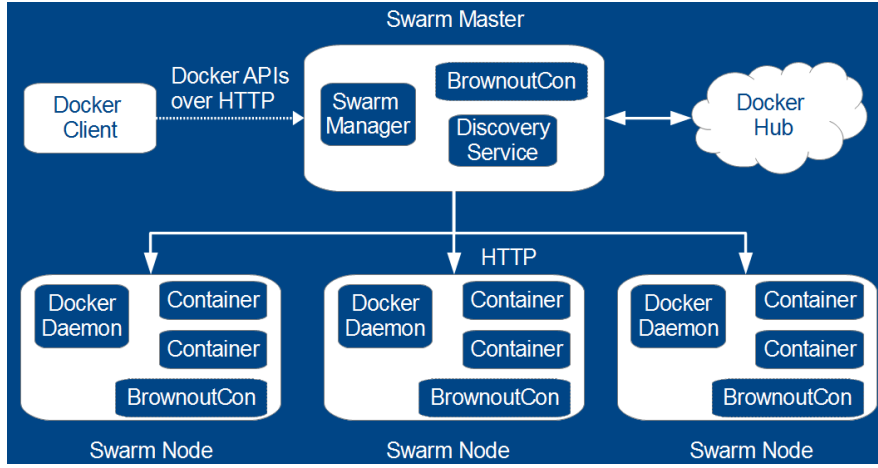


Figure 3: BrownoutCon integrated with Docker Swarm

BrownoutCon is installed on Docker Swarm node independently of Docker Swarm services. In addition, the activities of BrownoutCon are transparent to the Docker Swarm services, which means Docker Swarm does not need to reconfigure to fit with BrownoutCon and use its brownout feature. In other words, BrownoutCon can be installed on existing Docker Swarm cluster without modifying the configurations.

BrownoutCon achieves the transparency via the interactions with the public APIs of Docker Swarm cluster. BrownoutCon uses the APIs to obtain information about containers deployment, containers utilization, and containers properties. Although the operations of BrownoutCon will affect the system status and containers state by deactivating or activating containers, it is transparently processed by Docker Swarm public APIs.

The implication of this integration approach represents that the container deployment is handled by Docker Swarm, and BrownoutCon makes decisions

on deactivation or activation of containers. Figure 3 shows how BrownoutCon is integrated into Docker Swarm. In Docker Swarm, the nodes are categorized as two classes: swarm master node and swarm worker node. The master node is responsible for maintaining cluster state, scheduling services (containers) and serving swarm mode with Docker APIs over HTTP, while the purpose of worker nodes is executing containers. The respective BrownoutCon components are deployed on master and worker nodes.

3.5. Containers deployment with compose file

Docker provides compose² tool to deploy multiple containers, in which a configuration file is used to configure containers properties. With the compose file, the containers can be easily deployed and managed on clusters. In the compose file of our web application, to identify the recommendation engine microservice as optional, we labeled it as optional in the brownout feature. Moreover, as previously mentioned, the optional containers are only allowed to be deployed on the worker node, thus, we configure the placement constraint of this microservice as the worker node. More deployment properties can also be configured in the compose file.

3.6. Entity interaction diagram

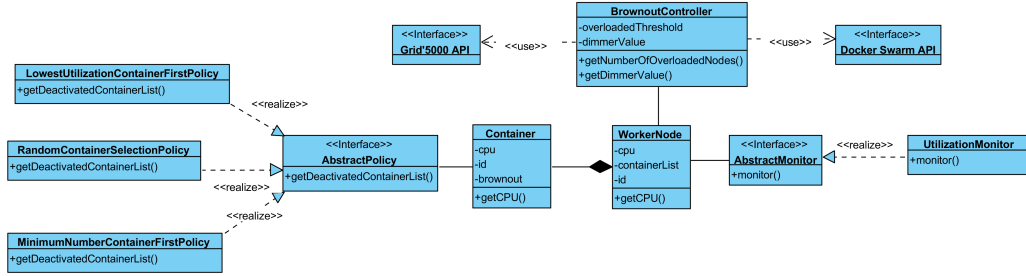


Figure 4: Entity interactions in BrownoutCon

To implement the aforementioned architecture and functionalities, we use Java to develop our software system. The main classes of BrownoutCon are depicted in Figure 4. The details of these classes are as below:

Docker Swarm API: This class wraps Docker Swam APIs and provides the interface for BrownoutController class to call. The Docker Swarm APIs

²See <https://docs.docker.com/compose/compose-file/> for more details.

offer the functions to fetch the information of containers and operate on containers, such as collecting containers utilization, containers id, containers property (optional or mandatory), deploying and updating containers with the compose file, deactivating and activating containers.

Grid’5000 API: This class uses Grid’5000 APIs to collect hosts energy consumption and switch status of hosts. Grid’5000 provides APIs to gather the total power at per second rate for all the hosts in data center. The APIs also allow BrownoutController class to switch the hosts into low power mode or turn the hosts on.

WorkerNode: This class models the host in the data center. Attributes of a WorkerNode include the CPU utilization and the containers deployed on the host. To be consistent with the status of real hosts, when the software system is running, the WorkerNode instances will keep updating their CPU utilization and container lists .

AbstractMonitor: It provides an interface to monitor the status system. With the monitored information, the system can know how many hosts are overloaded and make decisions based on this information. Other monitors, such as memory or network monitors can be extended if they implement the AbstractMonitor.

Container: The Container class models the containers deployed on hosts. The class defines the basic information of containers, including container id, CPU utilization and other information that can be fetched via Docker Swarm APIs.

AbstractPolicy: It is an interface that defines the functions that scheduling policies should implement. To deactivate some containers temporarily and reduce energy consumption, the policies that implement the AbstractPolicy interface are responsible for finding the containers that should be deactivated. The details of our implemented policies in BrownoutCon will be introduced in Section 4.

BrownoutController: This class is the core class of our software system. It assembles all the information from different sources and makes the decision for controlling hosts and the containers on them. BrownoutController knows system status from Docker Swarm APIs, Grid’5000 APIs and WorkerNode instances, and triggers brownout to handle overloads and reduce energy consumption via operations on hosts or containers.

3.7. Sequence diagram

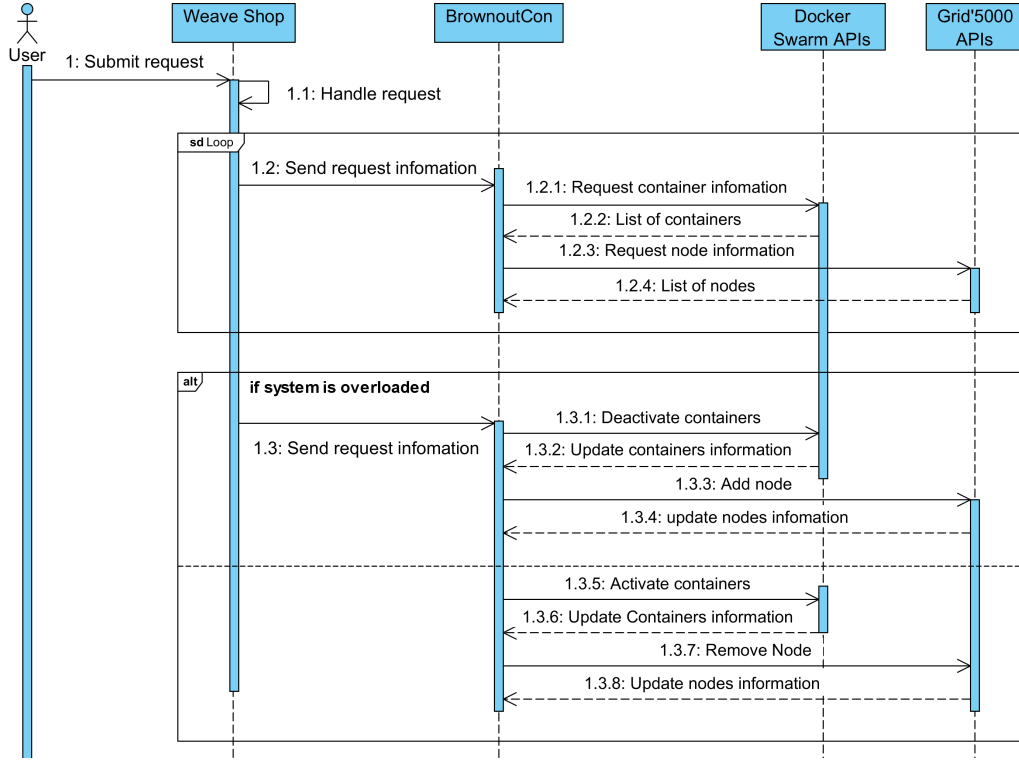


Figure 5: Sequence diagram of of handling requests by BrownoutCon

To provide an in-depth understanding of the working process of Brownout-Con, Figure 5 shows a sequence diagram of handling requests by our software system. Firstly, the users submit their requests to a web application called Weave Shop (more details about this application will be introduced in Section 5.2) Then the Weave Shop sends the information of requests to BrownoutCon, and BrownoutCon keeps collecting nodes and containers information periodically via Grid'5000 and Docker Swam public APIs respectively. When BrownoutCon is collecting information, if the system is overloaded, which means the Weave Shop cannot handle all the incoming requests, the BrownoutCon adds nodes to serve requests. The Brownout-Con also triggers brownout-based policies to deactivate containers to relieve overloads and reduce energy consumption. After these operations, the information of the nodes and containers are updated. Once the system is not overloaded, BrownoutCon activates the containers or removes the nodes from active nodes list (switching nodes into low-power mode). Upon the comple-

tion of operations on containers and nodes, the updated information is sent to BrownoutCon.

4. Policies Implemented in BrownoutCon

To demonstrate BrownoutCon software system capability, we plugged/incorporated some policies originally evaluated by simulations in [13]. As noted in Section 1, the scheduling problem can be divided into several sub-problems: (1) workload prediction; (2) overloaded status detection (3) brownout trigger; (4) deactivated containers selection; and (5) hosts scaling. In this section, we will introduce the implemented policies for reference. It is noted that the introduced policies are not the main focus of this paper. The focus of this work is designing and implementing the software system based on brownout and containers.

4.1. Workload prediction

We apply workload prediction to avoid overloads and improve system robustness. To predict the future workloads based on the previous workloads, we adopt the sliding windows as presented in Algorithm 1. The motivation of sliding windows is giving more weights to the request rates of recent time intervals. Let L_w to be the window size that is a constant integer value, e.g. 5, $num(k)$ to be the actual number of requests at time interval k , we estimate the number of requests at the time interval t as the average number of requests in the previous L_w windows as shown in Equation (1). To ensure enough historical data to be used for prediction, the time interval t for requests prediction should be no less than the window size L_w , for instance, if the window size $L_w = 5$, the time interval for requests predication can start from 5. The predicted number of requests $n\hat{u}m(5)$ at time interval 5 equals to the average number of actual requests of $num(0), num(1), \dots, num(4)$. The sliding window is moving forward along with the time. Based on the predicted workloads, the number of active hosts can be dynamically scaled in and out, which will be introduced in Section 4.5. The performance of Algorithm 1 will be evaluated in Section 5.1.

$$n\hat{u}m(t) = \frac{1}{L_w} \sum_{k=t-L_w}^{t-1} num(k) \quad (1)$$

Algorithm 1 Algorithm for predicting future workload based on sliding windows

Input: sliding window size L_w , the number of requests at previous L_w time intervals, the predicted time interval t ($t \geq L_w$)

Output: the predicted number of requests $\hat{num}(t)$ at time interval t

```

1: for  $k$  from  $t - L_w$  to  $t - 1$  do
2:    $\hat{num}(k + 1) \leftarrow \hat{num}(k) + num(k)$ 
3: end for
4:  $\hat{num}(t) \leftarrow \hat{num}(k + 1)/L_w$ 
5: return  $\hat{num}(t)$ 

```

4.2. Overloaded host detection

In our experiments, we use a predefined overloaded threshold to detect whether a host is overloaded or not. For instance, if the overloaded threshold is defined as 85%, the host is regarded as overloaded when its CPU utilization is above 85%. Currently, we only adopt CPU utilization to detect the overloaded host. Equations (2) and (3) show the way to calculate the number of the overloaded host. We use n_i^o to denote whether host i is overloaded or not, which is detected by the utilization u_i and overloaded threshold T_u . If u_i is no less than T_u , n_i^o equals to 1, otherwise it equals to 0. The total number of the overloaded host is denoted as n_o , which is the sum of n_i^o for all the hosts.

$$n_i^o = \begin{cases} 1, & \text{if } u_i \geq T_u \\ 0, & \text{if } u_i < T_u \end{cases} \quad (2)$$

$$n_o = \sum_{i=0}^{n-1} n_i^o \quad (3)$$

4.3. Brownout trigger

Once there are hosts detected as overloaded, the brownout mechanism will be triggered to handle the overloads as well as to reduce energy consumption. As noted in Section 1, firstly, the algorithm is required to calculate the dimmer value, which is the control knob to represent the probability to trigger brownout on hosts. The dimmer value θ_t at time t is calculated based on the number of overloaded hosts n_o as shown in Equation (4):

$$\theta_t = \sqrt{n_o/n} \quad (4)$$

Then the algorithm computes the expected utilization reduction on overloaded hosts. The expected utilization reduction u_i^r of host i is the product of dimmer value θ_t and the host utilization u_i as:

$$u_i^r = \theta_t \times u_i \quad (5)$$

4.4. Deactivated containers selection

Based on the expected utilization reduction, the policies select containers to deactivate based on different containers selection policies. In Brownout-Con, we have implemented three containers selection policies for deactivation. Based on the strategy design pattern³, these policies implement the AbstractPolicy interface in Figure 4 and can be selected independently at runtime.

4.4.1. Lowest Utilization Container First policy

The Lowest Utilization Container First (LUCF) policy selects a set of containers to reduce the utilization of overloaded hosts. The objective of LUCF is that the utilization after reduction is expected to be less than the overloaded threshold, and the difference between the expected utilization reduction and the sum of deactivated containers utilization is minimized. Thus, the host utilization is reduced and the reduced utilization is close to the expected reduction. The deactivated container list is defined in Equation (6). We use u_i' to denote the utilization of host i after the containers in the deactivated lists are deactivated, which equals to $u_i - u_i^{dcl}$. The utilization of all the containers in dcl_i is denoted as u_i^{dcl} . The $\min(|u_i^r - u_i^{dcl}|)$ represents the to minimize the absolute value of $u_i - u_i^{dcl}$.

$$dcl_i = \begin{cases} \{u_i' \leq T_u, \min(|u_i^r - u_i^{dcl}|)\}, & \text{if } u_i \geq T_u \\ \emptyset, & \text{if } u_i < T_u \end{cases} \quad (6)$$

Algorithm 2 presents the pseudocode of LUCF. The LUCF sorts the optional containers list ocl_i based on container utilization in ascending order so that the container with the lowest utilization is at the head of the list. The size of ocl_i is $ocl_i.size()$. The algorithm checks the hosts one by one, if the first container c_0 on host i has the utilization greater than u_i^r , c_0 is put into

³See https://en.wikipedia.org/wiki/Strategy_pattern for more details.

the deactivated container list dcl_i . Since we consider connected microservices, the policy also adds the container's connection tag T_0^c (a string value) that indicates how it is connected with other containers into a set S for recording connections. However, if the utilization of the first container is less than the expected utilization reduction, LUCF finds a containers sublist to deactivate more containers. The sublist is the one that has the sum of utilization that is closest to the expected utilization reduction than other sublists. Same as previous operations, these containers are put into the deactivated container list dcl_i and their connection tags are put into the set S . Then, the algorithm finds other connected containers and put them into the deactivated container list.

Algorithm 2 Lowest Utilization Container First policy (LUCF)

Input: the number of hosts n in data center, overloaded threshold T_u , deactivated container list dcl_i on host i , the optional container list ocl_i of host i , which is sorted based on utilization of containers u_j^c in ascending order, deactivated tag set S , connection tag T_j^c of container c_j

Output: deactivated container list dcl_i

```

1: for  $i \leftarrow 0$  to  $n - 1$  do
2:   if  $u_i > T_u$  then
3:     if  $u_0^c \geq u_i^r$  then
4:       add  $c_0$  into  $dcl_i$ 
5:       add  $T_0^c$  into  $S$ 
6:     end if
7:     for  $c_j$  in  $ocl_i$  ( $j = 0, 1, 2, \dots, ocl_i.size() - 1$ ) do
8:       if  $(u_j^c \leq u_i^r) \ \& \ (u_i^{dcl} \leq u_i^r)$  then
9:         add  $c_j$  into  $dcl_i$ 
10:        add  $T_j^c$  into  $S$ 
11:        minimize  $|u_i^r - u_i^{dcl}|$ 
12:      end if
13:    end for
14:    for  $c_j$  in  $ocl_i$  ( $j = 0, 1, 2, \dots, ocl_i.size() - 1$ ) do
15:      if  $T_j^c$  in  $S$  then
16:        add  $c_j$  into  $dcl_i$ 
17:      end if
18:    end for
19:  end if
20:  deactivate containers in  $dcl_i$ 
21: end for
22: return  $dcl_i$ 

```

4.4.2. Minimum Number of Containers First policy

As formalized in Equation (7), in order to deactivate fewer containers so that more optional functionalities can be provided, we also implement Minimum Number Containers First (MNCF) policy, which selects the minimum number of containers while saving the power consumption. Since it is quite similar to the LUCF, the pseudocode of MNCF is not provided here. The $\min(dcl_i.size())$ represents the objective to minimize the size of the deactivated container list.

$$dcl_i = \begin{cases} \{u'_i \leq T_u, \min(dcl_i.size())\}, & \text{if } u_i \geq T_u \\ \emptyset, & \text{if } u_i < T_u \end{cases} \quad (7)$$

4.4.3. Random Container Selection policy

Based on a uniformly distributed discrete random variable X that selects a subset of dcl_i randomly, the Random Container Selection (RCS) policy uses uniform distribution function $U(0, ocl_i.size() - 1)$ to randomly select a number of optional containers to reduce energy consumption, as presented in Equation (8).

$$dcl_i = \begin{cases} \{u'_i \leq T_u, X = U(0, ocl_i.size() - 1)\}, & \text{if } u_i \geq T_u \\ \emptyset, & \text{if } u_i < T_u \end{cases} \quad (8)$$

4.5. Hosts scaling

To scale the number of active hosts, we adopt the hosts scaling algorithm in [44] as shown in Algorithm 3, which is a predefined threshold-based approach. With profiling experiments, we set the overloaded requests threshold as the number of requests when the host cannot respond within an acceptable time limit. The algorithm computes the required hosts as the predicted number of request divided by the profiling number of requests of the overloaded threshold. If the required number of hosts is more than current active hosts, more hosts will be added to provide services, otherwise, if current active hosts are adequate, then the excess machines can be set as low-power mode to save energy consumption.

Algorithm 3 Hosts scaling algorithm

Input: number of hosts n in data center, number of active hosts n_a , number of requests when host is overloaded num_{T_u} , predicted number of requests $\hat{num}(t)$ at time t .

Output: number of active hosts n_a

```
1:  $n_a \leftarrow \lceil \hat{num}(t) \div num_{T_u} \rceil$ 
2:  $n' \leftarrow n_a - n$ 
3: if  $n' > 0$  then
4:   Add  $n'$  hosts
5: else if  $n' < 0$  then
6:   Remove  $|n'|$  hosts
7: else
8:   no host scaling
9: end if
10: return  $n_a$ 
```

5. Performance Evaluation

In this section, we evaluate our proposed software prototype system by conducting experiments under Grid'5000 infrastructure. The goals of this section include: 1) evaluating the behavior of the software system in an experimental environment, and 2) demonstrating suitability of the proposed system to enable experimental evaluations and scheduling policies in a practical setting.

5.1. Workload traces

To make the experiments reproducible, we use the real trace from Wikipedia requests on October 17, 2017 to replay the workload of Wikipedia users. The trace includes data on requests time stamp and their accessed pages. We filter⁴ the requests based on per second rate and generate the requests rate. The original request rate is around 1,500-3,000 per second. To scale the workload set to fit with our experiments, we use 10% of the original user requests size. Figure 6 shows the requests rate per second during the day. The blue line is the actual trace derived from Wikipedia and the red line is the predicted trace based on the sliding window (sliding window size is 5) as introduced in Section 4. We can observe some anomalies during intervals 600-800, which can be due to the unpredicted network congestion. While

⁴The details about how we filter the raw data are provided at: <https://github.com/Cloudslab/BrownoutCon>

during most time intervals, the variances between actual trace and predicted trace are small. We conduct statistical analysis for the actual and predicted traces with Root Mean Square Error (RMSE) metric, which has been widely used in statistical analysis to verify experimental results and measures the average spread of errors as $RMSE = \sqrt{\frac{1}{S} \sum_{s=1}^S (\hat{I}_s - I_s)^2}$, where S is the size of the trace, \hat{I}_s is the predicted value and I_s is the actual value. If RMSE is small, it means the predicted values are close to the actual values. In our experiments, the $RMSE = 2.67$, which represents the predicted trace is close to the actual one, and the predicted trace can serve a guide for host scaling strategy.




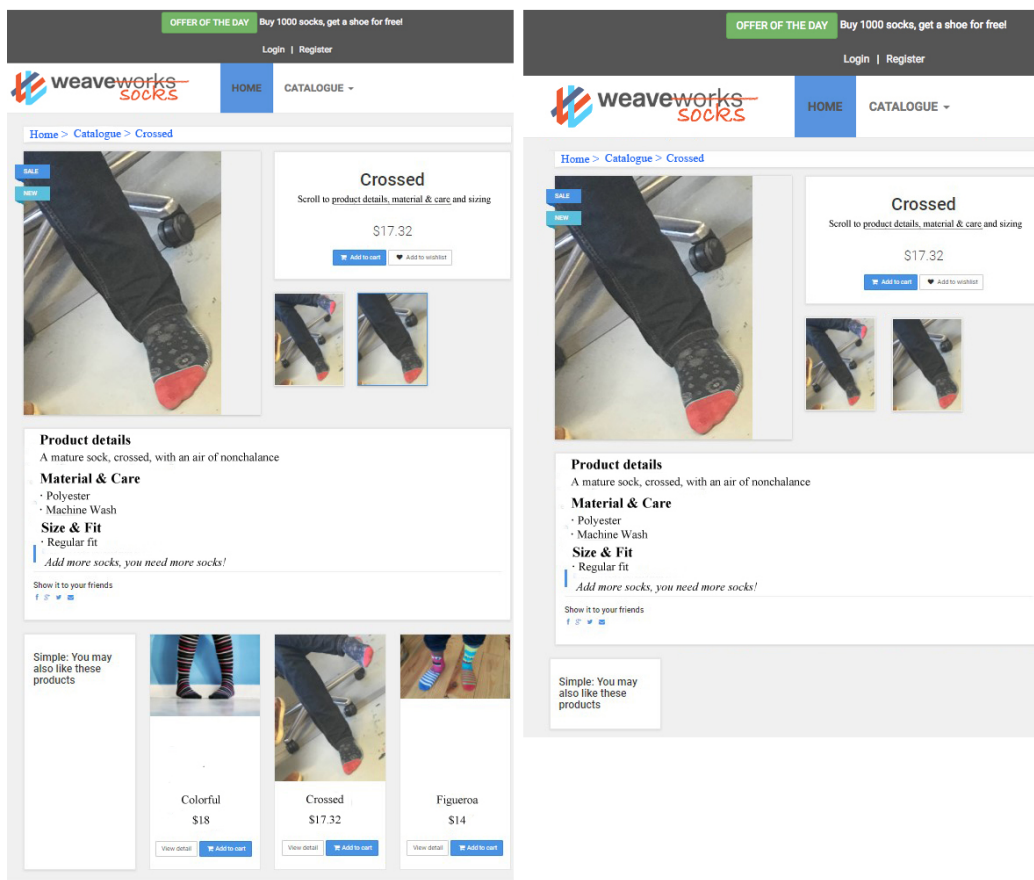
Figure 6: Requests rate of Wikipedia trace

5.2. Application deployment

We use the Weave Shop⁵ web application that implemented with containers as the application in our scenario. The Weave Shop is a shopping

⁵See <https://github.com/microservices-demo/microservices-demo> for more details.

system for selling socks online and has multiple microservices, including user microservice to handle user login, user database microservice for user information storage, payment microservice to process transactions, front-end microservice to show the user interface, catalog microservice for managing item for sale and etc. As these microservices are implemented independently, they can be deployed and controlled without impacting other microservices. The application is deployed by the compose file as introduced in Section 3.5, and part of the microservices are configured as optional, e.g. recommendation engine is noted as optional.



(a) Brownout is not triggered

(b) Brownout is triggered

Figure 7: Impact on user interface when brownout is triggered

The user interface may be influenced due to the deactivation of some microservices. Figure 7 shows the user interface of Weave Shop application. Figure 7(a) is the user interface when full services are provided during no resource saturated scenario, while Figure 7(b) illustrates the user interface when brownout is triggered and the recommendation engine service/-container is deactivated. As a result, other recommended products are not showed in Figure 7(b).

5.3. Performance metrics

We adopt total energy consumption, average response time and SLA violation ratio as our performance metrics, and their definitions are as follows:

Total Energy Consumption: The total energy consumption represents the amount of energy that is consumed by the software system. The key reasons for adopting this metric are: 1) one of the objectives of BrownoutCon is reducing energy consumption, and 2) this metric is widely used in research articles for energy efficient clouds. The total energy consumption $E(t)$ during time interval t is formed as the sum of all the host energy consumption in the data center as shown in Equation (9). Here, we only care about the physical server's energy consumption rather than other network devices or cooling equipment.

$$E(t) = \sum_{i=0}^{n-1} \int_t^{t+1} P_i(t) dt \quad (9)$$

where n is the total number of hosts in the data center, and $P_i(t)$ is the power at time t of host i . And the total energy consumption of data center during observation time period T can be represented as $\sum_{t=0}^T E(t)$. The Grid'5000 cluster provides APIs for the fine-grained power measurement of each physical server at per second rate. We use the APIs to collect the power measurement data (in watts unit) during our observation period and calculate the power consumption (in kWh unit).

Average Response Time: This metric measures the time between users send requests and receive the response on average. We choose this metric because 1) another objective of BrownoutCon is ensuring QoS, and 2) the average response time is a widely used metric of QoS. And the average response time is a widely used metric of QoS. We can use this metric to quantify the overall QoS of the system. We aim to ensure this metric to be below a specific value, e.g. 1 second.

SLA Violation Ratio: As another metric to measure QoS, this metric quantifies the ratio of requests that fail to satisfy the predefined SLA. The reason for choosing this metric is the same as the reason for using average response time as metric. The metric is formalized as:

$$SVR = \frac{num_{err}}{num_a} \quad (10)$$

where num_a is the total number of requests sent to the system, and num_{err} is the number of requests failing to get the response. SVR should also be optimized to be lower than a predefined value, e.g. 1%.

5.4. Experimental testbed

The testbed we used for evaluation is Grid'5000, and we adopt the cluster equipped with power measurement APIs at Lyon site. The hardware specifications are as below:

- 11 × Sun Fire V20z⁶ with AMD Opteron 250 CPU (2 cores, 2.4GHz) and 2GB memory, and all the hosts are running on Debian Linux operating system.

We choose the machines in the same site so that we can reduce the network influence of uncontrolled traffics from other sites. Among the 11 servers, nine are running as the Docker Swarm worker nodes, one is running as the Docker Swarm master node, and another node contains workload trace and installed with JMeter⁷ for sending requests to Docker Swarm cluster. The energy consumption of the workload trace node is not counted, as it is not regarded as a part of the cluster to provide services. All required softwares, such as Docker, Java, Ansible⁸ and JMeter have been installed on these machines to minimize the effects of CPU utilization and network delay.

5.5. Experimental design and setup

In our experiments, the overloaded threshold is configured as 85%, as this value has been evaluated in our previous work [40] that it can achieve a better trade-off between energy consumption and QoS than other values, e.g. below

⁶The maximum power of this model is 237 Watts, and the sleep mode consumes 10 Watts.

⁷See <http://jmeter.apache.org/> for more details.

⁸See <https://www.ansible.com/> for more details.

80% or above 90%, as small overloaded threshold triggers brownout too frequently and large overloaded threshold will not trigger brownout. Another configured parameter is the optional utilization percentage, which represents how much CPU utilization is allowed to be given to the optional containers. According to [13], the change of this parameter has an impact on energy consumption. We vary this parameter with 10%, 20%, and 30% respectively, as the large values, like 40%, can lead to non-negligible performance degradation, and the small values, like 5%, cannot reduce energy consumption effectively. The experiments are conducted by going through the algorithms as below:

1. NPA algorithm [3] — A baseline algorithm that does not consider overloads and optional containers, where the hosts are running all the time and containers are not deactivated.
2. HS [44] — Another baseline algorithm that applies the host scaling algorithm in Algorithm 3, while not applying brownout-based policies.
3. The LUCF, MNCF and RCS algorithms introduced in Section 4 are with the varied optional utilization percentages from 10% to 30% in increments of 10%.

We evaluate the energy consumption, average response time and SLA violation ratio for these algorithms. We run each experiment 5 times to deal with the variance resulted from random factors, such as initial containers deployment, network latency, and application running status.

5.6. Experimental results and analysis

Figure 8 depicts the energy consumption comparison of different algorithms. From the results, NPA has the highest energy consumption with 69.6 kWh, and HS reduces it to 43.9 kWh. For the brownout-based algorithms with varied parameters, the energy consumption of LUCF is from 40.5 kWh to 38.8 kWh when the optional utilization percentage is increased from 10% to 30%. For the MNCF, its energy is close to LUCF, which ranges from 41.1 kWh to 39.2 kWh when optional utilization percentage is varied. As for the RCS, it decreases the energy from 41.4 kWh to 38.8 kWh. All the brownout-based algorithm have shown a significant reduction around 40% to 44% in energy consumption than NPA, and they can also save about 6% to 12% power consumption than HS with different parameter settings.

We also compare the average response time in Figure 9. Although NPA consumes more energy than other algorithms, with the adequate resources, the average response time is the lowest as 174.3 ms. The average response




Figure 8: Energy Consumption Comparison

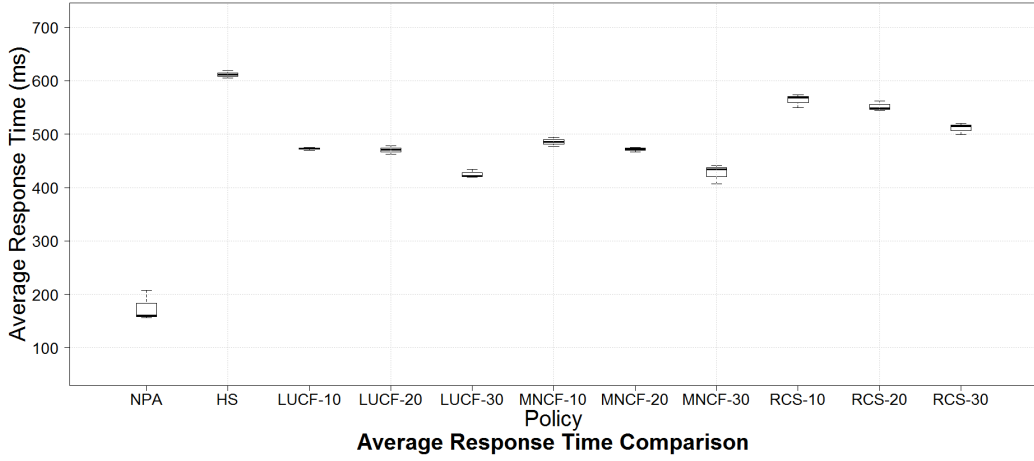


Figure 9: Average Response Time Comparison

time of HS is 611.6ms, while the other brownout-based algorithms decrease this value. LUCF lowers the average response time from 472.6 ms to 425 ms, MNCF reduces it from 485.6 ms and reaches 427.3 ms with 30% optional utilization percentage, and the average response time of RCS ranges from 564.0 ms to 511.3 ms.

In Figure 10, the SLA violation ratios are compared. As NPA has enough resources, it does not experience any SLA violation, while in HS, it has 4.3% SLA violation. Compared with HS, LUCF relieves the SLA violated

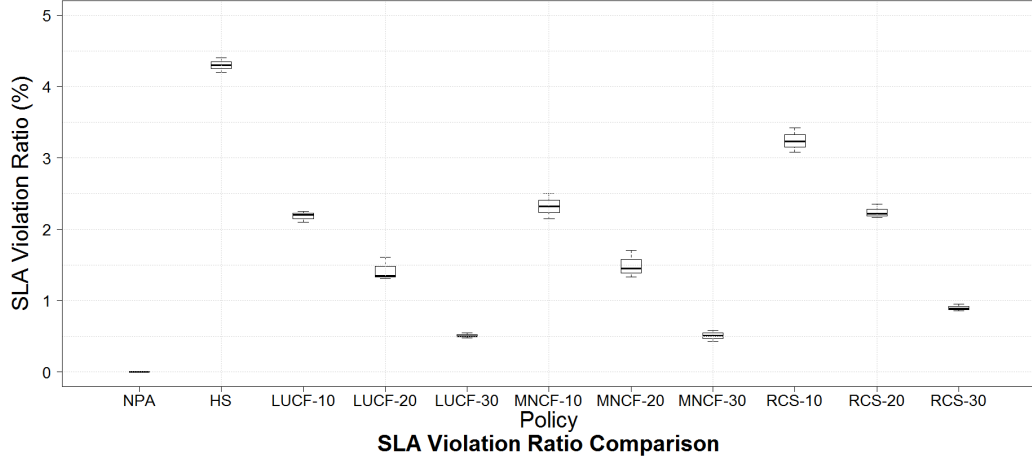


Figure 10: SLA Violation Ratio Comparison

Table 2: The experiment results

| Algorithm | Energy | Avg. Response Time | SLAVR |
|-----------|-----------------|--------------------|--------------|
| NPA | 69.6 kWh | 174.3 ms | - |
| HS | 43.9 kWh | 611.6 ms | 4.3 % |
| LUCF-10 | 40.5 kWh | 472.6 ms | 2.1 % |
| LUCF-20 | 39.5 kWh | 470.3 ms | 1.4 % |
| LUCF-30 | 38.8 kWh | 425.0 ms | 0.5 % |
| MNCF-10 | 41.1 kWh | 485.6 ms | 2.3 % |
| MNCF-20 | 40.4 kWh | 471.3 ms | 1.4 % |
| MNCF-30 | 39.2 kWh | 427.3 ms | 0.5 % |
| RCS-10 | 41.4 kWh | 564.0 ms | 3.2 % |
| RCS-20 | 39.8 kWh | 551.6 ms | 2.2 % |
| RCS-30 | 38.8 kWh | 511.3 ms | 0.9 % |

situation, reducing it from 2.1% to 0.5%. Similar to LUCF, MNCF also decreases the SLA violation to 0.5% from 2.3% when more optional utilization percentage is allowed. As for RCS, its SLA violation drops from 3.2% to 0.9%.

The mean values of obtained results are also displayed in Table 2. HS saves more energy than NPA because it dynamically turns hosts into low-power mode, however, since resources are limited and without brownout, HS also experiences higher average response time and SLA violation ratio. Assuming the average response time and SLA violation ratio of QoS constraints should be below 1 second and 1% respectively, we can conclude that the brownout-based algorithms, LUCF, MNCF and RCS, can save more en-

ergy than NPA and HS while ensuring QoS by reducing average response time and SLA violation ratio. The reason lies in that the brownout-based algorithms reduce energy by deactivating a set of containers and improve the QoS compared with the overloaded situation. And the performance differences of brownout-based algorithms are due to the different selections of deactivated containers. For the comparison of brownout-based algorithms, when more optional utilization percentage is provided, the algorithms perform better. Therefore, in practice, our software system works better if more containers are allowed to be deactivated, which also means that better performance of brownout-based approach can be achieved when more containers are configured as optional.

5.7. Scalability discussion

The design and implementation of BrownoutCon are based on Docker Swarm, therefore, the performance of BrownoutCon is relevant to the scalability of Docker Swarm. Scalability tests on Docker Swarm have been conducted in [45] with 1000 nodes and 30,000 containers, which shows that Docker Swarm is highly scalable system.

6. Conclusions and Future Work

In this paper, we proposed the design and development of a software system based on brownout and containers for energy-efficient clouds, called BrownoutCon. BrownoutCon is transparent system based on Docker Swarm for containers management and does not require to modify the default configurations of Docker Swarm via using its APIs. BrownoutCon can be customized for implementing brownout-based algorithms, which dynamically activates or deactivates containers to handle overloads and reduce energy consumption. The experiments conducted on Grid'5000 infrastructure show that the brownout-based algorithms in BrownoutCon are able to reduce energy consumption while ensuring QoS. The proposed software can be applied in the container-based environment as well as future research in brownout area.

As for future work, we would like to enable BrownoutCon to be available in other container environments, such as Kubernetes and Apache Mesos. We propose to extend the prediction algorithm to handle network congestion. We would also like to investigate memory-intensive workloads to en-

able BrownoutCon to be applied to more generic workloads. In addition, we plan to develop algorithms for integrated management of all resources of cloud data centers including cooling systems [46] to significantly reduce their energy consumption; and implement them in BrownoutCon software system.

Software Availability: We released BrownoutCon as open source, and it can be downloaded from: <https://github.com/Cloudslab/BrownoutCon>.

Acknowledgments

This work is partially supported by China Scholarship Council (CSC) and Australia Research Council (ARC) Discovery Project. We thank Marcos Assuncao from INRIA (France) for providing the access to Grid’5000 infrastructure. We thank Editor-in-Chief (Prof. Paris Avgeriou and Prof. David C. Shepherd), Area Editor (Prof. Helen Karatza), and anonymous reviewers for their excellent comments on improving the paper. We also thank Sukhpal Singh Gill and Shashikant Ilager for their comments on improving this paper.

References

- [1] T. Mastelic, A. Oleksiak, H. Claussen, I. Brandic, J.-M. Pierson, A. V. Vasilakos, Cloud computing: Survey on energy efficiency, *ACM Computing Surveys (CSUR)* 47 (2) (2015) 33.
- [2] B. Lavallée, Data Center Energy: Reducing Your Carbon Footprint — Data Center Knowledge, URL <http://www.datacenterknowledge.com/archives/2014/12/17/undertaking-challenge-reduce-data-center-carbon-footprint>, 2014.
- [3] A. Beloglazov, J. Abawajy, R. Buyya, Energy-aware resource allocation heuristics for efficient management of data centers for cloud computing, *Future Generation Computer Systems* 28 (5) (2012) 755–768.
- [4] Z. Liu, Y. Chen, C. Bash, A. Wierman, D. Gmach, Z. Wang, M. Marwah, C. Hyser, Renewable and cooling aware workload management for sustainable data centers, in: *ACM SIGMETRICS Performance Evaluation Review*, vol. 40, ACM, 175–186, 2012.

- [5] F. Kong, X. Liu, A survey on green-energy-aware power management for datacenters, *ACM Computing Surveys (CSUR)* 47 (2) (2015) 30.
- [6] R. Buyya, S. N. Srirama, G. Casale, R. Calheiros, Y. Simmhan, B. Varghese, E. Gelenbe, B. Javadi, L. M. Vaquero, M. A. S. Netto, A. N. Toosi, M. A. Rodriguez, I. M. Llorente, S. D. C. D. Vimercati, P. Samarati, D. Milojicic, C. Varela, R. Bahsoon, M. D. D. Assuncao, O. Rana, W. Zhou, H. Jin, W. Gentzsch, A. Y. Zomaya, H. Shen, A Manifesto for Future Generation Cloud Computing: Research Directions for the Next Decade, *ACM Comput. Surv.* 51 (5) (2018) 105:1–105:38, URL <http://doi.acm.org/10.1145/3241737>.
- [7] K. H. Kim, A. Beloglazov, R. Buyya, Power-aware provisioning of virtual machines for real-time Cloud services, *Concurrency and Computation: Practice and Experience* 23 (13) (2011) 1491–1505.
- [8] S. Newman, *Building Microservices*, ” O’Reilly Media, Inc.”, 2015.
- [9] M. Xu, R. Buyya, Brownout Approach for Adaptive Management of Resources and Applications in Cloud Computing Systems: A Taxonomy and Future Directions, *ACM Comput. Surv.* 52 (1) (2019) 8:1–8:27, URL <http://doi.acm.org/10.1145/3234151>.
- [10] C. Klein, M. Maggio, K.-E. Årzén, F. Hernández-Rodríguez, Brownout: building more robust cloud applications, in: *Proceedings of the 36th International Conference on Software Engineering*, 700–711, 2014.
- [11] Docker, Docker Documentation — Docker Documentation, URL <https://docs.docker.com/>, 2017.
- [12] Grid5000, Grid5000:Home, URL <https://www.grid5000.fr/mediawiki/index.php/Grid5000:Home>, 2017.
- [13] M. Xu, A. V. Dastjerdi, R. Buyya, Energy Efficient Scheduling of Cloud Application Components with Brownout, *IEEE Transactions on Sustainable Computing* 1 (2) (2016) 40–53.
- [14] M. Xu, R. Buyya, Energy Efficient Scheduling of Application Components via Brownout and Approximate Markov Decision Process, in: *Proceedings of the 15th International Conference on Service-Oriented Computing*, ISBN 978-3-319-69035-3, 206–220, 2017.

- [15] P. Delforge, Data Center Efficiency Assessment - Scaling Up Energy Efficiency Across the Data Center Industry: Evaluating Key Drivers and Barriers, URL <https://www.nrdc.org/sites/default/files/data-center-efficiency-assessment-IP.pdf>, 2014.
- [16] T. Bawden, Global warming: Data centres to consume three times as much energy in next decade, experts warn, URL <http://www.independent.co.uk/environment/global-warming-data-centres-to-consume-three-times-as-much-energy-in-next-decade-experts-warn-a6830086.html>, 2016.
- [17] T. Kaur, I. Chana, Energy efficiency techniques in cloud computing: A survey and taxonomy, *ACM Computing Surveys (CSUR)* 48 (2) (2015) 22.
- [18] A. Beloglazov, R. Buyya, Optimal online deterministic algorithms and adaptive heuristics for energy and performance efficient dynamic consolidation of virtual machines in cloud data centers, *Concurrency and Computation: Practice and Experience* 24 (13) (2012) 1397–1420.
- [19] Q. Chen, J. Chen, B. Zheng, J. Cui, Y. Qian, Utilization-based VM consolidation scheme for power efficiency in cloud data centers, in: *Proceedings of the IEEE International Conference on Communication Workshop*, 1928–1933, 2015.
- [20] Z. Han, H. Tan, G. Chen, R. Wang, Y. Chen, F. C. M. Lau, Dynamic virtual machine management via approximate Markov decision process, in: *Proceedings of the 35th Annual IEEE International Conference on Computer Communications*, 1–9, 2016.
- [21] X. Zhang, T. Wu, M. Chen, T. Wei, J. Zhou, S. Hu, R. Buyya, Energy-aware virtual machine allocation for cloud with resource reservation, *Journal of Systems and Software* 147 (2019) 147–161.
- [22] Z. Li, C. Yan, X. Yu, N. Yu, Bayesian network-based Virtual Machines consolidation method, *Future Generation Computer Systems* 69 (2017) 75–87.
- [23] P. Arroba, J. M. Moya, J. L. Ayala, R. Buyya, DVFS-aware consolidation for energy-efficient clouds, in: *Proceedings of the International*

- Conference on Parallel Architecture and Compilation, IEEE, 494–495, 2015.
- [24] F. Teng, L. Yu, T. Li, D. Deng, F. Magoulès, Energy efficiency of VM consolidation in IaaS clouds, *The Journal of Supercomputing* (2016) 1–28.
 - [25] M. Fan, Q. Han, X. Yang, Energy minimization for on-line real-time scheduling with reliability awareness, *Journal of Systems and Software* 127 (2017) 168–176.
 - [26] W. Dou, X. Xu, S. Meng, X. Zhang, C. Hu, S. Yu, J. Yang, An energy-aware virtual machine scheduling method for service QoS enhancement in clouds over big data, *Concurrency and Computation: Practice and Experience* 29 (14) (2016) 1–20.
 - [27] T. Adhikary, A. K. Das, M. A. Razzaque, M. Alrubaian, M. M. Hassan, A. Alamri, Quality of service aware cloud resource provisioning for social multimedia services and applications, *Multimedia Tools and Applications* 76 (12) (2017) 14485–14509.
 - [28] D. Bernstein, Containers and cloud: From lxc to docker to kubernetes, *IEEE Cloud Computing* 1 (3) (2014) 81–84.
 - [29] Z. Kozhimbayev, R. O. Sinnott, A performance comparison of container-based technologies for the Cloud, *Future Generation Computer Systems* 68 (2017) 175–182.
 - [30] C. Pahl, A. Brogi, J. Soldani, P. Jamshidi, Cloud Container Technologies: a State-of-the-Art Review, *IEEE Transactions on Cloud Computing* (2017) 1–14.
 - [31] M. A. Rodriguez, R. Buyya, Container-based cluster orchestration systems: A taxonomy and future directions, *Software: Practice and Experience* (2018) 1–22.
 - [32] K. Hightower, B. Burns, J. Beda, *Kubernetes: Up and Running: Dive Into the Future of Infrastructure*, " O'Reilly Media, Inc.", 2017.
 - [33] N. Naik, Building a virtual system of systems using Docker Swarm in multiple clouds, in: *2016 IEEE International Symposium on Systems Engineering (ISSE)*, IEEE, 1–3, 2016.

- [34] V. K. Vavilapalli, A. C. Murthy, C. Douglas, S. Agarwal, M. Konar, R. Evans, T. Graves, J. Lowe, H. Shah, S. Seth, et al., Apache hadoop yarn: Yet another resource negotiator, in: Proceedings of the 4th annual Symposium on Cloud Computing, ACM, 5, 2013.
- [35] M. Schwarzkopf, A. Konwinski, M. Abd-El-Malek, J. Wilkes, Omega: flexible, scalable schedulers for large compute clusters .
- [36] E. Boutin, J. Ekanayake, W. Lin, B. Shi, J. Zhou, Z. Qian, M. Wu, L. Zhou, Apollo: scalable and coordinated scheduling for cloud-scale computing, in: 11th {USENIX} Symposium on Operating Systems Design and Implementation ({OSDI} 14), 285–300, 2014.
- [37] K. Liu, K. Aida, S. Yokoyama, Y. Masatani, Flexible Container-Based Computing Platform on Cloud for Scientific Workflows, in: Proceedings of International Conference on Cloud Computing Research and Innovations, IEEE, 56–63, 2016.
- [38] L. Baresi, S. Guinea, G. Quattrocchi, D. A. Tamburri, MicroCloud: A Container-Based Solution for Efficient Resource Management in the Cloud, in: Proceeding of the IEEE International Conference on Smart Cloud, IEEE, 218–223, 2016.
- [39] E. A. Santos, C. McLean, C. Solinas, A. Hindle, How does Docker affect energy consumption? Evaluating workloads in and out of Docker containers, *Journal of Systems and Software* 146 (2018) 14–25.
- [40] M. Xu, A. Nadjaran Toosi, R. Buyya, iBrownout: An Integrated Approach for Managing Energy and Brownout in Container-based Clouds, *IEEE Transactions on Sustainable Computing* (2018) 1–14.
- [41] A. Beloglazov, R. Buyya, OpenStack Neat: a framework for dynamic and energy-efficient consolidation of virtual machines in OpenStack clouds, *Concurrency and Computation: Practice and Experience* 27 (5) (2015) 1310–1333.
- [42] Í. Goiri, W. Katsak, K. Le, T. D. Nguyen, R. Bianchini, Parasol and greenswitch: Managing datacenters powered by renewable energy, in: ACM SIGARCH Computer Architecture News, vol. 41, ACM, 51–64, 2013.

- [43] G. Dhiman, G. Marchetti, T. Rosing, vGreen: a system for energy efficient computing in virtualized environments, in: Proceedings of the 2009 ACM/IEEE international symposium on Low power electronics and design, ACM, 243–248, 2009.
- [44] A. N. Toosi, C. Qu, M. D. de Assunção, R. Buyya, Renewable-aware geographical load balancing of web applications for sustainable data centers, *Journal of Network and Computer Applications* 83 (2017) 155–168.
- [45] A. Luzzardi, Scale Testing Docker Swarm to 30,000 Containers - Docker Blog, URL <https://blog.docker.com/2015/11/scale-testing-docker-swarm-30000-containers/>, 2015.
- [46] S. S. Gill, R. Buyya, A Taxonomy and Future Directions for Sustainable Cloud Computing: 360 Degree View, *ACM Comput. Surv.* 51 (5) (2018) 104:1–104:33, URL <http://doi.acm.org/10.1145/3241038>.