

# Resource-sharing Policy in Multi-tenant Scientific Workflow-as-a-Service Cloud Platform

Muhammad H. Hilman\*, Maria A. Rodriguez, Rajkumar Buyya  
*Cloud Computing and Distributed Systems (CLOUDS) Laboratory  
School of Computing and Information Systems  
The University of Melbourne, Australia*

---

## Abstract

Increased adoption of scientific workflows in the community has urged for the development of platforms that provide workflow executions as a service. As a result, Workflow-as-a-Service (WaaS) concept emerged to address the future design of Workflow Management Systems (WMS). A traditional WMS is usually designed to execute a single workflow in a dedicated process while WaaS cloud platforms enhance the process by exploiting multiple workflows execution in a multi-tenant environment. In this paper, we explore resource-sharing policy to improve system utilization and to fulfil various Quality of Service requirements from multiple users in WaaS cloud platforms. We propose an **Elastic Budget-constrained resource Provisioning and Scheduling** algorithm for **Multiple** workflows that can reduce the computational overhead by encouraging resource-sharing to minimize workflows' makespan while meeting user-defined budget. Our experiments show that the algorithm achieves higher performance in terms of minimizing the makespan compared to the state-of-the-art budget-constraint scheduling algorithm.

---

## 1. Introduction

Scientific workflows have accelerated the triumph of scientific missions on today's multi-discipline sciences [1]. This technology orchestrates and automates scientific applications in a way that reduces the complexity of managing scientific experiments. These workflows are composed of numerous tasks that are interconnected by data or control dependencies. Furthermore, scientific workflows are widely known by their requirements of extensive computational resources. Therefore, these resource-intensive applications are deployed in distributed systems with a high capacity of storage, network, and computing power to achieve a reasonable processing time.

Cloud computing has become a beneficial infrastructure for deploying scientific workflows. Cloud services, especially Infrastructure as a Service (IaaS) offerings, provide a pseudo-infinite pool of resources that can be leased on-demand with a pay-per-use scheme. This pay-as-you-go model substantially eliminates the need for having an upfront investment for massive computational resources. IaaS provides virtualized computational resources in the form of Virtual Machines (VM) with pre-defined CPU, memory, storage, and bandwidth in various pre-configured bundles (i.e., VM types). The users can then elastically acquire and release as many VMs as they need, and the

providers generally charge the resource usage per time slot (i.e., billing period).

Designing algorithms for scheduling scientific workflow executions in clouds is not trivial. The clouds natural features evoke many challenges involving the strategy to decide what type of VMs should be provisioned and when to acquire and release the VMs to get the most efficient scheduling result. The trade-off between having a faster execution time and an economical cost is something that must be carefully considered in leasing a particular cloud instance [2]. Other challenging factors are performance variation of VMs and uncertain overhead delays that might arise from the virtualized backbone technology of clouds, geographical distribution, and multi-tenancy [3].

These problems have attracted many computer scientists into cloud workflow scheduling research to fully utilize the clouds' capabilities for efficient scientific workflows execution [4] [5]. The majority of those studies focus on the scheduling of a single workflow in cloud environments. In this model, they assume a single user utilizes a Workflow Management System (WMS) to execute a particular workflow's job in the cloud. The WMS manages the execution of the workflow so that it can be completed within the defined Quality of Service (QoS) requirements. Along with the growing trend of scientific workflows adoption in the community, there is a need for platforms that provide scientific workflows execution as a service.

Workflow-as-a-Service (WaaS) is an evolving idea that offers workflows execution as a service to the scientific com-

---

\*Corresponding author  
Email address: hilmanm@student.unimelb.edu.au  
(Muhammad H. Hilman)

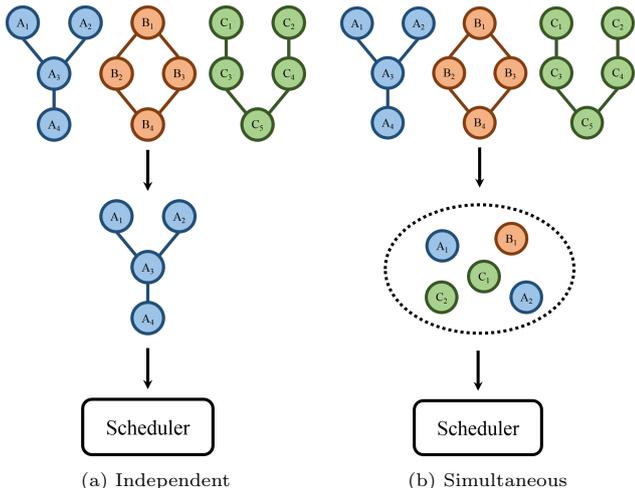


Figure 1: Two approaches on scheduling multiple workflows

munity. The platforms extend WMS technology that is commonly used for handling an individual execution of scientific workflow to serve a more significant number of users through a single point of service. Although some of the traditional cloud workflow scheduling algorithms can be extended for this problem, the specific nature of WaaS environments creates a set of particular needs that should be addressed. As an example, WaaS cloud platforms continuously receive many workflows’ jobs from different users with their unique QoS requirements. The providers must be able to process these requests in a way that each of the requirements is fulfilled. A traditional WMS may process the workflows individually in a dedicated set of VMs as depicted in Figure 1a. This approach, after all, is the simplest way to ensure the QoS fulfillment of each job.

In this dedicated service scenario, a WMS manages different types of tasks’ execution by tailoring their specific software configurations and requirements to a VM image. The VM containing this image then can be quickly deployed whenever a particular workflow is submitted. However, this model cannot easily be implemented in WaaS cloud platforms where many users with different workflow applications are involved. We cannot naively simplify the assumption where every VM image can be shared between multiple users with different requirements. Multiple workflow applications may need different software configurations, which implies a possible dependency conflict if they are fitted within a VM image. This assumption also creates a more complex situation where, at any given time, a new workflow application type needs to be deployed. This newly submitted job cannot reuse the already provisioned VMs as they may not contain its software configurations. Furthermore, dedicating a single workflow execution in a set of VMs is considered inefficient as it leads to the inevitable schedule gaps from inter-dependent workflow tasks that result in the VMs being underutilized.

Therefore, adopting an appropriate resource-sharing policy and at the same time scheduling multiple workflows simultaneously, as shown in Figure 1b, is consider-

ably preferred for multi-tenant WaaS cloud platforms. We argue that introducing this strategy creates a more efficient multi-tenant platform as it reduces the unnecessary overhead during the execution. The efficiency may be gained from sharing the same workflow application software for different users by tailoring a specific configuration in a container image instead of a VM image. This strategy enables (i) sharing and reusing the already provisioned VMs between users to utilize the inevitable scheduling gaps from intra-workflow’s dependency and (ii) sharing local cached images and datasets within a VM that creates a locality, which eliminates the need for network-related activities before the execution.

Based on these described problems and requirements, we propose EBPSM, an **Elastic Budget-constrained resource Provisioning and Scheduling** algorithm for **Multiple** workflows designed for WaaS cloud platforms. A further elaboration of the resource sharing policy that has been introduced in [6] by utilizing container technology to share computational resources (i.e., VMs). This new proposed algorithm focuses more on a budget-constrained scheduling scenario where its budget distribution strategy has been discussed in [7].

The EBPSM algorithm can make a quick decision to schedule the workflow tasks dynamically and empower the sharing of software configuration and reuse of already provisioned VMs between users using container technology. Our algorithm also considers inherent features of clouds that affect multiple workflows scheduling, such as performance variation of VMs [3] and provisioning delay [8] into the VMs auto-scaling policy. Furthermore, EBPSM implements an efficient budget distribution strategy that allows the algorithm to provision the fastest VMs possible to minimize the makespan and adopt the container images and datasets sharing policy to eliminate the need for network transfer between tasks’ execution. Our extensive experiments show that EBPSM which adopts the resource-sharing policy can significantly reduce the overhead, which implies the minimization of workflows’ makespan.

The rest of this paper is organized as follows: Section 2 reviews works that are related to this proposal. Section 3 explains the problem formulation of multiple workflow scheduling in WaaS cloud platforms, including the assumption of application and resource models. The proposed algorithm is described in Section 4 followed by the performance evaluation in Section 5. Finally, the conclusion and future work are elaborated in Section 6.

## 2. Related Work

Scheduling jobs in distributed systems is an extensive research area [9]. In general, the solution for specific jobs that comes from multiple users/sources utilized the online schedulers which implemented the lightweight approaches such as combinatorial search [10] and heuristics algorithm [11]. The similar online approaches are also used for managing the workflows jobs, which avoid the high computa-

tional complexity from the overhead and performance variation in dynamic environments. In this case, the problem of scheduling multiple workflows is facing both the challenges from the unique structure of the jobs and the opportunities to further utilize the systems which involves the resource-sharing of multiple users.

The majority of works in multiple workflows scheduling have pointed out the necessity of reusing already provisioned VMs to reduce the idle gaps and increase system utilization. Examples include the CWSA [12] that uses a depth-first search technique to find potential schedule gaps between tasks' execution. Another work is the CERSA [13] that dynamically adjusts the VM allocation for tasks in a reactive fashion whenever a new workflow job is submitted to the system. These works' idea to fill the schedule gaps between tasks' execution of a workflow to be utilized for scheduling tasks from another workflow is similar to our proposal. However, they merely assume that different workflow applications could be deployed into any existing VM available without considering the possible complexity of software dependency conflicts. Our work differs in the way that we model the software configurations into a container image before deploying it to the VMs for execution.

The use of the container for deploying scientific workflows has been intensively researched. Examples include the work by Qasha et al. [14] that deployed a TOSCA-based workflow<sup>1</sup> using Docker<sup>2</sup> container on e-Science Central platform<sup>3</sup>. Although their work is done on a single VM, the result shows a promising future scientific workflows reproducibility using container technology. A similar result is presented by Liu et al. [15] that convinces less overhead performance and high flexibility of deploying scientific workflows using Docker containers. Finally, the adCFS [16] is designed to schedule containerized scientific workflows that encourage a CPU-sharing policy using a Markov-chain model to assign the appropriate CPU weight to containers. Those solutions are the early development of containerized scientific workflows execution. Their results show high feasibility to utilize container technology for efficiently bundling software configurations for workflows that are being proposed for WaaS cloud platforms.

One of the challenges of executing scientific workflows in the clouds is related to the data locality. The communication overhead for transferring the data between tasks' execution may take a considerable amount of time that might impact the overall makespan. A work by Stavrinides and Karatza [17] shows that the use of distributed in-memory storage to store the datasets locally for tasks' execution can reduce the communication overhead. Our work is similar regarding the data locality policy to minimize the data transfer. However, we propose the use of cached datasets in VMs local storage to endorse the local-

ity of data. We enhance this policy so that the algorithm can intelligently decide which task to be scheduled in specific VMs that can provide the minimum execution time given the available cached datasets.

Two conflicting QoS requirements in scheduling (e.g., time and cost) have been a significant concern when deploying scientific workflows in clouds. A more relaxed constraint to minimize the trade-off between these two requirements is shown in several works that consider scheduling workflows within the deadline and budget constraints. They do not attempt to optimize one or both of the QoS requirements but instead maximizing the success rate of workflows execution within the constraints. Examples of these works include MW-DBS [18] and MW-HBDCS [19] algorithms. Another similar work is the MQ-PAS [20] algorithm that emphasizes on increasing the providers' profit by exploiting the budget constraint as long as the deadline is not violated. Our work considers the user-defined budget constraint in the scheduling, but it differs in the way that the algorithm aims to optimize the overall makespan of workflows while meeting their budget.

Several works specifically focus on handling the real-time workload of workflows in WaaS cloud platforms. This type of workload raises uncertainty issue as the platforms have no knowledge of the arriving workflows. EDPRS [21] adopts a dynamic scheduling approach using event-driven and periodic rolling strategies to handle the uncertainties in real-time workloads. Another work, called ROSA [22], controls the queued jobs—which increase the uncertainties along with the performance variation of cloud resources—in the WaaS cloud platforms to reduce the waiting time that can prohibit the uncertainties propagation. Both algorithms are designed to schedule multiple workflows dynamically to minimize the operational cost while meeting the deadline. Our algorithm has the same purpose of handling the real-time workload and reducing the effect of uncertainties in WaaS environments. However, we differ in that our scheduling objectives are minimizing the workflows' makespan while meeting the user-defined budget.

The majority of works in workflow scheduling that aim to minimize the makespan, while meeting the budget constraints, adopt a static scheduling approach. This approach finds a near-optimal solution of mapping the tasks to VMs—with various VM types configuration—to get a schedule plan before runtime. Examples of these works include MinMRW-MC [23], HEFT-Budg, and MinMin-Budg [24]. However, this static approach is considered inefficient for WaaS cloud platforms as it increases the waiting time of arriving workflows due to the intensive pre-processing computation time to generate a schedule plan.

On the other hand, several works consider scheduling budget-constrained workflows dynamically driven by the available user-defined budget. In this area, examples include BAT [25] which distributes the budget of a particular workflow to its tasks by trickling down the available budget based on the depth of tasks. Another work, MSLBL [26] algorithm distributes the budget by calculating a pro-

<sup>1</sup><https://github.com/ditrit/workflows>

<sup>2</sup><https://www.docker.com/>

<sup>3</sup><https://www.esciencecentral.org/>

portion of the sub-budget efficiently to reduce the unused budget. However, those solutions are designed for a single cloud workflow scheduling scenario. To the best of our knowledge, none of these type of budget-constrained algorithms, that aim to tackle the problem in multiple workflows—which resembles the problem in WaaS cloud platforms—has been proposed.

### 3. Problem Formulation

In this paper, we consider the resource provisioning and scheduling algorithm to process multiple workflows simultaneously. The workflows that are submitted for execution may have the same workflow application type and require the same datasets, but they are not necessarily related to each other. Specifically, the workflows might belong to different users; therefore, they may have different QoS requirements. We assume that these workflows are being executed in a WaaS cloud platform, and a proposed architecture for this particular system is depicted in Fig. 2.

#### 3.1. Workflow-as-a-Service (WaaS) Cloud Platform

WaaS cloud platforms can be placed either in the Platform as a Service (PaaS) or Software as a Service (SaaS) layer of the cloud service model. WaaS providers utilize distributed computational resources from Infrastructure as a Service (IaaS) services to serve the enormous need for computing power in the execution of scientific workflows. They provide an end-to-end service to scientists from the submission portal that allows users to define their requirements, automate their applications’ installation and configuration based on the available or newly created templates, and stage input/output data from the system. Meanwhile, the rest of the data and resource management, including workflow scheduling and resource provisioning, is transparent to the users. This platform is designed to process multiple workflows from different users as they arrive continuously. Therefore, WaaS cloud platforms must handle a high level of complexity derived from their multi-tenant nature, in contrast to a traditional WMS that is generally deployed to manage a single workflow execution. In this work, we focus on the resource provisioner and task scheduler aspects that are designed to deal with the dynamic workload of workflows arriving into WaaS cloud platforms.

#### 3.2. Application and Resource Model

We consider a workload of workflows that are modelled as DAGs (Directed Acyclic Graphs). Furthermore, a workload  $W = \{w_1, w_2, w_3, \dots, w_n\}$  is composed of multiple workflows, where a workflow  $w$  consists of a number of tasks  $T = \{t_1, t_2, t_3, \dots, t_n\}$  and edges  $E$ . An edge  $e_{ij}(t_i, t_j)$  represents the data dependency between task  $t_i$  and  $t_j$  where  $t_j$ , as a successor (i.e., child task), will only start the execution after  $t_i$ , as predecessor (i.e., parent task), is completed and output data  $d_{t_i}^{out}$  from  $t_i$  is available on a

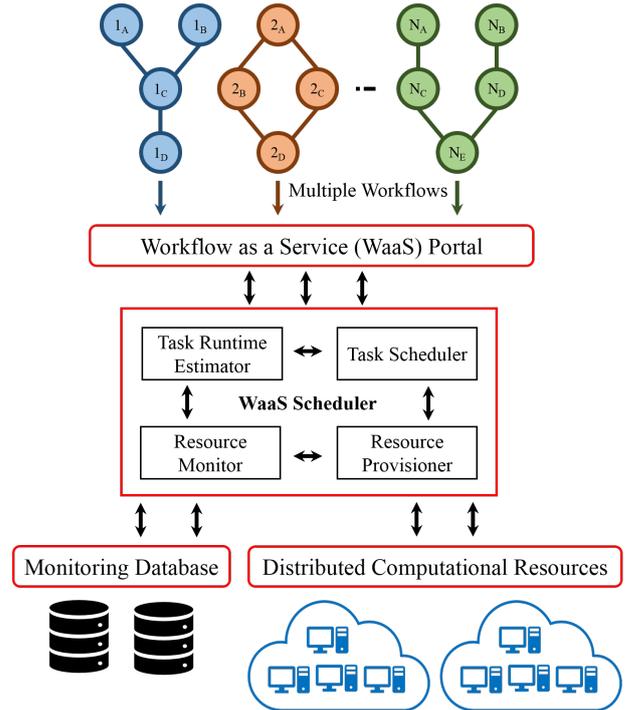


Figure 2: Architecture of Workflow-as-a-Service cloud platform

VM allocated for  $t_j$  as input data  $d_{t_j}^{in}$ . The workflow is completed the execution when the exit task  $t_{exit}$  (i.e., task without successors) is finished being processed. We also assume that each workflow  $w$  is associated with a budget  $\beta_w$  that is defined as a soft constraint of cost representing users’ willingness to pay for the execution of the workflows.

The task  $t$  is executed within a container—that bundles software configurations for a particular workflow in a container image—which is then deployed on VMs. A container provisioning delays  $prov_c$  is acknowledged to download the image, setup, and initiate the container on an active VM. In this work, we set aside the idea of co-locating several running containers within a VM for further study, as the focus of the work is mainly to explore the sharing policy of VMs in terms of its computing, storage, and network capacities. Therefore, we assume that only one container can run on top of a VM at a particular time. Furthermore, the same host VM performance of CPU, memory, and bandwidth, can be achieved in a particular container. Once the container is deployed, VM local storage maintains its image so it can be reused without the need to re-download the containers. We assume WaaS scheduler sends custom signals by using commands in the container (e.g., Docker exec) to trigger tasks’ execution within containers to avoid the necessity of container redeployment.

We consider a pay-as-you-go scheme in IaaS clouds, where VMs are provisioned on-demand and are priced per billing period  $bp$  (i.e., per-second, per-minute, per-hour). Hence, any partial use of the VM is rounded up and charged based on the nearest  $bp$ . In this work, we assumed a fine-grained per-second  $bp$  as it is lately being adopted by the majority of IaaS cloud providers including Ama-

zon EC2<sup>4</sup>, Google Cloud<sup>5</sup>, and Azure<sup>6</sup>. We model a data center within a particular availability zone from a single IaaS cloud provider to reduce the network overhead and eliminate the cost associated with data transfer between zones. Our work considers a heterogeneous cloud environment model where VMs with different VM types  $VMT = \{vmt_1, vmt_2, vmt_3, \dots, vmt_n\}$  which have various processing power  $p_{vmt}$  and different cost per billing period  $c_{vmt}$  can be leased. We consider that all types of VM always have an adequate memory capacity to execute the various type of workflows' tasks. Finally, we can safely assume that the VM type with a higher  $p_{vmt}$  has more expensive  $c_{vmt}$  than the less powerful and slower ones.

Each VM has a bandwidth capacity  $b_{vmt}$  that is relatively the same between different VM types as they come from a single availability zone. We do not restrict the number of VMs to be provisioned during multiple workflows execution, but we also acknowledge the delay in acquiring VMs  $prov_{vmt}$  from the IaaS provider. We assume that the VMs can be eliminated immediately from the WaaS cloud platform without additional delay. Furthermore, we consider performance variation of VMs that might come from a virtualized backbone technology of clouds, geographical distribution, and multi-tenancy [3] and that CPU of VM advertised by IaaS provider is the highest CPU capacity that can be achieved by the VMs. We do not assume another performance degradation of using the containerized environments as Kozhირbayev and Sinnott [27] reported its near-native performance.

A global storage system  $GS$  is modelled for data sharing between tasks (e.g., Amazon S3) with unlimited storage capacity. This global storage has reading rates  $GS_r$  and writing rates  $GS_w$  respectively. In this model, the tasks transfer their outputs from the VMs to the storage and retrieve their inputs from the same place before the execution. Therefore, the network overhead is inevitable and is one of the uncertainties in clouds as the network performance degradation can be observed due to the amount of traffic and the virtualized backbone [28]. To reduce the need for accessing storage  $GS$  for retrieving the data, VMs local storage  $LS_{vmt}$  is also modelled to maintain  $d_t^{in}$  and  $d_t^{out}$  after particular tasks' execution using FIFO policy. It means the earliest stored data will be deleted whenever the capacity of  $LS_{vmt}$  cannot accommodate new data needing to be cached. Furthermore, the time taken to retrieve the input data for a particular task's execution from global storage is shown in Eq. 1.

$$T_{vmt}^{d_t^{in}} = d_t^{in}/GS_r + d_t^{in}/b_{vmt} \quad (1)$$

It is worth noting that there is no need to transfer the input data from the global storage whenever it is available

<sup>4</sup><https://aws.amazon.com/about-aws/whats-new/2017/10/announcing-amazon-ec2-per-second-billing/>

<sup>5</sup><https://cloud.google.com/blog/products/gcp/extending-per-second-billing-in-google>

<sup>6</sup><https://azure.microsoft.com/en-us/pricing/details/virtual-machines/windows/>

in the VM as cached data from previous tasks execution. Similarly, the time needed for storing the output data on the global storage is depicted in Eq. 2.

$$T_{vmt}^{d_t^{out}} = d_t^{out}/GS_w + d_t^{out}/b_{vmt} \quad (2)$$

The runtime  $RT_{vmt}^t$  of a task  $t$  in a VM of type  $vmt$  is assumed as available to the scheduler as part of the scheduling process. The fact is that this runtime can be estimated using various approaches including machine learning techniques [29][30], but we simplify the assumption where it is calculated based on the task's size  $S_t$  in Million of Instruction (MI) and the processing capacity  $p_{vmt}$  of the particular VM type in Million of Instruction Per Second (MIPS) as shown in Eq. 3.

$$RT_{vmt}^t = S_t/p_{vmt} \quad (3)$$

It needs to be noted that this  $RT_{vmt}^t$  value is only an estimate and the scheduler does not depend on it being 100% accurate as it represents one of the uncertainties in clouds. Furthermore, a maximum processing time of a task in a VM type  $PT_{vmt}^t$  consists of reading the input data required from the global storage, executing the task, and writing the output to the storage which are depicted in Eq. 4.

$$PT_{vmt}^t = T_{vmt}^{d_t^{in}} + RT_{vmt}^t + T_{vmt}^{d_t^{out}} \quad (4)$$

From the previous equations, we can calculate the maximum cost  $C_{vmt}^t$  of executing a task  $t$  on a particular  $vmt$  as shown in Eq. 5.

$$C_{vmt}^t = [(prov_{vmt} + prov_c + PT_{vmt}^t)/bp] * c_{vmt} \quad (5)$$

The budget-constrained scheduling problem that is being addressed in this paper is concerned with minimizing the makespan (i.e., actual finish time of  $t_{exit}$ ) of each workflow while meeting the user-defined budget as depicted in Eq. 6.

$$\min AFT(t_{exit}) \quad \text{while} \quad \sum_{n=1}^T C_{vmt}^n \leq \beta_w \quad (6)$$

Intuitively, the budget  $\beta_w$  will be spent efficiently on  $PT_{vmt}^t$  if the overhead components  $prov_{vmt}$  and  $prov_c$  that burden the cost of a task's execution can be discarded. This implies to the minimization of  $AFT(t_{exit})$  as the fastest VMs can be leased based on the available budget  $\beta_w$ . Another important note is that, further minimization can be achieved when the task  $t$  is allocated to the VM with  $d_t^{in}$  available, so the need for  $T_{vmt}^{d_t^{in}}$  which is related to the network factor that becomes one of the sources of uncertainties can be eliminated.

However, it needs to be noted that there still exist some uncertainties in  $RT_{vmt}^t$  as the estimate of  $S_t$  is not entirely accurate, and the performance of VMs depicted in  $PC_{vmt}$  can be degraded at any time. Hence, there must be a control mechanism to ensure that these uncertainties do

---

**Algorithm 1** Budget Distribution

---

```
1: procedure DISTRIBUTE_BUDGET( $\beta, T$ )
2:    $S =$  tasks' estimated execution order
3:   for each task  $t \in T$  do
4:      $allocateLevel(t, l)$ 
5:      $initiateBudget(0, t)$ 
6:   end for
7:   for each level  $l$  do
8:      $T_l =$  set of all tasks in level  $l$ 
9:     sort  $T_l$  based on ascending Earliest Finish Time
10:     $put(T_l, S)$ 
11:  end for
12:  while  $\beta > 0$  do
13:     $t = S.poll$ 
14:     $vmt =$  chosen VM type
15:     $allocateBudget(C_{vmt}^t, t)$ 
16:     $\beta = \beta - C_{vmt}^t$ 
17:  end while
18: end procedure
```

---

not propagate throughout the tasks' execution and cause a violation of  $\beta_w$ . This control can be done by evaluating the real value of a task's execution cost  $C_{vmt}^t$  right after the task is completed. In this way, its successors can adjust their sub-budget allocation so that the total cost will not violate the budget  $\beta_w$ .

#### 4. The EBPSM Algorithm

In this paper, we propose a dynamic heuristic resource provisioning and scheduling algorithm designed for WaaS cloud platforms to minimize the makespan while meeting the budget. The algorithm is developed to efficiently schedule scientific workflows in multi-tenant platforms that deal with dynamic workload heterogeneity, the potential of resource inefficiency, and uncertainties of overheads along with performance variations. Overall, the algorithm enhances the reuse of software configurations, computational resources, and datasets to reduce the overheads that become one of the critical uncertainties in cloud environments. This policy implements a resource-sharing model by utilizing container technology and VMs local storage in the decision-making process to schedule tasks and auto-scale the resources.

When a workflow is submitted to the WaaS portal, its owner may define the software requirements by creating a new container image or choosing an existing template. Whenever the user selects the existing images, the platform will identify the possibly relevant information that is maintained from the previous workflows' execution, including analyzing previous actual runtime execution and its related datasets. Furthermore, the user then defines the budget  $\beta_w$  that is highly likely different from various users submitting the same type of workflow.

Next, the workflow is forwarded to the WaaS scheduler and is preprocessed to assign sub-budget for each task based on the user-defined  $\beta_w$ . This sub-budget along

with the possible sharing of software configurations and datasets will lead the decisions made at runtime to schedule a task onto either an existing VM in the resource pool or a new VM provisioned from the IaaS cloud provider. The first phase of the budget distribution algorithm is to estimate the potential tasks' execution order within a workflow. The entry task(s) in the first level of a workflow are scheduled first, followed by their successors in the next level until it reaches the exit task(s). In this case, we assign every task to a level within a workflow's structure based on the Deadline Top Level (DTL) approach, as seen in Eq. 7

$$level(t) = \begin{cases} 0 & \text{if } pred(t) = \emptyset \\ \max_{p \in pred(t)} level(p) + 1 & \text{otherwise} \end{cases} \quad (7)$$

Furthermore, to determine the tasks' priority within a level, we sort them based on their Earliest Finish Time (EFT) in an ascending order as shown in Eq. 8

$$eft(t) = \begin{cases} PT_{vmt}^t & \text{if } pred(t) = \emptyset \\ \max_{p \in pred(t)} eft(p) + PT_{vmt}^t & \text{otherwise} \end{cases} \quad (8)$$

After estimating the possible tasks' execution order, the algorithm iterates over the tasks based on this order and distributes the budget to each task. This budget distribution algorithm estimates the sub-budget of a task based on the cost  $c_{vmt}$  of particular VM types. At first, the algorithm chooses VMs with the cheapest types for the task. Whenever there is any extra budget left after all tasks get their allocated sub-budgets, the algorithm uses this extra budget to upgrade the sub-budget allocation for a faster VM type starting from the earliest tasks in the order. This approach is called the Slowest First Task-based Budget Distribution (SFTD) [7]. The detailed strategy is depicted in Algorithm 1.

Once a workflow is preprocessed, and its budget is distributed, WaaS scheduler can begin the scheduling process, this step is illustrated in Algorithm 2. The primary objective of this scheduling algorithm is to reuse the VMs that have datasets and containers—with software configurations available—in VMs local storage that may significantly reduce the overhead of retrieving the particular input data and container images from  $GS$ . This way, the algorithm avoids the provisioning of new VMs as much as possible, which reduces the VM provisioning delay and minimizes the network communication overhead from data transfer and downloading container images that contribute to the uncertainties in the WaaS environments.

Furthermore, WaaS scheduler releases all the entry tasks (i.e., tasks with no parents) of multiple workflows into the queue. As the tasks' execution proceeds, the child tasks—whose parents are completed—become ready for execution and are released into the scheduling queue. As a result, at any point in time, the queue contains all the tasks from different workflows submitted to the WaaS cloud platform

---

**Algorithm 2** Scheduling

---

```
1: procedure SCHEDULEQUEUEDTASKS( $q$ )
2:   sort  $q$  by ascending Earliest Start Time (EST)
3:   while  $q$  is not empty do
4:      $t = q.poll$ 
5:      $container = t.container$ 
6:      $vm = null$ 
7:     if there are idle VMs then
8:        $VM_{idle} =$  set of all idle VMs
9:        $VM_{idle}^{input} =$  set of  $vm \in VM_{idle}$  that have
10:         $t$ 's input data
11:        $vm = vm \in VM_{idle}^{input}$  that can finish  $t$  within
12:         $t.budget$  with the fastest time
13:       if  $vm = null$  then
14:          $VM_{idle} = VM_{idle} \setminus VM_{idle}^{input}$ 
15:          $VM_{idle}^{container} =$  set of  $vm \in VM_{idle}$  that have
16:           $container$  deployed
17:          $vm = vm \in VM_{idle}^{container}$  that can finish  $t$ 
18:          within  $t.budget$  with the fastest time
19:         if  $vm = null$  then
20:            $VM_{idle} = VM_{idle} \setminus VM_{idle}^{container}$ 
21:            $vm = vm \in VM_{idle}$  that can finish  $t$ 
22:            within  $t.budget$  with the fastest time
23:         end if
24:       end if
25:     else
26:        $vmt =$  fastest VM type within  $t.budget$ 
27:        $vm = provisionVM(vmt)$ 
28:     end if
29:     if  $vm \neq null$  then
30:       if  $vm.container \neq container$  then
31:          $deployContainer(vm, container)$ 
32:       end if
33:        $scheduleTask(t, vm)$ 
34:     end if
35:   end while
36: end procedure
```

---

that is ready for execution. The queue is periodically being updated whenever one of the two events triggered the scheduling cycle. Those are the arrival of a new workflow's job and the completion of a task's execution.

In every scheduling cycle, each task in the scheduling queue is processed as follows. The first step is to find a set of  $VM_{idle}$  on the system that can finish the task's execution with the fastest time within its budget. The algorithm estimates the execution time by not only calculating  $PT_{vmt}^t$  but also considering possible overhead  $prov_c$  caused by the need for initiating a *container* in case the available  $VM_{idle}$  does not have a suitable *container* deployed.

At first attempt,  $VM_{idle}$  with input datasets  $VM_{idle}^{input}$  are preferred. The  $VM_{idle}^{input}$  that has the datasets available in its local storage must also cache the *container* image from the previous execution. In this way, two uncertain factors  $T_{vmt}^{d_i^n}$  and  $prov_c$  are eliminated. In this case, the sub-budget for this particular task can be spent well on using the fastest VM type to minimize its execution time. This scenario is always preferred since the retrieval

---

**Algorithm 3** Budget Update

---

```
1: procedure UPDATEBUDGET( $T$ )
2:    $t_f =$  completed task
3:    $T_u =$  set of unscheduled  $t \in T$ 
4:    $\beta_u =$  total sum of  $t.budget$ , where  $t \in T_u$ 
5:    $sb =$  spare budget
6:   if  $C_{vmt}^{t_f} \leq (t_f.budget + sb)$  then
7:      $sb = (t_f.budget + sb) - C_{vmt}^{t_f}$ 
8:      $\beta_u = \beta_u + sb$ 
9:   else
10:     $debt = C_{vmt}^{t_f} - (t_f.budget + sb)$ 
11:     $\beta_u = \beta_u - debt$ 
12:   end if
13:   DISTRIBUTE_BUDGET( $\beta_u, T_u$ )
14: end procedure
```

---

---

**Algorithm 4** Resource Provisioning

---

```
1: procedure MANAGERESOURCE
2:    $VM_{idle} =$  all leased VMs that are currently idle
3:    $threshold_{idle} =$  idle time threshold
4:   for each  $vm_{idle} \in VM_{idle}$  do
5:      $t_{idle} =$  idle time of  $vm$ 
6:     if  $t_{idle} \geq threshold_{idle}$  then
7:       terminate  $vm_{idle}$ 
8:     end if
9:   end for
10: end procedure
```

---

of datasets from *GS* and downloading the *container* images through networks has become a well-known overhead that poses a significant uncertainty as its performance also may be degraded over time [28].

If VM has not been found in  $VM_{idle}^{input}$ , the algorithm finds a set of  $VM_{idle}$  that have *container* deployed. This set of  $VM_{idle}^{container}$  may not have the input datasets available as they may have been cleared from VMs local storage. If the set still does not contain the preferred VM, any VM from remaining set of  $VM_{idle}$  is chosen. In the last scenario, the overhead of provisioning delay  $prov_{vm}$  can be eliminated. It is still better than having to acquire a new VM. Whenever a suitable VM in the resource pool is found, the task then is immediately scheduled on it. If reusing an existing VM is not possible, then the algorithm provisions a new VM of the fastest VM type that can finish the task within its sub-budget. This approach is the last option to schedule a task on the platform.

For better adaptation to uncertainties that come from performance variation and unexpected delays during execution, there is a control mechanism within the algorithm to adjust sub-budget allocations whenever a task finishes execution dynamically. This mechanism defines a spare budget variable that stores the residual sub-budget calculated from the actual cost execution. Whenever a task is finished, the algorithm calculates the actual cost of execution using the Eq. 5 and redistributes the leftover sub-budget to the unscheduled tasks. If the actual cost ex-

Table 1: Characteristics of synthetic workflows

Workflow	Parallel Tasks	CPU Hours	I/O Requirements	Peak Memory
CyberShake	Very High	Very High	Very High	Very High
Epigenome	Medium	Low	Medium	Medium
LIGO	Medium High	Medium	High	High
Montage	High	Low	High	Low
SIPHT	Low	Low	Low	Medium

ceeds the allocated sub-budget, the shortfall will be taken from the sub-budget of unscheduled tasks. Therefore, the budget update (i.e., budget redistribution) will take place every time a task is finished. In this way, the uncertainties (e.g., performance variation, overhead delays) that befall to a particular task do not propagate throughout the rest of the following tasks. The details of this process are depicted in Algorithm 3.

Regarding the resource provisioning strategy, the algorithm encourages a minimum number of VMs usage by reusing as much as possible existing VMs. The new VMs are only acquired whenever idle VMs are not available due to the high density of the workload to be processed. In this way, the VM provisioning delays overhead can be reduced. As for the deprovisioning strategy, all of the running VMs are monitored every  $prov_{int}$  and all VMs that have been idle for more than  $threshold_{idle}$  time are terminated as seen in Algorithm 4. The decision to keep or terminate a particular VM should be carefully considered as the cached data within its local storage is one of the valued factors that impact the performance of this algorithm. Therefore, the  $prov_{int}$  and  $threshold_{idle}$  are configurable parameters that can lead to a trade-off between performance in term of resource utilization and makespan along with the VMs local storage caching policy.

## 5. Performance Evaluation

To evaluate our proposal, we used five synthetic workflows derived from the profiling of well-known workflows [31] from various scientific areas generated using the WorkflowGenerator tool<sup>7</sup>. The first workflow is CyberShake, that generates synthetic seismograms to differentiate various earthquakes hazards. This earth-science workflow is data-intensive with very high CPU and memory requirements. The second workflow is Epigenome, a Bioinformatics application with CPU-intensive tasks for executing operations that are related to the genome-sequencing research. The third workflow is an astrophysics application called Inspiral—part of the LIGO project—that is used to analyze the data from gravitational waves detection. This workflow consists of CPU-intensive tasks and requires a high memory capacity. The next workflow is Montage. An astronomy application used to produce a sky mosaics image from several different angle sky observation images. Most of the Montage tasks are considered I/O intensive while involving less CPU processing. Finally, we include

Table 2: Configuration of VM types used in evaluation

Name	vCPU (MIPS)	Storage (GB)	Price per second (cent)
Small	2	20	1
Medium	4	40	2
Large	8	80	4
XLarge	16	160	8

another Bioinformatics application used to encode sRNA genes called SIPHT. Its tasks have relatively low I/O utilization with medium memory requirements. The resume of these characteristics can be seen in Table 1.

The experiments were conducted with various workloads composed of a combination of workflows mentioned above in three different sizes: approximately 50 tasks (small), 100 tasks (medium), and 1000 tasks (large). Each workload contains a different number and various types of workflows that were randomly selected based on a uniform distribution, and the jobs’ arrival was modeled as a Poisson process. Every workflow in a workload was assigned a budget that is always assumed sufficient. Budget insufficiency can be managed by rejecting the job from the platform or re-negotiating the budget with the users. This budget was randomly generated based on a uniform distribution from a range of minimum and maximum cost of executing the workflow. The minimum cost was estimated from simulating the execution of all of its tasks in sequential order on the cheapest VM type. On the other hand, the maximum cost was estimated based on the parallel execution of each task on multiple VMs. In this experiment, we used the runtime generated from the WorkflowGenerator for the size measurement of the task.

We extended CloudSim [32] to support the simulation of WaaS cloud platforms. Using CloudSim, we modeled a single IaaS cloud provider that offers a data center within a single availability zone with four VM types that are shown in Table 2. These four VM types configurations are based on the compute-optimized (c4) instance types offered by the Amazon EC2, where the CPU capacity has a linear relationship with its respective price. We modeled the per-second billing period for leasing the VMs, and for all VM types, we set the provisioning delay to 45 seconds based on the latest study by Ulrich et al. [33]. On the other side, using the model published by Piraghaj et al. [34], the container provisioning delay was set to 10 seconds based on the average container size of 600 MB, a bandwidth 500 Mbps, and a 0.4 seconds delay in container initialization.

The CPU and network performance variation were modeled based on the findings by Leitner and Cito [3]. The CPU performance was degraded by a maximum 24% of its published capacity based on a normal distribution with a 12% mean and a 10% standard deviation. Furthermore, the bandwidth available for a data transfer was potentially degraded by at most 19% based on a normal distribution with a 9.5% mean and a 5% standard deviation. In this experiment, as mentioned earlier, each VM was modeled to maintain an  $LS_{vmt}$  that stores the cached data produced during the execution based on FIFO policy. The design

<sup>7</sup><https://confluence.pegasus.isi.edu/display/pegasus/WorkflowGenerator>

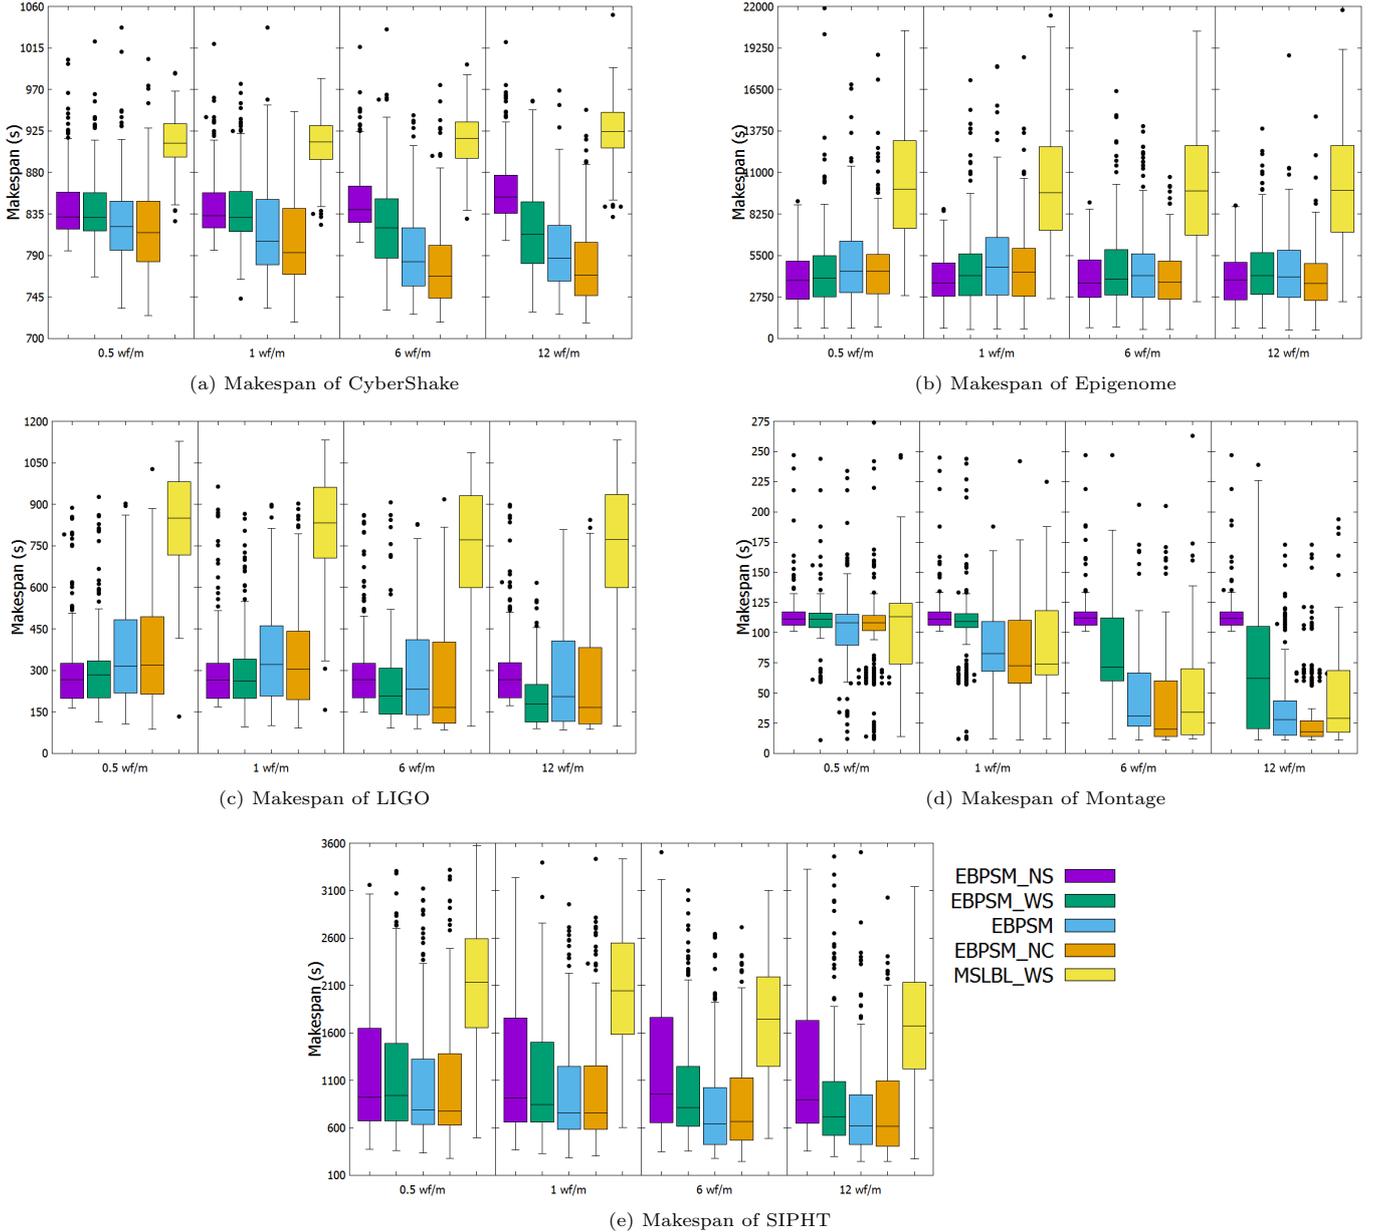


Figure 3: Makespan of workflows on various workloads with different arrival rate

for a more sophisticated strategy to maintain or terminate the lifetime of cached datasets within an  $LS_{vmt}$  is left for future work.

To create a fair baseline for EBPSM, we extended the MSLBL [26] algorithm for multiple workflows scheduling (MSLBL\_MW). MSLBL was designed for single workflow execution, so we added a function to handle multiple workflows by creating a pool of arriving workflows where the algorithm then dispatched the ready tasks from all workflows for scheduling. Furthermore, MSLBL assumed that a set of computational resources are available in a fixed quantity all over the scheduling time. Therefore, to cope up with a dynamic environment in WaaS cloud platforms, we added a simple dynamic provisioner for MSLBL\_MW that provisions a new VM whenever there is no existing VMs

available. This newly provisioned VM is selected based on the fastest VM type that can be afforded by the sub-budget of a particular task. This dynamic provisioner also automatically terminates any idle VMs to ensure the optimal utilization of the system. Finally, for MSLBL\_MW, we assumed that every VM can contain software configurations for every workflow application type and can be shared between any users in WaaS cloud platforms.

To demonstrate the benefits of our resource-sharing policy, we implemented three additional versions of EBPSM, which are EBPSM\_NS, EBPSM\_WS, and EBPSM\_NC. The EBPSM\_NS does not share the VMs between different users; it is a version of EBPSM that executes each workflow submitted into the WaaS cloud platform in dedicated service, as shown in Figure 1a. EBPSM\_WS tailors the

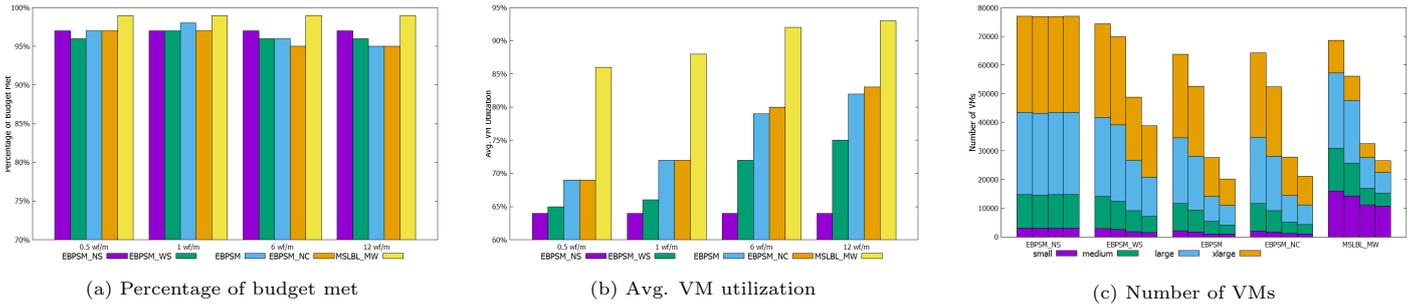


Figure 4: Percentage of budget met and VM usage on various workloads with different arrival rate

software configuration of workflow applications in a VM image instead of containers. Therefore, the algorithm allows only tasks from the same workflow application type (e.g., SIPHT with 50 tasks and SIPHT with 100 tasks) that can share the provisioned VMs during the execution. Meanwhile, EBPSM\_NC ignores the use of containers to store the configuration template and naively assumes that each VM can be shared between many users with different requirements. This version was a direct comparable case for MSLBL\_MW. Finally, the  $threshold_{idle}$  for EBPSM\_WS, EBPSM, and EBPSM\_NC was set to 5 seconds. It means the  $vm_{idle}$  is not immediately terminated whenever it goes idle to accommodate the further utilization of the cached data sets within the VM.

### 5.1. To Share or Not To Share

The purpose of this experiment is to evaluate the effectiveness of our proposed resource-sharing policy regarding its capability to minimize the workflows' makespan while meeting the soft limit budget. In this scenario, we evaluated EBPSM and its variants against MSLBL\_MW under four workloads with the different arrival rate of 0.5, 1, 6, and 12 workflows per minute. Each workload consists of 1000 workflows with approximately 170 thousand tasks of various size (e.g., small, medium, large) and different workflow's type generated randomly based on a uniform distribution. The arrival rate for these four workloads represents the density of workflows' arrival in the WaaS cloud platform. The arrival of 0.5 workflows per minute represents a less occupied platform, while the arrival of 12 workflows per minute models a busier system in handling the workflows.

Figures 3a, 3b, 3c, 3d, and 3e depict the makespan achieved for CyberShake, Epigenome, LIGO, Montage, and SIPHT workflows respectively. EBPSM\_NS that represents the traditional non-shared cloud resources paradigm, shows almost no difference in the algorithm's performance across different arrival rate. This version of the algorithm serves each of the workflows in dedicated and isolated resources. Therefore, it can maintain a similar performance for all four workloads. However, on the other hand, EBPSM\_NS shows the lowest percentage of average VM utilization due to this non-sharing policy, as seen in Figure 4b.

In contrast to EBPSM\_NS, the other three versions (e.g., EBPSM\_WS, EBPSM, EBPSM\_NC) exhibit a performance improvement along with the increasing density of the workloads. This further makespan minimization is the result of (i) the elimination of data transfer overhead between tasks' execution and (ii) the utilization of inevitable scheduling gaps between tasks' execution. In the case of data transfer elimination, we can observe that the improvement is relatively not significant for Epigenome workflows where the CPU processing takes the biggest portion of the execution time instead of I/O and data movement. On the other hand, the most significant improvement can be observed from the data-intensive workflows such as Montage and CyberShake applications. Furthermore, the superiority of EBPSM and EBPSM\_NC over EBPSM\_WS both in makespan and average VM utilization shows a valid argument for the schedule gaps utilization case. From the following result, we can conclude that the utilization of idle gaps between users from different workflow types can further minimize the makespan.

The naive assumption that every VM can be shared between any users in the platform explains the lower makespan and the higher utilization produced by EBPSM\_NC compared to EBPSM. Container usage generates additional initialization delays that affect EBPSM performance. However, the difference between them is marginal, and EBPSM still exhibits its superiority to the other versions.

We observe that in four out of five workflows cases, all versions of EBPSM overthrow MSLBL\_MW regarding the makespan achievement. This result comes from the different strategies of both algorithms in distributing the budget and avoiding the constraint's violation, which implies the type of VMs they provisioned. EBPSM prioritizes the budget allocation to the earlier tasks and leases the fastest VM type as much as possible. This approach is based on the idea that the following children tasks can utilize al-

Table 3: Cost/budget ratio for EBPSM budget violated cases

Percentile	0.5 wf/m	1 wf/m	6 wf/m	12 wf/m
10th	1.005	1.004	1.017	1.005
30th	1.017	1.018	1.032	1.023
50th	1.026	1.030	1.051	1.052
70th	1.046	1.053	1.065	1.069
90th	1.072	1.083	1.121	1.107

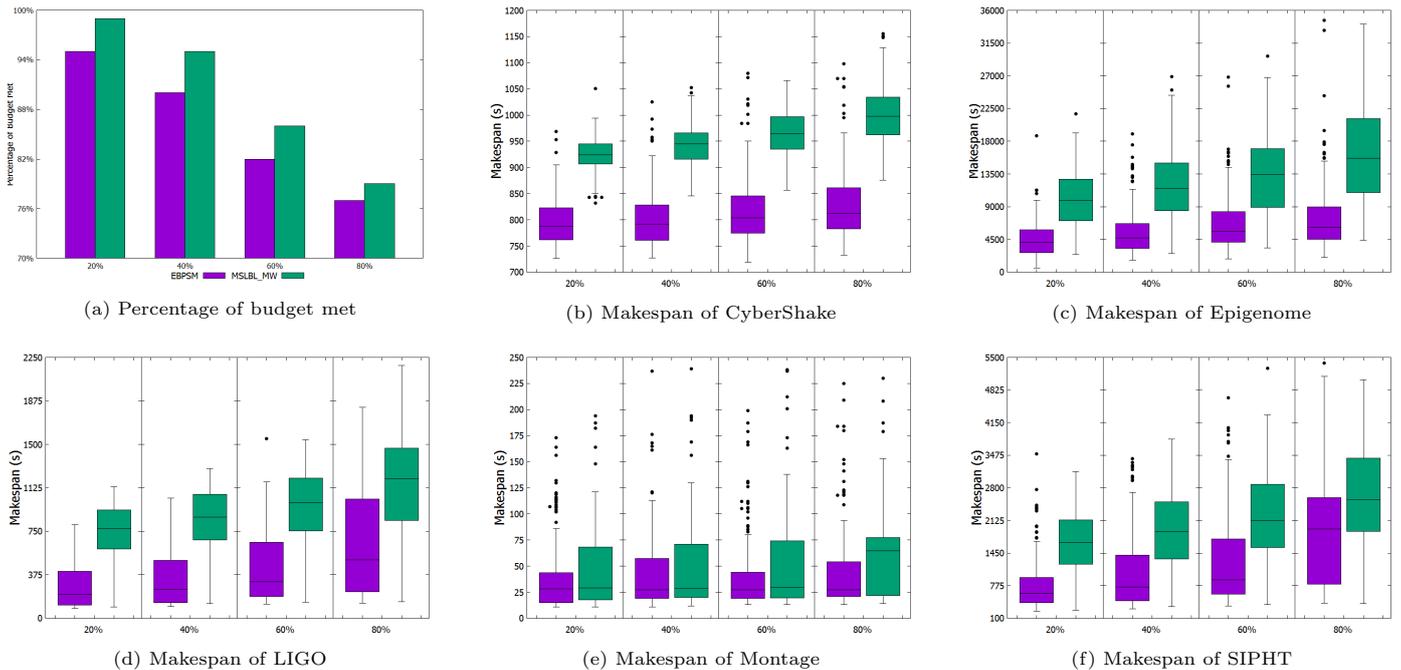


Figure 5: Percentage of budget met and makespan of workflows on various CPU performance degradation

ready provisioned VMs while maintaining the capability of meeting the budget by updating the allocation based on the actual tasks’ execution.

On the other hand, MSLBL\_MW allocates the budget based on the budget level factor, which creates a safety net by provisioning the VM that costs somewhere between the minimum and maximum execution cost of a particular task. In this way, MSLBL\_MW reduces the possibility of budget violation at the budget distribution phase. These two different approaches result in the different number of VM types used during the execution as can be seen from Figure 4c. MSLBL\_MW leases a lower number of faster VM types compared to EBPSM for all cases. The only case where the performance of MSLBL\_MW is relatively equal to EBPSM is in Montage workflow where the tasks are relatively short in CPU processing time, while the significant portion of their execution time takes place in the data movement. In this Montage case, the decision to lease which kind of VM type does not significantly affect the total makespan.

In Figure 4a, we can see that all of the algorithms can achieve at least 95% of the budget meeting for all cases. The margin between MSLBL\_MW and the four versions of EBPSM was never wider than 4%. MSLBL\_MW is superior to EBPSM regarding the average VM utilization. This result is caused by the difference in the VM deprovisioning policy. MSLBL\_MW eliminates any VMs as soon as they become idle, while EBPSM delays the elimination in the hope of further utilization and cached data for the following tasks on that particular idle VM. In this case, the configurable settings of *threshold\_idle* value may affect the VM utilization. However, from these two different approaches, a significant margin of makespan between

MSLBL\_MW and EBPSM can be observed in most cases.

We captured the cases where EBPSM failed to meet the budget and show the result in Table 3 to better understand the EBPSM performance. From the table, we can see that the cost/budget ratio for 90% of the EBPSM budget violation cases are lower than 1.12. This ratio means that the additional cost produced from these violations is never higher than 12%. This percentage is relatively small and may be caused by an extreme case of CPU and network performance degradation. It is not possible to eliminate the negative impact of these uncertainties in such a dynamic environment.

### 5.2. Performance Degradation Sensitivity

Adapting to performance variability is an essential feature for schedulers in multi-tenant dynamic environments. This ability ensures the platform can quickly recover from unexpected events that may occur at any given time, hence preventing a snowball effect that negatively impacts subsequent executions. In this experiment, we evaluate the sensitivity of EBPSM and MSLBL\_MW on the default

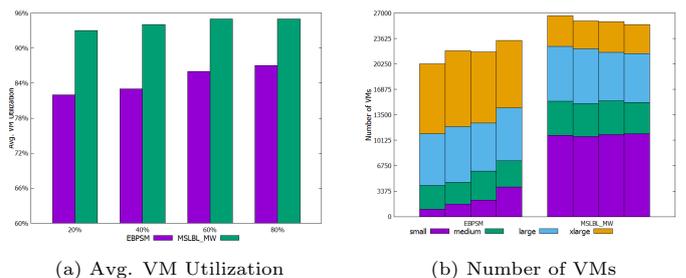


Figure 6: VM usage on various CPU performance degradation

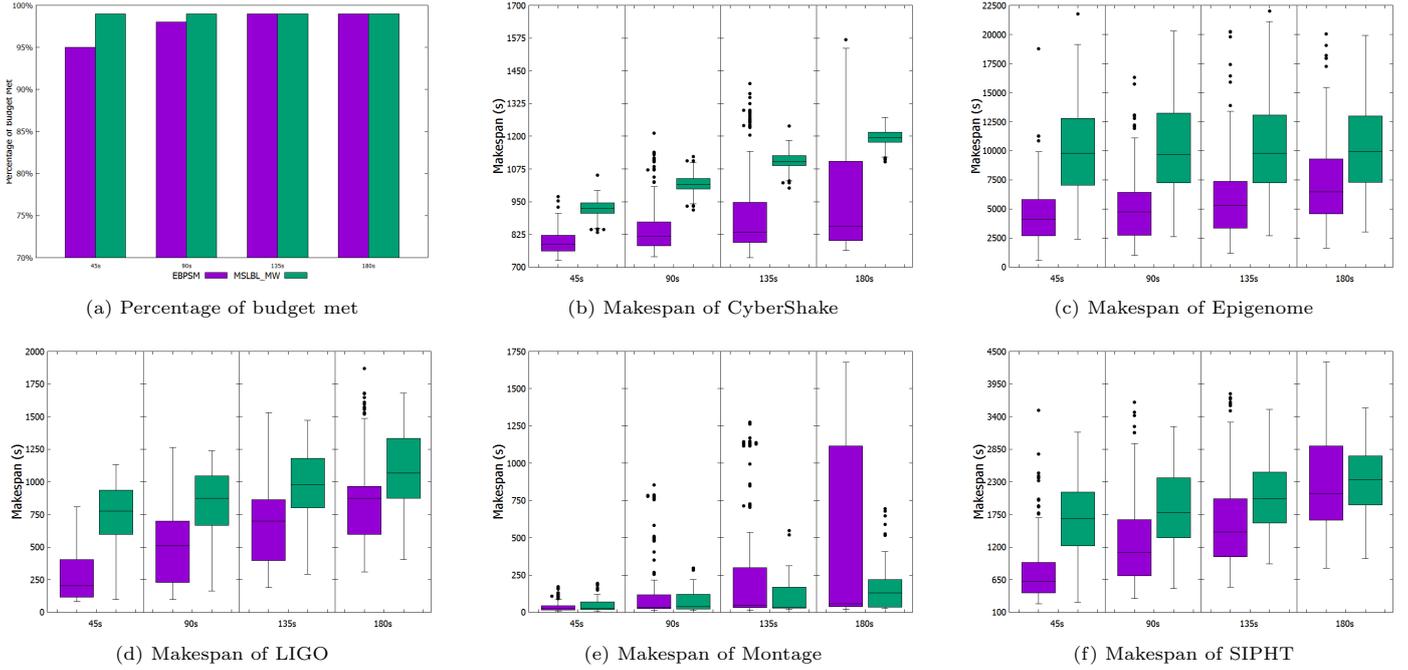


Figure 7: Percentage of budget met and makespan of workflows on various VM provisioning delay

environment—to CPU performance degradation by analyzing the percentage of budgets met, makespan, average VM utilization, and the number of VMs used on four different scenarios. We model CPU performance degradation using a normal distribution with 1% variance and different average and maximum values. The average value is defined as half of the maximum of the CPU performance degradation which ranges from 20% to 80%.

All algorithms are significantly affected by CPU performance degradation as their percentage of budget met decreases along with the increased maximum degradation value. However, MSLBL\_MW suffers the most as its performance margin with EBPSM is getting smaller, as seen in Figure 5a. This suffering also can be observed from Figures 5b, 5c, 5d, 5e, 5f. The increasing makespan as a response to the performance degradation for EBPSM was relatively lower than its effect on MSLBL\_MW. EBPSM can perform better than MSLBL\_MW because of its capability to adapt the changes by evaluating a particular task’s execution right after it was finished.

Meanwhile, MSLBL\_MW only relies on the spare budget from its safety net of budget allocation that limits the number of faster VM types at the budget distribution phase. When the maximum degradation value increases, there is no extra budget left from this safety net. Hence, MSLBL\_MW ability to meet the budget drops faster than EBPSM. This reason is also in line with the average VM utilization results where MSLBL\_MW cannot increase VMs utilization as it reaches the top limit of its capabilities, as seen in Figure 6a.

Another perspective can be observed from the number of VMs used by each algorithm, as depicted in Figure 6b. EBPSM recovers from the CPU performance degradation

by continually increasing the number of slower VM types, while maintaining as much as possible the faster VM types. This adaption is being made through the budget update process after a task finished being executed. Therefore, the earlier tasks are still able to be scheduled onto faster VMs, while the children tasks in the tail of the workflow are inevitably allocated to the slower ones due to the budget left from the recovering process. Different behavior is observed from MSLBL\_MW usage of VMs. We cannot see any significant effort to recover from the number of VM type used. We argue that MSLBL\_MW’s way to deal with CPU performance degradation is by highly relying on the safety net of budget allocation from the budget distribution phase.

### 5.3. VM Provisioning Delay Sensitivity

A large volume of arriving workflows results in a vast number of tasks in the scheduling queue to be processed at a given time. Regardless of the effort put into minimizing the number of VMs used by sharing and re-using already provisioned VMs, provisioning a large number of VMs is inevitable in WaaS cloud platform. This provisioning will

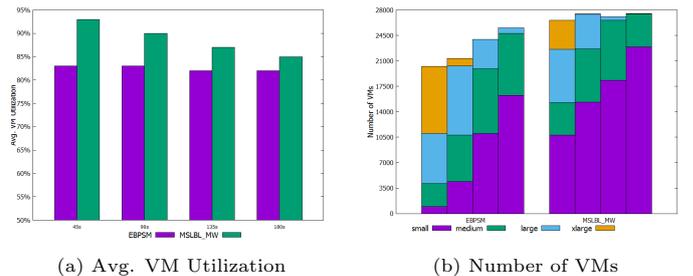


Figure 8: VM usage on various CPU provisioning delay

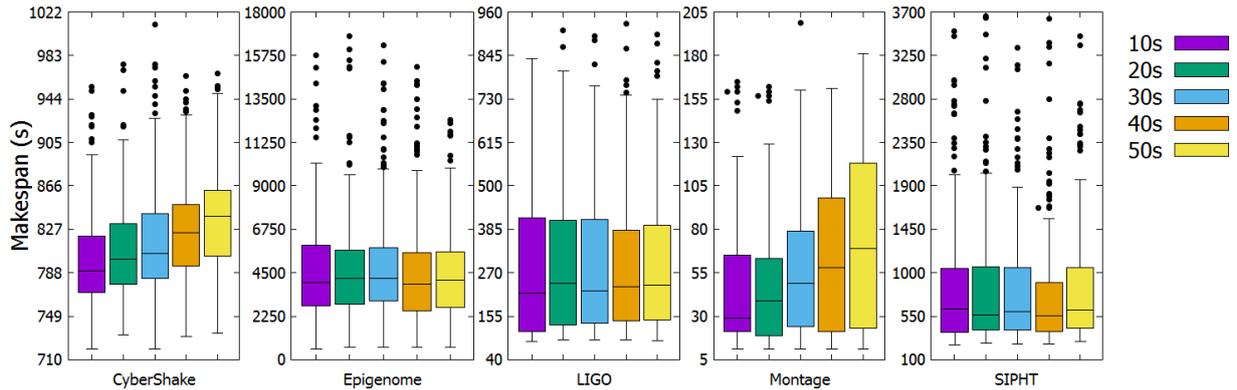


Figure 9: EBPSM performance on various container initiating delay

become a problem if the scheduler is not designed to handle the uncertainties from delays in acquiring the VMs. In this experiment, we study the sensitivity of EBPSM and MSLBL\_MW algorithms—on the default environment—toward VM provisioning delay by analyzing the percentage of budget met, makespan, average VM utilization, and the number of VMs used on four different VM provisioning delay scenarios. The delays range from 45 to 180 seconds.

An interesting point can be observed from Figure 7a that shows the percentage of budget met with the algorithms. All of the algorithms perform well, facing increasing VM provisioning delays. However, it comes with a greater trade-off of the workflows’ makespan as seen in Figures 7b, 7c, 7d 7e, and 7f, especially for the workflow applications that highly depend on CPU processing. We can observe that the spectrum for Montage and CyberShake makespan in EBPSM that results for 180 seconds delays is quite wide although its overall performance is still superior to MSLBL\_MW. In contrast to this situation, the makespan spectrum for MSLBL\_MW for those two workflows is very narrow.

One of the reasons that may cause this variation in the spectrum of EBPSM’s makespan is the variation of the number of shared VMs and the VM type used during the execution. Notably, in the higher VM provisioning delays scenarios. On the other hand, MSLBL\_MW adopts a simple sharing policy. This algorithm does not care whether a particular VM contains a cached data for a future scheduled task or not; it just merely scheduled the task to any idle VM available. In this way, the algorithm can reduce the variation of the VM type used during the execution.

When EBPSM leases faster VMs for the earlier tasks, the VM provisioning delays increase the actual cost for those VMs. Therefore, the algorithm cannot afford those already provisioned VMs if the budget left for the following tasks is not enough to provision the same VM type. In this case, EBPSM must provision a new VM with a slower type. Hence, a workflow with a tight budget will suffer unexpected longer makespan. This condition happens because the algorithm tries to allocate a faster VM type for the earlier tasks but turns out that the budget left is not sufficient. Therefore, to avoid the budget violation, the

algorithm automatically recovers during the budget update phase by provisioning a new slower VM type for that particular task. Figure 8b confirms this situation where EBPSM uses a broader variety of VM types compared to MSLBL\_MW. However, in general, the VM provisioning delays does not directly affect the average VM utilization as the VM initiating process is not counted through the total utilization, but still, these delays are being charged. So that is why the delay profoundly affects the total budget. Finally, this analysis does not hinder the fact that EBPSM, in general, is still superior to MSLBL\_MW in terms of makespan for most of the cases.

#### 5.4. Container Initiating Delay Sensitivity

Various applications consist of a different number of tasks and software dependencies that may impact the size of the container images. Therefore, in this experiment, we study the sensitivity of EBPSM algorithm—on the default environment—toward container initiating delay by analyzing the percentage of budget met and makespan with five different container initiating delays. The delays range from 10 to 50 seconds.

In this experiment, we found that the periodic budget correction of EBPSM can maintain budget-met compliance of 94-95% in all scenarios. Therefore, we do not plot the budget-met percentage and present the impact to the makespan instead. From Figure 9, we can see that the container initiating delay affects the workflows with a high number of tasks and I/O requirements. In this case, CyberShake and Montage makespan increases along with the delays, while the results for the other workflows are relatively indistinguishable.

From these two experiments, we can observe that the container initiating delays may have a higher performance impact to the workflows with a high number of parallel tasks. This condition may be caused by the necessity of deploying a high number of container images to cater to the parallel execution. In this case, a decision to store particular container images of these workflows must be carefully considered. Therefore, a more sophisticated strategy of maintaining container images locality is a crucial aspect to be explored.

### 5.5. Experimental Validation

We have implemented the EBPSM scheduling algorithm in the WaaS cloud platform prototype which extended the functionality of Cloudbus Workflow Engine [35] to support multiple workflows execution. We deployed the platform and utilized the Amazon EC2 cloud resources to carry out similar multiple workflows execution scenarios of bioinformatics workflow applications.

Given the experimental settings, the real system evaluation presented an actual impact of the overhead and performance variation of cloud resources—the results shown the fluctuated makespan and budget spending which confirmed the dynamicity of cloud computing environments. Overall, the performance evaluation of this platform exhibited similar results to the one presented in this work as a simulated environment. The implementation of the EBPSM algorithm in the WaaS cloud platform prototype along with the performance evaluation can be referred to the report by Hilman [36].

## 6. Conclusions and Future Work

The growing popularity of scientific workflows deployment in clouds drives the research on multi-tenant platforms that provides utility-like services for executing workflows. As well as any other multi-tenant platform, this workflow-as-a-service (WaaS) cloud platform faces several challenges that may impede the system effectiveness and efficiency. These challenges include the handling of a continuously arriving workload of workflows, the potential of system inefficiency from inevitable idle time slots from workflows' tasks dependency execution, and the uncertainties from computational resources performances that may impose significant overhead delays.

WaaS cloud platforms extend the capabilities of a traditional Workflow Management System (MWS) to provide a more comprehensive service for larger users. In a traditional WMS, a single workflow job is processed in dedicated service to ensure its Quality of Service (QoS) requirements. However, this approach may not be able to cope up with the WaaS environments. A significant deficiency may arise from its conventional way of tailoring workflows' software configurations into a VM image, intra-dependent workflows' tasks inevitable schedule gaps, and possible overhead delays from workflow pre-processing and data movement handling.

To achieve more efficient multi-tenant WaaS cloud platforms, we proposed a novel resource-sharing policy that utilizes container technology to wrap software configurations required by particular workflows. This strategy enables (i) the idle slots VMs sharing between users and (ii) gaining further makespan minimization by sharing the datasets and container images cached within VMs local storage. We implemented this policy on EBPSM, a dynamic heuristic scheduling algorithms designed for multiple workflows scheduling in WaaS. Our extensive exper-

iments show that the sharing scenarios overthrow a traditional dedicated approach in handling workflows' jobs. Furthermore, our proposed algorithm can surpass a modified state-of-the-art budget-constrained dynamic scheduling algorithm in terms of minimizing the makespan of workflows and meeting the soft limit budget defined by the users.

There are several aspects of our experiments that need to be further investigated. First, the budget distribution phase takes a vital role in budget-constrained scheduling. The decision to either allocate more budget for the earlier tasks so they can lease faster computational resources or maintain a safety net allocation to ensure the budget compliance must be carefully taken into account. A trade-off between having a faster execution time and meeting the allocated budget is always inevitable. In this way, defining the nature of execution, including strictness of the budget constraints can help to design an appropriate configuration between two approaches.

Second, the resource provisioning (and deprovisioning) strategy must consider the *quid pro quo* between having a higher system utilization (i.e., lower idle VM times) and an optimal data sharing and movement which utilizes the VM local storage. We observed that delaying a particular VM termination in the deprovisioning phase may improve the performance when the cached data stored within the VM is intelligently considered. In this work, we did not consider task failure either caused by the software or the infrastructure (i.e., container, VMs). Incorporating a fault-tolerant strategy into both EPSM and EBPSM algorithms is necessary to address the task failure that highly likely impacting the WaaS cloud platforms performance.

Finally, further investigation on how multiple container instances can be run and scheduled on top of a single VM is another to-do list. The delay in initiating a container image has reduced our algorithm's performance. There must be a way to counterbalance this issue by exploiting the swarming nature of container to gain benefits from this predetermined condition to enhance the WaaS cloud platforms efficiency further.

## Acknowledgments

This research is partially supported by LPDP (Indonesia Endowment Fund for Education) and ARC (Australia Research Council) research grant.

## References

- [1] E. Deelman, T. Peterka, I. Altintas, C. D. Carothers, K. K. van Dam, K. Moreland, M. Parashar, L. Ramakrishnan, M. Taufer, J. Vetter, The Future of Scientific Workflows, *The International Journal of High Performance Computing Applications* 32 (1) (2018) 159–175. doi:10.1177/1094342017704893.
- [2] E. Deelman, G. Singh, M. Livny, B. Berriman, J. Good, The cost of doing science on the cloud: The montage example, in: *Proceedings of The ACM/IEEE Conference on Supercomputing*, 2008, pp. 1–12. doi:10.1109/SC.2008.5217932.

- [3] P. Leitner, J. Cito, Patterns in the Chaos: A Study of Performance Variation and Predictability in Public IaaS Clouds, *ACM Transaction on Internet Technology* 16 (3) (2016) 15:1–15:23. doi:10.1145/2885497.
- [4] E. N. Alkhanak, S. P. Lee, R. Rezaei, R. M. Parizi, Cost Optimization Approaches for Scientific Workflow Scheduling in Cloud and Grid Computing: A Review, Classifications, and Open Issues, *Journal of Systems and Software* 113 (Supplement C) (2016) 1–26. doi:10.1016/j.jss.2015.11.023.
- [5] M. A. Rodriguez, R. Buyya, A Taxonomy and Survey on Scheduling Algorithms for Scientific Workflows in IaaS Cloud Computing Environments, *Concurrency and Computation: Practice and Experience* 29 (8) (2017) e4041–n/a. doi:10.1002/cpe.4041.
- [6] M. A. Rodriguez, R. Buyya, Scheduling Dynamic Workloads in Multi-tenant Scientific Workflow as a Service Platforms, *Future Generation Computer Systems* 79 (Part 2) (2018) 739–750. doi:10.1016/j.future.2017.05.009.
- [7] M. H. Hilman, M. A. Rodríguez, R. Buyya, Task-Based Budget Distribution Strategies for Scientific Workflows with Coarse-Grained Billing Periods in IaaS Clouds, in: *Proceedings of The 13th IEEE International Conference on e-Science, 2017*, pp. 128–137. doi:10.1109/eScience.2017.25.
- [8] M. Jones, B. Arcand, B. Bergeron, D. Bestor, C. Byun, L. Milechin, V. Gadepally, M. Hubbell, J. Kepner, P. Michaleas, J. Mullen, A. Prout, T. Rosa, S. Samsi, C. Yee, A. Reuther, Scalability of VM Provisioning Systems, in: *Proceedings of The IEEE High Performance Extreme Computing Conference, 2016*, pp. 1–5. doi:10.1109/HPEC.2016.7761629.
- [9] R. V. Lopes, D. Menascé, A taxonomy of job scheduling on distributed computing systems, *IEEE Transactions on Parallel and Distributed Systems* 27 (12) (2016) 3412–3428.
- [10] D. Tsafir, Y. Etsion, D. G. Feitelson, Backfilling using system-generated predictions rather than user runtime estimates, *IEEE Transactions on Parallel and Distributed Systems* 18 (6) (2007) 789–803.
- [11] A. Verma, L. Cherkasova, R. H. Campbell, Aria: Automatic resource inference and allocation for mapreduce environments, in: *Proceedings of the 8th ACM International Conference on Autonomic Computing, 2011*, p. 235–244. doi:10.1145/1998582.1998637.
- [12] B. P. Rimal, M. Maier, Workflow Scheduling in Multi-tenant Cloud Computing Environments, *IEEE Transactions on Parallel and Distributed Systems* 28 (1) (2017) 290–304. doi:10.1109/TPDS.2016.2556668.
- [13] H. Chen, J. Zhu, G. Wu, L. Huo, Cost-efficient Reactive Scheduling for Real-time Workflows in Clouds, *The Journal of Supercomputing* doi:10.1007/s11227-018-2561-9.
- [14] R. Qasha, J. Cala, P. Watson, Dynamic Deployment of Scientific Workflows in The Cloud Using Container Virtualization, in: *Proceedings of The IEEE International Conference on Cloud Computing Technology and Science, 2016*, pp. 269–276. doi:10.1109/CloudCom.2016.0052.
- [15] K. Liu, K. Aida, S. Yokoyama, Y. Masatani, Flexible Container-based Computing Platform on Cloud for Scientific Workflows, in: *Proceedings of The International Conference on Cloud Computing Research and Innovations, 2016*, pp. 56–63. doi:10.1109/ICCCRI.2016.17.
- [16] E. J. Alzahrani, Z. Tari, Y. C. Lee, D. Alsadie, A. Y. Zomaya, adCFS: Adaptive Completely Fair Scheduling Policy for Containerised Workflows Systems, in: *Proceedings of The 16th IEEE International Symposium on Network Computing and Applications, 2017*, pp. 1–8. doi:10.1109/NCA.2017.8171362.
- [17] G. L. Stavrinides, F. R. Duro, H. D. Karatza, J. G. Blas, J. Carretero, Different Aspects of Workflow Scheduling in Large-scale Distributed Systems, *Simulation Modelling Practice and Theory* 70 (Supplement C) (2017) 120–134. doi:10.1016/j.simpat.2016.10.009.
- [18] H. Arabnejad, J. G. Barbosa, Maximizing The Completion Rate of Concurrent Scientific Applications Under Time and Budget Constraints, *Journal of Computational Science* 23 (Supplement C) (2017) 120–129. doi:10.1016/j.jocs.2016.10.013.
- [19] N. Zhou, F. Li, K. Xu, D. Qi, Concurrent Workflow Budget- and Deadline-constrained Scheduling in Heterogeneous Distributed Environments, *Soft Computing* doi:10.1007/s00500-018-3229-3.
- [20] H. Arabnejad, J. G. Barbosa, Multi-QoS Constrained and Profit-aware Scheduling Approach for Concurrent Workflows on Heterogeneous Systems, *Future Generation Computer Systems* 68 (Supplement C) (2017) 211–221. doi:10.1016/j.future.2016.10.003.
- [21] H. Chen, J. Zhu, Z. Zhang, M. Ma, X. Shen, Real-time Workflows Oriented Online Scheduling in Uncertain Cloud Environment, *The Journal of Supercomputing* 73 (11) (2017) 4906–4922. doi:10.1007/s11227-017-2060-4.
- [22] H. Chen, X. Zhu, G. Liu, W. Pedrycz, Uncertainty-Aware Online Scheduling for Real-Time Workflows in Cloud Service Environment, *IEEE Transactions on Services Computing* (2018) 1–1doi:10.1109/TSC.2018.2866421.
- [23] H. Cao, C. Q. Wu, Performance Optimization of Budget-constrained MapReduce Workflows in Multi-clouds, in: *Proceedings of The 18th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing, 2018*, pp. 243–252. doi:10.1109/CCGRID.2018.00039.
- [24] Y. Caniou, E. Caron, A. K. W. Chang, Y. Robert, Budget-Aware Scheduling Algorithms for Scientific Workflows with Stochastic Task Weights on Heterogeneous IaaS Cloud Platforms, in: *Proceedings of The IEEE International Parallel and Distributed Processing Symposium Workshops, 2018*, pp. 15–26. doi:10.1109/IPDPSW.2018.00014.
- [25] V. Arabnejad, K. Bubendorfer, B. Ng, A Budget-Aware Algorithm for Scheduling Scientific Workflows in Cloud, in: *Proceedings of The 18th IEEE International Conference on High Performance Computing and Communications; 14th IEEE International Conference on Smart City; 2nd IEEE International Conference on Data Science and Systems, 2016*, pp. 1188–1195. doi:10.1109/HPCC-SmartCity-DSS.2016.0166.
- [26] W. Chen, G. Xie, R. Li, Y. Bai, C. Fan, K. Li, Efficient Task Scheduling for Budget Constrained Parallel Applications on Heterogeneous Cloud Computing Systems, *Future Generation Computer Systems* 74 (2017) 1–11. doi:10.1016/j.future.2017.03.008.
- [27] Z. Kozhimbayev, R. O. Sinnott, A Performance Comparison of Container-based Technologies for The Cloud, *Future Generation Computer Systems* 68 (Supplement C) (2017) 175–182. doi:10.1016/j.future.2016.08.025.
- [28] R. Shea, F. Wang, H. Wang, J. Liu, A Deep Investigation Into Network Performance in Virtual Machine Based Cloud Environments, in: *Proceeding of The IEEE Conference on Computer Communications, 2014*, pp. 1285–1293. doi:10.1109/INFOCOM.2014.6848061.
- [29] T. P. Pham, J. J. Durillo, T. Fahringer, Predicting Workflow Task Execution Time in The Cloud Using A Two-Stage Machine Learning Approach, *IEEE Transactions on Cloud Computing* PP (99) (2017) 1–1. doi:10.1109/TCC.2017.2732344.
- [30] M. H. Hilman, M. A. Rodríguez, R. Buyya, Task Runtime Prediction in Scientific Workflows Using an Online Incremental Learning Approach, in: *Proceedings of The 11th IEEE/ACM International Conference on Utility and Cloud Computing, 2018*, pp. 93–102. doi:10.1109/UCC.2018.00018.
- [31] G. Juve, A. Chervenak, E. Deelman, S. Bharathi, G. Mehta, K. Vahi, Characterizing and Profiling Scientific Workflows, *Future Generation Computer Systems* 29 (3) (2013) 682–692. doi:10.1016/j.future.2012.08.015.
- [32] R. N. Calheiros, R. Ranjan, A. Beloglazov, C. A. F. de Rose, R. Buyya, CloudSim: a Toolkit for Modeling and Simulation of Cloud Computing Environments and Evaluation of Resource Provisioning Algorithms, *Software: Practice and Experience* 41 (1) (2011) 23–50. doi:10.1002/spe.995.
- [33] M. Ullrich, J. Lässig, J. Sun, M. Gaedke, K. Aida, A Benchmark Model for The Creation of Compute Instance Performance Footprints, in: *Internet and Distributed Computing Systems, Springer International Publishing, 2018*, pp. 221–234.

- [34] S. F. Piraghaj, A. V. Dastjerdi, R. N. Calheiros, R. Buyya, ContainerCloudSim: An Environment for Modeling and Simulation of Containers in Cloud Data Centers, *Software: Practice and Experience* 47 (4) (2017) 505–521. doi:10.1002/spe.2422.
- [35] M. A. Rodriguez, R. Buyya, Scientific Workflow Management System for Clouds, in: I. Mistrik, R. Bahsoon, N. Ali, M. Heisel, B. Maxim (Eds.), *Software Architecture for Big Data and the Cloud*, Morgan Kaufmann, 2017, pp. 367–387.
- [36] M. H. Hilman, Budget-constrained Workflow Applications Scheduling in Workflow-as-a-Service Cloud Computing Environments, Ph.D. thesis, The University of Melbourne (2020).