

## An Analysis of Provisioning and Allocation Policies for Infrastructure-as-a-Service Clouds

David Villegas\*, Athanasios Antoniou<sup>‡</sup>, Seyed Masoud Sadjadi\*, and Alexandru Iosup<sup>‡</sup>

\*School of Computing and Information Sciences, Florida International University, Miami, USA  
 {dvill013,sadjadi}@cs.fiu.edu

<sup>‡</sup>Parallel and Distributed Systems Group, Delft University of Technology, Delft, The Netherlands,  
 A.Antoniou@student.tudelft.nl, A.Iosup@tudelft.nl

**Abstract**—Today, many commercial and private cloud computing providers offer resources for leasing under the infrastructure as a service (IaaS) paradigm. Although an abundance of mechanisms already facilitate the lease and use of single infrastructure resources, to complete multi-job workloads IaaS users still need to select adequate provisioning and allocation policies to instantiate resources and map computational jobs to them. While such policies have been studied in the past, no experimental investigation in the context of clouds currently exists that considers them jointly. In this paper we present a comprehensive and empirical performance-cost analysis of provisioning and allocation policies in IaaS clouds. We first introduce a taxonomy of both types of policies, based on the type of information used in the decision process, and map to this taxonomy eight provisioning and four allocation policies. Then, we analyze the performance and cost of these policies through experimentation in three clouds, including Amazon EC2. We show that policies that dynamically provision and/or allocate resources can achieve better performance and cost. Finally, we also look at the interplay between provisioning and allocation, for which we show preliminary results.

**Keywords**—Scheduling; Cloud computing; Provisioning policies; Allocation policies; Empirical performance analysis.

### I. INTRODUCTION

Recent advances [1], [2] in the high-speed yet low-cost interconnection of off-the-shelf computational and storage resources have facilitated the creation of data centers of unprecedented scale. As a consequence, a fundamental shift is beginning to occur in the way computational resources are provisioned and allocated by our society, from traditional ownership to Infrastructure-as-a-Service (IaaS) clouds—leasing and releasing virtualized resources. Although hundreds of commercial IaaS providers exist, to transition to IaaS clouds users still need detailed understanding of achieved performance and incurred cost. In particular, potential IaaS users need to understand the performance and cost of resource provisioning and allocation policies, and the interplay between them. To address this need, we conduct an empirical analysis of resource provisioning and allocation policies for IaaS clouds.

The basic operation of an IaaS cloud is to temporarily instantiate on-demand virtual machines (VMs)—of pre-agreed computing power and memory size, operating system and provided libraries, and, usually, some measure of Quality of Service (QoS). Providers host the resources shared by all users and can take advantage of economies of scale by leasing their physical infrastructure as virtualized resources

to many different classes of users. Users provision, that is, acquire and release, resources based on their actual needs, only when, where, and for how long needed; they can allocate the provisioned resources according to the specific requirements of the workloads at hand. This model makes IaaS clouds a flexible solution that reduces or completely eliminates the need for acquiring and managing costly physical resources, but also introduces the need to consider *online* the trade-off between performance (more resources) and cost.

The provisioning and allocation policies can have an important impact on the traditional performance metrics, from workload makespan to individual job slowdown [3]. Since instantiating a large number of VMs is simple, over-provisioning can incur a substantial yet unnecessary cost; when the allocation policy is inefficient, a dynamic provisioning policy may lease resources that remain largely unused [4]. The pricing scheme of the IaaS cloud, which may include hourly charging periods and discounts for first-time use, may lead to different cost gains than expected from actual resource consumption.

The performance-cost trade-off has been extensively studied in the context of grids, mainly from the perspective of users also acting as providers and looking for economic gain [5] or for sharing fairness [6]. Moreover, more traditional resource managers such as Condor support [7], [8], through versatile job specification languages, complex criteria for the selection of resources. Several studies [3], [9], [4] have approached this trade-off in simulation, in the context of clouds. However, until now no study has investigated in practice the impact of the provisioning and allocation policies that users can employ, and of the interplay between these policies, in the context of clouds. The need for empirical evaluation stems from recent results [10], [11], [12] in cloud performance evaluation, which show that cloud performance is much lower and more variable than considered in simulation models. In this work we propose a comprehensive investigation of provisioning and allocation policies for IaaS clouds. Our main contribution is threefold:

- 1) We identify eight provisioning and four allocation policies that can realistically be applied for managing workloads in IaaS clouds (Section III);
- 2) We conduct an empirical investigation using three IaaS clouds, including the services provided by the commercial IaaS Amazon EC2 (Section IV);

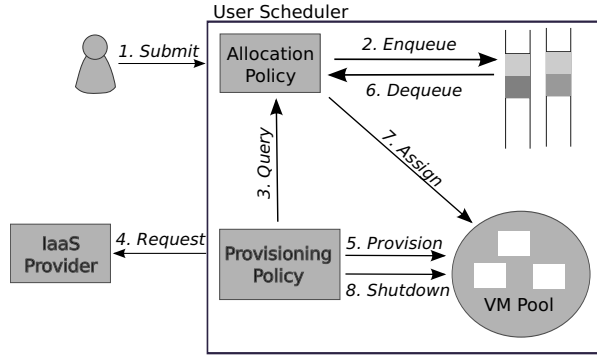


Figure 1. The cloud ecosystem.

- 3) We analyze empirically and, only for the otherwise expensive experiments, in simulation the performance and cost of resource provisioning and allocation policies, and the interplay between these two types of policies (Section V).

## II. SYSTEM MODEL

In this section we present the system model used throughout this work.

### A. Workload Model

Although desirable, it is not yet possible to define a realistic workload for IaaS clouds, due to the scarcity of public workload traces or common practice reports. The Grid and Parallel Workload Archives provide in total tens of workload traces, but it is yet unknown if the users of these environments will migrate to IaaS clouds [11]. We have already shown in our characterization of over fifteen grid workloads [13] two trends: the disappearance of tightly-coupled parallel jobs in favor of Bags of Tasks (BoTs), and the decrease of the amount of work (runtime) of each task.

For this work, we consider synthetic, BoT-based workloads with runtimes typical for data mining and semi-interactive processing, that is, several minutes [14], [15]. In our workload model, jobs are CPU-bound and their runtime is dependent on the speed of the (virtual) processor where they are executed and the amount of work requested, specified as an argument to the job. By using different workload distributions, we want to showcase the behavior of the policies under specific circumstances.

### B. Resource Model

In our system model, resources are provisioned exclusively from IaaS clouds. Although hybrid local-cloud systems still exist, we anticipate with this work the moment when buying and maintaining local resources will be anachronistic. The cloud ecosystem investigated in this work, which is comprised of the IaaS cloud provider(s) and user(s), is depicted in Figure 1. We assume that users send their workloads to a scheduler, which enqueues jobs and

Policy	Dynamic	Trigger	Job Duration	Increase Param.
Startup	No	—	—	—
OD-S	Yes	Queue size	No	1
OD-G	Yes	Queue size	No	n
OD-ExecTime	Yes	Exec. time	Yes	1
OD-ExecAvg	Yes	Exec. time	No	1
OD-ExecKN	Yes	Exec. time	No	1
OD-Wait	Yes	Wait time	No	1
OD-2Q	Yes	Queue size	Partial	1

Table I  
OVERVIEW OF PROVISIONING POLICIES.

assigns them to the pool of available resources based on an *allocation policy*. A system component manages the pool of resources via a *provisioning policy*, that is, a policy that decides when to lease and to release resources from IaaS clouds. This component can query the state of the allocation, such as the queue size or the average waiting time for jobs.

We model the operation of IaaS clouds based on Amazon EC2, a popular commercial IaaS. We assume that a provisioning request issued to the IaaS cloud will incur delays that depend on the process used by the IaaS to select, lease-and-boot, and release (shut down) a VM instance. Last, we assume that VMs have a cost associated to their operation, with a cost model that is proportional, possibly non-linearly, to the runtime of the VM instance, and not counting the time required to boot or shut down the instance.

## III. PROVISIONING AND ALLOCATION POLICIES

We present in this section the provisioning and allocation policies considered for this work. Although we envision that future policies will adapt to changing workloads, evolving resources, and complex Service Level Agreements, we focus in this work on the simple policies that may be realistically employed in practice, today. Our rationale was to select commonly used policies that have already been studied in the literature.

### A. Provisioning Policies

We consider for this work eight provisioning policies; we summarize their properties in Table I and describe them next. Overall, we identify two main classes of provisioning policies, static and dynamic. For the class of dynamic provisioning policies, we investigate three criteria for deciding when the policies are called (*triggers*): the current size of the queue, the accumulated execution time of the queued jobs, and the total waiting time among queued jobs. Some of the provisioning policies considered here require information about the duration of jobs, which is not always provided by the user; we consider, for these policies, alternatives that estimate this information from historical records. The last policy in Table I requires partial information, that is, it only requires classification information—a job can be either small or long. We now describe the provisioning policies, in turn:

1. *Startup*: is a provisioning policy that leases the maximum possible VM instances during the system startup. While this ensures that no delay will occur due to waiting for new resources to be leased, this policy is not flexible—whether jobs arrive or not in the system, VM instances are still leased and paid; also, this policy cannot cope well with bursty system overloads.

2. *On-Demand, Single VM (OD-S)*: is a naïve dynamic provisioning policy that leases a new VM instance for each job that cannot be assigned to a resource. VMs are shut down when they are not used for a certain duration, determined as a parameter; by setting this parameter to  $0$  seconds, VMs are released immediately after the job is completed. In general, this policy can lead to *thrashing*, that is, frequent leasing and releasing of VM instances.

3. *OD-Geometric (OD-G)*: is a dynamic policy that extends OD-S with geometric provisioning, that is, when new instances are needed this policy provisions  $n^0, n^1, n^2, \dots$  machines in successive leases, where  $n$  is the *increase parameter*. Similarly, this policy releases increasing amounts of instances. The decision of (re)leasing instances is taken periodically, when the number of available VMs falls below a threshold, or the number of tasks in the system queue exceed a limit; here, we use a period of 20 seconds and (re)lease again whenever there is at least one queued job/idle VM instance.

4. *OD-ExecTime*: is a dynamic provisioning policy that uses the future execution time of queued jobs, which is assumed to be known *a priori*, to decide on leasing or releasing VM instances. The decision to lease is adaptive to the cloud, in that the execution time of queued jobs must exceed the average time needed to provision and boot a VM instance (as observed for previously leased VMs) by a specified factor, which is set in this work to 5.

5. *OD-ExecAvg*: is similar to OD-ExecTime, but the execution time for each queued job is *estimated* as the average execution time of all the jobs that have already completed. An initial prediction of the average job run-time must be provided by the user.

6. *OD-ExecKN*: is similar to OD-ExecAvg, but uses a predictor based on [16]. For each job in the queue, OD-ExecKN acquires its  $k$ -nearest neighbors based on the job input parameter size. Then, the estimated execution time for a job is the average over this set of  $k$  neighbors.

7. *OD-Wait*: is similar to OD-ExecTime, but it considers the waiting times of queued jobs instead of their future execution time.

8. *OD-2Q*: is a provisioning policy that works in conjunction with the FCFS-2Q allocation policy (Section III-B). To minimize the trashing of VMs for short running jobs, we define a bi-queue on-demand provisioning policy that leases VM instances and assigns them to one of two pools. The two pools effectively implement two OD-S queues with separate idle time parameters; here, we use longer idle times for VMs

Policy	Uses Queue	Job Duration	Provisioning-Aware
FCFS	Yes	No	No
FCFS-NW	No	No	No
SJF	Yes	Yes	No
FCFS-2Q	Yes (2)	Partial	Yes

Table II  
OVERVIEW OF ALLOCATION POLICIES.

that will run short jobs, so that thrashing is reduced.

### B. Allocation Policies

We consider for this work four allocation policies; their properties are summarized in Table II. Most policies we investigate use one queue, but we also consider policies that use two or no queue. Similarly to our approach for provisioning policies, we consider here allocation policies that may require (partial) information about the job duration. The last policy in Table II is *provisioning-aware*, that is, it works in conjunction with the OD-2Q provisioning policy. We now describe the allocation policies, in turn:

1. *First-Come, First-Served (FCFS)*: is a traditional allocation policy that assigns one task per resource in the order in which the tasks have been submitted to the system.

2. *FCFS-NoWait (FCFS-NW)*: is an extension to FCFS where jobs are not queued when no available VMs exist. Instead, this policy assigns jobs to VMs that are already running other jobs, round robin. This policy eliminates the wait time, but may introduce bottlenecks in the execution of jobs.

3. *Shortest-Job First (SJF)*: is a traditional allocation policy that gives priority to shorter jobs, assuming there is some information about or estimation of their actual duration. Although it alleviates the FCFS problem of short jobs waiting for the allocation of longer jobs with earlier arrival time, it can lead to starvation for long jobs.

4. *FCFS-MultiQueue*: is an extension to FCFS where several FCFS queues, one for each range of job durations, are maintained. The simplest case considered here, FCFS-2Q, has two queues, one for short jobs and another for long jobs. Although an estimation of the runtime is necessary for this policy, it is enough to have partial knowledge of it to classify jobs. This policy can work in conjunction with a provisioning policy that divides VMs into pools for short and long running jobs, such as OD-2Q (introduced in Section III-A).

## IV. EXPERIMENTAL SETUP

In this section we discuss our experimental setup, in turn, the SkyMark empirical performance evaluation tool, the used testbeds, the workloads, and the employed metrics.

### A. The SkyMark Empirical Performance Evaluation Tool

Unless otherwise specified, we have conducted the experiments presented in this work in *real* environments. To this end, we have implemented SkyMark, a performance

System	Hardware Spec	VIM/Hypervisor	Max VMs
DAS4 Delft	8 dual quad-core 24 GB RAM	OpenNebula 3.0 / KVM (Full, HVM)	64
FIU	8 Pentium 4 5 GB Memory	OpenNebula 2.2 / XEN (Para., No HVM)	7
Amazon EC2 eu-west-1	-	- / XEN (Para., No HVM)	20

Table III  
OVERVIEW OF THE EXPERIMENTAL ENVIRONMENTS.

evaluation framework that enables the generation, submission, and monitoring of complex workloads to IaaS cloud environments. SkyMark extends C-meter[17], our prior IaaS performance evaluation framework, with provisioning and allocation policies, new workloads, and new analytical capabilities. SkyMark currently supports IaaS clouds that implement Amazon EC2’s interface, including all deployments using Eucalyptus, but interfaces to other clouds can be easily plugged-in to the tool. SkyMark also supports the XML-RPC system interface of OpenNebula.

The experimental process uses SkyMark as follows: The user provides a workload description file which is used to generate a real or synthetic workload. The workload is then submitted to a cloud, using pre-specified provisioning and allocation policies. Performance statistics are gathered throughout the execution of the workload and then stored in a database. Last, the database is used for post-experiment analysis. SkyMark effectively plays the role of both the user and the policies depicted in Figure 1.

Since using a real testbed is constrained by actual resource and budget availability, we have also developed a discrete event simulator that duplicates Skymark functionality. The simulator reads a workload description and generates events for job arrivals, VM lifecycle, and pluggable allocation and provisioning policies.

### B. Experimental Environments

We have performed experiments on three *real* IaaS clouds with different resources, middleware, and virtualization systems. The three used systems are: the Delft cluster of DAS-4<sup>1</sup>, a six-cluster wide-area distributed system in the Netherlands; a cluster at Florida International University (FIU); and the Amazon EC2 commercial IaaS. By running the experiments in these different scenarios, we want to test what is the behavior of the studied policies in various real deployments. The properties of the three IaaS clouds used in this work are summarized in Table III.

### C. Workloads

We have used the following workloads, each hour-long:

- 1) **Uniform:** A steady stream of tasks throughout the experiment; uses a Poisson arrival distribution. The average system load is around 70%.

<sup>1</sup><http://www.cs.vu.nl/das4/>

- 2) **Increasing:** The workload intensity increases over time, in steps. The average system load is around 50%.
- 3) **Bursty:** The workload features short spikes of intense activity amid long periods of mild or moderate activity. The average system load is around 15%; the maximum load is, for a few minutes, 170%.

For the simulated results in section V, we use the Uniform and Bursty workloads, and add a new one, Periodic, following a periodic increasing and decreasing arrival pattern. For these workloads we also change the job durations. More details are discussed in the corresponding section.

The jobs that comprise the workloads are synthetic and have an average execution time of 47 seconds with a standard deviation of 41.1, as measured on DAS-4 when jobs were run independently. The reason for selecting short job durations were discussed in Section II-A.

### D. Performance, Cost, and Compound Metrics

We use a variety of performance, cost, and compound metrics to analyze the impact of provisioning and allocation policies. For performance metrics, we look at the *average job slowdown* ( $JSD$ ), defined per job as the ratio of the actual runtime in the cloud and the runtime in a dedicated environment. We also look at the workload speedup, measured against a single node ( $SU$ ).

We use two cost metrics. The *actual cost* ( $C_a$ ) is defined as the sum of consumed resources, here, CPUtime. The *charged cost* ( $C_c$ ) is the price charged by the provider, here, using the Amazon EC2 pricing model of charging CPUtime consumption in increments of one hour.

We define two compound metrics, cost efficiency and utility. We define *cost efficiency* ( $C_{eff}$ ) as the ratio of the charged and actual cost; lower values are better, with a value of 1 indicating fair pricing, and values below 1 indicating dumping prices or economies-of-scale. We define *utility* ( $U$ ) as a compound metric that rewards low performance overheads and low cost:  $U = \frac{SU}{C_c}$ . The values for utility are normalized on the Startup provisioning policy.

## V. EXPERIMENTAL RESULTS

In this section we perform a set of experiments to explore the effects of different policies and their interactions on different systems and under varying conditions. We want to determine how policies perform, when it is better to use ones versus others, and which allocation and provisioning policies work better when used together.

We present here only representative results; for complete results, we refer to our technical report [18].

### A. Provisioning Policies

We first explore the effect of the provisioning policies. To this end, we use the same allocation policy, FCFS, coupled in turn with each one of the provisioning policies. We show the results in Figures 2, 3, 4, and in Table IV.




Figure 2. Workload speedup (SU) for Provisioning policies. OD-S was not tested on EC2, because of the significant cost that it would incur.

Figures 2 and 3 present the workload speedup (SU) and the job slowdown (JSD) respectively. `Startup` always achieves the best performance. OD-S has similar performance for the uniform workload, but is not as good for the variable workloads. From the threshold-based policies, `QueueWait` usually performs better than the rest, because it reacts faster to load variation. `ExecTime` and its variants have similar performance, with `ExecTime` usually performing better, since `ExecAvg` and `ExecKN` do not have exact job runtime information.

The charged cost ( $C_c$ ) is shown in Figure 4. OD-S incurs the highest cost, since VMs are started and stopped reactively to individual job arrivals. The group of threshold-based policies and especially the `Exec` family of policies significantly reduce the cost of workload execution. The cost reduction becomes bigger for the increasing and bursty workloads.

We additionally provide the cost efficiency ( $C_{eff}$ ) and utility (U) metrics for EC2, in Table IV. The dynamic policies hold on to resources for shorter periods of time than `Startup`, especially for bursty workloads. This leads to a worse cost efficiency value. However, they do achieve better utility scores, which means that they provide a better performance-cost trade-off.

### B. Allocation Policies for Static Resources

In this experiment we want to study the performance of different allocation policies and the static provisioning policy, `Startup`. Resources are acquired at the beginning of the experiment, and then jobs are sent to the system.

We use the `FCFS`, the `SJF`, and the `FCFS-NW` allocation policies in the three testbeds. Figure 5 lists the results for the job slowdown metric. The experiment shows that `SJF` gives a lower slowdown, since shorter jobs are processed first, which means jobs in wait less time in the queue. Overall, `FCFS` performs similarly to `SJF` for the uniform and increasing workloads, however its performance degrades when




Figure 3. Job Slowdown (JSD) for Provisioning policies.




Figure 4. Charged cost ( $C_c$ ) for Provisioning policies.

under a bursty load. Lastly, the `FCFS-NW` policy, which assigns jobs to VMs with round-robin, creates resource competition, and thus has worse results in all experiments.

### C. Effects of job size distribution in on-demand policies

In this section, we investigate the impact of the job size distribution. To this end, we create an artificial workload with periodic arrival intervals with different ratios of short and long jobs. For this case, jobs have runtime averages of 10 seconds for short ones and 1 hour for long ones. We consider workloads composed of 25%, 50% and 75% short jobs (SR25%, SR50% and SR75%, respectively).

Figure 6 shows that, compared to OD-S with `FCFS` and OD-S with `SJF`, the new combination of policies achieves a lower cost for the workload. Also, the slowdown is reduced for workloads with a high number of short jobs. The reason for this is that long jobs are no longer kept in the queue for long time. However, we can see that for workloads with a low number of short jobs, the new policies result in higher slowdowns. This is due to the fact that the workload cannot take advantage of differentiating the two types of

Workload	Charged Cost			Cost Efficiency			Utility		
	Uniform	Increasing	Bursty	Uniform	Increasing	Bursty	Uniform	Increasing	Bursty
Startup	40	40	40	1.9	1.9	1.9	1.0	1.0	1.0
ExecTime	26 (-35%)	15 (-62%)	24 (-40%)	1.6 (-15%)	1.7 (-12%)	4.0 (+111%)	1.4 (+41%)	2.4 (+137%)	1.5 (+51%)
ExecAvg	34 (-15%)	17 (-57%)	26 (-35%)	2.2 (+17%)	1.9 (-1%)	4.3 (+126%)	1.0 (-1%)	2.2 (+116%)	1.4 (+40%)
ExecKN	30 (-25%)	16 (-60%)	25 (-38%)	2.0 (+7%)	1.8 (-5%)	4.1 (+115%)	1.2 (+20%)	2.3 (+133%)	1.5 (+45%)
QueueWait	29 (-28%)	18 (-55%)	38 (-5%)	2.0 (+3%)	2.0 (+4%)	6.2 (+227%)	1.3 (+29%)	2.1 (+113%)	1.0 (+2%)

Table IV  
CHARGED COST, COST EFFICIENCY AND UTILITY FOR POLICIES ON EC2. FOR RESULTS ON DAS4 AND FIU, WE REFER TO OUR TECHNICAL REPORT [18]




Figure 5. Average Job Slowdown for Allocation policies.

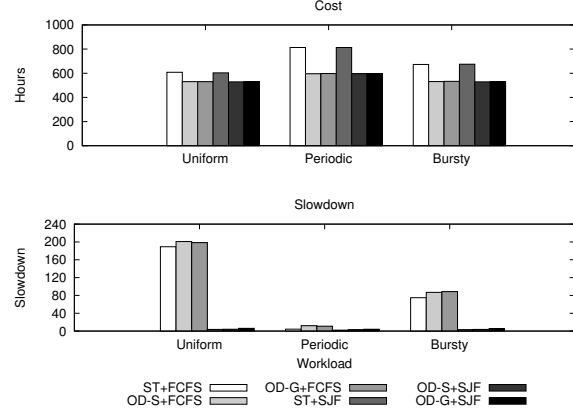


Figure 7. Slowdown and cost for groups of policies




Figure 6. Slowdown and cost for job runtime ratios

jobs, and additionally partitioning the pool of VMs increases the possibility of jobs from one group not having enough resources because the VMs are used by the other jobs, while other policies have a common pool of resources that can be used by any job.

#### D. Policy interactions

We study here the interactions between allocation and provisioning policies. We use the simulator to test six pairs of policies that comprise three provisioning, *Startup*, *OD-S*, and *OD-G*; and two allocation policies, *FCFS* and *SJF*. We use three workloads with one thousand jobs where half of the jobs have a runtime average of 10 seconds and the

rest of one hour. Jobs of the first workload, *Uniform*, have an interarrival time of 10 seconds. The second one, *Periodic*, has four periods of increasing and decreasing arrival times starting and ending at 60 seconds and peaking at 5 seconds. The last one, *Bursty*, alternates five periods of 200 jobs with high (30 seconds) and low (2 seconds) interarrival times.

Figure 7 shows the results of this experiment. The top half of the figure shows the cost of running the workloads, and it can be seen how on-demand provisioning policies are much more sensitive to variations in comparison to *Startup*. This is especially noticeable for the periodic workload, which has the highest variability; For the uniform workload, the system is at full utilization most of the time, minimizing the benefits of dynamic provisioning policies. The bottom half of the Figure 7 shows the average job slowdown, and it illustrates how the *SJF* allocation policy reduces the overall overhead for jobs by executing the shorter jobs first, and therefore minimizing the time that jobs wait in general. Another conclusion is that *Startup* results in lower slowdown in comparison to dynamic provisioning policies, due to the lack of overhead for VM booting and shutdown. Additionally, the figure illustrates how the *OD-G* policy results in slightly higher slowdown. The reason for this is that the geometric increase of VMs needs some time to ramp up to reach the required number of VMs to run all jobs in the queue, while the *OD-S* policy instantiates one VM for each waiting job.

## VI. RELATED WORK

Much effort [4], [9], [3] has been recently put in investigating algorithms and policies for scheduling jobs on dynamic resources such as clouds. In contrast to our work, which focuses on a system model where all resources are provided from IaaS clouds, most of the related work [19], [20], [21], [22] considers the use of cloud resources only as an extension to an existing, non-cloud system. Moreover, ours is the first comprehensive investigation of provisioning and allocation policies conducted in *real* IaaS clouds.

Closest to our work, Genaud et al. [4] *simulate* various algorithms to provision VMs and assign jobs to them, but consider different policies than in this work and it is unclear if their simulation results are realistic. Assuncao et al. [3] study in simulation three allocation policies and different provisioning policies for extending the capacity of a local private infrastructure with public cloud resources. Murphy et al. [23] discuss an algorithm to provision clusters of VMs for Condor jobs. Kijisiponse et al. [9] consider the FIFO allocation policy and study several VM provisioning policies. In [19], the local site is extended with IaaS resources based on the decision of provisioning policies that react to current load. Lu et al. [20] address the problem of idle instances when executing BLAST workloads by dynamically scaling in Microsoft's Azure cloud. Mao et al. [24] present an algorithm for automatic provisioning to host jobs with deadlines, taking VM startup time and different instance types into consideration. However, they concentrate on the provisioning decisions rather than on job allocation. Candeia et al. [21] propose a greedy allocation policy to manage bursts of bags-of-tasks. Ostermann et al. [25] design a hybrid resource manager that dynamically provisions cloud resources and adds them to a grid. Deelman et al. [8] and Wang et al. [26] explore how scientific communities can use cloud resources, where VM instances are leased based on the ratio of waiting jobs to total available VMs.

Other approaches consider complex policies, for example policies [27], [28], [29], [30] that try to predict the workload to determine when to provision new VMs. Cost is another important parameter when provisioning: Henzinger et al. [31] describe a model where a cloud presents different schedules and costs. Other related work [32], [33] uses market approaches to determine when to provision new resources. Our study complements these approaches with a more realistic investigation focusing on simpler policies.

Some authors have considered the effects of decoupling policies: Ranganathan et al. [34] consider job and data scheduling separately to improve data availability. Kee et al. [35] describe an approach where resource discovery and acquisition are integrated. Song et al. [36] perform multi-level scheduling in order to allocate jobs to VMs and VMs to Physical resources. Xu et al. [37] implement a two-level scheduling mechanism to manage resources based on SLAs.

Sotomayor et al. [38] consider the overhead impact of transferring a VM to the target host when scheduling a job, and plan such transfers accordingly. [39] discuss the performance of gang scheduling in a cloud, where VMs are acquired on demand. They contemplate the processes of scheduling jobs and handling machines dynamically separately and simulate the behavior of both types of policies. Zhang et al. [40] integrate resource consumption of a PaaS application and provisioning decisions.

## VII. CONCLUSION AND FUTURE WORK

To manage their workloads, current and near-future users of IaaS clouds need a better understanding of the performance and cost of provisioning and allocation policies. In this work we have conducted a comprehensive study of these two types of policies, and of the interaction between them.

Overall, we have investigated eight provisioning and four allocation policies, and their interplay. We have developed SkyMark, a framework for IaaS performance evaluation, and conducted with it empirical research in three IaaS clouds, including Amazon EC2. Due to actual resource and budget availability constraints, we have also duplicated SkyMark functionality in a discrete event simulator. Based on results obtained in real and simulated IaaS clouds, we conclude that none of the tested (combined) policies is consistently better than the others across all cases. Our five main findings are summarized below:

- 1) OD-ExecTime and its two variants (especially OD-ExecKN), are a good performance-cost trade-off among the investigated provisioning policies;
- 2) Geometric on-demand provisioning policies, such as OD-G, are worse than single-VM on-demand policies;
- 3) The combined OD-2Q-FCFS-2Q policy, which we are the first to investigate in the context of clouds, is a good slowdown-cost trade-off for workloads with a significant ratio of short to long jobs;
- 4) Naïve static provisioning policies deliver better performance when there are enough resources, but incur up to 5 times higher cost;
- 5) Allocation policies with information about job run-times achieve significantly better performance than uninformed allocation policies, when coupled with dynamic provisioning policies.

In the future, we plan to extend this work to consider new provisioning and allocation policies that adapt to changing workload, evolving resources, and complex Service Level Agreements. We will also investigate more diverse types of workloads, such as the typical workloads of grids [13].

## ACKNOWLEDGMENTS

This publication is based upon work supported by the National Science Foundation under Grant No. OISE-0730065, the Dutch national program COMMIT and STW/NWO Veni grant @larGe (11881), and by the EU EIT project EUROPA.

## REFERENCES

- [1] C. Guo, H. Wu, K. Tan, L. Shi, Y. Zhang, and S. Lu, "Dcell: a scalable and fault-tolerant network structure for data centers," in *SIGCOMM*, 2008, pp. 75–86.
- [2] A. G. Greenberg, J. R. Hamilton, N. Jain, S. Kandula, C. Kim, P. Lahiri, D. A. Maltz, P. Patel, and S. Sengupta, "V12: a scalable and flexible data center network," in *SIGCOMM*, 2009, pp. 51–62.
- [3] M. D. de Assuncao, A. di Costanzo, and R. Buyya, "Evaluating the cost-benefit of using cloud computing to extend the capacity of clusters," in *HPDC*, 2009, pp. 141–150.
- [4] S. Genaud and J. Gossa, "Cost-wait trade-offs in client-side resource provisioning with elastic clouds," in *IEEE CLOUD*, july 2011, pp. 1–8.
- [5] J. Sherwani, N. Ali, N. Lotia, Z. Hayat, and R. Buyya, "Libra: a computational economy-based job scheduling system for clusters," *Softw., Pract. Exper.*, vol. 34, no. 6, 2004.
- [6] R. Buyya, D. Abramson, J. Giddy, and H. Stockinger, "Economic models for resource management and scheduling in grid computing," *Concurrency and Computation: Practice and Experience*, vol. 14, no. 13-15, pp. 1507–1542, 2002.
- [7] D. Thain, T. Tannenbaum, and M. Livny, "Distributed computing in practice: the condor experience," *Concurrency - Practice and Experience*, vol. 17, no. 2-4, pp. 323–356, 2005.
- [8] E. Deelman, G. Singh, M. Livny, G. B. Berriman, and J. Good, "The cost of doing science on the cloud: the montage example," in *SC*, 2008, p. 50.
- [9] E. Kijispongse and S. Vannarat, "Autonomic resource provisioning in rocks clusters using eucalyptus cloud computing," in *MEDES*, 2010, pp. 61–66.
- [10] D. Jiang, G. Pierre, and C.-H. Chi, "Ec2 performance analysis for resource provisioning of service-oriented applications," in *ICSO/ServiceWave Workshops*, 2009, pp. 197–207.
- [11] A. Iosup, S. Ostermann, N. Yigitbasi, R. Prodan, T. Fahringer, and D. H. J. Epema, "Performance analysis of cloud computing services for many-tasks scientific computing," *IEEE TPDS*, vol. 22, no. 6, pp. 931–945, 2011.
- [12] A. Iosup, N. Yigitbasi, and D. H. J. Epema, "On the performance variability of production cloud services," in *CCGRID*, 2011, pp. 104–113.
- [13] A. Iosup and D. H. J. Epema, "Grid computing workloads," *IEEE Internet Computing*, vol. 15, no. 2, pp. 19–26, 2011.
- [14] B. Hindman, A. Konwinski, M. Zaharia, A. Ghodsi, A. D. Joseph, R. H. Katz, S. Shenker, and I. Stoica, "Mesos: A platform for fine-grained resource sharing in the data center," in *NSDI*, 2011.
- [15] Y. Chen, A. Ganapathi, R. Griffith, and R. H. Katz, "The case for evaluating mapreduce performance using workload suites," in *MASCOTS*, 2011, pp. 390–399.
- [16] M. A. Iverson, F. Özgüner, and G. J. Follen, "Run-time statistical estimation of task execution times for heterogeneous distributed computing," in *HPDC*, 1996, pp. 263–.
- [17] N. Yigitbasi, A. Iosup, D. Epema, and S. Ostermann, "C-meter: A framework for performance analysis of computing clouds," in *CCGRID*, may 2009, pp. 472–477.
- [18] D. Villegas, A. Antoniou, S. M. Sadjadi, and A. Iosup, "An analysis of provisioning and allocation policies for IaaS clouds," TU Delft, Tech. Rep. PDS-2011-009, Nov. 2011, <http://www.pds.ewi.tudelft.nl/fileadmin/pds/reports/2011/PDS-2011-009.pdf>.
- [19] P. Marshall, K. Keahey, and T. Freeman, "Elastic site: Using clouds to elastically extend site resources," *CCGRID*, pp. 43–52, 2010.
- [20] W. Lu, J. Jackson, J. Ekanayake, R. Barga, and N. Araujo, "Performing large science experiments on azure: Pitfalls and solutions," in *CloudCom*, 30 2010-dec. 3 2010, pp. 209–217.
- [21] D. Candêia, R. Araujo, R. Lopes, and F. Brasileiro, "Investigating business-driven cloudburst schedulers for e-science bag-of-tasks applications," in *CloudCom*, 2010, pp. 343–350.
- [22] O. A. Ben-Yehuda, A. Schuster, A. Sharov, M. Silberstein, and A. Iosup, "ExPERT: Pareto-efficient task replication on grids and clouds," Technion CS Dept., Tech. Rep. CS-2011-03, Apr 2011. [Online]. Available: <http://www.cs.technion.ac.il/users/wwwb/cgi-bin/tr-info.cgi/2011/CS/CS-2011-03>
- [23] M. Murphy, B. Kagey, M. Fenn, and S. Goasguen, "Dynamic provisioning of virtual organization clusters," in *CCGRID*, may 2009, pp. 364–371.
- [24] M. Mao, J. Li, and M. Humphrey, "Cloud auto-scaling with deadline and budget constraints," in *Grid Computing (GRID)*, 2010 11th IEEE/ACM International Conference on, oct. 2010.
- [25] S. Ostermann, R. Prodan, and T. Fahringer, "Resource management for hybrid grid and cloud computing," in *Cloud Computing*, N. Antonopoulos and L. Gillam, Eds., 2010.
- [26] L. Wang, J. Zhan, and W. Shi, "In cloud, can scientific communities benefit from the economies of scale?" *TPDS*, vol. PP, no. 99, p. 1, 2011.
- [27] A. Quiroz, H. Kim, M. Parashar, N. Gnanasambandam, and N. Sharma, "Towards autonomic workload provisioning for enterprise grids and clouds," in *Grid*, 2009.
- [28] J. Tirado, D. Higuero, F. Isaila, and J. Carretero, "Predictive data grouping and placement for cloud-based elastic server infrastructures," in *CCGRID*, may 2011, pp. 285–294.
- [29] T. J. Hacker and K. Mahadik, "Flexible resource allocation for reliable virtual cluster computing systems," in *SC*, 2011.
- [30] E.-K. Byun, Y.-S. Kee, J.-S. Kim, and S. Maeng, "Cost optimized provisioning of elastic resources for application workflows," *FGCS*, vol. 27, no. 8, pp. 1011–1026, 2011.
- [31] T. Henzinger, A. Singh, V. Singh, T. Wies, and D. Zufferey, "Flexprice: Flexible provisioning of resources in a cloud environment," in *IEEE CLOUD*, july 2010, pp. 83–90.
- [32] L. Wu, S. Garg, and R. Buyya, "Sla-based resource allocation for software as a service provider (SaaS) in cloud computing environments," in *CCGRID*, may 2011, pp. 195–204.
- [33] M. Salehi and R. Buyya, "Adapting market-oriented scheduling policies for cloud computing," in *ICA3PP*, 2010.
- [34] K. Ranganathan and I. Foster, "Decoupling computation and data scheduling in distributed data-intensive applications," in *HPDC*, 2002, pp. 352–358.
- [35] Y.-S. Kee, K. Yocum, A. A. Chien, and H. Casanova, "Improving grid resource allocation via integrated selection and binding," in *SC*, 2006.
- [36] Y. Song, H. Wang, Y. Li, B. Feng, and Y. Sun, "Multi-tiered on-demand resource scheduling for vm-based data center," in *CCGRID*, may 2009, pp. 148–155.
- [37] J. Xu, M. Zhao, J. Fortes, R. Carpenter, and M. Yousif, "On the use of fuzzy modeling in virtualized data center management," in *ICAC*, june 2007, p. 25.
- [38] B. Sotomayor, K. Keahey, and I. Foster, "Overhead matters: A model for virtual resource management," in *VTDC*, nov. 2006, p. 5.
- [39] I. Moschakis and H. Karatza, "Evaluation of gang scheduling performance and cost in a cloud computing system," *The Journal of Supercomputing*, pp. 1–18, 2010.
- [40] Y. Zhang, G. Huang, X. Liu, and H. Mei, "Integrating resource consumption and allocation for infrastructure resources on-demand," in *IEEE CLOUD*, july 2010, pp. 75–82.