

On the Cluster Admission Problem for Cloud Computing

Ludwig Dierks

University of Zurich

DIERKS@IFI.UZH.CH

Ian A. Kash

University of Illinois at Chicago

IANKASH@UIC.EDU

Sven Seuken

University of Zurich

SEUKEN@IFI.UZH.CH

Abstract

Cloud computing providers face the problem of matching heterogeneous customer workloads to resources that will serve them. This is particularly challenging if customers, who are already running a job on a cluster, scale their resource usage up and down over time. The provider therefore has to continuously decide whether she can add additional workloads to a given cluster or if doing so would impact existing workloads' ability to scale. Currently, this is often done using simple threshold policies to reserve large parts of each cluster, which leads to low efficiency (i.e., low average utilization of the cluster). We propose more sophisticated policies for controlling admission to a cluster and demonstrate that they significantly increase cluster utilization. We first introduce the cluster admission problem and formalize it as a constrained Partially Observable Markov Decision Process (POMDP). As it is infeasible to solve the POMDP optimally, we then systematically design admission policies that estimate moments of each workload's distribution of future resource usage. Via extensive simulations grounded in a trace from Microsoft Azure, we show that our admission policies lead to a substantial improvement over the simple threshold policy. We then show that substantial further gains are possible if high-quality information is available about arriving workloads. Based on this, we propose an information elicitation approach to incentivize users to provide this information and simulate its effects.

1. Introduction

Cloud computing is a fast expanding market with high competition where small efficiency gains translate to multi-billion dollar profits (Microsoft, 2018). Like many other markets (e.g., ridesharing platforms, kidney exchanges, online labor markets, and display advertising), the efficiency of this market relies on the performance of a matching algorithm (Ashlagi et al., 2019; Ma & Simchi-Levi, 2019; Behnezhad & Reyhani, 2018; Assadi et al., 2017). In the cloud computing case, the matching algorithm matches incoming requests for virtual machines to the hardware that will be used to satisfy them.

Despite the importance of this matching, most cloud clusters currently run at low efficiency. In the cloud domain, low efficiency means low *average utilization* of the cluster (i.e., only a relatively small fraction of resources are actually used by customers at any given time). There are many reasons for this (Yan et al., 2016). These include technical limitations (such as the need to reserve capacity for node failures or maintenance), inefficiencies in scheduling procedures (especially if virtual machines (VMs) might change size or do not use all of their requested capacity), as well as factors that are external to the cluster (such

as fluctuations in overall demand). Another important cause is the nature of many modern workloads: highly connected tasks running on different VMs that should be run on one cluster to minimize latency and bandwidth use (Cortez et al., 2017). In practice, this means that different VMs from one user are bundled together into a *deployment* of interdependent workloads. When the workload of a user changes, his deployment can request a *scale out* in the form of additional VMs or shut some of its active VMs down.

In this paper, we pay special attention to these size changes. Changing deployment sizes mean that providers face the difficult problem of deciding to which cluster to assign a deployment, as a deployment which is small today may, without warning, see a dramatic increase in size that must be accommodated. To get a sense of the difficulty of this problem, consider that, over time, the number of VMs needed by a specific deployment could vary by a factor 10 or even 100, and a request to scale out should almost always be accepted on the same cluster, as denying it would impair the quality of the service, possibly alienating customers. Furthermore, once a deployment is running in one cluster, it would be error prone to move it to a different cluster, making such migrations only feasible as a last resort (e.g., because of hardware failure). Providers consequently hold large parts of each cluster as idle reserves to guarantee that only a very low percentage of these scale out requests is ever denied, leading to relatively low average utilization.

1.1 Cluster Admission Control

We reduce the problem of determining to which cluster to assign a new deployment to the problem of determining, for a particular cluster, whether it is safe to *admit a deployment*, or if doing so would risk running out of capacity if some deployments scale (Cortez et al., 2017). While a lot of research has been done on scheduling *inside* the cluster (Schwarzkopf et al., 2013; Verma et al., 2015; Tumanov et al., 2016; Zhao et al., 2016), the admission problem has not been well studied before. Consequently, cloud providers are often still using simple policies like rejecting all new deployments once a cluster passes a fixed utilization threshold, effectively reserving a percentage of the cluster only for scale outs. These may seem reasonable at first glance, as the law of large numbers might seem to suggest that with many jobs in a large cluster the current utilization would be a good guide to future utilization. But as Cortez et al. (2017) have shown, a relatively small number of deployments account for most of the utilization. This suggests that the *types of deployments* (i.e., small/large, fast/slow scaling, short/long lived etc.) currently in a cluster have a larger impact on the failure probability than is apparent, and policies that only take the current utilization into account are suboptimal.

1.2 Overview of Contributions

We formalize the cluster admission problem as a constrained Partially Observable Markov Decision Process (POMDP) (Smallwood & Sondik, 1973) where each deployment behaves according to some stochastic process and the cluster tries to maximize the number of active compute cores without exceeding its capacity. Since the exact stochastic processes of individual arriving deployments are not known to the cluster, it has to reason about the observed behavior. The large scale of the problem as well as the highly complicated under-

lying stochastic processes make finding optimal policies infeasible, even for the underlying (fully observable) Markov Decision Process and with limited look-ahead horizon.

Since optimally solving this POMDP is not feasible, we next propose a strategy for constructing heuristic policies via a series of simplifying assumptions. These assumptions reduce the highly branching look-ahead space down to the approximation of a random variable using its moments. We then present the currently used threshold policy that does not take probabilistic information into account as well as two new policies that take successively higher moments into account. We fit our model to real-world data from a Microsoft Azure cloud computing center (Cortez et al., 2017) and, via simulations, show that our higher moment policies produce a 30% improvement over current practice, which would translate to hundreds of millions of dollars a year in savings for large cloud providers.

In our basic model, relatively little is known about arriving deployments, so the performance gains we observe from our more sophisticated policies are driven by being able to better condition admission decisions on the current state of the cluster. We next examine how the utilization of the cluster can be further increased if more precise *prior information* about arriving deployments is available. Prior work has explored similar opportunities in the context of resource planning and scheduling in analytics clusters (Jyothi et al., 2016; Rajan et al., 2016). To study the value of prior information, we introduce a simple framework which captures a notion of the quality of information available. Through additional simulations, we quantify how our policies benefit from this additional information. Depending on the quality of information available, the resulting gains increase to 50% – 65% relative to current practice.

Finally, given the importance of the quality of this information, we design a new information elicitation mechanism, with the goal of simultaneously improving the cluster’s utilization as well as the customers’ utility. This requires care to find a design that allows meaningful information to be elicited in an incentive compatible way while being simple for customers to use. To this end, we propose that rather than explicitly asking users to describe the behavior of their deployments, cloud providers instead provide them with the opportunity to group their deployments into customer-defined categories with similar characteristics. The cloud provider can then set a small portion of the fee for a deployment using a pricing rule based on the variance of resource demands of deployments in a category. We show that such variance-based pricing provides users with the right incentives to (a) label their deployments properly (into, e.g., high and low variance deployments) and (b) structure their workloads in a way that helps the cluster run more efficiently. We provide additional simulations to quantify the benefits of an accurate labeling.

In practice, the magnitude of the gains from our approach will depend on many details our simulations elide. However, we believe that our simulations provide a persuasive case that (a) there are substantial economic gains available from using our new admission policies for the process of matching deployments to clusters and (b) there are substantial further improvements possible by using our information elicitation approach to elicit relevant information from customers.

While our work focuses on a specific problem faced by cloud providers, our overall approach is fundamentally about managing the tail risks of a stochastic process. In our case, these are the rare events where the cluster runs out of capacity. Thus, our approach

may also be of interest in other domains where the management of tail risks is important, for example in finance.

1.3 Related Work

There is a large literature on *cluster scheduling and load balancing* (Schwarzkopf et al., 2013; Verma et al., 2015; Tumanov et al., 2016; Wolke et al., 2015). In addition, some work addresses a different notion of admission control to a cluster, namely how to manage queues for workloads which will ultimately be deployed to that cluster (Delimitrou et al., 2013). In our work, however, while studying which deployments to admit to a cluster, we abstract away from the question of exactly which resources should be used, so our research is orthogonal to this prior work on scheduling and load balancing.

Another strand of literature (Yuan et al., 2016; Yuan et al., 2017) studies yet another notion of admission control. While our work focuses on admitting deployments with scaling resource allocation, these works study admitting traffic requests to cloud applications in a way that seeks to optimally trade off expected revenue with energy and bandwidth costs.

There is also a literature that views scheduling through the lens of stochastic online bin packing (Cohen et al., 2019; Song et al., 2014). This literature also deals with issues of changing workloads on possibly overcommitted resources. However, the models in these papers operate at smaller scales and shorter time horizons. At these scales, the key phenomena we study are not present.

One of the core characteristics of the cluster admission problem is that the arrival of a new deployment into the cluster causes an increase in scale outs in the future (i.e., until the deployment dies). This effect of “work creating more work” is also broadly reminiscent of mutually-exciting (Hawkes) processes (Hawkes, 2018). These have also been studied in the context of queuing systems (Daw & Pender, 2018), though there are some important differences (e.g., in our problem only accepted deployments cause scale outs and the rate of additional scale outs varies greatly for different types of deployments).

There is a large literature on market design challenges in the context of the cloud (Kash & Key, 2016). Existing work has studied both queuing models where decisions are made online with no consideration of the future (Abhishek et al., 2012; Dierks & Seuken, 2019) and reservation models which assume very strong information about the future (Azar et al., 2015; Babaioff et al., 2017). Our work sits in an interesting intermediate position where users may have rough information about the types of their deployments. Furthermore, this literature focuses on using prices to determine which jobs should be served and which should not. While our problem is similarly about accepting or rejecting deployments, we do not want to ration through price discrimination. This is because (nearly) every request is ultimately served by the cloud computing provider and whether it goes into this or another cluster is of little consequence for the user.

Other market design work has looked at how multidimensional resources can be fairly divided among deployments. For example, *Dominant Resource Fairness* (Ghodsi et al., 2011) is an approach that has proven useful in practice (Hindman et al., 2011) and has inspired follow-up work more broadly in the literature on fair division (Parkes et al., 2015; Dolev et al., 2012; Gutman & Nisan, 2012; Kash et al., 2014). In our work, we assume that compute cores are the resource bottleneck and we do not model multi-dimensional resource

requirements. Therefore, the considerations studied in the above papers are not present in our work.

Solving POMDPs is a well-studied problem (Smith & Simmons, 2005; Russell & Norvig, 2016; Roy et al., 2005). Unfortunately, finding an optimal policy is known to be PSPACE-complete even for finite-horizon problems (Papadimitriou & Tsitsiklis, 1987). Even finding ϵ -optimal policies is *NP*-hard for any fixed ϵ (Lusena et al., 2001). In our case, the problem is further exacerbated by the existence of side constraints. While *constrained MDPs* are well studied (Altman, 1999; Ono et al., 2012; Haskell & Jain, 2015), comparatively little work has been done on *constrained POMDPs*. In their most general form, constrained POMDPs require that an optimal policy not only maximizes the objective function, but also satisfies bounds on one or more additional cost functions. General (approximation) strategies proposed in the past include linear programming (Poupart et al., 2015; Walraven & Spaan, 2018), piecewise linear dynamic programming (Isom et al., 2008), and point-based value iteration (Kim et al., 2011). Our SLA constraint can be seen as falling into a slightly more restricted sub-category, sometimes called *chance-constrained POMDPs*, where instead of bounding general cost functions, the probability to reach “bad” states has to be kept below some threshold. While such constraints can always be reduced to more general cost function constraints (Khonji et al., 2019), some more tailored approaches have also been proposed in the past, including a mix of online-look ahead and offline risk evaluation (Undurti & How, 2010) and forward search with pruning (Santana et al., 2016). Unfortunately, none of the above-mentioned general or specialized approaches are efficiently applicable when the state space of the underlying MDP is large or, as in our case, partly continuous. Khonji et al. (2019) recently proposed a fully polynomial time approximation scheme (FPTAS) for constant horizon constrained POMDPs. While their algorithm is polynomial in the size of the observation and action spaces, it is exponential in the number of time steps. This makes it not applicable in domains with long time horizons like cluster admission control. While some work has addressed continuous state space POMDPs (Porta et al., 2006; Duff & Barto, 2002; Brooks et al., 2006), none of this prior work is directly applicable to a constrained problem of the size we study in this paper.

2. Preliminaries

In this section, we formally model the cluster admission problem and then introduce a POMDP formulation to solve the provider’s control problem.

2.1 Formal Model

We consider a single cluster in a cloud computing center. A cluster consists of c *cores* that are available to perform work, also called the cluster’s *capacity*. These cores are used by *deployments*, i.e., interdependent workloads that use one or more cores. The set of deployments currently in the cluster is denoted by X , and each deployment $x \in X$ is assigned a number of cores C^x . Any core that is assigned to a deployment is called *active*,

while the remainder are called *inactive*.¹ We do not model the exact placement of cores inside the cluster and in consequence we also do not model the grouping of cores into VMs.

A deployment can request to *scale out*, i.e., increase its number of active cores. Each such request is for one or more additional cores and must be accepted whenever activating the requested number of cores would not make the cluster run over capacity. Following current practice, scale out requests must be granted entirely or not at all. Deployments may shut down some of their cores over time and these cores then become inactive. A deployment *dies* when its number of active cores becomes zero. This can happen in two ways. First, it can die by successively shutting down one core after another until reaching zero active cores. Second, it can die spontaneously by shutting down all of its cores at once; intuitively this models a decision by a user to kill the deployment.²

A deployment x is described by 4 deployment parameters $(C^x, \mu_x, \lambda_x, \sigma_x)$. We have already introduced the size of a deployment C^x . The remaining three parameters are drawn independently from population-wide distributions with PDFs $f_\lambda, f_\mu, f_\sigma$. We now explain how these parameters govern the behavior of the deployment.

At a high level, we assume that the deployments are memoryless (i.e., the basic processes governing a deployment’s behavior only depend on the current state, which results in all processes following Poisson/exponential distributions). This is common in the literature whenever arrival and departure processes are modelled (e.g., in queuing theory), and has been used in previous models of cloud computing (Abhishek et al., 2012; Dierks & Seuken, 2019). Memorylessness is reasonable at cloud scale and simplifies some calculations, but it is not essential for our approach and policies.

Specifically, we assume that each core’s *lifetime* is distributed according to an exponential distribution with parameter μ_x . The *maximum lifetime* of a deployment (i.e., the time between arrival and it spontaneously shutting down all of its remaining cores) is distributed according to an exponential distribution with parameter $\Delta\mu_x$, where Δ is a (population-wide) multiplicative factor. This effectively leads to an average maximum lifetime for the deployment of $\frac{1}{\Delta}$ average core lifetimes. The *number of scale outs* per time unit for the deployment x is distributed according to a Poisson distribution with rate parameter $\lambda_x(\mu_x)^\nu$, where ν is a population-wide parameter. This form of the rate parameter captures the empirical fact that deployments with longer-lived cores scale more slowly than those with short lived cores. The *size of a scale out* is distributed according to one plus a Poisson distribution with parameter σ_x . While this is an approximation on an individual level (VM sizes usually come in powers of 2), it is reasonable at the level of a cluster.

New deployment requests arrive over time and are accepted or rejected according to an *admission policy*.³ The policy must limit the admission of new deployments to ensure that the cluster is not forced to reject a higher percentage of scale out requests than is specified

1. We assume that inactive cores can become active at any time. This means that features such as hardware failure or capacity reserved for maintenance are not modeled. This is a reasonable simplification, as they do not significantly affect the relative utilization of policies.

2. We model death as permanent because with no active cores any future request could be assigned to a different cluster.

3. A rejected deployment is only rejected from this cluster, not from the cloud computing center as a whole. While outside of our model, in practice it then simply gets sent to the next cluster.

by an internal *service level agreement* (SLA) τ .⁴ If a scale out request cannot be accepted because the cluster is already at capacity, one failure for the purpose of meeting the SLA is logged. An optimal policy therefore maximizes the *utilization* of the cluster, i.e., the average number of active cores, while making sure the SLA is observed in expectation.

2.2 The Provider’s Control Problem: POMDP Formulation

The problem the provider is facing when deciding whether to admit a deployment is that the decision must be made under uncertainty regarding future arrivals and the future behavior of deployments. In addition, the provider cannot directly observe the parameters of each deployment’s processes. To understand how a provider can find a policy given this uncertainty, we model the problem as a Partially Observable Markov Decision Process (POMDP) $(\mathbb{S}, \mathbb{A}, \mathbb{R}, \mathbb{T}, \Omega, \mathbb{O})$ whose policy is constrained to meet the SLA τ .

For the POMDP formulation, we assume that time is discrete⁵ and that the problem has a finite time horizon⁶ denoted N . The state space, denoted \mathbb{S} , describes the space of all possible states of the cluster. A state $s \in \mathbb{S}$ contains all information about the cluster’s active deployments $X(s)$ (including, for each deployment x both its current size C^x and its scaling process parameters λ_x, μ_x and σ_x) as well as the deployments that arrived during the current time step. The action set \mathbb{A} consists of individually accepting or rejecting each of the deployments that arrived this time step. The reward function $\mathbb{R}(s) = \sum_{x \in X(s)} C^x$ is the number of active cores in a state s . The transition probability function is denoted $\mathbb{T}(s'|s, a) \forall s' \in \mathbb{S}, \forall a \in \mathbb{A}$. Given a state of the cluster and admission decisions, this function captures the distribution over scale outs, core deaths, and arrivals of new deployments that occur during the next time step. Ω is the set of possible observations and $\mathbb{O} : \Omega \times \mathbb{S} \rightarrow [0, 1]$ an observation model. In our case, the observation model \mathbb{O} is deterministic, but many states share the same observation. For state s , we always observe $\omega \in \Omega$ equal to the sizes of all deployments that are in the cluster and those that arrived with the last state transition.

As is standard, we further denote the cluster’s current knowledge about which state s it is in via a belief state $b \in \mathbb{B}$, i.e., a probability distribution over states. Specifically, a belief state b specifies, for each deployment x that is in the cluster or arrived with the last state transition, its current size C^x and (posterior) distributions $f_\lambda^x, f_\mu^x, f_\sigma^x$ over its scaling process parameters. For a given x , we let $\tilde{x} = (C^x, f_\lambda^x, f_\mu^x, f_\sigma^x)$, i.e., the provider’s belief about the deployment x . A policy π can now be defined as a mapping from belief states to actions.

Whenever the cluster obtains a new observation $\omega \in \Omega$ in time step $n + 1$, the belief state is updated according to the observation and transition models, i.e.,

$$b_{n+1}(s'|b_n, a, \omega) \propto \mathbb{O}(\omega|s') \sum_s \mathbb{T}(s'|s, a) b_n(s). \quad (1)$$

4. This is a cluster-level SLA and not a deployment-level SLA, as in practice, the probability of tail-events such as scale out failures cannot feasibly be evaluated for a single deployment.

5. While deployments can arrive at arbitrary times, it takes time to make the acceptance decision. Thus, there is little loss in discretizing time.

6. Our approach works for any choice of horizon (or even an infinite horizon with average or discounted rewards).

Given this, we can now define two auxiliary functions. We let $g_{n,\pi,b}$ denote the probability density function of the distribution over the states s_n and belief states b_n of the system in time step n , given policy π and starting belief b . Furthermore, we let $h(s_n, \pi(b_n))$ denote the expected percentage of scale outs that fail with the next state transition from a given state-action pair. We can now formulate the provider’s control problem as finding an optimal policy given an SLA.

Problem 1 (Cluster Admission Problem). *The cluster admission problem is to find an optimal policy π for the POMDP $(\mathcal{S}, \mathcal{A}, \mathbb{R}, \mathbb{T}, \Omega, \mathbb{O})$ subject to the following two constraints:*

$$\int_{(s_n, b_n)} g_{n,\pi,b}(s_n, b_n) h(s_n, \pi(b_n)) d(s_n, b_n) \leq \tau \quad \forall \text{ safe } b \forall 0 \leq n < N \quad (2)$$

$$\pi(b) = \text{reject all arrivals} \quad \forall \text{ unsafe } b \quad (3)$$

Here, we call belief state b safe if the policy π_0 which always rejects newly arriving deployments satisfies

$$\int_{(s_n, b_n)} g_{n,\pi_0,b}(s_n, b_n) h(s_n, \pi_0(b_n)) d(s_n, b_n) \leq \tau \quad \forall 0 \leq n < N \quad (4)$$

and unsafe otherwise.

Intuitively, we would like our SLA constraint (2) to hold in every belief state. However, even if we follow an optimal policy, we can reach belief states where (in retrospect) too many deployments have been admitted, such that, even if no new deployments are admitted ever again, the constraint (2) would be violated. Thus, if we would require Equation (2) to hold in all belief states, we would have an infeasible problem. To address this, we do not enforce Equation (2) in *unsafe* belief states (as defined in Problem 1). We instead require the policy to reject all arriving deployments until it reaches a *safe* belief state.⁷

Note that the current time step is not referenced in Equation (2) or (4). This is intentional to avoid horizon effects: a cluster should not aggressively start to accept new deployments close to the end of its lifetime.

3. A Tractable Problem Formulation

Optimal policies for the cluster admission problem (i.e., Problem 1) cannot be calculated in practice for three reasons. First, there is no simple closed form for the state transition probabilities. Second, the state space of the POMDP is very large: consider a cluster with 20,000 cores. It usually has hundreds of deployments, each described by 4 parameters, some of which are continuous. Even discretized, this results in a state space exponential in the size of the cluster. Third, even disregarding unlikely state transitions, the branching factor is large. This renders standard methods that rely on optimizing limited lookaheads infeasible. Therefore, we now present three carefully chosen simplifying assumptions under which we characterize an optimal policy. In Section 4, we use this characterization to design practical admission control policies. *These assumptions and the resulting characterization are only used to systematically design the heuristic policies. Any evaluation of our policies in future sections is still done on the full model without these assumptions.*

7. The requirement to reject *all* deployments is a design decision we revisit in Section 6.

Assumption 1 (No Future Arrivals). *No further deployments arrive after the current timestep.*

This assumption ensures that a policy does not reject deployments simply because better behaved deployments might arrive in the future. In the cloud domain, this behavior is desirable, as even customers with high demand variability must be served by some cluster in the data center.

Assumption 2 (Relaxed Capacity Constraints). *Deployments can scale out even if doing so exceeds the cluster capacity c . For the purpose of defining h , a scale out is considered to fail only if the cluster has already exceeded its capacity.*

With no future arrivals, the cluster's future state only depends on how the sizes of the currently active deployments change. However, if a cluster is full, further scale out requests by deployments are denied, introducing correlations between the future sizes of different deployments. Assumption 2 removes this correlation. In particular, let L_n^x denote the random variable that is the number of active cores of deployment x in time step n . With the first two assumptions, L_n^x is independent of $L_{n'}^{x'}$ for all $x \neq x'$ and all n, n' . The same holds for the random variable $L_n^{\tilde{x}}$ for the provider's belief. This is reasonable because the cluster being full should be rare if the SLA is being met.

Assumption 3 (At Most one Event per Timestep). *In any timestep, at most one event occurs (i.e., at most one deployment scales out, shuts down cores, or arrives to the cluster).*

Since the probability that more than one event occurs in a single time step approaches zero with increased granularity of the time discretization, it is reasonable to assume this.

Using these three assumptions, we can now simplify the problem of determining when the SLA constraint is met. Recall that $\tilde{x} = (C^x, f_\lambda^x, f_\mu^x, f_\sigma^x)$ specifies the provider's belief over a deployment x . In the following, we denote by $A_\pi(b)$ the set of beliefs \tilde{x} over the active deployments in belief state b and the deployments that are accepted with policy π in belief state b .

Proposition 1. *For all policies π , under Assumptions 1, 2 and 3, the following holds:*

$$\int_{(s_n, b_n)} g_{n, \pi, b}(s_n, b_n) h(s_n, \pi(b_n)) d(s_n, b_n) = Pr\left(\sum_{\tilde{x} \in A_\pi(b)} L_n^{\tilde{x}} > c\right) \quad \forall b \forall 0 \leq n < N. \quad (5)$$

Proof. To see that Equation (5) holds, note the following: By Assumption 1, it suffices to consider only the deployments that are currently in the cluster or arrive in the current time step, i.e., $\tilde{x} \in A_\pi(b)$. By Assumption 3, at most one scale out can fail per time step. Thus the left hand side of Equation (5) captures the following: if a scale out occurs in time step n , what is the probability that it fails. By Assumption 2, a scale out fails exactly when $\sum_{\tilde{x} \in A_\pi(b)} L_n^{\tilde{x}} > c$. \square

Using this result, it is straightforward to characterize an optimal policy for the simplified problem.

Corollary 1. *Under Assumptions 1, 2 and 3, an optimal policy π accepts an arriving deployment in belief state b if and only if*

$$Pr\left(\sum_{\tilde{x} \in A_\pi(b)} L_n^{\tilde{x}} > c\right) \leq \tau \quad \forall 0 \leq n < N. \quad (6)$$

Proof. By Assumption 3, it suffices to consider one arrival. By Assumption 1, if an arrival could be accepted without violating the constraint, doing so is optimal. By Proposition 1, the constraint is equal to Inequality (6). \square

Proposition 1 and Corollary 1 show that to implement an optimal policy for the simplified problem it suffices to evaluate the probability that a sum of independent random variables exceeds a threshold. The remaining question now is how to compute or approximate this probability for our complex processes fast enough to allow rapid responses to customer requests.

4. Designing New Admission Control Policies

In this section, we first define the complex random variables L_n^x in terms of simpler random variables that directly arise from the processes. The essence of our approach is to take this definition of L_n^x and use it to compute approximate moments of L_n^x (i.e., approximate summary statistics of the behavior of the random variable). We then use these approximate moments to design new policies.

We can describe L_n^x using the following random variables (which all have a superscript x that we generally omit for brevity). Note that many of these random variables are counterfactual (e.g., the lifetime distribution of a core that may or may not spawn in the first place).

- C is the variable denoting the number of active cores at time step 0.
- Y_i is the random variable denoting the number of scale outs that occur between time step $i - 1$ and time step i , assuming the deployment has not died. It holds that Y_i follows a Poisson Distribution with rate $\lambda(\mu)^\nu$, i.e., $Y_i \sim Pois(\lambda(\mu)^\nu)$.
- $S_{i,l}$ is the (counterfactual) size the l 'th scale out request would have if at least l scale out requests occur between time step $i - 1$ and time step i . It holds that $S_{i,l} \sim Pois(\sigma)$.
- $Z_{n,i,k}$ is the binary random variable denoting whether the k 'th core activated between time steps $i - 1$ and i would still be active in time step n , assuming at least k cores were activated and the deployment has not died. For $i = 0$, this instead refers to whether the k 'th core that is active at timestep 0 is still active at time step n . In either case, as core lifetimes follow an exponential distribution, the probability that a core is still alive after $n - i$ timesteps is given by the complementary CDF of the exponential distribution and thus $Z_{n,i,k} \sim Bernoulli(e^{(i-n)\mu})$.

- D_i is the random variable which is 1 if x would not have died due to a lack of active cores before time step i . It can be defined recursively as follows:

$$D_i = D_{i-1} \left(1 - \prod_{k=1}^C (1 - Z_{i,0,k}) \prod_{j=1}^{i-1} \prod_{k=1}^{\sum_{l=0}^{Y_i} S_{i,l}} (1 - Z_{i,j,k}) \right) \quad (7)$$

$$D_1 = 1 - \prod_{k=1}^C (1 - Z_{1,0,k}) \quad (8)$$

- B_n is the random variable denoting the number of cores that were active at time step 0 and are still active in time step n , which can be calculated as

$$B_n = \sum_{k=1}^C Z_{n,0,k}. \quad (9)$$

- Q_n is the random variable denoting the number of cores activated between time step 0 and time step n that are still active assuming no service termination, i.e.,

$$Q_n = \sum_{i=1}^n \sum_{k=1}^{\sum_{l=0}^{Y_i} S_{i,l}} Z_{n,i,k}. \quad (10)$$

- M_i is the random variable which is 1 if the maximum lifetime of the deployment is at least i and 0 otherwise. It holds that $M_i \sim \text{Bernoulli}(e^{(i-n)\Delta\mu})$.
- Finally, L_n can be calculated as $L_n = M_n D_n (Q_n + B_n)$.

While a ground truth for λ , μ and σ exists, the provider does not know it. Thus, any policy that acts on her belief L_n^x has to also treat the parameters as random variables with $\lambda \sim f_\lambda^x$, $\mu \sim f_\mu^x$, $\sigma \sim f_\sigma^x$. We now turn to the design of our approximate policies.

4.1 Baseline (Zeroth Moment Policy)

Before introducing our new policies, we state the baseline admission control policy that is widely used in practice. It is a myopic policy that simply compares the current number of active cores to a threshold. This policy does not use any information about the set of deployments besides the total number of active cores. It can be viewed as a degenerate case of our approach, as it does not take any probabilistic information about the random variables into account. We therefore also call it a *Zeroth Moment Policy*. Because it uses a limited amount of information, it must be conservative in how many deployments it accepts, since it does not know how often or fast they will scale out.

Definition 1 (Zeroth Moment Policy (Baseline)). *Under a zeroth moment policy π with threshold t , a newly arriving deployment is accepted if, after accepting the deployment, there would be less than t cores active.*

4.2 First Moment Policy

Our first policy approximates the probability of scale out failures (i.e., Equation (5)) by utilizing the first moments, i.e. the expected value of the deployment processes. By Markov's Inequality, for a non-negative random variable L and $c \geq 0$, it holds that

$$Pr(L \geq c) \leq \frac{E[L]}{c}. \quad (11)$$

Such a policy that utilizes Markov's Inequality therefore rejects an arriving deployment when the expected utilization lies above a chosen threshold and otherwise accepts it.

Definition 2 (First Moment Policy). *Under a first moment policy π with threshold t , a newly arriving deployment in belief state b is accepted if, after accepting the deployment, the expected number of active cores would be less than t in all future time steps, i.e.*

$$\sum_{\tilde{x} \in A_\pi(b)} E[L_n^{\tilde{x}}] \leq t \quad \forall 0 \leq n < N, \quad (12)$$

where $E[L_n^{\tilde{x}}]$ is approximated, for example using the approach described in Proposition 2.

As $L_n^{\tilde{x}}$ is a highly complex compound random variable without closed form, exactly evaluating $E[L_n^{\tilde{x}}]$ is infeasible in practice and we therefore employ the following approximation.

Proposition 2. *Assuming M_n , Q_n , B_n and the D_i for all i are uncorrelated and $Z_{i,j,k}$, Y_i , and $S_{i,l}$ are uncorrelated as constituents of D_i , the following holds:*

$$E[L_n] = E[M_n]E[D_n](E[Q_n] + E[B_n]) \quad (13)$$

$$E[Q_n] = \sum_{i=1}^n E[E[Y_1 | \lambda, \mu]E[S_{1,1} | \sigma]E[Z_{n,i,1} | \mu]] \quad (14)$$

$$E[B_n] = CE[Z_{n,i,k}] \quad (15)$$

$$E[D_i] \leq E[D_{i-1}](1 - (1 - E[Z_{i,0,1}])^C \prod_{j=1}^{i-1} (1 - E[Z_{i,j,1}])^{E[Y_i]E[S_{1,1}]}) \quad (16)$$

$$E[D_1] = (1 - (1 - E[Z_{1,0,1}])^C) \quad (17)$$

The proof is provided in Appendix A. It works by direct calculation and applying Jensen's Inequality to D_i . While ignoring some correlations introduces a nontrivial error into the approximation, this is done to ensure that the expectation can be evaluated in linear time. Additionally, any tail bound based on Markov's Inequality is inherently loose, especially for very small τ . This makes it necessary to calibrate t , as simply setting $t = \tau c$ would yield excessively conservative policies (e.g., potentially not admitting any deployment.). Nevertheless, as we will see in Section 5, this approximation still carries enough information for our policies to work well once tuned.⁸

4.3 Second Moment Policy

First moment policies do not take much information about the structure of deployments into account. In a sense they have to always assume the worst possible population mix and

8. Similar observations have been made in the literature on the use of effective bandwidth for admission control in queueing settings (Kelly, 1991; Berger & Whitt, 1998).

run the risk of accepting deployments with low expected size but high variance when close to the threshold. One way around this is to also take the second moment, i.e., the variance of L_n , into account. To address this, we propose to use Cantelli's inequality, a single-tailed generalization of Chebyshev's inequality, to approximate the probability of scale out failures (i.e., Equality (5)). Cantelli's inequality states that, for a real-valued random variable L and $\epsilon \geq 0$, it holds that

$$\Pr(L - E[L] \geq \epsilon) \leq \frac{\text{Var}[L]}{\text{Var}[L] + \epsilon^2}. \quad (18)$$

If we now set $\epsilon = (c - \sum_{x \in X} E[L_n^x])$, we obtain a bound for the probability of running over capacity that takes more information into account than a first moment policy.

Definition 3 (Second Moment Policy). *Under a second moment policy π with threshold ρ , a newly arriving deployment in belief state b is accepted if, after accepting the deployment, the estimated probability of running over capacity would be less than ρ in all further time steps, i.e.*

$$\sum_{\tilde{x} \in A_\pi(b)} E[L_n^{\tilde{x}}] \leq c \quad \forall 0 \leq n < N \quad (19)$$

$$\frac{\sum_{\tilde{x} \in A_\pi(b)} \text{Var}[L_n^{\tilde{x}}]}{\sum_{\tilde{x} \in A_\pi(b)} \text{Var}[L_n^{\tilde{x}}] + (c - \sum_{\tilde{x} \in A_\pi(b)} E[L_n^{\tilde{x}}])^2} \leq \rho \quad \forall 0 \leq n < N \quad (20)$$

where $E[L_n^{\tilde{x}}]$ is approximated using the approach described in Proposition 2 and $\text{Var}[L_n^{\tilde{x}}]$ is approximated using the approach described in Proposition 3.

Proposition 3. *Assuming M_n , Q_n , B_n and the D_i for all i are uncorrelated, the following holds:*

$$\text{Var}[L_n] = E[M_n]^2 \text{Var}[D_n(Q_n + B_n)] + \text{Var}[M_n] E[D_n(Q_n + B_n)]^2 \quad (21)$$

$$+ \text{Var}[M_n] \text{Var}[D_n(Q_n + B_n)] \quad (22)$$

$$\text{Var}[D_n(Q_n + B_n)] = E[D_n]^2 (\text{Var}[Q_n] + \text{Var}[B_n]) + (E[Q_n] + E[B_n])^2 \text{Var}[D_n] \quad (23)$$

$$+ \text{Var}[D_n] (\text{Var}[Q_n] + \text{Var}[B_n]) \quad (24)$$

$$\text{Var}[Q_n] = E[\text{Var}[Q_n | \lambda, \sigma, \mu]] + \text{Var}[E[Q_n | \lambda, \sigma, \mu]] \quad (25)$$

$$\text{Var}[Q_n | \lambda, \sigma, \mu] = \sum_{i=1}^n ((\text{Var}[Y_i | \lambda, \mu] E[S_{i,l} | \sigma]^2 + E[Y_i | \lambda, \mu] \text{Var}[S_{i,l} | \sigma]) E[Z_{n,i,1} | \mu]^2 \quad (26)$$

$$+ E[Y_i | \lambda, \mu] E[S_{i,l} | \sigma] \text{Var}[Z_{n,i,1} | \mu]) \quad (27)$$

$$\text{Var}[B_n] = C E[\text{Var}[Z_{n,i,k} | \mu]] + C^2 \text{Var}[E[Z_{n,i,k} | \mu]] \quad (28)$$

$$\text{Var}[D_n] = E[D_n] - E[D_n]^2 \quad (29)$$

The proof is provided in Appendix A. It works by direct calculation. Note that, since the expectation is approximated as given in Proposition 2, $\text{Var}[D_n]$ carries over any approximation errors from $E[D_n]$. As with first moment policies, the bound given by the inequality is again not tight enough to simply set it to $\rho = \tau$ and ρ has to be tuned.

COMPUTATIONAL OVERHEAD

The computational overhead of the second moment policy depends on the number of future time steps it evaluates and the chosen prior distributions for the provider’s belief state. As long as well-behaved priors are used (e.g., the Gamma priors we use in our simulations), each single rule application is fast. For such priors, updating the estimate for the second moment policy for a single deployment can be done in $O(n)$ where n is the number of evaluated time steps. Whenever a new deployment arrives, the estimate is updated for every active deployment. This leads to a worst case runtime of $O(|X|n)$ where $|X| \leq c$ is the number of active deployments. For multiple clusters this is fully parallelizable at the cluster level because each cluster has its own policy evaluation. Updating the prior of a deployment during runtime has negligible complexity ($O(1)$). A cloud computing center consisting of clusters of capacity c with an arrival rate of L new deployment requests per hour therefore has a computation overhead of at most $O(Lcn)$ each hour, parallelizable into jobs of size $O(n)$.⁹ This means that even relatively large look-ahead horizons n can easily be implemented in practice.

5. Empirical Evaluation

In this section, we evaluate the performance of our admission policies using a model fitted to the real-world data trace of Cortez et al. (2017).

5.1 Data Trace and Fitted Model

Cortez et al. (2017) published a data trace consisting of all deployments that populated a Microsoft Azure datacenter in one month. Since the data set is of limited size and only covers one month, we cannot directly evaluate the policies on the historical deployments. One month is too short to fully evaluate cluster admission policies as many effects only show up after months of usage. Instead, we fit processes to the data we do have, to simulate longer time periods (3 years, in our simulations). We defer evaluations against real deployments to future work.

An in-depth discussion of our fitting procedure can be found in Appendix B. The resulting model utilizes Gamma priors, which are a very general distribution (containing the Chi-squared, Erlang and Exponential distributions as special cases) and fit the data well. The fitted parameters are shown in Table 1. The moment approximation resulting from combining Propositions 2 and 3 with these priors is given in Appendix C. In the following we present the results of our simulations.

5.2 Simulation Setup

We simulate clusters with capacity $c = 20,000$ for a 3-year period with all three policies. An average of 1 new deployment per hour arrives according to a Poisson process. The parameters of each arriving deployment are drawn from the fitted distributions presented

9. If further machine learning is (optionally) employed to obtain an individual prior for arriving deployments (as discussed in Section 6), that computation time would need to be added and depends on the algorithm in question.

Priors	$\mu \approx \text{Gamma}(0.3107, 0.5778)$
	$\lambda \approx \text{Gamma}(0.4907, 0.4496)$
	$\sigma \approx \text{Gamma}(0.2616, 0.0552)$
Global Parameters	$\Delta = 0.119$
	$\nu = 0.673$

Table 1: Fitted processes

in Table 1. Furthermore, we set an SLA of 0.01%.¹⁰ We verify that the SLA is satisfied on average (over runs and months).

Evaluating our first and second moment policies with a three year time horizon and fine-grained time steps is fast enough to be done in real time. However, doing so would take too much computation power to simulate the thousands of years of cluster operation required for our experiments. Therefore, we use the following approach to simulate clusters with a three-year lifespan with a reasonable number of core-hours. We divide the first and second moment policies into 5 subpolicies and only accept a deployment if all subpolicies accept it. The subpolicies have increasingly fine-grained time steps, but each only evaluates a limited look-ahead horizon: 3 years, 1 year, 1 month, 1 week, and 24 hours. Each subpolicy discretizes its time into 600 timesteps. With this approach, simulating a single cluster over its complete lifetime of three years takes between 0.5 and 4 core hours on an Intel Xeon Processor E5-2680 v4.¹¹ We tune the threshold for each policy via binary search, performing 500 runs with each candidate policy parameter, resulting in an overall runtime on the order of 20,000 core hours per policy. Since failures to scale out are focused in the tail of the runs (e.g., with the tuned zeroth moment parameter only about 1% of runs contain any failures), we employ importance sampling to obtain sufficient samples from the tail to guarantee SLA satisfaction with high confidence. Details about our importance sampling approach can be found in Appendix D. The results we report are the average utilization over all runs of the fully tuned policies. To avoid misestimating confidence intervals with biased data, we report bias-corrected and accelerated 95% bootstrap confidence intervals instead of standard errors.¹²

5.3 Results

We now compare the utilization of our policies to the industry baseline zeroth moment policy. The results are summarized in Table 2. The zeroth moment policy obtains its best result with a threshold of $t = 8,864$, i.e., new deployments are accepted whenever less than 8,864 would be active in case of acceptance. This results in an average utilization of 50.45% over the lifetime of the cluster. The first moment policy with threshold $t = 14,223$ increases the utilization by 15.74 percentage points to 66.19%. This constitutes a relative increase in

10. This SLA is somewhat stricter than is typically used in practice, which helps counterbalance our model abstracting away complexities such as fragmentation and node failure.

11. As we simulate each individual deployment in the cluster, the exact runtime depends on whether a given run results in many small or few large deployments populating the cluster.

12. See Efron (1987) for details on the employed bias correction and acceleration techniques. We employ 100,000 re-samples.

Policy	Threshold	Utilization
Zereth Moment	$t = 8,864$	50.45% (48.2, 52.7)
First Moment	$t = 14,223$	66.19% (63.41, 68.94)
Second Moment	$\rho = 0.112$	67.32% (64.35, 70.26)

Table 2: Simulation results showing the performance of the three policies. 95% bootstrap confidence intervals are shown in parentheses.

utilization of 31.2% over the zeroth moment policy. Similarly, the second moment policy with threshold $\rho = 0.112$ achieves a utilization of 67.32%, a relative improvement of 33.44%.

At first sight, it may be surprising that the first and second moment policies achieve similar utilization. However, this can be explained as follows. Under both policies, the overwhelming number of simulated clusters never reject a scale out request. However, in a few runs, too many large, long-lived deployments are accepted in the beginning of a cluster’s lifetime. This leads to many rejections months or even years in the future. Since this happens early in a cluster’s lifetime when not much is known about deployments, the difference between the first and second moment policies is relatively small. This highlights the value of obtaining additional (probabilistic) information about arriving deployments. We study this in the next section.

6. The Value of Deployment-Specific Priors

So far, we have assumed that the cluster does not have any information about arriving deployments, except for the initial number of cores. The acceptance decision therefore had to primarily depend on the state of the deployments that are already in the cluster.

Intuitively, a policy could more precisely control whether accepting a deployment would risk violating the SLA if the policy had more information about the future behavior of the specific deployment. One way to obtain such information would be to use machine learning (ML) based on features of the arriving deployment and past deployment patterns of the submitting user (Cortez et al., 2017). While evaluating particular ML algorithms is beyond the scope of this paper, we evaluate the effect that different levels of available information have. To do this, we need to parameterize the level of knowledge. For this we assume that the cluster simply gets passed some number of observations from each true scaling process distribution of each arriving deployment.¹³

6.1 Improving the Handling of Short-Lived Deployments

Our moment policies as defined so far cannot yet make optimal use of this additional prior information. While an optimal policy with good prior information would balance the

13. As we have used conjugate prior distributions in our model, this approach matches the standard interpretation of parameters of the posterior distribution in terms of “pseudo observations.”

admission of long-lived and short-lived deployments to keep the utilization more stable over time, the moment policies always accept new deployments on a first-come first-served basis until their constraints are violated. This means that if many very long-lived, slow-scaling deployments arrive in the beginning, the cluster sometimes quickly reaches unsafe belief states in which it stops accepting any new deployments, but for which the critical event lies months or even years in the future. While stopping the admission of new deployments in such a situation is reasonable when no prior information about arriving deployments is available, *with* prior information the policy might know that some arriving deployments will almost surely be dead by the time the cluster has filled up. To make use of this, we now present a heuristic modification of our moment policies such that the resulting policy is allowed to accept deployments that only have a marginal impact on the possible SLA violation, even in unsafe states. As a simple condition for this, we call a deployment *marginal in timestep n* if its expected size is smaller than 10^{-5} (i.e., $E[L_n^x] < 10^{-5}$).¹⁴

Definition 4 (Marginal Heuristic). *Under a first or second moment policy π with the marginal heuristic, a newly arriving deployment x in belief state b is accepted if in each future time step $n < N$, after accepting the deployment, either the underlying moment policy’s condition is satisfied or the arriving deployment x is marginal, i.e., $E[L_n^x] < 10^{-5}$.*

Going forward, we use the marginal heuristic, unless explicitly noted. It should be pointed out that this heuristic does not have any effect when the cluster does not have good prior information about arriving deployments. With only the global prior, no deployment is marginal for any future timestep $n < N$.

6.2 Simulation Results

In this section, we present simulation results to demonstrate the value of deployment-specific priors. We simulate the first and second moment policies (with marginal heuristic), now with four different levels of prior information: 0, 1, 5, and 50 observations. Otherwise, we use the same simulation setup as in Section 5. The results are shown in Figure 1.¹⁵ We see that having prior knowledge equivalent to even a single observation improves utilization significantly, resulting in a utilization of 75% and 79.5% for the first and second moment policies, respectively. Better priors lead to even better utilization, with the second moment policy reaching a utilization of 83.8% with 50 observations.

While it is infeasible to calculate the utilization corresponding to an optimal solution of the POMDP, we have derived an *upper bound* of 92.1% by analyzing a policy that does not have to satisfy any SLA and accepts all arriving deployments.¹⁶ Thus, the second moment policy with good prior information achieves more than 90% of the theoretically achievable as given by this (unreachable) upper bound, while delivering a relative increase in utilization

14. Given that policy thresholds (i.e., t and ρ) are calibrated, the exact value here does not have much of an impact on performance. It is only important that the threshold is well below the SLA, yet large enough that the estimate for short lived jobs falls below this threshold even with a small number of samples.

15. We also simulated our policies without the marginal heuristic (see Appendix E), and we observe the same general patterns. As one would expect, without the marginal heuristic, the achieved utilization is somewhat smaller (especially with good priors).

16. Most of the remaining gap to 100% here arises from the cluster starting empty and requiring a few months to fill up, as well as from the fact that dying deployments cannot immediately be replaced.

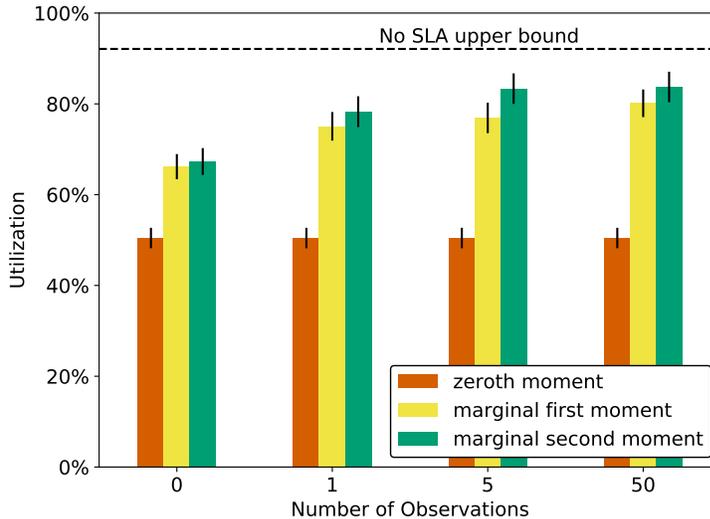


Figure 1: Performance of different policies depending on prior information (error bars indicate 95% bootstrap confidence intervals)

of 24.48% above the same policy without prior information and a 66.1% increase over the baseline (i.e., zeroth moment) policy. This shows both the power of our policies and the great importance of taking all available prior information about arriving deployments into account.

7. An Elicitation Mechanism to Improve Priors

Given the importance of the quality of prior information that we established in the last section, in this section, we use techniques from mechanism design to improve this quality. Our approach assumes that users do not typically submit deployments with arbitrary parameters. Instead, they may have a small number of different *types* of deployments. However, the typical mechanism design approach of using a direct revelation mechanism, where customers reveal their full type, seems problematic. First, it may be very cumbersome from a user interface perspective. Second, customers may not have such a detailed understanding of their deployments and thus would run the risk of being penalized for a “misreport.” Instead, we seek a design that allows meaningful information to be elicited in an incentive compatible way while being simple for customers to use. To this end, we propose that rather than explicitly asking users to describe the behavior of their deployments, cloud providers instead provide them with the opportunity to group them into customer-defined categories of roughly similar deployments. Learning priors for each individual category then results in more precise priors and higher utilization. To incentivize such grouping, the cloud provider can set a small portion of the fee for a deployment using a pricing rule based on the variance of resource demands of deployments in a category. This motivates users to group deploy-

ments together that have similar behavior (e.g., short-lived vs. long-lived deployments). We now first present and analyze such a variance-based pricing mechanism and then evaluate the potential utilization gains this mechanism may produce via additional simulations.

7.1 Variance-Based Payment Rule

Typically, users are charged a fixed payment per hour for each core their deployment uses. With a variance-based payment rule, we add a small additional charge based on the variance of the estimate for the deployment’s scaling process and allow users to *label* the type of their deployments, resulting in an hourly *variance-based payment rule* $q(x)$ of the form:

$$q(x) = \kappa_1 C^x + \kappa_2 \text{Var}(x), \quad (30)$$

where κ_1 and κ_2 are price constants and $\text{Var}(x)$ is an estimate of the variance of the deployment.¹⁷ A payment rule of this form incentivizes users to assign similar labels to similar deployments to minimize the estimated variance.¹⁸

To see this, consider a user who has two types of deployment, x and y , with true variances $\text{Var}(x)$ and $\text{Var}(y)$. He could now simply submit the deployments under a single label. For the provider, this means that each submitted deployment is of either type with a certain probability, which increases the variance of her prediction. But if the user would label each deployment appropriately with either “ x ” or “ y ,” then the provider would know for each arriving deployment which type it is, reducing variance and therefore the need to reserve capacity. The following proposition, which is immediate from the law of total variance, shows that, at least in the long run, labeling his deployments also reduces a user’s payments.

Proposition 4. *Let z be the mixture that results from submitting one of two types of deployments x, y chosen by a Bernoulli random variable $\alpha \sim \text{Bernoulli}(p_\alpha)$, i.e., such that z is of type x with probability p_α and of type y with probability $1 - p_\alpha$. Then it holds that*

$$p_\alpha \text{Var}(x) + (1 - p_\alpha) \text{Var}(y) \leq \text{Var}(z) \quad (31)$$

Proof. Since z has finite variance, the law of total variance states:

$$\text{Var}(z) = E[\text{Var}(z|\alpha)] + \text{Var}(E[z|\alpha]) \quad (32)$$

$$\geq E[\text{Var}(z|\alpha)] \quad (33)$$

$$= p_\alpha \text{Var}(x) + (1 - p_\alpha) \text{Var}(y) \quad (34)$$

□

Proposition 4 shows that the user would be better off by splitting the mixture and submitting the deployments under separate labels, directly resulting in the following corollary.

17. We suggest to define the “variance of a deployment” here as the maximum variance of the size in any future timestep, i.e., $\text{Var}(x) = \max_n \text{Var}[L_n^x]$. However, our concept of variance-based pricing is not limited to any particular notion of “variance of a deployment”: different notions of “variance of a deployment” simply carry different incentives for users to structure their deployments. An analysis of these incentives is beyond the scope of this paper.

18. Since the benefits come from improving the provider’s estimate about future deployments, there must be a sufficiently large number of deployments in each category.

Corollary 2. *Under any variance-based payment rule $q(x)$ of the form given in Equation (30) with $\kappa_2 > 0$, it is a dominant strategy for users with multiple deployment types to label deployments by type.*

Note that this corollary abstracts away issues of learning and non-stationary strategic behavior; but for reasonable learning procedures we expect a consistent labeling to lead to lower variance than a mixture while learning. Further, this approach not only gives the user correct incentives to reveal the desired information, but actually incentivizes him to improve the performance of the system. In particular, another way he can lower his payment under this scheme (outside the scope of our model) is to design his deployments in such a way that they have lower variance in their resource use. Since more predictable deployments would allow the policy to maintain a smaller buffer, this provides an additional benefit to the system’s utilization.

Remark. *The per core-hour payments of a user can be based on a-priori or a-posteriori estimates of a deployment’s variance. With an a-priori estimate, the user knows his payment (per core-hour) before he starts his deployment, which can be very important for certain users. On the other hand, such a payment rule invites strategic deployment submissions: a user could submit a number of small low-variance deployments before submitting a large high-variance deployment, with the goal of reducing the provider’s estimate and thus his payment for the large deployment. The provider could mitigate the potential gain of such a manipulation by carefully choosing the estimation procedure, so it is unclear how frequent and successful such manipulations would be in practice. With an a-posteriori estimate (i.e., the user’s hourly payment is based on the variance estimate of the deployment at the end of the hour), such strategic deployment submissions could be made unprofitable. However, now users do not know their exact payments in advance. To cater to users that require a fixed price before submitting a deployment, the provider might want to set an upper limit on prices and advertise lower prices as discounts. Thus, which type of estimate is optimal for a given provider depends on the requirements of her user base.*

How much any given user could ultimately save by labeling his deployments mostly depends on how different his deployment types are and on how high the provider sets the charge for variance. A user whose deployments are quite uniform will not save much, while a user with some deployments which never scale and some that scale a lot can potentially save a lot. Note that how much the provider should charge is not immediately clear. While she would want to set a high price to put a strong incentive on users, she also has to keep the competition from other providers in mind. At what point the loss of market share outweighs the gain in utilization is an intriguing problem we leave for future work.

7.2 Simulation Results

To illustrate the potential gains in utilization of such a variance-based payment rule, we consider a setting where all users have two deployment types, drawn independently from the same population distribution as in Section 5. We assume that each user only submits a single deployment and then departs, but that the provider has 5 prior observations each from every user’s two deployment types. Otherwise, the simulation setup is again the same as in Section 5. In this setting, we contrast the utilization of a provider employing a second moment

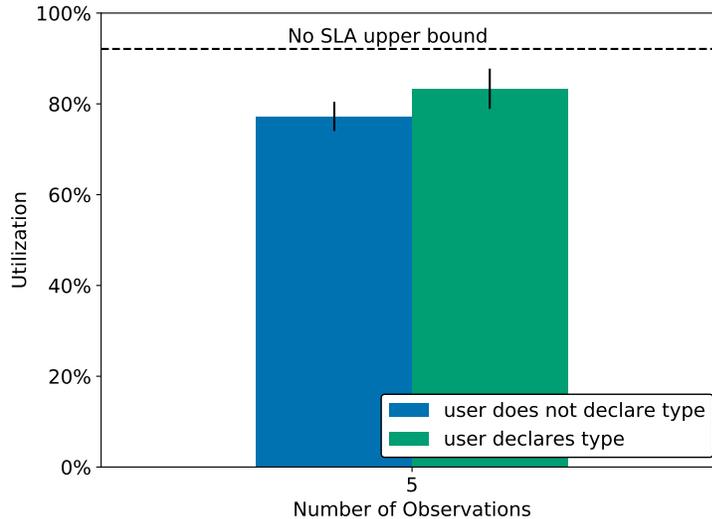


Figure 2: Performance of the second moment policy with two deployment types per user and 5 observations (error bars indicate 95% bootstrap confidence intervals)

policy with and without employing a variance-based payment rule. With the variance-based payment rule (and users consequently declaring their type), the setting becomes equivalent to the one presented in Section 6 with 5 observations. When the variance-based payment rule is *not* used (and users consequently do not label their deployments), we assume that the provider updates her belief for both types independently and evaluates her second moment policy on the mixture.

As we can see in Figure 2, when users do not label their deployments, this yields a utilization of 77%. In contrast, when users do label their deployments, then (as expected) the utilization increases to 83%. This shows that, from a cluster point of view, employing a variance-based payment rule leads to a sizable increase in utilization.

8. Conclusion

We have studied the problem of cluster admission control for cloud computing, where accepting demand now causes unrejectable demand in the future. The optimal policy is given as the solution to a very large constrained POMDP, which is infeasible to solve. In practice, simple threshold policies are employed for admission control. In contrast, we have proposed multiple more sophisticated policies. Our results demonstrate that the utilization can be increased by approximately 30% just from learning about deployments while they are active in the cluster. Furthermore, we have shown that this can be improved to a 50 – 65% gain if better prior information about arriving deployments is available, for example through learning or elicitation techniques. Even though the realized gains in practice are likely to be somewhat lower due to practical engineering constraints (e.g., the need to handle node

outages), they should still be sizable. At cloud scale, even savings of a few percent translate to many hundreds of millions of dollars, and any dollar saved directly translates to a gross profit increase for the cluster provider.

Our work points to a number of interesting future research directions. We have only looked at cluster admission policies at the level of a single cluster, abstracting away the question of which cluster should be chosen, implicitly assuming a first-fit or random-fit heuristic. Future research should look at the question whether filling all clusters with the same mixture of deployments is reasonable or if dedicating different clusters to different types of deployments could be used to further increase utilization. A related direction is that our model and policies assume that deployment behavior does not change during runtime. While this is a reasonable approximation for many deployments, some long running deployments might exhibit more involved life cycles in practice. One way for policies to account for this is to discount past observations. Similarly, while memorylessness is a reasonable assumption in cloud computing, a small subset of users might exhibit more involved scaling behavior. While it would likely be infeasible to learn individual process distributions for individual deployments (as opposed to only learning parameters), future work should study the merit and feasibility of instead employing mutually-exciting (Hawkes) processes to perturb process parameters whenever an event (i.e., a scale out or a core death) is observed. Importantly, with such an approach, additional parameters that determine the type and strength of a perturbation would also need to be learned.

There are also open questions regarding mechanism design in the cloud domain. In subsequent work, Dierks and Seuken (2020) have already analyzed the competitive effects of employing the variance-based payment rule we proposed in this paper in a duopoly model. They find that, in equilibrium, while using a variance-based payment rule weakly increases welfare, the effects on the providers’ profits are ambiguous. However, their analysis does not take into account that a variance-based payment rule can yield better priors about submitted deployments and thus improve the efficiency of the provider’s admission policy. Future work could explore an alternative economic design, where the provider offers a menu of two alternatives to her users: the standard alternative, where deployments can always scale out; and a cheaper alternative, where deployments are not allowed to scale out (or can scale only with a “best effort” guarantee). Such an approach could be viewed as implicitly selling finance-style options on the ability to scale out.

Acknowledgments

An early version of this work has appeared as a 6-page extended abstract in the proceedings of the 14th Workshop on the Economics of Networks, Systems and Computation (NetEcon’19). We thank the associate editor of JAIR and anonymous reviewers for their very helpful comments. A substantial portion of this work was done while Ian Kash was employed by Microsoft. This work was further supported by Microsoft Research through Grant 2015-058 and by the US National Science Foundation through Award CCF-1934915.

Appendix A. Omitted Proofs

Proof of Proposition 2. • If M_n , D_n , Q_n and B_n are uncorrelated, it holds by linearity and multiplicativity of the expected value for uncorrelated random variables:

$$E[L_n] = E[M_n]E[D_n](E[Q_n] + E[B_n]) \quad (35)$$

- Q_n : For the expectation of Q_n it holds:

$$E[Q_n] = E\left[\sum_{i=1}^n \sum_{k=1}^{\sum_{l=0}^{Y_i} S_{i,l}} Z_{n,i,k}\right] \quad (36)$$

$$= \sum_{i=1}^n E\left[\sum_{k=1}^{\sum_{l=0}^{Y_i} S_{i,l}} Z_{n,i,k}\right] \quad (37)$$

$$= \sum_{i=1}^n E\left[E\left[\sum_{k=1}^{\sum_{l=0}^{Y_i} S_{i,l}} Z_{n,i,k} \mid \lambda, \sigma, \mu\right]\right] \quad (38)$$

$$E\left[\sum_{k=1}^{\sum_{l=0}^{Y_i} S_{i,l}} Z_{n,i,k} \mid \lambda, \sigma, \mu\right] = E\left[E\left[\sum_{k=1}^{\sum_{l=0}^{Y_i} S_{i,l}} Z_{n,i,k} \mid \sum_{l=0}^{Y_i} S_{i,l}\right] \mid \lambda, \sigma, \mu\right] \quad (39)$$

$$= E\left[\sum_{l=0}^{Y_i} S_{i,l} \mid \lambda, \sigma, \mu\right] E[Z_{n,i,1} \mid \lambda, \sigma, \mu] \quad (40)$$

$$= E[Y_1 \mid \lambda, \mu] E[S_{1,1} \mid \sigma] E[Z_{n,i,1} \mid \mu]. \quad (41)$$

- B_n : By definition, it holds

$$E[B_n] = E\left[\sum_{j=1}^C Z_{n,0,k}\right] = CE[Z_{n,0,k}]. \quad (42)$$

- D_i : If all $Z_{i,j,k}$, Y_1 and $S_{i,l}$ are uncorrelated, it holds

$$E[D_i] = E[D_{i-1}(1 - \prod_{j=0}^{i-1} \prod_{k=0}^{\sum_{l=0}^{Y_j} S_{j,l}} (1 - Z_{i,j,k}))] \quad (43)$$

$$= E[D_{i-1}(1 - E[\prod_{j=0}^{i-1} \prod_{k=0}^{\sum_{l=0}^{Y_j} S_{j,l}} (1 - Z_{i,j,k})])] \quad (44)$$

$$= E[D_{i-1}(1 - E[E[\prod_{j=0}^{i-1} \prod_{k=0}^{\sum_{l=0}^{Y_j} S_{j,l}} (1 - Z_{i,j,k}) \mid Y, S]])] \quad (45)$$

$$= E[D_{i-1}(1 - E[\prod_{j=0}^{i-1} \prod_{k=0}^{\sum_{l=0}^{Y_j} S_{j,l}} (1 - E[Z_{i,j,k}])])] \quad (46)$$

$$= E[D_{i-1}(1 - E[\prod_{j=0}^{i-1} (1 - E[Z_{i,j,k}])^{\sum_{l=0}^{Y_j} S_{j,l}}])] \quad (47)$$

$$\leq E[D_{i-1}(1 - \prod_{j=0}^{i-1} (1 - E[Z_{i,j,k}])^{E[\sum_{l=0}^{Y_j} S_{j,l}]})] \quad (48)$$

$$= E[D_{i-1}(1 - \prod_{j=0}^{i-1} (1 - E[Z_{i,j,k}])^{E[Y_1]E[S_{1,1}]})] \quad (49)$$

$$E[D_1] = (1 - (1 - E[Z_{1,0,1}])^C) \quad (50)$$

where the third line follows by the law of total probability and the 6th by Jensen's Inequality. \square

Proof of Proposition 3. With M_n, D_n, Q_n and B_n uncorrelated, it holds for the variance of L_n :

$$\text{Var}[L_n] = \text{Var}[M_n D_n (Q_n + B_n)] \quad (51)$$

$$= E[M_n]^2 \text{Var}[D_n (Q_n + B_n)] + \text{Var}[M_n] E[D_n (Q_n + B_n)]^2 \quad (52)$$

$$+ \text{Var}[M_n] \text{Var}[D_n (Q_n + B_n)] \quad (53)$$

and further:

$$\text{Var}[D_n (Q_n + B_n)] = E[D_n]^2 (\text{Var}[Q_n] + \text{Var}[B_n]) + (E[Q_n] + E[B_n])^2 \text{Var}[D_n] \quad (54)$$

$$+ \text{Var}[D_n] (\text{Var}[Q_n] + \text{Var}[B_n]). \quad (55)$$

For the variance of Q_n it holds:

$$\text{Var}[Q_n] = E[\text{Var}[Q_n | \lambda, \sigma, \mu]] + \text{Var}[E[Q_n | \lambda, \sigma, \mu]] \quad (56)$$

and

$$\text{Var}[Q_n | \lambda, \sigma, \mu] \quad (57)$$

$$= \sum_{i=1}^n \left(\text{Var} \left[\sum_{l=0}^{Y_i} S_{i,l} | \lambda, \mu, \sigma \right] E[Z_{n,i,1} | \mu]^2 + E \left[\sum_{l=0}^{Y_i} S_{i,l} | \lambda, \mu, \sigma \right] \text{Var}[Z_{n,i,1} | \mu] \right) \quad (58)$$

$$= \sum_{i=1}^n \left((\text{Var}[Y_i] E[S_{i,l}]^2 + E[Y_i] \text{Var}[S_{i,l}]) E[Z_{n,i,1}]^2 + E[Y_i] E[S_{i,l}] \text{Var}[Z_{n,i,1}] \right) \quad (59)$$

by the law of total variance.

For B_n we can now use the law of total variance to obtain:

$$\text{Var}[B_n] = \text{Var} \left[\sum_{j=1}^C Z_{n,i,k} \right] \quad (60)$$

$$= E[\text{Var} \left[\sum_{j=1}^C Z_{n,i,k} | \mu \right]] + \text{Var} [E \left[\sum_{j=1}^C Z_{n,i,k} | \mu \right]] \quad (61)$$

$$= E[C \text{Var}[Z_{n,i,k} | \mu]] + \text{Var}[C E[Z_{n,i,k} | \mu]] \quad (62)$$

$$= C E[\text{Var}[Z_{n,i,k} | \mu]] + C^2 \text{Var}[E[Z_{n,i,k} | \mu]] \quad (63)$$

Lastly, for D_n , note that $E[D_n^2] = E[D_n]$ because $D_n \in \{0, 1\}$. It follows

$$\text{Var}[D_n] = E[D_n] - E[D_n]^2. \quad (64)$$

\square

Appendix B. Data Trace

To have a better understanding of the scaling behavior of real deployments and to create a model suitable for simulating clusters, we fitted the behavior of deployments to a real-world data trace. The particular data trace we use was published by Cortez et al. (2017). This dataset consists of one month of data of internal Microsoft Azure jobs. It contains 35,576 deployments,¹⁹ though only 29,757 of these deployments arrived during the observed time period. Since we want to fit distributions with the goal of simulating arriving deployments, only these 29,757 deployments can be used for most of our fitting. The deployments that arrived before the beginning of the observed period of time cannot be used when making maximum likelihood estimations, because for start times before the observed period of time, only longer lived deployments survived to be observed. Including them would strongly skew the fit. The 29,757 deployments activated 4,317,961 cores, out of which 4,211,926 became inactive again during the observed month. The exact lifetime of the remaining cores (i.e., the length of time between becoming active and then inactive again) is not known; instead we only have a lower bound on it (i.e., our observation is Type I censored: see for example (NIST, 2012)). Thus, for cores where we only have a lower bound on the lifetime we use the cdf in our likelihood function while for cores whose lifetime is known we use the pdf.

B.1 Fitting on the Deployment Level

We first fit arrival and departure processes for each individual deployment. In keeping with the Markov assumption, we fit a Poisson distribution to the scale out rate of each deployment, while we fit an exponential distribution to the lifetime of cores for each deployment for which at least one core became inactive during the observed time period. Note that while we model the cluster admission problem as a discrete-time POMDP, the processes are fit in continuous time. This is more general and avoids imprecisions introduced by time discretization. To fit the size of a scale out, we also used a Poisson distribution (plus 1, as scale outs must have at least one core).²⁰ We further assume that each deployment, had it lived forever, at some point would have made a scale out request for more than 1 core. Since we did not observe these scale-outs and therefore cannot make a direct likelihood fit, we introduce two parameters P_1 and P_2 to represent them. We assume that the scale out rate of deployments that never scaled out (some because they died, but many simply because the observation period of the dataset ended) is equal to the value for which not observing a scale out has probability P_1 . We equivalently set the scale out size for deployments that never increase their size by more than 1 core during one scale out event according to P_2 . We calibrated P_1 and P_2 by minimizing the (discrete) Cramér-von Mises distance of the size of deployments between samples drawn from our fitted model and the data set. The optimal distance is 0.1585 and an overlay of both cumulative distribution functions can be seen in Figure 3. Note that most of the remaining distance does not seem to be caused

19. In contrast to (Cortez et al., 2017) we did not consolidate all deployments a single user runs on a certain day into one. This is because cores that get requested as a new deployment do not need to be accepted on the same cluster.

20. As the Poisson distribution is single-parameter and its variance cannot be set independent of the average size, this is not a particularly good fit for users with large but consistent scale out sizes. However, its simplicity avoids overfitting on the often low number of samples per deployment and it results in a good fit on the population level.

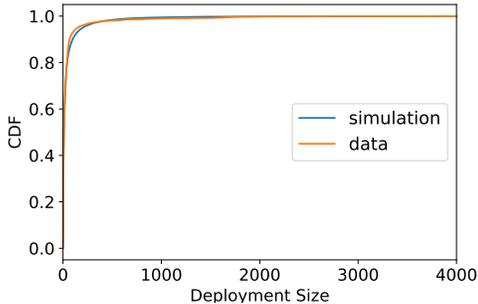


Figure 3: CDF over number of deployments of all sizes

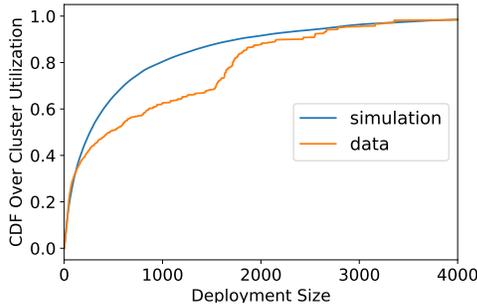


Figure 4: CDF over utilization percentage from deployments of all sizes

by limitations of our model or fitting procedure, but by limitations of the dataset. The dataset, while relatively large, still does contain a somewhat small selection of deployments from the tail. More importantly, it only contains *internal* Azure deployments, so the types of workloads are limited. As such, it contains few deployments of sizes between 100 and 1500, but a relatively large number of deployments of sizes between 1500 and 2000. This effect is visualized in Figure 4, which shows the CDF over the percentage of utilization in the cluster coming from deployments of different sizes for both our model and the dataset.

B.2 Fitting on the Population Level

With the distributions for each deployment in place, we now fit Gamma distributions for the population. The parameters of the processes for each arriving deployment are drawn from these populations. As the data was skewed, positive, and not really heavy tailed, a Gamma distribution is a natural and very general candidate (containing the Chi-squared, Erlang and Exponential distributions as special cases), with the added benefit of being conjugate prior to the deployment processes. The resulting model and parameters from our fits are shown in Table 3. While the scale out size is fit directly to the samples, scale out rate and core lifetime are highly correlated. The longer a deployment’s cores live, the lower the rate at which new cores arrive, as can be seen in Figure 5. This shows that deployments with long lived cores do not necessarily have more active cores. To account for this, we fit the power law relationship ν between scale out rate and lifetime, i.e., we fitted the prior distribution on scale out rates multiplied by the respective core lifetimes taken to the power

Priors	$\mu \approx \text{Gamma}(0.3107, 0.5778)$
	$\lambda \approx \text{Gamma}(0.4907, 0.4496)$
	$\sigma \approx \text{Gamma}(0.2616, 0.0552)$
Global Parameters	$\Delta = 0.119$
	$\nu = 0.673$

Table 3: Fitted processes

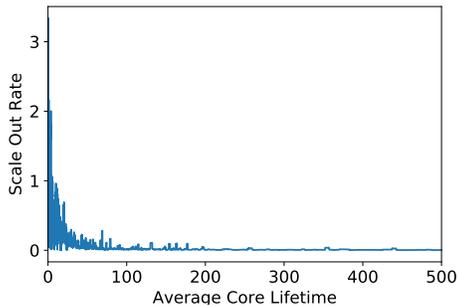


Figure 5: Scale out rate as a function of average core lifetime

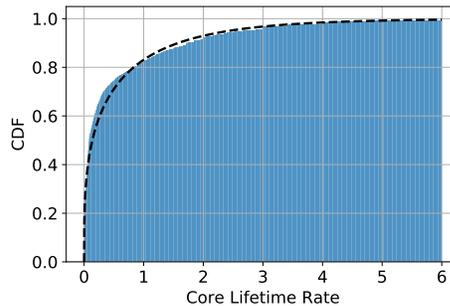


Figure 6: Distribution of the core lifetime rate parameter

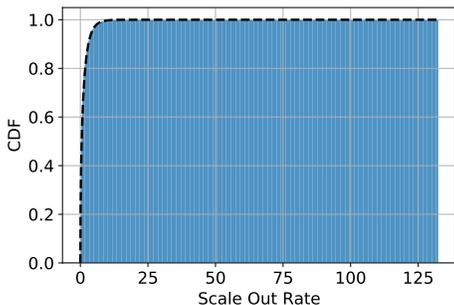


Figure 7: Distribution of the scale out rate parameter

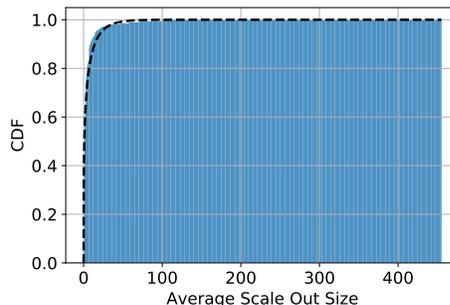


Figure 8: Distribution of the scale out size parameter

of ν . We have chosen ν such that the mean absolute distance between normalized scale out rate of each deployment and the average (normalized) scale out rate is minimized.

To visualize the fitted distributions, Figure 6 shows the CDF of the Gamma distribution for the lifetime parameter, overlaid over the normalized cumulative histogram of the fitted rates of the sample deployments.

Figure 7 shows the CDF for the normalized scale out rates over the relevant cumulative histogram. The actual scale out rate of a sampled deployment is now simply the normalized scale out rate multiplied by the average core lifetime. Figure 8 shows the fitted CDF for the scale out size parameter over its cumulative histogram.

Deployment Shutdown. While most deployments in the dataset die because they have zero active cores, 5,980 of the 22,241 deployments that both arrive and die during the observed period seem to get actively shut down. By this we mean that they had at least 3 VMs that all shut down simultaneously. This would be highly unlikely if deployments only die when cores or VMs become inactive independently. To capture such behavior we fit an exponential distribution over the number of expected core lifetime deployments lived. The

maximal lifetime of deployments that did not get shut down was assumed to be censored to their realized lifetime.

Appendix C. Moment Approximation with Gamma Priors

Proposition 5. *When $Y_i \sim \text{Pois}(\lambda\mu^\nu)$, $\lambda \sim \text{Gamma}(a, b)$, $Z_{n,i,j} \sim \text{Bernoulli}(e^{(i-n)\mu})$ (Bernoulli over complementary CDF of an exponential distribution), $\mu \sim \text{Gamma}(\mathbf{a}, \mathbf{b})$, $S_{i,l} \sim \text{Pois}(\sigma)$, $\sigma \sim \text{Gamma}(\alpha, \beta)$, and $M_i \sim \text{Bernoulli}(e^{(i-n)\Delta\mu})$ it holds:*

$$E[Q_n] = \frac{a}{b} \frac{\alpha + \beta}{\beta} \frac{\Gamma(\mathbf{a} + \nu)}{\Gamma(\mathbf{a})} \sum_{i=1}^n \frac{\mathbf{b}^{\mathbf{a}}}{(n + \mathbf{b} - i)^{\mathbf{a} + \nu}} \quad (65)$$

$$E[D_i] \leq E[D_{i-1}] \left(1 - \prod_{j=0}^{i-1} \left(1 - \left(1 + \frac{i-j}{\mathbf{b}}\right)^{-(\mathbf{a})} \frac{\mathbf{a}}{\mathbf{b}} \frac{\alpha}{\beta}\right)\right) \quad (66)$$

$$E[Z_{n,i,1}] = \frac{\mathbf{b}^{\mathbf{a}}}{(n + \mathbf{b} - i)^{\mathbf{a}}} \quad (67)$$

$$E[M_n] = \frac{\mathbf{b}^{\mathbf{a}}}{(\Delta n + \mathbf{b})^{\mathbf{a}}} \quad (68)$$

and

$$\text{Var}[Q_n] = \mathbf{b}^{\mathbf{a}} \frac{\Gamma(\mathbf{a} + \nu)}{\Gamma(\mathbf{a})} \left[\left(\frac{a}{b} \frac{\alpha}{\beta^2} + \frac{\alpha + \beta^2}{\beta} - 1 \right) \right] \quad (69)$$

$$\sum_{i=1}^n \frac{1}{(2n + \mathbf{b} - 2i)^{\mathbf{a} + \nu}} + \frac{a}{b} \frac{\alpha + \beta}{\beta} \sum_{i=1}^n \frac{1}{(n + \mathbf{b} - i)^{\mathbf{a} + \nu}} \quad (70)$$

$$+ \left(\frac{a^2}{b} \frac{\alpha + \beta^2}{\beta} + \left(\frac{a^2}{b} \frac{\alpha}{\beta^2} + \frac{a}{b^2} \frac{\alpha + \beta^2}{\beta} + \frac{a}{b^2} \frac{\alpha}{\beta^2} \right) \right) \quad (71)$$

$$\left[\mathbf{b}^{\mathbf{a}} \frac{\Gamma(\mathbf{a} + 2\nu)}{\Gamma(\mathbf{a})} \sum_{1 \leq i \leq j \leq n} \frac{1}{(2n + \mathbf{b} - i - j)^{\mathbf{a} + 2\nu}} \right] \quad (72)$$

$$- \left(\mathbf{b}^{\mathbf{a}} \frac{\Gamma(\mathbf{a} + \nu)}{\Gamma(\mathbf{a})} \right)^2 \sum_{1 \leq i \leq j \leq n} \frac{1}{(n + \mathbf{b} - i)^{\mathbf{a} + \nu}} \frac{1}{(n + \mathbf{b} - j)^{\mathbf{a} + \nu}} \quad (73)$$

$$+ \left(\frac{a^2}{b} \frac{\alpha}{\beta^2} + \frac{a}{b^2} \frac{\alpha + \beta^2}{\beta} + \frac{a}{b^2} \frac{\alpha}{\beta^2} \right) \left[\frac{\Gamma(\mathbf{a} + \nu)}{\Gamma(\mathbf{a})} \sum_{i=1}^n \frac{\mathbf{b}^{\mathbf{a}}}{(n + \mathbf{b} - i)^{\mathbf{a} + \nu}} \right]^2 \quad (74)$$

$$\text{Var}[Z_{n,i,1}] = \frac{\mathbf{b}^{\mathbf{a}}}{(n + \mathbf{b} - i)^{\mathbf{a}}} \left(1 - \frac{\mathbf{b}^{\mathbf{a}}}{(n + \mathbf{b} - i)^{\mathbf{a}}}\right) \quad (75)$$

$$CE[\text{Var}[Z_{n,i,k}|\mu]] + C^2 \text{Var}[E[Z_{n,i,k}|\mu]] = C \frac{\mathbf{b}^{\mathbf{a}}}{(n + \mathbf{b} - i)^{\mathbf{a}}} \left(1 - C \frac{\mathbf{b}^{\mathbf{a}}}{(n + \mathbf{b} - i)^{\mathbf{a}}}\right) \quad (76)$$

$$+ (C^2 - C) \frac{\mathbf{b}^{\mathbf{a}}}{(2n + \mathbf{b} - 2i)^{\mathbf{a}}} \quad (77)$$

$$\text{Var}[M_n] = \frac{\mathbf{b}^{\mathbf{a}}}{(\Delta n + \mathbf{b})^{\mathbf{a}}} \left(1 - \frac{\mathbf{b}^{\mathbf{a}}}{(\Delta n + \mathbf{b})^{\mathbf{a}}}\right) \quad (78)$$

Proof. • For Q_n it holds

$$E[Q_n] = \sum_{i=1}^n E[E[Y_1|\lambda, \mu]E[S_{1,1}|\sigma]E[Z_{n,i,1}|\mu]] \quad (79)$$

$$= \sum_{i=1}^n E[Y_1|\lambda, \mu]E[S_{1,1}|\sigma]E[Z_{n,i,1}|\mu] \quad (80)$$

$$= \sum_{i=1}^n \lambda \mu^\nu (\sigma + 1) e^{(i-n)\mu} \quad (81)$$

$$E[\lambda] = \frac{a}{b} \quad (82)$$

$$Var[\lambda] = \frac{a}{b^2} \quad (83)$$

$$E[\sigma + 1] = \frac{\alpha + \beta}{\beta} \quad (84)$$

$$Var[\sigma + 1] = \frac{\alpha}{\beta^2} \quad (85)$$

$$E[\mu^\nu e^{(i-n)\mu}] = \int_0^\infty \mu^\nu e^{(i-n)\mu} \frac{b^a \mu^{a-1} e^{-b\mu}}{\Gamma(a)} d\mu \quad (86)$$

$$= \frac{b^a}{\Gamma(a)} \int_0^\infty \mu^{a-1+\nu} e^{(i-n-b)\mu} d\mu \quad (87)$$

$$= \frac{b^a}{\Gamma(a)} (n + b - i)^{-a-\nu} \Gamma(a + \nu) \quad (88)$$

$$= \frac{b^a}{(n + b - i)^{a+\nu}} \frac{\Gamma(a + \nu)}{\Gamma(a)}. \quad (89)$$

It immediately follows:

$$E[Q_n] = \frac{a}{b} \frac{\alpha + \beta}{\beta} \frac{\Gamma(a + \nu)}{\Gamma(a)} \sum_{i=1}^n \frac{b^a}{(n + b - i)^{a+\nu}}. \quad (90)$$

Next we will calculate

$$Var[Q_n] = E[Var[Q_n|\lambda, \sigma, \mu]] \quad (91)$$

$$+ Var[E[Q_n|\lambda, \sigma, \mu]]. \quad (92)$$

Before we can do so, we need to collect a few easy supporting results:

$$Var[Y_1|\lambda, \mu] = \lambda \mu^\nu \quad (93)$$

$$Var[S_{1,1}|\sigma] = \sigma \quad (94)$$

$$E[Z_{n,i,1}|\mu]^2 = e^{((i-n)\mu)^2} \quad (95)$$

$$= e^{(2i-2n)\mu} \quad (96)$$

$$= E[Z_{2n,2i,1}|\mu] \quad (97)$$

$$\text{Var}[Z_{n,i,1}|\mu] = e^{((i-n)\mu)}(1 - e^{((i-n)\mu)}) \quad (98)$$

$$= e^{((i-n)\mu)} - e^{2((i-n)\mu)} \quad (99)$$

$$= E[Z_{n,i,1}|\mu] - E[Z_{2n,2i,1}|\mu] \quad (100)$$

We also need

$$E[\lambda^2] = \text{Var}[\lambda] + E[\lambda]^2 \quad (101)$$

$$= \frac{a}{b^2} + \frac{a^2}{b} \quad (102)$$

$$E[(\sigma + 1)^2] = \text{Var}[\sigma + 1] + E[\sigma + 1]^2 \quad (103)$$

$$= \frac{\alpha}{\beta^2} + \frac{\alpha + \beta^2}{\beta} \quad (104)$$

$$E[\mu^{2\nu} e^{(2i-2n)\mu}] = \int_0^\infty \mu^{2\nu} e^{(2i-2n)\mu} \frac{\mathfrak{b}^a \mu^{a-1} e^{-\mathfrak{b}\mu}}{\Gamma(\mathfrak{a})} d\mu \quad (105)$$

$$= \frac{\mathfrak{b}^a}{\Gamma(\mathfrak{a})} \int_0^\infty \mu^{a-1+2\nu} e^{(2i-2n-\mathfrak{b})\mu} d\mu \quad (106)$$

$$= \frac{\mathfrak{b}^a}{\Gamma(\mathfrak{a})} (2n + \mathfrak{b} - 2i)^{-a-2\nu} \Gamma(\mathfrak{a} + 2\nu) \quad (107)$$

$$= \frac{\Gamma(\mathfrak{a} + 2\nu)}{\Gamma(\mathfrak{a})} \frac{\mathfrak{b}^a}{(2n + \mathfrak{b} - 2i)^{a+2\nu}}. \quad (108)$$

This now allows us to calculate everything that is needed for the first half of the variance of Q_n , i.e., $E[\text{Var}[Q_n|\lambda, \sigma, \mu]]$. First note that

$$\text{Var}[Q_n|\lambda, \sigma, \mu] \quad (109)$$

$$= \sum_{i=1}^n ((\text{Var}[Y_i|\lambda, \mu] E[S_{i,l}|\sigma]^2 \quad (110)$$

$$+ E[Y_i|\lambda, \mu] \text{Var}[S_{i,l}|\mu]) E[Z_{n,i,1}|\mu]^2 \quad (111)$$

$$+ E[Y_i|\lambda] E[S_{i,l}|\sigma] \text{Var}[Z_{n,i,1}|\mu] \quad (112)$$

and

$$E[(\text{Var}[Y_i|\lambda, \mu] E[S_{i,l}|\sigma]^2 E[Z_{n,i,1}|\mu]^2] \quad (113)$$

$$= E[(\lambda)(\sigma + 1)^2 \mu^\nu e^{(2i-2n)\mu}] \quad (114)$$

$$= E[\lambda] E[(\sigma + 1)^2] E[\mu^\nu e^{(2i-2n)\mu}] \quad (115)$$

$$= \frac{a}{b} \left(\frac{\alpha}{\beta^2} + \frac{\alpha + \beta^2}{\beta} \right) \frac{\Gamma(\mathfrak{a} + \nu)}{\Gamma(\mathfrak{a})} \frac{\mathfrak{b}^a}{(2n + \mathfrak{b} - 2i)^{a+\nu}}. \quad (116)$$

Further, it holds

$$E[E[Y_i|\lambda, \mu] \text{Var}[S_{i,l}|\mu] E[Z_{n,i,1}|\mu]^2] \quad (117)$$

$$= E[\lambda \mu^\nu \sigma e^{(2i-2n)\mu}] \quad (118)$$

$$=E[\lambda]E[\sigma]E[\mu^\nu e^{(2i-2n)\mu}] \quad (119)$$

$$= \frac{a}{b} \frac{\alpha}{\beta} \frac{\Gamma(\mathbf{a} + \nu)}{\Gamma(\mathbf{a})} \frac{\mathbf{b}^{\mathbf{a}}}{(2n + \mathbf{b} - 2i)^{\mathbf{a} + \nu}} \quad (120)$$

and

$$E[E[Y_i|\lambda, \mu]E[S_{i,l}|\sigma]Var[Z_{n,i,1}|\mu]] \quad (121)$$

$$=E[\lambda\mu^\nu(\sigma + 1)(E[Z_{n,i,1}|\mu] - E[Z_{2n,2i,1}|\mu])] \quad (122)$$

$$=E[\lambda\mu^\nu(\sigma + 1)(e^{(i-n)\mu} - e^{(2i-2n)\mu})] \quad (123)$$

$$=E[\lambda]E[\sigma + 1](\mu^\nu E[e^{(i-n)\mu}] - E[\mu^\nu e^{(2i-2n)\mu}]) \quad (124)$$

$$= \frac{a}{b} \frac{\alpha + \beta}{\beta} \frac{\Gamma(\mathbf{a} + \nu)}{\Gamma(\mathbf{a})} \left(\frac{\mathbf{b}^{\mathbf{a}}}{(n + \mathbf{b} - i)^{\mathbf{a} + \nu}} - \frac{\mathbf{b}^{\mathbf{a}}}{(2n + \mathbf{b} - 2i)^{\mathbf{a} + \nu}} \right). \quad (125)$$

It follows:

$$E[(Var[Y_i|\lambda, \mu]E[S_{i,l}|\sigma]^2 E[Z_{n,i,1}|\mu]^2)] \quad (126)$$

$$= \mathbf{b}^{\mathbf{a}} \frac{\Gamma(\mathbf{a} + \nu)}{\Gamma(\mathbf{a})} \left[\left(\frac{a}{b} \left(\frac{\alpha}{\beta^2} + \frac{\alpha + \beta^2}{\beta} \right) + \frac{a}{b} \frac{\alpha}{\beta} - \frac{a}{b} \frac{\alpha + \beta}{\beta} \right) \right] \quad (127)$$

$$\left[\frac{1}{(2n + \mathbf{b} - 2i)^{\mathbf{a} + \nu}} + \frac{a}{b} \frac{\alpha + \beta}{\beta} \frac{1}{(n + \mathbf{b} - i)^{\mathbf{a} + \nu}} \right] \quad (128)$$

$$= \mathbf{b}^{\mathbf{a}} \frac{\Gamma(\mathbf{a} + \nu)}{\Gamma(\mathbf{a})} \left[\left(\frac{a}{b} \left(\frac{\alpha}{\beta^2} + \frac{\alpha + \beta^2}{\beta} - 1 \right) \right) \right] \quad (129)$$

$$\left[\frac{1}{(2n + \mathbf{b} - 2i)^{\mathbf{a} + \nu}} + \frac{a}{b} \frac{\alpha + \beta}{\beta} \frac{1}{(n + \mathbf{b} - i)^{\mathbf{a} + \nu}} \right]. \quad (130)$$

Finally for the second part of the variance, i.e., $Var[E[Q_n|\lambda, \sigma, \mu]]$, we need:

$$Var\left[\sum_{i=1}^n \mu^\nu e^{(i-n)\mu}\right] \quad (131)$$

$$= \sum_{1 \leq i \leq j \leq n} Cov[\mu^\nu e^{(i-n)\mu}, \mu^\nu e^{(j-n)\mu}] \quad (132)$$

$$= \sum_{1 \leq i \leq j \leq n} E[\mu^\nu e^{(i-n)\mu} \mu^\nu e^{(j-n)\mu}] - E[\mu^\nu e^{(i-n)\mu}] E[\mu^\nu e^{(j-n)\mu}] \quad (133)$$

$$= \sum_{1 \leq i \leq j \leq n} E[\mu^{2\nu} e^{(i+j-2n)\mu}] - E[\mu^\nu e^{(i-n)\mu}] E[\mu^\nu e^{(j-n)\mu}] \quad (134)$$

$$= \sum_{1 \leq i \leq j \leq n} \left(\frac{\mathbf{b}^{\mathbf{a}}}{(2n + \mathbf{b} - i - j)^{\mathbf{a} + 2\nu}} \frac{\Gamma(\mathbf{a} + 2\nu)}{\Gamma(\mathbf{a})} \right) \quad (135)$$

$$- \frac{(\Gamma(\mathbf{a} + \nu))^2}{\Gamma(\mathbf{a})} \frac{\mathbf{b}^{\mathbf{a}}}{(n + \mathbf{b} - i)^{\mathbf{a} + \nu}} \frac{\mathbf{b}^{\mathbf{a}}}{(n + \mathbf{b} - j)^{\mathbf{a} + \nu}} \quad (136)$$

$$= \mathbf{b}^{\mathbf{a}} \frac{\Gamma(\mathbf{a} + 2\nu)}{\Gamma(\mathbf{a})} \sum_{1 \leq i \leq j \leq n} \frac{1}{(2n + \mathbf{b} - i - j)^{\mathbf{a} + 2\nu}} \quad (137)$$

$$- \left(\mathfrak{b}^{\mathfrak{a}} \frac{\Gamma(\mathfrak{a} + \nu)}{\Gamma(\mathfrak{a})} \right)^2 \sum_{1 \leq i \leq j \leq n} \frac{1}{(n + \mathfrak{b} - i)^{\mathfrak{a} + \nu}} \frac{1}{(n + \mathfrak{b} - j)^{\mathfrak{a} + \nu}} \quad (138)$$

and

$$\text{Var}[\lambda(\sigma + 1)] \quad (139)$$

$$= E[\lambda]^2 \text{Var}[\sigma + 1] + \text{Var}[\lambda] E[\sigma + 1]^2 + \text{Var}[\lambda] \text{Var}[\sigma + 1] \quad (140)$$

$$= \frac{a^2}{b} \frac{\alpha}{\beta^2} + \frac{a}{b^2} \frac{\alpha + \beta^2}{\beta} + \frac{a}{b^2} \frac{\alpha}{\beta^2}. \quad (141)$$

Given this, we can write:

$$\text{Var}[E[Q_n | \lambda, \sigma, \mu]] \quad (142)$$

$$= \text{Var} \left[\sum_{i=1}^n E \left[\sum_{k=1}^{\sum_{l=0}^Y S_{i,l}} Z_{n,i,k} | \lambda, \sigma, \mu \right] \right] \quad (143)$$

$$= \text{Var} \left[\sum_{i=1}^n \lambda \mu^\nu (\sigma + 1) e^{(i-n)\mu} \right] \quad (144)$$

$$= \text{Var} \left[\lambda(\sigma + 1) \sum_{i=1}^n \mu^\nu e^{(i-n)\mu} \right] \quad (145)$$

$$= E[\lambda(\sigma + 1)]^2 \text{Var} \left[\sum_{i=1}^n \mu^\nu e^{(i-n)\mu} \right] + \text{Var}[\lambda(\sigma + 1)] \text{Var} \left[\sum_{i=1}^n \mu^\nu e^{(i-n)\mu} \right] \quad (146)$$

$$+ \text{Var}[\lambda(\sigma + 1)] E \left[\sum_{i=1}^n \mu^\nu e^{(i-n)\mu} \right]^2 \quad (147)$$

$$= \left(\frac{a^2}{b} \frac{\alpha + \beta^2}{\beta} + \left(\frac{a^2}{b} \frac{\alpha}{\beta^2} + \frac{a}{b^2} \frac{\alpha + \beta^2}{\beta} + \frac{a}{b^2} \frac{\alpha}{\beta^2} \right) \right) \quad (148)$$

$$\left(\mathfrak{b}^{\mathfrak{a}} \frac{\Gamma(\mathfrak{a} + 2\nu)}{\Gamma(\mathfrak{a})} \sum_{1 \leq i \leq j \leq n} \frac{1}{(2n + \mathfrak{b} - i - j)^{\mathfrak{a} + 2\nu}} \right) \quad (149)$$

$$- \left(\mathfrak{b}^{\mathfrak{a}} \frac{\Gamma(\mathfrak{a} + \nu)}{\Gamma(\mathfrak{a})} \right)^2 \sum_{1 \leq i \leq j \leq n} \frac{1}{(n + \mathfrak{b} - i)^{\mathfrak{a} + \nu}} \frac{1}{(n + \mathfrak{b} - j)^{\mathfrak{a} + \nu}} \quad (150)$$

$$+ \left(\frac{a^2}{b} \frac{\alpha}{\beta^2} + \frac{a}{b^2} \frac{\alpha + \beta^2}{\beta} + \frac{a}{b^2} \frac{\alpha}{\beta^2} \right) \left(\frac{\Gamma(\mathfrak{a} + \nu)}{\Gamma(\mathfrak{a})} \sum_{i=1}^n \frac{\mathfrak{b}^{\mathfrak{a}}}{(n + \mathfrak{b} - i)^{\mathfrak{a} + \nu}} \right)^2. \quad (151)$$

Inserting into Propositions 1 and 2 now yields the result.

- For D_i note the following: As an exponential distribution whose rate is drawn from a Gamma distribution with shape \mathfrak{a} and rate \mathfrak{b} is equal to a Lomax distribution with scale \mathfrak{b} and shape \mathfrak{a} , a single $Z_{i,j,k}$ is equal to a Bernoulli trial over the complementary CDF of the Lomax distribution, i.e.,

$$E[Z_{i,j,k}] = \left(1 + \frac{i-j}{\mathfrak{b}} \right)^{-(\mathfrak{a})}. \quad (152)$$

It therefore holds

$$E[D_i] \leq E[D_{i-1}](1 - \prod_{j=0}^{i-1} (1 - (1 + \frac{i-j}{\mathfrak{b}})^{-\mathfrak{a}})^{\frac{\alpha}{\mathfrak{b}}}) \quad (153)$$

$$E[D_1] = (1 - (1 - (1 + \frac{1}{\mathfrak{b}})^{-\mathfrak{a}})^C). \quad (154)$$

- It now directly follows

$$\text{Var}[Z_{n,i,1}] = E[\text{Var}[Z_{n,i,1}|\mu]] + \text{Var}[E[Z_{n,i,1}|\mu]] \quad (155)$$

$$= E[E[Z_{n,i,1}|\mu]] - E[E[Z_{2n,2i,1}|\mu]] + \text{Var}[e^{(i-n)\mu}] \quad (156)$$

$$= \frac{\mathfrak{b}^{\mathfrak{a}}}{(n + \mathfrak{b} - i)^{\mathfrak{a}}} - \frac{\mathfrak{b}^{\mathfrak{a}}}{(2n + \mathfrak{b} - 2i)^{\mathfrak{a}}} \quad (157)$$

$$+ \frac{\mathfrak{b}^{\mathfrak{a}}}{(2n + \mathfrak{b} - 2i)^{\mathfrak{a}}} - \frac{\mathfrak{b}^{\mathfrak{a}}}{(n + \mathfrak{b} - i)^{\mathfrak{a}}} \frac{\mathfrak{b}^{\mathfrak{a}}}{(n + \mathfrak{b} - i)^{\mathfrak{a}}} \quad (158)$$

$$= \frac{\mathfrak{b}^{\mathfrak{a}}}{(n + \mathfrak{b} - i)^{\mathfrak{a}}} (1 - \frac{\mathfrak{b}^{\mathfrak{a}}}{(n + \mathfrak{b} - i)^{\mathfrak{a}}}). \quad (159)$$

and

$$CE[\text{Var}[Z_{n,i,k}|\mu]] + C^2 \text{Var}[E[Z_{n,i,k}|\mu]] \quad (160)$$

$$= C(E[E[Z_{n,i,1}|\mu]] - E[E[Z_{2n,2i,1}|\mu]]) + C^2 \text{Var}[e^{(i-n)\mu}] \quad (161)$$

$$= C(\frac{\mathfrak{b}^{\mathfrak{a}}}{(n + \mathfrak{b} - i)^{\mathfrak{a}}} - \frac{\mathfrak{b}^{\mathfrak{a}}}{(2n + \mathfrak{b} - 2i)^{\mathfrak{a}}}) + C^2(\frac{\mathfrak{b}^{\mathfrak{a}}}{(2n + \mathfrak{b} - 2i)^{\mathfrak{a}}} - (\frac{\mathfrak{b}^{\mathfrak{a}}}{(n + \mathfrak{b} - i)^{\mathfrak{a}}})^2) \quad (162)$$

$$= C \frac{\mathfrak{b}^{\mathfrak{a}}}{(n + \mathfrak{b} - i)^{\mathfrak{a}}} (1 - C \frac{\mathfrak{b}^{\mathfrak{a}}}{(n + \mathfrak{b} - i)^{\mathfrak{a}}}) + (C^2 - C) \frac{\mathfrak{b}^{\mathfrak{a}}}{(2n + \mathfrak{b} - 2i)^{\mathfrak{a}}}. \quad (163)$$

- For M_i it holds by the same argument,

$$E[M_n] = \frac{\mathfrak{b}^{\mathfrak{a}}}{(\Delta n + \mathfrak{b})^{\mathfrak{a}}} \quad (164)$$

$$\text{Var}[M_n] = \frac{\mathfrak{b}^{\mathfrak{a}}}{(\Delta n + \mathfrak{b})^{\mathfrak{a}}} (1 - \frac{\mathfrak{b}^{\mathfrak{a}}}{(\Delta n + \mathfrak{b})^{\mathfrak{a}}}). \quad (165)$$

□

Appendix D. Importance Sampling

Importance sampling is a technique that, instead of drawing samples r from the *nominal sampling distribution* p in order to estimate the expected value of some feature of the samples f , it draws the samples from an *importance distribution* q . These samples are then weighted according to the ratio between both distributions in order to obtain an estimate of $E[f]$ with a lower variance. This can vastly reduce the number of samples required to make statements with high confidence. It is well known (see Kahn and Marshall (1953)) that the optimal importance density satisfies

$$q(r) = \frac{f(r)p(r)}{E[f(r)]} \quad (166)$$

While calculating this exactly would require knowledge about the very value we want to estimate, it can often be approximated reasonably well. In our case, where each simulation run r depends on tens of thousands of random variables, we define a heuristic measure that roughly indicates how likely a run is to fail and partition the set of all possible runs into buckets using this measure. We then approximate the optimal q on the level of buckets.

As a first step, we now define the heuristic measure we use:

Definition 5. For a deployment x , denote by

$$i^x = E[L_n^x] + \sqrt{(1 - 0.01)/0.01 * Var[L_n^x]} \tag{167}$$

the upper bound of the 99% confidence interval of a deployment's size in timestep n as given by Cantelli's inequality. For a given run r , pre-draw the parameters of all deployments that might arrive during the simulation run. Then consider the following extremely simplified simulation:

1. On the first of each month, 730 new deployments arrive.
2. Deployments only die at the end of the month after their maximum lifetime is reached. They do not die when they reach zero cores.
3. The cluster knows each deployment's exact type.
4. Deployment always take exactly their expected size.
5. Deployments are accepted into the cluster whenever $\sum_{\tilde{x} \in A_\pi(b)} i^x < 22000$.

Denote by $X(n)$, $n \in [1, 36]$ the set of deployments in the cluster at the beginning of each month during this simplified simulation. Then the badness measure BM of a run r is defined as

$$BM(r) = \max_n \sum_{\tilde{x} \in X(n)} i^x. \tag{168}$$

This is a reasonable (though highly heuristic) predictor of whether a run might produce a very large number failures. Most importantly, because it assumes away all randomness that occurs during the simulation run, it can be evaluated very quickly (< 1 second).

To properly utilize importance sampling, we now sort any simulation run r into one of three buckets based on their BM value: $I_1 = \{r : BM(r) \leq 25000\}$, $I_2 = \{r : 25000 \leq BM(r) \leq 30000\}$, $I_3 = \{r : 30000 \leq BM(r)\}$. Before we can apply importance sampling, we calculate the probability for a given run to be in each of the buckets. For this, we calculate BM for 100,000 runs. The resulting probabilities can be found in Table 4.

To sample runs from the different buckets with different weights, we employ a type of rejection sampling: Before starting a simulation run r , we evaluate $BM(r)$. Depending on the bucket I_i the run would result in, we then redraw with some probability $Pr_r(I_i)$ (i.e., all deployment parameters are discarded and redrawn) and block off all lower buckets (i.e., automatically rejecting any further redraws that would result in I_j , $j < i$). The highest bucket (in our case I_3) never gets redrawn, i.e., $Pr_r(I_3) = 0$. This scheme continues iteratively until we accept a run. This results in the following importance distribution.

Probabilities	I_1	I_2	I_3
$Pr(I_i)$	0.5699	0.4121	0.018
$Pr(I_i \cap_{k \geq i} I_k)$	0.88319	0.9582	1
$Pr_r(I_i \cap_{k \geq i} I_k)$	0.5369	0.8816	0

Table 4: Estimation of BM probabilities

Proposition 6. *For a run r with nominal probability $Pr(r)$ and BM such that $r \in I_i$, the above rejection scheme results in importance distribution q with*

$$q(r) \tag{169}$$

$$= Pr(r|I_i) \frac{Pr(I_i | \cap_{k \geq i} I_k)(1 - Pr_r(I_i | \cap_{k \geq i} I_k))}{1 - Pr(I_i | \cap_{k \geq i} I_k)Pr_r(I_i | \cap_{k \geq i} I_k)} \prod_{j < i} \frac{Pr(\cap_{k > j} I_k | \cap_{k \geq j} I_k)}{1 - Pr(I_j | \cap_{k \geq j} I_k)Pr_r(I_j | \cap_{k \geq j} I_k)} \tag{170}$$

Proof. We first show this for two buckets I_1 and I_2 . With only two buckets, since I_2 is the highest bucket, it has an acceptance probability of 1.

It follows that for a run that would be in I_1 , we redraw with probability $Pr_r(I_1)$ and otherwise accept. It thus holds

$$q(I_1) = Pr(I_1)((1 - Pr_r(I_1)) + Pr_r(I_1)q(I_1)). \tag{171}$$

This is a geometric series and therefore

$$q(I_1) = \sum_{k=0}^{\infty} Pr(I_1)(1 - Pr_r(I_1))(Pr(I_1)Pr_r(I_1))^k \tag{172}$$

$$= \frac{Pr(I_1)(1 - Pr_r(I_1))}{1 - Pr(I_1)Pr_r(I_1)} \tag{173}$$

Similarly, it holds

$$q(I_2) = \sum_{i=0}^{\infty} Pr(I_2)(Pr(I_1)Pr_r(I_1))^i \tag{174}$$

$$= \frac{Pr(I_2)}{1 - Pr(I_1)Pr_r(I_1)} \tag{175}$$

Iteratively applying this argument to more than two buckets by dividing the top bucket into two buckets then yields

$$q(I_i) = \frac{Pr(I_i | \cap_{k \geq i} I_k)(1 - Pr_r(I_i | \cap_{k \geq i} I_k))}{1 - Pr(I_i | \cap_{k \geq i} I_k)Pr_r(I_i | \cap_{k \geq i} I_k)} \prod_{j < i} \frac{Pr(\cap_{k > j} I_k | \cap_{k \geq j} I_k)}{1 - Pr(I_j | \cap_{k \geq j} I_k)Pr_r(I_j | \cap_{k \geq j} I_k)} \tag{176}$$

Finally, by Bayes' theorem it holds that $q(r) = q(r|I_i)q(I_i)$ and since it further holds that $q(r|I_i) = Pr(r|I_i)$, the statement of the proposition follows. \square

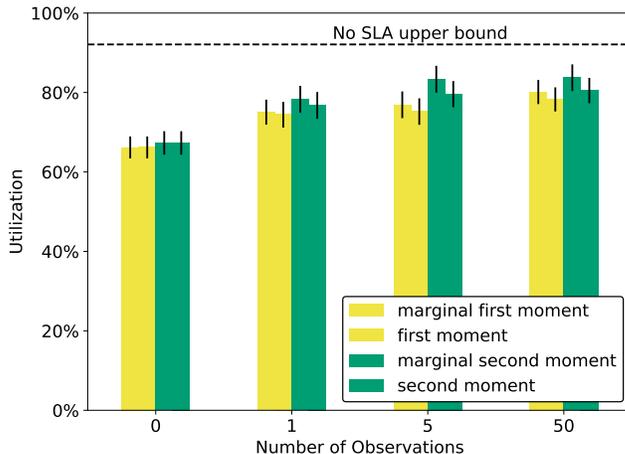


Figure 9: Ablation of policies with and without marginal heuristic (error bars indicate 95% bootstrap confidence intervals)

To find good rejection probabilities p_r that result in low sample variance, we did 500 runs for each bucket under the second moment policy with threshold 15000 to get a very rough estimate of f (i.e. the scale out failure probability) for each bucket. The rejection probabilities p_r are then calculated by combining Equation (166) and Equation (170). The resulting values can also be found in Table 4.

Appendix E. Ablation of Policies Without Marginal Heuristic

In this section, we ablate the simulation results for our policies with marginal heuristic with the same policies without these heuristic. Note again that without individual prior observations, no arriving deployment is ever marginal in the policy horizon. Thus, without prior observations, the marginal policies are equivalent to the policies without marginal heuristic.

We have simulated the policies without marginal heuristic for 1, 5, and 50 observations, with the same setup as in Section 5. In Figure 9, we contrast those results with the results for the marginal policies. We see that with just a few observations, the policies with and without marginal heuristic have a very similar performance (though the marginal heuristic still enables slightly higher utilization). This is not surprising, since relatively few arriving deployments are marginal at this level of prior information. Consequently, the utilization gap increases with more information. Both with 5 and 50 observations, the marginal second moment policy obtains a sizable utilization increase of more than 3% compared to the second moment policy without the heuristic.

References

- Abhishek, V., Kash, I. A., & Key, P. (2012). Fixed and market pricing for cloud services. In *7th Workshop on the Economics of Networks, Systems, and Computation (NetEcon)*, pp. 157–162.
- Altman, E. (1999). *Constrained Markov decision processes*, Vol. 7. CRC Press.
- Ashlagi, I., Burq, M., Dutta, C., Jaillet, P., Saberi, A., & Sholley, C. (2019). Edge weighted online windowed matching. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 729–742.
- Assadi, S., Khanna, S., & Li, Y. (2017). The stochastic matching problem: Beating half with a non-adaptive algorithm. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pp. 99–116.
- Azar, Y., Kalp-Shaltiel, I., Lucier, B., Menache, I., Naor, J., & Yaniv, J. (2015). Truthful online scheduling with commitments. In *Proceedings of the Sixteenth ACM Conference on Economics and Computation*, pp. 715–732.
- Babaioff, M., Mansour, Y., Nisan, N., Noti, G., Curino, C., Ganapathy, N., Menache, I., Reingold, O., Tennenholtz, M., & Timnat, E. (2017). Era: a framework for economic resource allocation for the cloud. In *Proceedings of the 26th International Conference on World Wide Web Companion*, pp. 635–642.
- Behnezhad, S., & Reyhani, N. (2018). Almost optimal stochastic weighted matching with few queries. In *Proceedings of the 2018 ACM Conference on Economics and Computation*, pp. 235–249.
- Berger, A. W., & Whitt, W. (1998). Extending the effective bandwidth concept to networks with priority classes. *IEEE Communications magazine*, 36(8), 78–83.
- Brooks, A., Makarenko, A., Williams, S., & Durrant-Whyte, H. (2006). Parametric pomdps for planning in continuous state spaces. *Robotics and Autonomous Systems*, 54(11), 887–897.
- Cohen, M. C., Keller, P. W., Mirrokni, V., & Zadimoghaddam, M. (2019). Overcommitment in cloud services: Bin packing with chance constraints. *Management Science*, 65(7), 3255–3271.
- Cortez, E., Bonde, A., Muzio, A., Russinovich, M., Fontoura, M., & Bianchini, R. (2017). Resource central: Understanding and predicting workloads for improved resource management in large cloud platforms. In *Proceedings of the 26th Symposium on Operating Systems Principles*, pp. 153–167.
- Daw, A., & Pender, J. (2018). Queues driven by hawkes processes. *Stochastic Systems*, 8(3), 192–229.
- Delimitrou, C., Bambos, N., & Kozyrakis, C. (2013). Qos-aware admission control in heterogeneous datacenters. In *Proceedings of the 10th International Conference on Autonomous Computing*, pp. 291–296.
- Dierks, L., & Seuken, S. (2019). Cloud pricing: The spot market strikes back. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, p. 593.

- Dierks, L., & Seuken, S. (2020). The competitive effects of variance-based pricing. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence*, pp. 362–370.
- Dolev, D., Feitelson, D. G., Halpern, J. Y., Kupferman, R., & Linial, N. (2012). No justified complaints: On fair sharing of multiple resources. In *Proceedings of the 3rd Innovations in Theoretical Computer Science Conference*, pp. 68–75.
- Duff, M. O., & Barto, A. (2002). *Optimal Learning: Computational procedures for Bayes-adaptive Markov decision processes*. Ph.D. thesis, University of Massachusetts at Amherst.
- Efron, B. (1987). Better bootstrap confidence intervals. *Journal of the American statistical Association*, 82(397), 171–185.
- Ghods, A., Zaharia, M., Hindman, B., Konwinski, A., Shenker, S., & Stoica, I. (2011). Dominant resource fairness: Fair allocation of multiple resource types. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation*, p. 323–336.
- Gutman, A., & Nisan, N. (2012). Fair allocation without trade. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems-Volume 2*, pp. 719–728.
- Haskell, W. B., & Jain, R. (2015). A convex analytic approach to risk-aware markov decision processes. *SIAM Journal on Control and Optimization*, 53(3), 1569–1598.
- Hawkes, A. G. (2018). Hawkes processes and their applications to finance: a review. *Quantitative Finance*, 18(2), 193–198.
- Hindman, B., Konwinski, A., Zaharia, M., Ghods, A., Joseph, A. D., Katz, R., Shenker, S., & Stoica, I. (2011). Mesos: A platform for fine-grained resource sharing in the data center. In *Proceedings of the 8th USENIX Conference on Networked Systems Design and Implementation*, pp. 295–308.
- Isom, J. D., Meyn, S. P., & Braatz, R. D. (2008). Piecewise linear dynamic programming for constrained pomdps.. In *Proceedings of the 23rd National Conference on Artificial Intelligence - Volume 1*, p. 291–296.
- Jyothi, S. A., Curino, C., Menache, I., Narayanamurthy, S. M., Tumanov, A., Yaniv, J., Mavlyutov, R., Goiri, I., Krishnan, S., Kulkarni, J., & Rao, S. (2016). Morpheus: Towards automated slos for enterprise clusters. In *12th USENIX Symposium on Operating Systems Design and Implementation*, pp. 117–134.
- Kahn, H., & Marshall, A. W. (1953). Methods of reducing sample size in monte carlo computations. *Journal of the Operations Research Society of America*, 1(5), 263–278.
- Kash, I. A., & Key, P. B. (2016). Pricing the cloud. *IEEE Internet Computing*, 20(1), 36–43.
- Kash, I. A., Procaccia, A. D., & Shah, N. (2014). No agent left behind: Dynamic fair division of multiple resources. *Journal of Artificial Intelligence Research*, 51, 579–603.
- Kelly, F. P. (1991). Effective bandwidths at multi-class queues. *Queueing systems*, 9(1-2), 5–15.

- Khonji, M., Jasour, A., & Williams, B. (2019). Approximability of constant-horizon constrained pomdp. In *Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence*, pp. 5583–5590.
- Kim, D., Lee, J., Kim, K.-E., & Poupart, P. (2011). Point-based value iteration for constrained pomdps. In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence - Volume Volume Three*, pp. 1968–1974.
- Lusena, C., Goldsmith, J., & Mundhenk, M. (2001). Nonapproximability results for partially observable markov decision processes. *Journal of artificial intelligence research*, *14*, 83–103.
- Ma, W., & Simchi-Levi, D. (2019). Tight weight-dependent competitive ratios for online edge-weighted bipartite matching and beyond. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 727–728.
- Microsoft (2018). Earnings press release FY18 Q2. <https://www.microsoft.com/en-us/Investor/earnings/FY-2018-Q2/press-release-webcast>. Accessed: 2021-04-22.
- NIST (2012). Sematech e-handbook of statistical methods. <http://www.itl.nist.gov/div898/handbook/apr/section4/apr412.htm>. Accessed: 2021-04-22.
- Ono, M., Kuwata, Y., & Balaram, J. (2012). Joint chance-constrained dynamic programming. In *2012 IEEE 51st IEEE Conference on Decision and Control*, pp. 1915–1922.
- Papadimitriou, C. H., & Tsitsiklis, J. N. (1987). The complexity of markov decision processes. *Mathematics of Operations Research*, *12*(3), 441–450.
- Parkes, D. C., Procaccia, A. D., & Shah, N. (2015). Beyond dominant resource fairness: Extensions, limitations, and indivisibilities. *ACM Transactions on Economics and Computation*, *3*(1), 3:1–3:22.
- Porta, J. M., Vlassis, N., Spaan, M. T., & Poupart, P. (2006). Point-based value iteration for continuous pomdps. *Journal of Machine Learning Research*, *7*(Nov), 2329–2367.
- Poupart, P., Malhotra, A., Pei, P., Kim, K.-E., Goh, B., & Bowling, M. (2015). Approximate linear programming for constrained partially observable markov decision processes. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence*, pp. 3342–3348.
- Rajan, K., Kakadia, D., Curino, C., & Krishnan, S. (2016). Perforator: eloquent performance models for resource optimization. In *Proceedings of the Seventh ACM Symposium on Cloud Computing*, pp. 415–427.
- Roy, N., Gordon, G., & Thrun, S. (2005). Finding approximate pomdp solutions through belief compression. *Journal of artificial intelligence research*, *23*, 1–40.
- Russell, S. J., & Norvig, P. (2016). *Artificial intelligence: a modern approach*. Malaysia; Pearson Education Limited.
- Santana, P., Thiébaux, S., & Williams, B. (2016). Rao*: an algorithm for chance constrained pomdps. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, p. 3308–3314.

- Schwarzkopf, M., Konwinski, A., Abd-El-Malek, M., & Wilkes, J. (2013). Omega: flexible, scalable schedulers for large compute clusters. In *Proceedings of the 8th ACM European Conference on Computer Systems*, pp. 351–364.
- Smallwood, R. D., & Sondik, E. J. (1973). The optimal control of partially observable markov processes over a finite horizon. *Operations research*, 21(5), 1071–1088.
- Smith, T., & Simmons, R. (2005). Point-based pomdp algorithms: Improved analysis and implementation. In *Proceedings of the Twenty-First Conference on Uncertainty in Artificial Intelligence*, pp. 542–549.
- Song, W., Xiao, Z., Chen, Q., & Luo, H. (2014). Adaptive resource provisioning for the cloud using online bin packing. *IEEE Transactions on Computers*, 63(11), 2647–2660.
- Tumanov, A., Zhu, T., Park, J. W., Kozuch, M. A., Harchol-Balter, M., & Ganger, G. R. (2016). Tetrisched: global rescheduling with adaptive plan-ahead in dynamic heterogeneous clusters. In *Proceedings of the Eleventh European Conference on Computer Systems*, p. 35.
- Undurti, A., & How, J. P. (2010). An online algorithm for constrained pomdps. In *2010 IEEE International Conference on Robotics and Automation*, pp. 3966–3973.
- Verma, A., Pedrosa, L., Korupolu, M., Oppenheimer, D., Tune, E., & Wilkes, J. (2015). Large-scale cluster management at google with borg. In *Proceedings of the Tenth European Conference on Computer Systems*.
- Walraven, E., & Spaan, M. T. (2018). Column generation algorithms for constrained pomdps. *Journal of Artificial Intelligence Research*, 62, 489–533.
- Wolke, A., Tsend-Ayush, B., Pfeiffer, C., & Bichler, M. (2015). More than bin packing. *Information Systems*, 52(C), 83–95.
- Yan, Y., Gao, Y., Chen, Y., Guo, Z., Chen, B., & Moscibroda, T. (2016). Tr-spark: Transient computing for big data analytics. In *Proceedings of the Seventh ACM Symposium on Cloud Computing*, p. 484–496.
- Yuan, H., Bi, J., Tan, W., & Li, B. H. (2016). Cawsac: Cost-aware workload scheduling and admission control for distributed cloud data centers. *IEEE Transactions on Automation Science and Engineering*, 13(2), 976–985.
- Yuan, H., Bi, J., Li, B. H., & Tan, W. (2017). Cost-aware request routing in multi-geography cloud data centres using software-defined networking. *Enterprise Information Systems*, 11(3), 359–388.
- Zhao, J., Yang, K., Wei, X., Ding, Y., Hu, L., & Xu, G. (2016). A heuristic clustering-based task deployment approach for load balancing using bayes theorem in cloud environment. *IEEE Transactions on Parallel and Distributed Systems*, 27(2), 305–316.