

Characterizing the Impact of the Workload on the Value of Dynamic Resizing in Data Centers

Kai Wang^{*†}, Minghong Lin[†], Florin Ciucu[‡], Adam Wierman[†] and Chuang Lin[§]

^{*}Chinese Academy of Sciences, [†]California Institute of Technology, [‡]T-Labs / TU Berlin, [§]Tsinghua University

Abstract—Energy consumption imposes a significant cost for data centers; yet much of that energy is used to maintain excess service capacity during periods of predictably low load. Resultantly, there has recently been interest in developing designs that allow the service capacity to be dynamically resized to match the current workload. However, there is still much debate about the value of such approaches in real settings. In this paper, we show that the value of dynamic resizing is highly dependent on statistics of the workload process. In particular, both slow time-scale non-stationarities of the workload (e.g., the peak-to-mean ratio) and the fast time-scale stochasticity (e.g., the burstiness of arrivals) play key roles. To illustrate the impact of these factors, we combine optimization-based modeling of the slow time-scale with stochastic modeling of the fast time scale.

I. INTRODUCTION

Energy costs represent a significant, and growing, fraction of a data center’s budget. Hence there is a push to improve the energy efficiency of data centers, both in terms of the components (servers, disks, network [24], [6], [8], [14]) and the algorithms [4], [11], [10], [22]. One specific aspect of data center design that is the focus of this paper is dynamically resizing the service capacity of the data center so that during periods of low load some servers are allowed to enter a power-saving mode (e.g., go to sleep or shut down).

The potential benefits of dynamic resizing have been a point of debate in the community [18], [11], [25]. On one hand, it is clear that, because data centers are far from perfectly energy proportional, significant energy is used to maintain excess capacity during periods of predictably low load when there is a diurnal workload with a high peak-to-mean ratio. On the other hand, there are also significant costs to dynamically adjusting the number of active servers. These costs come in terms of the engineering challenges in making this possible [13], [27], [5], as well as the latency, energy, and wear-and-tear costs of the actual “switching” operations involved [7], [11], [16].

The challenges for dynamic resizing highlighted above have been the subject of significant research. At this point, many of the engineering challenges associated with facilitating dynamic resizing have been resolved, e.g., [13], [27], [5]. Additionally, the algorithmic challenge of deciding, without knowledge of the future workload, whether to incur the significant “switching costs” associated with changing the available service capacity has been studied in depth and a number of promising algorithms have emerged [22], [3], [11], [15].

However, despite this body of work, the question of characterizing the potential benefits of dynamic resizing has still not been properly addressed.

The perspective of this paper is that, apart from engineering challenges, the key determinant of whether dynamic resizing is valuable is the workload. In particular, a key observation,

which is the starting point for our work, is that there are two factors of the workload which provide dynamic resizing potential savings:

- (i) Non-stationarities at a slow time-scale, e.g., diurnal workload variations.
- (ii) Stochastic variability at a fast time-scale, e.g., the burstiness of request arrivals.

To this point, we are not aware of any work characterizing the benefits of dynamic resizing that captures both of these features. There is one body of literature which provides algorithms that take advantage of (i), e.g., [11], [10], [22], [3]. This work tends to use an optimization-based approach to develop dynamic resizing algorithms. There is another body of literature which provides algorithms that take advantage of (ii), e.g., [15], [16]. This work tends to assume a stationary queueing model with Poisson arrivals to develop dynamic resizing algorithms.

The first contribution of the current paper is to provide an analytic framework that captures both effects (i) and (ii). We accomplish this by using an optimization framework at the slow time-scale (see Section II), which is similar to that of [22], and combining this with stochastic network calculus and large deviations modeling for the fast time-scale (see Section III), which allows us to study a wide variety of underlying arrival processes. We consider both light-tailed models with various degrees of burstiness and heavy-tailed models that exhibit self-similarity.

Using this modeling framework, we are able to provide both analytic and numerical results that yield new insight into the potential benefits of dynamic resizing (see Section IV). Specifically, we use trace-driven numerical simulations to study (i) the role of burstiness for dynamic resizing, (ii) the role of the peak-to-mean ratio for dynamic resizing, (iii) the role of the SLA for dynamic resizing, and (iv) the interaction between (i), (ii), and (iii). The key realization is that each of these parameters are extremely important for determining the value of dynamic resizing. In particular, for any fixed choices of two of these parameters, the third can be chosen so that dynamic resizing does or does not provide significant cost savings for the data center. Thus, performing a detailed study of the interaction of these factors is important. To that end, Figures 10-12 provide concrete illustrations of which settings of peak-to-mean ratio, burstiness, and SLAs dynamic resizing is and is not valuable.

There are some interesting facts about these parameters individually that our case studies uncover. Two important examples are the following. First, while one might expect that increased burstiness provides increased opportunities for dynamic resizing, it turns out the burstiness at the fast time-scale actually reduces the potential cost savings achievable via

dynamic resizing. Second, it turns out the impact of the SLA can be quite different depending on whether the arrival process is heavy- or light-tailed. In particular, as the SLA becomes more strict, the cost savings possible via dynamic resizing under heavy-tailed arrivals decreases quickly; however, the cost savings possible via dynamic resizing under light-tailed workloads is unchanged.

In addition to detailed case studies, we provide analytic results that support many of the insights provided by the numerics. In particular, Theorems 1 and 2 provide monotonicity and scaling results for dynamic resizing in the case of Poisson arrivals and heavy-tailed, self-similar arrivals. Due to page limitation, all proofs are omitted in the paper and the full version is available as the technical report [28].

II. SLOW TIME-SCALE MODEL

In this section and the one that follows, we introduce our model. We start with the “slow time-scale model”. This model is meant to capture what is happening at the time-scale of the data center control decisions, i.e., at the time-scale which the data center is willing to adjust its service capacity. For many reasons, this is a much slower time-scale than the time-scale at which requests arrive to the data center. We provide a model for this “fast time-scale” in the next section.

A. The Workload

At this time-scale, our goal is to provide a model which can capture the impact of diurnal non-stationarities in the workload. To this end, we consider a discrete-time model such that there is a time interval of interest which is evenly divided into “frames” $k \in \{1, \dots, K\}$. In practice, the length of a frame could be on the order of 5-10 minutes, whereas the time interval of interest could be as long as a month/year. The mean arrival rate to the data center in frame k is denoted by λ_k , and non-stationarities are captured by allowing different rates during different frames. Although we could allow λ_k to have a vector value to represent more than one type of workload as long as the resulting cost function is convex in our model, we assume λ_k to have a scalar value in this paper to simplify the presentation. Because the request inter-arrival times are much shorter than the frame length, typically in the order of 1-10 seconds, capacity provisioning can be based on the average arrival rate during a frame.

B. The Data Center Cost Model

The model for data center costs focuses on the server costs of the data center, as minimizing server energy consumption also reduces cooling and power distribution costs. We model the cost of a server by the operating costs incurred by an active server, as well as the switching cost incurred to toggle a server into and out of a power-saving model (e.g., off/on or sleeping/waking). Both components can be assumed to include energy cost, delay cost, and wear-and-tear cost. See [22] and [21] for further discussion of the model.

Note that this model ignores many issues surrounding reliability and availability, which are key components of data center service level agreements (SLAs). In practice, a solution that toggles servers must still maintain the reliability and availability guarantees; however this is beyond the scope of the current paper. See [27] for a discussion.

The Operating Cost: The operating costs are modeled by a convex function $f(\lambda_{i,k})$, which is the same for all the servers, where $\lambda_{i,k}$ denotes the average arrival rate to server i during frame k . The convexity assumption is quite general and captures many common server models. One example, which we consider in our numeric examples later, is to say that the operating costs are simply equal to the energy cost of the server, i.e., the energy cost of an active server handling arrival rate $\lambda_{i,k}$. This cost is often modeled using an affine function as follows

$$f(\lambda_{i,k}) = e_0 + e_1 \lambda_{i,k}, \quad (1)$$

where e_0 and e_1 are constants [1], [6], [19]. Note that when servers use dynamic speed scaling, if the energy cost is modeled as polynomial in the chosen speed, the cost $f(\cdot)$ remains convex. In practice, we expect that $f(\cdot)$ to be empirically measured by observing the system over time.

The Switching Cost: The switching cost, denoted by β , models the cost of toggling a server back-and-forth between active and power-saving models. The switching cost includes the costs of the energy used toggling a server, the delay in migrating connections/data when toggling a server, and the increased wear-and-tear on the servers toggling.

C. The Data Center Optimization

Given the cost model above, the data center has two control decisions at each time: determining n_k , the number of active servers in every time frame, and assigning arriving jobs to servers, i.e., determining $\lambda_{i,k}$ such that $\sum_{i=1}^{n_k} \lambda_{i,k} = \lambda_k$. All servers are assumed to be homogeneous with constant rate capacity $\mu > 0$. Modeling heterogeneous servers is also possible, as in [21]; however we limit the discussion in this paper to the homogeneous setting for clarity.

The goal of the data center is to determine n_k and $\lambda_{i,k}$ to minimize the cost incurred during $[0, K]$, which is modeled as follows:

$$\min \sum_{k=1}^K \sum_{i=1}^{n_k} f(\lambda_{i,k}) + \beta \sum_{k=1}^K (n_k - n_{k-1})^+ \quad (2)$$

$$s.t. \begin{cases} 0 \leq \lambda_{i,k} \leq \lambda_k \\ \sum_{i=1}^{n_k} \lambda_{i,k} = \lambda_k \\ \mathbb{P}(D_k > \bar{D}) \leq \bar{\epsilon} \end{cases} \quad (3)$$

where the final constraint is introduced to capture the SLA of the data center. We use D_k to represent the steady-state delay during frame k , and $(\bar{D}, \bar{\epsilon})$ to represent an SLA of the form “the probability of a delay larger than \bar{D} must be bounded by probability $\bar{\epsilon}$ ”.

This model generalizes the data center optimization problem from [22] by accounting for the additional SLA constraint. The specific values in this constraint are determined by the stochastic variability at the fast time-scale. In particular, we derive (for a variety of workload models) a sufficient constraint $n_k \geq \frac{C_k(\bar{D}, \bar{\epsilon})}{\mu}$ such that

$$n_k \geq \frac{C_k(\bar{D}, \bar{\epsilon})}{\mu} \implies \mathbb{P}(D_k > \bar{D}) \leq \bar{\epsilon}. \quad (4)$$

Here, μ is the constant rate capacity of each server and $C_k(\bar{D}, \bar{\epsilon})$ is to be determined for each considered arrival

model. One should interpret $C_k(\bar{D}, \bar{\varepsilon})$ as the overall effective capacity/bandwidth needed in the data center such that the SLA delay constraint is satisfied within frame k .

Note that the new constraint is only sufficient for the original SLA constraint. The reason is that $C_k(\bar{D}, \bar{\varepsilon})$ is computed, in the next section, from upper bounds on the distribution of the transient delay within a frame.

With the new constraint, however, the optimization problem in (2) can be considerably simplified. Indeed, note that n_k is fixed during each time frame k and the remaining optimization for $\lambda_{i,k}$ is convex. Thus, we can simplify the form of the optimization problem by using the fact that the optimal dispatching strategy $\lambda_{i,k}^*$ is load balancing, i.e., $\lambda_{1,k}^* = \lambda_{2,k}^* = \dots = \lambda_k/n_k$. This decouples dispatching $\lambda_{i,k}^*$ from capacity planning n_k , and so Eqs. (2)-(3) become:

$$\begin{aligned} \min \sum_{k=1}^K n_k f(\lambda_k/n_k) + \beta \sum_{k=1}^K (n_k - n_{k-1})^+ \quad (5) \\ \text{s.t. } n_k \geq \frac{C_k(\bar{D}, \bar{\varepsilon})}{\mu} . \end{aligned}$$

Note that (5) is a convex optimization, since $n_k f(\lambda_k/n_k)$ is the perspective function of the convex function $f(\cdot)$.

Note that the key difference between the optimization above, and that of [22], is the SLA constraint. It is this constraint that provides a bridge between the slow time-scale and fast time-scale models. Specifically, the fast time-scale model uses large deviations and stochastic network calculus techniques to calculate $C_k(\bar{D}, \bar{\varepsilon})$.

D. Algorithms for Dynamic Resizing

Though the Data Center Optimization Problem described above is convex, in practice it must be solved *online*, i.e., without knowledge of the future workload. Thus, in determining n_k , the algorithm may not have access to the future arrival rates λ_l for $l > k$. This fact makes developing algorithms for dynamic resizing challenging. However, progress has been made recently [22], [23].

Deriving algorithms for this problem is not the goal of the current paper. Thus, we make use of a recent algorithm called Lazy Capacity Provisioning (LCP) [22]. We choose LCP because of the strong analytic performance guarantees it provides – LCP provides cost within a factor of 3 of optimal for any (even adversarial) workload process.

LCP works as follows. Let $(n_{k,1}^L, \dots, n_{k,k}^L)$ be the solution vector to the following optimization problem

$$\begin{aligned} \min \sum_{l=1}^k n_l f(\lambda_l/n_l) + \beta \sum_{l=1}^k (n_l - n_{l-1})^+ \\ \text{s.t. } n_l \geq \frac{C_l(\bar{D}, \bar{\varepsilon})}{\mu} , n_0 = 0 . \end{aligned}$$

Similarly, let $(n_{k,1}^U, \dots, n_{k,k}^U)$ be the solution vector to the following optimization problem

$$\begin{aligned} \min \sum_{l=1}^k n_l f(\lambda_l/n_l) + \beta \sum_{l=1}^k (n_{l-1} - n_l)^+ \\ \text{s.t. } n_l \geq \frac{C_l(\bar{D}, \bar{\varepsilon})}{\mu} , n_0 = 0 . \end{aligned}$$

Denote $(n)_a^b = \max(\min(n, b), a)$ as the projection of n into the closed interval $[a, b]$. Then LCP can be defined using $n_{k,k}^L$ and $n_{k,k}^U$ as follows. Informally, LCP stays “lazily” between the upper bound $n_{k,k}^U$ and the lower bound $n_{k,k}^L$ in all frames.

Lazy Capacity Provisioning, LCP

Let $n^{LCP} = (n_0^{LCP}, \dots, n_K^{LCP})$ denote the vector of active servers under LCP. This vector can be calculated online using the following forward recurrence relation:

$$n_k^{LCP} = \begin{cases} 0, & k \leq 0 \\ (n_{k-1}^{LCP})_{n_{k,k}^L}^{n_{k,k}^U}, & 1 \leq k \leq K . \end{cases}$$

Note that, in [22], LCP is introduced and analyzed for the optimization from Eq. (5) without the SLA constraint. However, it is easy to see that the algorithm and performance guarantee extends to our setting. Specifically, the guarantees on LCP hold in our setting because the SLA constraint can be removed by defining the operating cost to be ∞ instead of $n_k f(\lambda_k/n_k)$ when $n_k < C_k(\bar{D}, \bar{\varepsilon})/\mu$.

A last point to highlight about LCP is that, as described, it does not use any predictions about the workload in future frames. Such information could clearly be beneficial, and can be incorporated into LCP if desired, see [22].

III. FAST TIME-SCALE MODEL

Given the model of the slow time-scale in the previous section, we now zoom in to give a description for the fast time-scale model. By “fast” time-scale, we mean the time-scale at which requests arrive, as opposed to the “slow” time-scale at which dynamic resizing decisions are made by the data center. To model the fast time-scale, we evenly break each frame from the slow time-scale into “slots” $t \in \{1, \dots, U\}$, such that $\text{frame_length} = U \cdot \text{slot_length}$.

We consider a variety of models for the workload process at this fast time-scale, including both light-tailed models with various degrees of burstiness, as well as heavy-tailed models that exhibit self-similarity. In all cases, our assumption is that the workload is stationary over the slots that make up each time frame.

The goal of this section is to derive the value of $C_k(\bar{D}, \bar{\varepsilon})$ in the constraint $n_k \geq \frac{C_k(\bar{D}, \bar{\varepsilon})}{\mu}$ from Eq. (4), and thus enable an interface between the fast and slow time-scales by parameterizing the Data Center Optimization Problem from Eq. (5) for a broad range of workloads.

Note that throughout this section we suppress frame’s subscript k for n_k , λ_k , C_k , and D_k , and focus on a generic frame.

Our approach for deriving the SLA constraint for the Data Center Optimization Problem is to first derive an ‘aggregation property’ which allows the data center to be modeled as a single server, and to then derive bounds on the distribution of the transient delay under a variety of arrival processes.

A. An Aggregation Property

Note that, if the arrival process were modeled as Poisson and job sizes were exponential, then an “aggregation property” would be immediate, since the response time distribution only depends on the load. Hence the SLA could be derived by considering a single server. Outside of this simple case, however, we need to derive a suitable single server approximation.

The aggregation result that we derive and apply is formulated in the framework of stochastic network calculus [9], and so we begin by briefly introducing this framework.

Denote the cumulative arrival (workload) process at the data center's dispatcher by $A(t)$. That is, for each slot $t = 1, \dots, U$, $A(t)$ counts the total number of jobs arrived in the time interval $[0, t]$. Depending on the total number n of active servers the arrival process is dispatched into the sub-arrival processes $A_i(t)$ with $i = 1, \dots, n$. The cumulative response processes from the servers are denoted by $R_i(t)$, whereas the total cumulative response process from the data center is denoted by $R(t) = \sum_i R_i(t)$. All arrival and response processes are assumed to be non-negative, non-decreasing, and left-continuous, and satisfy the initial condition $A(0) = R(0) = 0$. For convenience we use the bivariate extensions $A(s, t) := A(t) - A(s)$ and $R(s, t) := R(t) - R(s)$.

The service provided by a server is modeled in terms of probabilistic lower bounds using the concept of a stochastic service process. This is a bivariate random process $S(s, t)$ which is non-negative, non-decreasing, and left-continuous. Formally, a server is said to guarantee a (stochastic) service process $S(s, t)$ if for *any* arrival process $A(t)$ the corresponding response process $R(t)$ from the server satisfies for all $t \geq 0$

$$R(t) \geq A * S(t), \quad (6)$$

where $*$ denotes the min-plus convolution operator, i.e., for two (random) processes $A(t)$ and $S(s, t)$,

$$A * S(t) := \inf_{0 \leq s \leq t} \{A(s) + S(s, t)\}. \quad (7)$$

We are now ready to state the aggregation property. The proof is deferred to the technical report [28].

Lemma 1. *Consider an arrival process $A(t)$ which is dispatched to n servers. Each server i is work-conserving with constant rate capacity $\mu > 0$. Arrivals are dispatched deterministically across the servers such that each server i receives a fraction $\frac{1}{n}$ of the arrivals. Then, the system has service process $S(s, t) = n\mu(t - s)$, i.e., $R(t) \geq A * S(t)$.*

B. Arrival Processes

Now that we can reduce the study of the multi-server system to the study of a single server system using Lemma 1, we can move to characterizing the impact of the arrival process on the SLA constraint in the Data Center Optimization Problem.

In particular, the next step in deriving the SLA constraint $n \geq \frac{C(\bar{D}, \bar{\varepsilon})}{\mu}$ is to derive a bound on the distribution of the delay at the virtual server with arrival process $A(t)$ and service process $S(s, t) = C(\bar{D}, \bar{\varepsilon})(t - s)$, i.e.,

$$\mathbb{P}(D(t) > \bar{D}) \leq \varepsilon. \quad (8)$$

It is important to observe that the violation probability ε holds for the *transient* delay process $D(t)$, which is defined as $D(t) := \inf \{d : A(t - d) \leq R(t)\}$, and which models the delay spent in the system by the job leaving the system, if any, at time t . However, the violation probability ε is derived so that it is time invariant, which implies that it bounds the distribution of the steady-state delay $D = \lim_{t \rightarrow \infty} D(t)$ as well. Therefore, the value of $C(\bar{D}, \bar{\varepsilon})$ can be finally computed by solving the equation $\varepsilon = \bar{\varepsilon}$.

In the following, we follow the outline above to compute $C(\bar{D}, \bar{\varepsilon})$ for light- and heavy-tailed arrival processes. Interested readers may refer to Figure 1 in the technical report [28], which depicts examples of the three types of arrival processes we considered in 1 frame.

1) *Light-tailed Arrivals:* We consider two examples of light-tailed arrival processes: Poisson and Markov-Modulated (MM) processes. The latter is particularly interesting since it enables the adjustment of the burstiness level.

Poisson Arrivals: We start with the case of Poisson processes, which are characterized by a low level of burstiness, due to the independent increments property.

Let $A(t)$ be a Poisson process with some rate $\lambda > 0$, and define

$$\theta^* := \sup \left\{ \theta > 0 : \frac{\lambda}{\theta} (e^\theta - 1) \leq C(\bar{D}, \bar{\varepsilon}) \right\}. \quad (9)$$

Then a bound on the transient delay process is given for all $t \geq 0$ by (see [12])

$$\mathbb{P}(D(t) > \bar{D}) \leq e^{-\theta^* C(\bar{D}, \bar{\varepsilon}) \bar{D}} := \varepsilon. \quad (10)$$

Solving for $C(\bar{D}, \bar{\varepsilon})$ by setting the violation probability ε equal to $\bar{\varepsilon}$ yields the implicit solution

$$C(\bar{D}, \bar{\varepsilon}) = -\frac{1}{\theta^* \bar{D}} \log \bar{\varepsilon}.$$

Using the monotonicity of the function $\frac{\lambda}{\theta} (e^\theta - 1)$ in $\theta > 0$, and setting $K = -\frac{\log \bar{\varepsilon}}{\lambda \bar{D}}$, we immediately get the explicit solution

$$C(\bar{D}, \bar{\varepsilon}) = \frac{K}{\log(1 + K)} \lambda. \quad (11)$$

The alternative to using Eq. (10) is to use an exact result for the steady-state delay distribution at an M/D/1 queue. The advantage of the bound from Eq. (10), however, is that it readily leads to qualitative properties on the scaling of the capacity $C(\bar{D}, \bar{\varepsilon})$ (see Theorem 1); in addition, the bound is tight as shown in [12]. For these reasons we shall use the bound from Eq. (11) throughout.

Markov-modulated Arrivals: Consider now the case of Markov-Modulated (MM) processes which, unlike the Poisson processes, do not necessarily have independent increments. The key feature for the purposes of this paper is that the burstiness of MM processes can be arbitrarily adjusted.

We consider a simple MM processes with two states. Let a discrete and homogeneous Markov chain $x(s)$ with two states denoted by 'low' and 'high', and transition probabilities p_h and p_l between the 'low' and 'high' states, and vice-versa, respectively. Assuming that a source produces at some constant rates $\lambda_l > 0$ and $\lambda_h > \lambda_l$ while the chain $x(s)$ is in the 'low' and 'high' states, respectively, then the corresponding MM cumulative arrival process is

$$A(t) = \sum_{s=1}^t (\lambda_l I_{\{x(s)='low'\}} + \lambda_h I_{\{x(s)='high'\}}), \quad (12)$$

where $I_{\{\cdot\}}$ is the indicator function. The average rate of $A(t)$ is $\lambda = \frac{p_l}{p_h + p_l} \lambda_l + \frac{p_h}{p_h + p_l} \lambda_h$.

To adjust the burstiness level of $A(t)$ we introduce the parameter $T := \frac{1}{p_h} + \frac{1}{p_l}$, which is the average time for




Fig. 1. Illustration of the traces used for numerical experiments.

the Markov chain $x(s)$ to change states twice. We note that the higher the value of T is, the higher the burstiness level becomes (the time periods whilst $x(s)$ spends in the ‘high’ or ‘low’ states get longer and longer).

To compute the delay bound let us construct the matrix

$$\Psi(\theta) = \begin{pmatrix} (1-p_h)e^{\theta\lambda_l} & p_h e^{\theta\lambda_h} \\ p_l e^{\theta\lambda_l} & (1-p_l)e^{\theta\lambda_h} \end{pmatrix},$$

for some $\theta > 0$ and consider its spectral radius

$$\lambda(\theta) := \frac{(1-p_h)e^{\theta\lambda_l} + (1-p_l)e^{\theta\lambda_h} + \sqrt{\Delta}}{2},$$

where $\Delta = ((1-p_h)e^{\theta\lambda_l} - (1-p_l)e^{\theta\lambda_h})^2 + 4p_h p_l e^{\theta(\lambda_l + \lambda_h)}$. Let also

$$K(\theta) := \max \left\{ \frac{p_h e^{\theta\lambda_h}}{\lambda(\theta) - (1-p_h)e^{\theta\lambda_l}}, \frac{\lambda(\theta) - (1-p_h)e^{\theta\lambda_l}}{p_h e^{\theta\lambda_h}} \right\}.$$

The two terms are the ratios of the elements of the right-eigenvector of the matrix $\Psi(\theta)$. Also, let

$$\theta^* := \sup \left\{ \theta > 0 : \frac{1}{\theta} \log \lambda(\theta) \leq C(\bar{D}, \bar{\varepsilon}) \right\}.$$

Then a bound on the transient delay process immediately follows from the corresponding backlog bound (see [9], pp. 340) because of the constant rate service assumption

$$\mathbb{P}(D(t) > \bar{D}) \leq K(\theta^*) e^{-\theta^* C(\bar{D}, \bar{\varepsilon}) \bar{D}} := \varepsilon.$$

Setting the violation probability ε equal to $\bar{\varepsilon}$ yields the implicit solution

$$C(\bar{D}, \bar{\varepsilon}) = -\frac{1}{\theta^* \bar{D}} \log \frac{\bar{\varepsilon}}{K(\theta^*)}. \quad (13)$$

C. Heavy-tailed and Self-similar Arrivals

We now consider the class of heavy-tailed and self-similar arrival processes. These processes are fundamentally different from light-tailed processes in that deviations from the mean increase in time and decay in probability as a power law, i.e., more slower than the exponential.

As a concrete example of heavy-tailed and self-similar arrivals, consider the case of a source generating jobs in every slot according to i.i.d. Pareto random variables X_i with tail distribution $\mathbb{P}(X_i > x) = (x/b)^{-\alpha}$ for all $x \geq b$, where $1 < \alpha < 2$. X has finite mean $E[X] = \alpha b / (\alpha - 1)$ and infinite variance. The bound on the transient delay is then [20]

$$\mathbb{P}(D(t) > \bar{D}) \leq K(C(\bar{D}, \bar{\varepsilon}) \bar{D})^{1-\alpha} := \varepsilon, \quad (14)$$

where $K = \inf_{1 < \gamma < \frac{C(\bar{D}, \bar{\varepsilon})}{\lambda}} \left\{ \left(\frac{C(\bar{D}, \bar{\varepsilon})}{\gamma} - \lambda \right)^{-1} \frac{\alpha \gamma^{\frac{\alpha-1}{\alpha}}}{(\alpha-1) \log \gamma} \right\}$.

Setting the violation probability ε equal to $\bar{\varepsilon}$, we get the implicit solution (let $C(\cdot) := C(\bar{D}, \bar{\varepsilon})$)

$$\inf_{1 < \gamma < \frac{C(\cdot)}{\lambda}} \left\{ \frac{\gamma}{C(\cdot)^{\alpha-1} (C(\cdot) - \gamma \lambda)} \frac{\gamma^{\frac{\alpha-1}{\alpha}}}{\log \gamma^{\frac{\alpha-1}{\alpha}}} \right\} = \bar{\varepsilon} \bar{D}^{\alpha-1}. \quad (15)$$

The alternative to using Eq. (14) is to use large deviations results for the steady-state delay distribution [26]:

$$\mathbb{P}(D > \bar{D}) \sim \frac{\lambda}{C - \lambda} \frac{b^{\alpha-1}}{\alpha} (C \bar{D})^{1-\alpha}, \quad (16)$$

where $f(x) \sim g(x)$ stands for $\lim_{x \rightarrow \infty} f(x)/g(x) = 1$. Such asymptotic approximations, however, can be inaccurate for typical \bar{D} of interest [2]. Hence, we limit our focus to Eq. (14), which provides the state-of-the-art upper bound on the delay distribution in non-asymptotic regimes, though the same qualitative results follow when using Eq. (16).

IV. CASE STUDIES

Given the model described in the previous two sections, we are now ready to explore the potential of dynamic resizing in data centers, and how this potential depends on the interaction between non-stationarities at the slow time-scale and burstiness/self-similarity at the fast time-scale. Our goal in this section is to provide insight into which workloads dynamic resizing is valuable for. To accomplish this, we provide a mixture of analytic results and trace-driven numerical simulations in this section.

It is important to note that the case studies that follow depend fundamentally on the modeling performed so far in the paper, which allows us to capture and adjust independently, both fast time-scale and slow time-scale properties of the workload. The generality of our modeling framework enables thus a rigorous study of the impact of the workload on value of dynamic resizing.

A. Setup

Throughout the experimental setup, our aim is to choose parameters that provide conservative estimates of the case savings from dynamic resizing. Thus, one should interpret the savings shown as a lower-bound on the potential savings.

Model Parameters: The time frame for adapting the number of servers n_k is assumed to be 10 min, and each time slot is assumed to be 1 s, i.e., $U = 600$. When not otherwise specified, we assume the following parameters for the data center SLA agreement: the delay upper bound $\bar{D} = 200$ ms, and the delay violation probability $\bar{\varepsilon} = 10^{-3}$.

The cost is characterized by the two parameters of e_0 and e_1 , and the switching cost β . We choose units such that the fixed energy cost is $e_0 = 1$. The load-dependent energy consumption

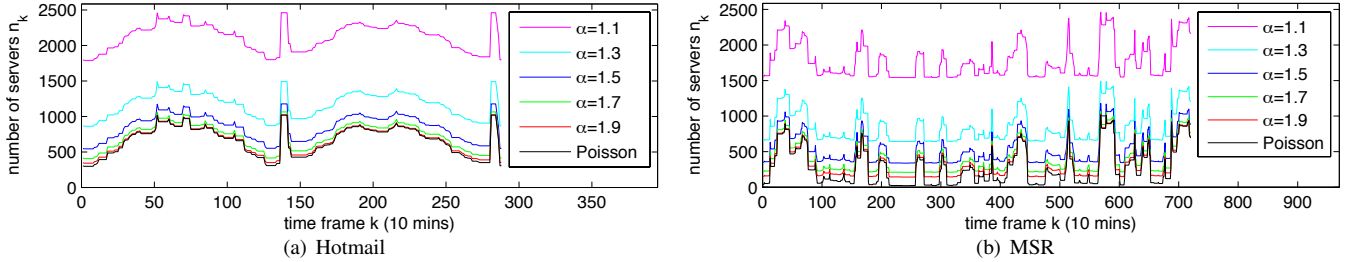


Fig. 2. Impact of burstiness on provisioning n_k for heavy-tailed arrivals.

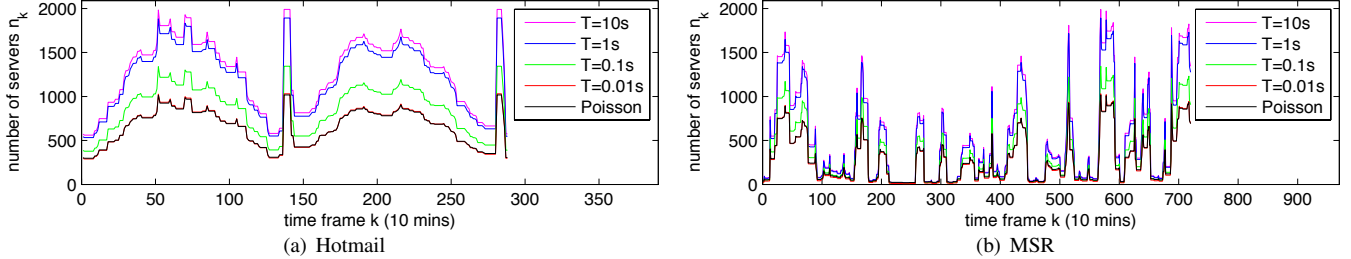


Fig. 3. Impact of burstiness on provisioning n_k for MM arrivals.

is set to $e_1 = 0$, because the energy consumption of current servers with typical utilization level is dominated by the fixed costs [1], [6], [19]. Note that adjusting e_0 and e_1 changes the magnitude of potential savings under dynamic resizing, but does not affect the qualitative conclusions about the impact of the workload. So, due to space constraints, we fix these parameters during the case studies.

The normalized switching cost β/e_0 measures the duration a server must be powered down to outweigh the switching cost. Unless otherwise specified, we use $\beta = 6$, which corresponds to the energy consumption for one hour (six frames). This was chosen as an estimate of the time a server should sleep so that the wear-and-tear of power cycling matches that of operating [7], [22].

Workload Information: The workloads for these experiments are drawn from two real-world data center traces. The first set of traces is from Hotmail, a large email service running on tens of thousands of servers. We used traces from 8 such servers over a 48-hour period, starting at midnight (PDT) on Monday August 4 2008 [27]. The second set of traces is taken from MSR Cambridge. The traced period was 1 week starting from 5PM GMT on the 22nd February 2007 [27]. Thus, these activity traces represent a service used by millions of users and a small service used by hundreds of users. The traces are normalized as peak load $\lambda_{peak} = 1000$, and are visualized in Figure 1. Both sets of traces show strong diurnal properties and have peak-to-mean ratios (PMRs) of 1.64 and 4.64 for Hotmail and MSR respectively. Loads were averaged over disjoint 10 minute frames.

The traces provide information for the slow time-scale model. To parameterize the fast time-scale model, we adapt the workload based on the mean arrival rate in each frame, i.e., λ . To parameterize the MM processes, we take $\lambda_l = 0.5\lambda$, $\lambda_h = 2\lambda$, and we adjust the burst parameter T while keeping λ fixed for each process. To parameterize the heavy-tailed processes, we adjust the tail index α for each process, while keeping the mean fixed at λ . Unless otherwise stated, we fix $\alpha = 1.5$ and $T = 1$.

Comparative Benchmark: We contrast three designs: (i) the optimal dynamic resizing, (ii) dynamic resizing via LCP, and (iii) the optimal ‘static’ provisioning.

The results for the optimal dynamic resizing should be interpreted as characterizing the potential of dynamic resizing. But, realizing this potential is a challenge that requires both sophisticated online algorithms and excellent predictions of future workloads.¹

The results for LCP should be interpreted as one example of how much of the potential for dynamic resizing can be attained with an online algorithm.

The results for the optimal static provisioning should be taken as an optimistic benchmark for today’s data centers, which typically do not use dynamic resizing. We consider the cost incurred by an optimal static provisioning scheme that chooses a constant number of servers that minimizes the costs incurred based on full knowledge of the entire workload. This policy is clearly not possible in practice, but it provides a *very conservative* estimate of the savings from right-sizing since it uses perfect knowledge of all peaks and eliminates the need for overprovisioning in order to handle the possibility of flash crowds or other traffic bursts.

B. Results

Our experiments are organized to illustrate the impact of a wide variety of parameters on the cost savings attainable via dynamic resizing. The goal is to understand for which workloads dynamic resizing can provide large enough cost savings to warrant the extra implementation complexity.

The Role of Burstiness: A key goal of our model is to expose the impact of burstiness on dynamic resizing, and so we start by focusing on that parameter. Recall that we can vary burstiness in both the light-tailed and heavy-tailed settings using T for MM arrivals and α for heavy-tailed arrivals.

A priori, one may expect that burstiness can be beneficial for dynamic resizing, since it indicates that there are periods of low load during which energy may be saved. However, this is

¹Note that short-term predictions of workload demand within 24 hours can be quite accurate [17], [19].




Fig. 4. Impact of burstiness on the cost savings of dynamic resizing for different switching costs, β .




Fig. 5. Impact of burstiness on the performance of LCP in the Hotmail trace. not actually true since resizing decisions must be made at the slow time-scale while burstiness is a characteristic of the fast time-scale. Thus, burstiness is actually detrimental for dynamic resizing, since it means that the provisioning decisions made on the slow time-scale must be made with the bursts in mind, which results in an larger number of servers needed to be provisioned for the same average workload. This effect can be seen in Figures 2 and 3, which show the optimal dynamic provisioning as α and T vary. Recall that burstiness increases as α decreases and T increases.

The larger provisioning created by increased burstiness manifests itself in the cost savings attainable through dynamic capacity provisioning as well. This is illustrated in Figure 4, which shows the cost savings of the optimal dynamic provisioning as compared to the optimal static provisioning for varying α and T as a function of the switching cost β .

Interestingly, though Figure 4 shows that the potential of dynamic resizing is limited by increased burstiness, it turns out that the relative performance of LCP is not hurt by burstiness. This is illustrated in Figure 5, which shows the percent of the optimal cost savings that LCP achieves. Importantly, it is nearly perfectly flat as the burstiness is varied.

The Role of the Peak-to-Mean Ratio: The impact of the peak-to-mean ratio on the potential benefits of dynamic resizing is quite intuitive: if the peak-to-mean ratio is high, then there is more opportunity to benefit from dynamically changing capacity. Figure 6 illustrates this well-known effect. The workload for the figure is generated from the traces by scaling λ_k as $\hat{\lambda}_k = c(\lambda_k)^\gamma$, varying γ and adjusting c to keep the mean constant.

In addition to illustrating that a higher peak-to-mean ratio makes dynamic resizing more valuable, Figure 6 also highlights that there is a strong interaction between burstiness and the peak-to-mean ratio, where if there is significant burstiness the benefits that come from a high peak-to-mean ratio may be diminished considerably.

The Role of the SLA: The SLA plays a key role in the provisioning of a data center. Here, we show that the SLA can also have a strong impact on whether dynamic resizing is valuable, and that this impact depends on the workload. Recall




Fig. 6. Impact of peak-to-mean ratio on the cost savings of the optimal dynamic resizing.

that in our model the SLA consists of a violation probability $\bar{\epsilon}$ and a delay bound \bar{D} . We deal with each of these in turn.

Figures 7 and 8 highlight the role the violation probability $\bar{\epsilon}$ has on the provisioning of n_k under the optimal dynamic resizing in the cases of heavy-tailed and MM arrivals. Interestingly, we see that there is a significant difference in the impact of $\bar{\epsilon}$ depending on the arrival process. As $\bar{\epsilon}$ gets smaller in the heavy-tailed case the provisioning gets significantly flatter, until there is almost no change in n_k over time. In contrast, no such behavior occurs in the MM case and, in fact, the impact of $\bar{\epsilon}$ is quite small. This difference is a fundamental effect of the “heaviness” of the tail of the arrivals, i.e., a heavy tail requires significantly more capacity in order to counter a drop in $\bar{\epsilon}$.

This contrast between heavy- and light-tailed arrivals is also evident in Figure 9, which highlights the cost savings from dynamic resizing in each case as a function of $\bar{\epsilon}$. Interestingly, the cost savings under light-tailed arrivals is largely independent of $\bar{\epsilon}$, while under heavy-tailed arrivals the cost savings is monotonically increasing with $\bar{\epsilon}$.

When is Dynamic Resizing Valuable?: Now, we are finally ready to address the question of when (i.e., for what workloads) is dynamic resizing valuable. To address this question, we must look at the interaction between the peak-to-mean ratio and the burstiness. Our goal is to provide a concrete understanding of for which (peak-to-mean, burstiness, SLA) settings the potential savings from dynamic resizing is large enough to warrant implementation. Figures 10–12 focus on this question. Our hope is that these figures highlight that a precursor to any debate about the value of dynamic resizing must be a joint understanding of the expected workload characteristics and the desired SLA, since for any fixed choices of two of these parameters (peak-to-mean, burstiness, SLA), the third can be chosen so that dynamic resizing does or does not provide significant cost savings for the data center.

Starting with Figure 10, we see a set of curves for different levels of cost savings. The interpretation of the figures is that below (above) each curve the savings from optimal dynamic resizing is smaller (larger) than the specified value for the curve. Thus, for example, if the peak-to-mean ratio is 2 in the Hotmail trace, a 10% cost savings is possible for all levels of burstiness, but a 30% cost savings is only possible for $\alpha > 1.5$. However, if the peak-to-mean ratio is 3, then a 30% cost savings is possible for all levels of burstiness. It is difficult to say what peak-to-mean and burstiness settings are “common” for data centers, but as a point of reference, one might expect large-scale services to have a peak-to-mean ratio similar to that of the Hotmail trace, i.e., around 1.5–2.5; and smaller scale services to have peak-to-mean ratios similar to that of the MSR trace, i.e., around 4–6. The burstiness also can vary

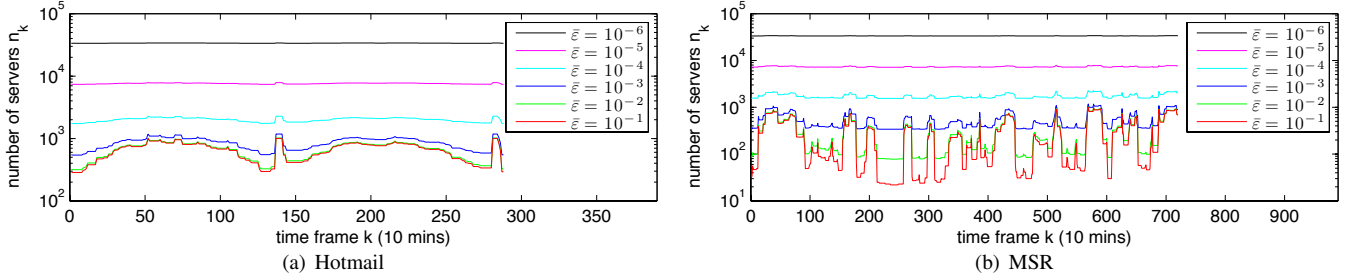
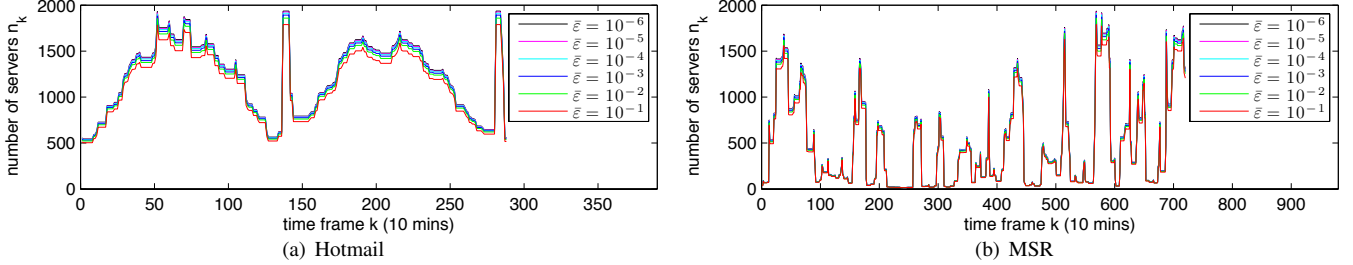


Fig. 7. Impact of $\bar{\varepsilon}$ on provisioning n_k for heavy tailed arrivals.Fig. 8. Impact of $\bar{\varepsilon}$ on provisioning n_k for MM arrivals.Fig. 9. Impact of $\bar{\varepsilon}$ on the cost savings of dynamic resizing.



Fig. 10. Characterization of burstiness and peak-to-mean ratio necessary for dynamic resizing to achieve different levels of cost reduction.

widely, but as a rough estimate, one might expect α to be around 1.4-1.6.

Of course, many of the settings of the data center will effect the conclusions illustrated in Figure 10. Two of the most important factors to understand the effects of are the switching cost, β , and the SLA, particularly $\bar{\varepsilon}$.

Figure 11 highlights the impact of the magnitude of the switching costs on the value of dynamic resizing. The curves represent the threshold on peak-to-mean ratio and burstiness necessary to obtain 20% cost savings from dynamic resizing. As the switching costs increase, the workload must have a larger peak-to-mean ratio and/or less burstiness in order for dynamic resizing to be valuable. This is not unexpected. However, what is perhaps surprising is the small impact played by the switching cost. The class of workloads where dynamic resizing is valuable only shrinks slightly as the switching cost is varied from on the order of the cost of running a server for 10 minutes ($\beta = 1$) to running a server for 3 hours ($\beta = 18$).

Interestingly, while the impact of the switching costs on the

Fig. 11. Characterization of burstiness and peak-to-mean ratio necessary for dynamic resizing to achieve 20% cost reduction as a function of the β .Fig. 12. Characterization of burstiness and peak-to-mean ratio necessary for dynamic resizing to achieve 20% cost reduction as a function of the SLA, $\bar{\varepsilon}$.

value of dynamic resizing is small, the impact of the SLA is quite large. In particular, the violation probability $\bar{\varepsilon}$ can dramatically affect whether dynamic resizing is valuable or not. This is shown in Figure 12, on which the curves represent the threshold on peak-to-mean ratio and burstiness necessary to obtain 20% cost savings from dynamic resizing. We see that, as the violation probability is allowed to be larger, the impact of the peak-to-mean ratio on the potential of savings from dynamic resizing disappears; and the value of dynamic resizing starts to depend almost entirely on the burstiness of the arrival process. The reason for this can be observed in Figure 7, which highlights that the optimal provisioning n_k becomes nearly flat as $\bar{\varepsilon}$ increases.

Supporting Analytic Results: To this point we have focused on numerical simulations, and further we provide analytic support for the behavior we observed in the experiments above. In particular, the following two theorems characterize the impact of burstiness and the SLA (\bar{D} , $\bar{\varepsilon}$) on the value of

dynamic resizing under Poisson and heavy-tailed arrivals. The proofs are given in the technical report [28].

Theorem 1. *The service capacity constraint from Eq. (11) increases as the delay constraint \bar{D} or the violation probability $\bar{\varepsilon}$ decrease. It also satisfies the scaling law*

$$C(\bar{D}, \bar{\varepsilon}) = \Theta \left(\frac{\bar{D}^{-1} \log \bar{\varepsilon}^{-1}}{\log(\bar{D}^{-1} \log \bar{\varepsilon}^{-1})} \right),$$

as $\bar{D}^{-1} \log \bar{\varepsilon}^{-1} \rightarrow \infty$.

This theorem highlights that as $\bar{\varepsilon}$ decreases and/or \bar{D} decreases $C(\bar{D}, \bar{\varepsilon})$, and thus the cost of the optimal provisioning, increases. This shows that the observations made in our numeric experiments hold more generally. Perhaps the most interesting point about this theorem, however, is the contrast of the growth rate with that in the case of heavy-tailed arrivals, which is summarized in the following theorem.

Theorem 2. *The implicit solution for the capacity constraint from Eq. (15) increases as the delay constraint \bar{D} or the violation probability $\bar{\varepsilon}$ decrease, or the value of α decreases. It also satisfies the scaling law*

$$C(\bar{D}, \bar{\varepsilon}) = \Theta \left(\left(\frac{1}{\bar{\varepsilon} \bar{D}^{\alpha-1}} \right)^{\frac{1}{\alpha}} \right) \quad (17)$$

as $\bar{\varepsilon} \bar{D}^{\alpha-1} \rightarrow 0$ for any given $\alpha \in (1, 2)$.

A key observation about this theorem is that the growth rate of $C(\bar{D}, \bar{\varepsilon})$ with $\bar{\varepsilon}$ is much faster than in the case of the Poisson (polynomial instead of logarithmic). This supports what is observed in Figure 9. Additionally, Theorem 2 highlights the impact of burstiness, α , and shows that the behavior we have seen in our experiments holds more generally.

V. CONCLUSION

Our goal in this paper is to provide new insight into the debate about the potential of dynamic resizing in data centers. Clearly, there are many facets of this issue relating to the engineering, algorithmic, and reliability challenges involved in dynamic resizing which we have ignored in this paper. These are all important issues when trying to *realize the potential* of dynamic resizing. But, the point we have made in this paper is that when *quantifying the potential* of dynamic resizing it is of primary importance to understand the joint impact of workload and SLA characteristics.

To make this point, we have presented a new model that captures the impact of SLA characteristics in addition to both slow time-scale non-stationarities in the workload and fast time-scale burstiness in the workload. This model allows us to provide the first study of dynamic resizing that captures both the stochastic burstiness and diurnal non-stationarities of real workloads.

REFERENCES

- [1] SPEC power data on SPEC website at <http://www.spec.org>.
- [2] J. Abate, G. L. Choudhury, and W. Whitt. Waiting-time tail probabilities in queues with long-tail service-time distributions. *Queueing Systems*, 16(3-4):311–338, Sept. 1994.
- [3] F. Ahmad and T. N. Vijaykumar. Joint optimization of idle and cooling power in data centers while maintaining response time. In *Proc. of the 15th International Conference on Architectural support for programming languages and operating systems (ASPLOS)*, 2010.
- [4] S. Albers. Energy-efficient algorithms. *Communications of the ACM*, 53(5):86–96, 2010.
- [5] H. Amur, J. Cipar, V. Gupta, G. R. Ganger, M. A. Kozuch, and K. Schwan. Robust and flexible power-proportional storage. In *Proc. of the 1st ACM symposium on Cloud computing (SoCC)*, 2010.
- [6] L. A. Barroso and U. Hözlze. The case for energy-proportional computing. *Computer*, 40(12):33–37, 2007.
- [7] P. Bodik, M. P. Armbrust, K. Canini, A. Fox, M. Jordan, and D. A. Patterson. A case for adaptive datacenters to conserve energy and improve reliability. Technical Report UCB/EECS-2008-127, University of California at Berkeley, 2008.
- [8] E. V. Carrera, E. Pinheiro, and R. Bianchini. Conserving disk energy in network servers. In *Proc. of the 17th annual international conference on Supercomputing (ICS)*, pages 86–97. ACM, 2003.
- [9] C.-S. Chang. *Performance Guarantees in Communication Networks*. Springer Verlag, 2000.
- [10] J. S. Chase, D. C. Anderson, P. N. Thakar, A. M. Vahdat, and R. P. Doyle. Managing energy and server resources in hosting centers. In *Proc. of the 8th ACM symposium on Operating systems principles (SOSP)*, pages 103–116, 2001.
- [11] G. Chen, W. He, J. Liu, S. Nath, L. Rigas, L. Xiao, and F. Zhao. Energy-aware server provisioning and load dispatching for connection-intensive internet services. In *Proc. of the 5th USENIX Symposium on Networked Systems Design & Implementation (NSDI)*, 2008.
- [12] F. Ciucu. Network calculus delay bounds in queueing networks with exact solutions. In *Proc. of the 20th International Teletraffic Congress (ITC)*, 2007.
- [13] C. Clark, K. Fraser, S. Hand, J. G. Hansen, E. Jul, C. Limpach, I. Pratt, and A. Warfield. Live migration of virtual machines. In *Proc. of the 2nd USENIX Symposium on Networked Systems Design & Implementation (NSDI)*, pages 273–286, 2005.
- [14] B. Diniz, D. Guedes, W. Meira, Jr., and R. Bianchini. Limiting the power consumption of main memory. In *Proc. of the 34th annual international symposium on Computer architecture (ISCA)*, pages 290–301, 2007.
- [15] A. Gandhi, V. Gupta, M. Harchol-Balter, and M. A. Kozuch. Optimality analysis of energy-performance trade-off for server farm management. *Performance Evaluation*, 67(11):1155 – 1171, 2010.
- [16] A. Gandhi, M. Harchol-Balter, and M. A. Kozuch. The case for sleep states in servers. In *Proc. of the 4th Workshop on Power-Aware Computing and Systems (HotPower)*, pages 2:1–2:5, 2011.
- [17] D. Gmach, J. Rolia, L. Cherkasova, and A. Kemper. Workload analysis and demand prediction of enterprise data center applications. In *Proc. of the 10th IEEE International Symposium on Workload Characterization (IISWC)*. IEEE Computer Society, 2007.
- [18] A. Greenberg, J. Hamilton, D. A. Maltz, and P. Patel. The cost of a cloud: research problems in data center networks. *SIGCOMM Comput. Commun. Rev.*, 39:68–73, December 2008.
- [19] D. Kusic, J. O. Kephart, J. E. Hanson, N. Kandasamy, and G. Jiang. Power and performance management of virtualized computing environments via lookahead control. In *Proc. of the 5th IEEE International Conference on Autonomic Computing (ICAC)*, 2008.
- [20] J. Liebeherr, A. Burchard, and F. Ciucu. Delay bounds in communication networks with heavy-tailed and self-similar traffic. *IEEE Transactions on Information Theory*, 58(2):1010–1024, Feb. 2012.
- [21] M. Lin, Z. Liu, A. Wierman, and L. L. Andrew. Online algorithms for geographical load balancing. In *Proc. of the 3rd International Green Computing Conference (IGCC)*, 2012.
- [22] M. Lin, A. Wierman, L. L. H. Andrew, and E. Thereska. Dynamic right-sizing for power-proportional data centers. In *Proc. of the 30th IEEE International Conference on Computer Communications (Infocom)*, 2011.
- [23] M. Lin, A. Wierman, L. L. H. Andrew, and E. Thereska. Online dynamic capacity provisioning in data centers. In *Proc. of the 49th Annual Allerton Conference on Communication, Control, and Computing (Allerton)*, 2011.
- [24] D. Meisner, B. T. Gol, and T. F. Wenisch. PowerNap: Eliminating server idle power. *ACM SIGPLAN Notices*, 44(3):205–216, Mar. 2009.
- [25] D. Meisner, C. M. Sadler, L. A. Barroso, W.-D. Weber, and T. F. Wenisch. Power management of online data-intensive services. In *Proc. of the 38th annual international symposium on Computer architecture (ISCA)*, 2011.
- [26] A. Pakes. On the tails of waiting-time distributions. *Journal of Applied Probability*, 3(12):555–564, Sept. 1975.
- [27] E. Thereska, A. Donnelly, and D. Narayanan. Sierra: a power-proportional, distributed storage system. Technical Report MSR-TR-2009-153, Microsoft Research, 2009.
- [28] K. Wang, M. Lin, F. Ciucu, A. Wierman, and C. Lin. Characterizing the impact of the workload on the value of dynamic resizing in data centers. Technical Report ISCAS-2012-GD1, available at <http://arxiv.org/abs/1207.6295>, 2012.