

# QPRED: Using Quantile Predictions to Improve Power Usage for Private Clouds

## *UCSB Computer Science Technical Report 2014-06*

Rich Wolski

Computer Science Department, University of  
California, Santa Barbara  
rich@cs.ucsb.edu

John Brevik

Department of Mathematics, California State  
University, Long Beach  
John.Brevik@csulb.edu

### Abstract

In this paper we describe a new, efficient predictive scheduling methodology for implementing computing infrastructure power savings using private clouds. Our approach, termed “QPRED,” estimates the quantiles on the distribution of future machine usage so that unneeded machines may be powered down to save power. A cloud administrator sets a bound on the probability that all available machines will be powered down when a cloud request arrives. This target probability is the basis of a Service Level Agreement between the cloud administrator and all cloud users covering start-up delay resulting from power savings. Our results, validated using activity traces from several private clouds used in commercial production, indicate that QPRED successfully reduces power consumption substantially while maintaining the SLAs specified by the cloud administrator.

**Categories and Subject Descriptors** D.4.7 [*Organization and Design*]: Distributed Systems

**Keywords** private cloud, load prediction, power optimization

### 1. Introduction

Cloud computing, in the form of “Infrastructure as a Service” (IaaS), has emerged as a new methodology for organizations to manage digital assets and the physical computing infrastructure that hosts them. Public clouds, such as Amazon’s AWS [17] and Google Cloud Platform [14], rent virtual machines (VMs), network connectivity, and storage

via web services APIs over the Internet. Customers of public clouds use an e-commerce-style interface to obtain these rentals in a way that is fully automated and self-service.

On the other hand, private clouds built using technologies such as Eucalyptus [6, 21], OpenStack [22], and CloudStack [1] operate in private datacenters, each under the control of an organization’s Information Technology (IT) staff. They offer the same automated self-service interfaces as public clouds but to employees under a quota-controlled charge-back accounting system rather than to billed customers. Thus private clouds are a way of using e-commerce technologies to automate and streamline IT management of private datacenters through e-commerce-style self-service.

In this paper, we describe a scheduling methodology that is designed to save electrical power in private cloud settings using on-line, non-parametric predictions of future demand. Private cloud operators must be able to offer Service Level Agreements (SLAs) to their users so that these users can reason about how applications will behave when they are hosted, just as they do when these applications are run on physical infrastructure. Our approach allows the cloud administrator to make a probabilistic guarantee regarding the impact that power saving will have on user experience.

Clouds, by their very nature, obscure the specific infrastructure characteristics from the infrastructure users in the form of abstractions. Users reason about cloud use in terms of SLAs associated with its abstractions and experience the cloud in terms of delivered performance. Powering off idle servers carries with it the potential for a user-perceived delay during virtual machine (VM) start-up if a physical server needs to be powered on before the VM can start. Deskside-and laptop-class hardware can “hibernate,” thereby minimizing this delay, but full hibernation support is not available for many commercial-grade servers. As a result, unused machines must be fully powered down to save power. With large memories and disk subsystems, the power-up delay associated server class machines can be significant (tens of minutes in some cases). Thus our methodology attempts to save as much power as possible subject to a probabilistic SLA that

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

Submission to SoCC '14, October, 2014, Seattle, WA, USA.

the cloud administrator sets with respect to the additional power-up delay a user might experience which is advertised to her or his users.

The key to our approach is the ability to make a conservative prediction of an arbitrary quantile from the distribution of machines that will be needed a short time into the future. Inspired by our earlier work with QBETS [19, 20] to predict bounds on batch-queue delay, this work uses a new fast, non-parametric prediction algorithm to estimate quantile bounds from measurement samples over fixed time epochs. Our methodology monitors cloud activity and uses the quantile prediction to estimate how many “hot spares” will be needed to host VMs that will be requested in the next time epoch. All other machines in the cloud not in use are then powered down.

The methodology is novel in that it does not rely on the *a priori* assumption that the “random” quantities of interest obey well-behaved and simple statistical distributions (*e.g.*, that process lifetimes are exponentially distributed). Comparable approaches [8, 9] employ sophisticated statistical models for queue wait times, workload, *etc.* based on distributional assumptions that enable computational tractability. In our experience [30, 31], in a sufficiently general private cloud setting these simple models are insufficient to describe the underlying distributions. As an alternative, our method uses a computationally simple non-parametric technique to estimate the statistical quantities that are necessary to make a prediction of future demand with results that appear to be comparable or better.

When a request to start a VM is initiated, the methodology first looks for a machine that is already powered up to host the VM. If no such machine is available, the scheduler will delay the VM request pending the power up of a dormant machine and the power-up time is experienced by the user as additional start-up delay. The quantile estimate allows the cloud administrator to set the maximum probability that no machine will be powered up and ready when a user request for service is initiated. The result is that the user experience is perturbed by a predictable fraction of the total request population. That is, the probability of finding no machine powered-up and available (defined by the quantile the administrator chooses) determines the maximum fraction of total user requests that will experience some form of delay. We term these quantile-predicting schedulers *QPRED* schedulers.

We validate the overall *QPRED* scheduling methodology using Eucalyptus [6, 21], an open-source platform for implementing private clouds in production datacenters. Eucalyptus allows the cloud administrator to partition the infrastructure resources into “Availability Zones” (AZs), each governed by a separate scheduler. The zones are then composed into a scalable cloud using an eventually consistent [27] data model that is shared among the schedulers. Our solution works at the AZ level (*i.e.*, under the assumption that

each AZ is characterized by its own Service Level Agreements) and is compatible with the performance and reliability specifications for Eucalyptus zone schedulers. In well-designed private clouds, AZs represent pools of equivalent resources governed by a single set of Service Level Agreements (SLAs). Thus while the cloud itself may be large, the number of resources in a AZ is typically small to allow administrators to configure a large number of different SLA pools. In this study, the node counts are modest with respect to overall cloud size but typical of production private cloud deployments in terms of AZ size.

Eucalyptus is used commercially, and several of its commercial customers have made traces of their respective workloads available (under the condition of anonymity) for the purposes of evaluating our approach. Thus the results we present herein depict effects that are observed from “real-world” production private cloud settings. In addition to understanding the degree to which power management could benefit this category of private cloud usage, we are also interested in an algorithm that can be made to work “on-line” as part of the resource scheduling implementation. *QPRED* uses a short history of single-valued measurements (typically no more than 1000) that it must consult in sorted order and an incrementally updated running calculation of the average cloud request interarrival time. Thus its implementation can be made highly efficient with respect to computational and required memory state.

Our results indicate that the *QPRED* methodology can result in substantial power savings in the form of powered-off node time while allowing the cloud administrator to ensure that the impact on user experience is minimal. We explore the boundaries of its capabilities, including the impact on efficacy on variations in the time and power requirements associated with powering on dormant machines.

Thus, while the problem of power management in data centers has been extensively studied [2, 3, 8, 26, 32] our work is the first to detail the efficacy of an efficient on-line statistical prediction strategy using production commercial private cloud workloads. It is unique in its use of an algorithm that can be implemented with minimal computational and storage requirements.

While our validation uses Eucalyptus only, we note that many of the popular private cloud platforms such as OpenStack, CloudStack, vCloud Director [13], Nimbus [11] and OpenNebula [16] share architectural characteristics with Eucalyptus, particularly with respect to zone scheduling. Thus we believe that these results generalize to other private cloud platforms as they are implemented today.

The remainder of this paper details *QPRED* scheduling, outlines the experimental methodology we have used to validate it, and investigates both the power savings and user impact *QPRED* would have had in several production cloud settings had it been available.

## 2. VM Scheduling and Power Consumption

Because private clouds must be able to manage workloads scalably, their scheduling algorithms must be efficient. Eucalyptus, for example, only includes schedulers (Greedy and Round-robin) that assign a VM to a node at the time the request for the VM arrives at the cloud from the user. Further, because VM migration can require substantial intra-cloud bandwidth, each scheduler makes only a single placement decision for a VM at the beginning of a VM’s lifetime.

Using only the ability to power machines on and off, the problem of optimizing power usage in this scheduling scenario without denying access (*i.e.*, turning away VM requests when machines are available but powered down) can be solved trivially: All machines are powered down until they are needed to run a VM. When a VM request arrives at the scheduler, the scheduler attempts to assign it to a node <sup>1</sup> that is already powered up. If no node is located, the scheduler chooses a node that is powered down, sends it a power-up signal, and launches the VM on the node once it has been successfully powered on. If the scheduler uses a “greedy” strategy – one that “fills” nodes with incoming VM requests before selecting a new node – this methodology is optimal with respect to power consumption under the constraints that

- each VM is considered once
- the scheduler makes only one placement decision for each VM at the time the VM start request arrives, and
- no additional information beyond what is needed to determine the capacity required for each VM is provided.

The scheduling complexity of such schedulers is  $O(n * m)$  for  $n$  VMs and  $m$  machines (each of  $m$  machines might need to be considered for each of  $n$  VMs worst case). Because  $n \gg m$  in most cloud settings, we consider this to be  $O(n)$  complexity.

Because this strategy waits to power up nodes until they are needed, VMs that cannot be started until a node has been fully powered on must also wait; this added delay is experienced by users until their VMs become available for use. Machine power-up times, particularly for server-class machines, can be lengthy: Depending on the machine’s configuration, it may require as much as 30 minutes to go from a powered-off state to one in which a VM can be started. Moreover, Eucalyptus makes heavy use of caching and copy-on-write techniques to reduce VM launch times. A cached 10-gigabyte VM can be launched in under a minute if the local disks are server class. Thus a scheduling strategy that tries to optimize power usage may also introduce VM start-up overhead that is dramatic and may be unacceptable for some applications or users.

<sup>1</sup> We will use the terms “machine” and “node” interchangeably to refer to a machine configured into a cloud that is running a hypervisor and can host a VM that is started and terminated by a user making requests to the cloud.

We formulate the problem of moderating power consumption in terms of a tradeoff between the probability that a VM (and its user) will experience a start-up delay and the power saved by having machines powered down. Well-written cloud applications are typically prepared for variation in VM start-up delay as long as the delays occur relatively infrequently. Thus our approach is to allow the cloud administrator to set a maximum target probability for any given VM to experience a start-up delay because a machine needs to be powered up. The scheduler must then keep enough “extra” machines (“hot spares”) powered on so that the probability that a VM start request will arrive while no powered-on machines are available is at or below the target. At the same time, the scheduler must maintain  $O(n)$  complexity to avoid introducing unacceptable or unpredictable overhead if the load scales.

Notice that this formulation of the scheduling problem prioritizes user experience in the form of minimized VM start-up delay over power savings. That is, we investigate schedulers that are designed to implement a Service Level Agreement (SLA) in terms of VM start-up times between the cloud’s users and the cloud’s operators while at the same time minimizing power usage subject to the SLA. This user-centric approach based on SLAs is typical for private cloud deployments.

Notice also that simply keeping a single or a fixed number of hot spares may not provide enough additional powered-on capacity if VM arrivals fill and exceed the capacity of the spares before a new spare can be fully powered on. As an example, suppose that the scheduler attempts to keep a single hot spare available, that each node in a cloud can host 8 single-core VMs, and that the machine power-up delay is 600 seconds. If 16 single-core requests arrive in a 600-second interval and there is only one hot spare, at least one VM will experience a start-up delay. Further, the cloud administrator cannot predict nor control the rate at which VMs (and users) experience start-up delay with this approach, making it difficult to provide a reliable SLA.

Thus, we investigate  $O(n)$  scheduling methods that make a *prediction* of the number of additional machines that must be powered on at any moment so that the maximum target probability specified by the cloud administrator (*i.e.*, the SLA) for VM start-up delay will not be exceeded.

### Scheduler Operation

Eucalyptus (in a way that is similar to other private clouds) divides machines that are capable of hosting VMs into pools (called “Availability Zones” Each Availability Zone (AZ) in a particular Eucalyptus installation has its own VM scheduler (a Eucalyptus cloud may be configured with multiple AZs).

When a VM start request arrives the scheduler makes a placement decision only that determines the specific node within the AZ that should host the VM. Eucalyptus does support VM migration, but only under the control of the cloud

administrator to keep user-induced migrations from causing inter-VM communication to degrade due to migration traffic. Further, if the machine hosting the scheduler fails, and if the cloud is configured for highly available operation, a backup scheduler will detect the failure and take over scheduling responsibilities for the AZ.

For Eucalyptus (and other well-designed cloud deployments) scheduling at the AZ level is important for two reasons. First, the scheduling methodologies we investigate must have internal state that is small and easily transferred or reconstructed in the event that a scheduler failure occurs (*i.e.*, to support “hot” fail over of the scheduler. The scheduler state must be replicated on the hot spare and the network load induced by this replication must be low.

Second, for private clouds, the AZ corresponds to a pool of equivalent resources described by a single set of SLAs. The cloud administrators configure different AZs according to the performance and reliability SLAs that they wish to expose to their users. For this reason, scheduling is implemented at the AZ level (so that different AZs can have different utilization profiles).

Note that for the purpose of this study we have been given log data corresponding to a single AZ in each case. While clouds may be large, AZs are often modest in size, since they represent different SLA partitions to the cloud users and there may be many such partitions in a private cloud. Eucalyptus is designed to support larger AZs than the ones for which we have trace data (hence the emphasis on efficient scheduling in its implementation). However, for the commercial production private clouds we have been able to observe, the AZ size, at present, is modest.

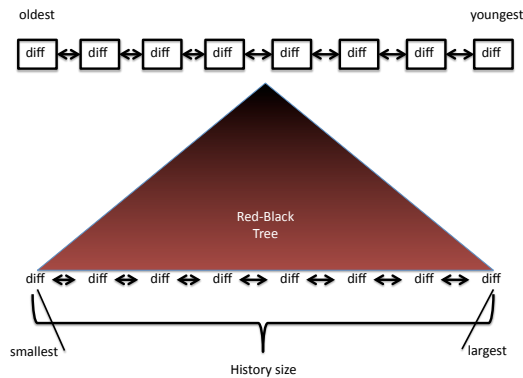
### The Prediction Method

The goal is to use the information provided by the history of node occupancy to predict the number of nodes that will be required going forward and therefore the number of hot spares to keep on hand. To this end, we poll the system at regular intervals. To be sure, over a particular time interval there will likely be points in time when there are fewer or more nodes occupied; the number with which we will be concerned is the *maximum* number of nodes simultaneously occupied during the time interval.

In principle, the time-series information used for this inference will consist of an  $(N + 1)$ -by- $(N + 1)$  matrix of transition probabilities (for a cloud with  $N$  nodes) between all possible numbers, including 0, of occupied nodes. This formulation defines a large number of transition probabilities, even for a modest-sized cloud, to estimate from a sample of any reasonable size. (On the other hand, the majority of the transitions are vanishingly improbable, as the transitions themselves will tend to be small, provided that the sampling interval is short enough, so the number of useful probabilities to be estimated, while still substantial, is not extremely large.)

In the data sets we have studied, the transition probabilities are almost completely captured by the probabilities of *differences* from one interval to the next. That is, for example, given that there are 4 occupied nodes in one interval, the probability of going to 6 at the next interval is very nearly identical to the probability of going from 1 to 3 or 5 to 7. This behavior allows us to use a much simpler time series, namely that of the difference in the number of occupied nodes from one time step to the next. This simplification, in fact, removes the time-series character from the problem entirely: If we would like to be, say, 95% certain that we will have enough hot spares to handle the incoming jobs for the next time interval, we need only look at some estimate (in the statistical sense) for the 0.95 quantile of the set of differences. (In this work, for quantile inference, we simply use the percentile from the current measurement history although it is possible to use confidence upper bound, easily calculated from the order statistic via binomial means [19, 20], as a conservative estimate if necessary.) As a simple example, if we have inferred that this quantile is less than +2, this reflects the belief that there is at least a 95% probability that the number of nodes required at the next time step will be no more than two greater than the number needed in the current time step. Thus, keeping two hot spares on hand will supply us with the desired confidence of having enough resources ready to handle incoming work without delay.

We implement the QPRED prediction methodology using a doubly linked list and a red-black tree, each holding the maximum difference in busy machines recorded over an epoch. Figure 1 depicts these data structures graphically. At



**Figure 1.** Data structures implementing QPRED. Doubly linked list holds fixed history of maximum differences. Red-black tree sorts current history of maximum differences.

the end of each epoch, the latest (youngest) difference of maxima (henceforth called simply the “difference”) is added to one end of the linked list and the oldest difference is removed. Similarly, the youngest difference is added to the red-black tree (so that differences are kept sorted) and the one that is removed is also deleted from the tree.

The history size (number of entries) is a fixed parameter supplied during configuration of a predictor. The total *time* covered by the history is the product of the number of entries in the history size and the epoch length.

To compute a prediction of the  $q^{\text{th}}$  quantile of the differences with a history size of  $H$ , the methodology extracts the entry corresponding to the  $(1 - q) \cdot H$  largest value in the red-black tree. For example, if  $H = 100$ , and  $q = 0.95$ , then the  $5^{\text{th}}$  largest value in the red-black tree is the prediction of the 0.95 quantile of the current history of differences.

This implementation is simple and speed efficient. Each addition and deletion to the linked list is constant time, the addition and deletion of a value to the red-black tree is  $O(\log(H))$ , and the scan for the quantile takes  $(1 - q) \cdot H$  operations (if  $q < 0.5$ , and  $(1 - q) \cdot H$  operations if  $q \geq 0.5$  since the sorted list can either be scanned from largest to smallest or *vice versa*). The implementation is also space efficient since only the current list of historical entries is needed<sup>2</sup>. Note that the original QBETS prediction methodology on which this method is based includes a change-point detector that implements history trimming in the event that conditions change suddenly. Such an enhancement is possible for QPRED at the cost of additional predictor state and complexity. As our results indicate, however, for the current state of the practice with production private clouds represented in the traces we have examined, the additional complexity associated with change-point detection appears to be unwarranted.

In addition, we wish to concern ourselves not with the fraction of *time intervals* in which there is a delay but rather the fraction of instances themselves that experience a time delay at startup. Notice that if the VM interarrival times are larger than the measurement time interval, there could be many intervals in which the difference is zero. QPRED would make correct quantile predictions for each interval, but not necessarily for each VM start.

For example, consider a hypothetical situation in which the interarrival time between VMs is five times larger than the polling interval length. In four out of every five intervals, the difference is zero because no VMs will attempt to start. The SLA provided to the cloud user, however, describes the probability that a VM will be delayed *in each interval*. Thus we need to adjust the quantile to be a factor of five *smaller* to account for the fraction of intervals that actually contain VM starts.

In order to make this adjustment, we compare the mean interarrival time between instances to the polling interval. If the latter is greater, we reduce the targeted probability of delay within a given time interval by the appropriate factor. (We could certainly make an adjustment in the other direc-

tion if the polling time is greater than the interarrival time, but we choose not to in order to keep our estimates conservative.) Returning to the example, suppose that the polling interval is 1000 seconds but that the instance interarrival time is 5000 seconds and that we want to maintain a probability of less than 0.05 of startup delay. In this case, we only want a delay once in every  $20 \cdot 5000 = 100000$  seconds, so that only a fraction of 0.01 of the intervals should see a delay. Thus we infer for the 0.99 quantile on the difference set as above.

In summary, suppose given a history of maximum occupancy numbers, a historic mean interarrival time  $I$ , polling interval  $t$ , and desired fraction  $\alpha$  of jobs delayed at startup. At the beginning of each time interval:

- Calculate the target fraction  $\beta = \min(\alpha, \frac{t}{I} \cdot \alpha)$  of time intervals experiencing a delay.
- Find a suitable upper bound  $M$  on the  $(1 - \beta)$  quantile for the differences of the maximum occupancies.
- Adjust the number of hot spares so that there are a total  $M$  machines powered on above the maximum number occupied at any point in the previous time interval.

Note that the methodology adjusts the number of powered up or down at the beginning of each time interval. If the number of hot spares is inadequate during any interval, Eucalyptus will immediately initiate the power-up of a machine, but the VM requests that arrive before the machine is operational will be delayed.

### Scheduling Methodologies

In Section 3 we compare the performance of four different scheduling methodologies. The performance of each methodology is characterized by the fraction of total power it uses, and the fraction of VMs that experience a start-up delay. The methodologies are defined as follows.

- **Power-greedy** – This scheduler results in the optimal power usage by a feasible implementation that considers VMs in the order they arrive (an  $O(n)$  algorithm) without regard for the number of VMs that will experience a start-up delay. It uses a “greedy” selection strategy that always chooses a node that is in use and has sufficient capacity over one that is “empty” when making a VM assignment decision. It also keep nodes powered off until they are needed and powers them off immediately when they become idle.
- **QPRED-greedy** – This scheduler makes greedy assignment decisions like Power-greedy, but it uses the quantile predictions to anticipate how many idle “hot spares” are needed at any moment to ensure that the probability a VM will be delayed falls below a target threshold.
- **Power-RR** – This algorithm is similar to Power-greedy in that it considers VMs in arrival order and only makes a single placement decision for each VM. However, instead

<sup>2</sup>The implementation actually maintains both the time-sorted linked list and the value-sorted red-black tree to improve speed efficiency. However, for the purposes of state exchange in the event of a fail-over, only the time-sorted list is needed – the red-black tree is reconstructed.

of attempting to keep nodes “empty” so that they can be powered down, it uses a round-robin rule to assign VMs to nodes that are powered up when each VM arrives.

- **QPRED-RR** – Like Power-RR, this scheduling algorithm chooses among powered-up nodes when a VM arrives and must be assigned to a node. However, it uses the quantile predictions to anticipate the number of idle “hot spares” need to be available to meet a target VM delay probability threshold.

Earlier versions of Eucalyptus included a version of Power-greedy that used the Ubuntu Power Nap [25] facility to allow machines to “sleep” until they were needed. This version of the Eucalyptus VM scheduler sends a message to each node instructing it to put itself to sleep whenever that machine becomes idle. When the scheduler needs to start a VM, it consults an internal record of node state and selects a node that currently has the capacity to run the VM and is also currently powered on. If no node is found, it then considers nodes that are in the process of “waking up” and chooses one that will have sufficient capacity once it is fully power on. Finally, if no “on” or “waking” nodes are located, it selects a node that is powered off, sends that node a *wake-on-lan* message [28] thereby putting it in the “waking” state, assigns the VM to the node, and waits until the node is fully powered up before starting it and any other VMs that are waiting. To keep VMs “packed” onto powered-up nodes, Power-greedy gives the nodes an arbitrary order and then always considers nodes in this order when making a placement decision. Power-RR is an alternative to Power-greedy that goes through each class of node (“on,” “waking,” and “off” in round-robin order (*i.e.* the scheduler starts with the next node in order when a new placement decision is needed).

The QPRED schedulers predict a bound on the maximum number of machines that will be required to start all VMs in a fixed time epoch such that the probability of a VM incurring a power-up delay is no greater than a fixed target probability supplied to the algorithm. QPRED-greedy uses the same greedy approach to making placement decisions as does Power-greedy, but it also attempts to power on enough hot spares (based on the quantile prediction) to control the probability that a future VM start will experience a start-up delay. Thus, compared to Power-greedy, QPRED-greedy trades additional speculative power usage for the ability to provide a statistically valid SLA. Alternatively, QPRED-RR is comparable to Power-RR except that it too uses the quantile prediction to forecast the number of additional nodes that must be powered up to meet a specific target SLA.

The difference between the greedy and round-robin versions of these schedulers is the degree to which the exploit multi-tenancy. The greedy schedulers will attempt to use a few machines as possible, thereby increasing the degree to which VMs will share nodes. As a result, they are more power-efficient than their round-robin counterparts; how-

ever, because of the greater potential sharing, VMs under a greedy schedule may experience greater I/O interference. Eucalyptus does not currently provide a way for the schedulers to access VM-specific performance information, nor do the schedulers “trust” the users to provide metrics that would allow the schedulers to determine VM affinity. Instead, administrators can use the round-robin scheduling disciplines to minimize inter-VM interference. These round-robin versions are necessarily less power efficient than their greedy counterparts.

## Data Sets

We present data from four separate commercial private clouds implemented using Eucalyptus<sup>3</sup>. The commercial entities operating these private clouds have allowed us to monitor their respective installations over an extended period and have agreed to have these results data made publicly available in an anonymized form. All four clouds support the commercial activity of their operators (they are not operated for, *e.g.*, evaluation or investigative purposes). Each data set contains start and stop time stamps for VMs that were launched and terminated in a Eucalyptus cloud, the name or IP address of the node to which each VM was assigned by the scheduler (either the Eucalyptus Greedy or Round-Robin scheduler) that was in use when the dataset was generated, and the number of CPU cores that each VM was given.

The first data set (DS1) is taken from an organization with several large-scale software development efforts. While the private cloud is used for some company-wide service hosting, its primary use is to support software testing and development. DS1 captures private cloud VM activity that combines software development with service hosting, with an emphasis on development.

The second data set (DS2) is taken from an IT organization that “sells” time on a re-charge basis to other organizational units in its umbrella company. The accounting charges translate to operating budget for the following fiscal year, making the economic incentives similar to those driving a public cloud. Thus the usage of this cloud is not known (*i.e.*, the cloud does not have a specific purpose other than to host the workloads of its paying customers). The function of the umbrella company, however, makes it likely that much of the activity is generated by software development.

The third data set (DS3) is taken from a private cloud used to allow business partners to integrate their respective software products with the products made by the company operating the cloud. It also supports user and customer trials of the company’s software products. Finally, these partners often use the cloud for demonstration or sales purposes. Thus the workload is a mixture of software development with on-demand hosting activities.

<sup>3</sup>We will make the data sets we have used in this study, in anonymized form, at the web address given in [29] once [30] appears in print or if this submission is accepted for presentation at SOCC.

Finally, the fourth data set (DS4) comes from a cloud used exclusively for software development and testing at a software start-up company that uses an Agile [23] engineering process. The Agile process makes heavy use of testing during development so the workload in this data set represents a mixture of user-controlled VMs and VMs that are launched and terminated by an automatic testing system.

Table 1 provides summary descriptions of the cloud deployments from which we have gathered these data sets.

Each data set measures the workload from a single Eucalyptus Availability Zone. Our goal in using these data sets is to measure workload across a spectrum of commercial activity. However, note that each zone is relatively small and is operated by a small IT staff. Also, we do not have measurements of the size of the user pool accessing each cloud. Thus the zone scale does not necessarily reflect the size of the user community that generated the workload captured in each data set.

All four data sets span several months of continuous usage. During the monitoring periods, each of the hosting organizations upgraded their respective Eucalyptus clouds, in one case multiple times.

### Experimental Methodology

The results presented in the Section 3 are generated from a faster-than-real time simulator that is able to “replay” each data set described in Table 1 using different scheduling methodologies. The simulator is able to replay each dataset as it was gathered (*i.e.* using the scheduling information in the data set). It also implements the different scheduling policies (both based on QPRED and otherwise), reports machine statistics such as node utilization, core utilization, power consumption, and the fraction of VMs that were required to wait for a node to power-up (also termed the “miss fraction” since the VM “missed” having an available node powered-on to start it). Because the results are simulated using datasets from machines that were not instrumented for power usage when the datasets were gathered, power consumption is reported as a fraction of the total power that was used. That is, the simulator records the ratio of time each node is powered-up using power-saving scheduler to the time that all node were powered up. We explain this method of measuring power consumption in greater detail in the next section.

## 3. Results

We begin by detailing the tradeoff between overall power usage and the probability that a VM’s start will be delayed while a machine is powering up to host it. We compare Power-greedy, QPRED-greedy, Power-RR, and QPRED-RR in terms of both power usage and VM delay fraction. In what follows, we will use the term “miss fraction” interchangeably with the term “VM delay fraction” because from the perspective of a scheduler (particularly the predictive sched-

ulers) a VM that experiences a VM start-up delay is a “miss” with respect to finding a machine powered on and ready to accept the VM.

### Power Usage and VM Delay Fraction

Table 2 compares the performance of these four schedulers using the data set described in Subsection 2. Each boldfaced number in the table denotes the fraction of maximal power that the scheduling methodology used. That is, the simulations compute the the total number of node-seconds used for each data set as a hardware-independent measure of the power that would have been consumed in the absence of power-aware scheduling. The boldfaced numbers are the fraction of this maximal usage number for each data set (thus, *e.g.*, a lower fraction represents greater power savings).

The italicized numbers show the fraction of VMs that incurred a delay as a result of having to wait for a machine to reach a fully powered on state. For this experiment, we used a power-on interval of 600 seconds, taken to represent a typical amount of time it takes a server-class machine to start up, and a target delay probability of 0.05 for the SLA given to the user. Thus, when the predictor is accurate, the total fraction of VMs experiencing a delay for either QPRED method should be less than or equal to 0.05.

As an example, consider the results for DS3 in Table 2. The boldfaced number in the second column (0.22) indicates the fraction of total power used with all of the machines powered up for the duration of the trace that Power-greedy scheduler would have used to complete the workload. Put another way, Power-greedy (which is the most power-efficient of the schedulers we examine) would use 22% of the power that was used by the system when it executed the workload originally, with all of its machines on and fully powered. At the same time, Power-greedy for DS3 generates a miss fraction of 0.27 (italicized number in column 2), indicating that 27% of the VMs would experience a start-up delay. In sum, Power-greedy for DS3 would use just 22% of the power that was used for the work load, but 27% of the user requests would incur a delay while waiting for a machine to power up.

In the third column for DS3, we show the power fraction (boldfaced) and VM delay fraction (italics) for QPRED-greedy. These data indicate that QPRED-greedy would have used 37% of the total power used originally, but only 2% of the VMs would have been delayed waiting for a machine to power up. Thus, QPRED-greedy would have used 15% more power (relative to the maximum) than Power-greedy while maintaining the 0.05 target probability (since 0.02 is less than 0.05) specified in the SLA.

Finally, in the fourth and fifth columns of the row for DS3, we show the results for Power-RR and QPRED-RR respectively. Power-RR uses 41% of the original power, but 60% of the VMs experience a start-up delay. Meanwhile, QPRED-RR uses 59% of the original power (18% more

Data Set	Nodes	Cores/Node	Time Period	Description
DS1	13	24	Aug. 2012 to Oct. 2012	Large company with 50,000 to 100,000 employees
DS2	7	12	Aug. 2012 to Apr. 2013	Medium sized company with 2,000 to 5,000 employees
DS3	7	8	Aug. 2012 to May 2013	Small company with 50 to 100 employees
DS4	12	8	May 2013 to Sep. 2013	Start-up company with 5 to 10 employees

**Table 1.** Summary of Private Cloud Dataset Characteristics

Data Set	Power-greedy	QPRED-greedy	Power-RR	QPRED-RR
DS1	<b>0.56</b> <i>0.08</i>	<b>0.62</b> <i>0.02</i>	<b>0.92</b> <i>0.06</i>	<b>0.87</b> <i>0.00</i>
DS2	<b>0.33</b> <i>0.45</i>	<b>0.51</b> <i>0.05</i>	<b>0.76</b> <i>0.20</i>	<b>0.83</b> <i>0.01</i>
DS3	<b>0.22</b> <i>0.27</i>	<b>0.37</b> <i>0.02</i>	<b>0.41</b> <i>0.60</i>	<b>0.59</b> <i>0.02</i>
DS4	<b>0.35</b> <i>0.46</i>	<b>0.56</b> <i>0.04</i>	<b>0.43</b> <i>0.67</i>	<b>0.67</b> <i>0.01</i>

**Table 2.** Comparison of Scheduler performance. Boldfaced numbers are fraction of maximal power. Italicized numbers are fraction of VM’s delayed. QPRED target delay fraction is *0.05*

that Power-RR relatively speaking) while respecting the 0.05 miss fraction specified in the SLA ( $0.02 < 0.05$ ).

### Predictor Efficacy

The data in Table 2 used an SLA with a target VM delay probability of 0.05 for all experiments. As described in Section 2, the predictor used in both QPRED schedulers attempts to estimate the quantile of the distribution of the maximum number of nodes occupied during each time epoch corresponding to this target probability. Thus for a target probability of 0.05, the quantile estimator attempts to choose the number of nodes that correspond to the 0.95 quantile of the distribution of the maximum node occupancies across epochs. From the table, the predictor is correct for a 0.05 VM delay probability since all of the observed delay fractions are less than or equal to 0.05.

In Table 3 we show the VM delay fraction for QPRED-greedy that results from parameterizing the predictor with different target quantiles corresponding to different SLAs. In each experiment, we use an epoch interval of 1000 seconds and a power-up delay of 600 seconds (the same as for the results in Table 2). In each column except the first we show the fraction of original power usage in boldfaced type and the miss fraction in italics for the target quantile  $q$  shown in the first row. We underline the entries where the observed VM delay fraction is greater than the target quantile (*i.e.* an SLA violation) indicating that the predictor failed to achieve a conservative bound.

As expected, the fraction of maximal power increases as the target quantile decreases. That is, smaller the fraction of VMs that can miss according to the SLA, the more power the cloud must use to ensure that the SLA is met. For example,

for DS4, an SLA of 0.01 uses 65% of the original power. If an SLA of 0.25 is chosen, the true miss fraction rises to 0.11 but the cloud uses only 45% of the original power. Thus the price of a 0.01 SLA guarantee versus a 0.25 SLA guarantee is 20% in terms of power usage for DS4.

For all target quantiles except  $q = 0.01$  for DS2 the predictor’s bound on VM delay fraction holds, although it appears quite conservative in many cases (*e.g.* the VM delay fraction for DS1 is 0.07 for a target of 0.25). The predictor misses outright with a miss fraction of 0.03 for DS2 with a target quantile of  $q = 0.01$ , however. This failure illustrates the effect that autocorrelation in the interarrival time series can have on our methodology. Specifically, the DS2 data set contains periods of time when few VM starts occur and also short intervals when a large number of VMs arrive. The prediction methodology does not take this “burstiness” into account. Thus, when a burst of VMs occurs in the DS2 data set, QPRED-greedy does not have enough machines ready and idle to absorb the burst such that at most only 1% of the VMs will experience a delay per the terms of the SLA.

Note that this effect increases with the length of the machine spin-up delay. QPRED-greedy initiates a machine spin-up whenever there is no powered-up machine capable of starting the VM *and* there is no machine in the process of powering-up that could start it once the power-up sequence has completed. Thus, there are two types of “misses” with respect to VM start-up delay.

- A **Full Miss** occurs when a VM start request arrives and only machines that are fully powered off are available to host it.



Data Set	$q=0.01$	$q=0.05$	$q=0.10$	$q=0.15$	$q=0.20$	$q=0.25$
DS1	<b>0.68</b> <i>0.01</i>	<b>0.62</b> <i>0.02</i>	<b>0.58</b> <i>0.03</i>	<b>0.57</b> <i>0.04</i>	<b>0.55</b> <i>0.07</i>	<b>0.55</b> <i>0.07</i>
DS2	<b>0.53</b> <u><i>0.03</i></u>	<b>0.51</b> <i>0.05</i>	<b>0.52</b> <i>0.05</i>	<b>0.47</b> <i>0.08</i>	<b>0.46</b> <i>0.11</i>	<b>0.45</b> <i>0.15</i>
DS3	<b>0.45</b> <i>0.01</i>	<b>0.37</b> <i>0.02</i>	<b>0.36</b> <i>0.03</i>	<b>0.35</b> <i>0.03</i>	<b>0.34</b> <i>0.04</i>	<b>0.32</b> <i>0.05</i>
DS4	<b>0.65</b> <i>0.01</i>	<b>0.56</b> <i>0.04</i>	<b>0.52</b> <i>0.04</i>	<b>0.49</b> <i>0.06</i>	<b>0.47</b> <i>0.07</i>	<b>0.45</b> <i>0.11</i>

**Table 3.** Power usage fraction in boldface and VM delay fraction in italics for different target quantiles using QPRED-greedy. Underlined fraction exceeds the target SLA.

- A **Partial Miss** occurs when a VM start request arrives, there are no powered-up machines available to start the VM, but a machine that is in its powering-up phase has sufficient capacity to host the VM once it has completed its start-up.

Notice that the longer the machine spin-up delay, the more probable a VM will experience a partial miss. Indeed, in the extreme, if the start-up delay were equal to the length of time each data set covers, all VMs except the first to start on each node would experience a partial miss and the VM delay fraction would necessarily be 1.0 no matter what the target quantile.

To investigate this effect further, Table 4 shows the VM delay fractions using QPRED-greedy with a target quantile of 0.05 and different machine spin-up delays. For this experiment, we show the results for Power-greedy and vary the spin-up delay from 60 seconds to 1800 seconds while keeping the length of the time epoch and the history length both at 1000 (as they were in the previous experiments) and the target VM delay probability set at 0.05. The miss fractions are shown in italics and fractions that exceed the target SLA probability of 0.05 are underlined.

From this information, it is clear that miss fraction increases with spin-up delay; however, the rate of increase is slow. For example, the miss fraction goes from 0.02 when the spin-up delay is 60 seconds to 0.04 when it is 1800 seconds. Moreover, Power-greedy is able to meet the requirements of an SLA specifying 0.05 as the maximum fraction of VMs to experience a spin-up delay in all but two cases. For DS2, when the spin-up delay is either 900 seconds or 1800 seconds, Power-greedy generates a miss fraction of 0.06, which exceeds the target of 0.05 set for the SLA. It is unclear whether this small violation is a result of time-series effects or a mismatch between the time epoch of 1000 seconds and the spin-up delay.

Table 5 compares the average delay (measured in seconds) experienced by those VMs that are delayed for Power-greedy with those for QPRED-greedy. Note that these values are not estimates of the expected value of the delay over the entire data set. That is, they measure the average delay for only those VMs that were delayed by each method. For example, for DS4, Power-greedy resulted in 46% of the VMs experiencing a delay (*c.f.* Table 2) and among those the average delay is 578 seconds. By contrast, QPRED-greedy de-

Data Set	Power-greedy delay	QPRED-greedy delay
DS1	541	563
DS2	557	417
DS3	556	516
DS4	578	571

**Table 5.** Average VM start-up delay generated by Power-greedy and QPRED-greedy for target probability 0.05. The units are seconds

layed only 3% of the VMs, and the average delay among those was 571 seconds.

These results seem somewhat counterintuitive. Our expectation was that because QPRED-greedy was keeping additional machines powered up, most of the misses would be “full” misses and thus the average for QPRED-greedy would be higher. Instead, a careful examination of the trace data indicate that when QPRED-greedy experiences a full miss, (*i.e.*, it does not have enough hot spares provisioned) it is more likely to be followed by a number of partial misses in rapid succession.

In [3] and [9], the authors report that the savings benefits gained by powering down machines when they are not needed can be overshadowed by the use of additional “peak” power during the spin up phase. When a machine is powered on, it may use more power (*e.g.*, to accelerate disks to operational speed) relative to its steady-state or idle-state usage. The ratio of peak usage during start-up to steady-state usage varies by machine manufacturer and model as well as by configuration (*e.g.*, the number and type of disks attached). The authors of both works note that the additional usage during power-up can be as much as 60% more than steady-state.

In Table 6 we show the power savings for QPRED-greedy over a range of peak-to-steady state ratios. For these experiments, we use a target quantile of 0.05 and a spin-up delay of 600 seconds. A ratio of 1.0 shows the case when there is no difference between power consumption during spin-up and steady-state. Thus column 2 of Table 6 (marked as 1.0) corresponds to column 3 (marked as  $q = 0.05$ ) of Table 3 discussed previously. Even if the spin-up cost were five times steady-state, the additional power usage is no more than 1% with a spin-up time of 600 seconds.

In Table 7 we repeat the same experiment with a spin-up delay of 1800 seconds to determine the effect that a significantly longer delay might have. These data indicate that with a start-up delay as long as 1800 seconds, the additional over-

Data Set	60s	90s	120s	300s	600s	900s	1200s	1800s
DS1	0.01	0.01	0.01	0.02	0.02	0.02	0.02	0.02
DS2	0.02	0.02	0.02	0.04	0.05	<u>0.06</u>	0.05	<u>0.06</u>
DS3	0.02	0.02	0.02	0.02	0.02	0.02	0.02	0.02
DS4	0.03	0.03	0.03	0.03	0.04	0.04	0.04	0.04

**Table 4.** VM miss fraction only QPRED-greedy and target quantile of 0.05 as a function of increasing VM spin-up delays (first row, units are seconds). Underlined fractions exceed the target SLA probability.

Data Set	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0
DS1	0.62	0.62	0.62	0.62	0.62	0.62	0.62	0.62	0.62
DS2	0.51	0.51	0.51	0.51	0.51	0.51	0.51	0.51	0.52
DS3	0.37	0.37	0.37	0.38	0.38	0.38	0.38	0.38	0.38
DS4	0.55	0.56	0.56	0.56	0.56	0.56	0.56	0.56	0.56

**Table 6.** VM power usage fraction for QPRED-greedy, target quantile of 0.05, and spin-up delay of 600 seconds as a function of increasing peak power ratio during spin-up.

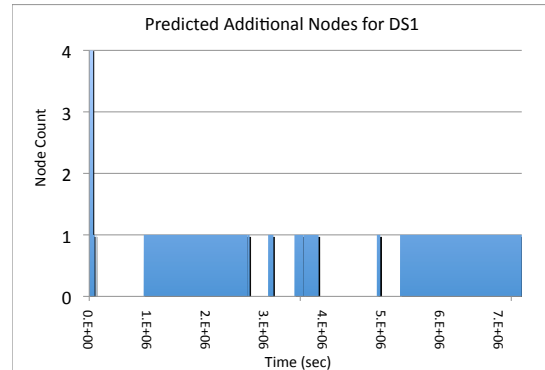
Data Set	1.0	1.5	2.0	2.5	3.0	3.5	4.0	4.5	5.0
DS1	0.62	0.62	0.62	0.62	0.63	0.63	0.63	0.63	0.64
DS2	0.51	0.51	0.51	0.51	0.51	0.52	0.52	0.52	0.52
DS3	0.39	0.39	0.39	0.40	0.40	0.40	0.41	0.41	0.41
DS4	0.57	0.57	0.58	0.58	0.58	0.59	0.59	0.60	0.60

**Table 7.** VM power usage fraction for QPRED-greedy, target quantile of 0.05, and spin-up delay of 1800 seconds as a function of increasing peak power ratio during spin-up.

all usage caused by peak consumption during spin-up will be no more than 3% compared to the theoretical case (*i.e.*, a ratio of 1.0) in which there is no difference between spin-up and steady-state usage.

These results do not contradict the previously published work in [3] and [9]. Rather, they indicate that with a target SLA quantile of 0.05 and production cloud workloads, QPRED-greedy does not generate enough spin-up events for a large peak-to-steady state ratio to have a substantial effect on power savings. For example, QPRED-greedy generates 252 power-up events for DS1 over the course of the roughly 8 million seconds that the trace covers. With 13 nodes, if each power-up event corresponds to a 600 second delay, the total time spent in power-up is 0.1% of the total  $13 * 8000000 = 104000000$  node-seconds.


Finally, we show two different examples of the predictions made by QPRED-greedy over the course of each trace. Figure 2 shows the time series of predictions for DS1 using QPRED-greedy and a target quantile of 0.05 for the SLA. In this experiment, the spin-up delay was 600 seconds the length of the time epoch was 1000 seconds, and the predictor used a history length of 100 (as in the previous experiments). The y-axis depicts the number of hot spares that the predictor determined were required to meet the target VM delay fraction of 0.05 specified in the SLA and the x-axis shows the number of seconds since the beginning of the trace. This



**Figure 2.** Time series of predictions made by QPRED-greedy for DS1 with a target quantile of 0.05 for the SLA. y-axis are predicted hot spares and x-axis is time from the start of the trace in seconds.

figure indicates that for DS1 a simple strategy of maintaining a single hot spare would likely have met the 0.05 SLA but would have been even more conservative (*i.e.*, would have used more power) than QPRED-greedy: Note the periods of time during which QPRED-greedy chose not to power up a hot spare while achieving an overall miss fraction (from Table 2) of 0.02.

Figure 3 shows the time series of predictions made by QPRED-greedy for DS4 under the same experimental conditions. Again, The y-axis depicts the number of hot spares



**Figure 3.** Time series of predictions made by QPRED-greedy for DS4 with a target quantile of 0.05 for the SLA. y-axis are predicted hot spares and x-axis is time from the start of the trace in seconds.

that the predictor determined were required to meet the target VM delay fraction of 0.05 specified in the SLA and the x-axis shows the number of seconds since the beginning of the trace. In this example, a single hot spare over the lifetime of the trace would *not* have met the 0.05 target probability for the SLA: Note the substantial time periods during which more than one additional machine was needed to achieve the desired miss fraction.

#### 4. Discussion

From the results we are able to make two assertions. First, for real enterprise private cloud workloads, substantial power savings are possible using simple scheduling algorithms and the ability to power machines up and down dynamically. Second, using an efficient non-parametric method for predicting quantiles on the distribution of future machine usage, it is possible to achieve power savings while at the same time meeting the requirements of a probabilistic SLA on the additional delay user’s will experience when machines must be powered up.

From an engineering perspective, these results are encouraging. The prediction methodology maintains relatively little state (a sorted list of 1000 numbers in these experiments) and the time necessary to make a prediction is constant. Most Linux systems support some form of power-down or hibernation through the ACPI interface [15] and an automatic power-up capability via “wake-on-lan.” Thus implementing this methodology should be straight-forward and non-invasive for most private cloud platforms.

Also, the quantile provides a single tuning parameter that can be used to control the tradeoff between power usage and induced delay. Cloud administrators can set and revise

the target quantile to control observed miss fraction dynamically. Even in the cases when the value of the quantile does not match the observed miss fraction (say at a particular point in time due to autocorrelation) a lower quantile will correspond to a lower miss fraction and vice versa. We believe that this new type of control parameter will prove useful to private cloud administrators.

From a user-experience perspective, the addition of power savings via the QPRED schedulers is also relatively non-invasive. For example, a 0.02 miss fraction achieved by QPRED-greedy for DS3 (the fraction obtained for the 0.05 target quantile) corresponds to an average 7 delayed VMs per month across all users. Thus, because the methodology is not complex, uses ubiquitous Linux functionality, is private-cloud neutral, and has little impact on user experience while at the same time saving power, we believe these results should influence future cloud scheduler designs.

Finally, our results support the trend toward higher core densities per machine as one that can reduce power consumption in data centers substantially. In particular, the effects of any autocorrelation in the VM interarrival time series is mitigated by higher core densities by a greedy scheduler. When a number of VM requests arrive in a burst, more cores allow a powered-on machine to absorb the burst more often, thereby lowering the miss fraction. Indeed, for the only example of where QPRED-greedy failed to achieve the SLA terms in Table 3 (the 0.01 target probability for DS2) increasing the core count to 32 in this cloud would have resulted in a successful miss fraction of 0.01.

#### 5. Related Work

Both because clouds aggregate usage and also because they commoditize compute and storage capabilities, they are especially well-suited for the implementation of automatic power optimization. In [9], [10], and [8] the authors discuss the efficacy of various “sleep state” formulations’ for data-center-hosted processors. This work, like ours, involves the application of Markov/time-series methods to the problem of power management. In these papers the Markovian approach appears in service of an  $M/M/k$  queuing model. Our work, which focuses on clouds rather than data centers, also uses a Markov-based approach to workload but in a much more direct way: we use a fast algorithm to make sample-based estimates of confidence bounds on transition probabilities (indeed with some further simplifying hypotheses). While their use of time series is in some sense more sophisticated than ours, we have found that for our purposes nonparametric and model-agnostic methods yield better results. For example, their work models arrivals as a Poisson process (which has the advantage of making the solution of an  $M/M/k$  queuing system computationally feasible) and job lengths as exponentially distributed. In the production workload data sets available to our study, however, arrivals are not well modeled as a Poisson process so the use of an

$M/M/k$  queuing model is not warranted; nor are job lengths typically exponentially distributed [30, 31]. Additionally, our work examines the SLA that a cloud must provide with respect to VM start-up delay; their work (perhaps because it focuses on workloads in data centers, where start-up delay is not typically subject to an SLA) does not consider start-up delay guarantees. Judging from the results reported for the reactive scheduler (which appears to be similar to Power-greedy described herein) our methodology is at least competitive if not more efficient in some cases.

In [4] the authors formulate the problem in terms of multi-dimensional optimization and then explore a set of heuristics for improving power usage. Our work differs in several respects. First, their approach uses measurements of VM activity to determine migration policies. They then explore the efficacy of their techniques using both a simulated workload and simulated cloud environments. Our efforts focus on predictive enhancements that augment cloud schedulers used in production today. Further, our investigation assumes no knowledge of VM behavior, making the results applicable to a wide range of private cloud settings.

In [18] the authors use a strategy for virtualizing the CPU power states supported by most server-class CPUs to implement power scaling on a per VM basis. Our work shares the goal of power optimization under the constraints of user-facing SLAs with this approach and also the use of measurement data gathered from engineered systems. However it complements this approach by relying strictly on the cloud scheduler and not the VMs themselves to implement power management cloud wide. Our work is substantially less complex in terms of apparatus necessary for implementation but significantly less rich in the spectrum of SLAs that can be supported.

The question of the effect that multiple VMs have on CPU power consumption is investigated in [24]. Again, our work is best construed as a complement to this work. Rather than looking at power consumption in terms of the level of draw that a CPU or machine requires, we view machines as either “on” drawing some variable amount of power or “off” drawing no power. It is certainly possible to optimize the power utilization from the “on” machines that our methodology requires.

The work in [32] investigates the power efficiency of the same scheduling strategy that has been implemented by Eucalyptus as the Power-greedy scheduler. In addition, they explore the effects of additional “hot spares” (called a “pool”) in this work. As described, our work prioritizes user experience in the form of an SLA and uses an on-line predictive methodology to predict how many hot spares are needed. Because of the similarity in base-line schedulers between OpenNebula [7] (the test platform for this work) and Eucalyptus, however, our approach should be directly applicable to their test environment.

In [12] and [5] the authors use a variety of statistical techniques including time series analysis and clustering to predict VM workloads from a virtualized data center that is intended to be used as a private cloud. Their study uses CPU utilization data gathered from each VM across a history of time intervals to predict aggregate load in the next time interval. Our work is similar in that we too discretize time into epochs and use time series of measurements to make a prediction for each epoch immediately before it begins. However, our methods uses measurements of overall cloud load rather than an aggregation of VM CPU utilization. Further, our approach predicts quantiles as a way of implementing user-facing SLAs whereas their method generates point-value predictions.

Finally, in [2] the authors propose to use economic characteristics of popular private cloud infrastructures (Eucalyptus and others). While private clouds may reach a scale for which our approach is inappropriate, it is sufficient to support production cloud computing as it is practiced today in the enterprise.

## 6. Conclusion and Future Work

This work shows that it is possible to use a simple, computationally efficient prediction methodology based on quantile estimation to improve cloud power usage while also implementing an SLA governing machine virtual machine start-up delay. The methodology predicts a conservative bound on the number of machines that must be powered on at any moment to ensure that the probability of having to power up a machine (*i.e.*, a miss) is at or below the target set by the cloud administrator. We illustrate the efficacy of the approach using VM activity traces gathered from four enterprise private clouds that were in production use at the time of their instrumentation. Our results show that QPRED (which is non-parametric and both computationally and space efficient) generates substantial power savings under settable probabilistic constraints on the tradeoff between power savings and degraded user experience.

We plan to carry this work forward in two ways. First, we will investigate ways in which the bounds (particularly on miss fraction for higher quantiles) can be improved. The current methodology, while correct, is quite conservative. It should be possible to improve the power savings while maintaining the accuracy of the predictions. Second, we developed this methodology to work in conjunction with Eucalyptus and other similar private cloud infrastructures that make a single placement decision for each VM when the VM starts. The current release of Eucalyptus includes a VM migration facility that allows the administrator to “evacuate” a target node so that it may be taken out of service without a disruption experienced by the users. We plan to investigate how quantile predictions can be used to optimize this process.

## References

- [1] Apache Cloudstack. <http://cloudstack.apache.org>, 2013.
- [2] R. Bahsoon. A framework for dynamic self-optimization of power and dependability requirements in green cloud architectures. In *Software Architecture*, pages 510–514. Springer, 2010.
- [3] L. A. Barroso and U. Hölzle. The case for energy-proportional computing. *IEEE computer*, 40(12):33–37, 2007.
- [4] A. Beloglazov and R. Buyya. Energy efficient resource management in virtualized cloud data centers. In *Proceedings of the 2010 10th IEEE/ACM International Conference on Cluster, Cloud and Grid Computing*, pages 826–831. IEEE Computer Society, 2010.
- [5] R. Birke, L. Y. Chen, and E. Smirni. Data centers in the cloud: A large scale performance study. In *Cloud Computing (CLOUD), 2012 IEEE 5th International Conference on*, pages 336–343. IEEE, 2012.
- [6] Eucalyptus Systems Inc. <http://www.eucalyptus.com>, 2013.
- [7] J. Fontán, T. Vázquez, L. Gonzalez, R. S. Montero, and I. Llorente. Opennebula: The open source virtual machine manager for cluster computing. In *Open Source Grid and Cluster Software Conference*, volume 86, 2008.
- [8] A. Gandhi. *Dynamic Server Provisioning for Data Center Power Management*. PhD thesis, Intel, 2013.
- [9] A. Gandhi, M. Harchol-Balter, and I. Adan. Server farms with setup costs. *Performance Evaluation*, 67(11):1123–1138, 2010.
- [10] A. Gandhi, M. Harchol-Balter, R. Raghunathan, and M. A. Kozuch. Distributed, robust auto-scaling policies for power management in compute intensive server farms. In *Open Cirrus Summit (OCS), 2011 Sixth*, pages 1–5. IEEE, 2011.
- [11] K. Keahey, M. Tsugawa, A. Matsunaga, and J. Fortes. Sky computing. *Internet Computing, IEEE*, 13(5):43–51, 2009.
- [12] A. Khan, X. Yan, S. Tao, and N. Anerousis. Workload characterization and prediction in the cloud: A multiple time series approach. In *Network Operations and Management Symposium (NOMS), 2012 IEEE*, pages 1287–1294. IEEE, 2012.
- [13] O. Krieger, P. McGachey, and A. Kanevsky. Enabling a marketplace of clouds: VMware’s vcloud director. *ACM SIGOPS Operating Systems Review*, 44(4):103–114, 2010.
- [14] A. Lenk, M. Klems, J. Nimis, S. Tai, and T. Sandholm. What’s inside the cloud? an architectural map of the cloud landscape. In *Proceedings of the 2009 ICSE Workshop on Software Engineering Challenges of Cloud Computing*, pages 23–31. IEEE Computer Society, 2009.
- [15] Linux ACPI. <http://acpi.sourceforge.net/documentation/>, 2013.
- [16] D. Milošević, I. M. Llorente, and R. S. Montero. Opennebula: A cloud management tool. *Internet Computing, IEEE*, 15(2): 11–14, 2011.
- [17] J. Murty. *Programming Amazon Web Services: S3, EC2, SQS, FPS, and SimpleDB*. O’Reilly Media, Inc., 2009.
- [18] R. Nathuji and K. Schwan. Virtualpower: coordinated power management in virtualized enterprise systems. *ACM SIGOPS Operating Systems Review*, 41(6):265–278, 2007.
- [19] D. Nurmi, J. Brevik, and R. Wolski. Qbets: Queue bounds estimation from time series. In *Job Scheduling Strategies for Parallel Processing*, pages 76–101. Springer, 2008.
- [20] D. Nurmi, R. Wolski, and J. Brevik. Probabilistic advanced reservations for batch-scheduled parallel machines. In *Proceedings of the 13th ACM SIGPLAN symposium on principles and practice of parallel programming*, pages 289–290. ACM, 2008.
- [21] D. Nurmi, R. Wolski, C. Grzegorzczak, G. Obertelli, S. Soman, L. Youseff, and D. Zagorodnov. The eucalyptus open-source cloud-computing system. In *Cluster Computing and the Grid, 2009. CCGRID’09. 9th IEEE/ACM International Symposium on*, pages 124–131. IEEE, 2009.
- [22] K. Pepple. *Deploying OpenStack*. O’Reilly, 2011.
- [23] A. Software Development. [http://en.wikipedia.org/wiki/Agile\\_software\\_development](http://en.wikipedia.org/wiki/Agile_software_development), 2013.
- [24] S. Srikantaiah, A. Kansal, and F. Zhao. Energy aware consolidation for cloud computing. In *Proceedings of the 2008 conference on Power aware computing and systems*, volume 10. USENIX Association, 2008.
- [25] Ubuntu power nap. <http://manpages.ubuntu.com/manpages/lucid/man8/powernap.8.html>, 2013.
- [26] H. N. Van, F. D. Tran, and J.-M. Menaud. Performance and power management for cloud infrastructures. In *Cloud Computing (CLOUD), 2010 IEEE 3rd International Conference on*, pages 329–336. IEEE, 2010.
- [27] W. Vogels. Eventually consistent. *Communications of the ACM*, 52(1):40–44, 2009.
- [28] wake-on lan. <http://en.wikipedia.org/wiki/Wake-on-LAN>, 2013.
- [29] R. Wolski and J. Brevik. <http://www.cs.ucsb.edu/rich/workload>, 2013.
- [30] R. Wolski and J. Brevik. Using parametric models to represent private cloud workloads. *IEEE Transactions on Service Computing (to appear)*, PP(99), October 2013.
- [31] R. Wolski and J. Brevik. Using parametric models to represent private cloud workloads. Technical Report 2013-05, University of California, Santa Barbara, August 2013. [http://128.111.41.26/research/tech\\_reports/reports/2013-05.pdf](http://128.111.41.26/research/tech_reports/reports/2013-05.pdf).
- [32] A. J. Younge, G. Von Laszewski, L. Wang, S. Lopez-Alarcon, and W. Carithers. Efficient resource management for cloud computing environments. In *Green Computing Conference, 2010 International*, pages 357–364. IEEE, 2010.