

Provisioning Multi-tier Cloud Applications Using Statistical Bounds on Sojourn Time

Upendra Sharma, Prashant Shenoy, Don Towsley
Dept. of Computer Science, University of Massachusetts
Amherst, MA, USA
{upendra,shenoy,towsley}@cs.umass.edu

ABSTRACT

In this paper we present a simple and effective approach for resource provisioning to achieve a percentile bound on the end to end response time of a multi-tier application. We, at first, model the multi-tier application as an open tandem network of M/G/1-PS queues and develop a method that produces a near optimal application configuration, i.e., number of servers at each tier, to meet the percentile bound in a homogeneous server environment – using a single type of server. We then extend our solution to a K-server case and our technique demonstrates a good accuracy, independent of the variability of service-times. Our approach demonstrates a provisioning error of no more than 3% compared to a 140% worst case provisioning error obtained by techniques based on an M/M/1-FCFS queue model. In addition, we extend our approach to handle a heterogenous server environment, i.e., with multiple types of servers. We find that fewer high-capacity servers are preferable for high percentile provisioning. Finally, we extend our approach to account for the rental cost of each server-type and compute a cost efficient application configuration with savings of over 80%. We demonstrate the applicability of our approach in a real world system by employing it to provision the two tiers of the java implementation of TPC-W – a multi-tier transactional web benchmark that represents an e-commerce web application, i.e. an online bookstore.

Categories and Subject Descriptors

C.4 [Performance of Systems]: Modeling techniques

Keywords

Performance modeling, cloud computing, multi-tier applications

1. INTRODUCTION

Cloud computing platforms are becoming increasingly popular for hosting enterprise applications due to their ability to support dynamic provisioning of virtualized resources to handle workload fluctuations and also because of usage based pricing. Enterprise

applications are known to observe dynamic workload and provisioning correct capacity for these applications remains an important and challenging problem. High variability in workload is caused by a variety of reasons, such as flash crowds, short term sustained surges, or long-term fluctuations based on change in business or underlying IT infrastructure etc. Predicting these workload fluctuations or the peak workload is challenging. Erroneous predictions often lead to under-utilized systems or in some situations cause temporarily outage of an otherwise well provisioned web-site; e.g. in November 2000 Amazon.com site suffered a forty-minute outage due to overload. Consequently, rather than provisioning server capacity to handle infrequent (and hard to predict) peak workloads, an alternate approach of dynamically provisioning capacity on-the-fly in response to workload fluctuations has become popular. Dynamic provisioning is especially well suited to the cloud due to the ability of cloud platforms to provision capacity when needed and charge for usage on pay-per-use basis.

There have been numerous efforts that have addressed the issue of dynamic provisioning of server capacity to distributed applications [8, 11, 20, 19]. These efforts fall into two categories - *proactive*, where a model of the application is used to compute the capacity needed to service a particular workload at a certain performance level, or *reactive*, where additional capacity is allocated after a workload spike arrives and causes significant performance degradation.

In case of proactive approaches, application models have been derived to predict how much capacity is needed to provide a certain *mean* response time for a given workload [19, 20]. However, typical service level agreement (SLAs) for the application are specified in terms of the worst case (or peak) response times [7] (e.g. 99% of the requests should see no more than a 1-sec response time). Consequently, there is a mismatch between the provisioning models which allocate capacity for a target mean response, time and the SLA, which dictates that the capacity should be allocated based on a high percentile (peak) response time.

Second, many enterprise applications are distributed or replicated with multi-tier architecture. Typically SLAs are specified on an end-to-end basis for the entire application. The few provisioning efforts that focus on allocating capacity for the tail of the work translate it to per-tier SLA metrics [20]; provisioning for per-tier SLA can result in large errors in provisioning the capacity if the tier response time estimates are incorrect.

Third, most provisioning techniques to-date are cost oblivious – they can determine *how much* server capacity to allocate but do not consider the *cost* of allocating the server capacity. In cloud platform, different server configurations are available at different prices. The capacity does not scale linearly across configurations and nor does the price. Since multiple combinations of servers can

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. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICAC'12, September 18–20, 2012, San Jose, California, USA.
Copyright 2012 ACM 978-1-4503-1520-3/12/09 ...\$15.00.

provision a certain capacity C for an application, a cloud specific provisioning scheme must take the cloud costs into account when making provisioning decisions.

In this work we present a new model driven provisioning approach targeted for cloud platforms. Our approach focuses on i.) allocating capacity based on peak (high percentile) of the workload, ii) takes a holistic view of the entire multi-tier application by considering bounds on end to end response times while making provisioning decisions and iii) takes cloud server configs and pricing models when determining the most cost effective config to provision a certain amount of capacity.

1.1 Research Contributions

The contributions of this paper are the following:

Cost aware provisioning subject to a percentile response time SLA. We present an algorithm for resource provisioning for a multi-tier cloud application, subject to an SLA expressed in terms of high percentile of end to end response time, that minimizes the total cost of compute resources required by the application. Our formulation models the application as an open tandem queue network of $M/G/1$ -PS queues.

Service time and response-time approximations. We present an approximation of the response time distribution of the $M/G/1$ processor sharing queue based on the distribution of conditional expected response times given the service times and show it to be accurate for our purposes. In addition, we present a new service time characterization based on a mixture of shifted exponential distributions.

Cost-efficient configuration with heterogeneous servers subject to percentile SLA. We extend the above approach to account for the presence of multiple types of servers with different costs and computational capabilities. This is achieved by formulating an integer optimization problem with the constraint that per-tier capacity should be at least as much as that computed by the queueing theoretic model.

Prototype implementation and experimentation. We have implemented our analytical model in MATLAB and tested it using a multi-tier application, i.e. java implementation of TPC-W, over a private cloud. For comparison, we also implemented a baseline case using $M/M/K$ -FCFS queues. Our experimental results show that our approach is able to provision the application to meet the SLA specified on 99 percentile of end-to-end response time with less than 3% provisioning error, while the baseline techniques provisioned with an error as large as 140%. In the case of heterogeneous provisioning, our approach shows, as high as, 81% savings in server cost as compared to that of the corresponding optimal homogeneous configuration. In case of private cloud experiments we found that heterogeneous approach showed around 11% cost saving (using Amazon EC2 pricing) over homogenous configurations.

2. BACKGROUND AND PROBLEM FORMULATION

In this section, we present the system model and a high level problem description. We describe the SLA performance metric, and thereafter formulate the provisioning problem that we address in this work.

Multi-tier Application: Modern large scale web applications are developed as multiple tiers for reasons pertaining to scalability. A multi-tier architecture offers flexibility for development as well as deployment of applications. Each application tier, typically, provides a specific functionality and the various tiers form a processing pipeline. In a typical multi-tier application various tiers participate

in the processing of an incoming request; each of the participating tier receives partially processed requests from the previous tier and feeds these requests into the next tier after local processing (see Figure 1). The tiers are replicated to scale according to the processing demand; a load balancer is used to distribute the load over all replicas of such a tier. Figure 1 depicts a two-tier application where both tiers are replicated. This is a commonly employed architecture by e-commerce web applications where, both, web-server and database tiers are clustered to scale up according to increase in the incoming workload.

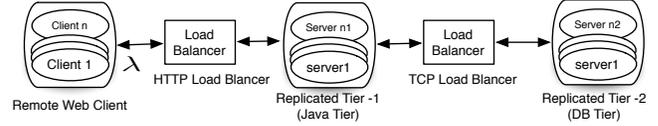


Figure 1: Topological configuration of a typical replicated two-tier web application

We assume that each tier is placed on a dedicated server and that replicating a tier essentially means replicating the server. Each clustered tier is also assumed to employ a protocol-session aware load balancer responsible for distributing requests to replicas in that tier. It is assumed that the each tier’s capacity (number of servers), can be varied dynamically without disturbing the application’s normal functioning, and that each tier can be independently provisioned for capacity.

Cloud Platforms: Cloud computing has emerged as a new IT delivery model. Infrastructure as a Service (IaaS) cloud-model is being seriously evaluated by enterprises to deploy their web applications that support dynamic capacity resizing. In this model, an organization/client can rent remote compute and storage resources to host networked applications and resources can be dynamically added or removed on an as-needed basis. We consider a cloud computing platform that allows compute servers to run hosted applications. We assume that the platform offers N heterogeneous server configurations for rent, each with a different rental-cost and configuration.

We assume that the cloud platform has an infinite pool of servers and that servers can be provisioned by invoking server-instance creation APIs; servers may be requested and terminated at any time and billing is based on the amount of time for which each server is used (e.g., based on the number of hours for which each server is used). We also assume that the cloud platform employs virtualization—each physical server is assumed to run a hypervisor that controls the allocation of physical resources on the machine and offers performance isolation to each of its virtual servers.

2.1 Problem Formulation

Let N and M denote the number of tiers and server-types respectively. Let tier j be jointly served by $\sum_{i=1}^M n_{ij}$ servers, where n_{ij} denotes the number of servers of type i present at tier j . Let $\bar{n}_j = [n_{1j}, n_{2j}, \dots, n_{Mj}]$ be a vector representing the server configuration of tier j and $\bar{p} = [p_1, p_2, \dots, p_M]$, where p_k denotes the cost of a server of type k . Let T be the end-to-end response time of requests to the multi-tier application and F_T be its CDF, i.e. $F_T(t) = P(T \leq t)$. Then for a given percentile bound θ , and response-time threshold T_D , the cost minimization problem becomes:

$$\text{minimize } \sum_{j=1}^N \sum_{i=1}^M n_{ij} p_i, \quad (1)$$

subject to the constraint

$$F_T(T_D) \geq \theta. \quad (2)$$

It should be noted that F_T , also depends on n_{ij} , since n_{ij} specifies the application configuration that determines the end-to-end response time of the application. In the next section we present a model of a multi-tier application which enables us to capture the effect of n_{ij} on F_T .

3. APPLICATION MODEL

In this section we model the multi-tier application as a network of queues. Our first model of multi-tier application is a chain of tiers where each tier is modeled as single M/G/1-PS queue (see Figure 2). Each tier carries out a specific function, for instance, a web-application server or a database server etc. In this work we assume single customer class.

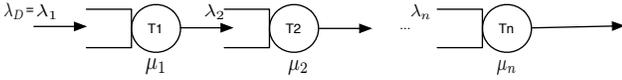


Figure 2: Multi-tier application model

Let A_i denote the i^{th} tier of the application, λ_i the average arrival rate of incoming requests at the i^{th} tier, and μ_i the average service rate $\forall i = 1 \dots N$. We define the total response time of a request as the time between when it enters the first tier and the time when it leaves the last tier. Note that different λ_i for each tier handles the case where one tier issues multiple requests to the lower tier.

Let T_j be a random variable representing the response time for tier j , then the end-to-end response time of a request is

$$T = \sum_{j=1}^N T_j. \quad (3)$$

Let $f_T(t)$ be the probability density function (PDF) of the response time T and $L_T(s) = \mathcal{L}(f_T(t))$ be the Laplace transform of the PDF of response time T then

$$L_T(s) = \prod_{j=1}^N L_{T_j}(s), \quad (4)$$

where $L_{T_j}(s)$ is the Laplace transform of the PDF of T_j . Thus the PDF of end-to-end response time, $f_T(t)$, can be computed by taking the Laplace inverse of (4)

$$f_T(t) = \mathcal{L}^{-1} \left(\prod_{j=1}^N L_{T_j}(s) \right). \quad (5)$$

To solve (5) we require the PDF of the random variable T_j . Unfortunately there are no exact formulas for response time distributions of an M/G/1-PS queue. We, therefore, present an approximation for the same in the next section.

4. ESTIMATING END-TO-END RESPONSE TIMES

In this section we describe our approach to estimate the PDF of end-to-end response time of a chain of M/G/1-PS queues. In order to do that we estimate the PDF of response time of a single M/G/1-PS queue and then leverage (5) to compute the end to end response time.

Section 4.1 describes our method of approximating the response time distribution of a M/G/1-PS queue. The result depends of the definition of the PDF of service-time distribution of the queue and we describe a mechanism to approximate the same for any real-life system in section 4.2. Section 4.3 provides a closed form equation of the end-to-end response time of the chain of queues.

4.1 Approximate Response Time Distribution

The exact form of the response time distribution for the M/G/1-PS is not generally known [23]. Thus we approximate it with the expected conditional response time distribution as described below. Let T denote the job response time, and X its service time; then the expected conditional response time, conditioned on the service time being x is

$$\tau = E[T|X = x] = \frac{x}{1 - \rho}, \quad (6)$$

where $\rho = \lambda/\mu$ is the average load.

We approximate T by τ . Since τ is a function of X ,

$$\begin{aligned} F_\tau(t) &= P[\tau \leq t] = P\left[\frac{X}{1-\rho} \leq t\right] = P[X \leq t(1-\rho)], \\ F_T(t) &\approx F_\tau(t) = F_X(t(1-\rho)), \end{aligned} \quad (7)$$

It has been observed in real-life systems that job service time distributions exhibit heavy tailed behavior [6]. Heavy tailed distributions have very high variance; high variance in service time distribution of jobs makes it a dominant factor in determining the behavior of response time distribution. Approximation proposed in (7) captures the variability of service time and will be particularly useful in such situations. We discuss the impact of variability of service time in section 7 and demonstrate that our approach shows significant improvement.

4.2 Approximate Service Time Distribution

In real systems, like computer clusters and web servers, there is a strong evidence that job service times are highly variable [6]. Some heavy tailed distributions do not have a closed-form Laplace transforms, e.g., the Pareto distribution, while those possessing convenient Laplace transforms might lead to an intractable complex function after undergoing an N^{th} order convolution in (4). We, thus, need a distribution function, which can closely approximate a service time distribution observed by a real world application and leads to an easily invertible Laplace transform even after undergoing higher order convolutions. In this section we describe such a distribution function and also present an algorithm to approximate the service time distribution from the service-time histogram; service time histograms can be easily collected from the server either through logs or through off-line profiling.

We express the service time distribution as a mixture of K shifted exponentials, as shown in (8). The motivation behind this is two fold: i.) the web application workload is a mix of different job types [12, 5]. Capturing the service time distribution as sum of shifted exponentials, essentially, means that job-size of each job-type is exponentially distributed but each job-type has a different mean job-size. ii.) The formulation leads to a Laplace transform that is easy to invert.

Formally, we want to fit a mixture of shifted exponentials,

$$f_X(x) = \sum_{k=1}^K \alpha_k \mathbf{1}\{x \geq t_k\} \mu_k e^{-\mu_k(x-t_k)}, \quad x \geq 0 \quad (8)$$

to data x_1, x_2, \dots, x_n , where $\mathbf{1}\{P\}$ is one if predicate P is true and zero otherwise. This involves inferring the number of shifted

exponentials, K , the shifts of each exponential, $\{t_k\}$, the mix proportion of the shifted exponential, $\{\alpha_k\}$, and their average rates $\{\mu_k\}$ from the data. Let us begin by assuming that K and t_1, \dots, t_K are already known. In other words we want to find the best fit for $\{\mu_k\}$ and $\{\alpha_k\}$; we perform maximum likelihood estimation using the expectation-maximization algorithm (EM).

4.2.1 EM algorithm for estimating mixture parameters

Suppose we know which shifted exponential distribution each observation x_i belongs to, in other words suppose we have $y_i \in \{1, \dots, K\}$ available to us where $y_i \in \{1 \dots K\}$ represents the particular shifted exponential distribution. Then the parameter values that maximize the log likelihood function can be computed as:

$$\alpha_k = \frac{1}{n} \sum_{i=1}^n \mathbf{1}\{y_i = k\} / n, \quad k = 1, \dots, K \quad (9)$$

$$1/\mu_k = \frac{\sum_{i=1}^n \mathbf{1}\{y_i = k\} x_i}{\sum_{i=1}^n \mathbf{1}\{y_i = k\}}, \quad k = 1, \dots, K \quad (10)$$

EM is an iterative algorithm that infers y_i as needed. Suppose μ_k^j and α_k^j are the estimates at the end of the j -th iteration. The next iteration consists of an expectation step followed by a maximization step as given below.

Expectation. Let $y_{i,k}$ denote the probability (expectation) that sample x_i belongs to the k -th shifted exponential. It is given as

$$\begin{aligned} y_{i,k} &= P[Y_i = k | X = x_i] \\ &= \frac{\alpha_k^j \mathbf{1}\{x_i \geq t_k^j\} \mu_k e^{-\mu_k^j (x_i - t_k^j)}}{\sum_{l=1}^K (\alpha_l^j \mathbf{1}\{x_i \geq t_l^j\} \mu_l e^{-\mu_l^j (x_i - t_l^j)}} \end{aligned} \quad (11)$$

$\forall i = i, \dots, n$ and $k = 1, \dots, K$. Note that $y_{i,k} = 0$ when $x_i < t_k^j$. *Maximization.* Having computed $y_{i,k}$, we now update our estimates of α_k and μ_k . This is done by using modified versions of (9) and (10).

$$\alpha_k^{j+1} = \frac{1}{n} \sum_{i=1}^n y_{i,k}, \quad k = 1, \dots, K \quad (12)$$

$$1/\mu_k^{j+1} = \frac{\sum_{i=1}^n y_{i,k} x_i}{\sum_{i=1}^n y_{i,k}}, \quad k = 1, \dots, K \quad (13)$$

This is referred to as the maximization step because the above estimates maximize the likelihood given the current values of $\{y_{i,k}\}$.

These steps are repeated until the parameters converge; $\{\alpha_k^0\}$ and $\{\mu_k^0\}$ are the initial values, which can be computed as mentioned in the section below.

4.2.2 Algorithm for approximating service-time distribution

We use an iterative approach to determine the best number of exponentials K , and then determine t_k, μ_k^0 , and α_k^0 , to initialize the EM algorithm, (11), (12) and (13).

The basic idea underlying the algorithm, as outlined in as mentioned in [12], is to iteratively run a k -means clustering algorithm for every value of $k = 1 \dots K_{max}$ and compute the following three metrics¹: coefficient of variation² of intra-cluster distance (C_{intra}), coefficient of variation of inter-cluster distance (C_{inter}), and ratio of intra-cluster to inter-cluster coefficient of variation (β_{cv}). The

¹the metrics are computed as mentioned in [12]

²Coefficient of variation or variation coefficient is defined as a ratio of the standard deviation to the mean, i.e. $C_v = \sigma/\mu$;

value of β_{cv} drops as number of clusters increase and will be minimum (i.e. zero) when number of clusters is equal to the total number of points. We find that K , where the rate of decrease of β_{cv} falls below a threshold (or the slope goes above a negative threshold value).

Having computed K , and the cluster centers e_k , we compute initial estimates of the mean service rate $\{\mu_k^0\}$ and mixture fraction (α_k^0) as follows:

$$\mu_k^0 = \frac{1}{e_k - t_k}, \quad \alpha_k^0 = \frac{\text{number of points in cluster}}{\text{total number of points}}. \quad (14)$$

We set the shifts to be equidistant from from two neighboring cluster centers, i.e., $t_i = (\mu_{i-1} + \mu_i)/(2\mu_{i-1}\mu_i), \forall i = 2 \dots K$. However, $t_1 = 0$, i.e., the shift for the first exponential is zero (details of the algorithm can be found in [17]).

4.3 Approximate Application Response Time Distribution

The PDF of the end to end response time of N -tier application is obtained using (8) and (7) in (5) as

$$f_\tau(t) = \mathcal{L}^{-1} \left(\prod_{j=1}^N \sum_{k=1}^{K_j} \frac{\alpha_{jk} \mu_{jk}' e^{-st_j'}}{(s + \mu_{jk}')}, \right), \quad (15)$$

where for each tier $j = 1, \dots, N$, service times are modeled as mixtures of K_j shifted exponentials and their density functions are expressed using (8); we rewrite the result for the j^{th} tier for the sake of completeness:

$$f_{X_j}(x) = \sum_{k=1}^{K_j} \alpha_{jk} \mathbf{1}\{x \geq t_{jk}\} \mu_{jk} e^{-\mu_{jk}(x-t_{jk})}. \quad (16)$$

After inverting (15), the final expression of $f_\tau(t)$ takes the following form:

$$\begin{aligned} f_\tau(t) &= \sum_{i_1=1}^{K_1} \dots \sum_{i_N=1}^{K_N} \left(\mathbf{1}\{t \geq t'\} \prod_{j=1}^N \alpha_{ji_j} \mu_{ji_j}' \times \right. \\ &\quad \left. \sum_{l=1}^N r_l e^{-\mu_{li}'(t-t')} \right), \end{aligned} \quad (17)$$

where $\mu_{ji_j}' = \mu_{ji_j}(1 - \rho_j)$, $t' = \sum_{j=1}^N t_{ji_j}/(1 - \rho_j)$, and $r_l = 1 / \left(\prod_{k \neq l} (\mu_{ki}' - \mu_{li}') \right)$.

Note that α_{ji_j} and μ_{ji_j} are the parameters of the k^{th} shifted exponential of the j^{th} -tier (as shown in (16)); ρ_j is the average utilization of the j^{th} tier, and r_j is the j^{th} residue, where $j = 1, \dots, N$.

Note that the expression in (15) does not involve higher order poles³ because none of the rates μ_{li} is ever equals any of the μ_{ji_j} . This becomes especially helpful in inverting the Laplace transform as absence of higher order terms in denominator leads to a simple computation of partial fractions.

The final expression of $f_\tau(t)$ in (17) is, essentially, a product of sums of the shifted exponentials, which is easily readable in (15). This means that the $f_\tau(t)$ will be expressed, in total, by $\prod_{j=1}^N K_j$ terms; for example let $K_j = a, \forall j = 1 \dots N$, then $f_\tau(t)$ will be expressed as a sum of a^N terms. It is easy to see that number of terms grow exponentially with number of tiers. Fortunately, real

³If for some l, j , $\mu_{li} = \mu_{ji_j}$, we slightly perturb the starting μ_{li}^0 for tier- l by adding a small random number and re-run the EM algorithm for that tier- l

life systems do not have more than three or at most four tiers and thus $f_\tau(t)$ is easily computable.

5. FINDING NEAR OPTIMAL HOMOGENEOUS CONFIGURATION

In this section we present a solution to the the resource optimization problem, as expressed by (1) and (2), but with only one type of server, $M = 1$ (homogeneous setting).

We substitute the approximate response time of an M/G/1-PS queue, i.e. $f_\tau(t)$ as shown in (17), in (2) to obtain:

$$F_\tau(T_D) = \sum_{i_1=1}^{K_1} \dots \sum_{i_N=1}^{K_N} \left(\mathbf{1}\{T_D \geq t'\} \prod_{j=1}^N \alpha_{j i_j} \mu_{j i_j}' \sum_{l=1}^N \frac{r_l (1 - e^{-\mu_{l i_l}' (T_D - t')})}{\mu_{l i_l}'} \right) \geq \theta, \quad (18)$$

where $\mu_{j k_j}'$ and r_j are the same as in (17) while $t' = \sum_{j=1}^N t_{j, i_j} / (1 - \rho_j)$.

Thus the problem of minimizing (1) reduces to the problem of maximizing ρ_j ($\forall j = 1, \dots, N$) such that $F_\tau(T_D) \geq \theta$, where $F_\tau(T_D)$ is given by (18). As this is an N -dimensional non-linear maximization problem, it is not easy to solve. However, the problem complexity is significantly reduced by assuming same utilization at each tier⁴, i.e.,

$$\rho_1 = \rho_2 = \dots = \rho_N = \rho.$$

It should be noted that it is desirable to have a balanced utilization at each tier in real-life systems. In practice, administrators often use a rule of thumb to bound the max utilization of servers of all tiers to avoid performance problems and outages [15].

Consequently, (18) reduces to an inequality in a single variable, namely ρ .

$$F_\tau(T_D) = \sum_{i_1=1}^{K_1} \dots \sum_{i_N=1}^{K_N} \left(\mathbf{1}\{T_D \geq t'\} \prod_{j=1}^N \alpha_{j i_j} \mu_{j i_j} \sum_{l=1}^N \frac{r_l' (1 - e^{-\mu_{l i_l}' (T_D - t')})}{\mu_{l i_l}'} \right) \geq \theta, \quad (19)$$

where, $t' = \sum_{j=1}^N t_{j, i_j} / (1 - \rho)$, and $r_l' = 1 / \prod_{k \neq n} (\mu_{k i_k} - \mu_{l i_l})$. We solve for the maximum value of ρ , say ρ^* , by numerically solving (19) as an equality.

5.1 Computing the Application Configuration

In practice, large scale applications have each of their tier replicated for scalability as depicted in Figure 3. The idea is to be able to handle increasing number of requests while conforming to the SLA. In an ideal situation an application-tier's ability to process the number of requests increases linearly with number of its replicas, which means that if an application or application-tier with a single replica had a service rate of μ then K replica version of application-tier will have a request rate of $K\mu$. We have assumed a linear scaling in this work but that is not a limitation and any kind of scaling function can be used in the technique to obtain the number of replicas at each application-tier. We have used replicas and servers interchangeably because we have assumed dedicated hosting model.

We use ρ^* to compute the number of servers at each tier, i.e. n_{i_j} . In the homogenous setup $i = 1$ and thus we solve for n_j , $j = 1 \dots N$. Let λ_j and μ_j be arrival and service rates respectively, at

⁴The constraint reduces the solution search space and thus the final solution is not guaranteed to be an optimal solution as it could result into a slightly over-provisioned system.

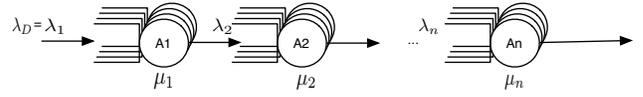


Figure 3: Multi-tier application model

tier j then $n_j = \lceil \lambda_j / (\zeta \rho^* \mu_j) \rceil$, where ζ is the scale factor, which can be chosen heuristically and

$$\mu_j = \sum_{i=1}^{K_j} \alpha_{j i} \frac{(1 + \mu_{j i} t_{j i}) e^{-\mu_{j i} t_{j i}}}{\mu_{j i}}. \quad (20)$$

The pseudo code of the algorithm for finding the application configuration in homogenous setup is outlined in [17].

6. COST EFFICIENT HETEROGENEOUS CONFIGURATION

We extend the solution approach described in Section 5 to be able to generate a cost efficient configuration in a heterogeneous setting.

The basic idea underlying our approach is to greedily search for a low cost configuration which has a high utilization. At a high level the algorithm is iterative involving the following three steps at each iteration: 1.) creating a single hybrid-server from a given hybrid-configuration for each tier, 2.) solve the homogeneous configuration problem for the hybrid-server, 3.) translate the solution for hybrid-server into a heterogeneous configuration, and the iterations are used to search for new hybrid-configuration with lower cost and higher utilization. Figure 4 shows the block diagrammatic representation of the cost effective heterogeneous configuration algorithm.

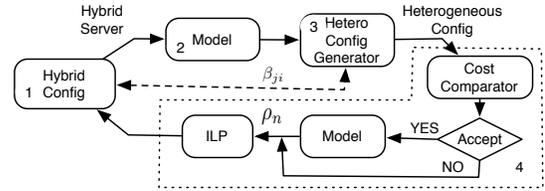


Figure 4: Functional block diagram of heterogeneous configuration algorithm

Hybrid server: In order to reuse our methodology for finding the near optimal number of servers in homogeneous setting, it is imperative that we approximate each hybrid configuration at each tier by a single server; we call it a hybrid-server. We construct the service time distribution of the hybrid-server for each tier as a proportional mixture of the service time distributions of the servers involved in the heterogeneous configuration. Let $\bar{n} = \{n_i\}$ denote the hybrid-configuration where n_i , $i = 1, \dots, M$, is the number of servers of type i . Then the hybrid-server's service-time distribution function for tier- j is expressed as

$$f_j' = \sum_{m=1}^M \beta_{j m} f_{j m}, \quad (21)$$

where $f_{j m}$ is the service-time probability density function (PDF) of the m^{th} -server-type at j^{th} -tier and f_j' is the PDF of the hybrid-server for tier- j ; $\beta_{j m}$ is the mixing proportion of the component server m for tier- j and is computed using the formula

$$\beta_{j m} = \frac{n_m \mu_{j m}}{\sum_{j=1}^M n_m \mu_{j m}}. \quad (22)$$

We explain our procedure of creating a hybrid-server with the following example: suppose we have two servers, say s_1 and s_2 , with corresponding average service rates at tier- j as $\mu_{j1} = 50$ and $\mu_{j2} = 100$, respectively. We construct a single hybrid-server, say s_n , by proportionally mixing the component shifted-exponentials of each s_1 and s_2 . Let the configuration be one-server of each type, i.e. $\bar{n} = [1, 1]$; then the mixing proportions using (22) is $\beta_{j1} = 1/3$ and $\beta_{j2} = 2/3$, and the final service-time distribution of the hybrid-server for the j^{th} tier is $f_j' = \left(\frac{f_{j1}}{3} + \frac{2f_{j2}}{3}\right)$.

Heterogeneous configuration: Once we obtain the optimal configuration for a given hybrid-server, and given workload and percentile, we translate this solution configuration to the corresponding heterogeneous server configuration; this is done by reversing the steps of creating the hybrid-server. Let us assume that the servers are indexed in increasing order of their average service rate; i.e. $\mu_1 \leq \mu_2 \leq \dots \leq \mu_M$; let n_j' be the number of hybrid-servers at tier- j , then the number of servers of type- i for tier- j is $n_{ji} = \beta_{ji}n_j' / (\mu_i / \mu_1)$.

Searching for a new hybrid-configuration: The cost of the new heterogeneous configuration, computed in the step above, is evaluated using the prices of the servers. If the cost is less than that of the current solution configuration, then this new configuration is accepted else it is dropped. The new configuration is again fed to the model, and its utilization ρ^* is evaluated for the desired arrival rate λ_D . We then try to search for a new hybrid-configuration which has higher utilization but lower cost than the current-configuration; the new utilization $\rho_n = (\rho_{max} + \rho_l)/2$, where ρ_{max} is maximum utilization of the hybrid-server and $\rho_l = \rho^*$. The new hybrid configuration is searched for using the following ILP solved for each tier:

$$\text{minimize } \sum_{i=1}^M n_{ji} p_i, \quad (23)$$

subject to the constraint

$$\sum_{i=1}^M n_{ji} \mu_i > \lambda_D / \rho_n. \quad (24)$$

Note that if the currently suggested configuration is not accepted we continue to search for higher ρ^* . The algorithm stops when $\rho_n - \rho_{max}$ is less than a pre-decided threshold; the pseudo code is outlined in [17].

7. EXPERIMENTAL EVALUATION

In this section we demonstrate the efficacy of our approach. We have implemented our analytical method using MATLAB[®]. For solving the ILP, we have used lpsolve version 5.5.2.0 and have used mxlpsolve MATLAB Interface version 5.5.0.6 for calling lpsolve from within the MATLAB environment.

We begin by showing the effectiveness of the service-time approximation algorithm on lognormal⁵ distribution with different coefficient of variations (C_v). Thereafter we evaluate the goodness of the approximation of the response-time distribution for a 1-tier and a 2-tier system by comparing the response times computed using (17) with those obtained using a multi-tier application-simulator described below. Finally we do a case-study of provisioning of a two-tier application for a SLA specified as a threshold on the 99th percentile of response time. We evaluate the effectiveness of our approach by computing the 99 percentile of response

⁵PDF of a log normal distribution is expressed as $f(x, \mu, \sigma) = \frac{1}{x\sigma\sqrt{2\pi}} e^{-((\ln(x)-\mu)/(\sqrt{2}\sigma))^2}$, where mean is $e^{\mu+\sigma^2/2}$

times obtained using a two-tier application-simulator configured according to the capacity decisions provided by our approach; note that the simulator depicts an ideal version of a multi-tier application which we analytically model as a chain of M/G/1-PS queues. We also evaluate the effectiveness of our approach, using a metric called provisioning error (described in Section 7.4), by comparing against the two other baseline approaches, which model the multi-tier application as an open tandem network of M/M/K-FCFS queues.

7.1 Multi-tier Application Simulator

We implemented a simulator for the PS queue in MATLAB[®]. It takes as input an array of request arrival instants and size of each request (in terms of service time) and outputs the request departure instants. We used this queue simulator to simulate a multi-tiered application by feeding the output of first queue to the input of the next queue.

To simulate an application with replicated tiers, we have implemented a loadbalancer, as shown in Figure 1, which takes the incoming requests from the previous tier and distributes it to the next tier according to a specific load distribution policy. It also does the necessary book-keeping to track each request across various tiers for computing the end-to-end response time. We have implemented a random loadbalancing policy, i.e. loadbalancer distributes the requests at random but ensures that each server gets the same load, i.e. ρ^* as computed in section 5. We have assumed an ideal loadbalancer, which means that it introduces no queuing and processing delay. Note that this is not a limitation of our approach, as our approach can easily account for loadbalancer by considering it as another tier and its capacity can also be computed, which is often needed in a real setup.

7.2 Service Time Approximation

We have implemented the EM algorithm (in MATLAB) for finding the parameters of mixture of shifted exponentials, namely α_i, μ_i in (8), using the E and M steps mentioned in Section 4.2.1. The shifts and initial values of parameters are estimated using the algorithm outlined in Section 4.2.2. We use MATLAB's implementation of KMeans and have kept $K_{max} = 20$ in all our experiments, which means that we search for the number of shifted exponents from 3 till 20. We evaluate the accuracy of CDF approximation using relative percentage error defined as $\epsilon(x) = (F_{approx}(x) - F_{sim}(x)) / (F_{sim}(x))$, where $F_{approx}(x)$ and $F_{sim}(x)$ are the values of approximate and actual CDFs, respectively, evaluated at x .

To evaluate the effectiveness of our approach in approximating highly variable distribution, we approximated the PDF a log-normal distribution with same mean rate of 20 but with a coefficient of variation of 100 ($C_v = 100$); it was expressed using 10 shifted exponentials. Figure 5a shows the CCDF of the actual distribution in red while our approximated distribution is shown in blue.

The CCDF in Figure 5 highlight the approximation of the tail of the distribution by plotting the $1 - F(x)$ in a log log scale. We observed that the approximation shows a relatively high error at low percentiles (as high as 21%) but displays low errors at the tail, with errors less than 1% at 95 percentile. This is because, that at low percentiles the number of exponentials available to approximate the distribution are less but as we approach the tail of the distribution a large number of exponentials contribute towards the approximation of the PDF and thus we observe much greater accuracy.

Another aspect of our algorithm is K , i.e. the number of exponentials required to approximate a distribution. We conducted a large number of experiments on various data sets with C_v ranging from 1 to 100. In our experiments we found that K its average

values starts at 14 for $C_v = 1$ and slightly decreases to an average value of 10 for $C_v = 100$. It should be noted that we are testing our approximation scheme for a smooth distribution function, but the scheme has been designed keeping a web application in vision, which has only a limited number of request types at each tier and our approach is tuned to estimate this number as K .

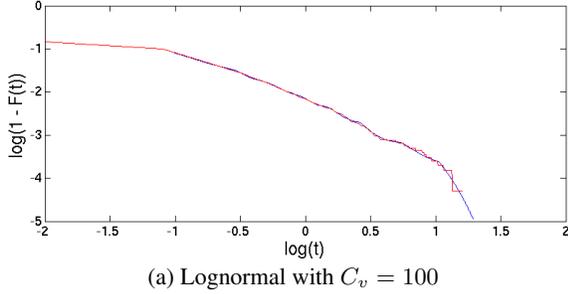


Figure 5: Figure shows the log log plot of 20,000 data points sampled from lognormal distribution with $C_v = 100$; the simulated CDF is shown in red and approximate in blue.

In summary, the service time approximation approach offers very low errors, i.e. less than 1%, in estimating the tail of a distribution.

7.3 Response Time Approximation

In this section we describe the effectiveness of our approach of approximating the end to end response time of an application modeled as a chain of M/G/1-PS queues.

We evaluated the goodness of the response-time approximation we compare the response time computed by our approach with that obtained from the simulator described above. We show the results for for a 2-tier setup by plotting the response time CDFs for our approximation and simulation. We have sampled service times from a lognormal distribution with a $C_v = 10$.

We generate the workload which has exponential inter-arrival times with $\lambda = 25$ and service times sampled from a lognormal distribution with $\mu = 50$ at each tier. The simulation results are considered exact since the simulation model is an exact representation of the queuing network under study.

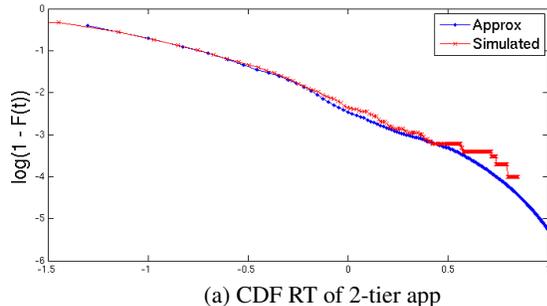


Figure 6: Figure shows the CDF plot of actual response time distribution in red and approximated using our approach in blue for a heavy-tailed service-time distribution with $\mu = 50$ and $c_v = 10$

The approximated response time, using our approach, exhibits high accuracy, as can be seen from Figure 6. The jagged tail of simulation result is because of less number of data points.

7.4 Provisioning in a Homogenous Setup

In this section we evaluate the effectiveness of our approach, outlined in section-5, in finding the homogenous configuration for a two-tier application, where each tier is replicated using same type of servers.

For a given SLA, expressed as a cutoff threshold T_D on the 99th percentile of end to end response time, we fix a service time distribution and for different arrival rates compute the number of servers required at each tier of the application. We, then, run the replicated application simulator with these number of servers and obtain the end to end response time distribution for the provisioned application. To evaluate the goodness of provisioning decisions made, we define a metric called *provisioning error*, which essentially calculates the error in the 99th percentile response time observed from the simulator, i.e. T_{scheme} , and T_D . Formally, $\epsilon_{scheme} = (T_{scheme} - T_D) * 100/T_D$. To do a comparative evaluation of our technique, we have implemented two baseline provisioning algorithms based on M/M/1-FCFS queues, namely *per-tier-exp* and *end-to-end-exp*. The schemes are described below:

per-tier-exp (pte) : In this scheme we assume the knowledge of average proportion of time spent by a request at each tier. In other words, let T be the total time spent by a request in the system and T_i be the time spent at tier i ; then, *pte* assumes the knowledge of $E[\delta_i]$, where $\delta_i = T_i/T$. We model each tier as an M/M/K-FCFS queue and again approximate multiple servers by a single server, thus each tier can be approximated by an M/M/1-FCFS queue. For this system the response time is exponentially distributed with parameter $\mu(1 - \rho)$. Finally, as in Section 5, for each tier j , we solve for ρ^* with $T_D = \delta_j T$ and compute $n_j = \lceil \lambda_j / (\rho^* \mu_j) \rceil$

end-to-end-exp (ete): We developed this scheme completely along the lines of our scheme, however assuming an M/M/K-FCFS queue based model instead of an M/G/K-PS queue based model. The corresponding version of (19) is:

$$F_T(t') = \sum_{j=1}^N r_j (1 - e^{-\mu'_j t'}) \geq \theta, \quad (25)$$

where $t' = T_D$, $r_j = 1 / \prod_{k \neq j}^N (\mu'_k - \mu'_j)$ and $\mu'_j = \mu_j(1 - \rho)$. The provisioning algorithm for homogenous setting is outlined in [17].

We ran the experiment with $T_D = 0.4s$, $\mu = 50$ and $C_v = 3$. We increased the workload from $\lambda = 40$ rps to 240 rps and for each λ we computed application capacity using each of the three algorithms. For *pte* we used $\delta_1 = \delta_2 = 0.5$. The results are shown in Table 1.

λ	% ϵ_{our}	% ϵ_{ete}	% ϵ_{pte}	Config _{our}	Config _{ete}	Config _{pte}
40	-3.63	16	15.2	[3;3]	[2;2]	[2;2]
80	-6.17	48.1	27.9	[5;5]	[3;3]	[4;3]
120	-0.235	94.3	38.5	[8;7]	[4;4]	[5;5]
160	-2.25	91.6	49.9	[9;9]	[5;5]	[6;6]
200	-2.17	140	40.7	[12;12]	[6;6]	[8;8]
240	2.57	91.3	53.7	[15;15]	[8;8]	[9;9]

Table 1: Homogeneous configuration suggested by the three schemes and their provisioning errors. Note that, unlike the positive error, negative value of ϵ is not an SLA violation.

A Positive value of ϵ means that some or all of the tiers of the application were provisioned with fewer servers than required (we call it under-provisioning); however, a negative value means the opposite (we call it over-provisioning). Thus a positive ϵ is an SLA violation, while a negative ϵ is not. However, a negative ϵ does suggests a possibility of finding a more cost efficient solution.

Note that our scheme reports a worst case provisioning error of 2.57% as opposed to the worst case under-provisioning of 140% by *ete* and 53.7% by *pte*.

In summary: for a single server type scenario (i.e. homogeneous setup), application provisioned by our scheme reports worst case provisioning error of 2.57%, while the baseline approaches shows as high as 140% provisioning error

7.5 Cost Efficient Server Configuration in a Multiple Server-type Environment

Here we demonstrate the effectiveness of our heterogeneous provisioning algorithm in finding a cost-efficient solution when multiple types of servers are available. We have kept the time threshold $T_D = 0.4$ -sec and varied the desired load from $\lambda = 40$ -rps to $\lambda = 240$ -rps. We have considered four types of servers, namely small (S), medium (M), large (L), and extra-large (XL), with their corresponding average service rates being 50, 100, 150 and 200 rps, respectively. The coefficient of variation of service times for requests at each of the tiers is $C_v = 9$.

ServerType	Small	Medium	Large	XLarge
Price	0.02	0.04	0.06	0.08

(a) server prices

λ	$\% \epsilon_{homo}$	$\% \epsilon_{hetero}$	Config _{homo}	Config _{hetero}	%Saving
40	1.63	-40.9	[9;9]	[0 1 0 0;0 1 0 0]	77.78
80	1.01	-35.7	[17;15]	[0 0 0 1;0 0 1 0]	78.13
120	1.16	-22.9	[26;23]	[0 0 2 0;0 1 1 0]	77.55
160	1.06	-23.5	[34;30]	[0 0 1 1;0 0 2 0]	79.69
200	1.09	-21.5	[43;47]	[0 0 3 0;0 1 2 0]	81.11
240	1.04	-9.82	[51;45]	[0 0 2 1;0 0 3 0]	80.21

(b) 99-percentile provisioning and cost benefit

Table 2: Heterogeneous configuration suggested by the three schemes and provisioning error of each scheme. Note that a negative ϵ only means over-provisioning and is not an SLA violation

We assume linear pricing as depicted in Table 2a. The results of provisioning algorithms in homogenous and heterogenous settings are shown in Table 2b. We call the computed capacity configurations in the homogenous and heterogeneous settings as $Config_{homo}$ and $Config_{hetero}$, respectively. Only the “small” server-types were used in $Config_{homo}$, while all the available server types were used to obtain $Config_{hetero}$. As in previous evaluations, we again test the computed configuration using the multi-tier application simulator.

Each configuration is $N \times M$ dimensional matrix depicting the number of servers of each type; each row j depicts the configuration of the j^{th} tier, while each column tells the number of servers for each type: for e.g. $Config_{homo} = [9; 9]$ means 9-small servers at both the tiers, while $Config_{hetero} = [0 1 0 0; 0 1 0 0]$ means 0-small, 1-Medium, 0-large and 0-x-large server at both the tiers. The “%Cost Saving” is computed as a percentage of cost of homogeneous configuration, i.e. $\frac{Cost(Config_{homo}) - Cost(Config_{hetero})}{Cost(Config_{homo})}$.

We make following important observations: 1) the percentage provisioning error for the heterogeneous scheme is as low as -41%, which means that not-only is this configuration cost-efficient but it also provides low average response-times (because negative provisioning error means the system is probably over-provisioned). The small positive error in the case of homogeneous configurations is because of approximation used in section 5.1 and can be easily corrected by setting $\zeta < 1$, i.e. a sublinear scaling. 2) it is better to use larger servers that fit the same cost and average service-rate; in

other words its better to use a small number of large servers instead of a large number of small servers.

In summary, it is better to use a small number of large servers instead of a large number of small servers for high percentile provisioning ii) Cost efficient heterogenous algorithm offers server configurations with cost savings as high as 81% and also offer a configurations with lower average response-times.

8. EVALUATION ON PRIVATE CLOUD

In this section we describe an experimental investigation for provisioning for a percentile SLA in a private cloud setup. Our goal is to evaluate our provisioning algorithm under situations which are typical to multi-tier web applications deployed in a datacenter or private/public cloud environment.

8.1 Private Cloud Setup

In this section we provide the necessary details of our experimental testbed, i.e private cloud, and necessary steps before we can perform server provisioning.

Web Application: We used TPC-W for our experiments. TPC-W is a multi-tier transactional web benchmark that represents an e-commerce web application – an online bookstore – comprising of a web server tier and a database tier. It simulates the activities of a retail store website using 14 different type of pages for web interactions; each of these pages are created dynamically by the web server using differing amounts of data stored in the database tables. TPC-W benchmark defines three different mixes of web interactions, namely browsing, shopping and ordering, each varying the ratio of inventory browsing related web pages and ordering related web pages. It applies the workload mixes via remote browser emulator (RBE).

We used the Java implementation of TPC-W [3]. The web application has following two-tiers: i.) Web server tier based on Apache Tomcat servlet container 5.5.26 ii.) database tier based on MySQL 5.0.77. We deployed each of the tiers on separate dedicated servers. We performed round robin load balancing between replicas of web server tier using a dedicated loadbalancer server on HAProxy [9] on a server as a dedicated load balancer. We used round robin load balancing at the database tier by setting up a master-slave replication configuration of MySQL servers; we instrumented TPC-W to use the replication aware MySQL JDBC connector version 3.1.12.

Private Cloud: We constructed private cloud using OpenNebula [14] on Xen/linux-based cluster consisting of two types of servers: 8-core 2GHz AMD Opteron 2350 servers and 4-core 2.4 GHz Intel Xeon X3220 systems. All machines run Xen 3.3 and Linux 2.6.18 (64bit kernel). Our platform is assumed to support *small* and *large* servers, comprising 1 and 4 cores, respectively. These are constructed by deploying a Xen VM on the above mentioned servers and dedicating the corresponding number of cores to the VM (by pinning the VM’s VCPUs to the cores)

Profiling servers for web server tier: For profiling the servers for the first tier, i.e. web-server tier, we instrumented Tomcat such that it reports per-request service times, along with the other default stats. We profile each server type (e.g. *small* and *large*) by provisioning an instance of that server-type and deploying the first tier of TPC-W on it. We then connect it to an already installed TPC-W database installed on a *large* server type instance. We then issue the browsing workload using the TPC-W clients (i.e. RBEs) for a duration of 35 mins and collect the service times from the tomcat server logs.

Profiling servers for database server tier: Profiling the servers for the second tier of TPC-W (i.e. the database tier) was in two steps: firstly, we collect the 35-min query logs from MySQL server,

executing the TPC-W workload; then for each server type we slowly replay each of the SQL query and record their execution time as service times.

8.2 Percentile Based Capacity Provisioning on Private Cloud

Given λ_D and T_D , we outline the high level steps required to compute application capacities for both homogenous and heterogeneous setup. In both the cases we assume to require an SLA where 99th percentile of the end-to-end response time must be less than 0.5 seconds. We follow the following sequence of steps

Step 1: Estimating service time distributions: We use the service times collected during the offline profiling step and use the service time approximation algorithm – as outlined in [17].

Step 2: Estimating capacity in a homogenous/Heterogeneous setup: We used the single core virtual machines (i.e. *small*) in our homogenous setup. Load across multiple web-server replicas was distributed using a HAProxy based load-balancer, however, in the case of database tier, we used the master-slave setup. In this setup all the writes are sent to the master, whereas the reads are load-balanced.

We test our approach for both homogenous and heterogeneous environment. For homogenous setup, we choose *small* server type for this case and assume $T_D = 0.5\text{sec}$. To test the provisioning setup for large change in workload, we varied λ_D from 15 rps to 90 rps. For each λ_D , using our approach, we computed server capacities for each of the tier of TPC-W. We ran the setup for 35-mins and in the end we collected the end-to-end response times from the first tier (i.e. web-server tier). We ran our heterogeneous provisioning algorithm on the same setup and found that it gave a different configuration, only for $\lambda_D = 90$. Table 3a provides the details of the final configuration and also the 99th-percentile of the end-to-end response time details of the experiment. We compute the percentage provisioning error, ϵ_{our} , as mentioned in 7.4.

λ_D	99 th %	% ϵ_{our}	Config _{our}
15	0.361	-27.8	[1:1]
30	0.459	-8.2	[1:2]
45	0.488	-2.4	[1:3]
90	0.512	2.4	[2:7]
90	0.46	-8.0	[2,0:2,1]

(a) Server Provisioning

Server Type	<i>small</i>	<i>large</i>
Prices (\$)	0.085	0.34

(b) Server prices

Table 3: Homogenous and heterogeneous provisioning decisions. Note that a -ve ϵ_{our} only means that the system is over-provisioned and thus SLA will not be violated

We found that server provisioning by our approach keeps provisioning error below 3%. The positive 2.4% error at $\lambda = 90$ for homogenous setup could be because database tier does not scale linearly as the master database server gets overloaded by replicating the updates to each of the 6 slaves. We see that the server provisioning for the heterogeneous environment, is not only 11.11% cheaper than the corresponding homogenous server setup but also has a lower 99th response time.

In summary, our algorithm effectively accurately captures the service time distributions and provisions the two-tier implementation of TPC-W with the worst provisioning error of 3%. Also, we, again, find that its better to use bigger server for high percentile provisioning.

9. RELATED WORK

A number of efforts have modeled internet applications. Modeling single tier has gotten much of the attention. Doyle *et al.*

propose a queuing model for static content [8], Menasce uses a queuing model to model the web servers [11], while Abdelzaher *et al.* in [1] use classical feedback control theory to model the bottleneck tier for providing performance guarantees for web applications serving static content, while Chen *et al.* in [4] use a machine learning technique for provisioning the database tier.

Ranjan *et al.* [16] use a G/G/N queuing model to compute the number of servers necessary to maintain a target utilization level. This strategy is shown to be effective for sudden increases in request arrival rate. Other efforts have employed M/G/1 queuing models in conjunction with offline profiling to model service delay and predict performance [18] but they do not provision for response time percentile and neither do they address the problem in heterogeneous environment. The approach in [21] formulates the application tier server provisioning as a profit maximization problem and models application servers as M/G/1/PS queuing systems; the approach only considers the impact of different number of end-clients (and thus, request volumes)

Benanni *et al.* in [2] employ approximate mean-value analysis (MVA) to develop an online provisioning technique for multiple request classes. Urgaonkar *et al.* in [19] develop a queuing network model for multi-tier Internet applications having request classes with differentiated QoS. Zhang *et al.* [24] use a multi-class model to capture the dynamics of workload by employing a fixed set of 14 predefined transactions-types and leverage it to predict the performance of a multi-tier system.

There has been some work for finding the pdf of response time, for e.g. Muppala *et al.* in [13] derive the response time for a closed queuing network using pteri-nets and sojourn time distribution was calculated for large Markov chains in [10]. The approach leads to an inversion of a complex Laplace transform. Xiong *et al.* in [22] perform the provisioning of a multi-station setup for a given percentile bound. The model the system as an open tandem network of M/M/1-FCFS queues and compute the response time PDF by numerical inversion of its Laplace transform; they assume that each station is serviced by same type of servers.

In contrast to these efforts, our work automatically characterizes service time distribution as a mixture of shifted exponentials and leverages this to estimate the response time distribution. The estimated distribution is used to estimate the capacity of the system which assists in finding a near optimal solution to the provisioning problem in homogeneous environment. Further, while most of these efforts have focused on a single server type environment (i.e. homogeneous), we extend our approach for a cloud specific heterogeneous environment as well. We developed a full prototype implementation and our experiments were conducted on an actual private cloud.

10. CONCLUSION

Multi-tier architecture is a preferred architecture for enterprise web applications and high response time percentile provisioning is the more meaningful than mean response time based ones. We present an approach of optimizing server allocation for a multi-tier application to achieve a percentile bound on the end to end response time. We model the application as an open tandem network of queues and model each tier as an M/G/1-PS queue. We have developed an approximate model to compute the response time distribution and have also developed a technique to estimate the service time distribution from the service time histograms. We have developed an algorithm to compute per tier server allocation of the application and in a homogenous setup. We also have extended the homogenous setup solution to solve the server allocation problem in a heterogeneous setup. We have tested the efficacy

of our approach using a multi-tier application simulator and also compared it against two other baseline approaches developed using models based on M/M/K-FCFS queue. We have demonstrated superior performance of our approach as compared to the baseline approaches. Our experiments indicated that its better to use small number of large servers than large number of small servers. Finally we tested our approach using the multi-tier implementation of TPC-W benchmark over private cloud created using Xen over Linux.

11. ACKNOWLEDGMENTS

We would like to thank our anonymous reviewers. This research was supported in part by NSF grants CNS-1117221, CNS-0916972, CNS-0855128 and OCI-1032765. Upendra Sharma was supported by an IBM PhD fellowship.

12. REFERENCES

- [1] T. F. Abdelzaher, K. G. Shin, and N. Bhatti. Performance Guarantees for Web Server End-Systems: A Control-Theoretical Approach. *IEEE Transactions on Parallel and Distributed Systems*, 13(1):80–96, 2002.
- [2] M. N. Bennani and D. A. Menasce. Resource allocation for autonomic data centers using analytic performance models. In *ICAC '05*, pages 229–240, Washington, DC, USA, 2005. IEEE Computer Society.
- [3] H. W. Cain and R. Rajwar. An architectural evaluation of Java TPC-W. In *In Proceedings of the Seventh International Symposium on High-Performance Computer Architecture*, pages 229–240, 2001.
- [4] J. Chen, G. Soundararajan, and C. Amza. Autonomic Provisioning of Backend Databases in Dynamic Content Web Servers. In *ICAC*, pages 231–242, June 2006.
- [5] L. Cherkasova and P. Phaal. Session based admission control: a mechanism for peak load management of commercial web sites. *IEEE Transactions on Computers*, 51(6), June 2002.
- [6] M. Crovella. Performance evaluation with heavy tailed distributions. In *Proceedings of the 11th International Conference on Computer Performance Evaluation: Modelling Techniques and Tools*, TOOLS '00, pages 1–9, London, UK, 2000. Springer-Verlag.
- [7] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, and W. Vogels. Dynamo: amazon's highly available key-value store. *SIGOPS Oper. Syst. Rev.*, 41:205–220, October 2007.
- [8] R. Doyle, J. Chase, O. Asad, W. Jin, and A. Vahdat. Model-based resource provisioning in a web service utility. In *Proceedings of Fourth USENIX Symposium on Internet Technologies and Systems (USITS '03)*, Seattle, WA, March 2003.
- [9] Haproxy the reliable, high performance tcp/http load balancer. <http://haproxy.1wt.eu/>.
- [10] P. G. Harrison and W. J. Knottenbelt. Passage time distributions in large markov chains. In *SIGMETRICS '02: Proceedings of the 2002 ACM SIGMETRICS international conference on Measurement and modeling of computer systems*, pages 77–85, New York, NY, USA, 2002. ACM.
- [11] D. Menasce. Web Server Software Architectures. In *IEEE Internet Computing*, volume 7, November/December 2003.
- [12] D. A. Menascé, V. A. F. Almeida, R. Fonseca, and M. A. Mendes. A methodology for workload characterization of e-commerce sites. In *EC '99: Proceedings of the 1st ACM conference on Electronic commerce*, pages 119–128, New York, NY, USA, 1999. ACM.
- [13] J. K. Muppala, K. S. Trivedi, V. Mainkar, and V. G. Kulkarni. Numerical computation of response time distributions using stochastic reward nets. In *Annals of Operations Research*, pages 155–184, 1994.
- [14] Opennebula. <http://www.opennebula.org>.
- [15] G. Pacifici, W. Segmuller, M. Spreitzer, M. Steinder, A. Tantawi, and A. Youssef. Managing the response time for multi-tiered web applications. In *IBM, Technical Report*, January 2005.
- [16] S. Ranjan, J. Rolia, H. Fu, and E. Knightly. Qos-driven server migration for internet data centers. In *Proceedings of the Tenth International Workshop on Quality of Service (IWQoS 2002)*, May 2002.
- [17] U. Sharma, P. Shenoy, and D. F. Towsley. Provisioning Multi-tier Cloud Applications Using Statistical Bounds on Sojourn Time. Technical Report UM-CS-2012-009, Dept. of Computer Science, Univ. of Massachusetts, March 2012.
- [18] C. Stewart and K. Shen. Performance Modeling and System Management for Multi-component Online Services. In *Proc. USENIX Symp. on Networked Systems Design and Implementation (NSDI)*, May 2005.
- [19] B. Urgaonkar, G. Pacifici, P. Shenoy, M. Spreitzer, and A. Tantawi. An Analytical Model for Multi-tier Internet Services and Its Applications. In *Proc. of the ACM SIGMETRICS Conf.*, Banff, Canada, June 2005.
- [20] B. Urgaonkar, P. Shenoy, A. Chandra, P. Goyal, and T. Wood. Agile dynamic provisioning of multi-tier internet applications. *ACM Transactions on Adaptive and Autonomous Systems (TAAS)*, Vol. 3, No. 1, pages 1–39, March 2008.
- [21] D. Villela, P. Pradhan, and D. Rubenstein. Provisioning Servers in the Application Tier for E-commerce Systems. In *Proceedings of the 12th IWQoS*, June 2004.
- [22] K. Xiong and H. Perros. Qrp01-6: Resource optimization subject to a percentile response time sla for enterprise computing. In *Global Telecommunications Conference, 2006. GLOBECOM '06. IEEE*, pages 1–6, 27 2006-dec. 1 2006.
- [23] S. F. Yashkov. Processor-sharing queues: some progress in analysis. *Queueing Syst. Theory Appl.*, 2(1):1–17, 1987.
- [24] Q. Zhang, L. Cherkasova, N. Mi, and E. Smirni. A regression-based analytic model for capacity planning of multi-tier applications. *Cluster Computing*, 11(3):197–211, 2008.