

# Quorum: Flexible Quality of Service for Internet Services

Josep M. Blanquer, Antoni Batchelli, Klaus Schauer and Rich Wolski  
*Department of Computer Science, University of California Santa Barbara*  
{blanquer,tbatchel,schauser,rich}@cs.ucsb.edu

## Abstract

In this paper we describe Quorum, a non-invasive approach to scalable quality-of-service provisioning that uses traffic shaping, admission control, and response monitoring at the border of an Internet site to ensure throughput and response time guarantees.

We experimentally compare an implementation of Quorum both to hardware over-provisioning and to leading software approaches using real world workloads. Our results show that Quorum can enforce the same QoS guarantees as either of the compared approaches, while achieving better resource utilization than over-provisioning and without the application rewriting overhead required by intrusive software approaches. We also demonstrate that our implementation can successfully handle extreme situations such as sudden traffic surges, application misbehavior and node failures. Furthermore, we demonstrate the flexibility of Quorum by providing QoS guarantees for a complex and heterogeneous Internet service that cannot be implemented by other current software approaches.

## 1 Introduction

The current commercial importance of Internet services makes it imperative for companies relying on web-based technologies to offer and guarantee predictable, consistent, and differentiated quality of service (QoS) to their consumers. For example, e-commerce companies often want to provide faster response times for purchasing than for catalog browsing to ensure that no sale is lost due to the perception of an unresponsive transaction. Differentiated QoS also enables more general and flexible application hosting environments. For example, a service provider that hosts a personalized webmail portal for several companies wants to guarantee different levels of service to its customers and to ensure that these service guarantees are provided to each customer independently, regardless of overload or misbehavior of the others.

To meet large demand, scalable Internet services are commonly hosted using clustered architectures where a number of machines, rather than a single server, work

together in a distributed and parallel manner to serve requests. Delivering reliable service quality guarantees in this distributed setting is the difficult challenge that our work addresses.

Both research and commercial Internet service communities have explored hardware-based and software-based approaches to QoS provisioning. The state-of-the-practice in current commercial settings is to deploy independent clusters for each service (hardware partitioning), each of which comprises enough capacity to service worst-case load conditions (over-provisioning). Unfortunately, because load fluctuations can be substantial, hardware partitioning and over-provisioning incurs a potentially high cost (sufficient resources must be available in each partition to handle load spikes) and low resource utilization (the extra resources are idle between spikes), making this approach inefficient.

As a result, software-based approaches have been proposed and developed to make better use of the resources employed to host Internet services. These approaches focus on embedding QoS logic at different levels of the site's internal software, including operating system [4, 6, 10, 37], middleware [33, 34, 41], and application code [2, 8, 40]. It is the function of this logic to distribute, effectively, the workload among the cluster resources as a way of improving both resource utilization and client experience. Low-level techniques have been shown to provide a tight control on the utilization of resources (e.g., disk bandwidth or processor usage) while techniques that are closer to the application layer are able to satisfy QoS requirements that are more directly experienced by clients. However, these software solutions require the hosted application services and/or the hosting operating system to be customized for QoS provisioning, thereby limiting flexibility and extensibility. Furthermore, most current Internet sites include a myriad of different hardware and software platforms which are constantly evolving and changing. An invasive QoS solution that requires the reprogramming of hosted service code carries with it high development and testing costs when new services are introduced, or the existing

site components (hardware and software) are reconfigured, upgraded, extended, etc. More problematically, the source code for many service components hosted at a site may not be available for proprietary reasons. This lack of source code makes the necessary software reprogramming remarkably difficult. Thus the growing complexity associated with Internet service hosting in commercial settings makes intrusive software QoS strategies less attractive as the need for extensibility and flexibility increases.

To address these needs, we propose a new approach to QoS provisioning for Internet services. Our approach offers reliable QoS guarantees at a lower cost than state-of-the-practice techniques, while giving the service providers the much needed flexibility that they require to rapidly reconfigure, upgrade and extend their complex set of services. In this paper we present *Quorum*, a non-invasive software approach that treats the cluster and the services it is hosting as a “black-box” system and uses only feedback-driven techniques to control dynamically which and when each of the requests from the clients is forwarded into the cluster. Because traffic shaping and admission control is done at the entrance of the site, and the system uses only the observed request and response streams for its control algorithms, new services can be added, old ones upgraded, and resources reconfigured without re-engineering the necessary QoS mechanisms into the services themselves or the system software that supports them.

We report on an implementation of Quorum and its experimental comparison with the state-of-the-practice (i.e., over-provisioning) and state-of-the-art (e.g., Neptune [33, 34]) software solutions using realistic services, client request traces and clustered machines. Neptune is a research and now commercially successful middleware system that implements QoS for Internet services, but which requires the services themselves to be re-written to use Neptune primitives. Using the *Teoma* [35] search engine, which is explicitly programmed so it can use Neptune, we show that Quorum can enforce the same QoS guarantees as Neptune for Neptune-enabled services, but without the additional engineering overhead associated with modifying the services that it supports. Furthermore, we illustrate Quorum’s ability to handle extreme situations such as sudden traffic surges, or internal application misbehavior – capabilities that are necessary for a successful deployment in large-scale, realistic settings. We also demonstrate the flexibility of Quorum by showing how it can provide QoS guarantees for complex heterogeneous Internet services which cannot be modified – a capability that none of the published, pre-existing software approaches is capable of achieving at present.

## 1.1 Contributions

This paper makes five main contributions:

- We present Quorum as a novel approach to QoS provisioning for large-scale Internet services that uses only observed input request and output response streams to control the load within the site so that quality guarantees are met.
- We describe a working implementation and demonstrate its viability using a large cluster system hosting commercial and community benchmark Internet services.
- We compare Quorum with the state-of-the-practice and state-of-the-art approaches in terms of efficiency and the degree to which they maintain QoS guarantees for both throughput and response times.
- We show the robustness of Quorum in successfully overcoming extreme situations (i.e., sudden traffic surges, application misbehavior and node failures) which arise in current commercial settings.
- We demonstrate that the flexibility provided by Quorum enables more efficient deployments of complex, heterogeneous Internet services than can currently be supported by existing approaches.

The remainder of this paper is organized as follows. Section 2 defines the models and assumptions employed by our approach and formally states the terms in which the QoS challenge is defined. Section 3 introduces Quorum’s approach and further describes its architecture. Section 4 experimentally compares Quorum to the best of the known approaches. In Section 5 we demonstrate the robustness of Quorum under extreme situations and also show its flexibility in providing reliable QoS guarantees in complex heterogeneous services. In Section 6 we discuss related work, and we conclude in Section 7.

## 2 Defining the QoS Challenge

Before describing the architecture of Quorum we define the models and assumptions employed by our approach, as well as detail the terms by which the QoS challenge is defined. We begin by outlining the model of Internet service transactions we use. We treat Internet services (see Figure 1) as a stream of requests coming from clients that are received at the entrance of the site, computed by the internal resources, and returned back to the clients upon completion. In the case of system overload or internal error condition, requests can be dropped before completion and thus may not be returned to the client. Each of the

requests that are received from the clients can be classified or grouped into different *service classes* according to a combination of service type and client identity.




Figure 1: System model for Internet services.

The computation of requests is modeled by treating the cluster as a parallel and multi-level resource system that processes requests in a time-shared manner. More specifically we model the cluster site as a black-box system that has the following two properties: 1) unbiased treatment: any request entering the cluster will be computed with the same priority, and 2) time-multiplexed: the internal computation is done in a multiplexed way where requests interleave the usage of resources in time intervals that are short, relative to the response time guarantees. The simplicity of our model allows cluster systems to be treated analytically, while remaining powerful enough to capture the behavior of most time-shared systems and thus the majority of existing sites. Note that these properties are defined in terms of the overall cluster behavior and may not necessarily hold true for each of the internal resources individually.

We define the QoS challenge as the ability to guarantee, at all times, a predefined quantitative characterization of the traffic in each service class as measured at the *output* of the cluster. Such traffic characterization is expressed through a QoS policy, which contains the desired QoS guarantees for each of the participating service classes. Such quality guarantees are defined at the boundary of the cluster and do not extend to requests traversing the internet back to the end users. Our goal is not to provide end-to-end guarantees as we see network QoS as a complementary function. We consider the quality characteristics defined in the QoS policy to be specified in terms of statistical (or soft) guarantees.

Finally, our system requires that the specified QoS policy that must be enforced is, in fact, *feasible* for the expected workloads and cluster capacity. A QoS policy is feasible if the existing cluster can meet the QoS guarantees without requiring any QoS mechanisms (i.e., a simple load-balancer) subject to the following two conditions: 1) incoming rates for each class are always kept below their guaranteed throughput and 2) resource de-

mands of incoming requests do not surpass their expected computation requirements. The expected computation requirements for a request stream are dependent on the type of service offered and must be agreed upon, *a priori*, between the provider and the consumer. In fact, feasibility is an implicit test or calculation that any service provider must perform when dimensioning their cluster for a given expected demand. Note feasibility has already taken into account the software and hardware configuration of the cluster, as well as any possible internal bottlenecks that may occur for the expected workloads. Since Quorum depends only on the feasibility condition, it can continue to operate correctly regardless of the presence of internal cluster bottlenecks.

### 3 The Quorum Architecture

In this section we describe the architecture of Quorum that follows our previously stated model and assumptions. In Quorum, the QoS policy is specified as a list of *QoS classes* describing the quality that must be ensured for each class of service. We define a QoS class as a tuple that describes: 1) how to identify requests of this class (*classification rules*) and, 2) what type of QoS must be enforced (*output guarantees*). In the same way as level-7 load-balancers [19, 20, 29], Quorum classifies requests based on a combination of parameters such as IP address, port, URL and path. Output guarantees are specified in terms of guaranteed minimum throughput and maximum response time. For example, Table 1 describes a QoS policy containing two QoS classes for a service provider hosting webmail portals for two different companies. In the example, *BigCorp* has a much higher guaranteed throughput due to an expected higher traffic volume and *SmallCorp* requires much tighter response time guarantees for its users. Notice that the definition of output guarantees includes both throughput and response time requirements. While it is often possible to meet one type of guarantee at the expense of the other, our solution accommodates both. Additionally, Quorum allows throughput and response time guarantees to be expressed using either percentiles or averages since the way in which each customer wishes to view a guarantee varies. In both cases, however, the time frame over which the average or percentile is computed is substantially longer than the time required to service an individual request.

Quorum uses a single-policy enforcement engine to intercept and control in-bound traffic at the entrance of the site hosting the services. By tracking the responses to requests that are served within the site, our system automatically determines when new requests can be allowed entry such that a specified set of QoS guarantees will be enforced. No knowledge of the internals of the

QoS Class	Classification	Output Guarantees	
	Layer 7 pattern	Throughput (Avg. req/s)	Response Time (95% ms)
BigCorp	http://bigcorp.com/mail/*	800	1000
SmallCorp	http://smallcorp.com/mail/*	50	200

Table 1: Example QoS policy for a service provider hosting webmail portals for two different companies.

site are needed and no instrumentation is required. In other words, to make an Internet site capable of providing QoS guarantees it is enough to deploy Quorum at its entrance point and define the desired QoS policy to be enforced.




Figure 2: The architecture of Quorum.

Figure 2 depicts the architecture of Quorum, consisting of four different modules each of which implements part of the functionality that is necessary to enforce a QoS policy. The *Classification* module categorizes the intercepted requests from the clients into one of the service classes defined in the QoS class. The *Load Control* module determines the pace (for the entire system and all client request streams) at which Quorum releases requests into the cluster. The *Request Precedence* module dictates the proportions with which requests of different classes are released to the cluster. The *Selective Dropping* module drops requests of a service class to avoid introducing work accumulation that would cause a QoS violation. In the next sections we detail further the implementation of the Quorum modules. We explicitly exclude the details associated with Classification since it is a well understood problem that has already been studied in the literature [22].

### 3.1 Load Control

The Load Control module has two primary functions. First, it prevents large amounts of incoming traffic from overloading the internal resources of the cluster. When the internal resources become overloaded, the internal software (i.e., operating system, web servers, applications, etc.) will delay or drop requests without regard for their QoS classification. Second, it maintains the resources within the cluster at a high level of utilization to achieve an overall good system performance (for the given cluster configuration). The goal of the Load Control module is to have the cluster operate at maximum

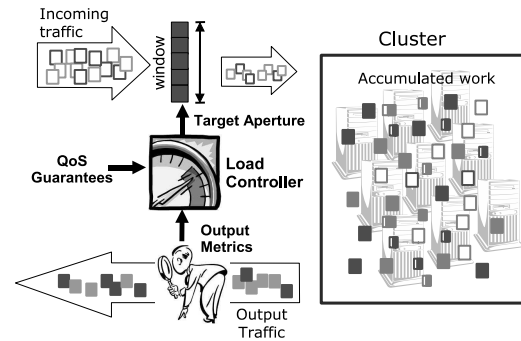


Figure 4: Structure of Load Control module.

capacity so that the largest possible capacity guarantees can be met, while also preventing overload conditions that would cause response time guarantees to be violated.

Based on the previously described time-shared model, our implementation exploits the direct correlation between the amount of work accumulation inside the cluster and the time required for requests to be computed by the hosted services. In general, more work introduced into the cluster corresponds to longer compute times for each service (given a fixed amount of resources) once the number of parallel requests exceeds the number of resources. With this in mind, the Load Control module can directly affect the amount of time that requests take to be computed inside the cluster (i.e., compute time) by controlling how much traffic is “in progress” at any time.

Similar to TCP, our implementation uses a sliding window scheme that defines the maximum number of requests that can be outstanding at any time (see Figure 4). The basic operation of the Quorum engine consists of successively incrementing the size of the window until the compute times of the QoS class with the most restrictive response times approaches the limits defined by its guarantees. Our current implementation uses a simple algorithm (see Figure 3.a) that increments (or decrements) the window linearly until the currently observed computing times are considered too high by the Selective Dropping module (see Section 3.3 for details on how this threshold is chosen). Our implementation updates the window size every 500ms, a compromise between having fast reaction times and allowing enough requests to finish within a period such that the collected computing times are significant. A more sophisticated (and reactive) version of the algorithm using non-linear variation of the window sizes is under study.

### 3.2 Request Precedence

The main function of Request Precedence is to virtually partition the cluster resources among each of the service classes. Resource isolation is a necessary functional-

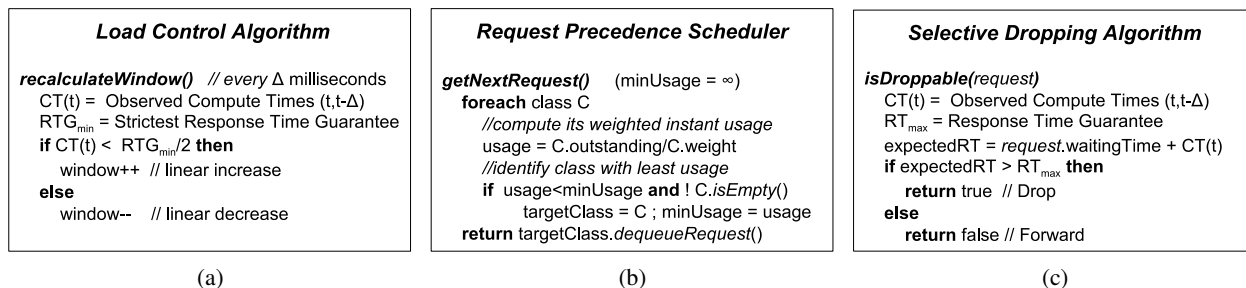


Figure 3: Simplified pseudo-code for the algorithms of the three main modules of Quorum.

ity that allows each service class to enjoy a minimum amount of processing capacity, independent of potential overload or misbehavior of others. This module is able to partition externally the service delivered by the cluster, by controlling the proportions in which the input traffic for each class is forwarded to the internal resources. Thus, the goal of this module is to ensure that the fraction of the overall cluster capacity devoted to each class is large enough to satisfy its throughput guarantees at all times.

The Request Precedence module also attempts to maximize performance in overload situations without allowing guarantees to lapse. It reassigns unclaimed resources to other QoS classes demanding more processing power than they have been granted. Reassigning unutilized capacity allows the QoS engine to take full advantage of the available cluster resources allowing some service classes to enjoy a level of service that is higher than what they have been guaranteed. At the same time, the Request Precedence module ensures that those classes that are not using their maximum allowable share of the overall capacity none-the-less receive enough capacity to meet their guarantees. By continually calculating and adjusting the fraction of cluster capacity that is given to each class, Quorum differs from an approach that relies on physical partitioning of the resources where temporary reassignment cannot be implemented.

Under Quorum, Request Precedence is implemented by a scheduling algorithm that logically partitions the window of outstanding requests (as dictated by the Load Control module) according to the throughput guarantees specified in the QoS classes. This method exploits (and depends on) the time-shared nature of clusters which assign resources equally amongst all running tasks. As a result, it is possible to increase the share of the cluster resources for a particular service class by increasing the number of tasks that are devoted to computing its requests. In particular, our model defines that each request that is being computed in the cluster will get  $\frac{1}{N}$ th of the total utilization, when  $N$  is the number of requests that are concurrently being processed in the cluster. Therefore, the aggregated utilization of  $M$  requests

will amount to  $\frac{M}{N}$ th of the total cluster usage. In the case where all of the  $M$  requests belong to the same service class, that service class is effectively enjoying a  $\frac{M}{N}$ th of the cluster's resources.




Figure 5: Function of Request Precedence module.

To that effect, our scheduler assigns a weight  $\Phi_i$  to each of the QoS classes and uses this weight to partition proportionally the window accordingly (see Figure 5). Instead of allocating a fixed number of slots of the window per class, our algorithm (see Figure 3.b) uses a dynamic method that achieves similar characteristics to Weighted Fair Queuing disciplines [11, 17, 31] in terms of proportional rate guarantees and reassignment of surplus. However, the guarantees in our case apply to window sizes instead of service rates (i.e., throughput). The reason for this choice is that throughput for a given service class can only be guaranteed when the computing requirements of the requests are known. In other words, the capacity necessary to achieve a given throughput is directly related to the computational complexity of the requests. On the other hand, assigning a particular window size corresponds to guaranteeing a portion of the cluster capacity, independent of the computing complexity of the incoming request stream. Therefore, by working with a capacity measure (i.e., proportions of outstanding requests), Quorum can provide effective isolation between classes when their computing requirements are not known *a priori* or can change dramatically.

Capacity can be seen as a fungible metric that links output throughput and computing requirements such that an increase in one can be made to force a decrease in the

other. For example, a capacity equivalent to 10 nodes may correspond to an output throughput of 500 req/s at a compute cost of 20ms/req, but also to 1000 req/s if the compute cost is only 10ms/req. The internal capacity allocated for a class is calculated from the *nominal* guaranteed throughput (as expressed in the QoS class) and the *expected* computation requirements of the requests (as agreed upon between the provider and the consumer). In the cases where the computation complexity is violated (i.e., higher than agreed upon) for a particular class, instead of dropping the traffic of the faulty class, Quorum will gracefully degrade its throughput to maintain the same internal capacity allocation.

### 3.3 Selective Dropping

The function of Selective Dropping is to discard the excessive traffic received for a QoS class in the situations where there is not enough available capacity to fulfill its incoming demands. A dropping module is necessary to prevent large delays from occurring in overloaded situations where requests would otherwise accumulate unboundedly in the engine and violate the QoS guarantees. The goal of the Selective Dropping module is to ensure that the response time guarantees of each class will be met for all requests that can be serviced.

The basic operation of the Selective Dropping module is, in essence, very simple, since it can leverage from properties that are already provided by the Load Control and Request Precedence modules. The Selective Dropping module independently observes each of the QoS queues of the engine and discards the requests that have been sitting in the queue for so long that the deadline for their service cannot be met. In our implementation (Figure 3.c), a request will be dropped if the time left for meeting the deadline once it gets at the head of the queue is less than the expected time of computation of its class. In other words, a request will be dropped if we expect it to miss its deadline according to how other requests of the same class are currently performing (Figure 6). The current computation times for a class can be considered a reliable estimation of their expected computation since they are stabilized by the feedback loop of the the Load Control module.

The Selective Dropping module leverages the queuing inside Quorum to absorb safely peaks of traffic during transient overload conditions without violating the response time guarantees. To this effect, it works closely with the Load Control module by signaling ahead of time when a service class is likely to become overloaded. In our implementation we signal the Load Control module to stop increasing the load of the cluster when the observed computing time of the most restrictive service class reaches half of its response time guarantee. By




Figure 6: Structure of Selective Dropping module.

closely working with the Load Control module, Selective Dropping can ensure that there is an available queuing time that is at least half the maximum allowed response time. Note, however, the queuing time is independent for each class, therefore classes with looser response time guarantees can support longer queuing periods and thus absorb of much larger transient peaks of traffic without violating the guarantees. The choice of ‘half’ is a compromise motivated by the tradeoff between maintaining cluster occupancy and allowing the necessary queuing space to absorb peaks of traffic. We are currently working on an optimized version that can dynamically adapt this threshold to allow more queuing without adversely affecting overall system performance.

Finally, the Selective Dropping module must also ensure that dropping requests of a particular service class does not incur in a violation of its throughput guarantees. To ensure that there are no invalid drops, our module relies on an important property of the Request Precedence module which states that the forwarding rate of requests for a class into the cluster will be no lower than its guaranteed throughput. This property is derived from the capacity guarantees and the windowing system of the Request Precedence module which allows a new request to be forwarded immediately after one of the same class finishes. Therefore, this property ensures that requests will only accumulate in the engine if the incoming rate of a class surpasses its guaranteed throughput, in which case drops can safely be executed since they would not violate any throughput guarantees. Note that these properties do not hold true for misbehaving classes where the computation requirements of incoming requests are higher than expected. However, this is not a problem since we have already discussed that QoS guarantees do not need to be met for such classes, which can be penalized both in terms of throughput and response times. The implementation of independent dropping techniques, coupled with the guarantees given by Load control and Request Precedence, allow this module to provide response time guarantees and isolate one class against misbehavior of others.

Combined, the functions of all four Quorum modules (Classification, Load Control, Request Precedence and Selective Dropping) enable cluster responsiveness, efficient resource utilization, capacity isolation and delay differentiation, thus guaranteeing capacity and response times for each independent service class.

## 4 Experimental Performance Comparison

In this section we demonstrate that the four modules of Quorum can provide QoS guarantees under realistic conditions even though they treat the cluster resources and Internet services as a “black-box”. We have performed extensive studies of each of the presented modules, both in isolation as well as operating together. Due to space constraints we do not include them in this paper, but the details of these studies can be found in [12]. Instead, in this section we focus on examining the performance of Quorum as a complete system, and study how it compares to the best of the known approaches. Our investigation is empirical and is based on the deployment of an Internet search service used by *Teoma* [35] using a 68-CPU cluster. We analyze how five different techniques (representing both state-of-the-practice and state-of-the-art) offer differentiated quality to distinct groups of customers using generated message traffic based on web-search traces. We then quantify the observed quality of service delivered by each method.

### 4.1 Experimental Methodology

Our experimental setup consists of several client machines accessing a cluster system through an intermediate gateway/load-balancer machine. Accessing the services through a load balancer machine is the most commonly used architecture in current Internet services. For example, Google [21] funnels traffic through several Netscaler [29] load-balancing systems to balance the search load presented to each of its internal web servers [23].

To perform our experiments in the most realistic possible manner, we have deployed a commercial-grade Internet service on a 68-CPU cluster system and replayed real traffic traces from its commercial operation [33]. The service deployed is the index search component of the Teoma commercial search service [35]. The index search component consists of traversing an index database and retrieving the list of URLs that contain the set of words specified in the search query. The total size of the index database used is 12GB and is fully replicated at each node. The index search application from Teoma is specifically built for the Neptune middleware [34], a cluster-based software infrastructure that provides replication, aggregation and load balancing for




Figure 7: Experimental test-bed used for our benchmark using Teoma’s search service.

network-based services. The version of Neptune we use also provides QoS mechanisms allowing the specification of proportional throughput guarantees and response times constraints through the definition of yield functions [33]. As it is the case with commercial search engines, our system accesses the service through a set of front-end machines that transform the received URLs into internal queries that are then forwarded to the middleware servicing the search database for processing. To mimic the environment at Teoma, we implement the front-end with an Apache web server [3] and a custom-built Apache module that interfaces with the Neptune infrastructure. This module is necessary to utilize the middleware functionality to locate other Neptune-enabled nodes and appropriately balance the requests based on the current load of the available servers. The cluster configuration used in our experiments is depicted in Figure 7. The hardware configuration of the cluster consists of 2.6 MHz Intel Xeon processors each with 3 gigabytes of main memory organized into nodes with either two or four processors per node. The network interconnect between processors is switched gigabit Ethernet and the host operating system is RedHat Linux/ Fedora Core release 1, using kernel version 2.4.24.

Our gateway node is a 4-CPU dedicated machine that can function in two different modes: as a load-balancer or as the Quorum engine. When running in load-balancer mode, the machine is configured to implement the typical (Weighted) Round Robin and maximum connections options available in most commercial hardware [19, 20, 29]. When running as Quorum engine, the gateway is configured to enforce the QoS policy defined for the experiment. Both the load-balancer and Quorum engine are entirely implemented in user-level software. The gateway is implemented as an event-driven Java application which makes extensive use of the new libraries for improved I/O performance [24]. We use Sun’s 1.5 Java virtual machine with low-latency garbage collection settings. Our performance tests show that our implementation can achieve a peak performance of 12Kreq/s (i.e., around 70K packets/sec) for certain client work-

QoS Class	Classification Layer 7 pattern	Output Guarantees		Experimental Workload	
		Throughput (Avg. req/s)	Response Time (Avg. ms)	Input (Avg. req/s)	Service Status
A	Host: A	375	200	185	Not Overloaded
B	Host: B	937	600	1718	Overloaded
C	Host: C	562	300	557	Fully Utilized

Table 2: QoS guarantees and traffic workload of the Teoma search engine benchmark.

loads. Thus the performance of our base-level system is high enough to be used in load levels that are comparable to current commercial systems (e.g. Google reports around 2500 req/sec [23], Ask Jeeves around 1000 req/sec [5]). Both our implementation of a load-balancer and the Quorum engine are based on the same core software for fielding and forwarding HTTP requests.

For this experiment our methodology consists of using the previously described test-bed to recreate search traffic and to explore the effectiveness with which five different approaches can enforce a particular QoS policy for a single service with multiple client groups. The five compared approaches are:

**Load Balancer** The gateway machine is configured as a load balancer and tuned to match common high performance settings of Internet sites. Specifically, we configure it to use the least connections load-balancing algorithm and limit the maximum number of open connections for each front-end to match their configured maximum (i.e., 250 processes for Apache server and 150 for the Tomcat engine).

**Physical Partitioning** A separate group of machines are dedicated for each of the existing QoS classes. We configure the load-balancer to forward requests of a particular class only to its restricted set of reserved nodes.

**Overprovisioning** The size of each physical partition is increased such that the resulting capacity and response time guarantees can be achieved as specified by the QoS policy (possibly at the expense of under utilized resources).

**Neptune QoS** The gateway is configured as a load balancer and the QoS mechanisms of Neptune are enabled to implement the QoS policy under study.

**Quorum QoS** The gateway runs the Quorum engine which implements QoS and the internal cluster resources implement only the Internet service. (i.e., QoS functionality in Neptune is disabled)

In order to benchmark Quorum and the other considered QoS methodologies, client requests are replayed

Throughput (req/s)							
Class	Input	Guarantee	Load Balancer	Physical Partitioning	Overprovisioning	Neptune QoS	Quorum
A	185	375	148 (8)	185 (9)	185 (10)	185 (9)	185 (10)
B	1718	937	1333 (33)	986 (41)	1718 (130)	1133 (68)	1059 (68)
C	557	562	443 (16)	557 (26)	557 (29)	548 (23)	557 (27)

Response Times (ms)							
Class	Guarantee	Load Balancer	Physical Partitioning	Overprovisioning	Neptune QoS	Quorum	
A	200	13335 (3175)	40 (43)	40 (22)	45 (50)	121 (46)	
B	600	13395 (3201)	18937 (3734)	75 (116)	256 (180)	600 (37)	
C	300	13487 (3106)	305 (218)	19 (19)	62 (49)	134 (53)	

Table 3: Experimental results for Teoma search engine.

from a request trace supplied by Teoma that spans 3 different days of commercial operation [33]. We also use Teoma-supplied traces of word sequences to generate real search queries. The levels of incoming traffic are designed so that the input demands of the different clients are far below (class A), far above (class B) and coinciding with (class C) the capacity constraints specified in their respective QoS classes. Clients for each QoS class use different inter-arrival times, corresponding to one of the three different days of the original traces. Table 2 further depicts the details of the QoS policy and input workload used in the experiment, including the capacity and response time guarantees for each QoS class.

## 4.2 QoS Results

Figure 8 presents the results in terms of achieved average throughput and average response times for the five QoS methodologies using the same input request streams. The upper portion of the figure shows how the totality of incoming traffic for a class (represented by the height of a bar) has been divided into traffic that is served and traffic that is dropped. Horizontal marks delimit the minimum amount of traffic that has to be served if the QoS guarantees are met. Note that a resulting throughput below the horizontal marks still meets the QoS guarantee for a class if the totality of its incoming traffic is successfully served (i.e., the system cannot serve more traffic than it has received). The lower part of Figure 8 presents the results in terms of response times. For response times, we use horizontal marks to denote the maximum response times allowed by the QoS policy and denote with a darker color the classes that do not meet the guarantees. We present these response time results using a logarithmic scale for better visual comparison since the delays differ substantially. Table 3 summarizes these results in tabular form (including standard deviations in parenthesis) to further aid their comparison.

We begin by analyzing the quality of the service achieved by a load-balancer-only technique. Throughput results show that the amounts of traffic served in this case are directly dependent on the levels of incoming traffic

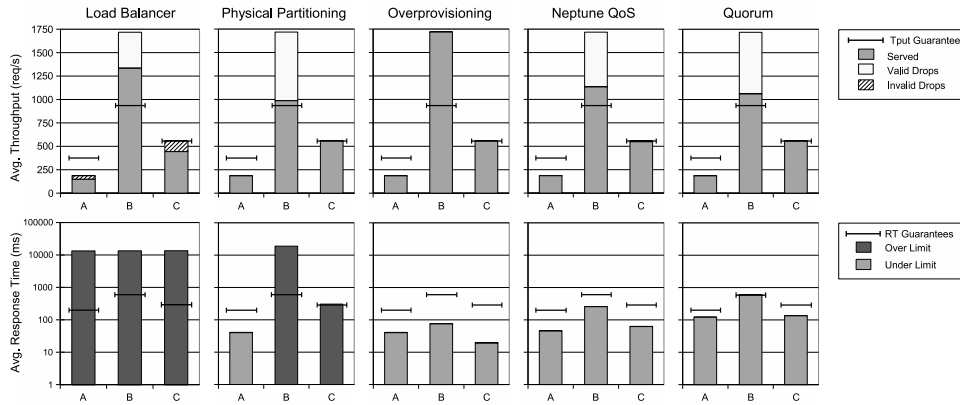


Figure 8: Experimental comparison of current approaches and Quorum, using Teoma's search engine.

rather than driven by the specified QoS policy, thus isolation between classes is not achieved. In this case we see that the dominance of class B traffic induces drops in A and C, even though the demands for these classes are always below (in the case of class A) or never exceed (for class C) the guaranteed capacity for each class. At the same time, the large response times shown in the lower figure, demonstrate that simple connection limiting techniques employed by the load-balancer are not enough to prevent large delays in response times (e.g. up to 14 seconds per request), rendering this technique inadequate to provide QoS guarantees.

When resources are physically dedicated through Physical Partitioning, the system is able to serve the expected amount of traffic for each of the classes and drop requests only in the cases when the demands of incoming traffic exceed the allocated capacity. Throughput guarantees are met, however, if we observe the results in terms of response time, we see that the overloaded partition B experiences a delay more than 30 times higher than the maximum allowed by the QoS policy. Thus while physically partitioning resources is able to provide capacity guarantees, it fails to ensure response times constraints for arbitrary incoming demands. It is worth noting that the reason for partition B serving more throughput than its guarantee is that the raw performance of the partition is slightly higher than the QoS guarantee defined in the policy.

When each of the partitions is augmented with enough resources (i.e., over-provisioning) all requests are successfully served. The response times are also reduced below the maximum allowed delay. In this case, class B and class C require an additional 10 and 2 CPUs respectively in order to meet the specified response time guarantees. Thus over-provisioning is the first of the techniques that can successfully provide both throughput and response time guarantees. However, meeting the QoS guarantees through over-provisioning comes with a high cost. In

our experiment, the increase in cost of overprovisioning was 60% (i.e., from 20 to 32 CPUs) with resource utilization declining to 80%. Further, these numbers represent the *minimum* amount of over-provisioning that allowed us to achieve the QoS goals. In general, between load spikes the extra resources needed to serve surges in load lay idle. Thus, given the wide load fluctuations that most commercial Internet services can experience (i.e., 3-10 times the normal amount [16]), we expect the resource utilization of over-provisioned systems *in situ* to be worse than what we observe in this experiment.

Neptune QoS and Quorum both meet the specified throughput and response time guarantees. Both techniques serve at least the necessary amount of traffic and are able to keep response time below the maximum delays associated with each guarantee. Furthermore, both techniques are able to successfully reassign the capacity not utilized by class A to the greedy clients of class B. We observe that direct control of the resources and services in the cluster (due to its invasiveness) allows Neptune to achieve a slightly better throughput than Quorum (i.e., 3%). This slight performance penalty can be seen as the cost that an external solution such as Quorum has to pay for not modifying any of the software internals. However, given the completely non-invasive nature of Quorum, we were surprised by how closely it matched the performance achieved by the invasive and commercially developed Neptune system. Figure 8 also shows that the resulting response times from Neptune are somewhat lower than Quorum. This difference is because Quorum is only designed to enforce maximum delay constraints and it is not concerned about minimizing the overall delay of service times. We are currently working on a prototype that can both ensure response time constraints and lower response delays when possible.

Summarizing, this experiment demonstrates the effectiveness of Quorum empirically, using a commercial Internet service and commercial traffic levels. Quorum in

this setting is competitive with the best of the current approaches in its ability to enforce both response time and throughput QoS guarantees. In particular, Quorum has less cost and achieves better resource utilization than over-provisioning techniques due to its ability to reassign unutilized capacity to those service classes that need it. At the same time, it achieves comparable QoS guarantees to an integrated and commercially available system such as Neptune, incurring only a small performance cost (i.e., 3%). In the next section (Section 5.4) we illustrate its flexibility by showing how it can provide reliable QoS guarantees in a complex and heterogeneous site running three different services.

## 5 Robustness under Extreme Conditions

In this section we investigate the robustness of Quorum and its QoS enforcement capabilities under scenarios that emulate the extreme conditions experienced by many current Internet services. To do so, we first study the reaction of Quorum to three circumstances: sudden traffic fluctuations (Section 5.1), sudden changes in computing requirements (Section 5.2) and node failures and recoveries (Section 5.3). We then present a larger-scale experiment in which we detail its response to the same conditions in a substantially more complex Internet hosting scenario (Section 5.4).

To conduct the initial set of isolated robustness studies we use two service classes: A and B. Service class A is a misbehaving class that begins with an input load that can be fully serviced with its allocated capacity, and then changes its demands to surpass the capacity required to meet its guarantees as well as to drive the overall system into overload. Service class B is a well-behaved class that receives a constant demand of traffic that is always below the traffic level that can be serviced under its guarantees. For each of the experiments, we detail how well Quorum insulates the quality of service experienced by the well-behaved class B from the fluctuations introduced by class A. We also investigate how the quality of service given to class A degrades gracefully during the periods when its demands exceed the capacity allocated to meet its guarantees. In particular, our goal is to provide as much capacity to A as possible without violating the guarantees made to either A or B. As described in subsection 3.2, however, the capacity allocated to A and B is fungible and constantly adjusted by Quorum as it responds to changes in load conditions.

To run these experiments we use a system consisting of 4-CPU's for client machines accessing a 16-CPU cluster through a gateway machine implementing the Quorum engine. Each of the servers runs the Tomcat application server [36], providing a "CPU-loop service" consisting of a servlet that loops a number of times so that

it utilizes a certain amount of CPU (as specified in the HTTP parameters of each incoming request). This artificial emulation of a true web service allows precise control of the CPU load requirements associated with each request. Requests received from the clients are classified into QoS classes according to the host field name found in the HTTP header of the request (i.e., host: A or B).

QoS Class	Classification	Output Guarantees	
	Layer 7 pattern	Throughput (Avg. req/s)	Response Time (95 <sup>th</sup> % rts)
A	Host: A	900	400
B	Host: B	900	400

Table 4: QoS policy used in the studies.

The QoS policy defined for the experiments allocates the same guarantees for both classes of service (Table 4). Note that unlike the previous experiments, the response time guarantees are expressed in terms of 95th percentiles and not averages – a much more challenging but potentially more desirable metric to enforce, especially given the range of conditions to which we subject the cluster. All figures in this section depict the resulting average of the observed throughput (upper graph) and the 95th percentile of response times (lower graph) over two-second sampling intervals.

### 5.1 Sudden Traffic Fluctuations

In this experiment we show how Quorum manages wide fluctuations of incoming traffic. To demonstrate this property we subject the service for class A to a sudden-but-sustained impulse of incoming traffic that is four times its normal rate. This sudden increase in demand is enough to bring the cluster to full utilization. Figure 9 shows the results from the experiment. In the Figure, the traffic fluctuation (labeled as "Input Class A") increases instantly from 600 req/s to 2400 req/s 120 seconds after the experiment has begun. Despite the sudden and sustained increase in A's traffic the degree to which service class B meets its guarantees is isolated from the change in input conditions. B's throughput is virtually unaffected and its response times, while they climb, are always kept below the maximum guaranteed delay. In response to the traffic surge, Quorum quickly shifts any uncommitted resources to class A. Strictly speaking, it is consistent with the guarantee given to class A simply to cap throughput at 900 req/s for that class. However, by automatically sensing the degree to which it can slow down B's response times (without violating B's guarantees) and committing additional resources to A, Quorum is able to give A as much throughput as can be spared while remaining within the constraints of both guarantees.




Figure 9: Quorum’s reaction to extreme fluctuation of incoming traffic.

We should note that the slight spike in response times occurring in second 120 appears a consequence of our short sampling period. We wish to depict circumstances that stress the capabilities of Quorum and as such, we calculate the percentiles with a two-second periodicity. In practice, it is unlikely that a commercial system will need to ensure QoS guarantees on such a fine-grained time scale, especially when using percentiles to specify guaranteed performance levels.

## 5.2 Computing Requirements Overload

In this experiment we investigate how Quorum handles wide variations in the computing requirements associated with a request stream. These types of variations can occur in situations such as application misbehavior (e.g., software bugs that cause excessive resources to be used in computing a request) or changes in the workload characteristics (e.g., requests incurring in unusually long and expensive database queries). We induce this anomaly by suddenly increasing the computing requirements for class A from 8ms to 40ms of exclusive CPU time. Again, the goal is to protect the performance of class B while degrading the throughput given to class A to a level that is both maximal and consistent with the guarantees for both classes. To better observe the expected service for class A we include the throughput guarantees normalized to its incoming computing requirements (i.e., the normalized throughput is five times lower than the nominal when requests are five times more difficult to compute).

Results from the experiment are depicted in Figure 10. As in the previous experiment the throughput given to class B remains virtually unaffected by the increase in computing requirements (seconds 120-180), and its response times are always kept below the guarantees. At the same time, in response to the increase in computing demands for the misbehaving class A, Quorum immedi-

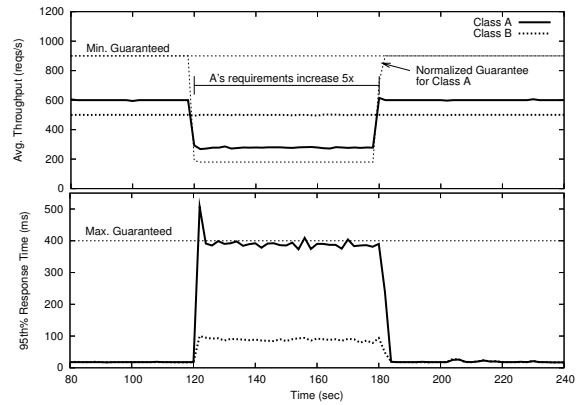


Figure 10: Behavior of Quorum when requests of class A suddenly require five times more resources for their computation.

ately decreases A’s throughput. Although degraded, A’s throughput is always maintained above the normalized guarantee corresponding to the internal capacity allocation Quorum made for this guarantee.

Recall from Section 3.2 that the Request Precedence module guarantees enough resources to class A to fulfill the nominal throughput guarantee of 900 req/s assuming 8ms of computing time. When the computing requirements increase to 40ms/req the throughput must be lowered to 180 req/s to preserve enough capacity for B’s guarantees. Thus we expect the system to enforce a throughput guarantee of 180 req/s for class A during the period in which its requests require 40ms of CPU time, as shown by the normalized guarantee line. However, between seconds 120 and 180 of the experimental period, class A is receiving a throughput of 280 req/s, which includes a surplus of 100 req/s corresponding to the resources that class B is not utilizing. If B’s requirements were to suddenly increase, Quorum would reduce A’s throughput to 180 req/s and change the proportion of B’s requests admitted to reallocate more resources to B. Note also that this constant allocation and reallocation of capacity is sensed by the Quorum engine automatically based on the observed responses leaving the cluster, and not based on predefined parameters or instrumentation describing the CPU requirements for each type of request. As is the case with the previous experiment, the short time scale over which each percentile is computed causes a single “spike” in response time during the two-second interval spanning second 120.

## 5.3 Node Failures and Recoveries

In this experiment, we depict Quorum’s response to significant node failures and recoveries. At second 120, we induce the failure of 2 out of the 8 nodes and then recover

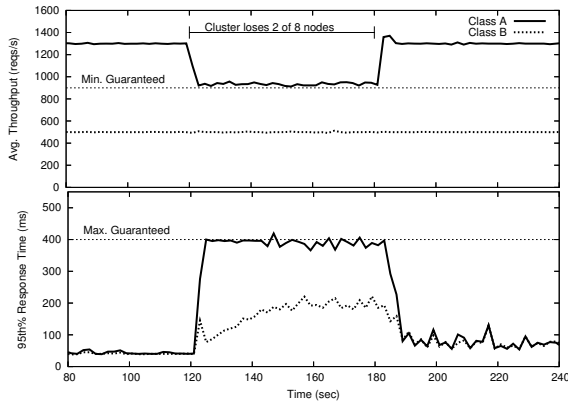


Figure 11: Quorum’s reaction to a failure of 2 nodes.

the nodes 60 seconds later. To introduce these failures we program our load-balancer module to stop forwarding traffic to the “failed” nodes. We have also increased the incoming traffic rate for class A to 1300 req/s in order to make the resulting change in throughput more visible.

We show the results of the experiment in Figure 11. When the nodes fail, Quorum rapidly reduces the throughput given to class A to its 900 req/s guarantee. Notice that this adjustment, again, does not violate the quality of the service guarantees given to class B. As with the previous two experiments, the throughput for B is unaffected while the response times grow to a level well below their maximum guaranteed delay.

We should note that in this example it was possible to enforce the QoS policy, even under the degraded operation, because there was enough spare capacity that B was not utilizing which could successfully be reassigned to A. In the cases where there are not enough resources to fulfill the guarantees across all classes (i.e., QoS policy is not feasible), Quorum reacts by degrading the service of each class proportionally to the guarantee associated with that class. For example, if the input demands for class B had been above the guaranteed 900 req/s, the Quorum would have evenly assigned a throughput of 700 req/s for each class since the degraded capacity of the system would support 1400 req/sec in total, and the guarantees for both A and B are the same. We believe that other non-proportional mechanisms for reapportioning fungible capacity when QoS policies become infeasible are highly desirable and we plan to investigate them further in our future work.

## 5.4 Complex Heterogeneous Services

Through the previous set of controlled experiments we have shown that Quorum can both enforce service isolation as well as gracefully degrade the service of misbehaving classes even under extreme operating condi-




Figure 12: Setup of the complex, heterogeneous Internet site.

QoS Class	Classification	Output Guarantees	
	Layer 7 pattern	Throughput (Avg. req/s)	Response Time (95 <sup>th</sup> % ms)
Teoma	Port=8888	450	600
CPU-Loop	Port=9999	1350	3000
RUBiS	Port=10000	700	3500

Table 5: QoS policy for the complex and heterogeneous Internet site.

tions. We now show how Quorum reacts to the same three severe circumstances for a larger-scale and substantially more complex Internet site that hosts three different services. Additionally, this experiment illustrates the flexibility of Quorum’s “black-box” approach: its ability to provide QoS guarantees using heterogeneous hardware configurations and multi-tiered software architectures where the source code of the applications cannot be modified. At present, we know of no other published infrastructure that can provide QoS for this complex Internet hosting scenario.

To perform this experiment we host the Teoma search and CPU-loop services (described previously) together with a third service called RUBiS [32] using shared set of cluster resources. RUBiS is a publicly available auction site modeled after eBay that has been used by several researchers for evaluating application server performance scalability [14, 15]. We use the version of RUBiS that is implemented using Enterprise Java Beans (EJB) deployed on top of JOnAS application server (v3.3.6) and Tomcat (v4.1) servlet engine. The Tomcat servers are configured with session replication and the JOnAS application server is configured to balance the execution of EJBs across each of its nodes according to their respective loads. The auction data is stored using a mySQL database with the same configuration and size as the benchmark described in [15]. Traffic for the RUBiS auction is generated by the client emulator supplied with




Figure 13: QoS results for a complex, heterogeneous Internet site.

the RUBiS software which performs typical actions of an auction user such as browsing, bidding or buying items. This type of service also allows us to illustrate the efficiency of Quorum when providing QoS guarantees under highly variable workloads. In this particular case the distribution of computation requirements resulted in ( $median=101$ ,  $mean=149$ ,  $95^{th}\%=457$ ,  $max=3088$ )ms which can be approximated with a two-phase hyper-exponential with the first mode on the mean.

Figure 12 depicts the hardware and software configuration used for this experiment. Notice that we include both nodes that are dedicated to a single service as well as nodes that are shared by more than one service. In particular, the CPU-loop service shares 7 of the 8 nodes used by the Search component of Teoma, and also with 2 of the 5 nodes running the RUBiS auction. Our intention is to capture both the fluid sharing of cluster resources as well as the static capacity planning that we believe will always be present in a commercial system.

Also for this experiment we program our Quorum engine with the QoS policy defined in Table 5, deploy it at the entrance of the site (with no other information than the QoS policy), and observe how well it performs in response to the same three types of changes explored in the previous subsections. Similarly, we generate three types of input load. For the Teoma service, we introduce incoming traffic that exceeds what can be completely serviced under the constraints of its guarantee. Alternatively, for the RUBiS service, we keep the incoming traffic load below the maximum serviceable level. We then vary the input for the CPU-loop service to create

a peak of demand during the period from seconds 140 to 220 and to increase its computing requirements from 8ms to 40ms during the period between seconds 300 and 420. Finally we kill one of the Teoma back-end nodes at second 475 and restart it 120 seconds later.

Figure 13 shows the evolution of throughputs (above) and response times (below) for each of the three different services during the 11 minute run, in which a total of 1.1 million requests were served. Vertical lines separate the three different conditions (input increase, computation increase, node failure) to which Quorum must respond. Throughput guarantees are again normalized to the expected computing requirements. Only CPU-loop service shows a deviation from the nominal throughput guarantees since it is the only service that suffers a change in its computation requirements. From the first segment of the figure, it is evident that Quorum protects the RUBiS service and also reassigns the available resources such that the two overloaded classes during the peak period are served according to the QoS policy. As we observed in Section 5.1, the amount of surplus service received by Teoma during the peak period, is given back to the CPU-loop service so that both classes can operate at their limits of throughput and response times.

In the second segment of the figure, the computing requirements of CPU-loop service increase to 5 times their original levels. In this case we induce a change in the computing requirements that is more gradual than the sharp change shown in Section 5.2 to better emulate how a true Internet site might degrade. Quorum reassigns capacity not needed to meet Teoma's guarantees to the

CPU-loop service while maintaining the guarantees for RUBiS. Also, the CPU-loop service suffers a degradation in throughput that is inversely proportional to the increase in its computing requirements, thus maintaining the fungible capacity described by its guarantee. In this case, there are no extra resources to be used in aiding the overloaded CPU-loop class, thus its resulting throughput is capped exactly at its normalized guarantee.

In the third segment of the experiment the dedicated search back-end from the Teoma service fails. In this case we induce a true failure by killing the server process of Neptune and use the fail-over and recovery capabilities of the middleware to detect the change. Note that the failure of the node only has an effect in reducing the spare capacity that Teoma service is enjoying. Both the throughput and response times of CPU-loop and RUBiS are, once more, unaffected.

The results from these experiments illustrate several important points. First, our prototype implementation of Quorum is able to provide *robust* QoS guarantees even in the presence of the extreme conditions which service providers are currently facing. Second, the QoS guarantees are provided in very *fine-grained* time scales even when using a strict metric such as the 95th percentile of the response times. Third, Quorum is a *flexible* QoS solution that can provide performance guarantees in heterogeneous Internet sites without requiring any prior knowledge of their internal hardware architecture or software configuration. Forth, Quorum is capable of handling *complex service types* which can exhibit wide (legitimate) variations in the computation requirements of their requests (e.g., RUBiS auction). In summary, our empirical evaluation shows that Quorum is a viable solution to QoS provisioning for Internet services, that has the robustness and flexibility that current service providers seek without requiring the modification of any of the existing software infrastructure of the sites.

## 6 Related work

There are many approaches to providing QoS for Internet services, but relatively few that combine flexibility and extensibility with response time and throughput performance. In this section we briefly introduce some of the most relevant work and compare it to the Quorum approach.

QoS for network communication is typically defined in terms of reliable communication between two endpoints with performance guarantees. Protocols such as diffserv [9] and intserv [13] or trunk reservation schemes [28] leverage the existing routing infrastructure and network knowledge to provide bandwidth allocation and packet delay guarantees over the Internet. At a higher level, approaches such as Content Distribu-

tion Networks [1] provide similar features by appropriately managing an overlay network to content closer to the end-user. These approaches focus on the communication component and do not address the computational requirements associated with the servicing of Internet requests. In contrast Quorum works at the boundary of the cluster hosting the services and, as such, complements approaches that ensure quality of network service between the client and the cluster.

**Load balancers** [19, 20, 29] are perhaps one of the most closely related approaches to Quorum. Properly tuned, load-balancers can greatly enhance the overall quality of the service offered by a cluster system. Products such as Packeteer [30] offer traffic shaping functionality such that minimum bandwidth guarantees can be allocated to distinct clients or applications. More sophisticated products such as Netscaler [29] apply intelligent connection management that protects the internal cluster nodes from overload in response to large bursts of incoming traffic. However, existing solutions are not aimed at providing throughput and response time guarantees, but are mainly designed to enhance the overall system performance. Furthermore, these techniques rely on the proper configuration of the load-balancers by an expert operator who knows and understands the internal operation of the site to be protected. As such, these are static configurations that are highly tuned for specific settings and that must be repeated for any change occurring in the site's internals. Quorum differs from these approaches in that it guarantees QoS in terms of both throughput and response times. At the same time Quorum does not need to be configured explicitly or tuned by an expert for the specifics of the hardware or software of the site.

At the **operating systems** level, the QoS challenge is typically addressed in terms of resource management. Many research operating systems [6, 10, 37] achieve tight control on the utilization of resources as a way of enforcing capacity isolation between service classes. Although these techniques have proven to be effective in terms of capacity isolation, they are not designed to provide response time guarantees. Furthermore, these techniques control the resources within a single machine and thus cannot be easily extended to clustered environments. One notable exception is Cluster Reserves [4] – a single-node approach that has been scaled to span clustered resources. Although this technique is shown to provide resource isolation at the cluster level, like its single-machine counterparts, it does not provide response time guarantees. Quorum is also a cluster-wide QoS solution that provides both capacity and response time isolation as well as throughput and response time guarantees. It also differs from systems such as Cluster Reserves in that it does not require customization of the operating system used by the cluster's internal nodes.

**Middleware systems** such as Neptune [34, 33] or Application Server [38, 7] include QoS functionality as part of a distributed and potentially scalable infrastructure. By programming the applications to use these primitives it is possible to construct distributed services that offer cluster-wide QoS guarantees. However in order for these frameworks to be effective each of the constituents of a service must be integrated with the middleware infrastructure. This often poses a very restrictive constraint given the heterogeneity and proliferation of current Internet services. Similar approaches that embed the QoS logic directly at the **application level** have also been proposed. For example, the approach presented in SEDA [40] advocates the use of a specific framework for constructing well-conditioned scalable services and [39] shows the effectiveness of this framework when explicit QoS mechanisms are built to prevent overload in busy Internet servers. Rather than building an application with QoS support, other work has modified existing applications to include QoS capabilities [2, 26]. For example, the work done in [2] shows how it is possible to modify the popular Apache web server to provide differentiated services without the use of resource management primitives at the operating system level. However, as is the case with middleware approaches, the large cost of modifying the application code to include QoS mechanisms is only effective if the entirety of the software deployment is able to function in a concerted way towards providing QoS. With Quorum, the applications hosted in an Internet site do not need to be modified or designed for any particular operating system or middleware infrastructure and can directly be used in their native non-QoS state.

Some recent work has investigated resource management techniques using **non-invasive approaches**. Façade [27] is a prototype implementation of a storage controller that throttles I/O requests to a (black-box) disk array. Similar to Quorum, it provides response time isolation (but no throughput isolation) for different I/O streams. However, response time guarantees can only be enforced as long as the total incoming load is below the capacity of the disk array (i.e., no dropping mechanism is implemented). In [25], Jin et al. analyze the effectiveness of several share-based scheduling techniques for differentiating service quality in networked servers. Some of the project goals are similar in nature to Quorum, however the analysis is done only through simulation, focuses only on storage server facilities and does not include a performance study in dynamic scenarios. Furthermore, the devised method is somewhat invasive since it requires offline profiling of the workload and more importantly assumes that the cost of every single request can be known at scheduling time. Other work such as Gatekeeper [18] proposes a proxy system, much like Quorum, that implements admission control for e-

commerce applications. However, Gatekeeper is not designed to provide any QoS guarantees, but targeted to reduce the overall response times and improve the performance of the system. Furthermore, it has only been tested in reduced size systems, it targets database backends and relies on extensive profiling of the service applications.

## 7 Conclusions and Future Work

Commercial Internet service provisioning depends increasingly on the ability to offer differentiated classes of service to groups of potentially competing clients. In addition, the services themselves may impose minimum QoS requirements for correct functionality. However, providing reliable QoS guarantees in large-scale Internet settings is a daunting task. Simple over-provisioning and physical partitioning of resources can be effective but inefficient. Invasive software approaches overcome the inefficiency problem but at the expense of reprogramming and/or re-engineering of the services within a site to implement QoS functionality.

In this paper we present an alternative, non-invasive software approach called Quorum that provides efficient QoS provisioning for Internet services while allowing new levels of flexibility that current service providers require. The presented system functions at the border of an Internet site and uses traffic shaping, admission control, and response feedback to treat the site as a “black-box” control system. Quorum intercepts the request and response streams entering and leaving a site to gauge how and when new requests should be forwarded to the hosted services to ensure throughput and response time guarantees.

We demonstrate the capabilities of our Quorum implementation by experimentally comparing it to the best state-of-the-practice and state-of-the-art approaches. Our results show that, despite being non-invasive, Quorum can enforce the same QoS guarantees as either of the compared techniques, while achieving better resource utilization than over-provisioning and without the application rewriting overhead required by intrusive software approaches. We also demonstrate that our implementation can successfully handle extreme situations such as sudden traffic surges, application misbehavior or node failures. Further, we also demonstrate the powerful flexibility of Quorum by providing QoS guarantees for a complex and heterogeneous Internet service that suffers the same type of harmful conditions. At present, we know of no other published infrastructure that can provide QoS under these challenging conditions. Encouraged by the performance of our results we are currently working on both enhancing the performance and scalability of the Quorum engine as well as improving our algorithms with

more sophisticated control mechanisms. Also we are interested in deploying Quorum on a wider array of Internet services including real commercial sites.

## Acknowledgements

This research was supported by the National Science Foundation (grant ITR-0220139), University of California MICRO (grant 02-067) and Sun Microsystems.

## References

- [1] Akamai Technologies, Inc. Content Delivery Network. <http://www.akamai.com>.
- [2] J. Almeida, M. Dabu, A. Manikutty, and P. Cao. Providing Differentiated Quality-of-Service in Web Hosting Services. In *Proceedings of the First Workshop on Internet Server Performance*, June 1998.
- [3] Apache HTTP server. <http://httpd.apache.org>.
- [4] M. Aron, P. Druschel, and W. Zwaenepoel. Cluster Reserves: A Mechanism for Resource Management in Cluster-based Network Servers. In *Proceedings of the ACM Sigmetrics Conference on Measurement and Modeling of Computer Systems*, Santa Clara, California, June 2000.
- [5] Ask Jeeves Search Engine. <http://www.ask.com>.
- [6] G. Banga, P. Druschel, and J. Mogul. Resource Containers: A New Facility for Resource Management in Server Systems. In *Proceedings of the Third Symposium on Operating Systems Design and Implementation*, 1999.
- [7] BEA Weblogic Server. <http://www.bea.com>.
- [8] N. Bhatti and R. Friedrich. Web Server Support for Tiered Services. *IEEE Network*, September 1999.
- [9] S. Blake, D. Black, M. Carlson, E. Davies, Z. Wang, and W. Weiss. An Architecture for Differentiated Services. RFC 2475, Dec. 1998.
- [10] J. Blanquer, J. Bruno, E. Gabber, M. McShea, B. Özden, and A. Silberschatz. Resource Management for QoS in Eclipse/BSD. In *Proceedings of the First FreeBSD Conference*, Berkeley, California, Oct. 1999.
- [11] J. Blanquer and B. Özden. Fair Queuing for Aggregated Multiple Links. In *Proceedings of the ACM SIGCOMM*, San Diego, CA, August 2001.
- [12] J. M. Blanquer, A. Batchelli, K. Schauer, and R. Wolski. Quorum: Providing Flexible Quality of Service for Large-Scale Internet Services. In *UCSB Computer Science, Technical Report TR-2004-27*, 2004.
- [13] R. Braden, D. Clark, and S. Shenker. Integrated Services in the Internet Architecture: An Overview. RFC 1633, July 1994.
- [14] E. Cecchet, A. Chanda, S. Elnikety, J. Marguerite, and W. Zwaenepoel. Performance Comparison of Middleware Architectures for Generating Dynamic Web Content. In *ACM/IFIP/USENIX International Middleware Conference*, Rio de Janeiro, Brazil, Jun 2003.
- [15] E. Cecchet, J. Marguerite, and W. Zwaenepoel. Performance and Scalability of EJB Applications. In *Proceedings of the ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications*, Seattle, Washington, Nov 2002.
- [16] CNN.com Article: Internet Proves Vital Communication Tool. <http://archives.cnn.com/2001/tech/internet/09/12/attacks.internet/>.
- [17] A. Demers, S. Keshav, and S. Shenker. Design and Simulation of a Fair Queuing Algorithm. In *Proceedings of the ACM SIGCOMM*, Austin, Texas, September 1989.
- [18] S. Elnikety, E. Nahum, J. Tracey, and W. Zwaenepoel. A Method for Transparent Admission Control and Request Scheduling in E-Commerce Web Sites. In *Proceedings of the International World Wide Web Conference*, New York, May 2004.
- [19] F5 Networks. Big/IP Load Balancer. <http://www.f5.com>.
- [20] Foundry Networks, Inc. Server Iron Internet Traffic Management. <http://www.foundrynet.com>.
- [21] Google. Internet Search Engine. <http://www.google.com>.
- [22] P. Gupta and N. McKeown. Algorithms for Packet Classification. *IEEE Network Special Issue*, March 2001.
- [23] P. Hochmuth. Speedy Returns are Google's goal. *Network World Fusion article*, Jan 2003. <http://www.nwfusion.com/news/2003/0901google.html>.
- [24] Java NIO. <http://java.sun.com/j2se/1.5.0/docs/guide/nio>.
- [25] W. Jin, J. Chase, and J. Kaur. Interposed Proportional Sharing for a Storage Service Utility. In *Proceedings of the ACM Sigmetrics Conference on Measurement and Modeling of Computer Systems*, New York, June 2004.
- [26] Y. Lu, T. F. Abdelzaher, and A. Saxena. Design, Implementation, and Evaluation of Differentiated Caching Services. *IEEE Transactions on Parallel and Distributed Systems*, May 2004.
- [27] C. Lumb, A. Merchant, and G. Alvarez. Façade: Virtual Storage Devices with Performance Guarantees. In *Proceedings of the USENIX Conference on File and Storage Technologies*, San Francisco, CA, March 2003.
- [28] D. Mitra, R. Gibbens, and B. D. Huang. State-Dependent Routing on Symmetric Loss Networks with Trunk Reservations I. *IEEE/ACM Transactions on Communications*, February 1993.
- [29] Netscaler, Inc. Request Switching Technology. <http://www.netscaler.com>.
- [30] Packeteer, Inc. Application Traffic Management System. <http://www.packeteer.com/>.
- [31] A. K. Parekh and R. G. Gallager. A Generalized Processor Sharing Approach to Flow Control in Integrated Services Networks—the Single Node Case. *IEEE/ACM Transactions on Networking*, June 1993.
- [32] RUBiS Rice University Bidding System. <http://rubis.objectweb.org>.
- [33] K. Shen, H. Tang, T. Yang, and L. Chu. Integrated Resource Management for Cluster-based Internet Services. In *Proc. of the 5th USENIX Symposium on Operating Systems Design and Implementation*, Boston, MA, Dec. 2002.
- [34] K. Shen, T. Yang, L. Chu, J. Holliday, D. Kuschner, and H. Zhu. Neptune: Scalable Replication Management and Programming Support for Cluster-based Network Services. In *Proceedings of the 3rd USENIX Symposium on Internet Technologies and Systems*, San Francisco, CA, Mar 2001.
- [35] Teoma. Internet Search Engine. <http://www.teoma.com>.
- [36] Apache Tomcat Server. <http://jakarta.apache.org/tomcat>.
- [37] T. Voigt, R. Tewari, D. Freimuth, and A. Mehra. Kernel Mechanisms for Service Differentiation in Overloaded Web Servers. In *Proceedings of the 2001 USENIX Annual Technical Conference*, Boston, Massachusetts, June 2001.
- [38] WebSphere Extended Deployment. <http://www.ibm.com/software/websphere/>.
- [39] M. Welsh and D. Culler. Overload Management as a Fundamental Service Design Primitive, Saint-Emilion, France Sept 2002.
- [40] M. Welsh, D. Culler, and E. Brewer. SEDA: An Architecture for Well-Conditioned, Scalable Internet Services. In *Proceedings of the Eighteenth ACM Symposium on Operating Systems Principles*, Banff, Canada, Oct 2001.
- [41] H. Zhu, H. Tang, and T. Yang. Demand-driven Service Differentiation in Cluster-based Network Servers. In *Proceedings of the IEEE INFOCOM*, Anchorage, Alaska, April 2001.