Size-based Scheduling to Improve Web Performance

Mor Harchol-Balter* Nikhil Bansal Bianca Schroeder Mukesh Agrawal

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213
⟨harchol,nikhil,bianca,mukesh⟩@cs.cmu.edu

Abstract

Is it possible to reduce the expected response time of *every* request at a web server, simply by changing the order in which we schedule the requests? That is the question we ask in this paper.

This paper proposes a method for improving the performance of web servers servicing static HTTP requests. The idea is to give preference to those requests which are short, or have small remaining processing requirements, in accordance with the SRPT (Shortest Remaining Processing Time) scheduling policy.

The implementation is at the kernel level and involves controlling the order in which socket buffers are drained into the network.

Experiments are executed both in a LAN and a WAN environment. We use the Linux operating system and the Apache and Flash web servers.

Results indicate that SRPT-based scheduling of connections yields significant reductions in delay at the web server. These result in a substantial reduction in mean response time, mean slowdown, and variance in response time for both the LAN and WAN environments.

Significantly, and counter to intuition, the *large requests* are only negligibly penalized or not at all penalized as a result of SRPT-based scheduling.

Keywords: scheduling, system performance and design, web servers, computer networks, conservation laws, SRPT.

1 Introduction

A client accessing a busy web server can expect a long wait. This delay is comprised of several components: the propagation delay and transmission delay on the path between the client and the server; delays due to queueing at routers; delays caused by TCP due to loss, congestion, and slow start; and finally the delay at the server itself. The aggregate of these delays, i.e. the time from when the client makes a request until the entire file arrives is defined to be the *response time* of the request.

In this paper we focus on what we can do to improve the delay at the server. Research has shown that in situations where the server is receiving a high rate of requests, the delays at the server make up a significant portion of the response time [9], [8], [30]. More specifically, [9], [8] find that even if the network load is high, the delays at a busy web server can be responsible for more than 80% of the overall response time of a small file, and for 50% of the overall response time of a medium size file.

Measurements [29] suggest that the request stream at most web servers is dominated by *static* requests, of the form "Get me a file." The question of how to service static requests *quickly* is the focus of many companies *e.g.*, Akamai Technologies, and much ongoing research. This paper will focus on *static* requests only.

Our idea is simple. For static requests, the *size of the request* (i.e. the time required to service the request) is well-approximated by the size of the file, which is well-known to the server. Thus far, almost no companies or researchers have made use of this information. Traditionally, requests at a web server are scheduled independently of their size. The requests are time-shared, with each request receiving a *fair share* of the web server resources. We call this FAIR scheduling. We propose, instead, *unfair scheduling*, in which priority is given to *short* requests, or those requests which have *short remaining time*, in accordance with the well-known scheduling algorithm Shortest-Remaining-Processing-Time-first (SRPT). The expectation is that using SRPT scheduling of requests at the server will reduce the queueing time at the server.

It is well-known from queueing theory that SRPT

*This research is funded by Cisco Systems via a grant from the Pittsburgh Digital Greenhouse 00-1 and by NSF-ITR 99-167 ANI-0081396. Equipment was also provided by the Parallel Data Lab.

scheduling minimizes queueing time, [36]. Allowing short jobs to preempt long jobs is desirable because forcing long jobs to wait behind short jobs results in much lower mean response time than the situation where short jobs to wait behind long jobs. Despite the obvious advantages of SRPT scheduling with respect to mean response time, applications have shied away from using this policy for fear that SRPT "starves" big requests [11, 38, 39, 37]. A primary goal of this paper will be to investigate whether this fear is valid in the case of web servers serving typical web workloads.

It is not immediately clear what SRPT means in the context of a web server. A web server is not a single-resource system. It is not obvious *which* of the web server's resources need to be scheduled. As one would expect, it turns out that scheduling is only important at the *bottleneck resource*. Frequently this bottleneck resource is the bandwidth on the access link out of the web server. "On a site consisting primarily of *static content*, *network bandwidth* is the most likely source of a performance bottleneck. Even a fairly modest server can completely saturate a T3 connection or 100Mbps Fast Ethernet connection."[27] (also corroborated by [14], [4]). There's another reason why the bottleneck resource tends to be the bandwidth on the access link out of the web site: Access links to web sites (T3, OC3, etc.) cost thousands of dollars per month, whereas CPU is cheap in comparison. Likewise disk utilization remains low since most files end up in the cache. It is important to note that although we concentrate on the case where the network bandwidth is the bottleneck resource, all the ideas in this paper can also be applied to the case where the CPU is the bottleneck — in which case SRPT scheduling is applied to the CPU.

Since the network is the bottleneck resource, we try to apply the SRPT algorithm at the level of the network. Our idea is to control the order in which the server's socket buffers are drained. Recall that for each (non-persistent) request a connection is established between the client and the web server. Corresponding to each connection, there is a socket buffer on the web server end into which the web server writes the contents of the requested file. Traditionally, the different socket buffers are drained in Round-Robin Order, each getting a fair share of the bandwidth of the outgoing link. We instead propose to give priority to those sockets corresponding to connections for small file requests or where the *remaining data* required by the request is small. Throughout, we use the Linux OS.

The implementation of SRPT scheduling was more involved than we initially anticipated. The implementation details are described in Section 3.

The goal of this paper is to compare FAIR scheduling with SRPT scheduling. These are defined as follows:

FAIR scheduling This uses standard Linux (fair-share draining of socket buffers) with an unmodified web server.

SRPT scheduling This uses modified Linux (SRPT-based draining of socket buffers) with the web server modified only to update socket priorities.

Experiments are executed first in a LAN, so as to isolate the reduction in queueing time at the server. Response time in a LAN is dominated by queueing delay at the server and TCP effects. Experiments are next repeated in a WAN environment. The WAN allows us to incorporate the effects of propagation delay in understanding more fully the client experience.

In the LAN setting, we experiment with two different web servers: the common Apache server [21], and the Flash web server [33] which is known for speed. Our clients use a request sequence taken from a web trace. Experiments are also repeated using requests generated by a web workload generator (see Section 4.1.2). This request sequence is controlled so that the same experiment can be repeated at many different server loads. The *server load* is the load at the bottleneck device – in this case the network link out of the web server. The load thus represents the fraction of bandwidth used on the network link out of the web server (for example if the requests require 8Mb/sec of bandwidth, and the available bandwidth on the link if 10Mb/sec, then the network load is 0.8).

The results under the Apache web server are quite similar to those under the Flash web server. So as not to overwhelm the reader we typically only show the results in the case of Apache. The Flash results, which are quite similar, are in the associated technical report [26].

We obtain the following results in a LAN:

- SRPT-based scheduling decreases mean response time in a LAN by a factor of 3 – 8 for loads greater than 0.5.
- SRPT-based scheduling helps small requests a lot, while negligibly penalizing large requests. Under a load of 0.8, 80% of the requests improve by a factor of 10 under SRPT-based scheduling. Only the largest 0.1% of requests suffer an increase in mean response time under SRPT-based scheduling (by a factor of only 1.2).

- The variance in the response time is far lower under SRPT as compared with FAIR, in fact two orders of magnitude lower for most requests.

- *There is no negative effect on network throughput or CPU utilization from using SRPT as compared with FAIR.*

Next we consider a WAN environment. Again we use the Apache web server and again run at different loads.

We obtain the following results in a WAN:

- The improvement in mean response time of SRPT over FAIR under a server load of 0.9 ranged from a factor of 8 (for clients with a Round-trip-time of 100 ms) to a factor of 20 (for clients with an RTT of 20ms). On the other hand there was hardly any improvement in SRPT over FAIR for a server load of 0.5.

- The improvement of SRPT over FAIR in a WAN can actually be greater than in a LAN, for the case of high load at the server.

- Unfairness to large requests is nonexistent in a WAN setting. All request sizes have higher mean response time under FAIR in a WAN environment. We provide theoretical justification for this highly counter-intuitive result in Section 6.

The poor performance of FAIR scheduling throughout encourages us to consider several enhancements to FAIR involving modifications to the Linux kernel. Some of these modifications have been suggested by previous literature and some are new. We find that while some enhancements help somewhat, they don't improve the performance of FAIR to anywhere near the performance of SRPT, in either the LAN setting or the WAN setting.

It is important to realize that this paper is a prototype to illustrate the power of using SRPT-based scheduling. In Section 9, we elaborate on broader applications of SRPT-based scheduling, including its application to other resources, and to non-static requests. We also discuss SRPT applied to web server farms and Internet routers.

2 Previous Work

There has been much literature devoted to improving the response time of web requests. Some of this literature focuses on reducing *network latency*, e.g. by caching requests ([22], [13], [12]) or improving the HTTP protocol ([20], [32]). Other literature works on reducing the *delays at a server*, e.g. by building more efficient HTTP servers ([21], [33]) or improving the server's OS ([19], [5], [28], [31]). Recent studies show that *delays at the server* make up a significant portion of the response time [9], [8]. Our work focuses on reducing delay at the server by using size-based connection scheduling.

In the remainder of this section we discuss only work on *priority-based* or *size-based* scheduling of requests. We first discuss related implementation work and then discuss relevant theoretical results.

Almeida et. al. [1] use both a user-level approach and a kernel-level implementation to prioritizing HTTP requests at a web server. The *user-level* approach in [1] involves modifying the Apache web server to include a Scheduler process which determines the order in which requests are fed to the web server. This modification is all in the application level and therefore does not have any control over what the OS does when servicing the requests. The *kernel-level* approach in [1] simply involves setting the priority of the process which handles a request in accordance with the priority of the request. Observe that setting the priority of a process only allows very coarse-grained control over the scheduling of the process, as pointed out in the paper. The user-level and kernel-level approaches in this paper are good starting points, but the results show that more fine-grained implementation work is needed. For example, in their experiments, the high-priority requests only benefit by up to 20% and the low priority requests suffer by up to 200%.

Another attempt at priority scheduling of HTTP requests which deals specifically with SRPT scheduling at web servers is that of Crovella et. al. [16]. This implementation does not involve any modification of the kernel. The authors experiment with connection scheduling at the *application level* only. They design a specialized Web server which allows them to control the order in which `read()` and `write()` calls are made, but does not allow any control over the low-level scheduling which occurs inside the kernel, below the application layer (*e.g.*, control over the order in which socket buffers are drained). Via the experimental Web server, the authors are able to improve mean response time by a factor of up to 4, for some ranges of load, but the improvement comes at a price: *a drop in throughput by a factor of almost 2*. The explanation, which the authors offer repeatedly, is that scheduling at the application level does not provide fine enough control over the order in which packets enter the network. In order to obtain enough con-

trol over scheduling, the authors are forced to limit the throughput of requests. This will not be a problem in our paper. Since the scheduling is done at the kernel, we have absolute control over packets entering the network. Our performance improvements are greater than those in [16] and do not come at the cost of any decrease in throughput.

The papers above offer coarser-grained implementations for priority scheduling of connections. Very recently, many operating system enhancements have appeared which allow for finer-grained implementations of priority scheduling [23, 34, 3, 2].

Several papers have considered the idea of SRPT scheduling in theory.

Bender et. al. [11] consider size-based scheduling in web servers. The authors reject the idea of using SRPT scheduling because they prove that SRPT will cause large files to have an arbitrarily high *max slowdown*. However, that paper assumes a worst-case adversarial arrival sequence of web requests. The paper goes on to propose other algorithms, including a theoretical algorithm which does well with respect to max slowdown and mean slowdown.

Roberts and Massoulie [35] consider bandwidth sharing on a link. They suggest that SRPT scheduling may be beneficial in the case of heavy-tailed (Pareto) flow sizes.

Lastly, Bansal and Harchol-Balter [7] investigate SRPT scheduling analytically for an M/G/1/SRPT queue (Poisson arrivals and general service times). We will discuss these theoretical results in Section 6.

3 Implementation of SRPT

In Section 3.1 we explain how socket draining works in standard Linux. In Section 3.2 we describe how to achieve priority queueing in Linux (versions 2.2 and above). Section 3.2.1 describes some pitfalls we came across in trying to modify the kernel.

Section 3.3 describes the implementation end at the web server and also deals with the algorithmic issues such as how to choose good *priority classes* and the setting and updating of priorities. Here too there are some subtleties. For example, one problem with size-based queueing is that for small requests, a large portion of the time to service the request is spent *before* the size of the request is even known. Section 3.3.3 describes our solution to this problem.

3.1 Default Linux configuration

Figure 1 shows data flow in standard Linux.

There is a socket buffer corresponding to each connection. Data streaming into each socket buffer is encapsulated into packets which obtain TCP headers and IP headers. Throughout this processing, the packet streams corresponding to each connection is kept separate. Finally, there is a *single*[1] "priority queue" (*transmit queue*), into which *all* streams feed. In the abstract, these flows take equal turns feeding into the priority queue. Although the Linux kernel does not explicitly enforce fairness, we find that in practice, TCP governs the flows so that they share fairly on short time scales.

This single "priority queue," can get as long as 100 packets. Packets leaving this queue drain into a short Ethernet card queue and out to the network.

3.2 How to achieve priority queueing in Linux

To implement SRPT we need more priority levels. To do this, we first build the Linux kernel with support for the user/kernel Netlink Socket, QOS and Fair Queueing, and the Prio Pseudoscheduler. Then we use the `tc`[3] user space tool to switch the device queue from the default 3-band queue to the 16-band prio queue. Further information about the support for differentiated services and various queueing policies in Linux can be found in [23, 34, 3, 2].

Figure 2 shows the flow of data in Linux after the above modification: The processing is the same until the packets reach the priority queue. Instead of a single priority queue (transmit queue), there are 16 priority queues. These are called bands and they range in number from 0 to 15, where band 15 has lowest priority and band 0 has highest priority. All the connections of priority i feed fairly into the ith priority queue. The priority queues then feed in a prioritized fashion into the Ethernet Card queue. Priority queue i is only allowed to flow if priority queues 0 through $i-1$ are all empty.

3.2.1 Why priority queueing may not work for you: Lessons learned

The following is a detailed account of some of our experiences in implementing priorities and some lessons learned. We hope that this will be useful to other implementors.

After making the above described modifications to Linux, we tried a simple experiment: We opened two TCP connections and flooded both connections with

[1] The queue actually consists of 3 priority queues, a.k.a. bands. By default, however, all packets are queued to the same band.

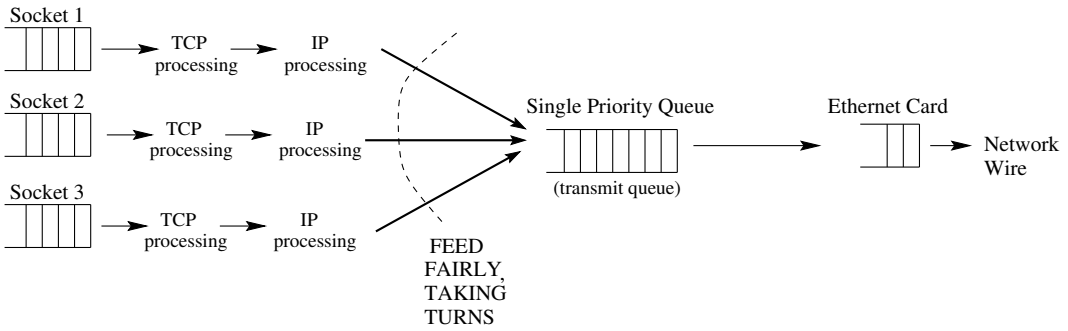

Figure 1: *Data flow in Standard Linux. The important thing to observe is that there is a* single *priority queue into which all connections drain fairly.*

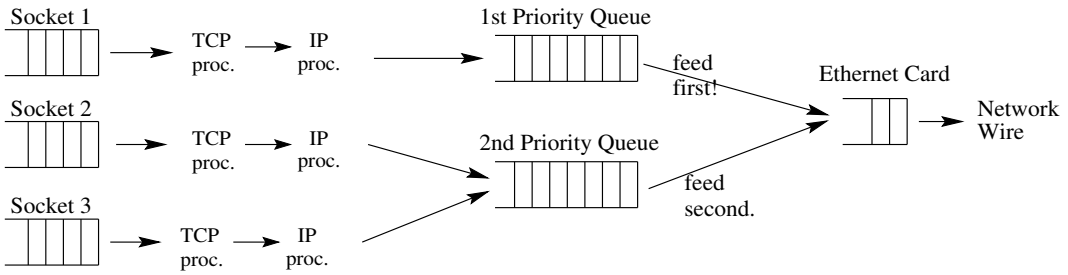

Figure 2: *Flow of data in Linux with priority queueing. It is important to observe that there are several priority queues, and queue i is serviced only if all of queues 0 through $i-1$ are empty.*

data. Specifically, we repeatedly wrote 1K of data into each socket within an infinite write loop (note that the writes are non-blocking). We gave high priority to connection 1 and low priority to connection 2. We expected that the packets on the wire would all be connection 1 packets. In truth, however, only 51% of the packets on the wire were connection 1 packets and 49% were connection 2 packets.

We repeated the above experiment, but this time with two UDP connections and saw a $60\%/40\%$ split between connection 1 packets and connection 2 packets.

After various other such frustrating experiments, and much discussion with others experiencing similar problems, we finally lucked upon a "fix": We found that if we increased the size of the sender's send socket buffer by a factor of 3, and also increased the size of the receive socket buffer a little, we were able to achieve *full* priority scheduling ($100\%/0\%$ ratio of connection 1 packets to connection 2 packets). The increase of the receive socket buffer wasn't always necessary.

The above "fix" did not seem satisfactory: It didn't seem that it would always be practical to increase the size of these socket buffers and it seemed to give SRPT a possible unfair advantage over FAIR.

We continued experimenting, but this time we tried running more than 2 simultaneous connections. We observed that when we increased the number of connections from 2 to 10, we always achieved the desired $100\%/0\%$ ratio, without requiring an increase in the send socket buffer.

In the experiments in this paper, we have hundreds of simultaneous connections. It turns out that in this case the `tc`-based modifications to the kernel allow us to achieve full priority queueing without modifying the size of the sender's send socket buffer or the receive socket buffer.

3.3 Modifications to web server and algorithmic issues in approximating SRPT

The modified Linux kernel provides mechanisms for prioritized queueing. In our implementation, the Apache web server uses these mechanisms to implement the SRPT-based scheduling policy. Specifically, after determining the size of a request, Apache sets the priority of the corresponding socket by calling `setsockopt`. As Apache sends the file, the remaining size of the request decreases. When the remaining size falls below the threshold for the current priority

class, Apache updates the socket priority with another call to `setsockopt`.

3.3.1 Implementation Design Choices

Our implementation places the responsibility for prioritizing connections on the web server code. There are two potential problems with this approach. These are the overhead of the system calls to modify priorities, and the need to modify server code.

The issue of system call overhead is mitigated by the limited number of `setsockopt` calls which must be made. Typically only one call is made. Even in the worst case, we make only as many `setsockopt` calls as there are priority classes (6 in our experiments).

A clean way to handle the changing of priorities totally within the kernel would be to enhance the `sendfile` system call to set priorities based on the remaining file size. We do not pursue this approach here as neither Apache nor Flash uses `sendfile`.

3.3.2 Size cutoffs

SRPT assumes infinite precision in ranking the remaining processing requirements of requests. In practice, we are limited to a small fixed number of priority bands (16).

Based on experimentation, we have come up with some *rules-of-thumb* for partitioning the requests into priority classes which apply to the heavy-tailed web workloads. The reader not familiar with heavy-tailed workloads will benefit by first reading Section 6.

Denoting the cutoffs by $x_1 < x_2 < \ldots < x_n$:

- The lowest size cutoff x_1 should be such that about 50% of requests have size smaller than x_1. The requests comprise so little total load in a heavy-tailed distribution that there's no point in separating them.

- The highest cutoff x_n needs to be low enough that the largest (approx.) .5% – 1% of the requests have size $> x_n$. This is necessary to prevent the largest requests from starving.

- The middle cutoffs are far less important. Anything remotely close to a logarithmic spacing works well.

In the experiments throughout this paper, we use only 6 priority classes to approximate SRPT. Using more improved performance only slightly.

3.3.3 Priority to SYNACKs

It seems that one should now be able to execute prioritized size-based scheduling by simply assigning sockets for larger requests to priority queues with lower priority (these have high band numbers) and updating priorities as described above.

It turns out that this is not sufficient however to get good performance improvement. The reason is somewhat subtle. An important component of the response time is the connection startup time. In SRPT scheduling, we are careful to separate the small requests from the large ones. However during connection startup, we don't yet know whether the request will be large or small. The packets sent during the connection startup might therefore end up waiting in long queues, making connection startup very costly. For short requests, a long startup time is especially detrimental to response time. It is therefore important that the SYNACK be isolated from other traffic. Linux sends SYNACKs, to priority band 0. It is important when assigning priority bands to requests that we:

1. Never assign any sockets to priority band 0.

2. Make all priority band assignments to bands of *lower* priority than band 0, so that SYNACKs always have highest priority.

Observe that giving highest priority to the SYNACKs does not negatively impact the performance of requests since the SYNACKs themselves make up only a negligible fraction of the total load.

Giving priority to SYNACKs is important in SRPT because without it the benefit that SRPT gives to small requests is not noticeable. Later in the paper (Section 4.3) we consider whether the FAIR policy might also benefit by giving priority to SYNACKs. We find that there is some improvement in FAIR as well, however that improvement is *nowhere close* to the improvement offered by SRPT.

3.3.4 The final algorithm

Our SRPT-like algorithm is thus as follows:

1. When a request arrives, it is given a socket with priority 0 (highest priority). This allows SYNACKs to travel quickly. This was explained in Section 3.3.3.

2. After the request size is determined (by looking at the URL of the file requested), the priority of the socket corresponding to the request is reset based on the size of the request, as shown in the table below.

Priority	Size (Kbytes)
0 (highest)	-
1	$\leq$ 1K
2	1K - 2K
3	2K - 5K
4	5K-20K
5	20K - 50K
6 (lowest)	$>$ 50K

3. As the remaining size of the request diminishes, the priority of the socket is dynamically updated to reflect the remaining size of the request.

4 LAN setup and experimental results

In Section 4.1 we describe the experimental setup and workload for the LAN experiments. Section 4.2 illustrates the results of the LAN experiments. Section 4.2.1 illustrates a simplification of the SRPT idea which involves only two priorities and yet still yields quite good performance. Section 4.3 proposes some enhancements to improve the performance of FAIR scheduling and describes the results of these enhancements.

4.1 Experimental Setup (LAN)

4.1.1 Architecture

Our experimental architecture involves two machines each with an Intel Pentium III 700 MHz processor and 256 MB RAM, running Linux 2.2.16, and connected by a 10Mb/sec full-duplex Ethernet connection. The Apache web server is running on one of the machines. The other machine (referred to as the "client machine") hosts the many (simulated) client entities which send requests to the web server. In the experiments in this section, about 200 client entities were needed to send the request sequence to the web server.

In Section 7, we repeat these experiments in a 100Mb/sec environment.

4.1.2 Workload

The clients' requests are generated either via a *web workload generator* (we use a modification of `Surge` [10]) or via a *trace*. Throughout this paper, all results shown are for a trace-based workload. We have included in the associated technical report [26] the same set of results for the `Surge` workload. The `Surge` workload is intended to be characteristic of web workloads. The `Surge` workload is a different request sequence than our trace workload – nevertheless, results are substantially similar under the `Surge` and trace-based workloads.

4.1.3 Traces

The trace-based workload consists of a 1-day trace from the Soccer World Cup 1998, from the Internet Traffic Archive [24]. The 1-day trace contains 4.5 million HTTP requests, virtually all of which are *static*.

An entry in the trace includes: (1) the time the request was received at the server, (2) the size of the request in bytes, (3) the GET line of the request, (4) the error code, as well as other information. In our experiments, we use the trace to specify *the time* the client makes the request and the *size in bytes* of the request.

The entire 1 day trace contains requests for approximately 5000 different files. Given the mean file size of 5K, it is clear why all files fit within main memory. This explains why the disk is not a bottleneck.

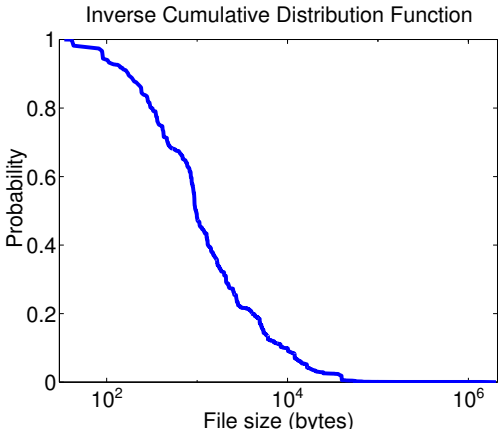

Figure 3: *Inverse Cumulative Distribution Function, $\overline{F}(x)$, for the trace-based workload.* $\overline{F}(x) = \Pr\{\textit{Job size } > x\}$

Each experiment was run using a busy hour of the trace (10:00 a.m. to 11:00 a.m.). This hour consisted of about 1 million requests, during which over a *thousand* files are requested.

Some additional statistics about our trace workload: The minimum size file requested is a 41 byte file. The maximum size file requested is about 2 MB. The distribution of the file sizes requested fits a heavy-tailed truncated Pareto distribution (with α-parameter ≈ 1.2). The largest $< 3\%$ of the requests make up $> 50\%$ of the total load, exhibiting a strong heavy-

tailed property. 50% of files have size less than 1K bytes. 90% of files have size less than 9.3K bytes.

The distribution of file sizes is shown in Figure 3.

4.1.4 Generating requests at client machines

In our experiments, we use `sclient` [6] for creating connections at the client machines. The original version of `sclient` makes requests for a certain file in periodic intervals. We modify `sclient` to read in traces and make the requests according to the arrival times and file names given in the trace.

To create a particular load, say 0.8, we simply scale the interarrival times in the trace's request sequence until the average number of bytes requested per second is 8Mb/sec. We validate the load both analytically and via measurement.

4.1.5 Performance Metrics

For each experiment, we evaluate the following performance metrics:

- *Mean response time*. The response time of a request is the time from when the client submits the request until the client receives the last byte of the request.

- *Mean slowdown*. The slowdown metric attempts to capture the idea that clients are willing to tolerate long response times for large file requests and yet expect short response times for short requests. The slowdown of a request is therefore its response time divided by the time it would require if it were the sole request in the system. Slowdown is also commonly known as *normalized response time* and has been widely used [11, 35, 18, 25].

- *Mean response time as a function of request size*. This will indicate whether big requests are being treated *unfairly* under SRPT as compared with FAIR-share scheduling.

4.2 Main Experimental results (LAN)

Before presenting the results of our experiments, we make some important comments.

- In all of our experiments the network was the bottleneck resource. CPU utilization during our experiments ranged from 1% in the case of low load to 5% in the case of high load.

- The measured throughput and bandwidth utilization under the experiments with SRPT scheduling is *identical* to that under the same experiments with FAIR scheduling. The same exact set of requests complete under SRPT scheduling and under FAIR scheduling.

- There is no additional CPU overhead involved in SRPT scheduling as compared with FAIR scheduling. Recall that the overhead due to updating priorities of sockets is insignificant, given the small number of priority classes that we use.

Figure 4 shows the mean response time under SRPT scheduling as compared with the traditional FAIR scheduling as a function of load. For lower loads the mean response times are similar under the two scheduling policies. However for loads > 0.5, the mean response time is a factor of 3 – 8 lower under SRPT scheduling.

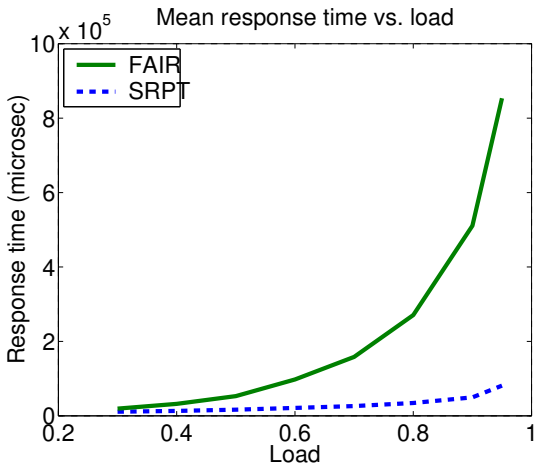

Figure 4: *Mean response time under SRPT versus FAIR as a function of system load, under trace-based workload, in LAN environment.*

The results are even more dramatic for mean slowdown. For loads 0.5, the mean slowdown improves by a factor of 4 under SRPT over FAIR. Under a load of 0.9, mean slowdown improves by a factor of 16.

The important question is whether the significant improvements in mean response time come at the price of significant unfairness to large requests. Figure 5 shows the mean response time as a function of request size, in the case where the load is 0.6, 0.8, and 0.9. In the left column of Figure 5, request sizes have been grouped into 60 bins, and the mean response time for each bin is shown in the graph. The 60 bins are determined so that each bin spans an interval $[x, 1.2x]$. It is important to note that the last bin ac-

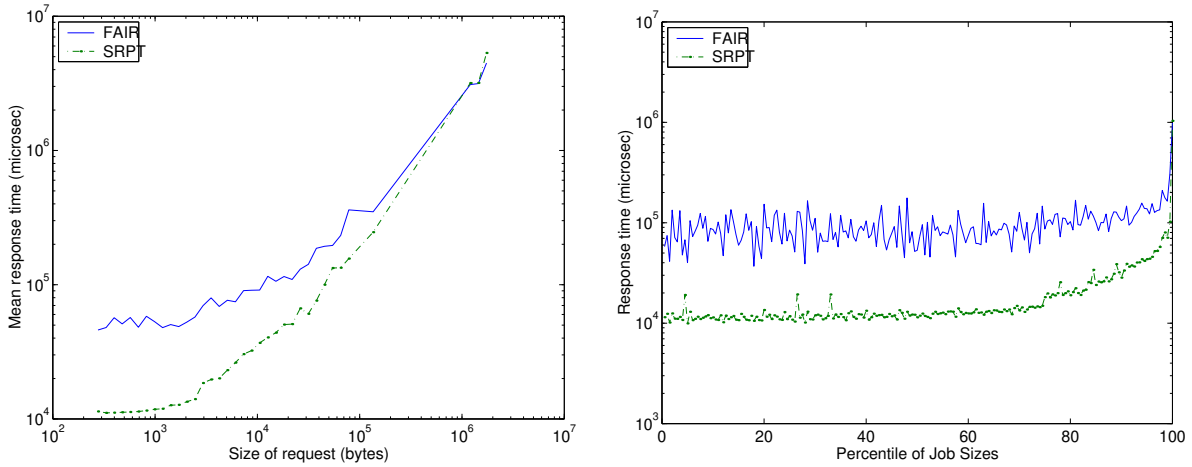

(a) load = .6

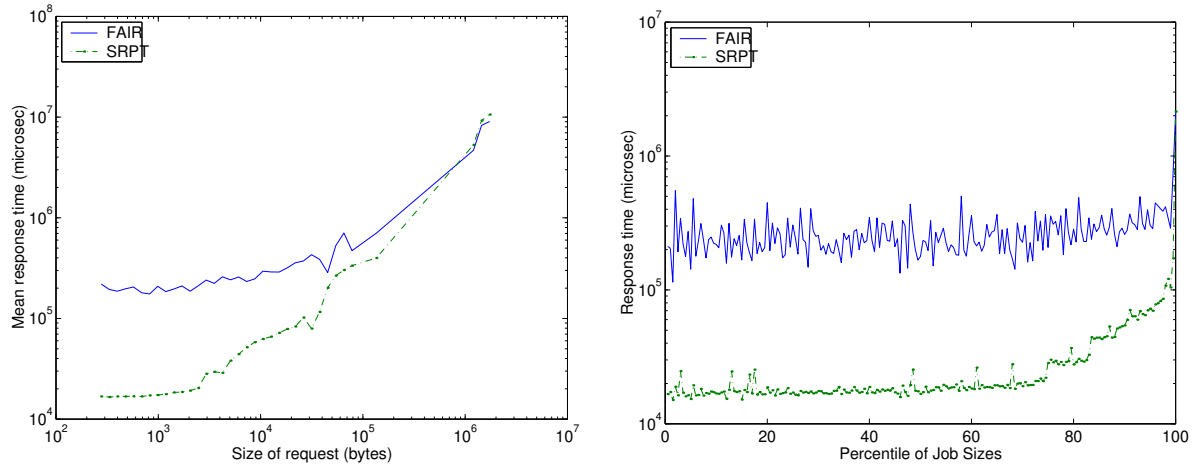

(b) load = .8

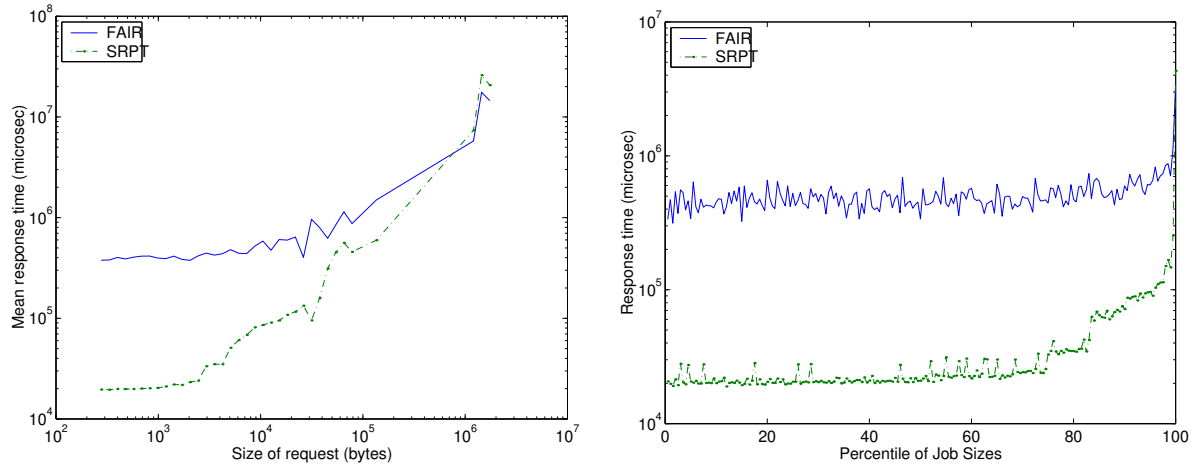

(c) load = .9

Figure 5: *Mean response time as a function of request size under trace-based workload, shown for a range of system loads, in a LAN. The left column shows the mean response time as a function of request size. The right column shows the mean response time as a function of the percentile of the request size distribution.*

tually contains only requests for the very biggest file. Observe that small requests perform far better under SRPT scheduling as compared with FAIR scheduling, while large requests, those $>$ 1 MB, perform only negligibly worse under SRPT as compared with FAIR scheduling. For example, under load of 0.8 (see Figure 5(b)) SRPT scheduling improves the mean response times of small requests by a factor of close to 10, while the mean response time for the largest size request only goes up by a factor of 1.2.

Note that the above plots give equal emphasis to small and large files. As requests for small files are much more frequent, these plots are not a good measure of the improvement offered by SRPT. To fairly assess the improvement, the right column of Figure 5, presents the mean response time as a function of the percentile of the request size distribution, in increments of half of one percent (i.e. 200 percentile buckets). From this graph, it is clear that at least 99.5% of the requests benefit under SRPT scheduling. In fact, the 80% smallest requests benefit by a factor of 10, and all requests outside of the top 1% benefit by a factor of > 5. For lower loads, the difference in mean response time between SRPT and FAIR scheduling decreases, and the unfairness to big requests becomes practically nonexistent. For higher loads, the difference in mean response time between SRPT and FAIR scheduling becomes greater, and the unfairness to big requests also increases. Even for the highest load tested though (.95), there are only 500 requests (out of the 1 million requests) which complete later under SRPT as compared with FAIR. These requests are so large however, that the effect on their slowdown is negligible.

Figure 6 shows the variance in response time for each request size as a function of the percentile of the request size distribution, for load equal to 0.8. The improvement under SRPT with respect to variance in response time is 2 – 4 orders of magnitude for the 99.5% smallest files. The improvement with respect to the squared coefficient of variation (variance/mean^2) is about 30.

4.2.1 SRPT with only 2 priorities (LAN)

To evaluate the importance of choosing precise cutoffs, we evaluate SRPT with only two priority classes. We define *small* requests as the smallest 50% of requests and *large* requests as the largest 50% of requests (note, this is not the same thing as equalizing load) The cutoff falls at 1K. We find that this simple algorithm results in a factor of 2.5 improvement in mean response time and a factor of 5 improvement in mean slowdown over FAIR. We also find that *all requests*, big and small, have lower expected response times under SRPT than under FAIR using this simple algorithm.

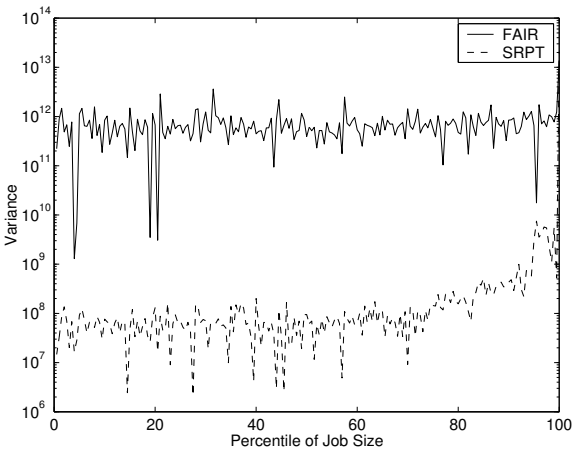

Figure 6: *Variance in response time as a function of the percentile of the request size distribution for SRPT as compared with FAIR, under trace-based workload with load* = 0.8*, in a LAN.*

4.3 Enhancements to FAIR

At this point it is natural to wonder why the FAIR policy performs so poorly as compared with SRPT. The obvious reason is that the time-sharing behavior of FAIR causes short requests to be delayed in part by long requests, whereas SRPT allows short requests to jump ahead of long requests.

However there is a second reason: In biasing towards short requests, SRPT also minimizes the number of requests in the system at any given time. The number of ongoing requests under FAIR is often much higher than that under SRPT. To understand the effect of this, we instrumented the web server's kernel to provide statistics including: the occupancy of the SYN and ACK (listen) queues, the number of incoming SYNs dropped due to the SYN queue being full, the number of times a client's acknowledgement of a SYNACK was discarded due to the ACK queue being full, and the number of outgoing SYNACKs dropped inside the kernel.

Under the newly instrumented kernel, we reran all the LAN experiments. Below we discuss our findings just for the case of load 0.9. For this case the response time was 452ms under FAIR and 38ms under SRPT, a dramatic difference.

Effect of length of SYN and ACK queues

Our first hypothesis was that perhaps the larger number of concurrent requests in the system under FAIR was causing the SYN queue or ACK queue to fill to capacity under FAIR, forcing client requests to timeout and have to be retransmitted. This would explain the occasional high response times under FAIR.

This hypothesis turned out to be *false*. We found that neither the SYN queue nor the ACK queue ever fills to capacity in the LAN experiments. Increasing the length limit on the SYN queue and ACK queue has no effect. (In these experiments the length of the SYN queue and the ACK queue are each set to the default value of 128.)

Effect of length of transmit queue

What did turn out to be an issue was not the SYN queue, but rather the transmit queue (see Figure 1). We found that the large buildup in connections at the server under FAIR, coupled with the fact that all packets drain into a single priority queue (transmit queue) under FAIR, was causing some packets to be dropped by the server at the transmit queue under high loads under FAIR, but *not* under SRPT. In fact, under load 0.9, our measurements indicate that (5%-10%) of connections experience some loss at the transmit queue of the server under FAIR.

We experimented with increasing the length limit for the transmit queue from 100 to 500, and then to 700. This helped a little — reducing mean response time from 452ms to 342ms. Still, performance was nowhere near that of SRPT – 38ms.

We also experimented with reducing the length of the transmit queue, but this only increased the mean response time, as expected, since loss at the transmit queue went up.

Effect of giving priority to SYNACKs

Recall (Section 3.3.3) that our implementation of SRPT includes giving priority to SYNACKS. We wondered if the same idea might help FAIR as well. We found that when SYNACKs were given priority the mean response time dropped from 452ms to about 300ms under FAIR – a decent improvement, but still nowhere close to the performance of SRPT – 38ms.

Lastly we combined all 3 of the above potential enhancements to FAIR. The performance remained at about 300ms for FAIR.

5 WAN setup and Experimental Results

In Section 5.1 we describe the setup for the WAN experiments. In Section 5.2 we describe the results of the WAN experiments.

5.1 Experimental setup (WAN)

We use the same server machine as for the LAN experiment. This time we have 6 client machines located throughout the Internet. The client machines each again host a couple hundred (simulated) client entities which generate the requests in the same way as before based on a trace. Again each experiment again spans 1 hour (about 1 million requests).

The client machines are located at various distances from the server (indicated by round trip times, RTT) and have varying available bandwidth, as shown in the table below.

Unfortunately, we were only able to get accounts for Internet2 machines (schools and some research labs). Companies were unable to offer us accounts on their machines for security reasons. The limitation in exploring only an Internet2 network is that loss and congestion may be unrealistically low.

Location	Avg. RTT	Avail. Bndwth
IBM, New York	20ms	8Mbps
Univ. Berkeley	55ms	3Mbps
UK	90-100ms	1Mbps
Univ. Virginia	25ms	2Mbps
Univ. Michigan	20ms	5Mbps
Boston Univ.	22ms	1Mbps

5.2 Experimental results for the WAN setup

Figure 7 shows the mean response time (in ms) as a function of load for each of the six hosts. This figure shows that the improvement in mean response time of SRPT over FAIR is a factor of 8–20 for high load (0.9) and only about 1.1 for lower load (0.5).

Figure 8(a) and 8(b) shows the mean response time of a request as function of the percentile the request size at a load of 0.8 for the hosts at IBM and UK respectively. It's not clear from looking at the figures whether there is any starvation. It turns out that *all* request sizes have higher mean response time under FAIR, as compared with SRPT. For the largest file, the mean response time is almost the same under SRPT and FAIR.

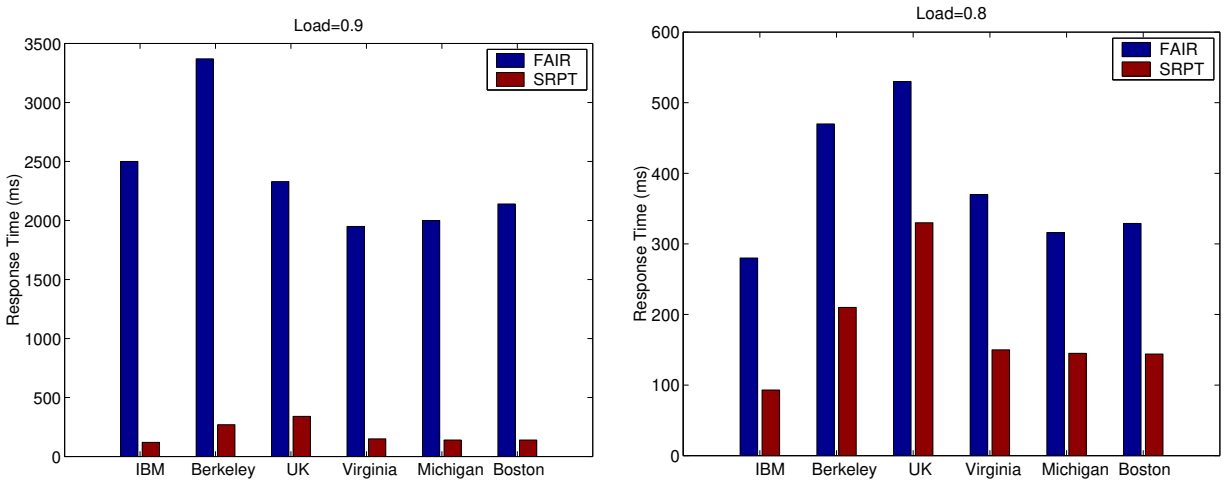

(a) load 0.9

(b) load 0.8

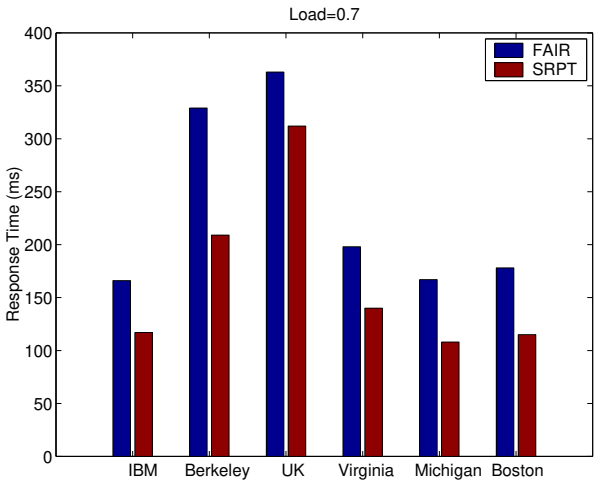

(c) load 0.7

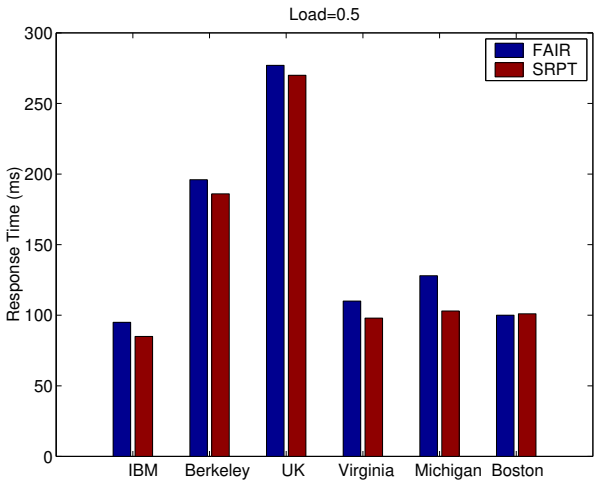

(d) load 0.5

Figure 7: *Mean response time under SRPT versus FAIR in a WAN under load (a) 0.9, (b) 0.8, (c) 0.7, and (d) 0.5.*

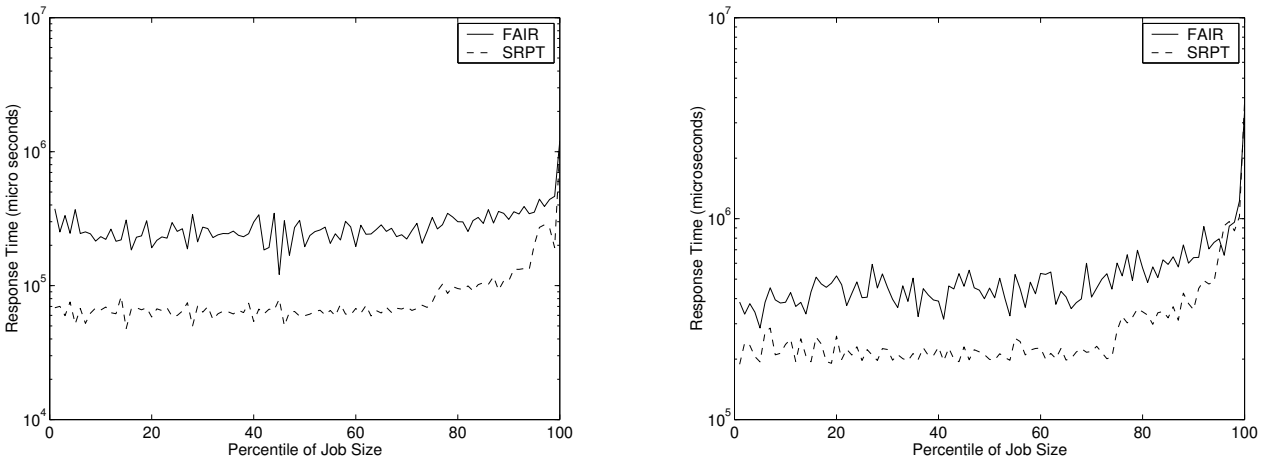

(a) IBM clients (b) UK clients

Figure 8: *Response time as a percentile of request size under SRPT scheduling versus traditional FAIR scheduling at load 0.8, measured for (a) the IBM host and (b) the UK host.*

We also measured the variance in response time in the WAN environment for a load of 0.8. While the variance for FAIR stayed the same under the LAN and WAN environments, the variance for SRPT increased somewhat in the WAN environment due to losses. Still, however, the variance in response time under SRPT remains over an order of magnitude below that in FAIR, for a load of 0.8.

We make the following observations:

Observation 1 *The improvement of SRPT over FAIR in mean response time is greater at higher loads.*

For example in Figure 7 the mean response time for the IBM host under SRPT improves over FAIR by a factor of 20 at load 0.9, a factor of 3 at load 0.8, a factor of 1.5 at load 0.7 and a factor of 1.1 at load 0.5. *Explanation:* We already saw in a LAN that under higher load, the difference between SRPT and FAIR is higher. This is coupled with the fact that at higher load the queueing delay at the server makes up a larger component of response time.

Observation 2 *The improvement of SRPT over FAIR is less for far away clients.*

For example in Figure 7, at load 0.8, the mean response time for the IBM host (RTT 20ms) improves by a factor of about 3 under SRPT over FAIR, whereas there is only a factor 1.6 improvement for the far away host at UK (RTT 90 ms). *Explanation:* The delays caused by *propagation* and *Internet congestion* mitigate the effect of the queueing delay on total response time.

Observation 3 *The improvement of SRPT over FAIR in a WAN environment can actually be greater than in a LAN environment, for the case of high load at the server.*

For example in Figure 7 the mean response time for the host at IBM is 2500 ms under FAIR scheduling, vs. 125 ms under SRPT scheduling, hence about 20 times better. However, in the LAN setup the mean response time improved by a factor of about 12 at load 0.9 (See Figure 4).

Explanation: This surprising observation is due to the fact that the WAN environment creates variance in the load, thus an average load of 0.9 translates to fluctuations ranging from 0.75 to 1.05, which means that there are moments of *transient overload*. Transient overload affects FAIR far worse than SRPT because the buildup in number of requests at the server during overload is so much greater under FAIR than under SRPT. Transient overload even occasionally results in a full SYN queue under FAIR, which means that SYN connection requests timeout ("SYN drops"), causing a 3-second delay. Recall that in the LAN environment where load could be better controlled, we never experienced SYN drops.

Observation 4 *The mean response times under SRPT are close to optimal even under high loads.*

For example, for the host at Berkeley (RTT 55ms), the mean response times are 186, 209, 210 and 270 ms at loads 0.5,0.7,0.8 and 0.9 respectively under SRPT scheduling. Observe that these are quite close to 170ms, which is the optimal mean response time

(i.e. when the load at the server is close to 0) for this host.

Observation 5 *While the penalty of SRPT to large requests is almost absent in the LAN setting (see Figure 5), we observe that it is even less of an issue in the WAN environment.*

As explained above now *all* request sizes have higher mean response time under FAIR, as compared with SRPT. *Explanation:* The reason is simply that the propagation delay in the case of a WAN mitigates the effect of the queueing delay (in particular the difference between the queueing delay under SRPT and that under FAIR).

6 How can *every* request prefer SRPT to FAIR in expectation? Theoretical Explanation

It has been suspected by many that SRPT is a very unfair scheduling policy for large requests. The above results have shown that this suspicion is false for web workloads. It is easy to see why SRPT should provide huge performance benefits for the small requests, which get priority over all other requests. In this section we describe briefly why the large requests also benefit under SRPT, *in the case of a heavy-tailed workload*.

In general a heavy-tailed distribution is one for which

$$\Pr\{X > x\} \sim x^{-\alpha},$$

where $0 < \alpha < 2$. A set of request sizes following a heavy-tailed distribution has some distinctive properties:

1. Infinite variance (and if $\alpha \leq 1$, infinite mean). In practice there is a finite maximum request size, which means that the moments are all finite, but still quite high.

2. The property that a tiny fraction (usually $< 1\%$) of the very longest requests comprise over half of the total load. We refer to this important property as the **heavy-tailed property**.

The lower the parameter α, the more variable the distribution, and the more pronounced is the heavy-tailed property, *i.e.* the smaller the fraction of long requests that comprise half the load.

Request sizes are well-known to follow a heavy-tailed distribution [15, 17]. Our traces also have strong heavy-tailed properties. (In our trace the largest $< 3\%$ of the requests make up $> 50\%$ of the total load.)

Consider a workload where request sizes exhibit the *heavy-tailed property*. Now consider a large request, in the 99%-tile of the request size distribution. This request will actually do much better under SRPT scheduling than under FAIR scheduling. The reason is that this big request only competes against 50% of the load under SRPT (the remaining 50% of the load is made up of requests in the top 1%-tile of the request size distribution) whereas it competes against 100% of the load under FAIR scheduling. The same argument could be made for a request in the 99.5%-tile of the request size distribution.

However, it is not obvious what happens to a request in the 100%-tile of the request size distribution (i.e. the largest possible request). It turns out that, provided the load is not too close to 1, the request in the 100%-tile will quickly see an idle period, during which it can run. As soon as the request gets a chance to run, it will quickly become a request in the 99.5%-tile, at which time it will clearly prefer SRPT. For a mathematical formalization of the above argument, in the case of an M/G/1 queue, we refer the reader to [7].

7 100 Mbps setup

7.1 LAN experiment

Throughout this paper during the LAN experiments, we have assumed the configuration described in Section 4.1 , which assumes a 10Mb/sec Ethernet leaving the server. For thoroughness, we now consider the case of a web server connected to a 100 Mb/sec Ethernet. Our new LAN configuration consists of 5 client machines, each running hundreds of (simulated) client entities. The 5 client machines are connected via a switch to a 100 Mbps full-duplex Ethernet which connects the switch to the server. At any moment there may be a thousand simultaneous ongoing connections. All machines are Intel Pentium III 700 MHz processors with 256 MB Ram, running Linux 2.2.16.

We use the same set of traces as described in Section 4.1.3 and the same method for generating requests described in Section 4.1.4. Again we experiment with load ranging from .3 to .95, which now translates to 30 Mb/sec up to 95 Mb/sec traffic through the network switch. The network bandwidth continues to be the bottleneck in all of these experiments. Results presented below are for the Flash web server.

Figure 9 presents the mean response time as a function of network load. We see that SRPT performs 5 times better than FAIR at a load of 0.8. This is com-

parable to the improvement achieved in the case of the 10 Mbit network.

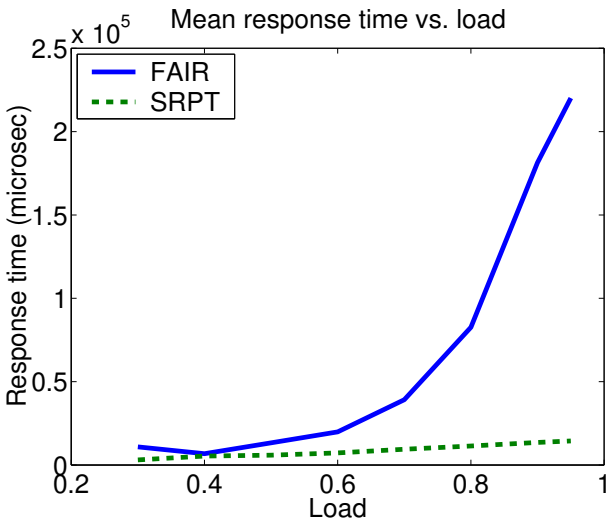

Figure 9: *Mean response time under SRPT versus FAIR as a function of system load, under trace-based workload, in LAN environment under 100Mb/sec connection.*

The breakdown of response time with respect to job size is also comparable. In the 100 Mb/sec experiment, *all jobs*, large and small, preferred SRPT scheduling in expectation.

In moving from a 10 Mbit network to a 100 Mbit network, observe that the arrival rate of jobs increases by a factor of 10, and the service requirement of jobs decreases by a factor of 10. Queueing theory thus tells us that we should expect mean response times to improve by a factor of 10. However, in comparing Figure 4 and Figure 9, we note that the improvement in practice for both FAIR and SRPT is not as high as predicted by theory (more like a factor of 4–6). We suspect that the performance limits of other system elements may be preventing the realization of the full 10 times improvement.

7.2 WAN experiment

We are unable to rerun our WAN experiment using a 100 Mb/sec Ethernet, because the client machines that we have access to do not have enough combined bandwidth to absorb 100 Mb/sec worth of data (we simply could not get access to enough far away machines). In practice, our server would be accessed by many client machines, each representing a single client, and these might easily fill a 100 Mb/sec link.

Although we can't assess the effect experimentally, we can approximate the effect of a 100 Mb/sec Ethernet connection in a wide-area setting. Recall that the 100 Mb/sec connection dropped the response time numbers for the LAN setting by at least a factor of 4 from those seen under the 10 Mb/sec connection. This brings the numbers into the range of 10 ms for low load and 100-200 ms for higher load. For a far away client, the average RTT might be as high as 100 ms. In that case, the difference in performance between SRPT and FAIR at the server might not be noticeable. For closer clients, where the average RTT is only 25ms, the decreased time at the server should be much more noticeable. Future trends indicate lower RTTs, due to faster backbone connections.

8 Conclusion

This paper demonstrates that the delay at a busy server can be greatly reduced by SRPT-based scheduling of requests at the server's outgoing link. We show further that the reduction in server delay often results in a reduction in the client-perceived response time.

Our SRPT-based scheduling algorithm reduces mean response time in a LAN setting significantly under high server loads, over the standard FAIR scheduling algorithm. In a WAN setting the improvement is similar for very high server loads, but is less significant at moderate loads.

Surprisingly, this improvement comes at no cost to large requests, which are hardly penalized, or not at all penalized. Furthermore these gains are achieved under no loss in byte throughput or request throughput.

9 Limitations and Future work

Our current setup involves only *static* requests. In future work we plan to expand our technology to schedule cgi-scripts and other *non*-static requests. Determining the size (processing requirement) of non-static requests is an important open problem, but companies are making excellent progress on predicting the size of dynamic requests. We propose additionally to deduce the size of a dynamic request as it runs. The request will initially be assigned high priority, but its priority will decrease as it runs.

Our current setup considers network bandwidth to be the bottleneck resource and does SRPT-based scheduling of that resource. In a different application (e.g. processing of cgi-scripts) where some other resource was the bottleneck (e.g., CPU), it might be desirable to implement SRPT-based scheduling of that resource.

Although we evaluate SRPT and FAIR across many server loads, we do not in this paper consider the case of overload. This is an extremely difficult problem both analytically and especially experimentally. Our preliminary results show that in the case of *transient overload* SRPT outperforms FAIR across a long list of metrics, including mean response time, throughput, server losses, etc.

More work is also needed in evaluating SRPT and FAIR under lossy and congested networks.

Our SRPT solution can also be applied to server farms. In this scenario the bottleneck moves from the outgoing link at each server to the access link for the server farm — thus scheduling needs to be applied at the access link. To achieve this, servers would mark packets to designate their priority. These priorities would be enforced by the router at the access link.

10 Acknowledgements

We would like to thank Mark Crovella for many productive discussions on the topic of scheduling in web servers and Linux flow control. Greg Kesden, Jan Harkes, and Srini Seshan also provided us with helpful information about Linux internals. We would also like to thank Alexey Kuznetsov, Steven Ives and Tim Canfield for helping us with the initial Linux Diffserv experimentation. Lastly, we thank our many readers who gave us feedback on this paper including Ion Stoica and M. Satyanarayanan and Srini Seshan.

References

[1] 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.

[2] W. Almesberger, J. Hadi, and A. Kuznetsov. Differentiated services on linux. Available at http://lrcwww.epfl.ch/linux-diffserv/.

[3] Werner Almesberger. Linux network traffic control — implementation overview. Available at http://lrcwww.epfl.ch/linux-diffserv/.

[4] Bruce Maggs at Akamai. Personal communication., 2001.

[5] G. Banga, P. Druschel, and J. Mogul. Better operating system features for faster network servers. In *Proc. Workshop on Internet Server Performance*, June 1998.

[6] Gaurav Banga and Peter Druschel. Measuring the capacity of a web server under realistic loads. *World Wide Web*, 2(1-2):69–83, 1999.

[7] Nikhil Bansal and Mor Harchol-Balter. Analysis of SRPT scheduling: Investigating unfairness. In *ACM Sigmetrics International Conference on Measurement and Modeling of Computer Systems*, 2001.

[8] Paul Barford and M. E. Crovella. Measuring web performance in the wide area. *Performance Evaluation Review – Special Issue on Network Traffic Measurement and Workload Characterization*, August 1999.

[9] Paul Barford and Mark Crovella. Critical path analysis of tcp transactions. In *SIGCOMM*, 2000.

[10] Paul Barford and Mark E. Crovella. Generating representative Web workloads for network and server performance evaluation. In *Proceedings of SIGMETRICS '98*, pages 151–160, July 1998.

[11] Michael Bender, Soumen Chakrabarti, and S. Muthukrishnan. Flow and stretch metrics for scheduling continuous job streams. In *Proceedings of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms*, 1998.

[12] Azer Bestavros, Robert L. Carter, Mark E. Crovella, Carlos R. Cunha, Abdelsalam Heddaya, and Sulaiman A. Mirdad. Application-level document caching in the internet. In *Proceedings of the Second International Workshop on Services in Distributed and Networked Environments (SDNE'95)*, June 1995.

[13] H. Braun and K. Claffy. Web traffic characterization: an assessment of the impact of caching documents from NCSA's Web server. In *Proceedings of the Second International WWW Conference*, 1994.

[14] Adrian Cockcroft. Watching your web server. The Unix Insider at http://www.unixinsider.com, April 1996.

[15] Mark E. Crovella and Azer Bestavros. Self-similarity in World Wide Web traffic: Evidence and possible causes. *IEEE/ACM Transactions on Networking*, 5(6):835–846, December 1997.

[16] Mark E. Crovella, Robert Frangioso, and Mor Harchol-Balter. Connection scheduling in web servers. In *USENIX Symposium on Internet Technologies and Systems*, October 1999.

[17] Mark E. Crovella, Murad S. Taqqu, and Azer Bestavros. Heavy-tailed probability distributions in the World Wide Web. In *A Practical Guide To Heavy Tails*, pages 3–26. Chapman & Hall, New York, 1998.

[18] Allen B. Downey. A parallel workload model and its implications for processor allocation. In *Proceedings of High Performance Distributed Computing*, pages 112–123, August 1997.

[19] Peter Druschel and Gaurav Banga. Lazy receiver processing (LRP): A network subsystem architecture for server systems. In *Proceedings of OSDI '96*, October 1996.

[20] Fielding, Gettys, Mogul, Frystyk, and Berners-lee. DNS support for load balancing. RFC 2068, April 1997.

[21] The Apache Group. Apache web server. `http://www.apache.org`.

[22] James Gwertzman and Margo Seltzer. The case for geographical push-caching. In *Proceedings of HotOS '94*, May 1994.

[23] A. Halikhedkar, Ajay Uggirala, and D.K. Tammana. Implemenation of differentiated services in linux (diffspec). Available at http://www.rsl.ukans.edu/ dilip/845/FAGASAP.html.

[24] Internet Town Hall. The internet traffic archives. Available at `http://town.hall.org/Archives/pub/ITA/`.

[25] M. Harchol-Balter and A. Downey. Exploiting process lifetime distributions for dynamic load balancing. *ACM Transactions on Computer Systems*, 15(3), 1997.

[26] Mor Harchol-Balter, Nikhil Bansal, Bianca Schroeder, and Mukesh Agrawal. Implementation of SRPT scheduling in web servers. Technical Report CMU-CS-00-170, 2000.

[27] Microsoft TechNet Insights and Answers for IT Professionals. The arts and science of web server tuning with internet information services 5.0. http://www.microsoft.com/technet/, 2001.

[28] M. Kaashoek, D. Engler, D. Wallach, and G. Ganger. Server operating systems. In *SIGOPS European Workshop*, September 1996.

[29] S. Manley and M. Seltzer. Web facts and fantasy. In *Proceedings of the 1997 USITS*, 1997.

[30] Evangelos Markatos. Main memory caching of Web documents. In *Proceedings of the Fifth Interntional Conference on the WWW*, 1996.

[31] J. Mogul. Operating systems support for busy internet servers. Technical Report WRL-Technical-Note-49, Compaq Western Research Lab, May 1995.

[32] V. N. Padmanabhan and J. Mogul. Improving HTTP latency. *Computer Networks and ISDN Systems*, 28:25–35, December 1995.

[33] Vivek S. Pai, Peter Druschel, and W. Zwaenepoel. Flash: An efficient and portable web server. In *Proceedings of USENIX 1999*, June 1999.

[34] S .Radhakrishnan. Linux – advanced networking overview version 1. Available at http://qos.ittc.ukans.edu/howto/.

[35] J. Roberts and L. Massoulie. Bandwidth sharing and admission control for elastic traffic. In *ITC Specialist Seminar*, 1998.

[36] Linus E. Schrage and Louis W. Miller. The queue M/G/1 with the shortest remaining processing time discipline. *Operations Research*, 14:670–684, 1966.

[37] A. Silberschatz and P. Galvin. *Operating System Concepts, 5th Edition*. John Wiley & Sons, 1998.

[38] W. Stallings. *Operating Systems, 2nd Edition*. Prentice Hall, 1995.

[39] A.S. Tanenbaum. *Modern Operating Systems*. Prentice Hall, 1992.