

Soft Timers: Efficient Microsecond Software Timer Support for Network Processing

MOHIT ARON and PETER DRUSCHEL

Rice University

This paper proposes and evaluates soft timers, a new operating system facility that allows the efficient scheduling of software events at a granularity down to tens of microseconds. Soft timers can be used to avoid interrupts and reduce context switches associated with network processing, without sacrificing low communication delays. More specifically, soft timers enable transport protocols like TCP to efficiently perform rate-based clocking of packet transmissions. Experiments indicate that soft timers allow a server to employ rate-based clocking with little CPU overhead (2–6%) at high aggregate bandwidths. Soft timers can also be used to perform network polling, which eliminates network interrupts and increases the memory access locality of the network subsystem without sacrificing delay. Experiments show that this technique can improve the throughput of a Web server by up to 25%.

Categories and Subject Descriptors: C.5.5 [Computer System Implementation]: Servers; D.4.1 [Operating Systems]: Process Management—*scheduling*; D.4.4 [Operating Systems]: Communications Management—*network communication*

General Terms: Design, Performance

Additional Key Words and Phrases: Timers, transmission scheduling, polling

1. INTRODUCTION

We propose and evaluate *soft timers*, an operating system facility that allows efficient scheduling of software events at microsecond (μ sec) granularity.

The key idea behind soft timers is to take advantage of certain states in the execution of a system where an event handler can be invoked at low cost. Such states include the entry points of the various OS kernel handlers, which are executed in response to system calls, exceptions (TLB

This work was supported in part by NSF grant CCR-9803673, by Texas TATP grant 003604, by an IBM partnership Award, and by equipment donations from Compaq Western Research Lab and from HP Labs.

Authors' address: Department of Computer Science, Brown School of Engineering, Rice University, Houston, TX 77251-1892; email: {aron;druschel}@cs.rice.edu.

Permission to make digital/hard copy of part or all of this work for personal or classroom use is granted without fee provided that the copies are not made or distributed for profit or commercial advantage, the copyright notice, the title of the publication, and its date appear, and notice is given that copying is by permission of the ACM, Inc. To copy otherwise, to republish, to post on servers, or to redistribute to lists, requires prior specific permission and/or a fee.

© 2000 ACM 0734-2071/00/0800–0197 \$05.00

miss, page fault, arithmetic), and hardware interrupts. In these “trigger states,” the cost of saving and restoring of CPU state and the shift in memory access locality associated with the switch to kernel mode have already been incurred; invoking an additional event handler from the trigger state amortizes this overhead over a larger amount of useful computation.

Of course, the times at which a system enters a trigger state are unpredictable and depend on the workload. Therefore, soft timers can schedule events only probabilistically: a soft timer event may be delayed past its scheduled time by a random but bounded amount of time. Under most practical workloads, trigger states are reached often enough to allow the scheduling of events at intervals down to a few tens of μ secs, with rare delays up to a few hundred μ secs. As we will show, soft timers allow the scheduling of events at these intervals with very low overhead, while the use of a conventional hardware interrupt timer at the same rate would result in unacceptable overhead on the system.

We explore the use of a soft-timers facility to perform two optimizations in the network subsystem. Soft timers enable a transport protocol like TCP to efficiently perform *rate-based clocking*, i.e., to transmit packets at a given rate, independent of the arrival of acknowledgment (ACK) packets. Rate-based clocking has been proposed as a technique that improves the utilization of networks with high bandwidth-delay products [Allman et al. 1997; Balakrishnan et al. 1997; Feng et al. 1999; Padmanabhan and Katz 1998; Visweswaraiah and Heidemann 1997]. Our experiments indicate that soft timers enable a Web server to employ rate-based clocking with low CPU overhead (2–6%) at aggregate bandwidths approaching 1Gbps.

A second optimization is *soft-timer-based network polling*. Here, soft-timer events are used to poll the network interface, thus avoiding interrupts. Experiments show that the performance of a Web server using this optimization can increase by up to 25% over a conventional interrupt-based implementation.

The rest of this paper is organized as follows. In Section 2, we provide background and motivation for this work. The soft-timers facility is presented in Section 3. Applications of soft timers are discussed in Section 4. We present empirical results obtained with a prototype implementation of soft timers in Section 5, discuss related work in Section 6, and conclude in Section 7. Background information on the motivation for rate-based clocking can be found in the Appendix.

2. BACKGROUND AND MOTIVATION

Modern CPUs increasingly rely on pipelining and caching to achieve high performance. As a result, the speed of program execution is increasingly sensitive to unpredictable control transfer operations. Interrupts and context switches are particularly expensive, as they require the saving and restoring of the CPU state and entail a shift in memory access locality. This

shift typically causes cache and TLB misses in the wake of the entry and the exit from the interrupt handler, or the context switch, respectively.

The cost of interrupts and context switches is generally not a concern as long as they occur on a millisecond (msec) timescale. For instance, disk interrupts, conventional timer interrupts used for time-slicing, and the associated context switches typically occur at intervals on the order of tens of msecs.

However, high-speed network interfaces can generate interrupts and associated context switches at intervals on the order of tens of μ secs. A network receive interrupt typically entails a context switch to a kernel thread that processes the incoming packet and possibly transmits a new packet. Only after this thread finishes is the activity that was originally interrupted resumed.

As we will show, these interrupts and context switches can have a significant impact on the performance of server systems performing large amounts of network I/O. Even a single Fast Ethernet interface can deliver a full-sized packet every 120 μ secs, and Gigabit Ethernet is already on the market. Moreover, many high-end Web servers already have backbone connections to the Internet at Gigabit speed.

2.1 Rate-Based Clocking

Achieving high network utilization on networks with increasingly high bandwidth-delay products may require transport protocols like TCP to perform *rate-based clocking*, that is, to transmit packets at scheduled intervals, rather than only in response to the arrival of acknowledgment (ACK) packets.

Current TCP implementations are strictly *self-clocking*, i.e., packet transmissions are paced by the reception of ACK packets from the receiver. Adding the ability to transmit packets at a given rate, independent of the reception of ACK packets (*rate-based clocking*), has been proposed to overcome several known shortcomings of current TCP implementations:

- Rate-based clocking can potentially allow a sender to skip the slow-start phase in situations where the available network capacity is known or can be estimated. This can lead to significant increases in network utilization and achieved throughput, particularly when traffic is bursty and the network's bandwidth-delay product is high. Such conditions arise, for instance, with Web (HTTP) traffic in today's Internet [Padmanabhan and Katz 1998; Visweswaraiah and Heidemann 1997].
- Rate-based clocking can overcome the effects of *ACK compression* [Zhang et al. 1991] and *big ACKs*. Either phenomenon may cause a self-clocked sender to transmit a burst of packets in close succession, which can adversely affect network congestion.
- Rate-based clocking allows a TCP sender to shape its traffic in integrated services networks [Feng et al. 1999].

Rate-based clocking requires a protocol implementation to transmit packets at regular intervals. On high-bandwidth networks, the required inter-

vals are in the range of tens to hundreds of μsecs . For instance, transmitting 1500-byte packets at 100Mbps and 1Gbps requires a packet transmission every 120 μsecs and 12 μsecs , respectively. Server systems with high-speed network connections transmit data at these rates even in today's Internet. As we will show in Section 3, conventional facilities for event scheduling available in general-purpose operating systems today cannot efficiently support events at this granularity. A more detailed discussion of the motivation for rate-based clocking can be found in Appendix A.

To summarize this section, interrupts and context switches are increasingly costly on modern computer systems. At the same time, high-speed network interfaces already generate interrupts and associated context switches at a rate that places a significant burden on server systems. Rate-based clocking in transport protocols, which has been proposed as a technique to increase network utilization and performance on high-speed WANs, necessitates even more interrupts when implemented using conventional timers.

In the following section, we present the design of the soft-timers facility, which enables efficient rate-based clocking and can be used to avoid network interrupts.

3. DESIGN OF THE SOFT-TIMERS FACILITY

In this section, we present the design of soft timers, a mechanism for scheduling fine-grained events in an operating system with low overhead.

Conventional timer facilities schedule events by invoking a designated handler periodically in the context of a hardware interrupt. For example, an Intel 8253 programmable interrupt timer chip is usually supplied with a Pentium-based CPU. The former can be programmed to interrupt the processor at a given frequency.

Unfortunately, using hardware interrupts for fine-grained event scheduling causes high CPU overhead for the following reasons:

- Upon a hardware interrupt, the system has to save the context of the currently executing program and, after executing the interrupt handler, restore the interrupted program's state.
- Hardware interrupts are usually assigned the highest priority in the operating system. Thus, irrespective of the process currently running on the CPU, the interrupt handler is allowed to interrupt the execution of the former. In general, the data and instructions touched by the interrupt handler are unrelated to the interrupted computation, which can adversely affect cache and TLB locality.

In summary, the overhead of saving state, restoring state, and the cache/TLB pollution associated with interrupts limits the granularity at which a conventional facility can schedule events. For instance, in Section 5 we show that the total cost of a timer interrupt in a busy Web server amounts to on average 4.45 μsecs on a 300MHz Pentium-II machine

running FreeBSD-2.2.6. This cost is insignificant when interrupts are being generated every msec but it is unacceptable when interrupts need to be generated (say) every 20 μ secs.

The key idea behind soft timers is as follows. During normal operation, a system frequently reaches states in its execution where an event handler could be invoked with minimal overhead. Examples of such opportune *trigger states* are

- at the end of executing a system call, just before returning to the user program,
- at the end of executing an exception handler, such as the ones triggered by a memory exception (e.g., TLB¹ or page fault) or an arithmetic exception,
- at the end of executing an interrupt handler associated with a hardware device interrupt, just before returning from the interrupt,
- whenever a CPU is executing the idle loop.

In these trigger states, invoking an event handler costs no more than a function call, and no saving/restoring of CPU state is necessary. Furthermore, the cache and TLB contents in these trigger states have already been disturbed due to the preceding activity, potentially reducing the impact of further cache pollution by the event handler. Performance results presented in Section 5 confirm this reasoning.

Whenever the system reaches one of the trigger states, the soft-timer facility checks for any pending soft-timer events and invokes the associated handlers when appropriate. As such, the facility can execute pending events without incurring the cost of a hardware timer interrupt. Checking for pending soft-timer events in a trigger state is very efficient: it involves reading the clock (usually a CPU register) and a comparison with the scheduled time of the earliest soft timer event.² As we will show, performing this check whenever the system reaches a trigger state has no noticeable impact on system performance.

A disadvantage of the soft-timer facility is that the time at which an event handler is invoked may be delayed past its scheduled time, depending on how much time passes between the instant when a soft-timer event becomes due and the instant when the system reaches a trigger state.

The maximal delay experienced by a soft-timer event is bounded, because the soft-timer facility still schedules a periodic hardware interrupt that is used to schedule any overdue events. The key point to notice is that as long as a system reaches trigger states with sufficient frequency, the soft-timer facility can schedule events at much finer granularity than would be feasible using a periodic hardware interrupt.

¹In some architectures (e.g., Pentium), TLB misses are handled in hardware; in these machines, TLB faults cannot be used as trigger states.

²A modified form of timing wheels [Varghese and Lauck 1987] is used to maintain scheduled soft-timer events.

Results presented in Section 5 show that a 300MHz Pentium II system running a variety of workloads reaches trigger states frequently enough to allow the scheduling of soft-timer events at a granularity of tens of μ secs.

The soft-timer facility provides the following operations.

- `measure_resolution()`. Returns a 64-bit value that represents the resolution of the clock (in Hz).
- `measure_time()` returns a 64-bit value representing the current real time in ticks of a clock whose resolution is given by `measure_resolution()`. Since this operation is intended to measure time intervals, the time need not be synchronized with any standard time base.
- `schedule_soft_event(T, handler)`: schedules the function *handler* to be called at least T ticks in the future (the resolution of T is specified by `measure_resolution()`).
- `interrupt_clock_resolution()`: gives the expected minimal resolution (in Hz) at which the facility can schedule events and equals the frequency of the system's periodic timer interrupt, which is used to "back up" soft timers.

The soft-timer facility fires an event (by calling the corresponding handler) when the value returned by `measure_time()` exceeds the value stored at the time the event was scheduled by at least $T + 1$ (the increment by one accounts for the fact that the time at which the event was scheduled may not coincide with a clock tick). If X is the resolution of the interrupt clock relative to the measurement clock (i.e., $X \equiv \text{measure_resolution}() / \text{interrupt_clock_resolution}()$), then the soft-timer facility puts the following bounds on the actual time (in ticks of the measurement clock) when the event fires:

$$T < \text{Actual Event Time} < T + X + 1$$

Figure 1 gives examples of the above bounds when $T = 1$ and $X = 2$. It is to be noted that the increment by one is negligible if the measurement clock is significantly finer than the interrupt clock (as is the case in most modern systems).

The reason for the upper bound is that the soft-timer facility uses a periodic timer interrupt to check for overdue soft-timer events. However, the actual time at which the handler is invoked is likely to be much closer to T . Expressed differently, if we assume that

$$\text{Actual Event Time} = T + d$$

where d is a random variable in the range $[0..X + 1]$, then the probability distribution of d would be uniform if a conventional timer interrupt-based facility was used.³ With typical values for the measurement resolution and

³It is assumed here that the event was scheduled at a random time not synchronized with the scheduling clock.

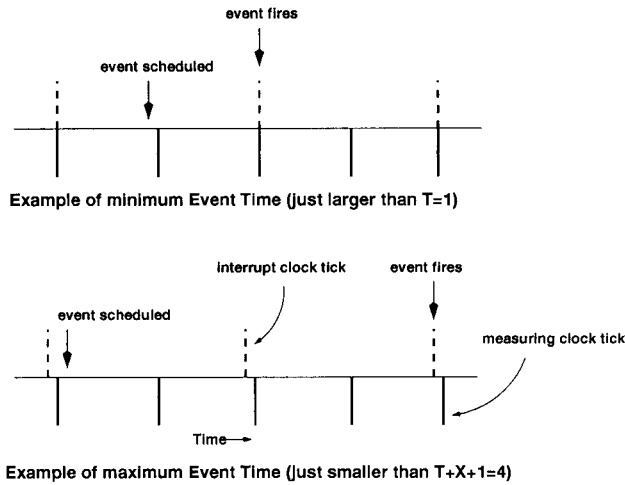


Fig. 1. Lower and upper bounds for event scheduling.

interrupt clock resolution of 1MHz (1 μ secs) and 1KHz (1 msec), respectively, X is 1000, and the maximal delay is 1001 μ secs.

With soft timers, the probability distribution of d is dependent on the system's workload, which influences how often trigger states are reached. Results shown in Section 5 show that among a variety of workloads, the worst-case distribution of d results in a mean delay of 31.6 μ secs and is heavily skewed toward low values (median is 18 μ secs). Therefore, applications that can benefit from fine-grained events on the order of tens of μ secs in the common case, but can tolerate rare delays up to the resolution of the system's interrupt clock (typically 1 msec), are well served by soft timers.

4. APPLICATIONS OF SOFT TIMERS

In this section, we describe two applications of soft timers: rate-based clocking and network polling. In Section 5, we will present empirical results to evaluate the use of soft timers in these applications.

4.1 Rate-Based Clocking

As discussed in Section 2.1, achieving high utilization in networks with large bandwidth-delay products may require transport protocols like TCP to perform rate-based clocking. In a conventional implementation of rate-based clocking, a periodic hardware timer event must be scheduled at the intended rate of packet transmissions. At network speeds of several hundred Mbps and a packet size of 1500 bytes (Ethernet), this would require timer interrupt rates of one every few tens of μ secs. Given the overhead of hardware timer interrupts (e.g., 4.45 μ secs), this would lead to unacceptable overhead.

We observe that transmitting multiple packets per timer event would lead to bursty packet transmissions and defeat the purpose of rate-based clocking, which is to transmit data at a relatively constant rate. However,

packet transmissions on different network connections that have separate bottleneck links could be performed in a single timer event.

Soft timers allow the clocked transmission of network packets at average intervals of tens of μ secs with low overhead. Due to the probabilistic nature of soft-timer event scheduling, the resulting transmission rate is variable. In Section 5, we will empirically show the statistics of the resulting transmission process.

An interesting question is how a protocol implementation should schedule soft-timer transmission events to achieve a given target transmission rate. Scheduling a series of transmission events at fixed intervals results in the correct average transmission rate. However, this approach can lead to occasional bursty transmissions when several transmission events are all due at the end of a long interval during which the system did not enter a trigger state. A better approach is to schedule only one transmission event at a time and let the protocol maintain a running average of the actual transmission rate. The next transmission event is then adaptively scheduled in the context of the previous event handler to smooth the rate fluctuations.

Our prototype implementation employs a simple algorithm for scheduling the next transmission. The algorithm uses two parameters, the target transmission rate and the maximal allowable burst transmission rate. The algorithm keeps track of the average transmission rate since the beginning of the current train of transmitted packets. Normally, the next transmission event is scheduled at an interval appropriate for achieving the target transmission rate. However, when the actual transmission rate falls behind the target transmission rate due to soft-timer delays, then the next transmission is scheduled at an interval corresponding to the maximal allowable burst transmission rate.

We will experimentally evaluate the use of soft timers for rate-based clocking in Section 5.

4.2 Network Polling

In conventional network subsystem implementations, the network interfaces generate a hardware interrupt to signal the completion of a packet reception or transmission.⁴ Upon a receiver interrupt, the system accepts the packet, performs protocol processing, and signals any blocked process that has been waiting to receive data. Upon a transmit interrupt, the system decreases the reference count on the transmitted packets' buffers, possibly deallocating them. In busy systems with high-speed network interfaces (e.g., server systems), network interrupts can occur at a rate of one every few tens of μ secs.

Another approach to scheduling network processing is polling, where the system periodically reads the network interfaces' status registers to test for completed packet receptions or transmissions. In a pure polling system, the

⁴Some interfaces can be programmed to signal the completion of a burst of packet transmissions or receptions.

scheduler periodically calls upon the network driver to poll the network interfaces.

Pure polling avoids the overhead of interrupts and can reduce the impact of memory access locality shifts by (1) testing for network events at “convenient” points in the execution of the system and by (2) aggregating packet processing. By performing polling when the scheduler is active, packet processing is performed at a time when the system already suffers a locality shift. By polling at an appropriate average rate, multiple packets may have completed since the last poll, thus allowing the aggregation of packet processing, increasing memory access locality.

However, the disadvantage of pure polling is that it may substantially increase communication latency by delaying packet processing. As a result, hybrid schemes have been proposed. Smith and Traw [1993] use periodic hardware timer interrupts to initiate polling for packet completions when using a gigabit network interface. This approach involves a trade-off between interrupt overhead and communication delay. Mogul and Ramakrishnan [1997] propose a system that uses interrupts under normal network load and polling under overload, in order to avoid receiver livelock. When processing of a packet completes, the system polls the network interface for more outstanding packets; only when no further packets are found are network interrupts reenabled.

Soft timers offer a third design choice. With soft-timer-based network polling, a soft-timer event is used to poll the network interfaces. As in pure polling, network interrupts are avoided and memory access locality is improved because network polling and processing is performed only when the associated soft-timer event expires and the system reaches a trigger state. However, since soft-timer events can be efficiently scheduled at μsec granularity, communications latency can be close to that achieved with interrupt-driven network processing in the common case.

In general, the soft-timer poll interval can be dynamically chosen so as to attempt to find a certain number of packets per poll, on average. We call this number the *aggregation quota*. An aggregation quota of one implies that one packet is found, on average, per poll.

We will experimentally evaluate the use of soft timers for network polling in Section 5.

5. EXPERIMENTAL RESULTS

In this section, we present experimental results to evaluate the proposed soft-timer facility. We quantify the overhead of a soft-timer facility and compare it to the alternative approach of scheduling events using hardware timer interrupts, using either on-chip or off-chip timer facilities. We also present measurements that show the distribution of delays in soft-timer event handling, given a variety of system workloads.

Next, we evaluate the performance of soft timers when used to perform rate-based clocking and network polling. Finally, we show how modern

on-chip timer facilities can be combined with soft timers to achieve high timer granularity with tight delays at low cost.

5.1 Base Overhead of Hardware Timers and Soft Timers

Our first set of experiments is designed to determine the base overheads of hardware interrupt timers and soft timers. Hardware timer interrupts are generated using two types of timer devices. The first are off-chip timers, such as the 8253 timer chip used in Intel-based personal computers. The second is an on-chip timer device, namely the APIC device found in Pentium III CPUs [Intel 2000]. The FreeBSD kernel, like most current OSes, uses the 8253 timer chip to drive its timing facilities. To be able to also measure APIC-based timers, the appropriate device support was added to the FreeBSD kernel.

We also implemented soft timers in the FreeBSD kernel. Trigger states were added in the obvious places described in Section 3. In practice, we found that the trigger interval distribution could be improved by adding a few additional trigger states to ensure that certain kernel loops contain a trigger state. Examples of such loops are the TCP/IP output loop and the TCP timer processing loop. Since Intel x86 CPUs handle TLB misses in hardware, TLB miss events cannot be used as trigger states on Intel-based PCs.

The idle loop checks for pending soft-timer events. However, to minimize power consumption, an idle CPU halts when either (a) there are no soft-timer events scheduled at times prior to the next hardware timer interrupt, or (b) another idle CPU is already checking for soft-timer events.

The experimental setup consists of a number of Compaq AlphaStation 500au (500MHz 21164 CPU), 500MHz Pentium-III (PIII), and 300MHz Pentium-II (PII) machines, connected through a switched 100Mbps Ethernet. We ran either the Apache-1.3.3 (<http://www.apache.org>) or the Flash [Pai et al. 1999] Web server on one of the machines while four PII machines ran a client program that repeatedly requested a 6KB file from the Web server. The number of simultaneous requests to the Web server were set such that the server machine was saturated.

The FreeBSD OS runs on the server machine (versions 2.2.6 and 4.0-beta runs on the Pentium and Alpha machines, respectively). The kernel uses its standard timer facilities to schedule all events in the system. To measure hardware timer overheads, an additional hardware timer interrupt was scheduled with varying frequency. To measure soft-timer overheads, we scheduled a periodic soft timer event such that a handler was invoked whenever the system reaches a trigger state. That is, we programmed the soft-timer facility to invoke a soft-timer event handler at the maximal frequency possible, given the Web server workload.

In the first experiment, a “null handler” (i.e., a handler function that immediately returns upon invocation) was invoked whenever a timer fires, to isolate the overhead of the various timer facilities alone. We measured the throughput of the Flash server in the presence of the additional timer

Table I. Per-Event Timer Costs with Null Event Handler

	Alpha-500	8253/PII-300	8253/PIII-500	APIC/PIII-500	Soft Timers
Overhead (μ sec)	8.64	4.45	4.36	0.8	≈ 0

events. By measuring the impact of timer events on the performance of a realistic workload, we are able to capture the full overhead of timers events, including secondary effect like cache and TLB pollution that result from taking an interrupt (in the case of hardware timers), executing the timer facility, and executing the (null) event handler.

The measured total overhead per event (as calculated from the slowdown of the Web server in the presence of the additional events) was found to be independent of the event frequency in this experiment.⁵ Table I shows the total overhead per (null) event handler invocation, for the various machines and timers.

The results show that the overhead per event handler invocation with the off-chip hardware timers is substantial, and ranges from 8.64 μ secs on the Alpha machine to 4.36 μ secs on the 500MHz PIII machine. Moreover, the results indicate that hardware timer overhead does not scale with CPU speed and suggests that the relative cost of timer interrupts increases as CPUs get faster. Finally, the result obtained with the Alpha machine indicates that the high overhead associated with off-chip timer interrupt handling is not unique to Intel PCs. The overhead of hardware timers based on the on-chip APIC device found on PIII CPUs is significantly lower than that of off-chip timers, but still substantial (0.8 μ secs).

The soft-timer handler invocations, on the other hand, caused no observable difference in the Web server's throughput. This indicates that the base overhead imposed by our soft-timer approach is negligible. This is intuitive because the calls to the handler execute with the overhead of a procedure call, whereas a hardware interrupt involves saving and restoring the CPU state. With soft timers, the event handler was called on average every 22.53 μ secs on the 300MHz PII machine and every 12.45 μ secs on the 500MHz PIII machine.

Note that the measured overhead of a timer interrupt can be lower on all platforms when the machine is idle, since the code, data, and TLB entries used during timer event handling may remain fully cached in this case. Our experiment tries to obtain a more meaningful measure of the overhead by evaluating the total impact of timer events on the performance of a real workload that stresses the memory system. The results show that hardware timers have a significant base overhead.

The above experiment still does not fully expose the cost of timer events in a real application, because the null event handler makes no memory references. That is, the overhead results are pessimal with regard to the cache and TLB pollution caused by a real event handler. To quantify the

⁵Measurements using performance counters to measure the average elapsed time spent in the interrupt handler confirm this result.

Table II. Timer Costs with Synthetic Event Handler, Scheduled Every 10 μ secs

	APIC/PIII-500	Soft Timers
Overhead (μ sec)	5.1	3.5
Icache-misses ($\times 10^6$)	153.2	149.7
Dcache-misses ($\times 10^6$)	551.4	377.9
ITLB-misses ($\times 10^6$)	18.25	17.00

impact of this pollution, we performed an additional experiment, this time with a synthetic event handler that touches 50 data cache lines and 2 instruction cache lines on 2 separate pages. (A different set of 50 data cache lines is touched during each invocation of the handler, sequentially cycling through a total of 4 data pages). Events were scheduled 10 μ secs and 20 μ secs, respectively, after the previous invocation of the handler. The results of this measurement are shown in Tables II (for events scheduled every 10 μ secs) and III (for events scheduled every 20 μ secs).

Each table shows the overhead per event (calculated from the slow-down of the Web server application), and the total number of data cache, instruction cache, and instruction TLB misses⁶ during the runtime of the experiment for each timer facility. Results for events scheduled every 10 μ secs could not be obtained for 8253-based timers due to the high overhead of that facility.

The results show that the memory accesses in the event handler cause substantial additional overheads, which dominate the base overhead measured in the previous experiment. The overheads are more pronounced at the lower event rate; this is intuitive, since a longer interval between successive handler invocations decreases the chance that data, instructions, and TLB entries used by the handler remain cached.

Soft timers are less affected by the cache/TLB pollution caused by the event handler than the hardware timers, resulting in less overhead. This confirms our reasoning that soft timers reduce cache pollution by invoking event handlers from trigger states in which the system already undergoes a shift in locality. The results show that the system suffers substantially less data cache misses (20–31%) and noticeably less instruction TLB misses (7–13%) when soft timers are used. The reduction in instruction cache misses is only slight—the likely reason is that our synthetic event handler only fetches two i-cache lines.

In summary, soft timers suffer from significantly less cache and TLB pollution than hardware-interrupt-based timers. The costs associated with cache and TLB pollution caused by the event handler can dominate the base overheads of all timer facilities. As a result, soft timers offer clear performance advantages, even when compared to the relatively efficient on-chip APIC timers.

⁶The PIII performance counters do not support the measurement of data TLB misses.

Table III. Timer Costs with Synthetic Event Handler, Scheduled Every 20 μ secs

	8253/PIII-500	APIC/PIII-500	Soft Timers
Overhead (μ sec)	9.3	6.7	4.8
Icache-misses ($\times 10^6$)	160.8	152.8	147.8
Dcache-misses ($\times 10^6$)	372.1	355.7	282.4
ITLB-misses ($\times 10^6$)	18.75	18.96	16.30

5.2 Soft-Timer Event Granularity Under Different Workloads

Recall that once a soft-timer event is due, the associated handler is executed at the earliest time when the system reaches a trigger state. The performance of a soft-timer facility, i.e., the granularity and precision with which it can schedule events, therefore depends on the frequency at which the system reaches trigger states.

We measured the distribution of times between successive trigger states for a variety of workloads. Figure 2 shows the cumulative distribution function of time between successive trigger states, for various workloads executing on a 300MHz PII machine.

The workloads are as follows. “ST-Apache” corresponds to the Apache Web server workload from the previous experiment. In “ST-Apache-compute,” an additional compute-bound background process is running concurrently with the Web server. “ST-Flash” is a Web server workload using a fast event-driven Web server called *Flash* [Pai et al. 1999]. “ST-real-audio” was measured with a copy of the RealPlayer (<http://www.realplayer.com>) running on the machine, playing back a live audio source from the Internet. “ST-nfs” reflects the trigger state interarrival times when the workload is a NFS fileserver. Finally, “ST-kernel-build” was measured while a copy of the FreeBSD-2.2.6 kernel was built on the machine from the sources.

Additional information about the distribution with each workload is given in Table IV. Two million samples were taken in each workload to measure the distributions.

The results show that under a workload typical of a busy Web server executing on a 300MHz PII machine, the soft-timer facility can schedule events at a mean granularity of tens of μ secs with negligible overhead and with delays over 100 μ secs in less than 6% of the samples. As shown below, this performance is sufficient to perform rate-based clocking of 1500-byte packets at several hundreds of Mb/s and allows effective polling of network interface events at the same rate.

In a busy Web server, it is intuitive that the many network packet arrivals, disk device interrupts, and system calls provide frequent trigger states. One concern is that the presence of compute-bound background computations may cause long periods where the system does not encounter a trigger state, thus degrading the performance of the soft-timer facility.

To measure this effect, we added a compute-bound background process to the Web server, which executes in a tight loop without performing system calls (“ST-Apache-compute”). The results show that the presence of background processes has no tangible impact on the performance of the soft-timer facility. The reason is that a busy Web server experiences frequent




Fig. 2. Trigger state interval (CDF), 300MHz PII.

network interrupts that have higher priority than application processing and yield frequent trigger states even during periods where the background process is executing.

“ST-nfs” is another example of a server workload. The NFS server is saturated but disk-bound, leaving the CPU idle approximately 90% of the time. The vast majority of samples indicate a trigger state interval around 2 μ secs on this workload.

The RealPlayer (“ST-real-audio”) was included because it is an example of an application that saturates the CPU. Despite the fact that this workload performs mostly user-mode processing and generates a relatively low rate of interrupts, it yields a distribution of trigger state intervals with very low mean, due to the many systems calls that RealPlayer performs.

Finally, we measure a workload where the FreeBSD OS kernel is built from the source code. This workload involves extensive computation (compilation, etc.) as well as disk I/O.

To determine the impact of CPU speed on the trigger interval distribution, we repeated the experiment with the “ST-Apache” and “ST-Flash” workloads on a machine with a 500MHz Pentium III CPU. The summary information about the resulting distribution is included in Table IV. The results show that the shape of the distribution is similar to that obtained with the slower CPU; however, the mean is reduced by a factor that roughly reflects the CPU clock speed ratio of the CPUs. This indicates that the granularity of soft-timer events increases approximately linearly with CPU speed.

While our selection of measured workloads is necessarily limited, we believe that the soft-timer facility can provide fine-grained event support across a wide range of practical workloads. The reason is that (1) most practical programs frequently make system calls, suffer page faults, TLB faults, or generate other exceptions that cause the system to reach a trigger

Table IV. Trigger State Interval Distribution

	Max (μ s)	Mean (μ s)	Median (μ s)	StdDev (μ s)	> 100 μ s (%)	> 150 μ s (%)
ST-Apache	476	31.52	18	32	5.3	0.39
ST-Apache-compute	585	31.59	18	32.1	5.3	0.43
ST-Flash	1000	22.53	17	20.8	1.09	0.013
ST-real-audio	1000	8.47	6	13.2	0.025	0.013
ST-nfs	910	2.13	2	3.3	0.021	0.011
ST-kernel-build	1000	5.63	2	47.9	0.038	0.033
ST-Apache (PIII-500)	1000	19.41	11	23	0.44	0.13
ST-Flash (PIII-500)	1000	12.45	9	12	0.02	0.01




Fig. 3. Trigger interval medians during 1 ms and 10 ms intervals, ST-Apache-compute workload.

state and (2) the soft-timer facility can schedule events at very fine grain whenever a CPU is idle.

In the most pessimistic scenario, all CPUs are busy, the executing programs make infrequent system calls, cause few page faults or other exceptions, and there are few device I/O interrupts. These conditions mark the absence of significant I/O or communication activity in the workload, and can arise, for instance, in scientific applications. However, observe that μ sec timers are used primarily in networking, and it is thus unlikely that any soft-timer events are scheduled under such conditions.

5.3 Changes in Trigger Interval Distribution Over Time

The trigger interval distributions shown in the previous section are aggregated over 2 million samples, corresponding to 4–64 secs. of execution time for the various workloads. A related question is how the trigger interval distribution changes during the runtime of a workload. For instance, it is conceivable that context switching between different processes could cause significant changes in the trigger interval distribution. To investigate this question, we computed the medians of the trigger interval distributions during intervals of 1 ms and 10 ms. Results are plotted in Figure 3 for a period of 10 secs. of the runtime of the “ST-Apache-compute” workload. The

Table V. Trigger State Sources

Source	Fraction of samples (%)
syscalls	47.7
ip-output	28
ip-intr	16.4
tcpip-others	5.4
traps	2.5

x-axis represents the runtime of the workload, the y-axis the median of the trigger interval distribution during a given interval (1 ms and 10 ms).

With 1 ms intervals, the bulk of the trigger interval medians are in the range from 14 to 26 μ secs. A few intervals (less than 1.13%) have medians above 40 μ secs. The medians for the 10 ms intervals (which corresponds to a timeslice in the FreeBSD system), on the other hand, almost all fall into a narrow band between 17 and 19 μ secs.

These results indicate that the dynamic behavior of the workload appears to cause noticeable variability in the trigger interval distribution over 1 ms intervals. However, there is little variability in the trigger interval distributions over 10 ms intervals.

5.4 Trigger Interval Distribution by Event Source

A related question is what fraction of trigger states is contributed by each event source and how that contribution affects the resulting trigger state interval distribution. To answer this questions, we separately accounted for trigger states by event source for the “ST-Apache” workload. Table V shows the fraction of trigger state samples contributed by each event source.

The sources “syscalls” and “traps” are self-explanatory. The source “ip-output” generates a trigger event every time an IP packet (e.g., TCP segment) is transmitted. The source “tcpip-others” represents a number of other trigger states in the network subsystem, such as the processing loop for TCP timers. Network interface interrupts are reflected in the “ip-intr” source.

Figure 4 shows the impact that each trigger source has on the trigger interval distribution. The graphs show the CDFs of the resulting trigger interval distributions when one of the trigger sources is removed. For instance, “no ipintr” shows the CDF of the resulting trigger interval distribution when there is no trigger state associated with network interrupts. “All” represents the original distribution for the “ST-Apache” workload from Figure 2. It is evident from the results that system calls and IP packet transmissions are the most important sources of trigger events in this workload.

5.5 Rate-Based Clocking: Timer Overhead

In this section, we evaluate the use of soft timers to perform rate-based clocking in TCP. We show results that compare the overhead of performing rate-based clocking with soft timers versus hardware timers, we evaluate




Fig. 4. Impact of event sources on trigger interval, CDF (ST-Apache workload).

the statistics of the packet transmission process and explore the potential for network performance improvements due to rate-based clocking.

Our first experiment is designed to explore the overhead of rate-based clocking in TCP using soft timers versus hardware timers. The experimental setup is the same as in the previous experiment except that the Web server's TCP implementation uses rate-based clocking using either soft timers or a conventional interrupt timer to transmit packets.

The soft timer was programmed to generate an event every time the system reaches a trigger state. One packet is transmitted whenever the handler is invoked and a packet is pending transmission. On a LAN, such as the one used in our testbed, FreeBSD's TCP implementation does not use slow-start. Thus, all packets are normally sent in a burst, as fast as the outgoing network link can transmit them. Since the transmission of a 1500-byte packet takes 120 μ secs on our 100Mbps network, the use of rate-based clocking has no observable impact on the network. Therefore, the experiment isolates the overhead of using soft timers versus hardware timers for rate-based clocking in TCP, but does not expose possible benefits of rate-based clocking.

Table VI shows the performance results obtained on the 300MHz Pentium II machine. We present results for both the Apache-1.3.3 Web server as well as the Flash server. For the results with hardware timers, the 8253 was programmed to interrupt once every 20 μ secs (50KHz frequency), causing the dispatch of a thread (BSD software interrupt) that transmits a packet. From the previous experiments, we know the base overhead for event dispatch at this rate is about 22%. The extra overhead indicated by the results is due to cache/TLB pollution caused by the event handler execution, since the computation performed by the handler is exactly the same as that performed during transmission of a packet in the original TCP implementation.

Table VI. Overhead of Rate-Based Clocking, 300MHz PII

	Apache	Flash
Base system, no rate-based clocking		
Throughput (conn/s)	774	1303
Icache-misses ($\times 10^6$)	532.3	339.6
Dcache-misses ($\times 10^6$)	278.6	171.6
ITLB-misses ($\times 10^6$)	32.12	14.98
8253 timer, target intvl = $20\mu\text{sec}$		
Throughput (conn/s)	560	827
Ovhd (%)	28	36
Icache-misses ($\times 10^6$)	584.8	388.5
Dcache-misses ($\times 10^6$)	308.3	206.9
ITLB-misses ($\times 10^6$)	36.99	19.00
Avg xmit intvl (μsecs)	31	35
Soft timer		
Throughput (conn/s)	756	1224
Ovhd (%)	2	6
Icache-misses ($\times 10^6$)	564.3	374.5
Dcache-misses ($\times 10^6$)	290.9	172.9
ITLB-misses ($\times 10^6$)	32.19	15.86
Avg xmit intvl (μsecs)	34	24

The results indicate that the effect of cache/TLB pollution with hardware timers is at least 4% ($28 - 22 - 2$) and 8% ($36 - 22 - 6$) worse than with soft timers for the Apache and the Flash server, respectively. The fact that Flash is more affected by the cache/TLB pollution can be explained as follows. First, the greater speed of Flash amplifies the relative impact of the fixed overhead associated with cache/TLB pollution on throughput. Second, Apache is a multiprocess server whose frequent context switching leads to relatively poor memory access locality, as indicated by the high number of cache and TLB faults. Flash, on the other hand, is a small, single-process, event-driven server with relatively good cache locality. It is intuitive, therefore, that the Flash server's performance is more significantly affected by the cache/TLB pollution resulting from the timer interrupts.

The results also show that the average time between transmissions with soft timers is only slightly higher than with the hardware timer when using the Apache server, and it is lower when using the Flash server. This result can be explained as follows. With hardware timers, the transmission rate is lower than the rate at which the 8253 chip was programmed because the transmission event handler may in general be delayed due to disabled interrupts. On the other hand, soft timers perform substantially better when the Flash server is used because that server is much faster than Apache and therefore generates trigger states at a higher rate. The combined effect is that soft timers with Flash result in a lower time between transmissions than the hardware timer.

Table VII shows the results of the same experiment obtained on the 500MHz Pentium III machine. The APIC on-chip timer facility available on

Table VII. Overhead of Rate-Based Clocking, 500MHz PIII

	Apache	Flash
Base system, no rate-based clocking		
Throughput (conn/s)	1302	2311
Icache-misses ($\times 10^6$)	199.8	124.8
Dcache-misses ($\times 10^6$)	289.4	154.0
ITLB-misses ($\times 10^6$)	30.98	15.15
APIC timer, target intvl = 20 μ sec		
Throughput (conn/s)	1132	1874
Ovhd (%)	13	19
Icache-misses ($\times 10^6$)	220.8	140.8
Dcache-misses ($\times 10^6$)	366.5	197.3
ITLB-misses ($\times 10^6$)	36.20	19.49
Avg xmit intvl (μ secs)	21.72	23.91
APIC timer, target intvl = 10 μ sec		
Throughput (conn/s)	1094	1798
Ovhd (%)	16	22
Icache-misses ($\times 10^6$)	218.7	144.2
Dcache-misses ($\times 10^6$)	327.5	207.7
ITLB-misses ($\times 10^6$)	42.05	18.99
Avg xmit intvl (μ secs)	11.67	14.41
Soft timer		
Throughput (conn/s)	1276	2173
Ovhd (%)	2	6
Icache-misses ($\times 10^6$)	208.7	135.8
Dcache-misses ($\times 10^6$)	305.5	160.4
ITLB-misses ($\times 10^6$)	31.68	16.86
Avg xmit intvl (μ secs)	19.95	13.07

this CPU was used, and measurements were made with the APIC timer programmed to interrupt every 10 μ secs and 20 μ secs, respectively. Qualitatively, the results are very similar to the ones obtained with the slower PII CPUs. Despite the much faster on-chip APIC timer facility available on the PIII, the overhead of hardware timer events is still substantially larger than that of soft-timer events (22% versus 6% and 13% versus 2% for comparable average xmit intervals with Flash and Apache, respectively).

In summary, the results of this experiment show that soft timers can be used to do rate-based clocking in TCP at rates that approach gigabit speed with very low overhead (2–6% in our experiment). Using a conventional off-chip interrupt timer at this rate has an overhead of 28–36% in our experiment and using the on-chip APIC on the Pentium III CPU still results in an overhead of 13–22%.

5.6 Rate-Based Clocking: Transmission Process Statistics

As discussed in Section 4, our implementation of rate-based clocking based on soft timers uses an adaptive algorithm for scheduling transmissions, in order to smooth variations in the transmission rate caused by the probabi-

Table VIII. Rate-Based Clocking (Target xmit Interval = 40 μ secs), 300MHz PII

Min interval (μ sec)	Soft timers		Hardware timers	
	Avg interval (μ sec)	Std Dev	Avg interval (μ sec)	Std Dev
12 (line speed)	40	34.5	43.6	26.8
15	48	31.6	-	-
20	51.9	30.9	-	-
25	57.5	30.9	-	-
30	61	30.5	-	-
35	65.9	30.1	-	-

listic nature of soft timers. The algorithm keeps track of the actual sending rate, and whenever this rate falls behind the target sending rate, the next transmission event is scheduled so as to achieve the maximal allowable burst sending rate, until the actual sending rate once again catches up with the target sending rate.

We performed an experiment to determine the actual achievable transmission rate and the resulting statistics of the transmission process, as a function of the maximal allowable burst transmission rate, assuming a target transmission rate of one packet every 40 μ secs and 60 μ secs, respectively. The workload in this experiment was that of the busy Apache Web server running on the 300MHz Pentium II machine (“ST-Apache” in Figure 2), which is among the two workloads with the largest mean trigger state interval (i.e., worst case).

We assume in this experiment that the bandwidth of the network link attached to the sender is 1Gbps, and the packet size is 1500 bytes. Therefore, the minimal interval setting of 12 μ secs reflects the maximal transmission rate of the network link. At this minimal interval setting, rate-based clocking is allowed to send packets at the link bandwidth whenever the actual rate is below the target transmission rate.

The results are shown in Tables VIII and IX for target transmission intervals of 40 μ secs and 60 μ secs, respectively. For comparison, results for hardware-timer-based rate-based clocking were also included. The hardware timer was programmed to fire regularly at the target transmission interval.

The results show that on the 300MHz PII machine running the “ST-Apache” workload, soft timers can support rate-based clocking up to rates of one packet transmission every 40 μ secs, if it is allowed to send bursts at the link speed of one packet every 12 μ secs. As the minimal allowable burst interval is increased, the soft timers can no longer maintain an average transmission interval of 40 μ secs, and drops to 65.9 μ secs at a minimal allowable interval of 35 μ secs.

At a target interval of 60 μ secs, soft timers can maintain the average interval up to a minimal allowable burst interval of 30 μ secs. The standard deviation is in all cases in the 30–35 μ secs range and improves as the minimal burst interval increases, as expected.

We note that these measurements apply to rate-based clocking on a single connection. Soft timers can be used to clock transmission on different

Table IX. Rate-Based Clocking (Target xmit Interval = 60 μ secs), 300MHz PII

Min interval (μ sec)	Soft timers		Hardware timers	
	Avg interval (μ sec)	Std Dev	Avg interval (μ sec)	Std Dev
12 (line speed)	60	35.9	63	27.7
15	60	33.2	-	-
20	60	32.3	-	-
25	60	31.2	-	-
30	61	30.5	-	-
35	65.9	30	-	-

connections simultaneously, even at different rates. (A server may perform many transmission simultaneously, resulting in large aggregate bandwidths.) In this case, multiple packets may be transmitted on different connections in a single soft-timer event (i.e., in the context of one trigger state).

With hardware timers, rate-based clocking falls short of the target transmission rate by 3 μ secs and 3.6 μ secs, respectively. The reason is that some timer interrupts are lost during periods when interrupts are disabled in FreeBSD. The hardware timers achieve a somewhat better standard deviation than soft timers, which is to be expected given the probabilistic nature of the latter.

5.7 Rate-Based Clocking: Network Performance

Our next experiment attempts to quantify the potential impact of rate-based clocking on the achieved performance of a Web server over network connections with high bandwidth-delay products. We emphasize that the presented results only give an envelope for the actual achievable performance with rate-based clocking. More research is needed to determine how rate-based clocking can be integrated into TCP; actual performance improvements attainable with a practical implementation are likely to be lower than indicated by our results.

In particular, our prototype implementation of rate-based clocking in TCP assumes that the available capacity in the network is known. Estimating the available capacity is not a trivial problem and is still the subject of on-going research. Furthermore, in order to protect network stability, practical implementations of rate-based clocking will likely have to be conservative, e.g., by starting with an initial sending rate that is only a fraction of the estimated capacity. Related work in this area is discussed in Section 6.

To show the potential effect of rate-based clocking on TCP throughput, we performed an experiment where a variable amount of data is transmitted over a network connection with high bandwidth-delay product. We model this connection in a lab environment by transmitting the data on a 100Mbps Ethernet via an intermediate Pentium II machine that acts as a "WAN emulator." This machine runs a modified FreeBSD kernel configured as an IP router, except that it delays each forwarded packet so as to emulate a WAN with a given delay and bottleneck bandwidth. In our

Table X. Rate-Based Clocking Network Performance (Bandwidth = 50Mbps, RTT = 100 msecs)

Transfer size (1448 Byte packets)	regular TCP		rate-based clocking		
	Xput (Mbps)	Response time (msecs)	Xput (Mbps)	Response time (msecs)	Resp. time reduction (%)
5	0.12	496	0.57	101.2	79
100	1.01	1145	9.36	123.7	89
1000	6.75	1714	34.07	340	80
10000	29.95	3867	46.33	2500	35
100000	45.54	25432	46.60	24863	2

experiment, we choose the WAN delay as 50 ms and the bottleneck bandwidth to be either 50Mbps or 100Mbps. As a result, the TCP connection between client and server machine has a bandwidth-delay product of either 5 or 10Mbits. Network connections with these characteristics are already available in vBNS (Very high performance backbone network service—<http://www.vbns.net/>) and will soon be available in the general Internet.

We performed HTTP requests across the “WAN” connection to an otherwise unloaded server. Either the standard FreeBSD TCP implementation was used, or alternatively our modified implementation, which avoids slow-start and instead uses soft-timer-based rate-based clocking at a rate corresponding to the bottleneck bandwidth, i.e., one packet every 120 μ secs (100Mbps) or 60 μ secs (50Mbps), respectively. Since a persistent connection is assumed to be already established prior to starting the experiment, there is no delay due to connection establishment. (In practice, HTTP requests to a server that has not recently been contacted by the browser require a TCP connection establishment, which increases the delay by one RTT). The results are shown in Tables X and XI.

We see that rate-based clocking can potentially lead to substantial improvements in throughput, response time, and network utilization on networks with high bandwidth-delay products. Maximal response time reductions due to rate-based clocking can range from 2% for large transfers to 89% for medium-sized transfers (100 packets or 141KBs) in our experiment. These improvements are the result of rate-based clocking’s ability to replace TCP slow-start, which tends to underutilize networks with large bandwidth-delay products on all but very large transfers.

Since the average HTTP transfer size is reported to be in the 5–13KB range [Arlitt and Williamson 1996; Mogul 1995], rate-based clocking can have a significant impact on the Web. We emphasize again, however, that performance improvements attainable by a practical implementation of rate-based clocking are likely to be more modest than those indicated by the results of our experiment.

5.8 Network Polling

Our final experiment evaluates the use of soft timers for network polling. We implemented network polling in the FreeBSD-2.2.6 kernel, using soft

Table XI. Rate-Based Clocking Network Performance (Bandwidth = 100Mbps, RTT = 100 msecs)

Transfer size (1448 Byte packets)	regular TCP		rate-based clocking		
	Xput (Mbps)	Response time (msecs)	Xput (Mbps)	Response time (msecs)	Resp. time reduction (%)
5	0.16	350	0.58	100.6	71
100	1.09	1056	10.34	112	89
1000	6.38	1815	51.94	223	87
10000	38.46	3012	86.77	1335	55
100000	81.37	14235	91.92	12601	11

timers to initiate the polling. The polling interval is adaptively set to attempt to find a given number of received packet per poll interval, on average (aggregation quota).

In this experiment, a 333MHz Pentium II machine with 4 Fast Ethernet interfaces was used as the server. Four 300MHz PII machines were used as the client machines, each connected to a different interface on the server.

We measured the throughput of two different Web servers (Apache and Flash), given a synthetic workload, where clients repeatedly request the same 6KB file. The throughput was measured on an unmodified FreeBSD kernel (conventional interrupt-based network processing) and with soft-timer-based network polling. Table XII shows the results for the two different servers, for aggregation quotas ranging from 1 to 15, and for conventional (HTTP) and persistent connection HTTP (P-HTTP).

The throughput improvements with soft-timer-based polling range from 3% to 25%. The benefits of polling are more pronounced with the faster Flash server, as it stresses the network subsystem significantly more than the Apache server and, owing to its better locality, is more sensitive to cache pollution from interrupts. With P-HTTP, amortizing the cost of establishing a TCP connection over multiple requests allows much higher throughput with both servers, independent of polling.

The difference between the results for the conventional interrupt-based system and network polling with an aggregation quota of 1 (i.e., one packet per poll on average) reflects the benefit of avoiding interrupts and the associated improvement in locality. The network polling results with aggregation quotas greater than one reflect the additional benefits of aggregating packet processing.

In general, aggregation of packet processing raises concerns about increased packet delay and ACK compression. However, we believe that aggregation is practical with soft-timer-based network polling, for two reasons. Firstly, soft-timer-based network polling is turned off (and interrupts are enabled instead) whenever a CPU enters the idle loop. This ensures that packet processing is never delayed unnecessarily. Secondly, when rate-based clocking is used, packet transmissions are not paced by incoming ACKs. With rate-base clocking, it is therefore no longer necessary to preserve the exact timing of incoming ACKs, i.e., ACK compression is of lesser concern.

Table XII. Network Polling: Throughput on 6KB HTTP Requests

Aggregation	Interrupt Xput (req/sec)	Soft Poll Xput (req/sec)				
	1	1	2	5	10	15
HTTP						
Apache	854 (1.0)	915 (1.07)	933 (1.09)	939 (1.10)	944 (1.11)	945 (1.11)
Flash	1376 (1.0)	1568 (1.14)	1620 (1.17)	1690 (1.23)	1702 (1.24)	1719 (1.25)
P-HTTP						
Apache	1346 (1.0)	1380 (1.03)	1395 (1.04)	1421 (1.06)	1439 (1.07)	1440 (1.07)
Flash	4439 (1.0)	4816 (1.08)	5071 (1.14)	5271 (1.19)	5376 (1.21)	5498 (1.24)

Finally, we observe that future improvements in CPU and network speeds will continue to increase the rate of network interrupts in conventional network subsystem implementations. Since the relative cost of interrupt handling is likely to increase as CPUs get faster (see Section 5.1), avoiding interrupts becomes increasingly important.

5.9 Using On-Chip Timers to Achieve Tighter Delay Bounds on Soft-Timer Events

Recall that the probabilistic delays experienced by soft-timer events are bounded by a periodic hardware timer interrupt. However, since the rate of the hardware timer interrupt is much lower than the typical soft-timer event granularity (e.g., 1 ms versus tens of μ secs), the resulting distribution of soft-timer delays can have a long tail, as shown in Figure 2. Here, we show how an on-chip timer facility, such as the APIC found on Pentium III CPUs, can be used to obtain much tighter delay bounds for soft-timer events at low cost.

Measurements presented in Section 5.1 show that on-chip timer facilities like the PIII APIC, when compared to conventional off-chip timers, have a much reduced but still substantial per-event cost. Another advantage of the APIC facility is that timer events can be scheduled and canceled at very low cost (i.e., cost of a register access), and substantial overhead is only incurred when a timer interrupt expires. This performance aspect of on-chip timers lends itself to an integration with soft timers, yielding much tighter worst-case delays for soft-timer events at a low marginal cost.

When a soft-timer event is scheduled, in addition to specifying a scheduled time, a *deadline* is specified for the event, corresponding to the latest time when the associated event handler must be invoked. The soft-timer facility schedules an APIC event at the earliest deadline of all scheduled soft-timer events. The associated APIC event handler invokes the handlers of all soft-timer events whose deadlines have expired. If a trigger point is reached at or after the scheduled time but before the deadline of the earliest soft-timer event, then the associated handler is invoked, and the APIC timer is canceled. It is important to note that as long as deadlines are not too tight relative to the average trigger state intervals, the APIC timer rarely fires and the associated overheads remain low.

Table XIII shows the results of an experiment that quantifies the effectiveness of using APIC timer events to back up soft timers. The

Table XIII. Overhead of APIC-Based Deadlines, 500MHz PIII, Flash

Deadline (μ secs)	APIC Ovhd (%)	APIC events fired (% of all ST events)
1000	1	0 (0)
100	1.1	47 (0.006)
50	1.3	3012 (0.4)
20	2.1	77147 (10)
10	4.6	202340 (26.3)

experiment is identical to that reported in Table VII for soft timers with Flash, except that an APIC timer was scheduled to enforce a deadline of between 10 and 1000 μ secs past the scheduled event times. The results show that there is a 1% base overhead for scheduling APIC events to ensure deadlines; this overhead is due to the necessary scheduling and canceling of APIC timer events. Enforcing deadlines of 10 μ secs past the scheduled event times results in an additional overhead of 4.6% over pure soft timers without APIC deadlines; in this case, 26.3% of all soft-timer events cause the APIC timer to expire.

In summary, using APIC timers to back up soft timers allows very tight upper bounds on soft-timer delays at low cost. Enforcing deadlines of 50 and 20 μ secs past the scheduled events times, for instance, incurs an additional overhead over plain soft timers of only 1.3 and 2.1%, respectively.

5.10 Discussion

Soft timers allow the efficient scheduling of events at a granularity below that which can be provided by a conventional interval timer with acceptable overhead. The “useful range” of soft-timer event granularities is bounded on one end by the highest granularity that can be provided by a hardware interrupt timer with acceptable overhead, and on the other end by the soft-timer trigger interval. On our measured workloads on a 300 MHz PII CPU, this useful range is from a few tens of μ secs to a few hundreds of μ secs. Moreover, the useful range of soft-timer event granularities appears to widen as CPUs get faster. Our measurements on two generations of Pentium CPUs (300MHz PII and 500MHz PIII) indicate that the soft-timer event granularity increases approximately linearly with CPU speed, but that the external hardware timer overhead is almost constant.

Soft timers can be easily integrated with an existing, conventional interval timer facility. The interval timer facility provides conventional timer event services, and its periodic interrupt is also used to schedule overdue soft-timer events. Conventional timers should be used for events that need to be scheduled at or below the granularity of the interval timer’s periodic interrupt. Soft timers should be used for events that require a granularity up to the trigger state interval, provided these events can tolerate probabilistic delays up to the granularity of the conventional interval timer.

Finally, soft timers can be integrated with on-chip hardware timer facilities like the one found in the Pentium III APIC to provide fine-grained events with tight deadlines at very low overhead.

6. RELATED WORK

The implementation of soft timers is based on the idea of polling, which goes back to the earliest days of computing. In polling, a main-line program periodically checks for asynchronous events, and invokes handler code for the event if needed.

The novel idea in soft timers is to implement an efficient timer facility by making the operating system “poll” for pending soft-timer events in certain strategic states. These “trigger states” are known to be reached very frequently during execution. Furthermore, these states are associated with a shift in memory access locality, thus allowing the interposition of handler code with little impact on system performance. The resulting facility can then be used to schedule events at a granularity that could not be efficiently achieved with a conventional hardware timer facility.

Smith and Traw [1993] use periodic hardware timer interrupts to initiate polling for packets completions when using a gigabit network interface. This approach involves a trade-off between interrupt overhead and communication delay. With soft-timer-based network polling, on the other hand, one can obtain both low delay and low overhead.

Mogul and Ramakrishnan [1997] describe a system that uses interrupts under normal network load and polling under overload, in order to avoid receiver livelock. Their scheme disables interrupts during the network packet processing and polls for additional packets whenever the processing of a packet completes; when no further packets are found, interrupts are reenabled.

In comparison, soft-timer-based network polling disables interrupts and uses polling whenever the system is saturated (i.e., no CPU is idle). That is, polling is used even when the packet interarrival time is still larger than the time it takes to process packets. Moreover, soft timers allow the dynamic adjustment of the poll interval to achieve a predetermined packet aggregation quota.

A number of researchers have pointed out the benefits of rate-based clocking of TCP transmissions [Allman et al. 1997; Balakrishnan et al. 1997; Feng et al. 1999; Padmanabhan and Katz 1998; Visweswaraiah and Heidemann 1997]. Our work shows that using conventional hardware timers to support rate-based clocking at high bandwidth is too costly, and we propose soft timers as an efficient alternative.

The use of rate-based clocking has been proposed in the context of TCP slow-start, when an idle persistent HTTP (P-HTTP) connection becomes active [Fielding et al. 1997; Mogul 1995; Padmanabhan and Mogul 1994]. Visweswaraiah et al. [1997] observe that an idle P-HTTP connection causes TCP to close its congestion window, and the ensuing slow-start phase tends to defeat P-HTTP’s attempt to utilize the network more effectively than

HTTP/1.0 [Berners-Lee et al. 1996] connections. A similar observation was made by Padmanabhan et al. [1998]. Soft timers can be used to efficiently clock the transmission of packets upon restart of an idle P-HTTP connection.

Allman et al. [1997] show the limiting effect of slow-start and congestion avoidance schemes in TCP in utilizing the bandwidth over satellite networks. Using rate-based clocking instead of slow-start addresses the former concern. Feng et al. [1999] propose the use of rate-based clocking in TCP to support the controlled-load network service [Wroclawski 1997], which guarantees a minimal level of throughput to a given connection.

Balakrishnan et al. [1997] have proposed *ACK filtering*, a mechanism that attempts to improve TCP performance on asymmetric network paths by discarding redundant ACKs at gateways. They observe that this method can lead to burstiness due to the big ACKs seen by the sender and suggest pacing packet transmissions so as to match the connection's sending rate.

Besides an efficient timer mechanism, rate-based clocking also depends on mechanisms that allow the measurement or estimation of the available network capacity. A number of techniques have been proposed in the literature. The basic packet-pair technique was proposed by Keshav [1991]. Hoe [1996] proposes methods to improve TCP's congestion control algorithms. She sets the slow-start threshold (ssthresh) to an appropriate value by measuring the bandwidth-delay product using a variant of the packet-pair technique. Paxson [1997] suggests a more robust capacity estimation technique called PBM that forms estimates using a range of packet bunch sizes. A technique of this type could be used to support rate-based clocking. Allman and Paxson [1999] compare several estimators and find that sender-side estimation of bandwidth can often give inaccurate results due to the failure of the ACK stream to preserve the spacing imposed on data segments by the network path. They propose a receiver-side method for estimating bandwidth that works considerably better.

7. CONCLUSIONS

This paper proposes a novel operating system timer facility that allows the system to efficiently schedule events at a granularity down to tens of microseconds. Such fine-grained events are necessary to support rate-based clocking of transmitted packets on high-speed networks and can be used to support efficient network polling.

Unlike conventional timer facilities, soft timers take advantage of certain states in the execution of a system where an event handler can be invoked at low cost. In these states, the saving and restoring of CPU state normally required upon a hardware timer interrupt is not necessary, and the cache/TLB pollution caused by the event handler is likely to have low impact on the system performance.

Experiments with a prototype implementation show that soft timers can be used to perform rate-based clocking in TCP at granularities down to a few tens of microseconds. At these rates, soft timers impose an overhead of

only 2–6% while a conventional timer facility would have an overhead of 13–36%.

Soft timers can also be used to perform network polling, thus avoiding network interrupts while preserving low communications delays. Experiments show that the performance of a Web server using this optimization can increase by up to 25% over a conventional interrupt-based implementation.

Furthermore, the performance improvements obtained with soft timers can be expected to increase with network and CPU speeds. As networks and CPUs get faster, so does the rate of network interrupts. However, the speed of interrupt handling does not increase as fast as CPU speed, due to its poor memory access locality. The relative cost of interrupt handling therefore increases, underscoring the need for techniques that avoid interrupts.

Soft-timer performance, on the other hand, appears to scale with CPU speed. Soft timers are cache friendly, and faster CPU speeds imply that trigger states are reached more frequently, thus improving the granularity at which soft timers can schedule events. Finally, soft timers can be integrated with the on-chip timer facilities found on modern CPUs to provide fine-grained events with very tight delay bounds at low overhead.

APPENDIX

A. RATE-BASED CLOCKING

In this appendix, we provide further motivation for rate-based clocking. We restrict ourselves here to a general discussion of how an appropriate timer facility can be used for rate-based clocking of transmissions. The details of how a specific protocols like TCP should be extended to add rate-based clocking requires further research and are beyond the scope of this paper.

A.1 ACK Compression and Big ACKs

Previous work has demonstrated the phenomenon of *ACK compression*, where ACK packets from the receiver lose their temporal spacing due to queuing on the reverse path from receiver to sender [Mogul 1993; Zhang et al. 1991]. ACK compression can cause bursty packet transmissions by the TCP sender, which contributes to network congestion. Balakrishnan et al. [1998] have observed the presence of ACK compression in a busy Web server.

With rate-based clocking, a TCP sender can keep track of the average arrival rate of ACKs. When a burst of ACKs arrives at a rate that significantly exceeds the average rate, the sender may choose to pace the transmission of the corresponding new data packets at the measured average ACK arrival rate, instead of the burst's instantaneous rate as would be dictated by self-clocking.

A related phenomenon is that of *big ACKs*, i.e., ACK packets that acknowledge a large number of packets or update the flow-control window

by a large number of packets. Upon receiving a big ACK, self-clocked senders may send a burst of packets at the bandwidth of the network link adjacent to the sender host. Transmitting such bursts can adversely affect congestion in the network. A detailed discussion of phenomena that can lead to big ACKs (i.e., ACKs that can lead to the transmission of more than 3 packets) in TCP is given in Section A.3.

Using rate-based clocking, it is possible to avoid sending packet bursts in the same way as was described above in connection with ACK compression.

A.2 Slow-Start

Self-clocked protocols like TCP use a *slow-start* phase to start transmitting data at the beginning of a connection or after an idle period. During slow-start, the sender transmits a small number of packets (typically two), and then transmits two more packets for every acknowledged packet, until either packet losses occur or the estimated network capacity is reached. In this way, the sender increases the amount of data transmitted per RTT exponentially until the network capacity is reached.

The disadvantage of slow-start is that despite the exponential growth of the transmit window, it can take many RTTs before the sender is able to fully utilize the network. The larger the bandwidth-delay product of the network, the more time and transmitted data it takes to reach the point of network saturation. In particular, transmissions of relatively small data objects may not allow the sender to reach the point of network saturation at all, leading to poor network utilization and low effective throughput.

The bulk of traffic in the Internet today consists of HTTP transfers that are typically short (between 5KB and 13KB) [Arlitt and Williamson 1996; Mogul 1995]. A typical HTTP transfer finishes well before TCP finishes its slow-start phase, causing low utilization of available network bandwidth and long user-perceived response times [Mogul 1995]. The magnitude of this problem is expected to increase as higher network bandwidth becomes available.

Slow-start serves a dual purpose. It starts a transmission pipeline that allows the sender to self-clock its transmission without sending large bursts of packets. At the same time, it probes the available network capacity without overwhelming the network. The key idea to avoid slow-start is the following. If the available network capacity is known or can be measured/estimated, then a TCP sender can immediately use rate-based clocking to transmit packets at the network capacity without going through slow-start [Padmanabhan and Katz 1998].

The problem of measuring available network capacity has been addressed by several prior research efforts, for instance packet pair algorithms [Brakmo and Peterson 1995b; Hoe 1996; Keshav 1991] and PBM [Paxson 1997]. Moreover, when starting transmission after an idle period, the network capacity during the last busy period can be used as an estimate for the current capacity [Fielding et al. 1997; Mogul 1995; Padmanabhan and Mogul 1994]. Finally, in future network with QoS support, the available network capacity may be known *a priori*.

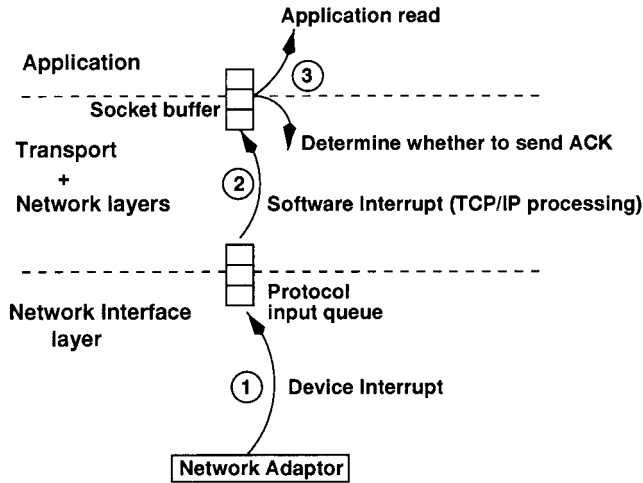


Fig. 5. Packet-processing path in OS.

A.3 Causes of Big ACKs

In the previous subsection, we discussed the effects of big ACKs on TCP connections. Here, we describe several phenomena that can cause big ACKs.

Figure 5 shows the processing of a packet, starting from its reception by the network adaptor to its delivery to the application. (1) A high-priority device interrupt places the packet into the input queues of the IP protocol, (2) TCP/IP processing is done in the context of a software interrupt, and the packet is placed in the application's socket buffer, (3) The application reads the data from its socket; in the context of this read, an ACK is sent back to the TCP sender if needed.

Upon reception of a packet acknowledging x packets, a TCP sender normally injects x new closely spaced packets into the network. In normal operation, x is 2 because TCP receivers usually delay every other ACK⁷ to take advantage of piggybacking opportunities. We now present some scenarios that cause a TCP receiver to send big ACKs (ACKs that acknowledge more than 3 packets), causing the sender to inject a burst of packets that can adversely affect congestion in the network.

Figure 5 indicates that an ACK is sent by the receiver when the application reads the data from the socket buffer (or when the delayed ACK timer fires). If the interarrival time of packets is smaller than the packet-processing time, then owing to the higher priorities of the interrupts as compared to application processing, all closely spaced packets sent by the TCP sender will be received before any ACK is sent. When the incoming packet train stops (due to flow control), the receiver will send a big ACK to the sender acknowledging all packets sent. The same happens if the

⁷The presence of TCP options causes TCP receivers to send an ACK for every 3 packets [Brakmo and Peterson 1995a].

delayed ACK timer fires first. The problem is self-sustaining because the TCP sender responds to the big ACK by sending a burst of closely spaced packets.

On a 300MHz Pentium II machine, the packet processing time can take more than 100 μ secs while the minimum interarrival time of 1500-byte packets on 100Mbps and 1Gbps Ethernet is 120 μ secs and 12 μ secs, respectively. This suggests that big ACKs can be prevalent in high-bandwidth networks.

The situation described above is not necessarily restricted to high-bandwidth networks. It can also happen when the receiver application is slow in reading newly arrived data from the socket buffers. This can happen, for example, when a Web browser (TCP receiver) is rendering previously read graphics data on the screen. During this time, ACKs for all packets from the Web server (TCP sender) shall be delayed until either the delayed ACK timer fires (once every 200 ms) or the browser reads more data from the socket buffer. The ACK packet when sent would acknowledge a large number of packets.

While high bandwidth is not yet widely available in WANs, we have analyzed TCP packet traces on a 100Mbps LAN and have observed big ACKs on almost every sufficiently long transfer. We have also analyzed packet traces from the Rice CS departmental Web server. Our results show that 40% of all transfers that were greater than 20KB showed the presence of big ACKs, thus confirming our hypothesis that big ACKs also occur on transfers over current low-bandwidth WAN links.

Brakmo and Peterson [1995a] have also observed these big ACKs in the context of recovery from large number of packet losses and reordering of packets. They propose to reduce TCP congestion window upon receiving a big ACK so that slow-start is used instead of sending packet bursts. Fall and Floyd [1996] propose to use a maxburst parameter to limit the potential burstiness of the sender for packets sent after a loss recovery phase (fast recovery). While these techniques can limit the burstiness, they adversely affect bandwidth utilization as the network pipeline is drained of packets.

ACKNOWLEDGMENTS

We are grateful to the anonymous TOCS and SOSP reviewers and our SOSP shepherd, Mendel Rosenblum, for their helpful comments. Thanks to Sitaram Iyer and Erich Nahum for their help with some experiments.

REFERENCES

- ALLMAN, M. AND PAXSON, V. 1999. On estimating end-to-end network path properties. In *Proceedings of the SIGCOMM '99 Conference*. ACM, New York, 263–274.
- ALLMAN, M., HAYES, C., KRUSE, H., AND OSTERMANN, S. 1997. TSP performance over satellite links. In *Proceedings of the 5th International Conference on Telecommunications Systems*. 456–469.

- ARLITT, M. F. AND WILLIAMSON, C. L. 1996. Web server workload characterization: The search for invariants. In *Proceedings of the ACM SIGMETRICS '96 Conference*. ACM, New York, 126–137.
- BALAKRISHNAN, H., PADMAHABHAN, V. N., AND KATZ, R. H. 1997. The effects of asymmetry on TCP performance. In *Proceedings of the 3rd ACM Conference on Mobile Computing and Networking*. ACM, New York, 77–89.
- BALAKRISHNAN, H., PADMAHABHAN, V. N., SESHAN, S., STEMM, M., AND KATZ, R. H. 1998. TCP behavior of a busy internet server: Analysis and improvements. In *Proceedings of IEEE INFOCOM '98*. IEEE, New York, 252–262.
- BERNERS-LEE, T., FIELDING, R., AND FRYSTYK, H. 1996. RFC 1945: Hypertext transfer protocol—HTTP/1.0. <ftp://ftp.merit.edu/documents/rfc/rfc1945.txt>.
- BRAKMO, L. AND PETERSON, L. 1995a. Performance problems in 4.4BSD TCP. *ACM Comput. Commun. Rev.* 25, 5 (Oct.), 69–86.
- BRAKMO, L. AND PETERSON, L. 1995b. TCP Vegas: End to end congestion avoidance on a global Internet. *IEEE J. Sel. Areas Commun.* 13, 8 (Oct.), 1465–1480.
- FENG, W. C., KANDLUR, D., SAHA, D., AND SHIN, K. G. 1999. Understanding and improving TCP performance over networks with minimum rate guarantees. *IEEE/ACM Trans. Networking* 7, 2 (Apr.), 173–187.
- FALL, K. AND FLOYD, S. 1996. Simulation-based comparisons of Tahoe, Reno, and SACK TCP. *Comput. Commun. Rev.* 26, 3 (July), 5–21.
- FIELDING, R., GETTYS, J., MOGUL, J., NIELSEN, H., AND BERNERS-LEE, T. 1997. RFC 2068: Hypertext transfer protocol—HTTP/1.1. <ftp://ftp.merit.edu/documents/rfc/rfc2068.txt>.
- HOE, J. C. 1996. Improving the start-up behaviour of a congestion controlscheme. In *Proceedings of the ACM SIGCOMM '96 Symposium*. ACM, New York, 270–280.
- INTEL. 2000. Pentium III processors—manuals. <http://developer.intel.com/design/pentiumiii/manuals/>.
- KESHAV, S. 1991. A control-theoretic approach to flow control. In *Proceedings of the ACM SIGCOMM '91 Symposium*. ACM, New York, 3–15.
- MOGUL, J. C. 1993. Observing TCP dynamics in real networks. In *Proceedings of the ACM SIGCOMM '92 Symposium*. ACM, New York, 281–292.
- MOGUL, J. C. 1995. The case for persistent-connection HTTP. In *Proceedings of the ACM SIGCOMM '95 Symposium*. ACM, New York, 299–313.
- MOGUL, J. C. AND RAMAKRISHNAN, K. K. 1997. Eliminating receive livelock in an interrupt-driven kernel. *ACM Trans. Comput. Syst.* 15, 3 (Aug.), 217–252.
- PADMAHABHAN, V. N. AND KATZ, R. H. 1998. TCP fast start: A technique for speeding up Web transfers. In *Proceedings of the IEEE GLOBECOM '98 Conference*. IEEE, New York, 41–46.
- PADMANABHAN, V. N. AND MOGUL, J. C. 1994. Improving HTTP latency. In *Proceedings of the 2nd International WWW Conference*. 995–1005.
- PAI, V. S., DRUSCHEL, P., AND ZWAENEPOEL, W. 1999. Flash: An efficient and portable Web server. In *Proceedings of the USENIX 1999 Annual Technical Conference*. USENIX Assoc., 199–212.
- PAXSON, V. 1997. End-to-end Internet packet dynamics. In *Proceedings of the ACM SIGCOMM '97 Symposium*. ACM, New York, 139–152.
- SMITH, J. M. AND TRAW, C. B. S. 1993. Giving applications access to Gb/s networking. *IEEE Network* 7, 4 (July), 44–52.
- VARGHESE, G. AND LAUCK, A. 1987. Hashed and hierarchical timing wheels: Data structures for the efficient implementation of a timer facility. In *Proceedings of the 11th ACM Symposium on Operating Systems Principles*. ACM, New York, 171–180.
- VISWESWARAIAH, V. AND HEIDEMANN, J. 1997. Improving restart of idle TSP connections. Tech. Rep. 97-661. Univ. of Southern California.
- WROCLAWSKI, J. 1997. RFC 2211: Specification of controlled-load network element service. <ftp://ftp.merit.edu/documents/rfc/rfc2211.txt>.
- ZHANG, L., SHENKER, S., AND CLARK, D. D. 1991. Observations on the dynamics of a congestion control algorithm: The effects of two-way traffic. In *Proceedings of the ACM SIGCOMM '91 Symposium*. ACM, New York, 133–148.