

Differentiated Predictive Fair Service for TCP Flows

IBRAHIM MATTA LIANG GUO

Computer Science Department
Boston University
Boston, MA 02215, USA

{matta, guo}@cs.bu.edu

Technical Report BU-CS-2000-012

May 2000

Abstract

The majority of the traffic (bytes) flowing over the Internet today have been attributed to the Transmission Control Protocol (TCP). This strong presence of TCP has recently spurred further investigations into its congestion avoidance mechanism and its effect on the performance of short and long data transfers. At the same time, the rising interest in enhancing Internet services while keeping the implementation cost low has led to several service-differentiation proposals. In such service-differentiation architectures, much of the complexity is placed only in access routers, which classify and mark packets from different flows. Core routers can then allocate enough resources to each class of packets so as to satisfy delivery requirements, such as predictable (consistent) and fair service.

In this paper, we investigate the *interaction* among short and long TCP flows, and how TCP service can be improved by employing a low-cost service-differentiation scheme. Through control-theoretic arguments and extensive simulations, we show the utility of *isolating* TCP flows into two classes based on their lifetime/size, namely one class of short flows and another of long flows. With such class-based isolation, short and long TCP flows have separate service queues at routers. This protects each class of flows from the other as they possess different characteristics, such as burstiness of arrivals/departures and congestion/sending window dynamics. We show the benefits of isolation, in terms of better predictability and fairness, over traditional shared queueing systems with both tail-drop and Random-Early-Drop (RED) packet dropping policies. The proposed class-based isolation of TCP flows has several advantages: (1) the implementation cost is low since it only requires core routers to maintain per-class (rather than per-flow) state; (2) it promises to be an effective traffic engineering tool for improved predictability and fairness for both short and long TCP flows; and (3) stringent delay requirements of short interactive transfers can be met by increasing the amount of resources allocated to the class of short flows.

Keywords: Differentiated service; TCP congestion control; class-based isolation versus sharing; packet dropping (tail-drop, RED); fairness; control theory; simulation.

I. INTRODUCTION

In order to improve Internet services, enhancements to the Internet's basic best-effort architecture have been recently proposed, most notably, the Intserv and Diffserv architectures. In the Intserv architecture [1], routers are stateful, i.e. they maintain state information about each communication flow of packets. Thus, per-flow delivery guarantees can be provided at the expense of complexity at routers. The Diffserv architecture [2], [3], [4], [5], [6], [7], [8], [9], [10], [11], [12], [13], [14], however, requires only access (border) routers at the edge of the network to maintain per-flow information, whereas routers in the core of the network are kept essentially stateless. In this paper, we concern ourselves with Diffserv-like architectures and how they can best manage TCP flows¹, which constitute the majority of the traffic (80-90%) on the Internet today [15].

Recent measurements of Internet traffic [16] show that the length (in terms of both lifetime and size) of TCP flows follows a heavy-tailed distribution, i.e. only a small percentage (e.g. less than 20%) of flows are long-lived (e.g. more than 20 packets), but they carry a large percentage (e.g. 85%) of the total traffic (or bytes). This calls for a careful design of the network so that these few long-lived TCP flows are managed well. Furthermore, many recent proposals (e.g. [5], [17]) have attempted to provide long-lived TCP flows with a predictive (or controlled-load) service, where a TCP flow is statistically allocated a certain rate/bandwidth. In these proposals, it has been shown that a TCP flow has to operate in a predictable mode, otherwise it may not be able to take advantage of its reservation. This predictability is hard to achieve in the presence of the many short-lived TCP flows, which have a more bursty flow arrival/departure process [18], [15] and more drastic sending window dynamics.

In this paper, we argue for an architecture that *isolates* short-lived and long-lived TCP flows into two classes. This can be implemented, for example, by using a class-based queueing (CBQ) scheme [19], or by routing these two classes of flows on (logically) separate communication paths [7]. See Figure 1 for an illustration of isolation control. Edge routers will be responsible for classifying flows and marking packets as belonging to long-lived or short-lived flow. Once a flow is classified into a long flow (e.g. after a time threshold expires), edge routers will be able to direct the recognized flow to a new path, for example, by establishing a label-switched path using MPLS [20], [21], [7], [22], [23].

We show using control-theoretic arguments and extensive simulations the utility of such class-based isolation scheme. We show that fairness and predictability are improved for both short and long TCP flows. Furthermore, by allocating more bandwidth to the

This work was supported in part by NSF grants CAREER ANI-0096045 and MRI EIA-9871022.

Corresponding author: Ibrahim Matta, Assistant Professor, Computer Science Department, Boston University, 111 Cummington Street, MCS-271, Boston, MA 02215, USA. Tel: (617) 358-1062, fax: (617) 353-6457, e-mail: matta@cs.bu.edu

¹We use the terms "flow" and "connection" interchangeably throughout the paper. A flow can be generally defined as a sequence of packets which share some common properties. In this paper, we define a flow as packet sequence between the same source host/network and destination host/network pair.




Fig. 1. Edge routers perform isolation control

class of short flows, which are usually interactive and delay-sensitive, we can significantly reduce their response time. Our work differs from previous work (reviewed in Section V) in that we study the benefits of class-based isolation to not only long TCP transfers², but also to short TCP transfers, which constitute a large percentage of TCP flows on the Internet today.

The rest of the paper is organized as follows. Section II uses control-theoretic arguments (similar to [24], [25]) to establish the benefits of isolation (or individual) control over traditional aggregate control, where flows of different characteristics (e.g. size) are mixed together to share the same resources (e.g. buffer, transmission capacity). Section III describes our simulation model and experiments. Section IV supports our claims by extensive simulations of TCP-reno flows. Section V discusses related work, and Section VI concludes the paper. The derivation of the stability condition for the aggregate control model of Section II is given in an Appendix, which also contains a brief description of the TCP congestion control algorithm [26].

II. AGGREGATE VERSUS INDIVIDUAL CONTROL

In this section, we formulate two discrete-time control models. In the first model (Section II-A), each flow controls its input traffic rate based on the *aggregate* state of the network due to all N flows. In the second model (Section II-B), each flow (or class of homogeneous flows) controls its rate based on its own *individual* state within the network. Here we do not model TCP flows whose control generally leads to oscillatory behavior [27], rather we assume PID control³ for which a stable solution exists under certain conditions [28], [25], [24]. The motivation behind this assumption is to keep the models simple while still gaining insights into the fundamental differences between aggregate and individual controls, and how these differences can affect the performance of TCP flows of different size.

In the aggregate control model, the number of new packets that a flow i sends into the network at time step k , denoted by w_i^k , is adapted based on the difference between a target *total* buffer space, denoted by B , and the current *total* number of outstanding packets, denoted by q^k . In the individual control model, w_i^k is adapted based on flow (or class) i 's target, denoted by B_i , and its current number of outstanding packets, denoted by q_i^k . We denote by c^k the total number of packets acknowledged at time step k , and by c_i^k the number of flow/class i packets acknowledged at step k . In what follows, for each control model, we determine conditions under which the system stabilizes. We then solve for the values of the state variables at equilibrium, and show whether fairness (or a form of weighted resource sharing) can be achieved. Table I lists all system variables along with their meanings.

TABLE I
TABLE DEFINING SYSTEM VARIABLES

Variable	Meaning
N	total number of flows (or classes of homogeneous flows)
w_i^k	the number of new packets of flow/class i at step k
q_i^k	the number of outstanding packets of flow/class i at step k
c_i^k	the number of flow/class i packets acknowledged (i.e. no longer outstanding) at step k
q^k	the total number of outstanding packets at step k
c^k	the total number of packets acknowledged (served) at step k
B	the target total buffer space
B_i	target buffer space allocated to flow/class i
α_i, β_i	parameters controlling the increase and decrease rates of w_i^k

A. Aggregate Control or Sharing

Under aggregate PID control, the evolution of the system state is described by the following difference equations:

$$\begin{aligned}
 w_i^k &= w_i^{k-1} + \alpha_i(B - q^k) + \beta_i[(B - q^k) - (B - q^{k-1})] \\
 q^k &= q^{k-1} + \sum_{i=1}^N w_i^{k-1} - c^{k-1}
 \end{aligned} \tag{1}$$

²Many TCP studies assume *infinite* transfers.

³PID control is a controller where the control signal is a linear combination of the error, its integral and derivative [28].

For simplicity, assume a constant acknowledgment (service) rate, i.e. $c^k = C$ for all k . Let α and β be the maximum of α_i and β_i , respectively. Then, it can be shown that this system is stable if the following condition is satisfied:

$$\alpha + 2\beta < \frac{4}{N} \quad (2)$$

Otherwise, the system does not converge to a stable state. This stability condition is derived in the Appendix. This indicates that the existence of flows which rapidly change their sending rates through high values of α_i and/or β_i can cause the system to become unstable. This suggests that short TCP flows, which aggressively change their sending windows in slow start phase, may affect the stability of long TCP flows, which change their sending windows cautiously in congestion avoidance mode, in a system that mixes both kinds of TCP flows. Furthermore, in such a system, the value of N may be high so as to cause instability.

We now derive the values of the state variables at equilibrium given that the system is stable, i.e. the system satisfies equation (2). Denote by $(w_i)_s$ and q_s the steady-state values of w_i^k and q^k , respectively. Then, at equilibrium, we have from equations (1):

$$\begin{aligned} (w_i)_s &= (w_i)_s + \alpha_i(B - q_s) + \beta_i[(B - q_s) - (B - q_s)] \\ q_s &= q_s + \sum_{i=1}^N (w_i)_s - C \end{aligned} \quad (3)$$

Thus, at equilibrium, $q_s = B$ and $\sum_{i=1}^N (w_i)_s = C$. In other words, the system converges to a state where the total input rate matches the total service rate, and the total buffer space is full.

We note that if $\alpha_i = \alpha$ and $\beta_i = \beta$ for all i , then w_i^k changes by the same amounts for every flow i . Consequently, if we start the evolution of the system with w_i^0 being the same for all flows, only then we have equal sharing of the network at steady-state, i.e. $(w_i)_s = \frac{C}{N}$, regardless of the initial values of q_i^0 . However, in general, when w_i^0 are not equal for all flows, the system converges to an *unfair* state, more precisely, to a state where

$$(w_i)_s = w_i^0 + \frac{C - \sum_{j=1}^N w_j^0}{N} \quad (4)$$

To summarize, controlling several flows by observing the resulting aggregate state of the network may lead to instability due to either the existence of flows which are rapidly adjusting their sending rates, or a high number of flows competing for the same *shared* resource. Furthermore, even if the system is stable, the system is highly likely to converge to an unfair state where flows receive unequal shares of the resource.

B. Individual Control or Isolation

Under individual PID control, the evolution of the system state is described by the following difference equations:

$$\begin{aligned} w_i^k &= w_i^{k-1} + \alpha_i(B_i - q_i^k) + \beta_i[(B_i - q_i^k) - (B_i - q_i^{k-1})] \\ q_i^k &= q_i^{k-1} + w_i^{k-1} - c_i^{k-1} \end{aligned} \quad (5)$$

Recall that under individual control, flow/class i regulates its input, w_i^k , based on its *own* number of outstanding and acknowledged packets. For simplicity, assume a constant acknowledgment (service) rate, i.e. $c_i^k = C_i$ for all k . Following the same stability analysis in the Appendix to derive equation (2), it is easy to see that flow/class i stabilizes if the following condition is satisfied:⁴

$$\alpha_i + 2\beta_i < 4 \quad (6)$$

Observe that, unlike aggregate control, flows/classes are *isolated* from each other. Therefore, the existence of flows/classes, which rapidly change their sending rates through high values of α_i and/or β_i (e.g. short TCP flows), does not affect the stability of other flows (e.g. long TCP flows). This isolation can be implemented using, for example, a class-based queueing (CBQ) discipline [19]. In such a CBQ system, each class of homogeneous flows can be allocated its own buffer space and service capacity.

We now derive the values of the state variables of flow/class i at equilibrium given that it stabilizes, i.e. flow/class i satisfies equation (6). Denote by $(w_i)_s$ and $(q_i)_s$ the steady-state values of w_i^k and q_i^k , respectively. Then, at equilibrium, we have from equations (5):

$$\begin{aligned} (w_i)_s &= (w_i)_s + \alpha_i(B_i - (q_i)_s) + \beta_i[(B_i - (q_i)_s) - (B_i - (q_i)_s)] \\ (q_i)_s &= (q_i)_s + (w_i)_s - C_i \end{aligned} \quad (7)$$

Thus, at equilibrium, $(q_i)_s = B_i$ and $(w_i)_s = C_i$. In other words, each flow/class i converges to a state where its input rate matches its allocated service rate, and its allocated buffer space is full. We note that if the allocated buffers B_i and service capacities C_i

⁴We set $N = 1$ in equation (2).

are equal, then every flow receives an equal share of the resources, *regardless of* the initial values of w_i^0 and q_i^0 . One can also achieve a weighted resource sharing by assigning different B_i and C_i allocations. Thus, a flow/class with higher priority (e.g. short interactive TCP flows) can be allocated more resources, so as to receive better throughput/delay service.

To summarize, controlling each flow (or class of homogeneous flows) by observing its own individual state within the network provides isolation between them. Thus, stability can be achieved for a flow/class regardless of the behavior and number of other flows/classes. Furthermore, the system can converge to a fair state where flows/classes receive a weighted share of the resources.

III. SIMULATION MODEL AND EXPERIMENTS

As pointed out in Section I, recent measurements of Internet traffic [16] show that the length (in terms of both lifetime and size) of TCP flows follows a heavy-tailed distribution, where few long-lived TCP flows carry most of the bytes. Short flows behave very differently compared to long flows. First, some measurement studies [18], [15] suggest that short-lived flows arrive to the network in a more bursty fashion than long-lived flows. In addition, for TCP traffic, which contributes most of the Internet traffic today, long-lived flows, which acquire enough knowledge about the congestion state of the network, spend most of their time in *congestion avoidance* phase, while short-lived flows mainly transmit in *slow start* phase [29]. In other words, short flows generally finish their transmission before they can adapt to their fair share of the bandwidth. Generally speaking, the TCP congestion window changes more drastically during slow start than during congestion avoidance. However, because short TCP flows generally attain smaller window sizes, they generate smaller packet bursts (albeit more variable in size) than long TCP flows.

These differences in the behavior of short and long TCP flows suggest that it may be beneficial to treat network flows differently based on their size. More precisely, isolating short flows from long ones promises to reduce the volatility (seen by long flows) of the state of congestion in the network created by the burstier arrivals of short TCP flows and their more drastic window dynamics (i.e. variation in packet burst sizes). Thus, the few well-behaving (more stable) long-lived TCP flows, which are carrying most of the bytes, can be protected and provided better service.

Furthermore, with isolation, short TCP flows are also protected from long TCP flows, which generally attain larger window sizes (or generate larger packet bursts). This reduces the chances that short TCP flows are completely shut off by long flows, which would increase their response time. This is clearly undesirable for short interactive/delay-sensitive flows. This undesirable situation has been observed when both shared tail-drop queues and Random-Early-Drop (RED) queues are deployed [30], [31], where sources generating smaller packet bursts (i.e. short flows with smaller window sizes) are penalized.

Our objective is to investigate the effect of isolating (or splitting) a set of short and long flows into two (size-homogeneous) classes, namely short class and long class, in the absence and presence of background traffic. We then measure various performance metrics, including fairness within each class. We obtain our simulation results using the UCB/LBNL/VINT Network Simulator—*ns* (version 2) [32]. All simulations are for the TCP-reno version [33], [34]. Simulations with TCP-tahoe support the same conclusions. In our experiments, we consider links with FIFO queues employing either tail-drop or RED packet dropping policy. Unless otherwise specified, packet queues are assumed to deploy a tail-drop policy.

A. Flow Assignment Policies

We compare three traffic management schemes employed at an edge router distributing flows over two parallel paths (cf. Figure 1): (1) a *load-balanced strategy* where incoming flows are randomly distributed with equal probability over the two paths; (2) a *size-based splitting strategy* where short flows are routed on one path (Path 1) and long flows on the other (Path 2); and (3) a *threshold-based splitting strategy* where each flow is first routed on Path 1, then if the flow is still active after some time or a certain number of packets (e.g. 30 packets are observed from that flow), that flow is considered long and is routed on Path 2.

When a splitting (isolation) policy is used, unless otherwise specified, a size-based splitting strategy is used. Note that size-based splitting assumes that TCP flows are classified based on the exact size of the transfer, rather than by classifying a flow as long-lived after observing it for some time or for a certain number of packets. The latter mechanism, used in threshold-based splitting, would likely be the one used in practice [16]. However, it is not straightforward to quantify the effect of a re-routed long flow on the ongoing group of long flows. Thus, in this paper, we also use size-based splitting to exclude these effects and examine isolation if we indeed have a “clean” split among short and long flows.

B. Performance Measures

We consider the following performance measures: (1) *average utilization and deviation* on each path to measure the amount and variability of the load; (2) *effective throughput (or goodput)* to measure the rate of successfully transmitted (i.e. acknowledged) packets; and (3) *fairness* within each class (i.e. among size-homogeneous flows). One way to measure fairness is in terms of the distribution of transmission (response) time (which is inversely proportional to the effective throughput). In particular, flows of about the same size should be transmitted in roughly the same time. Another way to measure fairness is using Chiu and Jain’s fairness index [27], which is defined as follows: if there are N connections competing for a bottleneck resource, and the goodput achieved by connection i is x_i , then the fairness index f is given by:

$$f = \frac{(\sum_{i=1}^N x_i)^2}{N \sum_{i=1}^N x_i^2}$$

The closer the value of f is to 1, the more fair the resource allocation is (i.e. the values of x_i 's are closer to each other). Throughout simulation lifetime, we measure the goodput and fairness index over 20-second intervals every 20 seconds, and either plot these instantaneous values as a function of time, or plot the average value.

IV. SIMULATION RESULTS

Simulation results in Sections IV-B and IV-C motivate isolating long TCP flows from short ones so as to protect the long class from the more bursty arrival/departure process and the more drastic changes in window sizes of flows of the short class. Sections IV-D and IV-E compare isolation and sharing for varying network pipe and buffer sizes in the presence of telnet background traffic. Section IV-F presents performance results with web background traffic.

A. General Observations

Before presenting the details of our simulation experiments, we summarize our observations:

- Isolating short flows from long ones can reduce the load variation on the route taken by long flows. This is because long flows are then protected from the more bursty arrivals/departures of short flows and their more drastic sending window dynamics. Thus isolation can improve the predictability of the service provided to long flows as well as fairness among them.
- Class-based isolation of short and long TCP flows also protects short flows from being completely shut off by long TCP flows, which generally generate larger packet bursts due to their larger sending windows. Thus isolation can also improve the predictability of the service provided to short flows as well as fairness among them.
- Class-based isolation provides short TCP flows with much better fairness and response time than that provided in a shared RED queueing system, since RED usually penalizes flows generating smaller packet bursts.
- Class-based isolation can improve service predictability and fairness without sacrificing the overall goodput and utilization of the network. This is especially true when network resources are relatively scarce—generally when the share of each connection is less than 10 packets per round-trip time.
- Class-based isolation can allow a service provider to more accurately predict performance and thus provide reliable guarantees to users, including throughput/delay guarantees to short interactive transfers.

B. Isolating Burstiness of Short Flows

Figure 2 shows simulations for (non-flow-controlled) UDP connections (flows) of different sizes. The goal is to examine the effect of the burstiness of the arrival/departure process of flows of different size, in the absence of the effect of flow/congestion control implemented by TCP flows. We assume that the length of each UDP flow follows a Pareto distribution with shape 1.25, and average flow length of 300 seconds. The inter-arrival time of flows also follows a Pareto distribution with shape 1.2 and average interarrival time of 1 second. We model each active flow as a constant-bit-rate traffic sending at 16 Kbps packets of size 125 bytes. We compare the three traffic management schemes, described in Section III-A, distributing flows over two parallel paths (cf. Figure 1). For size-based splitting, we simply assume that a flow is short if its length is less than the average. For threshold-based splitting, we define the threshold to be half of the average flow length. Table II shows the average utilization and deviation on each path.

TABLE II

UDP EXPERIMENTS: EFFECT OF BURSTY ARRIVALS/DEPARTURES OF SHORT FLOWS. UNDER A SPLITTING (ISOLATION) STRATEGY, LONG FLOWS ARE ROUTED ON PATH 2, WHILE SHORT FLOWS ARE ON PATH 1.

Distribution Strategy	Path 1		Path 2	
	average	deviation	average	deviation
Load Balancing	0.23089	0.07194	0.20789	0.07251
Size-Based Splitting	0.17765	0.09144	0.26113	0.06290
Threshold-Based Splitting	0.19679	0.10407	0.24199	0.05914

These experiments suggest that the arrival/departure process for long flows is less bursty than short flows; this is consistent with [18]. Therefore, splitting long and short flows can reduce the load variation on the route taken by long flows, thus improving the predictability of their service and fairness among them.

C. Isolating Window Dynamics of Short TCP Flows

Figure 3(a) shows the sending windows of TCP connections. Eight TCP connections of different sizes share a single 1.5 Mbps bottleneck link. A flow that terminates is replaced by another flow of the same size. The evolution of the window sizes of three of the TCP connections is shown—a long flow of size 20000 packets, a medium-size flow of 200 packets and a short flow of size 20 packets. We observe that longer connections are the most affected by packet losses, and are prevented from sending at the maximum window size of 64 packets or from stabilizing to their fair share of the bandwidth. Furthermore, a short flow can sometimes be completely shut off as seen by the window value staying at 1 packet. Thus, predictability and fairness can be much improved by isolating long-lived TCP flows from short-lived ones.

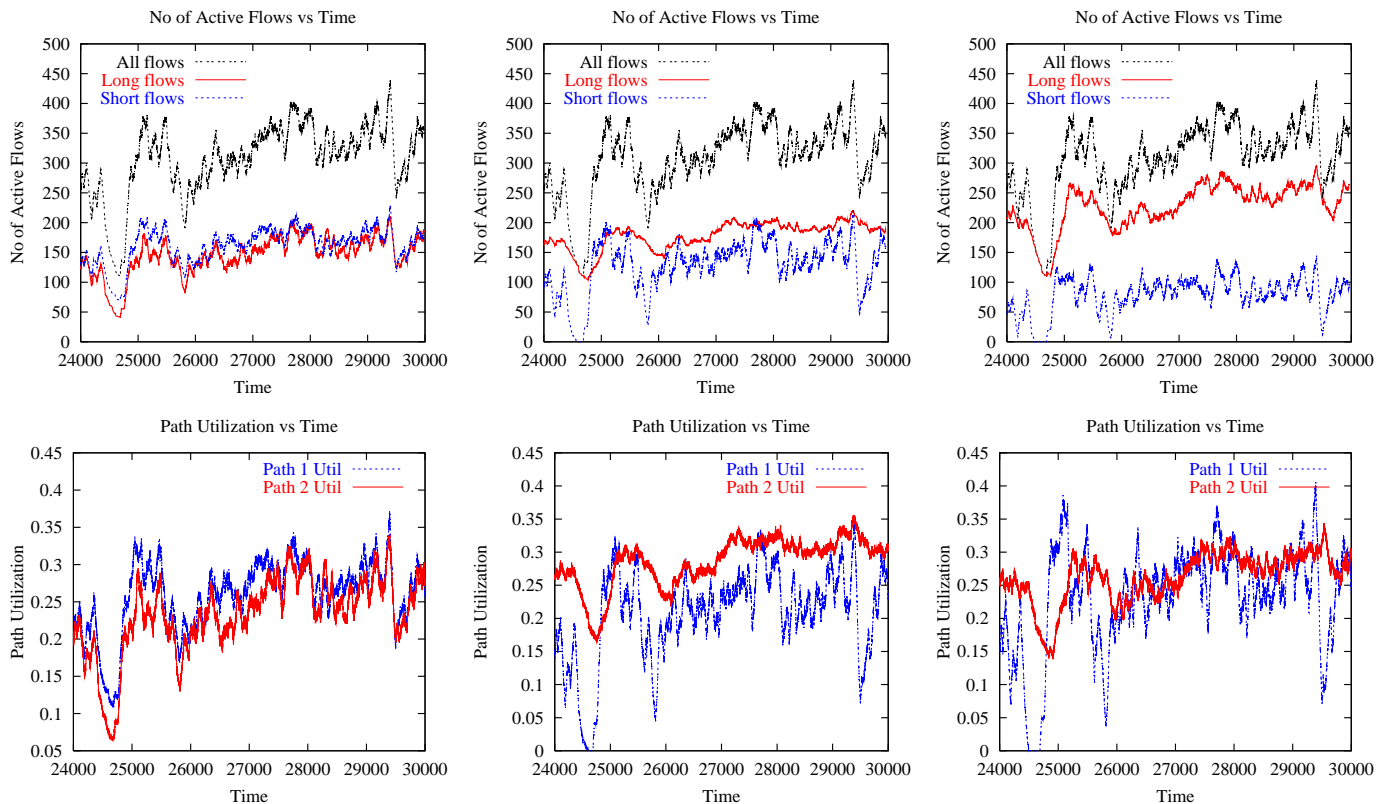


Fig. 2. Instantaneous # of active flows (top) and path utilization (bottom) for a load-balanced assignment policy of flows (left), an assignment policy that assigns short flows to one path and long flows to the other (middle), and an assignment policy that moves a flow from one path to the other once a flow length threshold is crossed.

Figure 3(b) shows the evolution of the TCP window of selected flows out of 600 TCP flows: 200 long flows whose size is uniformly distributed between 900 and 1100 packets, and 400 short flows whose size is uniformly distributed between 5 and 20 packets [35]. A short (long) flow that terminates is replaced by another short (long) flow. The packet size is 576 bytes and the maximum window size is 256 packets. All 600 TCP flows share a common bottleneck 10 Mbps link with 100 mseconds propagation delay (this gives a bandwidth-delay product of about 434 packets) and a buffer size of 500 packets. We observe, in this more loaded network, a more drastic change in window sizes of long flows, and a higher chance of short flows getting shut off. In this experiment, the throughput of long flows is observed to be around 60% of the total throughput.

Thus, in a second experiment, we split the two classes of flows, giving long flows 60% of the total link resources (buffer and capacity) and short flows the remaining 40%. Figure 4(top) shows selected windows for the long class (left), and selected windows for the short class (right). We can clearly observe that performance of long flows becomes more predictable. Moreover, the chance of a short flow getting shut off is reduced. Figure 4(bottom) illustrates that fairness is significantly improved when flows are classified, especially among short flows.⁵ In particular, the 99.9th percentile transmission time for long flows without isolation is 421 seconds and with isolation is 315 seconds — a 25% improvement. For short flows, the 99.9th percentile transmission time without isolation is 75 seconds and with isolation is 22 seconds — a 70% improvement.

The improvement in fairness for short TCP flows when all flows are mixed in a RED FIFO queue is smaller than that obtained by isolation — their 99.9th percentile transmission time is higher under shared RED (45 seconds compared to 22 seconds). Results are also shown when short flows are allocated 60% of the total link resources and long flows the remaining 40%. In this case, with isolation, one can achieve *both* better fairness and better response time for short flows, which are usually interactive/delay-sensitive. This is at the expense of relatively higher response time for long flows, however, fairness among them is still maintained.

D. Effect of Pipe Size

Figure 5 shows the topology used for this set of experiments. The link between nodes 0 and 1 is the bottleneck link. Each pair of sender S_i and receiver R_i sets up a session in which infinite amount of data is sent, however, the session data is spread over several connections. Each connection has a limited size and classified as either long if its size is 1000 packets or short if its size is 10 packets. Each session starts at a randomly chosen time in the first 5 seconds of simulation time. Whenever a connection finishes

⁵Note that a distribution curve that is more steep implies more fairness—a curve with a perpendicular increase implies ideal 100% fairness in the sense that all flows are experiencing the same transmission delay.

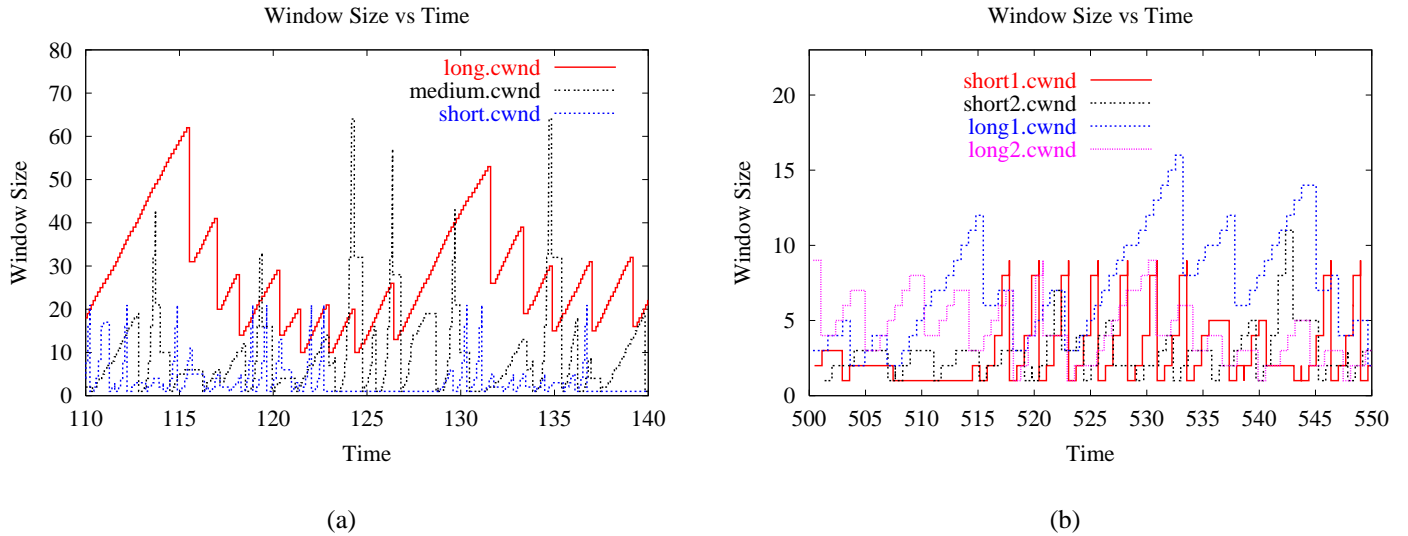


Fig. 3. TCP window size behavior when long and short flows share the same bottleneck link: (a) lightly loaded, and (b) more loaded.

its transmission, a new connection belonging to the same session starts. A total of six low-bandwidth telnet sessions (using `tcplip` in `ns` [32]) are used as background traffic to avoid deterministic (or synchronization) behavior. All the topology and protocol related parameters are listed in Table III.

TABLE III
SIMULATION PARAMETERS FOR STUDYING EFFECT OF PIPE SIZE

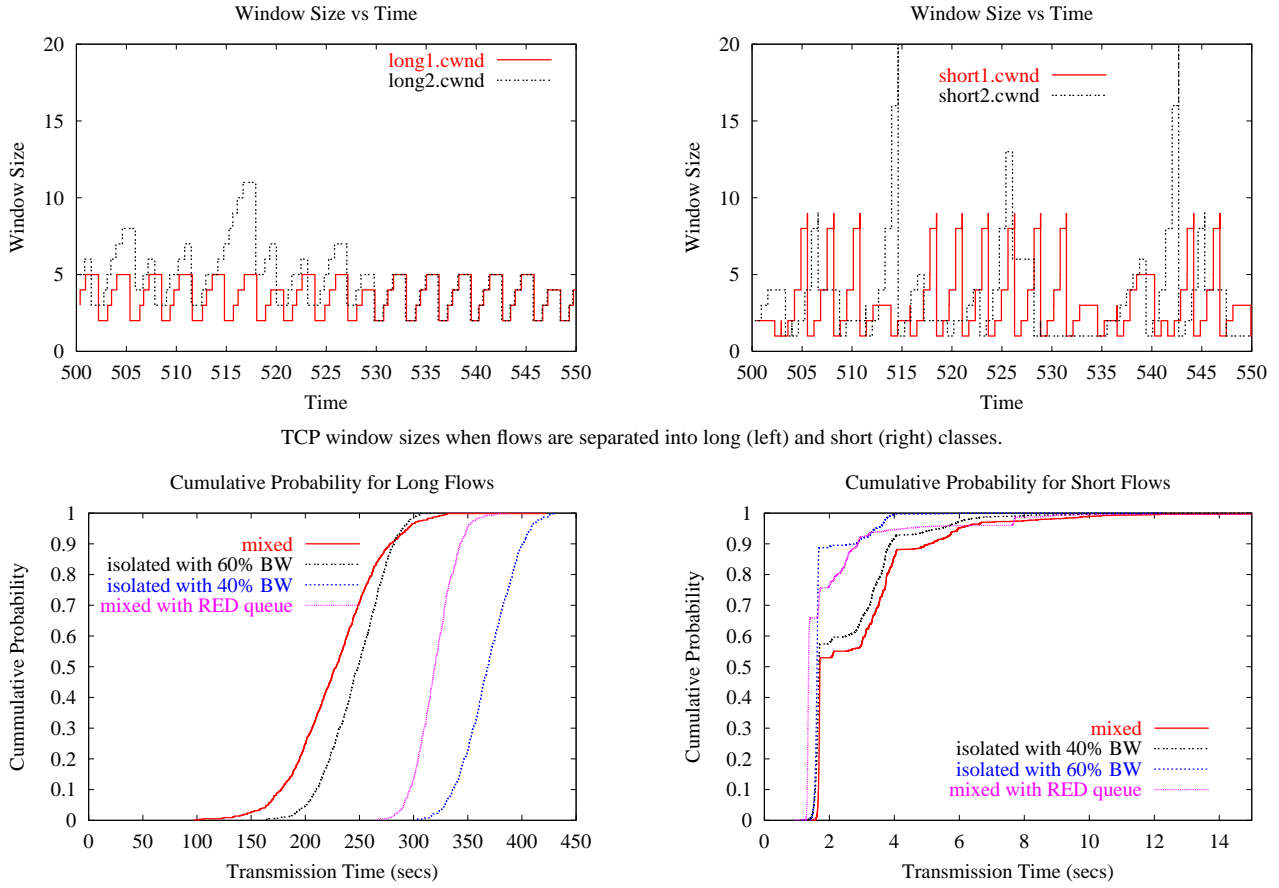
Description	Value
Packet size	576 bytes
Maximum TCP window size	128 packets
TCP timeout granularity	0.5 seconds
Bottleneck link delay	40 ms
Access link delay	5 ms
Capacity unit C_u	2.5 Mbps = 542 packets/s
Bottleneck capacity C	$k_c \times C_u$
Buffer unit B_u	50 packets $\approx C_u \times RTT$
Bottleneck link buffer B	$k_b \times B_u$
Access link capacity	$4 \times C_u$
Access link buffer	4000 packets
Long connection size	1000 packets
Short connection size	10 packets

We use k_c and k_b as scaling factors to control the size of the network pipe. From the simulation parameters, we can roughly compute the size of the network pipe and the average share that a connection should get from the network. We use the two-way propagation delay of 100 mseconds as an estimate of the round-trip time RTT . Therefore, the network pipe size, P , is computed as:

$$\begin{aligned}
 P &= C \times RTT + B \\
 &= k_c \times C_u \times RTT + k_b \times B_u
 \end{aligned}$$

Assume there are N concurrent connections, then each connection i should be able to transmit at rate $r_i = \frac{P}{N}$. In this section, we set $k_c = k_b = k$. If $k = 2$ and $N = 20$, we can compute P to be 208 packets and $r_i = 10.4$ packets/RTT.

We vary the value of N and k to study per-class fairness under different network congestion conditions for both isolation and sharing. In the *isolation case*, we assume class-based queuing is employed at the bottleneck link to essentially create two separate pipes: the first pipe, allocated to N sessions of short TCP flows, has one queue with $k = 1$ (and hence capacity of C_u and buffer space of B_u). The second pipe, allocated to N sessions of long TCP flows, has another queue with k ranging from 1 to 8 (and hence capacity ranging from C_u to $8C_u$ and corresponding buffer space ranging from B_u to $8B_u$). Each queue (pipe) is also used



Long (left) and short (right) flow transmission time distribution—long flows of size [980:1020], and short flows of size [9:11].

Fig. 4. Impact of isolation control.




Fig. 5. Topology to study effects of pipe and buffer size

by three of the six telnet sessions. In the *sharing case*, all $2N$ sessions, in addition to the six telnet background sessions, share a common queue (pipe) with combined capacity ranging from $2C_u$ to $9C_u$ and corresponding buffer space ranging from $2B_u$ to $9B_u$.

D.1 Large Pipe Case

In these experiments, we set $N = 30$, *i.e.* 30 sessions of long flows and 30 sessions of short flows. In this case, the size of the pipe allocated to short flows is 104 packets, thus each short connection can roughly transmit at a rate of $104/30=3.5$ packets/RTT. Since to transmit 10 packets, at least 4 round-trip times are needed, the maximum rate at which a short connection can transmit is 2.5 packets/RTT.⁶ Thus, this setup provides enough resources to short connections to transmit at the maximum rate. Figure 6(a) shows the fairness index for isolated short flows, isolated long flows, and for flows in the sharing (mixed) case. Figure 6(b) shows the goodput for each class of connections.

Since enough bandwidth is allocated to short connections that would allow them to transmit at maximum rate, short connections do not suffer any packet loss. Thus the fairness index for short connections is very close to 1. On the other hand, long connections

⁶Recall that the sending window of a short TCP connection, operating in slow-start phase, doubles every round-trip time, *i.e.* takes the values 1, 2, 4, 8, \dots . Thus, to send 10 packets, in the absence of packet losses, a TCP connection can send 1, 2, 4, and finally 3 packets over 4 round-trip times and finish its transmission.

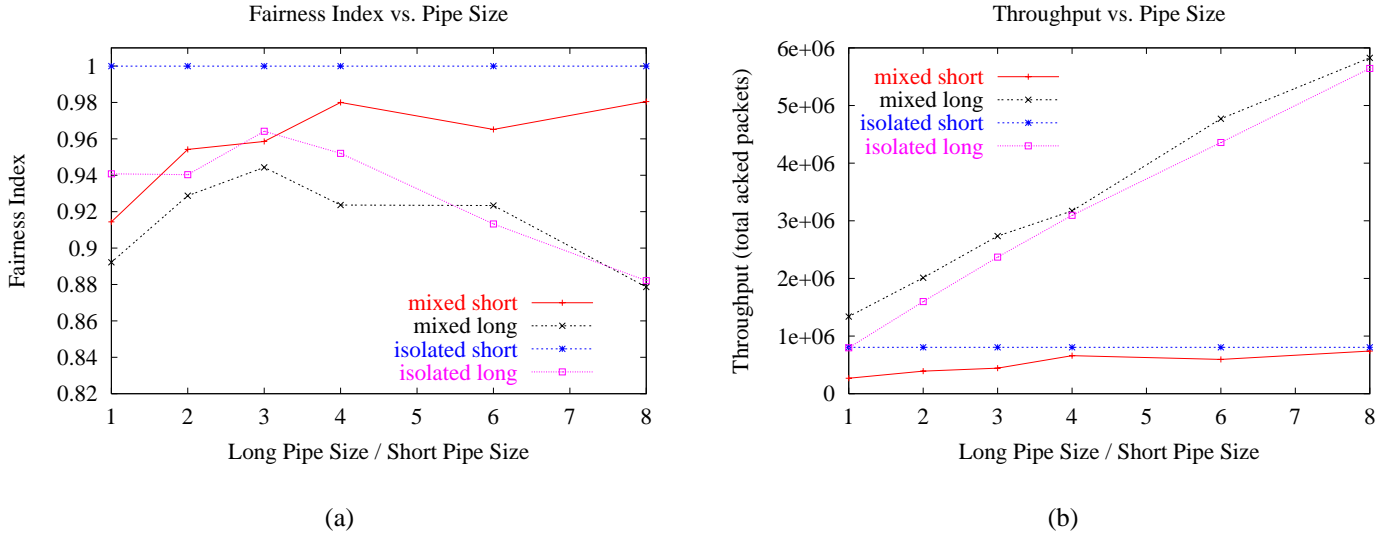


Fig. 6. Performance versus pipe size for large pipe case ($N = 30$): (a) Fairness index, and (b) Throughput

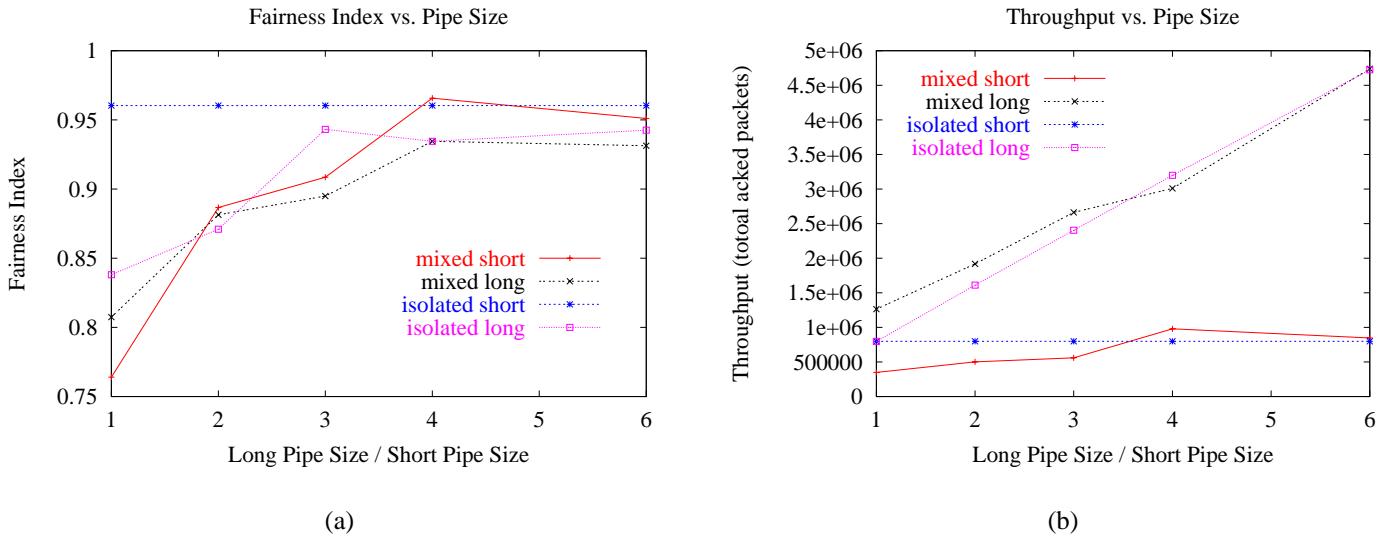


Fig. 7. Performance versus pipe size for small pipe case ($N = 60$): (a) Fairness index, and (b) Throughput

start to experience packet loss once the window size of a connection k approaches its fair share of $3.5k$ (unless $3.5k >$ the maximum window size of 128), resulting in unfairness among these long connections (consistent with [30]).

In the sharing case, when short and long connections are multiplexed in a common queue, short connections experience packet losses. Because long connections send larger windows (packet bursts), they can grab more bandwidth than their fair share (c.f. Figure 6(b)). Some of the short connections are also forced into timeout (due to insufficient number of duplicate ACKs to fast retransmit), resulting in unfairness among short connections as well. Therefore, with isolation in a well-engineered class-based network, each class of flows can enjoy a more friendly (fair) transmission environment without sacrificing the overall goodput and utilization of the network.

D.2 Small Pipe Case

In this set of experiments, we set $N = 60$, i.e. 60 sessions of long flows and 60 sessions of short flows. The results are shown in Figure 7. In this case, the share of each short connection becomes 1.75 packets/RTT. This results in short connections experiencing packet losses and timing out. Thus fairness among short connections is lower than in the previous large pipe case.

However, with isolation, long connections do not impede short connections, resulting in relatively high fairness (around 96%). On the contrary, when short and long connections are mixed, short connections suffer from severe unfairness ($f < 0.8$) when the network pipe is small (total capacity of $2C_u$). For long connections, isolation also results in improved fairness, yet less significant than that for short connections. Observe that isolation provides an overall goodput that is comparable to that obtained by sharing. More importantly, from Figure 6(b) and Figure 7(b), we observe a (almost) perfectly linear relationship between goodput and the

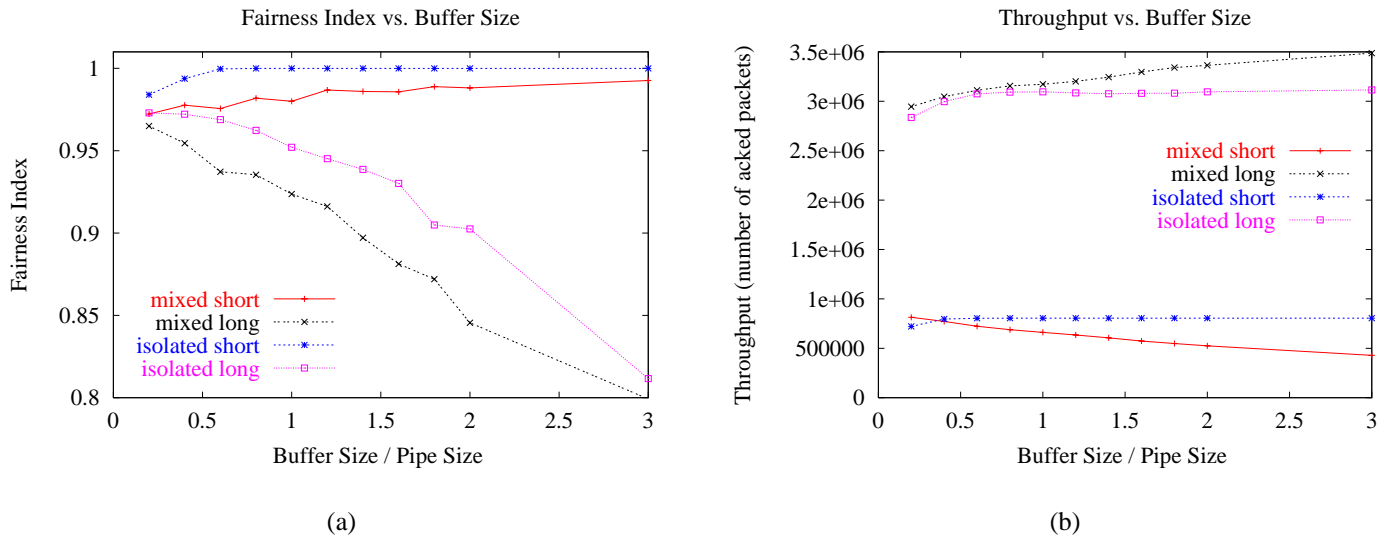


Fig. 8. Performance versus buffer/pipe size ratio: (a) Fairness index, and (b) Throughput

amount of resources allocated to long connections. This makes it possible for a service provider to reliably predict performance for each class of connections, which is hard to do when resources are completely shared.

E. Effect of Buffer Size

Buffer space provisioning is an important factor in TCP’s performance [36], [37]. Using the same setup as in Section IV-D, we investigate the effect of buffer space on isolation and sharing. We fix N at 30, *i.e.* 30 sessions of long flows and 30 sessions of short flows. Under isolation, we set $k_c = 1$ for the short class pipe and $k_c = 4$ for the long class pipe. Thus, $k_c = 5$ in the sharing case, *i.e.* a total capacity of $5C_u$. We control the ratio of the buffer space relative to the size of the pipe by varying the value of k_b from $0.2k_c$ to $3k_c$. Figure 8 shows the results.

In the sharing (mixed) case, as the buffer size increases, long connections are allowed to send bigger windows (packet bursts), thus they grab more bandwidth from short connections, and more short connections are forced to timeout. As a result, the overall goodput for short connections decreases. A side effect of this unfair treatment to short connections is an improvement in fairness among them! (They all start to equally see worse performance, consistent with [38], [39].) However, even with such a high “fairness” value for short connections, it’s very hard to predict the goodput of short connections, contrary to the isolation scenario where goodput is not affected very much by varying buffer space.

To summarize, when network resources are relatively scarce—generally when the share of each connection is less than 10 packets per round-trip time— isolation provides a more fair and more predictive service, as only “similar” connections (of the same class) compete for their allocated resources.

F. In The Presence of Web Background Traffic

In this set of experiments, we compare the three traffic management schemes, described in Section III-A, distributing flows over two parallel paths (cf. Figure 1), in the presence of Web background traffic on the topology shown in Figure 9.

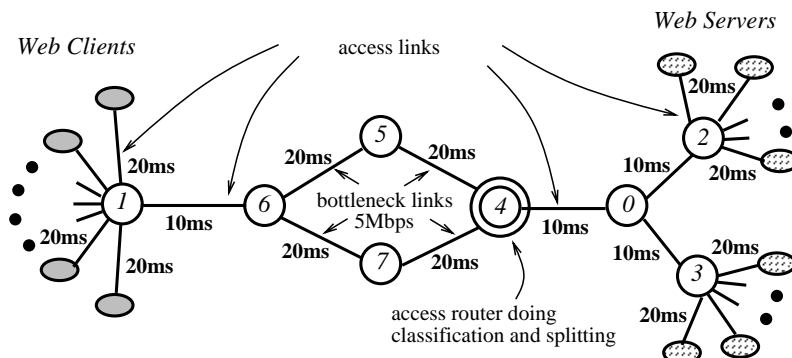


Fig. 9. Topology with Web background traffic

This model is similar to the “FLEXBELL” model used in [40]. Here, node 0 connects to the servers pool, while node 1 connects to the clients pool. There is a diamond topology between nodes 0 and 1, with two separate paths between nodes 4 and 6. This diamond topology represents an administrative domain where the edge node 4 is responsible for detecting and classifying flows. This diamond topology is also comprised of the bottleneck links/paths. We assume a Web session model, so a client from the clients pool sets up a *session* with a server. Within each session, the client can grab a *page*, which could contain several *objects*. For each object, a TCP connection is established for transmission. All parameters are given in Table IV. The thresholds used for size-based splitting and threshold-based splitting schemes were chosen so as to roughly divide the overall traffic equally among the two separate paths.

TABLE IV
SIMULATION PARAMETERS FOR WEB MODEL

Description	Value
Web Background Traffic	
Number of sessions	200
Number of connections per session	1900
Number of pages per connection	1
Number of objects per page	1
Interarrival of connections	exponential 2.4 seconds
Connection size distribution	Bounded Pareto [4:200000] with skew parameter of 1.2
Threshold for size-based splitting	75 packets
Threshold for threshold-based splitting	32 packets
Packet size	576 bytes
Maximum window size	128 packets
Bottleneck link capacity (C)	5 Mbps
Bottleneck link bandwidth-delay product	238 packets
Bottleneck link buffer size	250 packets
RED queue min_threshold	20 packets
RED queue max_threshold	245 packets
RED queue max drop rate	0.2 %
Access link capacity	$\geq 3C$
Access link buffer size	4000 packets
Foreground Traffic	
Number of long connections	10
Number of short connections	10
Long (foreground) connection size	1000 packets
Short (foreground) connection size	10 packets

Both foreground TCP connections and the Web background connections are routed according to the routing scheme employed at node 4. Performance measures are collected for the foreground TCP connections, and results are shown in Figure 10.

Results are also shown when RED shared queues are deployed in the routers (with a load-balanced strategy at node 4). The two isolation schemes, namely size-based and threshold-based splitting, can achieve fairness comparable to that of RED. However, with RED queues, fairness is achieved by sacrificing the goodput for short connections with smaller packet bursts (window sizes). On the contrary, with isolation (size-based splitting), short connections encounter less packet losses and achieve both high fairness and high goodput.

Without either RED or isolation, i.e. with tail-drop shared queues (with a load-balanced strategy at node 4), long connections can receive a much higher goodput. However, fairness among long connections is very low (as low as 20%) compared to those achieved by using RED shared queues or isolation (all above 80%). Low fairness values imply more variability in achieved goodput, which makes it very hard to predict the behavior of connections. In this sense, isolation provides an important engineering tool to service providers to more accurately predict performance and thus provide reliable guarantees to users. Furthermore, unlike RED shared queues, by allocating enough resources to short interactive TCP transfers, they can be provided throughput/delay guarantees.

V. RELATED WORK

A number of studies (e.g. [41], [42], [43]) examined isolation and sharing of network resources in the context of Intserv per-flow architectures. We concerned ourselves here with Diffserv-like architectures, where per-flow information is only maintained

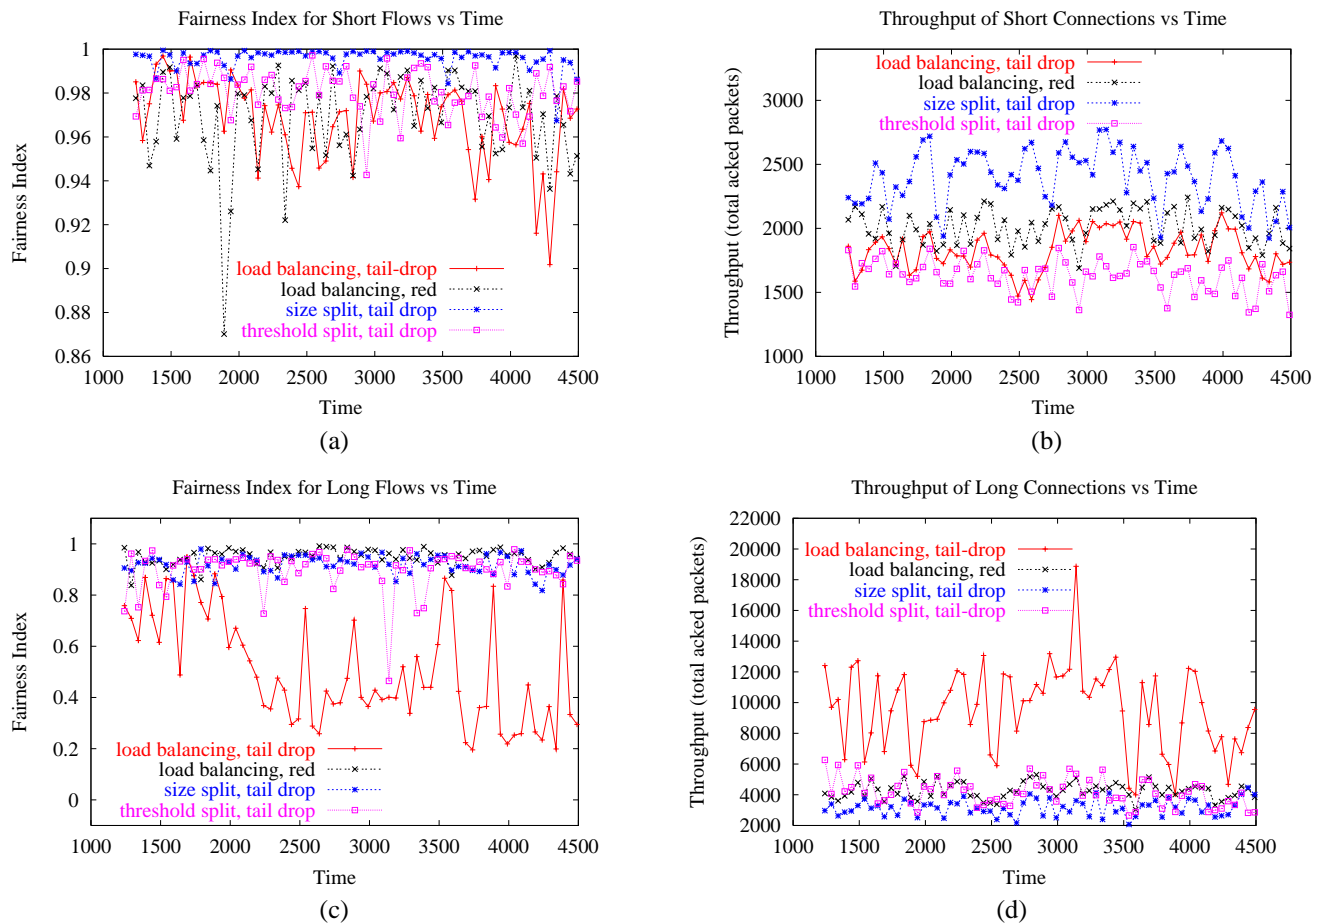


Fig. 10. Performance with Web background traffic: (a) Fairness index of short connections, (b) Goodput of short connections, (c) Fairness index of long connections, and (d) Goodput of long connections

at access routers. In Section II, we formulated discrete-time control models along the same lines as in [24]. However, Li and Nahrstedt in [24] do not address the benefits of isolation over sharing, particularly as related to fairness and stability of TCP flows.

A number of studies have proposed protocols that differentiate among flows based on their size or lifetime. For example, Shaikh *et al.* [16] investigate load-sensitive routing of only long-lived IP flows in order to improve routing stability and reduce overhead. However, the effect on flow-controlled sources (such as TCP) has not been studied. Other studies have considered the cut-through switching of long-lived IP flows (e.g. using IP/Tag Switching) [15], [18]. These studies only consider the reduced overhead due to switching, but not performance of flows.

Other studies have attempted to improve TCP fairness by either modifying TCP itself or by employing non-tail-drop buffer management at routers [44], [45], [46], [47], [48], [49], [50], [51], [30], [52], [53], [54], [35], [37], [25]. In particular, Morris [35], [37] proposes solutions that require per-flow information at all routers. In this paper, we advocate the use of a (less costly) class-based solution. Bonald *et al.* [38] and Nandy *et al.* [55] have shown that the well-known RED buffer management policy may have a strong loss bias against smooth UDP (e.g. audio) flows. Seddigh *et al.* [56], [57] propose a packet dropping mechanism to improve fairness among UDP and TCP flows. Our focus here has been on using isolation to improve predictability of the service and fairness of TCP flows of different size.

Most TCP studies consider very long (or *infinite*) TCP connections, and focus on characterizing the steady-state transfer throughput. Only few recent TCP studies [58], [29], [59] have started to investigate short flows, which comprise most of the current Internet flows [15]. However, to our knowledge, the interaction among short and long TCP flows has not been studied.

VI. CONCLUSIONS AND FUTURE WORK

Using control-theoretic arguments and extensive simulations, we have shown that service predictability and fairness can be much improved by isolating TCP flows based on their size. In this paper, we classified TCP flows into two classes: one of short flows and another of long flows. Each class is (logically) allocated separate resources (buffer space and capacity), for which size-homogeneous flows compete. Flow classification can be done only by an access (border) router, thus the complexity of implementing such a class-based solution is low. By appropriately allocating resources to each class, the many short (usually interactive/delay-sensitive) flows can not only enjoy better predictability and fairness, but also faster service. This is in sharp

contrast to traditional shared systems that do not differentiate among TCP flows with different characteristics, such as burstiness of arrivals/departures and congestion/sending window dynamics.

Future work remains on how to dynamically allocate resources to various classes, how to determine the best threshold values used to classify flows into classes, and how to account for characteristics other than flow size (e.g. various round-trip times).

ACKNOWLEDGMENTS

Thanks to Azer Bestavros for various discussions on aspects of this work. This work was supported in part by NSF grants CAREER ANIR-0096045 and MRI EIA-9871022.

APPENDIX

I. STABILITY CONDITION FOR AGGREGATE CONTROL

In this Appendix, we derive the stability condition in equation (2) for the aggregate control system of Section II-A. Without loss of generality, assume a constant acknowledgment (service) rate (i.e. $c^k = C$ for all k), all flows start with the same initial input state (i.e. w_i^0 is the same for all i), and that all flows adapt at the same rate (i.e. $\alpha_i = \alpha$ and $\beta_i = \beta$ for all i). Then, equations (1) can be rewritten as:

$$\begin{aligned} w_i^k &= w_i^{k-1} + \alpha(B - q^k) + \beta[(B - q^k) - (B - q^{k-1})] \\ q^k &= q^{k-1} + \sum_{i=1}^N w_i^{k-1} - C \end{aligned} \quad (8)$$

Since flows adapt their w_i^k at the same rate, then $w_i^k = \frac{\sum_{j=1}^N w_j^k}{N}$ for all i . Denote by e^k the error at time step k , i.e. $e^k = B - q^k$, and let $y^k = \sum_{j=1}^N w_j^k - C$. Equations (8) can then be rewritten as:

$$\begin{aligned} \frac{1}{N}y^k &= \frac{1}{N}y^{k-1} + \alpha e^k + \beta(e^k - e^{k-1}) \\ q^k &= q^{k-1} + y^{k-1} \end{aligned} \quad (9)$$

Taking the z-transform [28] of equations (9), we get:

$$\begin{aligned} \frac{1}{N}Y(z) &= \frac{1}{N}z^{-1}Y(z) + \alpha E(z) + \beta(E(z) - z^{-1}E(z)) \\ Q(z) &= z^{-1}Q(z) + z^{-1}Y(z) \\ E(z) &= B - Q(z) \end{aligned} \quad (10)$$

The relationships between $Q(z)$, $Y(z)$, $E(z)$ and B can be represented as shown in Figure 11. The ratio $\frac{Y(z)}{B}$ gives the so-called




Fig. 11. Aggregate control model.

transfer function of the whole system, and is given by:

$$\frac{Y(z)}{B} = \frac{D(z)}{1 + D(z)H(z)} \quad (11)$$

where $D(z)$ and $H(z)$ are given in Figure 11. To derive the stability condition of the system, we solve its so-called *characteristic equation* given by the denominator $1 + D(z)H(z)$. We then find the condition under which the roots (or poles) of this characteristic

equation lie within the unit circle in the z -plane, i.e. $|z| < 1$ [28]. After some algebraic manipulations, we obtain the characteristic equation:

$$z^2 + [N(\alpha + \beta) - 2]z + (1 - \beta N) \quad (12)$$

For $|z| < 1$, we must have:

$$\alpha + 2\beta < \frac{4}{N} \quad (13)$$

II. TCP CONGESTION CONTROL

We give here an outline of the TCP congestion control algorithm, in particular, how the congestion window is updated. See Figure 12. The reader is referred to [33] for more details.

On every ACK

```

if (window < ssthresh) // slow start phase
    window += 1 // double window every RTT
else // congestion avoidance
    window += 1 / window // increment by 1 every RTT

```

On timeout

```

ssthresh = window / 2
window = 1

```

On duplicate ACK

```

ssthresh = window / 2
window = window / 2 // fast recovery

```




Fig. 12. TCP congestion control.

REFERENCES

- [1] S. Shenker, C. Partridge, and R. Guerin, "Specification of Guaranteed Quality of Service," Internet RFC 2212, September 1997.
- [2] L. Zhang, R. Yavatkar, F. Baker, P. Ford, K. Nichols, M. Speer, and Y. Bernet, "A Framework for Use of RSVP with Diff-serv Networks," Internet Draft draft-ietf-diffserv-rsvp-00.txt, June 1998.
- [3] Y. Bernet and et al, "A Framework for Differentiated Services," Internet Draft draft-ietf-diffserv-framework-02.txt, IETF, February 1999.
- [4] K. Nichols, V. Jacobson, and L. Zhang, "A Two-bit Differentiated Services Architecture for the Internet," Internet Draft draft-nichols-diff-svc-arch-02.txt, April 1999.
- [5] D.D. Clark and W. Fang, "Explicit Allocation of Best-Effort Packet Delivery Service," *IEEE/ACM Transactions on Networking*, pp. 362–373, August 1998.
- [6] W. Feng, D. Kandlur, D. Saha, and K. Shin, "Adaptive Packet Marking for Providing Differentiated Services in the Internet," in *Proceedings of ICNP '98*, October 1998.
- [7] I. Stoica and H. Zhang, "LIRA: An Approach for Service Differentiation in the Internet," in *NOSSDAV'98*, London, UK, July 1998.
- [8] Z. Wang, "User-Share Differentiation (USD) Scalable Bandwidth Allocation for Differentiated Services," Internet Draft draft-wang-diff-serv-usd-00.txt, September 1999.
- [9] C. Dovrolis, D. Stiliadis, and P. Ramanathan, "Proportional Differentiation and Packet Scheduling," in *Proceedings of ACM SIGCOMM '99*, Cambridge, Massachusetts, August 1999.
- [10] N. Semret, R.R.-F. Liao, A.T. Campbell, and A.A. Lazar, "Peering and Provisioning of Differentiated Internet Services," in *Proceedings of IEEE INFOCOM 2000*, Tel-Aviv, Israel, March 2000.
- [11] T. Nandagopal, N. Venkitaraman, R. Sivakumar, and B. Vaduvur, "Relative Delay Differentiation and Delay Class Adaptation in Core-Stateless Networks," in *Proceedings of IEEE INFOCOM 2000*, Tel-Aviv, Israel, March 2000.
- [12] W. Aiello, Y. Mansour, S. Rajagopalan, and A. Rosen, "Competitive Queue Policies for Differentiated Services," in *Proceedings of IEEE INFOCOM 2000*, Tel-Aviv, Israel, March 2000.
- [13] S. Sahu, D. Towsley, and J. Kurose, "A Quantitative Study of Differentiated Services for the Internet," in *Proceedings of IEEE Global Internet '99*, Rio de Janeiro, Brazil, December 1999.

- [14] I. Yeom and A.L.N. Reddy, "Modeling TCP Behavior in a Differentiated Services Network," in *IEEE/ICNP'99*, November 1999.
- [15] K. Thompson, G.J. Miller, and R. Wilder, "Wide-Area Internet Traffic Patterns and Characteristics," *IEEE/ACM Transactions on Networking*, pp. 10–23, November 1997.
- [16] A. Shaikh, J. Rexford, and K.G. Shin, "Load-Sensitive Routing of Long-Lived IP Flows," in *SIGCOMM'99*, Boston, MA, September 1999.
- [17] W. Feng, D. Kandlur, D. Saha, and K. Shin, "Understanding TCP Dynamics in an Integrated Services Internet," in *NOSSDAV '97*, 1997.
- [18] A. Feldmann, J. Rexford, and R. Caceres, "Efficient Policies for Carrying Web Traffic over Flow-Switched Networks," *IEEE/ACM Transactions on Networking*, pp. 673–685, December 1998.
- [19] S. Floyd and V. Jacobson, "Link-sharing and Resource Management Models for Packet Networks," *IEEE/ACM Transactions on Networking*, pp. 365–386, August 1995.
- [20] R. Callon, P. Doolan, N. Feldman, A. Fredette, G. Swallow, and A. Viswanathan, "A Framework for MPLS," Ietf draft, Network Working Group, September 1999.
- [21] D. Awduche, J. Malcolm, J. Agogbua, M. O'Dell, and J. McManus, "Requirements for Traffic Engineering over MPLS," RFC 2702, Network Working Group, September 1999.
- [22] P. Newman, T. Lyon, and G. Minshall, "Flow Labelled IP: A Connectionless Approach to ATM," in *IEEE INFOCOM '96*, San Francisco, CA, March 1996, pp. 1251–1260.
- [23] Y. Rekhter et al., "Tag Switching Architecture Overview," Internet Draft, September 1996.
- [24] B. Li and K. Nahrstedt, "A Control Theoretical Model for Quality of Service Adaptations," in *IEEE International Workshop on Quality of Service (IWQoS 98)*, Napa, California, May 1998.
- [25] S. Shenker, "A Theoretical Analysis of Feedback Flow Control," in *ACM/SIGCOMM'90*, 1990.
- [26] V. Jacobson, "Congestion Avoidance and Control," in *ACM SIGCOMM'88*, Stanford, CA, August 1988.
- [27] D.M. Chiu and R. Jain, "Analysis of the Increase and Decrease Algorithms for Congestion Avoidance in Computer Networks," *Computer Networks and ISDN Systems*, pp. 1–14, 1989.
- [28] K. Ogata, *Discrete-Time Control Systems*, Prentice-Hall, Inc., 1987.
- [29] N. Cardwell, S. Savage, and T. Anderson, "Modeling TCP Latency," in *IEEE/INFOCOM'2000*, Tel-Aviv, Israel, March 2000.
- [30] S. Floyd and V. Jacobson, "Random Early Detection Gateways for Congestion Avoidance," *IEEE/ACM Transactions on Networking*, pp. 397–413, August 1993.
- [31] M. May, T. Bonald, and J. Bolot, "Analytic Evaluation of RED Performance," in *Proceedings of IEEE/INFOCOM'2000*, Tel-Aviv, Israel, March 2000.
- [32] UCB, LBNL, and VINT, "Network simulator - ns (version 2)," <http://www-mash.cs.berkeley.edu/ns/>, 1999.
- [33] W. Stevens, "TCP Slow Start, Congestion Avoidance, Fast Retransmit, and Fast Recovery Algorithms," Internet RFC 2001, January 1997.
- [34] W. Stevens, *TCP/IP Illustrated, Vol. 1: The Protocols*, Addison-Wesley, 1997.
- [35] R.T. Morris, *Scalable TCP Congestion Control*, Ph.D. thesis, Harvard University, Cambridge MA, The Division of Engineering and Applied Sciences, January 1999.
- [36] R.T. Morris, "TCP Behavior with Many Flows," in *IEEE International Conference on Network Protocols (ICNP'97)*, Atlanta, Georgia, October 1997.
- [37] R.T. Morris, "Scalable TCP Congestion Control," in *IEEE/INFOCOM'2000*, Tel-Aviv, Israel, March 2000.
- [38] T. Bonald, M. May, and J-C. Bolot, "Analytic Evaluation of RED Performance," in *IEEE INFOCOM'2000*, Tel-Aviv, Israel, March 2000.
- [39] A. Aggarwal, S. Savage, and T. Anderson, "Understanding the Performance of TCP Pacing," in *IEEE/INFOCOM'2000*, Tel-Aviv, Israel, March 2000.
- [40] A. Feldmann, A.C. Gilbert, P. Huang, and W. Willinger, "Dynamics of IP traffic: A Study of the Role of Variability and the Impact of Control," in *Proceeding of ACM/SIGCOMM'99*, Boston, Mass, September 1999.
- [41] D. Clark, S. Shenker, and L. Zhang, "Supporting Real-Time Applications in an Integrated Services Packet Network: Architecture and Mechanism," in *Proceedings of ACM SIGCOMM '92*, August 1992, pp. 14–26.
- [42] A. Demers, S. Keshav, and S. Shenker, "Analysis and Simulation of a Fair Queueing Algorithm," in *ACM SIGCOMM '89*, Austin, Texas, September 1989, pp. 1–12.
- [43] B. Suter, T.V. Lakshman, D. Stiliadis, and A.K. Choudhury, "Design Considerations for Supporting TCP with Per-Flow Queueing," in *Proceedings of IEEE INFOCOM'98*, San Francisco, CA, March 1998.
- [44] The PSC Networking group, "The TCP-Friendly Website," <http://www.psc.edu/networking/tcp-friendly.html>.
- [45] F.P. Kelly, A.K. Maullo, and D.K.H. Tan, "Rate Control for Communication Networks: Shadow Prices, Proportional Fairness and Stability," *Journal of the Operational Research Society*, 1998.
- [46] F.P. Kelly, "Proportional Fairness Website," <http://www.statslab.cam.ac.uk/frank/pf>.
- [47] P. Hurley, J.L. Boudec, and P. Thiran, "A Note on the Fairness of Additive Increase and Multiplicative Decrease," in *16th International Teletraffic Congress (ITC-16)*, Edinburgh, Scotland, June 1999.
- [48] S. Kunniyur and R. Srikant, "Fairness of Congestion Avoidance Schemes in Heterogeneous Networks," in *16th International Teletraffic Congress (ITC-16)*, Edinburgh, Scotland, June 1999.
- [49] P. Brown, "Resource Sharing of TCP Connections with Different Round Trip Times," in *IEEE INFOCOM'2000*, Tel-Aviv, Israel, March 2000.
- [50] L. Massoulié and J. Roberts, "Bandwidth Sharing: Objectives and Algorithms," in *Proceedings of IEEE/INFOCOM'99*, New York, 1999.
- [51] J. Mo, R.J. La, V. Anantharam, and J. Walrand, "Analysis and Comparison of TCP Reno and Vegas," in *IEEE/INFOCOM'99*, New York, March 1999.
- [52] W. Feng, D. Kandlur, D. Saha, and K. Shin, "Techniques for Eliminating Packet Loss in Congested TCP/IP Networks," Tech. Rep. CSE-TR-349-97, Univ. of Michigan, 1997.
- [53] T. Ott, T.V. Lakshman, and L. Wong, "SRED: Stabilized RED," in *IEEE/INFOCOM'99*, New York, 1999.
- [54] D. Lin and R. Morris, "Dynamics of Random Early Detection," in *ACM/SIGCOMM'97*, 1997.
- [55] B. Nandy, N. Seddigh, and P. Piedad, "Diffserv's Assured Forwarding PHB: What Assurance does the Customer Have?," in *Proceedings of NOSSDAV'99*, July 1999.
- [56] W. Fang, N. Seddigh, and B. Nandy, "A Time Sliding Window Three Colour Marker (TSWTCM)," Internet Draft, draft-fang-diffserv-tc-tswtcm-00.txt, October 1999.
- [57] N. Seddigh, B. Nandy, and P. Piedad, "Study of TCP and UDP Interaction for the AF PHB," Internet Draft, draft-nsnbp-diffserv-tcpudpaf-01.pdf, September 1999.
- [58] N. Cardwell, S. Savage, and T. Anderson, "Modeling the Performance of Short TCP Connections," Tech. Rep., Department of Computer Science and Engineering, Univ. of Washington, 1998, <http://www.cs.washington.edu/homes/cardwell/quals>.
- [59] C. Barakat and E. Altman, "Performance of Short TCP Transfers," in *Networking 2000*, Paris, May 2000.