Effect of Higher Priority EF Traffic on TCP Throughput and Fairness

Jasleen Kaur and **Harrick M. Vin**

Technical Report TR2001-1
Department of Computer Sciences
University of Texas at Austin
TAY 2.124, Austin, TX 78712-1188, USA

{jks,vin}@cs.utexas.edu

July 5, 2000

Abstract

In this paper, we study the effect on TCP of assigning higher-priority to traffic requesting *Expedited Forwarding* (EF) service in a *Differentiated Services* network. We analyze networks in which (1) EF traffic occupies different fractions of link bandwidth and is bursty at different time-scales; and (2) multiple TCP flows with heterogeneous round trip times share the network with the EF traffic. We find that even in the presence of bursty EF traffic, statistical multiplexing gains allow TCP to utilize most of the available bandwidth. Further, the presence of bursty EF traffic improves the fairness of bandwidth allocation among TCP flows; smaller more frequent bursts yield larger improvements in TCP fairness.

1 Introduction

The Internet has traditionally supported the *best-effort* service model in which the network offers no assurance about when, or even if, packets will be delivered. With the commercialization of the Internet and the deployment of inelastic continuous media applications, however, the best-effort service model is increasingly becoming inadequate. To facilitate the co-existence of these emerging applications with conventional elastic applications, the *differentiated services* architecture has been proposed [22]. In this architecture, traffic entering a network is classified and conditioned at the boundaries of the network, and is assigned to a small set of behavior (or flow) aggregates (also referred to as Per Hop Behaviors—PHB). Recently, several PHBs—such as the Expedited Forwarding (EF) and the Assured Forwarding (AF) PHB—and several end-to-end services—such as the Virtual Leased Line service [14, 22]—have been defined. However, very little is known about what end-to-end performance can be expected for flows that utilize a specific PHB, or how do the implementations for providing service differentiation among the PHBs impact the performance of best-effort flows. In this paper, we take a step towards addressing this question.

To formulate precisely the problem we investigate in this paper, consider the proposal for using the EF PHB to implement the Virtual Leased Line (VLL) service [14, 22]. The VLL service desires the network to provide *guaranteed rate* and *low delay* to flows. It is suggested that a differentiated services network can provide VLL service by following three steps: (1) shape the flows requesting the VLL service to *constant bit rate (CBR)*, and mark packets of the flow as belonging to the EF service class [14]) by appropriately setting the Type-of-Service (ToS) byte in the IP header of the packet [21]; (2) employ admission control algorithms at the routers to ensure that the aggregate rate of flows that request the Virtual Leased Line service does not exceed the capacity reserved for the EF PHB; and (3) provide higher priority to packets requesting EF PHB or implement a fair queuing algorithm to arbitrate access to link bandwidth among the different PHBs. It has been shown that providing higher priority to packets requesting EF PHB yields lower end-to-end delay and jitter to EF traffic [22], as required by the VLL service definition. Further, it is well-known that priority schedulers are simpler to implement than more sophisticated fair queuing algorithms. Hence, implementing VLL service by providing higher priority to EF traffic may be desirable – unless such an approach severely affects the performance of other traffic classes sharing the network with the EF traffic.

In this paper, we attempt to answer this very question: *what is the impact of providing higher priority to EF traffic on the throughput and the fairness of best-effort TCP flows sharing the differentiated services Internet?* We study—through simulations—the effect of different levels of burstiness in the EF traffic on the throughput and fairness of best-effort TCP flows. Our experiments show that:

1. The throughput of an isolated TCP flow is affected severely in the presence of higher priority bursty traffic; the loss in TCP throughput is higher when the EF traffic is bursty at short time-scales. However, the aggregate throughput of multiple TCP flows remains roughly unaffected by the higher priority bursty traffic.

2. Presence of bursty EF traffic improves the fairness of bandwidth allocation among TCP flows; smaller more frequent bursts yield larger improvements in TCP fairness.

We observe that these results hold when the EF traffic occupies several different fractions—in the 5%–30% range—of bottleneck link bandwidth.

Based on these observations, we conclude that providing high priority to EF traffic does not adversely affect the aggregate throughput TCP flows and does, in fact, improve the fairness of bandwidth allocation among the best-effort TCP flows.

The rest of the paper is organized as follows. Section 2 describes our simulation setup. In Sections 3 and 4, respectively, we discuss the effect of higher priority EF traffic on the throughput and fairness of TCP. Section 5 discusses the related work, and Section 6 summarizes our contributions.

2 Experimental Methodology

The objective of our study is to evaluate the effect of higher priority EF traffic on the throughput and fairness of best-effort TCP flows. We have conducted an extensive simulation study using the *NS-2* network simulator [1]. In what follows, we describe our simulation environment, the design of our experiments, and the measures for the performance evaluation.

2.1 Simulation Environment

2.1.1 Network Topology

We consider a network topology depicted in Figure 1; similar network topologies have been used in several prior TCP performance studies [9, 16, 20, 26, 27]. The topology contains two core routers, R_1 and R_2. All the network links have a bandwidth of 40Mbps. One of the input links of router R_1 carries the higher priority EF traffic[1]; this traffic occupies a fixed percentage of the bottleneck link bandwidth. Router R_1 is also connected to 8 other input links that carry the best-effort TCP traffic destined to one of more than 100 sinks connected to router R_2. Routers R_1 and R_2 provide higher priority with respect to link scheduling and link buffer occupancy to the packets belonging to the EF service class. In this setup, at the bottleneck link connecting routers R_1 and R_2, the presence of higher priority EF traffic would affect the queuing delay and packet loss rate experienced by TCP flows.

The topology models the heterogeneity in the round-trip propagation latencies for different TCP flows by assigning different deterministic propagation delays to each incoming link of router R_1 and each outgoing link of router R_2.

[1] Observe that this assumption ensures that the higher priority EF traffic does not occupy any link buffers. In practice, however, EF traffic would arrive at a core router on more than one input link, and would occupy link buffers. Hence, the results presented in this paper provide a conservative estimate on the impact of providing higher priority to EF traffic on the performance of best-effort TCP flows.

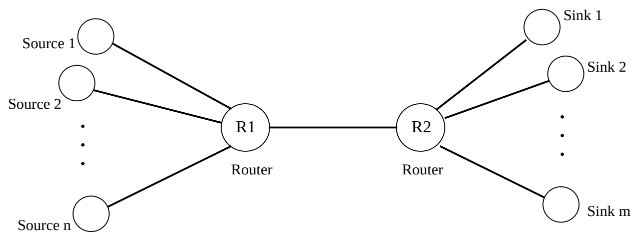

Figure 1: Network Topology

This allows us to model over 800 different TCP round-trip propagation latencies (RTPs)—ranging from around 6ms to 40ms. To maintain the network pipeline full, we provision link buffers in accordance with the *delay-bandwidth product* (i.e., the product of the link bandwidth and the maximum RTP) [6]. Finally, we assume *drop-tail* routers; drop-tail is the most widely deployed buffer management policy in today's Internet routers.

2.1.2 Network Traffic

Flows requesting the EF PHB are shaped to CBR at the ingress routers. However, it has been shown that the aggregation of CBR traffic entering a core router is bursty due to at least two reasons [8, 13, 24, 25]: (1) superpositioning of heterogeneous CBR flows yields inherently bursty traffic; and (2) CBR traffic gets distorted as it traverses through a multi-hop network. The burstiness of the EF traffic depends on several parameters including the heterogeneity in the bit rates and the packet sizes of the individual CBR flows being aggregated; the percentage of the link bandwidth available to the EF class; and the number of input ports in a core router.

We have generated and experimented with several traces of EF traffic for specific network configurations; however, to explore the design space thoroughly and for ease of parameterizing the traffic burstiness, in this paper, we present results obtained by modeling the EF traffic (i.e., the aggregation of CBR flows at core routers) as an *on-off* source, with exponentially distributed on- and off-durations. We experiment with different levels of burstiness in the EF traffic by selecting a wide range of values for the average on- and off-durations. During the on-durations, such a source transmits packets of size 1500B at the link speed. The fraction f of the bottleneck link bandwidth occupied by the EF traffic is given by:

$$f = \frac{T_{on}}{T_{on} + T_{off}}$$

where T_{on} and T_{off}, respectively, are the average on- and off-durations of the on-off source.

As for the TCP traffic, we use *TCP-Reno* [2], the most popular and widely deployed version of TCP in the Internet. TCP-Reno employs the *slow-start*, *congestion-avoidance*, *fast retransmit*, and *fast recovery* algorithms for congestion control [2]. The throughput achieved by a TCP flow is governed by the available network bandwidth, the maximum receiver-advertised window size, and the rate at which data is generated at the source for transmission. To isolate the effects on TCP performance of fluctuations in the available bandwidth and link buffers caused by higher priority EF traffic, we assume: (1) a large value for the receiver-advertised window size; and (2) back-logged TCP data sources, characterizing the long file transfers resulting from ftp and http on the Internet. This ensures that the TCP throughput is limited only by the available network bandwidth. Finally, we assume that the TCP flows send packets of size 536B, which is representative of a large number of TCP flows in the Internet [4].

2.2 Experimental Design

We conduct experiments to study the impact of higher priority EF traffic on the *throughput* and *fairness* of TCP.

1. **Throughput**: To gain basic understanding on the effect of higher priority EF traffic on the throughput of TCP, we first consider a simple network in which the EF traffic shares the bottleneck link with a single TCP flow.

We study the impact on the throughput of the TCP flow of different levels of burstiness in the EF traffic and different round-trip propagation latencies for the TCP flow (see Section 3.1).

We then consider a more realistic case where multiple TCP flows share the bottleneck link with the EF traffic (see Section 3.2). We consider TCP flows with equal and unequal round-trip propagation latencies. For both the settings, we measure the effect on the *aggregate* TCP throughput of (1) different levels of burstiness in the EF traffic and (2) increasing the number of TCP flows sharing the bottleneck link with the EF traffic.

2. **Fairness**: We study the impact of the higher priority EF traffic on TCP fairness in two network settings. First, we consider a network setting in which all best-effort traffic is carried by TCP flows. Second, we consider a network setting, similar to the current Internet, in which the best-effort traffic consists of a mixture of TCP and UDP flows, with UDP traffic occupying roughly 5% of the total best-effort traffic [5, 7]. In both settings, we compare the fairness of bandwidth allocation among the TCP flows in the presence and absence of EF traffic.

We have conducted these experiments in network environments where the EF traffic occupies different percentages—5%, 10%, 20%, and 30%—of the link bandwidth. Due to space constraints, in the following sections, we report only the results obtained from the set of experiments where the EF traffic occupies 30% of the link bandwidth; the conclusions and observations we report hold for all of the percentages.

2.3 Performance Measures

We measure throughput and fairness of TCP flows as follows.

1. *Throughput*: For an individual TCP flow, throughput is defined as the ratio of the total number of bytes received to the total time required for transmission. We compute the network utilization U achieved by TCP flows as the ratio of the aggregate throughput of TCP flows to the bottleneck link bandwidth available to the TCP flows.

   To measure the impact of burstiness in EF traffic on TCP throughput, we consider two types of networks: (1) a *TCP-CBR Network*, in which the TCP flows share the bottleneck link with a higher-priority CBR traffic; and (2) a *TCP-EF Network*, in which the TCP flows share the bottleneck link with a higher-priority bursty EF traffic. The CBR and the EF traffic impose the same average load; hence, in both networks, the bottleneck link bandwidth available to the TCP flows is the same. We then measure the loss in TCP throughput as the difference $(U_{CBR} - U_{EF})$, where U_{CBR} and U_{EF}, respectively, denote the network utilization achieved by TCP flows in the TCP-CBR and TCP-EF networks.

2. *Fairness*: The literature contains two measures—the *fairness index* F [15] and the *min-max ratio* M [18]—for measuring fairness. If $x_i (x_i \geq 0)$ denotes the throughput received by flow i, then:

$$F = \frac{\left(\sum_{i=1}^{n} x_i\right)^2}{n * \sum_{i=1}^{n} x_i^2}; \quad M = min_{i,j} \left\{ \frac{x_i}{x_j} \right\}$$

   Observe that F varies from $\frac{1}{n}$ (total unfairness) to 1 (total fairness), whereas M takes values between 0 (total unfairness) and 1 (total fairness). While the fairness index represents the fairness of resource allocation in general, the min-max ratio reflects fairness as perceived by individual users. For instance, if the throughput received by flow k is zero, and all other flows receive equal throughput $x > 0$, then $M = 0$ and $F = 1 - \frac{1}{n}$. When $n \to \infty, F \to 1$. Thus, the fairness index can be infinitely close to its optimal value even though from the perspective of flow k, the network is extremely unfair. Since our objective is to study the impact of higher priority EF traffic on the fairness perceived by individual flow, in the rest of this paper, we will use the *min-max* ratio as the fairness measure.