

A simple FIFO-based scheme for differentiated loss guarantees

Yaqing Huang ^{*}, Roch Guérin

Department of Electrical and Systems Engineering, University of Pennsylvania, 200 South 33rd Street, Philadelphia, PA 19104, United States

Received 29 July 2005; received in revised form 10 April 2006; accepted 13 July 2006

Available online 15 August 2006

Responsible Editor: E. Chong

Abstract

Today's Internet carries traffic from a broad range of applications with different requirements. This has stressed its original, one-class, best-effort model, and has been a major driver of the many efforts aimed at introducing QoS. These efforts have, however, been met with only limited success, in part because the complexity they add is often at odds with the scalability requirements of the Internet. This has motivated many investigations for solutions that offer a better trade-off between service differentiation and complexity. This paper shares similar goals and proposes a simple scheme, Bounded Random Drop (BRD), that supports multiple service classes and is implemented using a single FIFO queue and a basic random dropping mechanism. BRD focuses on loss differentiation, as although losses and delay are both important, the steady rise of Internet link speeds is progressively limiting the impact of delay differentiation. It offers strong loss differentiation capabilities, and does not require traffic profiles or admission controls. BRD guarantees each class losses that, when feasible, are no worse than a specified bound, while enforcing differentiation only when required to meet those bounds. The performance of BRD is investigated for a broad range of traffic mixes and shown to consistently achieve its design goals. © 2006 Elsevier B.V. All rights reserved.

Keywords: QoS; Queue management; Diff-Serv

1. Introduction

The traffic carried today over IP networks has evolved from a relatively homogeneous mix of basic data sources to a diverse set of applications with varying requirements and importance. This widening range of requirements has been behind the many efforts aimed at introducing service differentiation in

the Internet. However, the success to date of these efforts has been limited. This has been attributed by many to the intrinsic conflict that exists between the added complexity associated with service differentiation, and the scalability requirements of a continuously growing network. As a result, there have been a number of proposals aimed at offering some form of service differentiation while keeping complexity low. For example and of particular relevance to this paper, the Proportional Differentiated Services model [1,2] is one such effort.

This paper has a similar target, namely, providing different levels of service in IP networks while

^{*} Corresponding author. Tel.: +1 215 898 9351; fax: +1 215 573 2068.

E-mail addresses: yaqing@seas.upenn.edu (Y. Huang), guerin@ee.upenn.edu (R. Guérin).

introducing minimum additional complexity. We expand later on the various aspects of complexity when implementing service differentiation, but it broadly consists of implementation, deployment, and management complexity. Our goal is to develop a solution that while effective at enforcing different levels of service, introduces minimal added complexity along all above three dimensions and can be deployed incrementally in the network. Specifically, we are targeting a solution that from an implementation complexity perspective, requires little more than a simple FIFO queue. As we shall see, the only addition we consider is in the form of a random drop decision logic through which the different levels of service are enforced. This random drop logic calls for the *a priori* configuration of a single parameter for each offered service class, so that deployment complexity is also kept to a minimum. Finally, the system automatically adapts to the level of traffic in the different service classes, without the need for interactions between users and the network besides the *a priori* identification of the service class to which an user belongs. In other words, there is no need for active management of resources.

The mechanism we propose, called Bounded Random Drop (BRD), focuses on loss differentiation. There are two major sources of impairment in IP networks, packet loss, and queuing delay. Both are caused by network congestion that arises when the incoming traffic exceeds the network resources, i.e., link bandwidth and buffer space. However, over the last few years the speed of network links, including access links, has been steadily rising at a pace that exceeds that of the growth in buffer size [3]. As a result, the relative contribution of queuing delays to the end-to-end delay has been regularly decreasing. In contrast, losses are unaffected by the higher speed as they remain a function of the network load. This does not mean that delay has become irrelevant and that only losses matter, but points to losses as the increasingly dominant metric. This is the main motivation behind our focus on losses. Specifically, the paper investigates the possibility of providing per-hop differentiated loss guarantees without upstream policing, knowledge of traffic profiles, or exchange of signalling messages. The choice of per-hop guarantees, as opposed to end-to-end guarantees, is again motivated by our goal of minimum complexity and by the fact that most flows typically encounter only a few bottlenecks on their path. As a result, BRD per hop guarantees on bottleneck links

should offer a reasonable approximation of end-to-end guarantees.

There have been a number of previous works that share similar goals as ours. Several of these works originated from the proportional differentiated services model proposed in [1,2,4], and therefore share similar limitations in both performance and implementation complexity. More specifically, they focus primarily on long-term average loss performances and typically require more complicated implementations than what we consider in this paper. We will discuss these related schemes, particularly the schemes proposed in [5–9], and illustrate the differences that exist between them and our scheme in Sections 3 and 5.1.

The main contributions of this work are in proposing a *simple* FIFO-based scheme, BRD, that is effective at enforcing loss differentiation, and can be deployed relatively easily. BRD will gradually improve the overall loss performances if it is incrementally deployed across the network. The rest of the paper is organized as follows: Section 2 articulates more precisely the goals and requirements of BRD. Section 3 reviews a number of other works that share to different degrees some of our goals and discusses major differences. Section 4 is devoted to a more formal description of the algorithm on which BRD relies, while Section 5 evaluates BRD's performance through simulations.

2. Problem description

We assume that the network traffic can be categorized into N traffic classes, with the traffic intensity of each class unknown ahead of time. At a given hop, we assume there are *absolute loss guarantees* for each class, namely, each class specifies a bound on its loss rate LB_i , $i \in [1, N]$. Our definition of loss extends to both short-term and long-term loss performance. A significant portion of traffic flows in the current Internet, web traffic in particular, is of short duration [10,11]. Enforcing only long-term loss guarantees may, therefore, not be of much benefit to many applications. In addition, since some applications are more sensitive to losses or are simply deemed more important because their users are willing to pay more for better performances, it is natural to also require *relative loss guarantees* among the N classes. We adopt the definition of relative loss guarantees used in [1,2], i.e., *Class i always has better (or at least no worse) loss performance than Class j , $\forall j > i$, regardless of load*

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات