Available online at www.sciencedirect.com
ScienceDirect
Computer Networks 50 (2006) 3353–3374
Computer Networks
www.elsevier.com/locate/comnet

The QoSbox: Quantitative service differentiation in BSD routers ☆

Nicolas Christin [a,*], Jörg Liebeherr [b]

[a] *Information Networking Institute and CyLab Japan, Carnegie Mellon University, 1-3-3-17 Higashikawasaki-cho, Chuo-ku, Kobe 650-0044, Japan*
[b] *The Edward S. Rogers Sr. Department of Electrical and Computer Engineering, University of Toronto, 10 King's College Road, Toronto, ON, Canada M5S 3G4*

Received 11 March 2005; received in revised form 23 October 2005; accepted 4 January 2006
Available online 9 February 2006

Responsible Editor: G. Ventre

Abstract

We describe the design and implementation of the QoSbox, a configurable IP router that provides per-hop service differentiation on loss, delays and throughput to classes of traffic. The novel aspects of the QoSbox are that (1) the QoSbox does not rely on any external component (e.g., no traffic shaping and no admission control) to provide the desired service differentiation, but instead, (2) dynamically adapts packet forwarding and dropping decisions as a function of the instantaneous traffic arrivals and allows for temporary relaxation of some service objectives; also, (3) the QoSbox can enforce both absolute and proportional service differentiation on queuing delays, loss rates, and throughput at the same time. We focus on a publicly available implementation of the QoSbox in BSD-based PC-routers. We evaluate our implementation in a testbed of BSD routers over a FastEthernet network, and we sketch how the QoSbox can be implemented in high speed architectures.
© 2006 Elsevier B.V. All rights reserved.

Keywords: Quality-of-Service implementations; Service differentiation; PC-Routers; BSD; High-Speed networks

1. Introduction

The capacity of the network links of the Internet has steadily increased in the past decade, and has mostly resulted in an increase of the traffic that uses these links, without any improvement in the service received. Stated differently, increasing the capacity of the network far beyond what is needed to accommodate typical loads, a practice known as *overprovisioning*, does not always translate in improved

☆ Most of this work was done while both authors were with the University of Virginia. This work was supported in part by the National Science Foundation through grants ANI-9730103 and ANI-0085955.
* Corresponding author. Tel.: +81 78 360 6318; fax: +81 78 360 1618.
E-mail addresses: nicolasc@cmu.edu (N. Christin), jorg@comm.utoronto.ca (J. Liebeherr).

1389-1286/$ - see front matter © 2006 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2006.01.002

quality-of-service (QoS). Indeed, the service received by an end-to-end traffic flow is bounded by the service received at the link with the smallest capacity on the end-to-end path. Thus, augmenting the capacity of some links only moves the bottleneck to another part of the network, and consequently, only changes the location of the service bottleneck, e.g., from the core to the edge of a network. In particular, access links now appear to be one of the most likely points of congestion in the network [37].

An approach pursued by some Internet Service Providers (ISPs) to improve the average service received by the traffic they serve, is to simply filter out ''undesirable,'' resource-consuming traffic. However, network measurements (e.g., [43]) indicate that end users may actually be particularly interested in resource-consuming traffic such as IP telephony or peer-to-peer transfers, which provides an economic incentive for most ISPs to avoid traffic filtering, lest they lose an important segment of their customer base.

A more desirable solution, at least from the end users' perspective, is to have the network support service differentiation between different classes of traffic. Numerous service architectures for enhancing the best-effort service provided by the Internet have been proposed over the past decade (see [18] for an overview). However, only very few proposals have actually been deployed to some extent, in large part due to the complexity of the required arithmetic manipulations or the amount of state information to be maintained in network routers.

We have previously shown that one can offer relatively strong service differentiation, by solely relying on dynamically adaptive scheduling and buffer management at the bottleneck links of the network [11]. The architecture we proposed hinges on a specific feedback control-based packet handling algorithm to be implemented in routers, and does not require admission control, traffic policing, or network reservations, which could make it more easily deployable than some of the proposed alternatives (e.g., [7]).

The main contribution of the present paper is to demonstrate, through the description of a proof-of-concept implementation, that the router algorithms we propose are efficient enough for practical use at typical access link speeds. To that effect, we describe an implementation in BSD-based PCs of the QoSbox, an IP router providing per-hop service differentiation to classes of traffic. We show that this implementation, available to the public through the KAME [2] and ALTQ distributions, can be readily deployed in access networks with links in the order of a few hundreds of megabits per second, and we outline how the algorithms used can be further modified to be implemented in switch architectures governing high-speed links, in the order of several gigabits per second, at the expense of some reduced accuracy in the service differentiation.

Different from most QoS architectures, the QoSbox does not rely on any external component such as traffic shaping or admission control to enforce the desired service differentiation, which allows for deploying a network of QoSboxes in an incremental manner. Instead, the QoSbox dynamically adapts packet forwarding and dropping decisions dependent on the instantaneous traffic arrivals. The QoSbox can enforce both absolute bounds and proportional service differentiation on queuing delays, loss rates, and throughput at the same time. Since there is no admission control, it may not be feasible to enforce all absolute bounds at all times, in which case, the QoSbox relaxes some bounds according to an order specified a priori.

This paper is organized as follows. In Section 2, we present an overview of the proposed service architecture, and describe the mechanisms used by the QoSbox. In Section 3, we discuss implementation issues. In Section 4, we assess the efficiency of the QoSbox in a testbed of PC-routers. In Section 5, we present the related work. Finally, we summarize the current status of the project and draw brief conclusions in Section 6.

2. The QoSbox

In this section, we discuss the architecture of the QoSbox. To that effect, we first summarize the objectives and characteristics of the service the QoSbox offers. We then discuss the mechanisms implemented in the QoSbox to realize the proposed service.

2.1. Overview

Our objective is to provide a service architecture that can provide strong service differentiation in a large-scale network. More precisely, we want the service architecture to have the following properties: (1) it must be deployable in an incremental manner, without ever degrading the average level of service provided in the network, (2) it must scale to net-

Download English Version:

https://daneshyari.com/en/article/452533

Download Persian Version:

https://daneshyari.com/article/452533

Daneshyari.com