Available online at www.sciencedirect.com
ScienceDirect
Computer Communications 30 (2007) 3859–3869
computer communications
www.elsevier.com/locate/comcom

Hop-count based probabilistic packet dropping: Congestion mitigation with loss rate differentiation

Xiaobo Zhou *, Dennis Ippoliti, Terrance Boult

Department of Computer Science, University of Colorado at Colorado Springs, CO 80918, United States

Received 19 January 2007; received in revised form 14 September 2007; accepted 19 September 2007
Available online 26 September 2007

Abstract

Network applications and users have very diverse service expectations and requirements, demanding for provisioning different levels of quality of service on the Internet. As the speed of network links has been rising at a pace that exceeds that of the growth in the buffer size, packet loss rate differentiation has been an active research topic. However, none of the existing packet dropping schemes for loss rate differentiation considered an important issue, that is, the retransmission overhead of dropped packets. In this paper, we design a hop-count based probabilistic packet dropper (HPPD) for congestion mitigation and loss rate differentiation. HPPD aims to meet a unique two-fold objective by two-dimensional loss rate differentiation: the primary one is the congestion mitigation that aims to reduce congestion in the first place by dropping intra-class packets differently based on their maturity levels to reduce retransmission cost; the other is inter-class proportional loss rate differentiation. The maturity level of a packet, the number of hops it has travelled, is inferred from its time-to-live value in the IP header. We propose a novel intra-class nth-root proportional dropping scheme. The scheme reduces retransmission cost by giving higher dropping probabilities to less mature packets while all packets have their forwarding chances. The n is a controllable parameter trading off dropping fairness for congestion mitigation. It provides great controllability to network operators. Simulation results show that HPPD can significantly mitigate the congestion by reducing the retransmission overhead of dropped packets and achieve the proportional loss rate differentiation at the same time.

Keywords: Quality-of-service differentiation; Congestion mitigation; Proportional model; Loss rate; Hop count

1. Introduction

There is an increasing demand of provisioning different levels of quality of service (QoS) on the Internet to support different types of network applications and various user requirements. Differentiated Services (DiffServ) is one of the major efforts to meet the demand [1]. It aims to provide differentiated services between classes of aggregated traffic flows within a router, rather than to offer QoS guarantees to individual flows. To receive different levels of QoS, packets are assigned with different service types or traffic classes at the network edges. DiffServ-compatible routers in the network core perform stateless prioritized packet forwarding or dropping, so called "per-hop behaviors" (PHBs), to the classified packets. Due to its per-class stateless processing, the DiffServ architecture exhibits a good scalability. There are two basic schemes to DiffServ provisioning. Absolute DiffServ aims to provide statistical assurances for a class's received performance measures, such as a minimum service rate or maximum delay. Relative DiffServ is to quantify the quality spacings between different classes. The proportional differentiation model is popular due to its proportionality fairness and predictability [3,4].

Packet delay and loss rate are two key QoS metrics considered in the DiffServ context. Both are caused by network congestion that arises when the incoming traffic is close to or exceeds the network router resources, i.e., link speed and

* Corresponding author. Tel.: +1 719 2623493.
E-mail address: zbo@cs.uccs.edu (X. Zhou).

0140-3664/$ - see front matter
doi:10.1016/j.comcom.2007.09.010

buffer size. The likelihood of packet losses is a very important performance measure for most of Internet traffic, and it is especially important for peak workload control [3]. The past few years have witnessed that the speed of network links has been rising at a pace that exceeds that of the growth in buffer size [8]. As a result, packet loss rate differentiation in high workload situations has been an active research topic. There are a number of interesting differentiated buffer management and packet dropping schemes. PLR droppers in [3] aim to provide proportional loss rates to different traffic classes according to their pre-specified differentiation weights. JoBS in [15] extends the proportional loss rate model by providing both absolute loss and delay guarantees and proportional differentiations. The proportional differentiation constraint is relaxed to satisfy the absolute constraints when the two sets of constraints cannot be simultaneously satisfied. BRD dropper in [8] seeks to minimize the loss rate differences between traffic classes subject to the absolute loss constraints and the relative loss constraints. Those dropping schemes are able to effectively achieve their differentiation objectives. However, none of them considered an important issue, that is, the retransmission overhead of dropped packets.

A dropped packet might be retransmitted by protocols such as TCP or by end applications. Intuitively, dropping a packet which has travelled 20 hops results in more retransmission overhead and hence heavier congestion in networks than dropping a packet which has only travelled 2 hops. TCP's strategy is to control congestion once it happens, instead of trying to avoid congestion in the first place. An alternative to congestion control is congestion avoidance, which is to predict when congestion is about to happen and then to reduce the rate at which end nodes send data just before packets start being discarded. Previous studies have found that hop count distributions at gateways/routers are usually bell-shaped and the Gaussian distribution is a good first-order approximation; refer to [11] for the hop count distribution of a well-connected commercial server net.yahoo.com. To mitigate the congestion at the first place, our motivation is that packets should be given different dropping probabilities according to the number of hops they have travelled so as to reduce their retransmission overhead.

In this paper, we design a novel hop-count based probabilistic packet dropper (HPPD). The hop count of a packet is inferred from its time-to-live (TTL) value. HPPD provides two-dimensional loss rate differentiation. Within a traffic class, a less mature packet, which has a lower hop count value, has higher probability to be dropped than a more mature packet so as to reduce the retransmission overhead for dropped packets. We refer to this as *intra-class loss differentiation*. Between traffic classes, a packet from a low priority class has higher probability to be dropped than a packet from a high priority class. We refer to this as *inter-class loss differentiation*. Thus, HPPD aims to meet a two-fold objective: the primary one is the congestion mitigation that tries to reduce congestion in the first place by reducing retransmission overhead of dropped packets; the other is proportional loss rate differentiation for DiffServ provisioning. DiffServ is concerned with PHBs, since it is stateless. However, TTL brings some global state of a packet to the routers. The uniqueness of our work lies in the use of hop count information inferred from TTL in making differentiated packet dropping decisions so as to achieve congestion mitigation and loss rate differentiation at the same time.

There are two important considerations to the success of HPPD, i.e., *congestion mitigation*, and *fairness*. A simple way to reducing retransmission overhead is to drop packet strictly according to their maturity levels. A packet with a lower hop count will be always dropped before a packet with a higher hop count if there is backlogged one. However, this kind of strict priority dropping may lead to starvation for packets from neighbor routers. A feasible scheme should give packets with different hop count different chances to be forwarded. We propose a novel intra-class nth-root proportional dropping scheme for HPPD. The n is a controllable parameter trading off dropping fairness for congestion mitigation. Simulation results show that HPPD can significantly mitigate network congestion by reducing the retransmission overhead of dropped packets and achieve the proportional loss rate differentiation at the same time.

The structure of the paper is as follows. In Section 2, we review existing loss rate differentiation schemes. Section 3 presents the HPPD dropping strategies. Section 4 presents the design and implementation issues. Section 5 focuses on the performance evaluation. Section 6 concludes the paper with remarks in future work.

2. Related work

Congestion control and avoidance has been studied extensively in computer communications and networks. Early studies in packet networks included slow start [10], early random drop [7], and random early detection [6]. There are recent studies in buffe management for congestion mitigation and avoidance in wireless sensor networks [2,5,9]. There are also recent studies that propose to utilize the loss information to design robust active queue management [13,19]. Our approach is different that it utilizes the hop count information, inferred from TTL information, to execute loss rate differentiation so as to reduce the retransmission cost of dropped packets and mitigate the congestion in networks. It is complementary to previous work on congestion mitigation and control.

Loss rate differentiation in packet networks is an active research topic. Representative differentiated packet dropping schemes include PLR(∞), PLR(M), JoBS, and BRD. PLR(∞) and PLR(M) [3] were proposed for providing proportional loss rate differentiation. A packet class is assigned a loss differentiation parameter. The PLR schemes adjust the normalized loss rates so that they are eventually