Available online at www.sciencedirect.com
SCIENCE DIRECT
Computer Communications 29 (2005) 103–113
computer communications
www.elsevier.com/locate/comcom

Improvement of LRU cache for the detection and control of long-lived high bandwidth flows

Lichang Che*, Bin Qiu, Hong Ren Wu

School of Computer Science and Software Engineering, Monash University, Wellington Road, Melbourne, Vic. 3800, Australia

Received 12 October 2004; revised 3 May 2005; accepted 9 May 2005
Available online 2 June 2005

Abstract

The issue of long-lived high bandwidth flow detection and control with partial flow state is addressed in this paper. The starting point of this work is the least recently used cache proposal, which is generally effective in detecting long-term fast (LTF) flows. However, its performance can be impaired by the following two factors. The first one is the existence of a large amount of short-lived or slow (SLS) flows. Because randomly sampled SLS flows take up quite a lot cache spaces, long-lived high bandwidth flows are often expelled from the cache mistakenly, which leads to degraded performance. The second is the attack from malicious high-bandwidth on/off flows, which can avoid detection by injecting packets on and off alternatively. This paper proposes a novel cache update scheme, Landmark LRU scheme, to improve the performance of the LRU cache in the detection and control of long-lived high bandwidth flows. Firstly, the superiority of the Landmark LRU scheme over the original LRU cache is demonstrated by Markovian analysis. Then simulations driven by real Internet traces prove that the Landmark LRU scheme achieves higher accuracy in long-lived high bandwidth flow detection than the original LRU scheme. Finally, the Landmark LRU cache is incorporated into a well-known active queue management scheme, Random Early Detection, to show its effectiveness in the control of malicious on/off LTF applications. NS-2 simulation shows that the L-LRU-RED scheme outperforms the LRU-RED stably in regulating on/off flows.

Keywords: Communications Quality of Service (QoS); High bandwidth flow detection; Congestion control; Traffic management

1. Introduction

Over recent years, the detection and control of long-lived high bandwidth flows (LTFs)[1] at routers have received much attention [1–4]. The importance of detecting LTF flow is appreciated for a number of reasons. Firstly, it is understood that LTF flows are the main contributors of congestion, and there is little performance gain in congestion control by dropping packets from short-lived Web flows [5]. However, dropping packets preferentially from LTF flows can significantly improve the Quality of Service (QoS) perceived by end users [6]. Secondly, it has been widely reported that most Internet traffic is from LTF flows, i.e. the Internet traffic follows heavy-tailed distribution, which implies that traffic management and control that concentrates on LTF flows can greatly reduce the implementation cost of network management [4]. Thirdly, it has been shown that non-responsive high-bandwidth applications can consume most of link bandwidth during network congestion. In this situation, the detection and control of non-responsive LTF flows can effectively protect responsive flows from resource starvation [1].

In terms of the realization of LTF flow detection and control, it has been argued that traditional per-flow state based schemes tend to be unnecessarily expensive [2], since a great number of Internet flows are short-lived. This understanding has motivated many recent queue/buffer management schemes that utilize partial flow state to effectively identify and control LTF flows at routers [2,3]. Random Early Detection (RED) [7] with Preferential Dropping (RED-PD) [2] employs the history of the packets randomly dropped by RED to find high bandwidth flows, by comparing the number of dropped packets from a flow with a pre-defined threshold. In [3], a flow table and a virtual

* Corresponding author. Tel.: +61 3 9905 9007; fax: +61 3 9905 5157.
E-mail addresses: carski@csse.monash.edu.au (L. Che), bq@csse.monash.edu.au (B. Qiu), hrw@csse.monash.edu.au (H.R. Wu).

[1] A flow is differentiated by source/destination IP pair in this paper.

0140-3664/$ - see front matter © 2005 Elsevier B.V. All rights reserved.
doi:10.1016/j.comcom.2005.05.002

packet buffer is used by Pan et al. to find high bandwidth flows by counting the number of packets for flows with packet(s) in the buffer. While there exist some more sophisticated yet complex algorithms to identify LTF flows for the purpose traffic measurement, such as [4], this paper aims to find simple and light-weighted solution to identify LTF flows.

The Least Recently Used (LRU) cache scheme [1] is the one that achieves high accuracy in LTF flow detection with partial flow state [5]. It employs a cache of fixed size to accommodate flows recently sampled by a router with a constant probability. By constantly bringing the most recent incoming flow to the top of the cache and replacing least recently hit flow with newly sampled flow, only flows which are fast enough and last long enough can stay in the cache with high probability. To identify LTF flows, each flow in the cache maintains a packet/byte counter and the counter is updated for each incoming packet. When the counter exceeds a pre-defined threshold, the flow is reported as LTF flow.

While the LRU scheme works well when cache size is large, its performance suffers when cache gets smaller. This is because the number of SLS flows is so large that even with a small sample probability [5], the LRU cache still has to remove LTF flows frequently to make room for SLS flows to be added into the cache. As a result, many real LTF flows are expelled from the cache before its packet count exceeds the threshold, which can lead to increased false negative ratio, i.e. more LTF flows are not reported. Although some LTF flows might be detected at later stage, the prolonged detection process has adverse impact on traffic management and control.

Not only SLS flows, but also malicious on/off flows can cause trouble for the LRU cache. Because a flow will be flushed out of the cache soon after it stops sending traffic, a malicious flow that pumps traffic on and off alternatively can easily escape from the LRU cache. As a result, it can keep its packet count lower than the threshold and will never be reported as long-term high bandwidth flow. This weakness can be exploited by malicious flows to encroach a large bandwidth share when the LRU cache is applied to resource management [1].

Intuitively, those two issues above can be addressed by simply increasing the cache size. However, since the LRU scheme searches and updates the LRU cache at the order of line speed for each incoming packet, the implementation of the LRU cache requires expensive high-speed static RAM or Content Addressable Memory (CAM). Therefore, adopting large-size cache can cause scalability concern for the LRU scheme to be deployed at high-speed routers. Based on these observations, this paper proposes a novel cache update scheme, Landmark LRU (L-LRU) cache, to improve the performance of the LRU cache in the detection and control of LTF flow with small cache size.

The L-LRU scheme reduces the possibility of LTF flows being flushed out of the cache by placing a new flow at a fixed Landmark location, instead of at the top of the cache. By bringing flows that are fast and long enough only to the top of the cache, the L-LRU cache virtually allocates a very small proportion of cache space for those newly added flows. As a result, SLS flows are removed from the cache faster and LTF flows can be tracked in the cache much longer. Markovian analysis shows that the cache size required by the L-LRU cache to track high bandwidth flows at fixed probability in the cache is smaller than that required by the original LRU cache. Simulations driven by the Internet traffic traces further demonstrate that the L-LRU achieves significantly higher detection performance than the LRU scheme.

By reducing the number of flows added at the top of the cache, the L-LRU cache greatly increases the cache occupancy time for LTF flows. Therefore, it is much less likely for malicious on/off flows to escape from the cache, which implies that the L-LRU cache is applicable for resource management for better control of malicious on/off flows. As an example of such applications, in this paper the L-LRU cache is incorporated into the famous RED scheme for Active Queue Management (The new scheme is referred to as L-LRU-RED hereafter). NS-2 [8] simulation is conducted to evaluate the L-LRU-RED scheme. Simulation results prove that the L-LRU-RED regulates on/off flows effectively.

The remaining part of the paper is organized as follows. Section 2 briefly reviews the LRU scheme and identifies its weakness. The proposed L-LRU cache is presented in Section 3. Section 4 gives the detailed performance evaluation and related discussions. The paper is concluded in Section 5.

2. Background and motivation

Parameters used in the following discussions are listed below and these parameters will be explained in detail where applicable.

- C link bandwidth
- P packet size (bytes/packet)
- h probability that flows can be found in the cache
- m probability that flows cannot be found in the cache and $m=1-h$
- ψ the percentage of reference bandwidth over link bandwidth
- r_t reference bandwidth $r_t=\psi C$

The original LRU scheme [1] adopted a cache with a fixed size s to identify LTF flows. Fig. 1 gives the illustration of the LRU cache. Routers equipped with the LRU cache search the cache for the flow *id* (source/destination IP address pair) of all incoming packets. If the flow *id* of the incoming packet is found, which is referred to as *hit* event hereafter, the cache item for that flow