Computer Networks 53 (2009) 2444–2457

Contents lists available at ScienceDirect

Computer Networks

journal homepage: www.elsevier.com/locate/comnet

Transport-independent fairness

B. Behsaz, P. Gburzynski, M. MacGregor *

Dept. of Computing Science, University of Alberta, Edmonton, Alberta, Canada

ARTICLE INFO

Article history:
Received 26 August 2008
Received in revised form 13 January 2009
Accepted 28 April 2009
Available online 12 May 2009

Responsible Editor: G. Morabito

Keywords:
TCP
P2P
Congestion control
Optimization
Convex programming

ABSTRACT

The Internet relies on cooperative endpoints to react to signals from the network that congestion is occurring. In particular, TCP interprets packet loss as a signal of congestion. However there are many new non-cooperative protocols in use which attempt to exploit the network aggressively and do not reduce their demands when the network signals congestion. We propose the aggregate control of "fluxes" defined by policies at individual routers. Each router can then calculate an optimal allocation of bandwidth to each flux contending for a given output link. We propose a combined hill climbing and convex programming method for this optimization, which we call HCCP. HCCP is designed to punish greedy fluxes rather than just regulating them: such fluxes may find their bandwidth allocation reduced to zero if they are sufficiently aggressive. Our results show that HCCP is effective at regulating a wide range of rather generally characterized transport protocols. We explore the use of both throughput maximization and proportionally fair allocation and recommend the latter because the former often leads to the situation where one or more fluxes receive zero bandwidth.

© 2009 Elsevier B.V. All rights reserved.

1. Introduction

The network layer of the Internet as implemented by IP is connectionless. This means that all packets appear to it as independent datagrams. The concept of a session, understood as a sequence of related packets, is implemented at the edges of the network and formally the network layer can be completely oblivious of that concept. This simple paradigm has both advantages and drawbacks. Its obvious advantage is simplicity: routers can be made memory-less with respect to the packets they are forwarding. Their sole responsibility is to map destination addresses of incoming packets to outgoing links. The primary drawback is the lack of accountability of datagrams for their sessions. For example, during a congestion event a router may find it difficult to discard packets in a way that would reflect the extent to which particular sessions are responsible for the overload. In other words, the router may find it difficult to be fair.

In the early days of networking, when the networks were separate and confined to small communities with shared interests and goals the simple paradigm of IP was quite adequate. These days, however, the network is global and its users compete for limited bandwidth. One of the major problems facing Internet Service Providers (ISPs) is that a large proportion of the traffic in their networks is transported by aggressive, greedy protocols which serve the needs of disparate communities. Consequently there is a growing need for fairness mechanisms within the network core – mechanisms that can identify users exceeding their fair share of bandwidth and reduce the number of their packets in the network without hurting more considerate citizens. Most of the schemes attempting to address this issue focus on identifying transport layer sessions with the intention of treating them as the "users" whose bandwidth shares and behavior are subsequently assessed and policed [13,28,29,34]. In particular, TCP sessions are easy to identify and monitor because of their explicit connection-oriented nature.

The idea of viewing TCP sessions as the basic units which contribute to network load dates from the infamous

* Corresponding author. Tel.: +1 780 492 7434; fax: +1 780 492 1071.
E-mail address: mike.macgregor@ualberta.ca (M. MacGregor).

1389-1286/$ - see front matter © 2009 Elsevier B.V. All rights reserved.
doi:10.1016/j.comnet.2009.04.011

collapse of the Internet in 1988–1989 and its salvage by Jacobson [19]. The primary objective of any social (i.e., compliant to [19]) implementation of TCP is to navigate towards a state whereby random packet losses at a congested router will result in approximately the same fraction of bandwidth being received by each of the competing greedy sessions. This would have solved the fairness problem in an ideal world where:

(1) All TCP implementations are compliant and preferably identical.
(2) Users are synonymous with TCP sessions so that the fair treatment of individual TCP sessions translates into the fair treatment of actual users.
(3) There is no traffic in the network other than TCP sessions.

All three conditions would have to hold simultaneously; however, none of them in fact holds in today's Internet. First, as the behavior of a TCP session is solely up to the edge host there is little that a router can do to effectively enforce user compliance. The only possible approach is to apply some heuristics to the observed packet arrival process in order to detect misbehaving flows and then penalize those flows by dropping their packets [20]. But this is a cat and mouse game: a session familiar with those heuristics may be able to dynamically adjust the rate of its packets as to fool the router into giving it disproportionately more bandwidth.

Even more importantly, TCP sessions are not synonymous with users. In particular, an end-to-end application needing a lot of bandwidth can easily open multiple TCP sessions. And lastly, all the measures aimed at policing TCP sessions can be circumvented by resorting to UDP and implementing sessions within the application [14,16].

Network operators and users have adopted the idea of Service Level Agreements (SLAs) in which separate domains are individually responsible for enforcing the agreed allocation of resources to support the SLA. In this case, the requirements are known very specifically at the edge of the network, and rather than individual "hosts" the clients of the proposed scheme are other ISPs. Knowledge of the service requirements at an edge router alone, however, is not sufficient to ensure that the rest of the domain will respect the intent of an individual SLA, especially when presented with flow aggregates from other customers in other domains terminating on other edge routers whose resource requirements are supposedly guaranteed by some other SLA. Our proposal represents an effective mechanism that can be used to police the competing interests of multiple customers in a complex network.

For a router, the problem of identifying the actual *users* of the network is difficult and not even particularly well defined without introducing some additional concepts. The primary problem is that (as explained above) the perceptible attributes of packets, as seen by the router, do not allow it to authoritatively decide which "user" those packets belong to. Here by "user" we understand an entity that would be meaningful from the viewpoint of a globally acceptable fairness measure. The behavior of such an entity should also be controllable: a de-facto single user should not be able to fake multiple identities to receive extra bandwidth.

A truly workable mechanism for global fairness must be implemented within the network core, as opposed to at the edge. Such a mechanism must be completely transport-independent, in the sense that it cannot be based on formal transport-layer sessions, be they explicit TCP flows or some conceptual streams derived from packet parameters (e.g., source/destination addresses, ports, or patterns spotted in the IP payload). This is because those parameters can always be modified by the source in order to evade control. In particular, a collusion of hosts within a collaborating group of users may create a configuration whereby different parts of the same sessions follow different paths in the network, as in BitTorrent [25,33]. A reliable fairness enforcement scheme should be comfortable with such complications, which should not be perceived as cases of network abuse, but rather embraced as natural and creative ways of harnessing the power of the connectionless core of the network. If properly designed, such a scheme will also effectively obviate any need for explicit *calls* into the otherwise beautiful paradigm of connectionless operation and go a long way towards providing quality of service, thwarting DoS attacks, and generally solving the problems caused by disparity and unaccountability of bandwidth allocation.

In this paper, we propose a global scheme for fair allocation of network resources using mechanisms at the routers, rather than relying on the transport behavior of endpoint applications. We address this problem at a level of aggregated flows to avoid the usual fallacies associated with attempts to police individual transport layer sessions. The definition of our flows is dynamic and hierarchical so as to account for the different policies assumed by different hosts and routers. Its role is to meaningfully combine those policies and create a natural way of interpreting them as a definition of a globally fair network state.

One central issue is the identification of unfair flows, i.e., those demanding more bandwidth than their fair share, and differentiating their treatment. A second related issue is to make sure that those flows that consistently demand more than their fair share and behave in an uncooperative manner are not able to cheat the policing mechanism. We call such flows *greedy* to differentiate them from those that are only temporarily unfair. The idea is to discriminate more heavily against greedy flows than against those which are only temporarily unfair but which respond cooperatively to network regulation. The way the unfair flows are penalized encourages them to converge to their fair share. This does not assume any particular implementation of those flows at the hosts such as a compliant TCP implementation. The router does not try to guess at or emulate the mechanism used by the hosts in response to packet loss. It merely drops the excess packets of the unfair flow in a certain systematic manner so as to convey a certain universally meaningful *message* to the application. In a nutshell, the message says something like this: *"The more you try to exceed your fair share, the less of the actual bandwidth you are going to get."* The allocated share of bandwidth is not proportional to the flow's apparent demand, but instead it is a decreasing function of the demand

Download English Version:

https://daneshyari.com/en/article/451170

Download Persian Version:

https://daneshyari.com/article/451170

Daneshyari.com