Appeared in IEEE/ACM TRANSACTIONS ON NETWORKING, VOL. 12, NO. 1, FEBRUARY 2004

Network Border Patrol: Preventing Congestion Collapse and Promoting Fairness in the Internet

Célio Albuquerque, Member, IEEE, Brett J. Vickers, Member, IEEE, and Tatsuya Suda, Fellow, IEEE

Abstract—The Internet's excellent scalability and robustness result in part from the end-to-end nature of Internet congestion control. End-to-end congestion control algorithms alone, however, are unable to prevent the congestion collapse and unfairness created by applications that are unresponsive to network congestion. To address these maladies, we propose and investigate a novel congestion-avoidance mechanism called network border patrol (NBP). NBP entails the exchange of feedback between routers at the borders of a network in order to detect and restrict unresponsive traffic flows before they enter the network, thereby preventing congestion within the network. Moreover, NBP is complemented with the proposed enhanced core-stateless fair queueing (ECSFQ) mechanism, which provides fair bandwidth allocations to competing flows. Both NBP and ECSFQ are compliant with the Internet philosophy of pushing complexity toward the edges of the network whenever possible. Simulation results show that NBP effectively eliminates congestion collapse and that, when combined with ECSFQ, approximately max-min fair bandwidth allocations can be achieved for competing flows.

Index Terms—Border control, congestion control, congestion collapse, core-stateless mechanisms, end-to-end argument, Internet, max-min fairness.

I. INTRODUCTION

The fundamental philosophy behind the Internet is expressed by the scalability argument: no protocol, mechanism, or service should be introduced into the Internet if it does not scale well. A key corollary to the scalability argument is the end-to-end argument: to maintain scalability, algorithmic complexity should be pushed to the edges of the network whenever possible.

Perhaps the best example of the Internet philosophy is TCP congestion control, which is implemented primarily through algorithms operating at end systems. Unfortunately, TCP congestion control also illustrates some of the shortcomings of the end-to-end argument. As a result of its strict adherence to end-to-end congestion control, the current Internet suffers from two maladies: congestion collapse from undelivered packets, and unfair allocations of bandwidth between competing traffic flows.

The first malady—congestion collapse from undelivered packets—arises when bandwidth is continually consumed by packets that are dropped before reaching their ultimate destinations [1].

Nagle assigned the term “congestion collapse” in 1984 to describe a network that remains in a stable congested state [2]. At that time, the primary cause of congestion collapse was inefficient use of retransmission timers by TCP sources, which led to the unnecessary retransmission of delayed packets. This problem was corrected with more recent implementations of TCP [3]. Recently, however, a potentially more serious cause of congestion collapse has become increasingly common. Network applications are now frequently written to use transport protocols, such as UDP, which are oblivious to congestion and make no attempt to reduce packet transmission rates when packets are discarded by the network [4]. In fact, during periods of congestion some applications actually *increase* their transmission rates by introducing redundancy in the transmitted data in order to become less sensitive to packet losses [5]. The Internet presently has no effective way to regulate such applications.

The second malady—unfair bandwidth allocation to competing network flows—arises in the Internet for a variety of reasons, one of which is the existence of applications that do not respond properly to congestion. Adaptive applications (e.g., TCP-based applications) that respond to congestion by rapidly reducing their transmission rates are likely to receive unfairly small bandwidth allocations when competing with unresponsive applications. The Internet protocols themselves can also introduce unfairness. The TCP algorithm, for instance, inherently causes each TCP flow to receive a bandwidth that is inversely proportional to its round-trip time [6]. Hence, TCP connections with short round-trip times may receive unfairly large allocations of network bandwidth when compared to connections with longer round-trip times.

The impact of emerging streaming media traffic on traditional data traffic is of growing concern in the Internet community. Streaming media traffic is unresponsive to the congestion in a network, and it can aggravate congestion collapse and unfair bandwidth allocation. Recently, various researchers have documented and studied the problems of unfairness and congestion collapse due to unresponsive traffic, such as streaming media traffic [7]–[12]. This concern regarding the negative impact that streaming media traffic may bring has also been expressed in

END of the translation sample