به نام آن که جان را فکرت آموخت

دانشکده‌ی مهندسی کامپیوتر

دکتر یاشار گنجعلی

استادیار دانشکده‌ی علم کامپیوتر، دانشگاه تورنتو، کانادا

زمان: شنبه ۱۵ مهر ۱۳۸۵، ۱۳:۳۰ تا ۱۴:۳۰
مکان: تالار خوارزمی، دانشکده‌ی م. کامپیوتر

Buffer Sizing in Internet Routers

Abstract

Internet routers require buffers to hold packets during times of congestion. The buffers need to be fast, and so ideally they should be small enough to use fast memory technologies such as SRAM or all-optical buffering. Unfortunately, a widely used rule-of-thumb says we need a bandwidth-delay product of buffering at each router so as not to lose link utilization. This can be prohibitively large. Recently, Appenzeller et al. challenged this rule-of-thumb and showed that for a backbone network, the buffer sizes can be divided by $\sqrt{N}$ without sacrificing throughput, where N is the number of flows sharing the bottleneck.

In this talk, I will explain how buffers in the Internet backbone can be significantly reduced even more, to as little as a few dozen packets, if we are willing to sacrifice a small amount of link capacity. The argument is that if the TCP sources are not overly bursty, then with 20-50 packet buffers we can gain high throughput. Specifically, I will argue that $\mathcal{O}(\log W)$ buffers are sufficient, where W is the congestion window size of each flow. This claim is supported with analysis, variety of simulations, and experiments performed on real networks. The change we need to make to TCP is minimal - each sender just needs to pace packet injections from its window. Moreover, there is some evidence that such small buffers are sufficient even if we don't modify the TCP sources so long as the access network is much slower than the backbone, which is true today and likely to remain true in the future. The conclusion is that buffers can be made small enough for very simple single-chip routers, or even all-optical routers with small integrated optical buffers.

Biography

Yashar Ganjali has recently joined the Department of Computer Science at University of Toronto as an assistant professor. He received a BS degree in Computer Engineering from Sharif University of Technology, and an MS in Computer Science from University of Waterloo. He has a PhD in Electrical Engineering from Stanford University. His PhD dissertation focuses on buffer sizing for Internet routers. The goal is to determine the impact of reducing the buffer size in core routers from millions of packets to just tens of packets, and thus exploring the possibility of building all-optical networks. He has collaborated with Sprint Advanced Technology Labs, Level 3 Communications, Lucent Technologies, Internet2, and Verizon Laboratories on validation of the buffer sizing results. His other research interests include analysis and design of high performance switches, scheduling algorithms, congestion control, routing protocols, and network optimization.