NATIONAL UNIVERSITY OF SINGAPORE

School of Computing

PH.D DEFENCE - PUBLIC SEMINAR

Title:	A Rate-based TCP Congestion Control Framework for Cellular Data Networks
Speaker:	Mr Leong Wai Kay
Date/Time:	25 June 2014, Wednesday, 02:00 PM to 03:30 PM
Venue:	Executive Classroom, COM2-04-02

Supervisor: Dr Ben Leong, School of Computing

Abstract:

Modern 3G/4G cellular data networks have vastly different characteristics from other wireless networks such as Wi-Fi networks. It is also becoming more pervasive with the reducing cost of smartphones and cellular data plans. In this thesis, we investigate the major issues of cellular data networks and propose a radical TCP congestion control mechanism to overcome these problems.

Firstly, cellular data networks are highly asymmetric. Downstream TCP flows are thus affected by a concurrent uplink flow or a congested and slow uplink due to the ACK packets being delayed. Secondly, packet losses are very rare due to the hybrid-ARQ scheme used in the link-level protocol. Thus, this causes the cwnd in the TCP congestion control algorithm to grow until the buffer overflows. As ISPs typically provision huge buffers, this causes the bufferbloat problem where the end-to-end delay becomes very large. Thirdly, recent stochastic forecasting techniques have been used to predict the network bandwidth to prevent excessive sending of packets to reduce the overall delay. However, such techniques are complicated and require a long computation or initialization time and often overly sacrifice on throughput.

To address these issues, we propose a new rate-based congestion control technique and developed a TCP congestion control framework upon which various algorithms can be built on. In our rate-based framework, the sending rate is set by estimating the available bandwidth from the receive rate at the receiver. To achieve stability and to adapt to changing network conditions, we oscillate the sending rate above and below the receive rate which will fill and drain the buffer respectively. By observing the buffer delay, we can choose when to switch between the filling and draining of the buffer. By

controlling the various parameters, we can control the algorithm to optimize for link utilization by keeping the buffer always occupied, or for latency by keeping buffer occupancy low.

We implemented our framework into the TCP stack of the Linux kernel and developed two rate-based algorithms, RRE and PropRate. The algorithms were evaluated using the ns-2 simulator as well as using a trace-driven network emulator, and also tested on real cellular data networks. We show that by controlling the various parameters, the algorithms can optimize and tradeoff between throughput and delay. In addition, we also implemented two state-of-the-art forecasting techniques Sprout and PROTEUS into our framework and evaluated them using our network traces. We found that while forecasting techniques can reduce the delay, a quick reacting rate-based algorithm can perform just as well, if not better by maintaining a higher throughput.

Finally, our work advances the current TCP congestion control technique by introducing a new framework upon which new algorithms can be built upon. While we have showed that our new algorithms can achieve good tradeoffs with certain parameters, how the parameters can be chosen to match the current network conditions is room for further research. Similar to how many cwnd-based congestion control algorithms have been developed in the past, we believe that our framework opens new possibilities in the research community to explore a rate-based congestion control for TCP in emerging networks. In addition, because our framework is compatible with existing TCP stacks, it is suitable for immediate deployment and experimentation.