

Low-latency Datacenter Networks

Soudeh Ghorbani

Soudeh Ghorbani is a researcher in computer networks. She received her PhD from the University of Illinois at Urbana-Champaign in 2016 advised by Brighten Godfrey, and during 2017 was a postdoctoral research associate at the University of Wisconsin working with Aditya Akella. Her research has won a number of awards and fellowships including the VMware Graduate Fellowship (one of the 3 winners in 2015, worldwide), the best paper award at HotSDN, the Feng Chen Memorial Award, and the Gottlieb Fellowship.



Abstract

Datacenters host a wide range of today's low-latency applications. To meet their strict latency requirements at scale, datacenter networks are designed as topologies that can provide a large number of parallel paths between each pair of hosts. The recent trend towards simple datacenter network fabric strips most network functionality, including load balancing among these paths, out of the network core and pushes it to the edge. This slows reaction to microbursts, the main culprit of packet loss -- and consequently performance degradation -- in datacenters. We investigate the opposite direction: could slightly smarter fabric significantly improve load balancing? I will present DRILL, a datacenter fabric which performs micro load balancing to distribute load as evenly as possible on microsecond timescales. DRILL employs per-packet decisions at each switch based on local queue occupancies and randomized algorithms to distribute load. I will explain how we address the resulting key challenges of packet reordering and topological asymmetry and present results showing that DRILL outperforms recent edge-based load balancers, particularly under heavy load while imposing only minimal (less than 1%) switch area overhead. Under 80% load, for example, it achieves 1.3-1.4x lower mean flow completion time than recent proposals. Finally, I will discuss our analysis of DRILL's stability and throughput-efficiency. I will conclude by discussing some of the challenges and opportunities of applying cross-layer micro load balancing to improve the performance of some critical network applications.

MONDAY, MARCH 05, 2018

2:30 PM | LUDDY 4063

