



THE UNIVERSITY OF
CHICAGO

Department of
Computer Science

Muhammad Shahbaz, Stanford University

Fast and Flexible Next-Generation Data Centers



Modern datacenter networks need to be both fast (high throughput and low latency) and flexible (easy to program). With the demise of Moore's Law and the rise of Big Data, end-host CPUs alone cannot meet the performance demands of emerging workloads (e.g., machine learning). Network switches must now take on more responsibilities and offload CPUs---they can no longer be fixed-function devices for packet forwarding only. Doing so, however, may cause networks to sacrifice performance for more flexibility.

Exposing the right domain-specific abstractions and hardware primitives for networks can enable more flexibility with negligible (or zero) loss in performance. In this talk, I will show how a high-level language (P4) can let network programmers customize the behavior of software switches, like encoding state inside packets, without compromising performance. And, how new switch primitives can read and act on the encoded state to forward packets at hardware speed. These features enabled us to scale multicast to millions of groups---a long-standing problem in public clouds.

February 18

10:30 a.m.

Crerar 298

Host: Nick Feamster

Muhammad Shahbaz is a Postdoctoral Research Scholar in the Department of Electrical Engineering at Stanford University. His research focuses on the design and development of domain-specific languages, compilers, and architectures and their applications to emerging workloads (including machine learning). Shahbaz received his Ph.D. and M.A. in Computer Science from Princeton University and B.E. in Computer Engineering from the National University of Sciences and Technology (NUST). Before joining Princeton University, Shahbaz worked as a Research Assistant at the University of Cambridge, Computer Laboratory, and was a Senior Design Engineer at the Center for Advanced Research in Engineering (CARE). In the past, Shahbaz has built open-source systems like PISCES, SDX, and NetFPGA-10G that are widely used in industry and academia. He has earned (and shared) awards, including Internet2 Innovation Award, ACM SOSR Systems Award, Facebook Research Award, APNet Best Paper Award, and Outstanding Graduate Teaching Assistant Award.