COMPUTATIONAL RESEARCH in BOSTON and BEYOND SEMINAR

Microsoft's Production Configurable Cloud

DEREK CHIOU
Microsoft

ABSTRACT:

Microsoft has been building data centers with a programmable hardware in every server, creating a novel Configurable Cloud. The reconfigurable hardware, in the form of a field programmable gate array (FPGA) is cabled between the NIC and the data center network, as well as being attached to the CPUs via PCIe. This architecture enables an FPGA-centric, rather than CPU-centric, computational model since all communication in and out of the server is first processed by the FPGA that handles common tasks without CPU involvement and passes uncommon, complex tasks to the CPU that acts as a "complexity" offload engine. Microsoft has deployed a diverse set of applications, including deep neural networks and software defined networking acceleration, across its Configurable Cloud. I will describe the Cloud, some of its applications, and their performance.

BIO:

Derek Chiou is a Partner Architect at Microsoft where he leads the Azure Cloud Silicon team responsible for FPGAs and ASICs in Microsoft's data centers and an adjunct professor in the Electrical and Computer Engineering Department at The University of Texas at Austin. Before Microsoft/UT, Dr. Chiou was a system architect and lead the performance modeling team at Avici Systems, a manufacturer of terabit core routers.

FRIDAY, NOVEMBER 4, 2016
1:00 PM – 2:00 PM
Building 32, Room 141
(STATA)

Pizza and beverages will be provided.

http://math.mit.edu/crib/

Massachusetts Institute of Technology