北京大学海外名家讲学计划

BLUEDBM: A MUTI-ACCESS, DISTRIBUTED FLASH STORE FOR BIG DATA ANALYTICS

Arvind

Johnson Professor
Massachusetts Institute of Technology

2014年8月24日 星期日 08:45am
中关新园1号楼观湖厅

ABSTRACT: Complex analytics of the vast amount of data collected via social media, cell phones, ubiquitous smart sensors, and satellites is likely to be the biggest economic driver for the IT industry over the next decade. For many "Big Data" applications, the limiting factor in performance is often the transportation of large amount of data from hard disks to where it can be processed, i.e. DRAM. We will present BlueDBM, an architecture for a scalable distributed flash store which is designed to overcome this limitation in two ways. First, the architecture provides a high-performance, high-capacity, scalable random-access storage. It achieves high-throughput by sharing large numbers of flash chips across a low-latency, chip-to-chip backplane network managed by the flash controllers. Second, it permits some computation near the data via a FPGA-based programmable flash controller. We will present the preliminary results on accelerating complex queries using BlueDBM consisting of 20 nodes and up to 32 TB of flash.

BIOGRAPHY: Arvind is the Johnson Professor of Computer Science and Engineering at the Massachusetts Institute of Technology and a member of CSAIL (Computer Science and Artificial Intelligence Laboratory). From 1974 to 1978, prior to coming to MIT, he taught at the University of California, Irvine. Arvind received his M.S. and Ph.D. in Computer Science from the University of Minnesota in 1972 and 1973, respectively. He received his B. Tech. in Electrical Engineering from the Indian Institute of Technology, Kanpur, in 1969, and also taught there from 1977-78. He is a Fellow of the IEEE and the ACM, and he was elected to the National Academy of Engineering in 2008.

主办单位：北京大学国际合作部、北京大学高能效计算与应用中心