SORS: Big Data Applications on Flash Storage with Accelerators

Objectives

Speaker: Arvind, Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology

Abstract: Fast content-based searches and 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. Most such applications are performed on large clusters where the data resides in the collective DRAM of the cluster. Computing on such clusters is expensive and consumes enormous amount of power. We will present a cheaper and cooler alternative which provides high-performance, high-capacity, scalable random-access flash storage, and allows computation near the data via FPGA-based programmable flash controllers. We will discuss the preliminary results for a key-value store (KVS) implementation and several algorithms – Terabyte Sorts, PageRank and Breadth First Search, on BlueDBM consisting of 20 nodes and 20TB of flash.

Biography: Arvind is the Johnson Professor of Computer Science and Engineering at MIT. Arvind's group, in collaboration with Motorola, built the Monsoon dataflow machines and its associated software in the late eighties. In 2000, Arvind started Sandburst which was sold to Broadcom in 2006. In 2003, Arvind co-founded Bluespec Inc., an EDA company to produce a set of tools for high-level synthesis. In 2001, Dr. R. S. Nikhil and Arvind published the book "Implicit parallel programming in pH". Arvind's current research focus is to enable rapid development of embedded systems. Arvind is a Fellow of IEEE and ACM, and a member of the National Academy of Engineering and the American Academy of Arts and Sciences.

Barcelona Supercomputing Center - Centro Nacional de Supercomputación

Source URL (retrieved on *15 set 2024 - 11:55*): https://www.bsc.es/ca/research-and-development/research-seminars/sors-big-data-applications-flash-storage-accelerators