COMPUTATIONAL
SCIENCE
ZURICH

Distinguished Lecture Series

The Neo-Digital Age – When Moore's Law Died

Prof. Dr. Thomas Sterling, Indiana University
Professor of Electrical Engineering at the School of Informatics and Computing Department of Intelligent Systems Engineering (ISE) and Director of the IU Center for Research in Extreme Scale Technologies (CREST)

Wednesday, March 29, 2017 at 16:15 at the University of Zurich, Aula KOL G 201 (apéro after the lecture)

Even the highest scale contemporary conventional HPC system architectures are optimized for the basic operations and access patterns of classical matrix and vector processing. These include emphasis on FPU utilization, high data reuse requiring temporal and spatial locality, and uniform strides of indexing through regular data structures either dense or sparse. Such systems in the 100 Petaflops performance regime such as the Chinese Sunway Taihu-Light, and the US CORAL Summit and Aurora to be deployed in 2018 in spite of their innovations still are limited in these properties. Emerging classes of new application problems in data analytics, machine learning, and knowledge management demand very different operational properties in response to their highly irregular, sparse, and dynamic behaviors exhibiting little or no data reuse, random access patterns, and meta-data dominated processing. Close examination clearly suggests that at the core of these "big data" applications is dynamic adaptive graph processing which is in some ways diametrically opposite to conventional matrix computing. Of immediate importance is the need to significantly enhance efficiency and scalability as well as user productivity, performance portability, and reduce energy. Key to this is the introduction of powerful runtime system software for the exploitation of real-time system status information to support dynamic adaptive resource management and task scheduling. But software alone for runtime functionality will be insufficient for extreme-scale where near fine-grained parallelism is necessary and software overheads will limit efficiency and scalability. A new era of architecture research is beginning in the combined domains of accelerator hardware for both graph processing and runtime systems. This presentation will discuss the nature of the computational challenges, provide examples and describe experiments with a state-of-the-art runtime system software, HPX-5. Future directions in hardware architecture support for exascale runtime-assisted big data computation will be proposed and discussed. Questions and comments from the audience will be welcome throughout the talk.

For further information please visit www.zhcs.ch/news/events/distinguished-lecture-series/

ETH zürich University of Zurich[UZH]