Cambrian Explosion of Computing and Big Data in the Post-Moore Era

Satoshi Matsuoka
Riken Center for Computational Science & Tokyo Institute of Technology
matsu@is.titech.ac.jp

ABSTRACT

The so-called "Moore's Law", by which the performance of the processors will increase exponentially by factor of 4 every 3 years or so, is slated to be ending in 10-15 year timeframe due to the lithography of VLSIs reaching its limits around that time, and combined with other physical factors. We are now embarking on a project to revolutionize the total system architectural stack in a holistic fashion in the Post-Moore era, from devices and hardware, abstracted by system software and programming models and languages, and optimized according to the device characteristics with new algorithms and applications that exploit them. Such systems will have multitudes of varieties according to the matching characteristics of applications to the underlying architecture, leading to what can be metaphorically described as Cambrian Explosion of computing systems. The diverse elements of such systems will be interconnected with next-generation terabit optics and networks, allowing metropolitan-scale computing infrastructure that would truly realize high performance parallel and distributed computing. However, which algorithms and applications would benefit the most from such future computing, given that some physical constants, e.g., communication latency, cannot be improved? We speculate on some of the scenarios that would change the nature of current Cloud-centric infrastructures towards the Post-Moore era.

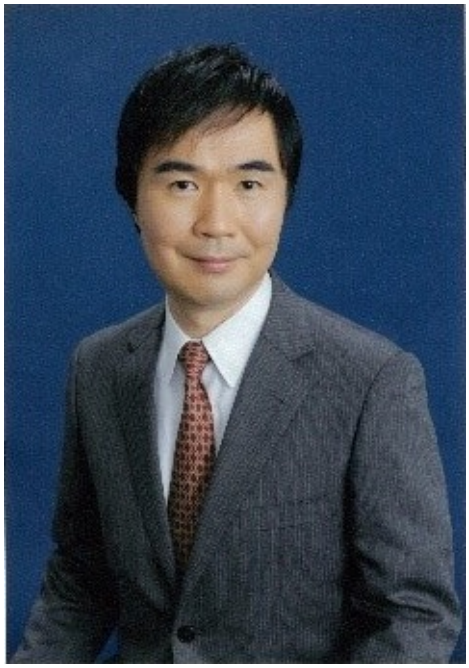

1 BIOGRAPHY

Satoshi Matsuoka is the director of Riken CCS, the top-tier HPC center that represents HPC in Japan, currently hosting the K Computer and developing the next generation Post-K machine, along with multitudes of ongoing cutting edge HPC research being conducted. He has been a Full Professor at the Global Scientific Information and Computing Center (GSIC), a Japanese national supercomputing center hosted by the Tokyo Institute of Technology, since 2000. He is a Fellow at the AI Research Center (AIRC), AIST, the largest national lab in Japan, as well as the head of the joint Lab RWBC-OIL (Open Innovation Lab on Real World Big Data Computing). He was the leader of the TSUBAME series of supercomputers, and has won the 2014 IEEE-CS Sidney Fernbach Memorial Award, the highest prestige in the field of HPC. Besides, he has an appointment as a Professor at the Department of Mathematical and Computing Sciences, Tokyo Institute of Technology, to continue his research activities spanning the three institutions in HPC and scalable Big Data & AI.

Satoshi's research is principally in system software for large scale supercomputers and similar infrastructures such as Clouds for HPC, and more recently, convergence of Big Data/AI with HPC, as well as investigating the Post-Moore Technologies towards 2025. Over the years he has been involved and lead a number of large collaborative projects that worked on basic elements that are now significant for the current and more importantly future exascale systems, such as fault tolerance, low power, strong scalability, programmability, as well as large-scale I/O.

Satoshi has led a number of major projects. His most recent awards include: 1) the Billion-way fault tolerance in supercomputers (2011-2015), which investigates the key technologies for reliable exascale supercomputing; 2) Ultra Green Supercomputing (2011-2015) and it successor Post-Green Project (2015-2019), where his group investigate low power control and cooling technologies to achieve exascale power requirements; and 3) Extreme Big Data (2013-2018), where his group converge HPC/AI and Big Data technologies to handle extreme-scale data that can only be processed with supercomputers.

HPDC '18, June 11–15, 2018, Tempe, AZ, USA

ACM ISBN 978-1-4503-5785-2/18/06.
https://doi.org/10.1145/3208040.3225055