Xiang-ke LIAO, Kai LU, Can-qun YANG, Jin-wen LI, Yuan YUAN, Ming-che LAI, Li-bo HUANG, Ping-jing LU, Jian-bin FANG, Jing REN, Jie SHEN, 2018. Moving from exascale to zettascale computing: challenges and techniques. *Frontiers of Information Technology & Electronic Engineering*, 19(10):1236-1244. https://doi.org/10.1631/FITEE.1800494

Moving from exascale to zettascale computing: challenges and techniques

Key words: High-performance computing; Zettascale; Micro-architectures; Interconnection; Storage system; Manufacturing process; Programming models and environments

Corresponding author: Kai LU
E-mail: kailu@nudt.edu.cn
ORCID: Kai LU, http://orcid.org/0000-0003-2284-7897

Motivation

1. The increasing performance of the No. 1 HPC systems slowed down around 2013.

2. The HPC systems are likely to reach the next milestone zettascale computing (10^{21} operations/s) by 2035.

3. We review challenges and techniques of moving from exascale computing to zettascale computing.

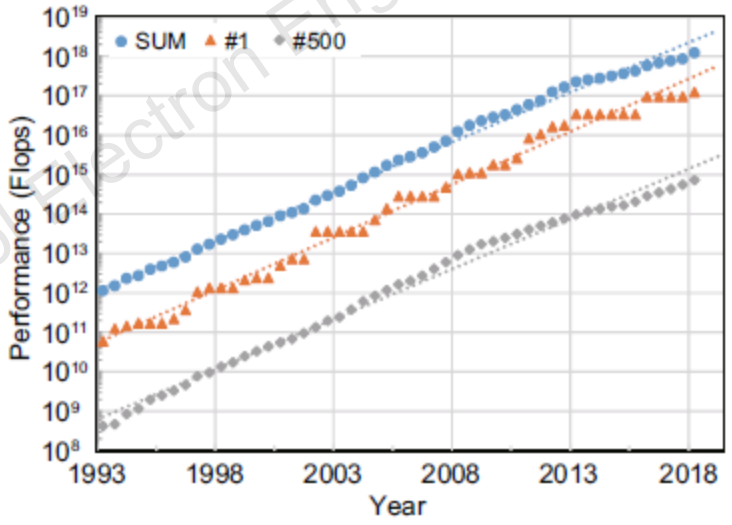

Fig. 1 Projected performance of the Top500 high-performance computing systems from 1993 to 2020. The plot is based on a data source publicly available on the Top500 website (https://www.top500.org/)

Zettascale Challenges

CH1: The traditional Moore's roadmap based on feature size reduction may end within three to four generations.

CH2: A sustainable future HPC system is expected to keep its power consumption budget within 20–30 MW.

CH3: We have to maintain the node bandwidth-performance ratio to be smaller than 0.04 for large-scale systems.

CH4: The huge performance gap between processors and memories has been greatly restricting the sustainable development of HPC.

CH5: The growing scale and complexity of future HPC systems makes reliability be a major concern.

CH6: Future programming models have to manage a large amount of parallelism in terms of programming, debugging, and tuning.

Zettascale Techniques (1/4)

1. Architecture:

(1) Dedicated processors will become the most attractive high-performance processing units for their high computing efficiency;

(2) The processor architectures will incorporate new advanced technologies such as 2.5D/3D integration to relieve the memory wall;

(3) The integration of inter-chip interconnection networks can reduce the communication latency and improve communication bandwidth;

(4) New computing devices based on the non-von Neumann structure will become an important branch of future architectures.

Zettascale Techniques (2/4)

2. High-performance interconnecting technology:

(1) A balanced, scalable, and low-cost HPC interconnect system will be designed to achieve flexible allocation of network resources based on silicon photonics interconnects;

(2) High-density opto-electronic integration enables deep integration of interconnects and computational memory, thus reducing latency and improving density power consumption.

(3) An all-optical interconnecting system will be realized, with an aim to increase the scalability of HPC systems and save interconnection costs.

Zettascale Techniques (3/4)

3. Emerging storage technology:

(1) Future HPC storage systems can have a much richer hierarchy:

Integrating storage and interconnection can significantly improve the efficiency of future storage systems;

The development of memristor and quantum computing enables a real integration of computing and storage.

4. New manufacturing process

(1) High-performance computers will continue to evolve towards higher density, in terms of computing, assembly, and energy consumption;

(2) The cooling granularity will be extended from the cabinet/board level to the chip level, thus achieving the accurate targeted heat dissipation;

(3) The optical technology will be widely used in high-speed signal design.

Zettascale Techniques (4/4)

5. Programming models and environments:

(1) Mining parallelism is still the most important way of enabling a sustainable high-performance computing system;

(2) MPI+X will still be the mainstream programming paradigm for future HPC systems and must be extended to realize performance portability;

(3) New programming paradigms will emerge and develop with new computing devices, e.g., quantum devices;

(4) Domain-specific frameworks and library tools will be the main ways to improve productivity;

(5) Techniques, such as machine learning, will be used to auto-tune various workloads during runtime.

Zettascale roadmap

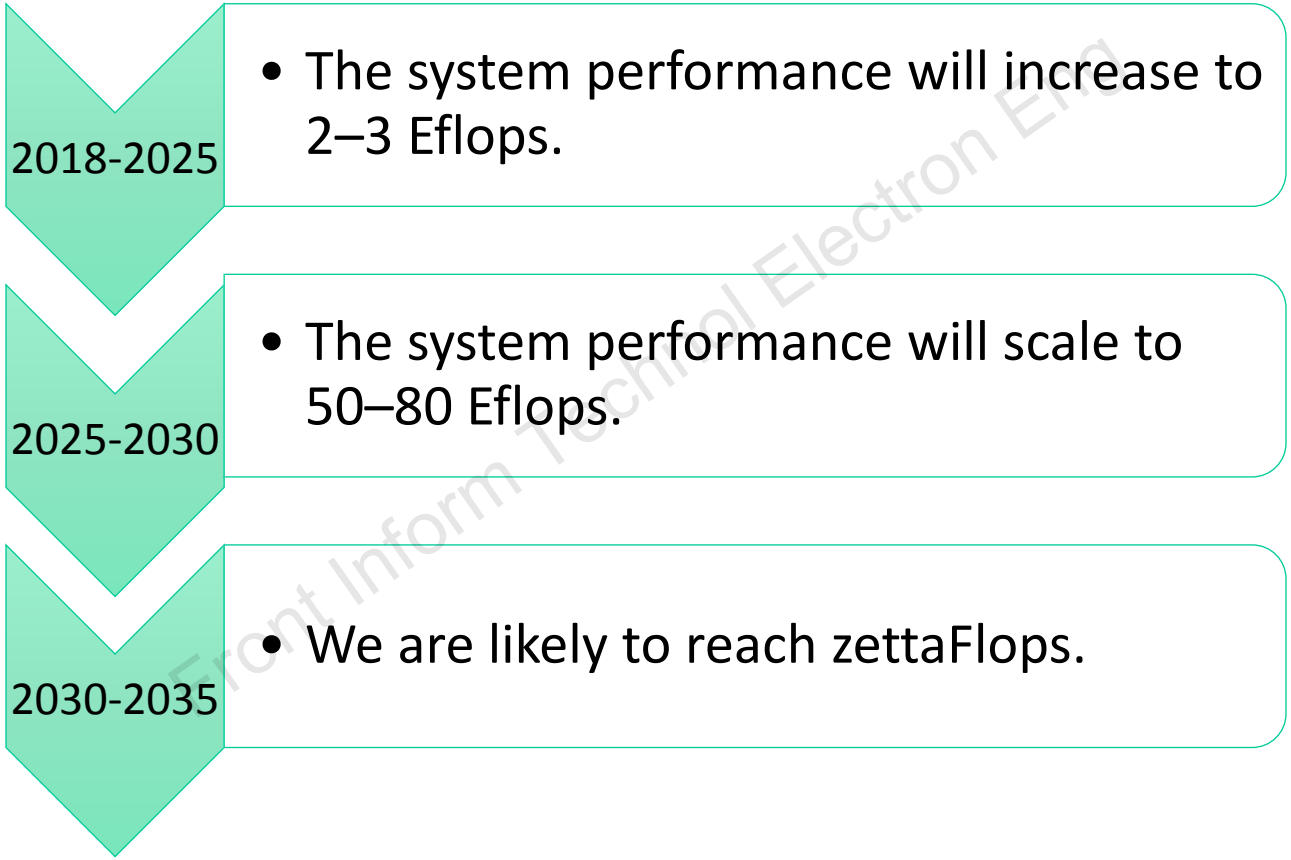

Zettascale System Metrics

Table 1 Zettascale metrics

Metric	Value
Peak performance	1 Zflops
Power consumption	100 MW
Power efficiency	10 Tflops/W
Peak performance per node	10 Pflops/node
Bandwidth between nodes	1.6 Tb/s
I/O bandwidth	10–100 PB/s
Storage capacity	1 ZB
Floor space	1000 m^2

Suggestions

1. We should develop collaborative design for heterogeneous architectures and develop customization technology for specific applications.

2. We should investigate new optoelectronic integration, key component design, and integrated network architecture in advance.

3. We should investigate new non-volatile storage technologies and explore the integration of computing and storage .

4. We should study high-density engineering technologies, low-energy-consumption targeted cooling technologies, etc.

5. We should develop transparent programming and software framework support to hide the changes brought by hardware changes .

6. We should develop a domestic processor architecture, build a sound, and robust high performance software environment.