Available online www.jsaer.com

***Journal of Scientific and Engineering Research*, 2021, 8(1):103-105**

Review Article

ISSN: 2394-2630
CODEN(USA): JSERBR

Exascale Computing: A Primer

Matthew N. O. Sadiku[1], S. R. Nelatury[2], Emad Awada[3], Sarhan M. Musa[1]

[1]Roy G. Perry College of Engineering, Prairie View A&M University, Prairie View, TX 77446
[2]School of Engineering, Pennsylvania State University, Erie, PA 16563-1701
[3]Electrical & Computer Engineering Department, College of Engineering, Applied Science University, Amman-Jordan 11931
Email: sadiku@ieee.org; srn3@psu.edu; emadawada@yahoo.com;
smmusa @pvamu.edu

Abstract Exascale computing is the term given to the next 50- to 100-fold increase in speed over the fastest supercomputers in use today. This super powerful machine is poised to transform modeling and simulation in science and engineering. It is hoped that the exascale machines will solve some or all of the major problems that are facing us today. This paper provides a brief introduction to exasclae computing.

Keywords exascale computing, supercomputers

Introduction

Computing technology greatly affects nearly every aspect of our life including education, entertainment, transportation, communication, economy, medicine, engineering, and science. The fastest supercomputers in the world today solve problems at the petascale, i.e. a quadrillion calculations each second. The next milestone in computing performance is the exascale. Exascale computing refers to supercomputers that are capable of at least one a quintillion calculations per second or 1,000,000,000,000,000,000 operations per second. It is considered to be a significant achievement in computer engineering.

Exascale computers represent powerful and scientific instruments and that amount of power would be phenomenal for a single computer machine. In addition to high speed, supercomputers need to have the ability to store and read vast quantities of data at high speeds. Exascale computer will require disruptive changes to hardware and software.

Various initiatives have been taken by the governments (US, EU, China, Japan, Taiwan, India, etc.) and industries (IBM, Intel, etc.) to build an exascale computer. Both the United States and China are competing to become the first nation to create exascale computer. It is estimated that each of these exotic computer machines will cost anywhere between $400 million and $600 million. The first US exascale system has been named Aurora 2021 (A21).

Exascale Computıng Basıcs

The speed of computers has traditionally been measured by the maximum number of floating point operations that can be performed in a second (FLOPS). An exascale system is a supercomputer that can solve problems 50X faster than on the 20PF

systems (Titan, Sequoia) of today or 100X faster than on Mira. High-performance computing is currently moving from the petaflops scale (10^15 FLOPS) towards exascale (10^18 FLOPS) computing. Exascale computing systems are 1,000 times faster than existing petaflop machines.

The main architectural elements of the exascale supercomputers include the following [1].

- *Processor*: Exascale systems are built using hybrid or heterogeneous platforms. The hybrid systems use both central processing units (CPU) and graphics processing units (GPU) to efficiently leverage the performances.
- *Memory*: Meeting the performance requirement of exascale computing requires increase in the memory bandwidth. This in turns increases the power consumption. The DRAM capacity of a system is basically limited by cost.
- *Algorithms*: The advances in the architecture of supercomputers raise the need for changing the programming systems approach. Exascale supercomputing systems need to run more efficiently in terms of scalability, reliability, and data movement.

Applications

Exascale computing systems that enable classical simulation and modeling applications to tackle complex problems that are currently out of reach. Such supercomputing systems will meet the computing needs of cutting-edge engineering work and scientific discovery. They will enable new types of applications such as machine learning, deep learning, and big data analytics. They will also have many applications in the research laboratories and industries.

Exascale computing will push the frontiers in a transformative fashion. It is expected to be applied in various computation-intensive research areas such as engineering, biology, materials science, cosmology, precision medicine for cancer, astrophysics, energy, climate science, renewable energy, biology, socioeconomic modeling, molecular modeling, astrophysical recreation, and national security.

Benefits

Exascale computers will simulate the processes involved in precision medicine, regional climate, additive manufacturing, nuclear physics, national security, and relationships behind many of the fundamental forces of the universe. They will potentially benefit society in a several ways. In addition to providing solutions to advances in healthcare, biology, and storm prediction, other benefits include [2]:

- ***Reducing Pollution:*** Exascale computing can reduce pollution caused by burning fossil fuels. It will be possible to increase the efficiency of combustion systems in engines.
- ***Advances in Healthcare:*** Exascale computing will accelerate cancer research by helping scientists understand the molecular basis of key protein interactions. It will also enable doctors to predict the right treatments for the patient by modeling drug responses.
- ***Predicting Severe Weather:*** Weather prediction models will be able to predict more accurately and quickly the timing and path of severe weather events such as hurricanes.
- ***Improving Quality of Life:*** The use of exascale computing in urban science promises to mitigate health hazards, reduce crime, and improve the quality of life in cities by optimizing infrastructure (such as transportation, energy, housing, healthcare) access and usage choices.
- ***Advances in Materials Science*:** Creating new technologies requires that we discover new materials with specific properties. Exascale machines can help researchers and companies to identify new useful materials very fast. They will help design, control, and manufacture advanced materials.

Extascale-scale computing will enable the solution of vastly more accurate predictive models and the analysis of massive quantities of data. It will have a broad and positive impact on U.S. industrial competitiveness.

Challenges

Going to the exascale is a challenging venture. These severe challenges arise both in the hardware realm and in the software [3,4]:

- *Hardware Challenges*: The architectural challenges for reaching exascale are dominated by power, memory, interconnection networks, and resilience. Exascale machines will require radical changes in hardware and in programming applications to effectively use tens of millions of cores. Hardware challenges include increased parallelism, reliability, energy consumption, and memory, network, and storage efficiencies.

- *Power Challenge:* To achieve an exascale system using current technology, the annual power cost to operate the system would be above $2.5 B per year. Achieving this new power goal for exascale systems creates serious research challenges. The cost of power can limit the scale of computers that can be deployed.
- *Memory Challenge:* The ratio of memory to processor is critical in determining the size of the problem that can be solved. External storage is far larger, but also operates at extremely slow rates relative to processor speeds. There are indications that memory will become the rate-limiting factor along the path to exascale.
- *Resiliency Challenge:* Resiliency will be one of the toughest challenges in exascale systems. This deals with the ability of a system to continue operation in the presence of faults. It is anticipated that exascale systems will experience various kind of faults daily.
- *Programming Challenge:* This requires inventing new programming environments that express massive parallelism, data locality, and resilience. Exascale systems will present programmers with tough challenges.

These challenges must be addressed by researcher in order to deliver capable exascale computing.

Conclusıon

Exascale computing is characterized by an infrastructure to support computational capability in the order of an ExaFLOP. Researchers on exascale supercomputers are already looking ahead to the next big 1,000x performance: zettascale computing, which a system capable of 10^21 double-precision 64-bit floating-point operations per second. They have the great opportunities to influence the direction of future architectures so that they meet DOE mission needs. The latest news and information on exascale computing can be found on the website of the Exascale Computing Project (ECP): https://www.exascaleproject.org/

References

[1]. R. Alshahrani, "The path to exascale computing," *Int'l Conf. Par. and Dist. Proc. Tech. and Appl.*, 2015, pp.123-126.

[2]. P. Messina, "Why we need exascale computing," https://www.huffingtonpost.com/entry/why-we-need-exascale-computing_us_58c94f59e4b0009b23bd94c4

[3]. U.S. Department of Energy, "The opportunities and challenges of exascale computing," Fall 2010.

[4]. D. A. Reed and J. Dongarra, "Exascale computing and big data," *Communications of the ACM*, vol. 58, no. 7, July 2015, pp. 56-68.

About Authors

Matthew N.O. Sadiku is a professor in the Department of Electrical and Computer Engineering at Prairie View A&M University, Prairie View, Texas. He is the author of several books and papers. His areas of research interests include computational electromagnetics and computer networks. He is a fellow of IEEE.

Sudarshan R. Nelatury is an associate professor at Penn State University, The Behrend College, Erie, Pennsylvania. His teaching and research interests lie in electromagnetics and signal processing.

Emad Awada is an associate professor in the Department of Electrical and Computer Engineering at Applied Science University, Jordan. He received his B.S., M.S., and Ph.D. in Electrical Engineering from Prairie View A&M University, Prairie View, Texas. He is the author of several papers. His areas of research interests include mixed signals analysis, signals computational, and electrical machine faults detection.

Sarhan M. Musa is a professor in the Department of Engineering Technology at Prairie View A&M University, Texas. He has been the director of Prairie View Networking Academy, Texas, since 2004. He is an LTD Sprint and Boeing Welliver Fellow.