

Extreme Computing: A Primer

Matthew N. O. Sadiku¹, Adedamola A. Omotoso², Sarhan M. Musa¹

¹Professor, ²Student

^{1,2}Roy G. Perry College of Engineering, Prairie View A&M University, Prairie View

How to cite this paper: Matthew N. O. Sadiku | Adedamola A. Omotoso | Sarhan M. Musa "Extreme Computing: A Primer" Published in International Journal of Trend in Scientific Research and Development (ijtsrd), ISSN: 2456-6470, Volume-3 | Issue-3, April 2019, pp.294-295, URL: <http://www.ijtsrd.com/papers/ijtsrd21723.pdf>



IJTSRD21723

ABSTRACT

The ever-increasing demand of computing power has led to the development of extremely large systems that consist of millions of components. Sustainable large scale computing systems can extend themselves to extreme scales. Both extreme and exascale computing defy the common wisdom of HPC and are regarded as unorthodox, but they could turn out to be indispensable necessities in the near future [1]. This paper provides a primer on extreme computing.

Keywords: extreme computing, extreme-scale computing, data-intensive computing

Copyright © 2019 by author(s) and International Journal of Trend in Scientific Research and Development Journal. This is an Open Access article distributed under the terms of the Creative Commons Attribution License (CC BY 4.0) (<http://creativecommons.org/licenses/by/4.0>)



INTRODUCTION

The history of computing is one of punctuated equilibrium, bringing new and unexpected changes. Mainframes gave birth to minicomputers, which gave birth to workstations, which gave birth to personal computers, which gave birth to smartphones [2]. Modern scientific computing is approaching the point where novel computational algorithms will become an indispensable tool. Existing supercomputers are not well suited for data-intensive computations [3]

Astronomy, weather forecasting, monitoring, and computer modeling/simulations are on the verge of generating petabytes of data. The same holds for large experimental or observational data. The ever-increasing demand of computing power has led to the development of extremely large systems that consist of millions of components. In such systems, cost, energy availability and consumption, and reliability are increasingly critical issues, which pose limits on the practicality of their deployment. Since the systems cost tens to hundreds of millions of dollars, making the best use of these valuable and scarce systems is important.

Extreme computing, also known as extreme-scale computing, involves massive data.

Exascale/extreme computers should have highly hierarchical architectures with nodes composed of multiple core processors (CPU) and accelerators (GPU). They promise

orders of-magnitude improvement in performance over current high-end computers.

APPLICATIONS

Developing extreme computing system requires a computing platform that supports billion-way parallelism. To develop extreme computers will include collaboration from industry, academia, and government, with expertise in algorithms, computer system architecture, operating systems, compilers, and applications.

An example of extreme computing system is the Blue Waters system at the University of Illinois. Blue Waters (www.ncsa.illinois.edu/BlueWaters), provides sustained petaflops performance on a wide range of applications. It is being sponsored IBM, NSF, and the University of Illinois [4].

Another example is the Dynamical Exascale Entry Platform (DEEP) project, which is a supercomputer one thousand times more powerful than existing machines. It is a collaborative project which started in October 2013 and coordinated by the Jülich Supercomputing Centre in Germany. It involves more than 15 European partners [5].

The Big Data Extreme Computing (BDEC) is a common initiative organized and funded by Europe, US, and Japan. The BDEC community has staged a series of workshops that will seek ways in which the major issues associated with extreme computing and big data intersect [6].

Software for extreme computing includes ADLB (Asynchronous Dynamic Load Balancing) library. ADLB hides communication and memory management from the application, providing a simple programming interface. Key to the functioning of ADLB is that when a process requests work, ADLB can find it in the system either locally or remotely on another server [7].

Extreme computing is permeating every modern application, from scientific R&D and business operations to medical practice and national security. All these communities depend in computer systems to process huge amount of data quickly and efficiently [8]. They are poised to utilize extreme computers. Without doubt, many applications will have to undergo some major changes, while others will be replaced entirely.

CHALLENGES

Moore's law is the most often quoted dictum by which the silicon revolution. The law essentially observes exponential growth of transistors in every new generation of integrated circuits. Developing applications for a transition from petascale to exascale systems requires a huge investment in R&D. The challenges facing the algorithm developer for extreme computing systems include targeting massive parallelism, reaching millions of cores, dealing with hybrid architectures, and managing complex memory hierarchies [1]. Other challenges include power consumption and associated carbon emissions. Programming extreme computers generates new challenges and algorithms issues. Find solutions to these challenges may lead to finding solutions to important societal challenges in the process.

CONCLUSION

Rapid growth in the usage and demands for huge data processing is leading to extreme computing. Energy efficiency is a key problem for extreme computing,

Traditional approaches do not work well in this environment. The United States, the European Union, and Japan have aggressively taken initiatives toward achieving exascale computing. This will provide the next generation of HPC software to support big data and extreme computing to aid scientific discovery.

REFERENCES

- [1] J. Dongarra et al., "With extreme computing, the rules have changed," *Computing in Science & Engineering*, May/June 2017, pp. 52-62.
- [2] D. Reed, "Extreme computing," December 04, 2017. https://www.hpccdan.org/reeds_ruminations/extreme-computing/
- [3] A. S. Szalay, "Extreme data-intensive scientific computing," *Computing in Science & Engineering*, November/December 2011, pp. 34-41.
- [4] J. Torrellas, "Architectures for extremescale computing," *Computer*, November 2009, pp. 28-35.
- [5] E. Suarez, "Extreme computing," International Innovation, p. 98 http://www.deep-project.eu/SharedDocs/Downloads/DEEP-PROJECT/EN/Publications/international-innovation-interview-07-2014.pdf?__blob=publicationFile
- [6] "The basis for BDEC," <https://www.exascale.org/bdec/>
- [7] E. L. Lusk, S. C. Pieper, and R. M. Butler, "More scalability, less pain: A simple programming model and its implementation for extreme computing," *SciDAC Review*, vol. 17, 2010
- [8] S. Rumley, "Optical interconnects for extreme scale computing systems," *Parallel Computing*, vol. 64, 2017, pp. 65-80.