

Chapter 52

The Future of High-Performance Computing (HPC)

Herbert Cornelius
Intel Corporation EMEA, Germany

ABSTRACT

For decades, HPC has established itself as an essential tool for discoveries, innovations, and new insights in science, research and development, engineering, and business across a wide range of application areas in academia and industry. Today high-performance computing is also well recognized to be of strategic and economic value – HPC matters and is transforming industries. This chapter will discuss new emerging technologies that are being developed for all areas of HPC: compute/processing, memory and storage, interconnect fabric, I/O and software to address the ongoing challenges in HPC such as balanced architecture, energy efficient high-performance, density, reliability, sustainability, and last but not least, ease of use. Of specific interest are the challenges and opportunities for the next frontier in HPC envisioned around the 2020 timeframe: ExaFlops computing. The authors also outline the new and emerging area of high-performance data analytics, big data analytics using HPC, and discuss the emerging new delivery mechanism for HPC – HPC in the Cloud.

INTRODUCTION

High-Performance Computing (HPC) is used to address and solve the world's most complex computational problems. For decades, HPC has established itself as an essential tool for discoveries, innovations and new insights in science, research and development, engineering and business across a wide range of application areas in academia and industry. It has become an integral part of the scientific method – the third leg along with theory and experiment.

Today, High-Performance Computing is also well recognized to be of strategic and economic value – HPC matters and is transforming industries (Osseyran & Giles, 2015).

High Performance Computing enables scientists and engineers to solve complex and large science, engineering, and business problems using advanced algorithms and applications that require very high

The Future of High-Performance Computing (HPC)

compute capabilities, fast memory and storage, high bandwidth and low latency throughput, high fidelity visualization, and enhanced networking.

Today, the IT industry is being transformed by cloud, big data, social media, artificial intelligence, and “Internet of Things” technologies and business models. All of these trends require advanced computational simulation models and powerful highly scalable systems. Hence, sophisticated HPC capabilities are critical to the organizations and companies that want to establish and enhance leadership positions in their respective areas.

Some industry verticals and application areas where HPC is used are as follows:

- Manufacturing, Computer Aided Engineering (CAE)
- Automotive Industry
- Aerospace Industry
- Weather Forecast and Climate Research
- Energy, Oil & Gas Industry, Geophysics
- Life-Science and Bio-Informatics (Genomics)
- Government Research Laboratories
- Universities (Academics), Machine Intelligence, Machine/Deep Learning, Artificial Intelligence (AI)
- Astrophysics, High-Energy Physics, Computational Chemistry, Material Science
- Financial Services Industry(FSI)
- Digital Content Creation (DCC)
- Defense
- Security and Intelligence

BACKGROUND

After its initial years of proprietary computer systems in the 1970/1980’s, HPC has evolved with industry standards that democratized Supercomputing, making advanced computing available to more users and wider application segments.

Today’s modern HPC solutions are utilizing high-performance server compute nodes connected with high performance fabrics connected to high-performance storage systems, mainly deployed on distributed cluster architectures running on Linux based operating systems with up to tens of thousands of processors. For specific workloads and applications with the need for large coherent shared memory capacity (terabytes of data: TB), more specialized solutions and systems are used based on cc:NUMA (Cache-Coherent Non-Uniform Memory Access) architectures. For example, the SGI UV system supports up to 256 CPU sockets and up to 64TB of cache-coherent shared memory in a single system.

While in the past chip designs used to be limited by space and the number of transistors available, now power consumption is becoming the main constraint for High-Performance Computing. With several new emerging technologies there will be multiple opportunities to address some of the ongoing challenges in HPC such as balanced architectures, energy efficiency, density, reliability, resiliency, sustainability, and last but not least ease-of-use.

As stated in a recent article (Bader, 2016), the trend in high- performance computing (HPC) to adopt more accelerators such as graphics processing units (GPUs), field-programmable gate arrays (FPGAs),

15 more pages are available in the full version of this document, which may be purchased using the "Add to Cart" button on the publisher's webpage:

www.igi-global.com/chapter/the-future-of-high-performance-computing-hpc/214655

Related Content

On Uplink Channel Estimation in WiMAX Systems

Yushi Shen, Pamela C. Cosman, Laurence B. Milstein and Eduardo F. Martinez (2010). *International Journal of Mobile Computing and Multimedia Communications* (pp. 67-77).

www.irma-international.org/article/uplink-channel-estimation-wimax-systems/43894

Securing EPR Data Using Cryptography and Image Watermarking

Youssef Zaz, Lhoussain El Fadil and Mohamed El Kayyali (2012). *International Journal of Mobile Computing and Multimedia Communications* (pp. 76-87).

www.irma-international.org/article/securing-epr-data-using-cryptography/66368

Did You See That?

Murray Crease and Joanna Lumsden (2008). *Handbook of Research on User Interface Design and Evaluation for Mobile Technology* (pp. 972-981).

www.irma-international.org/chapter/did-you-see/21876

Information-Centric Networking

Mohamed Fazil Mohamed Firdhous (2019). *Advanced Methodologies and Technologies in Network Architecture, Mobile Computing, and Data Analytics* (pp. 1237-1247).

www.irma-international.org/chapter/information-centric-networking/214696

Throughput-Delay Evaluation of a Hybrid-MAC Protocol for M2M Communications

Pawan Kumar Verma, Rajesh Verma, Arun Prakash and Rajeev Tripathi (2016). *International Journal of Mobile Computing and Multimedia Communications* (pp. 41-60).

www.irma-international.org/article/throughput-delay-evaluation-of-a-hybrid-mac-protocol-for-m2m-communications/148261