Technology Spotlight

The Exascale Era Has Arrived

Sponsored by NVIDIA

Steve Conway, Earl Joseph, Bob Sorensen, and Alex Norton
November 2018

EXECUTIVE SUMMARY

Earlier this year, scientists broke the exascale barrier by running multiple real-world codes for the first time at sustained speeds in excess of a billion billion (10^{18}) mathematical operations per second on the Summit supercomputer at Oak Ridge National Laboratory. Historically, the HPC community has associated performance legitimacy with 64-bit (double) precision. The Gordon Bell Prize-finalist runs on the Summit system demonstrate, however, that for a growing contingent of important codes, single- or half-precision rightfully complement double-precision. There is little reason not to recognize lower-precision workloads run at exaflop/exaop speed as exascale achievements in cases where double precision would add little or no additional insight. Hyperion Research agrees with the organizations that participated in the Summit runs–Oak Ridge, Lawrence Berkeley National Lab and NVIDIA–that these achievements mean the exascale era has begun.

This paper reviews the Summit runs, discusses other important mixed-precision applications that promise to benefit from sustained exascale performance, and then looks at the extensive effort NVIDIA and its partners made to develop the hardware-software ecosystem needed to enable real-world exascale performance.

SITUATION

Breaking the Exascale Barrier

In June 2018, Oak Ridge National Laboratory announced that researchers there had broken the exascale barrier for the first time ever, by running a genomics code (CoMET) on the IBM-NVIDIA Summit supercomputer at a top speed of 2.36 billion billion single and half-precision operations per second. In October 2018, Lawrence Berkeley National Laboratory, Oak Ridge National Laboratory, and NVIDIA teamed to apply mixed-precision deep learning to a climate science application that achieved peak performance of 1.15EF on the Summit system (Figures 1 and 2). Both achievements qualified as finalists for the prestigious ACM Gordon Bell Prize, which is given annually "for peak performance or special achievements in scalability and time-to-solution on important science and engineering problems."

FIGURE 1

Deep Learning on Climate Datasets

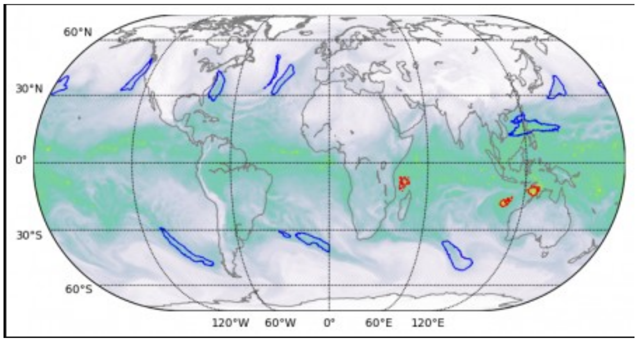

Network	Operation Count (TF/sample)	GPU Model (System)	Precision	Training Rate (samples/s)	Performance (TF/s)	% Peak
DeepLabv3+	14.41	V100 (Summit)	FP16	2.67	38.45	31
			FP32	0.87	12.53	80
Tiramisu	4.188	V100 (Summit)	FP16	5.00	20.93	17
			FP32	1.91	8.00	51
	3.703*	P100 (Piz Daint)	FP32	1.20	4.44	48

Fig. 2: Single GPU performance results from training the Tiramisu and DeepLabv3+ networks. Results are shown for all tested systems using FP32 and FP16 precision where relevant. Note that the operation count for Tiramisu on Piz Daint (marked with an asterisk) is computed from a modified network using 4 out of the 16 available input data channels.

Source: Hyperion Research, 2018

These pioneering achievements signaled that the exascale computing era had begun, even before the arrival of supercomputers designed to support exascale performance on 64-bit, real-world codes. Some pre-exascale systems can support exascale performance on lower- and mixed-precision workloads. For example, there are many applications in the molecular dynamics, physics, and geoscience domains that leverage single-precision math libraries. Some of the key single-precision HPC applications include: AMBER, NAMD, LAMMPS, GROMACS, CHROMA, MILC, and RTM. Hyperion Research agrees with the organizations that enabled the Summit runs—Oak Ridge, Lawrence Berkeley National Lab and NVIDIA—that lower-precision performance qualifies as exascale, especially in cases where double precision would add little or no additional insight.

Important Applications Poised for Exascale

Government agencies, especially in the U.S., Europe, China and Japan, have assembled lists of applications that represent "low hanging fruit" for exascale performance—primarily because of their high potential to benefit from being run at very large scale. Many of these applications are already being run today on one of the world's most powerful supercomputers, such as the Titan supercomputer at Oak Ridge National Laboratory, Piz Daint at the Swiss National Supercomputer Center (CSCS), and Blue Waters at the National Center for Supercomputing Applications (NCSA).

Following are just a few of the many economically and socially important applications poised for breakthroughs made possible by exascale performance. This list illustrates the broad scope of scientific and engineering disciplines that will be advanced in the exascale era:

- **Lithium-ion battery polymer discovery (3M node hours)** Lithium (Li)-ion batteries have strong performance but have safety issues. Goal: safe solid-state batteries with performance comparable to liquid batteries.
- **Mapping Earth's core (10M node hours)** Scientists have been searching for a way to map the entire Earth's interior to understand magma flows and tectonic boundaries. Goal: better understand earthquakes and volcanic activity.
- **Weather simulation (18M node hours)** Unmitigated global warming promises to increase floods, heat waves and droughts. Today's climate models have wide variability, making mitigation planning difficult. Goal: higher-resolution climate modeling to improve predictions of the impact of global warming.
- **HIV Capsid Structure Simulation (52M node hours)** Infection with HIV-1 is classified as a global pandemic by the WHO. Due to the high mutation rate of the virus, new drug treatments must be constantly developed. Goal: Understand the HIV capsid system structure and how it changes and responds to its environment.

NVIDIA and the Exascale Ecosystem

The Summit supercomputer that ran the first applications at exascale was developed by NVIDIA and IBM under the U.S. Department of Energy's CORAL program. The bulk of the processing power for the supercomputer is provided by NVIDIA Volta Tensor Core GPUs. Tensor cores excel at accelerating some important computing tasks, including matrix calls, that underpin many big data and AI tasks. For the 2.36EF breakthrough performance at Oak Ridge National Laboratory, the tensor cores in the NVIDIA GPUs alone provided a 4.5-fold speedup.

The other Gordon Bell finalist project that broke the exascale barrier explicitly called out the importance of ecosystem integration in their paper titled, "Exascale Deep Learning for Climate Analytics" (*Exascale Deep Learning for Climate Analysis. T. Kurth, et al. arxiv.org/pdf/1810.01993.pdf*).

In this paper it states:

"...while Deep Learning might appear to be a natural fit for existing petascale and future exascale HPC systems, careful consideration must be given towards balancing various subsystems (CPUs, GPUs/accelerators, memory, storage, I/O and network) to obtain high scaling efficiencies. Sustained investments are required in the software ecosystem to seamlessly utilize algorithmic innovations in this exciting, dynamic area."

NVIDIA has long recognized the HPC truism that leadership-class performance depends on integrating hardware and software to create an ecosystem that turns an HPC system into a highly coherent, productive resource for users. The HPC ecosystem includes not only hardware—processors and accelerators, networking and storage—but also a full software stack that extends from the operating system to the application kernels—libraries, compilers, tools, commands and more. All these elements need to be integrated within a system-level architecture that's designed to exploit their capabilities to the fullest extent.

Figure 3 shows the integrated NVIDIA hardware-software ecosystem, starting with the foundational hardware layer of accelerators and processors for use in systems located on premise or in cloud environments. The foundational Layer 1 is designed to be heterogeneous and flexible enough to deliver efficient computing power to a broad spectrum of HPC and AI customer use cases, such as those illustrated in the top row of Figure 3. In between the foundational layer and the use cases, and needing integration with both, are important layers of software.

- Layer 2 includes the NVIDIA software development kit (SDK) and an array of NVIDIA-developed libraries, tools and commands.
- Layer 3 consists of third-party (industry) frameworks and applications, including many from independent software vendors as well as open source and community software.

NVIDIA reports that it has spent the past 12 years optimizing this end-to-end ecosystem for over 580 applications in a broad range of domains, including scientific simulation applications, deep learning frameworks and, more recently, traditional machine learning workflows.

FIGURE 3

NVIDIA Universal Acceleration Platform for HPC and AI

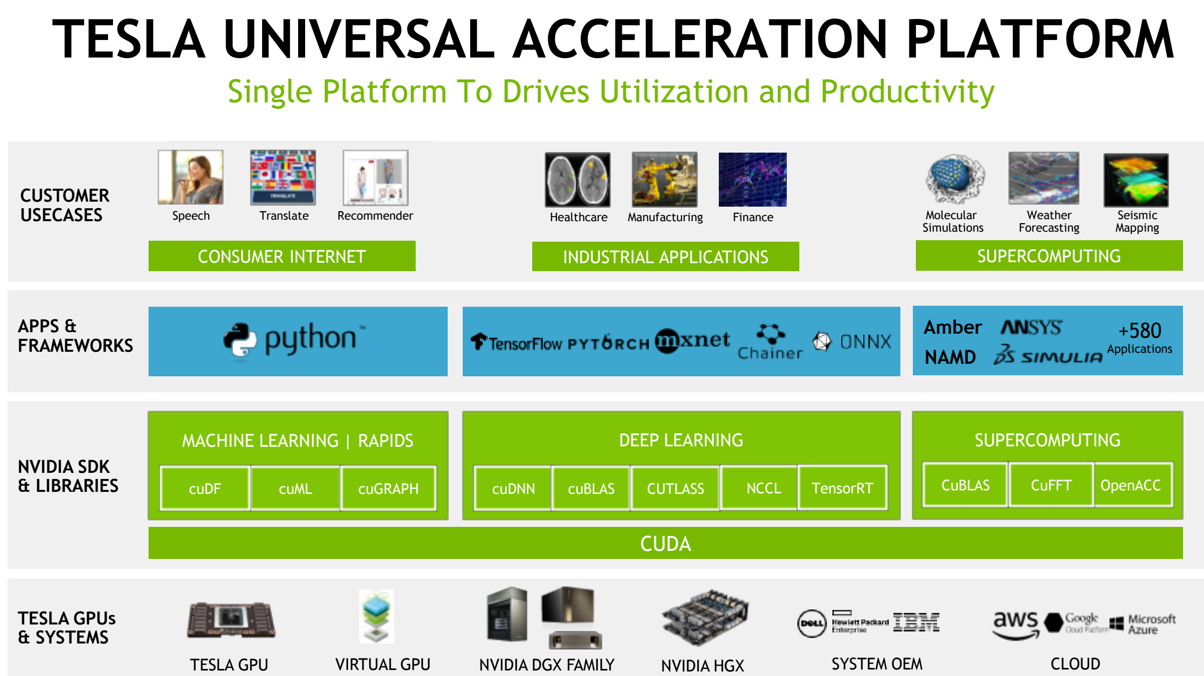

Source NVIDIA, 2018

FUTURE OUTLOOK

The real-world exascale performance achievements on the Summit supercomputer were strongly enabled by NVIDIA Volta Tensor Core GPUs and became finalists for the prestigious Gordon Bell Prize. These achievements confirm that the exascale computing era has arrived and that it will benefit a broad spectrum of important applications that may run optimally at half-precision, single precision, double precision or in mixed precision. Lists of applications representing "low hanging fruit" for early exascale performance have been drawn up by governments in the U.S., Europe, China, Japan and elsewhere.

NVIDIA today is the market-leading vendor of accelerators designed to handle data-intensive HPC and AI workloads. In collaboration with a large array of partner organizations, NVIDIA has developed a comprehensive software ecosystem that is designed to deliver the power of NVIDIA GPUs and other foundational hardware with high efficiency to HPC and AI end-user applications. Hyperion Research forecasts that the worldwide HPC market, including HPC-enabled AI and other advanced analytics applications, will grow robustly from $24 billion in 2017 to $38 billion in 2022. We believe NVIDIA is positioned well to continue benefiting strongly from this projected growth.

About Hyperion Research, LLC

Hyperion Research provides data-driven research, analysis and recommendations for technologies, applications, and markets in high performance computing and emerging technology areas to help organizations worldwide make effective decisions and seize growth opportunities. Research includes market sizing and forecasting, share tracking, segmentation, technology and related trend analysis, and both user & vendor analysis for multi-user technical server technology used for HPC and HPDA (high performance data analysis). We provide thought leadership and practical guidance for users, vendors and other members of the HPC community by focusing on key market and technology trends across government, industry, commerce, and academia.

Headquarters

365 Summit Avenue
St. Paul, MN 55102
USA
612.812.5798
www.HyperionResearch.com and www.hpcuserforum.com
