HPC and the Future of Seismic

I briefly profile the world's largest commercial computer installations, and consider their relevance for the future of high-end seismic imaging and AI pursuits by the oil and gas industry.

Introduction

When I began my career in seismic geophysics over 20 years ago we were regularly told that seismic processing and imaging was one of the largest users of computing, and indeed, it took developments in computing capability to mature much further before applications such as Full Waveform Inversion (FWI) and Least-Squares Migration (LSM) became commercially commonplace in the last 5-10 years. Although the oil and gas industry still has several representatives in the top-100 ranked supercomputers (discussed below), applications now more commonly include modeling and simulations essential for nuclear security, weather and climate forecasting, computational fluid dynamics, financial forecasting and risk management, drug discovery and biochemical modeling, clean energy development, exploring the fundamental laws of the universe, and so on.

The dramatic growth in High Performance Computing (HPC) services in the cloud (HPCaaS: HPC-as-a-Service) in recent years has in principle made HPC accessible 'on demand' to anyone with an appropriate budget, and key selling points are generally the scalable and flexible capabilities, combined with a variety of vendor-specific Software-as-a-Service (SaaS), Platform-as-a-Service (PaaS), and Infrastucture-as-a-Service (IaaS) offerings. This new vocabulary can be confronting, and due to the complexity of cloud-based solutions, I only focus here on the standalone supercomputer infrastructures. A future Industry Insights newsletter will share some experiences in transitioning to a balance of HPC on-premises and in the cloud from a seismic service company perspective.

Artificial Intelligence (AI) solutions are becoming abundant in every walk of life, and the oil and gas industry is embracing machine learning and Big Data applications in the relentless pursuit of lower costs, higher efficiencies, and the realization of greater value from all tangible assets. It is clear today is that the next waves of seismic imaging, reservoir modeling/forecasting and AI growth will be dependent on whether computing infrastructures are powerful enough to support the growing size and complexity of the use cases. **Figure 1** plots the (roughly) exponential growth in global computing capacity over the past 30 years. This trend needs to continue, and quite vast facilities are necessary to enable the growth. From the cloud HPC perspective, chip architecture explicitly influences cost per node hour (1 node-hour = the computational resource amount of 1 hour computation using 1 node), so aspects of the TOP500 list discussed below are relevant.

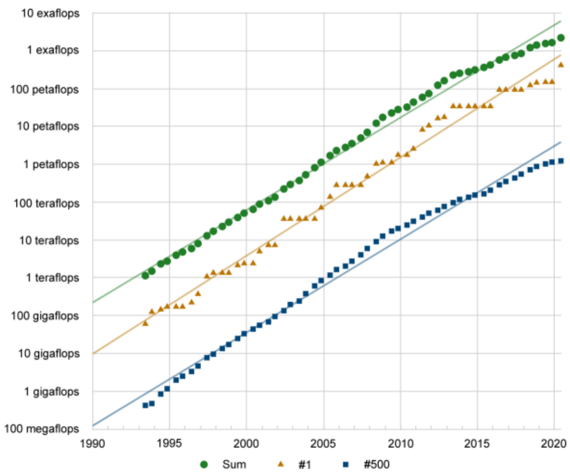

Figure 1. Actual vs. projected performance development of the world's largest supercomputers over the past 30 years (● = collective TOP500 performance, ▲ = the #1 supercomputer, and ■ = the average performance of the TOP500). Historical performance follows a generally exponential growth, although the improvements have slowed somewhat in the past few years.

Each unit on the vertical axis is defined as a number of floating point operations per second (mega = 10^6; giga = 10^9; tera = 10^{12}; peta = 10^{15}; and exa = 10^{18}).

.

The TOP500 Supercomputers: June 2020

The TOP500 list published at https://top500.org/ attempts to show the 500 most powerful commercially available computer systems, and is compiled by Erich Strohmaier and Horst Simon of Lawrence Berkeley National

Laboratory; Jack Dongarra of the University of Tennessee, Knoxville; and Martin Meuer of ISC Group, Germany. New entries wishing to be considered can be submitted at https://top500.org/project/call-for-participation/. The TOP500 list is updated every six months, with the June 2020 ranking at https://top500.org/lists/top500/2020/06/. Bragging rights are of course contentious, restricted to the installations acknowledged by the list, and continually being revised.

Representative of the rapidly changing order of supremacy, when PGS took delivery of the 4-petaflops (5.3 petaflops peak) Cray XC30 system named 'Abel' in 2015, it was the largest oil and gas industry computer on record. Abel was quickly overtaken by the 5.3 petaflops (6.7 peak) SGI 'Pangea' computer installed by Total in early-2016. As of June 2020, the oil and gas industry installations in the TOP500 list are as follows:

- Eni S.p.A. (#6, https://top500.org/system/179856/, 35 petaflops peak; as well as #19, https://top500.org/system/179444/; and #81, https://top500.org/system/178425/)
- Total Exploration Production (#15, https://top500.org/system/179689/, 17 petaflops peak; #49, https://top500.org/site/49546/)
- Petróbras S.A (#56, https://top500.org/system/179854/, 4 petaflops peak; and #82. https://top500.org/system/179681/)
- PGS (#62, https://top500.org/system/178556/, 4 petaflops peak). Note that PGS now has two XC-series Cray supercomputers in Houston, with a combined 8-petaflops capacity.

Ranking in the TOP500 list use the LINPACK benchmark because it is widely used and performance numbers are available for almost all relevant systems, and solves a dense system of linear equations. For the TOP500, the version of the benchmark used allows the user to scale the size of the problem and to optimize the software in order to achieve the best performance for a given machine. This performance does not reflect the overall performance of a given system, as no single number ever can. It does, however, reflect the performance of a dedicated system for solving a dense system of linear equations.

I avoid any discussion of chip design here, but it is observed in **Figure 2** that the evolution of supercomputer architecture has included many chip designs over the past decade; driven by both cost and the nature of the desired applications. GPUs have joined CPUs as general-purpose compute platforms, and now exceed CPU hardware in both memory bandwidth and floating point operations per second (FLOPs). There are many other factors to consider (available in the TOP500 listings); including the network interconnect, memory sharing, and so on. As might be expected, the TOP500 list is dominated by sites in China with 45.2% share, followed by the US with 22.6% share.

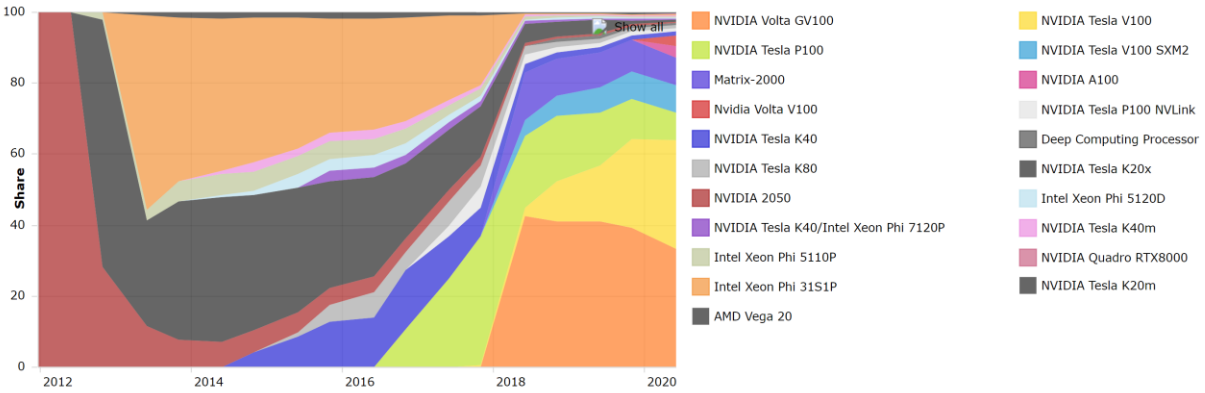

Figure 2. 2012-2020 performance share evolution of the TOP500 by Accelerator / Co-Processor design for the 21 most dominant architectures in 2020. More than 50 Accelerator / Co-Processor designs exist in the database. Several generations of Intel chips have been overtaken by NVIDIA chips (GPUs) in recent years, with NVIDIA now dominating the market by performance share. A variety of categories can be plotted over time using the resource at https://www.top500.org/statistics/overtime/.

A New Leader and the First to Exascale Capacity

The new top system, Fugaku (refer to **Figure 3**), turned in a High Performance LINPACK (HPL) result of 415.5 petaflops, besting the now second-place Summit system (Oak Ridge National Lab's Summit supercomputer in Tennessee, built by IBM) by a factor of 2.8x. Fugaku also beat the competition in the High Performance Conjugate Gradients (HPCG) benchmark used to test real-world application performance; the Graph500, a rating for dataintensive loads; and HPL-AI, a benchmark for rating artificial-intelligence workloads. A Fugaku prototype also took top spot for the most energy-efficient system on the Green500 list last November, achieving an outstanding 16.9-gigaflops-per-watt power efficiency during a 2.0 petaflops-per-second LINPACK performance run.

Fugaku is powered by Fujitsu's 48-core A64FX SoC, becoming the first number one system on the list to be powered by ARM processors. In single or further reduced precision, which are often used in machine learning and AI applications, Fugaku's peak performance is over 1,000 petaflops (1 exaflops).

Fujitsu claim their CPU is the world's first CPU to adopt Scalable Vector Extension—an instruction-set extension of Arm v8-A architecture for supercomputers. The 512-bit, 2.2-GHz CPU employs 1,024-gigabytes-per-second 3D-stacked memory and can handle half-precision arithmetic and (single-precision) multiply-add operations that reduce memory loads in AI and deep-learning applications where lower precision is admissible. The CPUs are directly linked by a 6.8-GB/s network Tofu D interconnect that uses a 6-dimensional mesh torus connection.

Figure 3. The Fugaku supercomputer facility at RIKEN Center for Computational Science (R-CCS) in Kobe, Japan.

The Rest of the Top 10

2. Number two on the TOP500 list is Summit, an IBM-built supercomputer that delivers 148.8 petaflops on HPL. Summit is running at Oak Ridge National Laboratory (ORNL) in Tennessee, and remains the fastest supercomputer in the US.
3. At number three is Sierra, a system at the Lawrence Livermore National Laboratory (LLNL) in California achieving 94.6 petaflops on HPL. Its architecture is very similar to Summit.
4. Sunway TaihuLight, a system developed by China's National Research Center of Parallel Computer Engineering & Technology (NRCPC) drops to number four on the list. Its HPL mark of 93 petaflops has remained unchanged since it was installed at the National Supercomputing Center in Wuxi, China in June 2016.
5. At number five is Tianhe-2A (Milky Way-2A), a system developed by China's National University of Defense Technology (NUDT), and deployed at the National Supercomputer Center in Guangzhou, China, with a HPL performance of 61.4 petaflops.
6. A new system on the list, HPC5, captured the number six spot, turning in an HPL performance of 35.5 petaflops. HPC5 is a PowerEdge system built by Dell and installed by Eni S.p.A, making it the fastest supercomputer in Europe.
7. Another new system, Selene, is in the number seven spot with an HPL mark of 27.58 petaflops. Selene is installed at NVIDIA in the US.
8. Frontera, a Dell C6420 system installed at the Texas Advanced Computing Center (TACC) in the US is ranked eighth on the list, with 23.5 HPL petaflops.

9. The second Italian system in the top 10 is Marconi-100, which is installed at the CINECA research center. Marconi-100's 21.6 petaflops earned it the number nine spot on the list.

10. Rounding out the top 10 is Piz Daint at 19.6 petaflops, a Cray XC50 system installed at the Swiss National Supercomputing Centre (CSCS) in Lugano, Switzerland.

HPC for Seismic Imaging (Actually, for Imaging by Inversion)

While there is much debate regarding the optimal platforms for large seismic applications, it can be observed that significant progress in HPC usage has been made in recent years due to the interplay between hardware and software advances:

- Much larger architectures becoming available; including shared memory architectures, commodity CPU / GPU, and cloud vendor solutions;
- Significant improvements in software for massively parallel applications, including domain decomposition and data compression schemes for different applications.

Regards high-end seismic imaging, conventional migration is being replaced by inversion-based and least-squares solutions, and FWI is moving towards high-frequency inversions and deep-model updating with ultra-long offsets; each solution with significant increases in associated shared memory demands.

As Bienati and Calcagni from Eni (2020) so elegantly describe, HPC technology, as a key enabler for the application of more accurate algorithms to richer datasets without impacting speed of execution, may provide the means to better characterize uncertainty, which remains inherent even in our best picture of the subsurface. Correspondingly, vast non-linear inversion challenges to optimize a range of exploration and production applications are expected to play a key part in future HPC requirements.

Summary

The growth in supercomputer capabilities over the past few decades has followed an exponential trend, and used a diverse variety of chip and system architectures. The transition to a balance between on-premise and cloud-based HPC is fast moving, far from trivial, and success is not yet measurable for the next wave of seismic imaging and AI growth. Computer architectures are constantly superseded by newer designs; just one of several important challenges to the adoption of cloud-based HPC, and a challenge for cloud-based HPC providers to stay competitive. Each supercomputer is built for selected applications, and overall HPL performance is a starting point only for determining the comparative usefulness to massive computational challenges.

How can we estimate the comparative capability of any supercomputer in the TOP500 to run specific seismic applications? For a given chip architecture (refer to the diversity of platforms in **Figure 2**), the supercomputer performance may scale relatively with the number of available nodes, but that performance is fundamentally dependent upon how the application software balances the load across all nodes while maximizing throughput. Overall, as more shared compute and memory becomes available, hitherto unachievable applications with global updates / network training / etc. can enable better management of uncertainty and risk, and that never-ending path towards increasing problem-solving sophistication is what drives investment in larger supercomputers.

Suggested Reading Material

- Bienati, N., and Calcagni, D., 2020, Leveraging high performance computing to enhance subsurface uncertainty management: Eni's experience. First Break, 38(7), 57-60. https://www.earthdoc.org/content/journals/10.3997/1365-2397.fb2020050
- Eni supercomputer pumps out reservoir simulations at record speed. https://www.top500.org/news/eni-supercomputer-pumps-out-reservoir-simulations-at-record-speed/
- Our Pangea III supercomputer (Total). https://www.ep.total.com/en/expertise/exploration/pangea-iii-world-first-commercial-computing-power
- Petrobras runs biggest supercomputer in Latin America. https://agenciabrasil.ebc.com.br/en/economia/noticia/2019-06/petrobras-runs-biggest-supercomputer-latin-america