

From Computational Science to Science Discovery: The Next Computing Landscape

Gilad Shainer, Brian Sparks, Scot Schultz, Eric Lantz, William Liu, Tong Liu, Goldi Misra

HPC Advisory Council

{Gilad, Brian, Scot, Eric, William, Tong, Goldi@hpcadvisorycouncil.com}

Computational science is the field of study concerned with constructing mathematical models and numerical techniques that represent scientific, social scientific or engineering problems and employing these models on computers, or clusters of computers to analyze, explore or solve these models. Numerical simulation enables the study of complex phenomena that would be too expensive or dangerous to study by direct experimentation. The quest for ever higher levels of detail and realism in such simulations requires enormous computational capacity, and has provided the impetus for breakthroughs in computer algorithms and architectures. Due to these advances, computational scientists and engineers can now solve large-scale problems that were once thought intractable by creating the related models and simulate them via high-performance compute clusters or supercomputers. Simulation is being used as an integral part of the manufacturing, design and decision-making processes, and as a fundamental tool for scientific research. Problems where high-performance simulation play a pivotal role include for example weather and climate prediction, nuclear and energy research, simulation and design of vehicles and aircrafts, electronic design automation, astrophysics, quantum mechanics, biology, computational chemistry and more.

Computational science is commonly considered the third mode of science, where the previous modes or paradigms were experimentation/observation and theory. In the past, science was performed by observing evidence of natural or social phenomena, recording measurable data related to the observations, and analyzing this information to construct theoretical explanations of how things work. With the introduction of high-performance supercomputers, the methods of scientific research could include mathematical models and simulation of phenomenon that are too expensive or beyond our experiment's reach. In turn, we can forecast weather conditions sooner, explore alternative energy sources, build safer vehicles and package consumed goods in a more economical way. In order to perform those numerical simulation effectively and productively, cost-effective or commodity based supercomputers architectures were created – high-performance clustering of computers.

High-performance computing (HPC) clusters are scalable performance compute solutions based on industry-standard hardware connected by a private system high-speed network. The main benefits of clusters are affordability, flexibility, availability, high-performance and scalability. A cluster uses the aggregated power of compute server nodes to form a high-performance solution for parallel applications. When more compute power is needed, it can be simply achieved by adding more server

nodes to the cluster. The Los Alamos National Lab (US) “Roadrunner” cluster (figure 1) was the first system to provide Petaflop (a thousand trillion CPU floating point operations or instructions per second) performance for scientific simulations (national nuclear weapons, astronomy, human genome science and climate change). Roadrunner was built using IBM Cell CPUs and AMD Opteron CPUs boards, and Mellanox InfiniBand to connect between them. Oak Ridge National Lab (US) “Spider” system is one of the world’s largest and fastest storage cluster file system that includes thousands of connections (based on InfiniBand interconnect) and over 10.7 PetaByte storage capacity to serve the high-performance systems at the lab. The National University of Defense Technology (China) “TianHe” system (figure 2) is the first Petascale system in Asia. The system is using thousands of Intel CPU and ATI GPUs, all connected via Mellanox InfiniBand networking.




Figure 1 – Los Alamos National Lab “Roadrunner” systems – the world’s first Petaflop system



Figure 2 – National University of Defense Technology “TianHe” system

With the creation of bigger and faster high-performance computing systems for scientific and engineering simulations, new generations of sensor-computer appliances have been created for specific applications. One example is the, the Australian Square Kilometre Array Pathfinder (ASKAP), an array of

radio telescopes that will comprise of 36 antennas each 12m in diameter, capable of high dynamic range imaging and using wide-field-of-view phased array feeds. ASKAP will be a telescope that can capture radio images with unprecedented sensitivity over large areas of sky. With a large instantaneous field-of-view ASKAP will be able to survey the whole sky vastly faster than is possible with existing radio telescopes.




Figure 3 – Illustration of the Australian Square Kilometre Array Pathfinder

Petaflop Supercomputers Create Exa-flood of Data

The ever increasing demands for computational power delivered by the ever increasing supercomputer capability and capacity produce an overwhelming flow of data. In one week the Australian Square Kilometre Array Pathfinder will generate more information than is currently contained on the whole World Wide Web, and in one month it will generate more information than is contained in the world's academic libraries. A Petaflop supercomputer equals 150,000 computations for every human on the planet per second, and a single day's usage world TOP500 supercomputers (according to the November 2009 list) is equal to 240 billion people armed with calculators for nearly 50 years.

With the increasing ramp of data generation from scientific and engineering simulations and observation targeted supercomputers, future technology development should be focused on creating scalable high-performance clusters of computers that can manage and process all of this data. The future premise of compute infrastructures should be aimed into building or providing tools and systems for “science discovery”, in which all of the computational science literature and databases can be available online and be shared by scientists, researchers and engineers around the globe. Distributed science can be seen as the fourth mode or paradigm where science becomes centralized throughout centralization of computing facilities, and those computing facilities are then targeted into managing,

visualizing and analyzing the data flood. Computational science drives the vast creation of data which is beyond our capabilities to analyze and understand, and the role of science discovery will extend to create the tools to extract the future science discoveries out of the data flood.

Furthermore, in many scientific fields of studies, the instruments are extremely expensive, and as such, the data must be shared. With this data explosion and as high-performance systems become a commodity infrastructure, the pressure to share scientific data is increasing. That resonates well with the emerging computing trend known as “the cloud” or “cloud computing”. While for the moment cloud computing appears to be a cost effective alternative for IT spending, or the shift of enterprise IT centers from capital expense to operational expense, research institutes have started exploring how cloud computing can create the desired compute centralization and an environment for researchers to share and crunch the flood of data. One example is the new system at the National Energy Research Scientific Computing Center (US), named “Magellan”. While Magellan’s initial target is to provide a tool for computational science in a cloud environment, it can be easily modified to become a center for data processing accessed by many researchers and scientists.

Centralized Data Crunching Compute Environment Throughout Cloud Computing

The concept of computing “in a cloud” is typically referred as a hosted computational environment (could be local or remote) that can provide elastic compute and storage services for users per demand. Therefore the current usage model of cloud environments is aimed for computational science. Future clouds can be served as environments for distributed science to allow researchers and engineers to share their data with their peers around the globe and allow expensive achieved results to be utilized for more research projects and scientific discoveries.

To allow the shift to the fourth mode of “science discovery” those cloud environments will need not only to provide capability to share the data created by the computational science and the various observations results, but also to be able to provide cost-effective high-performance computing capabilities, similar to that of today’s leading supercomputers, in order to be able to rapidly and effectively analyze the data flood. Moreover, an important criteria of clouds need to be fast provisioning of the cloud resources, both compute and storage, in order to service many users, many different analysis and be able to suspend tasks and bring them back to life in a fast manner. Reliability is another concern, and clouds need to be able to be “self healing” clouds where failing components can be replaced by spares or on-demand resources to guarantee constant access and resource availability.

The use of Grids for scientific computing has become successful in the last few years and many international projects led to the establishment of world-wide infrastructures available for computational science. The Open Science Grid provides support for data-intensive research for different disciplines such as biology, chemistry, particle physics, and geographic information systems. Enabling Grid for EScience (EGEE) is an initiative funded by the European Commission that connects more than 91 institutions in Europe, Asia, and United States of America, to construct the largest multi-science computing Grid infrastructure of the world. TeraGRID is an NSF funded project that provides scientists

with a large computing infrastructure built on top of resources at nine resource provider partner sites. It is used by 4000 users at over 200 universities that advance research in molecular bioscience, ocean science, earth science, mathematics, neuroscience, design and manufacturing, and other disciplines. While Grids can provide a good infrastructure for shared science and data analysis, several issues make the Grids problematic to lead the fourth mode of science – limited software flexibility, applications typically need to be pre-packaged, non elasticity and lack of virtualization. Those missing items can be delivered through cloud computing.

Cloud computing addresses many of the aforementioned problems by means of virtualization technologies, which provide the ability to scale up and down the computing infrastructure according to given requirements. By using Cloud-based technologies scientists can have easy access to large distributed infrastructures and completely customize their execution environment. Furthermore, effective provisioning can support many more activities and suspend or bring to life activities in an instant. This makes the spectrum of options available to scientists wide enough to cover any specific need for their research.

High-Performance Cloud Computing

In the past, high-performance computing has not been a good candidate for cloud computing due to its requirement for tight integration between server nodes via low-latency interconnects. The performance overhead associated with host virtualization, a prerequisite technology for migrating local applications to the cloud, quickly erodes application scalability and efficiency in an HPC context. The new virtualization solutions such as KVM and XEN aim to solve the performance issue by allowing native performance capabilities from the virtual machines by reducing the virtualization management overhead and by allowing direct access from the virtual machines to the network.

High-speed networking is a critical requirement for affordable high-performance computing, as clusters of servers and storage need to be able to communicate as fast as possible between them. A vast majority of the world top 100 supercomputers are using the high-speed InfiniBand networking due to this reason, and the interconnect allows those systems to reach to more than 90% efficiency, a critical element for effective for high-performance computing in any infrastructure, including clouds. National Energy Research Scientific Computing Center (NERSC, US) “Magellan” system is using InfiniBand as the interconnect to provide the fastest connection between servers and storage in order to allow the maximum gain from the system, highest efficiency and an infrastructure that will be able to analyze data in real time.

Power consumption is another important issue for high-performance clouds. As the HPC clouds become bigger, affordability of science discovery will be determined by the ability to save the costs of the power and cooling. Power management, which is implemented within the CPUs, the interconnect and the system management and scheduling will need to be integrated as a comprehensive solution. Non-utilized sections of the clouds need to be powered off or moved into power saving states and the scheduling mechanism will need to incorporate topology awareness.

The HPC Advisory Council (<http://www.hpcadvisorycouncil.com>) HPC|Cloud group is working to investigate the creation and usage models of clouds in HPC. Past activities on smart scheduling mechanisms have been published on the council's website, and future results will include the usage of KVM and XEN, many cores CPUs (such as AMD's Magny-Cours which includes 12 cores in a single CPU) and cloud management software (such as Platform ISF) will be published throughout 2010.

Science Discovery: The Next Computing Landscape

Science discovery through data-intensive analysis can be the next mode of science, after experimentation/observation, theory and computational science. This will be the mode in which high-performance cloud computing will connect the globe and provide the tool for researchers, scientists and engineers to share their experiments and to analyze the increasing data that is being gathered or created. Those cloud environments will be based on commodity servers and storage, all connected via high-speed networking with a comprehensive economical virtualization software management.

The HPC Advisory Council will continue to investigate the emerging technologies and aspects that will lead us into the fourth mode of science.

Acknowledgments

The authors would like to thank Cydney Stevens for her vision and guidance.