The Paradigm shift towards Exascale computing and its Applications

V.Rajaprakash, K.Jyothimurugan, J.Kalaiarasi,
UG Scholar, Department of Computer Science and Engineering,
Mahendra Engineering College, Namakkal, India

Abstract: Exascale computing is a new technique which is used in Big Data, Data Mining, etc. It is the next level after supercomputing. It is used to solve many complex mathematical operations. Its speed in exaFLOPS and hence it is called exascale computing.

***IndexTerms*—exaflop, ParallelX, HPX-5, brain science**

I. INTRODUCTION

Exascale computing refers to computing systems capable of at least one exaFLOPS or a billion calculations per second. Such capacity represents a thousand fold increase over the first petascale computer that came into operation in 2008. One exaflop is a thousand petaflops or a quintillion, 10^{18} floating point operations per second.

Going to the exascale will mean a radical change in computing architecture – basically, vastly increasing the levels of parallelism to the point of millions of processors working in tandem – which will force radical changes in how hardware is designed, in how we go about solving problems, and in how we link application codes to the underlying hardware. Understanding the advantages to be gained by going to the exascale, and evaluating the risks involved by going down this path, requires both an evaluation of past experiences in moving from the megaflop era to the present petaflop era, as well as an assessment of the readiness of advanced applications to take transformative advantages of exascale computing.

For the growing number of problems where experiments are impossible, dangerous, or inordinately costly, extreme-scale computing will enable the solution of vastly more accurate predictive models and the analysis of massive quantities of data, producing quantum advances in areas of science and technology that are essential to DOE and Office of Science missions. The exascale computing will push the frontiers of the following:

- Adaptation to regional climate changes such as sea level rise, drought and flooding, and severe weather patterns
- Reduction of the carbon footprint of the transportation sector
- Efficiency and safety of the nuclear energy secotr
- Innovative designs for cost-effective renewable energy sources such as batteries, catalysts and biofuels
- Certification of the U.S. nuclear stockpile, life extension programs and directed stockpile work
- Design of advanced experimental facilities, such as accelerators, and magnetic and inertial confinement fusion
- First-principles understanding of the properties of fusion and fusion reactions
- Reverse engineering of the human brain
- Design, control and manufacture of advanced materials

II. NEEDS FOR MOVING TOWARDSEXASCALECOMPUTING

The computing problems can be divided into three general categories [1].

A. Incrementally advanced computing

In traditional scientific computing, many computational problems are characterized by a benign connection between advances in computing capabilities and the benefits gained by such advances.

B. Voracious computing

In some problem areas of computational sciences, there is no a priori reason to believe that increased computing capabilities will lead to increased knowledge, even in an incremental way.

C. Transformational computing

In some cases, computational problems have the property that they can be fully solved with sufficiently large computations, where ‘sufficiently large’ means that one can specify quantitatively in advance what ‘large enough’ means. A sufficiently large computation will be capable of fully solving the problem. Such computing is transformational in the sense that a smaller calculation does not solve the underlying problem. Such computing is transformational in the sense that a smaller calculation does not solve the underlying problem, but a sufficiently large calculation can solve the problem completely.

III. GOALS OFEXASCALE COMPUTING

Goals and objectives for this project fall into three categories namely: i) Technology research and development, ii) technology transfer and commercialization, and iii) education [3].

A. Technology Research and Development

The goal is to create a paradigm shift and to develop enabling technology for accelerating progress in applied fields. The objectives for the core computer technology research and development are to: a) further refine innovation in execution models to guide the development of new system component layers and govern their interoperability in support of extreme scale application computing, b) fully structure an advanced introspective dynamic adaptive runtime system for guided computing, resource management and task scheduling that uses system and application runtime information for continued operational optimization through high efficiency and parallelism discovery, c) devise programming interfaces that expose powerful semantic constructs for exposure and exploitation of parallelism as well as interoperability with other program modules, and legacy codes, d) design hardware architecture components to accelerate runtime mechanisms for greater efficiency and scalability and implement these designs at IU and with industrial partners.

The objectives for the application technology research and development are to: a) cooperatively develop a select set of extreme-scale applications that will enable breakthroughs in strategic leading-edge fields – both to drive exascale advances and to demonstrate potential end-user achievement through future innovative exascale computing, b) support other IU Grand Challenge projects that may heavily rely on and benefit from exascale computing to better achieve their goals.

B. Technology Transfer and Commercialization

This goal is to accelerate adoption of the technology in the marketplace and to bring economic benefits to IU, Bloomington, and the state of Indiana. The objectives are to: a) derive plans for intellectual property stewardship and technology transfer from our results to industrial, commercial, academic, and government customers and implement these plans to yield market and licensing revenue, b) provide professional facilities to support and deploy IU exascale hardware and software products, leveraging established UITS capabilities, procedures, staff, and infrastructure management.

C. Education

This goal is to disseminate results of the research and development and further increase IU's reputation as a premier academic institution. The objectives are to: a) enhance IU education across a diversity of disciplines with scientific computing, b) support the nascent Intelligent System Engineering Department for graduate and undergraduate curriculum and research, so that it emerges as a leading program of its kind, c) provide cutting edge research opportunities for undergraduates from regional campuses.

IV. COUNTRIES INVOLVED IN EXASCALE COMPUTING

Japan was the only country in the world with a definitive road map to its first exascale supercomputer. There are specific plans for exascale systems in China, France, and the US has been revealed [2].

The U.S. Department of Energy is planning to bring two exascale machines to fruition by 2023. The 2023 date was formalized in the DOE's Exascale Computing Plan (ECP), which made public in April 2016. ECP hardware development will be done under the PathForward program, which will fund a number of vendors to conduct the necessary R & D.

The project called Flagship 2020 was launched in April 2014 by the Ministry of Education, Culture, Sports, Science and Technology by Japan. The RIKEN Advanced Institute for Computational Science will be the first recipient of the first system, as well as the development partner.

China revealed that a program to develop the country's first exascale supercomputers has begun. The R&D funding for these systems is being done under the 13th Five-Year Development Plan, which kicked off in 2016. The first systems and at this point is not clear how many are bring considered are scheduled to come online in 2020.

Although Europe has a number of exascale R&D programs in motion under the Partnership for Advanced Computing in Europe umbrella, only France and specifically French computer-maker Atos/Bull has a specific roadmap in place to develop and deliver an exascale computer.

V. EXASCALETECHNOLOGY STACK

Delivering exascale performance will require fundamental advances in multitier inter-related technology areas, from hardware architecture through application software design. Figure 1 shows the exascale technology stack.

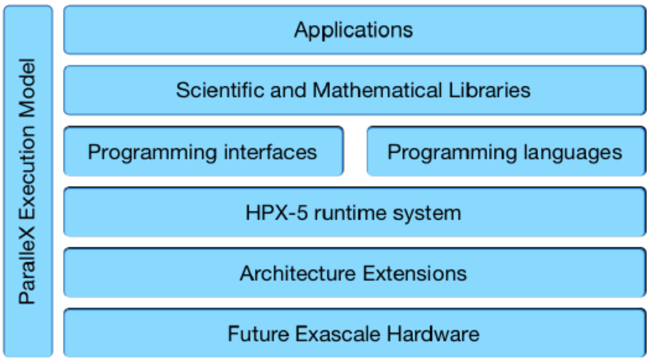

Fig. 1. Six layered architecture of IoT

A. ParallelX Execution model

TheParallelX execution model provides the conceptual framework for the cross-cutting design and operation of future exascale computing systems and in particular addresses the technical challenges through dynamic adaptive mechanisms and semantics. ParallelX incorporates and integrates the elements of local dataflow execution, distributed contexts, message-driven computation, global name spaces, and powerful synchronization mechanisms. It can support heterogeneous system architectures including GPUs and expose and exploit diverse parallel forms for scalability and adaptive load balancing.

B. HPX-5 runtime system

TheHPX-5 runtime system is a reduction to practice of the ParallelX concepts and delivers many of its benefits to large-scale conventional systems and provides a template for future advanced exascale computer architectures for reduced overheads and increased scaling.

C. Programming interfaces and languages

Access to runtime system capabilities will be provided directly through defined interfaces as well as through programming language extensions. The programming interfaces expose capabilities of the HPX-5 runtime system, whereas the programming language extensions make these capabilities first-class entities in the host programming language such as C and C++. Full support for the programming language will require toolchain development, including compilers, utility libraries, debuggers, and performance monitoring infrastructure.

D. Architectural extensions

For sufficient performance, some capabilities provided by the runtime system requires direct support by the underlying hardware. Support mechanisms can be realized directly using FPGAs, a technology specifically for hardware prototyping. In some cases, it may be necessary to incorporate hardware support mechanisms into the processor core, which will require vendor participation. Runtime system capabilities that are expected to benefit from hardware acceleration include thread and message scheduling, active global address translation, and fine-grained synchronization.

VI. APPLICATIONS OF EXASCALE COMPUTING

The development of exascale system software and architecture will be guided by a co-design process including key parallel applications, emphasizing both importance and intellectual content. These driver applications represent areas of brain science, machine learning for data analytics, materials and manufacturing, mesoscalemetereological phenomena, and computational science.

A. Brain Science

We are entering a golden age of data collection and computational modeling in neuroscience, which presents numerous opportunities for exascale computing. Data from instrumented neurons must be processed in order to understand the spatial and temporal relationships in the neuronal network. Inferring functional connections between only 500 neurons based on the time series data of their spikes might take 10-20 hours on a supercomputer. Future experiments will require rapid data analysis so that spike identification and functional network structure can be obtained within minutes. With this kind of rapid analysis, a first draft of network connectivity can be completed while the tissue is still viable, which can serve as a model to guide stimulation in maximally informative ways, allowing the model to be updated and refined. Additional variables like sensory inputs, motor output and behavioral state could also be added. To develop and refine predictive models of the brain, experiments of the future will involve a rapid loop between millions of neurons and an actively behaving animal.

B. Deep machine learning

The ecosystem surrounding Big Data includes a number of high-productivity systems that unfortunately suffer from performance limitations. Careful analysis of application characteristicshas enabled us to develop systems that provide the "best of both worlds" – a high-performanceruntime combined with the programmability of (e.g.) MapReduce. The IU Digital Science Center (DSC)has worked with NSF funding to look at multiple application communities such as bio-molecular simulations,network science and computational social science, computational epidemiology, computer vision,deep learning, et al. and so is well prepared to support a wide variety of IU needs. First, we will develop ahigh-performance “Infrastructure as a Service” that allows the same code to be deployed on cloud environmentsas well as HPC clusters. The systems to be built are insensitive to underlying technology changeand allow trade-offs to be made regarding performance, cost, security and ease of use. Second, we willintegrate our exascale runtime into streaming and batch programming models to allow a full range of dataanalytics to run with scalable performance. Third, we will develop the Scalable Parallel Interoperable DataAnalytics Library with core machine learning libraries combining parallel database backends and high-performanceruntimes.

C. Severe Weather prediction

Numerical weather prediction (NWP) exploits the cutting edge of supercomputing,yielding steadily improving forecasts of large-scale weather patterns on timescales of days toweeks. Through HPC, researchers have continually increased the resolution of NWP models, such thatthey can now resolve individual thunderstorms (on the order of 10 km or less) and produce useful forecasts on timescales of minutes to hours, even real-time operational forecasting unheard of 15 years ago.Severe thunderstorms cause numerous deaths, injuries, and property damage from flooding, hail, lightning,high winds, and tornadoes. Tornadoes constitute a particularly challenging problem due to their inherentlysmall size and short duration while exhibiting extremes of wind force. There is a great need tomeet this challenge, due to the high threat for tornadoes--which is maximized in the Plains and Midwest.Current operational NWP models can only resolve the larger scale mesocyclone circulations within thethunderstorms that spawn tornadoes. The ability to reliably resolve tornado-scale circulations in near realtimewould require exascale supercomputing technology with many fundamentally unsolved scientificproblems surrounding the development and behavior of tornadoes still remaining.

D. Materials by design for Advanced Manufacturing

A revolution in in manufacturing is occurring through additive fabrication commonly referred to as “3D printing”. Now being used as structural components incomplex systems like motors, engines, product bodies, and even human replacement elements, this coremethodology demands rapid simulation of strength, thermal, stress-strain, electrical, and other mechanicalproperties some of which are uniquely caused by the fabrication process itself. With each printed objectcustom made, complex simulation is in the critical path and therefore cost of production. Advances inHPC including exascale will become increasingly important to the success of additive manufacturing andthe emerging industries based on it. These related applications will support the nascent Intelligent SystemsEngineering (ISE) program at IUB. It arises from and will extend the successful large-scale Indianacollaboration among IU, Notre Dame and Purdue to design a complete solution to shock-wave productionof cubic boron nitride. Building on this existing capability, the proposed exascale application will presenta unique opportunity that will significantly enhance the attractiveness of the new engineering program atIU as well as lay the foundation for new opportunities in Indiana’s industrial base and skill force.

E. Computational Materials, Biology and Energy Science

Much of molecular science is based on a detailedunderstanding of quantum mechanics, dictated by the probabilistic interpretation of nature as expoundedby Schrödinger’s equations. A detailed understanding of these fundamental processes is criticalfor (a) designing artificial energy sources that use and store solar power through photosynthesis, (b) cancerresearch and vaccine design through the action of certain biological enzymes, (c) the design of newmaterials that will have impact on information technology through “quantum enabled” chips, and (d) impactair pollution and climate change. Detailed understanding of these processes is a major challenge ofcomputational science and applied mathematics. Solving such problems is extremely compute-intensive(NP-hard), requiring new heuristic-based algorithms, software frameworks, and exascale power to solve.

VII. CONCLUSION

There has been a lot of excitement regarding exascale computing and the extreme computing power it will make available to the scientific community as well as industry. In addition, there are considerable flow-down benefits that will result from meeting both the hardware and software high performance computing challenges.

Acknowledgement

We convey our thanks to our guide Dr.A.Neela Madheswari, Associate Professor, CSE Department for her effort to bring this work. We express our gratitude to our head of the department Dr.M.Kannan for his entire support in all our endeavors.

REFERENCES

[1] The Opportunities and Challenges of Exascale Computing, U.S. Department of Energy, Office of Science, 2010.

[2] https://www.top500.org/news/the-four-way-race-to-exascale/

[3] www.indiana.edu/~vpr/documents/.../Exascale_Computing_Systems_proposal.pdf