

Secretary of Energy Advisory Board

Report of the Task Force on Next Generation High Performance Computing

August 18, 2014



U.S. Department of Energy

Final Version for Approval

Charge to the SEAB High Performance Computing Task Force	4
Executive Summary	4
Key Findings	5
Summary of Recommendations	6
Historical Perspectives	7
The New Era of Supercomputing	8
Uncertainty Quantification	9
Systems-of-Systems	9
Internet-of-Things	9
Data Centric Systems Development	9
Needs for Next Generation High Performance Computing	10
Implications for Industry: Systems of Insight.....	10
Implications for Basic Science: Discovery through Modeling, Simulation and Analysis.....	12
Impact: Computational Scientific Discovery - enabling modeling and simulation of complex systems at a hitherto unattainable level of detail	13
Impact: Engineering Design and Optimization - enhancing our ability to incorporate science based analysis and simulation in engineering designs.....	13
Impact: Data Analytic Discovery - allowing us to analyze and interpret large data sets generated by large scientific instruments, ubiquitous sensors, and simulations	14
Fostering and Maintaining a Balanced Ecosystem	15
The Ecosystem	15
Hardware	15
Software	16
Application Algorithms	17
Data Analytics and Discovery Tools.....	17
Agent modeling tools	18
Cognitive Computing	19
Next Generation Neural Network Architectures	20
Partnerships and Markets	21
Path Forward	22
Acceleration of Advanced High Performance Computing Systems.....	22
A Robust Ecosystem	23
ISVs	24
Open Source Software	24

Custom Codes.....	25
Partnerships	25
References.....	27
Appendix 1: Task Force and Study Participants.....	28
Appendix 2: Charge to the Task Force	29

Charge to the SEAB High Performance Computing Task Force

On December 20, 2013, Secretary of Energy, Dr. Ernest J. Moniz, requested the co-chairs of the Secretary of Energy Advisory Board (SEAB), Professors John Deutch and Persis Drell, to form a Task Force “composed of SEAB members and independent experts to review the mission and national capabilities related to next generation high performance computing.”

Secretary Moniz requested that the Task Force look at the problems and opportunities that will drive the need for next generation high performance computing (HPC), what will be required to execute a successful path to deliver next generation leading edge HPC, make recommendations regarding if and to what degree the U.S. Government should lead and accelerate the development of next generation leading edge HPC, and make recommendations as to what specific role the DOE should take in such a U.S. Government program.

The Task Force was asked to deliver its report by June, 2014 and to discuss its report and its conclusion at the June 2014 SEAB meeting.

A copy of the full charge for the Task Force is shown in Appendix 2.

Executive Summary

For over 60 years, the federal government, partnering with the U.S. Computer Industry, has driven the state of the art in high performance computing. This has been spearheaded by the Department of Energy and the Department of Defense but largely led and driven by the Department of Energy -- primarily for the NNSA weapons development, and now, stockpile stewardship responsibilities. Advances in high performance computing have focused on computational capabilities in the solution of partial differential equations, as measured by the speed in floating point operations per second (FLOPs).

Current leadership machines across the national laboratory system, and in some premiere industrial applications, are delivering performance in the tens of petaflop range. These machines largely have been developed by following the historical path of the last several decades, taking advantage of Moore’s law progression to smaller and faster CMOS computing elements, augmented by the highly parallel architectures that followed the vector processing change at the pre-teraflop generation.

The computing environment has begun to change as the complexity of computing problems grows, and with the explosion of data from sensor networks, financial systems, scientific instruments, and simulations themselves. The need to extract useful information from this explosion of data becomes as important as sheer computational power. This has driven a much greater focus on data centric computing, linked to integer operations, as opposed to floating point operations. Indeed, computational problems and data centric problems are coming together in areas that range from energy, to climate modeling, to healthcare. This shift dictates the need for a balanced ecosystem for high performance computing with an undergirding infrastructure that supports both computationally-intensive and data centric computing. Plans are in place through the CORAL development and procurement program to deliver systems of

about 200 petaflops performance, with up to 5-10 petabytes of addressable and buffer memory in a data centric architectural context, with attendant focus on power efficiency, reliability and productive usability¹. In fact, the architecture of computing hardware is evolving, and this means that the elements of the backbone technology -- including memory, data movement, and bandwidth -- must progress together. As we move to the era of exascale computing, multiple technologies have to be developed in a complementary way, including hardware, middleware, and applications software.

Our findings and recommendations are framed by three broad considerations:

1. We recognize and recommend a “new” alignment between classical and data centric computing to develop a balanced computational ecosystem.
2. We recognize the DOE historical role and expertise in the science, technology, program management, and partnering, and recognize its vital role across USG, including in the National Strategic Computing Initiative (NSCI).
3. We examine and make recommendations on exascale investment but also on nurturing the health of the overall high performance computing ecosystem, which includes investment in people and in mathematics, computer science, software engineering, basic sciences, and materials science and engineering.

Key Findings

The following summarizes the key findings of our work:

1. **Investable needs exist for an exaX class machine.**
 - a. The historical NNSA mission (simulation for stewardship), multiple industrial applications (e.g., oil and gas exploration and production, aerospace engineering, and medicinal chemistry (pharmaceuticals, protein structure, etc.)) and basic science all have applications that demonstrate real need and real deliverables from a significant performance increase in classical high performance computing at several orders of magnitude beyond the tens of petaflop performance delivered by today’s leadership machines.
2. **Significant but projectable technology development will enable exascale level data centric computing.**
 - a. Optimization of current CMOS, highly parallel processing within the remaining limits of Moore’s law and Dennard scaling, together with data centric systems level innovations, will enable 1-10 exascale performance levels, within acceptable power budgets. Significant but projectable technology and engineering developments are needed to reach this performance level.
3. **“Classical” high end simulation machines are already significantly impacted by many of the data volume and architecture issues.**
 - a. The performance of many complex simulations is less dominated by the performance of floating point operations, than by memory and integer operations.

¹http://science.energy.gov/~media/ascr/ascac/pdf/meetings/20140331/CORAL_Update_for_ASCAC_March_31_2014_V31.pdf

- b. As the data sets used for classic high performance simulation computation become increasingly large, increasingly no-localized, and increasingly multi-dimensional, there is significant overlap in memory and data flow science and technology development needed for classic high performance computing and for data centric computing.
- 4. **Data centric at the exascale is already important for DOE missions.**
 - a. There is an evolution already underway in the DOE computing environment to one that supports more memory- and integer-operation dominated simulation for the NNSA security mission.
 - b. Applications of data centric computing for DOE, for other parts of the U.S. Government, and for the private sector, are rapidly scaling to and beyond levels of performance that are comparable to those needed for classic high performance floating point computation.
- 5. **Common challenges and under-girding technologies span compute needs.**
 - a. As the complexity of data centric problems increases, the associated calculations face the same challenges of data movement, power consumption, memory capacity, interconnection bandwidth, and scaling as do simulation-based computations.
- 6. **The factors that drive DOE’s historical role in leadership computing still exist and will continue to do so.**
 - a. The DOE National Labs are an important and unique resource for the development of next generation high performance computing and beyond.
 - b. The DOE partnering mechanisms with industry and academia have proven effective for the last several generations of leadership computing programs.
 - c. Because of its historical and current expertise in leading the development of next generation high performance computing, the DOE has a unique and important role to play in the National Strategic Computing Initiative.
- 7. **A broad and healthy ecosystem is critical to the development of exascale and beyond systems.**
 - a. Progress in leading-edge computational systems relies critically on the health of the research environment in underlying mathematics, computer science, software engineering, communications, materials and devices, and application/algorithm development.
 - b. A robust ecosystem requires a healthy vendor community, and the recognition of the importance to industry of commercial viability of HPC systems.
- 8. **It is timely to invest in science, technology, and human investments for “Beyond Next”.**
 - a. A number of longer term technologies will be important to “beyond next” generation high performance computing (superconducting, quantum computing, biological computation) but are not mature enough to impact the next leading edge capability investments at DOE.

Summary of Recommendations

- 1. DOE, through a program jointly established and managed by the NNSA and the Office of Science, should lead the program and investment to deliver the next class of leading edge machines by the middle of the next decade. These machines should be developed through a co-design process that balances classical computational speed and data centric memory and communications architectures to deliver performance at the 1-10 exaflop level, with addressable

memory in the exabyte range.

2. This program should be executed using the partnering mechanisms with industry and academia that have proven effective for the last several generations of leadership computing programs. The approximate incremental investment required is \$3B over 10 years. This would include a roadmap of DOE acquisitions, starting with the CORAL program. Such a roadmap would focus industry on key system level deliverables.
3. DOE should lead, within the framework of the National Strategic Computing Initiative (NSCI), a co-design process that jointly matures the technology base for complex modeling and simulation and data centric computing. This should be part of a jointly tasked effort among the agencies with the biggest stake in a balanced ecosystem.
4. DOE should lead a cross-agency U.S. Government (USG) investment in “over-the-horizon” future high performance computing technology, including hardware, software, applications algorithms, operating systems, data analytics and discovery tools, agent based modeling, cognitive computing, neurosynaptic systems, and other forward looking technologies, including superconducting computing.
5. DOE should lead the USG efforts to invest in maintaining the health of the underlying balanced ecosystem in mathematics, computer science, new algorithm development, physics, chemistry, etc. but also including ISV’s, the open source community, and other government entities.
6. The Path Forward requires operating in, and investing for, three timeframes and technology plateaus: (1) The greater Petascale timeframe (the next five years), (2) The Exascale timeframe (the next five to 10 years), and (3) Beyond Exascale.

We note that the combined DOE investment in maintaining a healthy ecosystem and pursuing over-the-horizon technology identification and maturation is in the range of \$100-150M per year.

Historical Perspectives

Over the past six decades, the U.S. government, spearheaded by the Departments of Energy and Defense, has sponsored the development and deployment of ever more capable HPC computing systems -- driving remarkable advances in the state of the art of high end computing, and establishing U.S. dominance in the area. The process has been characterized by a number of highly successful partnerships between government agencies and the U.S. computer industry, resulting in a continuously improving series of leadership systems to meet the government’s needs. Until now, the so-called supercomputer field was characterized by (a) an almost exclusive focus on computational capability for solving partial differential equations (i.e. FLOPS), (b) a handful of vendors with the technical and financial ability to participate, (c) little or no industrial and commercial demand for computation and simulation at the scales available, and hence (d) a market limited to government laboratories and a small number of research institutions, primarily in the U.S., but also in Europe and Japan.

This environment is rapidly evolving. As the complexity and sophistication of the problems they are required to address increases, and as these systems become more capable, the need to manage, analyze,

and extract useful information from the tremendous amounts of data they ingest and produce becomes commensurate and co-equal in importance to their computational power. This is the case across much of the government research enterprise, while the emerging confluence of Big Data and analytics capabilities with highly sophisticated modeling and simulation is promising to have a transformational effect on a number of major industries.

These changes are happening at a time when U.S. leadership in high end HPC is being seriously challenged by China and Europe. In order to address this changing environment, to continue to address national security requirements, and to realize the potential benefits to U.S. industrial competitiveness, the federal government, industry, and academia must partner in new ways to achieve the mutual goals of national security, economic security, and scientific leadership.

The New Era of Supercomputing

The government use of leading edge computing systems, developed by domestic computer manufacturers, goes back to the very dawn of the modern computer era with the application of the IBM/Harvard-developed Automatic Sequence Controlled Calculator (Mark I) in the Manhattan Project. It has continued, unabated, ever since. The Accelerated Strategic Computing Initiative (ASCI)² program, initiated in the early 1990's and brought about by the necessity to substitute modeling and simulation for physical testing of nuclear weapons, provided funding for and supported an industry/government partnership that greatly accelerated the pace of introduction of high end HPC technology. The successful introduction and exploitation of massively parallel systems and the software and messaging infrastructure that supports them are notable results of ASCI and its successor programs³. The systems they produced (at scale or in smaller versions) have been applied with great success to modeling and simulation phenomena in astrophysics, biophysics, materials science, combustion, climate modeling, weather forecasting, finance, oil and gas exploration, and a host of other fields.

As computer models of scientific phenomena have increased both in scale and in detail, the requirements for increased computational power, typically in the form of FLOPS, have increased exponentially, driving commensurate growth in system capability. The requirement for increasing FLOPS is not likely to slacken in the foreseeable future. However, the nature of the workloads to which these systems are applied is rapidly evolving. Even today, the performance of many complex simulations is less dominated by the performance of floating point operations, than by memory and integer operations. Moreover, the nature of the problems of greatest security, industrial, and scientific interest is becoming increasingly data-driven.

The highest performing computers of the future must be able to (1) quantify the uncertainty associated with the behavior of complex systems-of-systems (e.g. hurricanes, nuclear disaster, seismic exploration, engineering design) and thereby predict outcomes (e.g. impact of intervention actions, business implications of design choices); (2) learn and refine underlying models based on constant monitoring and past outcomes; and (3) provide real-time interactive visualization and accommodate “what if” questions in real-time (the “New Era of (Cognitive) Supercomputing”).

² Alex Lazelere has produced a comprehensive history of ASCI and its descendants, which can found at https://asc.llnl.gov/asc_history/

³ ASCI has been succeeded by the Advanced Simulation and Computing (ASC) program at DOE's weapons laboratories and the Leadership Computing programs at the science laboratories.

A more detailed examination of these requirements follows.

Uncertainty Quantification

Traditionally an entire high-end machine, with each increase in capability, has been devoted to simulating larger models of physical phenomenon at finer scale. Today, it is important to ask “what if” questions in many areas. This requires the study of a wide range of potential behaviors over a range of different conditions. A systematic approach to Uncertainty Quantification is becoming essential, and in itself can be a driver for exascale computing.

Systems-of-Systems

Increasingly, we want to better understand the behavior of coupled complex systems. For example, being able to simulate a combination of physical models for predicting the path of a hurricane with coastal topographic models, models of traffic patterns, and multimodal (text, image, audio) cell phone data about actual storm damage would enable local and state authorities to make more informed decisions as to the need and timing of evacuations, the allocation of disaster recovery resources, as well as in planning better evacuation strategies and routes.

Internet-of-Things

As the scope and sophistication of high-end HPC workloads increase, the proliferation of sensors of all kinds accelerates, the industrial deployment of solutions taking advantage of the massively connected and communicating elements (the Internet of Things), and the number and size of big science projects (e.g. the Hadron Collider, the Square Kilometer Array, the BRAIN Initiative Program, the Human Brain Project, etc.) grows, the resources and capabilities required to manage the data volumes involved begin to equal, if not surpass the capabilities and design challenges associated purely with computation. The ability to apply data analytics to harness these vast troves of information offers enormous potential not only to gain deeper scientific insights, but to identify and assess risks and threats, and to guide time critical, as well as strategic, decision making.

Data Centric Systems Development

Over the next five to ten years, in order to meet the continued computational and data-driven demands of emergent challenges and important problems in multiple domains, the highest performing computational systems must evolve to accommodate new data centric system architectures and designs, and an ever more sophisticated and capable software ecosystem. Evolving workflows will require the integration of a more diverse functionality in order to create a more flexible, data centric system design capable of efficiently handling data motion that will be highly variable in size, access pattern, and temporal behavior. Without close attention to these “data” issues, such systems will be hobbled by numerous data bottlenecks, and will fail to achieve their promised goals. Overall, it is important to recognize the importance of systems level innovation to satisfy requirements for data centric computation coupled to modeling and simulation.

For systems vendors, the tension between the strategic imperative of flowing technology and system components into the mainstream and the reality that mainstream (and commodity) markets drive different rates of technology and system adoption will create an ever-present design challenge. To operate at full capability, advanced HPC systems demand design elements, particularly in the areas of reliability, power efficiency, data movement, interconnect fabrics, storage, and I/O, that go beyond traditional market-

linked computational requirements and cost structures. The delayed adoption of some of the undergirding technologies into mainstream products could delay return on investment beyond financially acceptable levels for any given vendor. Consequently, the government has an important role to play both in continuing to invest in technology and systems development, and in promoting the application of HPC to industry.

Needs for Next Generation High Performance Computing

Implications for Industry: Systems of Insight

It is in the coupling of ever increasing capability for traditional modeling and simulation with the emerging capability for Big Data analytics that the potential for the significant impact on U.S. industry will be the greatest. In the commercial world, there is an emerging convergence of traditional systems oriented towards “back-office” functions, like transaction processing and data base management (“systems of record”), and systems focused on interactions that bring computing closer to the end user, like e-commerce, search, the cloud, and various social media (“systems of engagement”). This convergence, coupled with increasing HPC capabilities, will result in “systems of insight”, where modeling and simulation, analytics, big data, and cognitive computing come together to provide new capabilities and understanding.

Oil & Gas

The example of the petroleum industry provides insight into the promise and the challenges of the next generation of HPC systems. The upstream segment of the petroleum industry makes heavy use of HPC for exploration, and their data and computational requirements are growing exponentially. Many oil companies are predicting the need for exascale computing by the end of the decade. Some individual market players are already running data centers with over 60 petaflops of compute capacity and growing. Other players are contemplating data centers with hundreds of petaflops by the end of the decade. Unlike the integrated cutting edge HPC systems like those at DOE’s leading laboratories, this capacity is still typically in the form of huge clusters or multiclusters of generally available HPC servers, which are used to process large numbers of essentially independent computational and/or simulation tasks. This system organization and utilization pattern reflects today’s HPC driven oil exploration workflows, which are composed of many related, but distinct, high-level data processing stages. Individual stages are typically highly parallelizable, and, in some cases, multiple stages can be run concurrently. However, there is little or no automated integration of these stages – the stages stand as silos with user-based decisions only occurring when stages are complete. Stages frequently are rerun by hand, in an ad hoc fashion, when the output does not satisfy subjective criteria dictated by the experience of the operators.

The industry recognizes that the true exploratory power resides in the collective experience in the minds of their geophysicists; and that the full value of the data in this space will be achieved only when their geophysicists have the power to “play” with this process; to dynamically consider numerous, perhaps thousands or millions, of “what if” scenarios to leverage their knowledge to explore more effectively.

Additionally, companies are seeing the value of integrating various business areas for the added value the additional context provides. For example, imaging is being combined with reservoir simulation, which is

coupled to oil field management, which feeds into a long, complex supply chain that needs to be optimized for a variety of factors -- including market demand, weather, ship availability, reservoir production rates, etc.

Enabling this kind of coupled operation to unlock the value it offers requires the deep integration of currently siloed stages; it requires enabling dynamic visualization throughout all stages for analysis and computational steering of long complex processes. And it requires the incorporation of data analytics in various stages for estimating and managing both risk and value.

The Biospace

The desire to extract more value from growing data sets using increasingly complex algorithms can be seen in many other industries. Consider, for example, genomic medicine. It is currently economically feasible to generate a 1PB database of one million complete human genomes for about \$1B. Clearly traditional bioinformatics algorithms can be used to identify similarities and patterns between various individuals in the database. But there is much more value in combining bioinformatics with analytics applied to the medical histories of the individuals in the database to identify not only the patterns, but to correlate those patterns back to actual outcomes. Similarly, one can further increase the value of such datasets by extending this approach from genomics to proteomics and metabolomics. It is clear that a database of a million individuals would be just the start. As these databases grow, the data management and movement problems grow with them, again pushing the limits of today's systems and emphasizing the need for data centric system design.

Genomic Medicine is a specific example of a Biospace revolution that is underway where “omics”, Big Data analytics, modeling, and bioengineering will transform industries in agriculture and food, energy, environment, and natural resources, and chemical, pharmaceutical, and consumer products.

Financial Services

As a final example, the Financial Services industry currently deals with terabytes of new financial data daily, manages multi-petabyte databases, and must process hundreds of millions of tasks per day in under a millisecond each. These requirements plus data growth rates of 30% annually are driving the financial services industry to every larger, more capable, and more efficient HPC data centers. Additionally, the industry derives extensive value from monitoring worldwide news feeds to help inform and guide its decision-making. The challenge is to incorporate high fidelity, real time risk analytics that provide predictive actionable analysis combining asset portfolio data and external sources. Such an approach would provide value through improved risk management, better trading decisions, and enhanced regulatory compliance.

For the financial services industry, Big Data value is in finding the proverbial needle in a haystack. A company cannot know a priori what information will be important, so as many sources as possible must constantly be scoured. This drives the demand for growing data processing, analytics, and predictive stochastic modeling across numerous, disparate data sources -- both static and streaming. It is the growing aggregation of disparate data sources and corresponding different modeling techniques that drives the need in this industry for a more data centric system design to efficiently handle to processing of this growing data deluge.

Implications for Basic Science: Discovery through Modeling, Simulation and Analysis

Computational science – the use of advanced computing to simulate complex phenomena, both natural and human engineered – has become a complement to theory and experiment as a third method of scientific discovery. More recently, big data analytics has been called the “fourth paradigm,” allowing researchers and innovators to glean insights from unprecedented volumes of data produced by scientific instruments, complex systems, and human interaction.

This combination of advanced computing and data analytics will broadly and deeply impact science and engineering by: (1) enabling modeling and simulation of complex systems at a hitherto unattainable level of detail, (2) enhancing our ability to incorporate science-based analysis and simulation in engineering designs, and (3) allowing us to analyze and interpret large datasets generated by new, large scientific instruments, ubiquitous sensors, and simulations themselves. Beyond the scientific and engineering benefits, continuing development of advanced computing technology – both computation and data analytics – has deep and important benefits for U.S. national security and economic competitiveness.

The impact of next generation computing on science and engineering has been a subject of study, research, and scrutiny throughout the planning process and continued development of the DOE’s current exascale initiative, and its relationship to federal government interagency research and development efforts. Previous reports have summarized the multiple workshops, community input, and technical studies that have occurred over many years, beginning with context on computational science and big data. Current projects – including both the continued applications research in DOE and the work of the exascale Co-Design Centers - are refining and extending our understanding of these science and engineering impacts.

As noted in the 2010 Advanced Scientific Computing Advisory Committee (ASCAC) report, *The Opportunities and Challenges of Exascale Computing*,⁴ the most compelling impacts of the next generation computing initiative are those that are “transformational,” i.e., that will enable qualitatively new approaches and provide dramatic new insights, rather than simply incremental improvements in the fidelity and capability of current computational models. We will focus on such impacts here.

We note that there are additional impacts of the pervasive use and usability of next generation computing technologies that are significant but not transformational. Also, as with any new scientific instrumentation, there will likely be unexpected impacts that are transformational.

Three such transformational areas are delineated below:

As discussions of exascale computing began, the focus was on problems, such as modeling and simulation, where the large computational capability was most obviously needed. However, as noted above, it is crucial to address data intensive science, which is now an integral part of many fields. The complementary ASCAC report *Synergistic Challenges in Data-Intensive Science and Exascale Computing*⁵ studied how exascale computing and data intensive science interrelate and identified several areas of synergy between them. These findings are also reflected in the summary below.

⁴ http://science.energy.gov/~media/ascr/ascac/pdf/reports/Exascale_subcommittee_report.pdf

⁵ http://science.energy.gov/~media/ascr/ascac/pdf/reports/2013/ASCAC_Data_Intensive_Computing_report_final.pdf

Impact: Computational Scientific Discovery - enabling modeling and simulation of complex systems at a hitherto unattainable level of detail

- **Simulation of Materials in Extreme Environments:**
Will play a key role in solving many of today's most pressing problems, including producing clean energy, extending nuclear reactor lifetimes, and certifying the aging nuclear stockpile.
- **Simulation of Combustion in Turbulence:**
Enable current combustion research to make the critical transition from simple fuels at laboratory conditions to complex fuels in the high-pressure, turbulent environments associated with realistic engines and gas turbines for power generation. Combustion researchers will then be able to differentiate the properties of different fuels and capture their emissions characteristics at thermochemical conditions found in engines. This type of capability also addresses a critical need to advance the science base for the development of non-petroleum-based fuels.
- **Understanding Photovoltaic Materials:**
Will improve photovoltaic efficiency and lower cost for organic and inorganic materials. A photovoltaic material poses difficult challenges in the prediction of morphology, excited state phenomena, exciton relaxation, recombination and transport, and materials aging. The problems are exacerbated by the important role of materials defects, aging, and complex interface morphology.
- **Rational Design and Synthesis of Multifunctional Catalysts:**
Will help develop the fundamental understanding needed to design new multifunctional catalysts with unprecedented control over the transformation of complex feedstocks into useful, clean energy sources and high-value products. Computing with large-scale, high-throughput methods will play a central role because statistical mechanical sampling and free energies are fundamental concepts of this science
- **Astrophysics:**
Will include stellar modeling, galaxy formation, and collapse.
- **Computational Biology:**
Will allow seminal work on cell(s), organisms, and ecologies

Impact: Engineering Design and Optimization - enhancing our ability to incorporate science based analysis and simulation in engineering designs

- **Simulation of Advanced Reactors**
Enable an integrated simulation tool for simulating a new generation of advanced nuclear reactor designs. Without vastly improved modeling capabilities, the economic and safety characteristics of these and other novel systems will require tremendous time and monetary investments in full-scale testing facilities to assess their economic and safety characteristics.

- **Aerospace/airframes**
Will allow fully integrated, dynamic analysis of performance limits of gas turbine engines, next generation airframes, and launch / reentry vehicles.
- **Fusion**
Effectively model and control the flow of plasma and energy in a fusion reactor, scaling up to ITER-size.
- **Design for Resilience and Manufacturability**
Advanced manufacturing processes increasingly rely on predictive models of component wear and failure modes *in situ*. These combine structural dynamics, materials science, and environmental interaction. When combined with traditional (subtractive) manufacturing, as well as additive processes (3-D printing), computational models allow designers and manufacturers to reduce costs and improve customer experiences.
- **Biomass to Biofuels**
Enhance the understanding and production of biofuels for transportation and other bioproducts from biomass. The main challenge to overcome is the recalcitrance of biomass (cellulosic materials) to hydrolysis. Enable the design, from first principles, of enzymes and plants optimized for the conversion of biomass to biofuels to relieve our dependence on oil and for the production of other useful bioproducts.
- **Globally Optimized Accelerator Design**
Develop virtual accelerator modeling environment for the realistic, inclusive simulation of most relevant beam dynamic effects.

Impact: Data Analytic Discovery - allowing us to analyze and interpret large data sets generated by large scientific instruments, ubiquitous sensors, and simulations

- **Data Streaming and Accelerated Analysis for the Spallation Neutron Source (SNS) and other light sources**
Many of the technological advancements required in exascale computing are needed for the productive use of the data generated by the SNS. These advances span system architecture to advances in simulation and data analysis/visualization software. Explanatory and validated materials science simulation software optimized for time-to-solution is required in order to provide timely feedback during experiment. These improvements are non-trivial, requiring strong-scaling codes and corresponding scalable system architecture capable of providing time-to-solution improvements of up to 1000X. Advances in in-situ data processing, particularly in streaming data processing, will require lightweight, composable data analysis software optimized for use on next-generation systems.
- **Climate Science**
Develop validated models to enable understanding of the options for adapting to and mitigating climate change on regional space scales for an arbitrary range of emissions

scenarios. Fully integrate human dimensions components to allow exploration of socioeconomic consequences of adaptation and mitigation strategies. Quantify uncertainties regarding the deployment of adaptation and mitigation solutions.

- **Energy and Environment**

Understanding subsurface geophysics is key to environmentally friendly energy extraction and management. Correlation of data from new sensors and seismological instruments with geophysical models guides understanding of extraction locations and effects.

- **Instrumented Cities and Ecosystems**

Ubiquitous, inexpensive sensors now provide unprecedented levels of data for cities and engineered infrastructure. From electrical power grids through transportation systems to communication networks, insights from this rich data stream can help optimize designs and also reduce resource consumption.

Fostering and Maintaining a Balanced Ecosystem

To achieve a major advancement in very high performance computing capability requires advancing multiple, different technologies in a coherent and complimentary way, including hardware, software, and application algorithms. To be more specific, advancement is required in reducing power consumed by electronics, designing and implementing electronic parts and hardware architectures that deliver much higher processor and memory access rates, software systems that manage hardware resources to deliver useful computation in the presence of frequent hardware element failure, software development systems that support the development of application code, and new application algorithms that make cost effective use of the new memory and processor architectures. The challenge is not simply to design and build a faster processor/memory architecture. Without balanced progress on each of these dimensions, the desired computational capability will not be realized, and will not be cost-effective in advancing applications.

The Ecosystem

The committee views these multiple technologies that must be advanced in concert as defining an *ecosystem* because advancement in one technology must be made while taking into consideration the status of, and the route to advancement of, the others. This was necessary for the achievements of terascale and petascale computing systems as well. But the advancements needed are technically specific to today's technology requirements, including the major focus on power reduction. We take as the objective discussed in this report the goal of making a 100 fold to 1,000 fold improvement so that, in concert, the required technologies can attain 100 to 1,000 fold aggregated processor speed, and comparable cost effective movement of data.

Hardware

The greatest challenge in the hardware dimension will be a several hundred fold reduction in power consumption per operation as in today's petascale systems. Large petascale systems consume about 7 megawatts (MW) of power. DOE has budgeted roughly 20 MW of power for next generation systems.

This committee believes that a next generation of high performance computing advancement can be made with CMOS technology. As power is directly proportional to the square of the voltage supplied to the integrated circuits, reducing voltage reduces power consumption. For the next generation systems, further reduction in the voltage supplied to an integrated circuit can be made, however, the limit to voltage reduction is the minimum voltage needed to turn on a transistor, so there are clear limits to this lever. Further, as voltage gets closer to that threshold, integrated circuits behave more unreliably. Also, as transistor dimensions decrease, their performance characteristics are more variable. Variability and reduced power margins will cause circuits to fail.

Coping with stability and reliability of integrated circuits that will accompany lower voltages and smaller device dimensions is a major challenge. Circuit design can tolerate some variability and failure. The hardware architecture will likely need to offer an interface that informs the operating system about failures and permits the software to help manage errors that the hardware alone cannot detect and correct. Reliability management may rise to the level of application code, as the impact of a particular failure on a computation may only be understood in the context of the application algorithm being executed. New software techniques are needed that allow an application to adjust to failures localized within a computation, and to continue to make progress without results being contaminated by effects of failed elements.

Systems

Evolving workflows require the integration of widely diverse functionality and it will be increasingly critical to create a more flexible, modular, data-centric system design capable of efficiently handling data motion that will be highly variable in size, access pattern, and temporal behavior. Without close attention to these “data” issues, such systems will be hobbled by numerous data bottlenecks and will fail to achieve their promised goals. Similarly, usability, reliability, and productivity must be addressed over the entire system, not just at the component level.

Software

Effective management of hardware resources requires the development of an operating system tailored to the specific system architecture. The operating system schedules selected resources (memory and processors) for the one or more concurrently running applications. It must manage a hierarchy of memories with different performance characteristics as well as input/output devices and network connections. New algorithms to map data onto memories with predictable/known access patterns by processors are needed. Dynamic remapping may enhance performance. It is likely that the operating system, and possibly language compilers, will participate in energy management, as well as managing routinely failing hardware, and possibly software, elements.

Ideally, the operating system will monitor its own health and performance, reporting in terms that permit administrators to incrementally tune the operation of the system to attain higher performance and higher reliability. Building such an operating system for a new architecture will be challenging. It is unlikely that extant operating system software can be re-purposed to manage the resources of a new and novel architecture. To extract the potential speed from a novel system, the operating system software needs to be well matched to the hardware architecture in order to exploit its capabilities.

For application developers to cost-effectively program a system with 100 to 1,000 times more parallelism, a suite of software development and execution support tools will be needed. As with the operating system, the software tools that assist in application software development need to be built to exquisitely exploit the capabilities of the hardware.

These tools include programming languages in which the programmer can express concepts related to power consumption and the handling of failures. Language compilers may generate code that deliberately modulates power for different circuits. Compilers may need to generate code that supports adjustment in response to failures related to the execution of sections of code, either as directed by the programmer or in a background/automatic mode. Orchestrating millions if not billions of processor elements as well as the related data to memory mapping is challenging. Compilers will need to make it simple to instrument code at varying scales in order to gather, and meaningfully aggregate, performance data so that application software can be tuned to increase performance.

New paradigms for communicating information from one locale in a computation to others will be needed to prevent such communication from retarding processor cycle usage until the communication is accomplished.

Other software development tools include those that allow developers to be able to rigorously test code at scaling levels that span many orders of magnitude. In addition, there will be a need for simulators, test harnesses, test case generators, and performance analysis tools. While some adaptation of existing tools might suffice, to perform well such tools must be well matched to a novel architecture.

Application Algorithms

The majority of the applications of highest priority to the Department of Energy today are the same, or variants of those that were high priority in the past. And the majority of those seek to produce better understanding of physical phenomena, such as combustion, fluid flow, and nuclear activity as well as the interactions of materials at density and pressure extremes.

Until the architecture for the 100 to 1,000 fold more powerful systems is defined, it will be difficult to determine the extent to which extant codes can be re-purposed to that novel architecture. However, it is safe to assume that entirely new and innovative application algorithms will need to be invented, again to make cost-effective use of the more powerful system. There is a old adage that observes that whenever hardware performance is increased by an order of magnitude, new resource management algorithms need to be devised. These new algorithms will be both in the operating system and in applications.

Data Analytics and Discovery Tools

A major challenge for the next generations of computing will be analyzing and interpreting the massive data sets being generated by large scientific instruments, ubiquitous sensors, massive simulations, etc. There are a number of key challenges to be explored in developing the underlying science of these data analytics⁶ including:

- Data gleaned across multiple scales or with very large parameter spaces,

⁶ Hendler, J. and Fox, P. The Science of Data Science, *Big Data*, 2(2), 2014.

- Sparse systems with incomplete data or where systems being modeled may be highly non-linear, heterogeneous, or stiff, or
- Problems where we need to be able to do uncertainty quantification in open worlds, or use uncertain information (for example that processed from unstructured data).

One of the key technologies used in handling these problems today is the use of data-mining and machine-learning algorithms that either try to find non-obvious correlations across complex cohorts of data or which attempt to perform abductive processes to find parameter sets that can best provide predictions of future performance.

To date, most of the data analytic models that scale to very large datasets and/or datasets of high dimensionality, have been based on “support vector machines” (SVMs). Despite the name, SVMs are not actually based on a particular machine architecture, but rather they are a class of supervised learning algorithms that are particularly useful for classification and regression problems (particularly as extended for non-linear classification problems).

The modeling and simulation tools described earlier in this report were developed with high-performance computing in mind, and have been around long enough that significant libraries of software mapped to HPC architectures and languages have been developed. However, for machine learning tools, such as SVMs, robust libraries do not yet exist. Where these have been written, for example, they are heavily used within the Web-search and Web-mining industries. They have typically been done for server farm clusters, rather than for specialized architectures, sacrificing a level of performance for the advantages of horizontal scaling. As we move to new data scales, and to the sorts of data science challenges discussed above, the need for significantly higher performing technologies, driven by the much larger and more complex datasets of modern science and engineering, are needed

Agent modeling tools

One use of HPC systems to date has been in the area of “agent-based modeling”. These systems have primarily been used for two purposes. First, the technique has been used for the modeling of large numbers of similar entities responding to an environment (such as schooling of fish or the movement of invasive species into an environment). Such systems also have been used to model large numbers of humans reacting to events, such as escaping from a building in a disaster. But they have been based on idealized behaviors, assuming all agents act similarly in similar conditions. Second, these systems have been used to model economic behaviors, such as markets based on “rational” decision agents, or modeling of mechanisms for bidding or other such economic behaviors.

There also has been considerable work on the development of “intelligent agents” including significant DARPA investment in the area in the late 1990s. These systems have primarily been used to model small or moderate numbers of decision-making agents that can make complex decisions using logical processes⁷. Such systems allow for the modeling of more complex agents, including human decision makers that are motivated by beliefs and desires.

⁷ Helsing, A.; Thome, M.; Wright, T., Cougaar: A Scalable, Distributed Multi-Agent Architecture, *2004 IEEE International Conference on Systems, Man and Cybernetics (v2)*, October, 2004

Recent work has explored whether discrete event simulators, implemented on high performance computers⁸ can scale agent systems to much larger challenges. This is motivated by a growing need to model problems that include large numbers of humans where we cannot assume fully rational behaviors -- e.g. mathematically bounded resources or complex belief systems that cause seemingly irrational decisions. For example, in cases where actual disasters have been studied, the unusual behaviors of small numbers of agents acting contrary to global best interests have been shown to have significant impacts causing large divergences between the modeled and observed behaviors⁹.

Another use of such scalable agent modeling is to understand and predict the impacts of incentive systems on large-scale populations. For example, as energy providers explore smart grid technologies, an assumption is made that people will reduce energy consumption based on economic incentives. While it is clear this works to some degree, modeling it with any fidelity is extremely hard. Looking at the energy consumption of, for example, even a medium-sized city, would require modeling the behaviors of tens of thousands of consumers under complex conditions. Predicting how many would change a behavior based on what level of economic reward requires modeling tools (at high scales) not yet available.

Cognitive Computing

Another rapidly growing segment of the supercomputing ecosystem is the use of multiprocessing systems to process “unstructured” data, that is, the myriad of textual information available in machine-readable form or easily scanned in. The current generation of such processing is generally using cluster-based systems with various map-reduce and machine-learning algorithms. These systems are the backbone of many large Internet and Web providers, such as search engines, social networks, and e-commerce sites. The scale of data in these applications, especially in the large Web companies, is already into the petabytes per month range, and query results against the data are found in near real time (under 100 milliseconds is targeted). These systems currently are well-served by large server farms and commodity computers, although there is concern about the ability to continue to handle web growth with commodity scaling. New algorithms aimed at replacing map reduce with data flow algorithms, more similar to those used in high-performance supercomputing, are being explored.¹⁰

At the same time, new kinds of systems for handling text in a deeper way also are increasingly being explored. Referred to as cognitive computing systems, such systems perform deeper analysis of the textual data with a goal of creating important new applications that go beyond search and retrieval. The first such system to reach national prominence was the Watson system developed by IBM¹¹. In 2011, the

⁸ P. Barnes, C. D. Carothers, D. R. Jefferson, J. M. LaPre, “Warp Speed: Executing Time Warp on 1,966,080 Cores”, In Proceedings of the 2013 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation (PADS), Montreal, Canada, May 2013.

⁹ Helton, William S., Simon Kemp, and Darren Walton. "Individual Differences in Movements in Response to Natural Disasters Tsunami and Earthquake Case Studies." *Proceedings of the Human Factors and Ergonomics Society Annual Meeting*. Vol. 57. No. 1. SAGE Publications, 2013

¹⁰ cf. <http://www.informationweek.com/cloud/software-as-a-service/google-i-o-hello-dataflow-goodbye-mapreduce/d/d-id/1278917>

¹¹ <http://www.ibm.com/smarterplanet/us/en/ibmwatson/>

system was featured on the game show Jeopardy!TM where, without an internet connection, it was able to beat the two best human players of the televised question-answering game show. Since then, IBM has been working on the use of Watson in a number of areas, particularly the healthcare sector, with a primary focus on helping doctors with cancer diagnosis and treatment. Other applications are exploring how to couple such systems with more structured relational or graph data, with simulation and analytic systems, and with new capabilities for going beyond text to explore images and other non-textual data resources.

As cognitive computing systems are new and improving rapidly, it is hard to predict the exact architectural configurations such systems will use. End-users will likely access these systems through cloud-based application “fabrics” (sets of Application Program Interfaces of connected functionality supported through cloud computing resources). However, open questions remain as to what the backend systems for these new computing technologies will be. One thing we can predict for certain is that the increasing use of cognitive computing will call for increasing access to large datasets and/or data streams, especially as new applications are developed for interacting with the information generated by the growing “Internet of Things.” (For example, the recently announced “analysis in motion” program at Pacific Northwest National Laboratory is exploring how to use these emerging cognitive technologies for large scale DOE-linked experimentation on streaming data from next-generation scientific instruments.¹²) These systems will likely first be fielded on server-based clusters and special purpose hardware, but increasingly there will be a need for the government to be able to use specialized supercomputing systems more flexibly to interact with this new generation of application capabilities and specialized processors.

Next Generation Neural Network Architectures

This latter work crosses over with another approach to cognitive computing which grows out of earlier work in the neural network community. These systems use a combination of mathematical machine learning techniques and new architectures (usually referred to as neurosynaptic processors¹³) for processing data that includes scanned texts, images and videos, and streaming data. These systems can use massive amounts of processing, and it is projected that the supercomputing ecosystem will soon include hybrid machines that couple these advanced neural processors with other forms of high performance computing.

How neurosynaptic processors ultimately will impact the integration of perceptual information (particularly multimedia and video) with the information being mined from unstructured resources, social media, etc. is an emergent area for exploitation.

Experiments already are being conducted on integrating these processors with high-performance computing under the support of DARPA’s SyNAPSE program¹⁴. This includes use of HPC in the design of new generation of neural hardware, and the use of HP neural systems in large scale perceptual processing, etc.

¹² http://readthis.pnl.gov/marketsource/readthis/B3401_not_print_quality.pdf

¹³ *cf.* <http://www.research.ibm.com/cognitive-computing/neurosynaptic-chips.shtml>

¹⁴ http://www.darpa.mil/Our_Work/DSO/Programs/Systems_of_Neuromorphic_Adaptive_Plastic_Scalable_Electronics_%28SYNAPSE%29.aspx

Partnerships and Markets

Any major federal government investment in information technology directly affects (and is affected by) the computing industry – its current capability, future anticipated markets, the industry’s research and development (R&D) technology roadmaps, and investment plans to penetrate these markets. Any federal program to materially advance general purpose computation must start with the most promising technology of today, with goals for which there are some promising engineering approaches, and for which the industry can justify co-investment in what will be a major R&D program.

High-end computation for the U.S. government use alone is simply not large enough. Federal investment can be justified, if it leads to results that can be appropriated to develop vastly larger commercial markets, and for which the benefit will inure to many players, not just one corporation. The extremely high performance computing market alone is not sufficient to attract the computing industry’s main focus.

Historically, when major advances in high performance computation systems have been made, industry has refined the technology advances, made the technology more cost-effective to manufacture and package, and incorporated it into a multitude of products for a range of customers. Of particular importance are customers – both commercial and government, who are *fast followers*. They are customers who will buy highly capable computing systems that are somewhat less capable than the government’s most extreme systems. A robust fast follower market is critical to business investment decisions related to the subject of this report.

In past decades, government investment matured the hardware, the operating systems, the development environment, and selected application software-- for gigascale, terascale and petascale systems. Without government investment, these past generation systems would not have been built on the aggressive timelines on which they were developed. Fast follower customers materially contribute to the maturation of the software tools and algorithms discussed earlier. Today, systems capable of delivering some tens of petaflops can be found in many industrial sectors, such as pharmaceutical, aircraft, automobile, movie, and petroleum, as well in government agencies other than the Department of Energy (e.g. Department of the Army). These are examples of the fast follower market. We note that in some cases, a corporation does not publicly divulge its computational capability for proprietary reasons.

A number of both government and commercial fast followers naturally have applications that align well with those of the Department of Energy because many engineered products and services are improved only when the underlying physical processes are better understood. For example, a company that builds engines – for aircraft or for automobiles – must understand the process of combustion for their particular engines, if they are to optimize and extend performance. Oil and gas companies model geology to better understand how to locate and extract petroleum. The Navy has a critical need to predict weather, especially over the oceans.

Such fast follower users will adopt the operating system and development tools, and possibly some application codes developed with government investment, but only if the new architecture is well suited to their applications – as it has been in past generations.

Path Forward

Our recommendations are focused on three time frames:

1. The greater Petascale timeframe: roughly the next five years, and characterized by systems in the many tens to hundreds of petaflops and requiring up to a combination of 5-10 petabytes of addressable and buffer memory and over 100 petabytes of storage. With the CORAL program, the DOE has initiated a well-defined program with industry partners for development and deployment of hardware and software systems at the high-end of this spectrum, along with the provisions for application development. This should continue with a close cooperation between the vendors chosen and DOE labs.
2. The Exascale time frame: covering the next five to ten years, and characterized by systems in the 100's of petaflops to tens of exaflops, requiring tens of petabytes of memory and perhaps an exabyte of storage. Although some DOE and other government funding for needed technologies has begun to flow, government programs for addressing exascale system development and deployment are still in the formative stages. The CORAL program has stressed the desirability that systems developed in the petascale time frame be designed with the goal of leading naturally to exascale.
3. Beyond Exascale: The situation here is more uncertain. The extent to which CMOS technology and current architectural thinking can scale beyond a few tens or hundreds of exaflops is unclear, as are the likely candidates for replacing them. This is an area that will need additional attention and long term thinking over the next several years if we are going to sustain the growth and promise of HPC. This requires careful attention to overall HPC ecosystem needs.

Acceleration of Advanced High Performance Computing Systems

For the greater petascale time frame: Continued execution and funding of the CORAL development and procurement program, resulting in systems of about 200 petaflops performance in a data centric architectural context, with attendant focus on power efficiency, reliability, and productive usability. The resulting systems should achieve high performance on both throughput and scalable science applications that comprise the laboratories' actual production workloads; they should represent a next step in the evolution towards exascale; and they should embody an architecture that can naturally and easily be reconfigured and optimized to meet specific workflow and use requirements of individual laboratories. They should provide a flexible degree of heterogeneity to enable today's multi-scale applications to leverage the strengths of different computational elements (e.g. traditional CPUs vs GPUs), and they should provide the system attributes and distributed intelligence required to manage the huge data volumes of future workloads.

As planned, this development effort should unfold in close cooperation between the vendors chosen to develop the CORAL systems and the receiving laboratories. Receiving laboratories' personnel should be able to provide continuous feedback on and input into the design of both hardware and software, enabling productive co-design insuring design trade-offs are based on cost, performance, and mission impact. Emphasis should be placed on evolving a programming model that enables productive exploitation of the systems' advanced, data oriented architectures while preserving as much as possible existing investments in application software. Joint participation of computational scientists and domain experts from

government laboratories, academia, and system suppliers in application porting and enabling centers should begin early in the program to help insure that application workflows can fully utilize the new systems when they become available.

For the Exascale time frame: A timely commitment to fund and execute a full-scale exascale development program that is an integral extension of the CORAL program is required. The architecture and implementation attributes of CORAL must be viewed as part of a **Roadmap to Exascale**, so that investments in programming models, languages, and software development made by the CORAL community will be preserved for the future and that new optimized code will be positioned to take advantage of exascale features. Such a commitment will enable systems suppliers to provide to their customers a strategic roadmap that preserves their investments in software and applications; thus providing a predictable roadmap for both technology delivery and suppliers business models. Ideally, CORAL and exascale should be viewed as an integrated, long-term government/industry program.

Exascale systems will require advances in many areas, notable among these are circuit and optical technologies, photonics, power management and delivery, system cooling technologies, 3-D packaging, and non-volatile memory technology. As systems become larger and more capable, and workloads become more complex and sophisticated, continued evolution of system and applications development software will be needed. Because of the long lead times required to mature these new technologies and software capabilities, joint research and development programs with academia and industry should be nurtured and encouraged. Since most of these required technology and software advances will run far ahead of commodity market demand, the government programs to adequately fund their development must be in place. DOE's Fast Forward and Design Forward programs are steps in the right direction but need to be expanded. They should be periodically reviewed to make sure that the necessary technologies will be delivered in a timely fashion, and appropriate extensions and follow-up programs should be put in place as needed.

Beyond Exascale: A long-term, funded, research program with leading university and industrial research teams to explore and accelerate the potential applications of carbon nanotubes, cognitive computing, quantum computing, and other future technologies. Of these, carbon nanotube transistors are one technology that warrants early investments. Silicon technology scaling is expected to last another 3-4 technology generation nodes and silicon, as a transistor channel material, may need replacement. Carbon nanotube transistor technology has shown promise in this regime with several experimental demonstrations over the past few years. System simulations indicate that carbon nanotubes would outperform scaled silicon transistors by a factor of 3-5 in power/performance tradeoffs. The progress has been sufficiently promising to warrant a strong government funded project to develop a carbon nanotube based post-Si CMOS microprocessor technology. Other technologies, as well, are setting records on interesting development trajectories (e.g. 74K working Josephson Junction on 8-layer niobium superconductor chip).

A Robust Ecosystem

Key elements of a government program to help build this ecosystem are investments in partnerships with academia to help expand the volume of open source software and to develop algorithms for the new generations of high end HPC systems; support and encouragement of university programs in computational science; an effort to establish standards designed to insure software portability across

platforms; and the creation of monetary or other incentives to ISVs to modernize their software. Important to this ecosystem is the use of, and partnership with, DOE FFRDCs, in particular the NNSA Laboratories, and IARPA's C3 program, which now has foundry capability

The three principal sources feeding the contemporary software ecosystem for HPC are ISVs, the open source community, and custom software developed in-house by an HPC user. The state of this ecosystem is distressed: few ISV codes are scaled significantly either in core or node count, nor has there been much ISV activity to exploit newer architectural features, such as accelerators; there has been uneven activity within the open source community with most of the uptake occurring in national laboratories and universities but very little in commercial or industrial settings; and the creation of bespoke codes or tools have only parochial value to the institution of origin. The evolution of more ambitious system architectures incorporating data analytic attributes and capabilities is sure to introduce additional requirements and demands on the ecosystem and be a driver for bringing the Big Data, Analytics, and HPC communities together. As it stands now, these requirements and demands are not going to be well served by the community currently contributing software tools and applications to the existing HPC ecosystem unless material change takes place.

ISVs

HPC ISVs are commercial enterprises which have been evolving from the late 1980s to present day. In many cases these firms have achieved a dominant position in the market segments they serve and barriers to entry have been erected inhibiting innovative entrants. Most ISVs spend the bulk of the resource on serving their install base and, if the barriers to entry are sufficiently high, have no motivation to innovate around new hardware system features. This is an open market view characteristic of the developed economies of Europe and North America; China has no extant ISV community and, if motivated for strategic reasons to develop an indigenous capability independent of Western software companies, could make investments sufficient to create new ISVs. If so inclined, it is likely that the new software forthcoming would map to modern architectures, which would present the Chinese economy with a more modern software ecosystem than exists in the West: there is no install base to defend so software can be developed mapping to the attributes of scaled, accelerator based systems. Remedies to this situation are possible with concerted and focused effort. First, there must be investment in the development of new algorithms with explicit consideration for how modern system attributes can be exploited; many of the current ISVs have software based on algorithms developed more than thirty years ago when concepts of scaling and accelerators were inconceivable. Second, there should be direct investment made available to existing ISVs to compensate them for modernizing their software. With effort their codes can be made to scale and exploit modern system features; the investment is required to overcome the absence of market based incentives to modernize. Third, there should be effort focused on establishing software standards designed to insure portability across different system platforms or investments to facilitate code porting from the least capable but least expensive commodity systems to the more innovative systems likely to spring forth as simulation, modeling, and data analytics coalesce.

Open Source Software

Software development has received relatively small investment compared to hardware development in the HPC framework. An important area in this regard is open source software. Open source software has made its way into use within the HPC commercial community, where great emphasis is placed on serviceability and security. In particular, Linux; MPICH and OpenMPI; Lustre; Netlib (BLAS,

LAPACK), PETSc, and other math solver libraries; the Clang/LLVM compilers at the core of Cray, Apple, NVIDIA, and Intel compiler suites; and even applications, such as NWChem, LAMMPS, OpenFOAM, etc. are used in industry, and in some cases, are pretty fundamental to the HPC community, including industry.

There is an excellent track record in data analytics recently, with map/reduce as a notable example. Serviceability and security are lesser issues for universities or non-weapons national laboratories, but these entities represent no more than about 10%-15% of all HPC usage. Of course, one cannot “force” the adoption of open source software, but one should also not plan on it being a panacea to any ecosystem shortcoming. A focused investment effort within universities could expand the volume of open source software, and increase the chances that more of the software output could become commercialized. While the most significant consumption of open source software is in China, and it is also the case that the Chinese are rare contributors to open source software, there nonetheless should be an increase in open source software investment. This is important for two reasons: (1) Open source software development is critical to overall progress in software advances, and (2) Since much of open source software development occurs in universities, it is a key educational and research tool as part of a healthy HPC ecosystem.

Custom Codes

Custom codes are created within the confines of the institution having needs that are unmet by ISVs and open source. Depending on the nature of the institution, systematic investment in this activity could prove strategically important to the interests of that institution. We see many commercial firms giving serious thought and action to bringing software development back “inside”, either because it is a function that the ISVs or open source community are unwilling or unable to supply or because the cost of acquiring ISV software is too high. These considerations are based on standard cost-return tradeoffs but also need to factor in the availability of in-house skills to create the software to make this transition. Investment within universities to develop more computational science skills could help institutions acquire sufficient in-house capability to embark on more innovative exploitation of HPC. Beyond nurturing and investing in ISV, human capital development, and software/algorithmic advances, the overall ecosystem for HPC must take account of local environmental concerns and power systems issues. This includes the robustness of such systems and restrictions that local/regional governments and utilities may impose.

Partnerships

Today’s most advanced HPC systems are extremely complex to architect, program, and deploy, and the cost of solution development for even the largest companies in the world is a very significant barrier to the deployment of very high-end HPC solutions for industrial and business use. Nevertheless, there is real opportunity in building on ramps to this technology through government/academic/industry partnerships, with a focus on and investment in rapid transfer of algorithmic and application innovation/impact from government to industry, at scale, thus creating real competitive advantage for U.S. industry, while lowering the barriers to deployment of high-end HPC systems.

DOE's successful Innovative and Novel Computational Impact on Theory and Experiment (INCITE)¹⁵ program is an example of such a partnership. This program helped to familiarize scientists and engineers from industry and academia with the capabilities and potential of some of the most powerful HPC systems available, while at the same time enabling them to do the modeling and simulation critical to their work.

The establishment of HPC innovation centers along the lines of the Deep Computing Solutions center initiated by LLNL and IBM is another way to foster the required partnership. The focus of an HPC innovation center should be the development and deployment of HPC solutions for industrial, business, and research use. Such centers should possess the most advanced HPC systems available, along with a highly skilled staff of computational scientists and domain experts drawn from the National Labs, academia, and industry. Industrial partners would participate in these centers in order to gain access to the most advanced computing facilities, as well as to the expertise embodied in the center's staff and to leverage that expertise to develop the software, tools, infrastructure, etc., needed to exploit state of the art high end HPC systems. In return, the industrial partners would provide financial support and contribute domain expertise in the form of participation of their own personnel. In addition to supporting individual partners, the center as a whole would be structured to deliver a common shared scientific and computational infrastructure (tools, middleware, etc.) in support of the separate HPC business and research activities of its partners. It would also necessarily participate in the creation of the general software ecosystem for advanced HPC systems. Thus, an important feature of the center's operation should be that the systems and a significant fraction of the expertise required could be shared in a collaborative (noncompetitive) manner across many different industrial sectors.

Clearly, DOE's National Laboratories with their abundant and extraordinary scientific and computing research and development skills could have a major role lowering the barriers to the use of high-end HPC by U.S. industry through partnerships. However, for that to happen, contractual barriers to industry partnerships also must be lowered. For example, an industrial partner's rights to intellectual property developed as a result of the partnership must be appropriately respected, and risks associated with government indemnification policies must be bounded and manageable. Impediments to these types of partnerships¹⁶ need to be identified and dealt with.

¹⁵ <http://www.doeleadershipcomputing.org/incite-program/>

¹⁶ See, for Example, "Turning the Page: Reimagining the National Labs in the 21st Century Innovation Economy", a non-partisan white paper published by The Information Technology and Innovation Foundation, The Center for American Progress, and The Heritage Foundation

References

- Computational Science: Ensuring America's Competitiveness, President's Information Technology Advisory Committee, June 2005 (PITAC Report)
- The Fourth Paradigm: Data-Intensive Scientific Discovery, Microsoft Research, 2009
- Exascale Workshop Panel Meeting Report, January 2010 (Trivelpiece Report)
- "The Opportunities and Challenges of Exascale Computing", ASCAC , November 2010 (Rosner Report)
- "Synergistic Challenges in Data-Intensive Science and Exascale Computing", ASCAC March 2013 (Sarkar Report)
- "Ten technical approaches to address the challenges of Exascale computing", ASCAC March 2014 (Lukas Report)

Appendix 1: Task Force and Study Participants

Participants:

SEAB Members

Shirley Ann Jackson, Co-Chair, Rensselaer Polytechnic Institute
Michael McQuade, Co-Chair, United Technologies Corporation
Ram Shenoy, ConocoPhillips
Steve Koonin, NYU Center for Urban Science and Progress

External Participants

Roscoe Giles, Boston University
Jim Hendler, Rensselaer Polytechnic Institute
Peter Highnam, IARPA
Anita Jones, University of Virginia
John Kelly, IBM
Craig Mundie, Microsoft
Thomas Ohki, Raytheon BBN Technologies
Dan Reed, University of Iowa
Kord Smith, Massachusetts Institute of Technology
John Tracy, Boeing (Ted Colbert)

Process:

Co-chairs meeting with DOE team (DC)	October 8, 2013
Preliminary Task Force planning call	October 18, 2013
Fact finding meeting (DC)	November 8, 2013
Secure briefing for co-chairs (LLNL)	December 3, 2013
Task force meeting (DC)	January 10, 2014
Task force meeting (CUSP)	February 11, 2014
Task force meeting (DC)	March 11, 2014

Plus – three informal meetings with NSCI


Appendix 2: Charge to the Task Force



The Secretary of Energy
Washington, DC 20585

December 20, 2013

MEMORANDUM FOR THE CO-CHAIRS
SECRETARY OF ENERGY ADVISORY BOARD

FROM: ERNEST J. MONIZ 
SUBJECT: Establishing a Next Generation High Performance Computing Task Force

I request that you form a Secretary of Energy Advisory Board (SEAB) Task Force composed of SEAB members and independent experts to review the mission and national capabilities related to next generation high performance computing. The Task Force will examine the challenge problems and opportunities that drive the need for next generation high performance computing, as well as the advances and necessary steps to create and execute a successful path that will deliver next generation computational performance. The Task Force report should include recommendations on whether and to what degree the U.S. Government should lead and accelerate the development of next generation high performance computing applications and systems.

Purpose of the Task Force: The SEAB Next Generation High Performance Computing Task Force will examine and report on the following:

- The justification for an exascale computing capability initiative.
 - DOE missions
 - Fundamental research opportunities
 - Broader societal benefits from an open, non-classified exascale program and potential market barriers inhibiting private development of exascale computing
- Related basic research necessary to enable next generation high performance computing (e.g. mathematics, computer science, etc., including quantum and superconducting computing).
- The current state of technology and plans for an exascale program in the Department of Energy and other federal agencies.
- Role of the Department of Energy in leading the development of exascale computing – including its involvement and collaboration with industry, universities and other government agencies on high performance computing.
- Implications of data centric computing for exascale computing.

Designated Federal Official: Amy Bodette, Deputy Director, Office of Secretarial Boards and Councils

Schedule: The Task Force will submit a report by June 2014 and make a presentation at SEAB's June meeting.



Printed with soy ink on recycled paper