Secretary of Energy Advisory Board

Follow Up Response from the Task Force on High Performance Computing

November 30, 2014

On August 18, 2014, the Secretary of Energy Advisory Board approved the report from the Task Force on High Performance Computing, in response to the Secretary's December 20, 2013 charge. The report[1] delivered six 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 Independent Software Vendors (ISVs), 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.

Upon delivery of this report, the Secretary requested follow up from the Task Force on three topics:

1. The SEAB Task Force on High Performance Computing identified costs for a DOE exascale and beyond program and certain technology and engineering developments that would be met. Can

[1] Report of the Task Force on Next Generation High Performance Computing, August 18, 2014, http://energy.gov/sites/prod/files/2014/10/f18/SEAB%20HPC%20Task%20Force%20Final%20Report.pdf

the Task Force provide an additional level of granularity on the allocations against major technology areas and their timing that regulate success in getting to the 1-10 exascale range in a decadal time frame?
2. The SEAB Task Force on High Performance Computing recognized that barriers exist to broadening U.S. industrial adoption of high-end HPC. The Deep Computing Solutions Center (DCSC) represents an initial DOE laboratory effort to address this problem. Could the Task Force provide suggestions for what the Department could undertake to expand industrial high-end HPC use? This could include scaling up DCSC-like approaches or entirely different ways of using DOE laboratory expertise towards the stated goal.
3. Finally, any further thoughts on how a beyond-exascale research program (superconducting, quantum, neuromorphic…) might be structured would be appreciated.

I. The Exascale Investment

Two components go into high performance computing spending in the DOE. *R&D* spending includes the science and engineering research and development of future systems along with the programs to use such machines across the complex. Current annual baseline R&D spending on high performance computing is approximately $1.1B, split almost exactly in half between the Office of Science ($540M) and NNSA ($560M). Of this amount, 2015 base line spending on exascale computing is approximately $150M (SC: $90M, NNSA: $60M.)

Acquisition includes the periodic costs to acquire new leadership machines. In the Coral program three leadership machines in the 50-200 petaflop range will be delivered in roughly 2018 at a cost of approximately $160M per machine. This compares to the approximately $120M per machine acquisition cost of the most recent machines in the complex (Titan and Sequoia).

We envision a decade long program to mature technology, design and engineer systems, and deploy partial exascale prototypes in the 2021-2022 timeframe, leading to the delivery of a full machine in the 1-10 exaflop range with at least 1 exabyte of memory by 2025. Starting at $150M in 2015, proposed funding for the DOE's exascale program in the 2016 budget request is $240M. We expect program funding to peak in the 2018-2023 time frame, consistent with the delivery of three exascale prototypes (a few rack systems) to test technology and to be used for operating system and application software development.

We expect an exascale research and development program (i.e. program spend not including acquisition costs) will most likely be executed with three primary work streams:

1. Basic research aimed at maturing to Technology Readiness Level (TRL) 6 the technologies needed to commence an exascale systems design. As noted in our original report, this will include the broad national lab, academic and private sector co-design centers aimed at providing building block solutions to resolve:
    a. Memory and storage architectures, including the optimized system configurations that balance traditional highly parallel numerical processing with the need for large scale distributed data access;
    b. Energy budgets; and
    c. Architectures at the chip, subsystem and system level that are robust to faults and transient across a wide range of time scales.

This work stream begins immediately with the 2015 funding increase with the goal to provide technologies for initial system design commencing in 2017. We expect this work stream will account for roughly 20% of total program spend, with primary spending between 2015 and 2018.

2. System prototype design (2017-2019), prototype build (2019- 2022), and final system design for delivery in 2024-25. We expect this activity to be primarily undertaken by two primary partners through the prototype phase with a determination in the 2022-2024 time frame as to whether both designs will be carried into final systems or not. We expect this work stream to account for roughly 30% of total program spending, with partner selection in the 2016 time frame.
3. Systems technology development – software and operating systems, programming environments, upgrades and resetting of existing codes, and transformation to new code bases that take advantage of the chosen architecture. We estimate that this broad category, accomplished primarily at the national lab and academic user communities, along with the connection to the vendors selected to execute the hardware program, will account for roughly 50% of total program spend, spread roughly evenly across the entire program duration.

Figure 1 below shows an illustrative analysis of how the spending for such a program might progress.

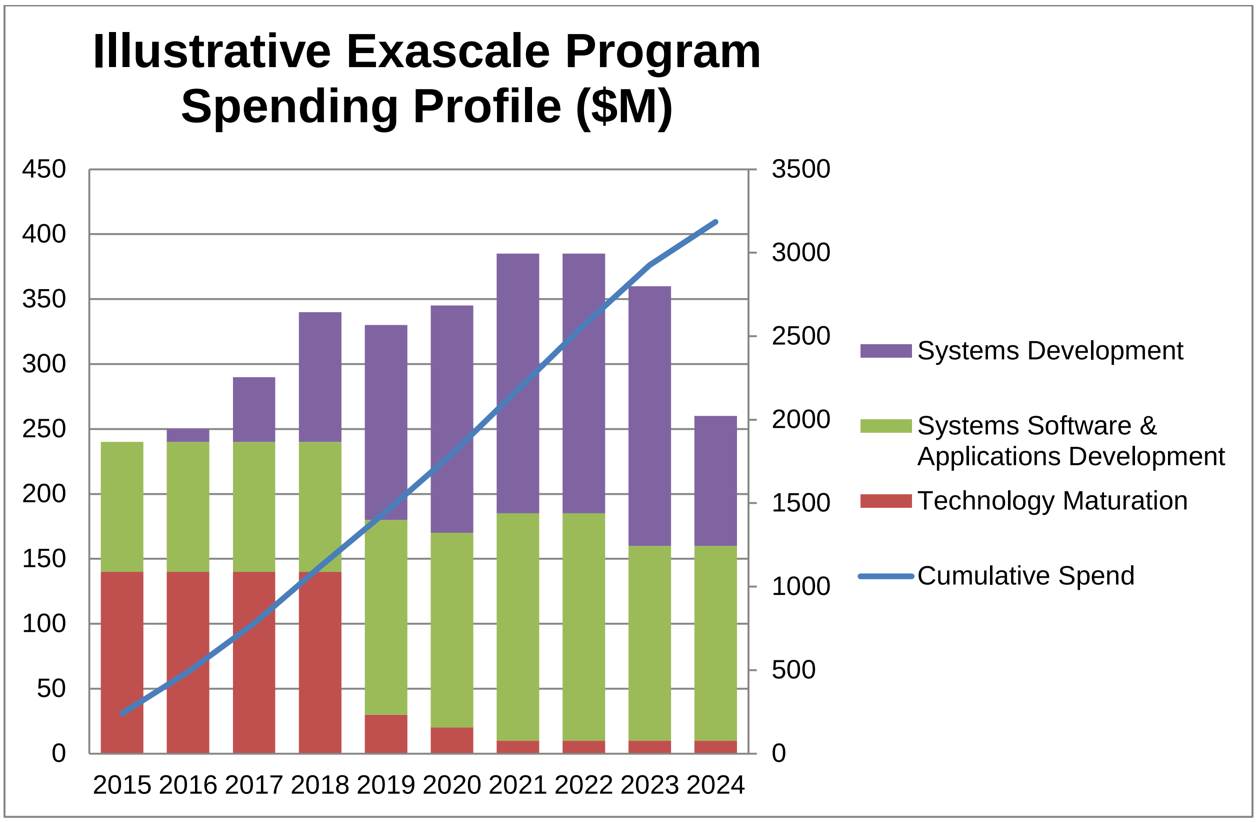

Figure 1. Example exascale program spend leading to final exascale machine in 2024-2025.

The final output of the program is a machine in the 1-10 exaflop performance range. Projecting from the baseline of current machine acquisition costs and the likely costs of the Coral machines, we estimate this machine to be in the $200-250M range, a $90M increment over what will then be the most recent generation machine.

II. Expanding Industrial Use of High-End HPC

Partnerships between the DOE and industrial users of high-end and leadership computing systems in the national laboratory system have a long and productive history. Programs such as the Innovative and Novel Computational Impact on Theory and Experiment (INCITE)[2] and the Deep Computing Solutions Center[3], located within LLNL's HPC Innovation Center, aim to expand the user base of HPC by providing both machine cycles and scientific expertise to the user community through competitive proposal awards.

In thinking about the expansion of the use of high-end HPC by the U.S. industrial base it is worthwhile to consider two, somewhat arbitrary classes of users. One group consists of large scale corporations, for example aerospace or oil and gas exploration firms. For such companies the use of large scale computational modeling, simulations, and data analysis is a standard part of daily engineering and scientific exploration and does not require "proof of concept" to justify its use. These companies usually invest in internal systems typically performing a generation (several orders of magnitude) below current leadership machines[4]. Drivers for this use, as detailed in our original report, include ever-more complex simulations to reduce design-prototype cycles or enabling shorter time to complete analysis of ever-more complex data sets from exploration. Often these users are more interested in running a large number of simulations and analyses with a productive workflow that allows full exploration of some relevant trade-space than they are performing one "hero" calculation. Thus workflow and data management, scheduling of resources, and reliability of machines in daily use are important.

As noted, this class of users does not, in general need to prove that the use of high-end HPC provides a business advantage. They do, however, benefit from the DOE leadership computing capability and support programs by:

- Being able to access leadership machines to develop new codes and new algorithms. Here the outcome is not "production" work but the creation of new "standard work" that can be deployed on the machines normally available to the company.
- Being able to access the combined computational science and domain expertise that exists in the national labs. As has been noted in the out-briefs of the HPC4Energy[5] access program at Lawrence Livermore National Lab, this internal resource has been extraordinarily important to the efficient use of the DOE's leadership computing capability by industrial partners.
- The continued development and maturation of the high-end codes that ultimately end up in the ISVs. The size of the user base for some of these codes is small enough that questions continually arise as to whether a business model exists for private companies to invest from beginning to end in new codes that can be made commercially available at a price the user base can afford. Thus, the initial development and derivation of these codes within the national laboratories is crucial.
- The role the national labs play in supporting the broad ecosystem of expertise across the HPC community. While an increasing percentage of engineers at large industrial companies possess deep computational modeling and analysis skills, the specific expertise at developing high-end applications codes and in optimizing the use of high-end HPC systems is often matured at the national labs. The continued care and feeding of this talent pool is a vital role played by the DOE lab community.

---

[2] http://www.doeleadershipcomputing.org/incite-program/

[3] https://www.llnl.gov/news/ibm-lawrence-livermore-researchers-form-deep-computing-solutions-collaboration-help-boost-us

[4] We note that large scale on-demand cloud computing is emerging as a viable option to investment in private HPC ownership.

[5] http://hpc4energy.org/

- Introducing small and medium enterprises who are either standalone businesses or, more importantly, critical members of a larger supply chains, to the value that high-end HPC can bring.

The second potential category who would benefit from expanded use of DOE HPC are those companies who, for reasons of size (people resources), capability (financial resources), or inexperience, do not currently make significant use of high performance computing in their operations. This group can benefit from the DOE leadership computing capability and support programs by:

- Executing test programs, modeling and simulations to solve a very specific problem for which internal resources are simply insufficient.
- Using such resources to a sufficient level so that data is available to make a determination of whether an investment in internal resources can be justified.

Critical to the success of such efforts is the availability of DOE staff expertise – computational and domain specific – as well as simple to execute administrative processes to allow utilization of the DOE leadership computing resources.

With these comments as background we recommend the following actions be undertaken by the DOE to enhance the utilization of high performance HPC by the industrial community:

1. Create and support an easy to navigate DOE portal describing all publicly available computing resources and programs to access them, tools to determine the best fit to the problems and opportunities presented by the private sector, and clear instructions and guidelines on how to access the resources and programs. Specifically, each DOE facility offering such access should identify and support a single individual to serve as an initial point of contact for a new company wishing to explore access.
2. Continue to support competitive programs that provide access to leading edge HPC computing at the DOE. Ensure that such awards include not only a designated amount of computing time but also ensured access to the computational and domain specific expertise in the labs that support those capabilities. Access to such programs should be through three categories:
    a. Initial awards at no charge to the outside party, with priority put on new users or new applications.
    b. Subsidized access to small and medium enterprises aimed at expanding the national user base of high performance computing.
    c. Follow on awards to large enterprises on a full or partially subsidized pay-as-you-go basis. As an added incentive, such programs might give “credit” to large enterprises that bring new partners to a computational program.
3. As part of the broader enhanced technology efforts underway at DOE, the Department should support programs leading to the commercialization of new or matured codes so that such codes are available through the ISV model to the public user community.
4. The DOE should be a key partner with the university community, the national academic accreditation bodies, and the private sector in enhancing engineering and science degree programs, to ensure that graduates have the necessary background and skills needed for a future in which effective use of high performance computing will be a standard expected capability for a STEM career.

III. The Beyond Exascale Research Program

As we look "over the horizon" beyond exascale computing, there are three promising areas of advanced high performance computing that are being researched: quantum computing, superconducting circuits, and neuromorphic computing.

Quantum computing uses quantum mechanics and the fact that microscopic particles exist in a superposition of quantum states to process information. Such a superposition of states is called a qubit. A qubit takes advantage of "quantum entanglement" - the ability of microscopic particles to be in multiple mutually exclusive states simultaneously. A qubit is based on a superposition of two states, and can be used to encode multiple possibilities. This allows, for example, the factoring of very large numbers, which is important for both encryption and de-encryption.

A number of companies are working in the quantum computing space. Researchers at IBM are making qubits from superconducting circuits, but the devices have high error rates, because there can be entanglement between a quantum computer and its environment. This makes quantum states very fragile. Another company, D-Wave Systems, has made special purpose quantum processors. But these processors are not faster than regular computers on a consistent basis. Finally, Microsoft is researching so-called "topological" qubits consisting of special quantum particles put into braid-like paths. Microsoft does not yet have a working computer.

As is the case with the work at IBM and elsewhere, superconducting circuits take advantage of the fact that superconductors have no electrical resistance. Because of resistance, standard conductors generate heat when voltage is applied to get current flow. In superconductors, there is coherent current flow at zero resistance. Therefore, there is no heat dissipation.

A group at MIT has fabricated a superconducting nanowire, three-terminal, electro-thermal device called a nanocryotron. Their device uses a geometry that allows superconducting flow to be switched off as a function of current. This switching is the basis of computation. Their device is smaller than, and requires less current than, Josephson junction devices, but large enough currents to allow appropriate cross-communication on a computer motherboard. These are three terminal devices - as are semiconductor transistors, can be integrated with magneto-resistive spintronic devices at room temperature, and have high impedance. This work is funded by NSF and by IARPA.

Superconducting circuits and architectures can be used for both classical and quantum computation.

Perhaps the most progressed of "over the horizon" computing is neuromorphic computing, which is brain-inspired computing. The newest programs on the road to exascale already are incorporating architecture for data-centric computing. But all of these computers still use the von Neumann architecture, which shuttles data back-and-forth between a CPU and memory chips in linear sequences of calculations. This design is fine for doing FLOPS, and for executing precisely written programs, but not for processing, for example, images or sound. A method for having a more data-centric architecture is through multicore processing, which uses multiple cores that work on data that is more proximate to each core. But trying to emulate the brain by using special software on conventional processors is too inefficient to be the basis of machines with greater intelligence.

Neuromorphic computing recognizes that biological circuits are more efficient and can process complexity in the form of images and sound. Therefore, work has begun on bio mimetic designs that employ a huge number of computational elements, each dedicated to a specific task or subtask. This is meant to avoid the excess energy involved in moving currents and transporting data. The design of such systems tries to emulate the synapses that occur between neurons.

A number of companies and universities are working on developing neuromorphic systems. Some are supported through a DARPA supported program called SyNAPSE - with the ultimate goal to develop a microprocessor that matches a human or animal brain in function, size, and power consumption. IBM and HRL Laboratories are the primary SyNAPSE contractors developing neuromorphic chips under the $100 million DARPA program which began in 2008 and will go until 2016. IBM and HRL subcontract parts of the research to various universities. Qualcomm is developing a neuromorphic chip modeled on biological brains that will respond to data in ways that are not programmed. In fact, Qualcomm's “Zeroth Program” represents the first large scale commercial platform for neuromorphic computing. Universities, including Cornell, Georgia Tech, MIT, and Stanford, all are developing their own approaches to neuromorphic computing.

There already are public and private investments in new computational models, architectures, and systems that are beyond the amount of investment the DOE can commit to through its normal funding channels, as it scales up its investments on the road to exascale computing. Rather than commit to specific technological bets in over-the-horizon computational systems, DOE should invest to maintain and strengthen the computational ecosystem, including work with universities. This would allow DOE to understand what already is underway, while focusing on more advanced elements of over-the-horizon computing, including software development.

If one asks how to structure a forward-looking program that supports the necessary advanced computational ecosystem beyond exascale, we have noted that, as DOE supports the road to exascale, a combined investment to sustain the advanced computing ecosystem, and to look "over the horizon" should be funded at $100-$150 million per year. Included in this amount should be $20-$25 million per year to enable DOE to stay abreast of developments being sponsored by others.

This would allow DOE to understand how technology is evolving as it supports work leading to exascale computing. Just as data-centricity already is being incorporated into the new computer system designs, likewise, the expectation would be that DOE would and could incorporate elements of "over the horizon" technologies into new computing architectures. For example, neuromorphic computing is less about components per se, than about how to compute in new architectures. Therefore, work in "over the horizon" computing will lead to big changes in software design for new computational capabilities.