Position Paper for 2nd Brainstorming HPCD Workshop

Barry I. Schneider

Feb 10-12, 2015

Dallas, Texas

I. *Preface*

In trying to address the issues being posed by the NAS report and to provide the most useful input to the workshop, I felt, as a former NSF Program Director in the ACI Division, my perspectives as an "insider", as well as a user of NSF HPC resources, are somewhat unique. However, they are also colored by that experience and as such may be viewed as biased by some. I have tried my best to keep my comments as objective as possible, given the circumstances.

II. *High Level View*

At the 50,000 foot level, I believe it is essential to the health of the open science community that the NSF continues to pursue an aggressive approach towards cyberinfrastructure (CI) at all levels as it has done for over 30 years. The establishment of the XSEDE program, the funding of Blue Waters at UIUC, Stampede at TACC, Gordon at SDSC, the new Comet resource also at SDSC, and some even newer hardware acquisitions such as Maverick at TACC and Bridges at PSC, are enlarging the NSF portfolio to encompass projects involving very large data analytics and/or visualization. These, as well as the increased funding of software projects, have all been quite positive developments. It should be mentioned that the funding of these projects occurred over a period of almost ten years, which is a few lifetimes for CI projects. However, with tighter budgets, it appears that the Foundation is faced with serious future challenges that could derail these valuable efforts. This unfortunately comes at a time when there is compelling evidence that computational and data intensive computing are a critical enabler for discovery in nearly all fields of science and engineering as well as for technological advancement. Asia and the EU have already recognized this potential and are increasing their investment in high performance computing.

The NSF is the sole agency responsible for basic science in the US. As such, it not only supports advanced computational resources for NSF grantees but provides resources to researchers supported by other agencies such as the DoE, NIH, NASA, etc. While some of these agencies also support researchers outside their domains (DoE the major player), others do not. At last count ~50% of all of the NSF supported compute cycles (including the cost of operations and maintenance) go to non-NSF supported researchers. In addition, the resources supported by NSF are for the most part located in university environments where education is a priority and overhead rates are comparatively low. The result is a high-quality environment for students being trained in computational science and scientists engaged in advanced computation. All of this is provided in an economical fashion. **This argues strongly for increased investment in NSF funded CI, not less.** Recent congressional discussions for increased investment in DoE supported HPC need to be vigorously enlarged to include NSF for all of the reasons stated above.

III. NSF Investments in HPC

The NAS report begins in much the same way as many other studies discussing the importance and impact of advanced CI, technology challenges, the declining success rates for proposals evaluated by the XRAC and the issue of double jeopardy for many researchers requiring advanced digital resources to accomplish their disciplinary research. The question of whether the NSF can continue to invest in future highest tier systems in the same class as those of DoE, DoD and other federal mission oriented is raised. This seems a curious statement. The NSF is the primary federal agency for the pursuit of non-mission driven science. Its very charter requires the agency to support research and researchers based on a competitive, peer review process. As such, there are strong arguments for increasing resources in strategic science and engineering areas. It is hard to argue against advanced CI not enabling advancement in many of these disciplines. For NSF this is a particularly difficult challenge as a consequence of the broad spectrum of science supported by the agency. Statements in the report that suggest that resources could be made available by purchasing support from other federal agencies or from commercial sources are, IMO, by and large misplaced as a consequence of government regulations, and/or comparative cost.

a. Action Item: The results of the workshops must provide compelling **quantitative** arguments as to the impact of NSF ACI on open science and the consequences of the NSF stepping back on a decades long commitment to HPC and advanced CI. Using other agency and/or commercial sources to fill the gap also need to be examined and the practical consequences exposed.

The omnipresent question of where should the money come from to enable the NSF to retain a leadership position in ACI is not simple. The current budget of the NSF ACI Division is about $210-220M. While it should be recognized that support for the broad spectrum of computational and data driven research comes from many parts of the foundation, the support of high end resources, such as the actual HPC engines and the support to maintain them and to provide users with the environment required to facilitate the use of these resources, such as XSEDE, comes from ACI, with the single exception of Yellowstone, which is a domain specific resource for the GEO Directorate. There have been arguments in the past that OCI (now ACI) is quite underfunded. That argument is even more the case today.

a. Action Item: The report should provide some possible scenarios of what would be needed in funding to keep ACI in a leading position in CI. This should be done realistically, looking at what other agencies are investing and what the investments are supporting. Keep in mind that budgets are not increasing and it might be necessary to provide arguments to the NSF Director that movement in funds are justified. The NSF is an $8B organization and has some discretion over its use of those funds. In addition, if the community makes a compelling case to the Director, based on what would be lost by retreating, that would be very helpful.

IV. *Balance*

It is worth examining how investments in ACI or its predecessors have been changing over the past ten years. At its outset, the Office of Scientific Computing and its immediate successors were organization almost entirely focused on HPC. Until a few years ago, these organizations reported to the NSF Director and were on an administrative par with other Directorates. That changed and now ACI is a Division in the CISE Directorate with a much broader agenda. There were periods where 80% of theorganizational budget was spent on HPC hardware and the support of center-like activities. This began to change and in recent years budget priorities shifted more towards software and big data and away from investment in hardware. This was done essentially under a flat budget (really decreasing) scenario. While this is not in and of itself without merit, there are always questions of proper balance between hardware investments, software investments and the investment in the shared services (read XSEDE) required to glue together the entire enterprise in order to produce an organic, functioning whole**.** This is a difficult challenge and one that requires input from all of the stakeholders.

a. Action Item: An effort should be made return what is now ACI to its previous status as an independent entity reporting directly to the Director. As a cross cutting, interdisciplinary organization, it deserves independence and direct access to the NSF Director. Also, the mission of ACI is quite different than the rest of CISE and I see few arguments to support its present location.
b. Action Item: The report should critically examine the current investments on NSF HPC hardware resources, software for both general and discipline specific software, big data, networking, education, etc. and provide some advice on whether those investments are meeting the needs of all the stakeholders or whether there needs to be some rebalancing. There is certainly strong evidence that there have been substantially increased demands by the scientific community for NSF HPC hardware resources. While the situation has always been competitive, it is now clear that many researchers are not receiving what they need to accomplish the scientific goals of their peer-reviewed proposals. Many applicants are simply turned away due to lack of resources.

V. A New Model for HPC Investment

The manner in which NSF currently supports advanced CI is worthy of discussion. Being primarily a basic research organization, the NSF regards competitive peer review as fundamental to its credo. This is generally totally appropriate for most of the NSF portfolio. One can question the validity of this approach for CI. CI is not research but it enables and supports scientific and engineering research across the entire NSF and de facto the open science community. Should there be some appropriate modification for infrastructure projects that are not research in and of themselves but catalyze research and discovery across a broad spectrum of the scientific community. Certainly one can envision alternative scenarios. Given that a functioning CI has some similarity to other large, shared and long-term research projects, perhaps some modification of the existing MREFC process would be appropriate. What is needed is something that is longer term, can serve across NSF disciplinary lines and is nimble enough to respond to the demands of different communities. Many large projects at NSF have substantial CI needs but have been reluctant to look outside their own projects for support as a consequence of a perceived lack of stability in NSF CI.

a. Action Item: The workshop should present a set of possible scenarios to better manage NSF CI for the community. The model needs to recognize the need for peer review and competition but also the need to sustain CI projects over longer lifetimes than a typical research project. The model should address CI needs across the Foundation in a way that tries to get around the "not invented here" issue and also the perception by a number of current MREFC projects that ACI funding is too short-term and not appropriate for MREFC projects which often span decades.

VI. The User Perspective

The group I have been associated with (PI-Klaus Bartschat of Drake University) is engaged in the ab-initio solution of the quantum-mechanical many-body problem, including charged-particle and (intense) laser interactions with atoms and molecules. This is basic research looking at numerical quantum mechanical correlations. Such calculations support, for example, very expensive experiments on coherent control, i.e., attempts to affect the outcome of chemical reactions by stirring the electrons (to the extent allowed by quantum mechanics) as well as plasma physics (atomic and molecular data for modeling of things like ITER). There are also many applications in materials science, astrophysics and even medical physics (e.g., radiation damage by light or charged particles). In a nutshell, many of the things that emerge from our computations not only help experimentalists better interpret their experiments but are often more accurate than the experiments. There are very few areas of physics which can make such a statement. The overall need is to solve the time-dependent or time-independent Schröedinger equation, depending on whether there is or is not an explicit time dependence in the Hamiltonian. In a very real sense, these problems bear strong similarities to those in quantum chemistry. They require the solution of high-dimensional partial differential equations in space and often time as well. The equations are discretized using grids or spectral methods (basis sets) leading to VERY large matrices which must either be diagonalized, employed to solve large linear systems or to propagate the known wavefunction \Psi(**r**,t) to \Psi(**r**,t=Delta t). In some cases, the matrices are sparse, but there are instances when they are dense. Thus, as the number of electrons grows, most of the ab-initio methods employed, scale quite poorly with the size of the system. There have been attempts, using variants of densityfunctional theory and effective core potentials, to reduce the scaling, but at present it is unclear whether these approaches are powerful or general enough to be predictive.

All of the codes we use have been written by members of the group. As the codes have matured, they have been made publically available through Computer Physics Communications. The codes are parallelized using MPI and in some cases MPI plus OpenMP. The current codes show strong scaling up to 5-10K cores. In one of the codes, extensive use is made of ScaLAPACK or Elemental to handle the large, dense eigenvalue problem. All of the eigenvalues and vectors are required to construct what is known as the R-matrix. This matrix provides the connection between the inner or short-range interaction region and outer region that is dominated by long-range multipole forces. Scattering information for a broad range of energies may be computed once the R-matrix is determined from the eigenvalues/eigenvectors of the large, dense matrix.

The laser-matter code is currently dominated by structured, sparse matrix-vector multiplies where the non-zero elements consist of a set of relatively small (~20 by 20), dense matrices. Effective use of say MKL for this problem is inefficient due to the size of the blocks but the blocks themselves are distributed to the processors and communication between processors is limited to nearest neighbors and then only at the edges of the boundary. This code is now being adapted to work with the Intel Phi accelerators on Stampede at TACC and we are looking at recasting the problem in terms of matrix-matrix operations.

Finally, our experience with using NSF resources via the XSEDE program has been quite positive. We have written competitive proposals beginning with very modest requests and as we learned and as our workloads increased expanded. The group is presently in the top 20-30 XSEDE users. We are using ~10M SUs a year. The people associated with the XSEDE program have always been quite helpful when we have had code issues (excellent ECSS support) and also in facilitating the movement of SU's from one machine to another. This would have been difficult in the absence of this infrastructure.