Critique of the Branscomb report

Anirudh Modi

February 3, 1998

1 Summary of Branscomb report

This report presented in August 1993 addresses an opportunity to accelerate the progress in almost every branch of science and engineering by dealing effectively with the development of high performance computing (HPC) resources and their availability to all kinds of organizations in the nation, viz. educational, governmental and industrial. The panel feels that HPC is the key to the nation's technological and economical success. It considers four main challenges that that need to be overcome and suggests suitable remedies for each of them.

1. Removal of technological barriers relating to the development of suitable parallel hardware, and implementation barriers relating to development of suitable algorithms (i.e. software) to take advantage of this kind of hardware. Also, National Science Foundation (NSF) need to weave together existing work in various developing scientific disciplines, as well as fostering new bridges between pure, applied and computational techniques.

2. Appropriate acquisition and distribution of HPC resources among the scientists and engineers depending on their need. Each level of institution foreseeing the need for HPC facilities should have atleast a minimum set amount of resources available locally to support their scientists and engineers.

3. Encouragement of usage of these HPC resources on a broad basis by expanding education and training and encouraging participation by other HPC institutions. This includes providing the right incentive for

all three levels of the computational institution pyramid by funding appropriate research.

4. Determination of a suitable role of NSF in creating the intellectual and management leadership for the future of HPC and its impact on other federal agencies. This mainly deals with the budgeting of the available funds by the NSF to promote critical sectors of High Performance Computing and Communications (HPCC) like networking, parallel architecture, etc.

The report describes four sets of interdependent recommendations to deal with these challenges.

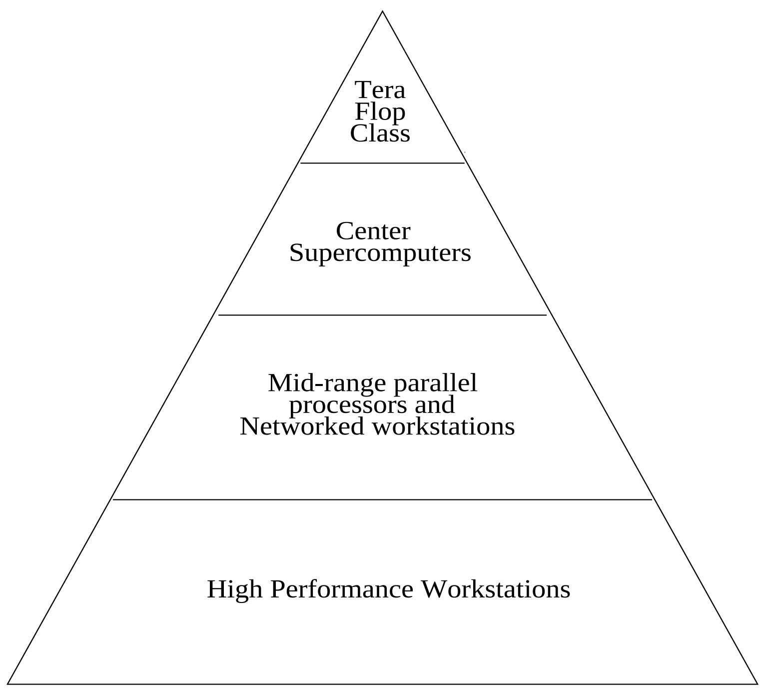

Figure 1: HPC pyramid

1. The report envisions the HPC *pyramid* environments and the first recommendation suggests the implementation of a balanced HPC *pyramid*, each element supporting the others resulting in a stable distribution of the resources. The apex of the pyramid (teraflop class computers) is the

need for national capability at the highest level of computing power to deal with "Grand Challenge" problems posed to us. The United States should lead the world in HPC by collaborating with big agencies like ARPA and DoE to expand access to all levels of the *pyramid* for all sectors of the nation. It is also aimed that the research universities should be assisted to acquire mid-range machines over the next five years, and that that NSF should double the current annual level of investment in scientific and engineering workstations for its 20,000 principal investigators. It was also recommended that NSF support the development of new parallel configurations for HPC applications.

2. To accelerate progress in developing the required HPC technology, NSF should create a challenge program in computer science with grant size and equipment access to support the systems and algorithm research. NSF should focus attention on support for the design and development of computational techniques, algorithmic methodology, and mathematical, physical and engineering models to make efficient use of the machines. NSF should also set up an agency-wide task force to develop a way to remove the imbalance in the HPC *pyramid* caused by the the under-investment in the emerging mid-range scalable and parallel computers and the inequality of access to stand-alone potentially networked workstations.

3. The Supercomputer Centers should be retained and their missions, as they have evolved, should be reaffirmed. An open re-competition of the four Supercomputer Centers should be avoided at this time, instead favoring periodic performance evaluation and competition for some elements of their activities. The NSF should continue to provide funding to support the Supercomputer Centers' HPC capacity. The NSF should review the administrative procedures used to allocate center resources to ensure that the burden on scientists applying for research support is minimized. The NSF should give strong emphasis to its education mission in HPC, and should actively seek collaboration with state-sponsored and other HPC centers not supported primarily on NSF funding.

4. The National Science Board should urge the Office of Science and Technology Policy (OSTP) to establish an advisory committee representing

the states, HPC users, NSF Supercomputer Centers, computer manufacturers, computer and computational scientists which should report to High Performance Computing, Communications, and Information Technology (HPCCIT) subcommittee.

2 Summary of role of NCSA

The National Center for Supercomputing Applications (NCSA) is the leading-edge site for the National Computational Science Alliance (Alliance), which is one of two recipients of the NSF's Partnerships for Advanced Computational Infrastructure (PACI) program. It anchors the partner teams and carries out much of the software development to support the Alliance. Led by NCSA at the University of Illinois at Urbana-Champaign (UIUC) with core funding from the NSF (upto $170M in the next 5 years) and the state of Illinois, the Alliance is building a National Technology Grid (Grid) which will integrate high-performance computers, advanced visualization environments, remote instruments, and massive databases via high-speed networks to form the most powerful problem-solving environment ever assembled. During the next decade, the Alliance aims at making the Grid as easy to use as the desktop is today. Researchers will work together and share ideas with ease through the use of collaboration technologies for the desktop or virtual reality. They will have at their disposal advanced searching and indexing capabilities, digital libraries, audio and video streaming, and virtual offices. It aims at creating a distributed computing grid populated with powerful computers that are individually scalable and that, when necessary, can be assembled into metacomputers. NCSA's focus is on the newly emerging distributed shared-memory (DSM) architecture. This platform combines the easy programmability of shared-memory symmetric multiprocessors (SMPs) with the scalability of the distributed-memory massively parallel processors (MPPs). NCSA already has a 512-processor SGI/CRAY Origin2000 which is supposed to be the largest unclassified CRAY Origin2000 in the world.

3 Summary of role of NPACI

The National Partnership for Advanced Computational Infrastructure (NPACI) is the second recipient of the NSF's PACI program. Led by University of

California at San Diego (UCSD), NPACI aims to revolutionize the nation's computational infrastructure by building on the foundation of San Diego Supercomputer Center (SDSC) and involving 37 of the nation's leading academic and research institutions. Driven by real application needs, NPACI is developing software infrastructure to link the highest performance computers, data servers, and archival storage systems to enable easier use of the aggregate computing power. Development work is focused in specific "thrust areas" that join applications scientists, computer scientists, and technology developers and leverage separately funded research projects to ensure rapid deployment and robustness of the resulting infrastructure. This work is complemented by an extensive education, outreach, and training program and collaborations with industry. In developing an integrated computing infrastructure (or metacomputing environment), NPACI will direct its efforts in four technology areas: metasystems, programming tools and environments, data-intensive computing, and interaction environments. The NPACI education and outreach thrust area plans to help society make productive use of computing technologies to understand and solve problems in science, business, government, and society by focusing on five areas: K12 education, undergraduate education, outreach to women and minorities, outreach to communities new to high performance computing, and evaluation of educational programs. NPACI's strategic plan calls for installing a teraflop-scale computer at SDSC within two years and a petabyte (1024 terabyte) archive with distributed caches within five years.

4 Recommendations of possible changes

Both NCSA and NPACI seem to be concentrating rather too much on the top level of the HPC *pyramid* (teraflop class computers) and in the process are neglecting the lowest level (high performance workstations). Little work seems to be going in that direction. I feel that they should search for better and more feasible ways to utilize the wasted power of the millions of workstation class PCs both in their organizations and nationwide. Also, focus should be primarily on extremely effective and efficient utilization of their resources than on acquiring more and faster resources. These will also lead to the optimum utilization of their funds. Also, little work is being done on exploring hardware possibilities (better processor, memory, etc.) for the future by these partnerships themselves, as heavily reliance is being placed on the hardware/chip industry, which seems to be having less and less players each day and seems poised for an unhealthy scenario soon. Currently being one of the primary consumer of these high performance hardware, both these alliances should also look forward to including some players from the hardware industry as their partners.

5 Conclusion

On the whole, the Branscomb report seems to provide a very accurate summary of the existing HPC status and doles out very powerful and intelligent guidelines to the NSF for continuing its HPC initiative. The two partnerships, NCSA and NPACI seem to be following these recommendations reasonably well and look poised for immense progress in the next five years. The growing amount of HPC resources seems to be spreading fast in every nook and corner of the nation and a rapid technological revolution seems imminent as we approach the 21st century.