

Impediments to Future Use of Petaflop Class Computers for Large-Scale Scientific/Engineering Applications in U.S. Private Industry

Myron Ginsberg

ACM Fellow and HPC Consultant
HPC Research and Education
Farmington Hills, Michigan 48335-1222 USA
m.ginsberg@ieee.org

Abstract. The environment in government/research HPC sectors is markedly different from that in private industry. Although both involve many of the same applications, the mindset and people/computer resources are significantly different. In this paper, we focus on the barriers to using future HPC machines in the private sector and some current actions and suggestions to overcome these problems. Experts generally agree that the realistic and/or real-time solution of industrial problems will require the use of computers about three orders of magnitude faster than current industrial machines. Impediments discussed include: limitations of ISV-based commercial software, inadequate benchmarking techniques for industrial size problems, limited access to the latest computer architectures and support facilities, paucity of computational science personnel, slow tech transfer of algorithms and modeling techniques from government/research facilities to private industry, and unexplored utilization of Blue Collar Computing™ in private industry.

1 Introduction

This paper focuses on the leading edge of the supercomputing frontier primarily in the U.S. industrial sector. Most of the applications involve the solution of realistic, three-dimensional simulations (modeling) of physical phenomena such as automotive/aerospace designs, climate and weather prediction, or a diverse collection of bioinformatics applications. All of these involve extremely large amounts of floating-point computations and data manipulation of massive databases. This work is typically performed in government/military research labs and/or in academic research facilities and increasingly often by a combination of such organizations. In the United States, such activities are performed in government facilities supported by the National Science Foundation, Department of Defense, NASA, Department of Energy, and other government agencies in cooperation with a variety of academic institutions. To obtain very accurate and fast and/or real-time three-dimensional solutions requires very effective utilization of computer resources several orders of magnitude faster than those typically now in use in private industry and most government labs. Experts

estimate that petaflop class machines (capable of performing 10^{15} floating-point operations/second) are needed to effectively solve such problems but current technology is about three orders of magnitude less in the teraflop range, 10^{12} floating-point operations/second.

2 Government Versus Industrial HPC Environments

There is a significant difference between the HPC environment in government/academic research facilities and that encountered in most private industry facilities. For example, in the former category there is a tacit tolerance to try to utilize the latest, most advanced computer systems, typically serial – 1, i.e. often a prototype or beta category system in which the system software is often relatively primitive, somewhat unstable, constantly in a state of transition and the application software is home-grown, in-house developed rather than commercial software. Despite this somewhat chaotic state, there is much excitement and creativity as well as patience trying to make significant leaps forward in the solution process for problems previously unsolvable and/or inadequately or incompletely formulated to provide realistic solutions. There is usually a large cadre of support people available on site to cope with a plethora of systems software, applications software and/or computational physics problems and/or computer hardware/networking difficulties which may arise during the solution process.

In sharp contrast to the government/academic research environment described above, most industrial facilities in the U.S. do not have significant in-house HPC research facilities nor a large support staff. Thus most industrial organizations are much more conservative, with little tolerance to cope with prototype/unstable systems in their industrial production environments. These organizations are almost entirely dependent upon the use of commercial independent software vendor (ISV) based applications. The ISVs are usually unwilling and financially unable to port and optimize their codes for a new architecture until they are relatively certain that there will be sufficient paying customers to ensure the port will be successful. The net effect of this understandable mindset is that industrial sites might often have to wait a year or more before new HPC architectures have available ported and optimized versions of their ISV-based codes, a significant bottleneck to the use of the latest HPC technology.

3 Accurate Benchmarking Needs for Industrial Customers

The situation is further complicated by the inability to accurately benchmark a new architecture until the ISV-based software is ported and optimized. At the present time there is no means of even getting good estimates of how industrial codes might perform on the new platform. Kernels of large codes can be initially tested but there is no way of accurately accessing overall code performance when the entire application is run. The overall performance can vary substantially depending upon which path through the code is traversed and which one can best utilize both the hardware and software on the new hardware platform; thus it is indeed possible to get drastically

different performance on the same hardware platform with the same commercial code applied to two very different applications. Current benchmarking strategies have to be considerably improved so that a potential industrial user of a new HPC platform can obtain a reliable indication of its value for a specific problem.

To help alleviate this situation there is a need for industry benchmark suites evaluated on diverse HPC hardware platforms using ISV-based software. Furthermore, individual companies need to devise and test their own benchmark codes and/or establish multi-industry benchmark tests. Independent benchmarking is needed rather than relying on testing conducted solely by the hardware and software vendors. A good example of such testing is the extensive efforts that DOE has expended in assessing upper end HPC machines such as the Cray X1. These benchmark tests [1] on in-house developed codes expose the good, bad, and the ugly traits of new HPC platforms with no noticeable vendor spin on the results. Such additional benchmarking is needed on an industry by industry basis. Furthermore, such benchmarks must also be able to provide some meaningful guidance to industrial users with their ISV-based codes.

4 Tech Transfer Between Government and Private Industry

In a previous paper [2] the author indicated the need in the U.S. auto industry for much faster tech transfer of algorithms from government labs to private industry. This can be accomplished in general by much closer interactions between the research community and the industrial ISVs. This could become a win win situation for both sides with U.S. industry the main beneficiary. This activity can be speeded up by increased industrial use of existing government subsidized math software libraries which could be embedded within their ISV-based codes. Increased industrial use of such libraries would also help in industrial benchmarking of new HPC hardware platforms on which those software libraries have already been ported.

5 Feedback of Recent Study of U.S. Industrial HPC Users

In addition to the obstacles to industrial HPC users mentioned in the previous sections above, a recent investigation of 33 U.S. industrial organizations [3], [4] offers its perspective on this subject. Here are some of the comments mentioned in the study: (1) 65% of respondents could not quantify direct benefit of HPC to bottom line of corporate expenses even though most survey responders admit that “high performance computing is essential to business survival”; (2) Security concerns are an important inhibiting factor to outsourcing HPC competitive sensitive problems; (3) There are current important HPC problems that are not being solved today because of one or more of the following factors: problems are too large and/or require too much computing time or memory with current in-house machines; (4) “Companies are failing to use HPC as aggressively as possible” because of corporate financial restrictions, management limited vision of HPC return on investment and/or limited trained technical personnel to deal with all aspects of HPC in-house; (5) “Most companies don’t have the HPC [hardware and software] tools they want and need;”

(6) “Dramatically more powerful and easy to use computers would deliver strategic competitive benefits”; (7) “High-performance computers are desired based on actual delivered results on end-users computational problems, but most {industrial} sites [especially small companies] cannot afford to purchase the fastest computers available in the market today.

6 Help Is Coming Now

Aid for many of the problems mentioned in the preceding sections is slowly coming via renewed government initiatives to help the U.S. to be more globally competitive and also prompted by concern from the U.S. Congress as well as from prominent scientists and engineers. For example, a recent report [5] by the Committee on the Future of Supercomputing of the National Research Council has outlined a series of recommendations to rejuvenate national supercomputing efforts partly in response to the Japanese success for the past few years with the Earth Simulator overshadowing performance of U.S. based supercomputers [6]; the recommendations address hardware, software and networking concerns with particular emphasis on assuring the success of multiple domestic suppliers. DARPA continues with its efforts with IBM, Cray, and Sun with government and university partners to establish a petaflop class supercomputer by the end of this decade [7], [8], [9], [10], [11]. Additional plans for government consolidation of HPC efforts across government agencies are outlined in the Federal Plan for High-End Computing Report [12]. Industry leader Steve Wallach and others point out concerns in developing software for Petaflop class machines [13].

Several efforts are focused on improving HPC benchmarking techniques. One of those efforts, HPC Challenge Benchmarks [14], is an expansion of the Top500 [15] criterion to include such machine attributes as sustainable memory bandwidth and latency as well as bandwidth of various simultaneous communication patterns and measures of the rate of integer random updates of memory. Such attributes when combined with LINPACK TPP benchmark gives a much better perspective of machine performance. Another benchmark test is the IDC balanced rating [16] system for comparing various machines.

An interesting HPC performance metric is system balance. This can be expressed as ratios comparing resources to CPU performance. For example, ratio of bytes of memory to flops, ratio of memory bandwidth to flops, ratio of interprocessor communications bandwidth to flops, or the ratio of disk I/O bandwidth to flops. For more details including ratios for some specific machines, see [17].

Several HPC performance measurements and new tools for more accurately assessing high performance computers are being developed at Lawrence Berkeley National lab [18] by several groups including The Performance Evaluation Research Center (PERC) [19] which is trying to develop a science for better comprehending HPC performance of scientific applications. The Berkeley Benchmarking and Optimization Group (BeBOP) [20] is focusing on the interaction between application software, compilers, and hardware as well as automating the performance tuning process.

One of the concerns mentioned in Section 6 above as an impediment to vigorous use of HPC in private industry is the lack of sufficient trained HPC personnel. This is

very noticeable when comparing the environments in government/research facilities vs. U.S. private industry. The former sector usually has large in-house, multidisciplinary HPC support staffs whereas in most industrial sites there are very few such people. This problem is becoming acute because as the complexity of HPC problems and computers grows, it is no longer possible for an application specialist to also be an expert in a growing variety of computer hardware, software, and networking directly impacting on the performance speed and correction of the application solution process. The long-term solution to the problem is training more people in multidisciplinary computational science programs. Too many academic departments, unfortunately, have established feudal domain mindsets which have greatly inhibited multidisciplinary projects across departments and colleges of the same or different universities. This mindset must cease if the U.S. is to be successful in the global competition arena. At present there are well over 30 computational science programs in the U.S. and Europe at both the undergraduate and graduate levels [21].

The current and future government/research HPC activities are likely to follow the pattern established by previous generations of HPC efforts (such as the transition to vector machines): innovations first accomplished by explicit tedious hand coding followed by development of math software libraries, explicit compiler directives, then automatic compiler optimizations of user source code. Thus we can expect adaptive techniques to effectively utilize the coming generation of HPC with heterogeneous multithreading processors utilizing both vector and scalar modes and incorporating automatic use of PIMS (processors in memory) and FPGA (field programmable gate arrays) technology. Even if this evolution occurs we will still need to develop new HPC programming languages [13], [22] to promote transportability and interoperability amongst HPC hardware platforms and to permit the movement of legacy codes to these new machines in a manner that will let such codes effectively utilize the new hardware and software with minimal explicit user provided modifications. This scenario will be critical to success of HPC in the private industry sector especially if this sector continues to lack a critical mass of trained computational science personnel. HPC problems have already become much too complex to expect the application specialist to correctly anticipate all factors which directly impact the correct numerical and computer optimization of the entire problem solution process.

7 Blue Collar Computing™

The HPC private industry community tacitly assumed in this paper has had to slowly emerge from observing and leaning from the government/research HPC community. Most of the companies that have immersed themselves in HPC have generally been large organizations with the financial and people resources to make the necessary commitment to HPC. These pioneers have slowly discovered that HPC definitely can have significant return on investment. For example, the U. S. automotive industry in 1980 had a 60 month (5 years) lead time between concept and production of a new vehicle; now in 2004 that lead time has shrunk to under 18 months in large part due to HPC efforts and math-based computer modeling techniques which drastically reduced

the amount of physical prototyping while at the same time allowing engineers and scientists to consider more design alternatives, improve the quality and safety of the new vehicles, and contribute to improving global product competitiveness [23], [24].

Now as Ohio Supercomputer Center (OSC) Executive Director Stan Ahalt pointed out in his SC2004 speech [25], there still remains many segments of U.S. private industry untouched by the innovative potential of HPC which could improve their products and services (especially in many manufacturing areas) as well as contribute to improving U.S. global competitiveness.

Many of the comments in this paper about HPC in industry apply to both the established HPC users as well as future novices such as those depicted as part of the Blue Collar Computing™ community. A few words of encouragement here for the latter group. In some ways your HPC journey will be both easier and harder than that of your previous HPC industry pioneers. For example, the latter group had to generally have access to multimillion dollar supercomputers which were rare in U.S. industry in the 1980's and 1990's but today the hardware prices of some Linux based cluster desktop machines are less than \$10,000 with deskside models under \$100,000 [26], [27]. Also thanks to growing availability of access to larger HPC machines in grid environments within government and research facilities, the Blue Collar Community will have additional non-in house resources for potential use. On the negative side, the most crucial issue facing the Blue Collar Computing™ community may well be the lack of sufficient computational science personnel; see comments in Section 6. The short-term solution will require financed interactions between the HPC government/research community and the Blue Collar Computing™ community. An example of one such interaction is NCSA's Private Sector Partner Program [28] which focuses on real-world industrial challenges and currently involves Allstate, Boeing Phantom Works, Caterpillar, IBM, and Motorola Labs and will be using NCSA's newly installed 7 teraflop, 512 node Dell Cluster which will enable these industrial partners to reap the benefits of early access to breakthroughs.

8 Summary and Conclusions

This paper has spotlighted the roadblocks to U.S. industry use of future petaflop class HPC including: (1) the need for more meaningful industrial benchmarking techniques to help in the selection process for new machines; (2) Faster and more effective tech transfer of application oriented algorithms between government/research facilities and private industry via closer relationships with commercial ISVs; (3) influx of more trained computational science personnel in the private sector to cope with increasingly more complex and multidisciplinary oriented HPC applications; (4) Government support for the creation of much faster and more efficient industrial friendly supercomputers (with easy to use software and hardware tools); (5) Much improved early industrial access to new HPC hardware platforms via improved techniques for fast porting and optimization of important industrial widely used commercial ISV-based codes.

Government and industrial attention and meaningful follow up to the issues in the previous paragraph will benefit both the established industrial HPC users as well as the next generation depicted in the Blue Collar Computing™ community. It would be

very helpful in the interim period if the former would help the latter via the creation of industrial mentoring activities.

References

1. Oak Ridge National Laboratory: Papers and Presentations on Cray X1 Evaluation. <http://www.csm.ornl.gov/evaluation/PHOENIX/index.html>
2. Ginsberg, M.: Influences on the Solution Process for Large, Numeric-Intensive Automotive Simulations. In: Alexandrov, V. N., Dongarra, J. J., Juliano, B. J., Renner, R. B., Tan, C.J. K. (eds.): Computational Science – ICCS 2001 International Conference Proceedings, Part 1, Lecture Notes in Computer Science, Vol. 2073. Springer-Verlag, Berlin Heidelberg New York (2001) 1189-1198
3. Joseph, E., Snell, A., Willard, C.G.: Council on Competitiveness Study of U.S. Industrial HPC Users. White Paper, IDC, Framingham, MA (July 2004) http://www.compete.org/pdf/HPC_Users_Survey.pdf
4. The Council on Competitiveness: First Annual HPC Users Conference: Supercharging U.S. Innovation & Competitiveness. Council on Competitiveness, Washington D.C. (2004)
5. Graham, S.L., Snir, M., Patterson, C.A. (eds.): Getting Up to Speed: The Future of Supercomputing. The National Academies Press, Washington, D.C. (2005)
6. The Earth Simulator. <http://www.es.jamstec.go.jp>
7. DARPA: DARPA Selects Three High Productivity Computing Systems (HPCS) Projects. http://www.darpa.mil/body/NewsItems/pdf/hpcs_phii_4.pdf (July 8, 2003)
8. Cray, Inc.: DARPA HPCS Cray Cascade Project. <http://www.cray.com/cascade/>
9. IBM, Inc.: DARPA HPCS IBM PERCS Project. http://www.research.ibm.com/resources/news/20030710_darpa.shtml
10. Sun Microsystems Inc.: DARPA HPCS Sun Hero Project. <http://www.ncsc.org/casc/meetings/CASC2.pdf>
11. Ricadela, A.: Petaflop Imperative. Information Week (June 21, 2004) <http://www.informationweek.com/story/showArticle.jhtml?articleID=22100641>
12. HECRTF: Federal Plan for High-End Computing: Report of the High-End Computing Revitalization Task Force (HECRTF). Office of Science and Technology Policy, Executive Office of the President, Washington, D.C. (May 10, 2004) http://www.itrd.gov/pubs/2004-hecrtf/20040510_hecrtf.pdf
13. Wallach, S.: Searching for the SOFTRON: Will We Be Able to Develop Software for PetaFlop Computing?. Keynote Talk, ISC2004, Heidelberg, Germany (June 23, 2004)
14. ICL: HPC Challenge Benchmark. The Innovative Computing Lab, U of Tennessee, Knoxville, TN, benchmarks and results available at <http://icl.cs.utk.edu/hpcc/index.html> and http://icl.cs.utk.edu/hpcc/hpcc_results.cgi
15. Top500 Supercomputer Sites. <http://www.top500.org>
16. HPC User Forum: IDC Balanced HPC Benchmark Rating Report. <http://www.hpcuserforum.com/benchmark/benchmarkresults.asp>
17. Cray, Inc.: Balance – the Key to Exceptional Application Performance. <http://www.cray.com/products/xd1/balance.html>
18. HPCWIRE: LBNL Revamps HPC Performance Measurements. Article 108895, HPCWIRE (Dec. 3, 2004)
19. The Performance Evaluation Research Center (PERC). <http://perc.nersc.gov/main.htm>

20. The Berkeley Benchmarking and Optimization Group (BeBOP). <http://bebop.cs.berkeley.edu/>
21. SIAM: List of Graduate and Undergraduate Programs in Computational Science. http://www.siam.org/cse/cse_programs.htm
22. Unified Parallel C. <http://upc.lbl.gov/>
23. Ginsberg, M.: An Overview of Supercomputing at General Motors Corporation. In: Ames, K.R., Brenner, A.G. (eds.): *Frontiers of Supercomputing II: A National Reassessment*. Volume in the Los Alamos Series in Basic and Applied Sciences, University of California Press, Berkeley (1994) 359-371
24. Ginsberg, M.: Supercomputers Help Auto Manufacturers Decrease Lead Time. In: Redelfs, A. (ed.): *HPC Contributions to Society*. Tabor Griffin Communications, San Diego, (November 1998) 74-79
25. Ahalt, S.C.: Towards a High Performance Computing Economy: Blue Collar Computing™. http://www.osc.edu/hpc/blue_collar/docs/stans_bc_speech_04.pdf
26. Cray, Inc.: The Cray XD1 High Performance Computer: Closing the Gap between Peak and Achievable Performance in High Performance Computing. White Paper, WP-0020404. (2004) <http://www.cray.com/products/systems/xd1/whitepaper.pdf>
27. HPCWIRE: Orion Brings Supercomputing to Your Desktop. Article 108305, HPCWIRE (September 3, 2004)
28. HPCWIRE: NCSA Adds Dell Cluster to Private Sector Resources. Article 108899, HPCWIRE (December 3, 2004)