FORSCHUNGSZENTRUM JÜLICH GmbH
Zentralinstitut für Angewandte Mathematik
D-52425 Jülich, Tel. (02461) 61-6402

Interner Bericht

Feeding Greedies on Meager Roadmaps

Friedel Hoßfeld

FZJ-ZAM-IB-2000-09

August 2000

(letzte Änderung: 18.08.2000)

Preprint: To be published in Proceedings of SGI'2000, Krakow, Poland, 11 - 14 October 2000

Feeding Greedies on Meager Roadmaps

Friedel Hossfeld
Central Institute for Applied Mathematics
John von Neumann Institute for Computing
Research Centre Juelich
Germany

Abstract

For the first time in computing history, in the 1990s we were able to establish a well balanced pyramid ranging from local-area and wide-area broadband networks, via the diversity of workstation and PC platforms in client-server structures supporting cooperative and even realtime computing, up to the layer of medium-sized innovative computer architectures capable to spread computational science and engineering as a key technology over science and industry and also to support methodologically and capacity-wise, finally, the apex of the pyramid, the infrastructure of top-level supercomputers with the mission of an often nation-wide resource. However, already H. H. Goldstine said: "The history of computers is littered with australopithecanes, the deviant apes that anthropologists keep finding: little evolutionary lines that don't lead anywhere", which is especially true for HPC. Also nowadays, long-range planning for supercomputer centres is a tough task which suffers from sudden deaths of innovative product lines and dead-ends in hailed, but finally meager supercomputer roadmaps. Simultaneously, false promises of cheap solutions may seduce funding agencies to stop enhancing supercomputing despite increasing needs of simulations. However, the time span when the performance level of today's Top500 supercomputers will be reached by the as well exponentially growing PC and workstation performance, amounts up to 14 years. This gap offers opportunities for innovative software product development even for countries where native hardware manufacturers are no longer existent who could act as focal points of national software achievements. Definitely, at least in Germany, missing (super)computer hardware industry demands supercomputer centres to play the role of crystallization kernels and attractors of competence in Computational Science and Engineering. The John von Neumann Institute for Computing (NIC) as one of the German Supercomputer Centres – whose primary supercomputing resources (two 512-processor CRAY T3E's, one 12-CPU CRAY T90 and two 16- and 12-CPU CRAY J90's) are provided by the Central Institute for Applied Mathematics (ZAM) at the Research Centre Juelich – acts along these strategic lines by offering computing capabilities to scientists and engineers in universities, research institutes, and industry, and by promoting competence and skills in scientific applications, mathematical algorithms, and visualization methods through research projects and educational programmes. The collective mission of supercomputer centres to satisfy also over the next decade the needs of ever more challenging applications, requires to dissolve stubborn hazes over the supercomputer roadmaps and to pave reliable paths into the foreseeable future.

1 On the Needs of Greedies

Several national initiatives focussed much attention and gave terrific technological and scientific impact to a research and development field which developed in parallel with the tremendous increase and ubiquitous distribution of computer capacity over the past five decades: Although born in the 1940s, it has been named *Computational Science* only in the mid-1980s by the Nobel Prize Winner Kenneth Wilson and has been termed in the 1990s *Computational Science & Engineering* /1/. Computer simulation has grown and established itself as the third category of scientific methodology. This ever-innovating discipline fundamentally supplements and complements theory and experiment, as the two traditional categories of scientific investigation, in a qualitative and quantitative manner while integrating these into the methodological tripod of science and engineering. Being comparable rather with an experimental discipline, Computational Science and Engineering wastly extends the analytical techniques provided by theory and mathematics; today, in a sense, it is synonymous with investigating complex systems. Its main instrument is the supercomputer; its primary technique is computer simulation. Unsolved complex problems in the areas of climate research and weather forecast, chemical reactions and combustion, biochemistry, biology, environment and ecological as well as economic and sociologic systems, order-disorder phenomena in condensed-matter physics, astrophysics and cosmology, quantum chromodynamics, and, in particular, hydrodynamics have been identified as “Grand Challenges” /2/.

The various strategic position papers in the 1980s /3-6/ and the government technology programs in the USA, in Europe, and in Japan in the early 1990s claimed that the timely provision of supercomputers to science and engineering and the ambitious development of innovative supercomputing hardware and software architectures as well as new algorithms and effective programming tools are an urgent research-strategic response to the grand challenges arising from these huge scientific and technological barriers. Scanning the history since the very birthday of Computational Science and Engineering, which may be dated back to 1946 when John von Neumann formulated the strategic program in his famous report on the necessity and future of digital computing together with H. H. Goldstine /7/, at that time complex systems were primarily involved with fluid dynamics. He expected that really efficient high-speed digital computers will “break the stalemate created by the failure of the purely analytical approach to nonlinear problems” and suggested fluid dynamics as a source of problems through which a mathematical penetration into the area of nonlinear partial differential equations could be initiated.

John von Neumann envisioned computer output as providing scientists with those heuristic hints needed in all parts of mathematics for genuine progress and to break the deadlock – the “present stalemate” - in fluid dynamics by giving clues to decisive mathematical ideas. In a sense, his arguments sound very young and familiar. As far as fluid dynamics is concerned, in his John von Neumann Lecture at the SIAM National Meeting in 1981 yet Garett Birkhoff came to the conclusion on the development of fluid dynamics that it be unlikely that computational fluid dynamics (CFD) would become a truly mathematical science in the near future, although computers might soon rival windtunnels in their capabilities; both, however, would be ever essential for research /8-10/. Despite significant progress in CFD and many other areas like condensed matter physics, astrophysics, theoretical chemistry, and quantumchromodynamics as well, it is fair to say that many important problems in fluid dynamics, in particular involved in turbulent systems, are still far from realistic computer modeling and simulation, and new challenges arise from biology, biotechnology, and genomics in the post-genome informatics era /11/ and from non-numerical applications in logistics, operations research, and knowledge processing.

2 On Supercomputing as Strategic Key Technology

The tripod of science and engineering, thus, has proved to provide scientific research and technology with the methodological basis and the instrumental laboratory to effectively approach the solutions of complex problems which are critical to the future of science, technology, and society. It will be a crucial factor for the industry in order to meet the requirements of international economic competition especially in the area of high-tech products. Despite the remarkable investments in research centers and universities in building up supercomputing power and skills and also some sporadic efforts in the industry concerning supercomputing in Europe, it took until the 1990s that the U.S. Government and the European Union as well as several national european governments started non-military strategic support programs like HPCC, HPCN, and HPSC /12-14/. Their goals were also to enhance supercomputing by stimulating the technology transfer from universities and research institutions into industry and by increasing the fraction of the technical community which gets the opportunity to develop the skills required to efficiently access the high-performance computing resources.

In recent years, computer simulation has reached even the highest political level, since, in 1996, the United Nations voted to adopt the Comprehensive Test-Ban Treaty banning all nuclear testing for military and peaceful purposes. Banning physical nuclear testing created a need for full-physical modeling and high-confidence computer simulation and, hence, unprecedented steps in supercomputer power. DoE's *Accelerated Strategic Computing Initiative (ASCI)* /15, 16/ aiming to replace physical nuclear-weapons testing by computer simulation, and NSF's *Partnerships for Advanced Computational Infrastructures* /17/ in the US targeting at the advancement of new computing and communication infrastructures for grid computing /18/ will definitely establish computer simulation as a fundamental methodology in science and engineering. The dedication of the Nobel Prize for Chemistry in 1998 to Computational Chemistry, in addition, confirmed its significance in the scientific community as well as in industry and politics.

The Research Centre Juelich (Forschungszentrum Jülich: FZJ) is one of the largest big-science centres in Europe carried by the German Federal Government and the local State Government of Northrhine-Westfalia, with the function and character of a national research laboratory with highly interdisciplinary research and manifold national and international interactions and cooperations with universities, research institutes, and industry. Its research and engineering activities are focussing on five areas: properties of matter, information technology, energy, environment, and life sciences /19/. Computational Science and Engineering has received here high recognition and priority since the 1960s. In 1987, as kind of a user meta-structure, the "High Performance Computing Centre" (in German named *Höchstleistungsrechenzentrum: HLRZ*) was established and jointly carried by the national labs FZJ, GMD, and DESY. In this then unique initiative, FZJ provided supercomputers to support projects in Computational Science and Engineering of the scientific community all over Germany. In 1998, after GMD had left this cooperation, HLRZ was replaced by the *John von Neumann Institute for Computing (NIC)* /20/ now carried by FZJ and DESY.

The Central Institute for Applied Mathematics (ZAM) at FZJ is responsible for the planning, installation, management, and operation of the central computer systems and of the campus-wide computer networks. Its mission as a central institute and the needs for scientific services at FZJ define ZAM's research and development projects in the fields of mathematics, computing, and communications /21/. ZAM as part of NIC also runs the supercomputer systems as provided to the science community by FZJ. At present, about 150 refereed and approved projects in Computational Science and Engineering are granted via NIC on the supercomputers at ZAM which in 1986 already opened these invaluable resources for scientists in universities and research institutions throughout Germany, as in recent years has been claimed by the German Science Council (Wissenschaftsrat) for the Stuttgart and Munich Supercomputer Centres as well /22/**.** Since the 1960s, ZAM has always run one of the most powerful scientific computing centres in Europe and worldwide as well. Until 1983, ZAM has been a pure IBM shop running always advanced mainframe systems for computer-based research and development. However, ZAM recommended already in 1978 to establish a german parallel computer centre in order to participate in and contribute to the upcoming new supercomputing technology. Yet, it was too early under german circumstances. In 1983, after exploring IBM's early attached vector devices 3838 and utilizing an installed Floating Point System FPS-device, ZAM entered vector supercomputing with the installation of the first CRAY X-MP/2 outside USA. In 1986, a CRAY X-MP/48 was added, and later on a CRAY Y-MP/864 replaced the first X-MP and a CRAY Y-MP/M94 the second one. This sequence of supercomputers successfully consolidated computer simulation at our lab and in the scientific community which we were serving outside FZJ, and much scientific progress was achieved due to this continuous growth of supercomputer power based on the solid ground of a prosperous and technologically extremely competent company Cray Research. The advancements of the Cray operating system Unicos as well as the tremendously increasing volume of valuable application software from the expanding number of customers, users, and software companies provided the fruitful soil for successful and cost-efficient supercomputing in those times.

3 On Hidden Costs along Roads to Dead-Ends

In the early 1990s, Cray vectorsupercomputers with shared memory architecture and proprietary bipolar CPU technology were foreseeable to run into difficulties to provide the giant steps in compute power needed by greedy users, then particularly in physics. Supercomputer architectures with massive parallelism and distributed memory emerged as the future, also more cost-effective, line for the very top end. In Germany, the government-funded *Suprenum* project was targeting already in the late 1980s towards the scientific market niche, however failed to deliver a competitive product. Besides the material loss, this dead-end of a hopeful road caused much damage to the climate of high performance computing in this country, although this project nourished the development of very early scientific know-how and technical competence in the massively parallel computing field at many active places. The Suprenum failure, finally, set a harsh end to all government endeavours in supercomputer hardware architectures, probably for decades. Hence, after installing an Intel iPSC-2 in 1991, an Intel Paragon X-PS with about 150 processing units with i860 chips was acquired by FZJ from the Intel daughter SSD as newcomer in the supercomputer business. The system was installed at the end of 1992, again as the first customer shipment ouside USA. Although we signed a cooperation contract with Intel SSD, it became a nightmare for us – and the Intel SSD experts as well – to make Paragon a reasonably stable system for production runs. Problem management und bug repair as well as time scheduling of software releases was a shere desaster measured at our experiences and standards which had been set through the many years of exceptionally good cooperation with Cray Research. It took us and SSD about two years to bring the failure rate down to a level which was acceptable to allow continuous and finally successful user work. Nobody dared to count, besides the apparent costs of purchasing, operating, and maintaining this system, the hidden costs – on both sides. While Paragon definitely became a useful parallel supercomputer at our site, Intel decided to dissolve SSD and go out of supercomputer business. Thus like all the other Paragon installations world-wide, we were confronted with the dead-end of the road-map for a finally promising product line of massively parallel computer systems for the high end. Almost all investments in application developments seemed to get lost due to the forced new orientation towards alternatives, although we have to admit that, with respect to education and competence of our experts in parallel supercomputers, software, and applications, we learned a lot.

Fortunately, after checking the application fields with their first effort in massive parallel computers, the T3D system, Cray Research stepped forward by the production of a novel parallel computer, the T3E, which was designed to provide the structure, functions and power which we as customers and the increasingly impatient greedies were waiting for, who had already moved there applications onto various parallel platforms like Paragon, Parsytec, KSR, CM-2, and IBM SP elsewhere /23/. However, T3E meant migration to another operating system which carried the well-known and highly acknowledged name *Unicos,* but with the short addendum */mk* for microkernel, which made it a whole new story. After replacing the old vector-supercomputers by a CRAY T90 with 16 CPUs in April 1996 – together with two J90 systems with 16 CPUs (as compute server for interactive and visualization work) and 12 CPUs (as file server to the supercomputers) –, we installed in August 1996 the first CRAY T3E based on alpha-chips with 300 Mhz. After an almost desastrous period of instabilities (and the detection of the malfunction of the stream buffers due to a design error) we were brave enough to close the 3D-torus of the 512 PEs in March 1997 to provide a full-size massively parallel system to our users. In fall 1997, we added a second 512-processor T3E system with 450 Mhz PEs which was upgraded later to 512 PEs with 600 Mhz and 512 MByte memory each which make it still today one of the most attractive powerful machines in the field. This second T3E installation benefited naturally from the intensive experiences with the first T3E. The almost chaotic initial phase of the first installation of that new parallel architecture, however, was much shorter than with the Paragon. Nevertheless, hardware deficiencies and immature software challenged – and sometimes frustrated – the best experts on both sides. These problems turned out to cost valuable time due to delays, system failures, and complex problem analyses and unforeseen investments on both sides to make these systems a productive parallel environment. Since quite some time, the T3Es at our supercomputer centre have become well received and highly estimated. The T3E systems and even still the T90 vectorcomputer are permanently overbooked; the ratio of user-requested to maximally deliverable compute time is swinging between three to five. This demonstrates that it is impossible, despite our extraordinary efforts in funding high performance computing at FZJ and to the benefit of NIC, to feed the greedy users to their level of expectation. Simultaneously, due to fundamental problems in the system, we ran into severe problems with the T90 which showed us that definitely the so far very successful technology of the Cray vectorsupercomputer line had been overstressed by the designers of the T90. We had to reduce the number of active CPUs from 16 to 10 (with two stand-by) to guarantee a relatively stable production operation. However, we are forced to regularly replace broken CPUs in order to achieve the goal of available ten active ones which, of course, puts high stress on the logistics to provide enough spare ones; this is "maintaining to death". Thus, also from a technical point of view this series of vectorsupercomputers has come to an end – another dead-end!

Shortly after these installations, the "merger" of Cray Research and Silicon Graphics happened, and a whole bunch of joint product-line roadmaps was developed and conveyed to the now common customers in iterated versions, which finally were not capable to totally dissolve the strange feelings of the notorious Cray customers on one side and of the quite differently coined SGI customers on the other side. It became clear very early that there was no hope for the users that the T3E product line would survive this merger or even become evolutionarily developed on the emerging SMP-node technology /24/ while further exploiting the potential of the powerful interconnection network and the meanwhile effective operating system. In the end, the divorce of SGI and Cray split the joint roadmaps again into separated bundles which both companies are presently engaged to consolidate and adjust to their potential and to the requirements and trends in the market. Thus, again we, like the many T3E customers worldwide, are confronted – and have to confront our users – with another dead-end of a promising roadmap in supercomputing, the end of T3E. Again a significant amount of investments in manpower will be lost, at least partially if we count the gain in experience and skill in massively parallel computing on the positive side of the account.

4 On "Computer Darwinism"

The obvious and the hidden costs will be high also on the manufacturers side in future, since due to the common trend they believe that they have to jump on the bandwagon to make the Linux operating system the general software platform even in high performance computing. This might be a reasonable management decision in order to enjoy the benefit not only of commodity chips, in particular of the Intel chip mass production, but also of the promises of unifying the whole scale and spectrum of servers from the desktop PC up to the top end of SMP-cluster based supercomputers. However, this might be a still hidden pitfall, because on the ground of Intel chips and the Linux operating system it may become very easy to exchange manufacturers and their products. The individual profiles of the companies then will vanish and, simultaneously, the strong bonds between customers and manufacturers in combining joint efforts to challenge and successfully exploit advanced technologies, architectures, and methodologies to enhance supercomputing will break. Thus, these meager roadmaps may lead high performance computing into a desert. On the other hand, there are the promises of thorough scalability. Anyway, what can be seen today, are frustrating time spans of procrastination until the still left key players will come up with products at the very apex of the performance pyramid, where the greedies are eagerly expecting them. This may offer quite an advantage for a manufacturer with a proprietary chip line and a remarkably powerful operating system. However, it is for sure that we will have to activate all our power in order to dissolve the stubborn hazes over the faintly structured roadmaps that reliable paths can be paved into the future.

During recent years, nearly thirty companies were offering massively parallel systems and others were planning to enter the market with new products, although many experts predicted that the market will not be able to sustain this many vendors /25/. In /23/, a chronological compilation of high performance computer history illuminates the “Darwinistic” forces affecting the supercomputer evolution lines. It demonstrates that the expected shake-out in the computer industry took and still takes place questioning the health and the future potential of this industry in total. Some went out of the parallel computer business – for quite different reasons –, others became just mergers. The dramatic survival battle in the supercomputer industry is also giving severe damage to the users and the customers in the supercomputing arena. Their investments into massively parallel computing may be definitely lost too often, and the establishment of new hardware and software platforms will require unexpected high investments of finances and manpower as well as psychological recovery from the frustration by unfulfilled soap-bubble promises. It is ubcertain if Intel commodities and Linux will become a remedy.

Quite often the critical situation of parallel computing has rigorously been analyzed with respect to the possible negative impacts on the future perspectives and the progress of this scientific discipline. Already Goldstine said /26/ that the history of computing is littered with “australopithecanes”, short computer lines which do not lead anywhere. It is also said /27/ that the history of computing is littered with failed long-term predictions and that it is right in claiming honest answers from the supercomputing arena to burning questions on the seriosity of predictions and promises concerning the reachability of the goals set by the computer industry.

Following in the wake of DoE’s ASCI program, in the forthcoming years powerful new supercomputers will be possibly brought into the market by the manufacturers participating in the high-performance computing race. It seems that only those supercomputer manufacturing companies will have a realistic chance to survive in the years to come who have the potential, capabilities, and favour to get involved in the ASCI or other significant US Government supported programs, with the exception of the Japanese companies. There are expectations that high performance computing at large can benefit from those ASCI-determined products. The large-scale computing facilities in the key research centers and industrial plants world-wide will soon have to surpass the teraflops performance barrier, too. The parallel architectures will be further extended to hierarchically clustered parallel computers mainly based on commodity-chip processors and SMP nodes tying together possibly tens of thousands of processing elements. The simulation of extremely complex systems will also determine future large-scale computing by interconnecting supercomputers of diverse architectures as giant supercomputer complexes. These developments will challenge not only system reliability, availability and serviceability to novel levels, but also interactivity of concurrent algorithms and, in particular, adaptivity, accuracy and stability of parallel numerical methods.

5 On Idiosyncratic Views beyond Teraflops

The requirements of the ASCI program reach far beyond the technology and architectures available in the market /15,16/. Therefore, the supercomputing centers nearly all over the world wait for the “ASCI machines” to get transformed into market so that they can benefit from the technology jumps in high-end computing achieved within the framework of the ASCI program to harness compute-based modeling, simulation, and virtual prototyping. The ASCI goal is to create the leading-edge computational capabilities. ASCI requests for near-time performance in the 10-to-30 Teraflops range until 2001, and for future supercomputer developments enabling 100 Teraflops platforms in the 2004 timeframe.

The Initiative’s applications require a threshold shift of 100 to 1000 times increase in computing capability in order to meet the mission target. The aggregation of new, mainly commodity-based building blocks for massively parallel supercomputers will challenge significant endeavours of integration and scaling technologies which are not currently driven by commercial markets. Therefore, ASCI is undergoing partnerships with various US manufacturers in order to accelerate the development of the supercomputers as required.

As is outlined in the ASCI program, achieving balanced systems at the 10 to 100 Teraflops scale will place stringent requirements on the processor power, the node architecture, the internode interconnect, the I/O systems, and the storage subsystems. Balanced ASCI systems are estimated to scale according to the following approximate ratios:

1 Teraflops peak performance/
1 Terabyte memory size/
50 Terabyte disk storage/
16 Terabyte per second cache bandwidth
3 Terabyte per second memory bandwidth/
0.1 Terabyte per second I/O bandwidth/
10 Gigabyte per second disk bandwidth/
1 Gigabyte per second archival storage bandwidth/
10 Petabyte archival storage.

The concept includes the key attributes: multiple high-performance commodity priced compute nodes, which represent regular commercial product lines and not special-purpose designs; hierarchical memory systems, including cache-only memory architectures and distributed shared memory systems with low-latency high performance memory access; very high performance storage and parallel I/O systems, scalable programming environments and operating systems; a universal programming paradigm; and much more. It is not obvious that, except for the vendors' hardware, the diversity of results of these ASCI developments will be easily and timely available worldwide to the non-ASCI scientific community. Therefore, whenever supercomputer centers outside the ASCI community expect to benefit from these forecast performance steps, significantly enhanced software for distributed operating systems as well as programming environments, tools, and libraries have to be developed and provided in order to enable the users to participate in these technological achievements /24/. Thus, national software initiatives outside ASCI will be urgent to get started now if the emerging technological gap in simulation capabilities and capacities between the ASCI community and the rest of the world shall be kept as narrow as possible.

There is no doubt that the new level of supercomputing then will put significant pressure also on the numerical methods. Unfortunately, so far no signs of preparing a response to the ASCI program can be recognized, for instance, in the european countries, except France where recently the *TERA Project* was started in the military research area.

The complex applications in the ASCI program, but also in the other innovative scientific and industrial environments which rely strongly on Computational Science and Engineering and, thus, on supercomputing, require significant upscaling of massively parallel architectures. The development of hierarchical parallel computers with clustered processing elements, e.g. SMP nodes, is on its way accelerated by ASCI. To meet the performance requirements with today's technology, the 10-to-30 Teraflops machines will need the interconnection of more than 10,000 processing elements, or equivalently of the order of 50 to 1000 compute nodes consisting of 256 to, respectively, 8 parallel processing elements.

With this upscaling, the user will have to face severe problems of drastically reduced system stability and reliability /28/. If it is assumed, referring only to hardware failures, that today the mean time between interrupt (MTBI) of such a (SMP) cluster of 128 processing elements, for instance, is around 500 hours and the total system consists of 100 clusters, the overall MTBI will be 500/100 = 5 hours. If a processor-memory module yields a mean time between failure (MTBF) of three years, the MTBF of the whole ensemble of 10,000 modules will end up with 3/10,000 years which corresponds to 2.6 hours.

Thus, the effective capability of a teraflops (10^{12} operations per second) computer is limited to an uninterrupted run of these few, say in the average three, hours; hence, a corresponding simulation will involve only about 10^{16} operations which is not much measured at the ASCI criteria of 3D real-physics applications.

The situation is getting worse when programming and other software bugs are taken into account, as has been awfully experienced in many large computational projects. Therefore, in order to exploit the performance of such a system, the software – and this means also numerical as well as non-numerical methods – will have to provide very sophisticated check-point and restart mechanisms. Intelligent hot-swap facilities will be crucial in order to cope with those interrupts and failures from the system side to avoid big capacity losses. New redundancy concepts will have to take care of these technological deficiencies.

6 On Prerequisites of Complex System Simulations

Dealing with Computational Science and Engineering problems of ASCI or other Grand Challenge scales, one has to face the fundamental responsibility to verify the results of numerical computer simulations, since these will be used as replacements of experimental explorations as in the ASCI context, but also in other areas like biochemistry, combustion, or crash tests. It can be rigorously shown how easy a numerically illiterate user can be trapped by the error pitfalls of brute force numerical simulations /28, 29/. One can also experience the low level of knowledge about numerical analysis and of numerical technical skills of computer scientists who nevertheless enter the field of Computational Science and Engineering. This refers to the issue of education; therefore also university curricula are challenged. For instance, a deep understanding of the ideas and measures for validation and verification of complex computer simulations as compiled and convincingly discussed in /30/, hence should be a prerequisite of dealing with complex modeling and simulation on high performance computers.

In many application fields, large-scale computer simulations are intended to reduce or even virtually supplant experiments in real world. However, while stimulated by the efficient methods of mathematical statistics developed since the 1920s, optimal experimental design has become a fruitful concept and practice /31/ even, if not primarily, in biological and industrial experimentation, so-called computer experiments almost totally ignored the progress and success of experimental design so far. Taking into account the high costs of complex system simulations on high-end supercomputers and the potential sources of errors and failures involved in such efforts, one should expect that in future an intense focus will be placed on the capabilities of optimal experimental design methods for the further improvement of efficiency and quality of large-scale computer experimentation as well, as has been discussed in /32, 33/.

References

1. Rice, J. R.: IEEE Computational Science & Engineering 1 (1994), 13
2. Special Double Issue: *Grand Challenges to Computational Science*, Future Generation Computer Systems 5 (1989), No. 2/3
3. Committee on Physical, Mathematical, and Engineering Sciences, Federal Coordinating Council for Science, Engineering, and Technology, Grand Challenges 1993: High Performance Computing and Communications, The FY 1993 U.S. Research and Development Program, Office of Science and Technology Policy, Washington, 1992
4. Board on Mathematical Sciences of the National Research Council (USA), The David II Report: Renewing U.S. Mathematics - A Plan for the 1990s, in: Notices of the American Mathematical Society, May/June 1990, 542-546; September 1990, 813-837; October 1990, 984-1004
5. Commission of the European Communities, Report of the EEC Working Group on High-Performance Computing (Chairman: C. Rubbia), February 1991
6. U. Trottenberg et al., Situation und Erfordernisse des wissenschaftlichen Höchstleistungsrechnens in Deutschland – Memorandum zur Initiative High Performance Scientific Computing (HPSC), Februar 1992; published in: Informatik-Spektrum 15 (1992), H. 4, 218
7. J. von Neumann and H. H. Goldstine, On the Principles of Large Scale Computing Machines, in: Collected Works Vol. V, Pergamon Press, 1963, p. 1-32

8. G. Birkhoff, Numerical Fluid Dynamics, SIAM Review 25 (1983), 1-34
9. K. Kuwahara, Flow Simulation on Supercomputers and its Visualization, Intern. J. High Speed Computing 4 (1992), No. 1, 49
10. P. J. Roache, Fundamentals of Computational Fluid Dynamics, Hermosa Publishers, Albuquerque, 1998
11. M. Kanehisa, Post-Genome Informatics, Oxford University Press, 2000
12. The Congress of the United States, Congressional Budget Office, Promoting High-Performance Computing and Communications, Washington, June 1993
13. High-Performance Computing Applications Requirements Group, High-Performance Computing and Networking, Report, European Union, April 1994; – High-Performance Networking Requirements Group, Report, European Union, April 1994
14. Bundesministerium für Forschung und Technologie, Initiative zur Förderung des parallelen Höchstleistungsrechnens in Wissenschaft und Wirtschaft, BMFT, Bonn, Juni 1993
15. ASCI: www.llnl.gov/asci/
16. ASCI-Pathforward: www.llnl.gov/asci-pathforward/
17. L. Smarr, Toward the 21st Century, Comm. ACM 40 (1997), No. 11, 28-32. – Ph. L. Smith, The NSF Partnerships and the Tradition of U.S. Science and Engineering, Comm. ACM 40 (1997), No. 11, 35-37
18. I. Foster and C. Kesselman (eds.), The Grid: Blueprint for a New Computing Infrastructure, Morgan Kaufmann Publishers, San Francisco, 1999
19. FZJ: www.fz-juelich.de
20. NIC: www.fz-juelich.de/nic
21. ZAM: www.fz-juelich.de/zam
22. Empfehlung zur künftigen Nutzung von Höchstleistungsrechnern, Wissenschaftsrat, Mainz, 12. Mai 2000
23. E. Strohmaier et al., The marketplace of high-performance computing, Secial Anniversary Issue (G. Joubert et al., eds.), Parallel Computing 25 (1999), No. 13/14, 1517-1544
24. F. Hossfeld et al., Gekoppelte SMP-Systeme im wissenschaftlich-technischen Hochleistungsrechnen – Status und Entwicklungsbedarf – (GoSMP), Analyse im Auftrag des BMBF, Förderkennzeichen 01 IR 903, Dezember 1999
25. The Superperformance Computing Service, Palo Alto Management Group, Inc., Massively Parallel Processing: Computing for the 90s, SPCS Report 25, Second Edition, Mountain View, California, June 1994
26. N. Macrae, John von Neumann, Pantheon Books, New York, 1992, p. 267
27. J. R. Gurd, Supercomputing: Big Bang or Steady State Growth?, Keynote Address, 7th International Conference on Supercomputing, ICS '93, Tokyo, in: ACM Computer Architecture News 22 (1994), No. 3, 3
28. J. Gustafson, Computational Verifiability and Feasibility of the ASCI Program, IEEE Computational Science & Engineering 5 (1998), No. 1, 36-45
29. A. B. Tucker Jr. (ed.), The Computer Science and Engineering Handbook, CRC Press, Boca Raton, 1997
30. P. J. Roache, Verification and Validation in Computational Science and Engineering, Hermosa Publishers, Albuquerque, 1998
31. A. C. Atkinson, The Usefulness of Optimum Experimental Designs, J. R. Statist. Soc. B 58 (1996), No. 1, 59-76
32. R. A. Bates et al., Experimental Design and Observation for Large Systems, J. R. Statist. Soc. B 58 (1996), No. 1, 77-94
33. Discussion of the Papers by Atkinson and Bates et. al., J. R. Statist. Soc. B 58 (1996), No. 1, 95-111