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

Interner Bericht

Teraflops Computing: A Challenge to Parallel Numerics?

Friedel Hoßfeld

FZJ-ZAM-IB-9901

Januar 1999

(letzte Änderung: 04.01.99)

Preprint:
To be published in Proceedings of "International Conference ACPC '99, Parallel Numerics",
Salzburg, Austria, 16 - 18 February 1999

Teraflops Computing: A Challenge to Parallel Numerics?

F. Hossfeld

John von Neumann Institute for Computing
Central Institute for Applied Mathematics
Research Centre Jülich
f.hossfeld@fz-juelich.de

Abstract

Following in the wake of the Accelerated Strategic Computing Initiative (ASCI) of the US Department of Energy, in the forthcoming years powerful new supercomputers will be brought into the market by the manufacturers participating in the high-performance computing race. Hence, the large-scale computing facilities in the key research centers and industrial plants worldwide will 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. In addition, heterogeneous computing and metacomputing will 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.

1 Computational Science & Engineering and the ASCI and PACI Impact

Crash simulations are the most intensive supercomputer applications in the automobile industry. In pharmaceutical research and industry, molecular modeling has exploited, and will continue to require, considerable computational capabilities. Aerodynamical flow optimization in car and airplane design is still belonging to the "Grand Challenges" /1/. 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 peaceful or military purposes. Banning physical nuclear testing created a need for full-physical modeling and high-confidence computer simulation and, hence, unprecedented steps in supercomputer power, since the Advanced Strategic Computing Initiative (ASCI) of the US Department of Energy (DoE) aims to replace physical nuclear-weapons testing with computer simulations by furthering the simulation technology on a nation-wide basis including the National Laboratories in Livermore and Los Alamos and the Sandia Labs, the industry and the universities /2, 3/. ASCI is accompanied by the NSF Partnership for Advanced Computational Infrastructure (PACI) centered around NCSA in Urbana-Champaign, Illinois, and SDSC, San Diego, in order to renew the US computing infrastructure by creating a National Technology Grid /4, 5/. These initiatives focus much attention and give terrific technological and scientific impact to an R&D 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 recently "Computational Science & Engineering" /6/.

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 sociological systems, order-disorder phenomena in condensed-matter physics, astrophysics and cosmology, quantum chromodynamics, and, in particular, hydrodynamics have been identified as "Grand Challenges".

The various strategic position papers in the 1980s /7-9/ and the government technology programs in the U.S., in Europe, and in Japan in the early 1990s claimed that the timely provision of super-computers 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 /10/. 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 /11/, at that time complex systems were primarily involved with flow 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 mechanics 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 mechanics by giving clues to decisive mathematical ideas. In a sense, his arguments sound very young and familiar. As far as fluid mechanics 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 analytical fluid dynamics that it be unlikely that numerical fluid dynamics 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 /12-14/.

The tripod of science and engineering, thus, has proved to provide scientific research and technology with the stable methodological basis and the instrumental laboratory to effectively approach the solutions of the complex problems which are crucial 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. Academia in the U.S. was pushing Computational Science and Engineering via a series of important strategic reports and initiatives. 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. and European as well as national governments started non-military strategic support programs like HPCC, HPCN, and HPSC /15-17/. 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. The ASCI and PACI initiatives will definitely establish computer simulation as a fundamental methodology in science and engineering; and the dedication of the Nobel Prize for Chemistry 1998 to Computational Chemistry will further support its position in the scientific community as well as in industry and politics.

2 Responses from Computer Technology

For the first time in computing history, we are able today to build a balanced pyramid of computing power in scientific and technical computation /18/. Whereas local area and wide area network systems with medium speed and bandwidth have been built up almost everywhere in research institutions and universities, and high-speed communication with broadband functionality is promoted in the U.S. on a large scale for scientific as well as commercial applications and also in some european countries strong efforts are made to provide the scientific community with broadband communication services, other european countries are still quite far from having access to broadband communications. There is consensus that the backlash in high-speed communications is a severe barrier to establishing a highly efficient nation-wide infrastructure which is capable to provide supercomputer capacity and functionality to the scientific community on a modern scale with transfer opportunities into the industry /19/.

The lack of high-speed communications is certainly an important reason for the retardation of the high-performance computing technology in these countries. In Germany, although still very expensive compared to the US, due to the beneficial activities of the German Research Net (DFN) the Broadband Science Network, B-WiN, is providing communication bandwidths up to 155 megabits per second. For the year 2000, the B-WiN is projected to migrate to gigabit per second bandwidth; the DFN Gigabit Testbeds South and West are pilot projects towards the future G-WiN backbone.

Strategically, for John von Neumann flow-dynamical phenomena have been the primary field where future efforts should have been invested to develop and establish the digital computer – as the ”digital windtunnel” – and, thus, by utilizing numerical methods, activate the mathematical penetration of the whole area of partial differential equations. Since then and up to now partial differential equations have been dominating in the advancement of high-speed computers and in the exploitation of their potential. The general solution methodology for such equations leads via discretization of space and time and via linearization into linear algebra and its numerical concepts and algorithms. The response of computer architecture to these early challenges of PDEs have been the vectorcomputers optimizing vector-pipeline processing and creating the effective instruments of vectorization /20/.

The exploration of the computing potential of the pipelining principle including programming and compiler techniques, tools, operating system functionality, and shared-memory organization and optimization resulted in the efficient arsenal of knowledge and experience about the strengths and the weaknesses of vectorcomputing. The highest class of vectorcomputers, e.g. the Cray, Fujitsu, and NEC systems, are still dominating many of the supercomputing production environments and the practice of Computational Science and Engineering /21/. Certainly, vectorcomputers will further develop in functionality and performance towards hundreds of gigaflops by exploiting the architectural and technological potential and expanding the “weak” parallelism well beyond hundred processors.

Although today the sustained performance of these systems, e.g. the NEC SX-4, Fujitsu VPP-700 and even CRAY T90, turns out to be competitive with mid-sized massively parallel systems for a wast majority of essential algorithms as well as large applications, the technological progress is tending to replace vectorcomputers as the very workhorses of Computational Science and Engineering by massively parallel computers. But workstations, however powerful they are or will become, cannot replace the potential of parallel computers which are basically built upon this technology of powerful microprocessor chips by tying them together via sophisticated broadbanded interconnection networks in order to support massive parallelism. Massively parallel computers are therefore undoubtedly considered as the – only – remedy to the needs of the demanding applications in the Grand Challenge category and maybe yet unrecognized applications which might soon emerge, for instance, from the expanding multimedia field.

Unfortunately, in the early 1990s the manufacturers of massively parallel systems promised that they would be capable to develop and deliver parallel supercomputers in 1995 which be able to reach the magical ”3 T’s” (i.e. 1 Teraflops in execution rate, 1 Terabyte in main memory, and 1 Terabyte/s interconnection bandwidth), thus indicating a revolutionary, rather than evolutionary, step of almost three orders of magnitude beyond the then state-of-the-art supercomputer performance. 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 /22/. In the meanwhile, the expected shake-out in the computer industry 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 in the supercomputing arena. Their investments into massively parallel computing may be definitely lost from time to time and the establishment of a new hardware and software platform will require new investments concerning finances and manpower as well as psychological recovery from the frustration caused by unfulfilled soap-bubble promises.

The critical situation of parallel computing has rigorously been analyzed /23/ with respect to the possible negative impacts on the future perspectives and the progress of this scientific discipline but also on the support which will be expected and requested from the politicians. The report states that "the history of computing is littered with failed long-term predictions"; it is right in claiming honest answers from the supercomputing arena to some burning questions on the seriosity of predictions concerning the reachability of the goals set in particular in the context of those national research initiatives.

A key issue in massively parallel computing is scalability. Parallelizing "dusty" decks from industry is certainly an important task to do in order to increase the acceptance of parallel computing in commercial environments. However, one cannot expect terrific performance gains in many of these programs from porting such originary sequential, in many cases also organically grown, codes to parallel systems. Therefore, scalability often breaks down when the number of parallel nodes is increased beyond sixty-four, sixteen or even eight. Even the benchmark results on naked algorithmic kernels stress the limiting factor of scalability. There is a big discrepancy between the peak rates of massively parallel systems and the sustained performance which can be reached with algorithmic kernels and, even more significantly, with real application programs and software packages /24/. With kernels, the state of the art of massively parallel computers may still deliver, together with a pretty large variance in performance, not more than 10% of the peak rate as sustained performance. This is disappointing. Since, so far, the microprocessor chips have been developed with a different market goal in mind, it is extremely difficult to exploit the performance hidden in the hardware design of these processors via high-level programming languages and compiler techniques /25/; very often this leads to a loss by a factor of five to ten referred to peak performance of the node /26, 27/. But it cannot be accepted as a reasonable software-technological approach to switch back to the very old times of assembler programming to reach reasonable performance levels. Convergence of hardware and compiler design together with the development of valuable programming tools must become the future development strategy.

Another important issue is programming models /28/. While Message Passing – with MPI – is widely and effectively used on distributed memory systems as the only efficiently implemented programming paradigm at present, one can hardly imagine that this programming model will carry all future efforts to introduce massively parallel computing as the overwhelming technology; especially large non-scientific applications will certainly suffer from this obstacle of explicit programming the data communication in message-passing style. Up to now, programming and software technology not only relied on sequential machines, but also on the shared-memory organizational concept. The dominant language in scientific programming has been Fortran with all its strengths and weaknesses. On this language basis, High Performance Fortran (HPF) seems to be an at least temporary platform to implement parallel applications. Urgent extensions towards more functionality are scheduled for HPF-2. There also seems to arise some renaissance of Cray's Multitasking by the OpenMP programming model /29/.

Despite the failure of the first commercially available massively parallel computer system which supported the programming paradigm of the Shared Virtual Memory (SVM), the efforts to explore this programming model should be tranferred into the product strategies of manufacturers. From a user's point of view this SVM paradigm /30/ seems to carry enough potential to overcome fundamental deficiencies which can be experienced with the Message-Passing paradigm, because SVM hides the physical local memories from the programmer and provides a virtual address space organized in pages which demand-driven move across the parallel processors. The advantage of this paradigm may be especially comfortable for the programmer if dealing with irregular data which are inherent in many scientific and engineering applications like, for instance, in the finite-element methods to treat partial differential equations /31/. Since in many cases data access and communication patterns are unknown prior to the parallel execution, data with efficient domain decomposition or applications with remeshing necessities cannot be realized in advance as is required by the Message-Passing model. However, this needs powerful tools to monitor the progress in the parallelization process and strong support on the hardware level which cannot be seen to be available in the near future due to the present lack of interest in this SVM paradigm on the manufacturers' side. In any case, together with genuinely parallel algorithms, powerful and user-friendly programming tools as well as performance-monitoring capabilities are key issues, too /32/.

3 System Aspects of Beyond–Teraflops Computing

The experiences with the strengths and weaknesses of the different architectures available for supercomputer applications in Computational Science and Engineering – the technological obstacles for major performance steps in vector-computing, the large variance in performance for algorithms on different parallel machines, and the very low average sustained performance in massively parallel processing relative to the peak rate, the present or even fundamental limitations to the scalability of systems and applications to reach and exploit massive parallelism – quite naturally lead to the concept of heterogeneous computing which requires the coexistence and cooperation of the different computer architectures. In heterogeneous computing /33-35/, the computational work of – parallel – programs can be split across different computers in order to achieve in total the fastest possible execution, where the individual portions of the work are sent to those computer systems in the heterogeneous ensemble which have been proved to be best for the specific characteristics of the work. This approach could generate results much faster than would be possible on any single system. It also might simplify the programming effort, since program components can be developed using diverse software environments which usually are not available on all machines. Heterogeneous computing is an attractive concept because it takes into account that the individual parallel machines, and vectorcomputers as well, spend much of their time on tasks for which they are unsuited. These effects lead to reduced performance, and also to scalability problems.

On the other hand, it is well known that a user generally invests tremendous efforts in order to extract even small improvements of performance out of an innovative computer system for his specific application, well knowing that the application principally implies a spectrum of heterogeneous requirements which cannot be efficiently satisfied by the single target system he is focussing on just because it is available to him in his specific computing environment. Since the performance of the known supercomputer architectures is a function of the inherent structures of the computations and the data communications involved, it is necessary to discriminate among types of code, algorithms, data, and communications in order to optimize the mapping of tasks onto computer structures.

Researchers in the field of innovative computing believe that there will be no single all-encompassing architecture which will be capable to satisfy heterogeneous requirements with equally optimal performance. Hence, the goal of heterogeneous computing is the efficiency of computation and thereby the effectiveness and cost-effectiveness of both computers and programmers. The price to pay are again grand challenges to respond by hardware and software as well as network designers. Certainly, cumbersome administrative and accounting problems involved in this cross-regional and inter-institutional computing concept will retard the potentiality of getting heterogeneous computing into efficient use. It becomes clear that high-speed networking is the fundamental technical requirement of heterogeneous computing on the way to metacomputing, which brings us back to broadband data communications as the very basis of the "technological pyramid of scientific computing", as discussed earlier.

Despite of taking metacomputing between the supersystems of the National Labs Los Alamos, Livermore, and Sandia into account according to the ASCI-Pathforward plans, the requirements of the ASCI program reach far beyond the available technology and architectures /2, 3/. Also, foreseeable trends in developing and implementing supercomputers fall well below the ASCI requirements. Therefore, the supercomputing centers 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. Thus, ASCI requests for near-time performance in the 10-to-30 Teraflops range in the late 1999 to 2001 timeframe, 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 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 following 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.

4 Meet the Needs of CS&E: Requirements Revisited

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. 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. Unfortunately, so far no signs of preparing a response to the ASCI program can be recognized, for instance, in the european countries. There is no doubt that the new level of supercomputing then will put significant pressure also on the numerical methods.

Besides heterogeneous computing and metacomputing, the complex applications in the ASCI program, but also in the other innovative scientific and industrial environments which rely strongly on Computational Science & Engineering and, thus, on supercomputing, require significant upscaling of massively parallel architectures /36/. The development of hierarchical parallel computers with clustered processing elements, e. g. SMP nodes, is on its way. To meet the ASCI performance requirements, the 10-to-30 Teraflops machines in the late 1999 to 2001 timeframe 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 /37/. 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 the 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.

Dealing with Computational Science & 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 /37, 38/. 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 & Engineering. This refers to the issue of education; therefore also university curricula are challenged.

5 Support to CS&E Competence “Grids”

But, finally, the new parallel computers seem to become a grand challenge for numerical mathematics itself.As has been already mentioned, governmental support programs are lacking today in order to face the grand challenges of methodological research and development in parallel software and parallel numerical methods and non-numerical algorithms as well; there is an urgent need to start an offensive development program. Its goal will not only be to explore fundamental structures in algorithms and numerical methods, which is undoubtedly necessary, too, if we consider the relatively poor results of just porting sequential methods and programs onto parallel computers. But the support program should also create the capabilities to transfer the results without delay into real and operating implementations on the parallel computer platforms emerging at the horizon.

In order to succeed in these efforts, the vitalization of interdisciplinary competence networks is a prerequisite. In this sense, the Grid as established by the US National Science Foundation within the framework of its PACI program can be a guiding model also for other countries, or even regions like Europe. The interconnection and tight cooperation of Supercomputer Centers would certainly enhance such efforts because these centers are acting as active crystallization kernels and attractors within the national or regional competence structures /39/. Thus, their interconnection and cooperation pattern is well suited to provide a stable skeleton for the growth of an interacting competence network for Computational Science & Engineering.

References

1. Special Double Issue: *Grand Challenges to Computational Science*, Future Generation Computer Systems **5** (1989), No. 2&3.
2. ASCI: www.llnl.gov/asci/.
3. ASCI-Pathforward: www.llnl.gov/asci-pathforward/.
4. 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. – D. A. Reed et al., *Distributed Data and Immersive Collab*oration, Comm. ACM 40(1997), No. 11, 39-48. – R. Stevens et al., *From the I-Way to the National Technology Grid*, Comm. ACM 40(1997), No. 11, 51-60. – K. Kennedy et al., *A Nationwide Parallel Computing Environment*, Comm. ACM 40(1997), No. 11, 63-72. – G. J. McRae, *How Application Domains Define Requirements for the Grid*, Comm. ACM 40(1997), No. 11, 75-83. – J. P. Ostriker and M. L. Norman, *Cosmology of the Early Universe Viewed Through the New Infrastructure*, Comm. ACM 40(1997), No. 11, 85-94.
5. I. Foster and C. Kesselman (eds.), *The Grid: Blueprint for a New Computing Infrastructure*, Morgan Kaufmann Publishers, San Francisco, 1999
6. Rice, J. R.: IEEE Computational Science & Engineering **1**(1994), 13
7. Committee on Physical, Mathematical, and Engineering Sciences, Federal Coor-dinating 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.
8. 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.
9. Commission of the European Communities, *Report of the EEC Working Group on High-Performance Computing* (Chairman: C. Rubbia), February 1991.
10. U. Trottenberg et al., *Situation und Erfordernisse des wissenschaftlichen Höchstleistungs-rechnens in Deutschland - Memorandum zur Initiative High Performance Scientific Computing (HPSC)*, Februar 1992; published in: Informatik-Spektrum **15** (1992), H. 4, 218.
11. 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.
12. G. Birkhoff, *Numerical Fluid Dynamics*, SIAM Review **25** (1983), 1.
13. K. Kuwahara, *Flow Simulation on Supercomputers and its Visualization*, Intern. J. High Speed Computing **4** (1992), No. 1, 49.
14. K. Hwang, *Advanced Computer Architecture: Parallelism, Scalability, Pro-grammability*, McGraw-Hill, 1993.
15. The Congress of the United States, Congressional Budget Office, *Promoting High-Performance Computing and Communications*, Washington, June 1993.
16. 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.
17. Bundesministerium für Forschung und Technologie, *Initiative zur Förderung des parallelen Höchstleistungsrechnens in Wissenschaft und Wirtschaft*, BMFT, Bonn, Juni 1993.
18. NSF Blue Ribbon Panel on High Performance Computing, *From Desktop To Teraflop: Exploiting the U.S. Lead in High Performance Computing*, Report, NSF-CISE, 19 October 1993.
19. R. Bayer, *Plädoyer für eine Nationale Informations-Infrastruktur*, Informatik-Spektrum **17** (1994), H. 5, 302.
20. F. Hossfeld, *Vector-Supercomputers*, Parallel Computing **7** (1988), 373.
21. J. J. Dongarra, H.-W. Meuer, and E. Strohmaier (eds.), *TOP500 Supercomputer Sites*, 11th Edition, University of Mannheim, RUM 54/98, and University of Tennessee, UT-CS-98-391, June 18, 1998.

22. 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.
23. 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.
24. D. H. Bailey et al., *NAS Parallel Benchmark Results 3-94; NAS Parallel Benchmarks Update 9-94*, RNR Technical Report RNR-94-006, NASA Ames Research Center, Moffett Field, California, March and September 1994.
25. H. Zima and B. Chapman, *Supercompilers for Parallel and Vector Computers*, ACM Press, 1990.
26. A. Mlynski-Wiese, *Leistungsuntersuchung des iPSC/860 RISC-Knoten-Prozessors: Architekturanalyse und Programmoptimierung*, Zentralinstitut für Angewandte Mathematik, Berichte des Forschungszentrums Jülich Jül-2766, Mai 1993.
27. A. Krumme, *Performance-Analyse des DEC Alpha-Chip: Moderne Compiler-Techniken zur Programmoptimierung*, Zentralinstitut für Angewandte Mathematik, Berichte des Forschungszentrums Jülich Jül-2912, Mai 1994.
28. D. B. Skillicorn and D. Talia, *Models and Languages for Parallel Computation*, ACM Computing Surveys **30**(1998), No. 2, 123-169.
29. L. Dagum and R. Menon, *OpenMP: An Industry-Standard API for Shared-Memory Programming*, IEEE CS&E **5**(1998), No. 1, 46-55.
30. K. Li, *Shared Virtual Memory on Loosely Coupled Multiprocessors*, PhD Thesis, Yale University, September 1986.
31. R. Berrendorf, M. Gerndt, Z. Lahjomri and Thierry Priol, *A Comparison of Shared Virtual Memory and Message Passing Programming Techniques Based on a Finite Element Application*, in: B. Buchberger and J. Volkert (eds.), *Parallel Processing: CONPAR 94 - VAPPVI*, Lecture Notes in Computer Science Vol. 854, Springer Verlag, 461-472.
32. W. E. Nagel and A. Arnold, *Performance Visualization of Parallel Programs: The PARvis Environment*, Proceedings 1994 Intel Supercomputer Users Group (ISUG) Conference, 24-31; Technical Report CCSF-47, Concurrent Supercomputing Consortium, Pasadena, California, May 1994.
33. Ch. Pancake, *Area Editor's Message: High-Performance Computing - Lighting the Way*, IEEE Computer **26** (1993), No. 6, 6.
34. R. F. Freund and H. J. Siegel, *Guest Editors' Introduction: Heterogeneous Processing*, IEEE Computer **26** (1993), No. 6, 13.
35. A. A. Khokhar et al., *Heterogeneous Computing: Challenges and Opportunities*, IEEE Computer **26** (1993), No. 6, 18.
36. L. Smarr, *Supercomputers: Directions in Technology, Architecture and Applications*, in: H.-W. Meuer (Hrsg.), *Supercomputer 1998, Anwendungen, Architekturen, Trends*, K. G. Saur, München, 1998, p. 2-16.
37. J. Gustafson, *Computational Verifiability and Feasibility of the ASCI Program*, IEEE CS&E **5**(1998), No. 1, 36-45.
38. A. B. Tucker Jr. (ed.), *The Computer Science and Engineering Handbook*, CRC Press, Boca Raton, 1997.
39. F. Hossfeld, *Verbund der Supercomputer-Zentren in Deutschland – Ansichten, Einsichten, Aussichten –*, in: H.-W. Meuer (Hrsg.), *Supercomputer 1998, Anwendungen, Architekturen, Trends*, K. G. Baur, München, p. 160-171.