APPENDIX A: BIBLIOGRAPHY

- -

LEGEND

YEAR

Title of article or book

Author(s) Reference source/publisher, etc. Other data on reference

Annotation -- summary, comments, observations, etc.

Reference Key Words:

- -

Notes:

1) Generally, references older than 1990 have been excluded.

2) This search (and study) included Petaflops and Exaflops computers as a goal and the search attempted to cover both. The literature found was not very extensive beyond about ten years hence and thus practically nothing was found on Exaflops computing – the factor of one thousand beyond Petaflops. At this time, at least, some physical limits of the present technology appear to be around one PetaFLOPS, with Exaflops beyond many present technical (and now economic) boundaries.

3) Some of the references use the term "Grand Challenges" which has been defined as ".....something useful that is possible within the next ten years". Thus in the context of these references, this implies the time period up to about 2003. Most of the references associate the term with a rate of one TeraFLOPS.

4) This bibliography includes Key Words assigned to each reference. These then are used as Categories to cross-index the references by topic or subject. An index of the Key Words as well as a listing of books by Categories is included as a part of this bibliography. The Key words were assigned by the author (Preston) as the references were included.

5) Each book reference has included an Annotation, written by the author (Preston) to aid him in analysis of the references. The annotations contain personal impressions *(usually in italics)* and no attempt was made to find published annotations.

6) Each reference has been assigned a short abbreviation that may be used to identify the specific reference elsewhere in text or charts. This consist of the first four letters of the author's name, followed by the year, thus: {Name96}

7) This Bibliography is very useful in the on-line form to search for information. A copy on a floppy, of via E-mail (etc.) may be available on request.

8) The assistance of David Adams, of the NASA Langley Research Center Library staff, in performing library reference searches is gratefully acknowledged.

- -

1989

Grand Challenges to Computational Science

Special Issue of FGCS (Future Generations Computer Systems) Vol. 5 No. 2-3 North Holland (Sept '89) Based on Conference on Grand Challenges held in Molokai, HI 3-6 Jan. 1989

Fourteen papers on Challenges to supercomputing. Significant papers are covered in the items immediately below. Most of the papers deal with results and do not cover the computational requirements.

FGCS89 Key Words: CHALLENGES

Grand Challenges to Computational Science

Wilson, K.G. Future Generations Computer Systems Vol. 5 No. 2-3 Page 171-89 North Holland (Sept '89)

This is a significant paper by the man that first applied the term Grand Challenges. He points out that research has been "crippled by compromises due to shortcomings in computing power, algorithms, software". The Grand Challenges should be defined as opportunities to overcome the compromises. "Supercomputer and other tools of computational science are in a quite primitive state now, compared with what they will become in the future." He defines three critical roles for supercomputers. 1) They can see into the future which instruments can't - example: weather forecasting. 2) They can resolve shorter time scales than experimental equipment can. 3) They can replace experiments as well as reach areas not accessible to experiment. To solve simple molecular problems requires "..... from 10^{12} to 10^{300} times the capabilities of current supercomputers" and protein folding simulations are "....limited to roughly 10^{-12} of the time period needed". He suggests that algorithm improvement could account for half of the needed increase and hardware technology the other half. Fortran to him is a fundamental long-term barrier to progress.

Wils89 Key Words: REQUIREMENTS; CHALLENGES

Computational Challenges in Aerospace

Peterson, V. Future Generations Computer Systems Vol. 5 No. 2-3 Page 243-258 North Holland (Sept '89)

An excellent analysis of the topic with speed an memory required for four major areas (and many subdivisions). Paper makes a case for needing in excess of ExaFLOPS (10^{18} FLOPs) and memories in excess of Pentawords (10^{15} words). Paper suggests that we proceed to the challenges in steps and that a teraflop machine can produce useful results. Also emphasizes that algorithm progress could match machine progress and shorten the time to the exa- and penta- time period. For some specifics, the following values were presented in both the graphs and text. The graphs make clear that problems cover a range of at least a factor of 1000 and sometimes more.

- Vision simulation -- 10^{14} FLOPs and 10^{12} words
- "Visional" guidance -- 10^{10} FLOPs and 10^{11} words
- Computational chemistry -- from 10^{11} to 10^{21} FLOPs and from 10^{9} to 10^{15} words
- Turbulence physics -- 10^{9} to 10^{20} FLOPs (200 hr runs) and from 10^{9} to 10^{16} words
- Aerodynamic calculations -- from 10^{10} to 10^{17} FLOPs and from 10^{8} to 10^{13} words
- Grand Challenge in aerodynamics:
  - Combined analysis -- 10^{15} FLOPs and 10^{10} words
  - Design -- 10^{18} FLOPs and 10^{10} words

Pete89 Key Words: REQUIREMENTS

Exploiting Parallelism in Computational Science

Johnson, G. Future Generations Computer Systems Vol. 5 No. 2-3 Page 319-337
North Holland (Sept '89)

Predicts that for the short-term, parallel processing has the greatest potential to close the gap between requirements and capability. At the time of the paper, many new areas were taking up the use of supercomputers and this was becoming a third mode of analysis to augment the traditional theoretical and experimental approaches in science and engineering.
Thirteen areas were examined and requirements presented. Five of these used data from NASA which is superseded by data from Peterson (in above reference); some of the other data is given below:

Aerospace structural analysis -- from 10^8 to 10^{13} FLOPs and from 10^7 to 10^{10} words
Chemistry -- from 10^7 to 10^{12} FLOPs and from 10^7 to 10^{11} words
Hypervelocity impact -- 10^{11} to 10^{12} FLOPs and 10^{10} to 10^{11} words
Large space structures -- 10^9 to 10^{14} FLOPs (200 hr runs) and from 10^6 to 10^{10} words
Laser Optics – from 10^8 to 10^{13} FLOPs and from 10^8 to 10^{11} words
Magnetohydrodynamics – 10^{14} FLOPs and 10^{12} words
Visualization – from 10^9 to 10^{12} FLOPs and from 10^8 to 10^{10} words
Strategic defense simulation – from 10^8 to 10^{11} FLOPs and from 10^8 to 10^{12} words

John89 Key Words: REQUIREMENTS

Some Analysis on Supercomputing Future Requirements -- Speed and Memory

Preston, Frank Report No. 9 Nov. 9, 1989 NASA - Langley Research Center

A review paper that pointed up the need for changes in direction in supercomputing uses and acquisition. The data came from Ames experience and unpublished figures which later became a part of the P. Kutler reference below. *(The data in this paper was largely not accepted at Langley.)*

Pres89 Key Words: REQUIREMENTS

Computational Fluid Dynamics - Current Capabilities and Directions for the Future

Kutler, Paul Proceedings Supercomputing '89 Nov. 1989 Page 113
ACM No. 415892 IEEE Computer Society 2021

This paper covers computational aerodynamics, some aircraft structures, propulsion, and space and earth modeling. This is one of the early papers on a grand challenge and one of the few that shows requirements of more than 1,000 TeraFLOPS (up to 10^{16} ops)

Kutl89 Key Words: REQUIREMENTS

1990

Physics Goals of the US QCD Teraflop Project

DeTar, C. Proceedings of the 25th International Conference on High Energy Physics
Publisher: World Scientific Page 1159-1161 Conference date: Aug. 1990

The goals of the project are discussed briefly. For a constrained solution (quenched QCD), which grows at a^{-6} (where a is resolution) a teraflop machine can begin to yield results. The full problem grows at a^{-8} and 1,000 teraflops does not appear to be sufficient with the algorithms used.

DeTa90 Key Words: REQUIREMENTS

An Examination of the Role of Parallel Computers in Scientific Supercomputing – 1991 - 1999

Preston, Frank Report No. 10 July 1990 NASA - Langley Research Ctr.

Report is dated but can serve as a benchmark for projections that are now at least six years old.

Pres(a)90 Key Words: FORECASTS

Impact of Larger Jobs on Output in TeraFLOP Computer ERA

Preston, Frank Report No. 12 Dec. 30, 1990 NASA - Langley Research Ctr.

The report looks at the storage problem as the complexity of the simulation increases. Considered in this is the growth in computations as a function of the number of nodes or elements in the problem are increased. Because the output per day does not increase as fast as the computations increase, the effect on storage is somewhat reduced. There is a greater impact on the computational speed required if the solution time is to be held fixed as job complexity grows.

Pres(b)90 Key Words: REQUIREMENTS

Semiconductor Manufacturing in the 21st Century: Capital Investment VS Technical Innovation

Losleben, Paul Stanford U. IEEE/CHMT Symposium 1990

Highlight the rapid growth in the cost of production facilities as resolution limits are approached. Makes point that the world can't afford many such plants.

Losl90 Key Words: SEMICONDUCTORS; MANUFACTURING

1991

Federal High Performance Computing and Communications Program. The Department of Energy Component

Office of Energy Research – Dept. of Energy DOE/ER-0489 June 1991 51 pg.

This document describes the DOE High Performance Computing and Communications (HPCC) Program. DOE has responsibility specifically for several of the Grand Challenges. They cite several simulations that generate 10+ TeraBytes of output for one problem. For chemical rate problems they list a requirement for 10,000 TeraFLOPS. For a QCD problem, the input lattice description requires 400 GigaWords (or about 3 TeraBytes of memory). For a global climate model the data base is from hundreds to thousands of TeraBytes.

DoE_91 Key Words: REQUIREMENTS

Lattice QCD in the 1990s

DeTar, C. AIP Conference Proceedings N260 Conference date 1991 Publication date 1992

The "QCD Teraflop Project" is described. Author states that the numerical approach to field theory is the method of choice because analytic techniques have failed. It is estimated that a teraflop machine will be suited to reaching the needed lattice spacing without quarks. Approximately 100 TeraFlops will be needed to include quarks. The QCD TeraFlop Collaboration (see paper in 1990) proposed to have the DOE procure (from Thinking Machines) a special purpose machine capable of a sustained teraflop and to go into operation in 1994!! The architecture is to be tuned to the requirements of QCD work.

DeTa91 Key Words: APPLICATIONS

1992

The NASA Computational Aerosciences Program – Toward TeraFLOPS Computing

Holst, T., Salas, M., Claus, R. AIAA - 92-0558 From 30th Aerospace Sciences Meeting of Jan. 6-9 1992 (Reno, NV)

Paper reports on the Ames, Langley, and Lewis HPCC program. A target of 1 TeraFLOPS sustained (10 peak) in AD2000 is expressed but due to funding constraints, a "scaled demo" of 100 Gigaflops is envisioned. *(Even this is a more rapid increase in speed than has been achieved for any time in the past 50 years. FSP)* A technical background is presented that ends in defining a baseline computation for today's technology. [Baseline requires 600K grid points, 25 megawords and 10 hours of cpu time on a 160 MFLOPS machine.] To handle this problem for a complex flow they estimate 30.4 hours on the 100 GFLOP machine. They state that to handle a full unsteady simulation on a complete engine to be beyond the scope of a teraflops machine. For a multi-disciplinary simulation — not including the engine — they estimate about 2,000 times the baseline time – for the simple aeroscience model. All estimates are worded as the "lower bounds" (*One would infer that a 1,000 TeraFLOPS is insufficient for a full aircraft simulation. FSP)*

Holz92 Key Words: REQUIREMENTS

Supercomputers and CFD (An Industry Perspective)

Miel, George Aerospace America (AIAA) Jan 1992 Page 32-35

States that software is 5 to 10 years behind the hardware technology. Also quotes Grosh's Law: "The cost of a computer in a given technology is proportional to the square root of its speed" *[However, this is not applicable to the shift to massively parallel and in any case, buyers are generally unable to fund even this cost growth. FSP]*

Miel92 Key Words: SOFTWARE; ECONOMICS

Semiconductor Lithography for the Next Millennium

Geppert, L (Assoc. Ed. IEEE) IEEE Spectrum April 1996 pg. 33-38

The Industry Roadmap, widely quoted in this paper, is based upon the assumption of continuing improvements of 4X every three years to 2010 (or 5+ generations). This takes the resolution beyond the now known limits of visual and X-ray lithography. The Roadmap makes this assumption clear but this paper fails to do so. *See the National Technology Roadmap for Semiconductors (1995).*

Gepp92 Key Words: SEMICONDUCTORS; TECHNOLOGY

Future Directions in Computing and CFD

Bailey, R. and Simon, H. AIAA paper 92-2734 From 10th Applied Aerodynamics Conference of June 22-24, '92 Page 149-160

Authors conclude that almost all Grand Challenge computing will be carried out on massively parallel computers. For a CFD simulation they develop the need for 60 GFLOPS sustained performance and approximately 960 gigaflops for a multidisciplinary analysis. Their reference #6 forecasts a single chip with a 1 GFLOPS speed by AD 2000. Then they go on to explain that an explicit algorithm may run 10 times faster on a parallel machine and yet take 100 times longer to converge. *{This demonstrates that MFLOPS is not a complete measure and algorithms are critical to parallel performance. FSP}*

Bail92 Key Words: FORECASTS

Physics Leads the Way at Edinburgh HPC Facilities

Bowler, Chantler, Heggie, Kenway, Tildesley, Trew, Wallace Computers in Physics Vol. 6 No. 4 July/Aug. 1992 Page 334-338

For astrophysics the N-body simulation of star clusters is Order N cubed which is very computationally demanding. N squared comes from the number of interactions and the third because the time scale. Gravitational interactions pose special problems for parallel computers.

Bowl92 Key Words: APPLICATIONS

Retire Fortran? A Debate Rekindled

Cann, D.C. Comm. of the ACM Aug. 1992 Vol. 35 No. 8 pg. 81-89

Paper makes a good case that our languages need to be machine independent and that FORTRAN fails to do so for a number or reasons, especially by using COMMON. A modern language should be functional whereas FORTRAN is imperative. Parallelism must be implicit and results must be deterministic. The author suggests SIZAL which he shows performs significantly better than FORTRAN on identical problems. The more complex the problem, the greater the advantage — by a factor of up to 5X. *{In my view, SISAL is not appropriate because it requires experts to write it. FSP}*

Cann92 Key Words: SOFTWARE

Ultracomputers: A Teraflop Before Its Time

Bell, Gordon Communications of the ACM Aug. 1992 Vol. 24 #8 pg. 27-45

This paper is strangely "out of character" for Gordon. His projections are not as wildly optimistic and his emphasis is more balanced rather than just addressing peak hardware speeds. He recognizes that although we may meet his 1997 date for a teraflop, users will not be able to afford to buy them. He emphasizes the need for a balanced approach, more software effort, better training, etc. His projection for the successful machines for 1995 were wrong – witness the failure of the Connection Machines and Intel dropping out of the market. He suggests a petaflops system by 2001 at $500 million and we had "better wait for lower prices". A few other quotes are noteworthy: "HPCC is a bureaucrat's dream"; "peak performance (rating) has replaced reality as a selection criteria"; "programming may not change"; *despite his view that is must FSP.*

Bell92 Key Words: TERAFLOPS

Modeling Reality

Grosman, M. IEEE Spectrum Sept. 1992 Page 56-60

Some examples of problems too big for present-day supercomputers were given. A biochemical reaction can now be simulated for a time measured in a fraction of a nanosecond whereas a fraction of a microsecond is required (1,000 times longer). The Human Genome file is expected to require data search and matching in a 4 terabyte library. The Magellan radar scan of Venus resulted in sending over 3 terabits of data to the earth. A recent global ocean model required 3,000 CPU hours of Cray-2 time for one simulation. An ozone-depletion model requires 10 supercomputer hours per day of activity whereas many years of activity must be studied to understand the problem. Researchers using the National Science Foundation supercomputers can on average get 24 hours of CPU time per year. As a result, most have turned to using high-performance PCs and workstations.

Gros92 Key Words: REQUIREMENTS

Summary of the Report of the NSF-Sponsored Purdue Workshop on Grand Challenges in Computer Architecture for the Support of High Performance Computing

Siegel, H.J. and Abraham, S. (Workshop Co-Chairs) IEEE Computer Society 1992
Proceedings of the Fourth Symposium on the Frontiers of Massively Parallel Computation (Oct. 19-21, 1992)

The title is in reality an abstract of the paper. The emphasis is on the grand challenge demands on the available architectures. The meat of the meeting is contained in the committees report.

The summary points out that teraflops computers are not sufficient for the grand challenges and points toward "useable peta-ops (10^{15} ops)". It also characterizes presently available massively parallel machines as "dismal" measured in comparison with effective performance and ease of use.

Sieg92 Key Words: CHALLENGES

1993

Modeling the World's Climate

Perry, T.S. IEEE Spectrum July '93 Pages 33-42

It is impossible to separate Climate Modeling from the Ocean Circulation. *(as is now done in listing of Grand Challenges. FSP)* Today's supercomputers frustrate researchers with their crudity. A 100 year simulation run takes 1,000 Cray Y-MP hours at course resolution. Desired resolution would run 125,000 + times as long. Researchers would like to be able to carry out one experiment in less than their lifetime. Local climate modeling should include topography, ground and surface water terrestrial ecosystems, marine biochemistry, atmospheric chemistry and clouds – all not now included. Chemical reactions increase computer demands exponentially.

Perr93 Key Words: APPLICATIONS; REQUIREMENTS

Supercomputing and the Transformation of Science

Kaufmann, W.J and Smarr, L.L. Scientific American Library 1993 238 pg.

A good basic book on supercomputing and the changing way research, analysis, and design are now being done. Some background on the methods used that are applicable to the Grand Challenges but nothing on computational requirements. The book forms an excellent introduction to the computational aspects of scientific computing, including the Grand Challenges. The authors do state that a "teraflop is too slow" and that the "exponential increase is supercomputer speed will continue well into the next century".

Kauf93 Key Words: BACKGROUND

1994

High-Performance I/O for Massively Parallel Computers

del Rosario, Juan M., Choudhary, Alok N. IEEE Computer March 1994 pg. 59-67

Discussion of requirements, performance, problems and prospects for I/O which includes networks, operating systems, memory, storage, and mapping of distributed data. A section on Grand Challenge applications lists memory, file, and bandwidth demands for four of the challenges. The greatest for storage was air-quality modeling that generates 1,000 TeraBytes per application studied. Communication bandwidths up to 2 GigaBytes/second can be demanded.

DelR94 Key Words: COMMUNICATIONS

Looking at the Grand Challenges Computing Requirements

Preston, Frank Report No. 21 March 31, 1994 NASA - Langley Research Ctr.

The report looks at the requirements for Grand Challenge problems that have been directly related to the next ten years or TeraFLOPS computers. Some of these Challenges actually require Petaflops and Exaflops computer systems. The bibliography in this report can supplement this Bibliography.

Pres94 Key Words: CHALLENGES

Molecular Computations to Combinatorial Problems
Adelman, L. M. Science Nov. 11, 1994 Vol. 266 pg. 1021-1023

Paper reports a demonstration of a potential speed advantage (over 1992) of 10^6 ops/sec., a potential energy improvement of 10^{10}, and a potential storage advantage of 10^{12}. The scheme is not applicable to efficient use for some problems.

Adel94 Key Words: PLANS

Proceedings of Workshop on Physics and Computation '94
IEEE Computer Society Nov. 17-20, 1994 at Dallas

The following topics were dealt with: nanoelectronics computing technology; computing with quantum devices; nanoelectronics parallel architectures; physics as combinatorial computation; Reversible logic; cellular automata; statistical mechanics and information; entropy and information; and parallel physics computations.

IEEE(a)94 Key Words: APPLICATIONS

The Stabilization of Quantum Computations
Berthiaume, A., Deutsch, D., Jozsa, R. Proceedings of Workshop on Physics and Computations '94 pub. by IEEE Computer Society

The problems with quantum computations, compared with classical, is that interference phenomena present new challenges and error detection is likewise an obstacle.

Bert94 Key Words: APPLICATIONS

Software Crisis for the Emerging MPP Industry
Gannon, Dennis Proceedings of the 8th International Parallel Processing Symposium 1994 IEEE Computer Society

"Unless we create a revolution in the way we think about and build software for parallel systems, there will be no way the MPP industry will survive." This opinion has been echoed by many writers in many publications.

Gann94 Keywords: SOFTWARE

1995

Toward Point One
Stix, Gary (editor) Scientific American Vol. 272 No. 2 Feb. 1995 pg. 90-95

Emphasis is on the great cost of facilities below 0.2 micron X-ray production in 2007 and later.

Stix95 Key Words: SEMICONDUCTORS

PetaFLOPS Frontier Workshop (February 1995)

Notes of Conference sponsored by NASA /Goddard for the HPCC Office and held in McClean VA on Feb. 6, 1995. The conference proceedings include eighteen papers with some as complete papers and some with only an abstract.

Significant papers are covered in the items immediately below.

NASA95 Key Words: PETAFLOPS

Processors-In-Memory (PIM) Chip Architectures for PetaFLOPS Computing

Kogge, P.M. (Extended Abstract only available)

....avoids silicon consuming speedup techniques, such as cache and extensive pipelining, ...upwards of 10X more efficient. The major economic constraint on reaching a petaflop system is the cost of the memory system to support it.

Kogg95 Key Words: SEMICONDUCTORS

Strategic Applications for PetaFLOPS Computational Systems

Stevens, R.L., and Taylor, V.E.

Outlines characteristics of a petaflops computing system and predicts availability by 2012.

Stev95 Key Words: FORECASTS

Taming Massive Parallelism: The Prospects of Opto-Electronic CRCW Shared Memory

Lukowicz, P., and Tichy, W.R.

Proposes a design for a scalable, concurrent read/write shared memory.

Luko95 Key Words: TECHNOLOGY

Some Applications Demonstrating the Existing Need for Petaflops/PetaOps Computing in Biomedical Research

Maizel, J. V.

A full comparison of the human and mouse genomes will require 10^18 to 10^19 operations. Protein structure prediction is of the order n^N where n is in the range of dozens and N is in the range of 1,000. For the design of drugs, a 10^12 increase in speed is needed over a YMP. To simulate a single heart beat on a C90 requires one cpu-week and 50 million words.

Maiz95 Key Words: REQUIREMENTS

Protein-Based Computers

Birge, R. R. Scientific American March 1995 pg. 90-95

This technology promises very high "resolution" leading toward dense memories and storage, and small computers. The author talks about approaching molecular size by the year 2030. Currently he envisions hybrid computers using electrical interconnection of modules. *Although perhaps technical feasibility is likely, we are a long way from something that can leave a laboratory. FSP*

Birg95 Key Words: TECHNOLOGY

The National Technology Roadmap for Semiconductors

Semiconductor Industry Association (SIA) April 1995 167 pages

A wealth of information from the semiconductor industry which skirts proprietary constraints. The major assumption is that Moore's Law will be followed up to at least 2010 AD. This takes the performance beyond the now accepted limits of resolution of X-rays (to 0.07 microns). Thus this Roadmap is not as much a forecast as an examination of what will be required to achieve the planned 4X improvement every three years. Emphasis is mainly on silicon CMOS. Continued cost improvements following past history, of 30% reduction per year per function, is assumed *{but not likely FSP}.* The Roadmap is long in goals and short on possible solutions.

See also the references to Solid State (Jan. 1996 IEEE Spectrum) and Semiconductor Lithography (April 1996 IEEE Spectrum) that quote from this report.

SIA_95 Key Words: SEMICONDUCTORS

Applications and Algorithm Challenges for PetaFlops Computing

Workshop at Bodega Bay, CA April 19-23, 1995 Proceedings of this Workshop do not seem to have been published. On-line documentation is available on the Argonne Home Page — PetaFLOPS On-line Resources — http://www.mcs.anl.gov/petaflops

Information covered in this can be found in later and more detail in these references: Fost96 Petaflops System Workshops ("PAWS" '96 and "PetaSoft" '96).

Argo95 Key Words: APPLICATIONS, ALGORITHMS

Realizable Universal Quantum Logic Gates

Sleator, T. and Weinfurter, H. Physical Review Letters Vol. 74 May 15, 1995 pg. 4087-90

"Quantum computing can be implemented with linear ion traps." "We believe they can be built with present or planned technology."

Slea95 Key Words: TECHNOLOGY

Quantum Computations with Cold Trapped Ions

Cirac, J.I. and Zoller, P. Physical Review Letters Vol. 74 May 15, 1995 pg. 4091-94

"....can build 2-bit quantum gate with present technology." From such a gate, quantum computational networks are possible. "We conclude that extended networks would be exceedingly difficult to build."

Cira95 Key Words: TECHNOLOGY

3-D Simulation of Deep-Submicron Devices

Zhou, Jing-Rong and Ferry, D.K. IEEE Computational Science and Engineering Summer 1995 pg. 30-37

Covers memory with density of 10^9 devices per chip and resolution in the 0.15 to 0.18 micron range. They state that the effect of impurities will be another limit on the downsizing of chips.

Zhou95 Key Words: SEMICONDUCTORS

Quantum-Mechanical Computers

Lloyd, Seth Scientific American Vol. 273 No. 4 Oct. 1995 pg. 140-145

"New technology must replace what we have." Logic devices demonstrate that NOT, COPY, and AND gates (all that is needed) can be made. "Quantum computers have been theoretically possible for about 20 years.

In Science Vol. 2611 pg. 1569-1571 (Sept 17, 1993 Lloyd describes physical systems that can function as quantum computers.

Lloy95 Key Words: TECHNOLOGY

Quantum Computation

DiVincenzo, David P. Science, Vol. 270, no. 5234 (Oct. 13, 1995) pg. 255-261

Report on "active work underway" at IBM. Quantum AND computer equals probabilities result. Possible potential but "...to think about it might be premature". Some limited special purpose algorithms could be implemented. No time projections included in paper. This article is heavy in physics and contains a good bibliography.

DiVi95 Key Words: TECHNOLOGY

AFM Fabrication of Sub-10-Nanometer Metal-Oxide Devices with in Situ Control of Electrical Properties.

Snow, E.S. & Campbell, P.M. Science, vol. 270, no. 5242 (Dec. 8, 1995), pg. 1649-1641

Sub 10 nanometer fabrication *{some lithography now reaches 10 nm FSP}* is possible using real-time control during processing – measuring parameters or device properties.

Snow95 Key Words: SEMICONDUCTORS

Device Fabrication by Scanned Probe Oxidation

Dagata, J.A. Science, Vol. 270, no. 5242 (Dec. 8, 1995), pg. 1625-1626

Using Scanned Probe Microscopy, at NIST, and direct oxidation techniques to form thick oxides (lines) 1 to 3 nanometers thick. Pattern can be used as an etch mask. Fairly simple to implement for research on small prototype devices, but costly and slow for production (now). Multiple arrays do "parallel" fabrication.

Daga95 Key Words: SEMICONDUCTORS

Enabling Technologies For Petaflops Computing

Sterling, T., Messina, P., and Smith, P. The MIT Press 1995 pg. 173

By the year 2000 it is estimated that the aggregate capacity of all of the computers in the U.S. will be about 1 PetaFLOPS. Thus a legitimate question to raise is, "Why would we want that much power in a single system?" This book attempts to consider this question and the potential of reaching this in two decades. This reference contains considerable background information, much general information on possibilities, and a limited amount of specific projections.

Ster95 Key Words: BACKGROUND; FORECASTS; TECHNOLOGY

The Computer in the 21st Century

Scientific American Staff Special Issue of Scientific American Pub. 1995 pg. 196

Twenty-two (22) articles reprinted from and updated from Scientific American issues of the prior several years, written by ranking "names" in the computer, technologies, civil and business world. Presents a good, broad overview of current status and future potential. The articles are long on text and short on specifics on forecasts.

Scie95 Key Words: BACKGROUND; FORECASTS

The Road Ahead

Gates, Bill Viking 1995 ISBN 0-670-77289-5 HE7572.U6G38

Bill Gates' (speaking for Microsoft), view of the future of personal computer, communications, entertainment in the home, etc. He major points are: 1) Continued technical growth at an accelerated rate; 2) Continued reductions in prices at the present rate; 3) High bandwidth in the US and to homes will be assured, accompanied with dramatic cost reductions; 4) Asynchronous communication and applications will become much more widespread; 5) major changes in how we live, act, buy & sell, etc. 6) More is better? *(Perhaps. FSP)*

I believe he is too optimistic on what will be available at affordable prices, and when new "killer" applications will be introduced. He projects items by 2005 that I think will take longer than getting Petaflops computers. FSP

Gate95 Key Words: LONG RANGE PLANNING

1996

Solid State - Technology 1996

Semiconductor Industry Association -- Quoted in IEEE SPECTRUM Jan. '96 page 53

The Industry projection is based upon Moore's Law which has held true for the past 20+ years. However, they project the feature size to smaller than is now regarded as the limit of X-rays, without any explanation in this paper. *See the National Technology Roadmap for Semiconductors (1995).*

Semi96 Key Words: SEMICONDUCTORS

Technology and Economics in the Semiconductor Industry

Hutcheson, G.D. & Hutcheson, J.D. Scientific American Vol. 274 No. 1 Jan. 1996 pg. 54-62

"Barriers are now being approached." "Industry is not likely to come to a screeching halt anytime soon." Authors forecast a continuation of present growth trends out to at least 2000. Good Bibliography for further study.

Hutc96 Key Words: SEMICONDUCTORS

Neural Computing (Special feature section of issue)

Computer (IEEE Computer Society) March 1996 (pg. 25 to 86)

Contents: Seven papers on Neural Networks — including a tutorial. Progress in Neural Computing keeps pace with advances in general purpose technology and is evolving along different lines. This can and does produce advances that may outpace the more conventional approaches. For this reason, neural computing progress should be tracked and "mined" for new concepts.

IEEE(a)96 Key Words: TECHNOLOGY

Trends in Nanotechnology: Waiting for Breakthroughs

Stix, Gary (editor) Scientific American April 1996 pg. 94-99

A summary of progress and status in a field with much promise but with many problems. It appears too early to forecast when this will impact the computer field but it clearly is some time off in the future.

Stix96 Key Words: TECHNOLOGY

Speed Gets a Whole New Meaning

From Business Week April 29, 1996 pg. 90

This is from the US Dept. of Energy's Accelerated Strategic Computing Initiative Supercomputer Program. The program announced is for a $ 46 million 1.8 teraflop system via a joint development between Intel and Sandia National Labs. The announced program is to be followed by a 10 teraflop and "even 100 and 500 teraflop systems". The DoE says their plan is to invest $ 940 million to pioneer new technologies. *This may be a real program, but the numbers are unreal measured by probable performance. FSP*

Busi96 Key Words: TERAFLOPS

The Next 10000_2 Years (i.e.: 1996 to 2012)

Lewis, Ted Computer (IEEE Computer Society) Part I: April 1996 (pg. 64- 70)
Part II: May 1996

Lewis predicts a breakdown of the current exponential growth of computer and semi-conductor performance — asserts that civilization is nearing the end of an "Age Change" which occurs only about every 200 years. Says software will have to change to adapt. Article is detailed in use of growth curves in forecasting – and the risk of extrapolating. A related column in the March '96 issue of IEEE Computer (pages 12 - 14) covers growth rates necessary "to succeed commercially" and specifically Software "The Dismal Science".

Lewi96 Key Words: FORECASTS

When Silicon Hits Its Limits, WHAT'S NEXT?

Thompson, Tom {tom_thompson@bix.com} BYTE April 1996 pg. 44-54

A survey article on three technologies: Holographic Storage; Molecular Memories: and, Quantum Processors. Summary: some promise, some problems. By the author's implying that these are technically possible but problematic, they fall in the time frame for after 2005.

Thom96 Key Words: TECHNOLOGY

Special report: Bioelectronic Vision

From IEEE SPECTRUM May 1996 Editor: Robert Braham pg. 21-69

Six articles: Toward an Artificial Eye; The Computational Eye; Neuromorphic Vision Chips; Ocular Implants for the Blind; Cortical Implants for the Blind; and Learning to see. Successful artificial seeing is forecast by the year 2015. Work in the past decade has moved away from the digital signal processing "television" approach toward simulation of the biological functions using electronic technology. As an example, an artificial retina silicon chip has been fabricated that apparently does a good job of simulation. Its performance is astounding: 10^{12} operations per second, a nine layer grid of 100 by 100 elements tied together with filters between each nearest neighbor, , on-chip memory, dynamic adaptation to lighting conditions, re-programming during operation in 1 millisecond, etc.

IEEE(b)96 Key Words: SOCIAL IMPACT; TECHNOLOGY

The Petaflops Systems Workshops *(Consisting of proceedings of two workshops as listed below)*

PetaFlops Architecture Workshop "PAWS '96" held April 22-25, 1996
PetaFlops System Software Summer Study "PetaSoft '96" held June 17- 21, 1996

Foster, Ian and Sterling, T. *No publication date* *Presumably by Cal. Tech* ~ 500 pg.

This excellent book provides probably the best summary/snapshot of the petaflops initiative as of mid-summer 1996. The material can serve as a source for additional references. Coverage includes; 1) the eight Point Design Studies — including PIM and Cache Only proposals; 2) Applications; 3) Architecture; 4) Software; and, 5) Findings and recommendations.

Fost96 Keywords: PETAFLOPS, SEMICONDUCTORS, TECHNOLOGY, TERAFLOPS

Next-Generation Compact Discs

Bell, A.E. Scientific American July 1996 pg. 42-46

The title says it clearly: this will be the storage means of choice for removable multimedia items. Total capacity growth in the next ten years or less is from the present limit of about 600 MB to 17 GB. Compatibility with present CDs will be retained. Much cheaper recording capability will emerge. Data rates will increase by about a factor of 10. The author does not see this as a replacement for magnetic on-line storage. *{Since it is much cheaper, I wonder. FSP}* *See below for blue laser which offers another factor of four possible improvement.*

Bell96 Key Words: SUPPORT SYSTEMS & COMPONENTS; FORECASTS

Blue-Laser CD Technology

Gunshor, R.L. and Nurmikko, A.V. Scientific American July 1996 pg. 48-51

Blue lasers are just now becoming feasible. The potential for a factor of four increased density is available and likely to develop. Improvements are needed in lower power, longer life, and manufacturing techniques before these will emerge from the laboratory.

Guns96 Key Words: SUPPORT SYSTEMS & COMPONENTS; SEMICONDUCTORS

Teraflops Computer Announcement (DoE)

Lawrence Livermore Nuclear Laboratory Via E-Mail 26 July 1996

LLNL awarded IBM a $93 million contract for “the world’s fastest supercomputer”, based upon RS6000s. The first of the staged deliveries is to be a 3 TeraFLOP system for demonstrations in Dec. ‘98. *(This is $300 per MegaFLOP and is thus right on the present cost forecast. FSP)* See also the reference Busi96, which is a related DoE procurement.

LLNL96 Keywords: TERAFLOPS

Redefining the Supercomputer

Taubes, Gary Science Vol. 273 20 Sept. 1996 pg. 1655 - 1657

Well prepared summary of the status, problems and the Governments activity aimed at computers 1,000 times, or more, as fast as the fastest today. In summary, the article emphasizes the magnitude of the task and that 50 to 70% of the effort will be in software — an effort no even started. Cites NASA Administrator as calling for petaflops speeds because “TeraFLOPS might be short-sighted”.

Taub96 Key Words: APPLICATIONS, PETAFLOPS, REQUIREMENTS

Frontiers ‘ 96 The Sixth Symposium on The Frontiers of Massively Parallel Computation

IEEE Computer Society Proceedings Oct. 27 - 31, 1996 372 pg.

A total of ten papers specifically on petaflops computing were included (Sessions 3B and 9B). [Approximately thirty papers covered the general subjects relating to massively parallel and high performance computing.] The ten papers in the petaflops area were reports on the progress of Federally sponsored “Point Design Studies”. These studies are aimed at up to 100 TeraFLOPS and specifically not at one PetaFLOPS or beyond. This represents some conservatism and possible doubts about this additional factor of ten step.

IEEE(c)96 Key Words: PetaFLOPS

1997

Solid State: Technology 1997 Analysis & Forecast

Geppert, L. (Ed.) IEEE Spectrum Jan. 1997 pg. 55-59

A good review article on the future and recent past of semiconductor computers and memory devices. This article contains some discussion of the physical, economic, and practical limits and possible new techniques to expand these.

Gepp97 Key Words: Semiconductors

Breaking the Teraflops Barrier

Clark, D. IEEE Computer Feb. 1997 pg. 12-14

Describes the DoE program (ASCI and ASAP) aimed at 100 teraflops for 2005. Quotes several experts that say we are "...now in the range of a few teraflops". Hints at some of the problems.

Clar97 Key Words: TERAFLOPS, PETAFLOPS

First Workshop on Hybrid Technology Multithreaded Architecture {HTMT} for Very High Performance Computing

Sterling, T. et al Conference Notes Cal. Tech. and JPL. (hosts) Feb. 24-25, 1997

Document consists of several hundred copies of visual aids used at the conference. As such, these provide a suitable summary of the conference coverage. For the expert in the field, these are nearly self-explanatory. However, they fail to provide the substantiation with hard data that may be needed to convince those not well grounded in the background. The technology proposed employs RSFS (Rapid-Single-Flux-Quantum} superconductors for memory elements. Another interesting presentation is on PIM {Processor In Memory}. At the very end of this document, in what may be a pre-publication version of a paper, is a very good summary of the technology and problems of developing a petaflops computer system; as well as a summary of the HTMT proposed system. This paper is separately referenced in the citation below.

Ster(a)97 Keywords: PETAFLOPS, SEMICONDUCTORS, TECHNOLOGY, TERAFLOPS

Steps to Petaflops Computing: A Hybrid Technology Multithreaded Architecture

Sterling, T. et al Included within: Conference Notes Cal. Tech. and JPL. (hosts) Feb. 24-25, 1997

A very good presentation of problems, potential, and the technology in moving from teraflops to petaflops systems. This paper says about 20 years to a petaflop using evolutionary and COTS development. The HTMT approach is presented as something possible within ten years. Also, this promises to solve the billions of watts problem posed by conventional components. Instead, it introduces whole orders of magnitude more latency. Bandwidth to memory requirements are extreme by today's standards.

Ster(b)97 Keywords: PETAFLOPS, SEMICONDUCTORS, TECHNOLOGY, TERAFLOPS

1997 Petaflops Algorithm Workshop (PAL '97)

Various Document is a compilation of visual aids used at conference held in Williamsburg, VA on April 13-18, 1997

This workshop was for PAL invitees. This publication is undated and does not constitute on open reference except to those who attended and received copies. The printing includes a number of papers as well as the visual aids. This source should be of interest to system and software personnel studying petaflops technology, problems, and software status. It is to be hoped this reference is replaced by a proceedings volume.

Vari97 Keywords: SOFTWARE

Computational Aerosciences Workshop 96

Feiereisen, W.J. NASA CD Conference Publication 20011 (Ames Research Center) May 1997 Conference held Aug. 13-15, 1996 pg. 324

The conference included about fifty four papers that present a snapshot of the state-of-the-art in the computational aerosciences — largely representing either specific applications or reports on system architecture and configurations. A paper on "The NAS Parallel Benchmarks 2.1 Results" showed test results of "out of the box" on four benchmarks for four popular (but largely superseded) processors. Efficiencies ranged from a high of about 30% down to about 10% for from 1 to 256 Processors. Extrapolating their results beyond the range tested seems to indicate that all of the processors converge at about 7% for 4K processors.

Feie97 Keywords: APPLICATIONS, TECHNOLOGY

When Caches Aren't Enough: Data Prefetching Techniques

Van der Wiell, S.P. and Lilja, D.J. COMPUTER (IEEE) July 1997 pg. 23-30

"...CPU performance has outpaced that of dynamic RAM, the primary component of main memory. This expanding gap ..." "It is not uncommon for scientific programs to spend more than half their runtimes stalled on memory requests." "Prefetching can nearly double the performance of some applications."

Vand97 Key Words: TECHNOLOGY, SEMICONDUCTORS

Chip-Scale Packaging

Thompson, Patrick Spectrum August 1997 pg. 36-43

Describes the status, pros, cons, etc. of chip-scale packaging of silicon chips. This relatively new technique can over improved performance for the same chip compared with the presently employed packaging. This technology is applicable in the computer area to memories at present.

Thom97 Keywords: MANUFACTURING

==

References below this line have been cited elsewhere and not yet studied because, although listed to be published later, the Langley Library was unable to obtain copies for this study (as of January 1997). They will remain below this line until available and reviewed.

Foresight: Proceedings of the Fourth Foresight Conference on Molecular Nanotechnology Nov. 9-11, 1995 Palo Alto, CA

Petaflops: Proceedings of the Petaflops Computing Workshop Aug. 14-23, 1995 Bodega Bay, CA
