SCCS-610
European High Performance Computing and Networking '94

Information Processing and Opportunities for HPCC Use in Industry

Geoffrey Fox and Kim Mills
Northeast Parallel Architectures Center
Syracuse University
111 College Place, Syracuse, NY 13244-41000
gcf@npac.syr.edu kim@npac.syr.edu

1 Introduction

The U.S. National HPCC effort, represented by a set of approximately fifty Grand Challenge problems, is largely focused on issues of scientific simulation [1]. In addition to advancing the mission of the federal agencies sponsoring this program, and the state of science in the field of physics, biology, and chemistry, and related fields such as atmospheric science, one can view the HPCC initiative as an effort to speed the process of technological innovation. Hardware systems are now relatively mature, and the HPCC program has successfully accelerated the development of new software technologies such as parallel compilers, scientific libraries, system software, and programming tools. HPCC has matured into a proven technology [2], but its success is mainly limited to applications in the research laboratory. Industry, and society more generally, have not yet adopted HPCC as a mainstream technology. Both technical and social obstacles limit industries' use of HPCC. The scientific community has focused on high performance levels, while industry places much more importance on reliability and ease of use. Perceived risk by industry executives of adopting "exotic" HPCC technologies, a general recommendation not to put "mission-critical" industry applications on HPCC systems, and pressure to maximize short term or end of the quarter earnings [3] are examples of the obstacles to HPCC application development in industry.

We use the term industry applications of HPCC to represent broad application of these technologies in diverse sectors such as manufacturing and health care, as well as education, consumer, and entertainment markets, since it is industry that is most likely to deliver technologies in these areas. Aside from issues such as performance vs. reliability, or ease of use, which are essential to the success of HPCC technologies in industry, there is a larger more pervasive issue holding back adoption of HPCC. We believe the most important opportunity for HPCC technologies lies in information processing applications. This conclusion is based on our interaction with industry over the past three years which we summarize in the form of a survey and describe in section 3 below. The importance of information processing applications in industry is becoming a central issue in many national HPCC meetings such as the industry panel discussions at the Pittsburgh HPCC Workshop [4], the first conference on Commercial Applications of Parallel Processing Systems [3], the National Information Infrastructure workshops at Supercomputing '93 [5], and the industry/HPCC consortia meetings [6].

For HPCC to deliver on its promise, it must be widely adopted and used in industry, and for this to happen, HPCC research and development must be more strongly integrated with the needs of industry. Over the past three years, the U.S. National HPCC effort has emphasized basic research or production of new knowledge, again largely in applications of scientific simulation. Policy experts point out that U.S. government stimulation of technical innovation has been most successful in space and defense, but is largely unsuccessful in civilian applications [7]. Research results alone are not useful to industry. Closely coordinated linkages are needed between a number of different groups which we believe must include: HPCC vendors; industry, government, and university research institutions; and potential HPCC technology consumers such as manufacturers, school districts, or software companies developing products for homes and small businesses. We do not elaborate on this process here but have designed InfoMall, our outreach program for delivering HPCC technologies to industry, in a set of steps to link researchers, vendors, small businesses, system integrators, and consumers [8,9].

This paper describes a classification of industry applications carried out over the past three years as part of our New York State funded industry outreach program [10]. We define four classes of applications based on types of information processing applications: information "production" or simulation; information analysis or data mining; information access and dissemination; and information integration or decision support.

We conclude that a core set of HPCC technologies —parallel architectures, software, and networks — combined with wide range of enabling technologies from partial differential equation solvers and mesh generators, to video compression and parallel databases, will form the basis of the National Information Infrastructure of the future. While scientific simulation will continue to be essential, information analysis applications will become the most important near term, and information integration applications will become the most important long term industry applications of HPCC technologies.

2 Trends in Information Technologies and Applications

Access to information, and the ability to organize and integrate information is playing an increasingly important role in education, business, world affairs, and society in general. This trend will result in applications such as global networks of 800 telephone numbers, instantaneous credit card verification systems, extremely large real-time databases, sophisticated scientific simulations in the classroom, decision support in the global economic war, early warning of political instabilities, and integration of Medicaid databases to streamline processes and uncover inefficiencies.

The success of the HPCC Grand Challenge program, which is focused almost entirely on scientific simulation, is built upon HPCC enabling technologies such as message passing libraries, high performance compilers, parallel operating systems, and numerical optimization software. Information processing applications will require development and integration of new HPCC enabling technologies such as parallel databases, high-speed networks, metacomputer integration software, multimedia support, and compression technology (Table 1).

Information based industries are rapidly expanding, growing two and a half times faster than the goods economy. Raw materials and labor now account for less than 25% of the cost of goods sold, while the cost of knowledge and information based services contributes to more than 50% of the costs of goods sold for *manufacturing* companies [11]. A competitive advantage will most certainly be captured by industries that can take advantage of HPCC technologies.

These trends in information, and observations about the role of information technologies in a range of applications prompted us to re-examine our initial industry survey of opportunities for HPCC in industry. This survey grouped industrial applications into 18 general areas, including for example computational fluid dynamics, electromagnetic simulation, particle transport problems, transaction processing, command and control and information processing, financial modeling, graphics, and flight simulation. For each application area, we defined example problems, and summarize the nature of parallelism inherent to the problem, and match this to appropriate machine architectures and software models [10]. Although a number of information processing applications were included, our initial survey strongly reflected the emphasis on simulation problems in the HPCC community, a trend which continues to the present day. In [10] we made the observation that information processing applications are the most important opportunity for HPCC application development in industry. We now define our industry survey of HPCC applications according to four types of information processing applications.

3 Classification of Industry Applications

We begin with a very broad definition of information, including both CNN Headline news, and insights gotten from new computational physics models. We define four classes of industry applications of HPCC based on information. These categories, listed in Table 2 HPCC Applications in Information Processing, include: information "production" or simulation; information analysis or "data mining"; information access and dissemination, and information integration. For each class, we outline example applications, and characterize the parallelism present in the problem and the natural hardware/software architectures for supporting the problem.

3.1 Information Production

Information "production" or simulation is currently the main focus of the national HPCC community. Computational fluid dynamics, structural analysis, and electromagnetic simulation are examples of simulation problems with relatively large federal research programs funded by U.S. agencies such as the National Aeronautics and Space Administration, and the Department of Energy. These applications have advanced the development of sophisticated parallel software technologies such as finite element, and finite difference partial differential equation solvers, and turbulence and mesh generators. A difficulty with this application, in terms of HPCC development in industry, is that many industries that use simulation technologies such as the aerospace industry, are in a period of contraction rather than growth. This makes transfer of new technology to industry extremely difficult. Simulation applications in industries experiencing growth, and therefore more amenable to funding new technological approaches, include mortgage pricing and stock option pricing in the financial industry, and systems simulations in defense (SIMNET), education (SIMCITY), and multimedia/virtual reality in the entertainment industry.

3.2 Information Analysis

Information analysis or "data mining" concerns the extraction of information from typically very large databases. A well established example is the petroleum industry, where oil extraction is based on analysis of terabytes of seismic data. Newer examples include extracting patterns of fraud or inefficiencies in health care, securities, and credit card transaction records, or extracting customer preferences from purchase data. For these information analysis applications, the owners of the database often consider the data itself as their real resource, and are concerned with ways to maximize access to the database, structuring information in ways that create meaning, and developing strategies for fast response to unique customer requests or new products [12]. Market segmentation is an information analysis application aimed at tailoring traditional mass marketing approaches down to the level of the individual, producing "mass customization" for products ranging from banking services, news media, and retail consumer products.

3.3 Information Access and Dissemination

Information access and dissemination covers a wide range of applications for which we define four important categories-- database information, video, imagery, and simulation on demand. These categories of information type form the basis of our InfoVision program (section 4.1 below).

Information access and dissemination includes text on demand or document retrieval from educational clearinghouses by school teachers, and accessing the world's news media and government reports by political analysts to predict political instabilities. Three dimensional geographic information systems linked with multimedia map databases are a form of simulation on demand, and could be used to deliver educational, tourism, or city planning information to students, the public, or government officials. Video on demand will soon become a large, new commercial market. Information access and dissemination applications are enabled by the National Information Infrastructure, which in the near future will link compute servers, and distributed databases over high-speed network links to the home and office settop box (soon to become a powerful PC). We believe that information access and dissemination will become the most important near term industry application of HPCC. Industry shake-ups, and billion dollar scale investments by the telecommunication and software industry in information sources, initially focused on capturing entertainment markets such as movie libraries and video games, serve to illustrate this opportunity.

3.4 Information Integration

Information integration combines information production, analysis, and access in a system of systems. One might view information integration as the dual-use equivalent of command, control, communications, and intelligence for business, manufacturing, and the consumer. By analogy to the aviation term "flying by wire," where the pilot is not flying a plane but an informational representative of the plane, the business leader or home shopper might "manage by wire" [14] or use decision support technologies to manage, optimize, and make tradeoffs between informational representatives of complex manufacturing processes or a global consumer market. Information integration represents the largest long term market for HPCC technologies.

4 Information Processing Applications: InfoVision, MADIC, and Electrical Power Networks

We expect information processing applications to become the important opportunity for HPCC in industry. In this section, we describe three application projects that NPAC is currently developing. InfoVision is a set of projects concerning information on demand technologies and applications. MADIC is a multidisciplinary analysis and design industry consortium for developing pre-competitive HPCC software. Transient stability analysis in power networks is an NPAC project to optimize the distribution of power in an electrical transmission system.

4.1 InfoVision

HPCC will play a pervasive role in the information-rich world of the future. HPCC hardware will implement both the networks and the enormous information repositories that will function as servers on the National Information Infrastructure (NII). We term this NII server role as InfoVision, for Information, Video, Imagery and Simulation On demand. Alternately, we can think of InfoVision as high performance multimedia, scalable HPCC algorithms and software used to manage the network packets, and implement new applications running on InfoVision servers.

Dedicated gigabit links between supercomputers will be only one use of the NII. In addition, the NII will support many millions of "personalized" ten megabit/sec streams supplying information on demand. InfoVision servers for video on demand will be classical parallel computers---whether clusters of workstations or integrated MPP's. We will have distributed clients with intelligent set top boxes and gigaflop video games in every home, school and business. But we will need production InfoVision servers in large data centers. HPCC techniques will be needed for the necessary hierarchical data storage, browsing, knowledge agents, text and video image content identification, and compression algorithms. The HPCC community needs to examine message passing standards and see that they not only span all homogeneous parallel machines but also the InfoVision metacomputer (100 million clients demanding real-time response from some 10,000 InfoVision superservers).

InfoVision is an InfoMall project which involves state-of-the-art demonstrations on NYNET, a regional ATM-based gigabit network in New York State, of distributed interactive intelligent information systems. InfoVision is a set of dual-use distributed multimedia projects which prototype near term civilian and military applications of HPCC. We believe that the demonstrations are rich enough to explore most of the essential issues for HPDC information systems---now called National Challenges---which will be implemented on the NII in the near future. These demonstrations need various key enabling technologies, and InfoVision supports the InfoTech activities needed to develop, test and evaluate these technologies. These technologies include multimedia database, transport protocols, network interoperability, architecture and management, collaboration, compression, security and privacy, graphics rendering, video browsing, image content and full text search techniques, distributed heterogeneous computing environments, user interfaces, and virtual reality.

Several InfoMall partners are needed to form the virtual organizations to both implement the demonstrations with component technologies and system integration as well as to explore the InfoTech activities needed by and opened up by the demonstrations. The partners could include Abrams Gentile Entertainment Inc., Columbia University, Cornell University, Digital Equipment Corporation, Eastman Kodak, IBM, Maspar, NYNEX, Oracle, the U.S. Air Force Rome Laboratory, Ultra Inc., and Syracuse University. At Syracuse University, we have involved several organizations including NPAC, Electrical and Computer Engineering (ECE), Computer Applications and Software Engineering (CASE) Center, Computer and Information Science (CIS), the School of Information Studies (IST), the School of Education, and the Maxwell School of Citizenship and Public Affairs.

As part of our InfoVision program, we are currently developing a focused set of application demonstrations which include simulation (using weather simulations and 3D terrain modeling of New York State), text (using full text databases from the Maxwell School and the AskERIC educational resource in IST), imagery (using Kodak's GIODE PhotoCD technology), and video (using real-time educational news clips from CNN as one possible initial example database) on demand. We are using these demonstrations to develop the necessary component technologies, perform the overall system integration, then test and evaluate these information on demand technologies [13].

4.2 MADIC

The Multidisciplinary Analysis and Design Industrial Consortium (MADIC) project addresses applications of high-performance computing and communications technologies to the analysis and design of large-scale, complex products, such as aircraft and automobiles. Complex systems typically require consideration of multiple engineering and scientific disciplines, including fluid dynamics, structural dynamics, heat transfer, combustion, and manufacturing consideration. HPCC technologies are expected to provide a new generation of design tools that make use of large-scale, high-resolution simulations executing simultaneously in a distributed, heterogeneous computing environment. Outreach to large industrial firms was initiated in December 1991. Interested firms organized MADIC in March 1992.

MADIC members collaborate on projects of mutual interest, with teams of members formed to address specific projects. Current projects include evaluation of NASA software for use by industry. Current members of MADIC include: Boeing, Ford Motors, General Electric Aircraft Engines, General Electric Corporate Research & Development, General Motors, Grumman, Lockheed, McDonnell-Douglas, Northrup, Rockwell, United Technologies Research Center, Vought, Northeast Parallel Architecture Center at Syracuse University, and Center for Research in Parallel Computation (Rice University).

Current industrial practice typically requires 12-18 months to develop new designs for large-scale, complex products, such as aircraft and automobiles. Global competitive pressures demand reductions in time-to-market as well as improved design quality for dual-use (commercial and military) applications. This project is intended to reduce significantly the time required to develop new designs of large-scale systems, while simultaneously improving the quality and reducing the cost of new products.

MADIC provides a mechanism for aerospace and automobile manufacturers to address current challenges in the analysis and design of complex products. MADIC operations are coordinated through the Center for Research in Parallel Computation (CRPC), which has its headquarters at Rice University and also includes the Northeast Parallel Architectures Center (NPAC) at Syracuse University. MADIC firms pay an annual membership fee, which provides primary funding for an Executive Director and support staff. A representative of each member firm sits on a Board of Directors, which provides oversight and coordination. MADIC members form teams to address specific projects, consistent with the objectives of the consortium. Meetings and workshops are arranged at least quarterly. MADIC projects are intended to develop, validate and deploy an integrated analysis and design environment operating on a heterogeneous network that includes workstations, vector supercomputers and massively-parallel processors. A proposal for the entire United States Multidisciplinary Analysis and Design Environments (USMADE) submitted to the ARPA TRP in July 1993 is pending at this writing. USMADE includes modules for geometry definition, grid generation, physical simulation (fluid dynamics, structural dynamics, electromagnetics), process simulation (manufacturing, maintenance), optimization and visualization. Work on selected aspects relevant to the USMADE project is currently underway with funding from other sources. NASA is funding a project to evaluate physical simulation software for use in industrial applications. In this project, industry teams are applying selected NASA codes to problems of current interest, including wings, wing-body interactions, nozzle flows, internal flows (turbines and compressors) and flexible wings. As part of this project, NPAC researchers are evaluating the suitability of the NASA codes to be implemented in High-Performance Fortran. CRPC researchers at Rice University are evaluating the use of the automatic differentiation system ADIFOR.

4.3 Transient Stability Analysis for Power Systems

One major NPAC project involves the utility industry where we are working with the Niagara Mohawk Power Corporation on their electrical transmission system. One can simulate this in real time using modern parallel machines with around 64 nodes. We have generated a prototype which Niagara Mohawk is now evaluating. Future steps involve the industry consortium EPRI, other utilities and the companies that currently build the sequential power grid software now used by electrical utilities. This application is relatively easy to parallelize as it involves forming a matrix (with embarrassingly parallel matrix element generation) followed by a sparse matrix solve. The latter uses a direct method on a roughly 10000 X 10000 matrix. Utilities must balance the cost of producing power against the cost of failures. Thus an accurate simulation of transition stability analysis -- the response of circuit to a disturbance -- allow better cost--

security tradeoffs and hence cheaper power. Power is very expensive and even modest savings from real-time simulation would warrant purchase of HPCC supercomputers by the utilities [15].

5 Summary and Conclusions

The National Information Infrastructure of the future will be built upon core HPCC technologies such as parallel architectures, software, and high-speed networks. In addition to the key enabling technologies associated with high-performance scientific computing, such as image processing algorithms and linear algebra libraries, information processing applications will require development and integration of new, enabling HPCC technologies such as collaboration services, and multimedia databases. These new technologies will find application in information industries such as home entertainment and agile manufacturing.

While simulation will remain an important HPCC application, the real opportunity for developing significant industrial applications of HPCC technologies lies in information processing. Information industries such as home shopping, electronic banking, distance education, and televirtuality are growing at a much faster rate than the product industries such as aerospace, or other manufactured goods. The current focus on scientific simulation in the national HPCC community must change in response to this trend.

Based on our survey of industrial applications of HPCC compiled over the past three years, we described four classes of information processing—production, analysis, access and distribution, and integration. We concluded that information analysis (e.g., data mining credit card holders' purchase patterns) is the most important near term opportunity, and information integration (e.g., decision support in the global economic war) is the most important long term opportunity for HPCC applications in industry.

Table 1.1 HPCC Industrial Applications: Simulation or Information Production

Item	Application Area and Examples	Problem Comments	Problem Architecture	Machine and Software
1	Computational Fluid Dynamics	PDE, FEM Turbulence Mesh Generation	Loosely Synchronous Adaptive Mesh is Asynchronous	SIMD, MIMD for irregular, adaptive HPF(+) Unclear for adaptive irregular mesh
2	Structural Dynamics	PDE, FEM Dominated by Vendor Codes (e.g. NASTRAN)	Loosely Synchronous mesh as in 1	MIMD as Complex geometry HPF(+)
3	•Electromagnetic Simulation •Antenna Design •Stealth Vehicles •Noise in high frequency circuits •Mobile Phones	PDE Moment method (matrix inversion) dominates ______________ Later FEM, FD?	Synchronous ______________ Loosely Synchronous	SIMD HPF ______________ SIMD, MIMD, HPF(+)
4	<u>Scheduling</u> •Manufacturing •Transportation (Dairy delivery to military deployment) •University Classes •Airline Scheduling of crew, planes in static or dynamic (midwest snowstorm) cases	Expert systems and/or Neural Networks, Simulated Annealing Linear Programming (hard sparse matrix)	Asynchronous Synchronous Loosely Synchronous	MIMD (unclear speedup) Asyncsoft SIMD HPF MIMD
5	Environmental Modeling - Earth/Ocean/ Atmospheric Simulation	PDE, FD, FEM Sensitivity to data	Loosely Synchronous	SIMD, MIMD for irregular, adaptive HPF(+) Unclear for adaptive irregular mesh

PDE	Partial Differential Equation
FEM	Finite Element Method
FD	Finite Difference
ED	Event Driven Simulation
TS	Time Stepped Simulation
VR	Virtual Reality
HPF	High Performance Fortran [16]
HPF+	Natural Extensions of HPF [16]
MPF	Fortran plus message passing for loosely synchronous software
Asyncsoft	Parallel Software System for (particular) class of asynchronous problems
CFD	Computational Fluid Dynamics

Note on Language: HPF, MPF are illustrative for Fortran: one can use parallel C, C++ or any similar extensions of data parallel or message passing languages

Table 1.1 HPCC Industrial Applications: Simulation or Information Production (continued)

Item	Application Area and Examples	Problem Comments	Problem Architecture	Machine and Software
6	Environmental Modeling - Complex systems e.g. lead concentration in blood	Empirical models Monte Carlo and Histograms	Embarrassingly Parallel plus global reductions	Some SIMD MIMD more natural
7	Basic Chemistry •Chemical Potentials •Elemental Reaction Dynamics	Calculate Matrix elements Matrix eigenvalue Multiplication, inversion	Embarrassingly Parallel Synchronous	MIMD (maybe SIMD) HPF
8	Molecular Dynamics	Particle Dynamics with irregular cutoff forces Fast Multipole methods Mix of PDE and Particles in PIC	Loosely Synchronous	HPF(+) or MPF for fast multipole
9	Economic Modeling •Real Time Optimization •Mortgage backed securities •Option Pricing	Individual (Monte Carlo) Full simulations of portfolios	Synchronous, Embarrassingly parallel Metaproblems	SIMD, HPF MIMD, SIMD Integration software
10	Network Simulations	Sparse matrices; Zero structure defined by network connectivity	Loosely Synchronous	MIMD MPF

PDE	Partial Differential Equation
FEM	Finite Element Method
FD	Finite Difference
ED	Event Driven Simulation
TS	Time Stepped Simulation
VR	Virtual Reality
HPF	High Performance Fortran [16]
HPF+	Natural Extensions of HPF [17]
MPF	Fortran plus message passing for loosely synchronous software
Asyncsoft	Parallel Software System for (particular) class of asynchronous problems
CFD	Computational Fluid Dynamics

Note on Language: HPF, MPF are illustrative for Fortran: one can use parallel C, C++ or any similar extensions of data parallel or message passing languages

Table 1.1 HPCC Industrial Applications: Simulation or Information Production (continued)

Item	Application Area and Examples	Problem Comments	Problem Architecture	Machine and Software
11	Particle Transport Problems	Monte Carlo methods as in neutron transport for explosion Simulations	Embarrassingly Parallel (Asynchronous)	MIMD HPF
12	Graphics (rendering) Hollywood Virtual Reality	Several operational Parallel Ray Tracers Distributed model hard	Embarrassingly Parallel (Asynchronous)	MIMD HPF Asyncsoft for distributed database
13	Integrated Complex Systems Simulations •Defense (SIMNET, Flight Simulators) •Education (SIMCITY) •Multimedia/VR in Entertainment •Multiuser virtual worlds •Chemical & Nuclear Plants	Event driven (ED) and Time Stepped (TS) Simulations. Virtual Reality Interfaces. Database backends. Interactive	Metaproblem Fully asynchronous if ED Loosely Synchronous for TS	Timewarp or other ED Asyncsoft HPF(+) Integration Software Database

PDE	Partial Differential Equation
FEM	Finite Element Method
FD	Finite Difference
ED	Event Driven Simulation
TS	Time Stepped Simulation
VR	Virtual Reality
HPF	High Performance Fortran [16]
HPF+	Natural Extensions of HPF [17]
MPF	Fortran plus message passing for loosely synchronous software
Asyncsoft	Parallel Software System for (particular) class of asynchronous problems
CFD	Computational Fluid Dynamics

Note on Language: HPF, MPF are illustrative for Fortran: one can use parallel C, C++ or any similar extensions of data parallel or message passing languages

Table 1.2 HPCC Industrial Applications: Information Analysis - "Data Mining"

Item	Application Area and Examples	Problem Comments	Problem Architecture	Machine and Software
14	Seismic and Environmental data analysis	No oil in NY State. Parallel Computer already important	Embarrassingly parallel as in many "event" analysis problems (high energy physics, astronomy, etc.)	SIMD, maybe MIMD needed HPF
15	•Image Processing •Medical Instruments •EOS (Mission to Planet Earth) •Defense Surveillance •Computer Vision	Commercial Applications of Defense Technology. Component of many Information Integration Applications e.g. Computer Vision in Robotics	Metaproblem Synchronous (low level) Loosely Synchronous (medium level) Asynchronous (expert system)	Metacomputer SIMD (low level) MIMD (medium/high level) HPF(+) Software Integration Asyncsoft Database
16	•Health Fraud Inefficiency •Securities Fraud •Credit Card Fraud	Linkage Analysis of database records for correlations	Synchronous if records "identical" otherwise Loosely Synchronous	SIMD MIMD
17	Market Segmentation	Sort and Classify records to determine customer preference by region (city --> house)	Loosely Synchronous	Aspects could be SIMD MIMD better for sorting?

PDE	Partial Differential Equation
FEM	Finite Element Method
FD	Finite Difference
ED	Event Driven Simulation
TS	Time Stepped Simulation
VR	Virtual Reality
HPF	High Performance Fortran [16]
HPF+	Natural Extensions of HPF [17]
MPF	Fortran plus message passing for loosely synchronous software
Asyncsoft	Parallel Software System for (particular) class of asynchronous problems
CFD	Computational Fluid Dynamics

Note on Language: HPF, MPF are illustrative for Fortran: one can use parallel C, C++ or any similar extensions of data parallel or message passing languages

Table 1.3 HPCC Industrial Applications: Information Access and Dissemination: InfoVision - Information, Video, Imagery and Simulation on Demand

Item	Application Area and Examples	Problem Comments	Problem Architecture	Machine and Software
18	Transaction Processing •ATM (automatic teller machine)	Database-most transactions short. As add "value" this becomes Information Integration	Embarrassingly Parallel	MIMD Database
19	Collaboratory •Telemedicine	Research Center or doctor(s) - patient interaction without regard to physical location	Asynchronous	High Speed Network
20	Text on Demand •Digital (existing) libraries •ERIC Education database, •United Nations - Worldwide newspapers	Multimedia database Full text search	Embarrassingly Parallel	MIMD Database
21	Video on Demand •Movies, News (CNN Newsource & Newsroom), •Current cable, •United Nations - Policy Support	Multimedia Database Interactive VCR, Video Browsing, Link of video and text database	Embarrassingly Parallel	MIMD Database Compression (SIMD) Video Editing Software
22	Imagery on Demand Kodak GIODE •"clip art" on demand •Medical images •Satelite images	Multimedia database Image Understanding for Content searching and (terrain) medical feature identification	Metaproblem Embarrassingly Parallel plus Loosely Synchronous Image Understanding	MIMD but much SIMD image analysis
23	Simulation on Demand •Education, Tourism, City planning, •Defense mission planning	Multimedia map database Generalized flight simulator Geographical Information System	Synchronous terrain rendering with Asynchronous Hypermedia	SIMD terrain engine (parallel rendering) MIMD database

PDE Partial Differential Equation
FEM Finite Element Method
FD Finite Difference
ED Event Driven Simulation
TS Time Stepped Simulation
VR Virtual Reality
HPF High Performance Fortran [16]
HPF+ Natural Extensions of HPF [17]
MPF Fortran plus message passing for loosely synchronous software
Asyncsoft Parallel Software System for (particular) class of asynchronous problems
CFD Computational Fluid Dynamics

Note on Language: HPF, MPF are illustrative for Fortran: one can use parallel C, C++ or any similar extensions of data parallel or message passing languages

Table 1.4 HPCC Industrial Applications: Information Integration

	These involve combinations of Information Production, Analysis and Access thus need software and machine architecture issues given for these "subproblems" • Systems of Systems
24	Command and Control • Battle Management, Command, Control, Communication, Intelligence and Surveillance (BMC^3IS) • Military Decision Support
25	SIMNET - Military Simulation with computers and people in the loop
26	Business decision support
27	Political decision support United Nationals uses video and multilingual newspaper (Maxwell School at Syracuse University)
28	Robotics
29	Electronic banking
30	Electronic shopping
31	Agile Manufacturing - Multidisciplinary Design - Concurrent Engineering •MADIC Industrial Consortium

6 References

[1] Grand Challenges 1993: High Performance Computing and Communications, The FY 1993 US Research and Development Program, Walter Massey, Chair. A Report by the Committee on Physical, Mathematical, and Engineering Sciences, Federal Coordinating Council for Science, Engineering, and Technology Committee, Office of Science and Technology Policy, 1993. (68 p.)

[2] G. Fox., P. Messina, and R. Williams. Parallel Computing Works!. Morgan Kaufmann, San Mateo, CA. 1994.

[3] Commercial Applications in Parallel Processing (CAPPS). Microelectronics and Computer Technology Corporation. Austin, Texas. October 19-20, 1993.

[4] Workshop on Grand Challenge Applications and Software Technology. May 4-6, 1993. Industry Panel Workshop. Hyatt Regency Hotel, Pittsburgh, PA.

[5] Supercomputing '93. National Information Infrastructure Testbed Panel. Portland, Oregon. November 15-19, 1993. IEEE Computer Society Press, Los Alamitos, CA.

[6] RCI North American Annual Member Executive Conferences on Technology Transfer. RCI, Ltd. Worldwide Headquarters, 1301 East 79th Street, Suite 200, Minneapolis, MN 55425.

[7] T. Pinelli, R. Barclay, A. Bishop, and J. Kennedy. Information Technology and Aerospace Knowledge Diffusion. electronic networking: research, applications and policy. 2:2 31-50. 1992.

[8] G. Fox, E. Bozucz, D. Jones, K. Mills, and M. Podgorny. InfoMall Program Plan: A Scalable Organization for the Development of HPCC Software and Systems. SCCS--531. (Available from NPAC, Syracuse University, Syracuse NY 13210-4100). 1993.

[9] K. Mills and G. Fox. InfoMall: An Innovative Strategy for High-Performance Computing and Communications Applications Development. to be published in electronic networking: research, applications and policy. 1994.

[10] G. C. Fox. Parallel Computing in Industry: An Initial Survey. in proceedings Fifth Australian Supercomputing Conference. World Congress Center. Melbourne, Australia. SCCS--302, Center for Research on Parallel Computation-TR92219. (Available from NPAC, Syracuse University, Syracuse NY 13210-4100). 1992.

[11] N. Johnson. The Nature of Information. in The Knowledge Economy. The Nature of Information in the 21st Century. Institute for Information Studies. 1993-1994.

[12] B. Ives and M. McKeown. The Promise of a New World Information Order. in The Knowledge Economy. The Nature of Information in the 21st Century. Institute for Information Studies. 1993-1994.

[13] G. Fox, S. Hariri, R. Chen, K. Mills, M. Podgorny. InfoVision--Information, Video, Imagery, and Simulation On Demand. SCCS--575. (Available from NPAC, Syracuse University, Syracuse NY 13210-4100). 1993.

[14] S. Haeckel and R. Nolan. The Role of Technology in an Information Age. in The Knowledge Economy. The Nature of Information in the 21st Century. Institute for Information Studies. 1993-1994.

[15] D. Koester, S. Ranka, and G. Fox. Power Systems Transient Stability - A Grand Computing Challenge. SCCS--549. (Available from NPAC, Syracuse University, Syracuse NY 13210-4100). 1992.

[16] High Performance Fortran Forum, (1993). High performance Fortran language specification, High Performance Fortran Forum, May 3, 1993, Version 1.0. 184 pp. HPFF92a. Copyright Rice University, Houston, Texas. (Available by anonymous ftp at minerva/npac.syr.edu HPFF/hpf-v10-final.ps)

[17] Alok Choudhary, Geoffrey Fox, Seema Hiranandani, Ken Kennedy, Chuck Koelbel, Sanjay Ranka and Joel Saltz, A Classification of Irregular Loosely Synchronous Problems and Their Support in Scalable Parallel Software Systems, in Proceedings from Darpa Software Technology Conference, 1992, pp. 138--149. Syracuse University Technical Report SCCS-255, June, 1991.