The Use of the National Information Infrastructure and High Performance Computers in Industry

Geoffrey C. Fox, and Wojtek Furmanski
Syracuse University
Northeast Parallel Architectures Center
111 College Place
Syracuse, New York 13244-4100
gcf@npac.syr.edu, furm@npac.syr.edu
http://www.npac.syr.edu
http://www.infomall.org

A presentation related to this paper will be found in [1]

Abstract

We divide potential NII (National Information Infrastructure) services into five broad areas: Collaboration and televirtuality; InfoVISiON (Information, Video, Imagery, and Simulation on Demand), and digital libraries; commerce; metacomputing; WebTop productivity services. The latter denotes the broad suite of tools we expect to be offered on the Web in a general environment we term WebWindows. We review current and future World Wide Web technologies, which could underlie these services. In particular, we suggest an integration framework WebWork for High Performance (parallel and distributed) computing and the NII. We point out that pervasive WebWork and WebWindows technologies will enable, facilitate and substantially accelerate such complex software processes on the NII.

We briefly analyze seven broad application areas: society; business enterprises; health care; defense command and control, and crisis management; education; collaboratory; manufacturing. We contrast their use of NII services with a more detailed examination of the manufacture of complex systems, such as aircraft and automobiles. This application will stress the NII, but there is a remarkable opportunity to develop new manufacturing practices

that offer cost savings and reduced time to market.

1 Introduction

We are at a remarkable time in high-performance computing and communications (HPCC) technology. Research and development over the last decade or so is maturing rapidly, and can be applied in several important ways to real world problems. Further, most progress is perhaps from the confluence of technologies and indeed academic disciplines, including parallel and distributing computing, networking, and information science. In Section 2, we review the current situation with current World Wide Web (WWW), National Information Infrastructure (NII), and HPCC. We set the standard multi-use NII vision with infrastructure driven initially by the entertainment and information industries, but applied broadly over a range of applications from education to manufacturing. In the speculative Section 3, we suggest an emerging World Wide Metasystem, where the user interacts with WebWindows—the operating system of the WWW. We describe how a network of combined compute and Web servers can be harnessed in WebWork as a distributed or parallel compute engine. We believe that Java, from Sun Microsystems, is (an example of) critical technology, and we illustrate its use in some examples, including dataflow computing, software engineering, and education. We discuss other critical emerging Web technologies, including the integration of relational and distributed database technology, and the unifying middleware, WebScript.

In the final section, Section 4, we tie the technologies to a set of base multi-use NII services whose use is illustrated by seven application areas, including manufacturing, which is discussed in a little more detail.

2 The WWW, NII, and HPCC today

2.1 The Current World Wide Web Technologies

Table 1 and Figure 1 summarize some important aspects of today's WWW [2], [3], and [4]. Clients have sophisticated display capabilities, and request and receive hyperlinked material from servers. Documents are typically written in HTML, which is a simple subset of SGML, which supports hyperlinked information via URLs [5]—universal resource locators—specifying location of related material. HTML is a dynamic standard and extensive

Table 1: Some early WorldWide Web (WWW) technologies and services

- **Clients** (such as Mosaic and Netscape) support browsing of hyperlinked documents, but have no internal interactive/compute capability [7]
- **Servers** read HTTP and deliver requested service to client
- **HTML**—a document format supporting hyperlinks
- **HTTP**—a Transport Protocol defining Interaction between Web servers and Clients
- **MIME**—a data format allowing agent-like (extended email) communication
- **CGI**—a standard interface allowing sophisticated server extensions
- **PERL**—a rapid prototyping language (script) aimed at text and file manipulation
- **Web Search engines**, such as YAHOO, HARVEST, WAIS, LYCOS—early distributed database access

and important new functionality is being added [6]. Servers and client use MIME as the data format, and this can be thought of as an extended email syntax with *headers* specifying format or method (program to be executed by remote procedure call, RPC) followed by the *body* of message, which is typically desired data. The Common Gateway Interface (CGI) is a standard way to add functionality to WebServers with software that typically uses PERL as a convenient rapid prototyping environment offering good access to system resources, files, and document manipulation.

2.2 Further Critical Technology Building Blocks for the NII

We envisage an architecture for the NII, shown in Figure 2 that builds applications on top of general (multi-use) services, which themselves ride on a collection of technologies, and middleware, which we term generically WebWindows, and discuss in Section 3.

In Table 2, we list some important pervasive technologies that will be

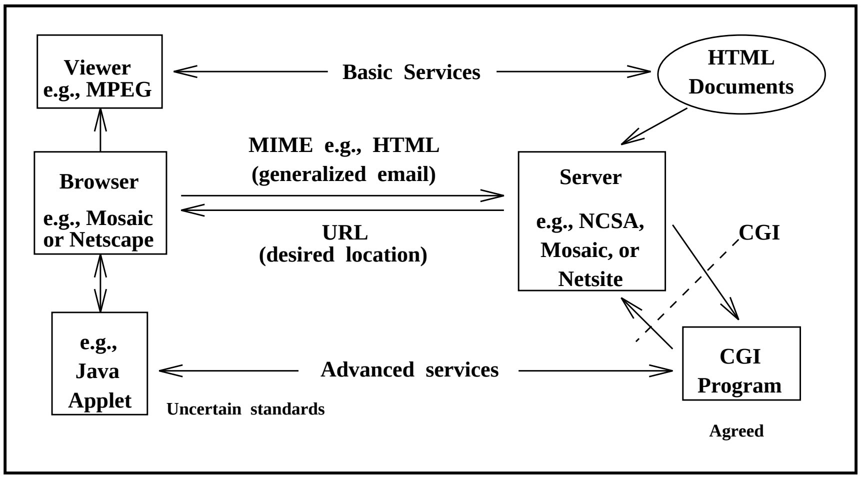

Figure 1: A simple representation of the software architecture of the WWW showing emerging advanced services, as well as the current technology

critical building blocks of the future NII. This includes the basic hardware architecture where there is general agreement on the importance of ATM. Currently, ISDN is being deployed quite widely, and it will be important to see if this performance (128 Kbits/second) is sufficient to enable many NII applications. One to two megabits/second performance is needed for NTSC quality digital video with MPEG compression. However, either better compression (such as Wavelet) or quality reduction (smaller images) is sufficient for many applications, such as education (where small clips—not two hour movies are needed) and collaboration. The latter illustrates a typical NII/Web issue. There are several excellent video conferencing systems, such as Intel's Proshare for the PC. However, a good collaboration environment requires not just video exchange of the participants, but also common whiteboards and other text and graphics information exchange. The latter are clearly best implemented with Web technology, and I see that the days of stand alone collaboration systems are numbered—video conferencing will be integrated with the Web and then will evolve to a full televirtual environment embodying such technologies as VRML and MOO's—the most elegant multi-user virtual worlds. Other integration challenges are seen in the message passing arena where we should reconcile parallel processing technologies, such as MPI and PVM with MIME and HTTP of the WWW.

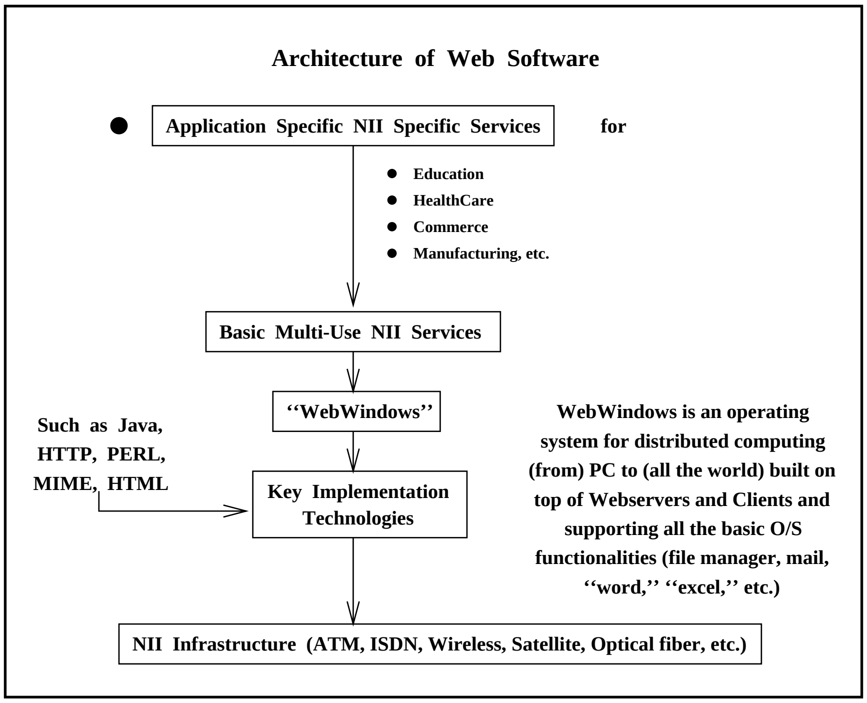

Figure 2: A layered view of Web (NII) software building applications on top of generic services that are in turn built on pervasive technologies

Again, current relational database systems and the powerful but less structured WWW search engines have complementary strengths, which need to be integrated [8].

2.3 An Entertainment and Telecommunication View of the Early NII

The NII will on one hand build on the Internet and the WWW. On the other hand, much of the infrastructure will be motivated by some early business opportunities, which roughly can be thought as digital versions of today's cable TV. As shown in Figure 3, several industry sectors are teaming and competing for this market, which also includes interactive shopping and the digital equivalent of the current video rental business.

In Figure 4, we try to quantify this vision assuming that some 100 million clients would be connected to the digital media information system. The year in which this becomes reality depends a lot on both technology and government—some regulations could stifle competition and slow digital deployment; others would speed the process through encouragement of provision of digital services to remote areas that might be uneconomic in a purely market driven process. The NII will link hybrid delivery technologies (including ISDN, cable, ADSL, wireless, mobile systems) for the off-ramps to ATM and satellite, which can be used for both off-ramps and the long distance trunks. Further, we expect large-scale Massively Parallel Machines (MPPs) to be part of a network of video servers, and more generally, WebServers that supply information to these clients. We term this scenario *InfoVISiON* for *Information, Video, Imagery*, and *Simulation on Demand*. This includes the storage, query, and dissemination of this wide range of multimedia data. There is surely at least 100,000 hours of interesting video material in the archives of Hollywood studios, CNN, Routers, and network TV. If compressed in MPEG format, this corresponds to some 100 Terabytes of needed storage capacity. MPEG2 or other format, such as motion JPEG, would require much more storage and are probably necessary to support editing and other video production applications. Note that the current WWW has only a few percent of its storage devoted to video—the future NII will be dominated by video data. Each NII offramp will, as shown in Figure 5, connect homes (offices, school desks) at the rate of 1–20 Megabits/sec to the set of NII InfoVISiON servers. This rate covers the range from compressed VHS to HDTV picture quality. Note that this performance is 100–1,000 times greater than today's conventional 14.4 Kbaud

Table 2: Some important real world technologies that are (to be) incorporated into the WWW.

- **Relational databases**—Oracle, DB2 have Web Interfaces
- **Collaboration** from Console Units (PIctureTel, CLI), Desktop (SGI Inperson) to MOOs
- **Compression** from MPEG with Wavelet to host of proprietary solutions—a faction of 20 to 200 saving in space and bandwidth
- **Geographical Information Systems**
- **Security** will enable commerce on the Internet—essential for Defence as well
- ATM, ISDN, Wireless, Satellite will be hybrid **physical implementation** of NII
- CORBA, Opendoc, OLE, SGML, Hytime are critical file and document **standards** [9]
- **High Performance Multimedia servers** to enable digital information delivery on demand
- **Data transport** from MPI/MSGWAY/PVM to AAL to CBR/VBR
- **Windows95/NT**—the last of the non-social (non-Web) operating systems—will follow dinosaurs (IBM mainframes) into extinction
- **Personal Digital Assistants**—WebNewtons done right—Telescript (agent based communication) and Magic Cap operating system

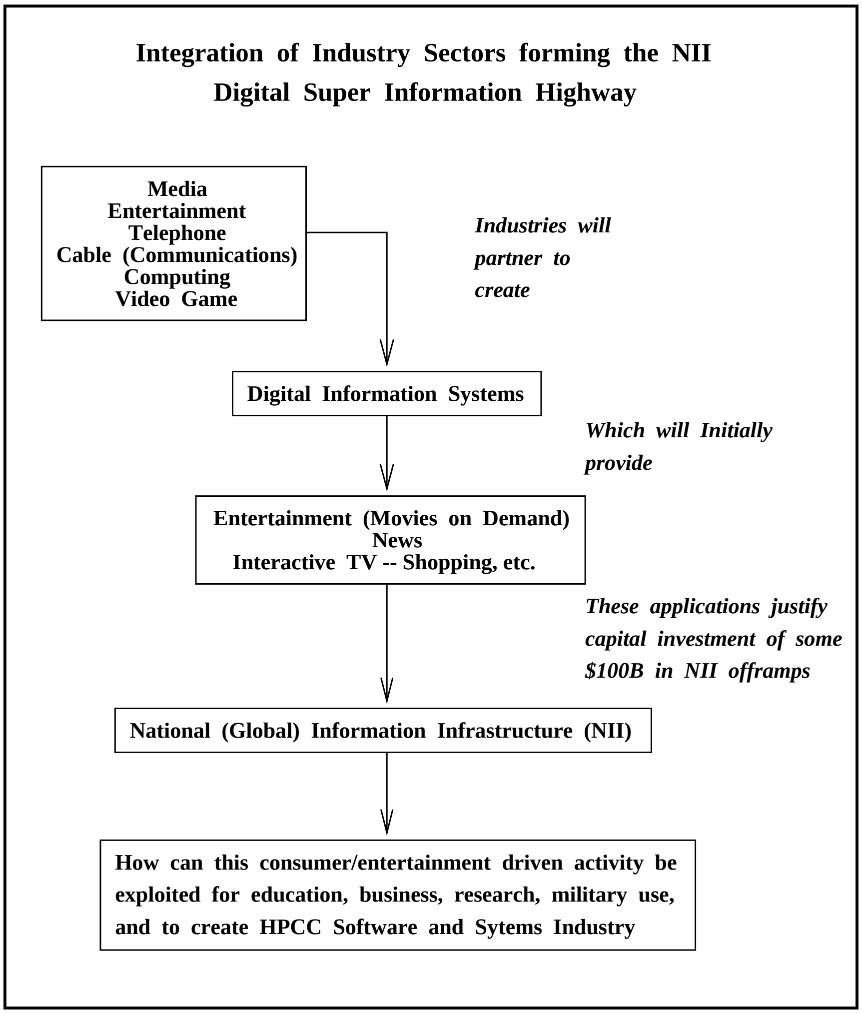

Figure 3: The convergence of industries and capital investment needed to deliver digital TV service

NII Compute & Communications Capability in Year 2005 - 2020

100 Supercomputers at a Teraflop each	**10^{14} (Fl)ops/sec at 100% Duty Cycle**
100 Million NII Offramps or Connections at realtime video speeds	**10^{14} bits to words/sec at about 10% to 30% Duty Cycle**
100 Million home PCs, Videogames or Settop Boxes at 100-1000 Mega(Fl)ops each	**10^{16} to 10^{17} (Fl)ops/sec at about 10% to 30% Duty Cycle**
1,000 to 10,000 High Performance Multimedia (parallel) servers each with some 1,000 to 10,000 nodes	**10^{15} to 10^{16} ops/sec at 100% Duty Cycle**

Each of three components (network connections, clients, servers) has capital value of order $10 to $100 billion

Figure 4: An estimate of the communication bandwidth and compute capability contained in the NII and supercomputer industries

modem on a twisted pair (plain old telephone service POTS) connection.

Returning to Figure 4, we estimate comparable investment in clients and servers and find an *InfoVISiON* or *WebServer* market that is at least an order of magnitude larger than that for supercomputers. This illustrates that it is the NII, and not large-scale number crunching, which is the best opportunity for parallel processing. Notice also that the compute power contained in this future NII is some 10–100 PetaFLOPS—far larger than the compute capability of an individual TeraFLOPS supercomputer. This motivates our interest in WebWork, described in Section 3, and aimed at extending Web to compute servers, and harnessing the power of the World Wide Metacomputer.

One could view this *InfoVISiON* scenario as the most amazing client-server application with 10^8 clients and 10^4 large servers. However, Vice President Gore has articulated NII democracy with everybody able to host

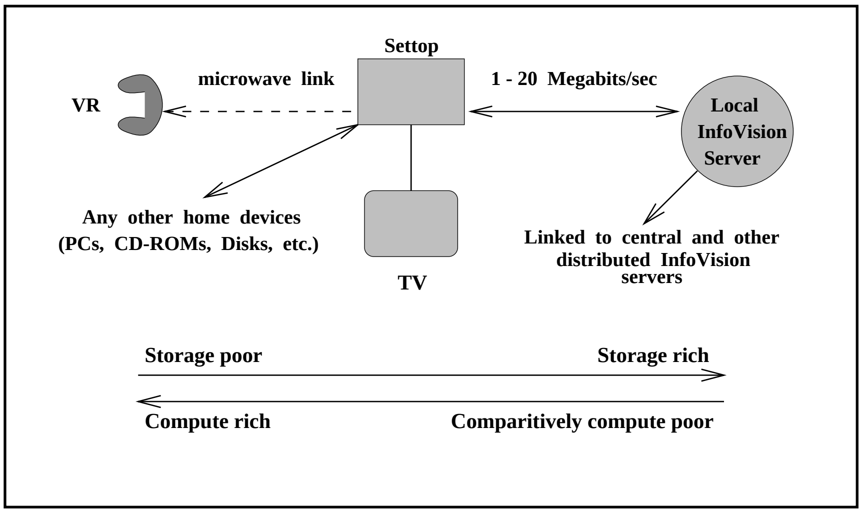

Figure 5: The basic *InfoVISiON* scenario as seen by a home in the year 2000 with an intelligent settop box interfacing the digital home to a hierarchical network of InfoVISiON servers

information. Further new technology, just as Java and WebTools (Section 3), suggests instead a heterogeneous server-server architecture with 10^8 distributed nodes. This can be viewed as a fascinating parallel computer with many more nodes than traditional tightly coupled systems (by a factor of 10^5 compared to typical large 1,000 node MPPs). It appears that although powerful, the communication backbone will not allow every "client" long distance simultaneous access to every other client or server. Rather, we must enforce the guiding principle of all computer architectures—namely, data locality. This is illustrated in Figure 6. When Jurassic Park VI is released on the Hollywood Server, one will not have everybody accessing it there at an average distance of some 1,500 miles. Rather, this "hot" movie will be cascaded down through the hierarchy of servers so that any individual will find it on a server a few miles away. This strategy reduces the needed fiber for the NII trunks by a factor of about 100.

Figure 7 illustrates a related point. We consider a WebServer as a system implementing the software model of Figures 1 and 2. This could be a server with full input and output access to the World Wide Information, or it could be aimed at a smaller organization with firewalls, as necessary,

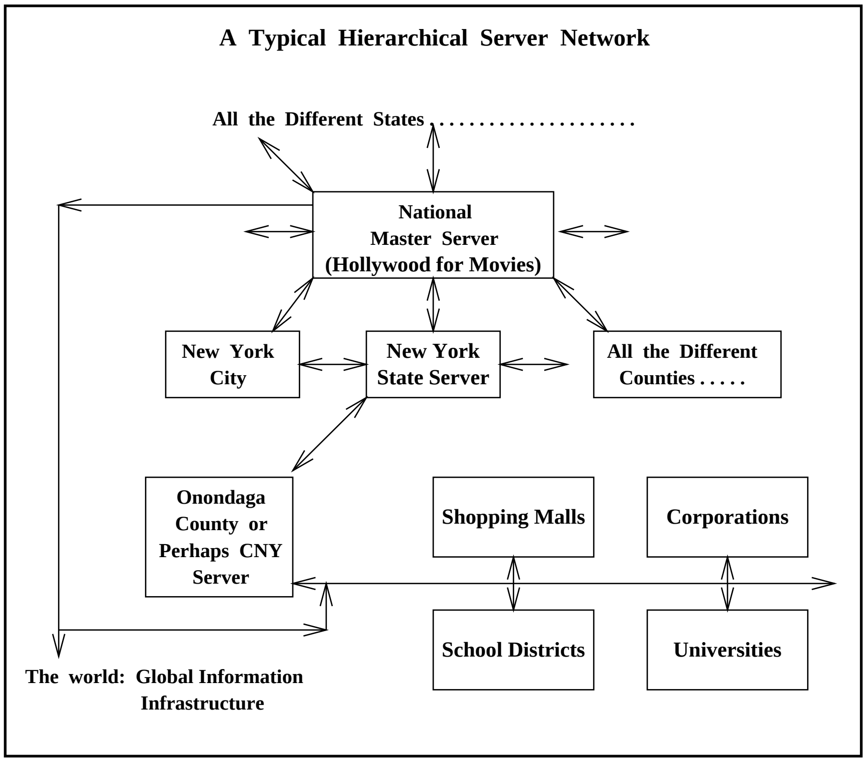

Figure 6: A typical hierarchical server network depicted for a master system in Hollywood cascading down with a fragment of node systems shown for central New York

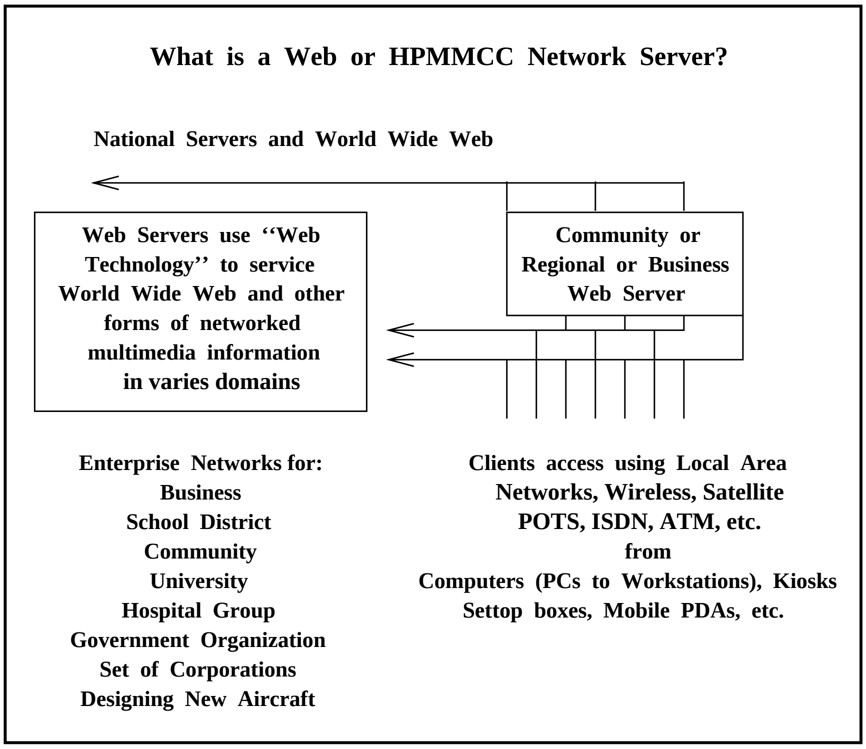

Figure 7: A Web server can be used at any level of organization from an individual home to the entire universe

to enforce security. We see that the mix of distributed and conventional (PC, Mainframe) software of Tables 1, 2, and 3 makes WebServers the most attractive technology to build a business enterprise support system.

3 Features of the Future WWW and Implications for HPCC

3.1 WebWindows and WebTools

We assume that current operating systems for individual computers—such as UNIX, Windows95, Apple Macintosh—will have less significance in the future. Rather, we will use the open non-proprietary WebWindows, which is the operating environment of the Web built from the collection of current

Table 3: Some emerging high-level WebWindows integration concepts

- **WebTools**—Early NPAC Prototype of WebWindows Equivalent to Program Manager with Navigation, File manipulation, Mail
- **WebDeskTop Publishing**—an early killer application under WebWindows supplanting Word, Wordperfect, LOTUS123, Persuasion, etc. Java allows clear powerful implementation
- **WebDBMS**—Integration of Relational and Distributed databases with both agent based heuristics, formal indices, and free text search
- **WebSpace**—Televirtual collaborative environment, such as implementation of full 3D MOO like environment building on LabSpace at Argonne for the virtual scientific laboratory
- **WebFlow**—NPAC prototype of Web based extended Khoros/AVS supporting dataflow linkage of computers for simulation and people and data for workflow management
- **WebScript**—the evolving Middleware of scripted languages including extended PERL5, Java, Telescript, MOVIE (NPAC compute oriented script), etc. [10], [11]

and emerging Web technologies [12]. WebServers will "run WebWindows" and so this environment can manage systems at all levels of granularity—from individual machines to the full Web metacomputer. We can illustrate the meaning of this concept with a typical productivity tool like Microsoft Word. Currently, one develops a separate version of Word for each machine—PC or Macintosh—and lament its unavailability for UNIX. In the future, one will develop a single product *WebWord*, which uses a suitable combination of client and server Web technology so that the machine dependence is isolated in Web servers and clients, and applications are architecture independent so that through a universal browser, they present the same interface on all machines.

Wojtek Furmanski, at NPAC, has illustrated the WebWindows vision with a set of CGI programs, *WebTools* that implement such universal tools for file management (create, delete, copy) and navigation—roughly equivalent to *finder* on the Macintosh or *program manager* on the PC; a sophisticated Webmail built on top of *mh*, the UNIX mail system; and HTML editor.

As shown in Table 3, we see that there will be a complete set of "Webtop" productivity tools supplanting and extending functionality of Word, Excel, Persuasion and similar tools. These will join capabilities indicated in Tables 1 and 2, and further applications to create a WebWindows environment that cannot be matched on any individual computer running conventional operating systems. WebWindows draws its power partly from the distributed computing simulation and information processing power—as recorded in Figure 4; more important is the creative energy of the "Web software engineers" (hackers) community interacting in the design and implementatio of WebWindows. Critical to success is open standards for interfaces and protocols—quality proprietary modules but with open interfaces are quite consistent with continued success of the Web. It is interesting to contrast the dynamic organization of the WebWindows development "team" with the more traditional static hierarchical structure used by IBM, Digital or Microsoft in the development of their major operating systems.

3.2 WebWork—A Pervasive Technology Base for HPCC

In many ways the HPCC program has been a great success and demonstrated as we like to say that "Parallel Computing Works!" [13]. However, MPPs remain rather difficult to use. This partly reflects inevitable problems in parallel programming, but also the immaturity of current MPP software development environments. We know how to do build much more powerful MPP software systems than are currently available. This is not surprising, as parallel (and distributed) computers are clearly the most complex computer systems. Thus, one would expect that it would require more effort to build a software environment for an MPP than for PC. Unfortunately, as shown schematically in Figure 8, the PC market is about two orders of magnitude larger than that for MPPs. Correspondingly, the software environment for PCs (and to a lesser extent) workstations must be much better than for MPPs. The situation can only improve if the size of the MPP market increases dramatically, and this currently appears unlikely in the technical computing area. We notice that the WWW is, as described in Section 2, a distributed computing environment associated with pervasive technology base and a corresponding large and vital software development infrastructure. In WebWork, we suggest, in collaboration with Boston University and Cooperating Systems, building a parallel programming environment on top of "Web Technology." This is not a terribly well defined statement as it implies and assumes an extension of the Web in many ways from information servers to combined compute-information servers. These will feature many enhancements from today, including integrated security and low latency HTTP-NG protocols connecting multi-threaded Web servers [14]. Much of this is discussed in [15]. Figure 8 shows a traditional HPCC strategy of porting MPP technologies so they can run on more available technologies, such as networks of PCs or workstations. WebWork takes the opposite approach of extrapolating from the base to the tip of the pyramid. The classic HPCC approach has the problem that it does not naturally produce technology used by and hence supported by a broad community. WebWork has a potentially equally serious problem that the extrapolation might "miss the top of the pyramid" i.e., produce a system that either did not meet the needs of parallel computing and/or produced very inefficient parallel code. We believe that careful design can avoid this problem. For instance, we are leading a group developing an open parallel compiler runtime that will support HPF (High Performance Fortran), and parallel C++ on several platforms [16], [17]. This runtime embodies key parallel computing synchronization and collective communication and computation primitives. As shown in Figure 9, Webwork will reuse such software, but provide an attractive front end and use suitable low latency Web compatible message passing systems. This concept when refined with a careful mix of interpretative and compiled environments leads to *WebHPL*—a general parallel language combining the lessons of expressing and implementing parallelism

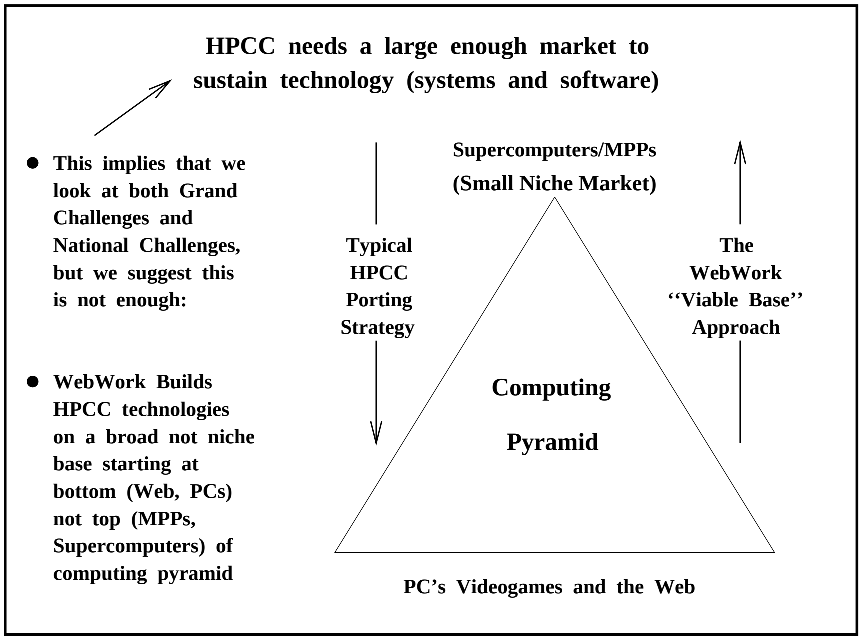

Figure 8: A contrast between WebWork's upward extrapolation from broad to niche markets with the traditional porting of MPP technologies

coming from HPCC research with the WebWork advantages of building on Web technologies.

Often, we have worried about the concept of parallel software engineering (PSE), but found no convincing approach. Now, I realize why I had difficulties. Good software engineering requires a good support (productivity) environment, and this is absent from all current parallel systems. WebWork has a clear mechanism for PSE by using the natural linkage of information and computing in the extend Web—this we term the Virtual Software Laboratory. We will illustrate this in Section 3.4 with our WebFlow concept implementing compute, information, and project management in a uniform dataflow framework.

3.3 Some Emerging Web Concepts and Technologies

In the near term, many new innovative technologies will be critical in building WebWindows as well as services and applications on top of it. These

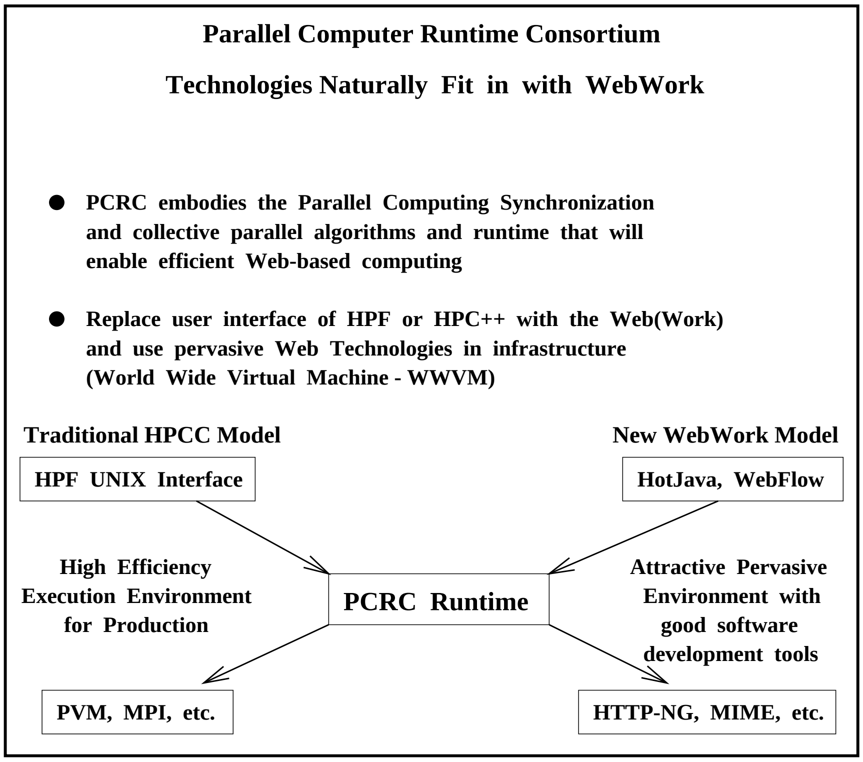

Figure 9: Integration of HPCC technologies with the Web illustrated for the parallel compiler runtime

include:

- **Java** [18]—This is a very powerful interpreted C++ like language (script), which, as in Figure 1, can be used to build fully interactive clients that allow one to build balanced client-server (server-server) systems. The CGI mechanism, as used in WebTools, is powerful but has limitations as extra functionality can only be added at the server.

- **VRML** [19], [20] can be thought of as a three-dimensional HTML that allows universal description of physical objects and graphics actions on them. There is some confusion as to relative importance and role of VRML and Java, but roughly one can consider VRML as the future data structure and Java the future language of the Web. VRML will be useful in such areas as telecollaboration, multiplayer gaming and distributed manufacturing, because it allows the interchange of virtual worlds and commercial product designs.

- **PERL5** is a useful extension of PERL4 with full object oriented characteristics and extended pointer (array) constructs. Again, there is some unclear competition with Java, but the languages are optimized in different ways. Java is partially compiled and will produce code that executes faster than that from PERL(5). PERL5 is better suited for rapid prototyping and further has excellent special capabilities built in for dealing with operating system functions and text processing.

- Performance of future combined Web-Compute Servers will be enhanced by good light weight **multi-threaded systems** combined with new protocols, such as HTTP-NG.

These technologies are linked together into loosely coupled integrating concepts summarized in Table 3. We have already described *WebTools* as an early prototype of WebWindows illustrating primitive Web operating system services. Again, *WebTop* publishing and productivity underlied our discussion of software engineering in *WebWork*.

As already mentioned in Section 2.2, we see Web database systems *WebDMBS* combining distributed (as in Harvest, Lycos, etc.), object (as in VRML) and relational (as an Oracle) capabilities. In Figure 10, we show three extensions of base PC multimedia services [21].

- HPCC with high-speed networks and MPP servers

- Immersion with virtual reality metaphor for spatial navigation

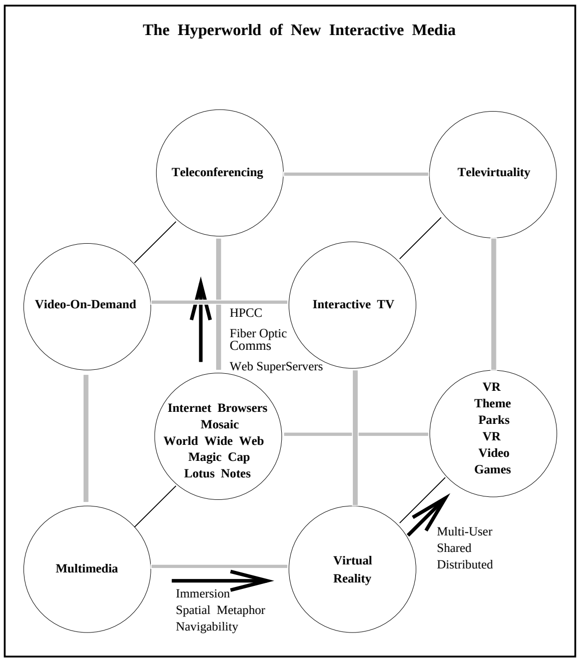

Figure 10: The three axes of multimedia extensions to give interactive services with different characteristics

- Multi-user shared distributed virtual worlds of information and simulation.

WebSpace denotes this combination leading to a full televirtual Web collaboration environment.

An important deeper and perhaps more controversial concept is *WebScript*. This denotes the complex NII middleware of scripted languages where we expect no universal solution but a loose federation where each component has different optimizations—VRML for three-dimensional objects, PERL5 for text, Telescript for agent based communication, Java, and MOVIE [22] for computation, etc.

3.4 WebFlow—A Simple Web Dataflow Interface

We consider WebWork as having three layers

- *World Wide Virtual Machine*—the basic network of Web-compute servers supporting MPI, HTTP, MIME, etc. data transport protocols and format.

- The intermediate integration layer, *WebScript*.

- At the highest level, we have many domain specific user interfaces optimized for different purposes. There are, of course, *Web browsers* for reviewing information and *WebHPL* for parallel programming. There will be several other such environments, and below we describe one example, *WebFlow*.

Another simple rather general high-level interface is *WebFlow* under development at NPAC. WebFlow is particularly interesting because the same framework can support both computing and the software engineering process that produces the software for the computation. Thus, it illustrates the power of WebWork's linkage of computing and information processing. Dataflow is a very useful programming paradigm first popularized with AVS and Khoros aimed at visualization and image processing. However, it has since been successfully applied to general coarse grain software decompositions [23], [24], [25], [26]. We are building this functionality into WebWork using Java with an early prototype of the visual *Computer-Web Editor* shown in Figure 11. This editor allows one to place compute modules and link them together. Figure 12 illustrates an example of WebFlow application in the area of workflow management—a software project management. Here, we suppose each software developer runs a personal *WebTools* server implementing the hyperlink enabled CASE tools *Hysource* supplied with WebTools. These personal servers are connected (via the *WWVM* in *WebWork*) to agent and manager servers also shown in Figure 12. The agent (database configuration) server receives automatic notification from developer servers on each software volume update. The agent server uses customizable thresholds to decide when to send a report to the manager or a deadline reminder to a developer. Note that the underlying database is stored in distributed fashion on the collection of WebTool servers.

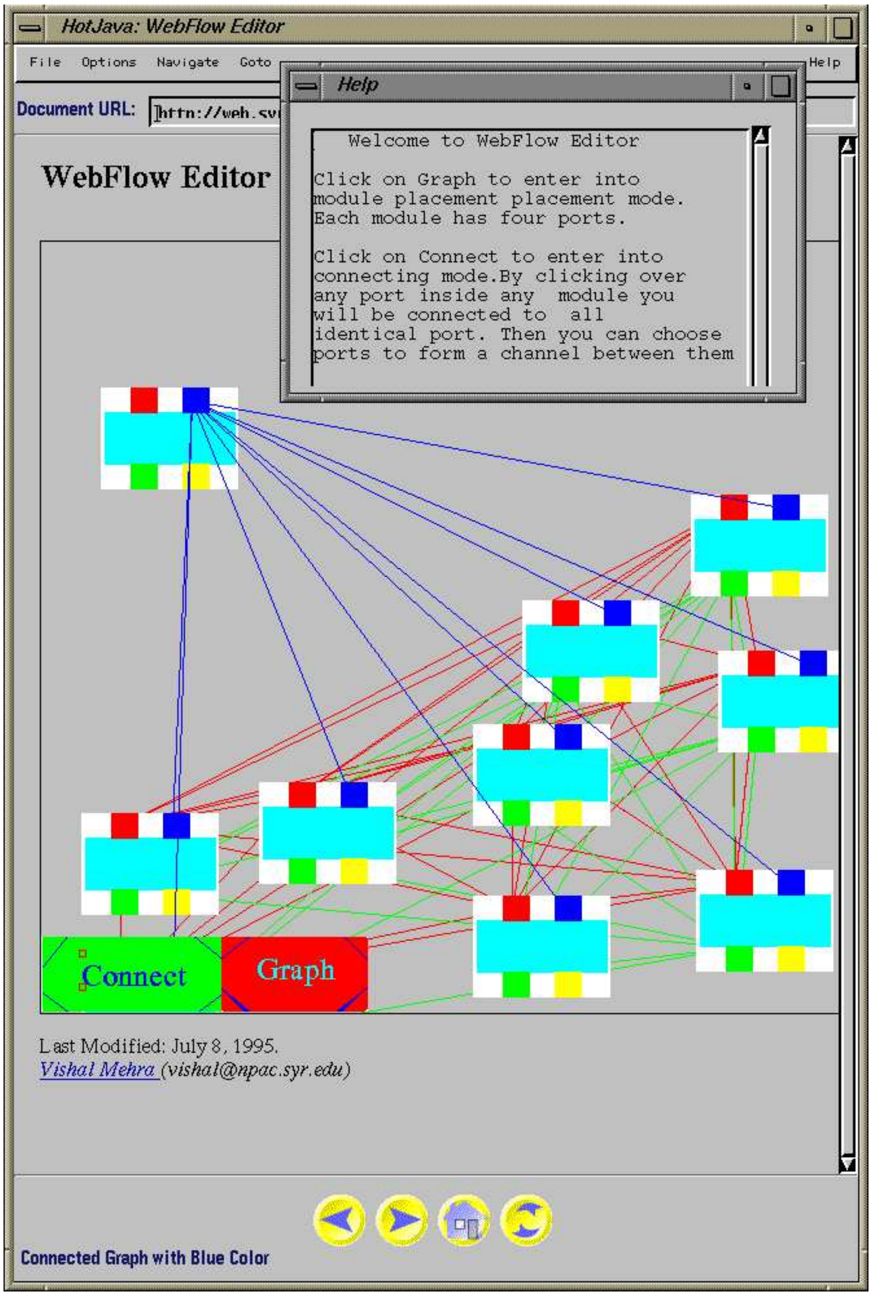

Figure 11: Early prototype of the Java-based WebFlow Compute-Web editor allowing interactive placement and linking of software modules

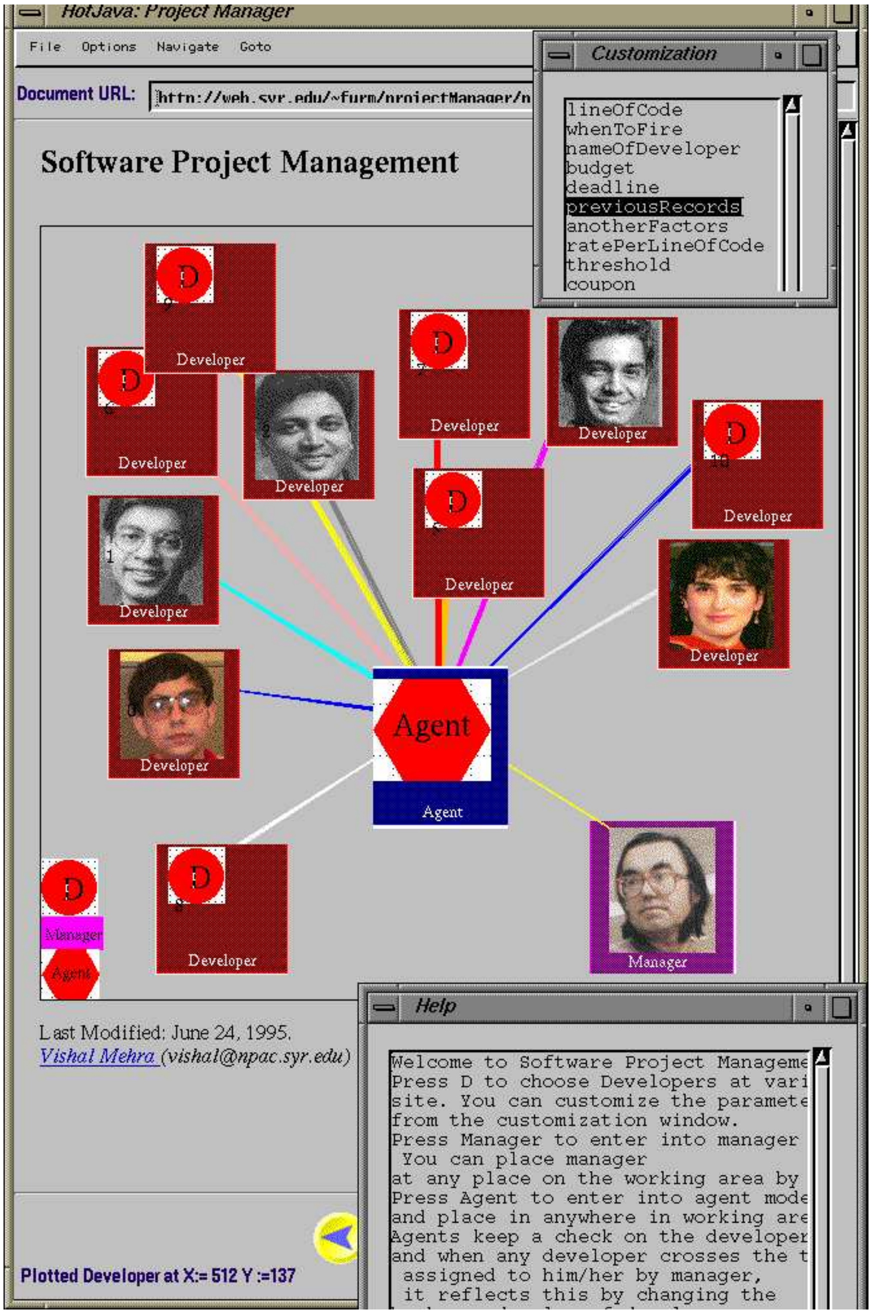

Figure 12: The Java front-end for a software project manager tool under development at NPAC. Developers, software configuration agents and project manager modules are supported

4 Some NII Applications and Services

4.1 Introduction

Returning to Web software architecture of Figure 1, we now discuss services and applications. These are not clearly defined for services are essentially generic applications, and most applications are complex metaproblems [27] built recursively from services and "sub-applications." Thus, there is a grey fuzzy line distinguishing services and applications. For this paper, we have chosen seven applications discussed in Sections 4.2 and 4.3, which we relate to the five services described below.

- **WebTop Services**—Publishing, Productivity, Software Engineering

  This includes all the base WebWindows functionalities that we have already discussed.

- **InfoVISiON**—Information, Video, Imagery, Simulation, ON demand

  This includes base database storage, management, query, and dissemination of the full range of multimedia archives of the World's distributed digital libraries. As already discussed, we can expect hundreds of Terabytes of available information dominated by video data. Note simulation—such as access on demand to a weather model—is included in this service.

- **Commerce**—Digital Cash, Security, Authentication, etc.

  This collection of services enables electronic commerce, including online banking and shopping. These services are also essential for the use of the WWW for processing and exchange of proprietary data.

- **Collaboration**—Real-Time Interactive and "Batch"

  This includes desktop video conferencing, three-dimensional graphics MOOs, geographically distributed CAVEs leading to full televirtual interactions. As discussed earlier, a wide variety of other types of interactive information exchange is necessary. This underlies the concepts of collaboratories (virtual research groups or scientific laboratories), and the virtual company of the next century's agile manufacturing environment. In the more static mode, we see workflow and configuration control (cf. Figure 12, which allows tightly integrated projects, such as those to build a complex system such as an aircraft (see Section 4.3) or a large software module with a distributed team (Section 4.2(f)).

- **Metacomputing**—The worldwide collection of computers organized in a subgroup as a single computational engine for simulation or information processing.

  This service can be used to control remote medical and scientific instruments; search the world for information; or link computers in different companies for a multi-disciplinary optimization of a new vehicle.

Some services listed above can be already prototyped in terms of today's Web technologies, for example base WebTop or early Collaboration services. Some other services are still waiting for their pervasive enabling technologies, such as physical infrastructure that will enable InfoVISiON or security that will enable Internet Commerce. Finally, the computationally extensive NII services, characaterized above broadly as "Metacomputing" require a major extension of the whole Web paradigm, currently still focused on static page services, but already gradually expanding towards computation and interactive simulation via technologies such as Java, or WebWork discussed in Section 3.4. Based on current and emergent Web technologies discussed in Section 3, we can start looking into the future to envision the coming generation of critical NII applications and high level services discussed in the following section.

4.2 Some Important Applications of the NII and Critical Services

a) Health Care and Telemedicine [28]

Critical to much NII use in health care are large-scale distributed databases to store patient records and medical instrument data. This should lead to more cost-effective health care with sophisticated database search techniques reducing fraud and waste while the large collection of uniformly prepared records will help emergency care, and the identification of "best of practice" care plans for particular medical problems.

Telemedicine (or remote medical care) naturally uses collaboration services, as well as the metacomputer capability for control of remote instruments. This could extend to use of a full televirtual environment for virtual-reality controlled remote surgery in battlefield or accident scenarios.

b) Education [29]

The NII is expected to make major impact on education both at the K–12, and University level. Indeed, the concept of the Virtual University [30] indicates that some fundamental concepts underlying the traditional residential college may need to be rethought. The bases of the educational use of the NII is collaboration and InfoVISiON. Rich collaboration environments including three-dimensional MOOs will support student-student and student-teacher interactions around the globe, and help students in rural areas and those home sick. InfoVISiON will allow full interactive and explorative multimedia curricula that can be expected to be more successful than current analog video delivery of distance education. We are currently working on Java and WebFlow based tools to support educational delivery systems.

Note that hyperlinking capabilities of the WWW are just as important as multimedia, for it enables student controlled exploration as opposed to current rather rigid model of learning embodied in a book. Geographical Information Systems will allow virtual field trips using the NII.

c) Society

The initial motivation for the installation of commercial NII communication instrastructure is digital delivery of current CATV and video rental services. This InfoVISiON capability will naturally generalize to a full range of multimedia products on demand, but it is unclear what novel services will be profitable. The NII will also enable advanced community networks that will given an interesting new view of local society, and could have a very important impact on local government. Commerce services will enable shopping and banking on the NII. InfoVISiON services would include multimedia digital yellow pages and online catalogs. The yellow pages could use an NII implementation of geographical information systems while virtual reality interfaces could allow you to try out new products such as cars and clothes. Collaboration services would support gossip as in today's computer forums and generalize the popular talk shows on radio and TV.

d) Business Enterprise Systems

Current OLTP (Online Transaction Processing) systems include these used in approving and recording credit card transactions. This is a type of InfoVISiON service which is operational today. It has major demands for security and reliability, but is comparatively undemanding on computer and network performance. On the other hand, decision support will be an important NII (InfoVISiON) service which could, as in Section 2, involve the search of a single giant relational database or a galaxy wide search of heterogeneous distributed database. Decision support is of growing interest in areas such as fraud detection, inventory planning, and market segmentation.

Both distributed business enterprises and scientific collaboratories can use workflow support as in Lotus Notes, which integrates collaboration, project management, and distributed databases. We can expect enterprise system support based on Web technology to be a major opportunity for early successful WebWindows applications.

e) Defence (Command and Control) and Crisis Management

A critical feature of this class of application is the need to make quick decisions in the presence of incomplete inaccurate data. We can term this scenario as judgment (as opposed to decision) support with the less precise context leading to the need for an InfoVISiON exploratorium. This will support multiple, but related video streams (say from different reporters of a given event) that need to be linked in time. Critical features of the network used to support crises is that it links a real-time, "come as you are" adaptive mix of computers and people. This is a complex metasystem mixing dynamic metacomputers and collaboration.

There are major needs to search community and government databases for immediate information on resource location and availability. As these databases are likely to be unfamiliar to judgment makers, metadata summaries in common formats are essential.

Collaboration will be required to link commanders in the field, specialized anchor disks (domain experts) and those in the field. A special issue will be the importance of excellent user interfaces as users will be tired and searching unfamiliar data. Geographical information systems will be very helpful as judgments will depend critically on spatial data.

f) Collaboratory—The Virtual Science or Engineering Laboratory

As desired by Wulf [31], the NII and collaboration services could enable much richer forms of scientific and engineering collaboration. This could link the national laboratories into a single superlab; it could enable innovative multidisciplinary research—another future trend—between researchers in industry and multiple universities; it could enable partnerships between major universities and developing minority institutions. Clearly, collaborative research and learning brings demanding requirements to future NII collaboration services.

Scholarly work traditionally involves the search of distributed databases and supported on the NII by InfoVISiON (digital library) services. Metacomputing will be needed for multidisciplinary multi-institution simulations, and for control, data storage, and analysis associated with remote instruments, such as satellites, telescopes, and accelerators.

The collaboratory will need workflow support, as in business enterprises (Section 4.2(d)) with special needs for distributed (over institutions) software engineering support for the production of distributed (over computers) software.

4.3 Use of the NII in the Manufacture of Complex Systems

We describe the application of the NII to the manufacture of aircraft, automobiles, and similar complex systems in rather more detail than the previous examples. This analysis stems from a NASA sponsored analysis of the NII requirements for a future concurrent engineering concept called ASOP (Affordable Systems Optimization Process). This involves an industry consortium MADIC (Multidisciplinary Analysis and Design Industrial Consortium) with a team from Rockwell, Northrop, Grumman Vought, McDonnell Douglas, General Electric, and General Motors. Interesting parameters specifying the scope of the design of the next major aircraft include:

- Construction will be led by a consortium of some six major companies and 20,000 smaller subcontractors.
- The number of engineers involved could be about:
  - 50 at conceptual design
  - 200 at preliminary design
  - 2,000 in final design
  - up to 10,000 for manufacturing and development

- ASOP involves multidisciplinary (multicomponent) optimizations (MDO) involving 10,000 separate programs that would be run in linked clusters (e.g., 10 programs at one time) for a set of specific design decision optimizations. These programs vary over a wide range from the full airflow simulation around a plane to a simple expert system to plan the best location of an inspection port to minimize maintenance life cycle costs.

  Such a manufacturing enterprise is an exciting and demanding challenge for the NII. First, note that the NII and its associated services is effectively essential for this application because the expertise and infrastructure needed for the design and manufacture of new aircraft is spread geographically through the country and perhaps globally. This expertise needs to be linked (by NII) to perform collaborative and coordinated design and simulation. If the NII did not exist, equivalent capabilities would need to be supplied by the involved companies, and indeed this has happened, using private lines as infrastructure, on some earlier projects.

  What are some key NII services needed by ASOP?

- Compared to the previous NII application's discussed in Section 4.2, manufacturing requires a close and integrated coupling of the very many people and computers involved. We are linking them in the design and manufacture of a very precise entity as opposed to the looser coupling required by, say, collaborative scientific research. Remote surgery in Section 4.2(a) is an example of another such close integration requirement for the NII.

- Metacomputing and distributed database support will have strong requirements to support the large number of linked programs connected to a logically central, but physically distributed design database.

- Workflow support must include configuration management and strong coordination with structured updates of the design database.

- Standards and security will be needed to link people and software from different organizations. In particular, efficient security for large files will be needed.

- Clearly NII and novel optimization techniques must present a good evolutionary path to allows re-use and incremental upgrading of existing software and people infrastructure. This implies good "wrapper" technology to support use of existing software modules with new interfaces.

- Finally, the NII collaborative services will be stressed by the close coordination of the large number of engineers needed in design process.

  We can use this application to briefly discuss the use of parallel computing in industrial simulations [32], [27], and [33]. This has not been as successful as many people hoped. For ASOP, we can see some aspects of the problem.

- Parallel processing can certainly be of value in simulations for propulsion, aerodynamic and probably other areas. However, these are a small fraction of the tasks (remember we mentioned 10,000 programs in ASOP) needed to design a new aircraft. Thus, we find a variant of Amdahl's law—parallel processing can effectively reduce needed computational fluid dynamics (CFD) simulation time. However, if this is only a fraction of x ($x \sim 0.1$ perhaps) of total endeavor, we have a speedup of at best $1 + x \sim 1.1$!

- As the aerospace industry adjusts to reduced DoD spending, and the construction of fewer military aircraft, it is hard for the industry to invest in new technologies with an unclear return on investment.

Thus, we see parallel processing on its own, insufficient to develop approaches to manufacturing. Rather, we need the integration of high speed networks and computers envisaged by the NII. Further, we require several basic NII services—security, metacomputing, collaboration, wrappers and agents, workflow with configuration management—to be well developed. Thus, we can expect the NII and, in particular, parallel processing to have a profound compact on manufacturing that will be great value to the National Enterprise. However, this is not an easy or short task, and we can expect significant government investment to be needed in basic precompetitive technologies and services. Industry is not likely to be able to make the necessary long term investments on its own. Correspondingly parallel computing will be used in a major fashion in manufacturing, but not in the near future, and not without continued thoughtful investment by industry, government, and academia.

4.4 Summary

We reviewed above the most plausible near term NII services and applications. Our claim is that this vision can be accomplished most realistically in the framework of pervasive, truly open Web technologies such as WebWork and WebWindows discussed previously in Section 3. Indeed, the traditional corporate computing failed to establish even the most elementary software standard, namely the cross-platform desktop publishing system. In contrast, Web already enjoyed a smashing success in stabilizing the base electronic information service—a homepage—and continues the process of developing and disseminating more advanced computing standards such as Java based graphics user interface etc. New Web technologies fully utilize and build upon already established standards, for example NPAC develops WebTools using WebTools CASE tools, or in a similar way Sun delivers Java together with complete on-line HTML documentation for all object classes. We expect the anticipated Web technologies discussed here such as WebWork and WebWindows to continue this bootstrap process and result in software environments of unprecendented power, functionality and quality, and yet fully open and pervasive. WebWork standards for module encapsulation and intermodular communication will enable complex manufacturing processes such as required by ASOP and will effectively turn our vision into the NII reality over the period of the next 5–10 years.

References

[1] Geoffrey C. Fox. Overview of NII for ASOP, July 1995. NASA Langley Workshop, http://www.npac.syr.edu /users/gcf/virtuniv95/fullindex.html.

[2] Tim Berners-Lee, R. Cailliau, J. F. Groff, and B. Pollermann. World-Wide Web: the information universe. *Electronic Networking: Research, Applications and Policy*, 2(1):52–58, 1992.

[3] A. Richmond. The Web developer's virtual library. http://www.stars.com/index.html, http://www.stars.com/Vlib/ (contains all relevant Web technology references).

[4] Gary Wolf. The curse of Xanadu, 1995. HotWired article WIRED 3.06, http://www.hotwired.com /wired/3.06/features/xanadu.html.

[5] Tim Berners-Lee. Internet RFC 1738: uniform resource locators (URL), December 1994. http://www.cis.ohio-state.edu /htbin/rfc/rfc1738.html.

[6] Tim Berners-Lee and Daniel W. Connolly. Hypertext markup language: A representation of textual information and metainformation for retrieval and interchange, 1993. http://info.cern.ch/hypertext /WWW/MarkUp/HTML.html.

[7] M. Andreessen. NCSA Mosaic: technical summary. Technical report, NCSA, May 1993.

[8] Harvest Team. The Harvest information discovery and access system. http://harvest.cs.colorado.edu/.

[9] S. Weibel, J. Godby, E. Miller, and R. Daniel. OCLC/NCSA metadata workshop: The essential elements of network object description. http://www.oclc.org:5046 /conferences/metadata/.

[10] Antony Courtney. Phantom: An interpreted language for distributed programming. In *Proceedings of the USENIX Conference on Object-Oriented Technologies (COOTS)*, June 1995.

[11] Kaleida Labs. Script X examples. http://www.kaleida.com /technical/exmplaes/index.html.

[12] General Magic, Inc. Magic cap concepts. http://www.genmagic.com/MagicCapDocs /Concepts/introduction.html.

[13] Geoffrey C. Fox, Paul C. Messina, and Roy D. Williams, editors. *Parallel Computing Works!* Morgan Kaufmann Publishers, San Francisco, CA, 1994. http://www.infomall.org/npac/pcw/.

[14] S. Spero. HTTP-NG architectural overview. http://www.w3.org/hypertext/WWW /Protocols/HTTP-NG/http-ng-arch.html.

[15] Geoffrey C. Fox, Wojtek Furmanski, Marina Chen, Claudio Rebbi, and James H. Cowie. WebWork: integrated programming environment tools for national and grand challenges. Technical Report SCCS-715, Syracuse University, NPAC, Syracuse, NY, June 1995. Joint Boston-CSC-NPAC Project Plan to Develop WebWork.

[16] J. Cowie and T. Haupt. Common runtime support for high performance Fortran. In *Parallel Compiler Runtime Consortium*, April 1995. http://cooperate.com/PCRC/DDDF.

[17] G. C. Fox, S. Ranka, M. Scott, A. D. Malony, J. Browne, M. Chen, A. Choudhary, T. Cheatham, J. Cuny, R. Eigenmann, A. Fahmy, I. Foster, D. Ganno, T. Haupt, M. Karr, C. Kesselman, C. Koelbel, W. Li, M. Lam, T. LeBlanc, J. OpenShaw, D. Padua, C. Polychronopolous, J. Saltz, A. Sussman, G. Wigand, and K. Yelick. Runtime support for high performance parallel languages. In *Supercomputing '93*, 1993. http://aldebaran.npac.syr.edu:1955.

[18] Java/HotJava Programming language and Web browser, April 1995. Alpha Release, http://java.sun.com.

[19] Gavin Bell, Anthony Parisi, and Mark Pesce. The virtual reality modeling language. Version 1.0 Specification, http://vrml.wired.com /vrml.tech/vrml10-3.html.

[20] Virtual reality modeling language forum. http://vrml.wired.com.

[21] G. Fox, W. Furmanski, P. Hornberger, J. Niemiec, and D. Simoni. Implementing televirtuality. *Applications of Virtual Reality*, 1994. British Computer Society, Computer Graphics and Displays Group, http://www.npac.syr.edu/PROJECTS /PUB/wojtek/hpsin/doc/tvr.ps.

[22] W. Furmanski. *MOVIE Multitasking Object-oriented Visual Interactive Environment.* Morgan-Kaufmann, 1994. http://www.npac.syr.edu/PROJECTS /PUB/wojtek/hpsin/movie.html.

[23] Gang Cheng, Chris Faigle, Geoffrey C. Fox, Wojtek Furmanski, Bin Li, and Kim Mills. Exploring AVS for HPDC software integration: Case studies towards parallel support for GIS. Technical Report SCCS-473, Syracuse University, NPAC, Syracuse, NY, March 1992. Paper presented at the 2nd Annual International AVS Conference *The Magic of Science: AVS '93*, Lake Buena Vista, Florida, May 24–26, 1993. http://www.npac.syr.edu/PROJECTS/PUB /wojtek/hpsin/doc/avs93.ps.

[24] Gang Cheng, Yinghua Lu, Geoffrey C. Fox, Kim Mills, and Tomasz Haupt. An interactive remote visualization environment for an electromagnetic scattering simulation on a high performance computing system. Technical Report SCCS-467, Syracuse University, NPAC, Syracuse, NY, March 1993. Proceedings of Supercomputing '93, Portland, Oregon, November 15–19.

[25] Kim Mills, Michael Vinson, Gang Cheng, and Finucane Thomas. A large scale comparison of option pricing models with historical market data. In *Proceedings of The 4th Symposium on the Frontiers of Massively Parallel Computing*. IEEE Computer Society Press, October 1992. Held in McLean, VA. SCCS-260.

[26] Kim Mills, Gang Cheng, Michael Vinson, Sanjay Ranka, and Geoffrey Fox. Software issues and performance of a parallel model for stock option pricing. In *Proceedings of the Fifth Australian Supercomputing Conference*, pages 125–134, December 1992. Held in Melbourne, Australia. SCCS-273b.

[27] Geoffrey C. Fox. Software and hardware requirements for some applications of parallel computing to industrial problems. Technical Report SCCS-717, Syracuse University, NPAC, Syracuse, NY, July 1995.

[28] L. Davis, J. Saltz, and J. Feldman. NSF workshop on high performance computing and communications, and health care. Report dated April 15, 1995 from Washington D. C. meeting, held December 8–10, 1994. http://www.umiacs.umd.edu /users/lsd/papers/nsfwork.html.

[29] K. Mills, G. Fox, P. Coddington, B. Mihalas, M. Podgorny, B. Shelly, and S. Bossert. The living textbook and the K–12 classroom of the future. Technical Report SCCS-731, Syracuse University, NPAC, Syracuse, NY, 1995.

[30] G. Fox and W. Furmanski. The virtual university presentation at conference tutorial HPDC-4, August 1995. Pentagon City, http://www.npac.syr.edu/users/gcf /virtuniv95/fullindex.html.

[31] Wm. A. Wulf. Collaborators. *Science*, 261:854, 1993.

[32] Geoffrey C. Fox. Involvement of industry in the national high performance computing and communication enterprise. Technical Report

SCCS-716, Syracuse University, NPAC, Syracuse, NY, May 1994. *Developing a Computer Science Agenda for High Performance Computing*, edited by U. Vishkin, ACM Press.

[33] G. Fox. HPC at the Crossroads! Academic Niche or Economic Development Cornucopia, Keynote talk at HPCS95 conference at Montreal, July 12, 1995. http://www.npac.syr.edu/users/gcf /hpcs95/fullindex.html.