Introduction to Web Technologies

Geoffrey C. Fox *

Abstract

This paper discusses basic Web technologies, PERL, Java, JavaScript, VRML, databases, and digital multimedia in contact of a layered WebWindows software architecture in which applications are built on top of multi-use services. We assume a community technology approach to both hardware and software. Perhaps the hardware model will be ATM connected PC's while the community software approach will not involve applications written for Windows NT/95 or UNIX, but rather for WebWindows, with interfaces defined by the standards of Web servers and clients. This universal environment will support WebTop productivity tools, such as WebWord, WebLotus123, and WebNotes built in modular dynamic fashion, and will undermine the business model for large software companies. We illustrate these ideas with some examples including business enterprise systems (intraNets), health care, financial services, manufacturing, and education.

1 Introduction

This article is meant to introduce and illustrate the use of six central Web technolgoies, which are described in Section 3. We assume the continuing trend to pervasive computers and networks implementing the National Information Infrastructure (NII) as the melding of today's Web, consumer PC, computer and communication infrastructure. Perhaps this will embody Gordon Bell's SNAP concept (Scalable Networks And Platforms), which suggests universal commodity PC and ATM network hardware with probably a growing use of Windows NT as the dominant operationg system. In Section 2.1, we describe the classic Web client server architecture and Section 2.2 describe sthe standard turn of the century ($\sim$ 2000+) NII scenario. Section 2.3 is devoted to WebWindows, which describes the new universal API (applicatio programmers interface) offered by the world viewed as a collection of Web servers and clients. We believe WebWindows is so much productive and powerful than conventional software interfaces that it will become the dominant and hardware independent "operationg system". In the final part of Section 2, we discuss NII applications and services briefly to show how they are built in terms of Web technologies and the WebWindows interfaces.

Section 3 discusses and illustrates our six chosne technologies—PERL, Java, JavaScript, VRML, databases, and digital multimedia.

Finally, in Section 4, we describe the role of the five applications—IntraNets and Business Enterprise Systems, Healthcare, Education, the Financial Industry, and finally, Manufacturing.

*Syracuse University, Northeast Parallel Architectures Center, 111 College Place, Syracuse, New York 13244, gcf@npac.syr.edu, http://www.npac.syr.edu

Figure 1:

Figure 2: Simple representation of “old” (pre-1996) WWW software architecture showing emerging advanced services, as well as basic technology

2

2.1 Client-Server Computing

Figure 1 shows a classic Web set up indicating it is “just” a client server application, and its implementation requires appropriate client, server, and network technologies. We will learn mainly abou the client and server software issues in this paper, and only describe enough about network hardware and protocols to put the other issues in contact. The original Web architecture, shown in Figure 2 was comparatively simple with clear functionality in each component. However, as the Web evolves, we move to the richer confusing scenario of Figure 3 with many overlapping concepts. Actually, the natural Web model is probably a server-server architecture, as all machines need capabilities embodied in servers. This fits with Gore’s view of the NII where everybody is both a consumer and producer of information. We can, of course, implement this (crudely) already by running a Web server and Web client on every machine (see Figure 4(b)). The growing functionality in clients is natural because, although individual servers are much more power than individual clients, the total compute power (and dollar investment) is an order of magnitude for clients than server (Figure 5)

2.2 The National Information Infrastructure (NII)

The SNAP vision recognizes that the World Wide Web (NII), parallel supercomputers, and your organization’s enterprise network will be built from the same pervasive communication infrastructure—twisted pair, coaxial cable, optical fiber, satellites and cellular phones connecting everybody and every institution together. Of course, different parts of this hardware will perform with different speeds and will be upgraded at different times in different parts of the country. However, a similar software base, Web technology, will dominate applications on this diverse infrastructure because it is built for the largest market, and therefore, will best leverage the software development effort into lower costs and higher functionality. Obviously, we now have an imperfect realization of this SNAP vision today, but there is enough infrastructure in pace to create the opportunity now being explored by America Online, Netscape, Microsoft, and others.

In Figure 5, we estimate the compute and communication capability of the NII and find it 100 to 1,000 times that of conventional supercomputers. However, we expect large-scale parallel processors to be used not only in classic supercomputers, but also in the network of video servers, and more generally, WebServers that supply information to these clients. We term this scenario *InfoVISiON* for *Info*rmation, *V*ideo, *I*magery, and *Si*mulation *ON* Demand.

Figure 3: 1996 Java/Netscape client-server model with evolving/confusing/overlapping capabilities

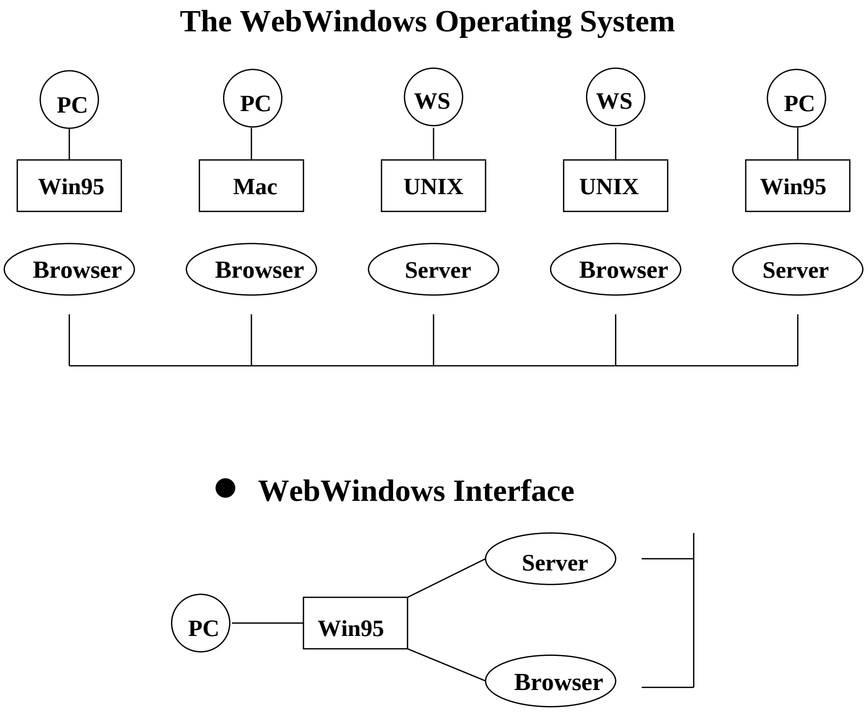

Figure 4: WebWindows for one PC or the World

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

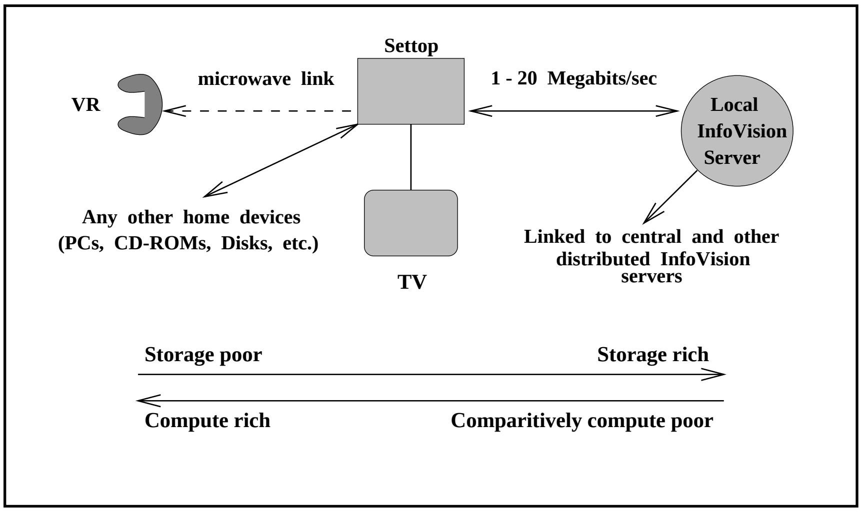

Figure 6: Basic InfoVISiON home scenario in the year 2000: Intelligent settop box interfacing the digital home to a hierarchical network of InfoVISiON servers

InfoVISiON 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 formats, 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 small percentage of its storage devoted to video, whereas the future NII will be dominated by video data.

Each NII offramp will, as shown in Figure 6, 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 28.8 Kbaud modem on a twisted pair (plain old telephone service POTS) connection.

Returning to Figure 5, we estimate comparable investment in clients and servers and find an nfoVISiON 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, that 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 is explored in more detail in [***reference "Concurrency: Practice and Experience/SIAM" articles, "Exaop/Petaop"***] that discuss extending the Web to compute services, 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 information. Further new technology 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 fac-

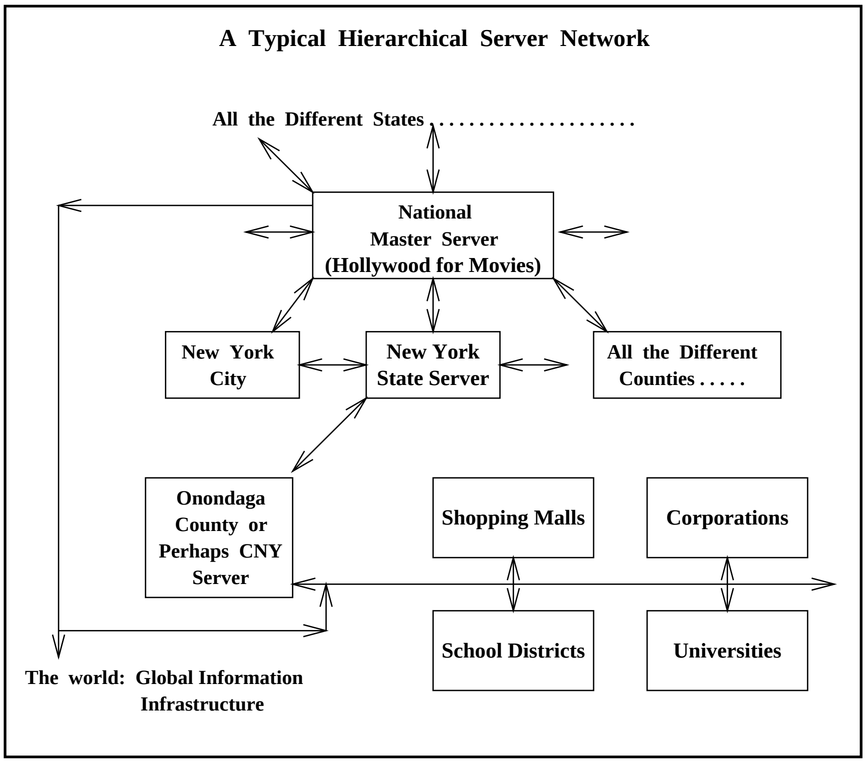

Figure 7: Typical hierarchical server network for Hollywood master system cascading down through a fragment of node systems shown for central New York

tor of 10^5 compared to typical large 1,000 node parallel processors). 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 7. 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.

2.3 WebWindows

Figures 2 and 3 illustrate the overall architecture of Web software linking multimedia servers and clients with standard interfaces and technologies, such as HTML, VRML, and Java—some key concepts are defined in Table 1.

We believe that the current trend to build applications for this architecture will accelerate. Now, we see WebMail, WebDatabase (see Figure 8), WebEditor, WebFoil, WebChat, but soon we see a set of more advanced products, such as "WebWord," which will offer the functionality (and more) of Microsoft Word (and similar products) but be built with Web technologies. WebWord is quite different (in architecture and hence implementation) from linking Word to the Web (Weberizing Word) and so Microsoft has no special advantage in this emerging

Table 1: WorldWide Web Technologies

Clients (such as Mosaic and Netscape)	support browsing of hyperlinked documents, but have no internal interactive/compute capability [Andreessen:93a]
Servers	read HTTP and deliver requested service to client
HTML	document format supporting hyperlinks
HTTP	transport protocol defining interaction between Web servers and clients
URL	universal resource locator is addressing scheme for items on the Web
MIME	data format allowing agent-like (extended e-mail) communication
CGI	standard interface allowing sophisticated server extensions
PERL	rapid prototyping language (Script) aimed at text and file manipulation
Java **JavaScript**	semi-interpreted C++ like language supporting an applet model of distributed computing interpreted language for integrating Web components
VRML 1.0	universal deScription for static three-dimensional objects
VRML 2.0	dynamic 3D deScriptor for virtual environments
JDBC (Java Database Connection)	set of methods linking Java to general relational databases

WebTop Productivity field. WebWord will use Web standards for all internal and external representations, and so allow easy integration of new functionality, and customization of WebTop environments for particular markets.

The Web will not only allow new technologies to develop better versions of current applications (such as Word) but also produce new uses of computers. The current Web roughly implements a client-server computing model, but as client-side Web technologies improve one sees an evolution to a new paradigm—the server-server computing model described earlier. This does not say that clients will disappear, but rather they will evolve to incorporate much of the functionality now in servers. In fact, current clients are much larger software packages than most servers!

Already, the Web links some 10,000 servers together, but technologies, such as Java and a growing use of interactive and collaborative applications will either directly (run a server on client machine) or indirectly (enhanced browsers) turn all clients into servers. Furthermore, JavaVM (Virtual Machine) supports client side multithreading, and so differences between (multitasking, e.g., UNIX) servers, and (multithreading, e.g., JavaVM) client technologies are, indeed, fading away. This, again, implements the democratic Web where all of us can be consumers and producers of information for the Web.

Remember that when we talk of products using Web technology, it does not have to be applied to the full world-wide Internet. Rather, WebWord could run a single PC which has both server and client capabilities. Again, Web linked relational or object databases (Figure 8 and Section 3.5), and a future WebLotusNotes product could use Web Technology and Webserver systems to support a Business Enterprise Information system or IntraNets on a closed corporate network.

The Web naturally supports virtual organizations with several small entrepreneurial companies linked dynamically together. It will be hard for the large monolithic systems and software companies to be competitive. There will, however, be an interesting systems integration business, which will put together and support linked sets of Web technology modules for particular services in targeted markets and organizations.

The above implicitly defines WebWindows as the target architecture for the future replacing UNIX, Apple Macintosh or the Win32 interface to Windows 95/NT. Bell argues that Windows NT will be the operating system of SNAP; we say that this may be correct but perhaps irrelevant as the real "action" will be at the higher level defined by the Web. In fact, the role and structure of operating systems could change for they need not support users directly, but just Web servers and clients. This changing architecture is shown in the comparison of Figure 9 and Figure 10 where the latter also illustrates the blurring of the distinction between (Web) servers and clients.

2.4 Generic Services on the NII

We can define a simple layered architecture for Web (NII) applications which are built in terms of multi-use services as shown in Figure 11. "Multi-use" extends the well known dual-use civilian-military interplay to a set of capabilities shared by many different applications. Note the Web is an excellent implementation technology for the COTS (customer off the shelf) choice used in many new defense software systems. There is no precise definition of services and their difference from applications, for services are essentially generic applications, and most applications are complex metaproblems built recursively from services and "sub-applications"

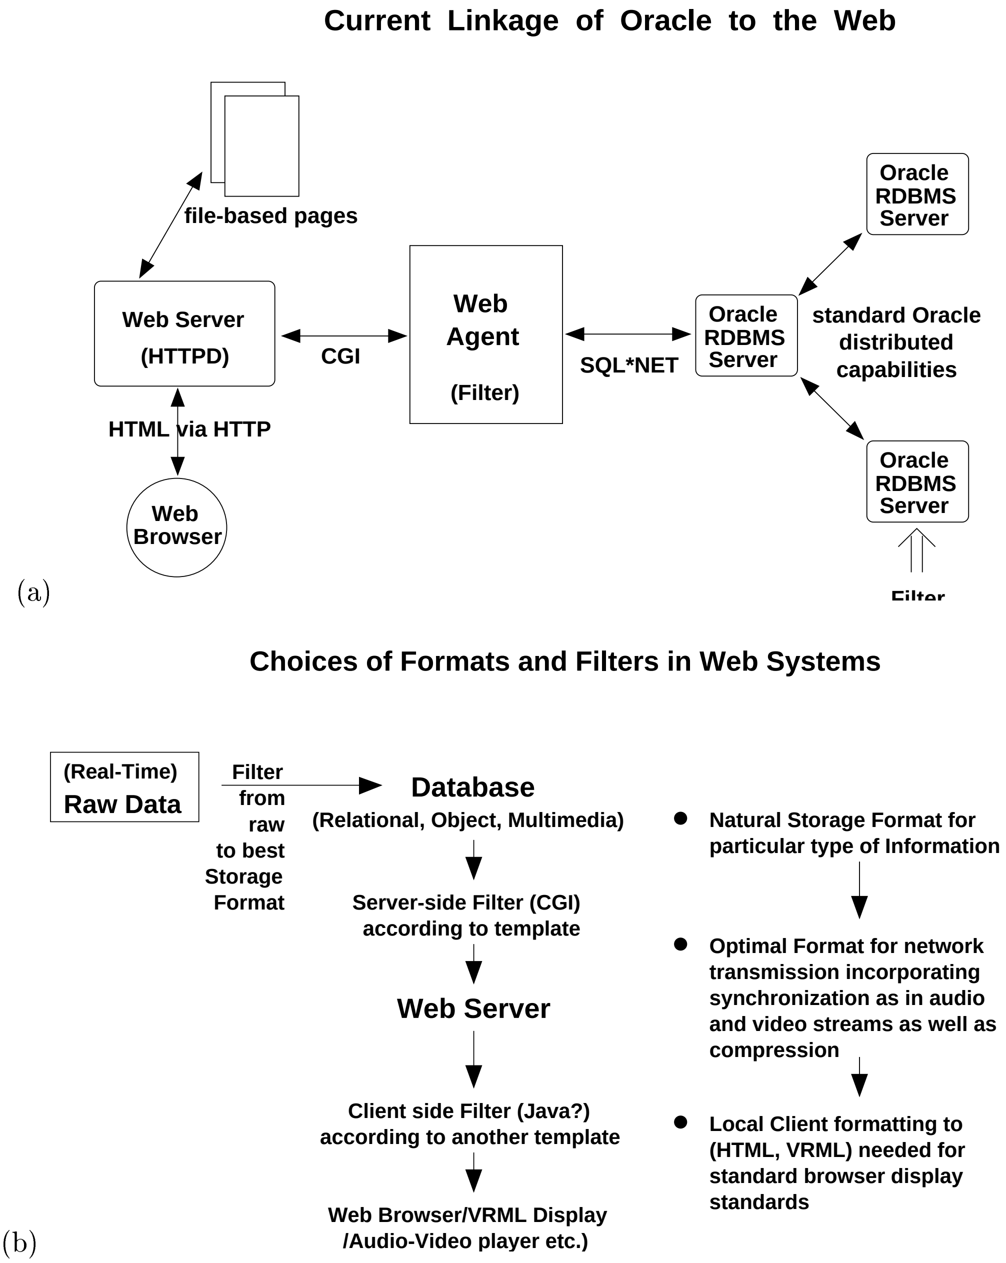

Figure 8: (a) Oracle-Web Integration Architecture; (b) Choices of Formats and Filters in Web Systems

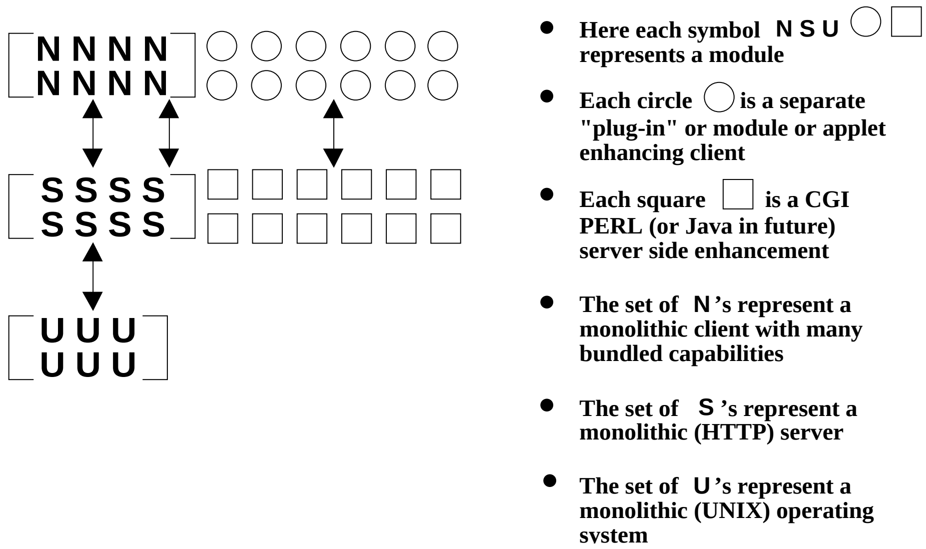

Figure 9: Architecture of Web Client-Server Software—Mosaic/Netscape Today

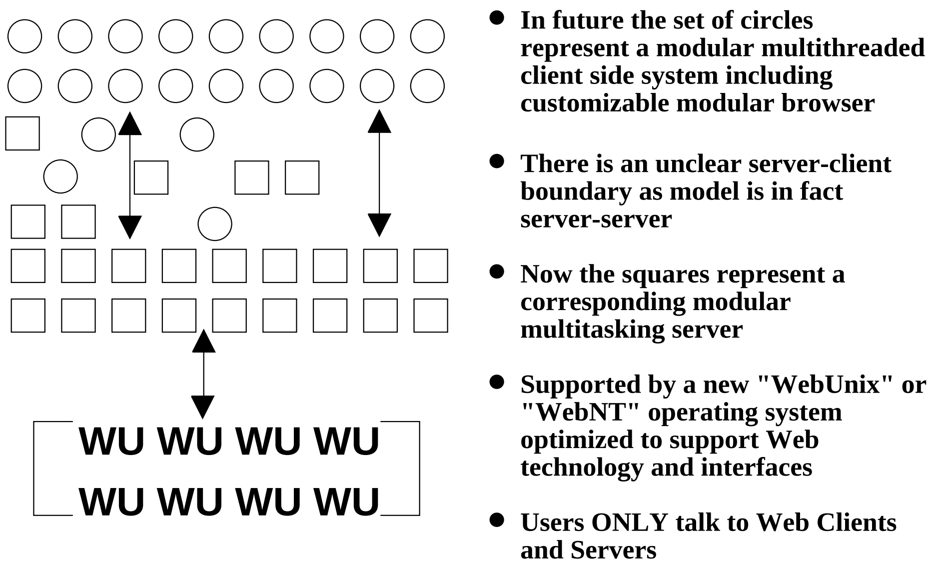

Figure 10: Architecture of Web Client-Server Software—The Future as suggested by Hotjava?

Table 2: Services used by seven NII applications

Applications	WebTop Productivity	Info VISiON	Commerce/ Security	Collab- oration	Meta Computing
Society	X	X	X	X	
Education	X	X		X	
Enterprise Systems	X	X	X	X	
Health Care	X	X	X	X	
Command and Control/ Crisis Management	X	X	X	X	X
Manufacturing	X	X	X	X	X
Collaboratory	X	X		X	

[Fox:95c]. Table 2 maps services onto seven application areas.

Thus, there is a grey fuzzy line distinguishing services and applications. Five possible NII services include:

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

  Here, we put base productivity tools including "WebWord," "WebLotus123," Web linkage to relationship databases, etc., an interesting prototype of such a capability. Essentially, all applications use this service.

- **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. We can expect the "digital TV" application to produce hundreds of Terabytes of available information dominated by video data. Note delivery of results of a simulation—such as access on demand to a weather model—is included in this service.

  In manufacturing, InfoVISiON corresponds to delivery of data in a configuration controlled database. In command and control, this service allows commander to index and access videos of battle damage videos or real-time engagements.

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

  This collection of services enables electronic commerce, including on-line banking and shopping. These services are also essential for the use of the WWW for processing and exchange of proprietary (manufacturing), and classified (military) data.

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

  This includes desktop video conferencing, three-dimensional graphics MOOs, geographically distributed CAVEs leading to full televirtual interactions. The emerging VRML 2.0

standards will be very important in building virtual environments. 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, which allows tightly integrated projects, such as those needed to build a complex system including an aircraft (see Section 4.4) or a large software module with a distributed team.

- **Metacomputing**—A worldwide collection of computers organized together as a single computational engine for simulation or information processing [Fox:95a] ***ref Petaop and Exaop***.

  This service can be used to control remote medical and scientific instruments; search the world for information; simulate the weather expected in a military engagement, 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 are now becoming available. 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, characterized 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 and VRML.

3 Some Specific Web Technologies

In this section, we discuss some of the specific technologies shown in Table 1 and Figures 2 and 3. We give simple examples of their use, and for languages, their syntax. We cover, in the following sub-sections, PERL, Java, JavaScript, VRML, Web-linked databases, and digital video.

3.1 PERL4 and PERL5—CGI Scripts

PERL is a very well-know interpreted welanguage that was developed (before the Web) as a powerful systems programming language. Although aimed at UNIX, PERL is also available on PC Windows platforms.

PERL is commonly used in so-called CGI Scripts that conveniently extend the capability of servers, as shown in Figures 2 and 3. In theory, one can use any language on a CGI Script, including higher-level shell Scripts or high-performance languages, such as C or C++. In fact, as shown in Figure 3, one can expect Java to be increasing used for server side extensions for both CGI and more sophisticated mechanisms. However, PERL currently has the most powerful text processing and (UNIX) file system handling capabilities of any common language, and is clearly currently superior to Java and JavaScript in this regard. The most common use of CGI Scripts is processing of HTML form input (see Figure 11) and this is dominated by text and file manipulation. Further, such work is often not compute intensive, and a flexible

Figure 11:

Figure 12:

interpreted implementation has adequate performance. Thus, PERL remains the language of choice for CGI programs. Further, many Web applications require substantial use of (off-line) utilities to reform text, and prepare information. Here again, PERL current is the best implementation language. Note that hybrid approaches for CIG Scripts are quite common with a PERL interpreted "wrapper" invoking special applications, such as database engine or an optimized C or Fortran numerical simulation.

When we describe JavaScript, it provides an alternative way of processing HTML forms solely on the client side. However, this is only appropriate for some Web applications, and many require server side processing either due to the complexity of the computation, or the need to access information stored on the server file systems. Thus, PERL and JavaScript offer complementary capabilities.

Java is already the preferred implementation of customized servers, such as those described in Section 4 for supporting collaboration. However, until Java has the same convenient text processing and flexible interpreted implementation, PERL will continue to be used extensively for CGI Scripts.

The most recent version, PERL5, has important extensions—especially in the area of multi-dimensional arrays and object-oriented features. However, although the new arrays features are very helpful, we find the object-oriented implementation much less elegant than Java, and it has not been broadly adopted.

In Figure 12, we show a simple "Hello World" program in PERL, showing the use of variable interpolation in character strings. In Figure 13, a HTML page allows the user to input a name that generates a "Hello World" HTML page using the PERL CGI Script in Figure 14.

3.2 Java and Its Use on the Web

Java is an object-oriented language that has similarities to C and C++. Its introduction, in 1995, had a major impact on the Web, as it allowed one to incorporate customized computer programs on the client (Web browser) side, as shown in the complicated Figure 3. Thus, Java, for the first time, allowed one to build balanced client-server applications. Note, Figure 5 points out that there is substantially more compte capability on clients than servers. Further, any server-based system is likely to have nontrivial latency (delay in response) and so Java can be used to build user interfaces where many relatively simple computations are involved, and the low-latency client-side approach is clearly preferable. However, Java is turning out to be of even greater importance than these Web architecture issues suggest. In particular, one expects it to challenge C++ and Fortran (***ref CPandE special issue***) as the future language of choice for software engineering and scientific computation, respectively. This trend is due to a variety of reasons. For instance, the WebWindows approach suggest future large software

Figure 13:

Figure 14:

Figure 15:

engineering projects will build to Web standards, which implies Java will often be a natural implementation language. Again, Java offers an attractive object-oriented environment with the potential of optimizing Java compilers being able to general comparable performance to Fortran. Further, we can expect a growing trend to learn Java as a first language, as children are exposed to it from Web browsing and find using Java to enhance their home pages as an inducement to learn programming at school. These positive indicators are tempered by the immaturity of technology, which is only to be expected, but implies today (1997) that one should use traditional languages to build robust high-performance languages.

Java is used on the Web in a semi-interpreted mode shown in Figure 15. Here, the *javac* compiler converts Java language (java) into JavaVM (class files), which can be thought of as universal machine code. The JavaVM files and libraries are downloaded as needed from the server to client. At the client, the "bytecodes" of the JavaVM are interpreted, checked for security issues, and converted to native machine code.

Java comes with a powerful runtime that is aimed in particular at the graphics, networking, and user interface features that one needs for interactive Web pages. There is, in particular, a comprehensive AWT (abstract windowing toolkit) that supports such things as buttons, forms, windows, and their layout.

In Figures 16 and 17, we give, respectively, the Java source, and invoking the HTML page that will dispaly "Hello World!" Note this ues built-in (runtime) classes, Applet and Graphics.

We can expect Java technology to improve in many ways that will increase its performance and functionality. We will see better "Just in Time" compilers—essentially, clever interpreters of JavaVM that generate reusable compiled code on the fly. We will see a wide variety of new libraries (runtime) supporting areas as diverse as commerce and scientific computing. Again, as well as the self-distributing software model of Figure 15, one will see native Java compilers that will get high performance by direct compilation of the Java language into the native machine code for particular machines.

3.3 JavaScript

JavaScript is a language that has some similarities, but many critical differences, from Java. One can view Java as an extension of a computer language (C++) and JavaScript as an extension of a text formatter HTML. Both are aimed at "animating Web pages" that is a superficial description of an interactive client-side application.

Figure 18 shows the simplest JavaScript "Hello World!" application. However, more interesting is the example of Figures 19 and 20, which shows a new HTML attribute ONCLICK being used to specify a JavaScript routine *compute* to process data specified in an HTML Form. We discussed, in Section 3.1, how this can be used to build very responsive client-side applications.

Figure 16:

Figure 17:

Figure 18:

Our final example, in Figure 21 shows how JavaScript can be used to generate parameterized HTML where JavaScript variables—possibly input in forms—are used to specify dynamic pieces of HTML pages. JavaScript is particularly useful in windows built from multiple frames as it allows actions in one frame to control the URL location and other properties of another frame.

Java was developed by Sun and is an industry standard. JavaScript was developed by Netscape and partially supported by other browser vendors including Microsoft. We believe that although it is no such a well established standard. JavaScript fulfills a clear role, and one can expect one or more technologies like this to be part of the Web "bag of tricks." In Section 3.2, we mentioned the AWT as the basic Java windowing toolkit. One can view JavaScript as the equivalent programmable AWT, but with the (Netscape) browser as the runtime. One can use the Java AWT to mimic JavaScript and have the advantage of a much more powerful and better defined language. However, JavaScript has a major advantage in that the current browsers essentially define a far more powerful runtime than that currently available to the Java AWT. In particular, the Hot Java browser, written entirely in Java, is currently not competitive with the standard non-Java browsers. JavaScript has a built-in advantage (independent of browser capabilities) that is fully interpreted, and it is much easier to rapidly prototype and modify new Web capabilities in JavaScript than in Java.

JavaScript is a more limited language than Java; JavaScript has object, but no iheritance as in Java; Javascript relies implicitly on the undefined (to the user) browser multitasking threads whereas Java offers the user substantial control over the execution environment with well-developed thread support in the language. JavaScript is used for relatively simple client-side document processing; Java can be used much more extensively including image processing, simulation, and networking on client and server.

3.4 VRML for the Real World on the Web

VRML was inspired by the simple observation that the world is three dimensional, and one needs to manipulate and view on the Web three-dimensional objects, as well as the images and text supported so well by the earlier technologies. The initial VRML 1.0 was not a programming language, but rather a data descriptor allowing one to specify arbitrary static three-dimensional scenes with arbitrary dynamic viewpoint allowing navigation. The three-dimensional scenes could embed clickable regions (spawning arbitrary Web locations) in a natural generalization of the two-dimensional image maps available in HTML. VRML was first proposed in Spring 1994 with the initial 1.0 version available a year later. There is now an extended standard VRML 2.0, which we will discuss briefly later on. Note that VRML is an acronym for Virtual Reality (VR) Modeling Language, but VRML is applicable to any application involving three-dimensional objects, and is more general than VR.

Figure 19:

Figure 20:

Figure 21:

We illustrate VRML in Figure 21, which shows six lines defning two unit size objects separated by three units in the x direction. VRML is based on SGI's Open Invetor graphics object database format with some restrictions and extensions allowing linking objects with particular Web sites. VRML three-dimensional worlds are built in a hierarchical fashion in terms of nodes grouped together by separators. *Cube* and *Sphere* are types of node as is *material* shown in Figure 22. Nodes have attributes, such as *diffusecolor* and *translation* in Figure 22 grouping is illustrated in Figure 23 where the extra separator limits the effect of the *Material* node to the cube.

A sophisticated use of VRML is shown in Figure 24 with a terrain overlayed by clouds from a weather simulation. However, the use of this technology has been handicapped by its need for high performance for both CPU and network. Some of this is inevitable—it takes more data to specify and more CPU processing to render three-dimensional objects than process HTML text. However, much of the performance degradation was due to simplistic language design (data was transmitted as uncompressed ascii) and immature VRML browsers. We expect a combination of improved hardware and refinements in language and software to increase the use and usefulness of VRML.

The new version VRML 2.0 addresses some of the performance issues, but more importantly adds the concept of a "behavior" that is essentially a computer program (script) that can be attached to the VRML nodes. This allows scenes with objects moving with respect to each other, and interacting in ways defined by the scripts. This will allow one to generate true virtual worlds and has natural application to both VR and Web-based computer games. Although there is an agreed VRML 2.0 standard, there are other approaches. One can extend the (*Java 3D*) Java runtime to support three-dimensional objects in a similar way to its current two-dimensional image support. Microsoft has an interesting active VRML alternative technology. Clerly, the ideas present in VRML 1.0 and VRML 2.0 will be part of critical base Web technologies; however, the particular implementation is much less clear than for Java, and we can only follow the Web community and see what paths are chosen.

3.5 Web-Linked Databases

This represents a different Web resource—namely, databases have been around a long time, but what's now is that they are now much cheaper and easier to use. For three years, NPAC used essentially no commerical databases, except for a single research activity using parallel Oracle. Now, we have some dozen activities using them. These include

1. newsgroup and "chat" multimedia messages stored in databases giving a searchable record of collaboration or discussion

Figure 22:

Figure 23:

Figure 24:

2. enterprise intranets with a Fortune 500 product catalog mode available to the public through NPAC's Web database interface 3) support of financial modelling "on demand" prototype
3. set of New York images indexed in a database for educational use
4. commercial CD-roms stored in databases for school reference
5. storage of close-captioned text and metadata to index video
6. electronic mail server with database replacing ink
7. storage of full text of books to allow rapid search
8. complete index of text in NPAC Web site
9. index of customized chemistry information
10. object database storage of three-dimensional maps on New York State
11. publication and other information from NPAC's technical reports.

All these applications use Web-linked relational (Oracle) or object (Illustra) databases. This explosion of use is easy to understand. Historically, we would have used flat files (or UNIX file system structure as, in fact, in a Web site information organization or many UNIX mail systems) store such data.

The trend to substantially lower database prices (especially on PC platforms) and the convenience of Web interfaces has charged this, and given us systems that combine the indexing and secure storage capabilities of traditional databases with the friendly HTML or Java interfaces of the Web. Figure **??**(a) shows a typical setup with a simple CGI Script linking Oracle to the Web. We have most experience with WOW, which has a very simple CGI Script, and the user query (input through an HTML form) is interpreted by Oracle's PL/SQL language. This has libraries HTP and HPF that make it relatively easy to produce HTML pages that are passed from the database server to the client.

PL/SQL is a convenient high-level language for generating standard SQL queries, but it is Oracle specific. We expect the so called Java Database Connection (JDBC) with a host of Java methods for accessing any SQL database to be very important in the future.

In Figure **??**(b), we emphasize that a complete system could well involve many filters as data is stored in natural application oriented form, but transformed to suit network (e.g., compression) or browser (e.g., conversion to VRML). Java is a convenient language for building such filters.

Figure 25 illustrates the Pl/SQL function to access a simple phone number database with simple HTML form (Figure 26 and resultant HTML (Figure 27).

Figure 25:

Figure 26:

3.6 Remarks on Digital Audio and Video

Multimedia is of growing importance for the Web. Given the expectation that the Web will subsume current TV and cable services we can expect increasing amount of multimedia information. This implies important consequences for both the underlying network, and for the client and server. Initially, the Web offered downloaded (MPEG compressed) video, but this is not an appropriate approach except for short clps. Rather, one must use streaming servers where data is passed continuously from server to client. This implies important quality of service constraints or the network as one cannot afford delays. The current internet does not have the necessary bandwidths or reliable service except for audio (where the well-known Real Audio commercial product can deliver AM quality audio over modern phone lines. Of course, this audio and video delivery depends on compression technology, such as fractal, MPEG, and Wavelet whose details are closely guarded corporate secrets.

New networking approaches, such as ISDN, ATM, and cable modems are very important for digital video as they offer high enough performance to support reasonable quality. However, the time scale and nature of their deployment is still unclear. One uncertainty is just how important to consumers, digital video is!

4 Five Typical Web Windows Application Areas

4.1 IntraNets and Business Enterprise Information Systems

Intranets have been popular recently. These correspond to building enterprise information systems using WebWindows technology. WebWindows can encompass the world (clients talking to large servers around the globe) or just a single machine (a PC Web browser linked to a server on the same machine). Further, we can apply these ideas flexibly to any enterprise of intermediate size and use Web technology to build general information services. These will use traditional concepts (e.g., databases) but link these to the rich WebTop environment. Note that this means one should not reject Web technologies because of today's low bandwidth or insecure Internet. If we build, today, a WebWindows Intranet, it can have whatever bandwidth and security the enterprise network supports. One of the applications explored by NPAC is a politics information system built around Newt Gingrich's recent visit to the University. This illustrates how WebWindows systems can be built for information storage and on-line discussion of any loosely knit enterprise from political parties through special interest organizations.

Figure 27:

4.2 WebMed: Web Technologies in Health Care

The Intranet ideas above can, of course, be applied to medical enterprises where one can use Web-linked databases combined with several different user interfaces. Here, we use Java so that we can customize the 'look and feel' and access privileges so that doctors, hospital administrators, insurance companies, nurses, and patients can share the same information. NPAC has two interesting local projects—one with the University College of Nursing with a prototype system to help K–12 school nurses; the second uses simple virtual reality technology from David Warner [http://www.ponfar.com/i3] to allow disabled persons to both navigate the Web and to transmit medical sensor data from their homes to Web servers where it can then be analyzed conveniently by doctors.

4.3 Web Technology in the Financial Industry

An area of interest to NPAC is use of WebWindows in the financial industry. Here, Java can be used to provide excellent PC (home, office) interfaces to real time trading data. Java will allow downloading of simple modeling packages to aid on-line trading. However, today's complex financial instruments require computationally intensive Monte Carlo algorithms for accurate simulation. Here, we use Web technology to allow investors access to stock data, and simulation systems running on the classic MPPs developed in the HPCC program.

4.4 Background Issues

4.4.1 Infrastructure issues

Bandwidth The basic principle of multiple-use says that for distance education, we will use the hardware infrastructure, software, and other technology produced for the suite of NII applications including business, health care, and entertainment.

End-user access to the Web can be based on 28.8 kilobits/second or similar speed modems using POTS (Plain Old Telephone Service) connections, but in practical terms, limits us to displaying audio, text and graphics. Applications involving quality video delivery require ISDN access (128 kilobits/second) or ATM service (scales to Gigabit speeds, individual user typically requires a maximum of a few megabits/second).

Full screen NTSC video display requires about 1.5 megabits/second of bandwidth assuming MPEG1 compression. By comparison, a quad speed CDROM provides transfer rates of 0.6 megabits/second. We note that ISDN provides 10but for educational applications we need perhaps only quarter screen displays with 15 frames per second (typical educational clip is one to two minutes). Improved compression schemes such as wavelets are becoming available but are not standardized and built into PCs as is MPEG.

The bandwidth required for digital video transmission (1.5 megabits/second) also supports the following applications relevant to distance education:

- Collaboration with real-time video of participants, supporting student-teacher and student-student and other individual interactions such as networked world-wide collaborations of students doing real science including remote control of instruments (telescopes, microscopes) and fetching data over web [***ref 13,14 of Ed***];

- Simulations based on video animation produced either in batch mode or in real-time by a computer, including display of physical phenomena with tagged annotations linking beginning students with the insights of experts;
- Navigation in Geographical Information Systems using transmission of rendered scenes of 3-dimensional terrain (VRML based vector representation of video is more efficient than transmitting pixels), used for homework and student projects including electronic laboratories and field trips;
- Text-indexed video clips selected from a large (100,000 hour or 70 terabyte) digital archive, supporting pedagogical (electronic books and similar material, also InfoVISION servers) and exploratory learning using the World Wide Web as a "virtual library" for projects.

Role of caching Caching content for use in distance education applications has important performance and quality of content implications. Education is typically carried out in a classroom setting, therefore a group of students browsing related content will naturally provide efficient use of cached resources. In a school experiment conducted by NASA Langley, 95% of the material accessed by students was found in cache, allowing a 28.8 kilobaud connection purchased by a school to look like a dedicated T1 line [*** ref 1 of Ed***].

Caching is also important to content quality. Pre-selecting material for later use in class can be used to ensure students browse relevant and appropriate material, and maintain efficiency in teacher preparation. Many schools limit Internet access due to the perception of easily accessible inappropriate material. Datamining the World Wide Web for resources relevant to K–12 classroom adds great value to Web content. KidsWeb [*** ref 15 of Ed***] was initially developed to support an educational program at NPAC for eighth graders in the Syracuse Area [*** ref 16 of Ed***], and is NPAC's most popular Internet access point.

Caching is also needed to support a hierarchical Web server/InfoVISION server scenario where content in high demand is replicated and stored close to end-user populations for bandwidth efficiency, and migrated to centralized archives for storage efficiency when demand lowers. For example, a popular movie released on a future Hollywood server would be watched throughout the U.S.A. with an average distance of 1,500 miles between server and client. If the movie is automatically (caching) or by hand (replication) downloaded to local servers in each town, one would reduce client server distance by a factor of at least 100. This will reduce the needed NII (Internet) bandwidth by a corresponding factor.

In a distance education experiment in New York State, we intend to distribute a CD-ROM to each student. Lessons will be delivered synchronously (all students are on the same page) by transmitting the same relative URL to each client. This approach uses modern (Java) collaboration technology such as Habanero [*** ref 13 of Ed***] or Tango [*** ref 14 of Ed***] to fetch data from the local CD-ROM, to display the same page on each client, for example a 100 kilobyte image. This effectively saves a factor of 1,000 in needed bandwidth and allows simple 28.8 kilobaud lines to support a graphically rich curricula.

4.4.2 Human interaction issues

Collaboration Collaboration is the aspect of the virtual K–12 school or university which is most sensitive to teacher-pupil separation. Distance education implies a distribution of students, an obvious obstacle to learning. However, Web technologies also allow school children to interact with their peers internationally. Distance education holds the promise of individualized learning on the users timeframe, but two-way communication between student and teacher must be supported. Collaboration between the learner and leading experts is technically possible through a hierarchical structure linking learners and a mix of distance university and leading expert partnerships. Learning remains a group activity even if conducted electronically, therefore multi-user collaboratory tools are needed to support group navigation, sharing of group experiences, and group authoring of content.

Collaboration is a particularly demanding distance education application because it cannot use caching and a hierarchical server structure, approaches which serve information access and delivery applications so well. Collaboration is sensitive to available bandwidth and requires the best compression and transport technology. As illustrated, a hybrid scheme involving a mix of locally cached (CD-ROM) material with high value interactive access at a distance is quite promising.

Our project experience with videoconferencing suggests that 56 kilobaud is the lowest practical speed, ISDN (128 kbaud) is adequate, but full screen video at 30 frames per second requires the standard 1.5 megabits/sec bandwidth with MPEG1 compression. We are experimenting with the use of audio (which is practical at 28.8 kilobaud) as a form of annotation to enhance on-line course material. This can be supplemented by images or new video codecs such as Java H263 applets on low speed lines.

Authoring The ubiquity of the Web suggests that we will soon see Web operating systems replacing proprietary operating systems, and the growth of a Webtop (rather than proprietary desktop) personal productivity environment. Content authoring in a Web environment will become a central function in distance education.

Homework will be done in an electronic environment as part of the learning experience. Web-based content must therefore be interactive, and editable. Students must be ably to create, modify, and edit on-line hypermedia material as part of the learning process. As part of our computer science classes we have developed a virtual programming laboratory where students can develop, compile, and run (homework) problems [*** ref 17 of Ed***]

Active simulation Simulation support is needed to provide two-way interactivity that enables real-time interaction with a simulated world. Students must be able to interact with active simulations by guiding the process (e.g., setting parameters, triggering events), navigating regions of interest, and accessing explanatory annotations (written either by teachers, application experts, or other students).

Virtual reality could be a revolutionary technology to improve the learning environment by immersing student in physical or more abstract information spaces.

Navigation The Web is an excellent source of educational content (see Kids Web [*** ref 15 in Ed***], and ThinkQuest [*** ref 18 in Ed***]) but the growing volume and complexity makes it too easy for students to get lost in cyberspace. Disciplined navigational tools are need, perhaps incorporating multi-sensory feedback [19] to create a coherent integration model of various multimedia modalities (hypertext, sound, video, simulations).

Structured navigation and free exploration must be balanced to allow the teacher to efficiently guide the investigation and the student to employ his or her own optimal model of learning. Weve experimented in existing NPAC projects with unstructured exploration (e.g., KidsWeb, a video archive) and structured courses and tutorials (e.g., on-line computational science courses [***ref 20,2 in Ed***].

Assessing the impact of new technologies and approaches on learning is itself an application where Web technologies are useful. Tools are needed to track students navigation pathways and their reactions to selected sequences, then we must use this information for evaluation and individual-student course refinement. It is straightforward to log the URLs of the pages accessed by teachers and students. We are using this today to link recorded audio of lectures to particular pages being discussed. Related research on navigational patterns termed cognitive throughput by disabled learners is designed to anticipate the response of users in an information environment. Tracking navigational pathways and inter-target intervals allows us to anticipate future navigational direction (to pre-download content and improve performance), and support cognitive studies of user needs for improved information design [*** ref 19 in Ed***].

The unit of information In traditional models, we teach and learn using lectures supported by books, notes, and presentations in the form of bulleted items and summary lists in a linear sequence. In the Web-based environment we envision, where a distributed hierarchy of information-on-demand servers are available, and students interactively explore active simulations, we must consider what is an appropriate unit of information.

The hyperlinked environment of the Web makes "sequential" books unnaturalpeople are better than computers at flipping sequentially through electronic pages. This may suggest an electronic encyclopedia model as the optimal model for the Web. This model would likely contain short, linked, modular entries produced by multiple authors having several different levels of description for a given subject.

The natural size of an information nugget is roughly a computer screenfulit is interesting to note that this is similar to the amount of data displayed on a viewgraph on an overhead projector. Thus we can expect presentation technology and techniques to be an important source of ideas for Web education.

4.4.3 Some Base Technologies

Text-indexed video NPACs text-indexed video project combines basic research in video-on-demand technology, software design and implementation, and systems to provide an operational VoD service for educational applications. All development in this project is Web-based.

In contrast to entertainment oriented video dialtone systems where the content stream is nearly 100integrates video server technology within a larger information system. A video server management layer is linked to full text searchable databases and metadata used to access video content. This approach allows us to deliver digital video news-on-demand, educational documentary content, and multimedia coverage of events with educational relevance (e.g., a national political campaign event [*** ref 21 in Ed***]).

Examples of distance education applications of our VoD system include:

- Searching a database on topics such as siege of Stalingrad, arrest of Mahatma Ghandi, or threat to the Apollo 13 moon mission to retrieve the original, unedited film clips (in digital video format) from a Reuters historical video archive;
- Searching a documentary transcript database on technology in the classroom to retrieve segments of The Discovery Channel Assignment Discovery series on video graphics technology;

- Searching a transcript database on streamlining government bureaucracies to retrieve a segment of video from House Speakers Newt Gingrichs presentation and dialog with political science graduate students during a visit to Syracuse University.

NPAC has focused on developing operational video server systems for education and related applications (e.g., decision support), while examining tradeoff s between image quality and bandwidth requirements. Transporting variable bit rate video streams (naturally occurring in video content) over a constant bit rate network channel (required in digital network delivery) requires a trade-off between image quality and bandwidth. Video compression techniques attempt to minimize the information content transmitted in each successive frames by carrying only the difference in scenes forward. The data requirements of displaying an action scene such as a sports event are much more demanding than for example a talk show [*** ref 4 of Ed***].

4.4.4 WebWisdom

Web Wisdom is a hierarchical tree of knowledge built from JavaScript and aimed at Netscape 3.0 Navigators [*** ref 5 of Ed***]. Web Wisdom is set up to display worlds of self defining Web Pages where the header of documents have appropriate JavaScript calls to define recursive, hierarchical structures.

This structure has natural applicability in education and in mail archives. Web Wisdom can handle several different worlds (limited only by the strength of JavaScript). We have currently implemented "Foilworld" which archives foils produced by NPAC, Family PhotoWorld and Administration.

FoilWorld is a part of WebWisdom devoted to accessing basic materials sorted by different criteria such as:

- Foilsets arranged in a simple alphabetic sort. This is for users who know the Foilset collection and want access to a particular set.
- Foilsets arranged by topics. This is for users who wish to browse topic areas. The user first accesses a list of areas, then an index to resources in each topic area.
- Following is a small sample of Foilset topic areas concerning Web technologies and education:
  - Education: Computational Science Information Track Base Technology Course Spring 1996;
  - Education: Computational Science Information Track Systems and Applications Spring 96;
  - Undergraduate Web Technology Class and Laboratory;
  - Education: Tutorial on Parallel Processing;
  - Education: Tutorial on Web Technologies for distance education at Supercomputing '95;
  - Education: Undergraduate level using Web Technologies
  - Education: Use of Web Technologies including K–12.

- A collection of tutorials and other offerings in the Web technology education area Here is a sample list of tutorials:
- Computational Science Courses 615/616 Fall/Spring Semester, 1996 (graduate courses at Syracuse University) [*** ref 20 in Ed***];
- One day Tutorial at Supercomputing December '95, San Diego [*** ref 2 in Ed***];
- One day Tutorial at High Performance Distributed Computing Workshop, August 1995 in Washington DC [*** ref 1 in Ed***].

FoilWorld has full support for Persuasion, Powerpoint, collections of screendumps, and general HTML pages. Family PhotoWorld uses the same technology to manage a collection of family photos. The digital images may come from scanning original prints, disks now available from commercial photofinishers, and images directly from digital cameras. The Administration area of WebWisdom allows users to select useful parameters, log URLs, and read and writing parameter sets.

The page syntax used in WebWisdom uses JavaScript to store metadata on each page which defines associated pages (children in the hierarchical structure). The page syntax also accommodates quick HTML versions of graphically rich pages, notes, associated illustrative (Java) programs, and multimedia material. We also have linked WebWisdom with an Oracle server which allows the user to search all pages at a given level or below a particular node.

A key facility of WebWisdom is a Java front-end which can create play-lists which produce new composite foilsets (the set of pages at any level of the WebWisdom world) with a mix of pages from different foilsets.

4.5 Manufacturing

HPCC (High-Performance Computing and Communication) and the NII are used today, and can be expected to play a growing role in manufacturing, and more generally, engineering. For instance, the popular concept of agile manufacturing supposes the model where virtual corporations generate "products-on-demand." The NII is used to link collaborating organizations. Powerful computers are needed to support instant design (or more accurately redesign or customization) and sophisticated visualization and virtual reality "test drives" for the customer. At the corporate infrastructure level, concurrent engineering involves integration of the different component disciplines—such as design, manufacturing, and product life cycle support—involved in engineering. These general ideas are tested severely when they are applied to the design and manufacturing of complex systems such as automobiles, aircraft, and space vehicles such as shuttles. Both the complexity of these products, and in some sense the maturity of their design, places special constraints and challenges on HPCC and the NII.

High-performance computing is important in all aspects of the design of a new aircraft. However, it is worth noting that less than 5% of the initial costs of the Boeing 777 aircraft were incurred in computational fluid dynamics (CFD) airflow simulations—the "classic" Grand Challenge in this field. On the other hand, over 50% of these sunk costs could be attributed to overall systems issues. Thus, it is useful but not sufficient to study parallel computing for large scale CFD. This is "Amdahl's law for practical HPCC." If only 5% of a problem is parallelized, one can at best speed up and impact one's goals—affordability, time to market—by this small amount. HPCC, thus, must be fully integrated into the entire engineering enterprise to be effective. Very roughly, we can view the ratios of 5% to 50% as a measure of ratio of 1:10 of the relevance of parallel (classic MPPs) and distributed computing (i.e., the NII or Web!) in this case.

The maturity of the field is illustrated by the design criterion used today. In the past, much effort has been spent on improving performance—more speed, range, altitude, size. These are still critical under extreme conditions, but basically these just form a given design framework that suffices to buy you a place at the table (on the short-list). Rather, the key design criteria is competitiveness, including time to market, and total affordability. Although the design phase is not itself a major cost item, decisions made at this stage lock in most of the full life cycle cost of an aircraft with perhaps 80% of total cost split roughly equally between maintenance and manufacturing. Thus, it certainly would be important to apply HPCC/NII at the design phase to both shorten the design cycle (time to market) and lower the later ongoing costs of manufacturing and maintenance.

We take as an example the design of a future military aircraft—perhaps 10 years from now. This analysis is taken from a set of NASA sponsored activities centered on a study of ASOP—Affordable Systems Optimization Process. This involved an industrial team, including Rockwell International, Northrop Grumman, McDonnell Douglas, General Electric, and General Motors. ASOP is one of several possible approaches to multidisciplinary analysis and design (MAD) and the results of the study should be generally valid to these other MAD systems. The hypothetical aircraft design and construction project could involve six major companies and 20,000 smaller subcontractors. This impressive virtual corporation would be very geographically dispersed on both a national and probably international scale. This project could involve some 50 engineers at the first conceptual design phase. The later preliminary and detailed design stages could involve 200 and 2,000 engineers, respectively. The design would be fully electronic and demand major computing, information systems, and networking resources. For instance, some 10,000 separate programs would be involved in the design. These would range from a parallel CFD airflow simulation around the plane to an expert system to plan location of an inspection port to optimize maintainability. There is a corresponding wide range of computing platforms from PCs to MPPs and a range of languages from spreadsheets to High-Performance Fortran. The integrated multidisciplinary optimization does not involve blindly linking all these programs together, but rather a large number of suboptimizations involving at one time a small cluster of these base programs. Here we see clearly, an essential role of the NII to implement these set of optimizations, which could well need linking of geographically separated compute and information systems. An aircraft is, of course, a very precise system, which must work essentially flawlessly. This requirement implies a very strict coordination and control of the many different components of the aircraft design. Typically, there will be a master systems database to which all activities are synchronized at regular intervals–perhaps every month. The clustered suboptimizations represent a set of limited excursions from this base design that are managed in a loosely synchronous fashion on a monthly basis. The configuration management and database system are both critical and represent a major difference between manufacturing and command and control, where in the latter case, real time "as good as you can do" response, is more important than a set of precisely controlled activities. These issues are characteristic of the high-performance distributed computing (HPDC) where, although loosely coupled, the computers on our global network are linked to "solve a single problem."

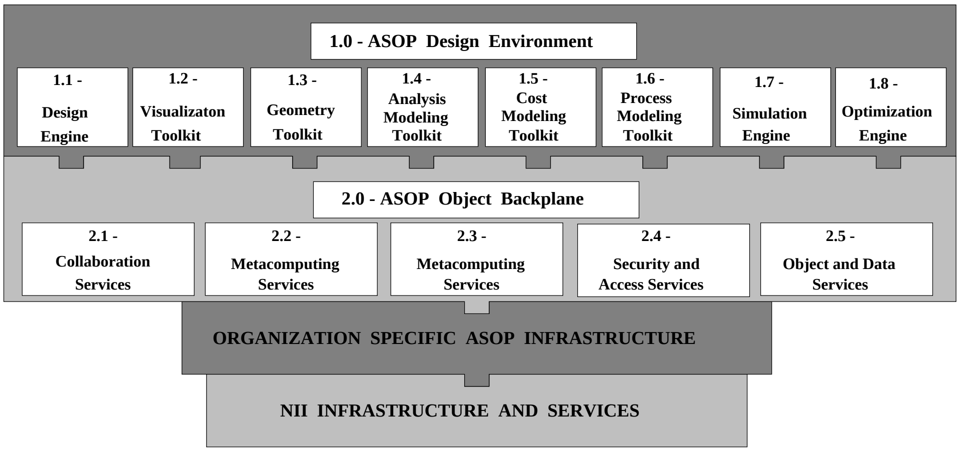

Figure 28: Affordable Systems Optimization Process (ASOP) Implemented on the NII for Aeronautics Systems

ASOP is designed as a software backplane (the NII or Web) linking eight major services or modules shown in Figure 28. These are design (process controller) engine, visualization, optimization engine, simulation engine, process (manufacturing, productibility, supportability) modeling toolkit, costing toolkit, analytic modeling toolkit, and geometry toolkit. These are linked to a set of databases defining both the product and also the component properties. Parallel computing is important in many of the base services, but HPDC is seen in the full system.

4.6 Command and Control or Crisis Management on the NII

Command Control (sometimes adding in Computing, Communications, Intelligence Surveillance, and Battle Management with abbreviations lumped together as BMC^{4}IS) is the task of managing and planning a military operation. It is very similar to the civilian area of Crisis management, where the operations involve combating effects of hurricanes, earthquakes, chemical spills, forest fires, etc. [***ref NRC report***] Telemedicine was once considered a video conferencing application, but now this is becoming like command control in the interventional informatics concept [Warner:95a] where the doctor (cf. commander) needs a rich interactive information to win the "battle" i.e., treating the patient. Both the military and civilian cases have computational "nuggets" where parallel computing is relevant. These include processing sensor data (signal and image processing) and simulations of such things as expected weather patterns and chemical plumes. One also needs large-scale multimedia databases with HPDC issues related to those described for InfoVISiON. We have built a powerful Web based command and control prototype on top of TANGO [***ref SIAM Tango and Tango-long***].

The NII is needed to link military planners and decision makers, crisis managers, experts at so-called anchor desks, workers (warriors) in the field, information sources such as cable news feeds, and large-scale database and simulation engines.

A key characteristic of the required NII support is adaptivity. Crises and battles can occur anywhere and destroy an arbitrary fraction of the existing infrastructure. Adaptivity means making the best use of the remaining links, but also deploying and integrating well mobile enhancements. The information infrastructure must exhibit security and reliability or at least excellent fault tolerance (adaptivity). Network management must deal with the unexpected capacity demands and real time constraints. Priority schemes must allow when needed critical information (such as chemical plume monitoring and military sensor data) precedence over less time critical information, such as background network video footage.

Needed computing resources will vary from portable handheld systems to large backend MPPs. As there will be unpredictable battery (power) and bandwidth constraints, it is important that uniform user interfaces and similar services be available on all platforms with, of course, the fidelity and quality of a service reflecting the intrinsic power of a given computer. As with the communications infrastructure, we must cope with unexpected capacity demands. As long as the NII is deployed nationally, computational capacity can be exploited in remote sites. The Department of Defense envisages using the basic NII (GII) infrastructure for command and control, augmented by "theater extensions" to bring needed communications into critical areas. The "take it as it is" characteristic of command and control requires that operating systems and programming models support a general adaptive mix (metacomputer) of coordinated geographically distributed but networked computers. This network will adaptively link available people (using perhaps personal digital assistants) to large-scale computation on MPPs and other platforms. There are large computational requirements when forecasting in real-time physical phenomena, such as the weather effects on a projected military action, forest fires, hurricanes, and the structure of damaged buildings. On a longer time scale, simulation can be used for contingency planning and capability assessment. Training with simulated virtual worlds supporting televirtuality, requires major computational resources. In the information arena, applications include datamining to detect anomalous entries (outliers) in large federated multimedia databases. Data fusion including sensor input and processing, geographical information systems (with perhaps three-dimensional terrain rendering), and stereo reconstruction from multiple video streams are examples of compute intensive image processing forming part of the needed HPDC environment.

A critical need for information management involves the best possible high-level extraction of knowledge from databanks—the crisis manager must make judgments in unexpected urgent situations—we cannot carefully tailor and massage data ahead of time. Rather, we need to search a disparate set of multimedia databases. As well as knowledge extraction from particular information sources, the systematic use of metadata allowing fast coarse grain searching is very important. This is a specific example of the importance of standards in expediting access to "unexpected" databases. One requires access to databases specific to crisis region or battlefield, and widespread availability of such geographic and community information in electronic form is essential. There are very difficult policy and security issues, for many of these databases need to be made instantly available in a hassle-free fashion to the military commander or crisis manager—this could run counter to proprietary and security classification constraints. The information system should allow both network news and warriors in the field to deposit in near real-time, digital versions of their crisis and battlefield videos and images.

As mentioned, we expect that human and computer expertise to be available in "anchor desks" to support instant decisions in the heat of the battle. These have been used in a set of military exercises called JWID (Joint Warrior Interoperability Demonstrations). We note that this information scenario is a real-time version of the InfoVISiON service needed to support the society of the Information Age.

Command and Control has historically used HPDC as the relevant computer and communication resources, are naturally distributed, and not centralized into a single MPP. We see this HPDC model growing into the standard information support environment for all the nation's enterprises, including business, education, and society, as indicated in InfoVISiON in Table 2.

5 Glossary of Concepts and Acronyms

Applets An application interface where referencing (perhaps by a mouse click) a remote application as a hyperlink to a *server* causes it to be downloaded and run on the *client*.

ASOP (Affordable Systems Optimization Process) refers to a process using multidisciplinary optimization to produce more affordable systems

CAVE A room with stereo images on the walls designed to create a televirtual environment.

CFD (Computational Fluid Dynamics) refers to computational solutions of differential equations, such as the Navier Stokes set, describing fluid motion.

COTS (Customer Off The Shelf)—An important concept in Defense Systems, labelling the use of commercial software snf system components rather than specialized limits which are expensive and hard to upgrade, even if initially of higher capability.

Data Fusion A common *command and control* approach where the disparate sources of information available to a military or civilian commander or planner, are integrated (or fused) together. Often, a *GIS* is used as the underlying environment.

Distributed Computing The use of networked heterogeneous computers to solve a single problem. The nodes (individual computers) are typically loosely coupled.

Geographical Information System (GIS) A user interface where information is displayed at locations on a digital map. Typically, this involves several possible overlays with different types of information. Functions, such as image processing and planning (such as shortest path) can be invoked.

Global Information Infrastructure (GII) The GII is the natural world-wide extension of the *NII* with comparable exciting vision and uncertain vague definition.

High-Performance Computing & Communications (HPCC) Refers generically to the federal initiatives, and associated projects and technologies that encompass *parallel computing*, *HPDC*, and the *NII*.

High-Performance Distributed Computing (HPDC) The use of distributed networked computers to achieve high performance on a single problem, i.e., the computers are coordinated and synchronized to achieve a common goal.

Hypertext Markup Language (HTML) A syntax for describing documents to be displayed on the *World Wide Web*.

Hypertext Transport Protocol (HTTP) The *protocol* used in the communication between *Web Servers* and *clients*.

InfoVISiON Information, Video, Imagery, and Simulation ON demand is scenario where *multimedia servers* deliver multimedia information to clients on demand—at the click of the user's mouse.

Integrated Service Data Network (ISDN) A digital multimedia service standard with a performance of typically 128 kilobits/sec, but with possibility of higher performance. ISDN can be implemented using existing telephone (*POTS*) wiring, but does not have the necessary performance of 1–20 megabits/second needed for full screen TV display at either VHS or high definition TV (HDTV) resolution. Digital video can be usefully sent with ISDN by using quarter screen resolution and/or lower (than 30 per second) frame rate.

Internet A complex set of interlinked national and global networks using the IP messaging protocol, and transferring data, electronic mail, and *World Wide Web*. In 1995, some 20 million people could access Internet—typically by *POTS*. The Internet has some high-speed links, but the majority of transmissions achieve (1995) bandwidths of at best 100 kilobytes/sec. the Internet could be used as the network to support a *metacomputer*, but the limited *bandwidth* indicates that *HPDC* could only be achieved for *embarrassingly parallel* problems.

Java A distributed computing language (*Web Technology*) developed by Sun, which has similarities with C++ but supports *Applets*.

MAD (Multidisciplinary Analysis and Design, or Multidisplinary Optimization) refers to the coupling of several areas, such as structural dynamics and fluid flow in a combine tradeoff to produce higher capability vehicles.

Massively Parallel Processing (MPP) The strict definition of MPP is a machine with many interconnected processors, where 'many' is dependent on the state of the art. Currently, the majority of high-end machines have fewer than 256 processors. A more practical definition of an MPP is a machine whose architecture is capable of having arbitrarily many processors—that is, it is scalable. In particular, machines with a distributed memory design (in comparison with shared memory designs) are usually synonymous with MPPs since they are not limited to a certain number of processors. In this sense, "many" is a number larger than the current largest number of processors in a shared-memory machine.

Metacomputer This term describes a collection of heterogeneous computers networked by a high-speed wide area network. Such an environment would recognize the strengths of each machine in the Metacomputer, and use it accordingly to efficiently solve so-called *Metaproblems*. The *World Wide Web* has the potential to be a physical realization of a Metacomputer.

Metaproblem This term describes a class of problem which is outside the scope of a single computer architectures, but is instead best run on a Metacomputer with many disparate designs. These problems consist of many constituent subproblems. An example is the design and manufacture of a modern aircraft, which presents problems in geometry grid generation, fluid flow, acoustics, structural analysis, operational research, visualization, and database management. The Metacomputer for such a Metaproblem would be networked workstations, array processors, vector supercomputers, massively parallel processors, and visualization engines.

Multimedia Server or Client Multimedia refers to information (digital data) with different modalities, including text, images, video, and computer generated simulation. Servers dispense this data, and clients receive it. Some form of browsing, or searching, establishes which data is to be transferred. See also *InfoVISiON*.

Multipurpose Internet Mail Extension (MIME) The format used in sending multimedia messages between *Web Clients* and *Servers* that is borrowed from that defined for electronic mail.

National Information Infrastructure (NII) The collection of *ATM*, cable, *ISDN*, *POTS*, satellite, and wireless networks connecting the collection of 10^8–10^9 computers that will be deployed across the U.S.A. as set-top boxes, PCs, workstations, and MPPs in the future.

The NII can be viewed as just the network infrastructure or the full collection of networks, computers, and overlayed software services. The *Internet* and *World Wide Web* are a prototype of the NII.

PERL An interpreted language particularly targeted at text manipulations and systems programming.

POTS The conventional twisted pair based Plain Old Telephone Service.

Televirtual The ultimate computer illusion where the user is fully integrated into a simulated environment and so can interact naturally with fellow users distributed around the globe.

Virtual Reality Modeling Language (VRML) A "three-dimensional" HTML that can be used to give a universal description of three-dimensional objects that supports *hyperlinks* to additional information.

WebTop refers to the implementation of a set of standard desk top and personal computer tools, which are essential in any computing environment.

Web Clients and Servers A distributed set of clients (requesters and receivers of services) and servers (receiving and satisfying requests from clients) using *Web Technologies*.

WebWindows The operating environment created on the World Wide Web to manage a distributed set of networked computers. WebWindows is built from *Web clients* and *Web servers*.

WebWork [Fox:95a] An environment proposed by Boston University, Cooperating Systems Corporation, and Syracuse University, which integrates computing and information services to support a rich distributed programming environment.

World Wide Web and Web Technologies A very important software model for accessing information on the Internet based on hyperlinks supported by *Web technologies*, such as *HTTP*, *HTML*, *MIME*, *Java*, *Applets*, and *VRML*.

References

[Andreessen:93a] Andreessen, M., "NCSA Mosaic: technical summary," May 1993. Technical report, NCSA.

[Fox:95c] Fox, G. C. "Software and hardware requirements for some applications of parallel computing to industrial problems." Technical Report SCCS-717, Syracuse University, NPAC, Syracuse, NY, June 1995. Submitted to ICASE for publication (6/22/95); revised SCCS-134c.

[Fox:95a] Fox, G. C., Furmanski, W., Chen, M., Rebbi, C., and Cowie, J. H. "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 (unpublished).

[Warner:95a] Warner, D., Sale, J., and Viirre, E. "A systems approach for developing an integrative healthcare information distribution infrastructure." Technical Report SCCS-749, Syracuse University, NPAC, Syracuse, NY, November 1995.