NCSA

Where Commercial Services and the Web are Headed

WWW Conference Fall '94
Wednesday, October 19, 10:15 AM, Lincoln Room

Chair: Henry Houh, MIT Laboratory for Computer Science

The World–Wide Web has spurred a considerable amount of interest in using the Internet for commercial purposes. The current WWW commercial model, a "pull" model rather than a "push" one, suits the culture of the Internet well.

In the short amount of time that commercial Web sites have been enjoying popularity, they have evolved a great deal, where now many commercial sites are providing value–added services beyond providing static marketing information. Such value–added services now seem to be important to attracting regular visitors.

What is the current status of the commercial uses of and Web sites on the Internet, and where will commercial activities be heading? This panel of experts will discuss their efforts and visions.

Invited Panelists:

Jeffrey Dearth – New Republic Magazine/The Electronic Newsstand

Jeffrey Dearth has been President of The New Republic Inc. since May 1987 and held the title of Publisher from 1984 to 1988. He began his career in publishing with Time Incorporated in New York and spent six years with Time magazine on various marketing assignments in London, Amsterdam and Paris before returning to the U.S. in 1981 as Marketing Director of Smithsonian magazine.

Mr. Dearth is the founder of The Electronic Newsstand Inc. and an active member of the Magazine Publishers of America's Small Magazine Advisory and New Media Committees. He was recently nominated to the Board of Directors of the MPA. He has been described as one of the "commercial pioneers" of the Internet and has been a leader in promoting the Internet s capabilities to the publishing and business communities. A native of the Detroit area, Mr. Dearth is a graduate of The University of North Carolina where he was a Morehead Scholar and graduated Phi Beta Kappa.

The Electronic Newsstand

The Electronic Newsstand was founded in July 1993 to provide the Internet community with easy access to a wide range of interesting information furnished by the world's leading publishers.

Like traditional newsstands, The Electronic Newsstand is a place where you can browse –– for free –– through many publications and have your interest stimulated by a variety of subjects. The Newsstand provides a window on the world of computers, technology, science, business, foreign affairs, the arts, travel, medicine, nutrition, sports, politics, literature and many, many other areas of interest.

Every Newsstand publisher provides the table of contents and several articles from each current issue. The Newsstand, which archives previously featured material, is also searchable by keyword. Access to The Electronic Newsstand is free and available via gopher or telnet.

Russ Jones – Digital Equipment Corporation

Russ Jones is the Internet Program Manager for Digital Equipment Corporation and is part of the newly formed Internet Business Group. Russ has been with Digital for 9 years. He recently co–authored a book entitled, "Managing Internet Information Services."

Digital Equipment Corporation

Digital Equipment is the leading worldwide supplier of networked computer systems, software, and services. Digital pioneered and leads the industry in interactive, distributed and multivendor computing. An international company, Digital does more than half its business outside the United States, developing and manufacturing products and providing customer services in the Americas, Europe, Asia, and the Pacific Rim.

Digital offers a full range of desktop, client/server, production, and mainframe systems for multivendor computing environments. Applications include transaction processing, data management, telecommunications, finance, realtime data acquisition and control, vector processing, education, publishing, manufacturing, software development, and health care.

Jay M. Tenenbaum – Enterprise Integration Technologies/CommerceNet

Jay M. Tenenbaum received the B.S. and M.S. degrees in electrical engineering from M.I.T. in 1964 and 1966, respectively, and the Ph.D degree in electrical engineering and computer science from Stanford University in 1970. From 1972 to 1980, he led the research program in computational vision at SRI's Artificial Intelligence Center. In 1980, he co–founded the Fairchild Laboratory for Artificial Intelligence Research (FLAIR), a forerunner of Schlumberger Palo Alto Research, serving as its director form 1983 to 1986. In 1986, Dr. Tenenbaum returned to active research, as a Schlumberger Fellow and Professor of Computer Science (Consulting) at Stanford University.

From 1988 to 1990, he also served as Director of Advanced Research Projects for Schlumberger Technologies. At Stanford, his research has focused on applications of AI in design and manufacturing. In January 1990, Dr. Tenenbaum left Schlumberger to found Enterprise Integration Technologies Corporation, an R&D and consulting organization specializing in information technology for electronic commerce, collaborative engineering and agile manufacturing. EIT is a leading developer of software and services that help companies do business on the Internet.

Dr. Tenenbaum is a Fellow and former board member of the American Association for Artificial Intelligence (AAAI). He has served on the US Air Force Scientific Advisory Board and numerous other government and professional committees. He currently serves on the editorial boards of 4 technical journals dealing with applications of information technology in design and manufacturing. Dr. Tenenbaum is author or co–author of over 50 technical papers in leading journals and conferences, and holds 2 patents.

CommerceNet

CommerceNet is a consortium of companies and organizations whose goal is to create an electronic marketplace where companies transact business spontaneously over the Internet. CommerceNet will stimulate the growth of a communications infrastructure that will be easy–to–use, oriented for commercial use, and ready to expand rapidly. The net results for businesses in this region will be lower operating costs and a faster dissemination of technological advancements and their practical applications.

The CommerceNet marketplace will support all business services that normally depend on paper–based transactions. Buyers will browse multimedia catalogs, solicit bids, and place orders. Sellers will respond to bids, schedule production, and coordinate deliveries. A wide array of value–added information services will spring up to bring buyers and sellers together. These services will include specialized directories, broker and referral services, vendor certification and credit reporting, network notaries and repositories, and financial and transportation services.

Bill Washburn – MecklerWeb

Bill has a Masters degree and a Ph.D. degree from Stanford University in philosophy (ethics) and policy analysis respectively. He served as Registrar and Director of Institutional Research at the University of Santa Clara in the Silicon Valley. At Stanford he held two successive positions as Assistant Dean of Undergraduate Studies and as Assistant Provost for Academic Computing. He served as Director of the Office of Information Technology at Colorado State University.

Dr. Washburn was the first executive director of the Commercial Internet Exchange (CIX) Association. He is interested in the philosophies and policies that form the foundation of the commercial Internet provider industry. He left CIX and joined MecklerWeb as its chief operating officer at the end of September 1994. MecklerWeb is a large Internet–based corporate communication and marketing system enabling client companies to achieve a powerful presence with respect to the global markets currently emerging online.

During his two years leading CIX, he built it from a struggling start–up with 8 members to a highly visible international organization. He forged and maintained partnerships with nearly 90 new and continuing members of CIX including: SprintLink, Compuserve, MCI, Advantis, ANS CO+RE, NEARnet, NorthWestNet, SURAnet NetCom, PSI, BARRnet, and WestNet. International partners include: Hong Kong Supernet, Singapore Telecom, Fujitsu, NEC, AT&T Japan, Taiwan, Korea, Australia, South Africa, Costa Rica, Sovam Teleport, Nordic Carriers, Demon Systems, well as EUnet and PIPEX.

MecklerWeb

MecklerWeb Corporation is a wholly owned subsidiary of Mecklermedia Corporation based in Westport, CT. Mecklermedia is a leading provider of information on cutting–edge technologies and produces magazines, newsletters, database directories, trade shows, and professional books on virtual reality, CD–ROM, electronic books and the Internet. Among many other monographs and serials, the company publishes Internet World magazine.

Henry Houh – MIT Laboratory for Computer Science

Henry Houh, a doctoral student in the MIT Laboratory for Computer Science's Telemedia, Networks and Systems group, conducts research in distributed multimedia systems. Henry received his B.S. degree in electrical engineering, a B.S. degree in physics, and an M.S. degree in electrical engineering, in 1989, 1990, and 1991, respectively, all from MIT. After working a year at AT&T Bell Laboratories in the Optical Computing Group, he returned to MIT to join the TNS group. Henry is a member of IEEE, Sigma Xi, and a District Director for Tau Beta Pi. He is also president and founder of the Agora Technology Group, Inc., a company which specializes in using the Internet to provide communications solutions through products and consulting, and creator and former maintainer of the Internet's "Commercial Services on the Web" and "What's New in Commercial Services" lists.

Software Reuse Libraries with Mosaic

Jeffrey S. Poulin and Keith J. Werkman

Loral Federal Systems–Owego
MD 0210, Owego, NY 13827
Tel: (607) 751-6899/4889, fax: (607) 751-6025
Email: poulinj,keithw@lfs.loral.com

Abstract

This paper describes a Reusable Software Library (RSL) interface and search tool implemented using Mosaic. Mosaic provides a simple, easy-to-use method to find and extract reusable assets from a RSL, allows distributed access to assets from a variety of platforms, and can support most of the features of formal RSLs without any modifications. Through the use of HyperText Markup Language (HTML) forms, we implemented functions normally found in commercial-grade RSLs, such as component search, user registration, and problem reporting. Automatic generation of HTML pages and the use of command scripts further allowed us to provide different views of the RSL, such as search by subject. Finally, integrating the RSL with Wide Area Information Search (WAIS) provided a keyword search with minimal effort. Our Mosaic RSL cost less than 1% of the cost to develop a standard RSL and has quickly gained favor due to its intuitive interface and simple yet powerful information retrieval tools. [1]

Keywords: Software Reuse, Reusable Software Libraries (RSL), World Wide Web (WWW), Mosaic.

---

[1]Proceedings of the *2nd International World Wide Web Conference: Mosaic and the Web*, Chicago, Illinois, 17-20 October 1994.

1 Background

The quest for ways to improve the software development process has led many organizations to pursue the substantial benefits available through software reuse. To this end, these organizations have given a lot of attention to technologies that facilitate reuse; application generators, domain analysis techniques, formal methods, and application frameworks. Many organizations focus their reuse initiatives on a reuse library where members of the organization can both store reusable assets and retrieve assets when they need them. Traditional RSLs use specialized methods for component classification, search, and retrieval. Unfortunately, these formal tools and techniques require both a large investment to implement and substantial training to use. For these reasons, many organizations have seen little use of their RSLs even though they may contain a large number of quality assets.

This paper describes a software reuse library interface and search ability using Mosaic [3, 4]. We developed this interface for the Loral Federal Systems Group RSL, which we refer to as the *Federal Reuse Repository* (FRR). Mosaic provides an simple, easy-to-use method to search for and extract reusable assets from the FRR. With large organizations investing as much as 80 to 130 person-years to develop a formal RSL, the Mosaic interface cost less than 1% of the cost to develop and maintain a standard, commercial-quality RSL [2, 11]. Nonetheless, the Mosaic-based RSL has quickly gained favor due to its intuitive interface and powerful yet simple features.

The Mosaic interface allows us to take advantage of the existing AIX^R $(UNIX^R)$ file structure and to generate multiple views of the RSL, thereby allowing users to search for assets in several ways. First, we provide a hierarchical view based on the original source of the asset. Second, we provide an "asset by subject" view of the same information. Third, we have implemented a keyword search using WAIS. Finally, we have found that the tools and features of Mosaic such as forms and the Common Gateway Interface (CGI) allow us to implement a variety of common RSL functions such as registering users of RSL components, logging statistics of module usage, and automatically notifying users of the RSL (via e-mail) of component updates, problem reports, and other RSL-related information.

2 Reusable Software Libraries

The original Loral Federal Systems RSL provided a central repository for sharing, managing, and reusing software-related products across Loral Federal Systems sites. A copy of the RSL ran at each site and operated in cooperative fashion with the RSLs at the other sites. Together, the RSLs established a system of shared libraries; any organization could establish a library to service the needs of a department, project, business area, or higher organization.

The RSL ran on IBM's two major mainframe operating systems, Multiple Virtual Storage $(MVS)^R$ and Virtual Machine $(VM)^R$, because nearly every member of the company has access to these systems. Although the RSL interface complied with the IBM Systems Application Architecture $(SAA)^{TM}$ and Common User Access $(CUA)^{TM}$ interface standards, developers repeatedly expressed the desire for a workstation-style Graphical User Interface (GUI). This need increased in priority as development work moved almost exclusively to workstation platforms.

The search mechanism implemented in the RSL used a detailed classification scheme based on the work of Prieto-Diaz and Freeman [12, 13]. To locate a component for reuse, a user invoked the RSL tool and specified "facets," or pre-defined software attributes, and acceptable values for those facets. The tool executed a search of the RSL database for components classified with facet values equal to those specified by the user. If the user felt satisfied with any of the components located by the search, the user copied the component from the RSL onto the user's local disk space. A detailed discussion of the classification scheme and the issues surrounding this kind of classification appears in [10].

In theory, this extensive and formal mechanism provides detailed information upon which a user can search for and assess the usefulness of reusable components. However, quite often the quantity and formality of the information only serves to confuse the user. First, the up-front presentation of large amounts of classifiers makes it difficult to quickly extract the key bits of needed information. Second, having a formal classification scheme requires users to receive training in its use; untrained users will not effectively use the mechanisms so carefully provided to assist them [14]. Users of the RSL demanded a simple, easy to use keyword search that although might lack the precision of a formal mechanism, required little or no training to use.

3 Requirements for the RSL

Users of the existing RSL made the first two requirements for our replacement implementation explicit:

1. It had to have a friendly, GUI interface.

2. Its use had to come naturally and intuitively.

The remaining major requirements consisted of a set of constraints and desired features.

3.1 Constraints

Like many companies facing highly competitive markets, we sought the most efficient and cost-effective means to manage our information. The implied task involved porting the contents of the RSL tool from the mainframe to a workstation-based environment. Furthermore, we needed a way to not only locate reusable software but also a way to locate other items of interest such as key personnel, information about ongoing programs, the latest developments in various technologies, and trade studies previously conducted by our company. The desired information retrieval tool had to handle all these types of information.

Since many Loral Federal Systems customers require compliance with "Open Systems" platforms, much of our software development staff works in an AIX or UNIX environment. However, we also have a large population who work with and integrate local area networks based on the DOS or $OS/2^{TM}$ operating systems. Finally, we required compatibility with the host legacy systems, especially for support personnel. We sought an information retrieval tool accessible by users on all these platforms and one that did not restrict the user's access to information which happens to reside on any other platform.

Of course, we especially desired a system that would cost very little to build and maintain.

3.2 Desired Features

The detailed classification scheme described above brought with it the burdensome process of manually classifying every component a user entered into the RSL. To avoid this, we wanted the ability to automatically or semi-automatically index components in the RSL, even if this meant losing a detailed search ability for the RSL. However, users felt that keyword searches adequately met their search needs despite their poor precision relative to that available with facets. With all the power resting on their desktops, users felt they could afford several attempts at locating a suitable component.

Because the tool would provide a search mechanism for more than just code the tool needed to have the ability to launch other tools such as browsers, viewers, and programs that handle diverse kinds of information such as multi-media. From security reasons we also needed an access control ability to limit access to the RSL to authorized users. Finally, we considered several additional features that we wanted to have:

- *Problem reporting for reusable assets.* The ability for a user to send e-mail to the owner of a component or the librarian should the user require service on a component extracted from the RSL.

- *Registration of users of assets.* The ability for a library administrator to locate users of assets for the purpose of recording use of the RSL, to report possible problems with extracted components reported by other reusers, and to aid in configuration management of extracted components as newer versions become available.

- *Version control of assets.* The ability to manage several versions of the same assets.

Finally, everyone insisted that the tool have very good performance.

4 Possible Solutions

We reviewed a number of specialized reuse library tools against the above requirements and now discuss some of the results. Among the tools we considered, the $ReDiscovery^{TM}$ information retrieval tool from IBM allows users to create and search meta-databases of information about virtually any kind of database

or file system. However, ReDiscovery requires the user to create and maintain these meta-databases and it only runs on OS/2.

We briefly looked at two IBM internal use tools but dropped them for consideration when IBM sold their Federal Systems Division (to which we belong) to the Loral Corporation. We first considered an OS/2 LAN-based tool, called the Reuse Library System (RLS), which provides users with a easy-to-use GUI interface components stored on an OS/2 LAN. The second of these tools, called *XGuru,* automatically indexes any collection of text files using a highly intelligent information retrieval algorithm [7]. Created by IBM research and in use at several IBM locations, XGuru allows the user to ask natural language queries. By using the same algorithm used to index the information XGuru returns candidate files in a ranked order of preference.

The public domain contains a specialized tool called the STARS Reuse Library (SRL), one product of the United States Advanced Research Projects Agency (ARPA) Software Technology for Adaptable, Reliable Systems (STARS) program. The Asset Source for Software Engineering Technology (ASSET) program uses the SRL to manage its reuse library. The SRL has a friendly system of menus and built-in security, browsers, and librarian tools. However, it requires an underlying Oracle^{TM} database and has no GUI interface [1].

Finally, we considered the *InQuisiX^{TM}* reuse library from the Software Productivity Solutions Corporation. SPS sells the InQuisix tool to companies who want a full function library tool. It met our requirements for GUI interface, distributed data management, and integration of multiple media types, but only ran on a limited set of platforms. The greatest disadvantage of InQuisix came from its requiring a license for every user of the tool, thereby leading to a cost we did not want to incur [5].

Among the options for free, shareware, or public-domain tools, we considered numerous utilities that have recently become popular for resource discovery on the Internet [8]. As opposed to specialized library tools and commercial products each of these tools runs on a wide range of operating systems and platforms. The popularity of these products, as demonstrated by the their explosive growth and wide distribution, gave us confidence in their use despite the lack of official 'support.'

Anonymous FTP sites provided perhaps the simplest and cheapest overall solution. A common method to distribute software and one that requires no special tools, it lacks a user-friendly interface and requires some training to use in its basic form. FTP also has no search ability and works best only when the user knows exactly where and for what to look.

Archie and Gopher both provide client-server methods of traversing pre-established lists of information at distributed sites. Both use primarily character interfaces, even in their GUI versions. Although Gopher does not have a search ability, Archie searches files at anonymous FTP sites for a user-specified filename or part of a filename; it does not search file contents. We would have to use additional tools, such as VERONICA, to provide text search ability.

Finally, the tool that met nearly all of our requirements came in the form of Mosaic. Mosaic provides a client-server method of traversing pre-established menus (presented to the user as "pages") via hypertext links. Although Mosaic does not have a search ability, WAIS provides a keyword search that integrates tightly with the Mosaic interface. Mosaic runs on every platform we required, handles multi-media, has the ability to launch external applications and viewers. The HTML forms capability allows us to easily implement most of the reuse related features (such as user registration and problem reporting) that do not come by default with the stock Mosaic server (such as usage statistics).

5 Approach using Mosaic

Porting the contents of the RSL from the mainframe to distributed platforms demanded reorganization of the FRR assets from the flat file system on the mainframe to the hierarchical system used by AIX systems. Each FRR component normally comes in its own file, along with up to 15 files containing information supporting its use; e.g., design specifications, an abstract, and integration instructions. We made the natural choice to store an asset and its supporting information together in one AIX directory. We then chose to organize the groups of assets based on source; e.g., a parent directory for all components supplied by Program *A* or licensed from Company *B*. Although this organization may not necessarily make it easy for reusers to locate assets, it made control of the assets much easier. We decided to provide alternate views (or indices) of the FRR, such as grouping assets by function, using html pages. The use of the hierarchical file system not only allowed a very nice way to organize the FRR contents, but it obviated the need for an underlying database

system to manage the data and therefore required little investment and effort.

The second major consideration concerned security. Again we relied on the safeguards already in place in our environment; Andrew File System (AFS) Access Control Lists (ACLs) and standard AIX file permissions allowed us to grant and deny access to the FRR. We actually authorize users through the use of subnet masks based on Internet Protocol (IP) address; this allows us acceptable level of access control to users throughout Loral Federal Systems.

Loral FS Reuse

You've reached the WWW page for Loral Federal Systems Reuse, home of the Federal Reuse Repository (FRR). The FRR contains several hundred reusable software and document components; click here for information about the FRR.

FRR, hierarchical view (arranged by Language/Library)

FRR, arranged by Subject

Search the FRR for needed components (via WAIS)

Click here to find out more about the LFS Reuse group or to send them mail.

Figure 1: The FRR Home Page

The figures show the implementation of the FRR using Mosaic. Figure 1 shows the FRR Home Page. As shown, we currently provide three ways to browse the FRR; (1) by a hierarchical view which mirrors the organization in the AIX file system, (2) by subject based on the major function or service provided by the component, and (3) via keyword search using WAIS.

We sorted the hierarchical view based first on implementation language. If the user selects the Ada programming language, the mosaic page shown in Figure 2 appears. This page lists the various sources of reusable Ada software. Selecting "Circuit Card Assembly and Processing System (CCAPS)," one of the programs that produced software for the FRR, results in the page shown in Figure 3.

Note that *monitors* and a *semaphore* both appear as components supplied from the CCAPS program. This hierarchical approach shows one way a user can get to this information. The user can also get to these same components using the second method on the FRR Home Page; Figure 4 shows the page produced if the user elects to search the *FRR, arranged by Subject*. Note that *Synchronization Components* appears on the FRR component listing by subject. Since monitors and semaphores both provide synchronization functions, the user will find them on the *Synchronization Components* page along with other monitors, events, barriers,

Ada Reuse Collections

Select the library you need:

- Advanced Automation System (AAS)
- Ada Run Time Environment (ARTE)
- Booch – Ada
- Booch – Ada Enhanced
- Common Ada Missile Packages (CAMP)
- Circuit Card Assembly and Processing System (CCAPS)
- General Purpose Ada
- Karlsruhe
- Realtime And Distributed Ada Services (RADAS)

Figure 2: The Ada Language Home Page

locks, pulses, and related Ada language synchronization from sources other than CCAPS.

The third search option consists of a standard ISINDEX WAIS search form. Because WAIS integrates well with xmosaic, we selected it for simple keyword searches of the FRR. By using the classification data from the original RSL to index the WAIS database we feel we have achieved a reliable index for keyword searches.

6   Related Work

The Repository Based Repository for Software Engineering (RBSE) research and development group has developed the Multimedia Oriented Repository Environment (MORE) using Mosaic and the Web as its sole user interface. MORE provides client browsing, search, repository definition, and data entry through Web clients. Unlike the system in this paper, MORE uses a meta-data based repository, or a database that contains information about the reusable assets rather than the actual assets. With the exception of the system home page the MORE dynamically generates the entire user interface [6].

7   Future Work

Our future activities include providing several searching capabilities for retrieval of software modules and associated documentation based on the *Structured Abstract* concept proposed in [9]. The Structured Abstracts will allow users to provide a more detailed search than the simple keywords offered by WAIS by using the classification information previously developed for the original RSL. Using html forms for the Structured Abstracts will also allow us to keep the simple, easy-to-use GUI interface. To provide access to the FRR to users of ascii-only terminals such as 3270 protocol sessions, we plan to implement a Lynx interface.

We will complete implementation of html forms for the "check out/check-in" capabilities of the Software Reuse Library via a back-end database and the webserver CGI. Keeping in mind the desire for a simple, intuitive tool, we plan html forms for component user registration, problem reporting, and submission forms

CCAPS

Circuit Card Assembly and Processing System

Please review the legal information before using any of these components. If you would like to preview descriptions of the CCAPS components, here is a collection of the CCAPS abstracts.

Select the component you need:

- ada interface to ibm gddm
- generic timer facility
- convert integer to string
- line parse manager
- msgm asm interface
- monitors
- rscs io interface
- semaphore
- system commands

Figure 3: The CCAPS Home Page

for new software modules. In addition, we will also look at integrating documentation (authoring) tools into the RSL environment to allow users to make their own html hyper-linked documentation point to the relevant reusable software modules in the FRR.

8 Acknowledgements

We would like to acknowledge Jim McKinstry for his early work on this project and Allen Matheson of the Cimarron Corporation for his continued support of the FRR and the key role he has had in implementing the work described in this paper. We would also like to thank Gary Kennedy and Tom Loggia of Loral Federal Systems-Bethesda for supporting this project.

9 Cited References

1. Asset Source for Software Engineering Technology (ASSET), Building 2600, Suite 2, 2611 Cranberry Square, Morgantown, WV, 26505, e-mail: *info@source.asset.com.*

2. Berg, Klaus, "CLASSLIB - Class Management and Reuse Support on a MVS Mainframe," *Reusability Track of the 1994 ACM Symposium on Applied Computing (SAC'94),* Phoenix, Arizona, 6-8 March 1994, pp. 53-58.

Select the type of component you need:

- Bit/String Manipulation Components
- Command Line Components
- Communication Components
- Data Structures Components (long list!)
- File Services Components
- Graphics Components
- Input/Output Components
- Miscellaneous Utilities
- Numerics and Math Packages
- OS (POSIX) Interfaces
- Real–Time Components
- Sorting and Searching Routines
- Synchronization Components

Figure 4: The FRR, Arranged by Subject

3. Berners-Lee, Tim, Robert Cailliau, Jean-Francois Groff, and Bernd Pollermann, "World-Wide Web: The Information Universe," *Electronic Networking: Research, Applications and Policy*, Vol. 11, No. 2, Meckler, Westport, CT, Spring 1992.

4. Berners-Lee, Tim, Robert Cailliau, Ari Luotonen, Henrik Frystyk Nielen, and Arthur Secret, "The World-Wide Web," *Communciations of the ACM*, Vol. 37, No. 8, August 1994, pp. 76-82.

5. Donaldson, Cameron, "InQuisiX: An Electronic Catalog for Software Reuse," *SIGIR Forum*, Vol. 28, No. 1, Spring 1994, pp. 8-12.

6. Eichmann, David, Terry McGregor, and Dann Danley, "Integrating Structured Databases Into the Web: The MORE System," *First International Conference on the World Wide Web*, Geneva, Switzerland, 25-29 May, 1994.

7. Maarek, Yoelle S., Daniel M. Berry, and Gail E. Kaiser, "An Information Retrieval Approach for Automatically Constructing Software Libraries," *IEEE Transactions on Software Engineering*, Vol. 17, No. 8, August 1991, pp. 800-813.

8. Obraczka, Katia, Peter B. Danzig and Shih-Hao Li, "Internet Discovery Services," *IEEE Computer*, September 1993, pp. 8-22.

9. "Automatic Generation of a Rapid Assessment Aid for Reusable Components," Jeffrey S. Poulin *Research Proposal*, Loral Federal Systems-Owego, Owego, NY, 18 July 1993.

10. Poulin, Jeffrey S. and Kathryn P. Yglesias, "Experiences with a Faceted Classification Scheme in a Large Reusable Software Library (RSL)," *Seventeenth Annual International Computer Software and Applications Conference*, Phoenix, AZ, 3-5 November 1993, pp. 90-99.

11. Poulin, Jeffrey S., "Populating Software Repositories: Incentives and Domain-Specific Software," to appear, *Journal of Systems and Software*, Spring 1995.

12. Prieto-Diaz, Ruben and Peter Freeman, "Classifying Software for Reusability," *IEEE Software*, Vol. 4, No. 1, January 1987, pp. 6-16.

13. Prieto-Diaz, Ruben, "Implementing Faceted Classification for Software Reuse," *Communications of the ACM,* Vol. 34, No. 5, May 1991, pp. 88-97.

14. Yglesias, Kathryn P., "Limitations of Certification Standards in Achieving Successful Parts Retrieval," *Proceedings of the 5th International Workshop on Software Reuse,* Palo Alto, California, 26-29 October 1992.

10 Vitas

Jeffrey S. Poulin, Ph.D.

Dr. Poulin works with the Advanced Technology department of Loral Federal Systems-Owego where he serves as Principal Investigator (PI) for Open Systems Environment Independent Research and Development (IRAD). His past Loral Federal Systems-Owego responsibilities include:

- Reuse Strategy, Sustaining Base Information Services (SBIS) program.

- SBIS Engineering and Technology Team Lead, Integrated Software Engineering Development Environment, Open Systems Development group.

Dr. Poulin formally served with the IBM corporate Reuse Technology Support Center (RTSC) where his responsibilities included reuse standards, economics, and legal issues. As part of his reuse metrics work, he helped lead the development of the IBM reuse measurements and return on investment (ROI) model. His background includes semantic data modeling in object-oriented database systems with a focus on support for Computer Aided Software Engineering (CASE). He participates in the Association for Computing Machinery and the IEEE Computer Society. A Hertz Foundation Fellow, Dr. Poulin earned his Bachelors degree at the United States Military Academy at West Point and his Masters and Ph.D. degrees at Rensselaer Polytechnic Institute in Troy, New York.

Keith J. Werkman, Ph.D.

Dr. Werkman, formerly employed by IBM's Federal Systems Company in Owego, NY, now participates as a member of Loral Federal Systems-Owego's Advanced Technology department assisting as one of the sites resources in artificial intelligence (AI). One of his research tasks integrates AI into a variety of new and existing business areas. Dr. Werkman has worked on developing AI tools with multimedia interfaces to support these enhanced user environments at LFS-Owego. His research interests include using distributed AI to support a variety of user information environments including those for software reuse, agile manufacturing, concurrent engineering, group decision support systems and groupware. He participates in the American Association of Artificial Intelligence (AAAI), Association of Computing Machinery (ACM), ACM SIGART, and IEEE Computer Society.

Dr. Werkman earned his Bachelors and Masters degrees from Lehigh University in Bethlehem, PA. While employed as a National Science Foundation Engineering Research Center (NSF-ERC) Fellow at the ATLSS Center and Lehigh, Dr. Werkman earned a doctorate in Computer Science based on his distributed AI negotiations research. Dr. Werkman has over 30 publications in the areas of Distributed Artificial Intelligence (DAI), Concurrent Engineering (CE), AI in Design, Computer Supported Cooperative Work (CSCW), Design for Manufacturability (DFM) and Enterprise Integration (EI).

For more information, contact:

poulinj@lfs.loral.com or keithw@lfs.loral.com

NII Awards Sponsors

Primary Sponsors

!NTERPRISE
Networking Services from US West

AT&T

intel.

IBM

Price Waterhouse LLP

Lotus.

sun microsystems

BELLSOUTH
Business Systems

ALCATEL

ACCESS MEDIA

EnterActive
group

Sponsoring Organizations

Academy of Television Arts and Sciences
Alliance for Public Technology
The American Film Institute
American Health Information
    Management Association
American Library Association
American Medical Association
Americans for Indian Opportunity
Annenberg/CPB Project
AFFIRM
Center for Civic Networking
C. Everett Koop Institute
Consortium for School Networking
Corporation for Public Broadcasting
Council on Competitiveness
Financial Services Technology Consortium
IDC Government

IDG World Expo
IEEE-USA
Intl. Communications Association
Intl. Society for Technology in Education
League of Women Voters
Morino Institute
NASIRE
National Academy of Public Administration
National Association of Counties
National Civic League
National Coalition on Black Voter Participation
National Education Association
National League of Cities
Public Technology, Inc.
US Distance Learning Association
US Information Infrastructure Task Force
World Institute on Disability

Media Sponsors

AMERICAN CITY & COUNTY.

HEALTH MANAGEMENT

T·H·E JOURNAL

BUSINESS COMMUNICATIONS REVIEW

INFO WORLD

TELECOMMUNICATIONS.

educational technology

Modern Healthcare

RED HERRING

Electronic Learning

Nation's Cities WEEKLY

VARIETY

FEDERAL COMPUTER WEEK

NETWORK WORLD

Wilson Library Bulletin

GOVERNMENT TECHNOLOGY

AHIMA

Category Sponsors

HEWLETT PACKARD

TEXAS INSTRUMENTS

Produced by Access Media Inc. and Triax Enteractive
All logos are registered trademarks

100% recycled paper

THE NII AWARDS

THE NATIONAL INFORMATION INFRASTRUCTURE AWARDS

Your opportunity to

be recognized as

one of the brightest

lights on

the Information

Highway

INTRODUCING

THE FIRST

ANNUAL

NATIONAL

INFORMATION

INFRASTRUCTURE

AWARDS

Awards to Celebrate the Best Uses of the Information Highway

The National Information Infrastructure (NII) Awards have been created by more than 50 industry and community leaders in cooperation with the Clinton Administration's Information Infrastructure Task Force.

They have come together because they understand that the NII - the "information highway" - has tremendous potential to benefit all Americans: potential reflected in the achievements of people like you.

To realize that potential requires that all sectors of our society understand, use and participate in the NII. Just like our telephone system, the value of the information highway increases as more people connect to and use it.

This is where you can help and, at the same time, get the recognition you deserve. When you enter the NII Awards, you'll serve as an example of how people can improve their lives, businesses and communities by using the information highway. Whether or not you are selected for an Award, we will acknowledge your achievements by including your entry in a national database. You only need to provide a structured description of your use of the "information highway". There is no entry fee.

The National Information Infrastructure

The National Information Infrastructure is the combination of public and private networks, computers and electronics that connects people with people and people with information. It is designed to allow people to communicate, collaborate and access information anytime, anywhere. The networks of businesses, carriers, schools, communities and government agencies and the Internet all form part of the NII.

Through applications such as electronic commerce, telemedicine, distance learning, community and health networks, 'virtual' libraries and communities, collaborative work, online information services and interactive entertainment, the NII is being used today to improve people's productivity, health and well being.

The Awards

The NII Awards will recognize uses, applications and services of the information highway that powerfully demonstrate its capabilities and utility - real examples and success stories that illustrate the NII's potential to provide new benefits and encourage communication, collaboration and access to information beyond traditional boundaries. One Award and up to five Honorable Mentions will be presented in each of six categories:

- Arts & Entertainment
- Business
- Community
- Education
- Government
- Health

Winners

NII Award winners will receive national recognition for their achievements. They will be honored in a special awards ceremony in Washington, DC in the spring of 1995 and will receive a trophy and certificate indicating their status as NII Award winners. Winners will also be featured in a national awareness and education program.

Additionally, all entries will be included in a database made available to all those interested in how they can make use of the information highway in their own lives and work.

For More Information

If you are using networking and information technologies to do extraordinary things, this is the opportunity to get the recognition you deserve.

For your Entry Kit, or to get more information, send an email message to info@niiawards.org and include "subscribe" in your message or call 313-453-9137. Entries must be postmarked no later than December 15, 1994.

The NII Awards were initiated by Access Media Inc., in cooperation with industry and community leaders, as part of the National Infostructure Campaign - a public education program to accelerate the development and use of the NII. The Infostructure Campaign will build awareness of the ways people can use and benefit from the information highway and will provide them with tools to help them realize those benefits.

THE NII AWARDS

THE NATIONAL INFORMATION INFRASTRUCTURE INFRASTRUCTURE

AWARDS TO CELEBRATE BEST USES OF INFORMATION HIGHWAY

*More than 50 Industry, Government & Community Organizations Sponsor
National Information Infrastructure Awards*

WASHINGTON, October 4, 1994 - Business, community, education, entertainment, government and health care organizations will be honored next spring for their use of the "information highway" under a new national awards program, the National Information Infrastructure (NII) Awards.

Sponsored by industry and community leaders in cooperation with the Clinton Administration's Information Infrastructure Task Force, the NII Awards program seeks to recognize those organizations which have achieved concrete gains through use of the information highway. Winners will receive national recognition for their achievements at an awards ceremony in Washington in the spring of 1995 and will be featured in a national education program. All entries will become part of a national database to provide people with examples of how they can use and benefit from the NII.

"The NII Awards program is an outstanding way to demonstrate the enormous benefits that effective use of a national information infrastructure can bring to users in both the public and private sectors," said Richard McCormick, chairman, CEO and president of operations, U S WEST.

"We think networking and information technology will be as important to the 21st century as radio and television were to this century. The NII Awards will stimulate even more creative ways to educate our children and serve the public's interest," said Richard W. Carlson, president and CEO of the Corporation for Public Broadcasting.

"The Awards will help Americans understand and value the enabling capabilities of the National Information Infrastructure. Recognizing these accomplishments will fuel interest and investment in the applications that will enrich our lives today and for generations to come," said Alex Mandl, executive vice president and CEO of AT&T's Communication Services Group.

"The NII Awards will make an important contribution to raising awareness of the applications of the National Information Infrastructure in healthcare and will help accelerate successful implementations by the health and medical communities," said C. Everett Koop, M.D., senior fellow of the C. Everett Koop Institute.

The National Information Infrastructure is the combination of public and private networks, computers and electronics that connects people with people and people with information. It is designed to allow people to communicate, collaborate and access information anytime, anywhere. The networks of businesses, carriers, schools, communities and government agencies and the Internet all form part of the NII.

- more -

Through applications such as electronic commerce, telemedicine, distance learning, community and health networks, 'virtual' libraries and communities, collaborative work, online information services and interactive entertainment, the NII is being used today to improve people's productivity, health and well being.

The NII Awards will be given in six categories: arts and entertainment, business, community, education, government and health. Judges will select winning entries based on how well they demonstrate the capability of the NII, the practical benefits they produce and their ability to encourage broader use of the information highway. Price Waterhouse will administer and verify the judging process.

"The biggest barriers to realizing the full potential of the NII are cultural, not technical," said Jim Manzi, president and CEO of Lotus Development Corporation. "The NII Awards will help break down those barriers by building understanding of how people can use and benefit from the information highway."

"The NII Awards will highlight the benefits of the most innovative applications of networking and information technology in education," said Keith Geiger, president of the National Education Association. "And, at a time when budgets are under great pressure, the Awards will help schools and teachers justify why they should plug into the information highway."

"The reality is that 500-channel movies-on-demand into the home is not what will drive the NII. The real investment is already being made by businesses and organizations that recognize how improved communications-information systems will help them do their work," said Scott McNealy, president and CEO of Sun Microsystems. "The National Infostructure Awards Campaign is the best chance we have of setting the record straight by recognizing those who are contributing to these solutions."

"The NII Awards will help drive the development of practical real-world applications which will make information on demand as ubiquitous as the personal computer," said Craig R. Barrett, chief operating officer of Intel Corporation.

"The one-size-fits-all era is over. It's time that technologies begin to more closely fit our customers' unique needs," said F. Duane Ackerman, president and CEO of BellSouth Telecommunications. "We believe our participation in the NII Awards Campaign is an excellent way to reinforce our focused applications development and solutions delivery on today's real-world communications needs in communities of interest that include Arts & Entertainment, Business, Education, Government and Healthcare."

The NII Awards organization has been convened by Access Media Inc. as part of the National Infostructure Campaign - a public education initiative to accelerate the development and use of the NII. The Infostructure Campaign will build awareness of the ways people can use and benefit from the information highway and will provide them with tools to help them realize those benefits.

Information on the Awards may be obtained from info@niiawards.org or from Access Media at 313-453-9137. Entries must be postmarked by December 15, 1994. No Entry Fee is required.

- end -

Topic Info Server
for the
World-Wide Web

Internet information navigation
(alpha version 0.3)

Verity's Topic Info Server for the World-Wide Web brings powerful information tools to Internet information publishers and providers. The Internet's fastest-growing set of services, the World-Wide Web provides consistent, cross-platform servers and user interfaces. Verity tools add sophisticated search capabilities to the Web, including state-of-the-art concept search techniques necessary for sifting through the Internet's terabytes of accessible data.

• **Create virtual libraries of distributed documents.** Verity's network indexing tools create searchable groups of documents, creating virtual libraries of any size. Documents can be logically grouped without the need to maintain local versions, creating efficient distributed information systems.

• **Index and serve multimedia information.** Newspaper publishers, support operations and other information publishers and providers are using Verity's Web tools to index and distribute news, technical data, marketing, archives and other data. The Web's integration of multimedia in HyperText Markup Lanuage (HTML) allows Verity's Web tools to index and serve images, graphics, audio and other multimedia data embedded in Web documents.

• **Use Verity's comprehensive information tools.** Verity indexes, called *collections,* contain comprehensive information about each document's contents, location, retrieval method and other key attributes. Web-based collections can be used by Verity's entire line of information navigation products. The Web tools are add-ons to Verity's Topic Database Building V3/V4 or to the Verity Development Kit. Supported documents include more than 50 word processing, spreadsheet, image and database formats, eliminating the need for translation from non-HTML formats.

Web Search Server

Verity's Web server accepts users' queries as generated by standard or fully customizable HTML forms. The server the returns an HTML document containing a relevancy-ranked list of documents. Selecting a document from the list causes the server to return selected information about the document, along with a link to the original. Publishers can choose which fields to include in the

results list and the document information page. Local documents are returned with hyperlinked matching terms at the top, allowing easy navigation to the relevant sections.

Features:

- Platforms — SUN/OS 4.1.x, HP700 HPUX Release 9, Microsoft Windows NT.
- Up to 12 simultaneous search sessions.
- Search via Verity's Simple Query language, including support for words, phrases, operators, evidence weighting and combination and concept-based libraries of Topics®.
- Support of the World-Wide Web HyperText Transfer Protocol (HTTP) including multimedia document delivery and launch of document viewers.
- Customizable look, feel and user authorization.
- Personal information views for fast, high-precision searching.
- Security features — limit access based on user profiles, content, or both.
- Standard access log files that can be analyzed with Web log file statistical tools.

Indexing Spider

Verity's Web "spider" works with the Verity engine to index a set of Web documents. Give the spider a starting point and it automatically indexes all linked documents on the same server.
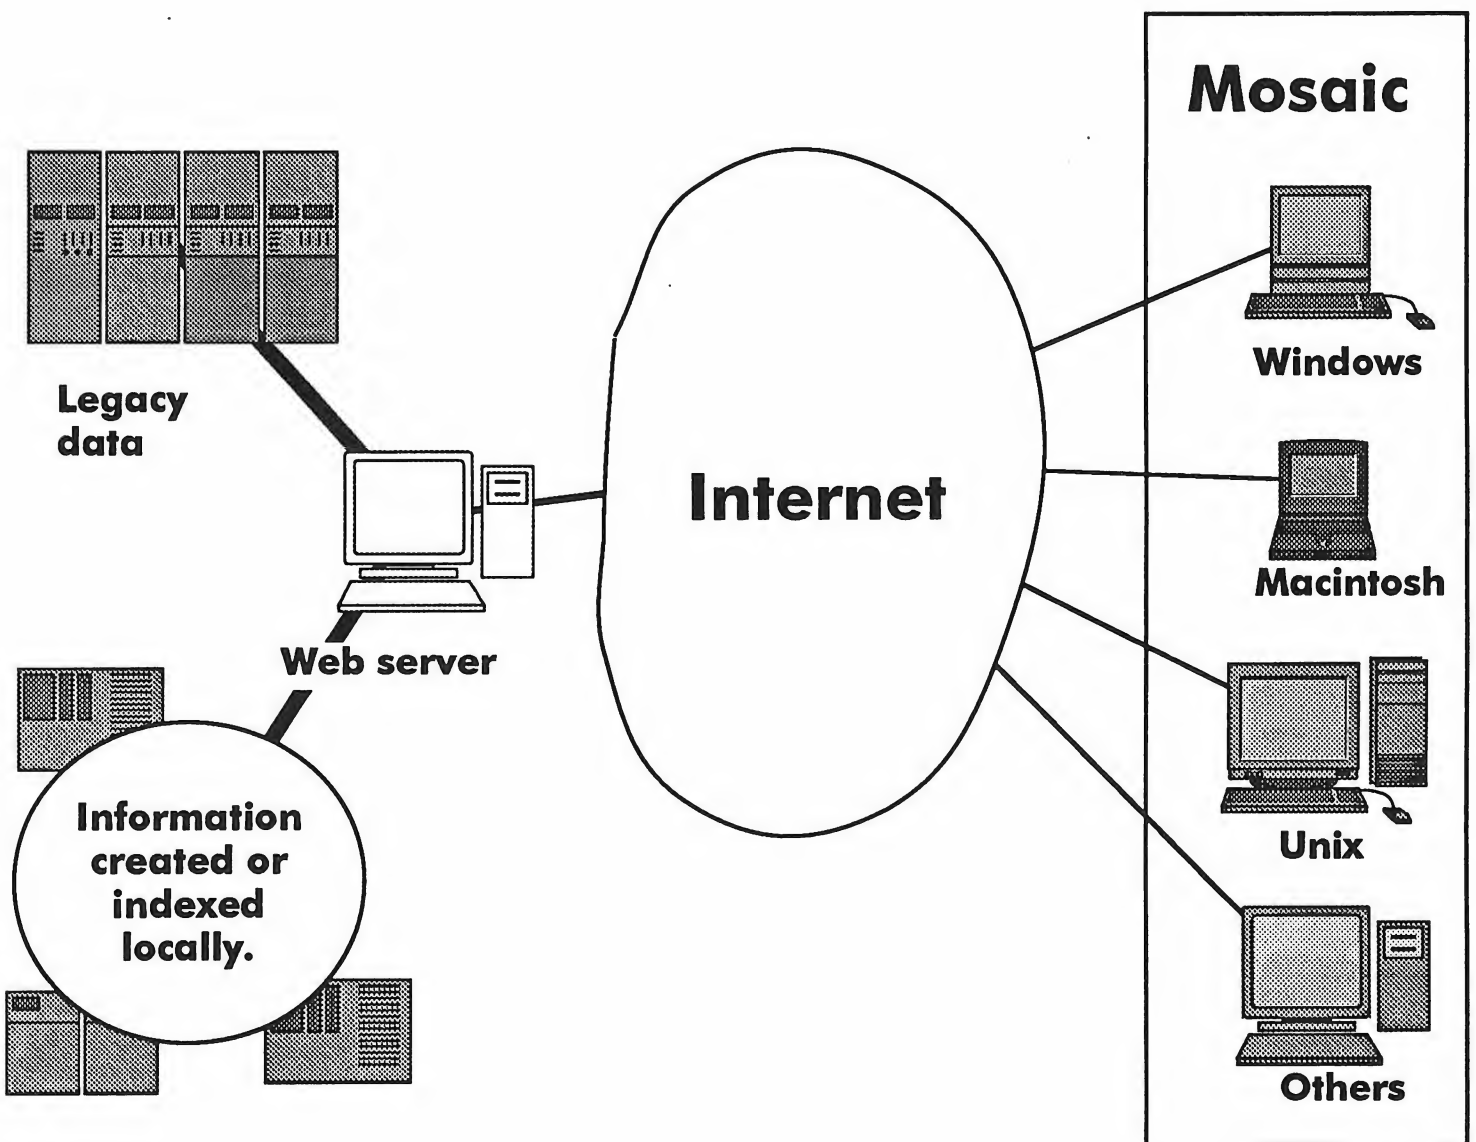
Features:

- Speed — Indexing typically is limited only by document latency.
- Source code is available for customization.
- Platforms — SUN/OS 4.1.x, HP700 HPUX Release 9, Microsoft Windows NT.

Contact Nick Arnett [narnett@verity.com], product manager.

Verity Inc.
1550 Plymouth St.
Mountain View, CA 94043
(415) 960-7600

Topic and the World-Wide Web

The World-Wide Web allows publishers, libraries, archives and others to create and maintain structured, organized documents that are easily browsed and searched by others around the world, via the Internet, regardless of their location or computer platform. Navigable collections can include "virtual libraries" of documents located on machines on any Internet-accessible computer, anywhere.

Mosaic

Legacy data

Internet

Windows

Macintosh

Web server

Unix

Information created or indexed locally.

Others

~~I AM NOT SURE~~

Wants: URL aliasing
CGI access to exact URL

SAM R

Short Docs, few graphics

Educational research of hypertext:
    good students do really well
    average students do worse : need more direction
    Students like hardcopy, which implies longer text
    table of context for multipage docs.
    Modification at beginning
    TWO tables of contents: one overview, one detailed
    Inform students about real URL.

    few actually used WWW for class
    Typical comment: more stuff there we could print

    URL: http://www.cs.dartmouth.edu/~sam

NCSA

    Simsurface: simulated analogy

    Internet services http://calliope.ncsa.~/ELM/InteResources/Edunet
                                                        .html

---

Steve Harly: Cardiff, Wales /
    < WWW on PARALLEL COMPUTING >

    msg passing facts wheam, mread, mwrite.

John Schmidt
    World Explorer System
    R&D Env for Web-based hypercurricula & hypertexts
        System Requirements:
            • Content base: Large, Global, Interdisciplinary, Multimedia, Format Uniformity
            • Data collection: hypertrail observation
            ○ Integrated: ...
            • Ecological: ready for real world apps.
                    • Whole language
                    • cognitive flexibility
                    • HOTS
                    • Demands of every day life

THE SECOND INTERNATIONAL
WWW CONFERENCE '94
Mosaic and the Web

SESSION EVALUATION

Please complete this evaluation form and drop it in one of the boxes located near the conference registration desk or near the message board (Francis I Room, 2nd Floor). You may also leave this form with this session's room monitor or mail it back using the address at the bottom of this page. Your feedback is important to us and will be used to improve future conferences.

SA = strongly agree   **A** = agree   **N** = neutral   **D** = disagree   **SD** = strongly disagree

Session title _____

Presenter #1 (name) _____

The session was well presented.	SA	A	N	D	SD
The content was relevant to my job.	SA	A	N	D	SD
The handout material was useful.	SA	A	N	D	SD
The presentation met my expectations.	SA	A	N	D	SD

Presenter #2 (name) _____

The session was well presented.	SA	A	N	D	SD
The content was relevant to my job.	SA	A	N	D	SD
The handout material was useful.	SA	A	N	D	SD
The presentation met my expectations.	SA	A	N	D	SD

Presenter #3 (name) _____

The session was well presented.	SA	A	N	D	SD
The content was relevant to my job.	SA	A	N	D	SD
The handout material was useful.	SA	A	N	D	SD
The presentation met my expectations.	SA	A	N	D	SD

Presenter #4 (name) _____

The session was well presented.	SA	A	N	D	SD
The content was relevant to my job.	SA	A	N	D	SD
The handout material was useful.	SA	A	N	D	SD
The presentation met my expectations.	SA	A	N	D	SD

Presenter #5 (name) _____

The session was well presented.	SA	A	N	D	SD
The content was relevant to my job.	SA	A	N	D	SD
The handout material was useful.	SA	A	N	D	SD
The presentation met my expectations.	SA	A	N	D	SD

Additional comments:

_____

_____

_____

PLEASE PRINT (optional information)

Name _____

Organization _____

Address _____

Phone _____

E-mail address _____

NCSA Mosaic

NCSA Mosaic is a window on the ever-expanding world of online information, bringing the most advanced multimedia networked information resources to any user's fingertips. As a *distributed hypermedia browser* designed for information discovery and retrieval, NCSA Mosaic provides a unified interface to the diverse protocols, data formats, and information archives used on the Internet.

The Internet provides access to a worldwide collection of information resources and services that has grown up over the last two decades. Each element of this global community has its own unique history and flavor, culture and methods. For a new user, this can be confusing; NCSA Mosaic draws these disparate pieces into a seamless picture that can be viewed effortlessly.

The NCSA Mosaic interface is based on the idea of *hypermedia*, where electronic links—known as *hyperlinks*—are embedded in richly formatted documents that can include full-color images and sound. These documents are presented to users like the pages of an interactive, scrollable, online book. NCSA Mosaic thus gives users the ability to move around in a complex document, and from document to document across the network, by clicking on the hyperlinks. Throughout the user's travels, this consistent and simple interface is used for both navigation and document viewing.

In addition to giving users access to the World Wide Web servers for which it was designed, NCSA Mosaic transparently integrates material from a wide variety of existing information servers on the Internet—including Gopher, WAIS, and anonymous FTP servers—without moving to a different application for each one. Such capabilities open a number of powerful new approaches for discovering, using, and sharing information in the network environment. For example, the user can take advan-

N C S A MOSAIC

X Window System • Microsoft Windows • Macintosh

tage of the full-text and index-based search tools provided by WAIS to locate a particular document. As development continues, the system's flexible design allows other information and data resources, such as relational databases, to be integrated into the NCSA Mosaic environment.

NCSA Mosaic and the World Wide Web

NCSA Mosaic is part of the World Wide Web environment—a system for maintaining distributed hypertext that originated at the European Laboratory for Particle Physics (CERN) in Geneva, Switzerland. Initially developed to keep track of researchers' information and to provide an easy method of sharing information among scientists, the Web has grown into one of the world's most widely used environments for information publishing, discovery, and retrieval.

The Web employs several established or developing standards to make it as universally usable and reliable as possible. The Web uses a standard naming convention to locate individual pieces of information sitting on machines anywhere across the network; this standard is known as the *Uniform Resource Locator*, or URL. Clients and servers use a standard transfer protocol to send and retrieve documents across the Web; this protocol is known as the *HyperText Transfer Protocol* (HTTP). Documents on Web servers are formatted with a standard markup language known as the *HyperText Markup Language* (HTML). HTML specifies the format for hyperlinks that allow the user to easily move around the Internet from one document to another.

Building on the Web's initial structures, NCSA Mosaic uses a client/server model for information distribution. A *server* sits on a machine at an Internet site answering queries sent by NCSA Mosaic *clients*, which may be located anywhere on the Internet. To a user, the client looks like any other machine application, only this one has immediate access to information all over the world! The pieces of information sent from servers to clients are known simply as *documents*.

N C S A MOSAIC

X Window System • Microsoft Windows • Macintosh

NCSA Mosaic is implemented for three types of platforms:
- X Window System on UNIX platforms
- Macintosh
- Microsoft Windows

Setting Up the Client

All three clients can be obtained for personal use from NCSA's anonymous FTP site: `ftp.ncsa.uiuc.edu`. (See Downloading Clients from the NCSA FTP Server toward the end of this document for specific information.)

Installing and running an NCSA Mosaic binary is straightforward. The procedures differ slightly across the X Window System, Macintosh, and Microsoft Windows platforms, but they are installed like any other application. The installation procedures are fully described in the readme files referred to in Downloading Clients from the NCSA FTP Server and in documentation available when downloading and uncompressing the files.

If NCSA Mosaic is to be used as a Web browser, the system on which it is installed must be operated in one of two modes: the system must be fully connected to the Internet or, if a dialup connection is used, it must be run through an account that provides full Internet connectivity through

either a SLIP or PPP connection (or an ARA connection on a Macintosh). If installed as distributed by NCSA, the first thing the program does when it is launched is download one of the NCSA Mosaic Home Pages (a startup document) from an NCSA server. If NCSA Mosaic executes but displays an error message instead of displaying a home page, the problem is most likely in the Internet connection. As an option, users who wish to avoid this initial contact with an NCSA server can configure NCSA Mosaic to load a local home page or any other startup document.

Should installation problems be encountered we suggest initially contacting your local systems administrator.

External Viewers

There are many types of data available on the Internet, many of which must be read and interpreted in a unique manner. It would be difficult, if not impossible, to design NCSA Mosaic itself to read and interact directly with all of them.

Therefore, NCSA Mosaic relies on *external viewers* to work with the widest possible variety of image, audio, video, and typeset data for-

mats. These external viewers are separate programs that are invoked when necessary to display certain types of data.

For example, the various NCSA Mosaic clients can use the following external viewers:

`lview` A Windows utility that displays GIF, JPEG, and TIFF images, as well as images recorded in several other formats (for the Microsoft Windows client)

`Ulaw` A Macintosh utility that plays audio data (for the Macintosh client)

`xv` An X utility that displays GIF, JPEG, and TIFF images, as well as images recorded in several other formats (for the X Window System client)

These are just samples of the available viewers. NCSA makes several standard viewers available, but the system can be configured to accommodate others. Information regarding available viewers can be found in the installation instructions for each platform. Other than providing information pointing to available software, NCSA does not provide technical support for external viewers.

Where to Go from Here

Once NCSA Mosaic has been successfully installed, spend some time becoming familiar with the interface it presents to the information available on the Internet. The NCSA Mosaic Demo Document, available via a hyperlink in the NCSA Mosaic Home Page or via a menu selection, is a self-contained overview of NCSA Mosaic's capabilities with hyperlinks to a wide variety of interesting information sources. Other menu selections and home pages give you access to a world of interesting material.

The Clients

NCSA has released the following versions of NCSA Mosaic:

- **NCSA Mosaic for the X Window System**
  This version can be used on almost any UNIX-based graphics workstation (e.g., Sun Sparc, IBM RS/6000, DEC 5000 or Alpha, Silicon Graphics IRIS).
- **NCSA Mosaic for the Macintosh**
  This version can be run on any Macintosh, including SEs, Classics, and Power Macintoshes, running System 7.0 or higher.
- **NCSA Mosaic for Microsoft Windows**
  Current versions can be used on any PC that matches or exceeds the following minimum configuration: Intel 80386SX microprocessor, Microsoft Windows release 3.1, and 4Mb of RAM. Current versions can also be used on any PC running Microsoft Windows NT, regardless of the microprocessor.

Email queries concerning these releases, as well as bug reports, can be sent to the following addresses:
Macintosh:
  mosaic-mac@ncsa.uiuc.edu
Microsoft Windows:
  mosaic-win@ncsa.uiuc.edu
X Window System:
  mosaic-x@ncsa.uiuc.edu

Downloading Clients from the NCSA FTP Server

NCSA Mosaic can be acquired from NCSA's anonymous FTP server. You must first execute the following ftp command:
  `ftp ftp.ncsa.uiuc.edu`
Enter `anonymous` at the login prompt and press Return. Enter your email address (e.g., `jdoe@business.com`) at the password prompt.

The executables are found in the directories noted below by platform. Downloading instructions are in an instructional file for each client.

NOTE: These clients are copyrighted; NCSA and the University of Illinois retain the copyright but will allow anyone retrieving the software from the FTP server to use it personally without any further license. Anyone who wishes to distribute the software or use it commercially must first obtain a license.

The X Window System Executables

Binary executables of the X Window System client for several common UNIX platforms are in the directory `/Mosaic/Unix/binaries`.
Retrieval instructions are in the file `/Mosaic/Unix/binaries/README-binaries`.
Source code for the X Window System client can be found in the directory `/Mosaic/Unix/source`. (This source code is provided only because some users must compile it specifically for their systems. Distribution and commercial use require a license.)

The Macintosh Executables

The executables for the Macintosh client are available in the directory `/Mosaic/Mac`.
Retrieval instructions are in the file `/Mosaic/Mac/QuickStart.Txt`.

The Microsoft Windows Executables

The executables for the Microsoft Windows client are available in the directory `/Mosaic/Windows`.
Retrieval instructions are found in the file `/Mosaic/Windows/readme.now`.

Setting Up a Server

NCSA has developed a freely available HTTP server for UNIX systems. It is available on NCSA's FTP server (`ftp.ncsa.uiuc.edu`) in the directory `/Mosaic/ncsa_httpd`. (See Downloading Clients from the NCSA FTP Server for instructions.)
Full configuration instructions are provided with each of the binary and source distributions available there. Send email to httpd@ncsa.uiuc.edu if you have any questions or problems.

## NCSA Mosaic for Windows, X Windows, and Macs

NCSA Mosaic is implemented for three types of platforms:

- X Window System on UNIX platforms
- Macintosh
- Microsoft Windows

Each implementation takes advantage of the strengths of its respective platform, but all three have been designed to preserve as much cross-system compatibility as possible.

An Integrated System

The NCSA Mosaic system is an integrated set of browsers, viewers, servers, gateways, and filters that allows the user to approach the Internet as one consistent information source. At the simplest level, NCSA Mosaic clients provide navigation and document viewing capabilities for browsing the information universe of the Internet. As a user's interest, needs, and skills develop, the configurable addition of external viewers allows easy and straightforward expansion to handle virtually any specific type of data. This flexibility lets NCSA Mosaic keep pace with the quickly evolving world of multimedia and to be customized for any working environment.
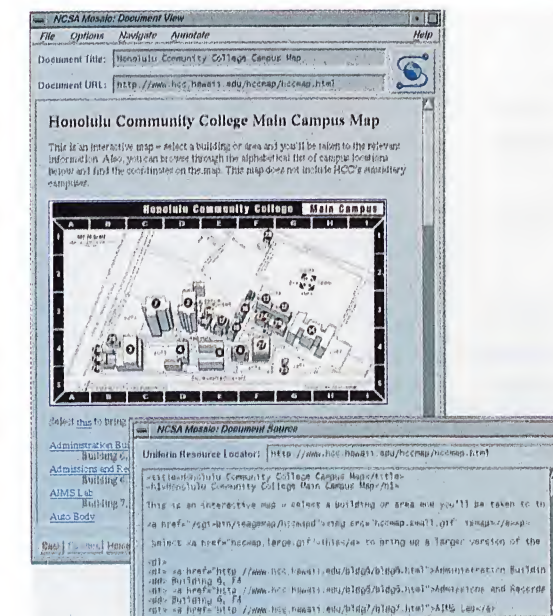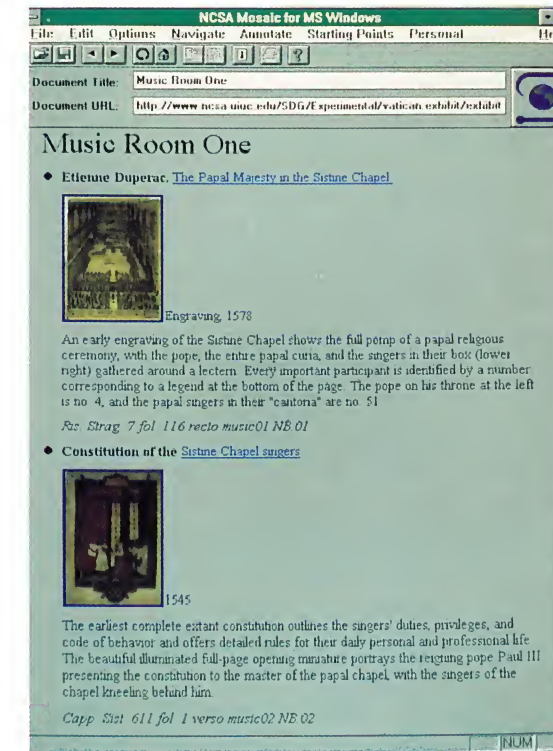
Top NCSA Mosaic for the Microsoft Windows client interface. Each underlined word or phrase is a hyperlink; clicking the mouse on a hyperlink causes the client to connect to the appropriate server, which can be located anywhere on the Internet, to retrieve and display the referenced document.

Bottom NCSA Mosaic for the X Window System client interface. The user can view document source code in a separate window, revealing the HTML formatting commands.

The NCSA Mosaic system is a completely open framework, allowing a user to enter at the most comfortable level, to move through various viewers, to create hypertext ,documents, to set up a personal server, to engage in multimedia collaboration with colleagues in distant locales, to develop scripts for specialized information filtering and presentation, to develop gateways to unique information resources, and to integrate them into the NCSA Mosaic information space for other users. Think of NCSA Mosaic and the Web as allowing the progressive customization of personal or group information space.

Setting Up a Server

To distribute information via the World Wide Web, a user will want to set up an HTTP server. The HTTP protocol is stateless, lightweight, and extremely fast, and provides capabilities not found in earlier protocols, such as FTP. HTTP server software is currently available from various Web sources for UNIX, Macintosh, and Microsoft Windows systems.

Learning about HTML, the SGML-based (Standard Generalized Markup Language) markup language used for formatted hypermedia documents in NCSA Mosaic, is a first step toward preparing information for publication on the Web.

A user trying to find or publish information will also want to learn about the URL scheme for consistently naming documents accessible on the Internet.

At the very end of the NCSA Mosaic Demo Document, which is accessible online through NCSA Mosaic, there are hyperlinks to primers on HTML and URLs. Those primers in turn point to more advanced sources.

For More Information

The best way to get more information on NCSA Mosaic, the World Wide Web, the Internet, and related technologies is to get online and start browsing. The online information is extensive and constantly updated.

NCSA Mosaic can be acquired from NCSA's anonymous FTP server at this FTP address: `ftp.ncsa.uiuc.edu` in the `Mosaic/` directory. Users without FTP access can order software by contacting `orders@ncsa.uiuc.edu` or (217) 244-4130.

Inquiries concerning the NCSA Mosaic clients can be emailed to `mosaic-x@ncsa.uiuc.edu` for the X client, `mosaic-mac@ncsa.uiuc.edu` for the Macintosh client, or `mosaic-win@ncsa.uiuc.edu` for the Microsoft Windows client.

NCSA Mosaic for the Macintosh client interface. Graphics can be displayed both within the NCSA Mosaic window and via external viewers.

SDG

Software Development Group

The Software Development Group (SDG) at the National Center for Supercomputing Applications (NCSA) is currently involved in several areas of software development:

- Scientific data analysis and visualization tools
- Collaborative tools and networked information tools
- Scientific data management tools

Until very recently, scientists had to search through pages and pages of numbers representing their experimental results, looking for trends to support their hypotheses or point them in new directions. The time needed for analysis meant time lost for further experimentation. And the available data analysis software required extensive customization for application to simulation analysis.

SDG has focused on improving scientists' productivity by providing software access to a distributed computing environment for scientific modeling, analysis, visualization, and output. With the constant improvement of desktop computing, it is vital to develop scientific tools that harness this power—not just for the few, but for the entire scientific community—no matter what the discipline or hardware platform.

SDG has made similar gains with information dissemination and retrieval tools and collaborative tools. Until recently, digital information was shared by sending disks and tapes back and forth among scientific labs and computing research groups. With the advent of the Internet, information could be sent to or retrieved from a distant system through the use of tools like FTP, bypassing the shipping of disks and tapes. But if several researchers wanted to work on the same data, they still had to either work independently then combine the results or work sequentially

on the same file; viewing and manipulating a file in unison was simply impossible. In such an environment, collaboration was awkward and information tended to be shared only among a limited set of interested parties.

The past several years have seen the development of tools that eliminate many of these limitations. SDG, working with research and development laboratories and commercial concerns around the world, has harnessed these developments and integrated them with the most advanced capabilities of the Internet to make information directly available, on a world-wide basis, to any desktop with a networked PC, Macintosh, or X workstation. NCSA Collage, NCSA Mosaic, and the NCSA HTTPd Server are all products of these efforts.

Through its work on visualization and analysis tools and in the development of collaborative information management tools, SDG has had to deal with the management of scien-

tific data in increasingly distributed, heterogeneous computing environments. These requirements have led to the development of the Hierarchical Data Format (HDF) and supporting software for managing scientific data. With these tools, related but dissimilar data, with metadata, can be stored in a single file; the same physical file can be examined and manipulated from any PC, Macintosh, or UNIX platform, including supercomputers; and a single data set can be built from multiple related files.

Areas of Focus

Scientific research has moved beyond the days of one person, one lab, one computer. Grand Challenge problems, as well as many daily tasks of scientists and engineers, demand collaborative technology. A prime example of this technology at work is the formation of the national MetaCenter—an integration of NCSA, Cornell Theory Center, Pittsburgh Supercomputing Center, and San Diego Supercomputer Center. The MetaCenter hardware environment is providing the foundation for extensive human collaboration. SDG is developing communications support and visualization data analysis software needed to enhance the effectiveness of the MetaCenter's joint efforts.

All NCSA-developed and supported software—tools for analysis and visualization, for communication, and for sharing and transferring data—are listed later in this brochure. Products that are especially useful for collaborative computer projects are described below.

Synchronous collaborative work sessions are now possible with NCSA Collage. This software combines communications and scientific visualization tools that let researchers conduct real-time collaborations from remote sites. NCSA Collage sessions can share data visualization, electronic whiteboard capabilities, text display and editing, and screen capture functions.

NCSA Mosaic has also been developed and is continually being refined at NCSA. With this software, researchers can browse networks, retrieve and exchange information, and conduct asynchronous collaborative work sessions. Based on the World Wide Web, a distributed hypertext-based information system developed at the Conseil Europeen pour la Recherche Nucleair (CERN—European Laboratory for Particle Physics), NCSA Mosaic uses standard existing formats and protocols to encourage widespread use and applicability. NCSA Mosaic supports Data Transfer Mechanism (DTM) that provides text information exchange with NCSA Collage and other DTM-aware applications. Future versions of NCSA Mosaic will support other DTM capabilities and provide for new methodology and technology in Internet-based networked information systems and data formats.

In addition to NCSA Collage and NCSA Mosaic, other SDG efforts continue to support research and institutional collaboration, including foundational systems such as NCSA HDF (Hierarchical Data Format) and DTM, as well as prototyping integration of digitized audio and

video. With improved networking speeds, it will be possible to handle all types of data in real time—including video and computer animations that appear on multiple computer screens around the country.

NCSA HDF is a multi-object file structure that lets people, projects, and machines on a network share scientific data such as images, annotations, multidimensional arrays, irregularly-gridded meshes, and multivariate graphical and scientific data. As networking environments brought data from different people and machines into more frequent contact, the disparity between file formats became frustratingly apparent. SDG recognized the need to facilitate transfer of information and created this foundational system. One of the greatest benefits of NCSA HDF is the interoperability it provides to all NCSA software. Most SDG-developed programs support reading and writing HDF files. NCSA Import2HDF for the Macintosh and NCSA Reformat/XReformat translate many non-HDF files into HDF format on Macintosh and UNIX systems, respectively.
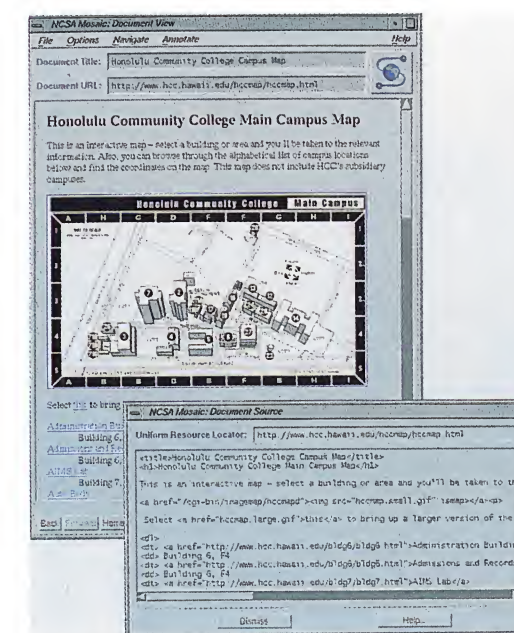
Computational support for scientists is also reflected in discipline-specific software developed by SDG to address particular research problems. This includes tools such as NCSA Mesher (a preprocessor for scientific computational codes requiring 3D meshes) and ALVIS (a 3D alpha shape visualizer).

Geometric visualization—an ongoing SDG project—incorporates mathematical algorithms into software that is robust and efficient. Research and development plans for the coming year address weighted alpha shapes and 3D regular triangulation with a focus on automatic and interactive grid generation.

Many scientists also require sophisticated rendering and analysis of complex datasets. Several tools focus on this need, including NCSA PolyView (an interactive 2D-3D rendering tool that uses HDF Vset files for input format), NCSA Collage

(a collaborative tool that allows multiple users to connect with each other over the network and share scientific data), NCSA Audible Collage for Microsoft Windows (the PC version of NCSA Collage lets a researcher interpret data through sound), and NCSA X DataSlice for the X Window System (a color imaging and data analysis tool for 3D, 32-bit floating-point scientific data stored in the NCSA HDF file format).

SDG is committed to continuing its user support through technical assistance, documentation, and anonymous FTP services. Online documentation and other support services are also being created.

Looking in Different Directions

NCSA is venturing into the social sciences with the addition of social science and human-computer interaction (HCI) research. In the development of collaborative tools, attention is being paid to how these tools affect and enhance relevant social processes. Computer-mediated communication and collaboration will have a tremendous effect on social science research and on how people function in the workplace. As SDG moves forward in the research and development of software tools, it continues its strong relationship with computational science—for both the high-performance computing and general computing communities. NCSA is committed to providing universally accessible resources that facilitate and expedite discoveries in science.

NCSA-developed Tools for Analysis and Visualization

ALVIS (3D ALpha-shapes VISualizer) is a general 3D geometric tool for modeling spatial point sets. Scientific visualization often requires the computation of the shape of a point set. Each shape is a polytope, derived from a Delaunay triangulation of the point set, with a parameter alpha controlling the desired level of detail. ALVIS was developed jointly by SDG and the Department of Computer Science at the University of Illinois at Urbana-Champaign (UIUC).

NCSA Audible Collage for Microsoft Windows is the IBM PC version of the collaborative visualization program NCSA Collage. NCSA Audible Collage helps researchers interpret data through the use of sound from a MIDI (Musical Instrument Digital Interface) device. A researcher can sonically overlay data variables, avoiding confusion with the visual image. (Sonification features are currently available only in the PC version of NCSA Collage.)

NCSA HDF-UCD (Unstructured Cell Data) is a high-level library that sits on top of NCSA HDF. Designed to handle structured and unstructured finite element grids, it also contains an efficient surface extraction algorithm.

NCSA Collage (Collaborative Analysis and Graphics Environment) combines communication and scientific visualization tools that allow researchers to conduct real-time, remote collaborations. Multiple users can connect with each other over the network and share scientific data in the form of raster images, palettes, animations, and scientific datasets. Sessions can share data visualization, electronic whiteboard capabilities, text display and editing, and screen capture functions.

NCSA Height-Color Visualizer for the Silicon Graphics4D Series Workstation provides a noninteractive batch utility that lets researchers easily create 3D animation from time-dependent data.

NCSA Import2HDF is a software tool that converts non-HDF files (FITS, TIFF, GIF, PICT, raw raster images, and ASCII) to HDF files so that researchers can use NCSA software or other tools that work with HDF format files. This program displays the contents of HDF files, RISs in the form of color raster images, and SDSs in spreadsheet fashion.

NCSA Isosurface Visualizer for the Silicon Graphics4D Series Workstation provides a noninteractive batch utility with which researchers can easily create and view animations of isosurfaces derived from 3D data.
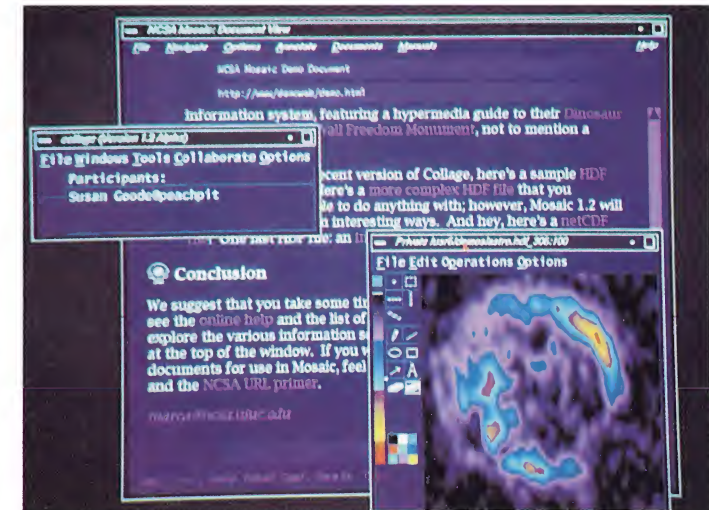

NCSA MinMaxer contains optimal 2D triangulation methods, including the Delaunay triangulation, and is used for grid generation.

NCSA PC Show is a color-imaging application that displays 8-bit binary scientific data as color raster images on an IBM PC or compatible with an EGA, VGA, or Revolution 512 graphics card from Number Nine Computer (N09). NCSA PC Show can read raw raster, raw palette, or HDF files; map images according to a user-defined palette; display multiple images in sequence; animate multiple image sequences; and rotate the palette.

NCSA PolyView for the Silicon Graphics4D Series Workstation is an interactive 2D and 3D rendering tool for polygonal data. NCSA PolyView is especially suited for unstructured grids and requires files formated with NCSA HDF for input.

*T*he Software Development Group is committed to providing universally accessible resources that facilitate and expedite discoveries in science.

other software that works with NCSA HDF format files. This program contains code for two distinct versions of Reformat—a command-line driven program (Reformat) and a graphical interface-driven program (XReformat).

NCSA X DataSlice for the X Window System provides a color imaging and data analysis tool for 3D, 32-bit floating-point scientific data stored in the NCSA HDF file format.

NCSA X Image for the X Window System provides a color-imaging and interactive data analysis tool for looking at 2D datasets, manipulating palettes, and animating multiple images. Features include a notebook (for describing and applying functions to datasets and making a distributed computing connection), color contour plots, profile plots, and both pixel and interpolated expansion of images.

NCSA-developed Tools for Communication

NCSA Collage supports synchronous collaborative work (see previous description).

NCSA Mosaic supports asynchronous collaborative work. This environment lets researchers browse networks, retrieve and exchange information, and conduct long-term collaborative work sessions. NCSA Mosaic's environment blends hypermedia and search functions for information discovery and retrieval

NCSA Shape2D is a weighted, 2D alpha-shapes visualization tool. NCSA MinMaxer and NCSA Shape2D were developed jointly by SDG and the UIUC Department of Computer Science.

NCSA Reformat/XReformat for UNIX-based Systems, similar to NCSA Import2HDF, converts certain non-HDF files (FITS, TIFF, GIF, Sun, raw raster images, and X Window System dumps) to HDF files. Once converted, the files can be processed with NCSA-developed software or

over wide-area networks. Based on standard formats and protocols to encourage widespread use and applicability, NCSA Mosaic speaks several TCP/IP network protocols and is fluent in NCSA's DTM protocol. This allows direct communication with NCSA Collage as well as other DTM sources (e.g., running simulations and remote instruments) and clients. NCSA Mosaic clients are available for Macintosh, Microsoft Windows, and X Window System platforms.

NCSA Telnet for the Macintosh is an ARPA standard telnet program with built-in standard FTP, configurable network parameters, Tektronix graphics emulation, interactive color raster graphics support, and VT100 emulation for multiple sessions. Newer features include an FTP client and the ability to use aliases for making connections. NCSA Telnet supports MacTCP, so it can be used simultaneously with other MacTCP products.

NCSA Telnet for Microsoft Windows provides much of the same functionality within the Microsoft Windows environment, including multiple VT100 sessions and a built-in FTP server.

NCSA Telnet for the PC provides ARPA standard telnet functionality with full color support, built-in standard FTP, configurable network parameters, Tektronix graphics emulation, and VT100 emulation for multiple sessions.

NCSA-developed Tools for Data Sharing and Transfer

Data Transfer Mechanism (DTM) is a message-passing library. It is designed to simplify the task of interprocess communication and to facilitate the creation of sophisticated distributed applications in a heterogeneous computing environment. To accomplish this, DTM provides a method of interconnecting applications at run time and reliable message passing, complete with synchronization and transparent data conversion. DTM has been optimized for large messages (100 Kbytes and up), but is also efficient for smaller messages. DTM is available on most platforms that support Berkeley sockets library, including Cray systems, CONVEX, TMC CM-2 and CM-5, most UNIX-based workstations, Macintoshes, and MS-DOS machines.

NCSA HDF (Hierarchical Data Format) allows researchers to manage scientific data with a data storage and retrieval system for scientific data and metadata on a range of hardware platforms. HDF-supported data and metadata include multidimensional gridded data, raster images, polygonal mesh data, multivariate datasets, finite-element data, non-Cartesian coordinate data, and text. HDF supports a hierarchical grouping structure called Vset that lets researchers organize data objects within HDF files to fit their views of how the objects go together. Most NCSA scientific visualization software, as well as many other applications, support HDF files.

How to Obtain NCSA-developed Software

NCSA software may be obtained via CD-ROM, anonymous FTP, an archive server, or U.S. mail.

The NCSA Digital Gallery CD-ROM

This CD-ROM contains NCSA-developed scientific software for several computer systems (Macintosh, IBM PC, Sun, SGI, and X Window System) plus nearly 400 megabytes of images and animation sequences produced by researchers around the world using these software products to study a wide range of scientific areas. Documentation for each software product is provided in either Microsoft Word or ASCII format. The NCSA Digital Gallery CD-ROM may be purchased through the *NCSA Technical Resources Catalog* (see U.S. Mail below).

Anonymous FTP

If you are connected to the Internet you can download software, documentation, and source code (if available) at no charge from an anonymous file transfer protocol (FTP) server at NCSA. The procedures are presented below. If you have any questions regarding the connection or procedure, consult your local system administrator or network expert.

1. Log on to a host at your site that is connected to the Internet and is running software supporting the FTP command.
2. Invoke FTP by entering the Internet address of the server: `ftp ftp.ncsa.uiuc.edu`
3. Log on using `anonymous` for the name.
4. Enter your local login name and address (for example, smith@ncsa.uiuc.edu) for the password.
5. Enter `get README.FIRST` to transfer the instructions file (ASCII) to your local host.
6. Enter `quit` to exit FTP and return to your local host.
7. Review the README.FIRST file for complete instructions concerning the organization of the FTP directories and the files you need to download for each application.

Archive Server

1. Email a request to: `archive-server@ftp.ncsa.uiuc.edu`
2. Include the word `help` in the subject or message line.
3. Press RETURN.
4. Send another email request to: `archive-server@ftp.ncsa.uiuc.edu`
5. Include the word `index` in the subject or message line.
6. Press RETURN.

U.S. Mail

NCSA software and manuals are also available for purchase through the *NCSA Technical Resources Catalog*—many are sold individually, some are available only on anonymous FTP reel or cartridge tapes. Orders can be processed only if accompanied by a check in U.S. dollars made out to the University of Illinois. To obtain a catalog, contact:

NCSA Documentation Orders
152 Computing Applications Building
605 East Springfield Avenue
Champaign, IL 61820-5518
(217) 244-4130
orders@ncsa.uiuc.edu

NCSA-supported Applications

Utility	IBM PC	Mac	SGI	Other UNIX
NCSA ALVIS			•	•
NCSA Audible Collage	•			
NCSA Collage	•	•	•	•
NCSA HDF	•	•	•	•
NCSA HDF-UCD			•	•
NCSA Height-Color Visualizer			•	
NCSA Import2HDF		•		
NCSA Isosurface Visualizer			•	
NCSA MinMaxer			•	•
NCSA Mosaic	•	•	•	•
NCSA PC Show	•			
NCSA PolyView			•	
NCSA Reformat/ XReformat			•	•
NCSA Shape2D			•	•
NCSA Telnet	•	•		
NCSA X DataSlice			•	•
NCSA X Image			•	•

This is NCSA

Today some of the country's most talented scientists, engineers, social scientists, artists, humanists, and educators come to the National Center for Supercomputing Applications (NCSA) to work in and experience its unique interdisciplinary environment. This distinctive atmosphere, along with the encouragement and assistance of over 180 staff members, inspires and nourishes creativity and imagination.

Established in February 1985 with a National Science Foundation (NSF) grant, NCSA opened to the national research community in January 1986. The state of Illinois, the University of Illinois at Urbana-Champaign (UIUC), corporate partners, and other federal agencies supply additional funding. The center has provided high-performance computing and communications (HPCC) resources for over 6,000 users at more than 380 universities and corporations.

These users are involved in computational science—an approach to scientific research that complements the traditional methods of laboratory experimentation and theoretical investigation. Computational science allows researchers to recreate numerically, or simulate, natural phenomena on a high-performance computer. In many cases, such phenomena cannot be investigated in the laboratory or fully evaluated theoretically due to such constraints as safety, cost, speed, or time. In other cases, the supercomputer is being coupled directly to laboratory instruments or observational facilities for extensive data computing.

The federal High Performance Computing and Communications Program, as well as the communications technology initiatives envisioned by the Clinton-Gore administration, support the creation of a permanent national infrastructure for high-performance computing. NCSA's direction reflects these national initiatives.

National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign

NCSA is committed to providing its diverse constituencies with production and experimental high-performance computing and communications resources.

The MetaCenter and the HPCC Environment

This decade will witness a fundamental change in the way researchers use high-performance computing systems. First, scientists will move away from stand-alone systems to a networked set of computers controlled at the desktop by distributed software. Second, networks will provide users with an increased variety of computing architectures. Anticipating these developments, the NSF supercomputing centers have formed the nucleus of a collaboration based on the concept of a *national metacenter for computational science and engineering*—a coalescence of intellectual and physical resources unlimited by geographical constraint, a synthesis of individual centers that will create a new resource greater than the sum of its parts.

The goals of the MetaCenter are to give scientists and engineers the ability to move their problems directly to appropriate computer architectures without regard for where the computers are located; to develop a national file system that gives researchers direct access to their files regardless of where they are located; and to design a common user interface that allows researchers to use the same commands on all systems at all centers.

Such diversity of resources will give researchers access to the best available technologies and intellectual resources through a single portal, a path to the computer resources and architectures of four NSF centers (NCSA, Cornell Theory Center, Pittsburgh Supercomputing Center, and the San Diego Supercomputer Center). Through that portal, scientists and engineers will be able to concentrate on solving their research problems at an accelerated pace.

No one metacenter can provide nationwide resources to support new technologies for economic growth. The National Consortium for High Performance Computing (NCHPC) was formed by the Advanced Research Projects Agency (ARPA) and NSF to respond to this need. Actually a consortium of consortia, five consortia were chosen by ARPA to comprise the NCHPC. The goal of the NCHPC is to advance massively parallel processing (MPP) systems and technology so that scalable computing is available nationwide. K–12 schools, and eventually communities, will join in over an expanded Internet. Interdisciplinary research partnerships are anticipated among industry, government, and academe.

Grand Challenge Research

Computational science is now being applied to Grand Challenge research—problems with potentially broad economic, political, or scientific impact that can be advanced by applying HPCC resources. Center researchers are at the forefront of addressing these problems in such areas as computational fluid dynamics, global change, semiconductor design, materials science, biomedical imaging, molecular dynamics, earth science, astrophysics, and relativity. Grand Challenge projects generate increasingly complex data, require realistic simulations of the processes under study, and demand intricate visualization of the results. This translates into a need for large amounts of machine resources, memory, disk space, and access to metacenter architectures across which the problems may be efficiently distributed.

NCSA's Supercomputing Systems

NCSA's plan for meeting the computational requirements of its users is constantly re-evaluated in response to advances in technology as well as changes in federal funding policy. NCSA is phasing out traditional vector processor platforms and moving to scalable, shared memory platforms constructed from microprocessors. Scalable computers are modular, upgradeable, and binary compatible from desktop to supercomputer, making them a flexible alternative to traditional architectures. These scalable platforms have an installed user base numbered in the hundreds of thousands with a rich software environment familiar to users on their local workstations. NCSA is deploying new scalable machines from Silicon Graphics Inc. (SGI) and Convex Computer Corporation.

At present, NCSA operates and maintains four production supercomputers:

- Connection Machine Model 5 (CM-5) with 512 nodes, 16 Gbytes memory, 120 Gbytes parallel disk (SDA)
- CONVEX C3880 (C3) with 8 processors, 4 Gbytes memory, 60 Gbytes disk space
- CONVEX Exemplar with 8 processors and 512 Mbytes of shared memory, with a planned upgrade to a 64-processor, 8-Gbyte system with 80 Gbytes RAID disk
- SGI Challenge with 32 processors, 2 Gbytes memory, 80 Gbytes RAID disk, with a planned upgrade to a 16-processor TFP Power Challenge

Mass storage at NCSA consists of UniTree running on a CONVEX C220 system with 100 Gbytes of disk and magnetic tape storage. Large datasets (10+ Gbytes) are handled with DD-2 helical scan tape technology capable of storing 25-165 Gbytes per tape cartridge at a peak transfer rate of 15 Mbytes per second.

The Andrew File System (AFS) has been implemented at NCSA. AFS is a distributed file system with a common name space. Data are stored in volumes in UNIX-like file structures on AFS file server machines and accessed through a cache manager on AFS client machines.

The computational environment at NCSA supports UNIX, Macintosh OS, and MS-DOS operating systems.

NCSA provides users with access to the center's facilities through a local-area network (LAN) using Ethernet, Fiber Distributed Data Interface (FDDI), and High Performance Parallel Interface (HIPPI) and through wide-area networks via the worldwide Internet. NCSA is a node on the 45 Mbps NSFNET backbone network.

Network Research

The center is actively involved in preparing the research and development groundwork necessary to implement and support the HPCC Program's National Research and Education Network (NREN). The center is currently participating in one of several national Gbps network testbeds coordinated by the Corpora-

tion for National Research Initiatives (CNRI) and funded by NSF, CNRI, and ARPA. The testbed—called BLANCA—includes NCSA; the universities of Illinois (Urbana-Champaign), California (Berkeley), and Wisconsin (Madison); AT&T Bell Laboratories; Ameritech; Pacific Bell; Bell Atlantic; Astronautics Corporation of America; and Lawrence Berkeley Laboratories. BLANCA is the only transcontinental testbed and includes the first very high-speed trunk (622 Mbps trunk between UIUC/NCSA and the University of Wisconsin). Links to Berkeley and Bell Labs run at least 45 Mbps, with higher rates to be determined as the experiment proceeds. Initial applications include atmospheric and climate modeling, radio astronomy imaging, and biomedical imaging as well as digital library and data management facilities. One of NCSA's priorities is to examine distributed collaboration tools that allow scientists to work together in real time on the network.

NCSA's strategy involves a three-tiered local area network architecture: Ethernet for low-end desktop machines, FDDI for medium-range systems, and HIPPI for high-performance computers. Installing prototypes and beta testing hardware technologies—such as FDDI in 1990 and HIPPI in 1991—are a pivotal part of these integration efforts. New local networking research has recently included asynchronous transfer mode (ATM) and shared-memory LANs. The center is participating in the definition of these emerging standards and is developing software to couple network hardware technologies with operating systems and user interfaces.

NCSA and the Microelectronics Center of North Carolina (MCNC) have proposed the creation of a packet-based audio and video teleconferencing solution integrated into a remote collaboration environment. This effort is specifically targeted to allow workstation video conferencing within NCSA Collage, a synchronous collaboration tool.

Applications Research

One of the most significant achievements of the NSF supercomputer centers program in the 1980s was to expand the number of researchers able to attack ambitious computational problems. In the 1990s, the centers are working together to meet the challenges presented by software development for advanced HPCC architectures. At NCSA the primary focus is on software and libraries for scalable architectures exemplified by NCSA's CM-5.

NCSA's strategy for advancing computational science on scalable architectures is to foster interdisciplinary research collaborations (see Programs and Services below). Every effort is made to include computer scientists in these collaborations so that the necessary interaction and feedback with computational scientists can occur.

NCSA recognizes that scientists need to collaborate across space transparently. They should be able to exchange and discuss simulations, text, graphics, animation sequences, and data analyses with colleagues across the network in real-time as easily as if they were in the same room. NCSA Collage software, developed by NCSA's Software Develop-

ment Group (SDG), is the first step toward accomplishing this. Workstation software developed by SDG is available via anonymous file transfer protocol (FTP) at ftp.ncsa.uiuc.edu.

Visualization and Virtual Environments

NCSA brings a state-of-the-art visualization environment to researchers. Distributed visualization capabilities allow remote users to link their desktop machines with high-performance, large-memory visualization servers. For example, AVS software is supported on a variety of platforms that make it feasible to distribute visualization subtasks to the appropriate machine. NCSA staff develop special software to take advantage of the high-performance environment, including tools to support volume rendering, remote medical instrumentation, and visualization of the massively large datasets often generated by high-performance computing.

As grand challenge teams develop, their members are becoming both users and co-architects of the emerging national metacomputer.

NCSA's visualization program rapidly adopts new technologies and develops new techniques in the service of computational science. Virtual Reality (VR) is the latest step in visualization technology. VR surrounds the user with a synthetic environment that emulates reality. NCSA's VR Laboratory provides a place where researchers can explore their data while experimenting with the latest equipment. NCSA has developed applications in product and drug design, manufacturing, cosmology, and education.

The newest addition to the NCSA VR Laboratory is the CAVE, a collaborative project between NCSA and the University of Illinois at Chicago's Electronic Visualization Laboratory (EVL). In the CAVE, a virtual environment is displayed on multiple walls of a room using rear-projection monitors. Because of the large space, several people can explore an environment together. The CAVE offers high-resolution and full-color imagery—significant advances over earlier VR installations. Full audio capabilities are also provided. The CAVE is powered by a Silicon Graphics Onyx supercomputer equipped with multiple RealityEngine2 graphics systems.

A High Definition Technologies (HDT) environment is being developed where scientists can compute and view HDT visualizations. This HDT environment is the visual complement to high-performance computers and large bandwidth networks. By merging the traditionally separate video and computer envi-

ronments, NCSA provides an opportunity for scientists to investigate a variety of emerging standards and their applications to Grand Challenge research.

NCSA's visualization facilities also include the Numerical Laboratory in the UIUC's Beckman Institute for Advanced Science and Technology. The laboratory provides access to a variety of high-end machines from Silicon Graphics Inc., Apple Computer Inc., IBM, and Sun Microsystems Inc. Included are an 8-processor Silicon Graphics Power Vision 360 VGX, capable of rendering over 1 million polygons per second, and SGI Indigo Elan workstations.

The Renaissance Experimental Laboratory (REL), also in the Beckman Institute, provides a place for teaching visualization skills. The lab is equipped with Silicon Graphics workstations and is open to faculty who incorporate computer graphics into their courses. Courses have been offered in art and design, biophysics, chemistry, computer science, and mathematics.

Cyberspace

The latest frontier technology is cyberspace—the web of interconnected information servers worldwide. This space offers a proliferation of information; so much information, in fact, that one of the challenges is how to organize this information to make it easy to find.

SDG has developed NCSA Mosaic, a software tool that responds to this challenge. This tool uses a blend of hypermedia, hierarchical organization, and search functionality for information discovery and retrieval. NCSA Mosaic is designed to be fully distributed across a wide-area network—a user of the system need never know the physical origin of any document.

Cyberspace is reached through information servers—the electronic publishing equivalent of the printing press. One of the most common systems used on the Internet is Gopher, which provides a hierarchical tree of documents. A more advanced system is the World Wide Web, which provides true hypertext—a click of the mouse on a highlighted word pulls up a related document on a different computer continents away.

NCSA currently runs both a Gopher server and a World Wide Web server. The former contains a range of NCSA documentation and MetaCenter information. The latter contains such information as recent issues of *access* magazine, descriptions of SDG software, and general NCSA information.

Programs and Services

NCSA's Interdisciplinary Research Center (IRC) enables the transfer of HPCC technology from the Meta-Center to computer users at universities, businesses, government agencies, and schools. The cornerstone of the center's approach is a reciprocal process: NCSA provides software, instruction, and information to the national communities; in turn, NCSA receives direction

and technical requirements from scientists, engineers, educators, and vendors.

NCSA staff foster collaborative research by mutidisciplinary, multi-institutional teams of computer scientists, research scientists, and engineers; postdoctoral research associates; and graduate students from the national and international community. These teams forge new approaches to previously insoluble research problems, develop community codes, and host workshops and seminars to transfer technology.

The IRC offers programs to support technology transfer:

- NCSA's Visitors Program provides working visits for faculty, industrial researchers, postdoctoral associates, educators, and students. While immersing themselves in the NCSA environment, the visitors develop new approaches to their research problems, spreading their newfound experiences with scalable systems to their home institutions. Visitors also enjoy the expertise of the many Beckman Institute and UIUC researchers associated with NCSA.

- The HPCC Partners Program—a collaboration between NCSA and computing and communications vendors—aims to advance the state of technology through grants, donations, personnel exchanges, joint research projects, and sharing of advanced technology.

NCSA's Industrial Program is a major research, service, and training endeavor that encourages corporate and academic researchers to work together on new problems of joint interest, thereby fostering technology transfer. The Industrial Program offers flexible partnerships based on computational experience and corporate needs. A unique strength of this program is the opportunity it offers corporations to actively participate in center and campus activities. Networked to NCSA, these partners are able to access the center's facilities from corporate locations, as well as from their offices at the center.

As an integral part of its mission, NCSA promotes science and mathematics education through a variety of activities aimed at developing and evaluating approaches to incorporate HPCC in the classroom. Kindergarten through graduate school students, teachers, and educators are offered supercomputing time and instruction. NCSA strives to change how students see, do, and feel about science and math by applying visualizations, simulations, and computational modeling tools in teaching.

The National Consulting Office helps users take full advantage of NCSA facilities. Its staff is well versed in supercomputing tools and techniques. Software and technology experts throughout the center provide additional support. Users may contact the Consulting Office in person or by phone at (217) 244-1144 weekdays from 8:30 a.m. to 6 p.m. central time or by electronic mail at consult@ncsa.uiuc.edu.

NCSA publicizes its efforts in scientific and technological areas through its *access* magazine and public information materials; both are available in print and through NCSA Mosaic. Users are given startup information and are kept informed of all available online and printed technical documentation.

NCSA provides training sessions on introductory and advanced topics, including the use of supercomputers, various applications and methodologies, and visualization techniques. Classes are held at the center or at other locations based on specialized needs of the host site.

The MetaCenter has implemented a videoconferencing architecture, tying together the NSF-sponsored centers, NSF, MCNC, the NSF Science and Technology Centers, and additional sites via commercial carriers. This system facilitates group meetings and training activities among the centers as well as individual staff meetings. Workstation-based videoconferencing is also being explored, especially concerning the use of collaborative research software.

To learn more about the programs, services, and facilities at the center, please call Jim Bottum at (217) 244-0072 or send electronic mail to bottum@ncsa.uiuc.edu.

Contacts Directory

General Information and Programs

Allocations
Radha Nandkumar
(217) 244-0650
allocations@ncsa.uiuc.edu

Vicki Halberstadt
(217) 244-5709
allocations@ncsa.uiuc.edu

Applications Group/Faculty Program
Melanie Loots
(217) 244-2921
mloots@ncsa.uiuc.edu

Chemistry User Group
Balaji Veeraraghaven
(217) 333-2754
balajiv@ncsa.uiuc.edu

Community Outreach Coordinator
Scott Lathrop
(217) 244-1099
scott@ncsa.uiuc.edu

Computing and Communications
Charles Catlett
(217) 333-1163
catlett@ncsa.uiuc.edu

Consulting Office
(217) 244-1144
8:30 a.m.–5:00 p.m. CST
consult@ncsa.uiuc.edu

Education and Outreach Program
John Ziebarth
(217) 244-1961
ziebarth@ncsa.uiuc.edu

Industrial Program
John Stevenson
(217) 244-0474

Media Relations
John Melchi
(217) 244-3049
jmelchi@ncsa.uiuc.edu

Media Services
Tony Baylis
(217) 244-5354
tbaylis@ncsa.uiuc.edu
media@ncsa.uiuc.edu (services)

NCSA Operations
Sue Lewis
(217) 244-0708
slewis@ncsa.uiuc.edu
(217) 244-0710 (services/help)

NCSA Receptionist
(217) 244-0072
FAX (217) 244-1987

NCSA Security Officer
Michael Smith
(217) 244-7714
msmith@ncsa.uiuc.edu

Networking
network@ncsa.uiuc.edu

Orders for Publications, NCSA Software, and Multimedia
Debbie Shirley
(217) 244-4130
orders@ncsa.uiuc.edu

Publications Group
Melissa Johnson
(217) 244-0645
melissaj@ncsa.uiuc.edu

Software Development Group
Jae Allen (information)
(217) 244-3473
jallen@ncsa.uiuc.edu

Terry McLaren (technical support)
(217) 244-3473
tmclaren@ncsa.uiuc.edu

Training Program
Alan Craig (information)
(217) 244-1988
acraig@ncsa.uiuc.edu

Deanna Spivey (registration)
(217) 244-1996
deannas@ncsa.uiuc.edu

User Services
Lex Lane
(217) 244-0642
lex@ncsa.uiuc.edu

Visitors Program
Jean Soliday
(217) 244-1972
jsoliday@ncsa.uiuc.edu

NOTE: All electronic mail addresses are via Internet.

National Center for Supercomputing Applications
University of Illinois at Urbana-Champaign

SGML Queries

The parser's transparent integration with OmniMark gives developers unique abilities. If desired, an OmniMark program can be completely aware of the SGML implications of the processing being performed, and ensure SGML conformance in the process. Integrated SGML parsing technology also reduces the effort required to perform SGML transformations. The SGML parser is always aware of the current state of the document structure, and can provide feedback to the programmer on context and surrounding elements. The presence of the parser also frees the developer from having to know anything about the syntactic details of SGML when processing SGML documents. Use of advanced features such as tag minimization is transparent to the OmniMark program developer, who need only know element names.

World Class Service, Support and R&D

OmniMark is backed with world-class service and support including free full-production demos, SGML and OmniMark training courses, and unlimited telephone support. Exoterica's commitment to heavy investment in ongoing research and development will ensure that OmniMark remains relevant in the future. Through regular feedback from OmniMark programmers worldwide we are continuing to identify and add capabilities to meet your expanding needs.

Need Help?

Exoterica's consulting group can help you get off the ground. By getting you started or delivering a complete system, our experienced consultants can have you up and running in record time.

But it doesn't stop with us. We can point you to a rapidly growing worldwide community of independent OmniMark consultants who specialize in particular information technology areas. We'll be happy to put you in touch.

Pricing and Platforms

OmniMark is available for a broad range of computing platforms including Intel PCs, Macintosh, most popular UNIX systems, and VAX/VMS. OmniMark programs you develop on one platform work unchanged on all other supported platforms (but a license is required on each system).

Exoterica's current price list contains complete details on supported platforms, pricing, and licensing methods.

EXOTERICA

Americas/Asia
Exoterica Corporation
1545 Carling Avenue, Suite 404
Ottawa, Ontario, CANADA K1Z 8P9
Tel: +1 613 722 1700
USA/Canada: (800) 565-9465
Fax: +1 613 722 5706
Internet: info@exoterica.com

Europe
Exoterica Corporation
149 Quai de Stalingrad
ISSY-LES-MOULINEAUX 92137 CEDEX
FRANCE
Tel: +33 1 45 29 18 00
Fax: +33 1 45 29 22 77
Internet: exoterica.eur@utopia.fnet.fr

Product Information

Exoterica's flagship product, OmniMark, is used daily by many of the world's largest Information Technology companies. They know its flexibility and power are unmatched. Here's why.

What is OmniMark?

OmniMark™ is a specialized programming language for analyzing, modeling and processing text-intensive information. It enables the translation of rich text information to and from SGML and its application across a wide range of software products, viewing and browsing technologies and composition systems.

Powerful

OmniMark is a powerful, intuitive and expressive programming language that can be applied to a wide variety of rich text computing and data-manipulation tasks. Increasingly, it is an indispensable part of any SGML programmer's product suite. OmniMark has a variety of features targeted at the unique problems encountered in electronic document composition, analysis and conversion. In particular, the pattern matching and text manipulation capabilities of OmniMark substantially outperform dedicated C programs. OmniMark is the most effective product available for a wide range of SGML and non-SGML applications.

Flexible

OmniMark is ideal for analyzing obscure data formats, rearranging data, performing computations, extracting information for reports, performing quality assurance, and for loading text and relational databases. OmniMark lets you work the way you want to work. It offers programmers concise and intuitive new ways to do jobs that are difficult to accomplish with the C language or other utilities.

Robust

OmniMark is engineered to automate the processing and manipulation of massive quantities of text. OmniMark is already in

Features

- ✓ A scalable, easy-to-learn, rule-based language which supports novices and expert professional programmers alike
- ✓ A built-in high speed validating SGML parser with graceful recovery from markup errors: no external SGML parser is needed
- ✓ A powerful context-sensitive SGML query language
- ✓ Resolve cross-references into multiple buffers and files; create hypertext links and indexes automatically
- ✓ Advanced function text and binary pattern recognition and analysis in an intuitive easy-to-read language
- ✓ Powerful string manipulation features
- ✓ Powerful stream I/O mechanism which supports binary and text files
- ✓ Associative arrays support powerful data structures
- ✓ Supports translations based on simple and complex pattern analysis and context-sensitive SGML parsing
- ✓ Standard recursive pattern matching on strings and files
- ✓ Standard arithmetic operators (plus, minus, times, divide, modulo)
- ✓ Binary arithmetic operators (mask, union, difference, complement, shift)
- ✓ String arithmetic operators (radix, base, binary)
- ✓ Optional declaration of global and local variables
- ✓ Optional "strong typing" of variables.
- ✓ Structured programming control structures (if-then-else, do, do-while, case)
- ✓ Rule grouping: switch banks of rules on and off dynamically to modularize programs
- ✓ Powerful string and file buffering features
- ✓ Automatic context saving for nested structures
- ✓ Flexible string formatting features
- ✓ Powerful macro feature which supports multiple arguments and can make programs more compact and readable

OMNIMARK

constant use in such critical path situations as extracting data from wire services and preparing data for multi-gigabyte on-line retrieval systems.

▶▶ Total Life Cycle Involvement

OmniMark is more than just conversion software. It is the only SGML product you need to successfully produce sophisticated multimedia and online applications.

OmniMark can analyze legacy data from any source and help you develop the SGML DTD while leading you to the trouble spots and anomalies in the data.

OmniMark can automatically mark and tag data in an iterative stepped refinement process that ensures rich, accurate detail in the final result. OmniMark lets you be in control.

OmniMark can establish a complex network of hypertext links, far beyond what can be economically accomplished by hand, and validate the links automatically. You get a more compelling product for less money.

OmniMark, used in conjunction with sound SGML practices, can completely automate the final production process. This includes establishment of formatting instructions, browse sequences, linking, and contextual search constraints which can be automatically interpreted by popular composition or viewing software.

OmniMark is a proven production solution for a wide range of viewing technologies including Apple HyperCard, Asymetrix Toolbook, EBT DynaText, Folio Views, InfoAccess Guide, Microsoft Multimedia Viewer, and NTergaid HyperWriter, to name a few.

▶▶ History of Success

OmniMark's unequalled SGML processing capabilities have been the key to the success of a rapidly expanding world of rich text, CD-ROM and on-line applications including multimedia, legal, dictionary, and reference work publication, and technical documentation systems including CALS.

Time and time again, organizations that use OmniMark are successful. Chances are, the next time you use an electronic book, from almanacs to dictionaries, from encyclopedias to the most popular multimedia movie guide, you are looking at an OmniMark success story.

OmniMark has the advantage of depth as well as breadth. Before OmniMark, problems became obstacles that could be negotiated successfully only at great expense — or not at all. OmniMark effectively breaks through barriers in all text-based environments. Spend your time fighting problems, not tools.

▶▶ Automate, Automate!

Increasingly, programmers and knowledge workers spend a lot of time doing mechanical text-intensive data manipulation— changing the format of data, checking its validity, searching for items with certain properties, performing quality assurance tests, printing reports, and many other similar kinds of tasks. All of these jobs ought to be automated, but it's a real nuisance to have to write a special-purpose program in a language like C or Pascal every time a requirement changes.

OmniMark is a special-purpose programming language that makes it possible to handle such tasks with easy to write, easy to test programs which can be very useful regardless of whether SGML is involved. In effect, OmniMark is a programming language that can replace a host of general tools.

▶▶ Event Driven Programming

In simple terms, an OmniMark program looks for the occurrence of prescribed events, associated with the input data stream, and produces certain actions depending on the results of processing the input data. For example, an OmniMark program might search a set of files for data matched by any prescribed pattern; when matching data is found, the appropriate set of actions is performed.

```
element city
    output "@CityStyle:%sc, "

element areacode
    set buffer areacode to "%sc"
    do scan "%g(areacode)"
        when attribute zone of
            ancestor country is equal "north-am"
        match ["2" to "9"] ["01"] [digit except "1"]
            ; area code is ok
        else
            put log "Invalid: %g(areacode) code."
    done
```

OmniMark's intuitive, english-like language provides content checking beyond that available from SGML parsers alone.

▶▶ Rapid Prototyping

OmniMark programs are usually much smaller than equivalent programs in a general programming language. Programs a line or two long can be composed at the keyboard, run once, then discarded or expanded to meet new requirements.

The same brevity of expression and convenience of operations make OmniMark valuable for prototyping larger programs. One can start with a few lines, then refine the program until it does the desired job, experimenting with designs by trying alternatives quickly. And although it isn't unusual for OmniMark programs to grow to thousands of lines, it's easy to get started, and easy to start over when experience suggests a different direction.

Even though it is easy to get started with OmniMark, you won't get to a point you can't get past. OmniMark doesn't quit before the project is finished.

▶▶ OmniMark and SGML

OmniMark is leading the way in making SGML the *de facto* standard for on-line and CD-ROM development. A key component of OmniMark is our SGML parser, the fastest and most conformant commercially supported implementation of SGML. It provides uniquely powerful fault tolerance and error recovery features.

▶▶ Validate SGML

The function of an SGML parser is to recognize and validate the structure and syntax of SGML documents. The purpose of the validation function is to verify that a document's content is organized and stored in conformance with the document's DTD. When SGML errors are encountered, OmniMark's parser provides informative helpful error messages.

WORD PROCESSOR
DATABASE SYSTEM
TYPESETTING ENVIRONMENT

OmniMark

SGML KNOWLEDGE BASE

OmniMark

PUBLISHING SYSTEM
DATABASE & ONLINE
CD-ROM VIEWER

Form continues from other side.

My organization plans to purchase the following types of computers:

(Select all that apply)

Microcomputer/PCs
01_____80386
02_____80486
03_____AST
04_____Compaq
05_____Dell
06_____Macintosh
07_____NEC
08_____Pentium/586
09_____Other (specify)_____
10_____None of the above

Workstations
11_____Acer/Altos
12_____CompuAdd
13_____Data General
14_____DEC
15_____HP 9000
16_____IBM RS/6000
17_____Intergraph
18_____NeXT
19_____Olivetti
20_____PowerPC
21_____PS/2
22_____Silicon Graphics
23_____Solbourne
24_____Sun
25_____Tatung
26_____Other (specify)_____

Midsize Server
27_____Data General
28_____DEC
29_____HP 1000
30_____HP 3000
31_____HP 9000/8X
32_____IBM AS/400
33_____IBM RS/6000
34_____Motorola
35_____NCR
36_____Sequoia
37_____Stratus
38_____Sun
39_____Unisys
40_____Other(specify) _____

Large Server
41_____Amdahl
42_____HDS
43_____Pyramid
44_____Sequent
45_____Tandem
46_____Other (specify)_____

Mainframe/Supercomputers
47_____Bull
48_____Cray
49_____DEC VAX 9000
50_____Hitachi
51_____IBM/PCM (370)
52_____NCR
53_____Unisys
54_____Other (specify)_____

I influence the purchase of the following:

(Select all that apply)

A_____Software
B_____Terminal
C_____Workstations
D_____Micros/PCs
E_____Minis
F_____Mainframes
G_____Disk drives
H_____Tape drives
I_____Printers/plotters/scanners
J_____Communication Devices
K_____Training/education services
L_____Multimedia
M_____Memory
N_____Storage devices
O_____Training/education services
P_____Other (specify)_____

These purchases cost:
A_____Under $25,000
B_____$25,000-$100,000
C_____Over $100,000

Number of employees at location:
A_____1-50
B_____51-500
C_____501-5000
D_____More than 5000

Please fold here. Seal with tape and mail to UniForum at address below.

FREE SUBSCRIPTION OFFER - See below

UniForum The International Association of Open Systems Professionals

UniNews

October 12, 1994
The Biweekly
Newsletter
For UniForum
Members
Volume VIII,
Number 17

FREE SUBSCRIPTION AND MOUSE PAD!

Bonus for
WORLD-WIDE WEB
CONFERENCE ATTENDEES

Get a FREE subscription to UniForum Monthly and UniNews. . .

plus a free mouse pad if you act now!

We're not kidding!

RETURN THE CARD TODAY. See inside for details ☞

We're Not Kidding!

Here's a unique offer for open system professionals, like you, who recently attended UNIX Expo.

We want to start a six month FREE SUBSCRIPTION for you to *UniForum Monthly* - the magazine written for open systems professionals, by open systems professionals and published by UniForum, the association of open systems professionals. In fact you're holding your first issue in your hands.

We also want to send you *UniNews*, twice a month. This informative newsletter will keep you up to date on important industry events and developments that can effect your company, and your career.

But there's more, too. **Act today - return the card on the back cover in the next 10 days and we'll send you a colorful UniForum Mouse Pad just to say thanks!**

What's Going On Here?

Each month *UniForum Monthly* carries thought provoking features and columns from some of the best industry writers around. The kind of reporting and analysis that will help you do your job better. In every issue you'll read in-depth case studies about companies that are implementing open systems solutions right now. The lessons you can learn - and apply to your own use - make *UniForum Monthly* a resource you'll want to have.

We want you to have this information as we believe you'll want to know more about The UniForum Association - the world's largest independent association of open systems professionals. There is *absolutely no obligation* on your part, just the opportunity to get a little bit smarter about the open systems industry.

Fill out and return the card today before we come to our senses! Free deals are rare these days, grab this one while you can.

UniForum
The International Association of Open Systems Professionals

The Road to Open Systems

UniForum 1994

UniForum The International Association of Open Systems Professionals

UniNews

Second World-Wide Web Conference Fills to Capacity
Four-day Chicago meeting will draw 1,000 for 'Mosaic and the Web'

An overflow attendance of more than 1,000 is expected at the Second International World-Wide Web Conference, "Mosaic and the Web," in Chicago Oct. 17-20. Originally limited to 600, registrations were sold out eight weeks following the July announcement of the event.

Attendance is expected to triple the number who went to the first international Web conference, held last May at CERN, the European Particle Physics Laboratory, near Geneva, Switzerland. Interest in the conference especially surprised the organizers since advertising was done only on the Web itself and through contributed distributions by sponsors and supporters. More than 10,000 interested persons accessed the Open Software Foundation (OSF) server listing Web conference registration information.

Sponsors include the National Center for Supercomputing Applications (NCSA) at the University of Illinois at Urbana-Champaign, the OSF Research Institute, the National Science Foundation, and CERN, in cooperation with UniForum.

Keynote speaker is Dr. Larry Smarr, director NCSA. Smarr will speak on "Mosaic and the Future of the National Information Infrastructure."

Day one of the conference, Oct. 17, is devoted to eight tutorials, covering an introduction to Mosaic, advanced features of Mosaic, basic and advanced topics in Hypertext Markup Language (HTML), management of content on WWW servers, systems administration of Web servers, system administration of WWW servers, media issues, and security.

Days two and three, Oct. 18 and 19, will feature two days of conference sessions, including presentation of 178 papers in 50 panel sessions, each with between one and seven papers. Both technology and user experiences are included. For example:

• Building HTML Application Systems: Converting Existing MS-Windows Applications to HTML.
• Providing Data on the Web: from Examples to Programs.
• A Webmaster's Starter Kit
• An Architecture for Scholarly Electronic Publishing on the World-Wide Web.
• The Media Business on the WWW
• Organizing Information in Mosaic: a Classroom Experiment

Day four, Oct. 20, will include sessions for Mosaic and Web software developers.

The range of attendees is expected to be wider than at last May's conference. But according to Ira Goldstein, conference co-chairman and director of research and advanced development at the OSF Research Institute, it's primarily for people who know the Internet. "It is for end users who have heard of the Internet and maybe have started looking at the Web, who really want to get a perspective on what's happening in other user sites, what the technology suppliers are doing and, maybe, what's coming down the road next," says Goldstein.

Future conferences are being planned by the International World-Wide Web Conference Committee, at the rate of two per year. One is planned for next spring at the Fraunhofer Center for Research in Computer Graphics in Darmstadt, Germany. The site for the Fall 1995 conference has not been decided.

Complete information about the upcoming conference can be viewed on the Web. The access URL is:

http://riww.osf.org:8001/ri/announce-ments/WWW_Conf_F94.html ◆

The Web and How it Works

WWW: *World Wide Web. Also spelled WorldWide Web, World-Wide Web or World-Wide-Web. Also called Web or W3. Developed in 1989 by Tim Berners-Lee with sponsorship by the European Particle Physics Laboratory, known as CERN. The server program used with the client program, called Mosaic, to form a web, which is used to distribute information on the Internet. For example, an electronic paper about the history of UNIX found on one host may have hypertext references to an electronic paper about MULTICS found on another host half-way around the world. All of this information is linked together to form an intricate web of information.*
—from UNIX: An Open Systems Dictionary

From the user's perspective, the World-Wide Web is a non-linear sequence of pages, containing a combination of text and graphics. Any page may contain hypertext—highlighted or underlined words or phrases that the user can click on for more information, sending him or her deeper and deeper into the Web's structure.

To link to the Web, the user needs only an Internet connection and an interface tool such as Mosaic, acting as a client. Using an address line called a Uniform Resource Locator (URL), the user can link immediately to an organization's home page. From there, the possible links are infinite, giving the user

---

NEWS FLASH! **UniForum '95 Keynote**

UniForum is pleased to announce that Bruce Tognazzini, distinguished engineer in the strategic technology office at SunSoft, will deliver the keynote address at UniForum '95 on Thursday, March 16. "Tog" joined SunSoft two years ago after 14 years at Apple Computer, where he was human interface evangelist, and is best known for his breakthrough work on the Apple II and human interface efforts. Among his current projects is a movie, *Star Fire*, that predicts the future of the Workstation in the Year 2005.

Tim Berners-Lee: The One Who Started It All
World-Wide Web creator looks to the future

The World-Wide Web was created in 1989 by Tim Berners-Lee, an Oxford University graduate with background in real-time communications and text processing software development. Berners-Lee devised the basis of the Web as a method of enabling the people working on a project to easily link to each other's information.

Berners-Lee was a consultant at CERN, the European Particle Physics Laboratory, near Geneva, Switzerland, when he found himself having to figure out the relationships and interdependencies between people, programs, and pieces of hardware at CERN. Working from an earlier program he had written, he made a hypertext file system—one designed to lay a new image over the old one every time the user found a new avenue to explore and clicked on a given word or phrase. "CERN had very complex structures," Berners-Lee says. "I needed the power of hypertext and the unrestrained ability to represent any relationship. This sort of information about what is going on and what is related to what is very important when you're running a project."

Berners-Lee quickly began to see his program as helpful to any group of people working together. "The idea was, it would be a common knowledge base. First, I wanted the people to be able to get all my information. But more importantly, I wanted to know if they already had used it. There was a requirement that if two groups started using information independently, and they started to develop links between the groups, that they could represent these links. The knowledge bases could then fuse without having to make any big changes or to copy databases."

After producing a Web prototype, Berners-Lee acquired a Next computer, which had just become available. He then wrote a WWW browser and a WYSIWYG editor, both within a two-month period in late 1990. He credits the NextStep environment with the short completion time. Development time for a browser on the X Window System subsequently took a year, he notes.

Berners-Lee and a CERN collaborator, Robert Cailliau, continued to develop the Web working in the Next environment, sharing their ideas using hypertext as they worked. As the work went on, Berners-Lee realized the huge potential of the application. But he also realized that people might never use it if they couldn't find some useful information to link to in the first place. So he and Cailliau began putting in gateways to various databases, each with its own server giving it a continuous link to the Web. "When you say 'I designed a neat hypertext system,' you put yourself in the ranks of thousands of people who have done it before," Berners-Lee says. "You can explain to people until you're blue in the face that it's neat, but it won't wash because nobody will read it if there is no data in there. And nobody will put any data in there because there's nobody reading it. The way we found of bootstrapping it was by finding servers which provided entry to existing databases. That was a really powerful way of getting large amounts of data on-line very quickly."

But the expansion brought Berners-Lee some regret. In turning the Web into largely an information resource, he also found that the Web's collaborative element got lost. Those who maintain the servers must now be in the loop when changes are made, instead of allowing users to make immediate hypertext links in others' work. "It lost its immediacy, and that's something I'm pushing for very much—to bring the collaborative aspect of it back and to make it more interactive, so that the knowledge can be fresher," Berners-Lee says.

The bootstrap worked. The Web expanded slowly at first, then caught fire in 1993. The availability of the Mosaic interface tool, as well as companies and organizations eager to get their information servers in on the ground floor, led to an explosion in traffic. World-Wide Web volume on the National Science Foundation network backbone (NSFnet) service rose from about 5,000 megabytes in March 1993 to 750,000 megabytes in March 1994. During a nine-month period ending last March, the number of Web servers increased from about 130 to more than 1,260.

Although he still has ties with CERN, Berners-Lee moved earlier this year to Cambridge, MA. There, based at the Massachusetts Institute of Technology he directs the W3 Organization (W3O), an industry consortium now being formed, whose job it will be to coordinate the evolution of standards and contributions. Although not a standards body itself, the W3O will be able to work closely with such bodies during the development of standards. A list of founding members is expected in about a month, after contracts are signed. "We already have a long list of people who have been extremely excited about joining and pushing us to get going," Berners-Lee says.

An objective that Berners-Lee hopes the W3O will help with is insuring that the Web standards remain both open and common to all servers. He worries that "We could have companies going out and making sub-worlds of the Web, which only work with their software and which are not accessible by other people. I feel that would be a major step backward and I'm sure it's something the world doesn't need." Technological changes will be another item on the W3O's agenda, changes that include a more efficient protocol and adding of computer-readable semantics. In the future, Berners-Lee expects that computers will do a lot of the Web browsing that people now do themselves. However, that's possible only if the computer can read an object and determine exactly what it's looking at.

Sometime in the future, Berners-Lee speculates, "Maybe we'll have so much stuff [data] that it goes underground, where much of the stuff becomes automatically operated on, not designed for presentation to humans. Maybe there's only a certain amount of stuff they can handle." ◆

CommerceNet Helps Companies Feel Out the Web
Fifty companies test the waters with self-help trade group

The increasing number of commercial ventures on the Internet has quickly spread to the World-Wide Web, where companies are beginning to offer catalogs, directories, indexes and other information. Although business transactions via the Web are not widespread, the infrastructure is evolving to make that possible.

Commercialization has been welcomed by Web and Mosaic developers, including Tim Berners-Lee, creator of the Web, and Joseph Hardin, associate director for software development at the National Center for Supercomputing Applications (NCSA) at the University of Illinois. "We've been arguing for the last two years that the Internet needs to become more commercial, in the sense of being open to commercial activity," Hardin says, "because that's the only way that we're really going to get the bandwidth we want for the kinds of things we want to do. People need to be able to make money over the Net."

One of the highest-profile commercial endeavors has been CommerceNet, the first large-scale market trial of electronic commerce on the Internet, which was formally launched last April. With headquarters in Santa Clara, CA, CommerceNet is a non-profit, federally funded organization with a three-year mission to help bootstrap the use of the Web by private companies. CommerceNet now has about 50 member companies, most of whom have established servers on the Web containing information available to the public.

The CommerceNet Web server is accessible at the URL **http://www.commerce.net/**, where a directory of member companies is available.

CommerceNet's sponsoring organizations are BARRNet, the San Francisco Bay Area Internet access provider; Enterprise Integration Technologies of Palo Alto, CA, research and development consultants; and Stanford University's Center for Information Technology.

The organizations felt Web commerce needed encouragement, says Allan Schiffman, chief technical officer of EIT. "When technology's in its early days, nobody's very confident about taking on a new market alone," Schiffman says. "But then when things start to become clearer, people start feeling their competitive oats. Big companies now are still feeling their way out on the Internet. What their role is they're not completely sure. They want that to be part of the future, but how they play in the marketplace they're not certain."

Company activities are largely limited to information now because of the lack of security on the Web and the Internet generally. "What you're seeing CommerceNet members do at this stage is bringing their product catalogs on-line and making their sales literature available," Schiffman says. They even have some facilities for doing ordering." Transactions are another problem. "If you can't do the most fundamental things, namely keep your information confidential between a buyer and a seller, and if you can't identify who the seller is or who the buyer is, then it's really hard to get very far," Schiffman says. "Big companies are really worried. When they send stuff to a customer, they want to make sure that people can't steal it and can't see what they're doing. They want to be able to supply confidential information and they want to be able to get paid for things. They can't do any of these things without some kind of a security mechanism."

A step toward allowing secure transactions was taken in August with the beta release of a secure version of NCSA Mosaic, the most popular Web interface. The software was developed as a cooperative effort among EIT, RSA Data Security, and NCSA. Full release of the product is expected sometime in October to CommerceNet members, and a general NCSA distribution will come later, Schiffman said. Secure Mosaic is supposed to make secure commercial transactions possible.

The secure version of NCSA Mosaic allows users to affix digital signatures that cannot be repudiated, as well as time stamps to contracts so that they become legally binding and auditable. In addition, sensitive information such as credit card numbers and bid amounts can be securely exchanged under encryption.

In the meantime, CommerceNet is helping its member companies establish their servers, get up on the Web and make their product information available. Some members are starting to work together on cooperative products to do more things, like electronic data interchange over the Internet and building multivendor catalogs.

Now midway through the first year of its three-year run, CommerceNet is still feeling the waters. "If it eventually becomes the case that doing business on the Internet is considered routine, then the need for this kind of self-help trade association might be less necessary," Schiffman says. "That's certainly not the case now." ◆

The Web and How it Works

(continued from page 1)

access not only to that organization's server, but to other servers as well, where hypertext links are present.

Organizations and companies with Web servers have obtained a permanent Internet connection, allowing clients to access their home pages via the organization's own URL (usually beginning **http://**). Their servers contain documents coded with HyperText Markup Language (HTML), a tagging code whose functions include designating the links to other pages. The Web's standard protocol is Hypertext Transport Protocol (HTTP). Organizations maintain their own documents and servers and update the information as needed. Some documents may have restricted access, requiring a password, or allow the user to send information back to the server using a form.

As an abstract information space, the Web is independent of all networks and administration is completely decentralized. "Anybody can publish, anybody can make links, and so to a certain extent, it's self-organizing or self-disorganizing," says Web creator Tim Berners-Lee. "The saving grace is that you make extra links only if you think they're worth reinforcing. Only rich and useful areas become reference points.

"The Web is a reflection of society and of all the relationships between people," Berners-Lee notes. "People refer to other people and they quote each other's work, and they quote what each other said yesterday on the bus, and they point people at places and things. If you want to find something out, you probably go through a sequence of documents and people, to try and find things. Within that basically amorphous space of people and pieces of paper, we make structures. Librarians impose their own structure to try to help people find books, for example. The idea of the Web is to be able to map into whatever structures we need for our society." ◆

NCSA's Joseph Hardin Heads Conference Planning
Mosaic developer has high hopes for the future

If the World-Wide Web is the Emerald City of the '90s, Mosaic is the Yellow Brick Road. Mosaic was created by software developers at the National Center for Supercomputing Applications (NCSA) at the University of Illinois as an access tool for Web servers and other sources of information on the Internet. NCSA's original project lead for Mosaic was Joseph Hardin, associate director for software development at NCSA and co-chairman of the Second International World-Wide Web Conference, "Mosaic and the Web."

The conference, the first in the United States to be devoted to Web development, is the result of months of planning by Hardin, co-chairman Ira Goldstein of the Open Software Foundation Research Institute, and other members of the International WWW Conference Committee. The huge growth interest in the Web in 1993 led to discussions between NCSA and CERN, the European Laboratory for Particle Physics where the Web originated, last winter, about holding a conference. As a result, two were scheduled: the first at CERN, in Switzerland, last May, and the second in Chicago Oct. 17-20. The conference committee also was formed to schedule and plan future conferences and to provide continuity.

"We've been working closely with CERN, since the technologies and our interests are so closely tuned," Hardin says. "We're both scientific research institutions and associated with the academic and research engineering communities, so it was just a natural liaison."

The Chicago conference, which has been sold out for weeks, will include "everything of interest to users and organizations interested in learning about the Web, from both a public and a commercial perspective," Hardin says. "We think it's going to be fascinating."

Because of the high interest in the conference, the organizers are looking into, but not promising, the use of remote sites with conference sessions multicast to faraway locations via the MBone, the Internet Protocol Multicast Backbone,' or via phone lines. The MBone was developed by the Internet Engineering Task Force for experiments with broadcasting messages.

"This is a global phenomenon," Hardin says. "One of the things we want to do is allow a lot of people to participate without necessarily having to come here. Since we

Joseph Hardin

are constantly pushing the boundaries of what can be done over the Net, the next natural step is for us to have remote sites participate in an increasingly interactive way. Broadcast via satellite is a future possibility, Hardin believes. "Using existing teleconferencing technology is too expensive and it violates the ethic of the Net anyway. The goal, and the joy of the Net, is to collapse that kind of thing down into the Web. So we're trying out different ideas. It's just a dream right now."

Development of Mosaic
Mosaic development began after David Thompson, an NCSA developer, introduced the Web to NCSA's software development group. Subsequently, Marc Andreessen and Eric Bina of NCSA developed most of the application. "When we looked at the World-Wide Web and the first browsers, we all looked at each other—Marc and Dave and I—and said, 'Wow, we can do a lot with this. Let's try it out and see what happens.'"

The objective was to build a graphical interface to the Web, having it apply across the Macintosh Windows, X Window and Microsoft Windows platforms. All three platforms are now supported. A forms functionality was added, allowing clients to interact with information on the Web by sending information back to the server. "The idea was for this to be a collaborative technology more than just an information delivery technology," Hardin says. "There's a real danger of passivity on the Net—of people staying too much in the traditional perspective of

broadcast technology, where you sit passively in front of a television. The fun is when you actually get involved."

Now NCSA has licensed Mosaic technology to several companies, hoping to see the technology develop commercially even as the public version is developed in parallel and remains available free of charge. In addition, NCSA has assigned future commercial licensing rights for Mosaic to Spyglass, Savoy, IL. "We were looking for some company that essentially shared our vision that the public version of the software was not competition for a commercial version, but rather helped create markets for it," Hardin says. "It's a good thing that a public version remain viable and vibrant and continue to be developed. Spyglass will work with all the existing licensees and license to people like IBM and DEC. That frees us from this incredible amount of effort that was directed toward trying to figure out what a commercial approach should be."

Hardin sees encouragement in the move of Encyclopedia Britannica to make its product available on Mosaic and the Web. "They have confidence that millions of people will be able to read their material because the standards of HTML and Mosaic are there. Otherwise, they would be at the mercy of providers like Prodigy and America OnLine and CompuServe, whom they originally approached. This is an example of an opening that I think will lead to diversity, proliferation, and creativity—and in the future, as we look back in history, one of the revolutionary characteristics of the next generation of global communications."

*NCSA Mosaic for the X Window System is available at **ftp.ncsa.uiuc.edu** in /Mosaic. Both source code and binaries (for Sun, SGI, IBM RS/6000, DEC Alpha OSF/1, DEC ULTRIX, and HP-UX) are available. You don't need to have Motif installed on your system to run NCSA Mosaic if you pick up a precompiled binary. However, you do need Motif 1.1 to compile Mosaic.*

*For NCSA Mosaic for X support, send e-mail to **mosaic-x@ncsa.uiuc.edu**. Macintosh-related questions should be directed to **mosaic-mac@ncsa.uiuc.edu** and Microsoft Windows questions to **mosaic-win@ncsa.uiuc.edu**.* ◆

Web Technology's a Paradigm Shift, Says Co-Chairman
OSF's Goldstein calls the Web 'a revolution in access technology'

Leaders in the Open Software Foundation (OSF) decided a year ago that the World-Wide Web would be an important area of computing, and one that OSF members would need to know about. As a result, the OSF Research Institute decided to co-sponsor the upcoming Second International World-Wide Web Conference in Chicago. Ira Goldstein, director of research and advanced development, is co-chairman of the conference.

"We concluded that Web activity was extremely significant to how computing would be used in the rest of the '90s and beyond," Goldstein says. "So we made commitments in our research program, in the kind of information we will try to communicate to the OSF membership, and commitments to facilitate this whole area."

The technology surrounding the Web, Goldstein believes, is "a technology that's going to be absolutely ubiquitous." After talks between OSF and representatives from CERN, the European Laboratory for Particle Physics, where the Web was developed, and the National Center for Supercomputing Applications (NCSA) at the University of Illinois, where Mosaic was developed, it was the sense of the group that both technology and usage are evolving so quickly that conferences need to be held frequently.

Since the first conference was held in May, "A variety of companies have reached the marketplace with their enhanced Mosaic," Bernstein says. "I think we'll see improved tools for placing material on the Web. I think the biggest thing is the continued usage of the Web, so there will be lots of talks by people who are using it in government and industry, the arts, and education. I think from the growing body of experience, we'll distill what's working well, what needs attention, and what might be less important."

As an experiment, OSF decided to advertise mostly via the Web and through UniForum, which is supporting the conference. The conference sold out in eight weeks. Based on that experience, "I would say that the Web is a pretty significant kind of communication medium," Bernstein says. "A couple of years from now, you really would reach a very significant percentage of the knowledge worker community."

What interests Bernstein most about Mosaic and the Web is that "This is all becoming a real paradigm shift in how we publish and communicate knowledge. There was the paradigm shift of the printing press that moved us from an aural culture to a very small number of written manuscripts, very expensive and only for the very special, to a culture in which the printed book was common. With the Web, what's happening is, first of all, the diversity of knowledge, the multimedia character of it is broader and richer, and your ability to publish it individually is enormously increased.

"So when I think what it would be like to be a knowledge worker in the 21st century, I think living in this Web, providing knowledge into it, browsing for knowledge, is really going to become a pretty profound part of one's professional life. I think this is not a revolution in computing. It's a revolution in a fundamental technology for society—in the fundamental information access technology." ◆

Standardization Effort Surrounds Web Language
HTML specification to be published by IETF

An ongoing effort to standardize the World-Wide Web's text-tagging language is now centered on the Internet Engineering Task Force (IETF). The standardization effort concerns the HyperText Markup Language (HTML), the standard set of tags used to prepare text for use on the Web, which is used by anyone preparing documents to be placed in a Web server and accessed by Web clients.

HTML is one of a set of languages comprised by Standard Generalized Markup Language (SGML), the data encoding system that allows information in documents to be shared—either by other document publishing systems or by applications performing electronic delivery and other functions. SGML is a vendor-neutral, formal international standard, as HTML is intended to be.

Tim Berners-Lee originated HTML as he was developing the Web. Since then, a leader in standardization efforts for HTML has been Dan Connolly, software engineer for Hal Software Systems, Austin, TX. As the Web became popular, Connolly discovered HTML on USENET, the worldwide UNIX bulletin board, where Berners-Lee had posted it. "I tried to come up with a document type definition—an SGML specification, basically—and it fit like a square peg in a round hole," Connolly says. Eventually, others wanted to know what was legal and what was not using HTML, so Connolly continued his standardization effort, trying to retrofit HTML to SGML. A draft HTML specification was published for the first time on the Internet in June 1993.

As items such as footnotes, superscript, and text centering needed to be added to the basic tags for headers, lists, paragraphs and links, a new specification called HTML Plus came into use, beginning in Bristol, England. Since March of this year, Connolly and Web and Mosaic developers, have been trying to produce a published specification through a working group of the IETF. "We kept spinning iterations and trying to get it just right," Connolly says. The specification is scheduled to come out as an Internet request for comment (RFC) from the IETF late this year. Once the specification is published, the IETF will be responsible for approving new features. "We're doing the best we can just to get this standard, that describes what's been going on for months, out the door," Connolly says.

Developing the standard has been a job of delicately balancing actual coding practices with the SGML standard itself. Both are important. "What people are doing is important so that when somebody goes to implement a new browser, he'll be able to use all the data that's out there," Connolly says. "The other side is to try to use SGML because it's an ISO [International Standards Organization] standard and there are benefits to be gained. For example, since we used SGML, the documents can use readers for blind people where you can take an SGML document and feed it in and it will read to them."

A separate IETF working group, the Uniform Resource Identifier Working Group, is attempting to standardize the Uniform Resource Locator (URL) address lines that point to documents on Web servers. ◆

Recruitment and Help Wanted

POSITIONS WANTED

Vice President/Director of Marketing
Seeks position as vice president/director of marketing for a dynamic, growing computer software or hardware company—a position utilizing expertise in sales and marketing, a strong technical background and familiarity with a broad range of markets, sales channels and business applications.

Experience: 19+ years experience, 11 years of marketing and sales of PC UNIX and DOS software and hardware products. Director of marketing for communications and networking manufacturer. Strong technical background. Former director of software development.

Personal: M.B.A., University of Rochester; prefer Santa Cruz or Santa Clara counties, CA; salary base plus commission/bonus $125-150K; available immediately.

Rick Freedman, (408) 438-5873.

Marketing/Technical Marketing
Seeks position in marketing/technical marketing for progressive high-tech company with vision and determination to succeed.

Experience: 10+ years in marketing, design, technical writing, public relations, product/project management, and trade show management.

Personal: B.S., computer science, University of California; prefer Chapel Hill-Raleigh-Durham, NC area; salary $45-65K; available in December.

Dave Schroeder, e-mail 73543.2035@compuserve.com; fax (919) 938-8862; voice (919) 933-0333.

System Integration/Technical Sales Support
Seeks position in system integration and support or technical sales support.

Experience: 17 years technical experience including hardware, software, operating systems, and networking; 8 years UNIX support and development; 4 years PC-UNIX LAN/WAN integration and development support. Strongest language: C. Informix on-line background. AT&T/GIS SVR4 SMP systems, Sun/Solaris, UTS, and DOS/Windows. TCP/IP, StarLAN, LanManager (UNIX-based). Significant understanding and experience with the AT&T Toll Network.

Personal: B.S., EE technology, DeVry, Kansas City; M.S., computer science, Illinois Institute of Technology; prefer Kansas City metropolitan area; salary $55K to $70K; available two weeks after accepting offer.

Curtis K. Olinger, AT&T BCS SMTS, 1100 Walnut, RM 812-3, Kansas City, MO 64106; cko@devildog.utsd.att.com; (816) 654-4047; fax (816) 654-2076. ◆

Receive *UniNews* On-line

UniNews is now available electronically, as a service to UniForum members, in lieu of your printed copy.

To receive the newsletter in your electronic mailbox, send e-mail to **pubs@uniforum.org** and include in the message body your UniForum membership number, your name and your e-mail address.

You will receive *UniNews* in ASCII text.

UniNews Recruitment

For inclusion in the *UniNews* Classified Section, please provide the following information, being as specific as possible. **If you do not want your name printed, please indicate in item No. 1 and UniNews will receive replies and forward them to you.** Please type or write legibly. Your classified may be edited for length or clarity. *UniNews* "Positions Wanted" classifieds are available FREE OF CHARGE to UniForum members only. Upon receipt of your material, we will publish your classified in the next TWO available issues of the biweekly *UniNews*.

YOU MUST BE A MEMBER OF UNIFORUM TO PARTICIPATE

1 Your name _____

Shall we print your name in *UniNews*? _____

Your UniForum Membership # (if available) _____

_____

2 Where Hiring Companies May Reach You (include phone, fax and e-mail) _____

3 Title and Description of the Job You Want _____

4 Geographical Preference _____

5 Professional Experience and Qualifications_____

6 Highest Grade or Degree Achieved, and Where:
_____

7 Salary Range $_____

8 Availability _____

You may mail or fax this form to: Sandy Parker, UniForum, 2901 Tasman Drive, Suite 205, Santa Clara, CA 95054; (408) 986-1645. GOOD LUCK!

UniForum

2901 Tasman Drive, Suite 205
Santa Clara, CA 95054

Tell Colleagues About UniForum Membership

☎ ☎ ☎

**Phone (800) 255-5620
or (408) 986-8840
for more information**

UniForum

Richard H. Jaross, *Executive Director*
Richard R. Shippee, *Director of Publications*
Jeffrey Bartlett, *Executive Editor*
Cheri DeBusk, *Manager of Art and Production*
Mary Margaret Peterson, *Managing Editor*
Don Dugdale, *Senior Editor*
Leslie Zimmerman, *Production Assistant*

The International Association of Open Systems Professionals

Publications, Conferences, Discounts and More
UniForum Member Benefits

For more than 12 years, UniForum has provided its members with the best open systems information and services. Purchased separately these benefits cost more than twice the price of general membership. Don't be left out—if you're not already a general member, can you really afford NOT to be?

General Member benefits:
($125 per year)

- *UniForum Monthly* magazine
- The biweekly *UniNews* newsletter
- The annual *Open Systems Products Directory*
- UniForum technical publications including:
  - *Electronic Mail Demystified*
  - *Your Guide to POSIX*
  - *POSIX Explored: Shell & Utilities*

- *Network Substrata*
- *Network Applications*
- *MOSES Whitepapers*

- Free "Positions Wanted" ads in *UniNews*
- Discount rates worldwide from Hotel Reservations Network
- Discount on Avis car rentals
- Discounts on hardware and software from UniForum corporate sponsors
- Discounts on products, training, and publications from more than 15 companies.

Call today for a complete list.
1-800-255-5620
or 1-408-986-8840
(outside US and Canada)

Comments or questions:
contact Susan J. Hoffmann,
Membership Services manager

Benefit of the Month

Solve all your UNIX training needs with one phone call. Order low-cost, high-content "Training On Video™" today! Train your entire staff for the price of one off-site seminar. With more than 11 years teaching and writing about UNIX, instructor and noted expert Ray Swartz brings UNIX down to earth and right onto your video screen. Special discounts available for UniForum members. For more information or to order contact Berkeley Decision Systems, 803 Pine St., Santa Cruz, CA 95062. Phone (800) 408-8649. Fax (408) 458-2721.

Don't miss out - Subscribe today!

UniForum Monthly Subscription/Application Form *(please print)*

☐ **YES!** I Want to join UniForum as a six month Trial Member (no obligation or cost, at all) and receive *UniForum Monthly, UniNews,* and my FREE Mouse Pad. (Mouse Pad offer expires 11/30/95)

Signature _____ Date _____

Name _____ MR/MS
     First   *Last*   *Initial*

Title _____

Company _____

Address _____

Mail Stop _____

City _____ State _____ Zip _____

Country _____ Telephone _____

E-mail *(where we will send UniNews)* _____

Fax _____

To allow us to process your subscription, please complete the following section. Don't forget to sign the form, fold and seal with tape. Mail it to us today!

World-Wide Web Conference Offer expires June 30, 1995

Please fill out the section below

My Organization's primary business: (Select only one)	My Job function:	Operating systems used:	
A_____ Government	A_____ Executive Mgmt.	01_____ AIX	21_____ NT
B_____ Finance	B_____ Finance/Admin. Mgmt.	02_____ A/UX	22_____ OSF
C_____ Manufacturing (non-computer)	C_____ Sales and Marketing	03_____ BSD UNIX	23_____ OS/2
D_____ Utilities	D_____ Systems/Operations Mgmt.	04_____ DG/UX	24_____ OS/400
E_____ R&D Laboratories	E_____ Systems/Operations Staff	05_____ DOS/VSE	25_____ RTU
F_____ Educational Associations	F_____ Technical Management	06_____ DYNIX	26_____ SCO Open Desktop
G_____ Wholesale/Distribution	G_____ Technical Staff	07_____ ESIX	27_____ SCO UNIX/Xenix
H_____ Retail/Hospitality	H_____ Consultant	08_____ FlexOS	28_____ SunOS/Solaris
I_____ Medical Research/Health Care	I_____ Educator	09_____ HP/UX	29_____ System 7 (Macintosh)
J_____ Professional Services	J_____ Publisher	10_____ Interactive UNIX	
K_____ Transportation	K_____ Other (specify)	11_____ IRIX	30_____ ULTRIX
L_____ Petroleum		12_____ LINUX	31_____ UNICOS
M_____ Agriculture		13_____ Macintosh	32_____ UNIXWare
N_____ Reseller (computer)	**Years using UNIX:**	14_____ MPE/XL/IX	33_____ UNIX System V4.2
O_____ Software Developer	A_____ None	15_____ MS-DOS	34_____ VM
P_____ Manufacturer (computer)	B_____ 1 to 5 years	16_____ MS-Windows	35_____ VMS/Open VMS
Q_____ Systems Integrator	C_____ More than 5 years	17_____ MS-Windows NT	36_____ Other (specify)
R_____ Other (specify)		18_____ MVS (ESA/XA/SP)	
		19_____ NeXT Step	
		20_____ Netware	

UniForum
The International Association of Open Systems Professionals