SLAC-PUB-7757
February 1998

The World Wide Web and High Energy Physics

Bebo White

Stanford Linear Accelerator Center
Stanford University
Stanford, California 94309

Submitted to *Annual Review of Nuclear Particle Sciences*

Stanford Linear Accelerator Center, Stanford University, Stanford, CA 94309

Work supported by Department of Energy contract DE–AC03–76SF00515.

CONTENTS

ABSTRACT

High energy physics and the World Wide Web (WWW) share a rich history. The Web, developed at CERN as a collaboration tool and quickly adopted by the Internet community, has become a communications phenomenon. This article reviews early WWW development and its basic technology. I also summarize some significant applications of Web technology past and present. Prospects for future use are discussed.

1. INTRODUCTION

The conferences on Computing in High Energy Physics (CHEP) offer a regular occasion for people active in the field to meet, to discuss progress and evaluate technologies, and to make projections for the future. The goal of the conference series has been to determine how advanced computing and networking can effectively enable high energy physics. CHEP92 held in September 1992 in the Alpine setting of Annecy, France was no exception. The technical program included papers in areas such as "Triggering and Data Acquisition," "QCD Calculations, Architecture and Experience," "Computer-Aided Detector Design," and "Mass Storage Technology."

Yet, at the conference wrap-up session, Terry Schalk of the Santa Cruz Institute for Particle Physics (SCIPP) indicated that a significant conference highlight for him was the discussion of a networking technology recently developed at CERN auspiciously named the World Wide Web or more simply WWW (1). CHEP92 was held before the introduction of Mosaic by the National Center for Supercomputing Applications (NCSA) at the University of Illinois which took the Internet world by storm in February 1993. It was certainly before WWW became the household word and have the ubiquitous presence that it has today. At the time of CHEP92, the number of Websites internationally was probably less than fifty, many of which were at high energy physics institutions. The first Website in the United

States at SLAC (the Stanford Linear Accelerator Center) had only recently "come on-line."

It may be difficult to think of the 1990-1993 time period as "early days," but the phenomena that is the World Wide Web was developed during that so recent period. Since that time, the HEP community has played a vital role in the development and creative usage of the Web and has helped to shape what has become known as the "killer application of the Internet." The demands of increasingly complex HEP experiments will continue to push the capabilities of the Web and provide imaginative applications usable not only by HEP but by the entire Internet community.

2. THE CERN WORLD WIDE WEB PROJECT

The background of WWW in high energy physics illustrates the need for collaboration in HEP. Most of the features of the early Web servers and browsers were developed to address that need. The Web's beginnings were far simpler than the enormous technological and commercial presence into which it has recently evolved. Many of today's dedicated "Web-surfers" would hardly recognize some of the early implementations of the technology.

An evaluation of the needs of the HEP community at CERN and its experimental collaborators and a review of current technological trends such as FTP and Gopher led Tim Berners-Lee and Robert Cailliau of the CERN Computing Division to write a project proposal to their management in March 1989. This proposal was entitled, quite simply, "World Wide Web: Proposal for a Hypertext Project"(2).

In the preface to their proposal, Berners-Lee and Cailliau write,

> "The current incompatibilities of the platforms and tools make it impossible to access existing information through a common interface, leading to waste of time, frustration and obsolete answers to simple data lookup. There is a potential large benefit from the integration of a variety of systems in a way which allows a user to follow links pointing from one piece of information to another one. This forming of a web of information nodes rather than a hierarchical tree or ordered list is the basic concept behind Hypertext.
>
> The project [WWW] will aim:
>
> - to provide a common (simple) protocol for requesting human readable information stored at a remote system, using networks;
> - to provide a protocol within which information can automatically be exchanged in a format common to the supplier and the consumer;
> - to provide some method of reading at least text (if not graphics) using a large proportion of the computer screens in use at CERN;

- to provide and maintain at least one collection of documents, into which users may (but are not bound to) put their documents. This collection will include much existing data;
- to provide a keyword search option, in addition to navigation by following references, using any new or existing indexes. The result of a keyword search is simply a hypertext document consisting of a list of references to nodes which match the keywords;
- to allow private individually managed collections of documents to be linked to those in other collections;
- to use public domain software wherever possible, or interface to proprietary systems which already exist;
- to provide the software for the above free of charge to anyone."

In many ways, the timing of this proposal was ideal. CERN had become the largest Internet site in Europe. In his "A Short History of Internet Protocols at CERN," Ben Segal recalls "...an entire culture had developed at CERN around 'distributed computing', and Tim [Berners-Lee] had himself contributed in the area of Remote Procedure Call (RPC), thereby mastering several of the tools that he had needed to synthesize the Web such as software portability techniques and network and socket programming. " Segal also contends "...the Web could have emerged considerably earlier if CERN had been connected earlier to the Internet" (3).

In October of 1990, Berners-Lee and Cailliau submitted a revised proposal. By November, the initial implementation of a Web client for the NeXT platform had been developed. This first browser was able to display documents using multiple fonts and styles and was even able to edit documents (capabilities only now becoming widespread in many browsers) , but access was limited to NeXT users.

March 1991 brought the development of the CERN "linemode" client/browser (Figure 1.), which was executable on numerous platforms but which displayed output only on character-based (VT100) terminals. The WWW server code was subsequently ported to the CERN VAX/VMS platform and the IBM VM/CMS mainframe. In December 1991, the CERN Computer Newsletter announced and described WWW to the HEP community.

```
                                                  Welcome to CERN
CERN Information

  CERN is the European Particle Physics Laboratory in Geneva, Switzerland.
  Select by number information here, or elsewhere (Return for more).

 Help [1]                 On this program, or the World-Wide Web project [2].

 Phone book [3]           People, phone numbers, accounts and email
                           addresses. See also the analytical Yellow Pages [4],
                           or the same index in French: Pages Jaunes [5]

CC Documentation[6]       Index of computer centre documentation,
                           newsletters, news, help files, etc...

News [7]                  A complete list of all public CERN news groups,
                           such as news from the CERN User's Office [8],
                           Cern computer center news [9], student news [10].
                           See also Private groups [11] and Internet news [12]

From other sites:-
1-24, <RETURN> for more, Quit, or Help:
```

Figure 1. The First CERN Linemode Browser

The early browsers developed at CERN were quickly followed by the first browsers designed for X-Windows - the Viola browser developed by Pei Wei at the University of California, Berkeley, and the Midas browser developed by Tony Johnson at the Stanford Linear Accelerator Center (SLAC). Both of these browsers pre-dates the more well known Mosaic browser at NCSA and were freely available to Internet users. Features provided in the Viola and Midas browsers appear to have had an impact on the Mosaic development process.

3. WWW CONCEPTS

The World Wide Web was designed to incorporate two fundamental concepts: hypertext documents and network-based information retrieval.

3.1 *Hypertext*

In 1965, Ted Nelson devised the term "hypertext" to describe non-sequential or non-linear text. It was defined as "a body of written or pictorial material interconnected in a complex way that it could not be conveniently represented on paper" (4). Nelson's intention was to define a mechanism for presenting information which more closely resembled the human thought process which involves free association rather than the linear limitations which he felt that the written word imposes.

While there are classic examples of non-sequential text applications, computer and networking technologies made it possible to realistically implement Nelson's vision of a " document universe" or "docuverse." An essential feature of computer-based hypertext is the concept of computer-supported links (both within and between documents). It is this linking capability which supports the non-sequential organization and presentation of text so fundamental to the hypertext paradigm.

Simple documents become hypertext documents when they contain words or phrases act as links to other documents. Typically hypertext documents are presented to a user with the text that can act as a link highlighted in some manner, and the user is able to access the linked document by selecting the highlighted area. The manner by which the link is selected is dependent upon the Web browser implementation. Some browsers have text interfaces (e.g., the CERN linemode browser and the Lynx browser from the University of Kansas) where links are specified via menus or highlighted text. Other browsers such as Midas, Mosaic, the Netscape Navigator and the Microsoft Internet Explorer support graphical user interfaces (GUIs) which allow a user to point to a link (usually with a mouse) and activate the transfer to the linked document (usually with a click).

HTML Hypertext pages and documents are defined within the WWW environment using the HyperText Markup Language (HTML). HTML is an implementation of SGML, the Standard Generalized Markup Language, a text-formatting system widely used in the HEP community. Quite simply, HTML is a small collection of document markup tags which can be used to define hypertext documents which are portable from one computing platform to another. HTML contains elements which are characteristic of descriptive and referential markup. These markup tags permit a page or document author to specify how information will be presented and which page objects will operate as hypertext links. The portability of HTML marked-up documents permits their independent use by Web browsers (5).

Numerous HTML references, standards documents, and style guides are widely available on the Web. HTML books and guides are available from most well-known and lesser-known technical publishers. The number of these books in print rivals that of books on any other computer programming language or application.

3.2 *Network-Based Information Retrieval*

The World Wide Web accomplishes network-based information retrieval by using a client-server architecture. A WWW user executes a WWW client program (usually referred to as a browser) on their local computer which is connected to a network. The client fetches documents from remote network nodes by connecting to a server on that node, requesting the document to be retrieved, and displaying that document to the WWW user.

3.2.1 THE WWW CLIENT-SERVER MODEL The client-server model offers advantages to both the information provider and the information user. The information provider is able to maintain control of their documents since they reside locally on the provider's computer. Furthermore the documents can be maintained by the information provider in any format, provided that format is recognized by the WWW server and user's client. This capability means that the minutes of a HEP collaboration meeting can be available to the entire collaboration as soon as the secretary makes them available to a WWW server. In addition, these minutes are maintained at a single location eliminating the problem of having multiple copies stored at different collaborator sites.

This client-server model can be naturally extended to allow documents to be dynamically created in response to a request from users. For example, the user's request can query a database and the result of that query returned as a hypertext document. The WWW interface to the SPIRES HEP preprints database is an implementation of this capability.

From the information user's perspective, all the documents on the Web are consistently presented as hypertext. If the information provider has done their job satisfactorily, the user will have access to a usable, consistent, reliable and accurate on-line resource. That user will not be required to know or understand any of the technical details surrounding the acquisition or even the location of the document provided.

3.2.2 HTTP A number of existing protocols and one new protocol are used in the WWW model. The HyperText Transport Protocol (HTTP) is a new protocol for file retrieval (not just hypertext) designed to operate as quickly as is needed in response to a request resulting from a hypertext link selection. HTTP is a very simple Internet protocol, similar in implementation to FTP or NNTP. A WWW client (browser) sends a document identifier with or without search words, and the WWW server responds with hypertext or plain text. The protocol runs over TCP, using typically one connection per document request. The client/browser renders the document as it is being received. The Multipurpose Internet Mail Extensions (MIME) standard is used to specify document types that the client supports, typically a variety of video, audio, and image formats in addition to plain text and HTML. Browser interpretation of these document types is known as "format negotiation."

3.2.3 UNIVERSAL RESOURCE LOCATORS (URLS) The power of the WWW model lies not in a complex protocol, but in the Universal Resource Locator (URL). [Note: some references define URL as Uniform Resource Locator or Unique Resource Locator.] The URL, designed to be compact, unambiguous, and printable, contains elements identifying the retrieval protocol to be used, the server to which the request is directed, and the document to be retrieved from that server. The URL is actually only one instance of a larger family of addressing mechanisms named URIs (Universal Resource Indicators). In general, the URI scheme addresses more specific issues such as guaranteeing access to documents in the event they are relocated.

The three elements are expressed in a URL in the following form:

<protocol>://<node>/<location>

The first element specifies the protocol to be used to access the document (e.g., HTTP, FTP, Gopher, NNTP, etc.) The second element specifies the network address of the server from which the document is to be obtained. The third element specifies a file path of the requested document on the WWW server. Therefore, the URL

http://www.slac.stanford.edu/home.html

defines a request for a file named "home.html" from a server with the Internet address "www.slac.stanford.edu." It specifies that this file should be retrieved using the HTTP protocol.

URLs may be transparently provided in a hypertext document using HTML markup or may be explicitly specified by a reader within a WWW browser. Expanded URL formats allow for non-specific document retrieval such as those resulting from the execution of a server-side command (e.g., common gateway interface scripts), a database query and for servers operating on non-standard port addresses.

Figure 2. illustrates how the URL addressing scheme, the HTTP protocol, and format negotiation server to integrate the client-server model used by WWW.

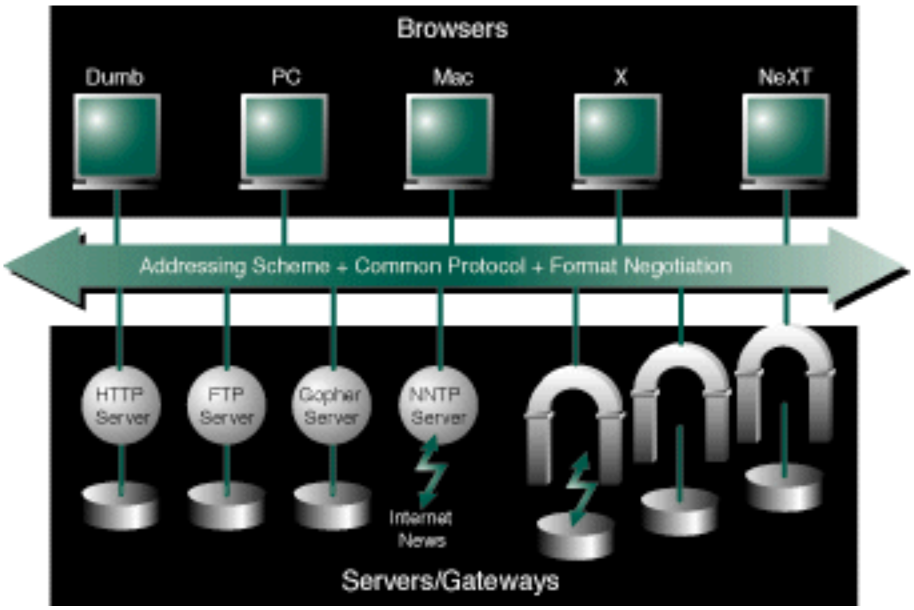

Figure 2. The World Wide Web Client-Server Model

4. EARLY HEP WWW APPLICATIONS

The World Wide Web did not immediately take the CERN physics community by storm. Hansjorg Klein of the Delphi collaboration recalls that a Hypertext Colloquium given by Tim Berners-Lee at CERN in November of 1990 was "overwhelming to many physicists." Early efforts to get Delphi experiment "on the Web" were not extremely successful. The adoption of WWW typically occurred (and still occurs) as a result of a particular need or application recognized by an experimental collaboration.

4.1 *SPIRES*

One of the first examples of a specific application of the World Wide Web to an HEP requirement was its ability to provide a readily-accessible, easy to use interface to the SPIRES databases at SLAC.

SPIRES (the Stanford Public Information REtrieval System) is used to support a set of databases covering a wide range of topics of relevance to HEP including experiments, institutes, publications, and particle data.

The largest of the SPIRES databases is the HEP preprints database, containing over 335,000 entries. Prior to its implementation on WWW, the only way to access the SPIRES databases was to log in to the IBM VM/CMS system at SLAC where the database resides, or to use the QSPIRES interface which operated only from remote BITNET nodes. With either method, in order to successfully access information in the database, a rudimentary knowledge of the somewhat esoteric SPIRES query language was required.

The Web, coupled with the widespread adoption of Bulletin Boards as the primary means of distributing computer-readable versions of HEP preprints, has revolutionized the ease of access and usefulness of the information in the SPIRES databases (6).

The SPIRES WWW server was one of the very first WWW servers established outside CERN and certainly one of the first to illustrate the power of interfacing WWW to an existing database, a task greatly simplified by WWW's distributed client-server design. Using this interface it is now possible for researchers to locate papers within the database without any knowledge of the SPIRES query language, by using simple fill-out forms (Figure 3.). Experienced SPIRES users are still able to use the SPIRES query language through this interface.

Access to some of the more advanced features of SPIRES, such as obtaining citation indexes and lists of top-cited papers, are also available via hypertext links. Since the access to the database uses WWW it can be used remotely from site on the Internet.

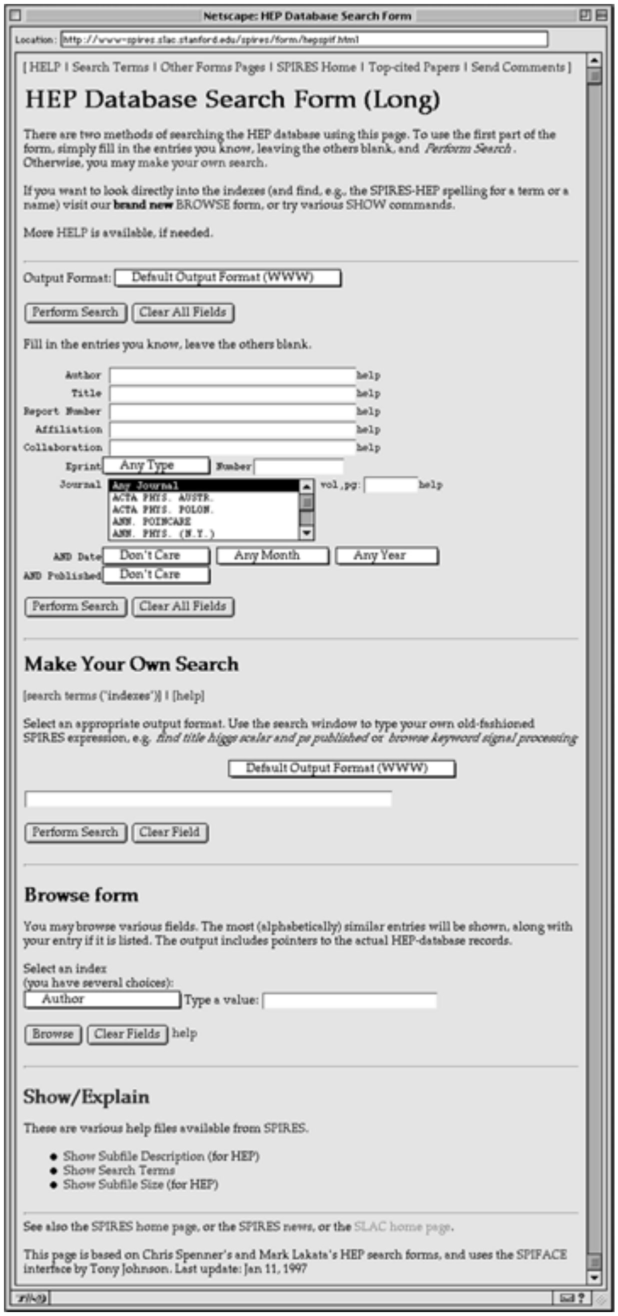

Figure 3. HEP Database Search Form

By linking the entries in the SPIRES databases to the computer-readable papers submitted to electronic Bulletin Boards at Los Alamos and elsewhere, it is also possible to follow hypertext links from the database search results and access either the abstract of a particular paper, or the full text of the paper, which can be viewed on-line and/or sent to a networked printer. A typical day may have over 10,000 accesses to the HEP preprints database from hundreds of institutions.

The WWW interface to SPIRES has now been extended to cover other databases including experiments in HEP, conferences, software, institutions, and information from the Lawrence Berkeley Particle Data Group.

4.2 *Home Pages*

Many HEP experiments and laboratories quickly realized how Web documents could be used to describe both their mission and results. Home pages and welcome pages for these laboratories and experiments were usually provided their Web authoring experience. Experiment "working pages" provide detailed information about an experiment to insure that collaborators are current in matters of data collection, analysis and software changes. Collaboration archives on the Web provide ready access to agendas, minutes, and schedules.

Home pages also became an integral part of a common computing environment. The ability of WWW to integrate with existing on-line systems permitted, for the first time, a common implementation-independent interface. Bjorn Nilsson recalls that in June 1993, ALEPH was able to generally provide information that was previously only available on-line on ALWS. A Web version of ALWHO provided accurate phone list and e-mail data. ALWS HELP and detector status were also provided.

The SLD collaboration incorporated its data monitoring system into a Web application available from its home page (Figure 4). This facility uses WWW forms to provide interactive access to databases containing up-to-date information on the performance of the detector and the event filtering and reconstruction software. Information can be extracted from the databases and used to produce plots of relevant data as well as displays of reconstructed events. Using these tools collaborators at remote institutes can be directly involved in monitoring the performance of the experiment on a day-by-day basis.

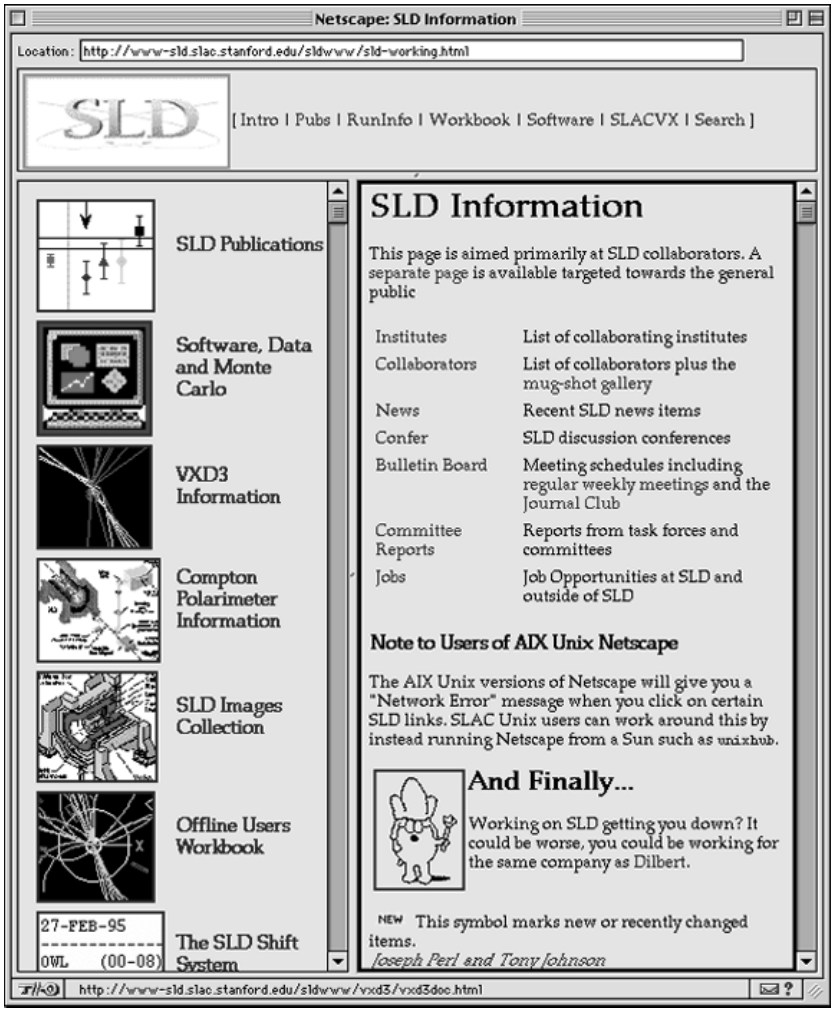

Figure 4. SLD Home Page

Major physics organizations such as the American Institute of Physics (AIP) and The American Physical Society (APS) also quickly recognized the potential of WWW as a communications medium. The APS first established a Website in January, 1994. By November, 1994 a second Web server was necessary to provide the desired services. The AIP debuted on WWW in April, 1994. Visitors to these organizational Websites (Figure 5.) have access to a wide variety of information - meetings, publications (including research journals), career services, fellowship opportunities, etc.

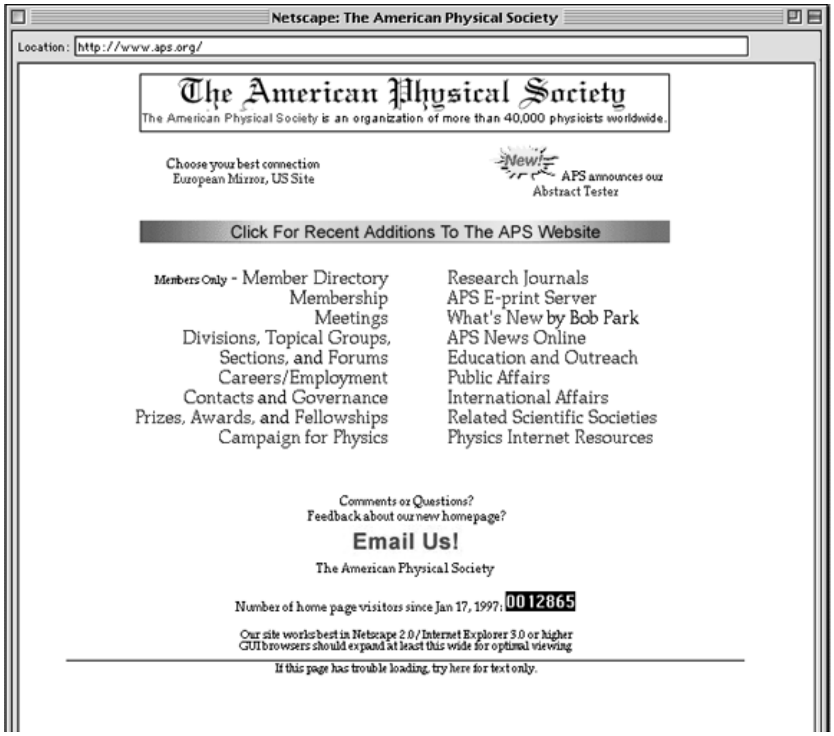

Figure 5. American Physical Society (APS) Home Page

The Particle Data Group (PDG) is an international collaboration that reviews Particle Physics and compiles and analyzes data on particle properties. PDG products are distributed to an estimated 30,000 physicists, teachers, and other interested parties. The Review of Particle Physics is one of the often most cited publications in HEP. The PDG home page, hosted by Lawrence Berkeley Laboratory (LBL), has made this distribution more accessible using WWW.

Institution and experiment resources on WWW (and other Internet services) have made research in high energy physics more accessible to the general public. Many HEP Websites now specifically contain information and resources targeted at non-HEP audiences such as the educational community (Figure 6.) and the press. A future commitment to provide such resources could play a major role in the design of Internet-based physics curricula.

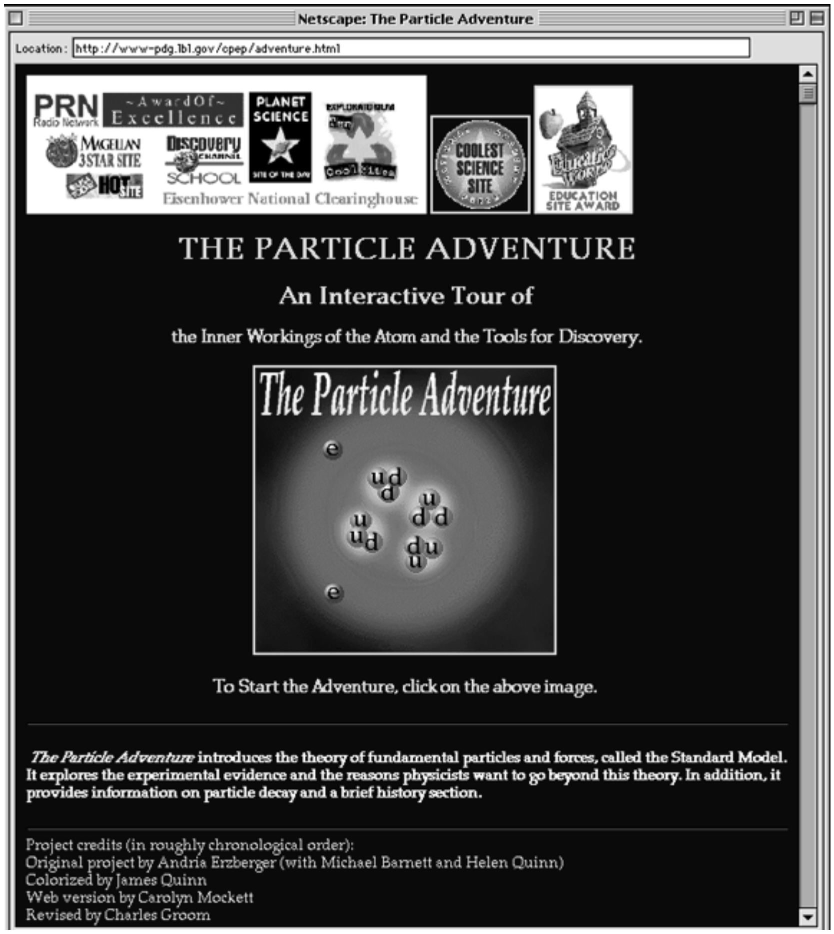

Figure 6. The Particle Adventure

Web pages provide an excellent mechanism for exceptional public announcements. Just hours after Prof. Martin Perl of SLAC was awarded the 1995 Nobel Prize, Web pages containing background information on the tau discovery, and Perl's first post-Nobel news conference (including photographs taken with a digital camera) were available to the world via WWW. Web pages on the Fermilab server chronicle the discovery of the top quark and the generation of anti-hydrogen (Figure 7.).

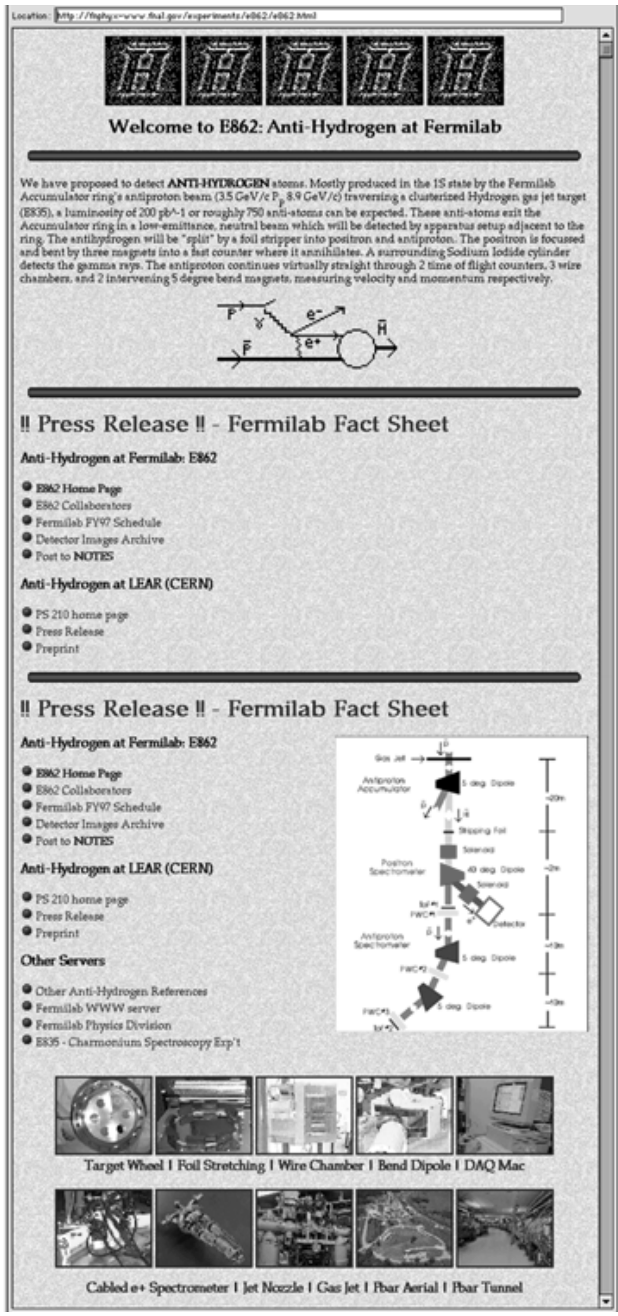

Figure 7. Welcome to E862 - Anti-Hydrogen at Fermilab

4.3 *Software Distribution*

WWW has played an important role in the success that HEP groups have had in maintaining central program source code bases while allowing distributed software development. It has also made a significant contribution to the process of distributing software and associated databases to remote sites. A decade ago, HEP researchers relied extensively on community wide software tools in addition to software developed for an individual experiment or research project. Today these same researchers are making increasing use of software developed outside of the HEP community including CAD/CAM systems, visualization software, symbolic mathematics packages, databases, software engineering packages, software for coarse and fine grained parallelism, C++ class libraries and many others. The Web provides a mechanism for keeping track of important software developments around the world in the many relevant fields or even within HEP itself.

FreeHEP consists of a collection of software and information about software which is useful in HEP and related fields (7). It was first proposed at the 1991 HEPLIB meeting to address the distribution of HEP and non-HEP software within the high energy physics community. FreeHEP consists of:

1. a global compilation of software useful in HEP from within HEP, other fields, and commercially;
2. tutorials on common subjects and reviews on the subject areas covered;
3. FTP and WWW access to information about software packages, documentation, source code, tutorials,

and reviews;

4. a mechanism for authors of software to distribute their packages;
5. a mechanism for users to communicate with software authors and/or with other users.

When first implemented, the FreeHEP compilation was available primarily via FTP. The compilation is now also imported into the SPIRES system which provides software database operations via hypertext.

5. CURRENT HEP WWW APPLICATIONS

5.1 *Conferencing Systems*

HyperNews is a Web-based application developed by the National Center for Supercomputing Applications (NCSA) that allows large groups of people to participate in electronic conferencing. It allows a user to read and contribute commentary, and supports various types of notification as activity in the area of discussion progresses. Each HyperNews conference begins with a "base article," which is followed by a nested list of postings by conference participants (8).

The BaBar experiment has adopted HyperNews as a mechanism for intra-collaboration communication. Figure 8. illustrates the wide variety of "conferences" in which BaBar collaborators are able to participate.

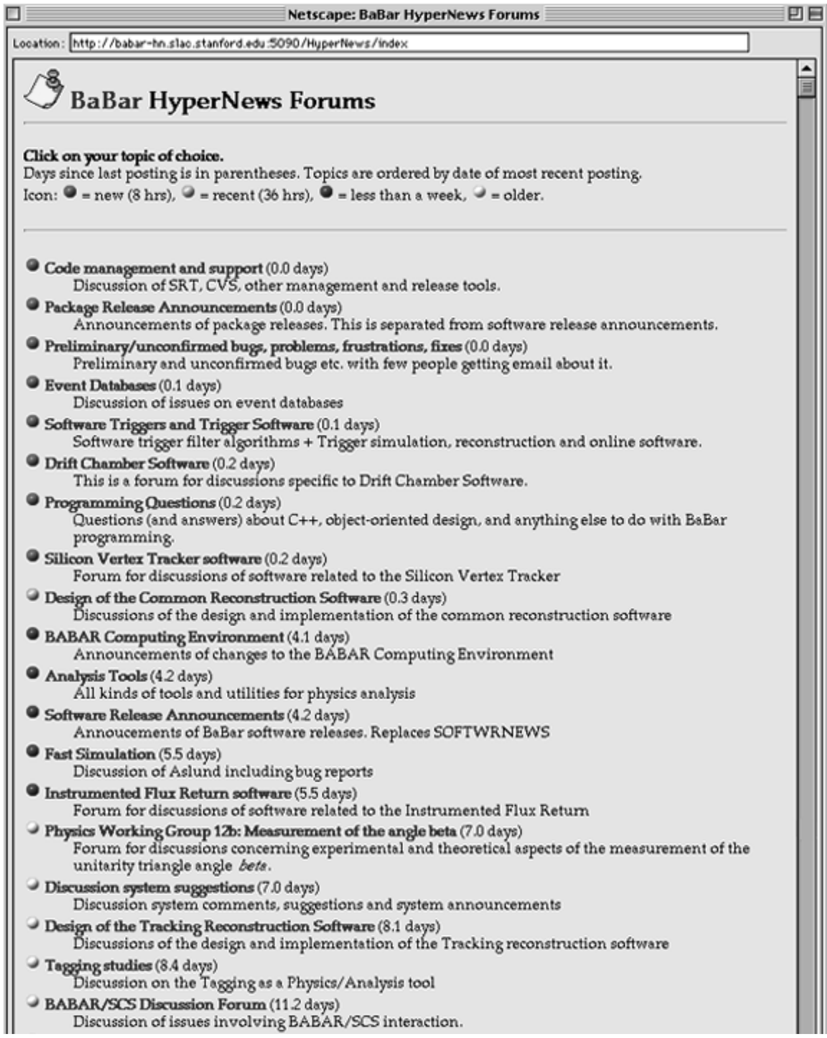

Figure 8. BaBar HyperNews Forums

While HyperNews is primarily text-based, the postings can contain multimedia elements. However, experimental Web-based conferencing systems available in the near future will support real-time audio and video, shared white boards, and other features now found in videoconferencing systems.

5.2 *Virtual Reality*

One of the most exciting research areas associated with WWW is the development of the Virtual Reality Modeling Language (VRML). The first discussions of incorporating virtual reality into the Web began at the First International WWW Conference at CERN in May, 1994. The VRML concept, like that of HTML, is to specify a very simple language capable of defining virtual reality scenarios. VRML documents can be translated into viewable objects on almost any computing platform, from small personal computers to very high-end graphic workstations. While the amount of graphical detail may vary between this range of platforms, the essential elements of the virtual world will be the same (5).

In 1994 the VENUS (Virtual Environment Navigation in the Underground Sites) project at CERN began an effort to develop a high performance 3D Web browser, particularly equipped to deal with large territorial and engineering models. This tool was named i3D and is intended to be used for LHC Engineering Design Integration. i3D will allow the entire HEP community to access and evaluate LHC's virtual prototypes on the Web (9, 16).

Components of the LHC detectors will be designed at many institutions located around the world. These conditions make design validation and integration an extremely difficult task. By its ability to load geometries via hyperlinks, i3D provides an ideal integration stage. Its data models, since i3D uses the WWW client-server model, can be maintained at the sites where they originate. Each object can be loaded on-line from wherever it is being designed, and a virtual world created in a users VRML-enabled browser.

The advantages of using i3D for LHC design integration are enormous. By using a database server, it is possible to manage centrally the hyperlinks and virtual worlds structure, in order to keep control of the overall project. For instance, the experiment's management controls naming of the components and associates URL's linking to geometry, drawings, home pages, etc. i3D then queries this database to obtain the appropriate information whenever a user selects a graphic hyperlink or "hyperdoor." Therefore navigating through an LHC detector with i3D allows not only an overall view (Figure 9.), but also the ability to obtain technical or organizational information on each component.

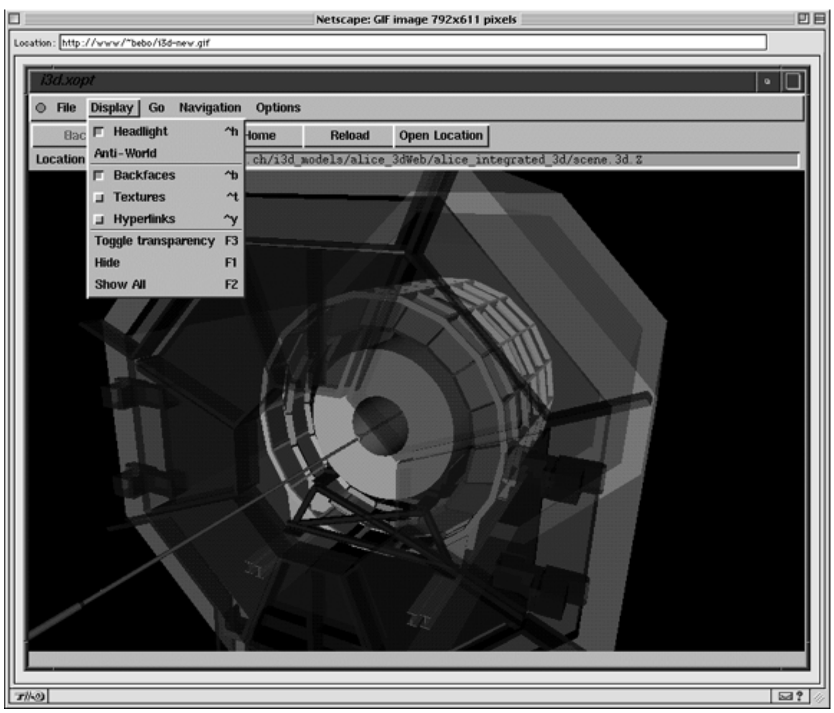

Figure 9. An i3D View of an LHC Detector

5.3 *Software Development and Maintenance*

Programmers who develop, use, maintain, and modify software are faced with the problem of scanning and understanding large quantities of documents, ranging from source code to requirements, analysis and design diagrams, user and reference manuals, etc. Their task is non-trivial and time consuming, because of the number and size of supporting documents, and the many implicit cross-references contained. In large distributed development teams, such as HEP collaborations where software and related documents are produced at various sites, the problem can be extreme. The fundamental concept of the LIGHT (LIfe cycle Global HyperText) system, developed by the Programming Techniques Group, ECP Division at CERN and the ALEPH experiment, is to provide access to all documentation associated with a software application or system, including source, via WWW with all cross-references automatically established (10). For example, links in a listing of a subroutine call written in FORTRAN refer to the subroutine's documentation; a link in a data element leads to the corresponding data definition, etc.

This concept and the LIGHT system has been applied to the JULIA reconstruction program of the ALEPH experiment at CERN. With JULIA/LIGHT, ALEPH programmers, documentation writers, maintainers of the data model and end user physicists are able to view through the Web the entire FORTRAN source code, data definition, and data design diagrams, as well as the JULIA, ALEPHLIB, CERNLIB, ADAMO, and BOS manuals (Figure 10.). All these documents are connected with hypertext links making it possible to inspect them in the widest variety of ways. The final production version of JULIA/LIGHT is expected to consist of approximately 6,000 HTML pages with 150,000 hypertext links.

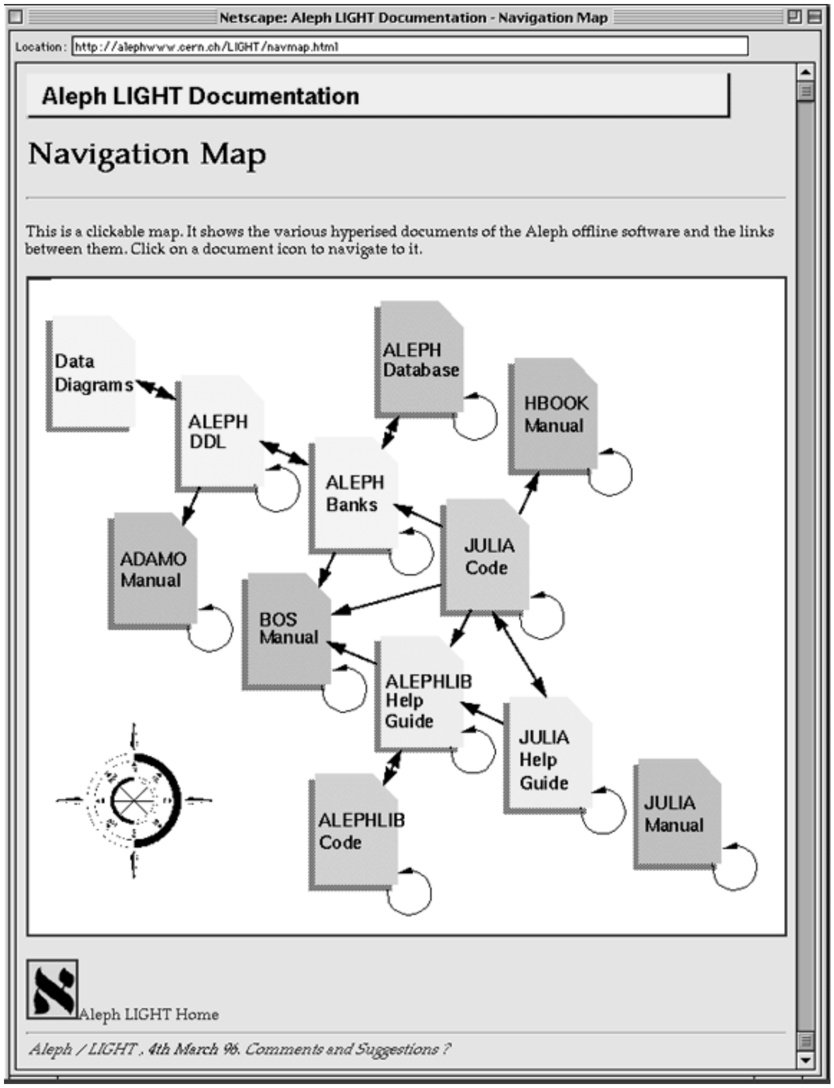

Figure 10. ALEPH LIGHT Documentation - Navigation Map

6. FUTURE HEP WWW APPLICATIONS

6.1 *Electronic Journals*

WWW has long been recognized as an effective medium for the distribution of documentation. In the near future it is likely to provide the basis for electronic journals in all disciplines not just HEP. The American Physical Society (APS) has a vision of the future of physics publishing, circa 2020. According to Burton Richter, former APS President and Director of the Stanford Linear Accelerator Center

(SLAC), "any physicist, any place in the country, can turn on his computer and for free browse through the table of contents of any APS journal." He adds that this browser "can select those things about which he wants to see an abstract, and then, after deciding what he might read, ask for the article itself and eventually pay for it like you pay your telephone bill" (11).

At first glance, Web-based publishing of scientific journals should be a relatively simple matter. HTML provides most of the required typesetting capability. Those capabilities absent (e.g., equation layout) can generally be accomplished with use of graphical elements.

The lack of a widespread availability of journals on the Web appears to be due to the fact that most parties involved are uncertain as to what "electronic publishing" actually means. Authors are primarily interested in electronic manuscript preparation and submission. Readers are interested in facilities provided for document/journal retrieval. Librarians focus on the delivery of information to users, but they often overlook archiving. Publishers concern themselves with the handling of electronic manuscripts, copyediting, formatting, typography and the production of versions suitable for delivery to the end users. They also worry about collecting the revenues needed to keep their operations financially viable. Authors and publishers are both concerned about the resolution of copyright and intellectual property rights associated with Web-published content (12).

Until technology advancements are able to address these issues, Web-based distribution of journals is unlikely to become commonplace. In the meantime, Web-based publishing and archiving of conference proceedings may provide a prototype. Figure 11. illustrates how the proceedings of the 1995 CHEP conference have been made available to Web users.

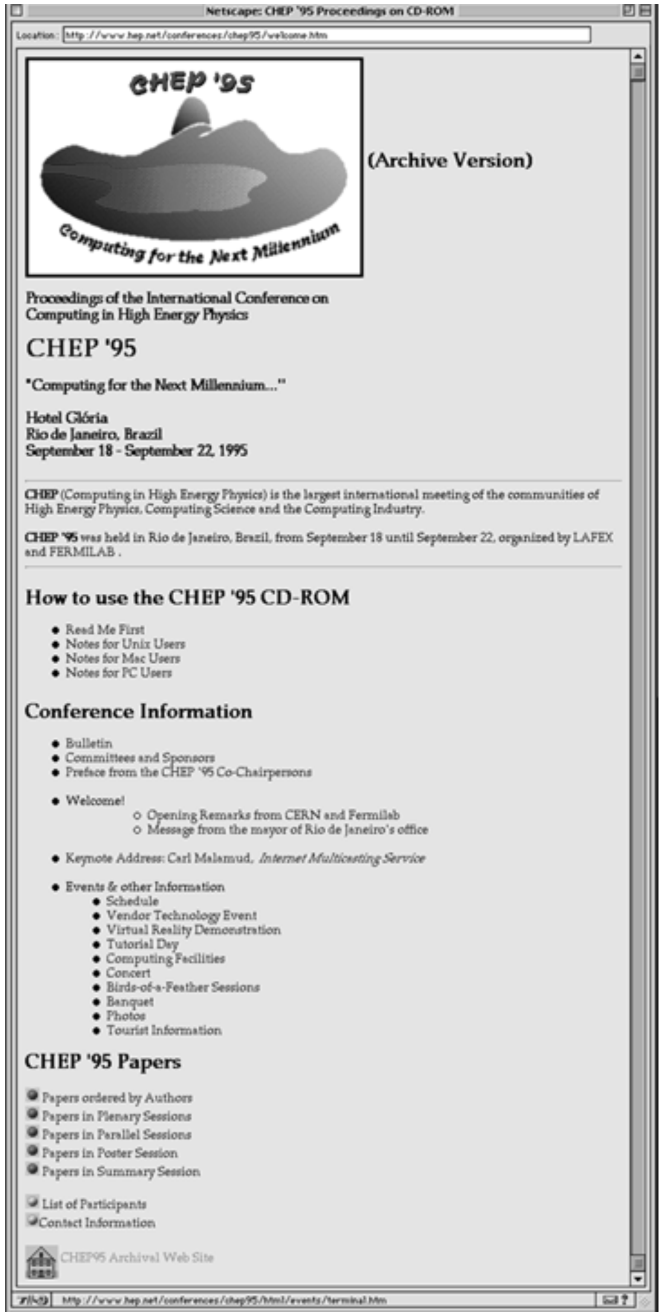

Figure 11. CHEP95 Proceedings

6.2 *Active Objects*

In addition to the virtual worlds demonstrated by the i3D LHC application, future Web-based HEP applications are likely to be enhanced by the inclusion of active objects. This technology, already demonstrated in a number of prototypes, will enable active objects such as spreadsheets or data plots to be embedded into objects such as HTML or VRML documents. While older browsers might only display these objects statically, the newer browsers will allow a user to interact with the object, perhaps by rotating a three dimensional plot, or expanding and rebinning a particular area of a data plot.

A current effort to provide similar functionality is EDWIN (Event Display WWW INterface) developed by DELPHI (13). EDWIN was designed to provide a generic method of interfacing event displays to the Web. Using an HTML form, a physicist can request a certain event and view (angle, zoom factor). The server returns the requested view as a bitmapped graphic. The requester iterates by resubmitting the form until satisfactory results are achieved. The EDWIN model is effective, but is very slow. Active object technologies will provide similar data interactivity but with faster response times.

6.3 *Collaboratories*

Scientific collaborations such as those in HEP currently rely heavily on face-to-face interactions, group meetings, individual action, and hands-on experimentation. Group size and attendance can vary widely. Technologies such as videoconferencing have begun to address some of these issues. A future solution may be found in "collaboratories."

The concept of a "collaboratory" came from William Wulf while working at the US National Science Foundation (NSF). Wulf merged the words "collaboration" and "laboratory" and defined a collaboratory as a ' ...center without walls, in which the nation's researchers can perform their research without regard to geographical location - interacting with colleagues, accessing instrumentation, sharing data and computational resource, and accessing information in digital libraries' (14)

A collaboratory facilitates scientific interaction within a group by creating a new, artificial environment in which group members can interact. This new environment must be socially acceptable to those who participate and improve their ability to work. Many computing tools must be brought together and integrated to allow seamless interaction. Some of these tools are already in wide use, such as WWW and e-mail, while others, like telepresence - the immersive electronic simulation of "being there" - are still being developed.

To facilitate scientific work, collaboratory systems must support the sharing of secure data, analysis, instruments, and interaction spaces. Videoconferencing improves communication and helps reduce the likelihood of major misunderstandings between collaboration members. Drawings, budgets, schedules and status reports can be shared in a collaboratory system insuring that engineering and manufacturing efforts are to specification, within budget and on schedule.

Because it is impractical to have all the critical personnel permanently located at the experiment site, collaboratory technology could even be used to decentralize the traditional central control room. Control of most HEP detectors is now done by networked computers, and permitting control over the Web from remote sites is a logical and easy extension of this trend.

The US Department of Energy has initiated a series of five major testbed projects known as the Distributed Collaboratory Experiments Environments (DCEE) Program. One of these testbeds is the Advanced Light Source (ALS) collaboratory at Lawrence Berkeley Laboratory. The model for this collaboratory is a loosely integrated set of Internet capabilities that appear as extensions to the Web. The ALS collaboratory has established a set of requirements that are similar to those of the other virtual laboratory DCEE projects. The requirements define a wide range of cross-platform functions that would allow researchers to interact with remote colleagues in a rich, in-process style. These requirements include:

* audio/video conferencing;
* chatting;
* shared computer display and whiteboard;
* shared on-line electronic notebook;
* file sharing with access security, safety, and data confidentiality;
* on-line instruments, remote experiment monitoring, computation and visualization;
* Web browser synchronization.

It is easy to recognize how the implementation of a collaboratory integrates many of the Web technologies discussed earlier in this chapter. The collaborative working environment can be presented as a virtual world. Web-based library technologies provide prototypes for electronic notebooks. Even the efforts supporting commerce on the Web are likely to make valuable contributions in areas such as security and privacy. Developing Web technologies such as Java will likely play a major role in the implementation of collaboratory requirements.

7. CONCLUSIONS

The World Wide Web was originally developed to address the needs of collaborative efforts in high energy physics. Since that time, the Web has undergone many changes, become a major force in the growth of the Internet, and has been used for dissemination of every conceivable form of information. Its widespread acceptance and modifiability are proof that its underlying design based on hypertext and client-server technology was a sound one.

The high energy physics community has adopted WWW as a critical component of its research effort. Web-based applications now support HEP in ways which the original WWW designers could never have imagined. Ongoing HEP developments using the Web closely parallel advances being made in overall WWW and Internet technologies. It is clear that WWW will continue to shape the way that HEP does physics and that HEP will continue to shape the development of WWW.

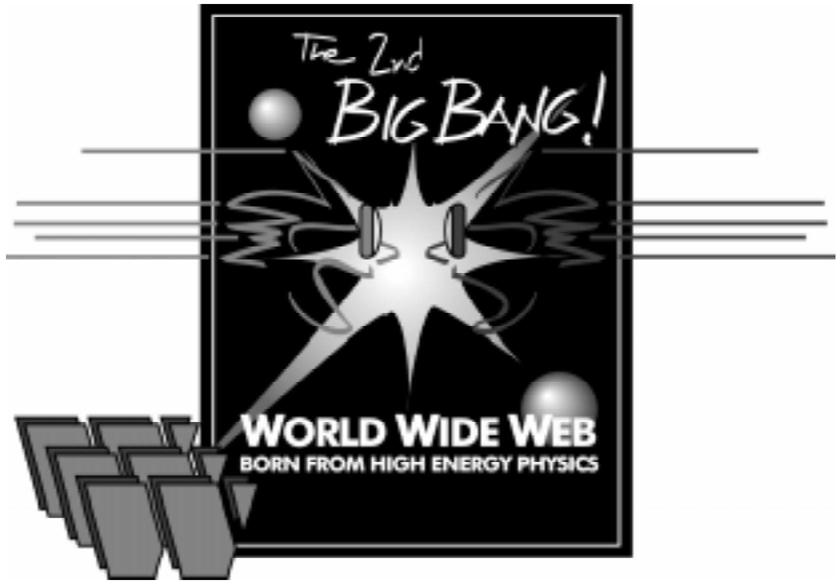
The 2nd
BIG BANG!
WORLD WIDE WEB
BORN FROM HIGH ENERGY PHYSICS

ACKNOWLEDGMENTS

I wish to thank Robert Cailliau, Hansjorg Klein, Paolo Palazzi, and Bjorn Nilsson for helpful discussions. Kathryn Henniss and Terry Anderson were invaluable in helping me to prepare this manuscript. This work was supported by the US Department of Energy Contract No. DE-AC03-76F00515.

Literature Cited

1. Verkerk, C, Wojcik, W. (edit), *Proc. Int. Conf. on Computing in High Energy Physics, Annecy, France, 1992*
2. Berners-Lee T, Cailliau R., *Proc. Software Engineering, Artificial Intelligence, and Expert Systems for High Energy and Nuclear Physics, La Londe-les-Maures, France, 1992*
3. Segal, B. *A Short History of Internet Protocols at CERN*, CERN PDP-NS, April 1995
4. Nelson, T. *Literary Machines*. Sausalito: Mindful Press (1992)
5. White, B. *HTML and the Art of Authoring for the World Wide Web.* Boston: Kluwer Academic Publishing (1996)
6. Johnson, T. 1994. *SLAC Beamline*. 24 (3):2-9
7. Johnson, T et al., [WWW document]. URL http://heplibw3.slac.stanford.edu/FIND/FHMAIN.HTML
8. LaLiberte, D et al., [WWW document]. URL http://union.ncsa.uiuc.edu/HyperNews/get/hypernews.html
9. de Gennaro, S. *Proc. Int. Conf. on Computing in High Energy Physics, Rio de Janeiro, Brazil, 1995* (1996)

10. Aimar, A et al., *Proc. Int. Conf. on Computing in High Energy Physics, Rio de Janeiro, Brazil, 1995* (1996)

11. Taubes, G. 1993. *Science*. 259:1246

12. Boyce P, Dalterio H. 1996. *Physics Today*. 49(1):42

13. Dönszelmann, M et al., *Proc. Int. Conf. on Computing in High Energy Physics, Rio de Janeiro, Brazil, 1995* (1996)

14. Kouzes, R, Myers, J, Wulf, W. 1996. Computer. 29(8):40-46

15. Agarwal, D et al.,*Proc. Int. Conf. on Computing in High Energy Physics, Rio de Janeiro, Brazil, 1995* (1996)

16. Balaguer, J-F. *Proc. 3rd Eurographics Workshop on Virtual Environments, Monte Carlo, 1996* (1996)