THE WORLD WIDE WEB AND HIGH-ENERGY PHYSICS

The Web—what you may now be using to buy an airline ticket or look up Miles Davis's discography—began life as a tool for improving the flow of information at CERN.

Bebo White

In his 1998 State of the Union address, President Bill Clinton told Congress:

We should enable all the world's people to explore the far reaches of cyberspace. Think of this: The first time I made a State of the Union speech to you, only a handful of physicists used the World Wide Web—literally, just a handful of people. Now, in schools, in libraries, homes and businesses, millions and millions of Americans surf the Net every day.

Who were the physicists that Clinton was referring to, and why did the physics community provide the fertile ground from which the Web first sprouted? Although the Web has come to mean many things to many people—electronic commerce, digital libraries, mass advertising, for instance—its roots lie in the simple need for scientists living and working far apart to collaborate with one another on common interests and purposes.

The nursery

CERN, located on the Franco-Swiss border near Geneva, is one of the most successful high-energy laboratories in the world. Its success can be measured in terms of both the number of Nobel prizes its staff have garnered and its proven record of carrying out huge projects. Experiments at CERN and other high-energy physics laboratories use complex particle detectors, which cost millions of dollars and typically involve hundreds of physicists from institutions spread around the world.

Two major elements of successful collaboration are communication and documentation. Researchers working together in different countries must have easy access not only to experimental data, computer program code and drafts of scientific papers, but also to more mundane administrative information, such as meeting minutes and agendas and telephone and e-mail lists. In general, computing and networking technologies provide useful tools for collaborations, but the implementation of specific technologies can also create roadblocks, since collaborating institutions tend to use different computer operating systems and document preparation systems. It was in this environment and to address these kinds of problems that the Web was conceived at CERN.

BEBO WHITE is a member of the research staff and the Webmaster at the Stanford Linear Accelerator Center in Stanford, California.

The early years

The development of the Web at CERN was actually the convergence and synthesis of three separate technologies—computer networking, document/information management and software user interface design. Applying these technologies to the needs of the experimental collaborations led to the solution that became the Web. However, as described below, the Web was not immediately recognized as the solution. CERN's Ben Segal remembers:

In the beginning was—chaos. In the same way that the theory of high-energy physics interactions was itself in a chaotic state up until the early 1970s, so was the so-called area of Data Communications at CERN. The variety of different techniques, media and protocols used was staggering; open warfare existed between many manufacturers' proprietary systems, various home-made systems, and the then rudimentary efforts at defining open or international standards. There were no general purpose Local Area Networks (LANs); each application used its own approach.¹

Network experimentation at CERN continued through the late 1970s and the 1980s. CERNET, which was started in 1976, linked a number of mainframes and minicomputers. From 1981 to 1983, the STELLA Satellite Communication Project used a satellite channel to link CERNET, the Istituto Nazionale di Fisica Nucleare in Pisa, Italy, and the Rutherford–Appleton Laboratory in Didcot, England. Despite the introduction of Ethernet to the CERN community in 1983, most mainframe and minicomputers still used proprietary networking systems. And although the chaos noted by Segal had been somewhat alleviated, an overall networking strategy was still missing.

While CERN was busy seeking solutions to its networking problems, the global phenomenon known as the Internet was maturing in the US. At the time, no one foresaw that these two efforts would collide and revolutionize the world of communications.

The birth of the Internet

Work on the Internet was initiated in the late 1960s by the US Department of Defense, which wanted to develop a communications network that could survive a nuclear war. In particular, messages had to be able to reach their destinations, even if significant communication routes were destroyed. The solution, the "packet switching"

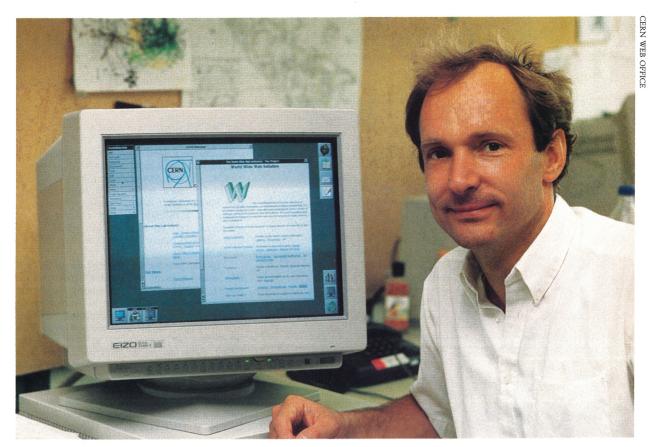


FIGURE 1. TIM BERNERS-LEE, the Oxford-educated physicist-turned-CERN programmer who invented the World Wide Web.

network, was designed so that the particular route a message takes to reach its destination is unimportant. This ability of a message to travel by other routes assures the sender that it will eventually be received. The original Internet was named Arpanet, after its sponsor, the Advanced Research Projects Agency.

Arpanet led to the development of numerous important networking technologies. Among them was TCP/IP (Transmission Control Protocol/Internet Protocol), which is a method for converting messages to data streams, directing them to a particular address on the network and reassembling the message at the destination. Arpanet (and later the Internet) was designed to be a network of networks that contain different types of computers and different networking technologies. To be able to communicate, all the computers and networks simply had to "speak" TCP/IP.

Arpanet's first "killer application" (a software program whose popularity surpasses its competitors to the point of all but killing them off) was e-mail, thanks to which Arpanet's potential was quickly realized by many people outside the US Defense agencies.

Formally decommissioned in 1989, Arpanet became, more generically, the Internet. Its users scarcely noticed, for Arpanet's functions not only continued, but also steadily improved. The use of TCP/IP standards for computer networking is now global.

Between 1985 and 1988, CERN local networks converged on TCP/IP. In 1989, CERN opened its first external connections to the Internet, and, by 1990, it was the largest Internet site in Europe. CERN quickly became a major force in the development and support of the Internet

both in Europe and worldwide. In his account entitled "A Short History of Internet Protocols at CERN," Segal writes:

A key result of all these happenings was that by 1989 CERN's Internet facility was ready to become the medium within which Tim Berners-Lee would create the World-Wide Web with a truly visionary idea. In fact an entire culture had developed at CERN around 'distributed computing. . . .' It is my belief that the Web could have emerged considerably earlier if CERN had been connected earlier to the Internet.¹

The networking experience of CERN provided a foundation for the development of the Web. Missing, however, was a way to describe information that could easily be delivered to the user community over the nascent Internet.

Hyperdocumentation

Back in 1983–84, a working group was set up to design an electronic document handling system that would facilitate the storage and retrieval of CERN's engineering and computing documents. At that time, document preparation at CERN was centered on the use of Waterloo Script on the IBM central mainframe computer. In 1985, some groups began experimenting with SGML (the Standard Generalized Markup Language). The first prototype of this documentation system was introduced on the IBM mainframe and named CERNDOC. It was one of the first electronic documentation systems to use a client–server model. As Eric van Herwijnen (one of the architects of CERNDOC) told me:

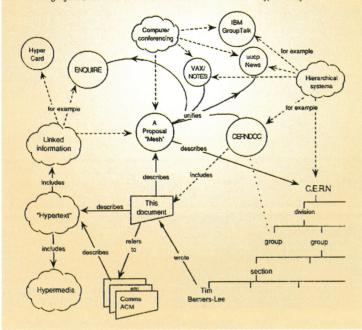
Tim [Berners-Lee] was very interested in CERN-

Information Management: A Proposal

Tim Berners-Lee, CERN

March 1989, May 1990

This proposal concerns the management of general information about accelerators and experiments at CERN. It discusses the problems of loss of information about complex evolving systems and derives a solution based on a distributed hypertext system.



DOC and spent quite some time talking about possible extensions to CERNDOC. The way that images were handled was particularly unsatisfactory. PostScript was just becoming popular at that time, and the system couldn't automatically translate from IBM's EBCDIC to ASCII for VAX clients. I believe that Tim got the client—server idea from CERNDOC as well as the idea of format negotiation through a special purpose protocol that relied on something looking like SGML as the document format.

With the Internet providing the transfer mechanism and CERNDOC providing a prototype for content, the Web lacked only a user interface. FTP (File Transfer Protocol) and Gopher were accepted Internet information retrieval systems, but neither provided a uniform, platform-independent user interface. And perhaps worst of all from the users' point of view, both systems required knowledge of a specific set of commands. The solution to the interface problem was found in the pointing and linking features of hypertext.

Proposing the Web

Tim Berners-Lee (see figure 1) was a programmer working on networking in CERN's computer division. In a 1989 paper entitled "Information Management: A Proposal," he recalled:

In 1980, I wrote a program for keeping track of software with which I was involved. Called Enquire, it allowed one to store snippets of information, and to link related pieces together in any way. To find information, one progressed

FIGURE 2. TIM BERNERS-LEE'S prototype Distributed Hypertext System for CERN, as sketched in the original proposal for the World Wide Web.

via the links from one sheet to another, rather like in the old computer game Adventure. I used this for my personal record of people and modules. It was similar to the application HyperCard produced more recently by Apple for the Macintosh. A difference was that Enquire, although lacking the fancy graphics, ran on a multiuser system, and allowed many people to access the same data.²

In his proposal, Berners-Lee addressed the management of general information about accelerators and experiments at CERN and the problems of inconvenient access to current and archived information. The solution, as he described it, was "the integration of a hypertext system with existing data, so as to provide a universal system, and to achieve critical usefulness at an early stage." Figure 2 shows a diagram from the proposal, which illustrates how the proposed system would integrate such diverse information sources as VAXNOTES, USENET, IBM GroupTalk (which I implemented at CERN) and CERNDOC.

Independently of Berners-Lee, Robert Cailliau was also looking at hypertext as a solution to CERN's information retrieval problems. He was especially interested in

platform-dependent hypertext systems such as HyperCard and in hypertext scripting languages such as HyperTalk.

In 1990, Berners-Lee and Cailliau decided to collaborate on a revision of Berners-Lee's 1989 proposal to CERN management. Cailliau recalls his discussion with Berners-Lee (held at the CERN cafeteria) over what to name the proposed project:

We agreed that Greek names were out—too old-fashioned. Then one of us had the idea: "What about the World-Wide Web?" "That'll do," said the other. "Let's put that down, and think of something better later."

The new proposal was entitled "World-Wide Web: Proposal for a Hypertext Project." Later never came.

In October 1990, the two programmers submitted the revised proposal, whose key goals were stated as being

▷ 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;

b 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

CERN WEB OFFICE

search is simply a hypertext document consisting of a list of references to nodes which match the keywords:

be 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:

be to provide the software for the above free of charge to anyone.

With the Berners-Lee-Cailliau proposal, which got the go-ahead in September 1990, all the pieces contributing to the development of the Web at CERN were complete. However, implementing the described model was not simple. It required the development of server software (to provide content) and client software (to receive and interpret content). In addition, the method that allows the client and server to communicate with one another-that is, the Hypertext Trans-

port Protocol (HTTP)—had to be invented, as did the Hypertext Markup Language (HTML), which enables content to be displayed effectively on a client (or browser). HTML is based on SGML, revealing again the Web's CERNDOC roots.

By November 1990, the initial implementation of a Web client for the NeXT computer (see figure 3) had been developed. This first client was able to display documents using multiple fonts and styles and could even be used 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" browser, which is shown in figure 4. It could run on numerous platforms, but displayed output only on character-based terminals, such as the VT100. Soon af-

terward, the Web server software was successfully installed on the CERN VAX/VMS platforms and the central IBM mainframe.

In December 1991, the CERN Computer Newsletter first introduced the Web to the high-energy physics community. At that point, Cailliau estimates, there were 10 Web servers worldwide.

Contrary to popular belief, the Web did not initially take the CERN and high-energy physics community by storm. Hansjorg Klein of the Delphi collaboration recalls that a Hypertext colloquium given by Berners-Lee at CERN in November 1990 before the Web's formal announcement was over-

FIGURE 4. A SAMPLE SCREEN from the first World Wide Web browser.

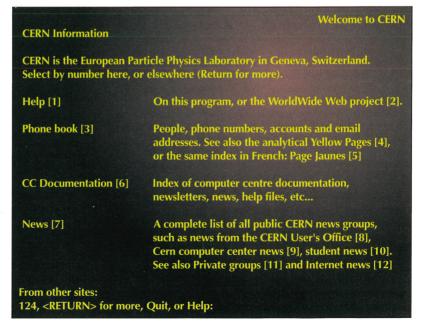


FIGURE 3. THE ACTUAL NeXT COMPUTER that ran the original World Wide Web server and browser at CERN.

whelming to many physicists, who viewed it as an esoteric computer science exercise and did not immediately recognize it as an improvement over existing methods for helping physicists to do physics. Early efforts to get the Delphi experiment on the Web were not successful. Van Herwijnen remembers:

As for initial user reactions, I remember that no one believed in the Web at CERN in the beginning. Many people thought that Tim [Berners-Lee] was not right in inventing yet another protocol (HTTP)—now they take part of the claim of having inventing it. The Web was only taken seriously here after Mosaic came along.

And Paul Kunz, a mainstay of the high-energy physics computing community, who was instrumental in establishing the first US Web site, at the Stanford Linear



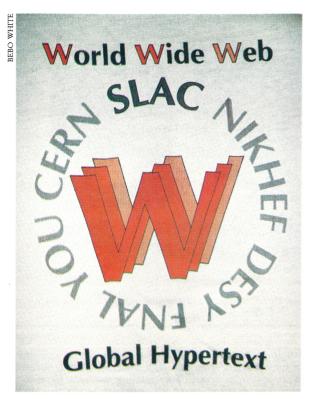


FIGURE 5. THE FIRST WORLD WIDE WEB T-SHIRT. Dating from 1992, it bears the names of the five high-energy physics laboratories on the Web at the time—the Stanford Linear Accelerator Center (SLAC), the National Institute for Nuclear Physics and High-Energy Physics (NIKHEF) in Amsterdam, the German Electron Synchrotron Laboratory (DESY) in Hamburg, Fermilab (FNAL) and CERN. The designers of the shirt joked: "Maybe some day we'll have enough sites to go around the wearer's chest."

Accelerator Center (SLAC), recalls:

The first I heard about the WWW was an article in a comp.sys.next newsgroup. Tim announced to NeXTStep (the NeXT operating system) users the availability of his application. My thoughts were something like: "what are these CERN people up to now?" I had absolutely no interest as the application was pushed as a way to distribute documentation, and what developer finds documentation fun?

Others, however, were quick to see the Web's potential. In early 1992, Ruth Pordes and Jonathan Streets of the Fermilab computing division's on-line systems department (OLS) were considering the problem of providing high-energy physics experimenters with information about on-line data acquisition systems. Seeing the Web presentation to the Artificial Intelligence in High Energy Physics conference (AIHEP '92) at La Londe, France, in February 1992, Streets recommended the Web as being "the best thing around," and OLS decided to adopt it. Kunz attended the same presentation and recalls that the Web demonstration included a connection to the SLAC server, which was already in production at that time.

At the wrap-up session of the Computing in High Energy Physics conference (CHEP92) held at Annecy, France, in September 1992, Terry Schalk of the Santa Cruz Institute for Particle Physics told the attendees that

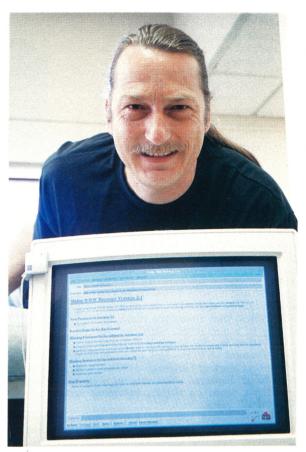


FIGURE 6. TONY JOHNSON and the Midas browser that he devised.

the conference highlight for him had been the short presentation about the Web. (The commemorative t-shirt is shown in figure 5.) Ever since then, Web applications have played a major role at the CHEP conferences. (CHEP97 in Berlin held joint sessions via videoconferencing with the Sixth International WWW Conference in Santa Clara, California.)

According to Cailliau, the number of Web servers worldwide at the end of 1992 was approximately 50.

CERN's Bjorn Nilsson recalls that in June 1993, the ALEPH experiment at CERN was able for the first time to provide information previously limited to users of ALWS (the Aleph offline system). A Web version of ALWHO provided an accurate phone list and e-mail data. ALWS HELP and detector status were also provided.

The first browsers

The Web is often described as the killer application of the Internet. However, in those early days, it needed its own killer app to catch on within the high-energy physics community. That application may well have been the Web's 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 high-energy physics, including experiments, institutes, publications and particle data.

Since 1974, the SLAC library has participated in providing SPIRES-HEP, a 300 000-record bibliographic da-

The Future Impact of the Web on the Physics Community

The usefulness of the World Wide Web to the high-energy physics community is no longer questioned, and the operation of a major experiment or laboratory without the Web is almost unimaginable. Not just a tool, the Web has also made research in high-energy physics more accessible to the general public through offering pages written for the educational community and the press. And the Web provides institutions with an excellent mechanism for making exceptional public announcements about such critical events as Nobel Prizes and important discoveries.

Although, for many, the Web has become an advertising and commercial medium whose credibility should not be taken for granted, the high-energy physics and other scientific communities continue to use it in many of the same ways that prompted its conception. There are, however, at least three major areas in the scientific community that could be dramatically changed by some of the new Web technologies—namely, advanced forms of software management, collaborative tools and electronic publishing.

Software management

High-energy physics experiments depend on computer software that is written, maintained and run by large numbers of people. The Web has played an important role, helping experimental groups to manage central databases of program source code while allowing programmers at different locations to develop the source code. It has also significantly eased the process of distributing software, documentation and associated databases to remote sites. Good examples of current applications in this area are the Web interface to the FreeHEP collection and the CERN LIGHT (Life cycle Global HyperText) system.

But in the near future, Web-based software systems will be quite different and exciting. For example, think of a physicist writing platform-independent analysis software in Java on a local personal computer. The software would be compiled, tested and debuggged

locally, but then uploaded to a remote server. Software users would run the application in their Web browsers after downloading the program code from the server. The server could also provide data sets of the most current experimental data. Such a system could potentially eliminate the need for local program libraries and guarantee access to the latest data. Where local program and data libraries are still necessary, Web "push" technology could provide version control by installing the latest software fixes. Current, or especially significant data sets could also be pushed to experiment collaborators for analysis.

Collaborative tools

Physicists must collaborate. The first Web-based conferencing system, WIT, was developed at CERN in 1994 as a quick hack. The BaBar experiment makes extensive use of HyperNews, a conferencing system developed by the National Center for Supercomputing Applications. Various other commercial and home-grown tools are also being used in other high-energy physics collaborations.

But the current text-based conferencing systems—and even videoconferencing—provide only tantalizing pieces of the col-

laborative systems of the near future. DOE2000 is a new initiative⁷ that will fundamentally change the way in which scientists work together and how they will address the major challenges of scientific computation. Web technology plays an important role in the implementation of the initiative's concept of the "collaboratory," which promotes itself as a "center without walls, in which researchers can perform their research without regard to geographical location—interacting with colleagues, accessing instrumentation, sharing data and computational resources, and accessing information in digital libraries."⁸

Electronic publishing

In a PHYSICS TODAY article (January 1996, page 42), Peter Boyce and Heather Dalterio analyzed the electronic publishing of scientific journals and concluded that, ultimately, "it creates new capabilities that extend far beyond what paper journals can provide. It is precisely these new capabilities that will make electronic publishing such a powerful tool for scientists." However, their enthusiasm was tempered by their concern that Web-based electronic publishing suffered from at least three major shortcomings: (1) Hypertext Markup Language (HTML) is not sufficiently structured or versatile for scientific literature; (2) papers are useless unless they can be promptly found; (3) electronic journals must collect enough revenue to support the infrastructure, but not in a way that inhibits the use of electronic tools.

Recent advances in Web technology have addressed these three issues head on. The eXtensible Markup Language (XML) presently being defined by the World Wide Web Consortium (W3C) promises to combine the power of SGML with the

compactness of HTML. With XML, Web authors will be able to create sets of data element tags and structures that define and describe the information contained in a document. XML emphasizes the separation of document structure and content. In conjunction with supplementary style sheets

mentary style sheets (for mathematics, say), XML will provide high-energy physics authors with the tools and functionality previously available only with TeX/LaTeX or SGML.

Continuing advances in search engine technology will make finding material on the Web easier, but will only partly address the frustrating resemblance Web searching bears to indiscriminate drift-net fishing. Document metadata will be used to ensure document applicability, quality and stability. Efforts such as the Dublin Core initiative seek to identify a core set of metadata elements—that is, indexing keywords—for finding things more effectively on the Web. Originally conceived for author-generated descriptions, this initiative has also attracted the attention of archivists at museums and libraries.

Document subscriptions work well in the paper world, but are unlikely to work on the Web. Micropayment technology¹⁰ addresses the revenue generation issues of electronic publishing. Under a micropayment scheme, readers of Web documents will pay cents (or fractions of cents) for access. According to Jakob Nielson, "Micropayments lower the threshold and do not require a big decision before users get their initial benefits."¹¹



Now three years old, the journal "New Astronomy" was launched simultaneously in Web and print versions.

tabase, to the world's high-energy physics community. Before its implementation on the Web, the only way to access SPIRES-HEP was to log in to the IBM system at SLAC, where the database resides, or to use the QSPIRES interface, which operated only from remote nodes using BITNET. With either method, a rudimentary knowledge of the somewhat esoteric SPIRES query language was required. According to Berners-Lee, the mounting of SPIRES-HEP on the Web was a vitally important factor in the rapid acceptance and utilization of the Web by the international high-energy physics community.⁵

Although ready access to SPIRES may have facilitated the acceptance of the Web, "It wasn't until NCSA [the National Center for Supercomputing Applications] made Mosaic that the average person could see why this was really a good thing," recalls David Ritchie, who chaired the working group that designed the first home page for Fermilab. Mosaic brought forth the Web from physics and computer science laboratories and set the path for the role it plays today. The first version of Mosaic emerged in February 1993, but before its release, several other significant Web browsers had already been developed.

The early browsers developed at CERN in 1991–92 had been quickly followed by the first browsers designed for the X Windows computer environment—namely, the Viola browser developed by Pei Wei of the University of California, Berkeley (and later O'Reilly and Associates) and the Midas browser developed by Tony Johnson at SLAC. (See figure 6.)

Like the Midas and Viola browsers before it, Mosaic was freely available on the Internet. Features provided by Midas and Viola undoubtedly influenced the course of Mosaic's development. Fermilab's Pordes remembers:

In 1993 I invited Tim [Berners-Lee] to Fermilab. I contacted NCSA and suggested that we go down and meet them. It was clear they were really going to market Mosaic. That was the first time we realized that he [Berners-Lee] was really going to have a big success.⁶

Through 1993 and 1994, CERN and NCSA continued to be the leaders in Web development. NCSA provided a rapid succession of improvements (and new features) for Mosaic and developed the NCSA server. Meanwhile, the small but dedicated Web group at CERN continued its efforts on the CERN server. More important, however, the CERN group developed and supported Libwww, the general software library—the nuts and bolts—for building Web browsers and servers. Pioneering work on critical Web features, such as access authorization, resulted from the Libwww effort.

In May 1994, the First International WWW Conference was held, at CERN. Described by one press pundit as the Woodstock of the Web, it was rapidly oversubscribed.

Cailliau estimates that the number of Web servers worldwide went from 2500 at the end of 1994 to 50 000 at the end of 1995. This growth was due largely to the availability of server software developed by CERN and NCSA.

Berners-Lee had long argued that the standardization of the Web be open and nonproprietary. His model was that of the X Consortium, which had been formed in 1988 to further the development of the X Window system and whose major goal was the promotion of cooperation within the computer industry in the creation of standard software interfaces at all layers in the X Window system environment. MIT had for many years provided the vendor-neutral architectural and administrative leadership required

to make that organization work. In December 1994, after much negotiation, Berners-Lee left CERN for the MIT Laboratory for Computer Science (LCS) to become director of the newly formed World Wide Web Consortium (W3C).

Ultimate irony

Ironically, it was high-energy physics that brought the end to Web development at CERN (and effectively at any other high-energy physics laboratories). In late December 1994 and early January 1995, the CERN Council unanimously approved the construction of the Large Hadron Collider, the world's most ambitious high-energy physics accelerator to date. Because the LHC would cost so much, the council decided not to fund further Web development, and, in negotiations with the European Commission and INRIA (the Institut National pour la Recherche en Informatique and Automatique in France), transferred the CERN Web project to INRIA. It was also agreed that INRIA, along with MIT/LCS, would host the W3C. (Keio University in Japan became the third host of the W3C in 1996.)

By the end of 1997, Cailliau estimates the number of Web sites worldwide had reached 650 000—on every continent and, possibly, in every country. The number of Web users is now certainly in the millions. Web addresses have become almost as widespread as telephone numbers, while Web sites provide every conceivable form of content or service. And, as described in the box on page 35, the Web promises to dramatically change how scientists work and interact.

In reflecting on the current size and scope of the Web, it is important to remember that this revolutionary technology grew from the needs and imagination of the high-energy physics community and that scientific research—even on the ultimate constituents of matter—often leads to significant technical spin-offs.

I wish to thank Tim Berners-Lee, Robert Cailliau, Tony Johnson, Hansjorg Klein, Paul Kunz, Jennifer Masek, Bjorn Nilsson and Eric van Herwijnen for their help and recollections. Work for this article was supported by the US Department of Energy Contract No. DE-AC03-76F00515.

References

- B. Segal, "A Short History of Internet Protocols at CERN," in-house CERN document (1995). Also available at http://www.c.cern.ch/pdp/ns/ben/TCPHIST.html.
- 2. T. Berners-Lee, "Information Management: A Proposal," inhouse CERN document (1989; rev. ed. 1990). Also available at http://www.w3.org/History/1989/proposal.html.
- T. Berners-Lee, R. Cailliau, "World-Wide Web: Proposal for a HyperText Project," CERN in-house document (1990).
- P. Garrett, D. Ritchie, Collaborating over the Web: Libraries and Laboratories or the Librarian and the Webmaster, Fermilab, Batavia, Illinois (1995).
- T. Berners-Lee, keynote address, History/Developers' Day, at Sixth International WWW Conference, Santa Clara, Calif., 11 April 1997. Also available at http://www.w3.org/ Talks/9704WWW6-tbl.
- 6. E. Berger, FermiNews, vol. 19, no. 16, 16 August 1996.
- 7. Available at http://www.mcs.anl.gov/DOE2000/.
- 8. R. Kouzes, J. Myers, W. Wulf, Computer 29 (8), 40 (1996).
- D. Connolly, ed., "XML: Principles, Tools, and Techniques. World-Web Web Journal," vol. 2 (4), O'Reilly & Assocs., Cambridge, Mass. (1997).
- A. Herzberg, "Safeguarding Digital Library Contents: Charging for Online Content." D-Lib Magazine, January 1998.
- 11. J. Neilsen, "The Case for Micropayments," available at http://www.useit.com/alertbox/980125.html.