A Brief History of Personal Computing, Part IV:

"How the Web Was Won"

By Retired Major Dale J. Long, USAF

This is the final installment in the "History" series. In the first two installments we looked at the development of personal computing hardware and software. In the last issue we retraced the development of the Internet. In this issue we will close by tackling the World Wide Web (WWW), the graphical overlay of the Internet that represents the current state-of-the-art for collaborative personal computing. The WWW has become a ubiquitous part of the computing landscape in an amazingly short time. It has only been 10 years since Marc Andreessen and others at the University of Illinois National Center for Supercomputing Applications (NCSA) released the first alpha version of "Mosaic for X" in February 1993. Mosaic was a Unix application with a graphical browser. This is most people's earliest memory of the WWW. It was not, however, the beginning of the Web.

Before we launch this discussion we should review two things: first, what is the basic difference between the Internet and the World Wide Web? The Internet is a backbone system. It connects us, and our computers at several basic levels (physical, data, application, transport, etc.), but it does not normally operate at higher levels of information cognizance. The Internet, at its core, is still a hardware-intensive, hierarchical, text-based environment. The WWW, on the other hand, was developed as a way to connect pieces of information using the Internet as a transport medium. It has added functionality to the Internet by enabling and expanding non-hierarchical functions like indexing, cross-referencing and complex page design through the use of graphical browsers and related development tools.

For the most part it is a symbiotic relationship. The WWW could not have come into being without the Internet; the Internet would never have achieved its current level of explosive growth without the WWW. While some may argue that this growth has been more intrusive than useful, the fact remains that the WWW has become a dominant force across the planet in far less time than any other major technological innovation since cavemen discovered that fire was hot and wheels roll.

Our second review includes the three basic ways by which humans navigate through information: hierarchy, index and cross-reference. If we think of this as a "city" metaphor (and we will look at a virtual implementation of this later in the article), information pathways become like the streets of a large city. Finding things by following a hierarchy (e.g., table of contents) is like knowing the city's streets well enough to navigate by memory. Finding things by knowing an address within a grid coordinate and zeroing directly on your destination is akin to using an index. And arriving at a location that presents you with choices of other interesting places to go resembles cross-indexing. This last method covers both preset links that page authors include and lists of items generated by search engines.

As we found with the first three History topics, the roots of the Web run much deeper than just the release of a single application like Mosaic, which was only made possible by technologies, concepts and standards that took many years of development. Also, as with the other articles, we will spend more time on roots and causes than on modern issues that have been talked to death in recent news. So once again let's wind up the Way Back Machine for a trip back to the 19th century to retrace the steps that brought us to today's hyperlinked Web.

Hypertext in the Paper Age

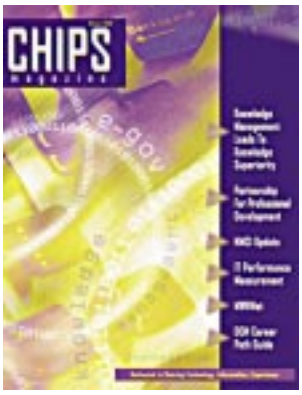

We start in the mid-1880s, where we find an ancient (by technology standards) device known as the telautograph. After Alexander Graham Bell patented the technology for transmitting voice over the telephone, inventors became interested in the idea of transmitting handwriting by wire. The telautograph reproduced handwriting and drawings by transmitting the movements of an electromagnetically-controlled pen along a line to a similar pen at the receiving end. Elisha Gray, the man who lost the telephone patent to Bell, was the first to develop and patent a practical version of a telegraphic writing machine, and it was Gray who coined the term telautograph. Telegraphic writing was allegedly quite a sensation at the 1893 World's Fair in Chicago, and two years later an improved version of the machine transmitted handwriting a stunning 431 miles from Cleveland to Chicago.

Using a telautograph, it was technically possible for one person writing from a single location to replicate his output simultaneously at any number of receiving points, limited only by how many connections were available. This is the earliest known practical implementation of a function we now see daily in Web-based chat rooms. The telautograph managed to find a niche next to the rapidly expanding telegraph and telephone industry, mainly in areas where high noise levels made using a telegraph or telephone impractical. For example, telautographs were often used in railroad stations to keep baggage and mail handlers informed of train movements.

The telautograph also had military applications. An early high point was its selection by the U.S. Army in the late 1890s for fire-control communication in America's coastal defense system. Before the advent of air power and submarines, the main defenses America relied upon to protect against enemy attack were the coastal artillery batteries. The guns were aimed on the basis of

data received from observers stationed some distance away, so the Army needed a reliable method to transmit the data. However, the noise in the gun pits was, as you may expect, quite deafening when the batteries were firing. Using a telephone or telegraph under these conditions was not practical. So, telautographs were installed in most important American coastal forts on both coasts. The military version of the telautograph was designed for ruggedness and reliability. The receivers were enclosed in heavy brass, waterproof cases suspended on shockproof mounts. Messages appeared behind a plate glass window, allowing the operator to read the messages without opening the case. An electric bulb inside the case allowed night reading. However, none of these coastal guns were ever fired at an enemy, so this first implementation of hypertext was never tested in battle.

OK, you're probably thinking, "That's an interesting bit of trivia, but what does it have to do with watching QuickTime movies of Super Bowl commercials on my Web browser?" While the telautograph may not be the epitome of graphical systems, it was the first. And it did produce some ripples in the technology pond decades later during World War II.

Hypertext in the Pre-Internet Age

The basic concepts of what we now consider hypertext were first formally proposed during World War II, long before computers were an integral part of our landscape. In an article entitled, "*As We May Think*," in the July 1945 issue of *Atlantic Monthly*, Dr. Vannevar Bush (http://www.iath.virginia.edu/elab/hfl0034.html), outlined his ideas for a machine that would be capable of storing information in a way so that any piece of information could be linked to any other piece. Bush, who was serving at the time as President Roosevelt's science adviser, called his system "Memex." His plan included references to associative indexing, which he defined as a process, "whereby any item may be caused at will to select immediately and automatically another." He also wanted the ability to create a trail of traveled links that a user could later retrieve, much like today's browser history files.

What sparked Bush's ruminations on linking information? One of his inspirations was apparently the telautograph machine. Bush believed that some future version of the telautograph would allow people to comment and make notes on documents without regard to distance. Perhaps he had a vision of the modern whiteboards used today for virtual two-way collaboration, even though the technology was not available at the time. However, though Bush did outline the important concepts, he did not specifically use the term hypertext. Theodor "Ted" Nelson, first coined the terms "hypertext" and "hypermedia" in a paper for the 1965 Association of Computing Machinery National Conference. In later writing, Nelson explained: "*By 'hypertext' I mean non-sequential writing — text that branches and allows choice to the reader, best read at an interactive screen.*"

Nelson later went on to found the Xanadu project, a wonderful idea about building a *"magic place of literary memory where nothing is forgotten."* It was supposed to become a universal, public, hypertext library that would spur the next phase of human evolution. Unfortunately, it became a 30-year vaporware project that never quite lived up to its lofty ideals. Xanadu was the great hacker dream to provide a universal library and collaborative editing with the ability to trace the changes in documents through successive versions. It was to have a means to track and credit authorship, a royalty system and non-sequential writing. Some of these capabilities exist in some form in other systems today, but the combined capability of the original plan has never been achieved completely by any single system. However, like climbing Mt. Everest or flying to the moon, Xanadu gave people a goal to shoot for, as we will see later.

In 1967, IBM sponsored development of the first actual hypertext-based system at Brown University. A team of researchers led by Dr. Andries van Dam developed the Hypertext Editing System, which ran on an IBM/360 mainframe. IBM later sold the system to the Manned Spacecraft Center in Houston, which reportedly used it for the Apollo space program documentation.

The next significant advance came in 1972, when researchers at Carnegie Mellon University developed ZOG. ZOG, by the way, is not an acronym and I did not find any specific references for why the name was used. The only historical reference I could find that might have inspired the name was: Ahmet Zogu, Zog I, King of the Albanians prior to Albania's subjugation by Italy during WW II. It's a long shot, though.

ZOG, the system, was a large database designed for a multi-user environment. It was a text-based system that used a basic, frame-based style sheet format that included a title, a description, a line with standard ZOG commands, and a set of menu items (called selections) leading to other frames. ZOG was a collaborative work tool that allowed users to modify the contents of a frame and make the changes visible immediately to other users through dynamically updated links. The U.S. Navy deployed ZOG in 1982 on the nuclear-powered aircraft carrier USS Carl Vinson (using PERQ workstations) to help automate certain management functions on the ship.

The epitome of pre-WWW hypertext applications was Apple Computer's HyperCard released in 1987. Arguably the most popular desktop hypertext application of all time, HyperCard was bundled free with all Macintosh computers. It was the first truly popular hypertext application to find a home on personal computers and contributed a great deal to the popularization of the hypertext model just in time to get people ready for the World Wide Web.

Spinning the Web

As with many great advancements in human history, the World Wide Web was created because of a desire for convenience. Tim Berners-Lee is credited by almost everyone as the Father of the World Wide Web. Trained in physics at Oxford, Berners-Lee was working at the Swiss-based European Particle Physics Laboratory (CERN - "*Conseil Europeen pour la Recherche Nucleaire*") in 1980 when he began a nine-year journey toward the specifications for a "global hypermedia system." His vision started out modestly enough: make his daily schedule planner, list of telephone numbers and documents all available through a single interface. At the time, CERN used a variety of platform dependent and proprietary information storage and retrieval methods. In addition, there were "in-house" systems that

were unique to CERN. As with most other organizations of that time, data was stored and manipulated in isolated machines with no real direct interaction or connectivity.

Berners-Lee's data was scattered over several such systems. What he wanted was a system that could store random associations between pieces of information based on his perceptions of their actual working relationship to each other. For example, he wanted to be able to pull mailing addresses out of his address list directly into letters on his word processor. We take mail merge for granted today, but Berners-Lee had to launch an entire technological revolution to get it. However, CERN didn't have Internet connectivity in 1980, so Berners-Lee's first attempt at connecting data was simply an attempt to unify personal data on the CERN systems. His initial system, developed around 1980, was called "Enquire." The name was allegedly based on an 1856 Victorian-era *how-to* book titled, "*Enquire Within Upon Everything*." Enquire met the modest functional goals he set for it, but when Berners-Lee left CERN shortly after building the system, it fell into disuse — one more homegrown system that died without its champion.

While Berners-Lee was elsewhere musing about data linkage and Ted Nelson's Xanadu project, CERN embraced TCP/IP (Transmission Control Protocol/Internet Protocol) and connected to the Internet in 1984. By 1989, when Berners-Lee returned, CERN was the largest Internet site in Europe and was heavily invested in both distributed and object-oriented computing. In short, they had finally caught up to where Berners-Lee had been seven years earlier.

It was at this point that Berners-Lee developed his concept of a "World Wide Web" (he coined the term in 1990) that would allow far-flung researchers (he saw this as a boon to scientists, not Web advertisers) to collaborate on large problems. The result of his work brought Vannevar Bush's ideas for Memex to life and fulfilled much of what Ted Nelson tried to achieve with Xanadu. In March of 1989, Berners-Lee submitted his first paper outlining his strategy for a global information system to his bosses at CERN titled: "*Information Management: A Proposal*." A later paper, titled "*World Wide Web: An Information Infrastructure for High-Energy Physics*," is even more specific about his intentions and motivations. Note that the title indicated his belief that the endeavor of building the Web was primarily in support of scientific study.

Berners-Lee believed that the motivation for this system arose "*from the geographical dispersion of large collaborations, and the fast turnover of fellows, students, and visiting scientists.*" He wanted to create an information environment where these transient people could use the Web to quickly integrate into projects and make lasting contributions by contributing to the assembled knowledge. In his original "*Information Management: A Proposal*," Berners-Lee described the deficiencies of the hierarchical information delivery systems in use at the time and argued the advantages of hypertext-based systems. He proposed that CERN incorporate several different servers of machine-stored information and a distributed hypertext system to provide "*a single user interface for many large classes of stored information such as reports, notes, databases, computer documentation and online systems help.*"

The key capabilities of his system would include a protocol for requesting readable information stored in remote systems. This was fulfilled by the development of Hypertext Transfer Protocol (HTTP). The Web also needed a common format for information exchange between information suppliers and consumers, and some method for reading text and graphics at the same time. The solution for this was Hypertext Markup Language (HTML), developed as a subset of Standard Generalized Markup Language (SGML). SGML was a well-developed publishing standard already in existence, but it was originally considered too complex for the Web.

Shared document libraries were also a key part of the plan, and Berners-Lee expressly wanted users to be able to add to individual libraries, as well as their own. Tying all the collected knowledge together would be links between documents (or even from within documents) in one library to documents in any other connected library. The enabler for this was the Uniform Resource Locator (URL). A URL syntax explicitly describes the unique location of every site, library, document and element of information with an independent existence on the Web. Finally, having all that information available would be fruitless without some ability to find what you wanted. URLs may describe everything, but unless you already know how to navigate hierarchically to the exact street you want, navigating the Web without an index would be like navigating an unfamiliar city without a map.

Berners-Lee envisioned the ability to index all documents in all libraries for retrieval by keyword search. Inverse indexing, essentially the ability to find and record the location and frequency of every word in any document, evolved rapidly on the Web as search engines like Lycos and WebCrawler provided new Web users with quick ways to find and retrieve information. The initial development of the Web was a two-phase project. In the first phase, CERN made use of existing software and hardware as well as implementing simple browsers for the user's workstations.

However, there wasn't this type of software already in common use, or was there? Most people remember NeXT Software, Inc. as what Steve Jobs did while he was waiting for Pixar to become a leading computer animation company so he could regain the reins of Apple Computer. Jobs, the creative force behind the Macintosh computer, founded NeXT (after being ousted from Apple) to bring object-oriented, high-end computing to the masses. The masses, and even most computer professionals, weren't ready for it and the NeXT brand became one more failed system that was too far ahead of its time.

However, NeXT survived on the Web. CERN developed the first hypertext graphical browser in November 1990 using NeXT's object-oriented technology. In addition to viewing HTML instances, it was also a "*what you see is what you get*" Web document editing application. The first WWW server was also developed and implemented on a NextStep computer. Web software was ported to other platforms in 1991 and released to the public. Berners-Lee and CERN changed the face of the Internet with their server, browser, HTTP, HTML and URLs.

The Birth of the Browser

Once CERN placed the WWW concepts and the protocols in the public domain, programmers and software developers worldwide began making their own contributions. Marc Andreessen was one of them. Andreessen was a graduate student at the University of Illinois National Center for Supercomputing Applications (NCSA). Andreessen led a team of graduate students who released the alpha version of Mosaic for X in February 1993. As I mentioned at the beginning of the article, Mosaic was originally a "point-and-click" graphical Web browser that ran on Unix. What really helped open the floodgates, however, was six months later in August 1993 when Andreessen and his team released free versions of Mosaic for both the Macintosh and Windows operating systems. This was a milestone in Web development. It was the first time a WWW client with a relatively consistent and easy to use point-and-click graphical user interface (GUI) was implemented on the three most popular operating systems available at the time.

Andreessen later helped found the Mosaic Communications Corporation which eventually became known as Netscape. By May 1994, practically all the members of the original Mosaic development team at NCSA had joined Netscape. There, the team decided to completely rewrite the underlying code of the Mosaic system, creating an entirely new browser that quickly became the most popular Web application in the world. The Microsoft Internet Explorer (IE) Browser later supplanted Netscape as the No. 1 browser on the Web. Ironically, the core code for IE was from Andreessen's original NCSA Mosaic project, purchased from NCSA by Microsoft.

The main remaining barrier to Internet entry fell when America Online, the most popular single online service in history, included a Web browser with their service. Swarms of "newbies," their modems plugged into RJ-11 outlets, mice in hand and armed with GUI browsers, rapidly entered this strange and wonderful new world. Much like the Old West, once the pioneers paved the way, civilization invaded and the Internet changed forever as the Web blanketed it like kudzu. Given the success of the Web, it was only a matter of time before it grew to a size that warranted some form of management thus the World Wide Web Consortium (W3C) was formed in December 1994. W3C's main objective is "*to promote standards for the evolution of the Web and interoperability between WWW products by producing specifications and reference software.*"

The Web Today

The issues facing the World Wide Web today are surprisingly similar to the ones that spawned its creation. Ideally, it is an interface that will unify all our data, information, and knowledge sources on desktops without regard to what type of computer, software or operating system you use. Of course, there are many companies that do care quite a bit about what type of computer, software, or operating system you use — preferably theirs. Unless you've been living in a hermitage, you have read or heard about things like the browser war between Netscape and Microsoft, the competition for dominance between Java and ActiveX, what's happening with Extensible Markup Language (XML), and whether the Web will remain in the public information space or become a commercial entity.

Unfortunately, I'm about out of space here, so we'll have to wait until some later date to move on to a "*Future of the Web*" article. If you want more information about the current state of the Web, I recommend you visit the W3C Web site at http://www.w3.org/.

Final Thoughts

In 1900, the military transmitted firing coordinates to gunners using telautographs sending electronic handwriting. Today, the crew on the destroyer USS McFaul uses PDAs (personal digital assistants) continuing the military revolution in hypertext started in those gun pits 100 years ago. If you are lucky enough to be part of the McFaul project, take a moment to consider how those turn-of-the-century troops in their gun pits felt when the quartermaster delivered those thick brass telautograph cases with wires sticking out one end and a newfangled light bulb illuminating the interior. It may have been the same feeling we get when someone hands us a PDA, the screen flickers on, and we realize that we are now connected to something a lot bigger than we are. That's what the leading edge feels like. Get used to it, because at the rate networking and the Web are evolving we will be riding on the edge for the foreseeable future.

Finally, I find it fascinating that the birth of the Web was simply a side effect of the elite culture that developed over four decades ago at CERN because of a desire for a better way to do particle physics research. There was nothing in the plan about most of what we use the Web for today — and they had no intention to profit monetarily from their innovations. It was science for science's sake — altruism at its finest. While the rest of us may never achieve the level of intellectual acumen of a nuclear research facility, we have managed to do some pretty useful and interesting things with Web technology. I just wish, however, I could get the image out of my mind of prehistoric monkeys staring at the big, rectangular monolith from "*2001: A Space Odyssey*."

Oh well, it's time for this monkey to go Web surf for some more movie trailers.

Happy Networking!

Long is a retired Air Force communications officer who has written for CHIPS since 1993. He holds a Master of Science degree in Information Resource Management from the Air Force Institute of Technology. He is currently serving as a telecommunications manager in the Department of Homeland Security.

To become a CHIPS subscriber log on to the CHIPS Web site at www.chips.navy.mil or e-mail us at chips@spawar.navy.mil.