Chapter 1: Living Ground

Cyberspace is a new world.

As our own space age began, we built craft and sent them forth into the entirely uncharted regions of the space above our heads. These craft carried cameras and antennas - literally the eyes and ears of a humanity hungry for news from these unknown lands. First we reached toward the Moon - Earth's own child - then our neighbors, Venus and Mars, then out, beyond the asteroids, to Jupiter, Saturn, Uranus and Neptune. At the start we didn't even know what to look for - we were just looking. But as we grew rich in images and dense with information, we could begin to cast our eyes backward, and imagine times past - based upon times present.

Quite often, our imaginations fixed upon the planetary neighbor most nearly like ours - Mars. Possessing an atmosphere, reasonable sunshine, and even polar ice caps, it seemed well within plausibility that life could gain a foothold there. So, in 1975, we sent out the *Viking* spacecraft, to land on Mars' surface and scout for life - only to find a dry, rocky, barren surface, absolutely sterile. The Mars of modern times did not teem with life; but perhaps, thought scientists, an earlier Mars had.

In August of 1996 - in the same week that the VRML 2.0 specification was released - NASA, in a maneuver worthy of *The X-Files*, announced that a meteor thrown from Mars to Earth as the result of some great planetary impact indicated the "strong probability" of life on Mars - in some distant past. The world next door suddenly assumed a different history - once the twin of Earth's earliest days, but long since drained of atmosphere, and hence, now unable to support life.

On the heels of this announcement, another major discovery rocked the scientific world. Biologists indicated they had identified a *third* basic family of living organisms on Earth. Previously, bacteria - that is, singled-cell organisms - and multi-celled organisms made up the two major families of life. To this, scientists now added *akaryotes*, first discovered in deep-ocean thermal vents, thriving in an superheated, corrosive, oxygen-depleted environment considered too hostile for any living organism.

Now scientists took another look at Mars; perhaps, they reasoned, these same forms evolved on Mars, and have secreted themselves away, in the deep soils of the Martian plains, safe from the elements we once thought so essential to their survival. We have recently learned that the Earth's crust ubiquitously teems with bacteria that somehow manage to live in solid rock, literally carving out an ecological niche from the bowels of the planet. Perhaps our neighbor enjoys a similar vitality.

With every discovery, another world has opened up to imagination, investigation and exploration; only our thinking had held us back, and - now freed from conceptions that seemed reasoned, but actually reflected our native prejudices - we begin to see the universe as teeming with life.

Just a week before I set these words down, NASA put forth another shocker - one which looks forward into the future of the solar system. It seems that the Jovian satellite of Europa has a planetary ocean of water, protected by a thin crust of ice. Some of the ice, broken into floes, reveals an endless sea - and this sea the spacecraft *Galileo* has found "indications" of life's earliest signatures. We can't know for sure until we land on Europa - something that couldn't happen for at least another ten years - but, in the most amazing reversal of modern scientific thought, we've gone from seeing ourselves as unique and alone in an empty universe to just one instance - albeit a highly evolved instance - of life, in an ocean of life. It's a new world, both here on Earth, and in the heavens above.

The explosive growth of the World Wide Web is like the discovery of a new world, populated by the content of our own teeming imaginations. With so much discovery and innovation in the world today, it's easy to become overwhelmed by the rapid progress of technology; we forget that in each of these technologies, we see a reflection of ourselves.

One of the latest of these innovations, the *Virtual Reality Modeling Language* (VRML) has itself achieved a sudden and broad popularity. Long before companies had released any VRML "browsers" or VRML editors, the press and the Web community began to act like kids on Christmas Eve, filled with expectations - visions of virtual plums dancing in their heads. To many it looked like the beginning of the cyberspace revolution, but that story stretches back through a far longer expanse of time. The story of cyberspace is the story of human communication.

Communicating Imagination

Since the dawn of consciousness in *Homo Sapiens*, we've been expressing our thoughts and feelings to each other. That is the largest part of our success as a species; because we communicate, our ideas can outlive ourselves. In the oral era - before the Sumerian invention of cuneiform - we kept our learning in mythic forms. Myths are a shorthand, encompassing a universe of meaning in just a few words. Myths are the stories before history.

In this age, we think of mythology as unscientific, irrational, unreasonable. Myths are none of these things - they are a human attempt to speak about things for which there are no words; the joy of birth, the valor and honor of heroic deeds, the mystery of death.

Finally, the dawn of human communication brought us imagination. Myths are never literal. The figurative power of mythology casts this *as* that, not this is that. The ancients understood their imaginations, and understood their myths as the essential forms of those imaginations.

Imagination Electrified

In the hundreds of thousands of years since that time, we've gotten better and better at our forms of communication. We invented writing, which extended our ability to retain information to vast quantities of information. Take a look at a cuneiform tablet or hieroglyphic monolith - you're looking at the high technology of the 20th century, B.C.E., and the ancients felt the same wonderment looking upon those marks that we get when we stare at the rainbow of exposed silicon on a memory chip or microprocessor.

It took hundreds of thousands of years to get to writing, but only a few thousand to get to the *Phoenician Alpha-beta* (phonetic alphabet), and a few thousand more to get to the printing press. When Gutenberg printed his first Bibles, he lit a fire that transformed all of civilization from a primarily oral and aural culture to a literate and visual culture. The human riches of poetry, drama, and song became the fertile grounds of mathematics, physics, and biology. These fields had existed before the invention of printing, but now any individual could share Newton's thoughts, or Harvey's or Linneas'. To stand on the shoulders of giants meant to prop oneself up on a shelf of books. These individuals might then write books; Plato and Democritus would influence Voltaire, who would then influence Rousseau, who would then influence Thomas Paine, Benjamin Franklin, and Thomas Jefferson. In this way, the explosion of knowledge in our modern era began.

Franklin himself studied physics avidly; his electricity experiments became known world-wide. Half a century later, Samuel Morse, who studied the works of Franklin, Volta, and Ampere, would develop a device that used electricity to send words, as fast as a beam of light. He called it the *telegraph*, Latin for "distant writing."

With the invention of the telegraph, human history changed suddenly and completely. For the first time, messages arrived instantaneously. With this, all media took on a new form; we saw the birth of the modern newspaper, which thrived on news reports from distant lands brought to one's home town through the magic of telegraphy. The telegram itself acquired a magical and foreboding quality - to receive one was akin to being struck by lightning, and could announce the visitation of a rich supply of either good or bad fortune.

The telegraph defines the Victorian era; Great Britain used the telegraph to give form and backbone to their colonial policies. India was mapped and railed and telegraphed by the British - this gave them an indomitable hold on the people of that nation. This hold continued unchecked until the press reported - via telegraph - the actions of a single man, Mahatma Ghandi. With little more than a pure heart and publicity, Ghandi brought the British empire to its end.

Electronic Computing

Even before the beginning of the electric era, we constructed enormous machinery, such as textile mills and locomotives. This machinery augmented human power, making possible tasks that would otherwise require an army of slaves or serfs, or wage laborers.

Despite this, the machinery was unreliable and often dangerous to the people who worked with it, because it had no understanding of itself, no mechanism to monitor its progress, and no sensors to detect problems. James Watt, who perfected the modern steam engine, invented a device called the "governor" - govern in its original sense means *regulate* rather than *control*. The governor regulated the output of the engine and kept it in balance, so that it would not destroy itself from its own fires. Watt created a system that took its own past results fed these outputs back into the system (hence the term *feedback*), to create a system which could self-regulate.

A few years later, British mathematicians Charles Babbage and Ada Lovelace designed the "difference engine," a mechanism which could perform numerical calculations and - embodying the same principle of self-regulation as Watt’s steam engine - feed intermediate calculations back into the engine to further modify its outputs. The difference engine is considered the forerunner of the modern computer.

A hundred years later, Alan Turing used this peculiar quality of the "thinking machine" to break the codes used by the German High Command in the Second World War. The Germans used a device, called ENIGMA, which could generate a "scrambling code" that was used to encrypt messages sent to submarines, spies, and so forth. The code could quickly be changed, and although the British captured one of the ENIGMA encoders, the Germans quickly modified the design to prevent reverse engineering. Turing, a mathematician and computing theorist, developed a set of logical steps that could, within a few hours, "break the code" of the ENIGMA and render the messages readable. It is believed that this breakthrough considerably shortened the duration of the war.

Turing relied upon the essential mutability of his computer - which was not much like what we consider a computer today - it could change itself, recharacterize its inputs and outputs, based upon its own inputs and outputs. The computer could make decisions, could change its behavior, and act as if it were, in some small way, conscious. At about the same time, working at Princeton, John Von Neumann developed the basics of the architecture used by almost all computers - an arithmetic unit, a decision unit, and a memory unit. These three units functioned together, each modifying the content and behavior of the other, creating the paradigm for modern computing.

After the war, the still-young field of electronics coupled with the brand-new discipline of computer science to create the electronic computer. These computers were used to automate huge decision-making tasks like the U.S. Census.

It's taken about fifty years, but computers have now become completely ubiquitous; we rely on a computer to keep the wheels of our automobiles from locking up in a skid - they talk to each other and adjust their own behavior. We use them in microwave ovens to cook our food just right. We use them in our bombs to make sure they hit their designated targets, and nothing else.

Computing Communication

The essential nature of the computer is as a simulator. The computer knows nothing of itself - it is nearly completely innocent, but, when filled with rules and data and sensations provided by a scientist, a nurse, or a video game player, the computer creates a *simulation* - an understanding of a situation - and then uses its rules to bring that simulation forward in time. Will this steel be strong enough? Will this patient survive? Will I make it to the next level of *QUAKE*?

Yet, isolated, cut off from the outside world, computers are very poor simulation engines. Simulation is based in reality, or an approximation of reality. The more communication you have with the real world, the more believable, the more accurate, and the more thrilling your simulation is likely to be.

Your weatherman works with simulations all of the time. A network of satellites observe current weather conditions, and pass this information along to supercomputers which attempt to simulate future weather conditions; working together, satellites and computers create the forecast you read in the newspaper or see on the television. In a vacuum - without that sophisticated infrastructure of electronic eyes and brains - that forecast wouldn't be very accurate. Even with the best information, the forecasts can be off base, but the accuracy of a near-term forecast has gone way up, primarily because we have a network of computers which communicate what they know (or predict) about the weather to each other.

To facilitate this coordination, to improve the quality of our simulations, we have taught the computers to talk to each other, so that they can act in concert, each modifying the others through a complex relationship of messages and behaviors. It's a *society of machines*, a society built upon communication and cooperation in a group context. In essence, we're imparting to our machines some of the basic attributes which make us social animals. We've only just started to do this - the Internet has been around for little more than a quarter-century - but, as this society of machines evolves into a social ecology, we'll no longer think of computers as isolated. We'll think of them in concert, each like a neuron in a brain.

There is a natural - and surprisingly circular - progression at work through these ages. First man communicates his imagination, then imagination is electrified with the telegraph. Next, mutability becomes the quality of electronic computing, and the loop closes with computer communication.

We can see that we're about to come full circle, back into the era of the communication of the imagination. That's what cyberspace is really all about. The content of cyberspace - what we put there - is bounded only by our imagination. Cyberspace is *shared imagination*, using as its ground electronic computing communications.

Cyberspace stands above this wheel of media, using all of them.

A Great Design Criterion

The Internet we know today started as a a fallout shelter for computers.

In the late 1960s, the U.S. Department of Defense began a detailed investigation of methodologies to protect military computing systems in the event of nuclear war. These systems - mostly very large and well-guarded - were rapidly becoming the backbone of the nation’s defense strategy. Using these machines in a coordinated fashion would produce a more accurate simulation of the defense posture of the United States, as well as the attack posture of its enemies.

This communications network, at once a great strength, had a very natural Achille's heel - cut the cables and the computers would stop talking. In the event of war - real war, with multimegatonnage blasts over major cities and military installations - the defense computing network would quickly collapse, and with it America's ability to simulate and monitor any threats.

Further, the burgeoning military-industrial complex had given the U.S. armed forces a vast array of incompatible hardware. Computers from IBM, Univac, Sperry, and Burroughs all fit into the defense infrastructure. They couldn't talk to each other, though, so the defense planners were faced with a choice: either all of the computers in America's defense infrastructure would have to be of the same design - and from the same company, most likely - or they would need a way to make many heterogeneous systems talk to one another. The planners opted for the latter solution, paving the way for a network of many types of machines, all speaking the same language or protocol.

The planners now moved from the design to the prototyping stage. On October 27, 1969, two computers began talking to each other across a line leased from the telephone company. This network was named ARPAnet after the agency that funded it - the Advanced Research Projects Agency. ARPAnet used an "abstract" protocol - tied to no particular computer hardware or software - that later became known as *Internet Protocol* (IP).

Internet Protocol got its name because it allowed sites that already had networks - rare except in high-end military environments - to provide a "gateway" into a "network of networks," or an Inter-network. The Internet has always been a collection of individual networks that have agreed to speak the same protocol with respect to each other, in much the same way that all of the states in the United States enforce the same laws with respect to the national government. A state has broad powers internally - as long as they do not violate conditions laid down in the Constitution - but is strictly regulated in how it deals with its neighbors. The Internet is a greater ocean composed of a sea of networks, island states.

The computers could send messages to each other across the Internet. A message on the Internet is called a *packet*, and it is the basic unit of communication. Computers communicate with each other seamlessly because these packets direct themselves - just as letters do in the Post Office - through the network. This is called *routing*. Using routing, the computer simply puts the packet onto the Internet, and the Internet ensures that it is delivered. The path between two computers is unimportant, as far as the computers are concerned. That path can be very complex, and it can even change in the middle of a conversation between two computers.

In the middle of a war, three computers - in Pasadena, Colorado Springs, and Cambridge, for example - might be assigned to tracking and targeting incoming missiles. If the computer in Colorado Springs suddenly went off net because of a 20 megaton explosion in its immediate vicinity, ARPAnet would dynamically adapt to the change in routing patterns, and send the packets for the computers in Pasadena and Cambridge (for as long as they lasted) over a different path, to other computers kept around for contingencies such as this. ARPAnet observed itself, monitored its own behavior, and fed its behavior back into itself, to recover from network failures.

To ARPAnet, nuclear war is no more than a bunch of routing errors that it could immediately correct.

Through years of trial and error, network planners have learned that whatever can happen to a network will happen. This means that disaster planning is the primary activity of a network designer, and nuclear war is clearly the ultimate disaster. When ARPAnet solved this worst-case in disaster planning, it gave itself the best possible recommendation as a reliable networking methodology. Internet Protocol, built on this foundation, slowly grew to become the dominant networking paradigm. The Internet ate network after network, adding them to its domain, giving all of the computers within its breadth the ability to talk to each other in a standardized, universal, fault-tolerant manner. These advantages prompted network administrators and planners worldwide to adopt IP as their long-term networking solution. Over the last decade, most networks have integrated into the Internet.

In short, nuclear war is a great design criterion for networks.

Computer-Centered Communication

While the Internet made it easy for computers to communicate with each other, its designers paid very little attention to the humans who would use it to communicate through those computers. Internet comes from an age where very few computers had interactive qualities; most of them ran in a "batch" mode - you gave the computer a set of commands, usually in the form of a stack of punched cards, and came back later - a few minutes, hours, or days later - to look at the results of your commands. Most of these commands had been designed around the computer, making it easy to "parse" or translate those commands into operations the computer could then execute. Very few people knew how to do this, and they were the white-coated "priesthood" of computing, the guardian of its terrible secret: computers had poorly designed interfaces, which made them extremely difficult to use. A bad interface disguised itself as sophisticated.

Over the last thirty years, we've moved from the sophisticated to the intuitive computer interface. Perhaps the most important thing we've learned is this: the easier it is to use a computer, the harder it is to design an interface for it. Developing applications for the Macintosh, which lead the ease-of-use movement in computing, was many times more difficult than for its competitor, the IBM PC. Every operation had to be thought through from the user's perspective, and that takes a lot of up-front design work.

The priesthood claimed computers were too sophisticated for most people to learn or use. In truth, the computers were spectacularly stupid, so much so that we would have to learn baby-speak - the computer's control language - to have any conversation with them at all. Most people found this unacceptable. The few who could tolerate the endless error messages, the burps and whines and diaper-changing of infantile computer interfaces, had good jobs managing the informational infrastructure of big companies or big governments.

The Internet inherited this user interface legacy; most commands were completely cryptic, like "ftp 192.100.81.101," "rlogin 92.2.3.1," or "ping 172.27.31.30." A lot of this was the legacy of UNIX, a cryptic operating system itself. As is commonplace in UNIX, most of these Internet-related commands had dozens of options - ways to massage the computers into successful, happy communication. Very few people knew them all, and even the most expert occasionally consulted a manual.

More than that, there was no Internet "road map." Many computers with many different types of resources sprouted up all over ARPAnet. Keeping track of any of these resources - which could move or change without warning - became almost impossible. The Internet then added the *Domain Naming Service* (DNS), which morphed the computer's address into its name; for example, "192.80.57.1" became the more sensible "shiva.com," and "192.100.81.101" became "ns.netcom.com." Service naming helped - you could move a computer around, but keep its name the same, and folks could find it again. You could also move the service to another computer, and then map the name to that new computer; no one would be the wiser.

Still, enormous feats of memory were called for, and a legion of "system administrators" answered the call. Although it's generally thought that "sysops" maintained the physical and software infrastructures on computers, they also functioned as the agents of community memory. The system administrator remembered where you put things, such as your accounting records, the test data, or your user surveys. Few others cared enough to learn Internet’s difficult interfaces. Only sysops knew exactly where things your files were, once they got sucked into the Internet. They had a tidy little empire, managing the memory of cyberspace.

Despite this, the Internet proved to be very useful for many people, and the amount of information which could be accessed through it grew dramatically through the second half of the 1980’s. Soon, even sysops lost track of the whole Internet universe, and began to specialize, each focused on their own areas of expertise. The threads of community memory had begun to unravel; but within a few years, they’d be sewn into a new mosaic.

The Web that Ate the Net

The ARPAnet proved very useful for the research organizations and universities which had access to it; so much so that by 1987, the U.S. Government divided ARPAnet into two zones, or domains. MILnet, as the name implies, handled secure military communications, while NSFnet, sponsored by the National Science Foundation, oversaw the development of Internet infrastructure for academia.

Before this decade, the Internet was used by only a few million people, but with the birth of NSFnet came many new commercial and educational Internet users. Most of these individuals used electronic mail systems, which provided an Internet without the unnecessary sophistication of Internet interfaces. Using electronic mail is easy - compared to file transfer or remote access - and often desktop software seamlessly integrates with it. People use electronic mail on the Internet because they don't have to think about the Internet to use it. The popular conception of the Internet as a big postal service has remained constant through the last several years.

Despite its difficulties, the total content of the Internet constantly grew into a small universe of documents, images, applications, and so forth that could be retrieved through it. Without any way to navigate or grasp any sense of the totality of the Internet, most of these resources lay underutilized, discovered by only a few people.

This problem confronted a man by the name of Tim Berners-Lee. In the 1980s, Berners-Lee worked as a software engineer at the European Center for Particle Physics (CERN) in Geneva, Switzerland. CERN has the world's largest atom smasher - a huge ring that looped out from CERN, under the Alps, and back again. Scientists from all over Europe reserve time on CERN's accelerator to conduct experiments on particle physics and the basic construction of the universe.

The scientists, located at universities all across Europe, rarely wanted to travel to CERN to perform there experiments, and then return to receive the results of those experiments. Using the Internet, it was possible to send those results back to the researchers electronically. Berners-Lee realized, however, that they lacked context - how did this week's experiment by researchers in Milan relate to last month's experiment by researchers in Copenhagen relate to next year's experiments by researchers at Oxford?

The idea of linking different sets of data to show their relations with respect to each other - something we call *hypertext* or, more often these days, *hypermedia* - has been around for over thirty years. Douglas Englebart, a researcher at SRI in Menlo Park, California, demonstrated the first practical hypertext system in 1969. With Englebart's system, is was possible to group various items and link them together. Englebart demonstrated a hypertext shopping list; using a mouse - something else he invented while inventing hypertext - he could reorganize his shopping list or connect it to other documents within his computer. The ability to connect two items together is the essential feature of hypertext, and is called *linking*. He could link "butter" to the document "2 sticks, Land-o-

Lakes, $0.79”, or group “butter”, “milk”, and “cheese” into a “dairy” group. The system was very basic, it demonstrated all of the features essential to hypermedia systems.

Hypermedia systems languished; a command-line is unfriendly to hypermedia, and until the late 1980s, most computers graced us with command-line interfaces. *Hypercard*, the first widely-popular hypermedia application, showed the power of a contextual environment - how sharing and linking documents can create a whole greater than the sum of the parts - but Hypercard lacked support for networking. You could read the hypermedia version of "Jurassic Park" years before the novel made it to the screen, but you couldn't click on a link and check out the collection of the paleontology department at the University of Montana. Static hypermedia - constructed once and stubbornly persistent in the face of a rapidly changing world - proved only marginally useful.

Berners-Lee suspected that a hypermedia system fit for the physicists of CERN would have to tie together all of the information related to their discipline, wherever it resided - at CERN, Princeton, Stanford, Oxford, or wherever. In 1989, he developed a prototype of a hypermedia system that could fit the entire Internet into its scope; this system came to be called the *World Wide Web* (WWW).

The World Wide Web effectively turns the Internet into the equivalent of a large computer disk drive, or series of drives. These "virtual disk drives" have names - they are the names of the computers on the Internet. Within these computers, the World Wide Web can access various documents - images, sounds, text, each of which has a unique name. Together these components create a unique name for any resource on the Internet, and are called *Universal Resource Locators* or URLs (pronounced "U-R-L" or "earl").

The innovations that Berners-Lee added to the Internet to create the World Wide Web had two fundamental dimensions: connectivity and interface. He invented a new protocol for the computers to speak as they exchanged hypermedia documents. This *Hypertext Transfer Protocol* (HTTP) made it very easy for any computer on the Internet to safely offer up its collection of documents into the greater whole; using HTTP, a computer that asked for a file from another computer would know, when it received the file, if it was a picture, a movie, or a spoken word. With this feature of HTTP, the Internet began to reflect an important truth - retrieving a file's data is almost useless unless you know what kind of data it is. In a sea of Web documents, it's impossible to know in advance what a document is - it could be almost anything - but the Web understands “data types” and passes that information along.

Once the data has been retrieved through the Web, Berners-Lee knew it would be necessary to display that data in a way that would be universal and consistent. Like the Internet itself, the Web needed to be independent of any computer upon it, yet all of these computers used differing systems to format data for presentation. The Macintosh used QuickDraw, Silicon Graphics' workstations used OpenGL, and IBM PCs used a number of incompatible solutions. The Web needed a standard interface language so that if a researcher created in Chicago created a Web document for use at a laboratory in Toronto, it wouldn't be necessary to consider what computer the document would be displayed upon. Postscript, used in most laser printers, runs on most computers, but is difficult to use to show how one set of data links to another. Something new was needed - something that could reflect the capabilities of the Web as a medium.

For this reason, Berners-Lee and others developed the *Hypertext Markup Language* (HTML). HTML is a subset of the *Structured Generalized Markup Language* (SGML), which has been in use for about 20 years. Using HTML, an individual can create a hypermedia document that looks pretty much the same on any computer, anywhere in the world. For example:

```
<Title>Learning VRML: Design for Cyberspace</Title>
Learning VRML: Design for Cyberspace
```

This very basic HTML file creates a document whose title is “Learning VRML: Design for Cyberspace”. The content of the document is simply a that same text.

HTML made it easy to reference or "link" documents together - even documents with widely varying data types like text and movies - so that a researcher could link a set of experimental test results with other relevant experiments, scientific papers, images, and so forth. We can change this document so that it links to my home page on the Web, by using the “anchor” tag, which links one document to another:

```
<Title>Learning VRML: Design for Cyberspace</Title>
<A HREF=“http://www.hyperreal.com/~mpesce”>
Learning VRML: Design for Cyberspace
</A>
```

Clicking on the text within a Web browser will cause the browser to “jump to” my home page.

HTML provided an interface to the Internet that hid the cryptic commands behind the URL; the Internet, following the rest of computing, thus became a comfortable journey of pointing and clicking around a set of hypermedia documents with a mouse. It wasn't perfect - its major drawback was a lack of places to go - but it provided so much, for so little, that it was an immediate success in the academic community.

In the United States of America, researchers at the *National Center for Supercomputer Applications* (NCSA) found the World Wide Web to be a potent tool for sharing information over the Internet. NCSA has a mandate from the U.S. Government to research and prototype tools for collaboration, developing interfaces that enable researchers to co-annotate and co-author documents. The Web, with its inherent features of linkage and supple interface, appealed to them immediately.

In early 1993, the developers at NCSA - Marc Andreesen, the founder of Netscape Communications Corporation, among them - took the World Wide Web and again extended it in the two dimensions that had proven so fruitful for Berners-Lee: interface and connectivity. For the interface, they added the ability to put images inside of an HTML page, and beyond this, they created the *image map*. An image map "knows" when and where a user's mouse-click occurs, and sends that information out across the Web as part of a request for more information. Imagine a map of the Earth, and clicking on Cambridge, Massachusetts, then finding yourself at MIT's home page. This was a small improvement that yielded enormous benefits. They also added forms. HTML forms, like normal business forms, have places where a Web user can enter data within a Web page - perhaps to supply a mailing address or credit card number - and send back to a Web server for processing. This opened the Web up to two-way communication; you could change the way a Web site behaved by information you provided on a form.

The engineers at NCSA also added a mechanism to make the Web more active. Very often, Web pages are static - they almost never change. Yet the real world changes all the time. To tie the static world of Web pages into the dynamic world of real data, NCSA created the *Common Gateway Interface* (CGI). CGI created a dynamic Web; a document doesn't have to exist before you see it on the screen - many are created on-demand, in response to a specific request. This means that databases, control programs, and other useful applications can be integrated into the World Wide Web with only very minor modifications.

In July of 1993, NCSA released their World Wide Web server, and a Web "browser," which they called NCSA *Mosaic*.

NCSA's Mosaic is without question the most influential application ever developed for a computer. It made the Internet accessible to anyone who knew how to use a computer; the few hundred thousand who knew the ins and outs of Internet suddenly became the tens of millions who knew how to use a word processor or spreadsheet. Mosaic was actually substantially easier to use than the other applications people had been trained to use, so people adopted it rapidly. At least two million people used it on a day-to-day basis in the first year after it was first released.

With more Web users, the Internet began to see an explosive growth in Web servers. Only about 300 Web servers existed world-wide in October of 1993. Within a year, that number had grown to about 10,000. It's believed that there are at least 1,000,000 in 1997.

Four years after its introduction, Mosaic or an equivalent Web browser is in day-to-day use on almost every computer in the world - perhaps as many as *sixty million* of them. The reason for this is easy to identify: Mosaic increased the usability of the Internet by several orders of magnitude - no more commands; just point and click and you're off.

The real lesson of NCSA Mosaic is simple: increased usability always translates into increased usage. So, over the course of 1994 and 1995, the ivory towers of the Internet came tumbling down, and the hordes of individuals who had only experienced electronic mail suddenly found themselves in an amazing universe of sights, sounds, and creative collaborations.

Something else happened - unexpected, but wonderful: people began to create their own corners of the Internet, graced with their own style, sensibilities, and ideas. Authoring documents in HTML was not rocket science - although it still required some knowledge of computers to do it well - so hundreds of thousands of people rushed into the Web to create a human presence there. No longer just corporate data or science reports, the Web better articulates the entire range of human experience than any communications technology that preceded it. People fall in love with the expressiveness of the Web, and so they use it, add to it, and talk about it endlessly. In less than two years, World Wide Web protocols became the dominant form of communication over the Internet - the Web ate the Net.

There's no *There* There

The World Wide Web has created the perception of a unified Internet; it's quite impossible to tell where Web data comes from - across the street or across the world. With HTML, it all looks pretty much the same. There are conventions of style and layout - many pioneered by Kevin Hughes when the Web was still very young - that have begun to look like a universal interface to the global information repository. It's not perfect, but the Web grows richer and increasingly more useful. Periodicals like *TIME, WIRED* and *OMNI* are accessible through the Web, and the United States Library of Congress, the largest book repository in the world, has begun to provide its texts across the Web.

The major issue confronting users of the Web concernes locating relevant data. In a sea of information, unbound to any system of organization, how can one find anything?

CERN started to tackle this problem by creating the WWW Meta-index, a Web page which contained a partial listing of the resources available on the Web, categorized by subject. CERN had a dedicated staff of Web “librarians” who kept the list of subjects and sites up to date in the Web’s early days. However, as the Web started growing uncontrollably, this list first became unusable - too big and only approximately accurate - and later became obsolete. There simply weren’t enough librarians to track the growth of the Web.

Other indexes, such as Stanford's *YAHOO*, sprang up, offering access thousands of categories of data. YAHOO offered two interfaces; a subject catalog, similar to a library’s, but with hyperlinks, and so much easier to use. YAHOO also provided a search interface, using HTML's forms capability. Just tell YAHOO what you wanted to look for and - presto - you’d have a page of links to relevant entries inside YAHOO’s database. But - with an exploding Web doubling every 53 days in 1995 - this job got quickly too big for Stanford, and YAHOO spun out, to become one of the Web’s powerhouse of directories.

As of mid-1997, the Web doubles in size every six months; any attempt to collate or map the *entire* content of the Web seems beyond the task of human beings - even YAHOO reviews only a small portion of the total. Computers have thus stepped in to do the dirty work of exhaustively searching the Web, charting its growth, and keeping indices of its content. *Alta Vista*, a "Web-crawler" developed by Digital Equipment Corporation’s

Network Research Center, uses fuzzy-logic-based query capabilities, combined with an extensive database of Web content (Alta Vista comes closest to "knowing" the entire content of the Web's tens of millions of document), to provide an approximately exhaustive search of the Web.

Presuming one could find the way to an interesting page within the Web, the next problem lay in finding the way back to that page. Early on, Mosaic adopted the "bookmark" interface, which maintained a list of URLs, but Web users quickly found that this list could grow into thousands of entries. Even at just a few additions a day, the bookmarks themselves would quickly outgrow the limited human abilities of searching and sorting. Bookmarks can grow so big that they need indexes into themselves!

All of this endless indexing and organizing points to a basic fact about the nature of the World Wide Web. The Web is "hyperspace;" that is, every "anchor" is directly linked to another point in the Web. You don't have to travel through anything as you cruise the Web; in short, there's no *there* there. While that's no problem for the computers - hyperspace comes naturally to them - it's a big problem for people, who have no sense of how to understand hyperspace. It eludes rational thought.

The best example of this is the ever-faithful URL. Perhaps the most useful construct in the history of the Internet, at its heart it's something built for computers and designed around what they can understand. To refer people to my home page in the Web, I tell them, "go to http://www.hyperreal.com/~mpesce" - and then I have to write it down for them, because they'll never remember it.

A URL is a message to the computer; as human beings, we have no hook to hang it on. Human beings don't think about things like URLs, or even lists. We think in very concrete terms, and always use the real-world as our guide. They're all we know. If I told someone, "Go to Market Street, down to Third, make a right, and take Third down to Bryant Street," that they'd understand. It's still a set of directions, but it makes sense, because it speaks to things human beings are very familiar with.

The Web creates an abstract knowledge space; useful, but hardly human. We can't grow to rely on bookmarks and search facilities forever - they aren't the way we work. We have to teach our computers to present information with a human focus: we search, explore, and accidentally stumble upon truth. Rarely can we search every possibility, and absorb an entire field of knowledge in search of one fact. We need a Web that we can stumble through, using our intuition *and* our intellect as guides.

Consider a trip to the library. How many of us have gone into the building looking for a specific title, and come out with another, more relevant text - one we never even knew we were looking for - that lay beside it or on another shelf nearby? Accident and chaos are critical to learning; we reinforce our own ignorance if we confine ourselves to the well-known.

Most of us cruise around the Web like drunks on a random walk. We often find things that are interesting, sometimes find things that are relevant, but we never know what we might be missing. The Web indexes, as useful as they are, will only tell you what the computer considers relevant, and that’s because the Web is still centered around the computer, not around the human. When humans can organize the Web *on their own terms*, we’ll have an information environment closer to the real world.

The Web needs to be more human. It’s good, but still makes us conform to the way it works. We need to turn the tables, and make it conform to the way we think and feel. The Web must become an accidental, explorative, intuitive environment; it's the single most important component in expanding individual human understanding.

Start Making Sense

With the development of multitasking operating systems in the early 1960s, which enabled the computer to execute several simultaneous tasks, scientists began to prototype new interfaces that could put these new capabilities to good use. One of these individuals was a graduate student at the Massachusetts Institute of Technology by the name of Ivan Sutherland. For his doctoral thesis, Sutherland developed the first interactive drawing and design program, *SketchPad*, within which users could draw, paint, and transform computer-generated images, all in real-time. Using a light pen, icons, and windows - all radical innovations to the human-computer interface - Sutherland's design completely revisioned how people thought of computers. They could now augment human creativity as effortlessly as they could do the payroll.

Sutherland moved from MIT to the University of Utah in the mid-60s, and, over several years, invented all of the major components of what we today call *virtual reality* (VR) - body tracking, the head-mounted display, and real-time three-dimensional graphics processors, to name but a few. It all seemed very "far out" in 1968, but the U.S. Defense Department understood the power of simulation - as generals have for thousands of years - and Sutherland's work laid the foundation for both the field of interactive computer graphics and real-time simulation systems.

It may well be that by the end of the millennium, Sutherland will be remembered as the individual who most helped to shape computing in the 20th century. The innovations that most of us will be using day-to-day in the 21st century - interactive computer graphics, human-centered computer interfaces and virtual reality - were first conceived of and created by Sutherland.

Systems capable of real-time simulation remained far too expensive for commercial or personal use through the 1960s and 1970s. The microprocessor revolution, which began to encroach on mainframe systems in the early 1980s, brought simulation into the high-end of commercial viability. NASA's Ames Research Center began the *Virtual Interface Environment Workstation* (VIEW) project in 1982, and Scott Fisher, Scott Foster, Elizabeth Wenzel, and Warren Robinett (among many others) brought Sutherland's work into the age of workstations. Capitalizing on new developments in low-power, lightweight Liquid Crystal Displays - the group created the modern head-mounted display. A gentleman named Jaron Lanier, working with Thomas Zimmerman of Stanford University, developed the *DataGlove*. The DataGlove is, as the name would suggest, a glove-like device which fits over the hand and then tracks its motion. The head-mounted display, inconjunction with the DataGlove, allows individuals to experience *immersion*, the sense of being placed inside a simulated environment. The head-mount and the DataGlove caused as great a revisioning of interface as Sutherland's work had twenty years before.

At the heart of this lay a complete paradigm shift in computing. Sensuality, which had never been thought of as a component of computing, suddenly became *the* important ingredient in interface design. Brenda Laurel, in her seminal work, *Computers as Theater,* structured an interface based upon Aristotelian poetics.

According to Laurel, the computer is like a stage set with actors. The dramatic arc (exposition, inciting incident, rising action, crisis, climax, falling action, and denoument) *appeals to our emotions*, rather than our intellect, and is the most effective way to design a user interface. The overwhelming popularity of games like *DOOM* and *Myst* clearly indicate that a well constructed stage and plot - even a simplistic one - can create a engaging, emotional experience.

In order to make computers clear to our minds, we have to teach them to speak to our hearts. Drama, music, architecture (the shape of a place) all have important places in computing.

Virtual reality has not lived up to the promises originally made a decade ago. Now that the hyperbole about virtual reality has abated, researchers understand that virtual reality is a methodology, not an end-point. The visions of disembodied travel through a sensually vacuous space have evolved into a set of techniques for making interfaces more human-centered.

What we've learned through all of this trial and error is that *sensuality makes sense*. If you can present data sensually, it will make more sense to the user. Technologies like three-dimensional displays, body tracking, spatialized audio, combined with dramatic narrative techniques create an experience which conforms to human expectations. Instead of looking at a list of numbers, why not turn it into a mountain range you can fly through, or a crescendo of sound? Instead of a keyboard, why not send the computer hand gesture? We’ve found that if we can use the computer in our heads - our cerebral cortex - we can often find things that the computer would never notice or consider significant.

Bringing this back to the Web, imagine an Internet interface where data sources - books, sounds, movies - could be represented naturally, as they are in the real world, with real-world metaphors. People *can* remember real-world metaphors, because they make sense. We have an intrinsic, biological understanding of the real world. If we didn't, when we got out of bed in the morning, we'd immediately fall and hit the floor. We organize our lives sensually - think of your record collection or your books - and need to bring that same technique to the Internet if we ever hope to be able to use it to our fullest capacity.

The last twenty-five years have seen a continuing set of improvements to the Internet in two areas; connectivity and interface. We've moved from a world of gurus and impossibly cryptic interfaces into an era of point-and-click browsing. That alone was enough to completely change how people thought of the Internet, and how they used it.

Immediately before us lies a transition as significant as all those that preceded it. We're about to make the Internet a human space, habitable, hospitable, intuitive and warm. The Internet has been a airy space of the intellect, we're rapidly to make it a place for drama, emotion, and sensuality - qualities of the heart.